UNIVERSITY
OF SKÖVDE

School of Information Technology

WRITTEN EXAMINATION

Course: Big Data Programming

Examination: Final written exam

Course code: IT739A

Credits for written examination: 5.5hp

Date: March 27, 2026

Examination time: 8:15 – 12:30

Examination responsible: Richard Senington

Teachers concerned: Juhee Bae, Richard Senington, Gunnar Mathiason

Aid at the exam/appendices

Other

Instructions

- ☐ Take a new sheet of paper for each teacher.
- ☒ Take a new sheet of paper when starting a new question.
- ☒ Write only on one side of the paper.
- ☒ Write your name and personal ID No. on all pages you hand in.
- ☒ Use page numbering.
- ☒ Don´t use a red pen.
- ☒ Mark answered questions with a cross on the cover sheet.

Grade points

Examination results should be made public within 18 working days

Good luck!

UNIVERSITY
OF SKÖVDE

IT739A, Exam, March 27th, 2026, at 8:15-12:30

Instructions that must be followed:

- The max number of points listed with each question indicates how thoroughly you are expected to answer the question.
- Number your answer clearly with the same numbering as the question and its sub-questions.
- Be clear in your writing. Sometimes fewer but more concise formulations are better than writing a lot of text.
- More points assigned indicate that a more detailed answer is expected.

The grading of your answers considers:

- The **correctness** of your answer/explanation
- The **clarity and logical cohesion** of your answer/explanation
- The **conciseness** of your answer/explanation

Grading levels:

- There are 5 questions (with sub-questions), which give a maximum 10 points each.
1) You need to pass 5 points per each of the 5 questions to pass the exam (25 points).
2) Accumulated points above that threshold linearly determines the exam grading (the exam maximum is 50 points).
- Exam grading range is A to F, which also determines the final course grading:
A: 45-50 points, B: 40-44 points, C: 35-39 points, D: 30-34 points, E: 25-29 points.

Examination goals:

This exam assesses the following Course Curriculum goals

- *critically evaluate the relevance and application of advanced AI methods,*
- *in an in-depth manner account for the development of advanced AI methods,*
- *apply and demonstrate appropriate analysis strategies using advanced AI tools and techniques,*
- *in an in-depth manner account for ethical use of AI to solve complex problems*

The exam questions begin on the next page.

UNIVERSITY
OF SKÖVDE

Exam questions:

1. Deep Learning, Keras and Tensorflow (10 points):

1. Describe how the *Loss function* and the *Optimizer* interact while training a neural network in a typical supervised learning setup. What are their differences? (2p)
2. Give two examples for how to avoid the vanishing gradients problem when stacking many layers in a deep neural network. (2p)
3. What is the underlying assumption about the data used with CNNs for being very effective in fitting a model to that data, using fewer weights than with a fully connected neural network? (2p)
4. Give the name of the computation architecture that makes TPUs much faster than GPUs at calculations that are involved in updating weights in neural networks. Also, why is this architecture much faster? (4p)

2. Anomaly Detection (10 points):

1. Describe the meaning of an anomaly, and hence anomaly detection. (2p)
2. What are the three primary components of an autoencoder? Explain how they work together, and describe one variation on an autoencoder and how it tries to improve over the basic design. (4p)
3. Describe how a GAN works and how to use it for anomaly detection, including advantages and disadvantages of this method. (3p)
4. Explain what YOLO is and how it differs from normal computer vision. (1p)

3. Recurrent Neural Networks (10 points):

1. Describe what a recurrent network is, how it differs from a standard neural network and why it is effective for working with data series. (3p)
2. What is the design difference between a GRU and a basic recurrent network? What is the benefit of this approach? (3p)
3. How can recurrent networks be used for anomaly detection? (2p)
4. What are the weaknesses of recurrent networks? (2p)

4. Concepts and ethics in Big Data (10 points):

1. One of the 5Vs of Big Data is Value. Please describe the Value aspect of big data and how it relates to Deep Learning. (3p)
2. Describe two different distributed data storage systems in general terms, and how they can be scaled to support big data. (3p)

3. When we perform deep learning on personal data (for example, medical data) we are looking for common patterns. Discuss the ethical implications of this action. (4p)

5. Large Language Models (LLM) (10 points):

1. Compare Prompting, RAG, and LoRA in terms of how they influence the behavior of a large language model. Justify your answers.
   a) Which of these methods directly modifies the model parameters, and which does not? (1p)
   b) A company needs to update model knowledge frequently without retraining. Which method is most appropriate and why? (1p)
   c) What is the main difference between prompting and RAG? (1p)
2. Explain how a Transformer determines which tokens in a sequence should influence the representation of another token in the self-attention mechanism. (3p)
3. Large language models can be efficiently adapted to new tasks using Low-Rank Adaptation (LoRA).
   a) Explain the primary objective of LoRA when adapting large language models. (1p)
   b) Describe the main idea behind the LoRA approach and how it enables efficient fine-tuning of large models. (2p)
   c) Identify one limitation of LoRA and suggest one possible approach to mitigate this limitation. (1p)