CS530 HOMEWORK 5, DUE 12/13

Justify your answers: if you do the problems by hand or calculator, show your intermediate formulas; if you do the problems by computer, show your code. Each problem in this homework is worth the same for grading.

Recall Waldo's 10-day experience at the intersection from homework 4.

L	r	g	r	r	r	g	g	y	g	r
T	f	f	f	f	t	f	f	f	f	t

In this homework, you will construct simple neural networks to help Waldo decide whether to stop on the 11th day.

Even without a specific decision to make in mind, we can use observations to train a Bayes net, which can then be used for a variety of inference tasks. For example, Waldo can use the same Bayes net from homework 4 to decide whether to stop given that the traffic light is yellow as well as to decide whether to stop without looking at the traffic light at all. (Perhaps the traffic light is so bright that it hurts his eyes.) By contrast, before we can train a classifier (be it a neural network or a decision tree), we must have identify some random variables as input and others as output. That is, the design of the neural network depends on the task that we want it to perform.

There are (at least) two ways to design a neural network to help Waldo decide whether to stop given the status of the traffic light. Both designs encode L as three input neurons, r, y, and g, such that exactly one of these neurons is at activation level 1, and the other two at activation level 0. (Note that we have switched from activation levels ranging between -1 and $+1$, as done in class, to activation levels ranging between 0 and 1, which seems a bit more natural to think about in the setting of a classifier.) The two designs differ in how they encode the presence or absence of traffic.

The first design is to use a single output neuron T to encode traffic: the activation level of T should be 0 if we are very sure there is no traffic, and 1 if we are very sure there is traffic. With no other (hidden) neuron, our network then looks like the following.

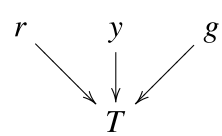

The activation level of the output neuron T is

$$L_T = g(W_{rT}L_r + W_{yT}L_y + W_{gT}L_g - W_{0T}),$$

where L_r, L_y, and L_g are the activation levels of the input neurons, W_{rT}, W_{yT}, and W_{gT} are the weights, and W_{0T} is the threshold of T. Also, the sigmoid function $g(x)$ is defined by

$$g(x) = \frac{1}{1 + e^{-x}}.$$

It is convenient to know that the derivative $g'(x)$ is equal to $g(x)(1 - g(x))$.

(1) Find a set of weights (including the threshold W_{0T}) that is optimal in the sense that it minimizes mean squared error. There is no need to use gradient descent. These weights include $+\infty$ or $-\infty$, because $g(x) = 0$ only when $x = -\infty$, and $g(x) = 1$ only when $x = +\infty$. Infinite weights are problematic for gradient descent—why? Therefore, we usually use

activation levels like 0.1 and 0.9 rather than 0 and 1 to encode training examples when training a neural network with gradient descent.

(2) The optimal weights are not unique. To show that you know why, give an example of an additional training example, consisting of input and output activation levels, that makes the optimal weights unique.

(3) Given that the traffic light is yellow on the 11th day, what is the activation level of the output neuron T?

Unfortunately, merely knowing the activation level of T does not lead to a rational decision on Waldo's part, because the activation level is not a probability, at least not when we minimize mean squared error.

Let us turn to the second design. It uses two neurons f and t to encode traffic: their activation levels should be 1 and 0 if we are very sure there is no traffic, and 0 and 1 if we are very sure there is traffic. With no other (hidden) neuron, our network then looks like the following.

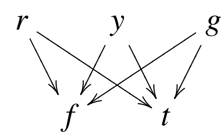

Analogous to the first design, the output activation levels are

$$L_f = g(W_{rf}L_r + W_{yf}L_y + W_{gf}L_g - W_{0f}),$$
$$L_t = g(W_{rt}L_r + W_{yt}L_y + W_{gt}L_g - W_{0t}).$$

(4) As in problem 1, find a set of weights (including the thresholds) that minimizes mean squared error. Given that the traffic light is yellow on the 11th day, what are the activation levels of f and t?

With this network, we can simply interpret the output neuron with the higher activation level as the one chosen by the network.

(5) Assuming that the network chooses correctly, should Waldo stop on the 11th day? Does this decision take into account the fact that Waldo has only seen the yellow light once before (explain why)?