

New York University Tandon School of Engineering
Computer Science and Engineering

CS-GY 6763 (CS-UY 3943)

Algorithmic Machine Learning and Data Science, Fall 2021

Professor Christopher Musco
Tuesdays 2:00pm-4:30pm, Rogers Hall, Room 707

CONTACT INFORMATION

Email: cmusco@nyu.edu.

Office: 370 Jay St., #1105 (Floor 11).

Virtual Zoom Office: <https://nyu.zoom.us/my/cmusco>

Ed Discussion Forum (for questions): <https://edstem.org/us/courses/13524/discussion/>

Office Hours: Weekly times TBA.

Course Webpage: <https://www.chrismusco.com/amlds2021/>

COURSE DESCRIPTION

This course gives a behind-the-scenes look into the algorithms and computational methods that make machine learning and data science work at large scale. How does a service like Shazam match a sound clip to a library of 10 million songs in under a second? How do scientists find patterns in terabytes of genetic data? How can we efficiently train neural networks with millions of parameters on millions of labeled images? We will address these questions and others by studying advanced algorithmic techniques like randomization, approximation, sketching, continuous optimization, spectral methods, and Fourier methods.

COURSE PREREQUISITES

We require a previous course in machine learning (for example, CS-UY 4563, CS-GY 6923, or ECE-GY 6143), a previous course in algorithm design and analysis (for example, CS-UY 2413, CS-GY 6033, or CS-GY 6043), and a previous course in linear algebra (for example, MA-UY 2034, 3044, or 3054). Experience with probability and random variables is also necessary. In general, *modern algorithm design uses a lot of math!* This is partially what makes it such an interesting subject to study. In this course, we rely most heavily on probability and linear algebra, but we will also learn some approximation theory, Fourier analysis, and high dimensional geometry (a student favorite in the past). Be prepared for theoretically rigorous problem solving, which may require you to refresh topics you covered in prior courses. In particular, you should know:

1. **Probability:** Random variables, discrete and continuous probability distributions, expectation, variance, covariance, independence, correlation, conditional and joint probability, Gaussian random variables, law-of-large-numbers.
2. **Linear Algebra:** Vector inner and outer products, matrix-vector and matrix-matrix multiplication, vector norms (e.g., Euclidean), matrix norms (e.g., Frobenius, operator), triangle inequality, Cauchy-Schwarz, solving systems of linear equations, linear subspaces, linear independence, projection, matrix rank, column/row span, null space, orthogonal matrices, basics of eigenvectors and eigenvalues.

Finally, you do a small amount of programming in this course. You can use whatever language you like (most students use MATLAB or Python) and will need to use basic linear algebra packages in these languages.

COURSE STRUCTURE AND GRADING

You will complete short weekly quizzes, work on problem sets at home, take one in-class midterm, and complete either a final exam or a final project. Details follow:

Class participation (10% of grade). This grade captures how much you contribute to your own learning and that of your peers. Since different students have different styles, there are *many ways* to earn full credit for this part of the course. You can actively participate in class. You can ask good questions, or answer those of your peers, on the class Ed forum. You can contribute to the reading group. You can attend and actively contribute to office hours.

Weekly quizzes (10% of grade). Every week a short Google Forms quiz will be posted after lecture as a way for both of us to assess how well you understood the class material. While designed to require no more than ~ 15 minutes, you can take as long as you want for the quiz. It just needs to be completed before the following week's lecture.

Written problem sets (45% of grade). These assignments are completed at home and involve the analysis and application of methods learned in class, with occasional programming exercises to further explore lecture content. I expect most of your learning to occur while working on these exercises, and investing time on them is the best way to prepare for the midterm exam and final. Assignments and their due dates will be posted on the course webpage. Late assignments will only be accepted if there are extenuating circumstances and you have obtained prior permission from me.

Midterm exam (15% of grade). In person for those who can attend, and we will make remote accommodations for others. The exam will take 1 hour, and we will have a short lecture afterwards. You will be allowed a cheat-sheet (a two-sided piece of paper with whatever information you want on it).

Choice of final project or final exam (20% of grade). You have the option of completing either a final project or final exam for this class. Typically preparing for the final exam is less work, so I suggest this as the "default option". However, any student interested in research, and especially in topics related to this class, might want to choose the project instead. The project can either be applied or theoretical, but must be related to algorithmic aspects of data science or machine learning. Typically, papers will center around 1 or 2 recent research papers. Expectations and a list of potential topics and papers to read will be posted on the course website.

Reading Group. It is an exciting time for research at the intersection of algorithm design and the data sciences. Many of the ideas covered in this course are still the subject of active research. Starting midway through the semester, we will hold a weekly reading group to discuss recent papers every week for 1.5 hours. The reading group is *optional for students taking the final exam* but *required for those completing a final project*. Students completing projects (or others who volunteer) will sign-up to present a paper related to their project during one week. Students will also be required to prepare discussion questions in advance of at least two of their peers' presentations.

LECTURES

There is one 2.5 hour class meeting per week, divided by a 15 minute break. The lecture (and virtual whiteboard) will be video recorded for those participating remotely. In accordance with NYU guidelines, you must wear a mask covering your face and nose whenever you are attending lecture.

COURSE POLICIES

Written problem sets: Turned in via NYU Brightspace/Gradescope. While not required, I encourage students to prepare written problem sets in LaTeX or Markdown (with math support.) Students will receive 10% extra credit on Problem Set 1 for preparing it in LaTeX or Markdown.

Homework collaboration policy:

- *Discussion of high-level ideas for solving problems sets or labs is allowed, but all solutions and any code must be written up independently.* This reflects the reality that algorithms and machine learning research are rarely done alone. However, any researcher or practitioner must be able to communicate and work through the details of a solution individually.
- *Students must name any collaborators they discussed problems with at the top of each assignment (list at the top of each problem separately).*

- *Do not write solutions or code in parallel with other students.* If problem ideas are discussed, solutions should be written or implemented at a later time, individually. **I take this policy very seriously. Do not paraphrase, change variables, or in any way copy another students solution.**

Grade changes: Fair grading is important to me. If you feel that you were incorrectly docked points on an assignment or exam, please let me know – it is very possible I or the TAs misunderstood your work. Do not wait until the end of the semester to bring up grading issues! If you notice something off, please ask ASAP.

Questions about performance: If you are struggling in the class, contact me as soon as possible so that we can discuss what’s not working, and how we can work together to ensure you are mastering the material. It is difficult to address questions about performance at the end of the semester, or after final grades are submitted: by Tandon policy no extra-credit or makeup work can be used to improve a students grade once the semester closes. So if you are worried about your grade, seek help from me *early*.

READINGS

There is no textbook to purchase. Many of the topics covered are new, and not sufficiently addressed in existing textbooks. Course material will consist of my written lecture notes, as well as assorted online resources, including papers, notes from other courses, and publicly available surveys. This year I will also be trying to post some type-written lecture notes. Please refer to the course webpage (<https://www.chrismusco.com/amlds2021/>) before and after lecture to obtain links.

COURSE SCHEDULE

The following schedule is tentative and subject to change.

THE POWER OF RANDOMNESS

1. 9/7, Random variables and concentration, Markov’s inequality, applications
2. 9/14, Chebyshev inequality, exponential tail bounds (Chernoff + Bernstein), and applications
3. 9/21, High dimensional geometry, the Johnson-Lindenstrauss Lemma
4. 9/28, Nearest neighbor search, locality sensitive hashing

OPTIMIZATION

5. 10/5, Convexity in machine learning, gradient descent and projected gradient descent
10/12, NO CLASS – MONDAY SCHEDULE DUE TO HOLIDAY.
6. 10/19, Online and stochastic gradient descent, coordinate descent, preconditioning
7. 10/26, Constrained optimization, center of gravity method, linear programming
8. 11/2, Discrete optimization: LP relaxation, greedy methods, submodularity

SPECTRAL METHODS AND LINEAR ALGEBRA

9. 11/9, Singular value decomposition, Krylov methods.
10. 11/16, Spectral graph theory, clustering, stochastic block models.
11. 11/23, Randomized linear algebra, sketching for linear regression, ϵ -nets arguments.

FOURIER METHODS

12. 11/30, The Fourier transform and applications, random Fourier features
13. 12/7, Compressed sensing, the restricted isometry property, basis pursuit
14. 12/14, Vote on a topic related to my recent research!

12/?? Final Exam Schedule, TBD

COURSE OBJECTIVES

1. Students will build experience with the most common algorithmic tools used in machine learning, including optimization, relaxation, spectral methods, Fourier methods, and Monte Carlo algorithms. They will also strengthen their understanding of the mathematical concepts behind these tools.
2. Through problem sets and in-class assignments, students will practice applying, combining, and modifying the methods learned in lectures to specific problems in machine learning and data analysis. The goal is to prepare students to use these tools in industrial or academic positions where they are developing computationally efficient machine learning methods.
3. Students will learn how theoretical questions (involving asymptotic complexity, communication complexity, convergence rate, approximation quality, etc.) can guide the design of new algorithmic methods in machine learning. We will practice framing and rigorously answering such questions.
4. The course will prepare students to contribute to research projects on computational methods for machine learning and data science.
5. Through assigned papers and discussions, students will improve their ability to read, summarize, and extract value from contemporary research papers in machine learning and data mining conferences like NeurIPS, ICML, KDD, ICLR and AAAI.

ADDITIONAL INFORMATION

MOSES CENTER STATEMENT OF DISABILITY.

If you are student with a disability who is requesting accommodations, please contact New York University's Moses Center for Students with Disabilities (CSD) at 212-998-4980 or mosescsd@nyu.edu. You must be registered with CSD to receive accommodations. Information about the Moses Center can be found [here](#).

NYU SCHOOL OF ENGINEERING POLICIES AND PROCEDURES ON ACADEMIC MISCONDUCT.

The complete Student Code of Conduct can be found [here](#).

- 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 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 has 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.

NYU SCHOOL OF ENGINEERING POLICIES AND PROCEDURES ON EXCUSED ABSENCES.

The complete policy can be found [here](#).

- A. Introduction: An absence can be excused if you have missed no more than 10 days of school. If an illness or special circumstance has caused you to miss more than two weeks of school, please refer to the section labeled Medical Leave of Absence.
- B. Students may request special accommodations for an absence to be excused in the following cases:
 1. Medical reasons
 2. Death in immediate family
 3. Personal qualified emergencies (documentation must be provided)
 4. Religious Expression or Practice

Deanna Rayment, deanna.rayment@nyu.edu, is the Coordinator of Student Advocacy, Compliance and Student Affairs and handles excused absences. She is located in 5 MTC, LC240C and can assist you should it become necessary.

NYU SCHOOL OF ENGINEERING ACADEMIC CALENDAR

The full calendar can be found [here](#). Please pay attention to notable dates such as Add/Drop, Withdrawal, etc. For confirmation of dates or further information, please contact Susana Garcia: sgarcia@nyu.edu.