


CS360 Homework 5– Solution

Artificial Neural Networks


- 1) Can a decision tree represent the Boolean function $f(P, Q) \equiv P \Rightarrow Q$? What about a single perceptron with a step function (whose step is at zero) as threshold function? What about a network of perceptrons with step functions (whose steps are at zero) as threshold functions? Now answer the same three questions for the Boolean function $f(P, Q) \equiv P \Leftrightarrow Q$.

Answer:


A decision tree to represent the Boolean function $f(P, Q) \equiv P \Rightarrow Q$:



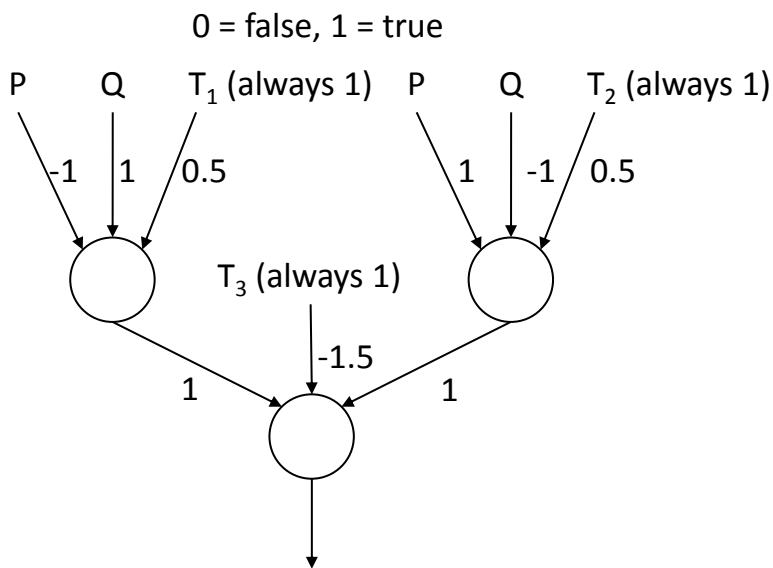
A decision tree that represents the Boolean function $f(P, Q) \equiv P \Leftrightarrow Q$:



A single perceptron (which is trivially a network of perceptrons) with a step function (whose step is at zero) that represents the Boolean function $f(P, Q) \equiv P \Rightarrow Q$:



A network of perceptrons with step functions that represents the Boolean function $f(P, Q) \equiv P \Leftrightarrow Q \equiv (P \Rightarrow Q) \wedge (Q \Rightarrow P)$:



A single perceptron cannot be used to represent the Boolean function $f(P, Q) \equiv P \Leftrightarrow Q$ because the four cases (where P and Q are true or false) is not linearly separable (similar to the XOR example in the Perceptron Learning slides).

- 2) Develop the training rule for a perceptron with a sigmoid function as threshold function and the sum of errors to the power of four (instead of the sum of squared errors) as error function.

Answer:

Terminology:

- g is the sigmoid function with $g(x) = 1/(1 + e^{-x})$ and its derivative $g'(x) = g(x)(1 - g(x))$.
- c_i is the desired value for the i th training example.
- o_i is the value computed by the perceptron as the value of the i th training example.
- w_j is the weight given to the j th feature.

- f_{ij} is the value of the j th feature of the i th training example.

We use the error function $E := 1/4 \sum_i (c_i - o_i)^4 = 1/4 \sum_i (c_i - g(\sum_j w_j f_{ij}))^4$

$$\begin{aligned} dE/dw_j &= \sum_i (c_i - o_i)^3 d(c_i - o_i)/dw_j \\ &= \sum_i (c_i - o_i)^3 d(c_i - g(\sum_j w_j f_{ij}))/dw_j \\ &= \sum_i (c_i - o_i)^3 (-g'(\sum_j w_j f_{ij})) d(\sum_j w_j f_{ij})/dw_j \\ &= \sum_i (c_i - o_i)^3 (-f_{ij} g'(\sum_j w_j f_{ij})) \end{aligned}$$

Thus, weight w_j gets updated by gradient descent as follows (batch update), where α is a small learning rate:

$$w_j := w_j - \alpha dE/dw_j = w_j + \alpha \sum_i f_{ij} (c_i - o_i)^3 g'(\sum_j w_j f_{ij})$$

If we want to update the weights one training example at a time, we use the following (for the i th example):

$$w_j := w_j + \alpha f_{ij} (c_i - o_i)^3 g'(\sum_j w_j f_{ij})$$

In this case, we need to cycle through all examples.

- 3) (Courtesy of Russell and Norvig) Suppose that a training set contains only a single example, repeated 100 times. In 80 of the 100 cases, the single output value is 1; in the other 20, it is 0. What will a neural network predict for this example, assuming that it has been trained on all training examples and reaches a global optimum? (Hint: To find the global optimum, differentiate the error function and set the resulting expression to zero.)

Answer:

We determine the weights that minimize the error function. The error function is

$$\begin{aligned} E &= 1/2 \sum_i (c_i - o_i)^2 \\ &= 1/2(80(1 - o)^2 + 20(0 - o)^2) \\ &= 50Output^2 - 80Output + 40 \end{aligned}$$

The derivative of the error with respect to the output is

$$dE/dOutput = 100Output - 80$$

Setting the derivative to zero, we find that $Output = 0.8$, the probabilistic prediction $P(Output = 1)$.

- 4) If we train a neural network for 1,000 epochs (one training example at a time), does it make a difference whether we present all training examples in turn for 1000 times or whether we first present the first training example 1000 times, then the second training example for 1000 times, and so on? Why?

Answer:

The resulting networks would be different. The error used to adjust the weights of the network during training is based on the training examples given to it during that epoch. In the case of all training examples presented in turn 1000

times, the error is based on the error over all examples because they are all present during each epoch. In the case where each training example is presented individually for 1000 times, one first reduces the error only with respect to the first training example, then only with respect to the second training example, and so on. The result is that the network will “forget” the earlier training examples. As later training examples are presented, the weights will move away from the weights necessary to classify the first training examples correctly.

- 5) You are given n numbers and have to determine their mean. Develop a gradient descent rule for this purpose.

Answer:

Let x_1, \dots, x_n denote the given numbers, \bar{x} denote their mean and \tilde{x} denote our current estimate of their mean. A possible error function is:

$$Error(\tilde{x}) = 1/2(\sum_i(\tilde{x} - x_i))^2.$$

This error function is non-negative and zero exactly at the mean (that is, when $\tilde{x} = \bar{x}$). It has exactly one global minimum, namely at the mean.

Update equation:

$$\tilde{x} := \tilde{x} - \alpha \frac{dError(\tilde{x})}{d\tilde{x}} = \tilde{x} - \alpha \frac{1}{2} 2 \sum_i (\tilde{x} - x_i) = \tilde{x} + \alpha \sum_i (x_i - \tilde{x})$$

We can use this update equation in two different ways.

Schema 1 (batch update):

$\tilde{x} :=$ random number

Repeatedly do:

$\tilde{x} := \tilde{x} + \alpha \sum_i (x_i - \tilde{x})$

Below is sample code to determine the mean of 2, 3 and 7 using Schema 1:

```
#include <stdio.h>

main()
{
    int i;
    double mean = 10.0;
    double alpha = 0.0001;

    for (i = 0; i < 100000; ++i)
    {
        mean = mean + alpha * (
            (2.0 - mean) +
            (3.0 - mean) +
            (7.0 - mean));
    }
}
```

```

        printf("%f\n", mean);
    }
}

```

Schema 2 (one training example at a time):

$\tilde{x} :=$ random number

Repeatedly do:

$\tilde{x} := \tilde{x} + \alpha(x_1 - \tilde{x})$

$\tilde{x} := \tilde{x} + \alpha(x_2 - \tilde{x})$

$\tilde{x} := \tilde{x} + \alpha(x_3 - \tilde{x})$

...

Below is a sample code to determine the mean of 2, 3 and 7, using Schema 2:

```

#include <stdio.h>

main()
{
    int i;
    double mean = 10.0;
    double alpha = 0.0001;

    for (i = 0; i < 100000; ++i)
    {
        mean = mean + alpha * (2.0 - mean);
        printf("%f\n", mean);
        mean = mean + alpha * (3.0 - mean);
        printf("%f\n", mean);
        mean = mean + alpha * (7.0 - mean);
        printf("%f\n", mean);
    }
}

```

- 6) Explain exactly why networks of perceptrons with linear activation functions are uninteresting (that is, networks of perceptrons where, for each perceptron, the output is some constant times the weighted sum of the inputs). Use equations if necessary.

Answer:

Networks of perceptrons with linear activation functions are uninteresting because a single perceptron can be used to represent a whole network of perceptrons with linear activation functions (since a weighted sum of weighted sums of feature values can be calculated simply as a weighted sum of feature values).

- 7) Is overfitting more or less likely when the training set is small or large? Is overfitting more or less likely when the number of parameters to learn (such as the number of weights in a neural network) is small or large?

Answer:

Overfitting is more likely when the training set is small and the number of parameters is large. The larger the number of parameters to learn, the more training examples are needed to estimate them well.