

PLACE AND TIME: Room DIA005, Mondays 11:15-13:00 and 14:15-16:00

INSTRUCTOR: Ali H. Sayed, Email: ali.sayed@epfl.ch


TEACHING ASSISTANT: Stefan Vlaski, E-mail: stefan.vlaski@epfl.ch

COURSE MATERIAL: Lecture notes authored and distributed by the instructor for exclusive use by students enrolled in the class.

PRE-REQUISITES: It is recommended that students have some familiarity with matrix theory, linear algebra, and probability. Supplemental material on these topics is provided by the instructor as needed.

GRADING: 5 homework assignments (40%) and one group project 60% (in groups of 2-3 students).

TOPICS: In this course, students learn to master tools, algorithms, and core concepts related to inference from data, data analysis, and adaptation and learning theories. *Emphasis is on the theoretical underpinnings and statistical limits of learning theory.* In particular, the course covers topics related to optimal inference, estimation theory, regularization methods, proximal methods, online and batch methods, stochastic learning, generalization and statistical learning theories, Bayes and naive classifiers, nearest-neighbor rules, self-organizing maps, decision trees, logistic regression, discriminant analysis, Perceptron, support vector machines, kernel methods, bagging, boosting, random forests, cross-validation, and principal component analysis. Project themes selected by students in consultation with instructor.



LECTURE	DATE	TENTATIVE TOPIC(S)
1	Feb. 19	Optimal Inference
2	Feb. 26	Bayesian and Maximum Likelihood Inference
3	Mar. 5	Regression Analysis
4	Mar. 12	Regularization
5	Mar. 19	Batch Learning
6	Mar. 26	Stochastic Gradient Learning
HH	Apr. 2	HOLIDAY (NO LECTURE)
7	Apr. 9	Stochastic Sub-gradient Learning
8	Apr. 16	Stochastic Proximal Learning
9	Apr. 23	Classification & Clustering: naïve Bayes, NN rule, k-means clustering, SOMs
10	Apr. 30	Decision Trees; Generalization Theory
11	May 7	Logistic Regression; Linear Discriminant Analysis; The Perceptron
12	May 14	Support Vector Machines; Kernel-Based Learning
HH	May 21	HOLIDAY (NO LECTURE)
13	May 28	Bagging and Boosting; Cross-Validation; PCA