

Course Information

Instructor: *Alex Andoni*

1 Basic Information

Lectures:

- Time: MW 2:40-3:55.
- Location: Mudd 1127.

Instructor:

- **Alex Andoni** (andoni@cs.columbia.edu)

Teaching Assistants:

- Nabiha Biviji (nyb2112@columbia.edu)
- Caden Lin (c14319@columbia.edu)
- Szymon Snoeck (sgs2179@columbia.edu)

Website: There is no textbook, but there's a website with scribes and my notes from a previous edition of the class, as well as links to additional resources:

<http://www.cs.columbia.edu/~andoni/advancedS26/materials.html>

Office hours: see calendar linked from Courseworks (requires LionMail login).

Self evaluation test: to confirm for yourself that you have the right background, you should complete the self-evaluation test asap (definitely, by the end of the first week). It will help you identify potential parts to brush up before the class starts. See the Courseworks for the test and the solution key.

Courseworks: Class announcements, including homework assignments will be posted on Courseworks.

Edstem: There is also a Edstem forum setup for the class (accessible through the Courseworks). You are encouraged to discuss class lectures and related topics, as well as ask questions or clarifications. (But you *cannot* discuss homework solutions on Edstem.)

2 Course Goals

The goals of the class are to introduce you to classic and modern algorithmic ideas that are central to many areas of Computer Science. The focus is on most powerful paradigms and techniques of how to design algorithms, and how to measure their efficiency. The intent is to be broad, covering a diversity of algorithmic techniques, rather than be deep. The covered topics have all been implemented and are widely used in industry. At the end of the class, expect to:

- know a diversity of classic algorithmic tools,
- be able to design algorithms for computational problems arising in your area,
- be able to read research-level papers in the field of algorithms.

Tentative topics to be covered (one bullet point corresponds to approx 0.5—2 lectures):

- **Hashing:**
 - universal hashing, perfect hashing
- **Sketching/Streaming:**
 - approximate counting, Morris algorithm
 - distinct elements count, impossibility results
 - heavy hitters, countSketch algorithm
 - frequency moments, dimension reduction.
- **Nearest Neighbor Search:**
 - sketching, nearest neighbor search
 - Locality Sensitive Hashing scheme.
- **Spectral Graph Theory:**
 - linear algebra overview, and graph Laplacian
 - spectrum of graph Laplacians
 - graph cuts, Cheeger inequality
 - expander graphs, applications.
- **Optimization / Linear Programming:**
 - introduction to Linear Programming (LP), structure of optima
 - strong duality, complementary slackness
 - simplex and ellipsoid algorithms
 - gradient descent, iterative methods
 - interior point method
 - multiplicative weights update, online algorithms.
- **Models for large-scale computation:**
 - external memory, cache-oblivious, GPU algorithms
 - parallel algorithms.
- **Other tentative topics:**
 - Algorithms for Transformers/LLMs;
 - Algorithms to beat exhaustive search algorithms (e.g., 3SAT).

3 Prerequisites

First and foremost, mathematical maturity is a must: the class is based on theoretical ideas and is proof-heavy. You are expected to be able to read and write formal mathematical proofs. Furthermore, some familiarity with algorithms and randomness will be assumed as well. COMS 4231 (Analysis of Algorithms) or equivalent is recommended, but not required if you have solid math background.

Here is a rough list of math/CS topics that you are expected to know or have background in:

- basic linear algebra (eigenvalues, eigenvectors);
- basics of probability theory (linearity of expectation, variance, Markov bound);
- asymptotic analysis of algorithms, runtime analysis;
- hashing, or binary search trees;
- graphs.

4 Evaluation and Grading

Your grade is based on the following three components:

- 4 homeworks: 20%;
- 4 short in-class quizzes (≤ 10 min at the start of the class, lowest grade dropped): 20%;
- Midterm (in class): 30%;
- Project: 30%.

5 Homeworks

Homeworks will be assigned roughly every two weeks and will be posted on Courseworks. They will be due in class on their due date before the lecture starts. Please follow the Homework Submission Guidelines below.

Late policy. You have a default 5 days of extension (fractions of a day are rounded up), over all the homeworks. Once you've used up the 5 days, late homeworks will be penalized at the rate of 10%, additively, per late day or part thereof (i.e. fractions of a day are rounded up), for up to 7 days. To allow us to distribute the solutions in a timely fashion, homeworks submitted more than 7 days after the deadline will not be accepted. Exceptions will be made only for exceptional unforeseen circumstances (e.g., serious illness), in which case you will need to provide some additional documentation (e.g., doctor's note).

You are strongly encouraged to start working on the homeworks *early*: some problems may require you to sit on the problem for a while before you get your “aha” moment. Starting early also gives you time to ask questions and make effective use of the office hours of the teaching staff.

Writing up solutions: precise and formal proofs. The goal of the class, in part, is for you to learn to reason about algorithms, precisely describe them, and formally prove claims about their correctness and

performance. Hence, it is important that you write up your assignments *clearly, precisely, and concisely*. Legibility of your write-up will be an important factor in its grading. When writing up (algorithmic) solutions, keep in mind the following:

- The best way for you to convey an algorithm is by using plain English description. A worked example can also help; but revert to pseudocode only if necessary. Generally, give enough details to clearly present your solution, but not so many that the main ideas are obscured.
- The analysis of the algorithm has to include both 1) proof of correctness, and 2) upper bound on performance (usually runtime, but sometimes space as well).
- You are encouraged (but not required) to type up your solutions using LaTeX (e.g., using overleaf).

Note that our lectures will generally be at a slightly lower level of formalism, in the interest of time.

Homework Submission Guidelines:

- Submit your homeworks electronically via GradeScope. You may write solution by hand, in which case you should either scan or photograph your solutions.
- Homeworks are due on the specified due date 40minutes before the class starts (2pm).

6 Collaboration and Academic Honesty

Collaboration: you are permitted to discuss the homework *assignments*. If you collaborate, you must write the solutions *individually* (without looking at anybody else's solutions), and acknowledge anyone with whom you have discussed the problems. It will be considered an honor code violation to consult solutions from previous years, from the web or elsewhere, in the event that homework problems have been previously assigned or solutions are available elsewhere.

You are expected to abide by the policies of academic honesty. The CS department web page lists the department's academic honesty policies: <http://www.cs.columbia.edu/education/honesty>.

7 Use of AI/LLMs

You are allowed to use AI/LLMs as a resource in the class. However, you should use it wisely. Majority of the class grade requires you to demonstrate essential understanding of class without any external aid (midterm, quizzes, project posters). So you shouldn't become too dependent on the LLM.

Treat the problem sets as a low-stakes way of getting practice with thinking and learning the material. We can't stop you from applying AI/LLMs to blindly solve the problem sets, but your grade may suffer on the other assignments.

Here are suggestions for how to use AI/LLMs in a smart way, especially for the problem sets.

Try to solve the problem yourself. Identify the relevant lecture material, and try to figure out an approach. Talk with a classmate (per collaboration above).

If you end up using the AI, don't just copy-paste the question/answer. First, you can enter a prompt like the following: *I'm a student trying to learn algorithms with theoretical guarantees, specifically [topic], and I'd like you to be a helpful study partner who doesn't just give me the answer, but helps me learn and solve problems myself.*

Then, you can use prompts like: “*Explain the analysis of the perfect hashing algorithm, like I’m a beginner.*”

“*I’m trying to solve the following problem: [insert general description]. However don’t tell me the answer. But my general approach is...*”

“*Here’s my derivation. Can you give me a hint as to the first line that breaks?*”

You must understand and verify everything you submit. A way to test whether you understand: could you explain your answers to a peer without relying on any external aid?

Ultimately, you’re taking this class because you want to learn how to design and think of modern algorithms. The AI already knows the material pretty well, so don’t let it rob you of the opportunity to learn.

8 Final Project

In the final project you will delve into a particular topic in more detail in a team of your own. The final projects can be of the following types:

- Research/Exploratory: investigate a research topic on your own (eg, develop an algorithm, and prove its properties; or prove an impossibility result). It may be more applied: e.g., perhaps in your area, certain theoretical algorithms can be modified to have even better performance, due to special properties of the datasets, etc.
- Implementation: implement some of the algorithms from the class (or from other theoretical literature), and perhaps apply to your area of interest/expertise, using real-world datasets. One aspect of such projects will be a comparison among a few algorithms.

Teams: you are allowed to have a team of 2–4 people in total per team. Single-person teams are possible but discouraged and need special permission from the instructor¹

You are encouraged to find a team early, and discuss with the instructor the potential topics. There will be a project proposal due. It will be of about 1 page long (exact details tbd). There will also be a poster presentation.

Topic: the topic of your project must be within the scope of Theoretical Computer Science, and preferably algorithmic. In particular, the focus is on algorithms with provable guarantees (for the implementation type, you may compare such theoretical guarantees with heuristics though). More details and suggestions will be given later in the class.

9 Callendar (tentative)

Below is the schedule of assignments. Information regarding what each lecture covers will be periodically updated on the course website.

¹The reason is that the topics are hard, and having a collaborating partner/s will qualitatively improve your understanding of the topic and experience overall.

Lecture	Date		HW out	HW/P/Q due
1	1/21			
2	1/26		HW1 out	
3	1/28			
4	2/2			
5	2/4		HW2 out	HW1 due
6	2/9			Quiz 1
7	2/11			
8	2/16			
9	2/18		HW3 out	HW2 due
10	2/23			Quiz 2
11	2/25			
12	3/2			
13	3/4			HW3 due
14	3/9			
15	3/11	Midterm		
	3/16	NO CLASSES: Spring Break		
	3/18	NO CLASSES: Spring Break		
16	3/23			
17	3/25			
18	3/30			
19	4/1		HW4 out	Project proposal due
20	4/6			
21	4/8			
22	4/13			Quiz 3
23	4/15			
24	4/20			HW4 due
25	4/22			Quiz 4
26	4/27			
27	4/29			
	5/1	Final Project posters/presentations		
	5/11	Final projects write-ups due		Project due