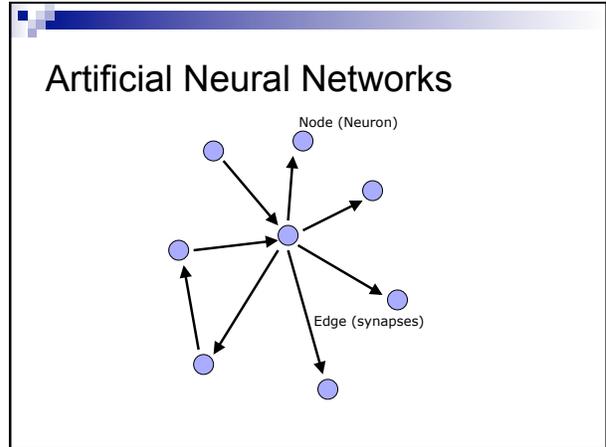


# Neural Networks 2

David Kauchak  
CS30  
Spring 2015

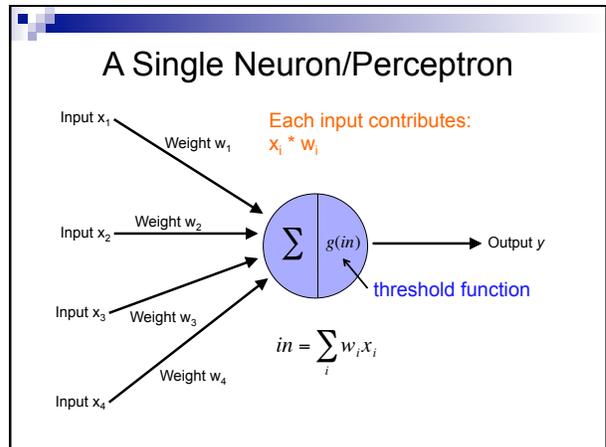


Node A (neuron) → Weight  $w$  → Node B (neuron)

$W$  is the strength of signal sent between A and B.

If A fires and  $w$  is **positive**, then A **stimulates** B.

If A fires and  $w$  is **negative**, then A **inhibits** B.



### Training neural networks

$x_1$	$x_2$	$x_3$	$x_1$ and $x_2$
0	0	0	1
0	1	0	0
1	0	0	1
1	1	0	0
0	0	1	1
0	1	1	1
1	0	1	1
1	1	1	0

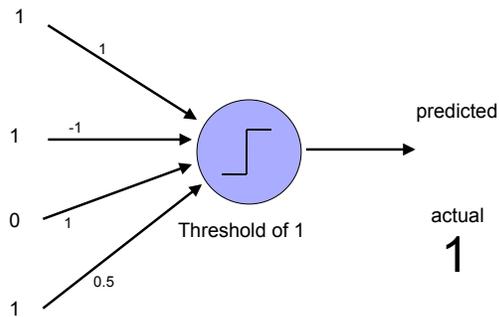
1. start with some initial weights and thresholds
2. show examples repeatedly to NN
3. update weights/thresholds by comparing NN output to actual output

### Perceptron learning algorithm

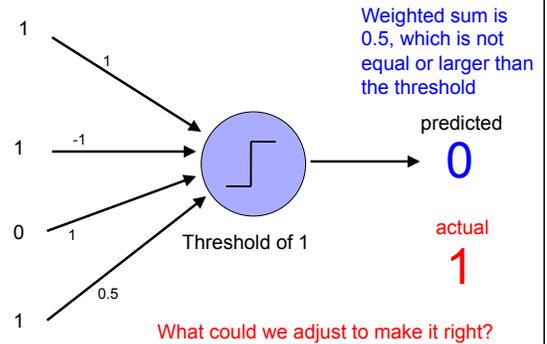
repeat until you get all examples right:

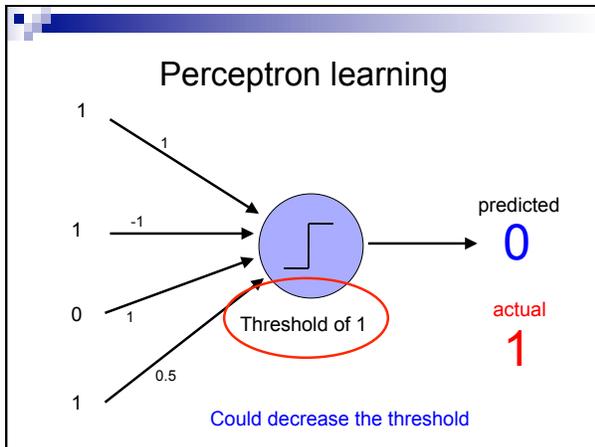
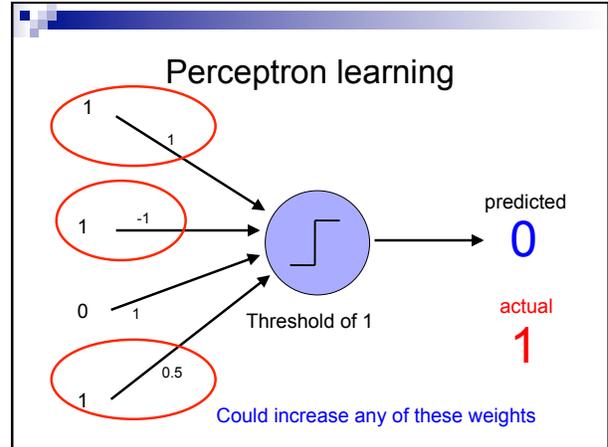
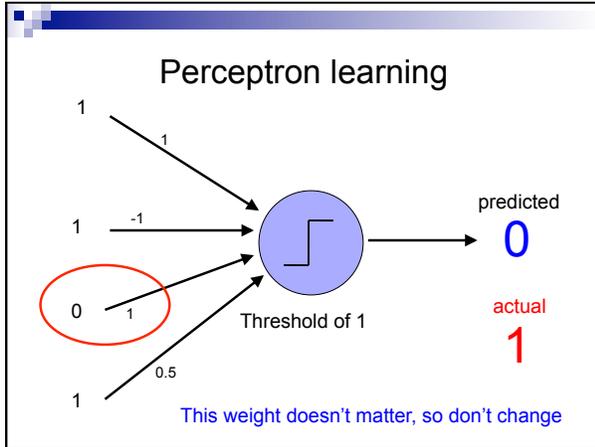
- for each "training" example:
  - calculate current prediction on example
  - if *wrong*:
    - update weights and threshold towards getting this example correct

### Perceptron learning



### Perceptron learning





### Perceptron update rule

- if *wrong*:
  - update weights and threshold towards getting this example correct

↓

- if *wrong*:
  - $w_i = w_i + \Delta w_i$
  - $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

What does this do in this case?

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

causes us to increase the weights!

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

What if predicted = 1 and actual = 0?

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

We're over the threshold, so want to decrease weights:  
 actual - predicted = -1

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

What does this do?

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

Only adjust those weights that actually contributed!

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

What does this do?

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

"learning rate": value between 0 and 1 (e.g 0.1)  
 adjusts how abrupt the changes are to the model

### Perceptron learning

$w_i = w_i + \Delta w_i$   
 $\Delta w_i = \lambda * (\text{actual} - \text{predicted}) * x_i$

**What about the threshold?**

$1 \text{ if } \sum_{i=1}^3 w_i x_i \geq t$

$1 \text{ if } w_4 + \sum_{i=1}^3 w_i x_i \geq 0$

$1 \text{ if } \sum_{i=1}^3 w_i x_i \geq t$

$1 \text{ if } w_4 + \sum_{i=1}^3 w_i x_i \geq 0$

equivalent when  $w_4 = -t$

### Perceptron learning algorithm

initialize weights of the model randomly

repeat until you get all examples right:

- for each "training" example (*in a random order*):
  - calculate current prediction on the example
  - if *wrong*:
    - $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

initialize with random weights

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:

$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Right or wrong?

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:

$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Wrong

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

1  $W_1 = 0.2$   
0  $W_2 = 0.5$   
1  $W_3 = 0.1$

sum = 0.3: output 1

Output y

new weights?

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

1  $W_1 = 0.1$   
0  $W_2 = 0.5$   
1  $W_3 = 0.0$

sum = 0.3: output 1

Output y

decrease (0-1=-1) all non-zero  $x_i$  by 0.1

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

1  $W_1 = 0.1$   
1  $W_2 = 0.5$   
1  $W_3 = 0.0$

Output y

Right or wrong?

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

1  $W_1 = 0.1$   
0  $W_2 = 0.5$   
1  $W_3 = 0.0$

sum = 0.6: output 1

Output y

Right. No update!

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

Right or wrong?

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

Wrong

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

new weights?

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

decrease (0-1=-1) all non-zero  $x_i$  by 0.1

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Right or wrong?

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Right. No update!

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Right or wrong?

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Wrong

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

Diagram description: A blue circle representing a neuron. Three arrows point into it from the left. The top arrow is labeled  $W_1 = 0.1$  and has a value of 0 next to it. The middle arrow is labeled  $W_2 = 0.3$  and has a value of 1 next to it. The bottom arrow is labeled  $W_3 = -0.2$  and has a value of 1 next to it. An arrow points out of the circle to the right, labeled "Output y". Above the output arrow, it says "sum = 0.3: output 1". Below the diagram, it says "decrease (0-1=-1) all non-zero  $x_i$  by 0.1".

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

Diagram description: A blue circle representing a neuron. Three arrows point into it from the left. The top arrow is labeled  $W_1 = 0.1$  and has a value of 1 next to it. The middle arrow is labeled  $W_2 = 0.3$  and has a value of 1 next to it. The bottom arrow is labeled  $W_3 = -0.2$  and has a value of 1 next to it. An arrow points out of the circle to the right, labeled "Output y". Above the output arrow, it says "sum = 0.2: output 1". Below the diagram, it says "Right. No update!".

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

Diagram description: A blue circle representing a neuron. Three arrows point into it from the left. The top arrow is labeled  $W_1 = 0.1$  and has a value of 0 next to it. The middle arrow is labeled  $W_2 = 0.3$  and has a value of 0 next to it. The bottom arrow is labeled  $W_3 = -0.2$  and has a value of 1 next to it. An arrow points out of the circle to the right, labeled "Output y". Above the output arrow, it says "sum = -0.2: output 0". Below the diagram, it says "Right. No update!".

$x_1$	$x_2$	$x_1$ and $x_2$	$\lambda = 0.1$
0	0	0	
0	1	0	if wrong:
1	0	0	$w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$
1	1	1	

Diagram description: A blue circle representing a neuron. Three arrows point into it from the left. The top arrow is labeled  $W_1 = 0.1$  and has a value of 1 next to it. The middle arrow is labeled  $W_2 = 0.3$  and has a value of 0 next to it. The bottom arrow is labeled  $W_3 = -0.2$  and has a value of 1 next to it. An arrow points out of the circle to the right, labeled "Output y". Above the output arrow, it says "sum = -0.1: output 0". Below the diagram, it says "Right. No update!".

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Are they all right?

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Wrong

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

decrease (0-1=-1) all non-zero  $x_i$  by 0.1

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

Are they all right?

$x_1$	$x_2$	$x_1$ and $x_2$
0	0	0
0	1	0
1	0	0
1	1	1

$\lambda = 0.1$

if wrong:  
 $w_i = w_i + \lambda * (\text{actual} - \text{predicted}) * x_i$

We've learned AND!

## Learning in multilayer networks

Similar idea as perceptrons

Examples are presented to the network

If the network computes an output that matches the desired, nothing is done

If there is an error, then the weights are adjusted to balance the error

## Learning in multilayer networks

Key idea for perceptron learning: if the perceptron's output is different than the expected output, update the weights

Challenge: for multilayer networks, we don't know what the expected output/error is for the internal nodes

perceptron

multi-layer network

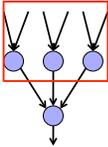
## Backpropagation

Say we get it wrong, and we now want to update the weights

We can update this layer just as if it were a perceptron

### Backpropagation

Say we get it wrong, and we now want to update the weights



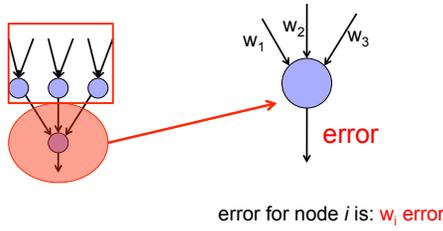
“back-propagate” the error:

Assume all of these nodes were responsible for some of the error

How can we figure out how much they were responsible for?

### Backpropagation

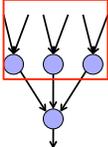
Say we get it wrong, and we now want to update the weights



error for node  $i$  is:  $w_i \text{ error}$

### Backpropagation

Say we get it wrong, and we now want to update the weights



Update these weights and continue the process back through the network

### Backpropagation

- calculate the error at the output layer
- backpropagate the error up the network
- Update the weights based on these errors

Can be shown that this is the appropriate thing to do based on our assumptions

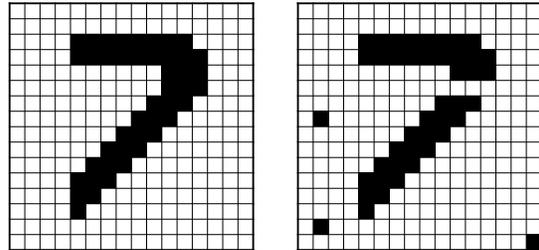
That said, many neuroscientists don't think the brain does backpropagation of errors

## Neural network regression

Given enough hidden nodes, you can learn *any* function with a neural network

### Challenges:

- overfitting – learning only the training data and not learning to generalize
- picking a network structure
- can require a lot of tweaking of parameters, preprocessing, etc.



Popular for digit recognition and many computer vision tasks  
<http://yann.lecun.com/exdb/mnist/>

## Cog sci people like NNs

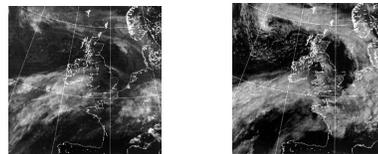
Expression/emotion recognition

- Gary Cottrell et al



Language learning

## Interpreting Satellite Imagery for Automated Weather Forecasting



### What NNs learned from youtube



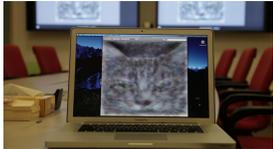
<http://www.nytimes.com/2012/06/26/technology/in-a-big-network-of-computers-evidence-of-machine-learning.html>

### What NNs learned from youtube

trained on 10M snapshots from youtube videos

NN with 1 billion connections

16,000 processors



### Summary

Perceptrons, one layer networks, are insufficiently expressive

Multi-layer networks are sufficiently expressive and can be trained by error back-propagation

Many applications including speech, driving, hand written character recognition, fraud detection, driving, etc.

### Our python NN module

Data:

$x_1$	$x_2$	$x_3$	$x_1$ and $x_2$
0	0	0	1
0	1	0	0
1	0	0	1
1	1	0	0
0	0	1	1
0	1	1	1
1	0	1	1
1	1	1	0

➔

```
table = \
[ ([0.0, 0.0, 0.0], [1.0]),
  ([0.0, 1.0, 0.0], [0.0]),
  ([1.0, 0.0, 0.0], [1.0]),
  ([1.0, 1.0, 0.0], [0.0]),
  ([0.0, 0.0, 1.0], [1.0]),
  ([0.0, 1.0, 1.0], [1.0]),
  ([1.0, 0.0, 1.0], [1.0]),
  ([1.0, 1.0, 1.0], [0.0]) ]
```

### Data format

list of examples

```
table = \
[ ([0.0, 0.0, 0.0], [1.0]),
  ([0.0, 1.0, 0.0], [0.0]),
  ([1.0, 0.0, 0.0], [1.0]),
  ([1.0, 1.0, 0.0], [0.0]),
  ([0.0, 0.0, 1.0], [1.0]),
  ([0.0, 1.0, 1.0], [1.0]),
  ([1.0, 0.0, 1.0], [1.0]),
  ([1.0, 1.0, 1.0], [0.0]) ]
```

→ ( [0.0, 0.0, 0.0], [1.0] )

input list      output list

example = tuple

### Training on the data

Construct a new network:

```
>>> nn = neuralNet(3, 2, 1)
```

constructor: constructs a new NN object

input nodes

hidden nodes

output nodes

### Training on the data

Construct a new network:

```
>>> nn = neuralNet(3, 2, 1)
```

3 input nodes      2 hidden nodes      1 output node

### Training on the data

```
>>> nn.train(table)
error 0.195200
error 0.062292
error 0.031077
error 0.019437
error 0.013728
error 0.010437
error 0.008332
error 0.006885
error 0.005837
error 0.005047
```

by default trains 1000 iteration and prints out error values every 100 iterations

### After training, can look at the weights

```

>>> nn.train(table)
>>> nn.getIHWeights()
[[-3.3435628797862624, -0.272324373735495],
 [-4.846203738642956, -4.601230952566068],
 [3.4233831101145973, 0.573534695637572],
 [2.9388429644152128, 1.8509761272713543]]
    
```

### After training, can look at the weights

```

>>> nn.train(table)
>>> nn.getIHWeights()
[[-3.3435628797862624, -0.272324373735495],
 [-4.846203738642956, -4.601230952566068],
 [3.4233831101145973, 0.573534695637572],
 [2.9388429644152128, 1.8509761272713543]]
    
```

### After training, can look at the weights

```

>>> nn.getHOWeights()
[[8.116192424400454],
 [5.358094903107918],
 [-4.373829543609533]]
    
```

### Many parameters to play with

`nn.train(trainingData)` carries out a training cycle. As specified earlier, the training data is a list of input-output pairs. There are four optional arguments to the `train` function:

- `learningRate` defaults to 0.5.
- `momentumFactor` defaults to 0.1. The idea of momentum is discussed in the next section. Set it to 0 to suppress the affect of the momentum in the calculation.
- `iterations` defaults to 1000. It specifies the number of passes over the training data.
- `printInterval` defaults to 100. The value of the error is displayed after `printInterval` passes over the data; we hope to see the value decreasing. Set the value to 0 if you do not want to see the error values.

You may specify some, or all, of the optional arguments by name in the following format.

```

nn.train(trainingData,
         learningRate=0.8,
         momentumFactor=0.0,
         iterations=100,
         printInterval=5)
    
```

## Calling with optional parameters

```
>>> nn.train(table, iterations = 5, printInterval = 1)
error 0.005033
error 0.005026
error 0.005019
error 0.005012
error 0.005005
```

## Train vs. test

TrainData

input	output
0.0	0.00
0.2	0.04
0.4	0.16
0.6	0.36
0.8	0.64
1.0	1.00

TestData

input	output
0.3	0.09
0.5	0.25
0.7	0.49
0.8	0.64
0.9	0.81

```
>>> nn.train(trainData)
>>> nn.test(testData)
```

<http://www.sciencebytes.org/2011/05/03/blueprint-for-the-brain/>