FLORENCE SCHOOL OF BANKING AND FINANCE

EXECUTIVE TRAINING COURSE

PROGRAMME

AI Methods in economics and finance research

25-27 May 2026

Ten years of excellence,
shaping knowledge for tomorrow

2016 - 2026

fbf.eui.eu

Co-funded by
the European Union

Introduction

Research methods have developed significantly in the last decade; first with the rise of machine learning (ML) approaches and more recently with the use of Generative AI tools employing, for instance, Large Language Models (LLM). Such tools are also gaining ground in financial research. This course will provide a solid foundation for setting up your own workflows.

The course starts with a whirlwind tour of 'traditional' AI. Then approaches to explain the sometimes opaque model outcomes are discussed. After this, neural networks, training and architectures are covered. This prepares us for looking into LLMs and AI agents. We close with discussing how to ensure solid workflows and the practical implementation of state-of-the-art LLM workflows.

Location | EUI Premises in Florence, Italy

Course Directors

- Iman van Lelyveld (Part-time Professor | Florence School of Banking and Finance)

Faculty

- Iman van Lelyveld (Part-time Professor | Florence School of Banking and Finance)
- Michiel Nijhuis (Data Scientist | De Nederlandsche Bank (DNB))

Teaching Associate

- Hugo Bourrousse (Research Associate | Florence School of Banking and Finance)

Format | Residential

Level | Intermediate, Advanced

Approach | Quantitative

Learning objectives:

- Understand the fundamentals of machine learning models and large language models
- Design and implement a basic machine learning pipeline, including data preparation, model training, and evaluation
- Apply large language models in practical research settings

Target audience:

This course is designed for researchers with a traditional background in economics and finance who wish to become acquainted with emerging AI tools. Rather than focusing on cutting-edge implementations, the course aims to provide a clear sense of what can be achieved with these techniques, along with initial hands-on practical experience.

MONDAY 25 MAY

09:30 - 10:30 **Lecture 1**

- **Introduction to the course**
- **Break-out to discover common interests**
- **Overview of tooling, Machine Learning and the LMM landscape including the impact**

10:30 - 11:00 **Coffee break**

11:00 - 12:30 **Lecture 2**

Machine Learning as a framework

- ML as loss minimization; OLS as a special case; where ML is applied (prediction, classification, anomalies)

How ML works in practice

- Supervised vs unsupervised learning; models, (hyper)parameters, optimization, preprocessing
- Model evaluation: overfitting, bias–variance tradeoff, validation strategies

Modern ML models and controlling complexity

- Dimensionality reduction (PCA) and regularization (Ridge, LASSO, Elastic Net)
- Nonlinear models: trees, ensembles (bagging, boosting, XGBoost), and applications (e.g. asset pricing, anomaly detection)

12:30 - 13:30 **Lunch break**

13:30 - 14:30 **Lecture 3**

Explainability is a must

Which techniques can be used to explain results from black box models?

- Permutation importance
- SHAP
- Counterfactuals

Explainable AI (XAI) in practice

14:30 - 15:30 **Hands on exercise 1**

Applying XAI yourself:

- Can you spot bias in a model
- Finding which variables matter the most

15:30 - 16:00 **Coffee break**

16:00 - 17:30 **Lecture 4**

Basic neural networks: from origins to workhorse models

- From regression to a neural network
- Backpropagation

Why neural networks work

- Universal approximation and representation power
- Flexibility vs interpretability tradeoff

Applications in finance

- Credit risk, asset pricing, forecasting, and anomaly detection

TUESDAY 26 MAY

09:30 - 10:30 **Lecture 5**

Design choices in neural networks

- Feature engineering vs. learning representations with many parameters
- Initialization and learning rate
- Neural network architectures

Feedforward networks

- Recurrent and convolutional networks
- Training deep learning models

Minibatch stochastic gradient descent

- Vanishing/exploding gradients and the role of regularization layers

10:30 - 11:00 **Lecture 6**

PyTorch as a deep learning framework

- Tensors, computation graphs, and automatic differentiation

Neural networks in PyTorch

- Defining models, loss functions, and optimizers

11:00 - 11:30 **Coffee break**

11:30 - 13:00 **Hands-on exercise 2**

Building your own neural network

- Implementing and training a simple NN in PyTorch
- Changing the architecture

Understanding model performance

- Exploring how hyperparameters affect convergence and accuracy

13:00 - 14:00 **Lunch break**

14:00 - 15:00 **Lecture 7**

Recurrent Neural Networks and LSTMs

- Sequential modelling and longterm dependencies
- How LSTMs can be used for time series and text

Transformers as the foundation of LLMs

- Attention mechanisms and parallel processing
- Why Transformers enable models like ChatGPT

15:00 - 16:00 **Hands-on exercise 3**

Training deep learning network

- Building a training loop
- Handling vanishing gradients

Solving training issues

- Can you identify and fix various training issues

16:00 - 16:30 **Coffee break**

16:30 - 18:00 **Lecture 9**

Large Language Models (LLMs): state of the art

- Overview of leading LLMs and their capabilities

Trends in foundation models

- Small Language Models (SLMs)
- Multimodal models
- Timeseries foundation models

From theory to AI agents

- Agent concepts and workflows
- Using LangGraph to structure, coordinate, and manage LLM-based agents

WEDNESDAY 27 MAY

09:30 - 11:00 **Hands-on exercise 4**

Building your own LLM multi agent system

- Applying the LangGraph framework

Designing agent workflows

- Roles, state, and coordination between agents

11:00 - 11:30 **Coffee break**

11:30 - 13:00 **Lecture 10**

From AI agents to data engineering

- Why reliable data flows are critical for agentbased systems

Key data engineering concepts

- Data ownership, data locations, and data access patterns
- Reproducibility, versioning, and data lineage

ML workflows in practice

- Setting up an end to end ML workflow
- Participants bring their own data and questions (default dataset provided)

13:00 - 14:00 **Lunch break**

14:00 - 16:00 **Lecture 11**

Why AI results need careful scrutiny

- Powerful models do not automatically imply reliable or fair outcomes

Key pitfalls of AI and ML models

- Bias in data and models
- Model drift over time as data and environments change

Managing risks in practice

- Monitoring, validation, and governance to maintain trustworthy models

Session timings are provisional and subject to change

fbf.eui.eu

Co-funded by
the European Union