Course Overview and Goals

The goal of Machine Learning is to make predictions based on data. "Classical" machine learning has much in common with inference from traditional statistics (e.g., Linear Regression); newer ("Deep Learning") machine learning ventures into territory associated with Artificial Intelligence.

Machine Learning is widely used in many domains outside of Finance. The use within Finance is largely identified (to date) with numerical data and Classical methods. Other domains are blessed with vast quantities of non-numerical data (images, text) that necessitate the use of New methods. But history is not destiny: these alternative types of data will no doubt become integral to Finance. Thus, with an eye on the future, the course will focus on Machine Learning broadly (not just for Finance), frequently drawing on inspiration and examples from other domains.

The first half of the course will focus on "Classical" Machine Learning methods (e.g., Regression, Bayesian methods) while the second half will focus on Deep Learning (e.g., Neural Networks).

The course philosophy is that Machine Learning is an experimental art in which the real learning comes from doing, i.e., conducting experiments. While we will discuss the theory and mathematics underlying Machine Learning, this course will have a heavy computational focus. Homework assignments and projects will all involve programming in Python, using Jupyter notebooks.

By the end of the course the student should have the ability to take a data set and make predictions, using a variety of methods.

Course objectives

- **Teach students the experimental process involved with Machine Learning**
- **Introduce students to Classical and Deep Learning methods**
- **Give students a solid background in the underlying mathematics and theory**
- **Become facile with the standard tools used in Machine Learning: Python, scikit-learn, Jupyter notebooks, TensorFlow.**
- **Give the student the ability to succeed in completing a practical machine learning assignment that a potential employer may use as part of the job interview process**

Course Structure

One lecture per week; each lecture will have both assigned readings and a Jupyter notebook (a combination of descriptive text and executable code) illustrating the concepts. Students will be expected to understand both the readings and the code, and have the ability to experiment/change the notebook's code to deepen their understanding.

Consistent with the philosophy of being an experimental art, the early weeks will be a sprint to get students up to speed with the programming tools involved. Only a small amount of class time

will be spent introducing the programming tools (Python, numpy, Pandas, scikit-learn, TensorFlow, keras) ; students will be expected to acquire these skills via self-directed learning. Students will rapidly gain the ability to experiment.

Course requirements

Graduate standing; Basic probability, statistics, including linear regression, calculus; linear algebra, solid Python programming skills

Course materials

Lectures

Lectures will be distributed in advance of each class as a Jupyter notebook.

Under "Resources" in the course site on NYU Classes, there is a folder "Getting started". There you will find instructions for

1. Setting up your Machine Learning environment, e.g. Jupyter
2. Obtaining lesson materials

Textbooks

1. **Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 2nd Edition (**https://www.oreilly.com/library/view/hands-on-machine-learning/9781492032632/)
2. **Python Data Science Handbook, by Jake VanderPlas**
    1. **Also available as a free, online book https://jakevdp.github.io/PythonDataScienceHandbook**
3. **Deep Learning, by Ian Goodfellow, Yoshua Bengio, and Aaron Courville**
    1. **Also available as a free, online book https://www.deeplearningbook.org/**

Grading

The entirety of the grade will be based on programming assignments:

1. **20% of final grade based on approximately 4 bi-weekly assignments**
2. **30% of final grade based on midterm project**
3. **50% of final grade based on final project**

Disability Disclosure Statement

Academic accommodations are available for students with disabilities. Please contact the **Moses Center for Students with Disabilities** (212-998-4980 or mosescsd@nyu.edu) for further information. Students who are requesting academic accommodations are advised to reach out to the Moses Center as early as possible in the semester for assistance.

Inclusion statement

The NYU Tandon School values an inclusive and equitable environment for all our students. I hope to foster a sense of community in this class and consider it a place where individuals of all backgrounds, beliefs, ethnicities, national origins, gender identities, sexual orientations, religious and political affiliations, and abilities will be treated with respect. It is my intent that all students' learning needs be addressed both in and out of class, and that the diversity that students bring to this class be viewed as a resource, strength and benefit. If this standard is not being upheld, please feel free to speak with me.