CS 4700: Foundations of Artificial Intelligence
Spring 2018
Homework 4
Due: Friday, April 13, 1:25PM

1. The outputs of a sigmoidal activation function range from 0 to 1. Imagine if you need a 0/1 output for a given problem you take the final output of the neural network and return 1 if its value is ≥ 0.5 and 0 otherwise.

   Consider using this approach to use a neuron with a sigmoidal activation function to represent an AND gate – where inputs are two variables x_1 and x_2, each either 0 (false) or 1 (true), and returns 0 or 1 using the threshold of 0.5 described above, so that the output equals what you would get for $x_1 \wedge x_2$. Give a set of weights for this neuron that would correctly implement this AND gate.

2. Imagine you use g(z) = z as the activation function in your neural network. Consider a network where there are d inputs $x_1, \ldots, x_d$, (each represented by level-one neurons $a_1, \ldots, a_d$ as described in class), two hidden units d+1 and d+2 receiving inputs from $a_1, \ldots, a_d$, and one output unit d+3 that receives as its inputs the activation from neurons d+1 and d+2. Show how this neural network can be represented by a single neuron using the same activation function. Your answer should give the weights of the new neuron as a function of the weights of the original network.

3. Imagine you use $g(z) = z^3$ as the activation function in your neural network. Consider the back-propagation algorithm found in Figure 18.24 of the textbook. What do the update rules for Δ[j] and Δ[i] in the lower third of the algorithm become for this activation function?

4. The function of a single perceptron does not change if you double all the weights. What happens to a single neuron using the sigmoidal activation function if you double all the weights? (This question is about the function that the neuron computes, not about the update rule.)