

COMPSCI 590OP: Applied Numerical Optimization

I. COURSE CATALOG

This course provides an overview of numerical optimization methods and covers how to implement them in Python. We will discuss common algorithms ranging from gradient descent to stochastic methods, with applications ranging from image processing to neural networks.

II. COURSE DESCRIPTION

This course provides an overview of the important topic of numerical optimization. In this introductory-level course, we will cover the basic concepts of optimization, the key algorithms, and their applications in image/signal processing, machine learning, and statistical estimation. Topics covered include, but are not limited to: i) the basic concepts in optimization, e.g., linear algebra overview, convex sets, norms, optimality conditions, duality, ii) common optimization algorithms, e.g., gradient descent methods, Newton and quasi-Newton methods, conjugate gradient, proximal methods, linear programming, interior point methods, stochastic methods, distributed methods, and iii) applications, e.g., fitting generalized linear models, neural networks, sparsity, recommender systems, image processing, network utility maximization. Required background: basic knowledge of Python programming, basic knowledge of probability and statistics, linear algebra, multivariate calculus.

III. COURSE DETAILS

Instructor: Andrew S. Lan

TA: TBD

Class hours: TTH 4:00–5:15pm, CS 142

Credits: 3 **Prerequisites:** CompSci 240 + Math 235 + Math 233, with grades C or above

Textbooks (Optional): Jorge Nocedal and Stephen Wright, “Numerical Optimization”, Springer-Verlag, 2006; Stephen Boyd and Lieven Vandenberghe, “Convex Optimization”, Cambridge University Press, 2004

IV. GRADING

- 10 total homeworks and programming assignments, 50%
- In-class midterm exam, 20%
- Final exam, 30%
- Tentative grading scale: A: 90-100%, A-: 85-90%, B+: 80-85%, B: 75-80%, B-: 70-75%, C+: 65-70%, C: 60-65%, C-: 55-60%, D+: 50-55%, D: 40-50%, F: 40% and below

V. TENTATIVE SCHEDULE

Homeworks (HWs) and programming assignments (PAs) are assigned on Fridays and due 1 week later.

- Week 1: Linear algebra review, norms, convex sets, HW1
- Week 2: Least squares, gradient descent, application: Poisson regression, PA1
- Week 3: Newton’s method, application: logistic regression, PA2
- Week 4: Quasi-Newton methods, conjugate gradient
- Week 5: Constrained optimization, optimality conditions, duality, HW2
- Week 6: Proximal methods, application: sparsity, PA3
- Week 7: Proximal methods, application: image processing, PA4
- Week 8: Review and midterm exam
- Week 9: Linear programming, HW3
- Week 10: Distributed methods, application: recommender systems, PA5
- Week 11: Interior point methods, application: network utility maximization, PA6
- Week 12: Stochastic methods, application: neural networks, PA7
- Week 13: Connections to sampling methods, review

VI. ACCOMMODATION STATEMENT

The University of Massachusetts Amherst is committed to providing an equal educational opportunity for all students. If you have a documented physical, psychological, or learning disability on file with Disability Services (DS), you may be eligible for reasonable academic accommodations to help you succeed in this course. If you have a documented disability that requires an accommodation, please notify me within the first two weeks of the semester so that we may make appropriate arrangements.

VII. ACADEMIC HONESTY STATEMENT

Since the integrity of the academic enterprise of any institution of higher education requires honesty in scholarship and research, academic honesty is required of all students at the University of Massachusetts Amherst. Academic dishonesty is prohibited in all programs of the University. Academic dishonesty includes but is not limited to: cheating, fabrication, plagiarism, and facilitating dishonesty. Appropriate sanctions may be imposed on any student who has committed an act of academic dishonesty. Instructors should take reasonable steps to address academic misconduct. Any person who has reason to believe that a student has committed academic dishonesty should bring such information to the attention of the appropriate course instructor as soon as possible. Instances of academic dishonesty not related to a specific course should be brought to the attention of the appropriate department Head or Chair. Since students are expected to be familiar with this policy and the commonly accepted standards of academic integrity, ignorance of such standards is not normally sufficient evidence of lack of intent. See <https://www.umass.edu/honesty/>.