

Artificial Neural Networks



Artificial Neural Networks

- **Artificial neural networks** (ANNs) provide a general, practical method for learning real-valued, discrete-valued, and vector-valued functions from examples.
- Algorithms such as **BACKPROPAGATION** gradient descent to tune network parameters to best fit a training set of input-output pairs.
- ANN learning is robust to errors in the training data and has been successfully applied to problems such as interpreting visual scenes, speech recognition, and learning robot control strategies.

Biological Motivation

- The study of artificial neural networks (ANNs) has been inspired in part by the observation that biological learning systems are built of very complex webs of interconnected neurons.
- Artificial neural networks are built out of a densely interconnected set of simple units, where each unit takes a number of real-valued inputs (possibly the outputs of other units) and produces a single real-valued output (which may become the input to many other units).
- The human brain is estimated to contain a densely interconnected network of approximately 10^{11} neurons, each connected, on average, to 10^4 others.
 - Neuron activity is typically inhibited through connections to other neurons.

ALVINN – Neural Network Learning To Steer An Autonomous Vehicle




Properties of Artificial Neural Networks

- A large number of very simple, neuron-like processing elements called **units**,
- A large number of weighted, directed connections between pairs of units
 - Weights may be positive or negative real values
- Local processing in that each unit computes a function based on the outputs of a limited number of other units in the network
- Each unit computes a simple function of its input values, which are the weighted outputs from other units.
 - If there are n inputs to a unit, then the unit's output, or activation is defined by $a = g((w_1 * x_1) + (w_2 * x_2) + \dots + (w_n * x_n))$.
 - Each unit computes a (simple) function g of the linear combination of its inputs.
- Learning by tuning the connection weights

Appropriate Problems for NN Learning

- Instances are represented by many attribute-value pairs.
- The target function output may be discrete-valued, real-valued, or a vector of several real-valued or discrete-valued attributes.
- The training examples may contain errors.
- Long training times are acceptable.
- Fast evaluation of the learned target function may be required.
- The ability of humans to understand the learned target function is not important.

Perceptron



$$o(x_0, \dots, x_n) = \begin{cases} 1 & \text{if } \sum_{i=0}^n w_i x_i > 0 \\ -1 & \text{otherwise} \end{cases}$$

Perceptron

- **Perceptron** is a **Linear Threshold Unit (LTU)**.
- A perceptron takes a vector of real-valued inputs, calculates a linear combination of these inputs, then outputs 1 if the result is greater than some threshold and -1 otherwise.
- Given inputs x_1 through x_n , the output $o(x_1, \dots, x_n)$ computed by the perceptron is:

$$o(x_1, \dots, x_n) = \begin{cases} 1 & \text{if } w_0 + w_1x_1 + w_2x_2 + \dots + w_nx_n > 0 \\ -1 & \text{otherwise} \end{cases}$$

each w_i is a real-valued constant, or weight, that determines the contribution of input x_i to the perceptron output.

- The quantity $(-w_0)$ is a **threshold** that the weighted combination of inputs must surpass in order for the perceptron to output 1.
 - To simplify notation, we imagine an additional constant input $x_0 = 1$

Perceptron - Learning


- Learning a perceptron involves choosing values for weights w_0, \dots, w_n .
- The space H of candidate hypotheses considered in perceptron learning is the set of all possible real-valued weight vectors

$$H = \{\tilde{w} \mid \tilde{w} \in \mathbb{R}^{(n+1)}\}$$


Representational Power of Perceptrons

- A perceptron represents a hyperplane decision surface in the n -dimensional space of instances.
- The perceptron outputs 1 for instances lying on one side of the hyperplane and outputs -1 for instances lying on the other side.
- The equation for this decision hyperplane is $\vec{w} \cdot \vec{x} = 0$
- Some sets of positive and negative examples cannot be separated by any hyperplane.
 - Those that can be separated are called **linearly separable** sets of examples.
- A single perceptron can be used to represent many boolean functions.
 - AND, OR, NAND, NOR are representable by a perceptron
 - XOR cannot be representable by a perceptron.

Representational Power of Perceptrons



Representable by a perceptron



NOT representable by a perceptron

Perceptron - Example


X	Y	Out
0	0	0
0	1	1
1	0	1
1	1	1

Give a perceptron to represent OR function.

Perceptron - Example

X	Y	Out
0	0	0
0	1	1
1	0	1
1	1	1

Give a perceptron to represent OR function.




Out is 1 if $w_0 + w_1 * X + w_2 * Y > 0$; 0 otherwise

What will be the weights?

Perceptron - Example

X	Y	Out
0	0	0
0	1	1
1	0	1
1	1	1

Give a perceptron to represent OR function.



Out is 1 if $w_0 + w_1 * X + w_2 * Y > 0$; 0 otherwise

$$w_0 = -0.5 \quad w_1 = 0.7 \quad w_2 = 0.7$$

Perceptron - Example


Attributes		Classification
x	y	
0	0	1
0	1	1
1	0	1
1	1	0

Give a linear threshold unit (a perceptron) that implements the following function by giving its weight values.

Perceptron - Example

Attributes		Classification
x	y	
0	0	1
0	1	1
1	0	1
1	1	0

Give a linear threshold unit (a perceptron) that implements the following function by giving its weight values.




Out is 1 if $w_0 + w_1 * x + w_2 * y > 0$; 0 otherwise

What will be the weights?

Perceptron - Example

Attributes		Classification
x	y	
0	0	1
0	1	1
1	0	1
1	1	0

Give a linear threshold unit (a perceptron) that implements the following function by giving its weight values.



Out is 1 if $w_0 + w_1 * x + w_2 * y > 0$; 0 otherwise

$w_0 = 1$ $w_1 = -0.6$ $w_2 = -0.6$

Perceptron - Example


Attributes			Classification
x	y	z	
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	1
1	0	1	1
1	1	0	1
1	1	1	1

Give a linear threshold unit (a perceptron) that implements the following function by giving its weight values.

Perceptron - Example

Attributes			Classification
x	y	z	
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	1
1	0	1	1
1	1	0	1
1	1	1	1

Give a linear threshold unit (a perceptron) that implements the following function by giving its weight values.




Out is 1 if $w_0 + w_1 * x + w_2 * y + w_3 * z > 0$; 0 otherwise
What will be the weights?

Perceptron - Example

Attributes			Classification
x	y	z	
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	1
1	0	1	1
1	1	0	1
1	1	1	1

Give a linear threshold unit (a perceptron) that implements the following function by giving its weight values.



Out is 1 if $w_0 + w_1 * x + w_2 * y + w_3 * z > 0$; 0 otherwise

$w_0 = -0.5$ $w_1 = 0.6$ $w_2 = 0.3$ $w_3 = 0.3$

Perceptron Training Rule

- To learn an acceptable weight vector is to begin with random weights, then iteratively apply the perceptron to each training example, modifying the perceptron weights whenever it misclassifies an example.
 - If the training example classifies correctly, weights are not updated.
- This process is repeated, iterating through the training examples as many times as needed until the perceptron classifies all training examples correctly.
 - Each pass through all of the training examples is called one **epoch**
- Weights are modified at each step according to **perceptron training rule**

Perceptron Training Rule

- Weights are modified at each step according to **perceptron training rule**.

$$\mathbf{w}_i = \mathbf{w}_i + \Delta \mathbf{w}_i$$
$$\Delta \mathbf{w}_i = \eta (\mathbf{t} - \mathbf{o}) \mathbf{x}_i$$

t is the target value

o is the perceptron output

η is a small constant (e.g. 0.1) called *learning rate*

- If the output is correct ($t=o$) the weights w_i are not changed
- If the output is incorrect ($t \neq o$) the weights w_i are changed such that the output of the perceptron for the new weights is *closer* to t .
- The algorithm converges to the correct classification
 - if the training data is linearly separable
 - and η is sufficiently small


Perceptron Learning Rule -Example

$$\begin{aligned}w_i &= w_i + \Delta w_i \\ \Delta w_i &= \eta (t - o) x_i \\ \eta &= 0.1\end{aligned}$$

red ones are positive examples
i.e. $t=1$

blue ones are negative examples
i.e. $t=-1$

$$\begin{aligned}[w_0, w_1, w_2] &= [0.25, -0.1, 0.5] \\ x_2 &= 0.2 x_1 - 0.5\end{aligned}$$




Perceptron Learning Rule -Example

$$\begin{aligned}w_i &= w_i + \Delta w_i \\ \Delta w_i &= \eta (t - o) x_i \\ \eta &= 0.1\end{aligned}$$

red ones are positive examples
i.e. $t=1$


blue ones are negative examples
i.e. $t=-1$

$$\begin{aligned}[w_0, w_1, w_2] &= [0.25, -0.1, 0.5] \\ x_2 &= 0.2 x_1 - 0.5\end{aligned}$$



Perceptron Learning Rule -Example

$$\begin{aligned}w_i &= w_i + \Delta w_i \\ \Delta w_i &= \eta (t - o) x_i \\ \eta &= 0.1\end{aligned}$$




$$\begin{aligned}[w_0, w_1, w_2] &= [0.25, -0.1, 0.5] \\ x_2 &= 0.2 x_1 - 0.5\end{aligned}$$

consider this example
 $(x, t) = (-1, -1)$
 $o = \text{sgn}(0.25 + 0.1 - 0.5) = -1$

Perceptron Learning Rule – Example

$$\begin{aligned}w_i &= w_i + \Delta w_i \\ \Delta w_i &= \eta (t - o) x_i \\ \eta &= 0.1\end{aligned}$$

$$\begin{aligned}[w_0, w_1, w_2] &= [0.25, -0.1, 0.5] \\ x_2 &= 0.2 x_1 - 0.5\end{aligned}$$




$$\begin{aligned}\Delta w &= [0.2, -0.2, -0.2] \\ [w_0, w_1, w_2] &= [0.45, -0.3, 0.3]\end{aligned}$$

Perceptron Learning Rule – Example

$$\begin{aligned}w_i &= w_i + \Delta w_i \\ \Delta w_i &= \eta (t - o) x_i \\ \eta &= 0.1\end{aligned}$$

consider this example
 $(x,t) = ([2,1], -1)$
 $o = \text{sgn}(0.45 - 0.6 + 0.3) = 1$

$$\begin{aligned}[w_0, w_1, w_2] &= [0.25, -0.1, 0.5] \\ x_2 &= 0.2 x_1 - 0.5\end{aligned}$$



$$\begin{aligned}\Delta w &= [0.2, -0.2, -0.2] \\ [w_0, w_1, w_2] &= [0.45, -0.3, 0.3]\end{aligned}$$

$$w_i = w_i + \Delta w_i$$

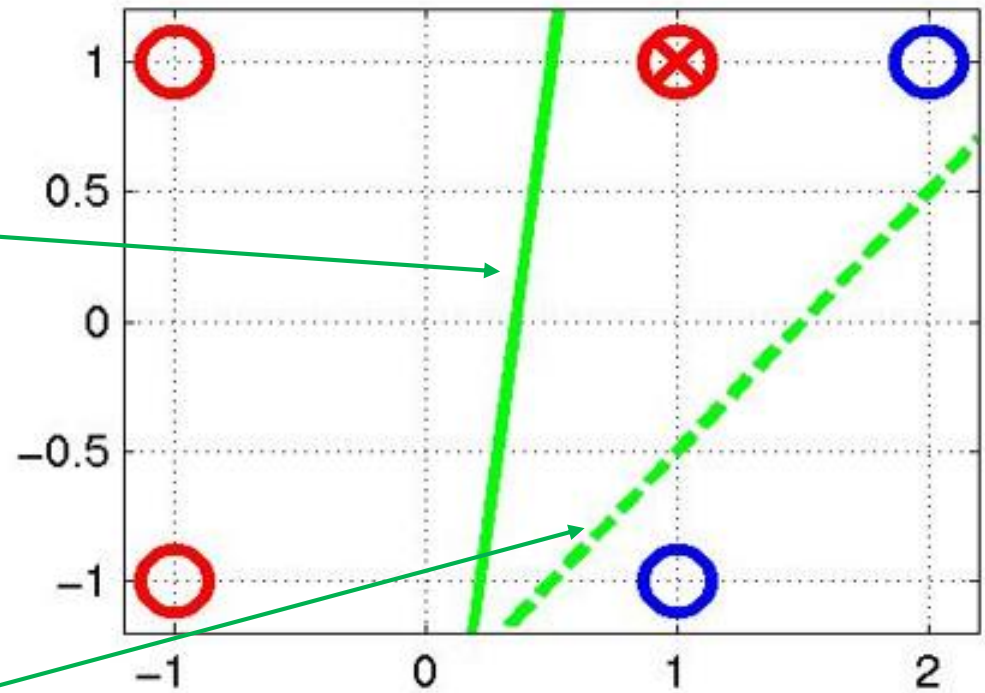
$$\Delta w_i = \eta (t - o) x_i$$

$$\eta = 0.1$$

Perceptron Learning Rule – Example

$$\Delta w = [-0.2, -0.4, -0.2]$$

$$[w_0, w_1, w_2] = [0.25, -0.7, 0.1]$$



$$[w_0, w_1, w_2] = [0.45, -0.3, 0.3]$$

Perceptron Learning Rule – Example

$$\begin{aligned}w_i &= w_i + \Delta w_i \\ \Delta w_i &= \eta (t - o) x_i \\ \eta &= 0.1\end{aligned}$$

consider this example

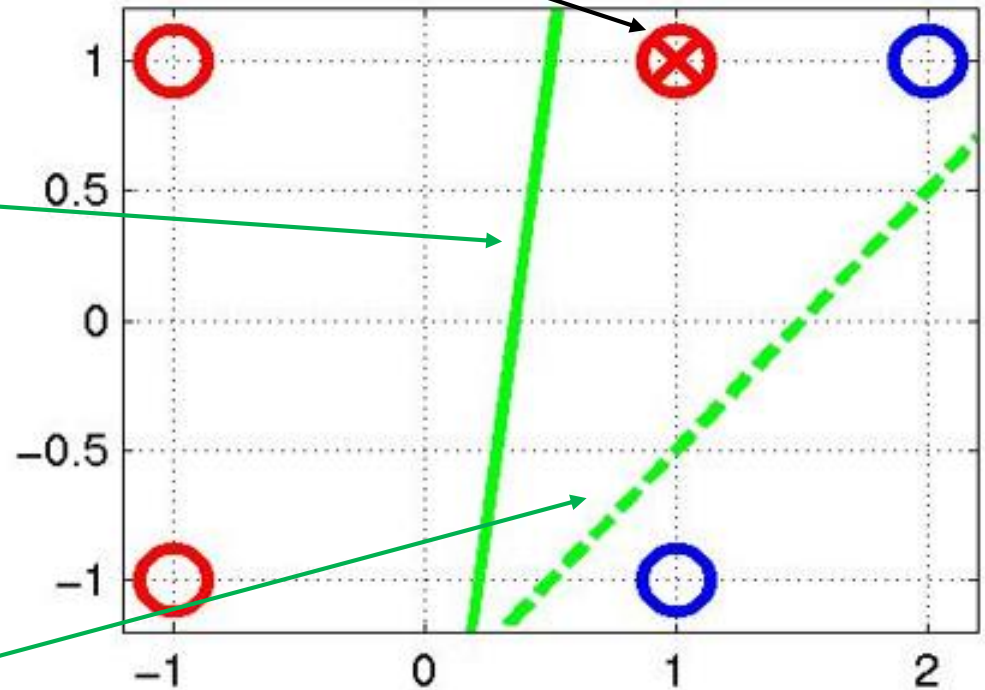
$$(x,t) = ([1,1],1)$$

$$o = \text{sgn}(0.25 - 0.7 + 0.1) = -1$$

$$\Delta w = [-0.2, -0.4, -0.2]$$

$$[w_0, w_1, w_2] = [0.25, -0.7, 0.1]$$


$$[w_0, w_1, w_2] = [0.45, -0.3, 0.3]$$



Perceptron Learning Rule – Example

$$\begin{aligned} \mathbf{w}_i &= \mathbf{w}_i + \Delta \mathbf{w}_i \\ \Delta \mathbf{w}_i &= \eta (t - o) \mathbf{x}_i \\ \eta &= 0.1 \end{aligned}$$

$$[\mathbf{w}_0, \mathbf{w}_1, \mathbf{w}_2] = [0.25, -0.7, 0.1]$$



$$\Delta \mathbf{w} = [0.2, 0.2, 0.2]$$

$$[\mathbf{w}_0, \mathbf{w}_1, \mathbf{w}_2] = [0.45, -0.5, 0.3]$$

Perceptron Learning Rule – Learning OR

x1	x2	T	O	$\Delta w1$	w1	$\Delta w2$	w2	$\Delta w0$	w0
-	-	-	-	-	.1	-	.5	-	-.8
0	0	0	0	0	.1	0	.5	0	-.8
0	1	1	0	0	.1	.2	.7	.2	-.6
1	0	1	0	.2	.3	0	.7	.2	-.4
1	1	1	1	0	.3	0	.7	0	-.4
0	0	0	0	0	.3	0	.7	0	-.4
0	1	1	1	0	.3	0	.7	0	-.4
1	0	1	0	.2	.5	0	.7	.2	-.2
1	1	1	1	0	.5	0	.7	0	-.2
0	0	0	0	0	.5	0	.7	0	-.2
0	1	1	1	0	.5	0	.7	0	-.2
1	0	1	1	0	.5	0	.7	0	-.2
1	1	1	1	0	.5	0	.7	0	-.2

Learning rate parameter is 0.2

The result of executing the learning algorithm for 3 epochs.

Perceptron Learning Rule – Learning AND

x1	x2	T	O	$\Delta w1$	w1	$\Delta w2$	w2	$\Delta w0$	w0
-	-	-	-	-	.1	-	.5	-	-.8
0	0	0	0	0	.1	0	.5	0	-.8
0	1	0	0	0	.1	0	.5	0	-.8
1	0	0	0	0	.1	0	.5	0	-.8
1	1	1	0	.2	.3	.2	.7	.2	-.6
0	0	0	0	0	.3	0	.7	0	-.6
0	1	0	1	0	.3	-.2	.5	-.2	-.8
1	0	0	0	0	.3	0	.5	0	-.8
1	1	1	0	.2	.5	.2	.7	.2	-.6
0	0	0	0	0	.5	0	.7	0	-.6
0	1	0	1	0	.5	-.2	.5	-.2	-.8
1	0	0	0	0	.5	0	.5	0	-.8
1	1	1	1	0	.5	0	.5	0	-.8

Learning rate parameter is 0.2

The result of executing the learning algorithm for 4 epochs.

1 epoch not shown in the table.

Gradient Descent and the Delta Rule

- The **perceptron rule** finds a successful weight vector when the training examples are linearly separable, it can fail to converge if the examples are not linearly separable.
- The **delta rule** overcomes this difficulty.
- If the training examples are not linearly separable, the delta rule converges toward a best-fit approximation to the target concept.
- The key idea behind the delta rule is to use **gradient descent** to search the hypothesis space of possible weight vectors to find the weights that best fit the training examples.
- The delta rule is important because gradient descent provides the basis for the **BACKPROPAGATION** algorithm, which can learn networks with many interconnected units.
 - The gradient descent can serve as the basis for learning algorithms that must search through hypothesis spaces containing many different types of continuously parameterized hypotheses

Gradient Descent

- Consider linear unit without threshold and continuous output o (not just $-1,1$)


$$o = w_0 + w_1 x_1 + \dots + w_n x_n$$

- Train the w_i 's such that they minimize the squared error

$$E[w_0, \dots, w_n] = \frac{1}{2} \sum_{d \in D} (t_d - o_d)^2$$

where D is the set of training examples t_d is the target output for training example d , and o_d is the output of the linear unit for training example d .

Gradient Descent



- The w_0, w_1 plane represents the entire hypothesis space.
 - The vertical axis indicates the error E relative to some fixed set of training examples.
 - The error surface summarizes the desirability of every weight vector in the hypothesis space (we desire a hypothesis with minimum error).
 - The arrow shows the **negated gradient** at one particular point, indicating the direction in the w_0, w_1 plane producing steepest descent along the error surface.
-
- **Gradient descent search** determines a weight vector that minimizes E by starting with an arbitrary initial weight vector, then repeatedly modifying it in small steps.
 - At each step, the weight vector is altered in the direction that produces the **steepest descent** along the error surface.
 - This process continues until the global minimum error is reached.

Gradient Descent

- How can we calculate the direction of steepest descent along the error surface?
- This direction can be found by computing the derivative of E with respect to each component of the vector w_0, \dots, w_n .
- This vector derivative is called the *gradient* of E with respect to w_1, \dots, w_n , written $\nabla E(w_0, \dots, w_n)$
- **Gradient:**
$$\nabla E[w_0, \dots, w_n] = [\partial E / \partial w_0, \dots, \partial E / \partial w_n]$$
- When the gradient is interpreted as a vector in weight space, the gradient specifies the direction that produces the steepest increase in E .
- The negative of this vector ($-\nabla E[w_0, \dots, w_n]$) gives the direction of steepest decrease.

Training Rule for Gradient Descent

- For each weight w_i
 - $w_i = w_i + \Delta w_i$
 - $\Delta w_i = -\eta \nabla E[w_i]$
- η is a small constant called *learning rate*

$$\begin{aligned} \Delta w_i &= -\eta \nabla E[w_i] \\ &= -\eta \left(\frac{\partial E}{\partial w_i} \right) \end{aligned}$$

·
·

$$\begin{aligned} \frac{\partial E}{\partial w_i} &= \frac{\partial}{\partial w_i} \frac{1}{2} \sum_{d \in D} (t_d - o_d)^2 \\ &= \frac{1}{2} \sum_{d \in D} \frac{\partial}{\partial w_i} (t_d - o_d)^2 \\ &= \frac{1}{2} \sum_{d \in D} 2(t_d - o_d) \frac{\partial}{\partial w_i} (t_d - o_d) \\ &= \sum_{d \in D} (t_d - o_d) \frac{\partial}{\partial w_i} (t_d - (w_0 x_{0d} + \dots + w_n x_{nd})) \end{aligned}$$

$$\begin{aligned} \Delta w_i &= -\eta \sum_{d \in D} (t_d - o_d) (-x_{id}) \\ &= \eta \sum_{d \in D} (t_d - o_d) (x_{id}) \end{aligned}$$

$$\frac{\partial E}{\partial w_i} = \sum_{d \in D} (t_d - o_d) (-x_{id})$$

Gradient Descent

Gradient-Descent(*training_examples*, η)

Each training example is a pair of the form $\langle(x_1, \dots, x_n), t\rangle$ where (x_1, \dots, x_n) is the vector of input values, and t is the target output value, η is the learning rate (e.g. 0.1)

- Initialize each w_i to some small random value
- Until the termination condition is met, Do
 - Initialize each Δw_i to zero
 - For each $\langle(x_1, \dots, x_n), t\rangle$ in *training_examples*, Do
 - Input the instance (x_1, \dots, x_n) to the linear unit and compute the output o
 - For each linear unit weight w_i Do
 - $\Delta w_i = \Delta w_i + \eta (t - o) x_i$
 - For each linear unit weight w_i , Do
 - $w_i = w_i + \Delta w_i$

Gradient Descent: Example

						t
$f(x_1, x_2) = 0.5 + 0.5 \cdot x_1 - 0.5 \cdot x_2$				x1	x2	f(x1, x2)
				1	1	0,50
				1	3	-0,50
				2	1	1,00
				2	4	-0,50
$f(x_1, x_2) = w_0 + w_1 \cdot x_1 + w_2 \cdot x_2$						
w0	w1	w2		eta		
0,400	0,400	-0,400		0,050		

Gradient Descent: Example

x1	x2	o	eta(t-o)	deltaw0	deltaw1	deltaw2
1	1	0,400	0,005	0,005	0,005	0,005
1	3	-0,400	-0,005	-0,005	-0,005	-0,015
2	1	0,800	0,010	0,010	0,020	0,010
2	4	-0,400	-0,005	-0,005	-0,010	-0,020
				0,005	0,010	-0,020
		w0	w1	w2		
		0,405	0,410	-0,420		
x1	x2	o	eta(t-o)	deltaw0	deltaw1	deltaw2
1	1	0,395	0,005	0,005	0,005	0,005
1	3	-0,445	-0,003	-0,003	-0,003	-0,008
2	1	0,805	0,010	0,010	0,020	0,010
2	4	-0,455	-0,002	-0,002	-0,004	-0,009
				0,010	0,018	-0,002
		w0	w1	w2		
		0,415	0,428	-0,422		

		t
x1	x2	f(x1,x2)
1	1	0,50
1	3	-0,50
2	1	1,00
2	4	-0,50

Gradient Descent: Example

x1	x2	o	eta(t-o)	deltaw0	deltaw1	deltaw2
1	1	0,420	0,004	0,004	0,004	0,004
1	3	-0,424	-0,004	-0,004	-0,004	-0,011
2	1	0,848	0,008	0,008	0,015	0,008
2	4	-0,419	-0,004	-0,004	-0,008	-0,016
				0,004	0,007	-0,016
		w0	w1	w2		
		0,419	0,435	-0,438		
x1	x2	o	eta(t-o)	deltaw0	deltaw1	deltaw2
1	1	0,415	0,004	0,004	0,004	0,004
1	3	-0,461	-0,002	-0,002	-0,002	-0,006
2	1	0,850	0,007	0,007	0,015	0,007
2	4	-0,464	-0,002	-0,002	-0,004	-0,007
				0,008	0,014	-0,001
		w0	w1	w2		
		0,427	0,449	-0,439		

		t
x1	x2	f(x1,x2)
1	1	0,50
1	3	-0,50
2	1	1,00
2	4	-0,50

Incremental (Stochastic) Gradient Descent

Gradient descent is a strategy for searching through a large or infinite hypothesis space that can be applied whenever

- the hypothesis space contains continuously parameterized hypotheses (e.g., the weights in a linear unit), and
- the error can be differentiated with respect to these hypothesis parameters.

The key practical difficulties in applying gradient descent are

- converging to a local minimum can sometimes be quite slow (i.e., it can require many thousands of gradient descent steps), and
- if there are multiple local minima in the error surface, then there is no guarantee that the procedure will find the global minimum.

Incremental (stochastic) gradient descent tries to solve these problems.

Incremental (Stochastic) Gradient Descent

- Batch mode : gradient descent

$w = w - \eta \nabla E_D[w]$ over the entire data D

$$E_D[w] = 1/2 \sum_d (t_d - o_d)^2$$

- Incremental mode: gradient descent

$w = w - \eta \nabla E_d[w]$ over individual training examples d

$$E_d[w] = 1/2 (t_d - o_d)^2$$

Incremental Gradient Descent can approximate Batch Gradient Descent arbitrarily closely if η is small enough

Incremental (Stochastic) Gradient Descent

Incremental (Stochastic) Gradient-Descent(*training_examples*, η)

Each training example is a pair of the form $\langle (x_1, \dots, x_n), t \rangle$ where (x_1, \dots, x_n) is the vector of input values, and t is the target output value, η is the learning rate (e.g. 0.1)

- Initialize each w_i to some small random value
- Until the termination condition is met, Do
 - ~~Initialize each Δw_i to zero~~
 - For each $\langle (x_1, \dots, x_n), t \rangle$ in *training_examples*, Do
 - Input the instance (x_1, \dots, x_n) to the linear unit and compute the output o
 - For each linear unit weight w_i Do
 - ~~// $\Delta w_i = \Delta w_i + \eta(t - o) x_i$ //~~ $w_i = w_i + \eta(t - o) x_i$
 - ~~For each linear unit weight w_i Do~~
 - ~~// $w_i = w_i + \Delta w_i$ //~~

Incremental (Stochastic) Gradient Descent

- In standard gradient descent, the error is summed over all examples before updating weights, whereas in stochastic gradient descent weights are updated upon examining each training example.
- Summing over multiple examples in standard gradient descent requires more computation per weight update step. On the other hand, because it uses the true gradient, standard gradient descent is often used with a larger step size per weight update than stochastic gradient descent.
- In cases where there are multiple local minima with respect to $E(w_0, \dots, w_n)$, stochastic gradient descent can sometimes avoid falling into these local minima because it uses the various $\nabla E_d(w_0, \dots, w_n)$ rather than $\nabla E(w_0, \dots, w_n)$ to guide its search.

Comparison Perceptron and Gradient Descent Rule

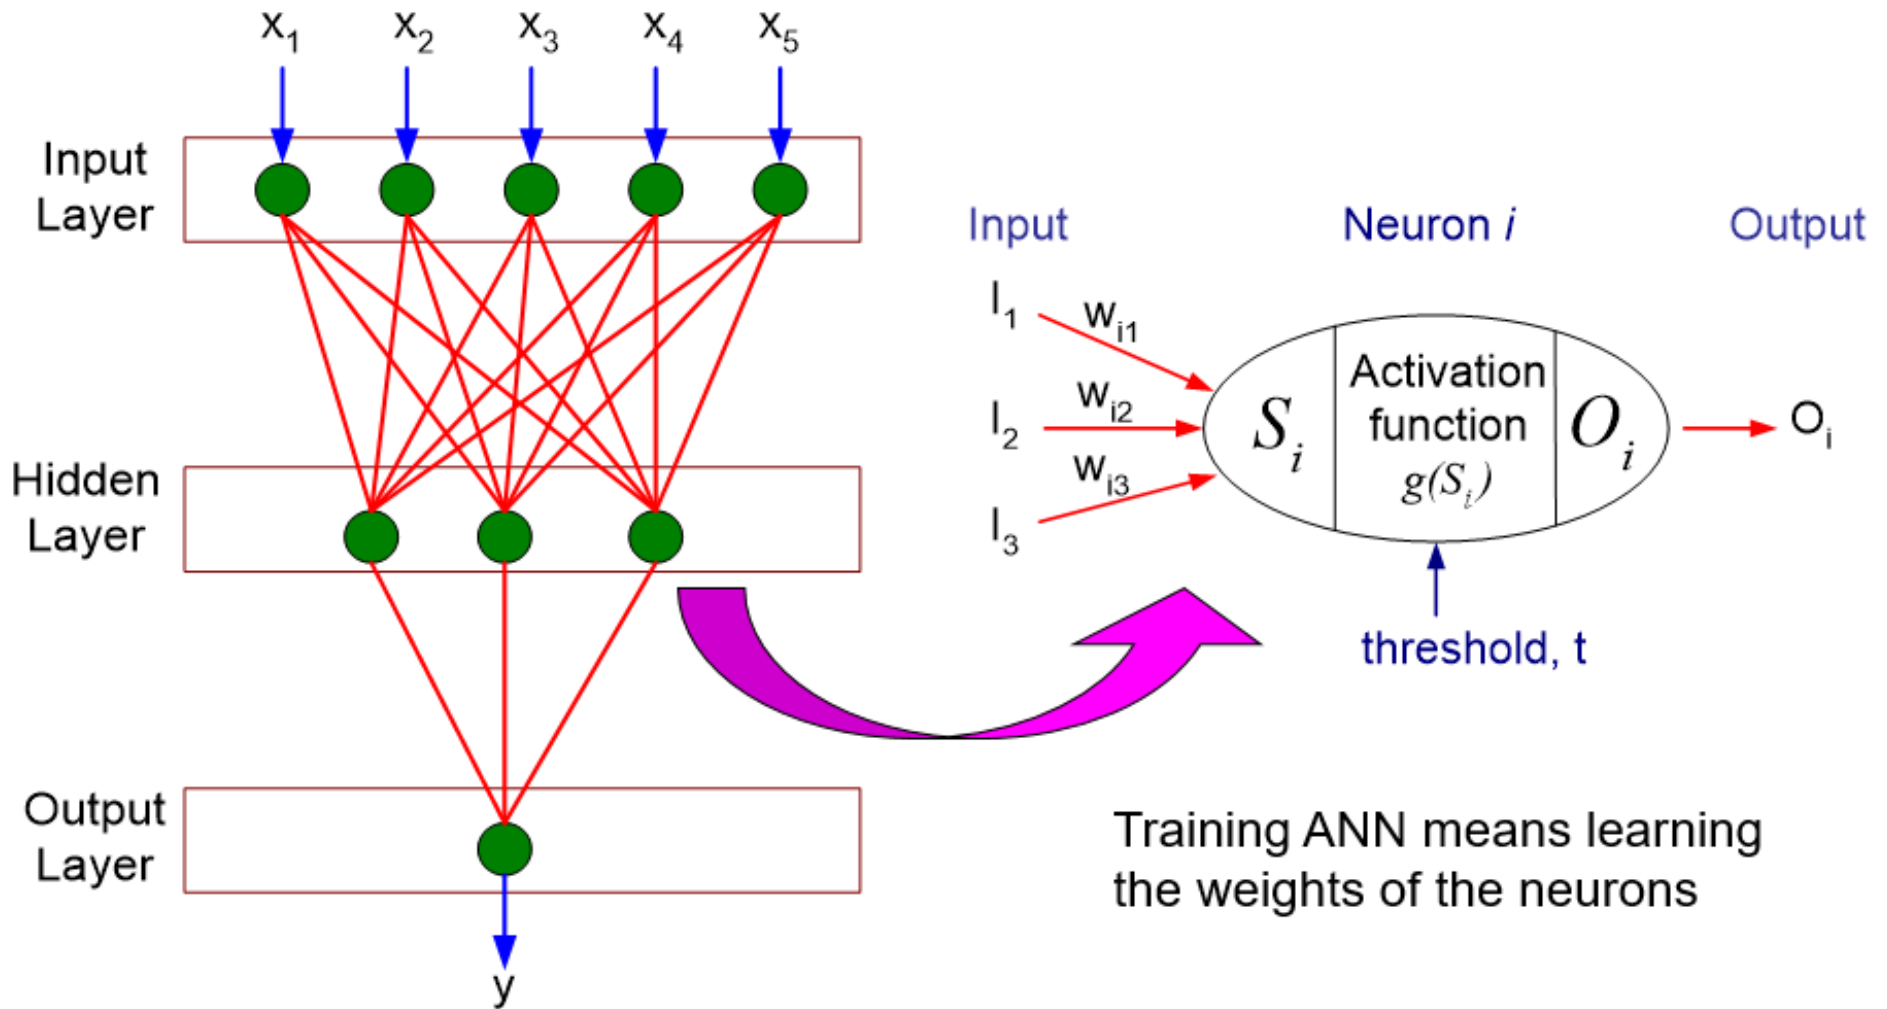
Perceptron learning rule guaranteed to succeed if

- Training examples are linearly separable
- Sufficiently small learning rate η

Linear unit training rules uses gradient descent


- Guaranteed to converge to hypothesis with minimum squared error
- Given sufficiently small learning rate η
- Even when training data contains noise
- Even when training data not separable by H

General Structure of ANN



Multi-Layer Networks

- Single perceptrons can only express linear decision surfaces.
- Multilayer networks are capable of expressing a rich variety of nonlinear decision surfaces.



Multi-Layer Networks with Linear Units


Ex. XOR

- Multiple layers of cascaded linear units still produce only linear functions.

$$\text{OR: } 0.5*x_1 + 0.5*x_2 - 0.25 > 0$$

$$\text{AND: } 0.5*x_1 + 0.5*x_2 - 0.75 > 0$$

$$\text{XOR: } 0.5*x_1 - 0.5*x_2 - 0.25 > 0$$



Multi-Layer Networks with Linear Units - Example


Give an artificial neural network for the following function. Make clear the structure of your ANN and the used weights.

Attributes		Classification
x	y	
1	1	0
1	2	1
2	1	1
2	2	0

Multi-Layer Networks with Linear Units - Example


Give an artificial neural network for the following function. Make clear the structure of your ANN and the used weights.

Attributes		Classification
x	y	
1	1	0
1	2	1
2	1	1
2	2	0



Multi-Layer Networks with Linear Units - Example

Attributes		Classification
x	y	
1	1	0
1	2	1
2	1	1
2	2	0



Unit1: $w10 = -1$ $w11 = 0.4$ $w12 = 0.4$

Unit2: $w20 = 1$ $w21 = -0.3$ $w22 = -0.3$

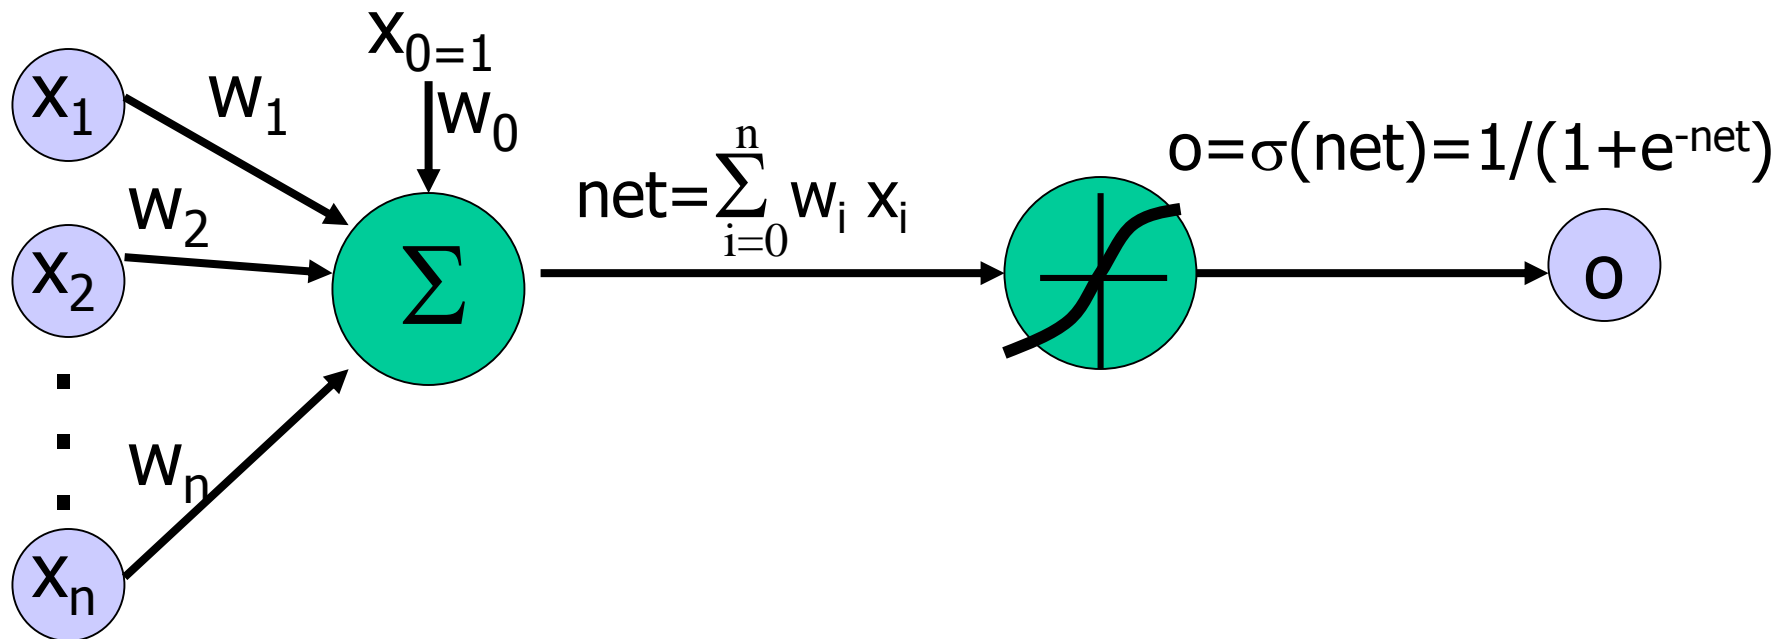
Unit3: $w30 = -0.6$ $w31 = 0.5$ $w32 = 0.5$ (and function)

x	y	unit1	unit2	unit3	Classification
1	1	0	1	0	0
1	2	1	1	1	1
2	1	1	1	1	1
2	2	1	0	0	0

Multi-Layer Networks with Non-Linear Units

- Multiple layers of cascaded linear units still produce only linear functions.
- We prefer networks capable of representing highly nonlinear functions.
- What we need is a unit whose output is a nonlinear function of its inputs, but whose output is also a differentiable function of its inputs.
- One solution is the **sigmoid unit**, a unit very much like a perceptron, but based on a smoothed, differentiable threshold function.

Sigmoid Unit



$\sigma(x)$ is the sigmoid function: $1/(1+e^{-x})$

$d\sigma(x)/dx = \sigma(x) (1 - \sigma(x))$ ← differentiable function

Sigmoid Unit

Derive gradient descent rules to train:

- one sigmoid function

$$\partial E / \partial w_i = -\sum_d (t_d - o_d) o_d (1 - o_d) x_{id}$$

- Multilayer networks of sigmoid units \rightarrow backpropagation

Backpropagation Algorithm

- Create a feed-forward network with n_i inputs, n_{hidden} hidden units, and n_{out} output units.
- Initialize each w_i to some small random value (e.g., between -.05 and .05).
- Until the termination condition is met, Do
 - For each training example $\langle(x_1, \dots, x_n), t\rangle$, Do
 - // Propagate the input forward through the network:
 - 1. Input the instance (x_1, \dots, x_n) to the network and compute the network outputs o_k for every unit

// Propagate the errors backward through the network:

2. For each output unit k , calculate its error term δ_k

$$\delta_k = o_k(1-o_k)(t_k-o_k)$$

3. For each hidden unit h , calculate its error term δ_h

$$\delta_h = o_h(1-o_h) \sum_k w_{h,k} \delta_k$$

4. For each network weight $w_{i,j}$, Do

$$w_{i,j} = w_{i,j} + \Delta w_{i,j} \quad \text{where} \quad \Delta w_{i,j} = \eta \delta_j x_{i,j}$$


Backpropagation

- Gradient descent over entire *network* weight vector
- Easily generalized to arbitrary directed graphs
- Will find a local, not necessarily global error minimum
 - in practice often works well (can be invoked multiple times with different initial weights)
- Often include weight *momentum* term




$$\Delta w_{i,j}(n) = \eta \delta_j x_{i,j} + \alpha \Delta w_{i,j}(n-1)$$

- Minimizes error training examples
- Training can be slow typical 1000-10000 iterations
- Using network after training is fast



Single Layer – Linear Problem



Multi-Layer Network – NonLinear Problem





Multi-Layer Network – NonLinear Problem2



$$\text{output} = f \tanh_1(x) + g \tanh_2(x) + a$$
$$\tanh_1(x) = \tanh(d \cdot x + b)$$
$$\tanh_2(x) = \tanh(e \cdot x + c)$$

Multi-Layer Network – NonLinear Problem2



Expressive Capabilities of ANNs

Boolean functions

- Every boolean function can be represented by network with single hidden layer
- but might require exponential (in number of inputs) hidden units

Continuous functions

- Every bounded continuous function can be approximated with arbitrarily small error by network with one hidden layer
- Any function can be approximated to arbitrary accuracy by a network with two hidden layers