

Information page for written examinations at Linköping University



Examination date	2021-10-25
Room (2)	FE246(2) <u>TER4(41)</u>
Time	14-18
Edu. code	TSBB19
Module	TEN1
Edu. code name Module name	Machine Learning for Computer Vision (Maskininlärning för datorseende) Written examination (Skriftlig tentamen)
Department	ISY
Number of questions in the examination	16
Teacher responsible/contact person during the exam time	Per-Erik Forssén
Contact number during the exam time	013-285654
Visit to the examination room approximately	14.30-15 and 16-16.30
Name and contact details to the course administrator (name + phone nr + mail)	Carina Lindström, 013-28 44 23
Equipment permitted	Dictionary Swedish-English-Swedish
Other important information	
Number of exams in the bag	

TSBB19 Machine Learning for Computer Vision

Written Exam 2021-10-25

Instructions: This exam consists of **16** questions that require description of terms, phenomena, relations, etc. Each question gives a score [**0,1,2**], for:

- 0:** Wrong/No Answer
- 1:** Almost Complete/Partially Correct Answer
- 2:** Complete and Correct Answer

Each exam question can give a maximum score of **2p**, and a total of **32p** for the whole exam.

In order to pass with grade **3**, at least **15p** are required.

In order to pass with grade **4**, at least **22p** are required.

In order to pass with grade **5**, at least **27p** are required.

All tasks should be answered on **separate sheets**. It is fine to answer multiple questions on one sheet. Be brief and to the point.

Write your AID-number and the date on all paper sheets that you hand in. In addition, these sheets should be numbered in consecutive order.

Good luck!

Per-Erik Forssén and Michael Felsberg

- Question 1:** In the triplet loss, an anchor(A), a positive(P), and a negative(N) sample are used together. How should these be chosen?
- Question 2:** When the DCF is generalized to multiple channels and multiple samples, an approximation is made if the optimal filter is computed incrementally with separate numerator and denominator. Explain which one and why.
- Question 3:** Constellation models for object detection are typically star shaped. What does this mean, and why is it preferable to an explicit/pairwise model?
- Question 4:** For a descriptor, both invariance, and discriminative power are important properties. Explain these concepts.
- Question 5:** In contrast to ResNet, very deep VGG networks suffer from the vanishing gradient problem. What is the main difference between the two approaches that causes this different behavior? Why do very deep networks suffer from vanishing gradients?
- Question 6:** Region proposals are sometimes used in object detection. Explain what these are, and why they are useful.
- Question 7:** What is the special case of the softmax activation for binary problems? Which activation function is more common for hidden layers?
- Question 8:** In descriptor matching, a distance ratio is often used instead of a distance. Explain *which* distance ratio, and why it is useful.
- Question 9:** What statistical assumption can be used to justify descriptor matching using scalar products?
- Question 10:** Define false positive rate(FPR), and explain what happens to this measure when positives are rare, as in e.g. object detection.
- Question 11:** With the common availability of GPUs, DCFs are more commonly optimized in the spatial domain. Which drawback of the Fourier domain approach is avoided and why are GPUs addressing the drawback of spatial convolutions?
- Question 12:** Why is it important to have accurate ground truth bounding boxes to measure accuracy differences of tracking methods? The Jaccard index can be computed without counting pixels if axis-aligned bounding boxes are used. How?
- Question 13:** What happens if all channels in the respective layer are initialized with the same weights? What happens if additionally dropout is used?
- Question 14:** Explain what a Bag of features(BoF) is, and what it is used for.
- Question 15:** We need to optimize the learning rate. Do we use test data or validation data for this purpose? We run 10 instances of a network (random initializations) on the test data. Which results must be reported?
- Question 16:** In the ATOM tracker, a bounding box regressor is combined with a classifier to improve accuracy. The classifier may use deep or shallow features in general. Which option is better for ATOM and why?