

Name:

NetID:

Note: This quiz will be a set of practice problems for the upcoming exam. Complete this at your leisure outside of class but please bring in an **INDIVIDUAL** copy of the quiz **ON PAPER** to turn in at the start of the exam. While doing the work feel free to work with your classmates, but I'd like you all to turn in a copy individually. You don't necessarily need to print and fill out this exact document, just writing the answers on your own piece of paper will be fine.

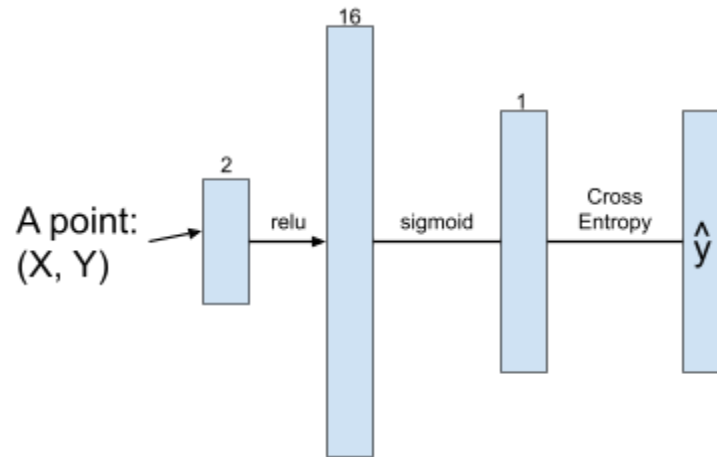
1.) Given the cost function $J(\theta_1, \theta_2, \theta_3) = \theta_1^5 - \theta_2^3 - \theta_3^2 - 5$. Perform 2 steps of gradient descent starting at $\theta_1 = 2$ and $\theta_2 = 1$ and $\theta_3 = 3$.

2.) Explain why are sigmoids particularly prone to vanishing gradients, especially in the case of RNNs

3.) What is the gradient we're descending when we use gradient descent? What are we trying to optimize and what do we take the partial derivatives with respect to to do so?

4.) What is grid search and what do we use it for?

5.) Given the network diagram, answer the following questions.



For reference:

The derivative of the sigmoid is: $\sigma(x)(1 - \sigma(x))$

The derivative of Cross-Entropy Loss is: $-(Y / \hat{y}) - ((1 - Y) / (1 - \hat{y}))$

a.) Draw the computation graph for the network. (5 pts.)


b.) Give the derivative chain for calculating the gradient of the bias in the hidden layer. (2 pts.)

$$\frac{\partial \mathcal{L}}{\partial b}$$

c.) Give the equation you'd use to actually calculate the gradients. (5 pts.)

6.) What's an example of sequential data and how do RNNs capture the context?

7.) Why do we call them **recurrent** neural networks? What makes them “recurrent”?



Bonus.) Next semester my current plan is to drop exam 03, I think it's too tight to fit everything in. What are your thoughts on that?

