



Deep Learning

Designing, Visualizing and Understanding Deep Neural Networks

EECS 182/282A

Gireeja Ranade
University of California, Berkeley

First lecture outline

- Introductions
- Administrivia
- Deep Learning

Faculty



Prof. Gireeja Ranade

4334 Gateway

ranade@eecs.berkeley.edu

Professor Office Hours: Right after lecture 11-12 pm on Tu/Th Gateway location

Course Staff

- GSI: Matteo Guarrera
- GSI: Mert Cemri
- GSI: Sizhe Chen


Resources

- Website: eecs182.org
- EdSTEM: edstem.eecs182.org

<https://edstem.org/us/courses/104046/discussion>

- Recommended textbook: Prince, Understanding deep learning

<https://udlbook.github.io/udlbook/>



Enrollment

- Class will not expand
- If you are on the waitlist and want to be in the class, please submit HW0 and complete all the work as though you are in the class.
- Students who do not submit HW0 will be dropped from the class / will not be allowed into the class.
- Lectures will not be recorded, lecture attendance is expected.

Discussion sessions

- Weekly discussion sections conducted by the course staff
 - Friday at various times
- Choose any one section to attend. All sections cover the same material.
- Discussion participation is required and graded.

Getting help

- Instructor office hours after class
 - GSI office hours Tuesday 4-6 pm Gateway B1009
 - Discussion sections
 - Ed forum
 - Each other
-
- For personal questions please make a private question on EdSTEM, faculty and GSIs will be monitoring this.
 - Email professor only for very private questions where you prefer GSIs not know about the question.
 - Ed is the fastest way to get help.

Accommodations

- We are happy to work with the DSP office to ensure all students have appropriate accommodations.
- Please reach out to us as needed.
- Homework extensions are automatic if you have a DSP letter, Gradescope will accept all submissions before the final exam.
- Project extensions cannot be offered due to the nature of the course. However, the work on the project is typically distributed over a long period of time.

Grading

- 0% homework
 - 60% Final exam
 - 15% Project
 - 15% Quizzes
 - 10% Participation
-
- Both the final exam and the project gate overall grades: not passing either would result in failing the class.
 - Fixed bins --- no curve.

Homework policies

- Weekly homeworks
 - Due Wednesday at 10:59 pm
 - Gradescope deadline will be set to the end of semester
 - Strongly recommend keeping up
 - Create accountability through a study group
 - Homework is formally worth 0% in the course
-
- Doing the homework is the best way to learn the material.
Historically, there have been exam questions very similar to the homework.

Quizzes

- 3-4 quizzes throughout semester
- Held in discussion section
- 20 minutes
- Dates announced next week
- Problems drawn exclusively from homework.

Exam policies

- Time-conflict enrollment is not permitted
- The final exam will be at the time set by the registrar.
- We will not be offering alternate exams under any circumstances.
- Do NOT enroll in a course with lecture scheduled at the same time as this class. The final exams will be at the same time. No accommodations will be made.
- Make a private post on EdStem if you have a DSP accommodation letter. No accommodations for DSP students due to time-conflict enrollment.

Should you take this course?

- We expect
 - Machine learning at the level of CS189/Stat 154
 - Linear Algebra/Optimization at the level of EECS 127
 - Probability at the level of EECS 126/Data 140/ Stat 134
- Given the low levels of staffing from the department, we will not be able to support having students catch up on background material.
- To get the most out of the course, it is in your interest to have solid understanding of these courses.

Final project

- More details will be announced later on Ed.
- Groups of 4: can mix/match undergrad and grad students
- There will be a couple of options:
 - Each open-ended but bounded scope that let you approach the research frontier and build on published papers.
 - Grad students from outside EECS will be permitted to lead projects outside of our specified topics.
- Bring-your-own compute: we are trying to secure some compute resources for you, but you might want to use more to maximize your learning.
- You will have to submit proposals, meet with staff, submit drafts, and individually perform peer reviews for other projects.
- Project grades will be the product of your individual peer reviewing grade and the group's project quality grade.

Participation/Citizenship

- Overall: making the class better for everyone involved
- Participation:
 - Attend and actively participate in all the weekly discussions
 - Be actively helpful in the online forum
 - Reach out if you need an accommodation for participation



What is Deep Learning?

Deep Learning is the use of (simulated) analog circuits for information processing, where the parameters associated with the circuits are set using an optimization algorithm driven by data.

These circuits are typically deep in that information flows through multiple circuit elements before the output.

The goal is that the learned circuit has approximately captured the relevant patterns underlying the data used to train it, so that we can use this circuit to solve problems involving new data for which that pattern is relevant.

Studying Deep Learning...



THIS IS YOUR MACHINE LEARNING SYSTEM?



YUP! YOU POUR THE DATA INTO THIS BIG
PILE OF LINEAR ALGEBRA, THEN COLLECT
THE ANSWERS ON THE OTHER SIDE.

WHAT IF THE ANSWERS ARE WRONG?


JUST STIR THE PILE UNTIL
THEY START LOOKING RIGHT.



Engineering or Alchemy?



Quelle: Deutsche Fotothek




1950

1950: Turing describes how learning could be a path to machine intelligence



1960

1957: Rosenblatt's **perceptron** proposed as a practical learning method
And LMS was developed in the adaptive signal processing/control field



1970

1969: Minsky & Papert publish book describing fundamental limitations of neural networks

1980

1986: Backpropagation as a practical method for training deep nets
1989: LeNet (neural network for handwriting recognition)

1990


2000

Huge wave of interest in ML community in probabilistic methods, convex optimization, but mostly in shallow models

2010

~2006: deep neural networks start gaining more attention

2012: Krizhevsky's AlexNet paper beats all other methods on ImageNet



[FIG2] Circuit of the first adaptive neuron.

Mathematics in the Age of AI

<https://arxiv.org/pdf/2608.16753>

MATHEMATICS IN THE AGE OF AI

TERENCE TAO

ABSTRACT. An essay, based on a public lecture delivered at the 2026 International Congress of Mathematicians, on how the mathematical community might respond to the arrival of artificial intelligence tools that are capable of performing research-level mathematical tasks. Rather than debating the capabilities of such tools, we condition on the hypothesis that these capabilities will arrive, and examine instead a question that is orthogonal to it: what the goals and values of mathematical research actually are. The problem-solving component of mathematics is used as a case study.

1. A HISTORICAL PROLOGUE

For centuries, mathematics operated successfully on “naive” foundations. Practicing mathematicians proved theorems about sets, numbers, and infinities without feeling any particular need to say precisely what these quantities actually were (or what a proof itself was, for that matter); such questions were largely delegated to philosophers, and the working mathematician was free to get on with the mathematics.

But in the early twentieth century, discoveries such as Russell’s paradox in 1901 [14] and the Gödel incompleteness theorems in 1931 [8] forced mathematicians to critically re-examine assumptions about their subject that had previously been left implicit. The axioms of naive set theory contradicted each other; and a formal system could not simultaneously be consistent, sufficiently expressive, and capable of proving its own consistency.

The resulting *crisis in foundations*, lasting roughly from 1900 to 1930, was a genuinely turbulent period for the subject. But the end product of that turbulence was extremely valuable: an explicit, rigorous, and standardized foundational framework, in which the ob-