Artificial Neural Networks and Deep Learning - 03/09/2020 Exam (270 Points)

During the whole exam you should have only one screen active connected to the Zoom meeting that you can access from the POLIMI portal. No external monitors or virtual screens are allowed.

You should only have a browser open and no other application besides a text editor open next to the browser to keep a safety copy of your answers. The browser should have only the exam tabs open. The exam is saved into the browser cookies so if you close it and reopen it nothing should happen, but we cannot guarantee.

Do not maximize the window of your browser.

You do not need to work on paper so keep you eyes on the browser and do not look around you. You cannot go out of sight. Your microphone should be on all the time.

You will have a fixed amount of time, after which the Form is automatically closed and there will no possibility to submit your answers any more. You will be noticed 15' in advance by the teacher not to miss the delivery time. Exams which are not submitted withing the given time will be considered as RITIRATO.

In case you do not see the IMAGES try to reload the page.

Section 1

QUESTION 1: MACHINE LEARNING / DEEP LEARNING

Deep learning uses very complex models rich in parameters. Keeping this in mind answer the following.

1

Is data representation, i.e., features, learned via deep learning always better than hand-crafted features? Why?
(10 Points)

Enter your answer

2

How data representation is learned via convolutional neural networks?
(20 Points)

Enter your answer

3

How data representation is leaned via recurrent neural network?
(20 Points)

Enter your answer

4

How data representation is learned via deep autoencoders?
(20 Points)

Enter your answer

Section 2

QUESTION 2: NEURAL NETWORKS TRAINING

Deep neural networks a like any other learning machine, they can be a benefit or ... they can overfit.

5

What is overfitting? What it is due to?
(20 Points)

Enter your answer

6

What is early stopping and why does it help with overfitting?
(20 Points)

Enter your answer

7

What is weight decay and why does it help with overfitting?
(20 Points)

Enter your answer

8

What is dropout and why does it help with overfitting?
(20 Points)

Enter your answer

Section 3

QUESTION 3: CONVOLUTIONAL NEURAL NETWORKS

With reference to the deep neural network in the image below, answer the following questions.

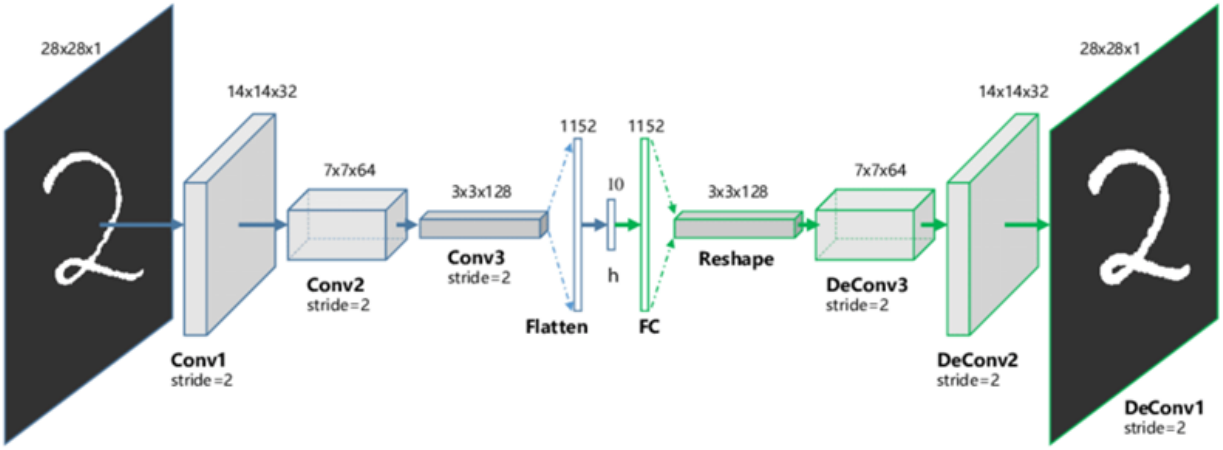

9

Detail the building blocks in each layer of the above network. Please indicate
- what kind of block is
- what is its size
- what is the overall number of parameters, together with a short description on how you compute them, e.g., 3x5x5=45 (not just 45!). Consider convolutions having a spatial 3x3 extent.

You can use the calculator, but we are more interested in the formulas than on numbers!
(20 Points)

Enter your answer

10

This network is called denoising autoencoder and it has the goal of cleaning or restoring a damaged image. How would you train it? What kind of loss function would you use?
(20 Points)

Enter your answer

Section 4

QUESTION 4: RECURRENT NEURAL NETWORKS

Sequences such as text, time series, and sensor streams are often processed via recurrent neural networks. With reference to a generic stream of input data, answer the following questions.

11

Do I really need a recurrent neural network to process a stream of input data? Why can't I use a standard feed forward architecture?
(10 Points)

Enter your answer

12

Classic recurrent neural networks suffer the vanishing gradient problem. What is it and what is it due to? Does it affects only recurrent networks?
(20 Points)

Enter your answer

13

What is ReLU and why it helps with vanishing gradient?
(15 Points)

Enter your answer

14

What is Xavier initialization and why it helps with vanishing gradient?
(15 Points)

Enter your answer

15

What is an LSTM and why it helps with vanishing gradient?
(20 Points)