Introduction to Machine Learning for Social Sciences

Instructor: Dr Jens Wäckerle, University of Cologne

Course outline

In this course, students will learn the basics of machine learning, specifically for social scientists. While the course will introduce specific models, such as support vector machines, decisions trees, or lasso regressions, we will also focus on general principles that apply to all machine learning algorithms. This will equip students with the tools to not only apply machine learning in their own research, but also judge the performance of such models and adapt to new developments in the field. The course includes practical applications with code and examples from different social science fields.

Prerequistes

All code examples and applications will be done in R. The course requires basic knowledge of R, which includes being able to load and work with data. The code examples are based on the tidyverse (e.g. filter, mutate, summarise functions), but can be adjusted to base R without much trouble. No special mathematical knowledge beyond master level statistics and quantitative methods courses is required, the course focuses on the concepts of machine learning algorithms rather than the mathematical basis. The examples and lab session don't require any in-depth knowledge of political science or any other field.

Time Schedule and Topics covered

Every day, classes will be split roughly evenly between a seminar session that introduces new methods, approaches, measures and applications and a lab session in which the content is applied to actual datasets from social science.

Monday 18.03.2024: General Introduction to Machine Learning and Classification Problems
- What is machine learning?

- What are some overarching themes to understand? Parameters, Hyperparameters, cross-validation, …
- What are possible applications?
- How is machine learning different from your intro to quantitative methods class?

Tuesday 19.03.2024: Classification Problems with SVMs and Trees

- How can we classify new data into pre-existing categories?
- How can we evaluate the performance of an algorithm?

Wednesday 20.03.2024: From trees to forest ensembles

- How can we expand on trees?
- How can we combine many good algorithms into one great algorithm?

Thursday 21.03.2024: Regression problems

- What does regression have to do with machine learning?
- How can we predict continuous outcomes with machine learning?

Friday 22.03.2024: Workflow, Implementing Machine Learning in Research, Wrap-Up

- How do we implement a machine learning workflow?
- How can we report results and get machine learning published?
- Machine Learning as methods papers, part of the research process, or outcomes.

About the Trainer

Jens Wäckerle is a post-doctoral researcher at the Cologne Center for Comparative Politics and currently the substitute professor at the Chair of European Politics. His research focuses on women in politics, legislative politics, quantitative text analysis and the European Union. He holds a Master degree from the University of Essex and a PhD from the University of Cologne.