

CS344 Exercise 7

Task 1: Convolutional Neural Networks

One advantage of convolutional layers as compared to fully connected layers is that convolutional layers normally have fewer parameters to learn during training.

TRUE

FALSE

A **valid** convolutional layer uses padding and its output has the same shape as its input.

TRUE

FALSE

What is the output shape of a **same** convolutional layer (filter size 3x3, stride of 1) whose input has shape (1200, 900)?

(1198, 898)

(1200, 900)

Every unit (filter) in a convolutional layer has the same number of parameters.

TRUE

FALSE

The values of the parameters are the same for every unit (filter) in a convolutional layer.

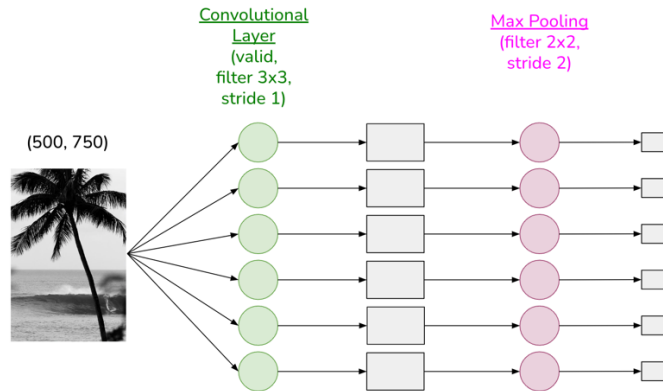
TRUE

FALSE

Convolutional layers don't use an activation function.

TRUE

FALSE



Consider the network architecture shown above where the input has shape (500, 750).
What is the shape of the output from each of the 6 Max Pooling units?

(249, 374)
 (498, 748)

(250, 375)
 (500, 750)

Consider the network architecture shown above where the input has shape (500, 750).
How many parameters does the network have that are learned during training?

54
 78

60
 84

Max pooling doesn't use an activation function.

TRUE

FALSE

Show the result (a 2D array filled with values) of the convolution below, assuming the convolution is valid and uses a stride of 1.

4	2	1	5
3	4	3	1
0	5	2	2
1	3	0	4

 \ast

-1	0	0
1	1	0
0	-1	1

 $=$

Show the result (a 2D array filled with values) of the convolution below, assuming a same convolution.

2	0	1
0	4	0
3	0	1

 *

-1	0	0
1	1	0
0	-1	1

 =

Show the result (a 2D array filled with values) of max pooling as shown below.

8	6	2
0	3	4
7	1	0

max

 =

Task 2: Coding with convolutional neural networks

Download the Jupyter Notebook for Exercise 7 from the course website. Open the Notebook in your web browser and work through it. As you work through the Notebook, answer the following questions.

Fashion MNIST

For which class did the model perform best (based on F1 score)? For which class did the model perform worst (based on F1 score)?

How many layers does this CNN have (only counting layers with trainable parameters)?

What percentage of the total parameters does the convolutional layer account for?

What is the accuracy of the trained model on the testing data?

Breast cancer ultrasound

Do you think the model is likely to be overfitting the training data? Why or why not?

For what class does the model perform worst?

Does the convolutional layer have the same number of parameters in both? Why or why not? Do the dense layers have the same number of parameters in both? Why or why not?

Flowers

For the model that performs best on the validation data, what is the model's F1 score on the validation data? How many parameters does this tuned model have?

What is your tuned model's accuracy on the testing data?

CS344 Exercise 7 Final Page

In the *TIME* column, please estimate the time you spent on this exercise. Please try to be as accurate as possible; this information will help us to design future exercises.

PART	TIME
Exercise	