

622 Exam 1

⚠ This is a preview of the published version of the quiz

Started: Jan 20 at 5:36pm

Quiz Instructions

This is your first midterm exam. You will have two hours to complete this exam once you have started. You are welcome to use your notes, textbook and lecture videos. You must work individually. You will want to have a calculator handy.

Question 1

5 pts

What does it mean to "learn" in the context of machine learning? Explain in 2 sentences or less.

Edit View Insert Format Tools Table

12pt ▾ Paragraph ▾ | **B** *I* U A ▾  ▾ τ^2 ▾ | ⋮

p



0 words



Question 2

20 pts

| # | x_1 | x_2 | x_3 | y |
|---|-------|-------|-------|-----|
| 1 | 0 | 0 | 0 | 1 |
| 2 | 0 | 1 | 0 | 0 |
| 3 | 1 | 0 | 0 | 0 |
| 4 | 1 | 0 | 1 | 1 |
| 5 | 1 | 1 | 0 | 1 |

Given the following training data, build the best depth-1 decision tree using the greedy information gain algorithm. Show your work.

Edit View Insert Format Tools Table

12pt Paragraph | **B** *I* U A  \top^2 | :

p   | 0 words |   

Question 3

10 pts

How could you adjust the K-Means algorithm such that it utilizes labels? Be clear, but concise.

Edit View Insert Format Tools Table

12pt ▾ Paragraph ▾ | **B** *I* U A ▾  ▾ τ^2 ▾ | ⋮

p



0 words



Question 4

5 pts

K-Means is guaranteed to converge. That is, it will always return a set of K cluster centers. It will never loop infinitely. How do we know this? Explain in 3 sentences or fewer.

Edit View Insert Format Tools Table

12pt ▾ Paragraph ▾ | **B** *I* U A ▾  ▾ τ^2 ▾ | ⋮

p



0 words



Question 5

10 pts

Is there a setting in which K-Means and K-NN will give the same result? Explain.

Edit View Insert Format Tools Table

12pt ▾

Paragraph ▾

B

I

U

A ▾

 ▾

T^2 ▾

⋮

p



0 words



Question 6

5 pts

True/False: There is no way to train a perceptron on data that is not linearly separable. Explain in one sentence.

Edit View Insert Format Tools Table

12pt ▾ Paragraph ▾ | **B** *I* U A ▾ ▾ T² ▾ | ⋮

p



0 words



Question 7

20 pts

I want to train a linear classifier with the logistic loss and squared norm regularizer:

$$L(w, b) = \frac{1}{\log(2)} \log(1 + e^{-y(w \cdot x + b)}) + \frac{\lambda}{2} \|w\|^2$$

What are the update equations for w and b using gradient descent on this loss?

Assume the logs are natural logarithms ($\log_e = \ln$). Remember that the derivative of the natural log is as follows

$$\frac{d}{dx} \ln(x) = \frac{1}{x}$$

Edit View Insert Format Tools Table

12pt ▾ Paragraph ▾ | **B** *I* U A ▾  ▾ T^2 ▾ | ⋮

p



0 words



Question 8

5 pts

True/False: Running K-Means with K=1 is an example of overfitting. Explain in 3 sentences or less.

Edit View Insert Format Tools Table

12pt ▾ Paragraph ▾ | **B** *I* U A ▾  ▾ T^2 ▾ | ⋮

p




0 words



Question 9

20 pts

Given the following data, run the perceptron algorithm and give the resulting weights and bias. Assume red circles are the negative class and blue crosses are the positive class. There are three of each. The data is also given in a table below. Use the ordering from the table to train your perceptron.



| # | x_1 | x_2 | y |
|---|-------|-------|-----|
| 1 | -3 | 0 | -1 |
| 2 | -2 | 0 | -1 |
| 3 | -1 | -1 | -1 |
| 4 | 1 | 2 | 1 |
| 5 | 2 | 2 | 1 |
| 6 | 2 | 0 | 1 |

Edit View Insert Format Tools Table

12pt ▾ Paragraph ▾ | **B** *I* U A ▾  ▾ T² ▾ | ⋮

p



0 words



Not saved

Submit Quiz