

DSC 190 - Homework 08

Due: Wednesday, May 31

Write your solutions to the following problems by either typing them up or handwriting them on another piece of paper. Unless otherwise noted by the problem's instructions, show your work or provide some justification for your answer. Homeworks are due via Gradescope at 11:59 PM.

Problem 1.

Consider the neural network H shown below. You may assume for simplicity that there are no bias weights. Assume that the hidden layers use the ReLU as their activation function, and that the output layer uses the linear activation.



Suppose that $\vec{x} = (2, 1)^T$ and that the weights of this network are:

$$W^{(1)} = \begin{pmatrix} 2 & -2 \\ 3 & 1 \end{pmatrix} \quad W^{(2)} = \begin{pmatrix} 2 & 1 & 2 \\ -3 & -1 & 0 \end{pmatrix} \quad W^{(3)} = \begin{pmatrix} 4 \\ 1 \\ 2 \end{pmatrix}$$


Carry out backpropagation in order to fill in the tables of partial derivatives below. Each entry should be a number. Show your work.

Hint: some of the entries of the tables have been filled in for you – use these to check your work.

$\frac{\partial H}{\partial W_{11}^{(1)}}$ $\frac{\partial H}{\partial W_{12}^{(1)}}$ $\frac{\partial H}{\partial W_{21}^{(1)}}$ $\frac{\partial H}{\partial W_{22}^{(1)}}$	$\frac{\partial H}{\partial W_{11}^{(2)}}$ $\frac{\partial H}{\partial W_{12}^{(2)}}$ $\frac{\partial H}{\partial W_{13}^{(2)}}$ $\frac{\partial H}{\partial W_{21}^{(2)}}$ $\frac{\partial H}{\partial W_{22}^{(2)}}$ $\frac{\partial H}{\partial W_{23}^{(2)}}$	7	$\frac{\partial H}{\partial W_{11}^{(3)}}$ $\frac{\partial H}{\partial W_{21}^{(3)}}$ $\frac{\partial H}{\partial W_{31}^{(3)}}$
13			14

Problem 2.

This problem will use the data shown below:



You find this data at:

https://f000.backblazeb2.com/file/jeldridge-data/008-noisy_sin/data.csv

- a) First, train a Gaussian RBF network H on this data using three Gaussians, centered at 2, 5, and 8, and each with a width parameter of $\sigma = 1$. Plot $H(x)$ for x from 0 to 10, on top of a scatter plot of the data to visualize the accuracy of its predictions.

Hint: See Discussion 06.

- b) Next, train a neural network H on the same data. The network should have two hidden layers: each containing 5 nodes. Use sigmoid activations for the hidden layers, and linear activation for the output layer. Train your network by minimizing the mean squared error.

Plot two things: 1) the mean squared error as a function of the training epoch; and 2) $H(x)$ for x from 0 to 10, on top of a scatter plot of the data.

Hint: See Discussion 07. Not covered in that discussion was how to retrieve the mean squared error after each epoch. It turns out that the `.fit()` method returns a `History` object, which itself has a `.history` attribute that is a dictionary. One of the keys of this dictionary is `loss`; its corresponding value is the empirical risk after each epoch of SGD (or whatever optimization algorithm is used).

You will need to determine how many epochs of SGD to run; you can use the plot to do this. If you choose the number of epochs large enough, the plot should be mostly flat towards the end – otherwise, you could decrease MSE substantially by running SGD longer. Conversely, if your plot of $H(x)$ does not appear to fit the data, you might need to run for more epochs (or just try running again – you might be stuck in a local minimum). You should be able to attain an MSE of less than 0.25.

You might find that your NN does not fit the data quite as well as the RBF network – that’s OK.