Homework #6 (due 6/3/2020)

1. How many weights and biases are in a network with $m = N_0 = 4$ inputs in each feature vector v_0 and $N = 6$ neurons on each of the 3 hidden layers? How many activation functions (ReLU) are in this network, before the final output?

2. (playground.tensorflow.org)
   Example 4 with blue and orange spirals is much more difficult! With one hidden layer, can the network learn this training data? Describe the results as N increases.