WEEKLY PARTICIPATION 8: THE EXPRESSIVITY OF RELU NEURAL NETWORKS

The key reason two layer neural networks with nonlinear monotonic activation functions have the universal approximation property is that a small number of such neurons can be used to express a spike function, then we can approximate the function using a superposition of those spikes (see Lecture 14).

In this participation, you will verify that you can form a spike using a two layer ReLU feedforward neural network.

Let $\sigma(x) = x_+$ denote the ReLU activation function. Write the function s given below using a one hidden layer ReLU neural network containing three neurons. That is, find appropriate parameters $\mathbf{W}_2, \mathbf{W}_1, \mathbf{b}_1, \mathbf{b}_2$ so that

$$s(x) = \sigma(\mathbf{W}_2 \cdot \sigma(\mathbf{W}_1 x + \mathbf{b}_1) + \mathbf{b}_2).$$

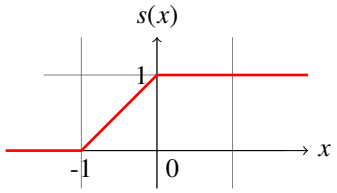

Express the spike f function below in terms of s, and use this observation to give a two layer ReLU neural network containing 5 neurons that computes f. That is, find appropriate parameters $\mathbf{W}_2, \mathbf{W}_1, \mathbf{b}_1, \mathbf{b}_2$ so that

$$f(x) = \sigma(\mathbf{W}_2 \cdot \sigma(\mathbf{W}_1 x + \mathbf{b}_1) + \mathbf{b}_2).$$

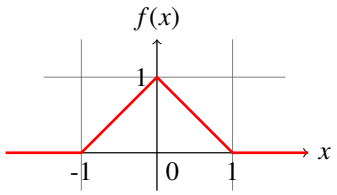

Note that this two layer ReLU neural network computed a piecewise linear function. In fact, *every* ReLU neural network, regardless of depth or width, computes a piecewise linear function. So another way to interpret the universal approximation property of ReLU networks is that one can approximate any nice function arbitarily well by using a piecewise linear function.