

Applied Machine learning with Python

Lecturer: **Andrea Giussani**

Course language

English

Course description and objectives

The purpose of this course is to give you a solid introduction to modern applied Machine Learning (ML) methods and pipelines that are generally available for practitioners in the field of Machine and Statistical learning. In particular, you will be guided through the construction of an end-to-end pipeline using both linear and nonlinear methods. In particular, each method will be briefly described and accompanied by hands-on, practical snippets written in Python.

At the end of this course, each participant will be able to independently apply modern machine learning techniques to elementary problems related to Economics and Social Sciences.

Audience

The course is aimed for people involved in applying ML methods in the industry, especially for those interested in working in data analytics, or for those who want to extend their knowledge in modeling and statistical inference.

Upon successful completion of the course (attendance of at least 75% of the scheduled hours and passing the final exam), students will get 2 credits and an Open Badge, sharable across the web (LinkedIn) or personal CV.

Prerequisites

The student must have a solid knowledge of Probability and Statistical Inference, and great knowledge of the Python language. All students are strongly encouraged to bring their own laptop.

It is strongly recommended that students have been exposed to the following libraries: Numpy, Pandas, Matplotlib. We are going to use them a lot, so please be sure you have a good understanding of them.

Duration

24 hours

Teaching mode

This course will take place exclusively in **synchronous mode** in the **classroom**. Online mode will not be provided.

The final test of the course will take place during the last day of class.

Calendar

Lesson	Date	Time	Room
1	Thu 06/02/2025	18.15 - 19.45	N12
2	Fri 07/02/2025	16.30 - 18.00	N12
3	Thu 13/02/2025	18.15 - 19.45	N12
4	Fri 14/02/2025	16.30 - 18.00	N12
5	Thu 20/02/2025	18.15 - 19.45	N12
6	Fri 21/02/2025	16.30 - 18.00	N12
7	Thu 27/02/2025	18.15 - 19.45	N12
8	Fri 28/02/2025	14.45 - 16.15	N12
9	Fri 28/02/2025	16.30 - 18.00	N12
10	Thu 06/03/2025	18.15 - 19.45	N12
11	Fri 07/03/2025	14.45 - 16.15	INFOAS04/05
12	Fri 07/03/2025	16.30 - 18.00	INFOAS04/05

Syllabus of the course

Lecture	Topics	Book reference
1	Introduction to Machine Learning <ul style="list-style-type: none"> - Objectives of the course - Introduction to Machine Learning (ML) - The Python ML toolkit: a gentle introduction 	

2 Modern Machine Learning Pipeline: Preprocessing Ch. 1

- Feature Distribution and Scaling
- Normalization

3 Modern Machine Learning Pipeline: Preprocessing (2) Ch. 1

- Imputation
- Dealing with Categorical Variables
- Model Selection

4 Modern Machine Learning Pipeline: Modeling Ch. 2

- A simple ML Pipeline
- Overfitting and Underfitting
- Model Selection

5 Modern Machine Learning Pipeline: Modeling (2) Ch. 2

- Linear vs Logistic Regression
- Shrinkage Methods
- Lab: Classification
- Lab: Regression

6 Modern Machine Learning Pipeline: Modeling (3) Ch. 3

- Nonlinear Models: why do we need them
 - Gradient Boosting
 - Random Forest
 - XGBoost
 - Explain Machine Learning Pipeline using Shap
 - Fraud Detection
-

7 Introduction to Natural language processing

- Preprocessing with unstructured data
 - Working with Embeddings
 - Standard NLP Pipeline with scikit-learn
 - Lab: Text Classification
-

Software used

Python 3.9.x (or greater)

Suggested bibliography

- Giussani A., *Modern Applied Machine Learning with Python*, EGEA, 2020
- Hastie T., Tibshirani R., Friedman J., *The elements of statistical learning*, Springer, 2009

Available seats

This activity is limited to **110** participants and reserved to **students of the Master of Science Programs**. Registrations cannot be carried out once this number has been reached or after closing of the registration period.