AMCS 211
Numerical Optimization
Course Syllabus, Spring 2023

Course description. The course presents practical methods for formulating and solving linear and nonlinear numerical optimization problems that arise in science, engineering, and business applications. The course covers unconstrained and constrained problems. Emphasis is on algorithmic description. The course includes an implementation component where students develop code and use state of the art numerical libraries for modeling and solving problems in machine learning, imaging, shape design, parameter identification, classification, financial planning, PDE-constrained optimization, and others.

Program guide description. Solution of nonlinear equations. Optimality conditions for smooth optimization problems. Theory and algorithms to solve unconstrained optimization; linear programming; quadratic programming; global optimization; general linearly and non-linearly constrained optimization problems.

Logistics. Class meets on Sun and Wed, 1:00–2:25pm in Room 9-3125. Additional lab sessions will also be scheduled. The class is taught by Prof. George Turkiyyah, office 1-0148, phone 808-0414, email `george.turkiyyah@kaust.edu.sa`. The course TAs will be announced on Blackboard shortly.

Learning Outcomes. At the end of the course students are expected to develop:
- a working knowledge of current state of the art algorithms for smooth optimization
- the modeling skills needed to formulate domain problems as numerical optimization problems
- a sufficient understanding of the methods, applications, and limitations, to be able to read the literature

Schedule.
- Preliminaries: Python, numerical linear algebra, multidimensional Taylor's series, etc. (1 week)
- Unconstrained optimization (3 weeks)
  - gradient methods, Newton's method, line search, convergence analysis
  - non-linear least squares problems, Gauss Newton methods, regularization
  - quasi-Newton methods
- Systems of nonlinear equations: Newton methods, globalization, continuation methods (1 week)
- Linear and convex optimization (3 weeks)
  - Lagrange multipliers, KKT optimality conditions
  - duality and dual problems
  - interior point methods
- Methods for local constrained optimization: trust region methods, augmented Lagrangian (2 weeks)
- Conic / semidefinite optimization (1 weeks)
- Global optimization: branch and bound framework (1 week)
- Large scale problems: modern first-order methods, inexact Newton methods, (2 weeks)

Grading. The course grade will be based on:
- Homework (25%). Homework will be assigned weekly—every Sunday due the following Sunday. Keeping up with the homework is the best way to make progress on the material.
- Course project (25%). The course project will take place during the last few weeks of the semester. It is essentially a mega-homework where you get to solve a substantial problem.
- Two exams: a midterm (25%) which will take place on March 19, and a final (25%).

Resources. Resources for the course include:
- J. Nocedal and S. J. Wright. Numerical Optimization, second edition. Springer Verlag, 2006. An e-book version is available through the library at `https://link.springer.com/book/10.1007/978-0-387-40065-5`
- S. Boyd and L. Vandenberghe. Convex Optimization. Cambridge University Press, 2004. An e-version of the book is available at `www.stanford.edu/~boyd/cvxbook/bv_cvxbook.pdf`.