

CS4700 Fall 2011: Foundations of Artificial Intelligence

Homework #5

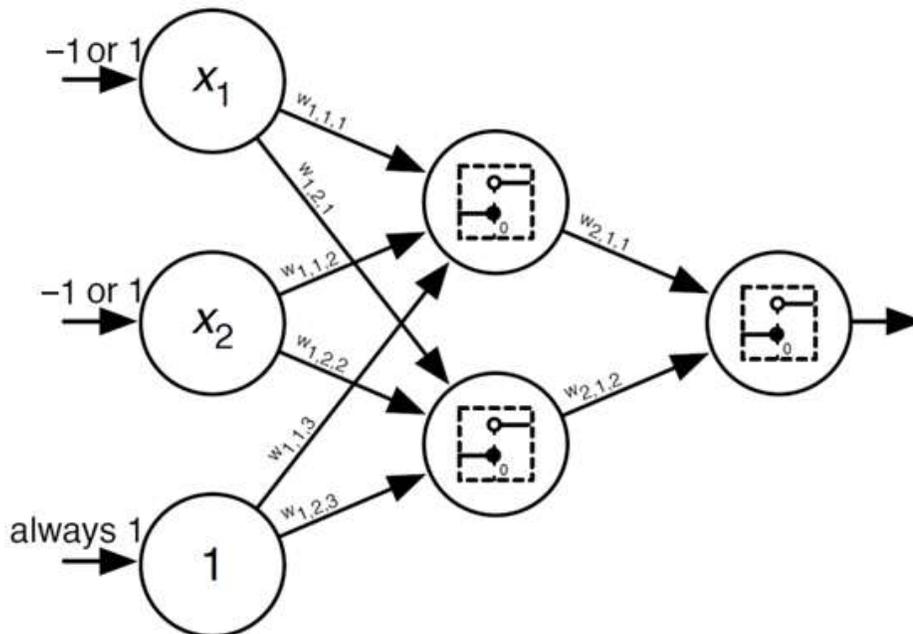
Due Date: Monday Nov 21 on CMS (PDF) and in class (hardcopy)

Thereafter 10% off each 24H period late until homework is graded on CMS.

Submission: Submit on CMS and paper copies at the beginning of class or to TA directly. Please include your **name**, **NetID**, the **CMS submission date and time**, **slack days** used and remaining. Code in 9pt courier single spacing with any blank lines removed and function names highlighted. Your assignments should reflect your individual work – ok to discuss strategies but do not compare your answers with anyone else.

Question 1: Stop XNORing (25 points)

Suppose we'd like to construct a simple ANN for the XNOR operator \oplus . For binary operands $x_1, x_2 \in \{-1, 1\}$, $x_1 \oplus x_2$ may be defined as follows: $-1 \oplus -1 = 1$, $-1 \oplus 1 = 0$; $1 \oplus -1 = 0$, $1 \oplus 1 = 1$ (intuitively, both are the same). This is a canonical "nonlearnable" function for perceptrons, but it is learnable with two level ANNs. Our neural network will have two layers (one hidden layer and the output layer). The hidden and output layers have two and one units respectively. Further, we have three inputs (two inputs for the XNOR operands, and one bias input always at 1). Our output function for the hidden units and the output unit will be a simple threshold function $f(x)$, where $f(x)$ is 1 if $x > 0$, and is zero otherwise. Give weights such that the network below outputs $x_1 \oplus x_2$. Label the figure.



Question 2: Very Perceptive (25 points)

Construct a support vector machine that computes the XNOR function. It will be convenient to use values of 1 and -1 instead of 1 and 0 for the inputs and for the outputs. So an example will look like $([-1, 1], -1)$ or $([-1, -1], 1)$. It is typical to map an input \mathbf{x} into a space consisting of five dimensions, the two original dimensions x_1 and x_2 and the three combinations x_1^2 , x_2^2 , and x_1x_2 . But for this exercise we will consider only the two dimensions x_1 and x_1x_2 .

- Draw the four input points in this space, and draw the maximal separator. What is the margin?
- If we use x_1^2 , x_2^2 , and x_1x_2 is it linearly separable? If it is, draw the separating line back in the original Euclidean space
- Consider a function that takes three binary inputs and that outputs 1 if exactly two of the inputs are 1 (0 otherwise). Decide whether a linear perceptron can learn this function. Either show why it cannot or construct a perceptron that can.

Question 3: Flowers (50 points)

In this exercise, we analyze the Iris flower data set which is available online¹. Remember to shuffle the data points before you use them. Normalize the data by subtracting its mean and dividing each dimension by its standard deviation.

We will try to classify the data using a set of hyperplanes. Use your perceptron training algorithm from last homework to train a single perceptron for each of the three flower classes. Each perceptron should determine whether a data point belongs to a class, or not. The three perceptrons can be used together to vote on a data point, using the distance from the margin to break ties. Use 5-fold cross validation to evaluate the performance of this algorithm, as well as any design choices you make (e.g. the learning rate parameter, early stopping).

Programming tips: Don't create anything more complicated than you need to. This is a simple program. Implement it simply. Help your grader help you: be clear with comments and naming.

Bonus Question 4: More Flowers (additional 60 points)

Bonus 60 points: Implement from scratch a neural network with back propagation to train a multi-layer network with four inputs, a hidden layer, and three outputs to perform the classification. Determine the best accuracy you can get and show a validation curve. Use the validation curve to do early stopping. Choose several sizes for the hidden layer, and show how performance changes for different numbers of hidden neurons. Did you perform better than the three perceptrons?

¹ http://en.wikipedia.org/wiki/Iris_flower_data_set