



University of Pavia

Ph.D. School of Electrical and Electronics Engineering and Computer Science

Deep Learning and TensorFlow

A short course for PhD students

May 4 - June 14, 2018

Instructor:

Marco Piastra, PhD, marco.piastra@unipv.it

Co-instructors:

Mirto Musci, PhD, mirto.musci@unipv.it

Andrea Pedrini, PhD, andrea.pedrini@unipv.it

Gianluca Gerard, gianluca.gerard01@universitadipavia.it

Deep Learning techniques are at the heart of the recent revolution in Artificial Intelligence that has opened many new perspectives for engineering applications. On a closer look, nonetheless, such techniques appear to be more evolutionary - in terms of substantial improvements over existing methods - than revolutionary. This short course is intended as a first and reasonably self-contained introduction to the topic, with the purpose of analyzing Deep Learning in a broader theoretical context and exploring the reasons of its unquestionable success, in several specific application cases. The short course includes also an introductory, hands-on tutorial on TensorFlow, one of the best known open-source deep learning frameworks which is supported and developed by Google. The tutorial aims to give the fundamental notions about this framework by working on specific examples.

The book “Deep Learning”, by I. Goodfellow, Y. Bengio and A. Courville, MIT Press, 2017, will be assumed as reference for these episodes. At present, an online, html version of the book is freely accessible online at <http://www.deeplearningbook.org/>.

For TensorFlow, a good introductory book is “Hands-On Machine Learning with Scikit-Learn & TensorFlow” by Aurélien Géron, O’Reilly, 2017

Online registration

Due to the limitations of working places for the hands-on tutorial, participants must register online (details to be defined). The registration is free of charge and reserved to the students and staff of University of Pavia.

Prerequisites (for the hands-on tutorial)

Previous knowledge of the topics below will be assumed as prerequisites:

- **Linux Ubuntu:** minimal survival skills
- **Python programming:** basic operations, lists, tuples, dictionaries, conditions, loops, functions, classes and instances, modules, exception handling (all code examples will be in Python 3)
- **Numpy basics:** array creation, basic operations and linear algebra (matmul), indexing, slicing, iteration, reshaping
- **PyCharm:** some experience is not mandatory but strongly recommended (with the debugger, in particular)

Detailed Syllabus

Friday, May 4, 2:15 pm - 4:30 pm, Aula ex_ELN floor D

Machine learning: Artificial Neural Networks

- An introductory example: linear regression
- Representation: artificial neural networks as universal approximators
- Evaluation: loss functions
- Optimization: gradient descent, stochastic gradient descent, mini-batch gradient descent

Friday, May 11, 2:15 pm - 4:30 pm, Aula ex_ELN floor D

Deep Learning for Neural Networks

- Deep network representation: fundamentals, potential advantages and major issues
- The choice of non-linear functions: ReLU
- More tricks: batch normalization, dropout
- Beyond gradient descent: momentum, Nesterov, AdaDelta, AdaGrad

Friday, May 18, 2:15 pm - 4:30 pm, Aula ex_ELN floor D

Deep Convolutional Neural Networks

- The ImageNet challenge
- Convolutional filters: representation and gradients
- Max pooling, normalizations: representation and gradients
- Advanced architectures: Inception and ResNet
- Transfer learning
- Representation in deep layers, adversarial generation, merging
- Beyond single-object classification: object location, segmentation (hints)

Friday, May 25, 2:15 pm - 5:30 pm, Aula ex_ELN floor D

TensorFlow Fundamentals (theory)

- Flow Graphs: constant, variables, placeholders and ops
- Tensor transformations: slicing, broadcasting, reshaping
- Sessions: computing flow graphs
- Optimization: loss functions, gradients, optimizers

Thursday, May 31, 2:15 pm - 6 pm, Aula B1 floor B

A first example with TensorFlow

- Dataset preparation (MNIST)
- Flow graph representation: fully connected neural networks
- Optimization: stochastic gradient descent in detail and with an optimizer
- Running a session and visualizing results

Thursday, June 7, 2:15 pm - 6 pm, Aula B1 floor B

TensorFlow for DCNN - Part 1

- Convolutional networks anatomy: all the building blocks
- Going practical: a working example with VGG-16
- Preparation: setup, dataset, pre-trained models
- Fine tuning and network surgery: telling cats from dogs

Thursday, June 14, 2:15 pm - 6 pm, Aula B1 floor B

TensorFlow for DCNN - Part 2

- The bag of tricks: batch normalization, dropout, improved optimizers
- Visualizing DCNN representations
- An higher level of abstraction: the Keras front end

Organizer
Prof. Virginio Cantoni

Ph.D. Coordinator
Prof. Paolo Di Barba