


Neural Networks

Prof. Ankur Sinha

Indian Institute of Management Ahmedabad

Gujarat India

A typical Neuron



Applications


- Speech recognition
- Handwriting recognition
- Driverless Cars
- Products: Google translate, Alexa

Neural Networks


©2016 Fjodor van Veen - asimovinstitute.org

- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probablistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool


Deep Feed Forward (DFF)




Perceptron (P)




Feed Forward (FF)




Radial Basis Network (RBF)




Recurrent Neural Network (RNN)




Long / Short Term Memory (LSTM)




Gated Recurrent Unit (GRU)



Auto Encoder (AE)




Variational AE (VAE)




Denosing AE (DAE)



Sparse AE (SAE)



Markov Chain (MC)




Hopfield Network (HN)




Boltzmann Machine (BM)



Restricted BM (RBM)



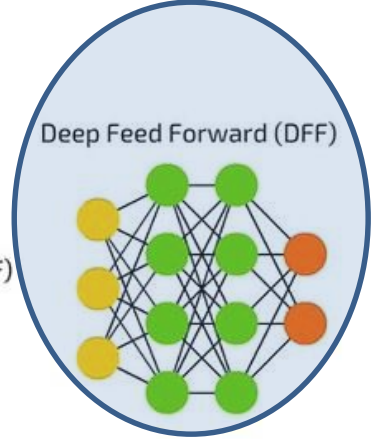
Deep Belief Network (DBN)



A mostly complete chart of


Neural Networks

©2016 Fjodor van Veen - asimovinstitute.org




- Backfed Input Cell
- Input Cell
- △ Noisy Input Cell
- Hidden Cell
- Probablistic Hidden Cell
- △ Spiking Hidden Cell
- Output Cell
- Match Input Output Cell
- Recurrent Cell
- Memory Cell
- △ Different Memory Cell
- Kernel
- Convolution or Pool


Perceptron (P)




Feed Forward (FF)




Radial Basis Network (RBF)




Recurrent Neural Network (RNN)




Long / Short Term Memory (LSTM)




Gated Recurrent Unit (GRU)



Auto Encoder (AE)




Variational AE (VAE)




Denosing AE (DAE)




Sparse AE (SAE)



Markov Chain (MC)



Hopfield Network (HN)




Boltzmann Machine (BM)




Restricted BM (RBM)



Deep Belief Network (DBN)



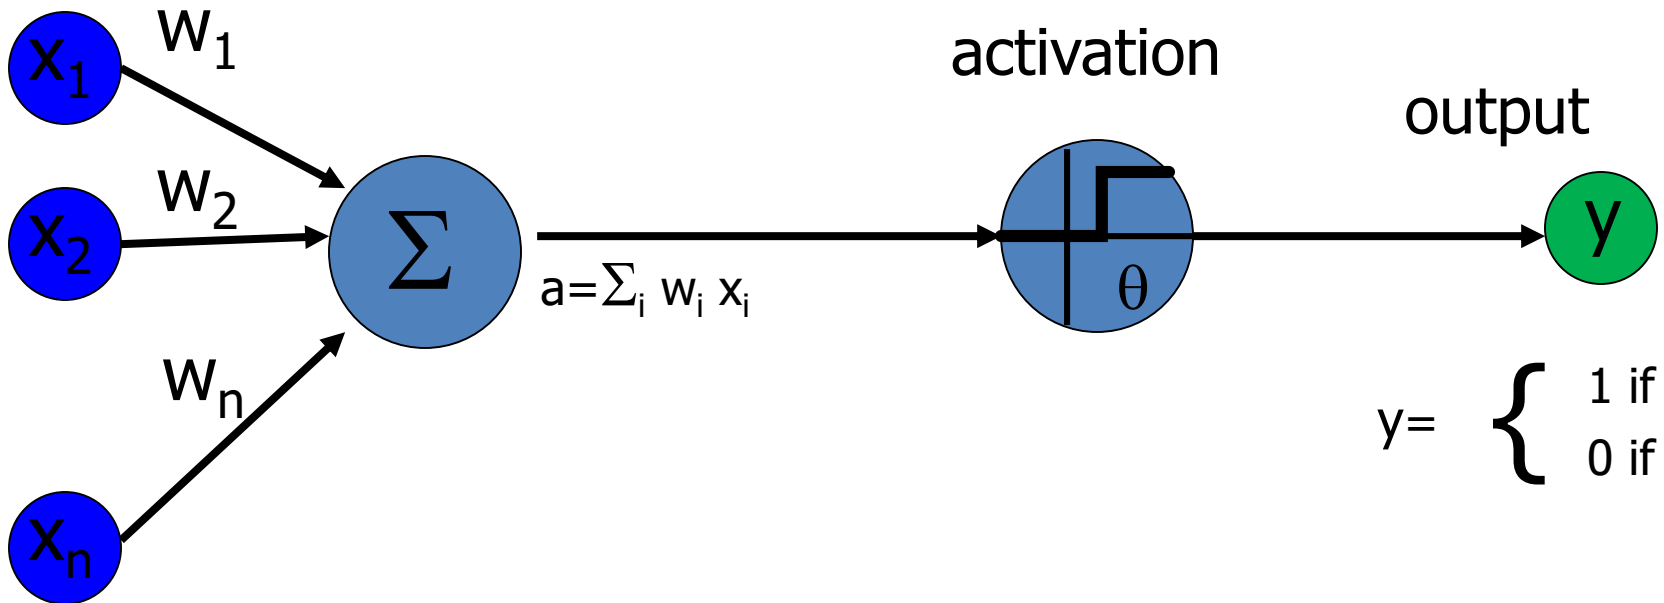
MLP Architecture




A Simple Architecture

A Threshold Logic Unit

inputs



Decision Surface of a TLU



Decision line




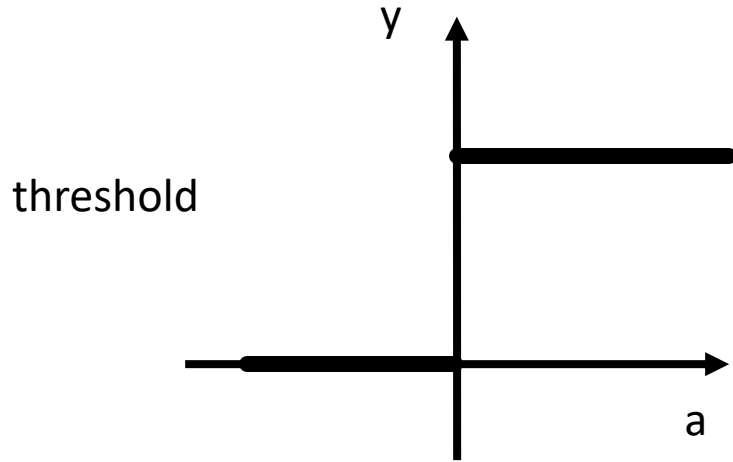
$$w_1 x_1 + w_2 x_2 = \theta$$

A TLU works as a linear classifier


Similar to SVM?

How do you identify the weights and threshold?

Types of Activation Functions



Types of Activation Functions




Training Neural Network

- A training set S of examples $\{\mathbf{x}, t\}$ is required
 - \mathbf{x} is an input vector
 - t is the desired target vector
- Finding acceptable values of w and θ
 - Assume some values for w and θ
 - For the training example \mathbf{x} , compute the network output y
 - Compare output y with targets t , a difference denotes error
 - Adjust w and θ so that the error can be reduced
 - Accept w and θ that leads to minimum error

A Linear Unit

inputs




Tries to give the best linear relationship between input and output
Similar to regression?

Neuron with Sigmoid Function

A Threshold Logic Unit


inputs



$$y = \sigma(a) = 1 / (1 + e^{-a})$$

Gradient descent rules are used to learn the parameters of the NN

Multiple Layers



Backpropagation approach is used to train the neural network

More about NN Parameters


- The weights of the neural network are determined by training data
- As more training data is obtained the weights should be updated

Neural Networks are Universal

- Any boolean function can be learnt by a neural network with single hidden layer
 - It might require a large number of hidden units
- Any mathematical function that is continuous and bounded can be approximated to an arbitrarily small accuracy using a neural network with one hidden layer
 - A large number of hidden units might be required if the error of approximation is very small


Be Careful!

- Neural network can easily lead to overfitting
- Try to minimize the generalization error than the training error



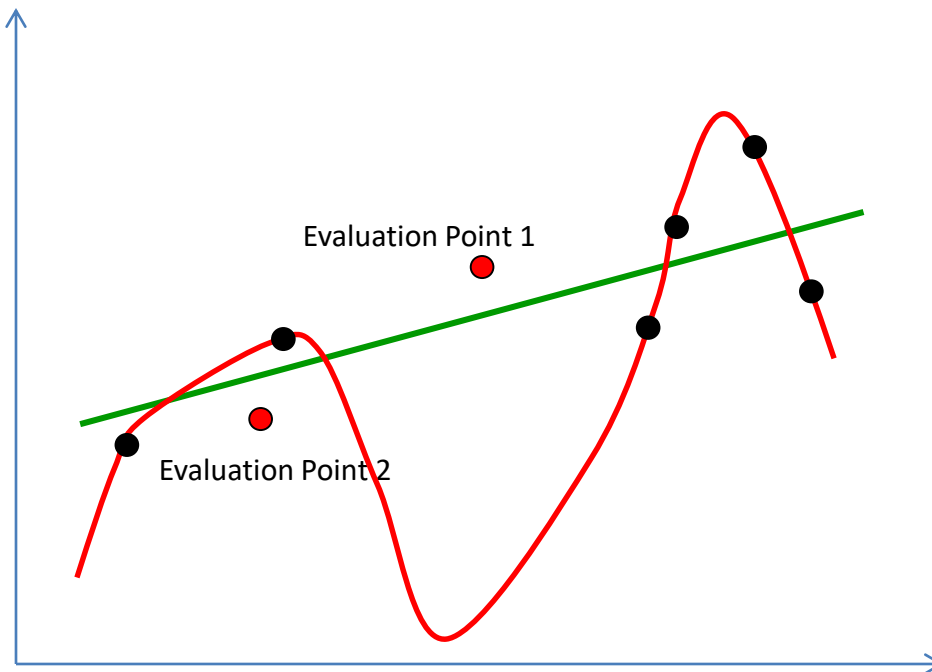
Be Careful!

- Neural network can easily lead to overfitting
- Try to minimize the generalization error than the training error




Be Careful!

- Neural network can easily lead to overfitting
- Try to minimize the generalization error than the training error



MNIST Database




The MNIST database contains 60,000 training images and 10,000 testing images.

MNIST Database



28x28
pixels

784 Pixels



Each of the output nodes fires a 0 or 1
(or the probability)