

Neural Networks

Part 1


Yingyu Liang

`yliang@cs.wisc.edu`

Computer Sciences Department
University of Wisconsin, Madison


Motivation I: learning features

- Example task



Indoor


Experience/Data:
images with labels



outdoor

Motivation I: learning features

- Featured **designed** for the example task




Motivation I: learning features

- More complicated tasks: hard to design
- Would like to **learn** features




Motivation II: neuroscience

- Inspirations from human brains
- Networks of **simple** and **homogenous** units



Motivation II: neuroscience

- Human brain: 100, 000, 000, 000 neurons
- Each neuron receives input from 1,000 others
- Impulses arrive simultaneously
- Added together*
 - an impulse can either increase or decrease the possibility of nerve pulse firing
- If sufficiently strong, a nerve pulse is generated
- The pulse forms the input to other neurons.
- The interface of two neurons is called a synapse



Successful applications

- Computer vision: object location

Microsoft Research

person : 0.998
person : 0.987
person : 0.947
person : 0.946
person : 0.935
chair : 0.677
chair : 0.985
dining table : 0.879
cake : 0.645
wine glass : 0.997
book : 0.830
wine glass : 0.982
knife : 0.685
knife : 0.997

Our results on COCO – too many objects, let's check carefully!

*the original image is from the COCO dataset

ICCV15 International Conference on Computer Vision

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. "Deep Residual Learning for Image Recognition". arXiv 2015.
Shaoqing Ren, Kaiming He, Ross Girshick, & Jian Sun. "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks". NIPS 2015.

Slides from Kaimin He, MSRA

Successful applications

- NLP: Question & Answer

I: Jane went to the hallway.
I: Mary walked to the bathroom.
I: Sandra went to the garden.
I: Daniel went back to the garden.
I: Sandra took the milk there.
Q: Where is the milk?
A: garden

Figures from the paper “Ask Me Anything: Dynamic Memory Networks for Natural Language Processing”, by Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Richard Socher

Successful applications

- Game: AlphaGo




Outline

- A single neuron
 - Linear perceptron
 - Non-linear perceptron
 - Learning of a single perceptron
 - The power of a single perceptron
- Neural network: a network of neurons
 - Layers, hidden units
 - Learning of neural network: backpropagation
 - The power of neural network
 - Issues
- Everything revolves around **gradient descent**

Linear perceptron

- Perceptron = a math model for a single neuron
- Input: x_1, \dots, x_D (signal from other neurons)
- Weights: w_1, \dots, w_D (dendrites, can be negative)
- We sneak in a constant (bias term) $x_0 = 1$, with some weight w_0
- Activation function: linear (for the time being)
$$a = w_0 + w_1 * x_1 + \dots + w_D * x_D$$
- This is the output of a linear perceptron



Learning in linear perceptron

- Training data $\{(X_1, y_1), \dots, (X_N, y_N)\}$
- X_1 is a vector: (x_{11}, \dots, x_{1D}) , so are $X_2 \dots X_N$
- y_1 is a real-valued output, so are $y_2 \dots y_N$
- **Goal:** learn the weights w_0, \dots, w_D , so that given input X_i , the output of the perceptron a_i is close to y_i
- Define “close”:
$$E = \frac{1}{2} \sum_i (a_i - y_i)^2$$
- E is the “error”. Given the training set, it is a function of w_0, \dots, w_D .
- Minimize E : unconstrained optimization with variables w_0, \dots, w_D . **Exactly linear regression.**

Learning in linear perceptron

- Gradient descent: $W \leftarrow W - \alpha \nabla E(W)$
- α is a small constant, “learning rate” = step size
- The gradient descent rule:


$$E(W) = \frac{1}{2} \sum_i (a_i - y_i)^2$$

$$\partial E / \partial w_d = \sum_i (a_i - y_i) x_{id}$$

$$w_d \leftarrow w_d - \alpha \sum_i (a_i - y_i) x_{id}$$

- Repeat until E converges.
- E is convex in W : there is a unique global minimum

Visualization of gradient descent




The (limited) power of linear perceptron

- Linear perceptron is just

$$a = WX$$

- where X is the input vector, augmented by $x_0 = 1$
- It can represent any linear function in $D + 1$ dimensional space... but that's it
- In particular, it won't be a nice fit to binary classification ($y = 0$ or $y = 1$)




Non-linear perceptron

- Change the activation function: use a **step function**

$$a = g(w_0 + w_1 * x_1 + \dots + w_D * x_D)$$

- $g(h) = 0$, if $h < 0$; $g(h) = 1$ if $h \geq 0$




- Can you see how to make logic AND, OR, NOT with such a perceptron?

Non-linear perceptron

- Change the activation function: use a **step function**

$$a = g(w_0 + w_1 * x_1 + \dots + w_D * x_D)$$

- $g(h) = 0$, if $h < 0$; $g(h) = 1$ if $h \geq 0$



- AND: $w_1 = w_2 = 1, w_0 = -1.5$
- OR: $w_1 = w_2 = 1, w_0 = -0.5$
- NOT: $w_1 = -1, w_0 = 0.5$


Now we see the reason for bias terms

Non-linear perceptron for AND

- Change the activation function: use a **step function**


$$a = g(w_0 + w_1 * x_1 + \dots + w_D * x_D)$$

- $g(h) = 0$, if $h < 0$; $g(h) = 1$ if $h \geq 0$



Example Question


- Will you go to the festival?
- Go only if at least two conditions are favorable



All inputs are binary; 1 is favorable

Example Question

- Will you go to the festival?
- Go only if Weather is favorable and at least one of the other two conditions is favorable




All inputs are binary; 1 is favorable

Sigmoid activation function: Our second non-linear perceptron

- The problem with LTU: step function is discontinuous, cannot use gradient descent
- Change the activation function (again): use a **sigmoid function**

$$g(x) = 1 / (1 + \exp(-x))$$




Sigmoid activation function: Our second non-linear perceptron

- The problem with LTU: step function is discontinuous, cannot use gradient descent
- Change the activation function (again): use a sigmoid function

$$g(x) = 1 / (1 + \exp(-x))$$

- Exercise: $g'(x) = ?$




Sigmoid activation function: Our second non-linear perceptron

- The problem with LTU: step function is discontinuous, cannot use gradient descent
- Change the activation function (again): use a sigmoid function

$$g(x) = 1 / (1 + \exp(-x))$$

- Exercise: $g'(x) = g(x)(1 - g(x))$



Learning in non-linear perceptron

- Again we will minimize the error:

$$E(W) = \frac{1}{2} \sum_i (a_i - y_i)^2$$

- Now $a_i = g(\sum_d w_d * x_{id})$

$$\partial E / \partial w_d = \sum_i (a_i - y_i) a_i (1 - a_i) x_{id}$$

- The sigmoid perceptron update rule

$$w_d \leftarrow w_d - \alpha \sum_i (a_i - y_i) a_i (1 - a_i) x_{id}$$

- α is a small constant, “learning rate” = step size
- Repeat until E converges

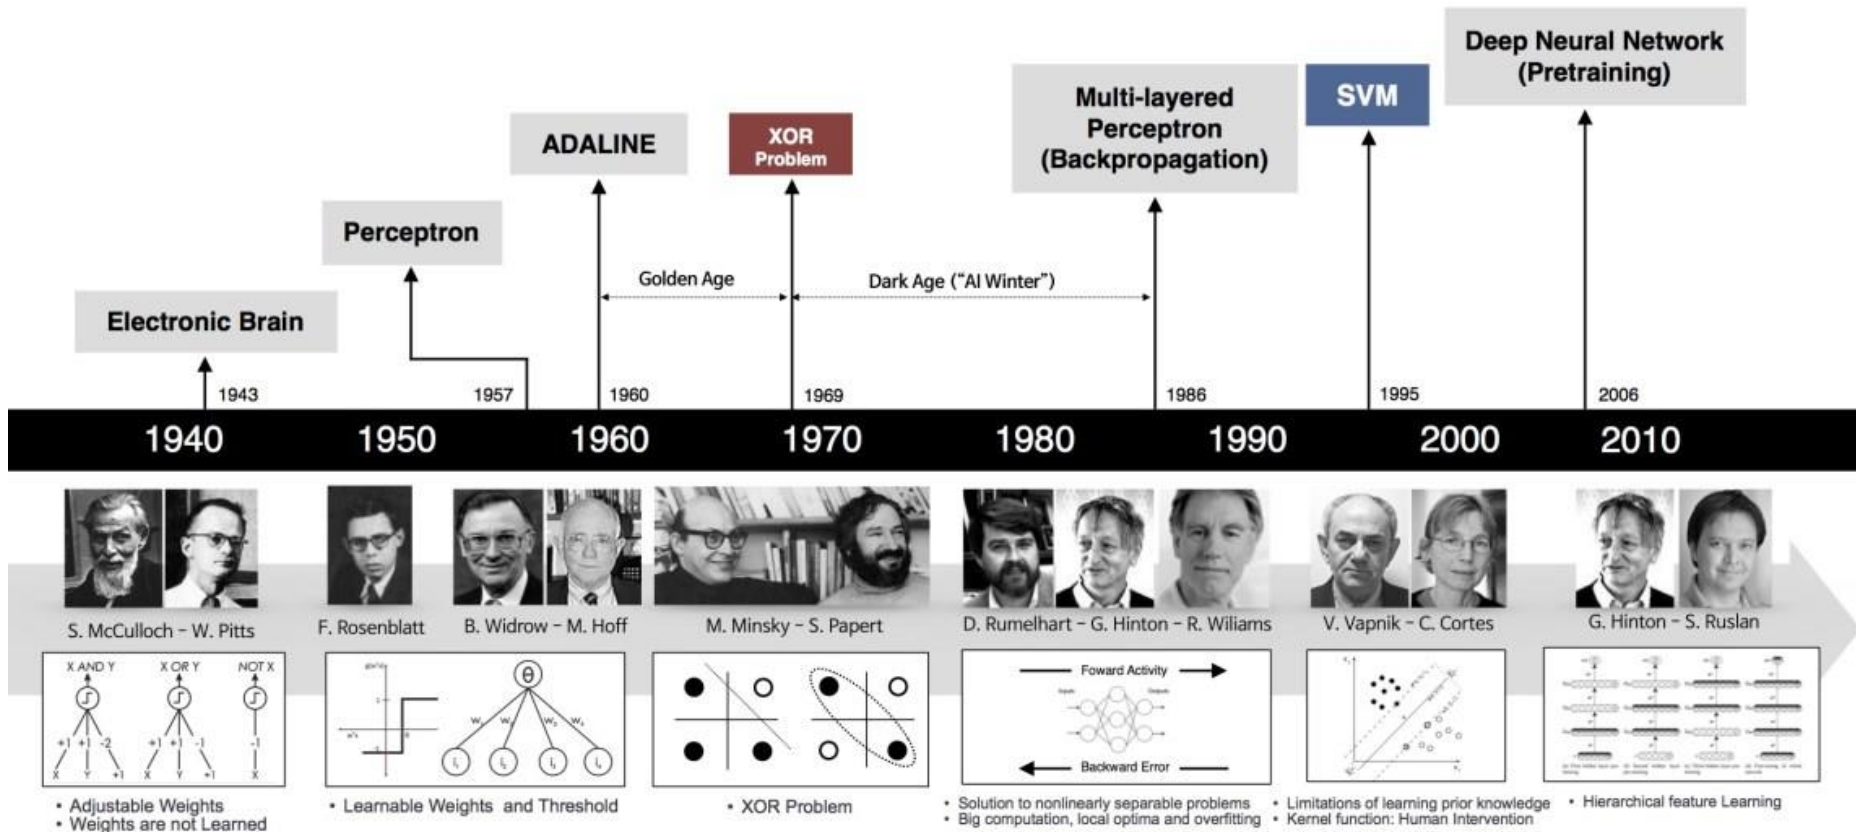
The (limited) power of non-linear perceptron

- Even with a non-linear sigmoid function, the **decision boundary** a perceptron can produce is still **linear**
- AND, OR, NOT revisited
- How about XOR?

The (limited) power of non-linear perceptron


- Even with a non-linear sigmoid function, the **decision boundary** a perceptron can produce is still **linear**
- AND, OR, NOT revisited
- How about XOR?
- This contributed to the first AI winter

Brief history of neural networks




(Multi-layer) neural network

- Given sigmoid perceptrons



- Can you produce output like



- which had non-linear decision boundaries

0

1

0

1

0