Syllabus

208424: Optimization for Statistical Learning

Semester 2/2025

1. Course Information

- **Course Code:** 208424
- **Course Title:** Optimization for Statistical Learning
- **Credits:** 3 (3-0-6)
- **Prerequisites:** I will assume that you are comfortable with linear algebra (specifically matrix-vector and matrix-matrix operations), probability and programming (preferably Python). I also recommend knowing some calculus as a tool for optimization, but it is not required.

2. Instructor Information

- **Instructor:** Donlapark Ponnoprat (https://donlapark.pages.dev)
- **Lecture Session:** Tuesday & Friday, 13:00 - 14:30 AM, at Room STB205.
- **Course Website:** https://donlapark.pages.dev/208424/ for course schedule and slides.
- **Canvas:** https://mango-cmu.instructure.com/courses/21987 for homework assignments.

3. Course Description

We will cover various classical and modern applications of optimization. Topics include: Optimization of Functions, Linear Programs, Quadratic Programs, Duality, Gradient Descent, Momentum and Accelerated Gradient Descent, Projected, Stochastic and Mirror Descent, Adaptive Learning Rate, Second-Order Optimization and Newton's Method, and Various Modern Applications of Optimization.

4. Course Schedule

The following is my tentative schedule for the course. I may make adjustments based on our class progress and your interests.

Week	Topics	Lecture Hours
1	Calculus Review & Matrix Calculus for Optimization	3
2	Introduction to Optimization & Linear Programming	3
3	Duality and Its Interpretations & Introduction to Convexity	3
4	Convex Optimization Problems in Statistics	3
5	The Lagrangian and KKT Conditions	3
6	Mixed-Integer Programming	3
7	Break week	
8	Midterm Exam	
9	Gradient Descent	3
10	Momentum, Accelerated, and Adaptive Gradient Methods	3
11	Projected, Stochastic, and Mirror Descent	3
12	Second-Order Optimization & Newton's Method	3
13	Application: Matrix Factorization and Recommender Systems	3
14	Variational Inference & Optimal Transport	3
15	Sampling as Optimization & Generative Modeling	3
16	Final Exam	

5. Assessment and Grading

5.1 Grading Scheme

I will determine your final grade based on the following components:

- **Homework (3-6 assignments):** 30%
- **Midterm Examination:** 35% Please fill out the form (link in Canvas) so that we can decide on the best time slot.
- **Final Examination:** 35%
  - **Date:** Friday, March 20, 2026
  - **Time:** 15:30 - 18:30 (3 hours)

5.2 Grading Policies

- **Mandatory Exams:** You must take both the midterm and final examinations. If you miss either exam, you will receive an automatic 'F' grade for the course.
- **Grading:** Your final grade may be curved at the end of the semester.

6. Course Policies

- **Academic Honesty:** I expect all work you submit to be your own. Plagiarism, cheating, or any form of academic dishonesty will not be tolerated and will result in severe penalties,

including a failing grade for the assignment or the course. Please refer to the university's academic integrity policies.

- **Attendance:** While I may not directly grade attendance, I highly encourage active participation in lectures as it is crucial for your success in this course. It is your responsibility to catch up on any missed classes.
- **Communication:** I will make all important announcements via email. You are responsible for regularly checking your inbox. For individual questions, please send a direct message via Canvas or direct email.
- **Disability Services:** If you have a disability and may need accommodations in this class, I encourage you to contact me as early as possible so that we can arrange reasonable accommodations.

7. Required and Recommended Resources

My lecture notes and materials will be provided on the course website.

Disclaimer: Please note that this syllabus is my guide for the course and may be subject to change at my discretion. I will announce any changes in class and on the course website.