USC Viterbi
School of Engineering

CSCI 599: Optimization for Machine Learning

Units: 4.0
Spring 2025 — MonWed — 4:00-5:50PM

Location: DMC 101
https://spkreddy.org/optml2025.html

Instructor: Sai Praneeth Karimireddy
Office: TBD
Office Hours: TBD
Contact Info: karimire@usc.edu; please be sure to include CSCI 599 in the title

Catalogue Description
Optimization methods in machine learning: convexity; gradient descent; non-convex optimization; SGD; adaptive optimization; min-max optimization

Course Description
Fast optimization algorithms which scale to massive datasets have been powering the rapid progress in machine learning. This course establishes the theoretical foundations of such methods, with a particular focus on the challenges arising in modern large-scale ML. We will cover the following topics:

- what is convexity and why is it useful (convex optimization)
- why can we train using only mini-batches (stochastic optimization)
- why Adam is typically preferred over SGD (preconditioning and adaptivity)
- how to train robust ML models (min-max optimization)
- explaining feature learning in neural networks (non-convex optimization)
- how to train privately on distributed datasets (federated optimization)

Note: this course is theoretical.

Learning Objectives
This course will prepare you to view machine learning through the formalism of optimization. You will be able to analyze optimization algorithms and derive their convergence rates. You will learn the plethora of algorithms used in modern ML, and understand the tradeoffs each of them is designed to achieve. You will also learn to implement them in practice.

Recommended Preparation
Probability (at the level of MATH 505a), Linear Algebra and Multi-Variable Calculus (at the level of MATH 225), Analysis of Algorithms (at the level of CSCI 570), and Machine Learning (at the level of CSCI 567).

Course Notes
Grading type: Letter or Credit No/Credit

Technological Proficiency and Hardware/Software Required
You will need a laptop which can run PyTorch for this course – bring your laptop to exercise sessions. Link information for the USC Computing Center Laptop Loaner Program.

Required Readings and Supplementary Materials
We will use material from multiple sources including parts of the the following text books:

- Course lecture notes (link)
- Bubeck, Sébastien. "Convex optimization: Algorithms and complexity." *Foundations and Trends® in Machine Learning* 8.3-4 (2015): 231-357 (link)
- Bach, Francis. “Learning Theory from First Principles”. *MIT Press* 2024 (link)

Optional Readings and Supplementary Materials
There are several related courses which could be useful supplementary references:

- Optimization for Machine Learning at EPFL (link)
- Lecture Notes on Convex Optimization, by Moritz Hardt at UC Berkeley (link)
- Large Scale Optimization for Machine Learning (ISE 633), by Meisam Razaviyayn at USC (link)
- OPTML yearly workshop: https://opt-ml.org/
- Francis Bach’s blog: (link)

Description of Assignments and How They Will Be Assessed

1. **Final Exam (50%)** Closed book exam consisting of theoretical questions similar to exercises. You are allowed to bring one cheat sheet (A4 size paper, both sides can be used).
2. **Team Course Project (40%):** The course includes a major project where you will **present** and **reproduce** a paper on optimization for machine learning in **teams of 2.**
    a. **Presentation (15%):** you will closely read a paper related to the course and make a 20min presentation. See this note on advice by Yiling Chen and evaluation criterion.
    b. **Report (25%):** you will then **reproduce** the main experiment/theorem of the paper. If it is a theoretical paper, you will rederive and present the proof of the main claims as you understand them. If it is an experimental paper, you will follow the ML reproducibility challenge where you will attempt to reproduce the experiment in the paper.
3. **Discussion and participation in exercise sessions (10%).** There will be a weekly exercise session consisting of a mix of theoretical and practical Python exercises for the corresponding topic each**. Coming prepared** and **actively participating** in the exercise session discussions counts for 10% of the grade.

Participation

Discussions and participation will count for 10%. There will be a weekly exercise session consisting of a mix of theoretical and practical Python exercises for the corresponding topic each**. Coming prepared** and **actively participating** in the exercise session discussions counts for 10% of the grade.

Grading Breakdown

Assessment Tool (assignments)	% of Grade
Final Exam	50
Project presentation	15
Project report	25
Discussions and participation	10
TOTAL	100

Assignment Submission Policy

Via Brightspace.

Attendance

There is no explicit need for attendance, but participation and discussion in exercise sessions accounts for 10% of the grade.

Academic Integrity for this Class

Unless otherwise noted, this course will follow the expectations for academic integrity as stated in the USC Student Handbook. The general USC guidelines on Academic Integrity and Course Content Distribution are provided later in this syllabus.

Please ask the instructor [and/or TA(s)] if you are unsure about what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

Class Recordings and Course Content Distribution: You may not record this class without the express permission of the instructor and all other students in the class. Distribution of any notes, recordings, exams, or other materials from a university class or lectures — other than for individual or class group study — is

prohibited without the express permission of the instructor; violations will be considered an intentional act to facilitate or enable academic dishonesty and reported to the university.

Use of Generative AI in this Course

Generative AI is not permitted: Since creating, analytical, and critical thinking skills are part of the learning outcomes of this course, all assignments should be prepared by the student working individually or in groups as described on each assignment. Students may not have another person or entity complete any portion of the assignment. Developing strong competencies in these areas will prepare you for a competitive workplace. Therefore, using AI-generated tools is prohibited in this course, will be identified as plagiarism, and will be reported to the Office of Academic Integrity.

Course Evaluation

Course evaluation occurs at the end of the semester university-wide. In addition, we will use continuous anonymous feedback to help shape the course to your best interests. Remember, this is an advanced level course and will be what you make of it.

Course Schedule

	Topics/Daily Activities	Recommended Reading	Deliverables/due dates

Week 1	**Introduction and Fundamentals of Convexity**
	• **Lecture**:
	  ○ Course overview and objectives
	  ○ Convex sets and convex functions
	  ○ Fenchel conjugate and optimality conditions
	• **Exercise Session**:
	  ○ Getting started with numpy programming
		• Chapter 1 of lecture notes.	Exercise 0 released

Week 2	**Gradient Descent I**
	• **Lecture**:
	  ○ Smoothness and Strong Convexity
	  ○ Gradient descent and its convergence properties
	• **Exercise Session**:
	  ○ Theoretical problems on properties of convex sets and functions (from chapter 1 of notes)
		• Chapter 2 of lecture notes.	Exercise 1 released

Week 3	**Gradient Descent II**
	• **Lecture**:
	  ○ Subgradient descent
	  ○ Proximal Gradient Descent
	• **Exercise Session**:
	  ○ Implementing gradient descent
	  ○ Theoretical problems from chapter 2
		• Chapter 3.4 and Chapter 4 of lecture notes.	Exercise 2 released

Week 4	**Stochastic Optimization and Non-convex optimization**
	• **Lecture**:
	  ○ Stochastic Gradient Descent (SGD)
	  ○ Non-convexity and local optimality
	  ○ Convergence to critical points
	• **Exercise Session**:
	  ○ Optimizing Lasso
	  ○ Theoretical problems from chapters 3 & 4.
		• Chapter 5 and 6.1 of lecture notes.	Exercise 3 released

Week 5	**Understanding Momentum and Acceleration**
	• **Lecture**:
	  ○ Nesterov Acceleration
	  ○ Momentum as smoothing non-convexity
	• **Exercise Session:**
	  ○ Comparing mini-batch and full batch methods
	  ○ Problems from Chapters 5
		• Goh, Gabriel. "Why momentum really works." *Distill* 2.4 (2017): e6.
		• Chapter 2.6 of lecture notes.
			Exercise 4 released

Week 6	**Second order and Adaptive Optimization Methods**
	• **Lecture**:
	  ○ Newton's method
	  ○ Adaptive preconditioning: AdaGrad, Adam
	  ○ (L0,L1)-smoothness and clipped SGD
	• **Exercise Session**:
	  ○ Understanding momentum methods in practice
		• Chapter 7 of lecture notes.
		• Blogpost by Sebastian Ruder on optimizers for deep learning
		• Zhang et al. "Why gradient clipping accelerates training: A theoretical justification for adaptivity." *ICLR 2020.*
			Exercise 5 released

Week 7	**Min-Max Optimization**
	• **Lecture**:
	  ○ Variational inequalities
	  ○ Gradient-descent-ascent
	  ○ Extragradient method
	  ○ Optimistic gradient method
	• **Exercise Session**:
	  ○ Comparative analysis of SGD and adaptive methods
	  ○ Problems from Chapter 7.
		• Sections 2.4 and 3 of Wadia et al. 2023. "A Gentle Introduction to Gradient-Based Optimization and Variational Inequalities for Machine Learning."
		• Lecture 15 by Gauthier Gidel
			Project paper selection due date.
			Exercise 6 released

Week 8	**Training Neural Networks I**
	• **Lecture**:
	  ○ Linear neural networks
	• **Exercise Session**:
	  ○ Adversarial robust training
	  ○ Discuss projects
		• Chapter 6 of lecture notes.	Exercise 7 released.

Material below will not be part of exam unless explicitly stated otherwise.			

Week 9	**Training Neural Networks II**
	• **Lecture**:
	  ○ Gradient flows
	  ○ Neural Tangent Kernel
	  ○ Feature learning in neural networks
	• **Exercise Session**:
	  ○ Visualizing feature learning in neural networks
		• Bach and Chizat. "Gradient descent on infinitely wide neural networks: Global convergence and generalization." *2021.*	

Week 10	**Transfer Learning and Pretraining (Time permitting)**
	• **Lecture**:
	  ○ Why does pre training help?
	  ○ When it harms: Feature distortion in fine tuning
	• **Exercise Session**:
	  ○ Discuss projects
		• Kumar, et al. "Fine-tuning can distort pretrained features and underperform out-of-distribution." *ICLR 2022.*	

Week 11	**Optimization in LLMs: Zeroth-order methods, RLHF**
	• **Lecture**:
	  ○ Overview of LLM training
	  ○ Zeroth-order (memory-efficient) optimization methods
	  ○ RLHF and Direct Preference Optimization
	• **Exercise Session**:
	  ○ Discuss projects
		• Rafailov et al. "Direct preference optimization: Your language model is secretly a reward model." *NeurIPS 2024.*	

Week 12	**Distributed and Federated Optimization**
	• **Lecture**:
	  ○ Heterogeneity in federated optimization
	  ○ Communication compression
	  ○ Byzantine robustness
	• **Exercise Session**:
	  ○ Simulating federated learning environments
	  ○ Discuss projects
		• Wang, et al. "A field guide to federated optimization." *arXiv preprint arXiv:2107.06917* (2021).
		• Stich and Karimireddy. "The error-feedback framework: Better rates for SGD with delayed gradients and compressed communication" JMLR 2020.
		• Karimireddy et al. "Learning from history for byzantine robust optimization." *ICML 2021*.	

Week 13	**Student presentations**		In class presentations.
Week 14	**Student presentations**		In class presentations.
FINAL	**Final Exam and reports due**		Will be conducted on the university-scheduled final exam date.

Academic Integrity

The University of Southern California is foremost a learning community committed to fostering successful scholars and researchers dedicated to the pursuit of knowledge and the transmission of ideas. Academic misconduct — which includes any act of dishonesty in the production or submission of academic work (either in draft or final form) — is in contrast to the university's mission to educate students through a broad array of academic, professional, and extracurricular programs.

This course will follow the expectations for academic integrity as stated in the USC Student Handbook. All students are expected to submit assignments that are their own original work and prepared specifically for this course and section in this academic term. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s). Students suspected of engaging in academic misconduct will be reported to the Office of Academic Integrity.

Other violations of academic misconduct include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

Academic dishonesty has a far-reaching impact and is considered a serious offense against the university. Violations will result in a grade penalty, such as a failing grade on the assignment or in the course, and disciplinary action from the university itself, such as suspension or even expulsion.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment or what information requires citation and/or attribution.

Course Content Distribution and Synchronous Session Recordings Policies

USC has policies that prohibit recording and distribution of any synchronous and asynchronous course content outside of the learning environment.

Recording a university class without the express permission of the instructor and announcement to the class, or unless conducted pursuant to an Office of Student Accessibility Services (OSAS) accommodation. Recording can inhibit free discussion in the future, and thus infringe on the academic freedom of other students as well as the instructor. (Living our Unifying Values: The USC Student Handbook, page 13).

Distribution or use of notes, recordings, exams, or other intellectual property, based on university classes or lectures without the express permission of the instructor for purposes other than individual or group study. This includes but is not limited to providing materials for distribution by services publishing course materials. This restriction on unauthorized use also applies to all information, which had been distributed to students or in any way had been displayed for use in relation to the class, whether obtained in class, via email, on the internet, or via any other media. Distributing course material without the instructor's permission will be presumed to be an intentional act to facilitate or enable academic dishonestly and is strictly prohibited. (Living our Unifying Values: The USC Student Handbook, page 13).

Statement on University Academic and Support Systems

Students and Disability Accommodations:
USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Student Financial Aid and Satisfactory Academic Progress:
To be eligible for certain kinds of financial aid, students are required to maintain Satisfactory Academic Progress (SAP) toward their degree objectives. Visit the Financial Aid Office webpage for undergraduate- and graduate-level SAP eligibility requirements and the appeals process.

Support Systems:
Counseling and Mental Health - (213) 740-9355 – 24/7 on call
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

988 Suicide and Crisis Lifeline - 988 for both calls and text messages – 24/7 on call
The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline consists of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL) – 24/7 on call
Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

Office for Equity, Equal Opportunity, and Title IX (EEO-TIX) - (213) 740-5086
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-2500
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

The Office of Student Accessibility Services (OSAS) - (213) 740-0776
OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

USC Campus Support and Intervention - (213) 740-0411
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity, Equity and Inclusion - (213) 740-2101
Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-1200 – 24/7 on call
Non-emergency assistance or information.

Office of the Ombuds - (213) 821-9556 (UPC) / (323-442-0382 (HSC)
A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

Occupational Therapy Faculty Practice - (323) 442-2850 or otfp@med.usc.edu
Confidential Lifestyle Redesign services for USC students to support health-promoting habits and routines that enhance quality of life and academic performance.