Neural Networks

## The Neuron Metaphor

Biological Neural Networks


## The Neuron Metaphor

- Inspired by attempts to simulate biological neural systems.
- Neurons
- accept information from multiple inputs,
- transmit information to other neurons.
- Multiply inputs by weights along edges
- Apply some function to the set of inputs at each node



## Artificial Neural Networks (ANN)

| $\mathrm{X}_{1}$ | $\mathrm{X}_{2}$ | $\mathrm{X}_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 1 | 0 | 0 | -1 |
| 1 | 0 | 1 | 1 |
| 1 | 1 | 0 | 1 |
| 1 | 1 | 1 | 1 |
| 0 | 0 | 1 | -1 |
| 0 | 1 | 0 | -1 |
| 0 | 1 | 1 | 1 |
| 0 | 0 | 0 | -1 |



Output Y is 1 if at least two of the three inputs are equal to 1 .

## Artificial Neural Networks (ANN)

Input

| $\mathrm{X}_{1}$ | $\mathrm{X}_{2}$ | $\mathrm{X}_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 1 | 0 | 0 | -1 |
| 1 | 0 | 1 | 1 |
| 1 | 1 | 0 | 1 |
| 1 | 1 | 1 | 1 |
| 0 | 0 | 1 | -1 |
| 0 | 1 | 0 | -1 |
| 0 | 1 | 1 | 1 |
| 0 | 0 | 0 | -1 |



$$
Y=\operatorname{sign}\left(0.3 X_{1}+0.3 X_{2}+0.3 X_{3}-0.4\right)
$$

$$
\text { where } \operatorname{sign}(x)=\left\{\begin{array}{cc}
1 & \text { if } x \geq 0 \\
-1 & \text { if } x<0
\end{array}\right.
$$

## Artificial Neural Networks (ANN)

- Model is an assembly of inter-connected nodes and weighted links
- Output node sums up each of its input value according to the weights of its links
- Compare output node against some threshold t (also named bias b)
- Bias $b=>$ the output of the transformation is biased toward being $\mathbf{b}$ in the absence of any input.


## Characterizing the Artificial Neuron

- Input/Output signal may be:
- Real value
- Unipolar $\{0,1\}$
- Bipolar [-1, +1]
- Weight (w or sigma): $\theta_{i j}$ - strength of connection from unit $j$ to unit $i$
- Learning amounts to adjusting the weights $\theta_{i j}$ by means of an optimization algorithm aiming to minimize a cost function, i.e., as in biological systems training a perceptron model amounts to adapting the weights of the links until they fit the input-output relationships of the underlying data.


## Characterizing the Artificial Neuron

- The bias $b$ is a constant that can be written as $\theta_{i 0} x_{0}$ with $x_{0}=1$ and $\theta_{i 0}=\mathrm{b}$ such that

$$
n e t_{i}=\sum_{j=0}^{n} \theta_{i j} x_{j}
$$

- The function $f\left(\right.$ net $\left._{i}(x)\right)$ is the unit's activation function for the output neuron.
- The simplest case, $f$ is the identity function, and the unit's output is just its net input. This is called a linear unit.
- Otherwise we can have other functions that we see later ....


## The Perceptron Classifier

## Perceptron

- Single layer network
- Contains only input and output nodes
- Activation function: $f=\operatorname{sign}(w \bullet x)$
- Applying model is straightforward

$$
\begin{aligned}
& Y=\operatorname{sign}\left(0.3 X_{1}+0.3 X_{2}+0.3 X_{3}-0.4\right) \\
& \text { where } \operatorname{sign}(x)=\left\{\begin{array}{cl}
1 & \text { if } x \geq 0 \\
-1 & \text { if } x<0
\end{array}\right.
\end{aligned}
$$

- $X_{1}=1, X_{2}=0, X_{3}=1=>y=\operatorname{sign}(0.2)=1$


## Learning Iterative Procedure

During the training phase the weight parameters are adjusted until the outputs of the perceptron become consistent with the true outputs of the training examples.

Initialize the weights ( $w_{0}, w_{1}, \ldots, w_{m}$ ) - randomly
Repeat
For each training example ( $x_{i}, y_{i}$ )
Compute $f\left(w^{(k)}, x_{i}\right)$
Update the weights: $w^{(k+1)}=w^{(k)}+\lambda\left[y_{i}-f\left(w^{(k)}, x_{i}\right)\right] x_{i}$
Leapping condition is met rate

## Perceptron Learning Rule

- Weight update formula:

$$
w^{(k+1)}=w^{(k)}+\lambda\left[y_{i}-f\left(w^{(k)}, x_{i}\right)\right] x_{i} ; \lambda: \text { learning rate }
$$

- Intuition:
- Update weight based on error: $e=\left[y_{i}-f\left(w^{(k)}, x_{i}\right)\right]$
- If $y=f(x, w), e=0$ : no update needed
- If $y>f(x, w), e=2$ : weight must be increased so that $f(x, w)$ will increase
- If $y<f(x, w), e=-2$ : weight must be decreased so that $f(x, w)$ will decrease


## The Learning Rate

- Is a parameter with value between 0 and 1 used to control the amount of adjustment made in each iteration.
- If is close to 0 the new weight is mostly influenced by the value of the old weight.
- If it is close to 1 , then the new weight is mostly influenced by the current adjustment.
- The learning rate can be adaptive: initially moderately large and the gradually decreases in subsequent iterations.


## The delta rule

- Whenever the perceptron makes a wrong classification $\rightarrow$ change weights and threshold in "appropriate direction".
- If the desired output is $\mathbf{1}$ and the perceptron's output is $\mathbf{0}$ lower the threshold and adjust the weights depending on the sign and magnitude of the inputs.
- If the desired output is $\mathbf{0}$ and the perceptron's output is $\mathbf{1}$ increase the threshold and adjust the weights depending on the sign and magnitude of the inputs.


## The delta rule

- The delta rule recommends to adjust the weight and the threshold values as:

$$
\begin{gathered}
w_{i}^{\text {new }}=w_{i}^{\text {old }}+\Delta w_{i} \\
b^{\text {new }}=b^{\text {old }}+\Delta b
\end{gathered}
$$

- $w_{i}$ : A weight of the perceptron
- $b$ : The threshold value of the perceptron
- $\left(x_{1}, x_{2}, \ldots, x_{n}\right)$ : An input vector
- $t$ : the desired output for input vector $\left(x_{1}, x_{2}, \ldots, x_{n}\right)$
- $y$ : the real output of the Perceptron for input vector $\left(x_{1}, x_{2}, \ldots, x_{n}\right)$
- $\lambda>0$ : the Learning rate


## The delta rule

- The delta rule recommends to adjust the weight and the threshold values as:

$$
\begin{gathered}
\Delta w_{i}=\left\{\begin{array}{lr}
0 & \text { if correct prediction } \\
+\lambda x_{i} & \text { if } y=0 \text { and } t=1 \\
-\lambda x_{i} & \text { if } y=1 \text { and } t=0
\end{array}\right. \\
\Delta b= \begin{cases}0 & \text { if correct prediction } \\
-\lambda & \text { if } y=0 \text { and } t=1 \\
+\lambda & \text { if } y=1 \text { and } t=0\end{cases}
\end{gathered}
$$

## Example: Learning the logical operator AND

- Training Data:

$$
\begin{equation*}
\{((0,0), 0),((0,1), 0),((1,0), 0),((1,1), 1)\} \tag{1}
\end{equation*}
$$

- Learning rate $\lambda=1$
- $\mathrm{t}=$ target

- Initialization: $w_{1}=w_{2}=b=0$

|  | $\boldsymbol{x}$ | t | y | $\Delta w_{1}$ | $\Delta w_{2}$ | $\Delta \mathrm{b}$ | $w_{1}^{n e w}$ | $w_{2}^{\text {new }}$ | $b^{\text {new }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |  | 0 | 0 | 0 |
| Epoch 1 | 00 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
|  | 01 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
|  | 10 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
|  | 11 | 1 | 0 | 1 | 1 | -1 | 1 | 1 | -1 |

$$
\begin{gathered}
\Delta w_{i}=\left\{\begin{array}{lr}
0 & \text { if correct prediction } \\
+\lambda x_{i} & \text { if } y=0 \text { and } t=1 \\
-\lambda x_{i} & \text { if } y=1 \text { and } t=0
\end{array}\right. \\
\Delta b= \begin{cases}0 & \text { if correct prediction } \\
-\lambda & \text { if } y=0 \text { and } t=1 \\
+\lambda & \text { if } y=1 \text { and } t=0\end{cases}
\end{gathered}
$$

## Example: Learning the logical operator AND

|  | $x$ | t | y | $\Delta w_{1}$ | $\Delta w_{2}$ | $\Delta \mathrm{b}$ | $w_{1}^{n e w}$ | $w_{2}^{n e w}$ | $b^{\text {new }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |  | 1 | 1 | -1 |
| Epoch 2 | 00 | 0 | 1 | 0 | 0 | 1 | 1 | 1 | 0 |
|  | 01 | 0 | 1 | 0 | -1 | 1 | 1 | 0 | 1 |
|  | 10 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 |
|  | 11 | 1 | 0 | 1 | 1 | -1 | 2 | 1 | 0 |
|  | $x$ | t | y | $\Delta w_{1}$ | $\Delta w_{2}$ | $\Delta \mathrm{b}$ | $w_{1}^{n e w}$ | $W_{2}^{\text {new }}$ | $b^{\text {new }}$ |
|  |  |  |  |  |  |  | 2 | 1 | 0 |
| Epoch 3 | 00 | 0 | 0 | 0 | 0 | 0 | 2 | 1 | 0 |
|  | 01 | 0 | 1 | 0 | -1 | 1 | 2 | 0 | 1 |
|  | 10 | 0 | 1 | -1 | 0 | 1 | 1 | 0 | 2 |
|  | 11 | 1 | 0 | 1 | 1 | -1 | 2 | 1 | 1 |

## Example: Learning the logical operator AND

|  | $x$ | t | y | $\Delta w_{1}$ | $\Delta w_{2}$ | $\Delta \mathrm{b}$ | $w_{1}^{\text {new }}$ | $w_{2}^{\text {new }}$ | $b^{\text {new }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |  | 2 | 1 | 1 |
| Epoch 4 | 00 | 0 | 0 | 0 | 0 | 0 | 2 | 1 | 1 |
|  | 01 | 0 | 0 | 0 | 0 | 0 | 2 | 1 | 1 |
|  | 10 | 0 | 1 | -1 | 0 | 1 | 1 | 1 | 2 |
|  | 11 | 1 | 0 | 1 | 1 | -1 | 2 | 2 | 1 |
|  | $x$ | t | y | $\Delta w_{1}$ | $\Delta w_{2}$ | $\Delta \mathrm{b}$ | $w_{1}^{\text {new }}$ | $w_{2}^{\text {new }}$ | $b^{\text {new }}$ |
|  |  |  |  |  |  |  | 2 | 2 | 1 |
| Epoch 5 | 00 | 0 | 0 | 0 | 0 | 0 | 2 | 2 | 1 |
|  | 01 | 0 | 1 | 0 | -1 | 1 | 2 | 1 | 2 |
|  | 10 | 0 | 0 | 0 | 0 | 0 | 2 | 1 | 2 |
|  | 11 | 1 | 1 | 0 | 0 | 0 | 2 | 1 | 2 |

## Example: Learning the logical operator AND

|  | $\boldsymbol{x}$ | t | $y$ | $\Delta w_{1}$ | $\Delta w_{2}$ | $\Delta \mathrm{b}$ | $w_{1}^{n e w}$ | $w_{2}^{\text {new }}$ | $b^{\text {new }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |  | 2 | 1 | 2 |
| Epoch 6 | 00 | 0 | 0 | 0 | 0 | 0 | 2 | 1 | 2 |
|  | 01 | 0 | 0 | 0 | 0 | 0 | 2 | 1 | 2 |
|  | 10 | 0 | 0 | 0 | 0 | 0 | 2 | 1 | 2 |
|  | 11 | 1 | 1 | 0 | 0 | 0 | 2 | 1 | 2 |

What classification problems can a perceptron solve?

## Linear Separability



The parameters $w_{1}, w_{2}$, define the line. All input patterns above this line are assigned to class 1, all input patterns below the line to class 0.

## Non-linearly Separable Data

- Since $f(w, x)$ is a linear combination of input variables, decision boundary is linear.
- For nonlinearly separable problems, the perceptron fails because no linear hyperplane can separate the data perfectly.
XOR Data

| $\mathrm{x}_{1}$ | $\mathrm{x}_{2}$ | y |
| :---: | :---: | :---: |
| 0 | 0 | -1 |
| 1 | 0 | 1 |
| 0 | 1 | 1 |
| 1 | 1 | -1 |

- An example of nonlinearly separable data is the XOR

$$
y=x_{1} \oplus x_{2}
$$

 function.

## What can a single Perceptron do?

Example: The exclusive OR (XOR) defines a classification task which is not linearly separable.


## Adding one more layer

The exclusive OR (XOR) can be solved adding one more layer


## Multilayer Neural Network

## Multilayer Neural Network

- Hidden Layers: intermediary layers between input and output layers.
- Different type of activation functions (sigmoid, linear, hyperbolic tangent, etc.).
- Multi-layer neural network can solve any type of classification task involving nonlinear decision surfaces.

- Perceptron is single layer.
- We can think to each hidden node as a perceptron that tries to construct one hyperplane, while the output node combines the results to return the decision boundary.



## General Structure of ANN



- The neurons perform a linear transformation on this input using the weights and biases.
- An activation function is applied to it
- The output moves to the next hidden layer


## Artificial Neural Networks (ANN)

- Various types of neural network topology
- single-layered network (perceptron) versus multi-layered network
- Feed-forward: connections only between nodes of level $\mathrm{L}_{\mathrm{i}}$ and the next one $\mathrm{L}_{\mathrm{i}+1}$
- Recurrent network: feedback connections in which outputs of the model are fed back into itself.


Recurrent Neural Network


Feed-Forward Neural Network

- Various types of activation functions (f)

$$
Y=f\left(\sum_{i} w_{i} X_{i}\right)
$$

## Activation function: Sigmoid

- Logistic function (Sigmoid): values in $[0,1]$

$$
f(n e t)=\sigma(n e t)=\frac{1}{1+e^{-n e t}}
$$



- Derivative:

$$
\sigma^{\prime}(n e t)=\frac{\partial}{\partial n e t}\left(\frac{1}{1+e^{-n e t}}\right)=\frac{e^{-n e t}}{\left(1+e^{-n e t}\right)^{2}}=\sigma(n e t)(1-\sigma(n e t))
$$

## Activation function: Hyperbolic Tangent

- Hyperbolic Tangent has values in [-1,1]
- More complex wrt Sigmoid
- Better than sigmoid because symmetric wrt 0 and leads to a faster convergence

- We can obtain it by Sigmoid

$$
f(n e t)=\tau(n e t)=2 \sigma(2 \cdot n e t)-1
$$

- Derivative: $\tau^{\prime}(n e t)=1-\tau(n e t)^{2}$


## Activation function: RELU

- Rectified Linear Unit: values in [0,inf]

$$
f(\text { net })=\left\{\begin{array}{cc}
0(\text { or } \epsilon) & \text { for net }<0 \\
\text { net } & \text { for net } \geq 0
\end{array}\right.
$$



- It will output the input directly if it is positive, otherwise, it will output zero
- Overcomes the vanishing gradient problem, allowing models to learn faster and perform better.


## Activation function: Softmax

- Each value in the output of the softmax function is interpreted as the probability of membership for each class.
- Generalization of the logistic function to multiple dimensions
- Useful for addressing a multiclass problem

$$
z_{k}=f\left(n e t_{k}\right)=\frac{e^{n e t_{k}}}{\sum_{c=1 \ldots s} e^{n e t_{c}}}
$$

## Activation Functions Summary



$$
f(x)=\left\{\begin{array}{l}
0 \text { for } x<0 \\
1 \text { for } x \geq 0
\end{array}\right.
$$



$$
f(x)=\frac{e^{x}-e^{-x}}{e^{x}+e^{-x}}
$$

$$
f(x)=x
$$





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

$$
f\left(x_{j}\right)=\frac{e^{x_{j}}}{\sum_{k} e^{x_{k}}}
$$

Softmax Function

## Training Multilayer NN



## Training Multilayer NN



How do we update these weights given the loss is available only at the output unit?


## Error Backpropagation



## Forward propagation



## Forward propagation



## Learning Multi-layer Neural Network

- Can we apply perceptron learning to each node, including hidden nodes?
- Perceptron computes error $e=y-f(w, x)$ and updates weights accordingly
- Problem: how to determine the true value of $\boldsymbol{y}$ for hidden nodes?
- Approximate error in hidden nodes by error in the output nodes
- Problems:
- Not clear how adjustment in the hidden nodes affect overall error
- No guarantee of convergence to optimal solution


## Gradient Descent for Multilayer NN

- Error function to minimize: $E=\frac{1}{2} \sum_{i=1}^{N}\left(y_{i} f\left(\sum_{j} w_{j} x_{i j}\right)\right)^{2} \longleftarrow \begin{gathered}\text { whadratic function from } \\ \text { which we fan find a global } \\ \text { minimum solution }\end{gathered}$
- Weight update: $w_{j}^{(k+1)}=w_{j}^{(k)}-\lambda \frac{\partial E}{\partial w_{j}}$
- The second term states:
- adjust the weights based on the gradient descent technique, i.e. proportionally to the gradient of the error function
- that weight should be increased in a direction reducing the overall error term
- The gradient descent learning rule moves a small $(\lambda)$ step in the negative gradient direction
- Gradient indicates the direction of growing of the function
- Activation function $f$ must be differentiable
- Error function is nonlinear GD method can get trapped in a local minimum


## Gradient Descent for Multilayer NN

- Weights are updated in the opposite direction of the gradient of the loss function
- Gradient direction is the direction of uphill of the error function
- By taking the negative we are going downhill
- Hopefully to a minimum of the error

$$
w_{j}^{(k+1)}=w_{j}^{(k)}-\lambda \frac{\partial E}{\partial w_{j}}
$$



## Gradient Descent for Multilayer NN

- For output neurons, weight update formula is the same as before (gradient descent for perceptron)
- For hidden neurons:

$$
\begin{aligned}
& w_{p i}^{(k+1)}=w_{p i}^{(k)}+\lambda o_{i}\left(1-o_{i}\right) \sum_{j \in \Phi_{i}} \delta_{j} w_{i j} \\
& \quad \text { Output neurons: } \delta_{j}=o_{j}\left(1-o_{j}\right)\left(y_{i}-o_{j}\right) \\
& \text { Hidden neurons: } \delta_{j}=o_{j}\left(1-o_{j}\right) \sum_{k \in \Phi_{j}} \delta_{k} w_{j k}
\end{aligned}
$$



Consider the Sigmoid $\sigma(x)$ the derivative is $\sigma(x)(1-\boldsymbol{o}(x))$

## Backpropagation in other words

- In order to get the loss of a node we multiply the value of its corresponding $f^{\prime}(z)$ by the loss of the node it is connected to in the next layer (delta_1), by the weight of the link connecting both nodes
- We do the delta calculation step at every unit, back-propagating the loss into the neural net, and finding out what loss every node/unit is responsible for.



## Backpropagation


$w_{(x 1) 1}^{\prime}=w_{(x 1) 1}+\eta \cdot \delta_{1} \cdot x_{1}$


## Learning a MLNN

Inizialize $n_{H}, \mathbf{w}, \lambda$, Nun_epoch, size ${ }_{\mathrm{mb}}$
epoch $=0$
do epoch $=$ epoch +1
random sort the Training Set ( $n$ istances)
for each mini-batch B of size $_{\mathrm{mb}}$
reset gradiente
for each $\mathbf{x}$ in B
forward step
backward step
update weights
Compute Loss
Compute accuracy on Training Set and Validation Set
while (not convergence and epoch < Num_epoch)

## Backpropagation: example

- All the weights are initialized to 0.5
- The learning rate is 1
- Activation function: $Z_{j}=\mathrm{f}\left(\right.$ net $\left._{j}\right)=\frac{1}{1+e^{- \text {net }_{j}}}$
- $f^{\prime}(n e t)=Z_{j}\left(1-Z_{j}\right)$
- $\Delta w_{i j}=\lambda \delta_{j} Z_{j}$
- Output node: $\delta_{j}=\left(T_{j}-Z_{j}\right) f^{\prime}(n e t)$
- Hidden node: $\delta_{j}=\sum_{k}\left(\delta_{k} w_{j k}\right) f^{\prime}(n e t)$



## Backpropagation: example

- net $_{4}=0.9 * 0.5+0.6 * 0.5+1 * 0.5=1.25$
- net $_{5}=0.9 * 0.5+0.6 * 0.5+1 * 0.5=1.25$
- $z_{4}=1 /\left(1+e^{-1.25}\right)=0.77$
- $z_{5}=1 /\left(1+e^{-1.25}\right)=0.77$
- net $_{7}=0.77 * 0.5+0.77 * 0.5+1 * 0.5=1.27$
- $z_{7}=1 /\left(1+e^{-1.27}\right)=0.78$
- $Z_{j}=\mathrm{f}\left(\right.$ net $\left._{j}\right)=\frac{1}{1+e^{- \text {net }_{j}}}$
- $f^{\prime}(n e t)=Z_{j}\left(1-Z_{j}\right)$
- $\Delta w_{i j}=\lambda \delta_{j} Z_{j}$
- Output node: $\delta_{j}=\left(T_{j}-Z_{j}\right) f^{\prime}(n e t)$
- Hidden node: $\delta_{j}=\sum_{k}\left(\delta_{k} w_{j k}\right) f^{\prime}(n e t)$



## Backpropagation: example

- net $_{4}=1.25, n e t_{5}=1.25, z_{4}=0.77, z_{5}=0.77$,
- $f^{\prime}(n e t)=Z_{j}\left(1-Z_{j}\right)$
$n e t_{7}=1.27, z_{7}=0.78$
- $\Delta w_{i j}=\lambda \delta_{j} Z_{j}$
- Output node: $\delta_{j}=\left(T_{j}-Z_{j}\right) f^{\prime}(n e t)$
- Hidden node: $\delta_{j}=\sum_{k}\left(\delta_{k} w_{j k}\right) f^{\prime}(n e t)$
- $\delta_{7}=(0-0.78) * 0.78 *(1-0.78)=-0.13$
- $\delta_{4}=(-0.13 * 0.5) * 0.77 *(1-0.77)=-0.11$
- $\delta_{5}=(-0.13 * 0.5) * 0.77 *(1-0.77)=-0.11$



## Backpropagation: example

- $Z_{j}=\mathrm{f}\left(\right.$ net $\left._{j}\right)=\frac{1}{1+e^{- \text {net }_{j}}}$
- $n e t_{4}=1.25$, net $_{5}=1.25, z_{4}=0.77$,
- $f^{\prime}(n e t)=Z_{j}\left(1-Z_{j}\right)$
$z_{5}=0.77$, net $_{7}=1.27, z_{7}=0.78$
- $\delta_{7}=-0.13$
- Output node: $\delta_{j}=\left(T_{j}-Z_{j}\right) f^{\prime}(n e t)$
- $\delta_{4}=-0.11$
- $\delta_{5}=-0.11$
- $w_{47}=0.5+(1 *-0.13 * 0.77)=0.39$
- $w_{57}=0.5+(1 *-0.13 * 0.77)=0.39$
- $w_{67}=0.5+(1 *-0.13 * 1)=0.36$
- Hidden node: $\delta_{j}=\sum_{k}\left(\delta_{k} w_{j k}\right) f^{\prime}(n e t)$



## Backpropagation: example

- $Z_{j}=\mathrm{f}\left(\right.$ net $\left._{j}\right)=\frac{1}{1+e^{- \text {net }_{j}}}$
- $\delta_{7}=-0.13, \delta_{4}=-0.11, \delta_{5}=-0.11$
- $w_{47}=0.39, w_{57}=0.39, w_{67}=0.36$
- $w_{14}=0.5+(1 *-0.11 * 0.9)=0.48$
- $w_{15}=0.5+(1 *-0.11 * 0.9)=0.48$
- $w_{24}=0.5+(1 *-0.11 * 0.6)=0.49$
- $w_{25}=0.5+(1 *-0.11 * 0.6)=0.49$
- $w_{34}=0.5+(1 *-0.11 * 1)=0.48$
- $w_{35}=0.5+(1 *-0.11 * 1)=0.48$
- $f^{\prime}(n e t)=Z_{j}\left(1-Z_{j}\right)$
- $\Delta w_{i j}=\lambda \delta_{j} Z_{j}$
- Output node: $\delta_{j}=\left(T_{j}-Z_{j}\right) f^{\prime}(n e t)$
- Hidden node: $\delta_{j}=\sum_{k}\left(\delta_{k} w_{j k}\right) f^{\prime}(n e t)$



## Learning Rate

Too low


Too high


Too large of a learning rate causes drastic updates which lead to divergent behaviors

## On the Key Importance of Error Functions

- The error/loss/cost function reduces all the various good and bad aspects of a possibly complex system down to a single number, a scalar value, which allows candidate solutions to be compared.
- It is important, therefore, that the function faithfully represents our design goals.
- If we choose a poor error function and obtain unsatisfactory results, the fault is ours for badly specifying the goal of the search.


## Design Issues in ANN

- Number of nodes in input layer
- One input node per binary/continuous attribute
- $k$ or $\log _{2} k$ nodes for each categorical attribute with $k$ values
- Number of nodes in output layer
- One output for binary class problem
- $k$ nodes for $k$-class problem
- Number of nodes in hidden layer
- Initial weights and biases


## Characteristics of ANN

- Multilayer ANN are universal approximators but could suffer from overfitting if the network is too large.
- Gradient descent may converge to local minimum.
- Model building can be very time consuming, but testing can be very fast.
- Can handle redundant attributes because weights are automatically learnt.
- Sensitive to noise in training data.
- Difficult to handle missing attributes.


## References

- Artificial Neural Network. Chapter 5.4 and 5.5. Introduction to Data Mining.

