FRE-GY 9073, Stochastic Systems and Modern Machine Learning Theory

Instructor Information

- Renyuan Xu, Assistant Professor, https://renyuanxu.github.io/
- Instructor office address: 1 MetroTech Center, 10th Floor
- Instructor office hours: 11:00-12:00am on Tuesday
- Instructor's email address: rx2364@nyu.edu

Course Information

- Course time: 8:00-10:30am on Tuesday
- Face-to-face class meeting days and times: in-person unless otherwise stated
- Classroom number and building: Rogers Hall, Room 216
- Virtual (online) meeting days and times, if any: October 22

Co-requisite or prerequisite

This course is designed for PhD students as a special topics course. Master students with a strong mathematical background and basic understanding of machine learning are also encouraged to enroll. The following courses are recommended as prerequisite but not required: ECE-GY 6253 Linear Systems and ECE-GY 6233 System Optimization Method.

Course Overview and Goals

This course provides a comprehensive introduction to the mathematical foundations of stochastic systems and stochastic controls in discrete time. The course also explores the applications of stochastic controls in the theoretical developments of modern machine learning, including reinforcement learning, generative diffusion models, deep neural network training, and fine-tuning large language models.

Upon completion of this course, students will be able to:

- **Model Real-World Problems**: Formulate real-world decision-making problems using stochastic control framework.
- **Apply Analytical Tools**: Solve control problems with appropriate analytical tools.

- **Understand Neural Network Training**: Gain insights into neural network training and fine-tuning model parameters from the perspective of stochastic controls.
- **Design Machine Learning Algorithms**: Develop efficient and scalable machine-learning algorithms using control theory principles.

Course Requirements

Class Participation

In-person participation is encouraged but not mandatory.

Assignments

There are three homework assignments (two before the mid-term exam and one after the mid-term exam).

Tests & Quizzes

The mid-term exam is in person during the lecture time in the week of Oct 28th. Sample exam sheets will be provided in early October.

There is no final exam. The final project is due on Dec 20th (the last day of the final exam week). The final project can take one of the following forms:

1. Theoretical developments,
2. A survey of several papers on an emerging direction, or
3. Implementation of course-related algorithms for a practical application.

All projects should be closely related to the course content

Assigned Readings

Readings are not mandatory and they will be provided every week to support the course materials.

Grading of Assignments

The grade for this course will be determined according to the following formula:

Assignments/Activities	% of Final Grade
Midterm Exam	40%

Homework Assignments	30%
Final Project	30%

Letter Grades

Letter grades for the entire course will be assigned as follows:

F	0
C	70
C+	76.67
B-	80
B	83.33
A-	90
A	93.33

View Grades

Grades for each homework assignment, midterm exam, and final project will all be available on NYU Brightspace.

Course Schedule

Topics and Assignments

Week/Date	Topic	Reading	Assignment Due
[Week 1, Sept 3]	Course introduction, Markov decision process (MDP): (1) state dynamics, reward process	[Puterman] Ch 1	HW 1 available

	(2) finite and infinite horizon problem		
[Week 2, Sept 17]	Finite-horizon MDP: (1) dynamic programming (2) optimality of Markovian policy (3) Bellman equation (3) value iteration algorithm	[Puterman] Ch 4	
[Week 3, Sept 17]	Infinite Horizon MDP: (1) Bellman equation (2) linear programming approach (3) value iteration and policy iteration	[Puterman] Ch 6	
[Week 4, Oct 1]	Financial Applications: (1) optimal investment and consumption problem (2) optimal execution	[Hambly & Xu & Yang]	HW 1 Due (before class) HW 2 Available
[Week 5, Oct 8]	Introduction to reinforcement learning (RL): (1) set-up and evaluation criterion (2) Monte Carlo (MC) estimation (3) policy evaluation using MC	[Barto & Sutton] Ch 5	
[Week 6, Oct 15]	Value-based method: (1) TD Learning (2) Q learning (3) convergence of stochastic approximation (SA) method	[Barto & Sutton] Ch 6	
[Week 7, Oct 22]	Foundations of gradient-based methods:	[Barto & Sutton] Ch 13	HW 2 Due (before class)

	(1) gradient descent in optimization (2) policy gradient theorem (3) actor-critic (AC) method (4) trust region policy optimization (TRPO)	[Szepesvári] Ch 4.4	Mid-term practice sheet and solution available
[Week 8, Oct 29]	Mid-term exam (in class)		HW 3 Available
[Week 9, Nov 5]	Introduction to neural network and deep learning	[François-Lavet et al] Ch 2.2	
[Week 10, Nov 12]	Deep neural network (DNN) as a dynamic system: (1) information propagation inside DNN (2) stability analysis & implication	[Liu &Theodorou]	HW 3 Due (before class)
[Week 11, Nov 19]	Training DNN with optimal control: (1) Mean-field optimal control derivation (2) Mean-field Hamilton-Jacobi-Bellman equation	[Liu &Theodorou]	Final project proposal due (one-page description)
[Week 12, Nov 26]	Introduction to large-language models (LLMs) and generative diffusion models	[Uehara et al.] [Rafailov et al.]	
[Week 13, Dec 3]	Fine-tuning LLM and generative diffusion models as entropy-regularized optimal control	[Uehara et al.] [Rafailov et al.]	

[Week 14, Dec 10]	Guest lectures: (1)Reinforcement learning in finance (2)Generative diffusion models for financial time series		

Tests and Quizzes

- Midterm exam: in class on October 28th

Course Materials

Required Textbooks & Materials

No requested textbooks. The following books and tutorials are useful:

- [Puterman] Markov Decision Processes: Discrete Stochastic Dynamic Programming, by Martin L. Puterman, John Wiley & Sons, 2014.
- [Barto & Sutton] Reinforcement Learning: An Introduction, by Andrew Barto and Richard S. Sutton, MIT press, 2018.
- [Szepesvári] Algorithms for Reinforcement Learning, by Csaba Szepesvári, Synthesis lectures on artificial intelligence and machine learning 4.1 (2010): 1-103.
- [François-Lavet et al] "An introduction to deep reinforcement learning." Foundations and Trends® in Machine Learning 11.3-4 (2018): 219-354.
- [Hambly & Xu & Yang] Hambly, Ben M., Renyuan Xu, and Huining Yang. "Recent Advances in Reinforcement Learning in Finance." Mathematical Finance 33.3 (2023): 437-503.
- [Liu &Theodorou] Liu, G. H., and Evangelos A. Theodorou. "Deep learning theory review: An optimal control and dynamical systems perspective. arXiv." Learning (2019).
- [Uehara et al.] Uehara, Masatoshi, et al. "Fine-Tuning of Continuous-Time Diffusion Models as Entropy-Regularized Control." arXiv preprint arXiv:2402.15194 (2024).
- [Rafailov et al.] Rafailov, Rafael, et al. "Direct preference optimization: Your language model is secretly a reward model." Advances in Neural Information Processing Systems 36 (2024).

Resources

- **Access your course materials**: NYU Brightspace
- **Databases, journal articles, and more**: Bern Dibner Library (library.nyu.edu)
  NYU Virtual Business Library (guides.nyu.edu/vbl)
- **Obtain 24/7 technology assistance:** Tandon IT Help Desk (soehelpdesk@nyu.edu, 646.997.3123)
  NYU IT Service Desk (AskIT@nyu.edu, 212-998-3333)

Policies

Academic Misconduct

A. Introduction: The School of Engineering encourages academic excellence in an environment that promotes honesty, integrity, and fairness, and students at the School of Engineering are expected to exhibit those qualities in their academic work. It is through the process of submitting their own work and receiving honest feedback on that work that students may progress academically. Any act of academic dishonesty is seen as an attack upon the School and will not be tolerated. Furthermore, those who breach the School's rules on academic integrity will be sanctioned under this Policy. Students are responsible for familiarizing themselves with the School's Policy on Academic Misconduct.
B. Definition: Academic dishonesty may include misrepresentation, deception, dishonesty, or any act of falsification committed by a student to influence a grade or other academic evaluation. Academic dishonesty also includes intentionally damaging the academic work of others or assisting other students in acts of dishonesty. Common examples of academically dishonest behavior include, but are not limited to, the following:
    1. Cheating: intentionally using or attempting to use unauthorized notes, books, electronic media, or electronic communications in an exam; talking with fellow students or looking at another person's work during an exam; submitting work prepared in advance for an in-class examination; having someone take an exam for you or taking an exam for someone else; violating other rules governing the administration of examinations.
    2. Fabrication: including but not limited to, falsifying experimental data and/or citations.
    3. Plagiarism: Intentionally or knowingly representing the words or ideas of another as one's own in any academic exercise; failure to attribute direct quotations, paraphrases, or borrowed facts or information.
    4. Unauthorized collaboration: working together on work that was meant to be done individually.

5. Duplicating work: presenting for grading the same work for more than one project or in more than one class, unless express and prior permission have been received from the course instructor(s) or research adviser involved.
6. Forgery: altering any academic document, including, but not limited to, academic records, admissions materials, or medical excuses.

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.