

i Front page

Trial examination in TEK5040

2020 FALL

Duration: 9.12.2020, 09:00 to 16:00

General information:

- It is important that you check the course's semester page regularly. Important messages during the exam will be posted on the semester page. The messages during the exam can be posted in Canvas instead of the semester page if your course uses Canvas.
- Remember that your submission need to be anonymous, do not write your name in your submission.
- All examination support materials are permitted. You need to gather information from available sources, assess the information quality, and put it together in a submission based on your own processing of the content. The submission must reflect your individual level of knowledge.
- Answers can be given in Norwegian or English. All questions should be answered. The questions are weighted differently. The maximum score is shown for each question.
- If a screenshot is small on your screen, try zooming:
Press the keys ctrl and + at the same time
or
Press ctrl and scroll on the mouse wheel

Collaboration

- It is not permitted to communicate with other people about the exam and its questions during the exam. Plagiarism is not allowed. Do not directly copy text from lecture slides, books or other sources.
- You can be selected for a control interview on your examination answer, in order to determine your ownership of the answer. This discussion will not affect the grade, but can lead to the Department issuing a suspicion-of-cheating case.
- UiO's routines for handling suspicion of cheating can be found here: <https://www.uio.no/english/about/regulations/studies/studies-examinations/routines-cheating.html>

Contact:













- [User support for exams](#)

Good luck!

1 Stride in Convolutional Neural Networks

Assume you have an input image of size 600×300 , and that you apply a convolutional layer with kernel size $[3, 3]$ and stride 2 in both directions. What is the (approximate) size of the output of the convolutional layer.

Fill in your answer here

Format | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  |  |  |  |  |  | 











Words: 0

Maximum marks: 12

2 Image Captioning

Recurrent Neural Networks can be trained to generate image captions with supervised learning, using cross-entropy loss on the prediction of a word. We may evaluate this loss on a validation set as well to make sure we are not overfitting. Low validation loss is not necessarily any guarantee for good performance however, what is an effect the validation loss is not able to capture?

Fill in your answer here

Format | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  |  |  |  | Σ | ABC | 

Words: 0

Maximum marks: 12

3 Proximal Policy Optimization

Assume π is the current policy. The task of finding the optimal policy $\tilde{\pi}$ may then be written as

$$\max_{\tilde{\pi}} \sum_s \rho_{\tilde{\pi}}(s) \sum_a \tilde{\pi}(a|s) d_{\pi}(s, a)$$

where $\rho_{\tilde{\pi}}$ is the discounted state-visitation frequencies of policy $\tilde{\pi}$ and $d_{\pi}(s, a) = q_{\pi}(s, a) - v_{\pi}(s)$ (where q_{π} and v_{π} are the action-value and state-value functions of policy π respectively). In order to get to the unclipped objective of PPO we needed to make *three* approximations to the above objective, what are they?

Fill in your answer here

Format
-
B
I
U
 x_2
 x^2
 \int_x

 Ω

 Σ
ABC











Words: 0

Maximum marks: 12

4 Generative Adversarial Networks

Assume that you have a set of color images x_1, x_2, \dots, x_N and you want to train a neural network to *colorize* grayscale images, i.e. so that given a grayscale image x' the network may generate different plausible color images. As an application, you could then use this network to possibly colorize old digitized grayscale photos. Briefly describe how you can formulate this as a conditional GAN problem.

Fill in your answer here

Format | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  |  |  |  | Σ | ABC | 

Words: 0

Maximum marks: 12

5 Learning Concepts

Meta-learning can be formulated as an optimization problem with the following objective function :

$$[\theta^*, \phi^*] = \arg \max_{\theta, \phi} \frac{1}{N} \sum_{n=1}^N \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{D}_{query}^{(n)}} p_{g_\phi(\theta, \mathcal{D}_{support}^{(n)})}(\mathbf{y}|\mathbf{x})$$

where $\mathcal{D}_{support}^{(n)}$ and $\mathcal{D}_{query}^{(n)}$ are respectively the n^{th} support set and query set of the meta training set, whereas \mathbf{x} and \mathbf{y} are respectively the input and label of any data sample. $p_{g_\phi(\theta, \mathcal{D}_{support}^{(n)})}$ is a probability distribution parameterized by the output of a function $g_\phi(\cdot, \cdot)$ which itself is parameterized by ϕ . Describe a meta-learning technique which makes use of this objective function.

Fill in your answer here

Format
-
B
I
U
x₂
x²
I_x
📄
📄
↶
↷
↺
≡
≡
Ω
🔍
Σ
ABC
✖

Words: 0

Maximum marks: 14











6(a) Bayesian Deep Learning

In variational inference, we try to approximate the posterior distribution $p(\mathbf{w}|\mathcal{D})$ with an variational distribution $q(\mathbf{w})$, where \mathbf{w} and \mathcal{D} are network parameters and training data respectively. This is achieved by minimizing the KL divergence between the two distributions defined by

$$\text{KL}(q(\mathbf{w})||p(\mathbf{w}|\mathcal{D})) = -\mathbb{E}_{q(\mathbf{w})} \ln \frac{p(\mathbf{w}|\mathcal{D})}{q(\mathbf{w})}.$$

Explain why it is difficult to minimize the KL divergence directly.

Fill in your answer here

Format | **B** | *I* | U | x_2 | x^2 | \int_x |  |  |  |  |  |  |  | Ω |  |  | Σ | ABC | 

Words: 0

Maximum marks: 4

6(b) **Bayesian Deep Learning**

Gradient of ELBO, $\nabla_{\lambda} \mathcal{L}(\lambda)$ can be written as

$$\nabla_{\lambda} \mathcal{L}(\lambda) = \int \left[\ln p(\mathbf{w}, \mathcal{D}) - \ln q(\mathbf{w}, \lambda) \right] \nabla_{\lambda} [q(\mathbf{w}, \lambda)] d\mathbf{w},$$

where λ parameterizes the distribution $q(\cdot, \cdot)$.

Briefly explain how do you convert this to the following form and the motivation for this conversion :

$$\nabla_{\lambda} \mathcal{L}(\lambda) = \mathbb{E}_{q(\mathbf{w}, \lambda)} \left\{ \left[\ln p(\mathbf{w}, \mathcal{D}) - \ln q(\mathbf{w}, \lambda) \right] \nabla_{\lambda} [\ln q(\mathbf{w}, \lambda)] \right\}$$

Fill in your answer here

Format | **B** | *I* | U | \times_2 | \times^2 | \int_x | | | | | | | Ω | | | Σ | ABC |











Words: 0

Maximum marks: 8

7(a) Deep Learning for Control

Inverse reinforcement learning (IRL) has been proposed to be a solution for the problem that manual design of reward functions for complex reinforcement learning problems is difficult. Give two challenges in IRL.

Fill in your answer here

Format | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  |  |  |  | Σ | ABC | 

Words: 0

Maximum marks: 4

7(b) **Deep Learning for Control**

The following table shows the equivalence between Generative Adversarial Network (GAN) and Generative Adversarial Imitation Learning (GAIL). Fill the missing entries in the column for GAIL.

GAN	GAIL
Generator(G)	?
Discriminator (D)	Discriminator (D)
Real data samples (x)	?
Synthesized data samples (z)	?
G maximizes probability of z	?
D maximizes probability of x	?
D minimizes probability of z	?

Fill in your answer here

Format
B
I
U
 x_2
 x^2
 I_x
↶
↷
↺
↻
☰
☷
Ω
☒
Σ
ABC
✖

Words: 0

Maximum marks: 8

8(a) **3D Processing**

Describe how the PointNet can be used for segmentation of a 3D point cloud.

Fill in your answer here

Format | **B** | *I* | U | x_2 | x^2 | I_x | | | | | | | Ω | | | Σ | ABC |










Words: 0

Maximum marks: 8

8(b) 3D Processing

There are n points in a point cloud and each point has three components (x, y, z) . Using the Edge Convolution based on a Multi-layer Perceptron (MLP) of 16 outputs, a graph convolution operation is performed on the point cloud. Briefly describe the operation. What are the dimensions of the resulting structure of this graph convolution?

Fill in your answer here


Format | **B** | *I* | U | x_2 | x^2 | I_x |  |  |  |  |  |  | Ω |  |  | Σ | ABC | 

Words: 0

Maximum marks: 6


9 New Question

Here you can upload a PDF file which contains items such as diagrams and equations etc. related to the previous questions. Remember to mark every item in the file with the relevant question number.



Upload your file here. Maximum one file.

All file types are allowed. Maximum file size is **1 GB**

 Select file to upload

Maximum marks: 0