Syllabus

208424: Optimization for Statistical Learning

Semester 2/2026

1. Course Information

- **Course Code:** STAT 424 (208424)
- **Course Title:** Optimization for Statistical Learning
- **Credits:** 3 (3-0-6)
- **Prerequisites:** MATH 324 (206324); and STAT 329 (208329) or DSCI 351 (229351); and CS 217 (204217) or DSCI 223 (229223). I will assume that you are familiar with basic calculus, linear algebra, statistics and programming in Python.

2. Instructor Information

- **Instructor:** Donlapak Ponnoprat (https://donlapark.pages.dev)
- **Lecture Sessions:** Tuesday & Friday, 13:00–14:30, Room STB205
- **Course Website:** https://donlapark.pages.dev/208424/ for course schedule, slides and notes
- **Canvas:** https://mango-cmu.instructure.com/courses/31782 for homework

3. Course Description

This course covers optimization for statistical learning and for decisions made from data. Topics include the calculus and linear algebra that optimization needs; estimation as optimization, gradient descent, Newton's method, proximal gradient methods and stochastic gradient methods; linear, second-order cone, semidefinite and integer optimization, with duality and shadow prices; decisions under uncertainty and sample average approximation; and decision-focused learning, which fits a model for the decisions it leads to. We will primarily use Python libraries such as `numpy`, `scipy` and `cvxpy`.

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	1. Introduction: Syllabus and Course Overview	1
1-2	2. Mathematics Review 2.1 Calculus 2.2 Linear Algebra	5
3-5	3. Optimization Tools for Statistical Learning 3.1 Estimation as Optimization (Maximum Likelihood) 3.2 Gradient Descent and Ridge Regression 3.3 Newton's Method 3.4 Proximal Gradient Method and the Lasso 3.5 Stochastic Gradient Descent and Adam	9
6-8	4. Optimization with Constraints 4.1 Linear Optimization and Duality 4.2 Second-Order Cone and Semidefinite Optimization 4.3 Integer Optimization	9
9	Midterm Exam	
10-11	5. Decisions under Uncertainty 5.1 Expected Value, Quantiles and Risk 5.2 Sample Average Approximation 5.3 Decisions that Use Features	6
12-15	6. Decision-Focused Learning 6.1 Predict, then Optimize 6.2 Learning the Policy Directly 6.3 Policies that Satisfy Constraints 6.4 Plug-in Policies and the Solution Map 6.5 Weighted Logistic Regression	12
16	Final Exam	

5. Assessment and Grading

5.1 Grading Scheme

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

- **Homework:** 30%
- **Midterm Examination:** 35%
  - **Date:** TBA
  - **Time:** TBA
- **Final Examination:** 35%
  - **Date:** SUN MAR 21
  - **Time:** 12:00–15:00

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.
- **Homework:**

- Submit each homework on Canvas.
- Each homework is due 23:59 on the next Wednesday.
- Your two lowest homework scores will be dropped from the total score calculation.

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, including the checkpoint questions, as it is crucial for your success in this course. It is your responsibility to catch up on any missed classes.
- **Late Submissions: Late submissions will not be accepted under any circumstances.**
- **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

7.1 Main Course Materials

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

7.2 Recommended Books and Online Resources

- **Mathematics for Machine Learning** By Marc Peter Deisenroth, A. Aldo Faisal and Cheng Soon Ong.
  - Book PDF available at: https://mml-book.github.io/
- **Convex Optimization** By Stephen Boyd and Lieven Vandenberghe.
  - Book PDF available at: https://web.stanford.edu/~boyd/cvxbook/
- **Numerical Optimization** By Jorge Nocedal and Stephen J. Wright.
- **Introduction to Linear Optimization** By Dimitris Bertsimas and John N. Tsitsiklis.
- **Lectures on Stochastic Programming: Modeling and Theory** By Alexander Shapiro, Darinka Dentcheva and Andrzej Ruszczyński.
- **From Predictive to Prescriptive Analytics** By Dimitris Bertsimas and Nathan Kallus, *Management Science*, 2020.
- **Smart "Predict, then Optimize"** By Adam N. Elmachtoub and Paul Grigas, *Management Science*, 2022.

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.