Syllabus

Predictive and Causal Machine Learning in R

Summer semester 2023

(as of April 17, 2023)

Prof. Martin Huber, Ph.D.

Phone: +41 26 300 82 74

Email: martin.huber@unifr.ch

Website: https://www.unifr.ch/appecon/de/lehrstuhl/prof/

1. General information

1.1 Course content

This course provides an introduction to predictive and causal machine learning based on the software "R". It is directed at students at Bachelor, Master, and PhD level.

Predictive machine learning aims at forecasting the value of an outcome of interest, e.g., sales or turnover, based on observing specific patterns of potentially relevant factors (or "predictors") like price, quality, weather, advertisement campaigns etc. That is, predictive machine learning "learns" from patterns among predictors in (past) data to forecast the value of the outcome in the future.

Causal machine learning aims at assessing the causal effect of some intervention or treatment, like offering or not offering a training program to jobseekers, on an outcome of interest, like employment. The assessment of a causal effect requires that groups receiving and not receiving a discount are comparable in background characteristics which also affect their employment (e.g., previous labor market history, education etc.). Causal machine learning can be used to generate such comparable groups in a data-driven way by estimating two separate models for how the characteristics affect the intervention and the outcome. Such approaches also permit detecting subgroups for whom the treatment effect is particularly large as a function of their observed characteristics (effect heterogeneity analysis). This is useful for optimally targeting specific subgroups by the treatment (optimal policy learning). Finally, by repeatedly assigning alternative treatments over time in an appropriate way, one may learn and converge to the assignment of the most effective treatment (reinforcement learning).

This course first discusses the underlying assumptions, intuition and usefulness of machine learning for forecasting and causal analysis. It then introduces various machine learning algorithms and discusses their application for prediction/forecasting and causal analysis. Using the statistical software "R" and its interface "R Studio", these methods are applied to various real-world data sets.

- Introduction to the concepts and purposes of predictive and causal machine learning
- Basics of predictive machine learning: Model tuning (cross-validation) and performance evaluation (out-of-sample testing)
- Prediction based on penalized regression (lasso and ridge regression)
- Prediction based on tree-based approaches (trees, bagging, random forests)

- Further predictive machine learners: boosting, support vector machines, neural networks (deep learning), and ensemble methods
- Causal analysis based on penalized regression (lasso and ridge regression)
- Causal analysis using tree-based approaches (causal trees and causal forests)
- Causal analysis based on double machine learning
- Assessing effect heterogeneity across subgroups
- Optimal policy learning to maximize treatment effectiveness using tree-based approaches
- Reinforcement learning to learn the most effective treatment (among several alternatives) by repeated treatment assignment over time
- Application of all methods to real world data using the statistical software "R" and its interface "R Studio"

1.2 Course goals

- To understand the ideas, goals, and differences of machine learning for prediction and for causal analysis
- To understand the intuition, advantages, and disadvantages of alternative methods
- To be able to apply predictive and causal machine learning to real world data using the software "R" and its interface "R Studio"

1.3 Prerequisites

- Introductory statistics (probability theory, conditional means, linear regression), basic command of the statistical software "R".

1.4 Credit points

This course is an additional offer; no credit points can be earned. However, there is the option to obtain a certificate of participation upon the successful completion of the final examination.

1.5 Registration

Please register via Stud.IP (course ID: 801873) between **24 April 2023, 00:00** and **12 June 2023, 23:59**. The number of participants is restricted to 25. Students will be admitted on a first-come-first-serve basis. In addition, there will be a waiting list.

2. Course overview

2.1 Scheduling

The course takes place on campus on 13–16 June 2023. Please **bring your own computer** as we will not have access to the computer labs. Ideally, have R and R Studio installed by the time the course starts: https://posit.co/download/rstudio-desktop/

2.2 Meetings

See Stud.IP/EXA for rooms and potential updates to the schedule.

Tuesday (June 13) and Wednesday (June 14):

08:30–10:00	First lecture
10:00–10:20	Coffee break
10:20–11:50	Second lecture
11:50–13:10	Lunch break
13:10–14:40	Third lecture
14:40–15:00	Coffee break
15:00–16:30	Fourth lecture

Thursday (June 15) and Friday (June 16):

08:30–10:00	First lecture
10:00–10:20	Coffee break
10:20–11:50	Second lecture
11:50–12:00	Coffee break
12:00–12:45	Third lecture

Approximate schedule in terms of topics covered:

Day 1

- Purpose of statistical modelling and machine learning for predictive analysis (in contrast to causal analysis).
- Basic concepts underlying all machine learning algorithms like performance measurement, the variance-bias tradeoff, sample splitting (into training, validation, and test data), and cross-validation.

- Machine learning algorithms based on penalized regression or shrinkage methods, which shrink the importance of weak predictors, namely lasso regression, ridge regression, and elastic nets.
- Machine learning algorithms that based on decision trees, namely regression/classification trees, bagged trees, and random forests.

Day 2

- Further machine learning algorithms: Boosting, support vector machines, neural networks, and ensemble methods.
- Purpose of causal analysis, potential outcome notation, identifying assumptions for treatment evaluation (like the selection-on-observables assumption).
- Causal analysis based on lasso estimation (penalized regression): double selection or partialling out procedures to condition on important control variables, estimator's properties with and without sample splitting/cross-fitting.

Day 3

- Causal analysis based on random forests: effect interpretation in nonparametric models with binary and non-binary treatments, average treatment effect estimation, estimator's properties, and heterogeneity analysis.
- Causal analysis based on double machine learning: concept of doubly robust estimation (can be combined with a range of machine learning algorithms), average treatment effect estimation, estimator's properties, and heterogeneity analysis.

Day 4

- Optimal policy learning: Combining double machine learning and tree based-methods for (1) a data-driven segmentation into subgroups and (2) an optimal treatment assignment across subgroups for maximizing treatment effectiveness.
- Reinforcement learning: Repeated random assignment of alternative treatments over time to iteratively converge to the treatment assignment which is (on average) most effective.

2.3 Examination and grading of the module

Participants have the opportunity to take an exam in order to obtain a certificate of participation. The exam will be graded as pass or fail. The exam format will be a take-home exam with empirical exercises using R to be submitted to martin.huber@unifr.ch until 15 August 2023.

2.4 Course materials

Lecture slides, R code, and data files will be made available to the course participants.

Textbooks:

- *For predictive machine learning:* G. James, D. Witten, T. Hastie, and R. Tibshirani (2021): An Introduction to Statistical Learning with Applications in R, Springer, New York. Freely available at: https://www.statlearning.com/
- *For causal machine learning:* M. Huber (2023): Causal analysis - Impact evaluation and causal machine learning with applications in R, forthcoming in MIT Press, Cambridge. A draft version in pdf format will be made available to the course participants.