

COMPSCI 690RA: Randomized Algorithms and Probabilistic Data Analysis


Prof. Cameron Musco

University of Massachusetts Amherst. Spring 2022.

Lecture 1

Motivation

Randomized algorithms take steps that depend on both their inputs and on the outcomes of random coin flips. I.e., they are algorithms that make random decisions during their execution.




This is sort of weird when you think about it – the algorithm can't do anything better than just randomly guess?

Motivation

In many settings randomized algorithms give big advantages over deterministic ones:


- Can be much faster than the best known deterministic algorithms (polynomial identity testing, approximation algorithms, linear algebraic computation and data analysis)
- Often very simple and elegant (randomized quicksort, Karger's min-cut algorithm)
- In many cases, by using random sampling and related ideas, they can achieve things that are simply impossible for deterministic algorithms (sublinear time algorithms, efficient communication protocols, small-space streaming algorithms)

Section 1: Probability Foundations & Random Hashing





- Review of probability tools and concentration inequalities, that will be used throughout the course.
- Applications to the analysis of random hashing algorithms. E.g., analysis of chaining, linear probing, and cuckoo hashing.
- Hashing for efficient communication and lookup. Fingerprints, pattern matching, communication complexity.
- Streaming and distributed graph algorithms. ℓ_0 sampling and graph sketching.

Section 2: Random Sketching and Randomized Numerical Linear Algebra (RandNLA)



- Sampling and sketching for approximate matrix multiplication, low-rank approximation, and trace estimation.
- Dimensionality reduction. The Johnson-Lindenstrauss lemma, subspace embedding, and sketching for linear regression.
- Importance sampling and leverage scores. Connections to matrix concentration bounds and spectral graph theory.

Section 3: Markov Chains and Other Topics



- Introduction to Markov chains and basic tools for their analysis.
- Algorithms based on Markov chains: randomized methods for 2-SAT and 3-SAT, Markov Chain Monte Carlo methods (MCMC).
- Convex relaxation and randomized rounding for approximating combinatorially hard problems.
- Mathematical proofs via randomized algorithms: the probabilistic method and the Lovász Local Lemma.

Style of the Course

This is a **theory** course.

- Assignments will emphasize algorithm design, correctness proofs, and asymptotic analysis (minimal coding).
- The homework will be challenging. You will design and analyze algorithms using the tools taught in class, but the solutions will require significant thought and novel ideas.
- A strong algorithms and mathematical background **are required** (particularly in linear algebra and probability).
- If you are an undergraduate or Master's student I would not recommend taking this course unless you have previously taken either 514 or 611 and done very well – the course will be a significant step up in difficulty from 514.

Who Should Take this Course?

- You are a Ph.D. student that wants to do research in algorithms/theory.
- You are a Ph.D. student in another area that wants to apply randomized algorithms or probabilistic analysis in your work.
- You are an undergrad/Master's student excited about algorithms, who potentially wants to pursue a Ph.D.
- If you are unsure if you are prepared for the course, or if it will fulfill your goals for taking it, shoot me an email or Piazza message.

Course Logistics

See course webpage for logistics, policies, lecture notes, assignments, etc.

<http://people.cs.umass.edu/~cmusco/CS690RAS22/>

Visit Moodle or my homepage for this link if you lose it.

Professor: Cameron Musco

- Email: cmusco@cs.umass.edu
- Office Hours: Weds, 2pm-3pm (directly after class). CS 234 (directly upstairs). Starting next week.
- I encourage you to come as regularly as possible to ask questions/work together on practice problems.
- If you need to chat individually, please email to set up a time.

TA: Weronika Nguyen

- Email: thuytrangngu@cs.umass.edu
- Office Hours: Mon, 3:30pm-4:30pm. CS 207.

Piazza and Participation

We will use Piazza for class discussion and questions.

- See website for link to sign up.

You may earn up to 5% extra credit for participation.

- Asking good clarifying questions and answering questions during the lecture or on Piazza.
- Answering other students' or instructor questions on Piazza.
- Posting helpful/interesting links on Piazza, e.g., resources that cover class material, research articles related to the topics covered in class, etc.

Textbooks and Materials

I will post optional readings and references for each class. A lot of the content is covered in the following two textbooks which may be helpful for reference:

- *Probability and Computing*, by Mitzenmacher and Upfal
- *Randomized Algorithms*, by Motwani and Raghavan

Lecture notes will be posted before each class, and annotated notes posted after class. Recordings should be available via Echo360 – may take a couple of classes to work out any kinks.

Homework

We will have 4 problem sets, which you may complete in **groups of up to 3 students**.

- We strongly encourage working in groups, as it will make completing the problem sets much easier/more educational.
- You **should not** simply split up the problem set and complete different parts independently.
- Collaboration with students outside your group is limited to discussion at a high level. You may not work through problems in detail or write up solutions together.
- See Piazza for a thread to help you organize groups.

Problem set submissions will be via Gradescope.

- See website for a link to join and entry code.

Weekly Quizzes

I will release a multiple choice quiz in Moodle each Wednesday after lecture, due the next Tuesday at 8pm.

- Designed as a check-in that you are following the material, and to help me make adjustments as needed.
- Will take around 15-30 minutes per week, open notes.
- Will also include free response check-in questions to get your feedback on how the course is going, what material from the past week you find most confusing, interesting, etc.

Grade Breakdown:

- Problem Sets (4 total): 40%, weighted equally.
- Weekly Quizzes: 10%, weighted equally, lowest score dropped.
- Midterm (March 9th, in class): 20%.
- Final OR Final Project: 30%.
- Extra Credit: Up to 5% for participation, and lots more available on problem sets, for questions asked in class, etc.

Final project

Optionally, instead of taking the final exam, you can complete a final project.

- Identify a topic of current research, formulate a research problem, and make an effort to tackle that problem.
- Submit a 10 page final report.
- Completed in groups of two – if you would like to work alone, please email the instructor to request permission.
- More details + deadlines will be posted shortly on the course webpage.
- Much more work than if you just took the final, but valuable if you are interested in doing research in the area.

Academic Honesty

- A first violation cheating on a homework, quiz, or other assignment will result in a 0 on that assignment.
- A second violation, or cheating on an exam will result in failing the class.
- For fairness, I adhere very strictly to these policies.
- All students in a problem set group are responsible for violations, even those that they were not aware of. So I emphasize – make sure you actually work on the problems together and don't just split up the work.

Disability Services and Accommodations

UMass Amherst is committed to making reasonable, effective, and appropriate accommodations to meet the needs to students with disabilities.

- If you have a documented disability **on file with Disability Services**, you may be eligible for reasonable accommodations in this course.
- If your disability requires an accommodation, please email me by next Friday 2/4 so that we can make arrangements.

I understand that people have different learning needs, home situations, etc. If something isn't working for you in the class, please reach out and let's try to work it out.

Questions?

Background on Randomized Algorithms

Types of Randomized Algorithms

Las-Vegas: Always correct, but the runtime is a random variable (if you get unlucky, the algorithm might be slow).



Monte-Carlo: Runtime is bounded deterministically, but the algorithm may be incorrect with small probability.



Complexity Theory

- **P**: Decidable by a deterministic polynomial time algorithm.
- **ZPP**: Decidable by a randomized algorithm which is always correct and runs in polynomial time **in expectation**.
- **BPP**: Decidable by a polynomial time randomized algorithm that is correct with probability $\geq 2/3$ on both 'yes' and 'no' instances (two-sided error).
- **RP**: Decidable by a polynomial time randomized algorithm that is correct with probability 1 on 'yes' instances, and $\geq 1/2$ on 'no' instances (one-sided error).

Think-Pair-Share 1: Why for BPP is the probability of success stated as $2/3$, rather than say $1/4$ or $9/10$? Why for RP is it stated as $1/2$ for 'no' instances rather than say $1/4$ or $2/3$?

Think-Pair-Share 2: How do these complexity classes compare? Which is the biggest? Which is the smallest?

Think-Pair-Share 1: Why for BPP is the probability of success stated as $2/3$, rather than say $1/4$ or $9/10$? Why for RP is it stated as $1/2$ for 'no' instances rather than say $1/4$ or $2/3$?

Think-Pair-Share 2: How do these complexity classes compare? Which is the biggest? Which is the smallest?

Complexity Theory

Major Open Question: $P \subseteq ZPP \subseteq RP \subseteq BPP$ but does $P = BPP$?

- Most think yes, that sufficiently strong pseudorandom number generators will eventually show that $P = BPP$.

Are there natural problems in BPP that we don't know are in P ?

- For a long time **primality testing** was one such problem.
- Randomized algorithms for primality testing (Miller-Rabin '76, '80 and Solovay-Strassen '77) were very important. They made RSA encryption possible – central in the rise of randomized algorithms
- In 2002, Agrawal-Kayal-Saxena finally gave a polynomial time deterministic test. In practice, randomized tests are still used.
- Currently, **polynomial identity testing** is probably the most important problem known to be in BPP but not P .

Polynomial Identity Testing

Given an n -variable polynomial $p(x_1, x_2, \dots, x_n)$, determine if the polynomial is **identically zero**. I.e., if $p(x_1, x_2, \dots, x_n) = 0$ for all x_1, \dots, x_n . E.g., you are given:

$$p(x_1, x_2, \dots, x_3) = x_3(x_1 - x_2)^3 + (x_1 + 2x_2 - x_3)^2 - x_1(x_2 + x_3)^2.$$

- Can expand out all the terms and check if they cancel. But the number of terms can be as large as $\binom{n+d}{d}$ – i.e., exponential in the number of variables n and the degree d .

Extremely Simple Randomized Algorithm: Just pick random values for x_1, \dots, x_n and evaluate the polynomial at these values. With high probability, if $p(x_1, \dots, x_n) = 0$, the polynomial is identically 0!

$$p(5, 2, \dots, -1) = -1(5 - 2)^3 + (5 + 2 \cdot 2 + 1)^2 - 5(2 - 1)^2 = 68.$$

Probability Review

(with some classic algorithmic applications)

Conditional Probability

Consider two random events A and B .

- **Conditional Probability:**

$$\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}.$$

- **Baye's Theorem:**

$$\Pr(A|B) = \frac{\Pr(B|A) \cdot P(A)}{P(B)}.$$

- **Independence:** A and B are independent if:

$$\Pr(A|B) = \Pr(A).$$

Using the definition of conditional probability, independence means:

$$\frac{\Pr(A \cap B)}{\Pr(B)} = \Pr(A) \implies \Pr(A \cap B) = \Pr(A) \cdot \Pr(B).$$

Independence

Sets of events: For a set of n events, A_1, \dots, A_n , the events are **k -wise independent** if for any subset S of at most k events,

$$\Pr\left(\bigcap_{i \in S} A_i\right) = \prod_{i \in S} \Pr(A_i).$$

For $k = n$ we just say the events 'are independent'.


Random Variables: Two random variables X, Y are independent if for all s, t , $X = s$ and $Y = t$ are independent events. In other words:

$$\Pr(X = s \cap Y = t) = \Pr(X = s) \cdot \Pr(Y = t).$$

Expectation and Variance

Consider a random X variable taking values in some finite set $S \subset \mathbb{R}$. E.g., for a random dice roll, $S = \{1, 2, 3, 4, 5, 6\}$.

- **Expectation:** $\mathbb{E}[X] = \sum_{s \in S} \Pr(X = s) \cdot s$.
- **Variance:** $\text{Var}[X] = \mathbb{E}[(X - \mathbb{E}[X])^2]$.



Exercise: Verify that for any scalar α , $\mathbb{E}[\alpha \cdot X] = \alpha \cdot \mathbb{E}[X]$ and $\text{Var}[\alpha \cdot X] = \alpha^2 \cdot \text{Var}[X]$.

Linearity of Expectation

$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$ for any random variables X and Y . No matter how correlated they may be!

Proof:

$$\begin{aligned}\mathbb{E}[X + Y] &= \sum_{s \in S} \sum_{t \in T} \Pr(X = s \cap Y = t) \cdot (s + t) \\ &= \sum_{s \in S} \sum_{t \in T} \Pr(X = s \cap Y = t) \cdot s + \sum_{t \in T} \sum_{s \in S} \Pr(X = s \cap Y = t) \cdot t \\ &= \sum_{s \in S} \Pr(X = s) \cdot s + \sum_{t \in T} \Pr(Y = t) \cdot t \\ &= \mathbb{E}[X] + \mathbb{E}[Y].\end{aligned}$$

Maybe the single most powerful tool in the analysis of randomized algorithms.

Linearity of Variance

$\text{Var}[X + Y] = \text{Var}[X] + \text{Var}[Y]$ when X and Y are independent.

Claim 1: (exercise) $\text{Var}[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$ (via linearity of expectation)

Claim 2: (exercise) $\mathbb{E}[XY] = \mathbb{E}[X] \cdot \mathbb{E}[Y]$ (i.e., X and Y are uncorrelated) when X, Y are independent.

Together give:

$$\begin{aligned}\text{Var}[X + Y] &= \mathbb{E}[(X + Y)^2] - \mathbb{E}[X + Y]^2 \\ &= \mathbb{E}[X^2] + 2\mathbb{E}[XY] + \mathbb{E}[Y^2] - (\mathbb{E}[X] + \mathbb{E}[Y])^2 \\ &= \mathbb{E}[X^2] + 2\mathbb{E}[XY] + \mathbb{E}[Y^2] - \mathbb{E}[X]^2 - 2\mathbb{E}[X] \cdot \mathbb{E}[Y] - \mathbb{E}[Y]^2 \\ &= \mathbb{E}[X^2] + \mathbb{E}[Y^2] - \mathbb{E}[X]^2 - \mathbb{E}[Y]^2 \\ &= \text{Var}[X] + \text{Var}[Y].\end{aligned}$$

Linearity of Variance

Exercise: Verify that for random variables X_1, \dots, X_n ,

$$\text{Var} \left(\sum_{i=1}^n X_i \right) = \sum_{i=1}^n \text{Var}(X_i),$$

whenever the variables are **2-wise independent** (also called pairwise independent).

Application 1: Matrix Product Testing

Matrix Product Testing

Given matrices $A, B, C \in \mathbb{R}^{n \times n}$, output 'yes' if $AB = C$ and 'no' otherwise. Equivalently, check if $AB - C = 0$.

- How fast can you solve this problem deterministically?
- Any thoughts on how to solve it faster with randomness?

Freivald's algorithm:


1. Let $\mathbf{x} \in \mathbb{R}^n$ be a vector with each entry set independently to 0 or 1 with probability 1/2.
2. Compute $\mathbf{y} = (AB - C)\mathbf{x}$.
3. Output 'yes' if \mathbf{y} is the all zeros vector, and 'no' otherwise.

Runs in $O(n^2)$ time and if $AB = C$, always outputs 'yes'.

Freivald's Analysis

Theorem: If $AB \neq C$, Freivald's algorithm outputs 'no' with probability at least $1/2$.

- Let $D = AB - C$. If $AB \neq C$, there is some i, j s.t $D(i, j) \neq 0$.



- $y(i) = D(i, 1) \cdot x(1) + \dots + D(i, j) \cdot x(j) + \dots + D(i, n) \cdot x(n)$.
- Let $s = \sum_{k \neq j} D(i, k) \cdot x(k)$. So $y(i) = D(i, j) \cdot x(j) + s$.

Upshot: Freivald's algorithm runs in $O(n^2)$ time, as compared to $O(n^3)$ time to deterministically check if $AB = C$. Satisfies:

- If $AB = C$, outputs 'yes'.
- If $AB \neq C$, outputs 'no' with probability at least $1/2$.

Repeating k times boosts the success probability to $1 - 1/2^k$.
I.e. ≈ 0.999 for $k = 10$.