



Webinar #12 Getting started with Deep Learning: A gentle introduction and working example



7 July 2020

The webinar will start at 1pm CEST

Presenter: Juan Durillo Barrionuevo (Leibniz Supercomputing Centre, LRZ)

Moderator: Katya Ahmad (University College London, UCL)



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 675451

The webinar series is run in collaboration with:









Webinar #12 Getting started with Deep Learning: A gentle introduction and working example



7 July 2020

Welcome!



Presenter: Juan Durillo Barrionuevo (Leibniz Supercomputing Centre, LRZ)

Moderator: Katya Ahmad (University College London, UCL)



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 675451

The webinar series is run in collaboration with:





Getting started with Deep Learning: A gentle introduction and working example 07.07.2020 | PD Dr. Juan J. Durillo

Agenda

- Introduction
- Introduction to (Deep) Neural Networks for Machine Learning
- Computer Vision as working example
- Introduction to Convolutional Neural Networks
- Deep Neural Network Architecture
- Data Augmentation
- Hands on session: Implementing a Convolutional Neural Network from scratch



Artificial + Intelligence





Observable Trend





(graph taken from Nvidia learning institute)

Perceptron – Artificial Neuron



Single artificial neurons work well for linearly separable datasets (indeed output is the activation effect on a linear combination of the input)



most popular activation functions



Neural Network



 Works well even when the data is not linearly separable





lrz

(Supervised) Learning

• Data domain Ζ: Χ×Υ

 $X \rightarrow$ domain of the input data

 $\Upsilon \rightarrow$ set of labels (knowledge)



truck, car, horse, bird, boat

- Data Distribution is a probability distribution over a data domain
- Training set $z_1, ..., z_n$ from Z assumed to be drawn from the Data Distribution D
- Validation set $v_1, ..., v_m$ from Z also assumed to be drawn from D
- A machine learning model is a function that given a set of parameters Θ and z from Z produces a prediction
- The prediction quality is measured by a differentiable non-negative scalar-valued loss function, that we denote ℓ(Θ; z)

(Supervised) Learning



- Given Θ we can define the expected loss as: $L(\Theta) = \mathbb{E}_{z \sim D}[\ell(\Theta; z)]$
- Given D, ℓ, and a model with parameter set Θ, we can define learning as:
 "The task of finding parameters Θ that achieve low values of the expected loss, while we are given access to only n training examples"
- The mentioned task before is commonly referred to as *training*
- Empirical average loss given a subset of the training data set $S(z_1, ..., z_n)$ as:

$$\hat{L}(\Theta) = \frac{1}{n} \sum_{t=1}^{n} [\ell(\Theta; z_t)]$$

 Usually a proxy function, easier to understand by humans, is used for describing how well the training is performed (e.g., accuracy)

(Supervised) Learning



- The dominant algorithms for training neural networks are based on mini-batch stochastic gradient descent (SGD)
- Given an initial point Θ_0 SGD attempt to decrease \hat{L} via the sequence of iterates

 $\Theta_t \leftarrow \Theta_{t-1} - n_t g(\Theta_{t-1}; B_t)$

$$g(\Theta; B) = \frac{1}{|B|} \sum_{z \in B} \nabla \ell(\Theta; z)$$

Definitions	B_t : random subset of training examples		
	n_t : positive scalar (learning rate)		
	epoch: update the weights after going over all training set		

Computer Vision



- Why? Focus on a kind of Deep Neural Network called Convolutional Neural Network (CNN)
- CNNs ability to extract multi-scale localized spatial features and compose them to construct highly expressive representations led to breakthroughs in almost all machine learning areas

COMPUTER VISION TASKS

Image Classification



predicting the type or class of an object in an image

Image Classification + Localization



predicting the type or class on an object in an image and draw a bounding box around it

Object Detection



predicting the location of objects in an image via bounding boxes and the classes of the located objects

Image Segmentation rz



predicting the class to which each pixel in the image belongs to

(inspired by a slide used in cs231n lecture from Stanford University)

On Input Representation





_dict=['EOS','a','my','sleeps','on','dog','cat','the','bed','floor'] sentence = ['a', 'dog', 'sleeps', 'on', 'the', 'floor', 'EOS'] [0. 0. 0. 0. 0. 0. 0. 0. 0. 1. 0. 0. 0. 0. Θ. 0. 0. 0. 0. 0. 0. 0. 0. 0. 1. 0. 0. 0. 0. 0. 0. 0.] 0. 0. 0. 1. 0. 0. 0. 0. 0. 0. 1.] 0. 0. 0. 0. 0. [1. 0. 0. 0. 0. 0. 0. 0. 0.]]

language

image

Fully Connected Neural Network

Neural Networks for Image Classification





Training Neural Networks





15 117 184 214 214 155 117 49 19 19 0 0 0 0 0 0 0 16 173 0 0 0 5 90 247 254 0 0 0 8 43 203 254 250 121 0 0 43 146 254 254 229 111 0 4 137 245 254 254 238 40 0 0 107 254 254 254 254 254 255 100 23 55 196 196 158 196 223 254 254 254 216 0 0 0 0 0 18 70 245 253 96 0 0 0 0 0 0 1 197 254 13 is a zero 0 0 0 0 14 200 254 0 0 0 0 0 88 246 253 0 0 0 0 0 0 0 0 67 254 254 170 0 0 0 0 11 20 140 248 254 204 9 254 229 85 79 78 161 176 218 254 254 254 168 6 0 -0 0 is a one 0 6 0 0 0 0 0 0 55 97 157 156 187 0 0 0 5 90 247 254 219 68 0 8 43 203 254 250 121 15 43 146 254 254 229 111 0 0 4 137 245 254 254 238 40 0 0 0 133 254 254 254 254 254 254 235 106 23 0 5 196 196 158 196 223 254 254 216 23 0 0 0 18 70 245 253 99 0 0 0 0 18 70 245 253 99 0 0 0 0 0 1 197 254 134 0 0 0 0 0 14 200 254 134 0 0 0 0 0 88 246 253 99 0 0 0 0 67 254 254 170 0 11 20 140 248 254 204 9 0 79 79 161 176 218 254 254 254 168 6 0 54 254 255 254 254 254 237 95 19 6 0 0 0 0 5 155 156 111 58 58 36 0 0 0 0 0 0 0 0 0 shift to the left is a five 0 0 0 0 37 55 128 254 156 254 229 79 79 86 29 203 254 254 254 255 254 254 254 23 0 0 0 0 0 0 is a nine

Neural Networks for Image Classification

.



Convolutional Neural Networks

No More Feature Engineering: Learning Features From Data









Learning features from data: Convolutions





 $\begin{array}{c}
0 \\
1 \times (-1) + 0 \times 0 + 1 \times 1 + \\
0 \times (-2) + 1 \times 1 + 0 \times 2 + \\
0 \times (-3) + 0 \times 0 + 1 \times 3 = 4
\end{array}$

Convoluted Image



Filter is convoluted with all the pixels of the image

How many units the filter moves horizontally or vertically is called **stride** and can be different in both dimensions

The stride defines the size of the convoluted image

1	-1	0	1	0	1	0	1
0	-2	1	2	1	0	1	0
0	-3	0	3	0	1	0	1
1	0	1	0	0	1	0	0
0	0	0	0	1	0	1	0
0	0	1	0	0	1	1	1
0	0	0	0	0	0	1	0
0	0	1	0	0	1	0	1

1	0	1	0	0	1	0	1
0	1	0	0	1	0	1	0
0	-1	0	1	0	1	0	1
1	-2	1	2	0	1	0	0
0	-3	0	3	1	0	1	0
0	0	1	0	0	1	1	1
0	0	0	0	0	0	1	0
0	0	1	0	0	1	0	1

1	0	1	0	0	1	0	1
0	1	0	0	1	0	1	0
0	0	1	0	0	1	0	1
1	0	1	0	0	1	0	0
0	0	0	0	1	0	1	0
0	0	1	0	0	-1	0	1
0	0	0	0	0	-2	1	2
0	0	1	0	0	-3	0	3



*The London skyline image is designed by Freepik

Convolutional Neural Networks (CNN)





A pooling layer down sample the feature maps produced by a convolution into smaller number of parameters to reduce the computational complexity.

It is a common practice to add pooling layers after each one or two convolutions layers in the CNN architecture.

Differences in the Workflow and Challenges





Deep Neural Networks have shown ability to outperform humans in many tasks, but with a high price tag:

- Computational cost of training networks, compound by identifying an optimal architecture
- Significant volumes of data, which might not always be available

CNN Architecture: A Common Pattern and its Influence





The execution time required during a forward pass through a neural network is bounded from below by the number of floating point operations (FLOPs).

This FLOP count depends on the deep neural network architecture and the amount of data.

LeNet Architecture





Architecture summary :

• 3 convolutional layers filters in all the layers equal to 5x5

(layer 1 depth = 6, layer 2 depth = 16, layer 3 depth = 120)

As activation function the tanh function is used

AlexNet and VGG Architectures





17

parameters



GoogleNet



- What is the best kernel size for each layer?
- Concatenating filters instead of stacking them for reducing computational expenses







Increasing complexity

2015 - Microsoft ResNet

Superhuman Image Recognition





2016 - Baidu Deep Speech 2 Superhuman Voice Recognition

2017 - Google Neural Machine Translation Near Human Language Translation

How much data?

- A non desired issue with NN is known as overfitting
 - high accuracy in the training set
 - poor generalization (low accuracy in the validation and test sets)
- Overfitting is possibly more noticeable in DNNs due to the large number of parameters
- Getting more data not always viable
- Data augmentation
 - Field dependent





texture-out



grey scale

edge enhanced







flip/rotate





Hands on: A quick code review of the LeNet Implementation using Keras and Tensorflow

Summary



- Brief introduction to Deep Learning with emphasis in Deep Convolutional Neural Networks
- Review of basic concepts: from perceptron to the learning task
- Debrief of most important concepts of neural network architectures
- Introduction to data augmentation
- Code review





To pose a question, you can write your question in the "Questions" tab



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 675451

The webinar series is run in collaboration with:





Thank you for participating!

...don't forget to fill in our feedback questionnaire...

Visit the CompBioMed website (<u>www.compbiomed.eu/training</u>) for a full recording of this and other webinars, to download the slides and to keep updated on our upcoming trainings



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 675451

The webinar series is run in collaboration with:



Webinar series