

Instructor and contacts: Stephen Becker, ECOT 231, (303) 492-0662, stephen.becker@colorado.edu. I will respond to most reasonable email, but (1) lengthy questions should be posed at office hours and not via email, and (2) I will not check email at odd hours or just before homework is due.

Meeting times: Monday, Wednesday and Friday, Fleming 102, 10:00 AM to 10:50 AM.

Office hours: Wednesdays 4–5 PM, Thursdays 1–2 PM and 4–5 PM, in ECOT 231.

Learning Goals: To understand the basic use cases and analysis for modern randomized methods, especially as applied to numerical linear algebra and low-rank approximation.

Prereqs: APPM 3310 “Matrix Methods” or similar; but APPM 4440 “Real Analysis” or equivalent (e.g., APPM 5440) is suggested. Basic familiarity with probability is essential.

Text: There is no textbook. See the website for some helpful monographs

Syllabus: See next pages for detailed syllabus

Recitations: There are no recitation sessions.

Exams: There are no exams (no midterms, no final)

Project: (25% of grade) The last few weeks, there will be no/less homework and instead there is a final project, which is open-ended. Students are encouraged to form groups (up to three people). It can be either theory or computation. Topics should be discussed with the instructor. The final week, groups will give brief (about 10 min) presentations on their topic and a short write-up; a very long write-up can be accepted in lieu of a presentation.

Grades/Homeworks: (75% of grade) Your grade will be entirely based on homeworks and a project. There will be about 10 homeworks due every week. We will mostly alternate between **normal homeworks** (some theory, some computation) and **reading assignment homeworks** (read part of a given article or book and then write a prompted response).

1. students taking the course for graduate credit (5720) may drop **one** homeworks
2. students taking the course for undergraduate credit (4720) may drop **two** homeworks
3. ... however, regardless of whether you take the course for undergrad or grad credit, you must complete **all** homeworks to be eligible for a grade of “A” (if you drop a homework, then your highest possible grade is “A-”).
4. Your final homework score, worth 75% of the grade, will be averaged over the homeworks you do not drop.

Late homework and cheating policy: Homework is due at the beginning of class.

Late homework will not be accepted or graded. Any instance of cheating or other violation of the honor code will result in a 0 grade on the relevant assignment and a referral to the honor council.

Course web page: (course website: <https://bit.ly/2FqjA4H>) It is your responsibility to check the web page on a regular basis. Here you will find detailed information such as homework assignments and solutions, office hours, and special announcements. In addition, it contains policies on illness, academic honesty, and special accommodations for religious holidays and documented special needs. We will use Canvas in addition to the main course website, mainly for grades.

Dropping the course: Advice from your department advisor is recommended before dropping any course. After Jan. 30, dropping a course results in a “W” on your transcript and you’ll be billed for tuition. After March 22, dropping the course is possible only with a petition approved by the Dean’s office.

Syllabus

Randomized methods have been a core part of computer science (CS) for a long time, but are only recently being used in scientific and numerical computing. This course focuses on these recent advances, and is distinguished from more traditional CS “randomized algorithm” courses by its focus on continuous (or at least floating point) math, as opposed to discrete math. The course will also discuss stochastic gradient methods and variance reduced versions like SVRG. Research in the field is ongoing and there is not yet a textbook, and it also means the course is not comprehensive but rather will present a biased selection of topics.

Learning Objectives After taking this course, students should be familiar with core **tools**, such as various types of randomized sketches, and **applications** to various matrix decompositions, least-squares and simple optimization problems, and modern stochastic gradient methods for complicated optimization problems.

Details

Part 1: Preliminaries

1. How a computer works according to a mathematician
2. Blocked computations, BLAS, memory hierarchy, single-pass methods
3. Applied linear algebra facts, matrix multiplication, norms, Eckart-Young etc.
4. Matrix decompositions: QR, SVD, LU, Cholesky
5. Iterative methods: power, CG, MINRES, etc.
6. Basic probability: Chebyshev, Markov and Bernstein/Chernoff inequalities

Part 2: Sketches and randomized dimensionality transforms

1. Gaussian and Haar distribution
2. Johnson-Lindenstrauss Lemma
3. Fast JL versions
4. Count sketch
5. Selection via leverage scores

Part 3: The Good Stuff

1. Least squares/regression, iterative Hessian Sketch, BLENDENPIK and LSRN
2. SVD decompositions
3. One-pass SVD decompositions
4. Nyström, CX, CUR and ID decompositions (CSSP)
5. Sketching for PCA
6. Tensor sketching (cf. Malik and Becker)
7. Randomized trace estimators for genomics (cf. Border and Becker)
8. Sketching covariance matrices (cf. Pourkamali-Anaraki and Becker)
9. Stochastic Gradient methods (SGD)
10. Variance reduced SGD, e.g., SVRG, SAGA

Possible additional topics

1. Monte Carlo and quasi-Monte Carlo
2. Analysis of K-means++
3. Sketching with conditional gradient/Frank-Wolfe: “sketchy decisions”
4. Smoothed analysis
5. CS-style randomized algorithms (sorting, shortest-path, etc.)

Software

You may program in whatever language you like. Programming in C, Fortran or C++, you may use standard libraries like NAG or GLPK, but these libraries are not always state-of-the-art. Many academic researchers are using high-level languages for proof-of-concept implementations; we recommend MATLAB, Python or Julia.

The following policies are standard CU-Boulder policies

Accommodation for Disabilities If you qualify for accommodations because of a disability, **please submit to your professor a letter from Disability Services** in a timely manner (for exam accommodations provide your letter at least one week prior to the exam) so that your needs can be addressed. Disability Services determines accommodations based on documented disabilities in the academic environment. Information on requesting accommodations is located on the [Disability Services website](#). Contact Disability Services at 303-492-8671 or dsinfo@colorado.edu for further assistance. If you have a temporary medical condition or injury, see [Temporary Medical Conditions](#) under the Students tab on the Disability Services website. Any student requiring [exam accommodations](#) should contact their instructor and the Help Room Coordinator, Rachel Cox (rachel.cox@colorado.edu; office: ECCR 241). Make sure to schedule arrangements with Ms. Cox 5 business days in advance of any exam requiring accommodation.

Religious Holidays Campus policy regarding religious observances requires that faculty make every effort to deal reasonably and fairly with all students who, because of religious obligations, have conflicts with scheduled exams, assignments or required attendance. In this class, there should be minimal conflict since there is no attendance grade. If you must miss the final exam due to religious reasons, talk to the professor at the beginning of the semester to make special arrangements. If the homework is due on the date of a religious holiday, you are expected to turn the homework in early. If you have a religious holiday that lasts longer than one week, and so therefore you would not receive the homework with enough time to turn it in early, talk to the professor at the beginning of the semester. See the [campus policy regarding religious observances](#) for full details.

Classroom Behavior Students and faculty each have responsibility for maintaining an appropriate learning environment. Those who fail to adhere to such behavioral standards may be subject to discipline. Professional courtesy and sensitivity are especially important with respect to individuals and topics dealing with differences of race, color, national origin, sex, pregnancy, age, disability, creed, religion, sexual orientation, gender identity, gender expression, veteranstatus, political affiliation or political philosophy. Class rosters are provided to the instructor with the student's legal name. I will gladly honor your request to address you by an alternate name or gender pronoun. Please advise me of this preference early in the semester so that I may make appropriate changes to my records. For more information, see the policies at [classroom behavior](#) and the [Student Code of Conduct](#).

Sexual Misconduct, Discrimination, Harassment and/or Related Retaliation The University of Colorado Boulder (CU Boulder) is committed to maintaining a positive learning, working, and living environment. CU Boulder will not tolerate acts of sexual misconduct (including sexual assault, exploitation, harassment, dating or domestic violence, and stalking), discrimination, and harassment by members of our community. Individuals who believe they have been subject to misconduct or retaliatory actions for reporting a concern should contact the Office of Institutional Equity and Compliance (OIEC) at 303-492-2127 or cureport@colorado.edu. Information about the OIEC, university policies, [anonymous reporting](#), and the campus resources can be found on the [OIEC website](#).

Please know that faculty and instructors have a responsibility to inform OIEC when made aware of incidents of sexual misconduct, discrimination, harassment and/or related retaliation, to ensure that individuals impacted receive information about options for reporting and support resources.

Honor Code All students enrolled in a University of Colorado Boulder course are responsible for knowing and adhering to the Honor Code. Violations of the policy may include: plagiarism, cheating, fabrication, lying, bribery, threat, unauthorized access to academic materials, clicker fraud, submitting the same or similar work in more than one course without permission from all course instructors involved, and aiding academic dishonesty. All incidents of academic misconduct will be reported to the Honor Code (honor@colorado.edu; 303-492-5550). Students who are found responsible for violating the academic integrity policy will be subject to nonacademic sanctions from the Honor Code Council as well as academic sanctions from the faculty member. Additional information regarding the academic integrity policy can be found at the [Honor Code Office website](#). Students are encouraged to work in groups, however all work turned in must be your own, and you are responsible and accountable for all group work associated with your name.