

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, ECCR 151, 1:00 PM to 1:50 PM.

Office hours: Three hours per week, times to be determined during the first week of class.

Learning Goals: The course investigates landmark convex optimization algorithms and their complexity results. We aim to cover both key theory, such as polynomial time algorithms for linear programming and the optimality of Nesterov's method, while also surveying current practical state-of-the-art methods.

Prereqs: APPM 3310 "Matrix Methods" or similar; but APPM 4440 "Real Analysis" or equivalent (e.g., APPM 5440) is suggested. A first-course on optimization like CSCI 5254 "Convex Optimization" is helpful but not necessary.

Text: See next pages for details

Syllabus: See next pages for detailed syllabus

Recitations: There are no recitation session.

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.

Grades/Homeworks: (75% of grade) Your grade will be entirely based on homeworks and a project. There will be about 6 double-homeworks (due every two weeks). Each "double-homework" will consist of two parts (e.g., one part may be theory, one part computation).

1. students taking the course for graduate credit (5720 instead of 4720) may drop **two** single homeworks (e.g., one full double-homework, or two parts from different double-homeworks)
2. students taking the course for undergraduate credit (4720) may drop **four** single 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 policy: Homework is due at the beginning of class.

Late homework will not be accepted or graded.

Course web page: (course website: <http://tinyurl.com/zyem3tx>) 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 may also use D2L in addition to the main course website.

Dropping the course: Advice from your department advisor is recommended before dropping any course. After Jan. 27, dropping a course results in a "W" on your transcript. After March 18, dropping the course is possible only with a petition approved by the Dean's office.

Syllabus

Since there is no established optimization curriculum at CU Boulder, we will attempt to cover a lot of material in just one semester. We can divide topics into four broad categories:

Analysis Basic theory (convex analysis, optimality, duality)

Methods Standard methods, both classical and modern, and standard classifications. The idea is to give students a toolkit of standard approaches to solve actual problems.

Convergence Convergence theory (ellipsoid, Nesterov, etc.). This involves some very beautiful mathematics (and is also useful!)

Applications Often in signal processing and machine learning

Details

Section 1: Background, duality, motivation

1. Intro
2. Convex Sets (ch 2 in [BV2004])
3. Convex Functions (ch 3 in [BV2004])
4. Convex Analysis, subdifferentials ([BauschkeCombettes])
5. Convex Optim. Problems (ch 4 in [BV2004])
6. Duality, optimality (ch 5 in [BV2004])
7. Some applications, ERM and machine learning, sparse recovery, image denoising, low-rank matrix recovery
8. CVX, CVXPY, `convex.jl` tutorials

Section 2: Conic programs & relaxations

1. Linear programming
 - a) standard LP theory, duality ([WrightIPM])
 - b) simplex method, complexity, smoothed analysis ([Vanderbei_LP], [Nemirovski_LP])
 - c) IPM ([WrightIPM], [BV2004])
 - d) ellipsoid method (TBD, [Nemirovski_LP])
2. Integer Linear programming ([Vanderbei_LP])
3. SDP ([BV2004], Pataki?)
4. QP, convex and non-convex (S-procedure, relaxations) ([Nemirovski_LP],[BN_Modern], [BV2004])
5. MAXCUT relaxation ([Nemirovski_LP],[BN_Modern])
6. Polynomial optimization, SOS ([Nemirovski_LP],[BN_Modern])

Section 3: Algorithms

1. Gradient Descent ([V_course])
2. Accelerated and projected/proximal gradient descent, optimality ([V_course])
3. Classical methods ([NW05], Bertsekas, [Nesterov2004])
 - a) Subgradient descent
 - b) Newtons method, self-concordancy
 - c) Non-linear conjugate gradient
 - d) quasi-Newton
 - e) Levenberg-Marquardt, Gauss-Newton

- f) Augmented Lagrangian
 - g) Sequential Quadratic Programming (SQP)
 - h) Step-size (Barzilai-Borwein), line search
4. Recent methods
- a) Douglas-Rachford and ADMM
 - b) Primal-dual methods (e.g., Condat)
 - c) Conditional Gradient and Frank-Wolfe
 - d) (Randomized) Coordinate Descent
 - e) Stochastic gradient descent
 - f) State-of-the-art variants (SDCA, SAGA, etc.)

Possible additional topics

1. Benders decomposition
2. Geometric Programming, linear-fractional
3. Online Convex Optimization
4. Mirror Descent
5. Bregman distances
6. Robust optimization
7. Randomized Kaczmarz method; Algebraic Reconstruction Technique (ART) for tomography
8. POCS (projection onto convex sets), best approximation, Dykstra's algorithm

Texts and references

We will not follow a single textbook. The following resources (mostly available online) will be the primary references:

[BV2004] S. Boyd and L. Vandenberghe, “Convex Optimization” (Cambridge U. Press, 2004). Free electronic version at www.stanford.edu/~boyd/cvxbook/. This is a standard text that gives some convex analysis background and focuses on interior-point methods and engineering applications.

[B_course] S. Boyd’s [course notes for EE364a](#), which follow his book.

[V_course] L. Vandenberghe’s [course notes for ee236b](#) follow the book, while his [course notes for ee236c](#) contain more details on proofs and advanced topics.

[NW05] J. Nocedal and S. Wright, “Numerical Optimization” (Springer, 2005). We have free electronic access at CU via [SpringerLink](#). This is the go-to reference for implementing a standard method.

————— *Some advanced topics will be pulled from these resources:* —————

[BubeckPDF] Sebastien Bubeck, “Convex Optimization: Algorithms and Complexity” (2015 monograph, free at <http://research.microsoft.com/en-us/um/people/sebubeck/Bubeck15.pdf>). This covers a very modern perspective

[BubeckBLOG] Sebastien Bubeck’s blog for his ORF523 course (similar to his monograph): <https://blogs.princeton.edu/imabandit/orf523-the-complexities-of-optimization/>

[Beck14] Amir Beck, “Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with MATLAB” (SIAM, 2014, \$89). This covers classical results but by a modern researcher aware of current research

[BN_Modern] Ben-Tal and Nemirovski, [Lectures on Modern Convex Optimization](#) (the PDF is 2013; they have a [SIAM book from 2001](#) as well)

[Nemirovski_LP] A. Nemirovski, [Introduction to Linear Optimization](#). More lectures notes and monographs of A. Nemirovski available at his [website](#).

[Myths] Harvey Greenberg (UC Denver, emeritus), [Myths and counterexamples in mathematical programming](#). Entertaining and informative

[Duals] A. Benjamin, [Sensible rules for remembering duals](#). SIAM Review (37)1 1995.

[Vanderbei_LP] R. Vanderbei, “Linear Programming.” (2008) Standard LP theory, including brief vignettes into integer programming. Free download at [Vanderbei’s website](#) and also via SpringerLink

[Wright_IPM] Stephen Wright, “Primal-Dual Interior-Point Methods” (SIAM 1997), available electronically from [SIAM](#). Good summary of LP theory, and more practical/thorough IPM discussion than in [BV2004] (e.g., includes Mehotra’s predictor-corrector method), though restricted to LP of course.

[Rockafellar1970] R. Tyrrell Rockafellar, “Convex Analysis”, Princeton 1970 (reprinted 1997). You can download via [John Dattorro’s website](#). Terry Rockafellar wrote the first, and definitive, book on the subject.

[BauschkeCombettes] H. Bauschke and P. Combettes, “Convex Analysis and Monotone Operator Theory in Hilbert Spaces”, Springer 2011, available electronically via [SpringerLink](#). More advanced, recommended as a reference to lookup results.

[Nesterov2004] Y. Nesterov, “Introductory Lectures on Convex Optimization”, Springer 2004. Concise book, presenting the material in a unique way, from one of the best researchers. Free access via [SpringerLink](#)

[BertsekasNedicOzdoglar] D. Bertsekas, A. Nedic, A.E. Ozdoglar, “Convex Analysis and Optimization” (Athena Scientific).

[Bertsekas] D. Bertsekas, “Convex Optimization Theory” (Athena Scientific). [Bertsekas has a lot of other books with useful information as well]

[BorweinLewis] J. M. Borwein and A. S. Lewis, “Convex Analysis and Nonlinear Optimization” (Springer). A bit more mathematical

[Hiriart-UrrutyLemarechal] J.B. Hiriart-Urruty and C. Lemarechal, “Convex Analysis and Minimization Algorithms” (Springer). More on analysis, less on optimization. Technical and a useful reference.

[LuenbergerYe] D. Luenberger and Y. Ye, “Linear and Nonlinear Programming” (Springer). A bit more mathematical than [BV2004], slightly different topics.

Research articles

Especially for the state-of-the-art topics, we will use journal articles as references. Here are a few we may use:

Splitting and Proximal methods

- “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers” by S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein (2010). Readable and popular survey on ADMM.
- “Proximal Algorithms” by N. Parikh and S. Boyd (2013). Popular survey article.
- “A Primer on Monotone Operator Methods” by E. Ryu and S. Boyd (2016). Easier introduction to advanced topic.
- P. L. Combettes and J.-C. Pesquet, “Proximal splitting methods in signal processing,” 2011. Good survey
- Laurent Condat “A primal-dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms”, 2011 (J. Optim. Theory and Appl. 2013)

Coordinate descent and stochastic methods

- “Coordinate Descent Algorithms” by S. Wright (2015). Survey article, easier to follow, good starting point.
- SDCA: “Stochastic Dual Coordinate Ascent Methods for Regularized Loss Minimization” by S. Shalev-Shwartz and T. Zhang (2012). Well-known algorithm, improved to mini-batch (Takac, Richtarik, Srebro) and proximal versions (Shalev-Shwartz and Zhang 2014) in later years
- Lin Xiao, “Dual Averaging Methods for Regularized Stochastic Learning and Online Optimization”, 2010 JMLR. Not coordinate descent, more like stochastic gradient descent (SGD) with averaging. Analyzes in an *online learning* framework.
- Defazio, Bach and Lacoste-Julien, “SAGA: A Fast Incremental Gradient Method With Support for Non-Strongly Convex Composite Objectives”, NIPS 2014. Recent variance-reduction technique (variant of SGD) for solving very large problems.

Software

You may program in whatever language you like, but we recommend one of the following three.

MATLAB We suggest [CVX](#)

Python We suggest [cvxpy](#)

Julia We suggest [Convex.jl](#)

Programming in C, Fortran or C++, you may use standard libraries like NAG or GLPK, but these are not always state-of-the-art. Most researchers are using high-level languages.

The following policies are standard CU-Boulder policies

Disability: 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. Contact Disability Services at 303-492-8671 or by e-mail at dsinfo@colorado.edu. If you have a temporary medical condition or injury, see Temporary Medical Conditions: Injuries, Surgeries, and Illnesses guidelines under Quick Links at [Disability Services website](#) and discuss your needs with your professor.

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 little conflict since the grade is based on homeworks and a single final exam. 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. Full details of the campus-wide policy are at www.colorado.edu/policies/fac_relig.html.

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, culture, religion, creed, politics, veteran's status, sexual orientation, gender, gender identity and gender expression, age, disability, and nationalities. Class rosters are provided to the instructor with the student's legal name. The instructor 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 we may make appropriate changes to my records. See policies at www.colorado.edu/policies/classbehavior.html and at www.colorado.edu/studentaffairs/judicialaffairs/code.html#student_code.

Harassment and discrimination: 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, discrimination, harassment or related retaliation against or by any employee or student. CU's Sexual Misconduct Policy prohibits sexual assault, sexual exploitation, sexual harassment, intimate partner abuse (dating or domestic violence), stalking or related retaliation. CU-Boulder's Discrimination and Harassment Policy prohibits discrimination, harassment or related retaliation based on race, color, national origin, sex, pregnancy, age, disability, creed, religion, sexual orientation, gender identity, gender expression, veteran status, political affiliation or political philosophy. Individuals who believe they have been subject to misconduct under either policy should contact the Office of Institutional Equity and Compliance (OIEC) at 303-492-2127. Information about the OIEC, the above referenced policies, and the campus resources available to assist individuals regarding sexual misconduct, discrimination, harassment or related retaliation can be found at the [OIEC website](#).

Honor code: All students enrolled in a University of Colorado Boulder course are responsible for knowing and adhering to the [academic integrity policy](#) of the institution. Violations of the policy may include: plagiarism, cheating, fabrication, lying, bribery, threat, unauthorized access, clicker fraud, resubmission, and aiding academic dishonesty. All incidents of academic misconduct will be reported to the Honor Code Council (honor@colorado.edu; 303-735-2273). Students who are found responsible of 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 honorcode.colorado.edu.

Note: sanctions for first-time offending **graduate students** are typically more severe than for first-time offending **undergraduate students**, since graduate students are held to higher standards.