Advanced Introduction to Machine Learning

This is an ambitious class that will provide a broad overview of a large variety of machine learning methods in a short amount of time. You will learn to understand the basics of: linear regression; logistic regression; k-nearest neighbors; PCA, LDA, and dimensionality reduction methods; feature selection and engineering; cross validation; the bootstrap, bagging, and boosting; decision trees and random forests; naive Bayes; generative vs. discriminative models; support vector machines and kernel methods; (deep) neural networks; Bayesian nonparametric methods; clustering; Gaussian mixtures; the EM algorithm; ensemble methods; reinforcement learning; representation learning; information theory; Gaussian processes; supervised, unsupervised, and semi-supervised learning; graphical models; sparsity and compressed sensing; planning and control; information retrieval; structured prediction; matrix factorization; Monte Carlo methods; time-series analysis and HMMs; multi-agent learning; transfer and multi-task learning; active learning; submodularity; and machine teaching. Along the way, we will motivate the above using applications in computational biology, networks, computer vision, speech recognition, and natural language processing. We will also touch on the philosophy of machine learning and artificial intelligence, and discuss if we can build a computer system having true artificial general intelligence.

The class will require programming in python and the use of python libraries (e.g., numpy, sklearn, and pytorch).

Previous knowledge of linear algebra, calculus, and basic probability theory and statistics is a must.