ONLINE LEARNING WITH APPLICATIONS TO DIGITAL MARKETS

The explosive growth of online markets has created complex ecosystems of algorithmic agents. To optimize their revenue, agents need to understand how the market works, and to do so they often resort to strategies that learn from past observations. In this course we introduce Online Convex Optimization, the main algorithmic framework for the study of sequential decision-making problems. The first part of the course describes the most important settings and algorithms for online learning. In the second part, we review some recent results characterizing the strengths and limitations of sequential decision-making approaches applied to various problems arising in digital markets. The analysis sheds light on how learning depends on the interplay between the form of the revenue function and the feedback provided during the learning process.

Schedule: **Mon-Fri 10:30-12:30** and **14:30-16:30**

Lecturers:

- Nicolò Cesa-Bianchi (10 hours)
- Pierre Laforgue (10 hours)

SYLLABUS

- The online learning framework: introduction and motivation
- Sequential risk and regret
- Follow the Leader: linear lower bound
- Follow the Regularized Leader (FTRL)
- The Euclidean case: Online Gradient Descent
- The entropic case: Exponentiated Gradient (EG)
- Proving the FTRL regret bound
- Parameter-free algorithms
- Euclidean vs entropic regularization
- EG with linear losses: experts, bandits, and feedback graphs
- Multitask online learning
- Regret minimization and equilibria
- Dynamic pricing, retail pricing, and adaptive taxation
- Auctions
- Bilateral Trading

Reference material:

- A Modern Introduction to Online Learning
- Additional material will be posted here