Introduction to Machine Learning

Reto Wüest
University of Geneva

Course Description

With ever more data available in electronic form, automated methods of data analysis become increasingly important also in the social sciences. Machine learning refers to a set of methods that can automatically detect patterns in data, or "learn" from data. The uncovered patterns can then be used by the analyst to make accurate predictions and decisions under uncertainty.

This course will introduce participants to the fundamentals of machine learning. Students will leave the course with a thorough understanding of the core issues in machine learning (prediction and inference, supervised and unsupervised learning, overfitting, bias-variance trade-off), knowledge of some of the most widely used machine learning methods, and the ability to apply these methods in their own research. All course materials are available at http://retowuest.github.io/ml2018/.

Software

The course will use the open-source software R, which is freely available for download at https://www.r-project.org/. We will interact with R through the user interface RStudio, which can be downloaded at https://www.rstudio.com/products/rstudio/download/.

Prerequisites

Participants are expected to have a solid understanding of linear and binary regression models. The course will also assume at least a basic familiarity with the R statistical programming language.

Schedule

Session 1: Introduction to Machine Learning

(March 6, 2018, 13:00-17:00)

The first session will provide an introduction to machine learning. We will discuss the goals of machine learning (prediction, inference, or both), the difference between supervised and unsupervised machine learning, the problem of overfitting, and the bias-variance trade-off. We will then get to know the first class of important supervised learning methods, namely shrinkage methods (Ridge regression and the Lasso).

Class Schedule

Time	Topic
13:00-13:30	Introductions and course overview
13:30-14:00	General introduction to machine learning (prediction and inference, supervised and unsupervised learning)
14:00-14:45	Assessing model accuracy (overfitting, bias-variance trade-off, cross-validation)
14:45-15:15	Break
15:15-15:45	Shrinkage methods I: Ridge regression
15:45-16:15	Shrinkage methods II: The Lasso
16:15-17:00	Application of Ridge regression and the Lasso

Main Readings

- James et al., *An Introduction to Statistical Learning*, Ch. 2 and 6

Recommended Readings

- Hastie et al., *The Elements of Statistical Learning*, Ch. 3 and 7
- Shalev-Shwartz and Ben-David, *Understanding Machine Learning*, Ch. 5, 13
- Bishop, *Pattern Recognition and Machine Learning*, Ch. 12

Session 2: Classification and Regression Trees (CART)

(March 27, 2018, 13:00-17:00)

The second session will deal with tree-based methods, which are another important and highly flexible class of supervised learning methods. After an introduction to the basics of decision trees and a general discussion of the advantages and disadvantages of tree-based models, we will look at three specific widely-used tree-based methods: bagging, random forests, and boosting.

Class Schedule

Time	Topic
13:00-13:30	Introduction to classification and regression trees
13:30-14:00	Advantages and disadvantages of trees
14:00-14:45	Bagging, random forests
14:45-15:15	Break
15:15-16:00	Boosting
16:00-16:30	Application I: Random forests
16:30-17:00	Application II: Boosting

Main Readings

- James et al., *An Introduction to Statistical Learning*, Ch. 8

Recommended Readings

- Hastie et al., *The Elements of Statistical Learning*, Ch. 9, 10, and 15
- Shalev-Shwartz and Ben-David, *Understanding Machine Learning*, Ch. 18
- Lantz, *Machine Learning with R*, Ch. 11

Session 3: Unsupervised Learning

(April 17, 2018, 13:00-17:00)

In the third session, we will move to unsupervised machine learning methods. We will cover two important unsupervised learning techniques: principal components analysis (PCA) and clustering analysis (K-means clustering and hierarchical clustering).

Class Schedule

Time	Topic
13:00-13:30	Introduction to unsupervised learning
13:30-14:15	Principal components analysis (PCA)
14:15-14:45	K-means clustering
14:45-15:15	Break
15:15-16:00	Hierarchical clustering
16:00-16:30	Application I: PCA
16:30-17:00	Application II: Clustering methods

Main Readings

- James et al., *An Introduction to Statistical Learning*, Ch. 10

Recommended Readings

- Hastie et al., *The Elements of Statistical Learning*, Ch. 14
- Shalev-Shwartz and Ben-David, *Understanding Machine Learning*, Ch. 22, 23
- Bishop, *Pattern Recognition and Machine Learning*, Ch. 12
- Barber, *Bayesian Reasoning and Machine Learning*, Ch. 15
- Lantz, *Machine Learning with R*, Ch. 9

References

Barber, David. 2016. *Bayesian Reasoning and Machine Learning.* New York: Cambridge University Press. Available for free as a PDF.
URL: *http://web4.cs.ucl.ac.uk/staff/D.Barber/pmwiki/pmwiki.php?n=Brml.HomePage*

Bishop, Christopher M. 2006. *Pattern Recognition and Machine Learning.* New York: Springer.

Hastie, Trevor, Robert Tibshirani and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* 2nd ed. New York: Springer. Available for free as a PDF.
URL: *https://web.stanford.edu/ hastie/ElemStatLearn/*

James, Gareth, Daniela Witten, Trevor Hastie and Robert Tibshirani. 2013. *An Introduction to Statistical Learning with Applications in R.* New York: Springer. Available for free as a PDF.
URL: *http://www-bcf.usc.edu/ gareth/ISL/*

Lantz, Brett. 2015. *Machine Learning with R.* 2nd ed. Birmingham: Packt Publishing.

Shalev-Shwartz, Shai and Shai Ben-David. 2014. *Understanding Machine Learning: From Theory to Algorithms.* New York: Cambridge University Press. Available for free as a PDF.
URL: *http://www.cs.huji.ac.il/ shais/UnderstandingMachineLearning/*