Unit 5 - Week 4

Assignment 4

The due date for submitting this assignment has passed.
As per our records you have not submitted this assignment.

Due on 2020-02-26, 23:59 IST.

1) Your first task is to perform image classification on MNIST data. Please visit this notebook for answering the questions 1 to 5. What are the shapes of train and test data in MNIST dataset? *1 point*

- (50000, 32, 32), (10000, 32, 32)
- (60000, 28, 28), (10000, 28, 28)
- (50000, 28, 28), (10000, 28, 28)
- (60000, 32, 32), (10000, 32, 32)

No, the answer is incorrect.
Score: 0
Accepted Answers:
(60000, 28, 28), (10000, 28, 28)

2) Which loss function would be appropriate here? *1 point*

- mean_squared_error
- mean_absolute_error
- binary_crossentropy
- sparse_categorical_crossentropy

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

3) Which of the following is true for total number of parameters 'P' in this model: *1 point*

- P < 100
- 101 < = P <1000
- 1001 < = P <10000
- 10001< = P <100000

No, the answer is incorrect.
Score: 0
Accepted Answers:
10001< = P <100000

4) Train the model for 10 epochs, for the final validation accuracy 'valid_acc', select the correct option: *1 point*

- 0.8 < A < 0.85
- 0.85 < A < 0.9
- 0.9 < A < 0.95
- 0.95 < A < 1.0

No, the answer is incorrect.
Score: 0
Accepted Answers:
0.95 < A < 1.0

5) Modify your model such that it has the following layers: *1 point*

- Flatten(Input)
- Dense(10, activation='softmax')

After training for 10 epochs, for the final validation loss 'valid_loss', select the correct option:

- 0.0 < valid_loss < 0.01
- 0.01 < valid_loss < 0.05
- 0.05 < valid_loss < 0.1
- 0.1 < valid_loss < 0.2
- 0.2 < valid_loss < 0.5

No, the answer is incorrect.
Score: 0
Accepted Answers:
0.2 < valid_loss < 0.5

6) Your second task is to perform regression on Boston Housing data. Please visit this notebook for answering the questions 6 to 10. Which of the loss functions would you use to measure the performance of the model? *1 point*

- [ ] Squared error
- [ ] Binary cross-entropy
- [ ] Absolute error
- [ ] Absolute percentage error

No, the answer is incorrect.
Score: 0
Accepted Answers:
Squared error
Absolute error
Absolute percentage error

7) Define all the feature columns to be numerical columns, build and compile the model based on the instructions in the notebook and train the model for 600 epochs. What is the range of the validation loss at the end of the training process? *1 point*

- 55-60
- 45-50
- 60-65
- 50-55

No, the answer is incorrect.
Score: 0
Accepted Answers:
50-55

8) Change the model architecture by a hidden layer of 3 units and use sigmoid activation. Train the model for 600 epochs. What is the range of the training loss at the end of the training process? Try to think of why this is happening. *1 point*

- Greater than 300
- 120-250
- 100-120
- 250-300

No, the answer is incorrect.
Score: 0
Accepted Answers:
120-250

9) We will now experiment by changing the loss function. First, remove the hidden layer that we added for Q8. Then, try the solutions of the Q6 as the loss function for the optimizer. What is the mean squared error on the test dataset for the best model? (approx) *1 point*

- 58
- 39
- 4
- 53

No, the answer is incorrect.
Score: 0
Accepted Answers:
39

10) We will now try to bucketize one of the feature columns and see its effect on the model's performance. Bucketize the 'RAD' column with the *boundaries* parameter as [2, 5] and retrain the model for 200 epochs. The mean squared error on the test dataset: (Use the best model from the previous question) *0 points*

- Increases.
- Decreases.

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