




“GHEORGHE ASACHI” TECHNICAL UNIVERSITY OF IASI

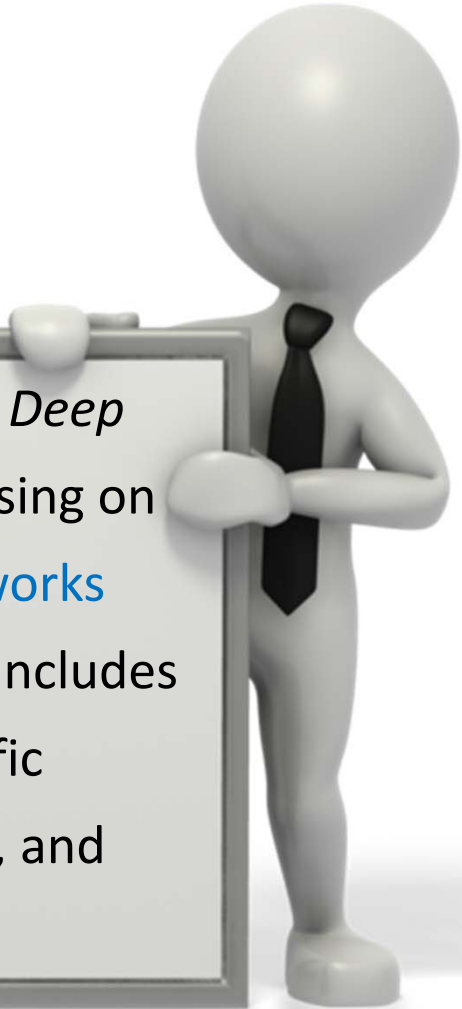
Faculty of Electronics, Telecommunications and Information Technology



CONVOLUTIONAL NEURAL NETWORKS

principles
architectures
applications

IULIAN B. CIOCOIU

A 3D white figure wearing a black tie, standing and holding a large rectangular sign with both hands. The figure is positioned to the right of the sign, facing left. The sign has a grey border and a white background with black text. The figure's shadow is cast on the ground to its right.

A historical perspective on *Deep Learning* approaches, focusing on **Convolutional Neural Networks** (CNN's). The presentation includes operating principles, specific terminology, architectures, and current applications.

CONTENTS



INTRODUCTION



Motivations, principles, terminology, historical notes

FOUNDATIONS



Architectures, algorithms, limitations

APPLICATIONS



Pattern recognition, biomedical, automotive, biometrics, arts

CONCLUSIONS



What's next ?


Introduction

Historical notes

A bit of history:

Gradient-based learning applied to document recognition

[LeCun, Bottou, Bengio, Haffner 1998]



Introduction

Historical notes

A bit of history:
**ImageNet Classification with Deep
Convolutional Neural Networks**
[Krizhevsky, Sutskever, Hinton, 2012]




Figure copyright Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012. Reproduced with permission.


“AlexNet”

Introduction

Historical notes

1998

LeCun et al.



of transistors




10^6

of pixels used in training

10^7 NIST

2012

Krizhevsky et al.



of transistors



10^9

GPUs




of pixels used in training

10^{14} IMAGENET

Figure copyright Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012. Reproduced with permission.


Introduction

Historical notes




Introduction

Historical notes



Introduction

Historical notes



Convolutional Neural Networks (CNN's)

Principles, architectures, terminology

Lecture 5: Convolutional Neural Networks

Fei-Fei Li & Justin Johnson & Serena Yeung

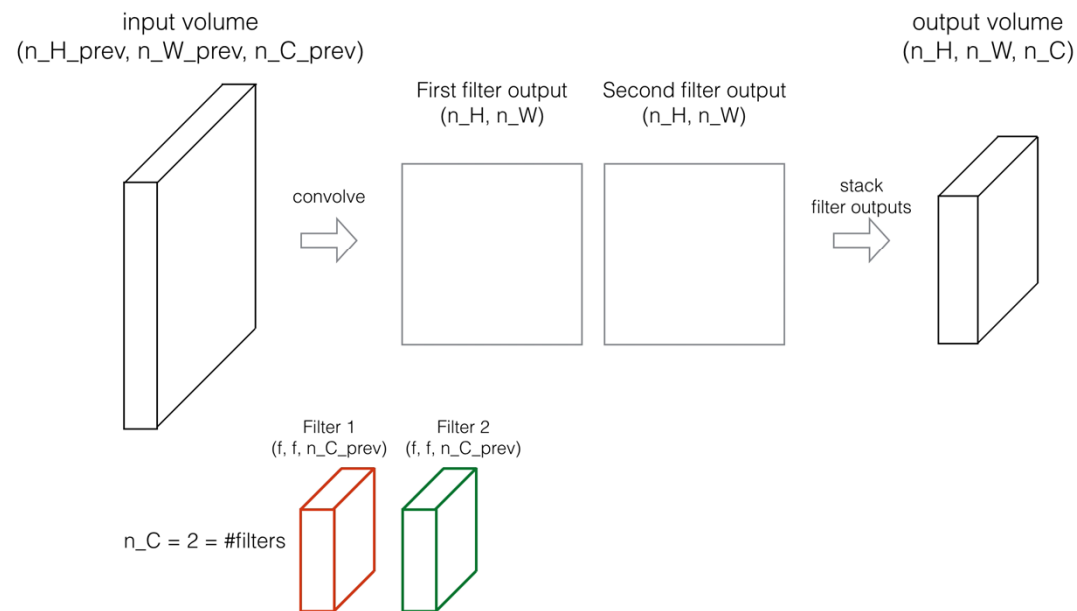
Lecture 5 - 1

April 16, 2019

Convolutional Neural Networks (CNN's)

Principles, architectures, terminology

How do convolutions work?




Convolutional Neural Networks (CNN's)

Principles, architectures, terminology

Last time: Activation Functions


Sigmoid

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




tanh

$$\tanh(x)$$




ReLU

$$\max(0, x)$$



Leaky ReLU

$$\max(0.1x, x)$$




Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

ELU


$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$



Convolutional Neural Networks (CNN's)

Principles, architectures, terminology


Last time: Weight Initialization



Initialization too small:

Activations go to zero, gradients also zero,


No learning =(



Initialization too big:

Activations saturate (for tanh),

Gradients zero, no learning =(



Initialization just right:


Nice distribution of activations at all layers,

Learning proceeds nicely

Convolutional Neural Networks (CNN's)

Principles, architectures, terminology


Last time: Data Preprocessing



Convolutional Neural Networks (CNN's)

Principles, architectures, terminology

First-Order Optimization




Convolutional Neural Networks (CNN's)

Principles, architectures, terminology

Second-Order Optimization

- (1) Use gradient **and Hessian** to form **quadratic** approximation
- (2) Step to the **minima** of the approximation



Convolutional Neural Networks (CNN's)

Principles, architectures, terminology

Today: CNN Architectures

Case Studies

- AlexNet
- VGG
- GoogLeNet
- ResNet

Also....

- SENet
- NiN (Network in Network)
- Wide ResNet
- ResNeXT
- DenseNet
- FractalNet
- MobileNets
- NASNet

Convolutional Neural Networks (CNN's)

Principles, architectures, terminology

Lecture 9: CNN Architectures

It is the most important computing development in the last 20 years, and [every major technology company is] going to have to race to make sure that AI's a core competency.

Jen-Hsun Huang
CEO Nvidia



HOW TO REDUCE TRAINING TIME/DATABASE DIMENSION ?



KEY QUESTIONS ABOUT DEEP LEARNING

Applications

Deep learning frameworks

Caffe



PYTORCH



theano



Applications

Object detection/classification





Image © Electronic Frontier Foundation



ImageNet Large Scale Visual Recognition Challenge

is the biggest object classification competition,
organized since 2010

Nobody phrases it this way, but I think that artificial intelligence is almost a **humanities discipline**. It's really an attempt to understand human intelligence and human cognition.

Sebastian Thrun

Professor, Stanford University