

Information Sheet

Lecturer: Aslan Tchamkerten (3.C51, aslan.tchamkerten@telecom-paris.fr)

Teaching assistants: Breno Skuk, Ibrahim Djemai

Language: English

Website: <https://www.tchamkerten.com/micas911-si221>

Lectures are Tu 1.30-4:45pm and consists of 3 periods of 55 minutes each, 2 lecture periods plus one exercise period. Between periods there is a 7 minute break.

Exercise periods (TD) consists of exercises for the first 3 sessions and lab work for the last 4 sessions. For the exercises, some will be presented in class, and some will be given as homeworks. From lecture 4 onwards homeworks should be handled back at the beginning of the next lecture and will be graded.

Grading policy: final grade = 85% final exam and 15% homeworks

To pass the course: your final grade must be $\geq 10/20$.

Office hours : 7/7

Bibliography:

1. Understanding Machine Learning, by S. Shalev-Shwartz and S. Ben-David, Cambridge University Press, 2014. Available on first author's webpage
2. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. by T. Hastie, R. Tibshirani and J. Friedman, J. (2009).
3. High Dimensional Statistics, lecture notes by P. Rigollet, MIT, available @ <http://www-math.mit.edu/~rigollet/PDFs/RigNotes17.pdf>
4. High-Dimensional Statistics: A Non-Asymptotic Viewpoint High-Dimensional Statistics, by M. Wainwright, Cambridge University Press, 2019

The course will mostly follow 1., referred to as UML.

In a nutshell: 1. provides a unifying view on machine learning problems (rather than a toolbox perspective), 2. good if you want to build intuition and like pragmatism, many examples. 3. and 4. are references for the (most) mathematically inclined students.

About timing: lectures will start right on time. Unless “un cas de force majeure” if you think you might be late, just skip the current period and ask the notes of a colleague. If you need to leave the class early, let me know before the course starts.