

PROBABILITY AND COMPUTING

ALPEREN A. ERGÜR

Imagine you have a huge vector, so huge that you don't even want to save a copy of it in your hard drive. Nevertheless, it turns out that you have to compute the distance of this huge vector to another huge vector. Yikes! Can you do this quickly? I mean, say the vectors have n bits, can you do the computation using $\log(n)$ bits? Of course, you can not. But wait, then why am I advertising this course to you?

Course Content

- Elementary Probability: Definitions of events and their probabilities. Algorithmic applications: Randomized identity testing for polynomials (PIT) and a randomized min-cut algorithm
- Random variables and Expectation: Bernoulli, Binomial, and Geometric distributions, expectation, conditional expectation. Algorithmic application: Average analysis of Quicksort
- Moments of random variables: First and Second Moment, Chebyshev's Inequality. Algorithmic application: A basic sampling algorithm to compute the mean
- Basic Concentration: Chernoff and Hoeffding Inequalities. Algorithmic application: there are so many applications, the instructor will decide one of them on the fly (with a coin toss of course)
- Basic Structures in Random Discrete Objects: Balls and Bins, Poisson Approximation. Algorithmic application: Bucket sort (an algorithm that breaks speed barriers)
- Basic Structures in Random Discrete Objects: First Moment Methods and Lovasz Local lemma. Algorithmic application: Finding independent sets, and satisfiability problems
- The Gaussian: Normal distribution and Central Limit theorem. Algorithmic application: dimensionality reduction (time permits)

Spoiler Alert The answer to the above question is yes, but you can't insist on being %100 sure of the correctness of the computation all the time.

Prerequisites There will be no formally listed prerequisites, but we assume mental fitness. A course in discrete mathematics (i.e. CS2233 or CS333) or a course in linear algebra (MAT2233) would help, but the knowledge of these courses will not be assumed. The main prerequisite is being ready to get challenged and grow, and be willing to question the lecture material before believing of its correctness.

Textbook Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis by Mitzenmacher and Upfal

Why this Course? This course will make you connected to the world of randomized computation which already operates in your cell phone. Intellectually, you will be empowered with the knowledge of probability theory-randomized computation, will be more prepared to go deeper in discrete mathematics and to tackle applications in data science and machine learning.

How do I know if this course is a good fit for me? If you have a taste for discrete mathematics, probability, or algorithms; this course is for you. If you are interested in gaining some randomized computing skills; this course is for you. If you are only interested in practical skills like SQL and TensorFlow proficiency; this course is not for you.