# Neural Networks

Y qt"ngf"'Vwvqt kcn'1

"

Question 1: Perceptron Learning

The chart below shows a set of two-dimensional input samples from two classes:

(a) It looks like there exists a perfect classification function for this problem that is linearly separable, and therefore a single perceptron should be able to learn this classification task perfectly. Let us study the learning process, starting with a random perceptron with weights $w_0 = 0.2$, $w_1 = 1$, and $w_2 = -1$, where of course w_0 is the weight for the constant offset $i_0 = 1$. For the inputs, just estimate their coordinates from the chart.

   Now add the perceptron's initial line of division to the chart. How many samples are misclassified? Then pick an arbitrary misclassified sample and describe the computation of the weight update (you can choose $\eta = 1$ or any other value; if you like you can experiment a bit to find a value that leads to efficient learning). Illustrate the perceptron's new line of division in the same chart or a

different one, and give the number of misclassified samples. Repeat this process four more times so that you have a total of six lines (or fewer if your perceptron achieves perfect classification earlier). You can do the computations and/or graphs either by hand or by writing a computer program. If you write a program, please attach a printout, and let the program run until the perceptron achieves perfect classification (after how many steps?).

(b) If your perceptron did not reach perfect classification, determine a set of weights that would achieve perfect classification, and draw the separating line for those weights.

(c) Now let us assume that less information were available about the samples that are to be classified. Let us say that we only know the value for i_1 for each sample, which means that our perceptron has only two weights to classify the input as best as possible, i.e., it has weights w_0 and w_1, where w_0 is once again the weight for the constant offset $i_0 = 1$. Draw a diagram that visualizes this one-dimensional classification task, and determine weights for a perceptron that does the task as best as possible (minimum error, i.e., minimum proportion of misclassified samples). Where does it separate the input space, and what is its error?

Question 2: Graduating from Kindergarten

Having solved the easy task from Question 1, let us tackle a more difficult classification challenge:

(a) As you certainly noticed, a single perceptron cannot do this classification task perfectly. Determine the minimum error that a single perceptron can reach, and show the dividing line in the input space for such a perceptron.

(b) Clearly, we need a multi-layer perceptron to do this task perfectly. Develop such a system of connected perceptrons and write it down, together with the required weights for each unit. Illustrate the dividing lines for these perceptrons in a copy of the chart above.

Question 3: A Matter of Dimensions

(a) As we have seen in Question 1, the same samples can be perfectly classified if we have two pieces of information about them (i_1 and i_2), while our perceptron (and in fact any classifier) fails if only one piece (i_1) is known. Do you think that generally, the more pieces of information (i.e., input dimensions) we add, the greater is the chance that your perceptron can do its task perfectly? Explain your answer, including a geometrical interpretation (one- vs. two- vs. three-dimensional input spaces, and so on). Under what conditions does additional information help with the classification, and under what conditions does it not?

(b) Is there a minimum number of input dimensions d for which it is guaranteed for any set of n input samples that there is a perceptron that perfectly classifies that set? Explain why this is or is not the case.