PY 895: Machine Learning for Physicists

Instructor: Pankaj Mehta

Time: TTh 9:30-11:00

Office Hours: TR11-12:30.

Prerequisites: PY 541 or permission of instructor

Description: Machine Learning (ML) is one of the most exciting and dynamic areas of modern research and application. The purpose of this class is to provide an introduction to the core concepts and tools of machine learning in a manner easily understood and intuitive to physicists. The class begins by covering fundamental concepts in ML and modern statistics such as the bias-variance tradeoff, overfitting, regularization, and generalization before moving on to more advanced topics in both supervised and unsupervised learning. Topics covered in the class include ensemble models, deep learning and neural networks, clustering and data visualization, energy-based models (including MaxEnt models and Restricted Boltzmann Machines), and variational methods. Throughout, we emphasize the many natural connections between ML and statistical physics. A notable aspect of the class is the use of Jupyter notebooks to introduce modern ML/statistical packages to readers using physics-inspired datasets (the Ising Model and Monte-Carlo simulations of supersymmetric decays of proton-proton collisions). We will conclude with an extended outlook discussing possible uses of machine learning for furthering our understanding of the physical world as well as open problems in ML where physicists maybe able to contribute.

The course follows are recent review available here:
https://www.sciencedirect.com/science/article/pii/S0370157319300766
(see TOC for list of topics)

Grading: There will also be a final group project in teams of 3. The grade will be based 25% on HWs and 15% on the final project and 60% on class participation/readings.

HW: The class will have a mix-of more theoretical problems and practical programming assignments based on Notebooks from class.

Final Project: The final project must be chosen in consultation with the instructor by the beginning of November. The final project will be a group project (teams of 2-3 members) and must have a significant practical programming/data analysis component.