TOPICS IN ACTIVE PORTFOLIO MANAGEMENT
FALL 2021
SYLLABUS

1. General Information

Instructors.

(1) E. Jerome Benveniste, PhD (ejb14@nyu.edu)
Office Hours: TBD

(2) Gordon Ritter, PhD (wgr2@nyu.edu)
Office Hours: virtually on Zoom every Thursday 1pm – 2pm, or by appointment.

Textbook. Sutton, Richard S., and Andrew G. Barto. *Reinforcement learning: An introduction.* 2nd edition. MIT press, 2020.

Online edition: http://incompleteideas.net/book/the-book.html

2. Overview

This semester, we will provide an introduction to dynamical programming and reinforcement learning, with a view to applications in portfolio construction and automated trading. We will generally follow the textbook by Sutton & Barto, *Reinforcement Learning: An Introduction*, which we will supplement with more specialized articles as well as hitherto unpublished work on applications to financial markets.

3. Prerequisites

Students should have a solid grounding in undergraduate mathematics, including calculus of several variables, linear algebra, probability, and statistics including multivariate regression. Moreover, there will be programming assignments in python. As a prerequisite, we assume that students have access to python 3.6 or later, and are familiar with the basics of writing simple python programs, loading files, and using numpy to perform standard linear algebra computations such as matrix multiplication and matrix inversion. Some familiarity with the stock market will be helpful, but very little prior finance knowledge is required, as we will cover the parts we need. Prior knowledge of machine learning is likewise *not* a prerequisite.

4. Grading

Your course grade will be based on your homework grades and grades on class projects, including a final data project.

5. Homework

Homework will be posted roughly every two weeks on Classes and solutions should be submitted to Classes on or before the due date.

You are free to discuss the homework problems among yourselves, but you must turn in your own work; students turning in substantially identical work will be penalized.

Programming assignments should be submitted as `.py` files, along with any notebooks or output files which must be in PDF format.

6. Course Outline

(1) Markov Decision Problems
(2) Basics of Dynamic Programming and Bellman's Principle
(3) Application: Multiperiod Portfolio Optimization Problems
(4) Basics of Reinforcement Learning
(5) Monte Carlo Methods
(6) Temporal Difference Learning
(7) Application: Optimal Execution in a Limit-Order Book
(8) Approximate Solutions: On-policy approximation of value functions
(9) Approximate Solutions: Off-policy approximation of value functions
(10) Policy approximation
(11) Application: Microstructure trading with long-term utility function

7. Diversity and Inclusion Statement

The NYU Tandon School values an inclusive and equitable environment for all our students. We 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 our 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 us.