

Unit 4 - Week 3

Course outline
How does an NPTEL online course work?
Week 1
Week 2
Week 3
<input checked="" type="radio"/> Introduction to Tensors <input checked="" type="radio"/> Mathematical Foundations of Deep Learning - Contd. <input type="radio"/> Building Data Pipelines for Tensorflow-Part 1 <input type="radio"/> Building Data Pipelines for Tensorflow-Part 2 <input type="radio"/> Building Data Pipelines for Tensorflow-Part 3 <input type="radio"/> Text Processing with Tensorflow <input type="radio"/> Quiz : Practice Assignment 3 <input checked="" type="radio"/> Quiz : Assignment 3 <input type="radio"/> Week 3 Feedback <input type="radio"/> Solution - Assignment 3
Week 4
Week 5
Week 6
Week 7
Week 8
Text Transcripts
Download Videos

Assignment 3

The due date for submitting this assignment has passed. **Due on 2020-02-19, 23:59 IST.**
 As per our records you have not submitted this assignment.

1) We have a 50 x 25 grid with 5 separate quantities for each entry of the grid. Each of the 5 quantities is further represented by 100 values. **1 point**
 If we want to represent the 50 such grids using a tensor, what is the best way to represent such a tensor (defined using NumPy), say X? Given below are the dimensions of 4 tensors. Choose the best option.

(50, 25, 5, 100, 50)
 (50, 50, 25, 5, 100)
 (25, 50, 100, 5, 50)
 (50, 50, 25, 100, 5)

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 (50, 50, 25, 5, 100)

2) From the data described in the above question, what will be the dimensions of the following tensor: **1 point**

Y = X[20:, 10:25, :50, :100, :50]

(5, 15, 5, 5, 50)
 (20, 15, 5, 100, 50)
 (30, 15, 25, 5, 50)
 Error
 None of the above

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 (30, 15, 25, 5, 50)

3) Let's take the same tensor X from Q1 and run the following code: **1 point**

Y = numpy.mean(X, axis = 3)

(25, 25, 5, 100)
 (50, 25, 25, 100)
 (25, 25, 100, 5)
 None of the above

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 None of the above

4) Given a tensor t, load the tensor into a tensorflow's tf.data.Dataset and run the reduce operation with parameters initial_state = 1 and reduce_func = lambda x, y: x + y on the dataset. Choose the result of the operations from the options below: **1 point**

t = [[1, 2, 3, 4, 5, 6, 7], [2, 3, 4, 5, 6, 7, 8], [3, 4, 5, 6, 7, 8, 9]]

[6 9 12 15 18 21 24]
 [7 10 13 16 19 22 25]
 [30 37 44]
 [28 35 42]

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 [7 10 13 16 19 22 25]

5) Using the same tensor from Q4, we load the tensor into a tensorflow's tf.data.Dataset and run the map operation with map_func = lambda x: (x + tf.reduce_mean(x)) . We then run the reduce operation with parameters initial_state = 5 and reduce_func = lambda x, y: 3*x - 4*y on the dataset. Choose the result of the operations from the options below: **1 point**

[-246, -298, -350, -402, -454, -506, -558]
 [-273, -325, -377, -429, -481, -533, -585]
 [-300, -352, -404, -456, -508, -560, -612]
 [-165, -217, -269, -321, -373, -425, -477]

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 [-165, -217, -269, -321, -373, -425, -477]

6) What will be the first, last elements and the size of the output array of the following code: **1 point**

```
x = tf.data.Dataset.range(1,10)
x_ = x.repeat(2).batch(20)
n = next(iter(x_))
print(n.numpy())
```

1, 2, 20
 1, 10, 20
 1, 9, 18
 1, 1, 10

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 1, 9, 18

7) For the given raw data, which of the following preprocessing steps do you think were necessarily applied before providing it for the training of a neural network? Hint: Only three are correct. **1 point**

```
time,status,age,month_year,col1,col2
10,A,76,Jul-1972,6.76,Yes
30,B,56,Jan-1968,,No
35,A,41,Nov-1977,1.34,
99,C,71,,2.9,No
185,B,52,Mar-1965,12.08,Yes
```

- Smooth out noisy columns
- Handle missing values
- Removing redundant columns
- Normalization
- Column month_year split
- Quantization
- One-hot encoding

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 Handle missing values
 Column month_year split
 One-hot encoding

8) Suppose the shape of a given data is (100, 15). Our task is to perform regression on the last 5 columns using others 10. We assigned 20% of this data for validation while training process, 20% for evaluating the model after it is trained. What are the correct dimensions of x_train, y_train, x_valid, y_valid, x_test, y_test respectively? **1 point**

(60, 5), (60, 10), (20, 5), (20, 10), (20, 5), (20, 10)
 (20, 5), (20, 10), (20, 10), (20, 5), (80, 5), (80, 10)
 (60, 10), (60, 5), (20, 10), (20, 5), (20, 10), (20, 5)
 (20, 10), (20, 5), (20, 10), (20, 5), (80, 10), (80, 5)
 (80, 5), (80, 10), (0, 0), (0, 0), (20, 5), (20, 10)
 (60, 10), (60, 5), (0, 0), (0, 0), (40, 10), (40, 5)

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 (60, 10), (60, 5), (20, 10), (20, 5), (20, 10), (20, 5)

9) We are building a two layer neural network with 7 units in the input layer, 4 units in the hidden layer and 1 unit in the output layer. The hidden layer uses ReLU activation and the output layer uses sigmoid activation. Find below the weights and biases of the network. w_{ij} is the weight of the j^{th} unit in the i^{th} layer. The bias terms b and output terms o follow the same pattern. **1 point**

$w_{11} = [-0.05, 0.1, 0.1, 0.2, 0.35, 0.6, -0.9, -0.1]$ $b_{11} = -0.8$
 $w_{12} = [-0.5, 0.1, 0.1, 0.02, 0.3, 0.36, 0.9, 0.1]$ $b_{12} = -0.1$
 $w_{13} = [-0.05, 0.1, 0.1, 0.2, 0.35, 0.6, 0.9, -0.1]$ $b_{13} = -0.1$
 $w_{14} = [-0.5, 0.1, 0.1, 0.02, 0.3, 0.36, -0.9, 0.1]$ $b_{14} = -0.1$

Given input $x = [1, 0.2, 0.4, 0.9, 1, 0, 0.6, 0.3]$, what are the outputs (rounded) of the hidden layer?

[-0.83, 0.348, 0.95, -0.732]
 [-0.86, 0.346, 0.90, -0.734]
 [0, 0.348, 0.95, 0]
 None of the above

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 [0, 0.348, 0.95, 0]

10) In continuation to the previous question, given the weights and bias for the output layer, what is the output of the neural network? **1 point**

$w_{21} = [-0.05, 0.1, 0.1, 0.2];$ $b_{21} = 0$

0.506
 0.532
 0.505
 None of the above

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 0.532

11) Which of the following will help in reducing variance in a neural network? **1 point**

Increasing the dropout rate
 Increasing lambda (parameter that controls L2 regularization)
 Increasing the number of units in a hidden layer
 Increasing the size of training data by using data augmentation

No, the answer is incorrect.
 Score: 0
 Accepted Answers:
 Increasing the dropout rate
 Increasing lambda (parameter that controls L2 regularization)
 Increasing the size of training data by using data augmentation