

Decision trees are used a little, neural Networks a used quite a bit

Brain works pretty well, we should be able to make an algorithm that emulates the brain that also works well!



These points can be *very* high dimensional. For, example the points above (GPA, SAT, HS) could be (GPA, SAT, HS, PlaysTennis, PlaysPiano, PlaysFootball, PlaysClarinet, LikedZombieland...).

The admissions department can learn to predict your GPA! Each admitted student is a vector of attributes, and the output is their GPA when they graduate.



A neuron. Has inputs (dendrites) and outputs (axon). Some inputs are positive (make it more likely the neuron will fire), and some are negative (make it less likely the neuron will fire)

Lots of neurons in your brain.



Wij is the weight from neuron i to neuron j

Activation function g is a function of the weighted sum of the input neurons

A0 doesn't come from another neuron, rather it is the bias. Bias is a constant factor and determines how much input it takes for the neuron to fire. "How hard is it to make the neruon fire"



The in3 is the weighted sum of all the inputs to the neuron





We can *always* center this at 0, because of the bias term. Note if we didn't have bias (b) we'd just have it centered at whatever the bias was.

Threshold activation functions are very common.

The sigmoid is continuous and differentiable. This is really nice, as we will see later.

These two functions are very similar...almost always practically equal

We are assuming all neurons use the same activation function



OR is like adding- you need at least one.



Here you need both.



Try to make an XOR. Actually, you can't, and we'll tell you why later.





Feed forward neural networks do not have any cycles.

Cycles make things much more difficult. With cycles whether 4 fires or not would potentially not converge. Luckily we almost always do feed forward neural networks









Only one layer of weights. Slightly confusing name, as there are two layers of neurons (the input layer and the output layer.



Only need to analyze with a single output unit.

What functions can be represented by a perceptron?



If this were a decision tree, it'd have to be really big. Isn't that nice?





A vector of weights dot product with a vector of inputs.



Linearly separable means all negative and positive examples can be separated by a line (in higher dimensions, separable by a hyperplane)

And *this* is why you can't do XOR, as it is not linearly separable. Therefore a single threshold perceptron cannot do XOR

However, if you have multiple perceptrons, you can do a lot more.

Learning in Perceptrons

Learning algorithm that will fit a threshold perceptron to any linearly separable function

Idea: adjust the weights to minimize some notion of error

We will use the notion of squared error

For a single example with input x and true output y:

 $E = (1/2)Err^2 = (1/2)(y - h_W(x))^2$

Where $h_w(x)$ is the output of the perceptron

We're doing local search on the weights, trying to minimize the squared error.



We can derive the equation to minimize the squared error. If the error (not squared error) is positive, we'll make the weight increase. If the error is negative (you're overshooting), you decrease the weights.

This is why using a g which has continuous partial derivatives in respect to the weights is useful.

The black box tells you how to update the weights when learning. g' is the derivative of g.

Multi-Layer Feed-Forward NNs

Can represent more complex functions

With a single sufficiently large hidden layer can represent any continuous function with arbitrary accuracy

BIG QUESTION: What structure to use for a given problem?

Hidden layer is the layer between the input layer and the output layer.

Just guess a structure.

Often people guess a structure, learn the weights, then go back and revise if the error is bad.

Structure learning is another interesting problem.













Okay, so since we know what the output of 5 and 6 should have been we can calculate that easily. But what about unit 3?





This is the back-propagation algorithm. The idea is that you are back-propagating the errors from the units in the end layers. Then once you can calculate the error, you can learn the weights.

