Hellenic Mediterranean University
School of Engineering
Department of Electrical & Computer Engineering

Exercise #4

1. You are given an input image (x), kernel (w) and bias (b). Your task is to evaluate the shaded pixel in the image after the convolution. The origin of the kernel is the shaded pixel. Assume we are using zero padding.

$x =$

-2	4	5	-4	1
2	1	7	5	6
7	3	2	-9	7
8	2	3	4	9
6	-1	0	1	6

$w =$

2	5	-6
0	2	4
5	1	3

$b =$ 2

2. After training a neural network, you observe a large gap between the training accuracy (100%) and the test accuracy (42%). Which of the following methods is commonly used to reduce this gap?

(i) Generative Adversarial Networks
(ii) Dropout
(iii) Sigmoid activation
(iv) RMSprop optimizer

3. Which of the following propositions are true about a CONV layer? (Check all that apply.)

(i) The number of weights depends on the depth of the input volume.
(ii) The number of biases is equal to the number of filters.
(iii) The total number of parameters depends on the stride.
(iv) The total number of parameters depends on the padding.

4. You want to solve a classification task. You first train your network on 20 samples. Training converges, but the training loss is very high. You then decide to train this network on 10,000 examples. Is your approach to fixing the problem, correct? If yes, explain the most likely results of training with 10,000 examples. If not, give a solution to this problem.

5. You have a dataset D1 with 1 million labelled training examples for classification, and dataset D2 with 100 labelled training examples. Your friend trains a model from scratch on dataset D2. You decide to train on D1, and then apply transfer learning to train on D2. State one problem your friend is likely to find with his approach. How does your approach address this problem?

6. Consider the following CNN classifier. For each layer, calculate the number of weights, number of biases and the size of the associated feature maps. The notation follows the convention:
• *CONV-K-N* denotes a convolutional layer with N filters, each of them of size *K x K*, Padding and stride parameters are always 0 and 1 respectively.
• *POOL-K* indicates a *K x K* pooling layer with stride K and padding 0.
• *FC-N* stands for a fully connected layer with N neurons.

Layer	Activation map dimensions	Number of weights	Number of biases
INPUT	$128 \times 128 \times 3$	0	0
CONV-9-32			
POOL-2			
CONV-5-64			
POOL-2			
CONV-5-64			
POOL-2			
FC-3			

7. Which of the following techniques can be used to reduce model overfitting?
(i) Data augmentation
(ii) Dropout
(iii) Batch Normalization
(iv) Using Adam instead of SGD

8. You are benchmarking runtimes for layers commonly encountered in CNNs. Which of the following would you expect to be the fastest (in terms of floating-point operations)?

(i) Conv layer (convolution operation + bias addition)

(ii) Max pooling

(iii) Average pooling

(iv) Batch Normalization

Good luck