

Syllabus

[cmcd.economia@fgv.br]

Course Statistical Learning

Professor

2026 FIRST SEMESTER

COURSE OUTLINE

This course describes the main techniques in statistical learning and their applications to economics and finance. We start with the definition of statistical learning relative to machine learning and high-dimensional econometrics, and then discuss big data in economics and finance. We then go through the details of regularization methods, describing not only their corresponding intuition, but also the challenges of carrying out statistical inference. Next, we introduce statistical learning algorithms for nonlinear models, including tree-based regressions and artificial neural networks, as well as for the analysis of text as data. We conclude with a discussion about causality.

METHODOLOGY

Students must read the references before lectures, in order to foster discussion.

PROGRAM

1. Introduction: Definitions and review of applications
2. Linear models: Regularization methods
3. Nonlinear models: Overfitting issues
4. Text as data: NLP techniques (time allowing)
5. Causality: DAG and causal trees (time allowing)

BIBLIOGRAPHY

Masini, Medeiros & Mendes (JES 2023), Athey (NBER 2019), Chetty (Harvard course 2019).

Bryzgalova (2015), Bryzgalova & Julliard (2019), Bryzgalova, Pelger & Zhu (2019), Callot, Kock & Medeiros (JAE 2017), Chen, Pelger & Zhu (2019), Chinco, Clark-Joseph & Ye (JF 2019), Engle, Ledoit & Wolf (JBES 2019), Feng, Giglio & Xiu (2017), Feng, Polson & Xu (2018), Fernandes, Medeiros & Scharth (JBF 2014), Fernandes & Vieira (JECD 2019), Gu, Kelly & Xiu (RFS 2019), Ledoit & Wolf (JEF 2003, JMA 2004, JPM 2004, AoS 2012, RFS 2017), Medeiros, Vasconcelos, Veiga & Zilberman (JBES 2019), Olivares-Nadal & DeMiguel (OR 2008).

GRADING

In-class presentation:	20%
Mock referee report:	30%
Empirical work:	40%
In-class participation:	10%

PROFESSOR