

i General information

- Standard duration of the examination: **3 hours**.
- Number of questions: **18**.
- Number of available points: **20**. Each correctly answered question earns you one point, except for two questions, revealingly titled "Convolutional and Pooling layer" and "Vanishing gradients", which may earn you two points each.
- You will have no access to external props: no books, slides, internet, calculators, and phones.
- From 15:00, the course coordinator will be present in the exam room to answer questions about the exam.
- Question types:
 - **Multiple choice (one correct answer): 7 questions**
 - **Multiple choice (multiple possible correct answers): 1 question** --> ATTENTION: selecting the wrong answer(s) zeroes the gain from that question!
 - **Free text: 3 questions** --> Show your knowledge/understanding of the issue at hand in a relatively short form.
 - **Basic calculations: 7 questions** --> No complex calculations! In case you have to report your calculations, consider utilizing the character * to represent the multiplication operators (i.e. both for the basic multiplication and the "dot product" between matrices).

Grading of the digital examination -- information cloned from "Canvas TIG122 VT24 > Syllabus" --

The Bonus points will be added to the examination score, after being transformed by the "ceiling" function, i.e. the Bonus points will be transformed into the least integer number greater than or equal to their value (e.g. 0.5 Bonus points are transformed into 1).

Total score = Examination score + ceil(Bonus points)

- Total score ≥ 14 --> VG
- $8 \leq$ Total score < 14 --> G
- Total score < 8 --> U

In case the Total score misses the pass threshold (i.e. Total score = 8), the Bonus points will be reconsidered during reexamination.

1 Turing test

Discuss the structure of the Turing test.

Is that a good way to tell if a computer program is "intelligent"? Why/why not?

Fill in your answer here

Format ▾ | ↺ | ↻ | ✎ | Σ | ✕

Words: 0

Maximum marks: 1

2 Decision trees

Which statement is correct?

Select one alternative:

- The last step in random forest classification is voting.
- The point of bootstrapping is to permute the order of the table rows.
- Decision trees is an ensemble method.

Maximum marks: 1

3 Regression trees

Which statement is correct?

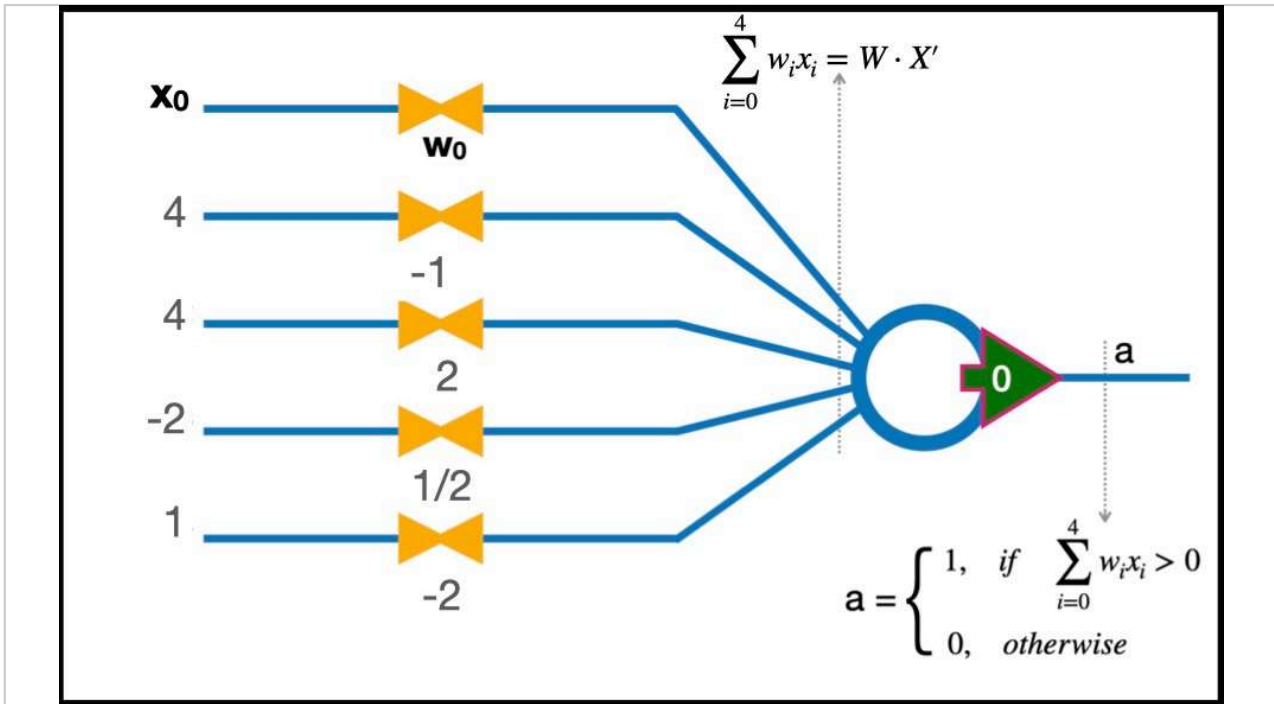
Select one alternative:

- When constructing regression trees, the goal is to minimize the number of Gini dimensions.
- When constructing regression trees, the goal is to minimize the Gini impurity.
- When constructing regression trees, the goal is to minimize the sum of the squared residuals.

Maximum marks: 1

Maximum marks: 1

6 Step activation with extra input



Consider the step function unit from the previous question, with its input and weight distribution, and a threshold $\theta = -2$.

Complete the equivalent description of the unit by indicating the values for x_0 and w_0 (see figure above). Calculate the weighted sum and its corresponding activation. (report your calculations)

Fill in your answer here

Format | ↺ | ↻ | ✂ | ✎ | Σ | ✖

Words: 0

Maximum marks: 1

7 Activation function

A unit (artificial neuron) generates an activation value of 2. Can its activation function be ReLU?

Select one alternative:

- Yes.
- No.
- It depends on the threshold.

Maximum marks: 1

10 Layout 2

```
model = Sequential()  
model.add(Dense(10, activation='sigmoid', input_shape=(100,)))  
model.add(Dense(10, activation='softmax'))
```

Consider the code snippet in the figure above.

Briefly describe the network layout. For each layer, specify the number of weights and biases.

Fill in your answer here

Format ▾ | ↺ | ↻

✎ | Σ | ✕

Words: 0

Maximum marks: 1

11 Backpropagation 1

Is the sign function a suitable activation function for training by backpropagation? Why?

Fill in your answer here


Format

▾

↶

|

|



|

Σ


|



Words: 0

Maximum marks: 1

12 Backpropagation 2



Refer to the figure. During training by backpropagation, what is the problem related to the argument of the sigmoid activation function falling in region A?

Select one alternative:

- Low activation value.
- Low value of the function's slope.
- Negative value of the function's argument.

Maximum marks: 1

13 Learning rate

What is the effect of the learning rate η ?

Select one alternative:

- To stabilize the process of weight convergence.
- To prevent overfitting.
- To slow down learning.

Maximum marks: 1

14 Batches

With reference to the use of mini-batches in backpropagation, which statements are correct?

Select one or more alternatives:

- During each epoch, each item in the training set is shown to the network once and only once.
- Splitting the training data into smaller batches can result in increased memory usage by the system.
- The sequence of training samples is different for each epoch.

Maximum marks: 1

15 Number of layers

Imagine you have trained a network for image classification with three hidden layers. Increasing the number of layers to four does not result in any increment in the performance. Based on that observation, what is the most reasonable action to take next?

Select one alternative:

- Test the performance after doubling the number of units per layer.
- Test the performance for five hidden layers.
- Test the performance for two hidden layers.

Maximum marks: 1

16 Dropout

After training a neural network while using dropout to facilitate model generalization, which measures must be taken during validation or use?

Select one alternative:

- Divide the weights at each layer by a factor that is consistent with the average dropout rate.
- Apply the input and obtain the output as if dropout was not used.
- Multiply the weights at each layer by a factor that is consistent with the dropout rate at that layer.

Maximum marks: 1

17 Vanishing gradients

What is the origin of the problem known as "vanishing gradients"? Is it more or less severe in deeper networks, i.e. networks with a higher number of layers? Name (at least) one possible solution to the problem and briefly explain why it can be effective.

Fill in your answer here

Format ▾ | ↶ | ↷ | ✎ | Σ | ✕

Words: 0

Maximum marks: 2

18 Convolutional and pooling layers

Consider an $8 \times 8 \times 3$ RGB image (i.e. an 8×8 image with three color channels) with no padding. Apply one 3×3 kernel with a stride of 1. What is the size and depth of the resulting activation map?

Then, apply onto such an activation map a 2×2 max pooling layer with a stride of 2. What is the size and depth of the final map?

Fill in your answer here

Format ▼ | ↺ | ↻ | ↶

✎ | Σ | ✖

Words: 0

Maximum marks: 2