Data 188: Introduction to Deep Learning

Introduction and Logistics

Speaker: Eric Kim
Lecture 01 (Week 01)
2026-01-20, Spring 2026. UC Berkeley.

Acknowledgement: slides are based on CMU's Deep Learning Systems course (10-414/714, Fall 2025). Thanks to: Prof. Zico Kolter, Prof. Tianqi Chen, Prof. Tim Dettmers.

Announcements

- No discussion sections or office hours during week 1
  - First discussion section: Wed Jan 28th 2026
  - First office hours: Mon Jan 26th 2026
- HW0 will be released today
  - Due: Feb 5th 2026 (11:59 PM PST)
  - Start early!

Outline

Course overview

Why study deep learning?

Course info and logistics

Outline

Course overview

Why study deep learning?

Course info and logistics

Course Learning Objectives

This course will provide you with a **solid foundation in deep learning**.

You will learn about the underlying concepts of modern deep learning, like: automatic differentiation, common neural network architectures, optimization, and efficient operations on systems like GPUs

Goal: after taking this course, you will be able to read an academic paper about deep learning and understand what's going on well enough to (with effort/research) re-implement the paper.

Quick Course Tour

- Softmax classification models
- Backpropagation, "autograd"
- Convolutional Neural Networks ("convnets")
- Sequence models (eg the Transformer model)
- Applications to computer vision, natural language processing ("NLP")
- ...and more!

Outline

Course overview

Why study deep learning?

Course info and logistics

A (very, very brief) History of Deep Learning

- In The Beginning: there were manually crafted input features

- Ex: Scale-Invariant Feature Transform ("SIFT") [David Lowe, 2004]
  - For its time, an effective way of calculating image keypoint descriptors that are (somewhat) robust across viewpoint and illumination changes.
  - Applications: 3D geometry, image representation, robotics.
  - Ex: SLAM + SIFT [link_slides_2007]

- ...followed by cottage industry of image descriptors: SURF, GIST, HoG, etc.

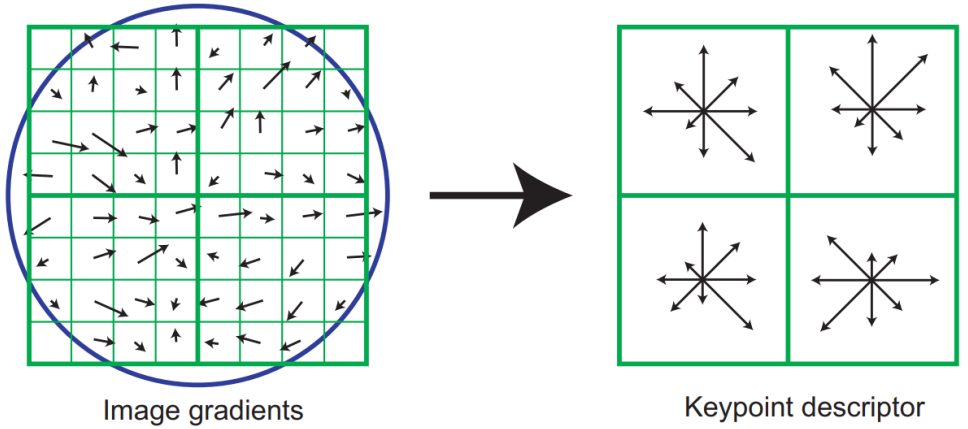

Figure 12: The training images for two objects are shown on the left. These can be recognized in a cluttered image with extensive occlusion, shown in the middle. The results of recognition are shown on the right. A parallelogram is drawn around each recognized object showing the boundaries of the original training image under the affine transformation solved for during recognition. Smaller squares indicate the keypoints that were used for recognition.

Deep Learning, aka Feature Learning

- **Idea**: rather than hand-crafting features (eg SIFT), can we learn effective feature representations **directly from the data**?
  - And: can machine-learned feature representations outperform our hand-crafted features derived from (decades of) domain expertise?
- **The answer: A resounding yes!**

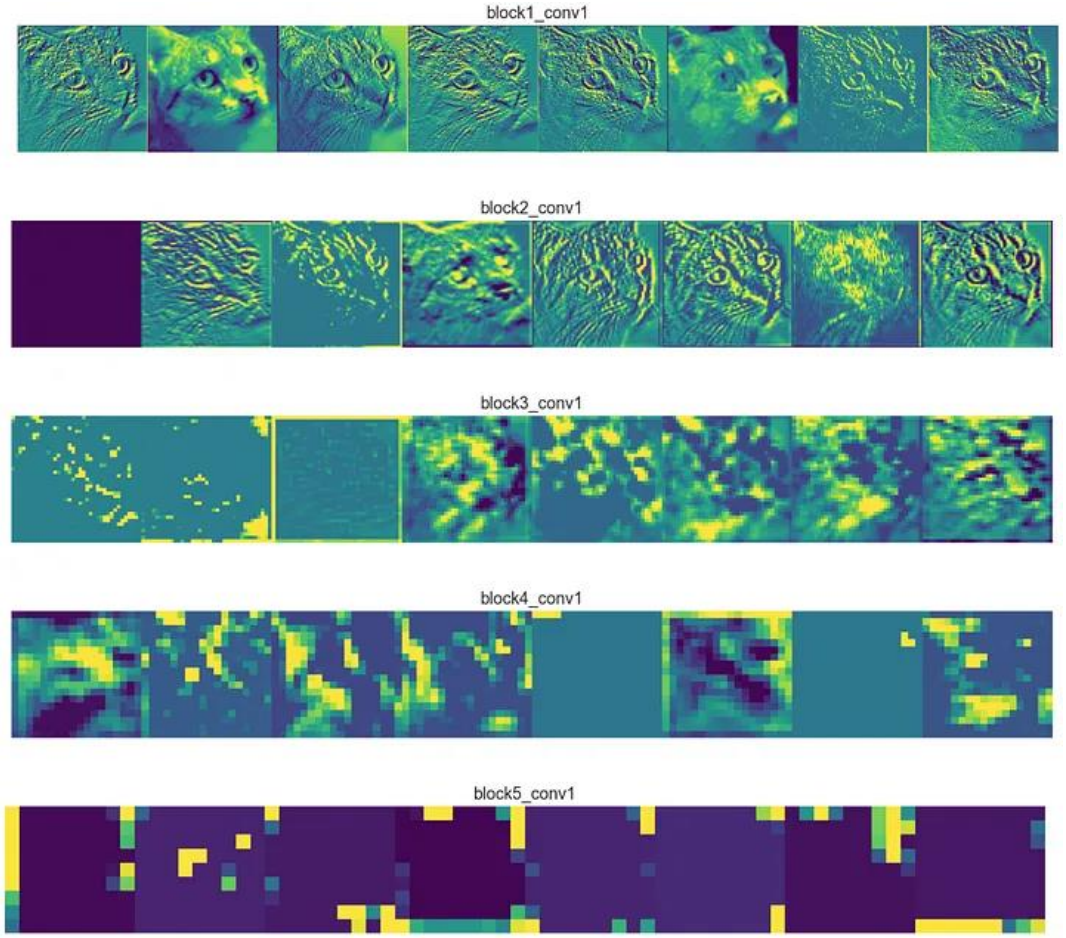

Low level: edge detectors

High level: (hopefully) object detectors.

Pictured: intermediate activations of a convnet on an image of a cat.

https://medium.com/@saba99/feature-map-35ba7e6c689e

Deep learning wins

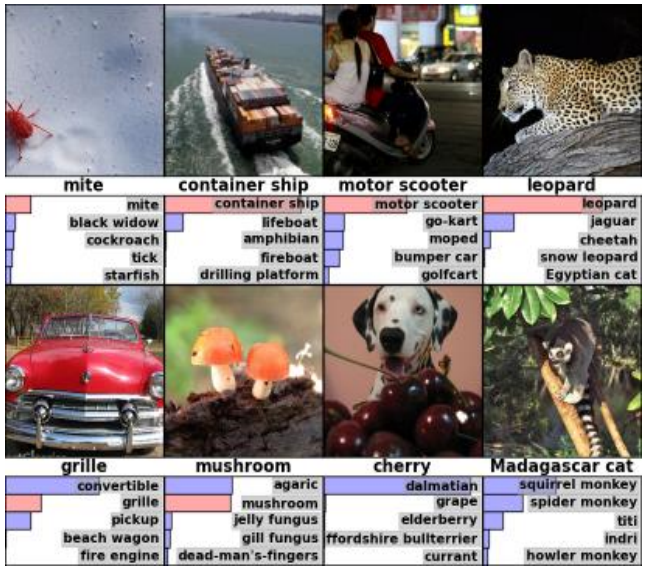

AlexNet (Krizhevsky et al., 2012)

AlphaGo (Silver et al., 2016)

StyleGAN (Karras et al., 2018)

One of the early "big wins" of deep learning (ConvNets, image classification). Beat out previous approaches using handcrafted features (eg SIFT).
Proved out idea that "deep learning + GPU compute" pays off!

Deep learning wins

Write a summary of the "Deep Learning Systems: Algorithms and Implementation" course at CMU.

As of my last update in September 2021, I don't have specific details about a course named "Deep Learning Systems: Algorithms and Implementation" at Carnegie Mellon University (CMU). However, given the title and based on the general topics covered in similar courses around the world, I can provide a hypothetical summary:

Deep Learning Systems: Algorithms and Implementation

Course Overview:
This course aims to provide students with a comprehensive understanding of the algorithms that power deep learning systems, as well as the practical skills needed to implement, optimize, and deploy these systems in real-world applications.

Course Objectives:

1. Understand the foundational algorithms behind deep learning models.
2. Gain hands-on experience in designing, training, and fine-tuning deep learning models.
3. Learn to implement efficient and scalable deep learning systems.
4. Understand the challenges and solutions for deploying deep learning models in various environments.

Course Topics:

1. **Introduction to Deep Learning Systems:** Overview of the landscape, challenges, and importance of efficient system design.

ChatGPT
(OpenAI et al., 2022)

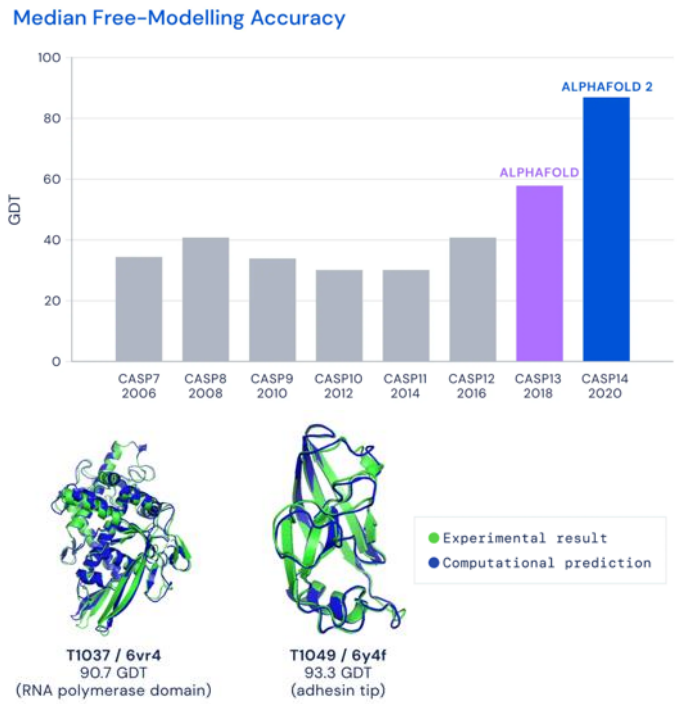

AlphaFold 2 (Jumper et al., 2021)

A dog dressed as a university professor nervously preparing his first lecture of the semester, 10 minutes before the start of class. Oil painting on canvas.

Stable Diffusion
(Rombach et al., 2022)

Multimodal models

As of 2026, models that operate on multiple input modalities (eg text and image) are extremely popular (image captioning, image generation, ...).

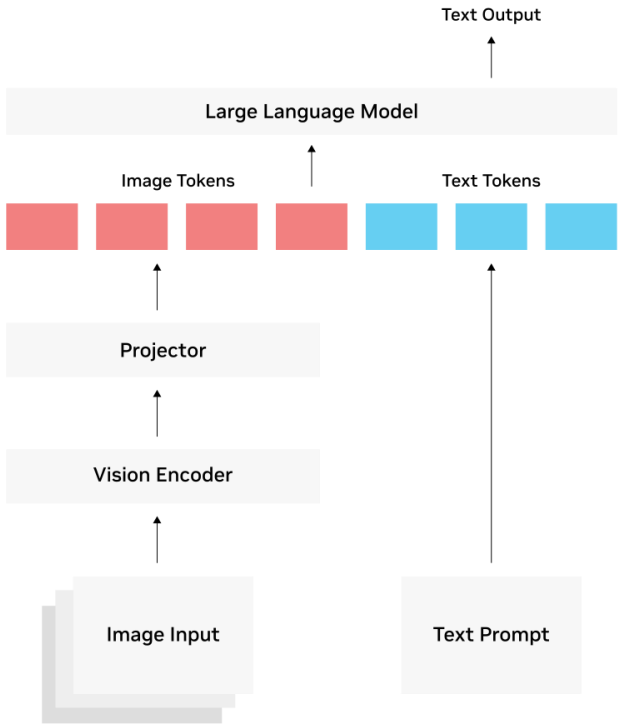

https://www.nvidia.com/en-us/glossary/vision-language-models/

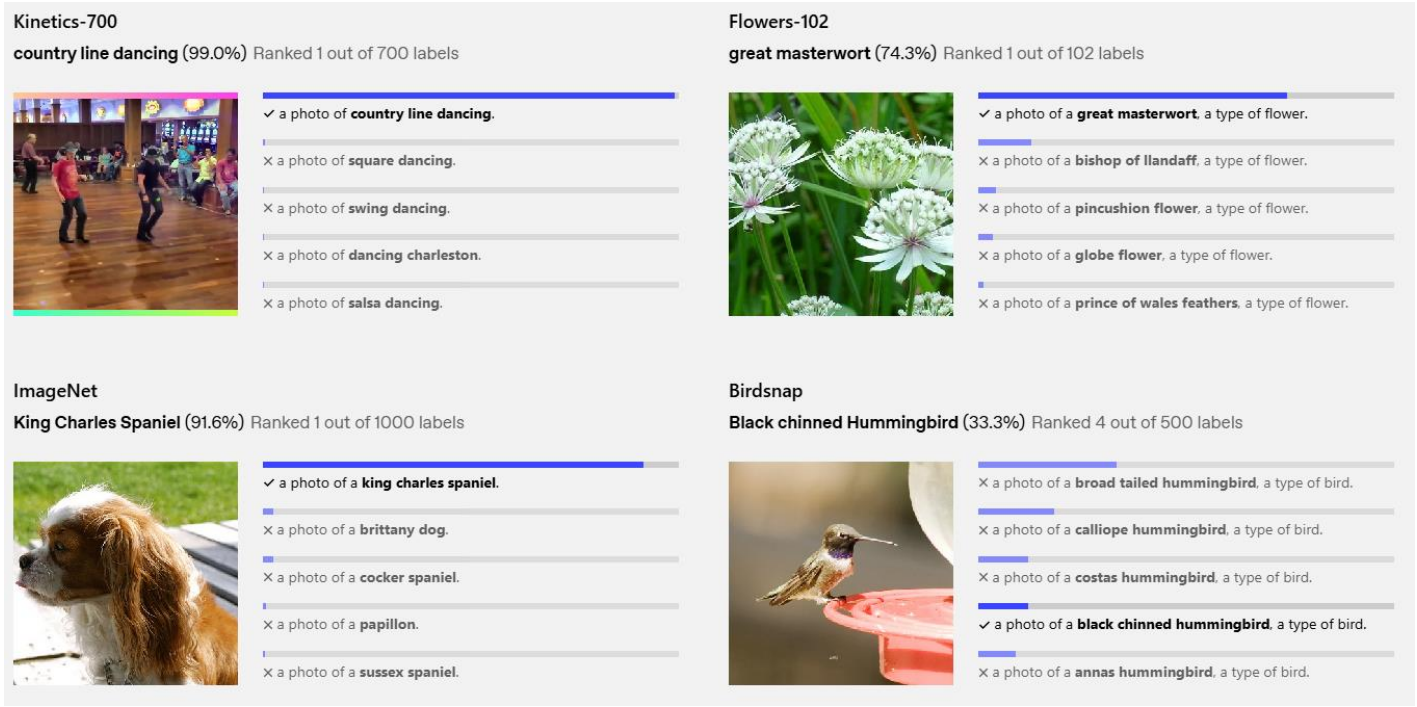

https://openai.com/index/clip/

The rise of deep learning

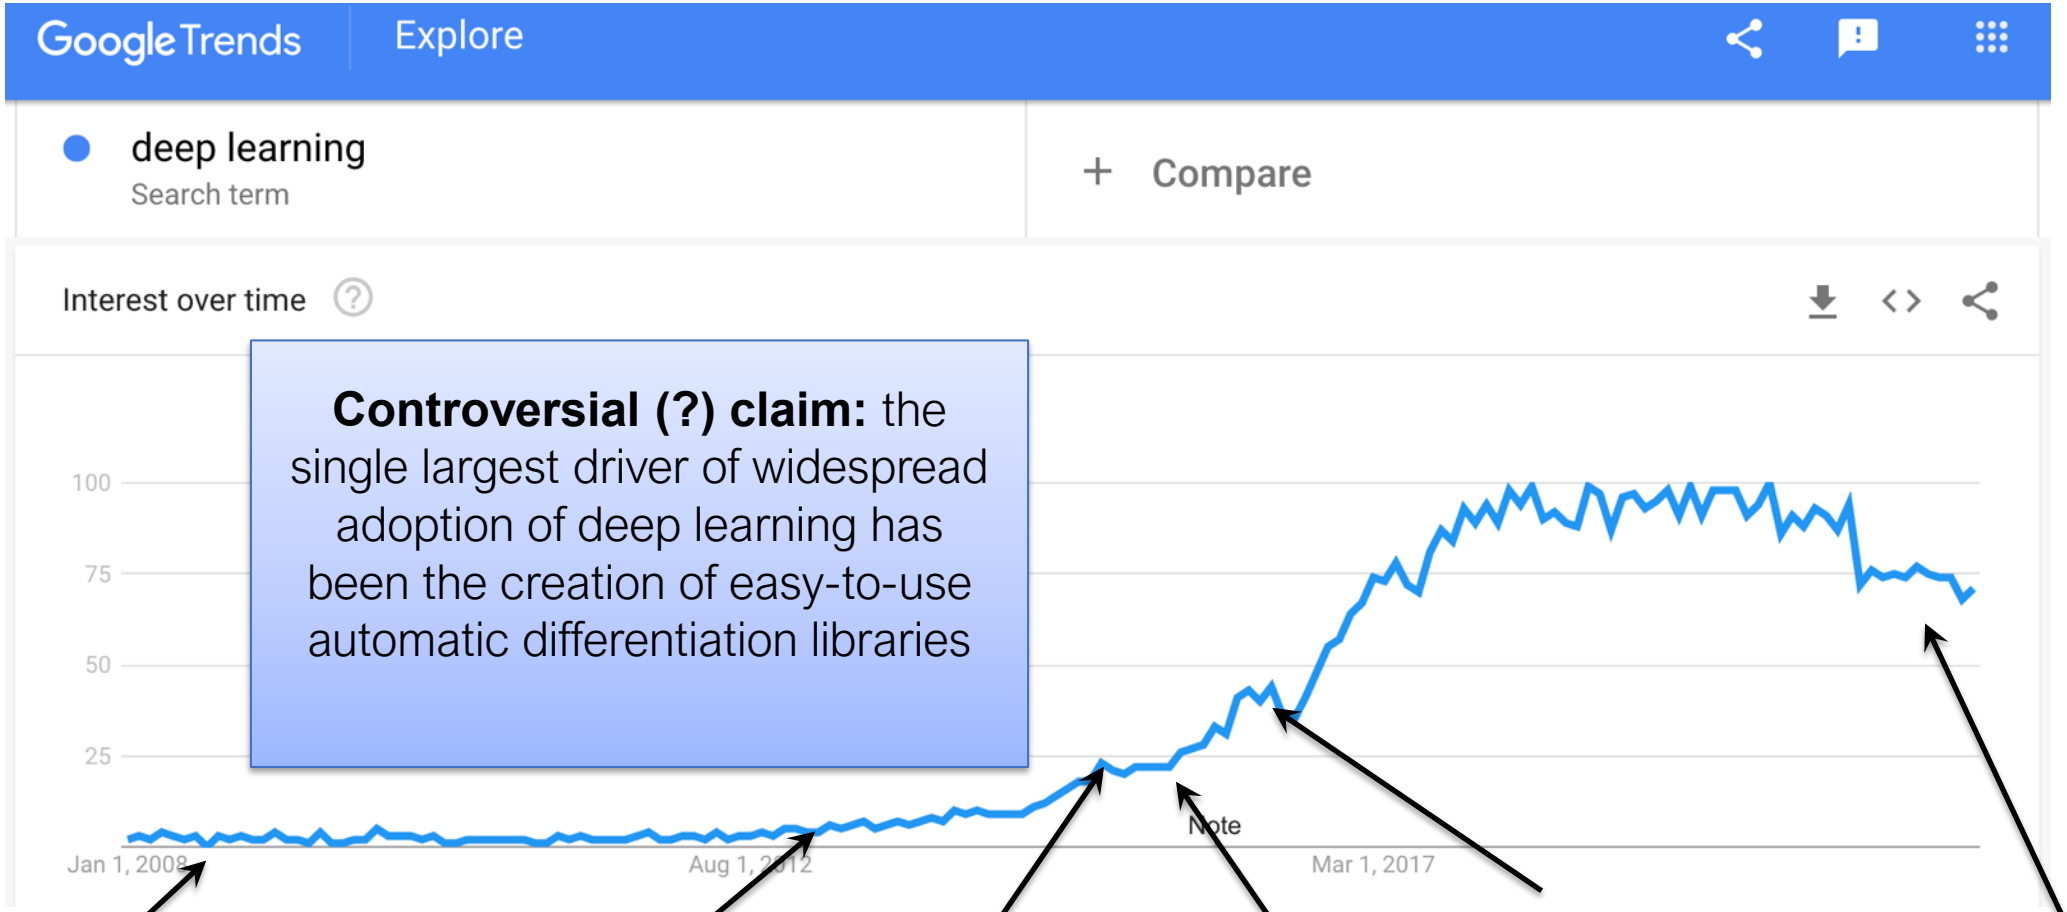

Controversial (?) claim: the single largest driver of widespread adoption of deep learning has been the creation of easy-to-use automatic differentiation libraries

Deep learning gains traction at NeurIPS

AlexNet

Keras released

TensorFlow released

PyTorch released

(Uh oh?)

Working on deep learning ten years ago

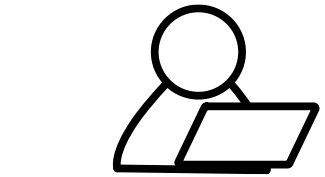

Researcher

AlexNet ... ...

ML Models

44k lines of code

Six months

IMAGENET

Data

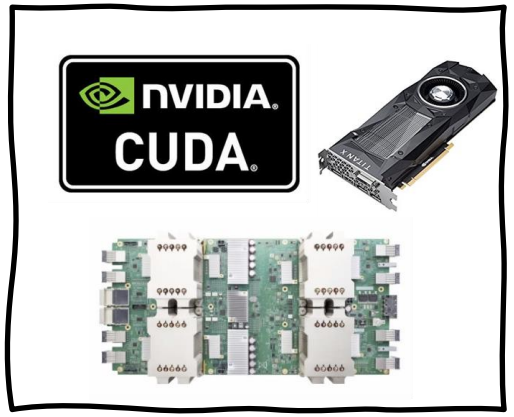

Compute

Based on real story

Working on deep learning now

Researcher

ResNet

Transformer

…

ML Models

100 lines of code

A few hours

Deep learning systems

IMAGENET

Data

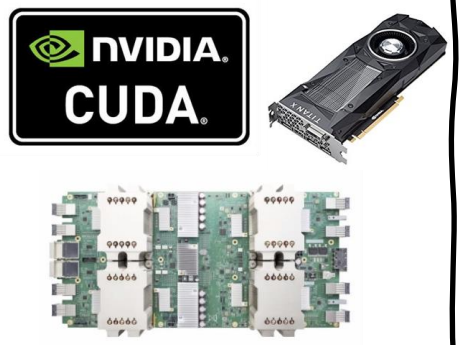

Compute

Based on real story

Working on deep learning with agents

Researcher

ResNet

Transformer

…

ML Models

500 lines of code

~~A few hours~~ Minutes

Deep learning systems

IMAGENET

Data

Compute

Based on real story

Working on deep learning (continuously evolving)

Researcher

FLUX deepseek GPT-5

Bigger models

Deep learning systems ?

LAION wikitext

IMAGENET internet

Large high-quality data

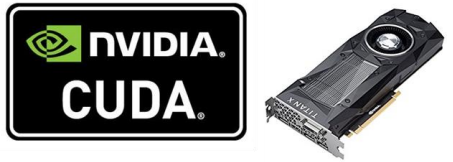

More diverse Compute

Outline

Course overview

Why study deep learning?

Course info and logistics

Course Instructor

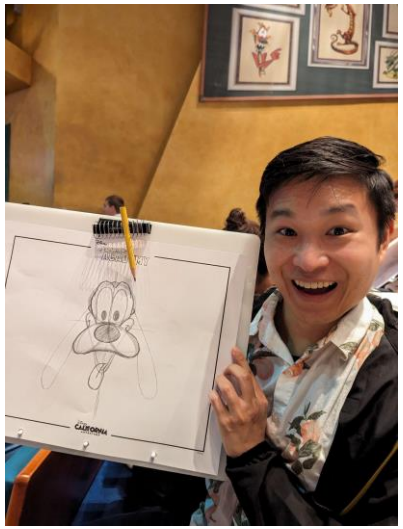

Eric Kim: ekim555@berkeley.edu (more about me here)

Brief history:
UC Berkeley (2011, L&S CS), UCLA (2016, MS with focus in AI/ML and computer vision).
ML engineer at Pinterest Visual Search team (2017 to 2026, aka now!).

What I do at work: ML modeling (object detection, representation learning, multimodal models, recommendation systems), ML serving infra (training/serving, GPU), big-data pipelines, and more…

What I teach: Data C88C, Data C182 / Data 188 (Intro to Deep Learning)

What I do for fun: Music (piano, guitar, bass, vocals. Rock, pop, blues, jazz, musicals, classical), video games, running (casually). I love Pixar movies, and animation/film in general!

Meet the staff!

TA's

Andria Xu SHE/HER

andria.xu@berkeley.edu

About Me: Hello! I'm Andria, a current junior from the Bay Area. I'm very excited to be teaching Data 188 this semester! Previously, I taught Data C88C. I also do research in reinforcement learning! In my free time, I enjoy hiking, playing board games, and trying random new things. Looking forward to getting to know you all! :D

Rebecca Dang SHE/HER

rdang@berkeley.edu

About Me: Hey there! I'm a 5th Year MS EECS student and I'm excited to teach DATA 188 for the first time! In the past, I've taught DATA C88C and DATA 101. In my free time, I love fangirling over the most recent book/TV show/movie I'm obsessed with, hiking, and playing guitar. Come talk to me about anything :)

Terry (Taehan) Kim

terry.kim@berkeley.edu

About Me: Hi! I'm Terry (Taehan) Kim, a senior majoring in Computer Science and a and lover. Super excited for this semester!

Zekai Wang HE/HIM/HIS

zekai.wang@berkeley.edu

About Me: Hi! I'm Zekai, a 4th year CS and Applied Math major. I am interested in machine learning and robotics. Looking forward to meeting you all!

Tutors

Ryan Michael Tom HE/HIM

ryantom@berkeley.edu

About Me: Hi everyone! I'm Ryan, a 3rd year EECS major who enjoys playing basketball, running, raving, and adding new places to Beli! Very excited to meet you all :)

Yihang Chen HE/HIM

yhc0720@berkeley.edu

About Me: Hi! I'm Yihang, a senior majoring in Computer Science and Data Science. I'm thrilled to be part of the Data 188 course staff and look forward to working with students throughout the semester. In my free time, I enjoy snowboarding and spending time outdoors.

Learn more about your TA's + Tutors here: https://data-188-berkeley.github.io/sp26/staff/

Enrollment, Waitlists

As of today (2026-01-20), the course is full (178/160 seats. Waitlist: 71/112).

The department is prioritizing enrollment into Data 188 for DSUS seniors graduating in SP26/SU26 that need to fulfill their MLDM major requirement (which Data 188 fulfills).

If you meet the above conditions, and are on the waitlist: please email ds-advising@berkeley.edu (and cc me, ekim555@berkeley.edu) to see about your chances of enrolling into the course.

If you don't meet the above, and are very high on the waitlist: unfortunately it's unlikely you will get into the course this time. I'm sorry!

For more info, see this Ed post: [link]

Course platforms: course website

- Course website: https://data-188-berkeley.github.io/sp26/
- Contains useful info like: syllabus, lecture slides, homework assignments, discussion pdfs, office hour times, etc
- Updated regularly.

Course platforms: Edstem

- Edstem: https://edstem.org/us/courses/91872/discussion

- Primary way for us to communicate with you
  - Keep an eye out for announcements!

- The primary way for you to ask questions!
  - Asking questions on Ed is the best (and fastest!) way to get an answer.
  - Students: you're also welcome (and encouraged!) to help answer questions, foster discussion etc. Let's build a community!

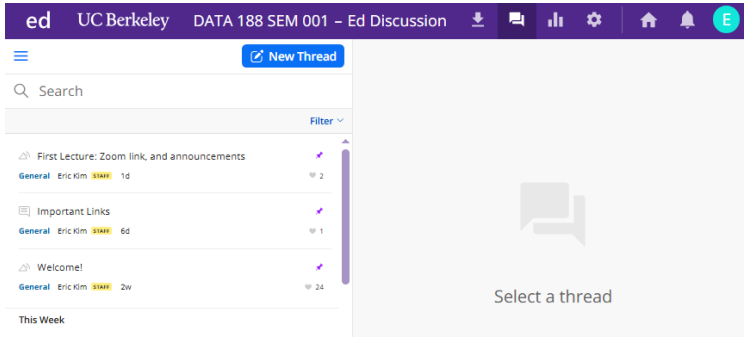

Course platforms: Gradescope

- Gradescope: https://www.gradescope.com/courses/1217889
- This is where you will submit your assignments
- Gradescope autograder: unlimited attempts allowed!

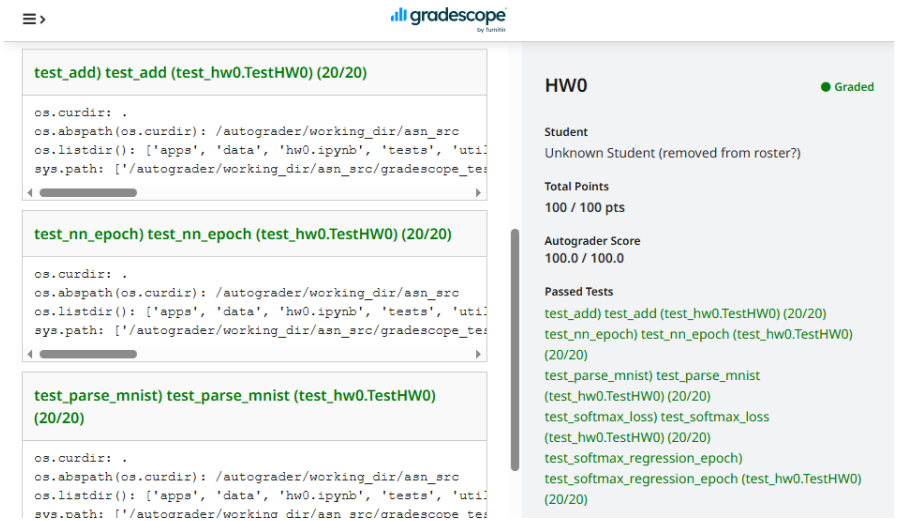

Course platforms: bCourses Media Gallery

- bCourses Media Gallery: https://bcourses.berkeley.edu/courses/1551684/external_tools/90481
- This is where you can find lecture recordings
- Other than this, we won't use bCourses for this course.

Prerequisites (more)

(roughly in order of priority)
Python programming. Course assignments will consist of a lot of Python programming, particularly with libraries like numpy and Colab/Jupyter.

Linear algebra. Primarily more applied/numerical things (eg matrix-vector calculations) rather than abstract things (ex: Math 110 is overkill for this course).

Basic ML knowledge. Ex: ideally you have studied things like: logistic regression and supervised training before (ex: Data 100).

Calculus. You should be familiar with taking simple (univariate) derivatives, and the common calculus techniques (ex: chain rule).

Tip: HW0 is a decent litmus test for whether you are ready to take the course. It's OK if you're not strong in all of the above, you can pick up what you need to during the course (just be prepared to put in extra time+effort!)

Course activities

This course will consist of the following:

1. Lectures (online, Zoom)
2. Discussion sections (in-person)
3. Office hours (Zoom, in-person)
4. Programming-based (individual) homework assignments
5. Exams (midterm, final)

Important to take part in all of these to get the full value from the course

Ex: one does not learn how to play an instrument by watching videos: one must practice, struggle, trial and error, etc. Same with ML!

Exams

- **Midterm exam**: Tuesday March 10th, 2026. 7 PM PST.
- **Final exam**: Thursday May 14th, 2026. 11:30 AM.
- Exams will be in-person, pencil and paper, and proctored by course staff. No electronic devices allowed.

Exams: extenuating circumstances

- **Alternate Midterm Exam**: if you can't make the primary midterm exam time, then there will be one alternate midterm time.
    - **Excused Midterm Exam**: if you can't make either the primary or the alternate midterm time, then your final exam score counts for your midterm.

- For both of the above: you must get explicit approval from the course instructor!

- **Final Exam**: all students must take the final exam to pass the course. There is no "Excused Final Exam" policy.
    - There will likely be one Alternate Final exam as well, but the time/location for this is not yet finalized. Stay tuned for details!

- To learn more, see the "Exams" section of the syllabus.

Grading

- Homework: 40%
- Midterm exam: 20%
- Final exam: 40%

No "exam clobber" policy.

Final grades are subject to a curve to meet department grading standards.

Participation

- Attendance for lecture/discussion will not be tracked, nor part of your grade
- **Extra credit**: Students that are particularly active in the course (eg Edstem, lecture, discussion) in a noteworthy, productive manner will be granted extra credit points.

Class lectures

Class lectures: 3:30 PM - 5:00 PM PST. Online, on Zoom

Please ask questions! Student-teacher interaction can be very fun and rewarding.

Lectures will be "live streamed" on Zoom. Video recordings for lectures will be uploaded to bCourses's "Media Gallery", generally within 24 hours after the end of lecture.

Slides for lectures will be posted to the course web page.

Discussion sections

Intent: go over course concepts in greater detail than covered in lecture, while having the benefit of small group learning headed by your discussion TA. This is a valuable opportunity to learn not only from your TA's, but also from your peers.

You may attend whichever, as many, and as few discussion sections you like. Attendance is not tracked, nor part of your grade.

Flexibility: you are welcome to attend a discussion section(s) other than the one you are officially enrolled in. However, if there are limited seats, please give up your seat to those that are officially enrolled in the discussion section.

Discussion section times/locations are here: [link]

Office hours

TA, Tutor office hours: in-person, all in Warren 101B (Area A or Area B).

Instructor office hours: online, on Zoom.

To view the OH schedule, see: https://data-188-berkeley.github.io/sp26/calendar/

Note: any changes to office hours will be announced on Edstem!

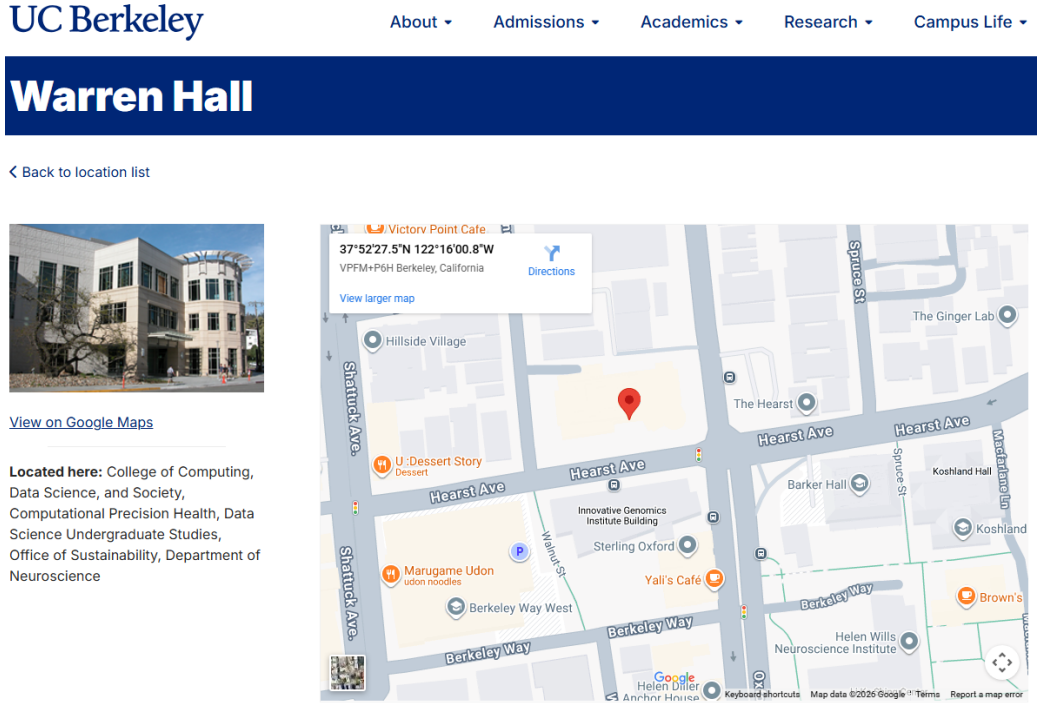

Programming homework assignments

The course will consist of several programming-based homework assignments, plus an additional Homework 0 meant as a review / test of your background

Homeworks are done *individually* (unless otherwise noted), see policies in a subsequent slide (and on the syllabus)

Homeworks are *largely* coding-based

Homeworks will be autograded using Gradescope's autograder feature. You will have unlimited submission attempts.

Late policy: no late submissions accepted. See syllabus for more details.

Collaboration policy

All submitted content (code and prose for homework assignments) should be your own content (or written by the group members, for group assignments)

However, you *may* (in fact are encouraged to) discuss the homework with others in the class and on the discussion forums

- This creates some room for undue copying, but please obey the reasonable person principle: discuss as you see fit, but don't simply share answers

For more info, see the "Collaboration" section of the syllabus.

Generative AI Policy

I'll define "AI tools" as: large language models like ChatGPT, and coding assistants like Cursor.

You **MAY USE** AI tools for purposes of **learning and understanding.**

You **MAY NOT USE** AI tools for **generating anything that you turn in**.

Litmus test: if it feels dishonest/unethical (eg you would be hesitant to tell the instructor about your AI usage), then please don't do it!

Violations of this AI policy will be treated as an Academic Integrity violation.

For more info, see "AI policy" in the syllabus.

Student well-being

University is a stressful environment.

In our experience, most academic integrity violations are the product of these environments and decisions made out of desperation

Please don't let it get to this point (or potentially much worse); contact the course staff ahead of time if you feel that issues are coming up that are interfering with your ability to participate fully in the course

Don't sacrifice quality of life for this course: make time to sleep, eat well, exercise, be with friends/family, socialize, etc

Expectations on Course Behavior

All participants in this course (students and course staff) must treat each other with respect and professionalism.

If you observe any violations of this, please feel free to contact course staff.

In the remaining time…

Log on to Ed https://edstem.org/us/courses/91872/discussion and say hello

Check out the course homepage https://data-188-berkeley.github.io/sp26/

Read the syllabus and policies: https://data-188-berkeley.github.io/sp26/syllabus/

Take a peek at HW0! https://github.com/data-188-berkeley/hw0

Tip: Upgrade to Google Colab Pro: [link]

Any questions?