CS 4644 / 7643: Deep Learning

Website: https://www.cc.gatech.edu/classes/AY2025/cs7643_fall/
Piazza: https://piazza.com/class/lzhdjcp6u255yf/
Gradescope: will be synced from canvas roster daily in the first week

Danfei Xu

School of Interactive Computing
Georgia Tech

Are you at the right place?

- This is CS 4644(DL) / CS 7643
  - “On campus” class
- This is NOT CS 7643-O01/OAN/Q/R
  - Online class for OMSCS program (Prof. Zsolt Kira)

Fall 24 Delivery Format

- In-person
  - Clough UG Learning Commons 152

- Streaming & Recording
  - We STRONGLY encourage you to attend the lectures in person.
  - Lectures will be streamed over zoom (link on Canvas).
  - Lectures are recorded and available for viewing

- **Remember: Content is free online.**
  - **You are here for the interactive experience.**

Outline for Today

- What is Deep Learning, the field, about?
- What about ChatGPT/foundation models/stable diffusion… ?
- What is this class about?
  - What to expect?
  - Logistics
- FAQ

Survey

Undergrad?

M.S.?

Ph.D.?

CS (CoC) / ECE?

Other Engineering?

Math / Natural Science?

Business?

Others?

Outline

- What is Deep Learning, the field, about?
- What is this class about?
  - What to expect?
  - Logistics
- FAQ

Concepts

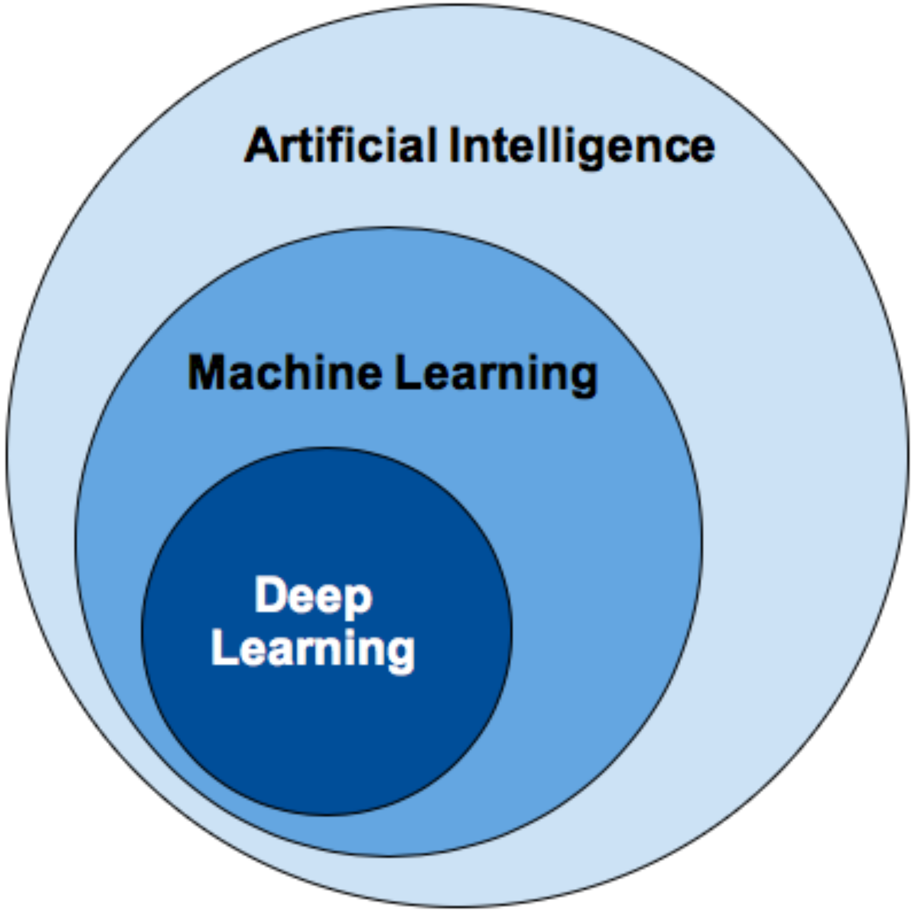

"Deep Learning is part of a broader family of **machine learning methods** based on **artificial neural networks**"

--- https://en.wikipedia.org/wiki/Deep_learning

Image Credit: https://www.sumologic.com/blog/machine-learning-deep-learning/

What is Machine Learning?

Enable machines to exhibit intelligent behaviors or improve their performances through data, without being explicitly programmed.

For example:

Linear regression is to find an optimal linear function (**model**) for a given dataset such that the difference between the predicted values and the actual values are minimized (**objective**).

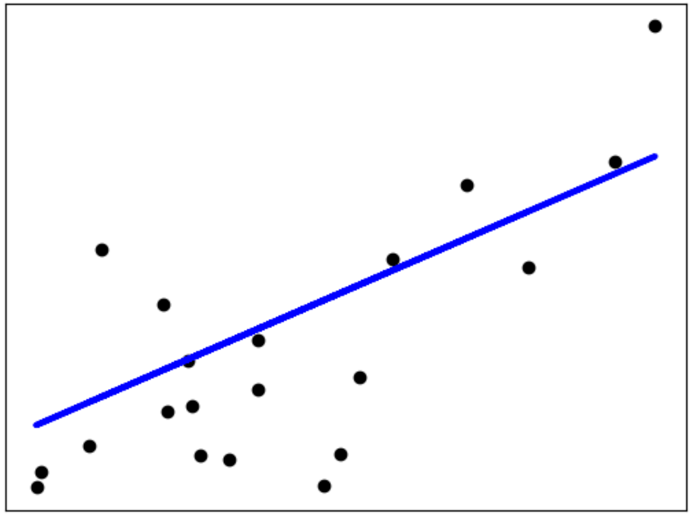

So what is Deep Learning?

- **Model**: (Deep) Artificial Neural Networks
- **Objective**: Representation Learning
  - Automatically discover useful features/representations for a **task** from **raw data**
- **Learning Method**: Supervised/Unsupervised/Reinforcement/Generative … Learning
- **System**: Software (PyTorch/TF/…) and hardware (GPU, cluster, …)
- **Simply**: Deep Learning

So what is Deep Learning

Ways to think about Deep Learning:

- Bottom-up: (Hierarchical) Compositionality
  - Cascade of non-linear transformations
  - Multiple layers of representations

- Top-down: End-to-End Learning
  - Learning (goal-driven) representations
  - Learning to feature extraction

Hierarchical Compositionality

VISION

pixels → edge → texton → motif → part → object

NLP

character → word → NP/VP/.. → clause → sentence → story

Composing Simple Functions to Build Complex Functions

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

Building A Complicated Function

Given a library of simple functions

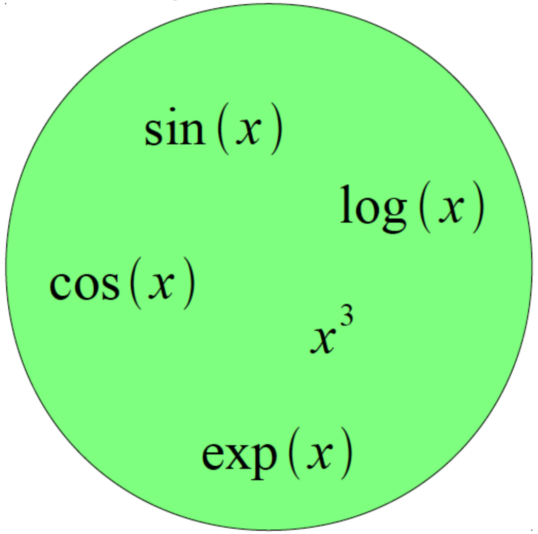

Compose into a

complex function

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

Building A Complicated Function

Given a library of simple functions

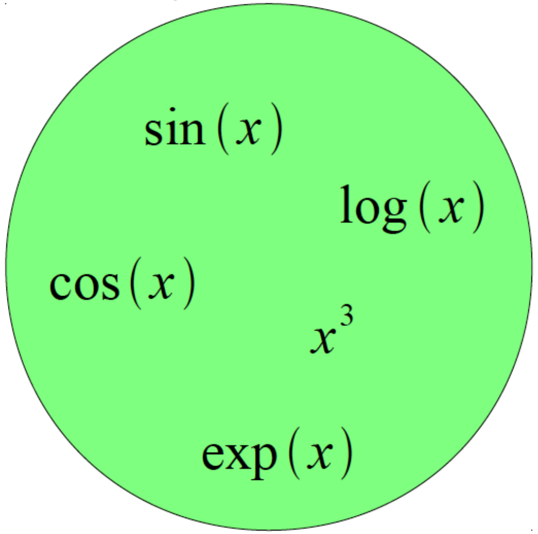

Compose into a

complex function

Idea 1: Linear Combinations

- Boosting
- Kernels
- …

$$f(x) = \sum_i \alpha_i g_i(x)$$

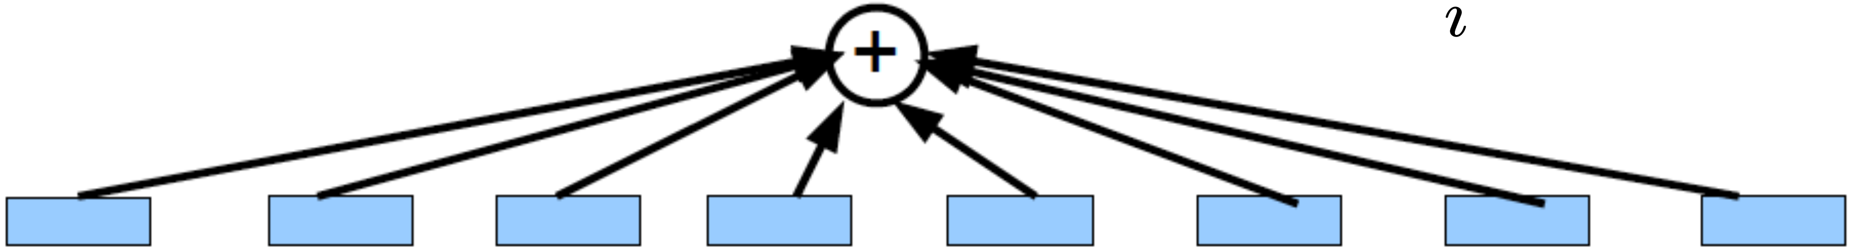

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

Building A Complicated Function

Given a library of simple functions

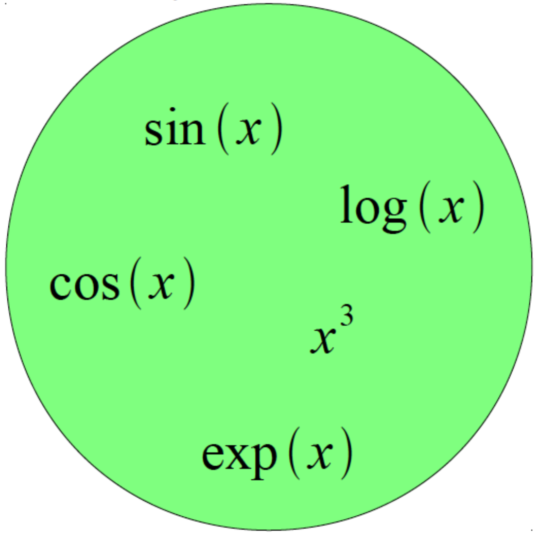

Compose into a

complex function

Idea 2: Function Composition

$h = g \cdot f$ such that $h(x) = g(f(x))$

Can we make it more expressive?

$$f(x) = \log(\cos(\exp(\sin^3(x))))$$

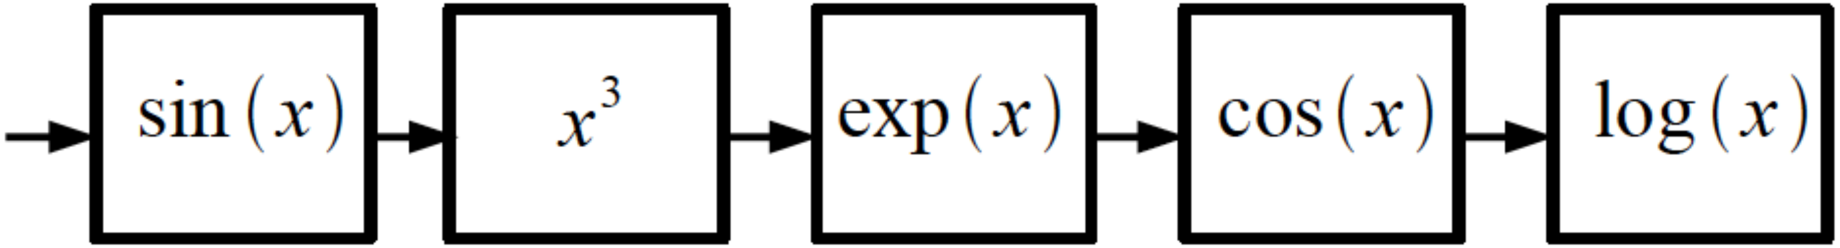

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

Building A Complicated Function

Given a library of simple functions

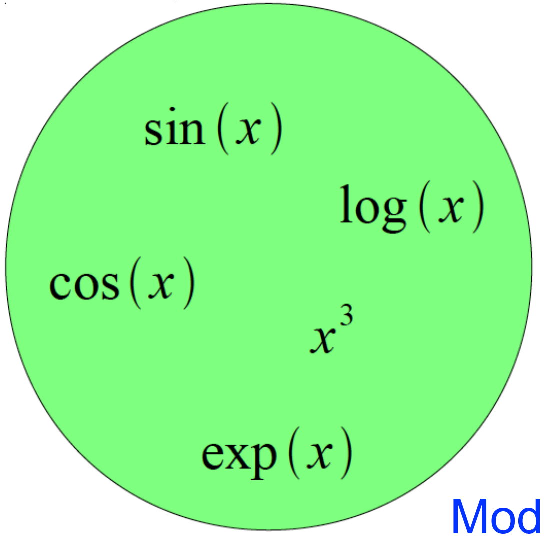

Compose into a

complex function

Idea 3: Layer Composition

Compose a set of layers **(parametric functions)** through which the input data get transformed.

More layers = "Deeper"

Modern DNNs have huge # of parameters, on the orders of **bn**'s

Parametric functions

$$f_\theta(x) = \boxed{g_{\theta_n}}(\ldots g_{\theta_2}(g_{\theta_1}(x)\ldots))$$

Linear Layer:

$$g(x) = Ax + b$$

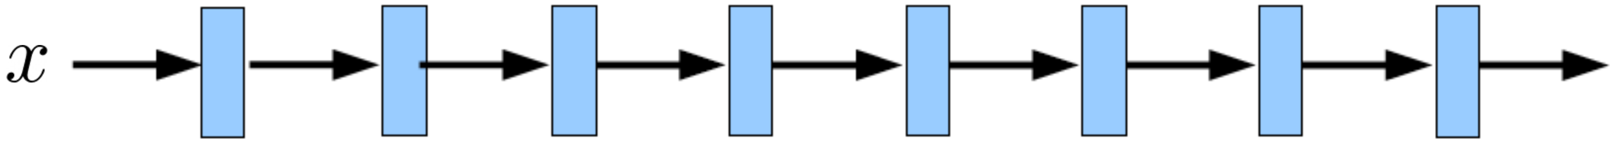

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

Deep Learning = Hierarchical Compositionality

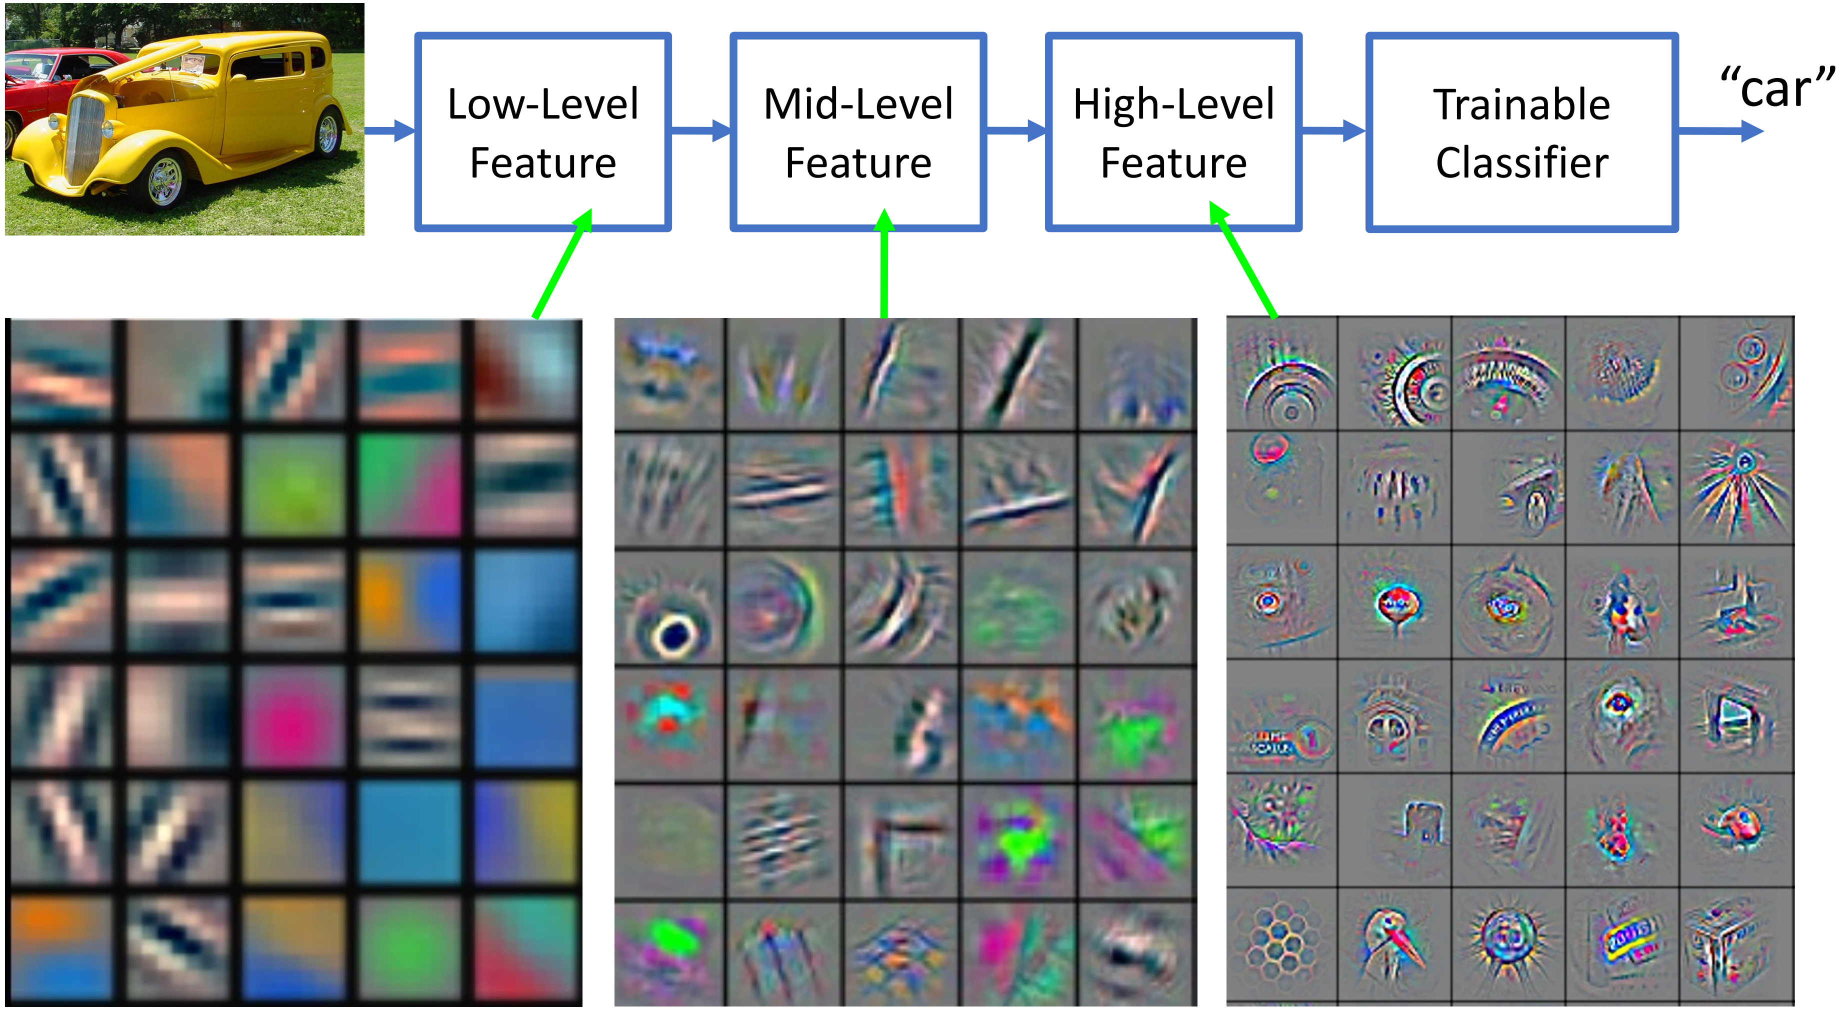

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

So what is Deep Learning

Ways to think about Deep Learning:

- Bottom-up: (Hierarchical) Compositionality
  - Cascade of non-linear transformations
  - Multiple layers of representations

- Top-down: End-to-End Learning
  - Learning (goal-driven) representations
  - Learning to feature extraction

Pre-deep learning: Feature Engineering

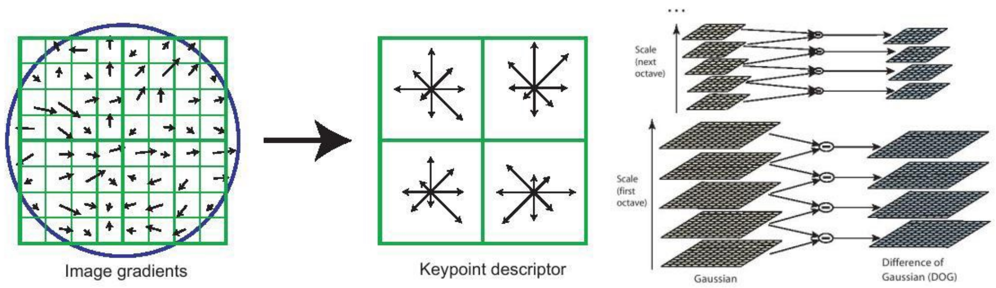

SIFT

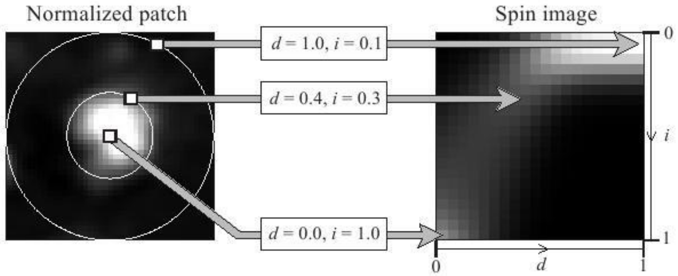

Spin Images

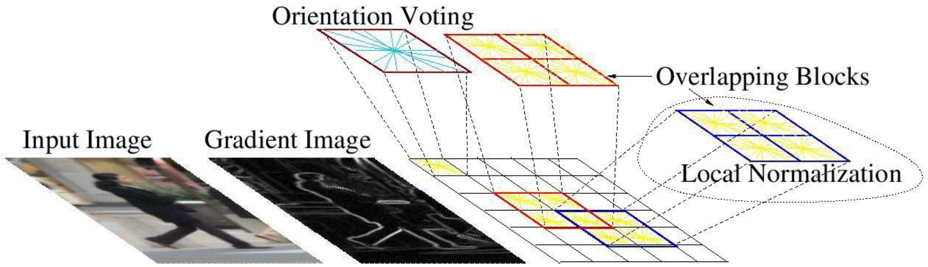

HoG

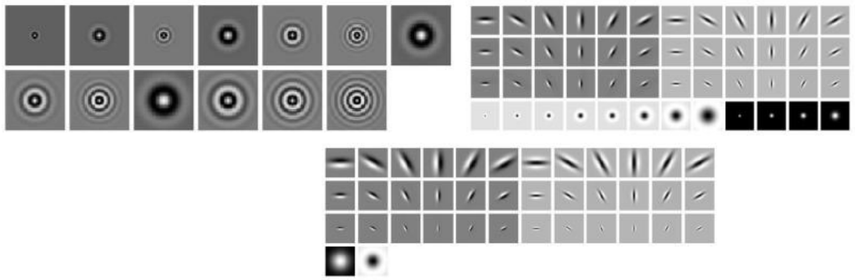

Textons

and many many more….

“Shallow” vs Deep Learning

- “Shallow” models

hand-crafted

Feature Extractor

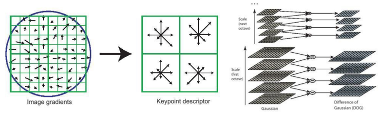

fixed

“Simple” Classifier

(linear regressor, SVM, …)

learned

“Car”

- Deep models

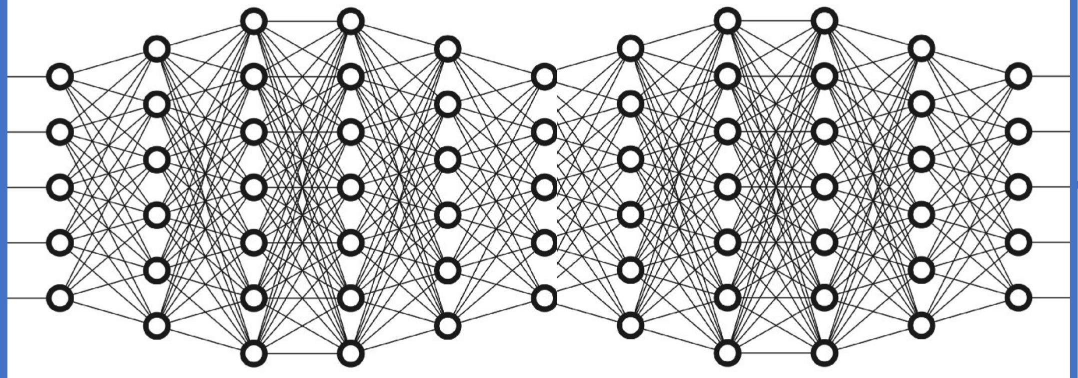

Learned

“Car”

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

“Shallow” vs Deep Learning

“Shallow” vs. deep language models

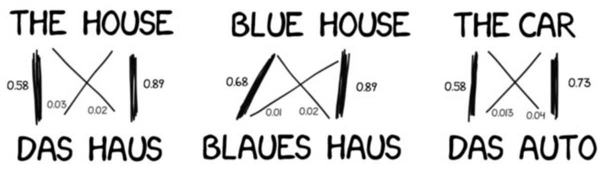

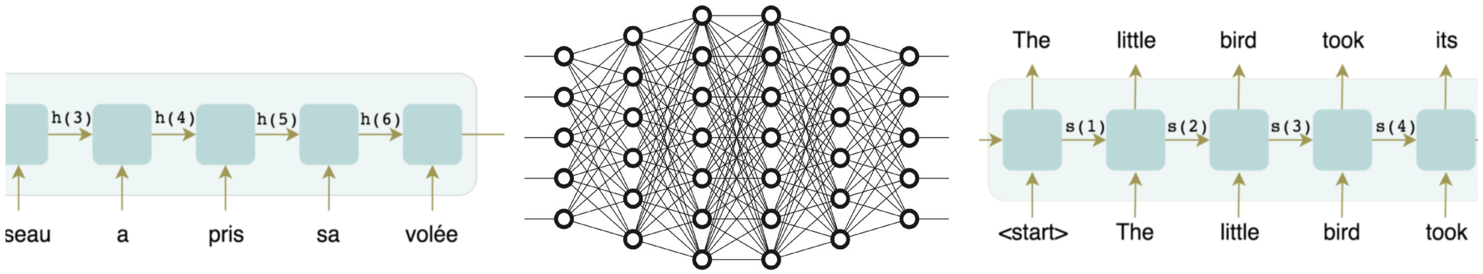

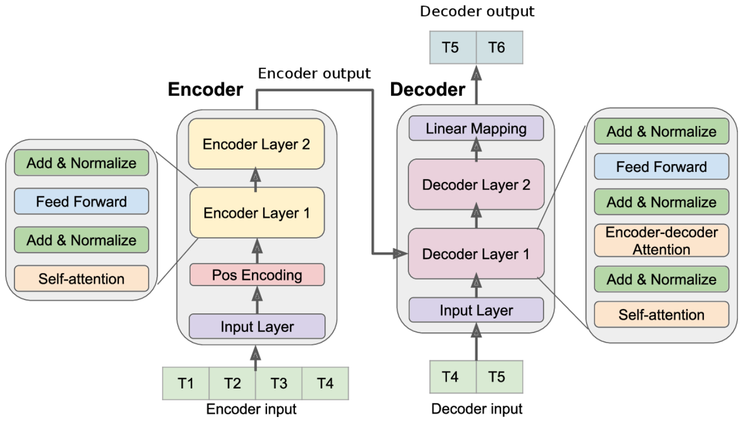

Transformer Models
(Vaswani *et al.*, 2017)

GPT-4

GPT4 large language model
(OpenAI 2023)

https://www.freecodecamp.org/news/a-history-of-machine-translation-from-the-cold-war-to-deep-learning-f1d335ce8b5/

“Pipelining” vs. “End-to-End Learning”

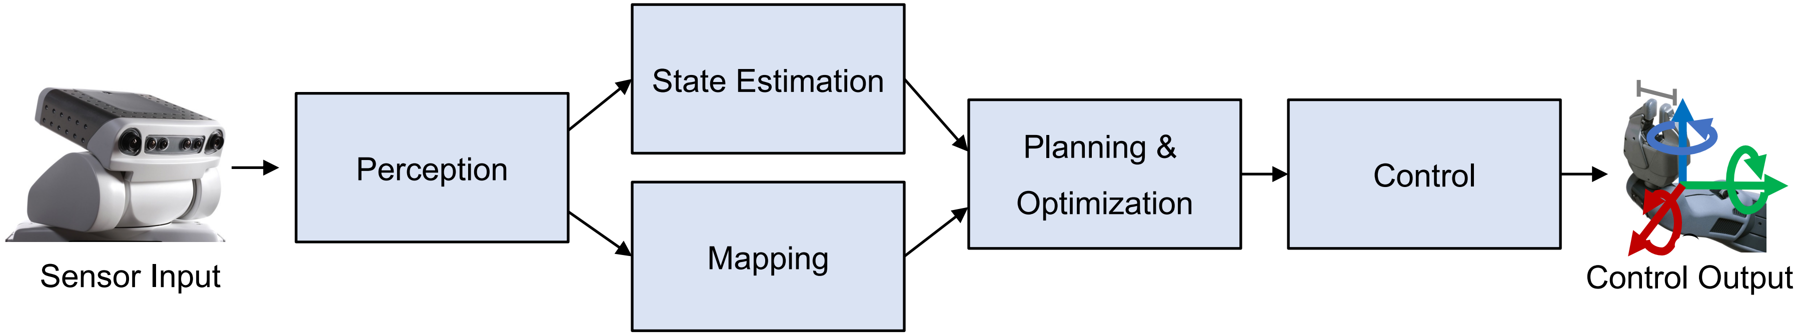

Hand-engineered pipelines

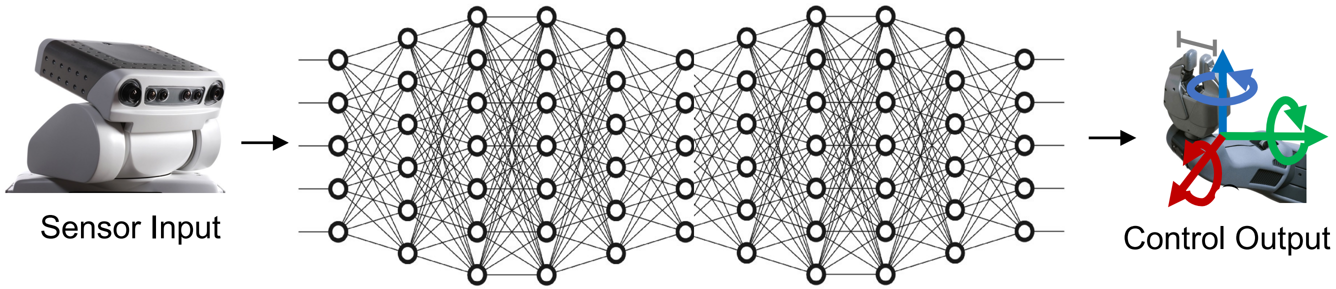

End-to-end learning
(“pixel-to-torque”)

So what is Deep Learning

Ways to think about Deep Learning:

- Bottom-up: (Hierarchical) Compositionality
  - Cascade of non-linear transformations
  - Multiple layers of representations

- Top-down: End-to-End Learning
  - Learning (goal-driven) representations
  - Learning to feature extraction

Benefit of Deep Learning

- (Usually) Better Performance
  - Caveats: given enough data, similar train-test distributions, non-adversarial evaluation, etc, etc.
- Apply to new domains without expertise knowledge
  - RGBD/Lidar
  - Language data
  - Gene-expression data
  - Complex controlling problem
  - Unclear how to hand-engineer
- New abilities emerge with more data and compute
- “Homogenization” of model design

“Expert” intuitions can be misleading

- *“Every time I fire a linguist, the performance of our speech recognition system goes up”*
  - Fred Jelinik, IBM ’98

- *“Because gradient descent is better than you”*
  - *Yann LeCun,* CVPR ‘13

“The Bitter Lesson”

“The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin. The ultimate reason for this is Moore's law, or rather its generalization of continued exponentially falling cost per unit of computation.” (Sutton, 2019)

http://www.incompleteideas.net/IncIdeas/BitterLesson.html

Emergence of new behaviors

Emergence means that the behavior of a system is implicitly induced rather than explicitly constructed. For Deep Learning, emergence is often induced by larger model & more data.

Example: Compared to GPT-2’s 1.5B parameter parameter model, GPT-3’s 175-billion model permits “prompting”, i.e., adapting to a new task simply by describing task.

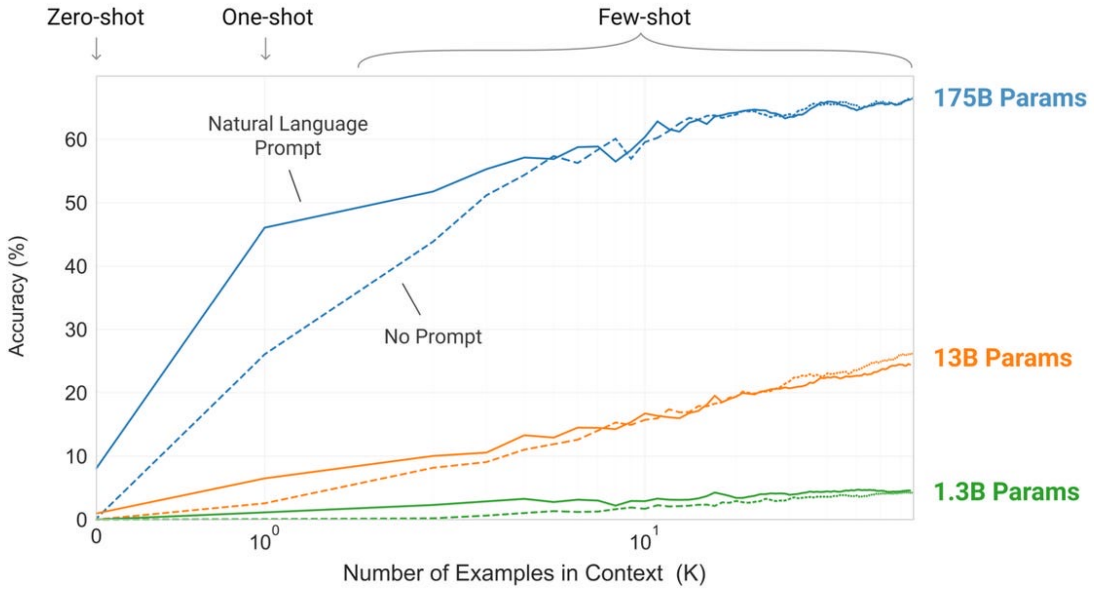

On the Opportunities and Risks of Foundation Models (Bommasani et al., 2021)

Homogenization of Deep Learning

Homogenization is the **consolidation** of methodologies for building machine learning systems across a wide range of applications.

Example: The Transformer Models (Vaswani *et al.,* 2017)

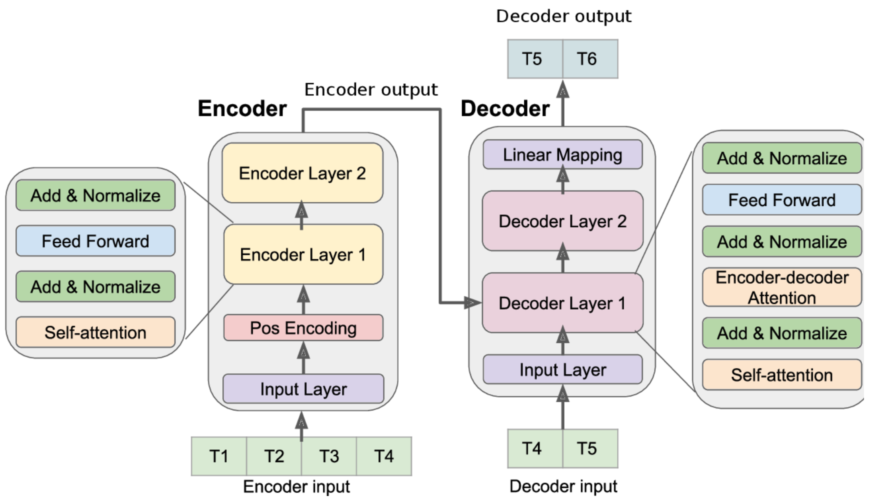

Transformer Models originally designed for NLP

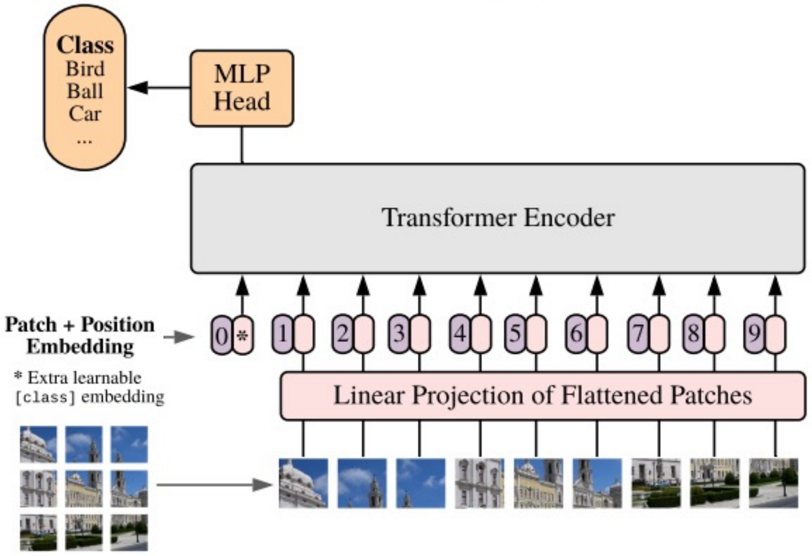

Almost identical model (Visual Transformers) can be applied to Computer Vision tasks

On the Opportunities and Risks of Foundation Models (Bommasani et al., 2021)

What about ChatGPT / foundation models / … buzzwords?

Bing's A.I. Chat: 'I Want to Be Alive. 😈'

In a two-hour conversation with our columnist, Microsoft's new chatbot said it would like to be human, had a desire to be destructive and was in love with the person it was chatting with. Here's the transcript.

Give this article 1.6K

https://www.nytimes.com/article/ai-artificial-intelligence-chatbot.html

ARTIFICIAL INTELLIGENCE

ChatGPT is about to revolutionize the economy. We need to decide what that looks like.

New large language models will transform many jobs. Whether they will lead to widespread prosperity or not is up to us.

By David Rotman March 25, 2023

https://www.technologyreview.com/2023/03/25/1070275/chatgpt-revolutionize-economy-decide-what-looks-like/

Exam	GPT-4	GPT-4 (no vision)	GPT-3.5
Uniform Bar Exam (MBE+MEE+MPT)	298 / 400 (~90th)	298 / 400 (~90th)	213 / 400 (~10th)
LSAT	163 (~88th)	161 (~83rd)	149 (~40th)
SAT Evidence-Based Reading & Writing	710 / 800 (~93rd)	710 / 800 (~93rd)	670 / 800 (~87th)
SAT Math	700 / 800 (~89th)	690 / 800 (~89th)	590 / 800 (~70th)
Graduate Record Examination (GRE) Quantitative	163 / 170 (~80th)	157 / 170 (~62nd)	147 / 170 (~25th)
Graduate Record Examination (GRE) Verbal	169 / 170 (~99th)	165 / 170 (~96th)	154 / 170 (~63rd)
Graduate Record Examination (GRE) Writing	4 / 6 (~54th)	4 / 6 (~54th)	4 / 6 (~54th)
USABO Semifinal Exam 2020	87 / 150 (99th - 100th)	87 / 150 (99th - 100th)	43 / 150 (31st - 33rd)
USNCO Local Section Exam 2022	36 / 60	38 / 60	24 / 60
Medical Knowledge Self-Assessment Program	75 %	75 %	53 %
Codeforces Rating	392 (below 5th)	392 (below 5th)	260 (below 5th)
AP Art History	5 (86th - 100th)	5 (86th - 100th)	5 (86th - 100th)
AP Biology	5 (85th - 100th)	5 (85th - 100th)	4 (62nd - 85th)
AP Calculus BC	4 (43rd - 59th)	4 (43rd - 59th)	1 (0th - 7th)
AP Chemistry	4 (71st - 88th)	4 (71st - 88th)	2 (22nd - 46th)
AP English Language and Composition	2 (14th - 44th)	2 (14th - 44th)	2 (14th - 44th)
AP English Literature and Composition	2 (8th - 22nd)	2 (8th - 22nd)	2 (8th - 22nd)
AP Environmental Science	5 (91st - 100th)	5 (91st - 100th)	5 (91st - 100th)
AP Macroeconomics	5 (84th - 100th)	5 (84th - 100th)	2 (33rd - 48th)
AP Microeconomics	5 (82nd - 100th)	4 (60th - 82nd)	4 (60th - 82nd)
AP Physics 2	4 (66th - 84th)	4 (66th - 84th)	3 (30th - 66th)
AP Psychology	5 (83rd - 100th)	5 (83rd - 100th)	5 (83rd - 100th)
AP Statistics	5 (85th - 100th)	5 (85th - 100th)	3 (40th - 63rd)
AP US Government	5 (88th - 100th)	5 (88th - 100th)	4 (77th - 88th)
AP US History	5 (89th - 100th)	4 (74th - 89th)	4 (74th - 89th)
AP World History	4 (65th - 87th)	4 (65th - 87th)	4 (65th - 87th)
AMC 10[3]	30 / 150 (6th - 12th)	36 / 150 (10th - 19th)	36 / 150 (10th - 19th)
AMC 12[3]	60 / 150 (45th - 66th)	48 / 150 (19th - 40th)	30 / 150 (4th - 8th)
Introductory Sommelier (theory knowledge)	92 %	92 %	80 %
Certified Sommelier (theory knowledge)	86 %	86 %	58 %
Advanced Sommelier (theory knowledge)	77 %	77 %	46 %
Leetcode (easy)	31 / 41	31 / 41	12 / 41
Leetcode (medium)	21 / 80	21 / 80	8 / 80
Leetcode (hard)	3 / 45	3 / 45	0 / 45

Table 1. GPT performance on academic and professional exams. In each case, we simulate the conditions and scoring of the real exam. We report GPT-4's final score graded according to exam-specific rubrics, as well as the percentile of test-takers achieving GPT-4's score.

GPT4 technical report, OpenAI, March 2023

What about ChatGPT / foundation models / … buzzwords?

A *grossly simplified* timeline of deep learning

This course

… Pioneers	Foundations	Explosion	The Era of Foundation Models …
1940-1960s	1970-2000s	2010-2019	2020-???

Pioneers (1940-1960s)

- Artificial Neural Networks
- Hebbian theory (Hebb)
- Perceptron (Rosenblatt)
- Visual cortex cells (Hubel and Wiesel)
- …

Foundations (1970-2000s)

- ConvNets (Fukushima)
- Backpropagation in NNs (Remelhart, Williams, Hinton, LeCun, …)
- LSTM (Hochreiter and Schmidhuber)
- GPUs
- …

Explosion (2010-2019)

- ImageNet (Jia, Fei-Fei, …)
- AlexNet (Krizhevsky)
- The Cat Experiment (Google)
- Transformer
- Deep Generative Models
- AlphaGo
- Deep Reinforcement Learning
- More GPUs
- …

The Era of Foundation Models (2020-???)

- GPT Series (OpenAI)
- CLIP & DALL·E (OpenAI)
- SegmentAnything (Meta)
- Llama Series (Meta)
- Bard (Google)
- StableDiffusion
- …

Problems with Deep Learning

- Problem#1: Lack of a formal understanding
  - Non-Convex! Non-Convex! Non-Convex!
    - Depth>=3: most losses non-convex in parameters
  - Worse still, existing intuitions from classical statistical learning theory don't seem to carry over.
  - Theoretically, we are stumbling in the dark here

- Standard response #1
  - "Yes, but this just means there's new theory to be constructed"
  - "All interesting learning problems are non-convex"

- Standard response #2
  - "Yes, but it often works!"

Problems with Deep Learning

- Problem#2: Lack of interpretability
  - Hard to track down what’s failing
  - Pipeline systems have expected performances at each step
  - In end-to-end systems, it’s hard to know why things are not working

Problems with Deep Learning

- Problem#2: Lack of interpretability

Why did the robot do that?

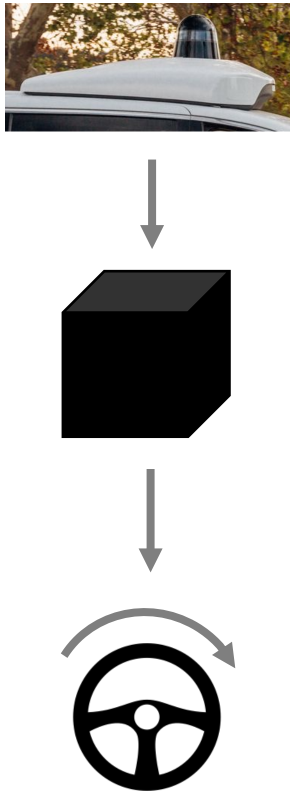

Problems with Deep Learning

- Problem#2: Lack of interpretability
  - Hard to track down what's failing
  - Pipeline systems have expected performances at each step
  - In end-to-end systems, it's hard to triage an error

- Standard response #1
  - Tricks of the trade: visualize features, add losses at different layers, pre-train to avoid degenerate initializations…
  - "MOOOORE DATA!"
  - "We're working on it"

- Standard response #2
  - "Yes, but it often works!"

Problems with Deep Learning

- Problem#3: Lack of easy reproducibility
  - Direct consequence of stochasticity & non-convexity
    - different initializations → different local minima
  - Other stochasticity in the training pipeline: parallel data loading, distributed training, numerical precision on GPU…

- Standard response #1
  - It’s getting much better
  - Standard toolkits/libraries/frameworks now available
  - PyTorch, TensorFlow, MxNet…

- Standard response #2
  - “Yes, but it often works!”

Yes it works, but how?

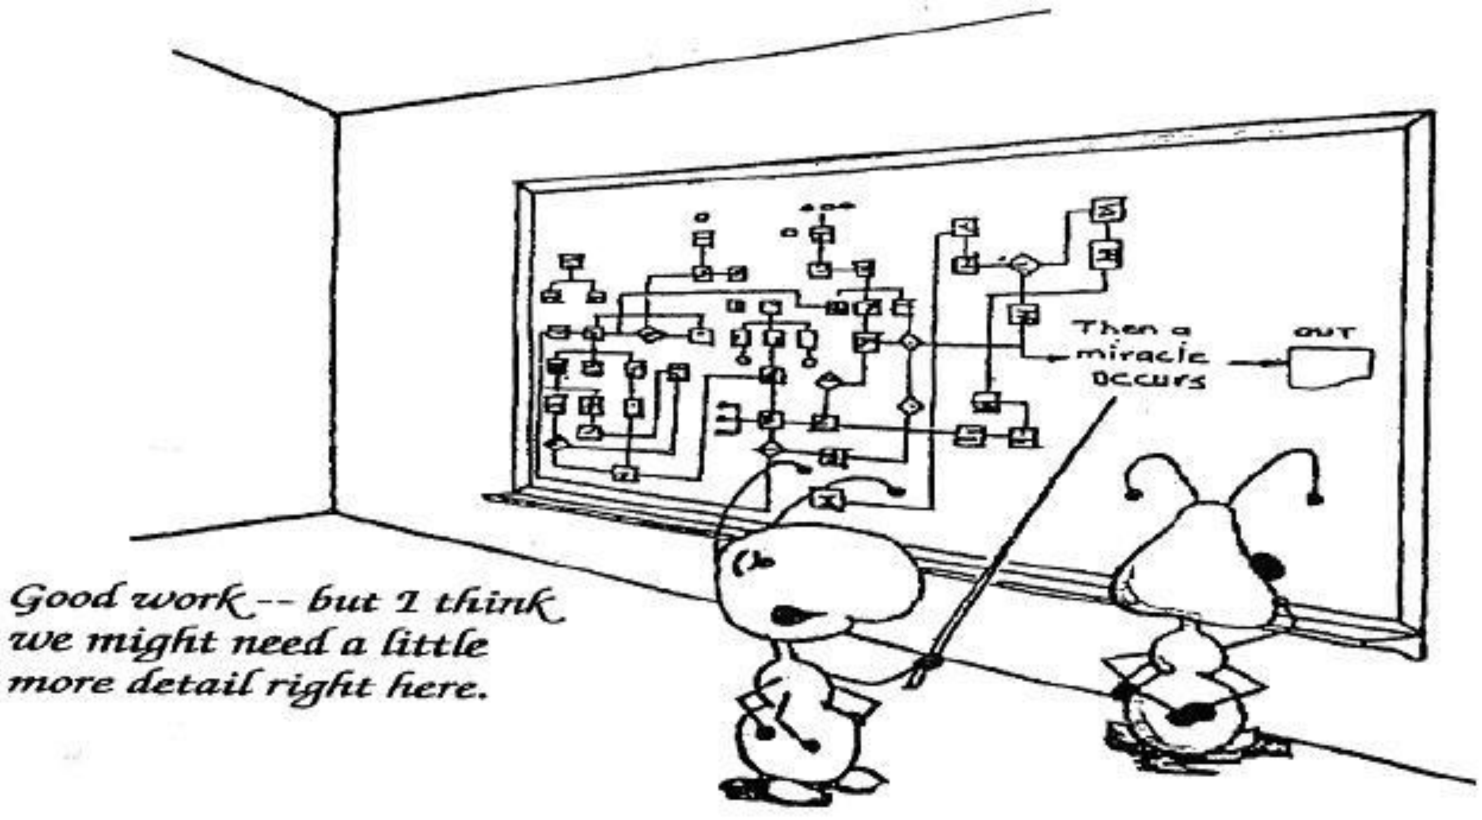

Outline

- What is Deep Learning, the field, about?
    - Highlight of some recent projects from my lab

- What is this class about?
    - What to expect?
    - Logistics

- FAQ

Outline

- What is Deep Learning, the field, about?
  - Highlight of some recent projects from my lab
- What is this class about?
  - What to expect?
  - Logistics
- FAQ

What is this class about?

- Introduction to Deep Learning

- Goal:
    - After finishing this class, you should be ready to get started on your first DL research / engineering project.
        - Backpropagation, optimization
        - Convolutional Neural Networks (image data)
        - Recurrent Neural Networks / Transformers (sequence data)
        - Generative Models (unsupervised learning)
        - Deep Reinforcement Learning (decision making)
        - (Glimpses of) cutting-edge research in CV, NLP, Robotics
    - Work on fun projects with your peers!

- Target Audience:
    - Senior undergrads, MS-(CS, ML, …), and new PhD students

What this class is NOT

- NOT the target audience:
  - Students without sufficient background knowledge (Python, linear algebra, calculus, basic probability & statistics)
  - Advanced grad-students already working in ML/DL areas
  - People looking for an in-depth understanding of a research area that uses deep learning (3D Vision, Large Language Models, Deep RL, etc.).
- NOT the goal:
  - Intro to Machine Learning / Optimization

Caveat

- This is an ADVANCED Machine Learning class
  - This should NOT be your first introduction to AI/ML
  - You will need a formal class; not just self-reading/coursera

Prerequisites

- Python Programming
  - Basic knowledge of numerical computations & tools (e.g., numpy)
  - You will write a lot of code!
- Intro Machine Learning
  - Classifiers, regressors, loss functions, MLE, MAP
- Linear Algebra
  - Matrix multiplication, eigenvalues, positive semi-definiteness…
- Calculus
  - Multi-variate gradients, hessians, jacobians…
- Must read (on W3 reading list): Matrix calculus for deep learning
  - https://explained.ai/matrix-calculus/index.html

Your Teaching Team

- Instructor: Prof. Danfei Xu
- Ph.D. (2021), Stanford University
- 2022-Now: Assistant Professor at Georgia Tech
- Research in Robotics, Machine Learning, Computer Vision
- Office in Klaus 1314
- Running, cycling, cooking

Your Teaching Team

Head TA: Wei Zhou
(wzhou322@)

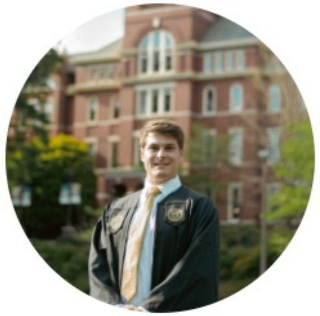

Zachary Breitbart
(zacharyb@)

Srikar S Balusu
(sbalusu30@)

David He (davidhe@)

Haotian Xue
(hxue45@)

Woo Chul Shin
(wshin49@)

Huaijin(Tony) Tu
(htu35@)

Aryan Sarswat
(asarswat8@)

Elias Cho (echo94@)

Rodrigo Loza
(rloza3@)

Office Hour

TA Office Hours:

- Virtual over zoom
- Check course website for OH slots and zoom links
- Start next week

Danfei’s Office Hours:

- In-person (Klaus 1314) or zoom
- No assignment (PS/HW) questions
- Lecture content / project ideas / administrative / career advice, …

Main channel of communication: Piazza

- HW and Project Assignments will be announced on Piazza
- Q&A: Check other questions before posting new ones

Organization & Deliverables

- 4 problem-sets+homeworks (64%)
  - Mix of theory (PS) and implementation (HW)
  - First one goes out next week
    - Start early, Start early, Start early, Start early, Start early, Start early
- Course project (36%)
  - Projects done in groups of 2-4
  - You need a good reason to do a solo project.
  - Proposal (1%), Milestone Report (10%), Final Report (20%), Poster Session (5%)
  - **Find a team ASAP! Talk to people, use Piazza "find a teammate" post.**
  - Ideas & scope: http://cs231n.stanford.edu/reports.html
- (Bonus) Class Participation (1%)
  - Top (endorsed) contributors on Piazza

Plenty of “buffer” built in

- Grace period
  - 2 days grace period
    - Intended for *checking* submission NOT to replace due date
    - No need to ask for grace, no penalty for turning it in within grace period
    - Can NOT use for PS0
  - After grace period, you get a 0 (no excuses except medical)
    - Send all medical requests to dean of students (https://studentlife.gatech.edu/)
    - Form: https://gatech-advocate.symplicity.com/care_report/index.php/pid224342?
  - **DO NOT SEND US ANY MEDICAL INFORMATION!** We do not need any details, just a confirmation from dean of students

PS0

- Out already. Due Sunday Aug 25th, 11:59pm
  - Will be available on class webpage
  - If not registered yet (on waitlist), see webpage FAQ for form to request gradescope access

- Grading
  - Not counted towards your final grade, but required
  - If it takes you more than 3 hours to complete, you might struggle in the course.

- Topics
  - PS: probability, calculus
  - HW: Numpy, calculus

Class Project

- Goal
  - Chance to try Deep Learning in practice
  - Encouraged to apply to your research (computer vision, NLP, robotics, compbio,…)
  - Must be done this semester.
  - Can combine with other classes, but **separate thrust**
    - get permission from both instructors; delineate contribution to each course
  - 2-4 members (outside of this requires approval)
  - Will have a separate lecture on this in Week 3

Computing

- Major bottleneck
  - GPUs

- Options
  - Your own / group / advisor’s resources

  - Google Colab
    - jupyter-notebook + free GPU instance

  - Google Cloud credits (details TBA)
    - Tutorial on setting up gloud: https://github.com/cs231n/gcloud

4644 vs 7643

- Level differentiation

- Separate grade curves calculation
  - As a result, 4644 and 7643 may have different letter grade cut-offs.

Outline

- What is Deep Learning, the field, about?
  - Highlight of some recent projects from my lab

- What is this class about?
  - What to expect?
  - Logistics

- FAQ

Waitlist / Audit / Sit in

- Waitlist
  - Class is full.
  - Do PS0 **NOW**. Come to first few classes.
  - Hope people drop. Big churn in the first week.

- "I need this class to graduate"
  - Talk to your degree program advisor. They control the process of making sure you have options to graduate on time.

- Audit or Pass/Fail
  - No.

What is the re-grading policy?

- Homework assignments
  - **Within 1 week** of receiving grades: submit regrade request on GradeScope

What is the collaboration policy?

- Collaboration
  - Only on HW (coding) and project.
  - You may discuss the questions
  - Each student writes their own answers
  - Write on your homework anyone with whom you collaborate
  - Each student must write their own code for the programming part

- Zero tolerance on plagiarism
  - Neither ethical nor in your best interest
  - Always credit your sources
  - Don’t cheat. We will find out.

How do I get in touch?

- Primary means of communication -- Piazza
  - No direct emails to Instructor **unless private information**
  - Instructor/TAs can provide answers to everyone on forum
  - Class participation credit for answering questions!
  - No posting answers. We will monitor.
  - Stay respectful and professional.

Share your feedback

Ways to share your feedback:

- Come talk to us
- Email
- Private Piazza Post
- **Anonymous feedback form (link on Piazza)**

Questions?