

# Machine Learning

*Tuesday, 24 July 2018 13:30 (1h 30m)*

Machine learning (ML) is a thriving field with active research topics. It has found numerous practical applications in natural language processing, understanding of speech and images as well as fundamental sciences. ML approaches are capable of replicating and often surpassing the accuracy of hypothesis driven first-principles simulations and can provide new insights to a research problem.

We here provide an overview about the content of the Machine Learning tutorials.

Although the theory and practice sessions are described separately, they will be taught alternating one to the other, during the four lectures. In this way, after we've introduced new concepts, we can immediately use them in a tailored exercise, which will help us absorb the material covered.

## Theory

We'll start with a gentle introduction to the ML field, introducing the 3 learning paradigms: supervised, unsupervised, and reinforcement learning. We'll then delve into the two different supervised sub-categories: regression and classification using neural nets' forward and backward propagation.

We'll face overfitting and fight it with regularisation.

We'll soon see that smart choices can be done to exploit the nature of the data we're dealing with, and introduce convolutional, spectral, recurrent, and graph neural nets.

We'll move on to unsupervised learning, and we'll familiarise with generative models, such as variational autoencoders and generative adversarial networks.

## Practice

We will introduce machine learning technology focusing on the open source software stack, namely PyTorch and Keras frameworks.

Brief introduction to PyTorch architecture, primitives and automatic differentiation, implementing multi-layer perceptron and convolutional layers, a deep dive into recurrent neural networks for sequence learning tasks.

Introduction to Keras.

Learn to debug machine learning applications and visualize training and validation process with pytorchviz or TensorBoard. Discuss ways to train multi-GPU and distributed models on a cluster with Horovod package.

All exercises will use PyTorch or Keras. Python programming experience is desirable, but previous experience with PyTorch and Keras is not required.

**Presenters:** SVYATKOVSKIY, Alexey (Princeton University); CANZIANI, Alfredo (NYU Center for Data Science)