PY 580: Machine Learning for Physicists

Instructor: Pankaj Mehta
pankajm@bu.edu
SCI 323
617-358-6303

Time: TBD

Office Hours: TBD.

Course Credits: 4 credits

Prerequisites: PY 541 or permission of instructor

Course website: https://physics.bu.edu/~pankajm/PY580.html

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 (including double descent), overfitting, regularization, and generalization before moving on to more advanced topics in both supervised and unsupervised learning. We will focus on discussing possible uses of machine learning for furthering our understanding of the physical world as well as open problems in ML where physicists may be able to contribute. For this reason, the second half of the class will explore new techniques including: self-supervised learning, transformers, and diffusion models. The course will be taught in a "flipped" format. Students will be expected to do reading before the class and then we will discuss results in class.

Books and Other Course Materials

The course follows are recent review available here:
https://www.sciencedirect.com/science/article/pii/S0370157319300766
as well as the corresponding notebooks available here:
https://physics.bu.edu/~pankajm/MLnotebooks.html

The goal of this class is to write a new set of lecture notes that update the review above.

Grading: There will also be a final group project in teams of 3. The grade will be based 50% on HWs and 10% on the final project and 40% on class participation/readings.

HW: The class will have a mix-of more theoretical problems and practical programming assignments in the 20 Notebooks (https://physics.bu.edu/~pankajm/MLnotebooks.html) as well as new assignments based on the topics we cover.

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.

Community of Learning: Class and University Policies

1) Course members are responsibility for ensuring a positive learning environment by being respectful of their fellow students.
2) **Attendance & Absences.** Class discussion and participation are a major part of the course. Class participation is 40% of the grade and absences will lower this score as appropriate except when students receive permission from the instructor or for religious observances.

3) Assignment Completion & Late Work. Due dates are 1 week after assignment. Late HWs will only be accepted with permission from instructor. Students **are encouraged** to use modern AI coding tools for HW and projects.

4) Academic Conduct Statement. All students are expected to follow GRS Academic Conduct Code: http://www.bu.edu/cas/students/graduate/grs-forms-policies-procedures/academic-discipline-procedures/

Weekly assignments: We will work through 1-2 chapter in the review each week. Students are expected to read the corresponding text in the review before class to facilitate discussion. The exact pace of readings depends on discussions/confusions and interest of the students.