Math 527- Machine Learning for Finance

Course Description for Bulletin:
The purpose of this course is to introduce students to the theory and practice of supervised and reinforcement learning to big data problems in finance. This course emphasizes the various mathematical frameworks for applying machine learning in quantitative finance, such as quantitative risk modeling with kernel learning and optimal investment with reinforcement learning. Neural networks are used to implement many of these mathematical frameworks in finance using real market data. (3-0-3)

Enrollment:
Elective for AMAT and other majors, such as Masters in Mathematical and Finance and Masters of Science in Finance.

Textbook(s):
Hastie, T., Tibshirani, R. & Friedman, J. H. (2009). The elements of statistical learning: data mining, inference, and prediction. 2nd ed. New York: Springer. Open Access.
Rasmussen, C. E. & Williams, C. K. I. (2006). Gaussian Processes for Machine Learning. MIT Press, Cambridge, MA, USA. Open Access.
Sutton, R.S. & Barto, A.G. (1998). Introduction to Reinforcement Learning (2nd ed.). MIT Press, Cambridge, MA, USA. Open Access,

Other required material:
Python

Prerequisites:
MATH 475 or equivalent at the consent of the instructor.

Objectives:
1. Students will learn the fundamentals of statistical learning theory
2. Students will learn the basics of Neural Networks for point estimation from financial data
3. Students will learn the basics of Gaussian Processes for financial risk modeling
4. Students will learn the basics of Reinforcement Learning for optimal stochastic control problems in finance
5. Students will gain hands on experience working with real market data and implementing machine learning methods in Python (or another similar software package such as R or Matlab).

Lecture schedule:
3 50 minutes (or 2 75 minutes) lectures per week

Course Outline:

Lectures	**Hours**
1. Fundamentals of statistical learning theory	8
a. Convergence and learnability	
b. Kullback-Leibler Information	
c. Model selection and the bias variance trade-off	
d. Cross-validation	
e. Regularization	
f. Generative vs Discriminative models	

2. Neural Networks 6

a. The Perceptron
b. Feed-Forward Neural Networks
c. Back-propagation and stochastic gradient descent
d. Regularization and drop-out
e. Application to investment management

3. Recurrent Neural Networks for Econometrics 6

a. Econometric models for time series prediction
b. Filtering for time series
c. Recurrence in neural networks and relation to ARIMA
d. Gated Recurrent Unit (GRU) and Long Short-Term Memory (LSTM) networks as a dynamic econometrics model
e. Application to forecasting models used in algorithmic trading

4. Bayesian Machine Learning 6

a. Bayesian inference, filtering and prediction
b. Kernel learning
c. Gaussian processes (GPs)
d. Multi-GPs
e. GPs for derivative pricing and risk management

5. Introduction to Reinforcement Learning 6

a. Markov Decision Processes (MDPs), with examples in finance
b. Partially Observable MDPs
c. Value and action-value functions
d. Bellman optimality
e. Policy iteration
f. Q-learning
g. Exploitation versus exploration

6. Introduction to Inverse Reinforcement Learning 3

a. Imitation learning
b. Constraints based inverse reinforcement learning
c. Maximum entropy inverse reinforcement learning
d. Applications in algorithmic trading

7. Reinforcement Learning for Investment Management 4

a. Merton's optimal consumption
b. Optimal hedging strategies
c. Robo-advisors for optimal allocation
d. Learning an investor's preferences

8. Machine Learning for Risk Management 4

a. Model-free derivative pricing
b. Value-at-risk estimation with GPs
c. Credit Value Adjustment with GPs

Assessment:

Homework	10-30%
Computer Programs/Project	10-20%
Quizzes/Tests	20-50%
Final Exam	30-50%

Syllabus prepared by:

Tomasz R. Bielecki, Igor Cialenco and Matthew F. Dixon

Date:

Nov.28, 2018