EE-559 Deep learning – Submission 1, Students' questions

A question is denoted by **Q**; the corresponding answer is denoted by **A**. The questions-related exercises are marked by their numbers in Marked_exercises_submission1_lecture$i$.pdf and Marked_exercises_submission1_lecture$i$.ipynb documents.

Q: My assert in exercise 1.2.2 for MSE (image, noisy_image) does not pass.
A: That might be because of the issue with the way Microsoft Windows reads the image file: (i) check other asserts, they should work or (ii) try running notebook in Gnoto.

Q: Will the notebooks for Submission 1 be released weekly?
A: The notebooks for *Submission 1* will be released weekly.

Q: Are we expected to write additional comments in the code? Are we allowed to re-use the code from the other cells in the notebook?
A: If we expect you to write the comments, we will explicitly mention this in the notebook. If needed, we will re-run cells in your notebook, consequently, so you can use the functions defined in earlier cells.

Q: When is the deadline for the first submission?
A: Please, refer to Lecture 2 slides, p.26, "Submissions of marked exercises". The deadline is April 8. 23:59. You will need to collect and submit all the notebooks related to the first submission. For the instructions for Submission 1, please refer to Week 5 Lab slides, p.3 and to Checklist for Submission 1 which is available under the Submission 1 page. You DO NOT need to submit graded exercises after each week.

Q: What is a Siamese network?
A: Siamese networks are a special class of models, which focus on measuring the similarity and dissimilarity between the pair of inputs. Typically, two identical copies of a network are used to obtain the embeddings of the pair of inputs: (x_{i1}, x_{i2}). The corresponding ground truth label y_i indicates whether the input samples 1 and 2 belong to the same ($y_i = 1$) or different ($y_i = 0$) classes. The network weights are then updated with respect to the calculated loss:

$$L(\boldsymbol{x_1}, \boldsymbol{x_2}, \boldsymbol{y}, m) = \frac{1}{N} \sum_{i=0}^{N-1} \Big(y_i \cdot d(x_{i1}, x_{i2})^2 + (1 - y_i) \cdot \max\Big(0, m - d(x_{i1}, x_{i2})\Big)^2 \Big),$$

where $d(\cdot)$ is a specified distance metric, and N – number of samples in a dataset.

Q: My asserts have passed, am I correct?
A: Yes, as long as you do not modify the assert and if the output checked is a result of a genuine computational process.

Q: Does my code need to be "pretty"?
A: No, you are not being evaluated on how efficient your code is. However, your code should be executable, compilable within reasonable time, and addressing the question in the practice or marked assignment.

Q: Am I allowed to change the hyperparameters for submission 1?
A: No, you should not be changing the values of hyperparameters for the Submission 1 exercises.