# U C L A

Computer Science Department

**Course number**: CS260D
**Course Catalog Title**: Large Scale Machine Learning
**Short Title**: Large Scale Machine Learning
**Units**: 4
**Grading basis**: Letter Grades
**Format**: lecture with discussions
**GE or Major/Minor requirement**: No
**Requisites**: M146
**Description**:
To alleviate the costs and improve the robustness and generalization performance of modern machine learning models, it becomes crucial to develop methods with strong theoretical guarantees to warrant efficient, accurate, and robust learning. The course discusses advanced topics and state-of-the-art research to improve efficiency, robustness, and scalability of machine learning algorithms on large data. Topics include: Advanced Optimization, Variance Reduction, Distributed Training, Federated Learning, Data Summarization, Robust Learning, Neural Network Pruning, Neural Architecture Search, Neural Network Quantization

**Justification**:
The great success of modern machine learning systems is contingent on exceptionally large computational resources that enable training complex models on abundant data. This necessitates expensive hardware, and incurs a substantial environmental cost due to the significant energy consumption. Moreover, as datasets grow, the commonly used data collection, crowd-sourcing, and automated data labeling techniques result in noisy labels and malicious data points being ubiquitous in large real-world datasets. Such examples have a drastic effect on the performance of the trained models. The course discusses various ways to reduce costs of learning from massive data and increase robustness of current machine learning methods against noisy and malicious examples. Such a course does not exist at UCLA.

**Supplemental Information**: --

**Grading structure**:

| | | |
|---|---|---|
| Homework | 15% | 3 homework assignments |
| Class Presentation | 25% | |
| Midterm | 30% | $8^{th}$ week of the class (in class; 90 minutes) |
| Final Project | 30% | Presentation and report |

**Effective Date**: Fall 2022

**Syllabus**:
**Title: Large scale machine learning**

**Course Description:** Discussion of advanced topics and state-of-the-art research to improve efficiency and scalability of machine learning algorithms to larger data. Topics include data compression, model (neural network) compression, distributed learning methods, and improved optimization methods.

**Schedule:**

| | | | |
|---|---|---|---|
| week 0 | Sep 23 | Introduction | |
| week 1 | Sep 28 | [Adv. opt.] Momentum & adaptive lr (review), Variance reduction | |
| | Sep 30 | [distributed opt.] distributed SGD, Hogwild | |
| week 2 | Oct 5 | [distributed opt.] distributed SGD, Hogwild | HW1 |
| | Oct 7 | [distributed opt.] federated learning | |
| week 3 | Oct 12 | [distributed opt.] federated learning | **HW1 due** |
| | Oct 14 | [data compression] submodularity | |
| week 4 | Oct 19 | [data compression] distributed/streaming algs | HW2 |
| | Oct 21 | [data compression] subsets for ML | |
| week 5 | Oct 26 | [data compression] robust subsets | **HW2 due** |
| | Oct 28 | [model compression] NN pruning | |
| week 6 | Nov 2 | [model compression] NN pruning | HW3 |
| | Nov 4 | [model compression] NAS | |
| week 7 | Nov 9 | [model compression] NAS | **HW3 due** |
| | Nov 11 | Veteran's day | |
| week 8 | Nov 16 | **Midterm** | |
| | Nov 18 | [model compression] quantization | |
| week 9 | Nov 23 | [model compression] quantization | |
| | Nov 25 | Thanksgiving | |
| week 10 | Nov 30 | Final project presentations | |
| | Dec 2 | Final project presentations | |

**Enrollment and student evaluation numbers from past course offering**

This course was offered as CS 269 in Fall 2021. 24 students took the course with an additional 30-40 who wanted to take but were unable to due to course capacity. The class went very well and most students felt they learned a lot from the course.

Summary of Reviews (please see the end of this file).

Value - You have learned something you consider valuable.
Median 9 Average 8.64

Overall - Your overall rating of the instructor.

Median 9 Average 8.73

Overall - Your overall rating of the course.
Median 9 Average 8.64

**Students' Comments** (a sample trying to demonstrate both positive and negative comments)

**Please identify what you perceive to be the real strengths and weaknesses of this instructor and course.**

I really enjoy the class, even though I have a hard time understanding the math part behind the theories. The professor used homework assignments to make me understand the math part more clearly since she turned the theories into practical problems, which is a lot easier to understand. She is also pretty helpful. When you have any questions related to the class materials, she tried her best to help you understanding all the concepts.

I learned a lot and enjoyed this class.

I thought the course overall was very interesting, especially submodular optimization, NAS, and pruning. I would have liked to have some practical coding assignments (nothing super complicated) along with the standard homework to get more intuition about some of the concepts.

Prof. Mirzasoleiman is easily one of the best professors I've had at UCLA. She is extremely knowledgeable in large-scale machine learning and is an excellent teacher. She really ensured that students were engaged (through real-world examples, jokes, etc.) and reviewed course content regularly so everybody felt comfortable with it. Her homework assignments were fantastic and really helped me understand the course content better. Her midterm was extremely fair. It was very clear that she cared about student success and understanding. It was apparent from the creativity and enthusiasm displayed in the final course projects that all the students in the class genuinely enjoyed and learned a lot from the class. This class has significantly shaped my research interests. I hope, for the benefit of all CS students, that this class becomes a regular offering.

Professor is incredible! Easily one of the best classes I've taken here at UCLA. I'm surprised that the CS department has not made this a permanent class yet, I'm not sure what they are waiting for. One of the best professors I've had--she is extremely clear, helpful, and supportive. Could not have asked for a better experience.

Strengths: Relevant practical concepts in ML taught, Prof Baharan is very helpful and is concerned about student learning. Weakness: Slides used in the course were sometimes not very informational, Prof went too fast in the beginning of the course in which most concepts were mathematical

Very good course!