6.S098: Introduction to Applied Convex Optimization
IAP 2023

Instructor Information

Theo Diamandis
Email: tdiamand@mit.edu

Class Information

Dates: T/Th 1/10/2023 – 2/2/2023
Time: (1.5hrs) TBD – TBD
Classroom: TBD
Office Hours: TBD

Course Description

Convex optimization problems appear in a huge number of applications and can be solved very efficiently, even for problems with millions of variables. However, recognizing what can be transformed into a convex optimization problem can be challenging. This course will teach you how to recognize, formulate, and solve these problems. We will briefly survey theoretical results in convex analysis, but the majority of the course will focus on formulating and solving problems that come up in practice. Applications will include signal processing, statistics & machine learning, finance, circuit design, mechanical structure design, control, power systems, and other areas based on student interest. This course is designed for advanced undergraduates and beginning graduate students.
Prerequisites: multivariable calculus (18.02), linear algebra (18.06 or 18.061), basic probability, programming, mathematical maturity (e.g., 6.042).

Course Objectives

After this course, you should be able to...

- recognize and formulate convex optimization problems that appear in various fields.
- use open source software to solve these optimization problems.
- decide which solver is best for your problem.

Resources

Textbook: *Convex Optimization*, by Stephen Boyd and Lieven Vandenberghe.

Software: `Convex.jl` and solvers including Hypatia and COSMO.

Course Schedule

The schedule below is tentative and subject to change.

Week 1:
Jan 10: **Introduction to Optimization**
- Overview of optimization: least squares $\rightarrow$ linear programs $\rightarrow$ convex programs
- Examples: portfolio optimization, mechanical design, machine learning, max flow problem
- Convex sets & their properties

Jan 12: **Convex Analysis & Optimization**
- Convex functions & their properties
- Linear, quadratic, and second-order cone programming
- Disciplined Convex Programming (DCP) & convex optimization software

Week 2:
Jan 17: **Convex Optimization Problems and Applications**
- Case study: model predictive control (how SpaceX lands rockets)
- Geometric programming
- Case study: circuit design
- Multiobjective optimization

Jan 19: **Duality (*i.e.*, understanding the output of your solver)**
- The Lagrange dual function
- Case study: two-way partitioning problem
- The dual problem

Week 3:
Jan 24: **Duality II**
- Optimality conditions
- Case study: assigning power to communication channels ("water-filling")
- Sensitivity analysis
- Case study: FX arbitrage and no-arbitrage bounds on prices

Jan 26: **Applications: Approximation, Signal Processing**
- Regression problems (regularized, robust, quantile)
- Case study: minimax polynomial fitting (how does my computer compute `sin`, `exp`, etc.)
- Case study: image & audio denoising

Week 4:
Jan 31: **Applications: Statistics & Machine Learning**
- Maximum likelihood & maximum a posteriori probability estimation (MLE & MAP)
- Case study: Logistic regression
- Case study: Support Vector Machine
- Nonparametric distribution estimation

Feb 2: **Advanced Topics**
- TBD based on class interest
- Possible topics: more application areas, mixed integer programming, semidefinite programming

Grading

This course is offered for 6 units of credit, and the grading is P/D/F. To receive credit, you must get 14 or more points. The main goal of this course is to learn about optimization and solve some fun problems.

Problem sets	12
Attendance	6
Project	4
Total	22

Problem Sets

There will be three problem sets, assigned on Wednesday and due the subsequent Wednesday. Grading will be very coarse–emphasis is put on understanding the concepts over correct answers. Each problem set will ask you to model and solve convex optimization problems that come from various application areas, which will be chosen based on interests of the class. Problem sets will require use of the Julia programming language, but no prior experience in Julia is required. Problem sets are expected to take around 8 hours each.

Project (Optional)

There will be an optional course project during the last week of IAP that will allow you to apply convex optimization to a problem of interest. The final product will be a short mathematical description of this problem (akin to the descriptions you see on the problem sets) and a solution. Course grading is designed such that this project can take the place of the final problem set.

Attendance

Attendance is encouraged and will reflected in your final grade. You will received 1pt for each lecture attended, up to a maximum of 6pts.

Acknowledgements

Much of the material for this course comes from the following courses:
- Stephen Boyd's Convex Optimization I (link)
- Lieven Vandenberghe's Convex Optimization (link)
- Ryan Tibshirani's Convex Optimization (link)