



**University of Pavia**

**Department of Electrical, Computer and Biomedical Engineering**  
**Computer Vision and Multimedia Lab**

# Deep Learning: an Introduction

*Artificial Intelligence Reading Group*

May 4 - May 18, 2018

**Instructor:**

**Marco Piastra**, PhD, [marco.piastra@unipv.it](mailto:marco.piastra@unipv.it)

**Co-instructors:**

**Mirto Musci**, PhD, [mirto.musci@unipv.it](mailto:mirto.musci@unipv.it)

**Andrea Pedrini**, PhD, [andrea.pedrini@unipv.it](mailto:andrea.pedrini@unipv.it)

**Gianluca Gerard**, [gianluca.gerard01@universitadipavia.it](mailto: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.

# Programme

Friday, May 4, 2:15 pm - 4:30 pm, Aula 8 Nuovo Polo floor 1

## **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 8 Nuovo Polo floor 1

## **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 8 Nuovo Polo floor 1

## **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)

All presentations will be made available at:

<https://vision.unipv.it/AI/AIRG.html>