


Introduction to Deep Neural Networks

M.Sc. Petr Nejedlý
nejedly@isibrno.cz

Outline

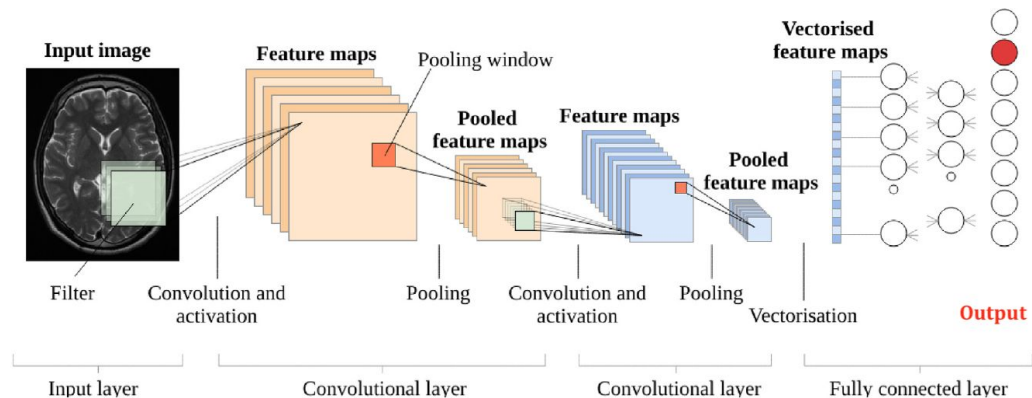
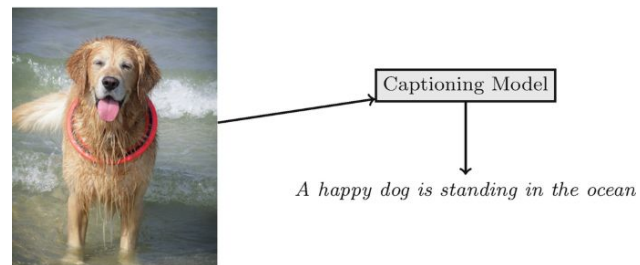

- Example of Deep Learning applications
- Machine Learning vs Deep Learning
- Single Perceptron/Neuron
- Basic Neural Network
- How to train?
- Convolutional Neural Networks
- Recurrent Neural Networks



What can we do with neural networks?

Image


- classification
- localization
- captioning
- segmentation



What we can do with neural networks?

- Self Driving Cars

Vision-based approach



No Lidar

No HD maps




What we can do with neural networks?

- Speaker classification, Speech2Text translation, Voice assistant (Siri, Alexa)



- Language translation (Google Translate)
- Sentiment analysis



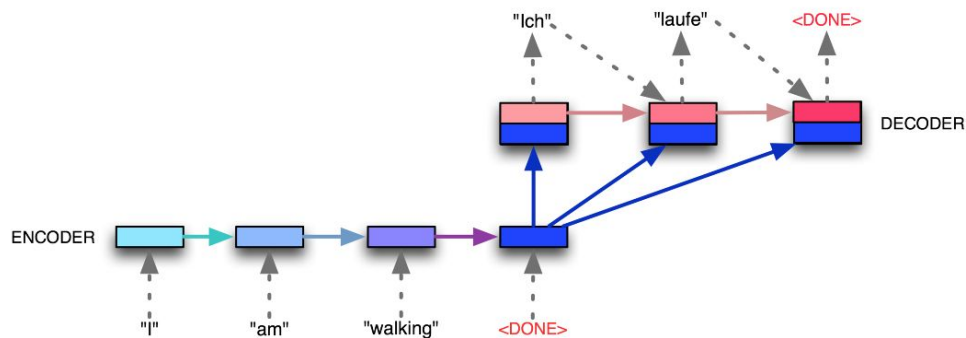
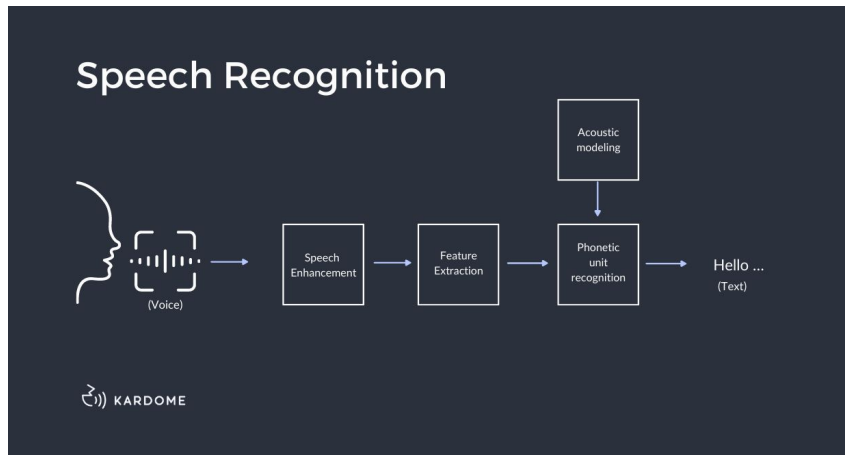
"I am happy with this water bottle."



"This is a bad investment."


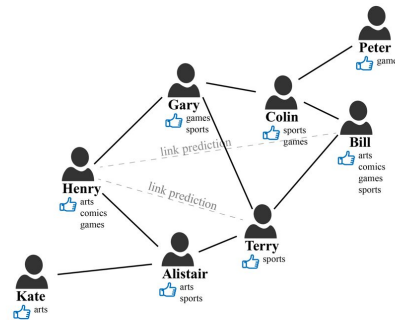


"I am going to walk today."




What we can do with neural networks?


- Drug discovery
- Social network analysis
- Movie recommender system (e.g. Netflix)
- And many others ...



Machine Learning vs Deep Learning

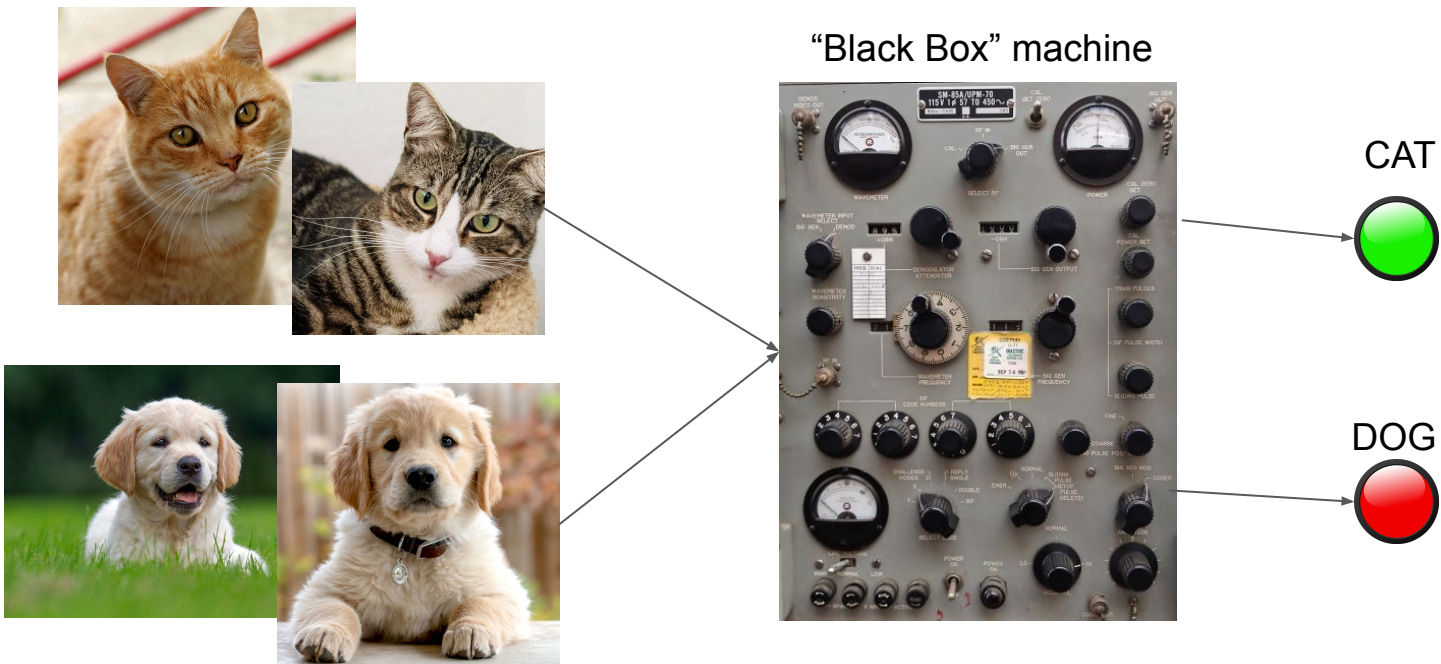


- requires manual feature engineering
 - + small datasets posible
-
- + NO manual feature engineering
 - requires large datasets !!!



Supervised learning


- Manual feature engineering is difficult! For example, how can you describe dog or cat algorithmically?
- Train a machine by showing examples instead of programming it
- When the output is wrong, tweak the parameters of the machine




How should we design the "BlackBox" machine?

Can we design the “Black Box” as a human brain?

- First neural networks were originally designed to approximate human brain/neurons (Perceptron, Rosenblatt 1957)




- More details in Neuroscience vs AI presentation



Basic Block of NN: Perceptron

Perceptron = simple mathematical model of biological neuron

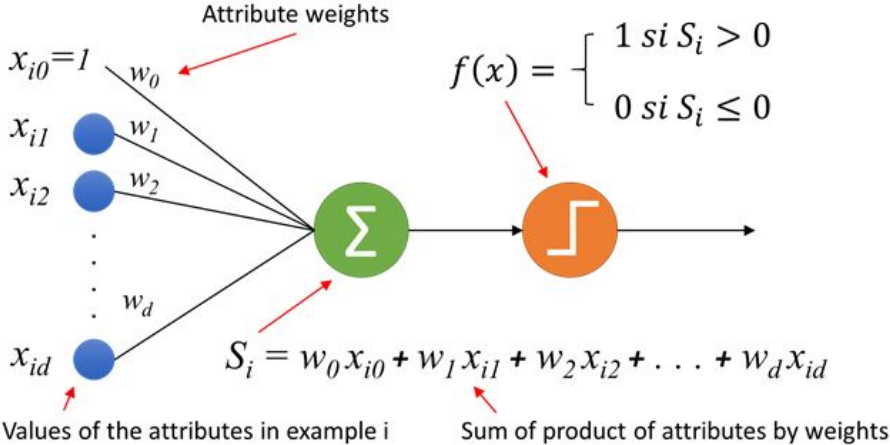


NEURON

$$\hat{y} = g \left(w_0 + \sum_{i=1}^m x_i w_i \right)$$

Labels for the equation:



- Output: \hat{y}
- Linear combination of inputs: $w_0 + \sum_{i=1}^m x_i w_i$
- Non-linear activation function: g
- Bias: w_0



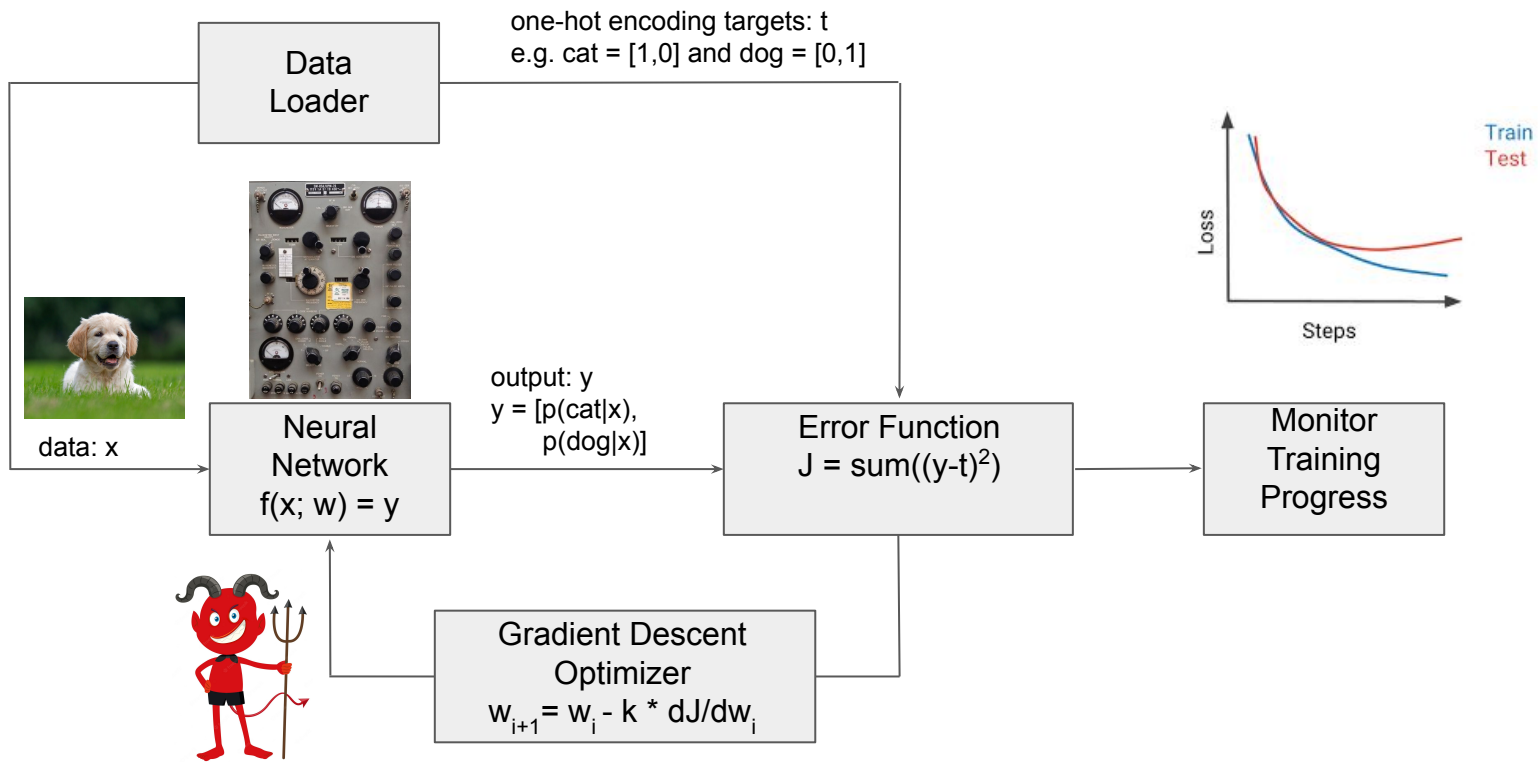
PERCEPTRON

Neural Networks

- Neural network = Collection of neurons organized in multiple layers
- Mathematically speaking, neural network allows finding of approximation of any function $y = f(x)$ that can map inputs (x) to outputs (y)
- In theory, large enough NN can be used to approximate any function => **Universal Approximation Theorem!**



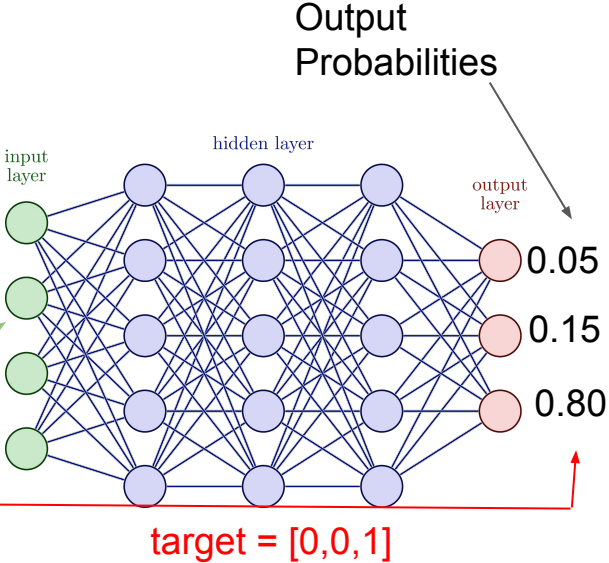
Basic building blocks




Example, shallow NN for Iris classification

Table with features

sepal		petal		class
length	width	length	width	
6.3	2.3	4.4	1.3	versicolor
6.2	3.4	5.4	2.3	virginica
5.2	3.4	1.4	0.2	setosa
6.9	3.1	5.4	2.1	virginica
5.7	4.4	1.5	0.4	setosa
5.4	3.7	1.5	0.2	setosa
5	3.3	1.4	0.2	setosa
6.4	2.8	5.6	2.1	virginica
6	3	4.8	1.8	virginica
5.5	2.5	4	1.3	versicolor



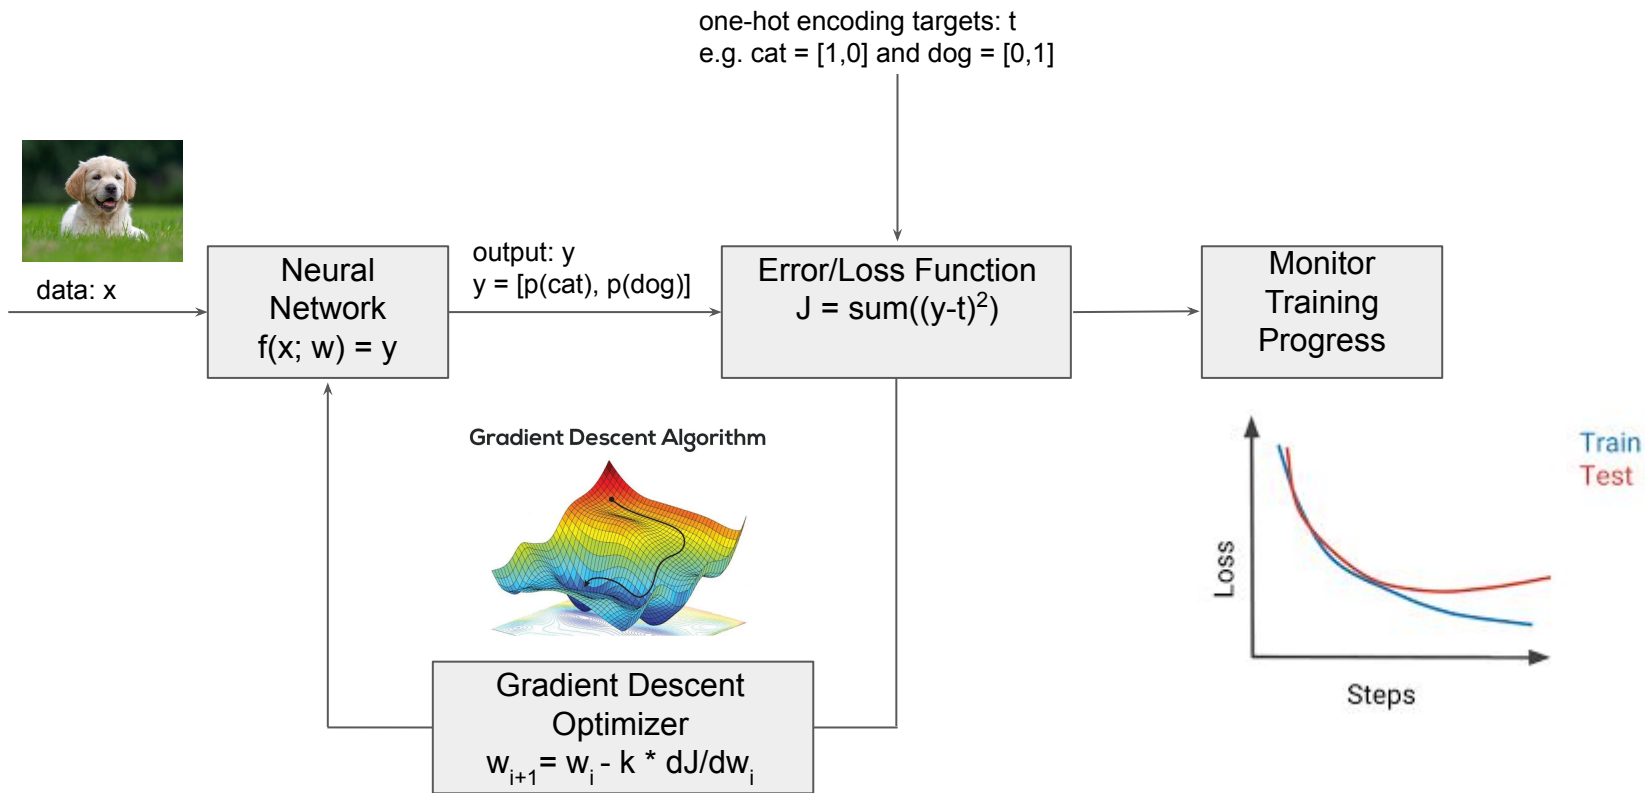
one-hot encoding:
 setosa = [1,0,0]
 versicolor = [0,1,0]
 virginica = [0,0,1]



- Feed line by line to your NN model
-

Training

- NN are trained by Backpropagation and Gradient descent (difficult math needed)
- Basically, iteratively adapting “Black Box” to minimize objective error function



Neural network playground

<https://playground.tensorflow.org/>

DATA


Which dataset do you want to use?



Ratio of training to test data: 50%



Noise: 0



Batch size: 10



REGENERATE

FEATURES

Which properties do you want to feed in?



+ - 2 HIDDEN LAYERS

+ -

4 neurons

+ -

2 neurons

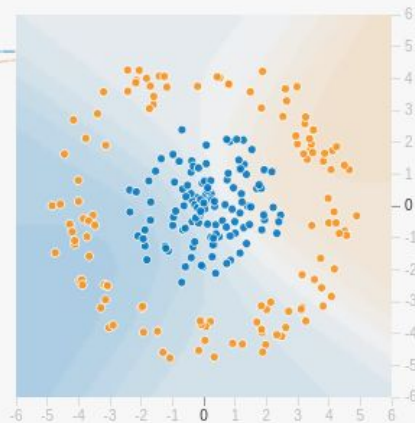
This is the output from one neuron. Hover to see it larger.

The outputs are mixed with varying weights, shown by the thickness of the lines.

OUTPUT

Test loss 0.548

Training loss 0.514



Colors shows data, neuron and weight values.

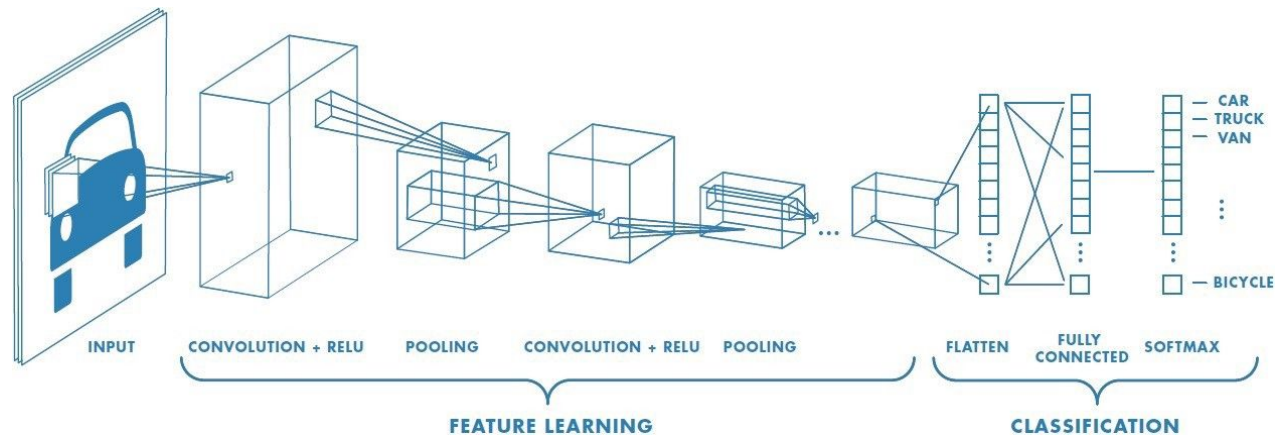
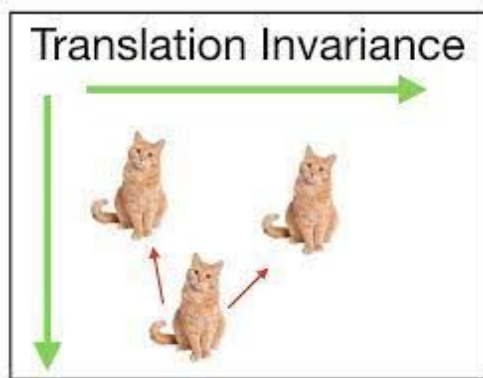


Show test data

Discretize output






Image classification => Convolutional Neural Network

- NN must be invariant for translation, i.e. it shouldn't matter where the cat is in the image
- Convolutions are used to extract translation invariant features



Convolution

- sliding the kernel “K” over the Image “I”
- Basically, multiply corresponding matrix cells and sum the results


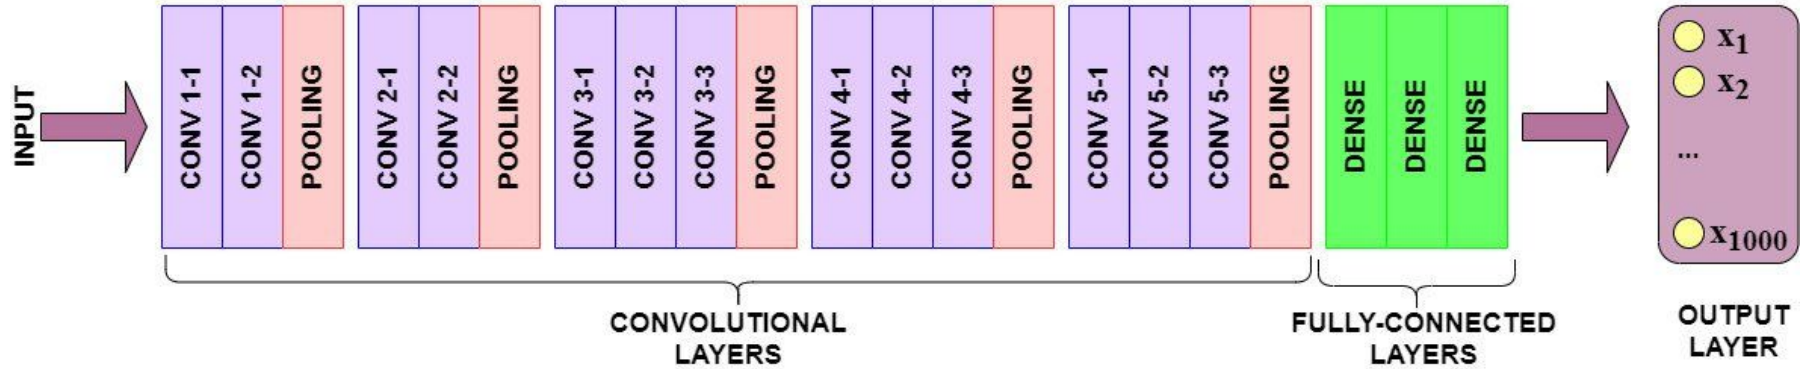


Vertical Edge

$$y[m, n] = x[m, n] * h[m, n] = \sum_{j=-\infty}^{\infty} \sum_{i=-\infty}^{\infty} x[i, j] \cdot h[m - i, n - j]$$

Convolutional Neural Network Architecture

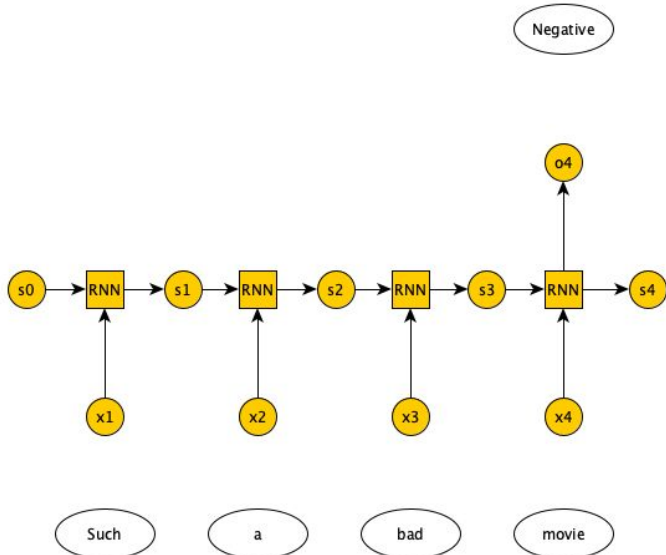

VGG16 MODEL ARCHITECTURE



Convolutional kernels gets more complex with increasing number of convolutional layers

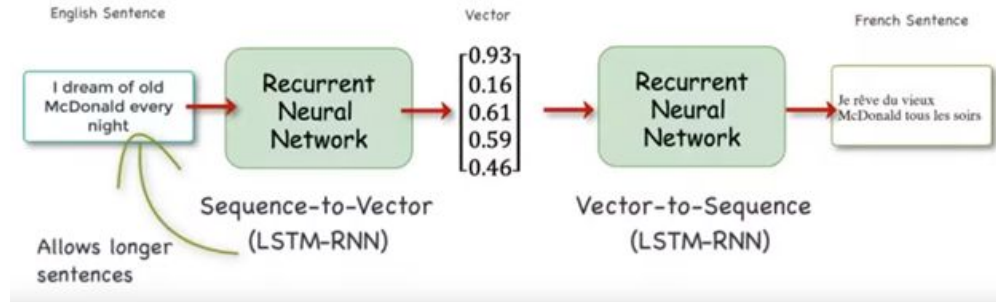
Recurrent Neural Networks

- Neurons have recurrent connections
- Works great for sequence
e.g. sentiment analysis, machine language translation



Language Translation

Encoder-Decoder Architecture



THANK YOU FOR YOUR ATTENTION

ANY QUESTIONS?

makeameme.org