

# 18-847F: Special Topics in Computer Systems

## Foundations of Cloud and Machine Learning Infrastructure



# Lecture 1: Introduction and Logistics

## Foundations of Cloud and Machine Learning Infrastructure



# Graduate Seminar Class

(Almost) no lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project (No Exams!)

# Learning Objectives

- Know the state-of-the-art frameworks in cloud and machine learning and their theoretical foundations
- Read and provide constructive criticism of research papers
- Present to an audience, and answer their questions
- Do creative, collaborate research

# Why study Cloud and ML infrastructure?



What are the largest words after 'Big Data'?

# Big Data Gold Rush



Who got rich in the  
California gold rush?

# Big Data Gold Rush



Who got rich in the California gold rush?

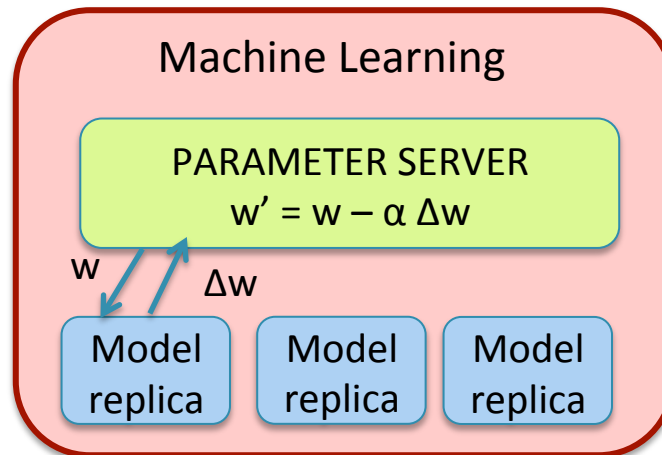




Google Compute Engine




In the Big Data rush, it's the infrastructure companies

# Topics Covered





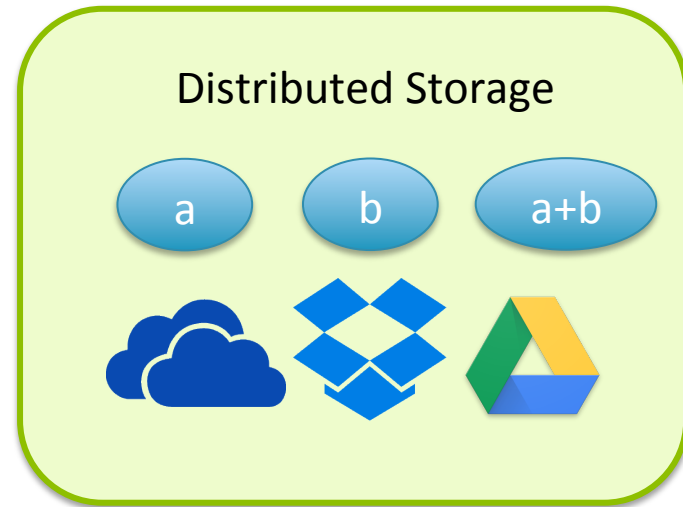
# Topics Covered



- Scheduling in Parallel Computing
  - MapReduce, Spark
  - Straggler Replication
- Task Replication in Queueing Systems
  - Coded MapReduce

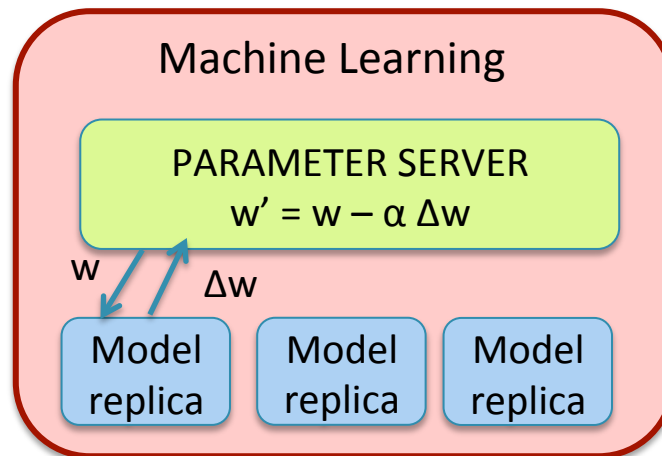
# Topics Covered

- Coding for locality/repair
- Systems implementation of codes
- Reducing latency in content download

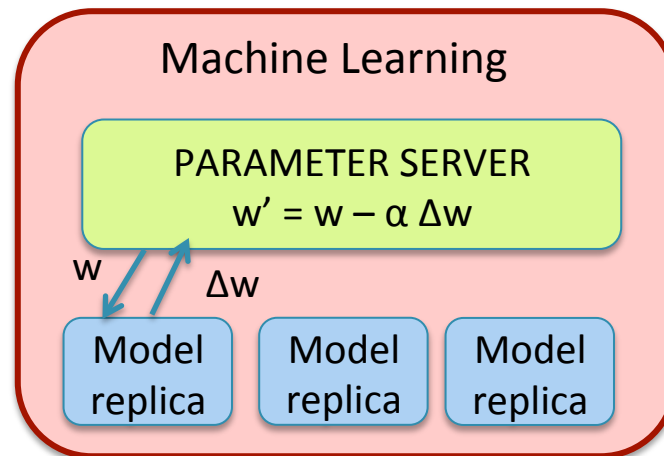




# Topics Covered

- SGD and Support Vector Machines
- Backpropagation, LeNet, AlexNet, GoogleNet
  - Distributed Gradient Descent
    - Hyper-parameter tuning
- GANs, Deep reinforcement learning



# Topics Covered



# Instructor: Gauri Joshi



B.Tech+M.Tech  
2005-2010



SM + PhD  
2010-2016



Research Staff Member  
2016-2017

**Carnegie  
Mellon  
University**

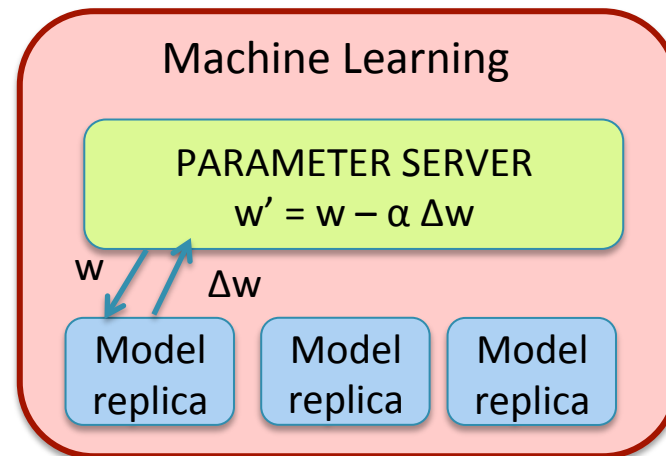
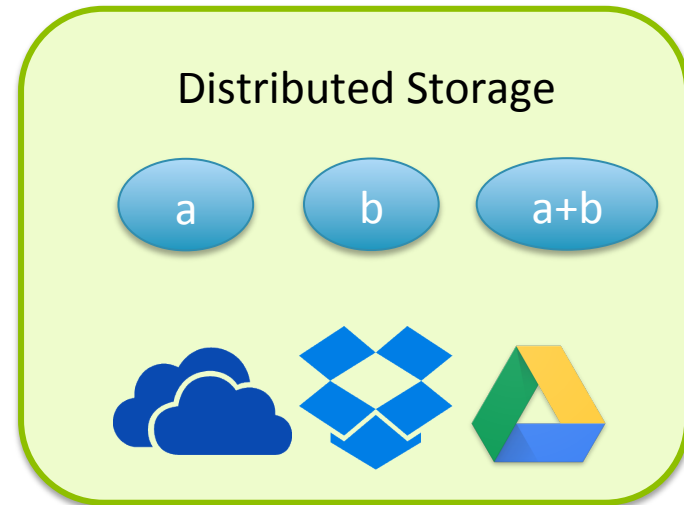

Assistant Professor  
Fall 2017 -

Internships



Bell Labs 

# Have worked in all these areas



# Student Introductions

- Name?
- Department?
- Masters or PhD?
- Previous related classes (if any)?
- What you are looking to learn from this class?

# Class Hours and Website(s)

- When: Mon, Wed 4:30-6:00 pm
- Where: Scaife Hall 222
- Class Website (Readings, Schedule):  
<https://www.andrew.cmu.edu/user/gaurij/18-847F-Fall-2017.html>
- Canvas Site (Readings, Assignments, Projects):  
<https://canvas.cmu.edu/>



# Reading Material

Papers will be posted on the class website or on Canvas

- Book chapters
- Survey papers
- Theory papers (Scheduling, Queuing, Coding, Optimization)
- Systems papers (Cloud, Machine Learning)

Additional reference books listed in the syllabus

# Instructor and Office Hours

Prof. Gauri Joshi, ECE Dept

**Email:** gaurij [AT] andrew.cmu.edu

**Office Location:** CIC 4105

**Office Hours:** Wed 2:00-3:00pm or by appointment

# Graduate Seminar Class

(Almost) no lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

# Lectures

- Next week: Deeper Overview of course topics
- 3-4 Guest lectures during the semester by authors of papers relevant to this class

# Graduate Seminar Class

(Almost) no lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

# Homeworks (45%)

- Class divided into two groups; read the paper assigned to your group
- Submit paper review (due 9:00 am before class)
  - Two reviews per week (advise: finish them together!)
- Discussion with your group is okay, but write reviews in your own words. List collaborators in the homework submission

# Paper Review Format

- Brief Summary of the paper
  - Reflects your understanding of the paper
  - No judgments (positive or negative) here
- High-level technical/writing comments
  - Significance & correctness of results (Don't be mean!)
  - Paper organization
- Low-level technical/writing comments
  - Smaller clarifications, corrections, typos
- Discussion Questions for Class (at least 2)
  - Confusions about the paper, open research directions

# Tentative Grading Rubric (Total: 10 pts)

- Clarity, Organization (1 pt)
- Understanding of the paper (5 pts)
- High-level comments (2 pts)
- Low-level comments (2 pt)



# Graduate Seminar Class

(Almost) no lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

# Class Presentations (20%)

- One student from each group signs up for presentation, at least 1 week in advance
- Each student will present ~2 times in the semester
- 20 min presentation, followed by 25 min discussion
  - Motivation and Related work
  - Summary of main results
  - Your views on the paper

# Tentative Grading Rubric (Total: 10 pts)

- Motivation (2 pts)
- Clarity (2 pts)
- Correctness (2 pts)
- Engaging the audience (2 pts)
- Extra research, going beyond the paper (2 pts)

# Graduate Seminar Class

(Almost) no lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

# Tentative Grading Rubric (Total: 5 pts)

- Discussion questions (2 pt)
- Attendance and attention (1 pt)
- Speaking up in class (1 pt)
- Insightful Questions/Comments (1 pt)

# Graduate Seminar Class

(Almost) no lectures

Reading research papers

Student presentations

Class Discussions

Final Research Project

# Research Project (25%)

- Groups of 2 or 3
- Original research on a topic of your choice
  - Topics aligned with your research allowed and encouraged
  - If you can't think of topics, come talk to me!
- Possible Project Types:
  - New theoretical analysis
  - Implementation using one of the frameworks discussed
  - In-depth literature survey of a particular topic

# Timeline

- 1-page proposal due Sept 27
- Publishable quality report (max 5 pg) in ACM format
  - Initial draft due: Nov 22
  - Final report due: Dec 8
- Peer-review 2 other reports
- Last 2 weeks of class: Presentations (20 min each)



# Tentative Grading Rubric (Total: 20 pts)

- Originality (1 pts)
- Review of Related Work (1 pts)
- Writing and Organization (2 pts)
- Technical Results (5 pts)
- Peer-Review of Other reports (1 pts)
- Final presentation (10 pts)

## In Summary..

- Paper Reading
- Submitting Reviews
- Class Presentations (~2 in the semester)
- Final Project

Might seem like a lot of work but..

- You will get fast and efficient at reading papers
- The project will be a fun, collaborative exercise
- No exams!

# TO DO

- Form groups for class readings
- Sign-up for presentation
- Form groups for class projects
- Start reading the papers
- Start thinking about projects

# Next Class

History and overview of cloud and machine learning infrastructure

Will give you additional time for the first paper reviews and presentations