Deep Learning and Applications

Jian Tang
HEC Montreal
Mila-Quebec AI Institute
CIFAR AI Research Chair
Email: jian.tang@hec.ca

Course Information

- Website:
  - https://deepgraphlearning.github.io/Math80648A_2025A/index.html
- Instructor: **Jian Tang**
  - Associate Professor, Department of Decision Science, HEC Montreal
  - Homepage: www.jian-tang.com
  - Email: jian.tang@hec.ca
- Teaching Assistants: Xinyu Yuan, Zhihao Zhan
  - Ph.D at Mila and University of Montreal
  - Email: xinyu.yuan@mila.quebec
- Discord: https://discord.gg/NBgrmeYy
  - Join discord for discussion

"Artificial Intelligence is the most revolutionary technology in decades, on par with computers, cellphones and the internet."

Artificial Intelligence: the Fourth Industrial Revolution

- Artificial Intelligence
  - "the term is often used to describe machines (or computers) that mimic "cognitive" functions that humans associate with the human mind, such as "learning" and "problem solving"." -- Wikipedia

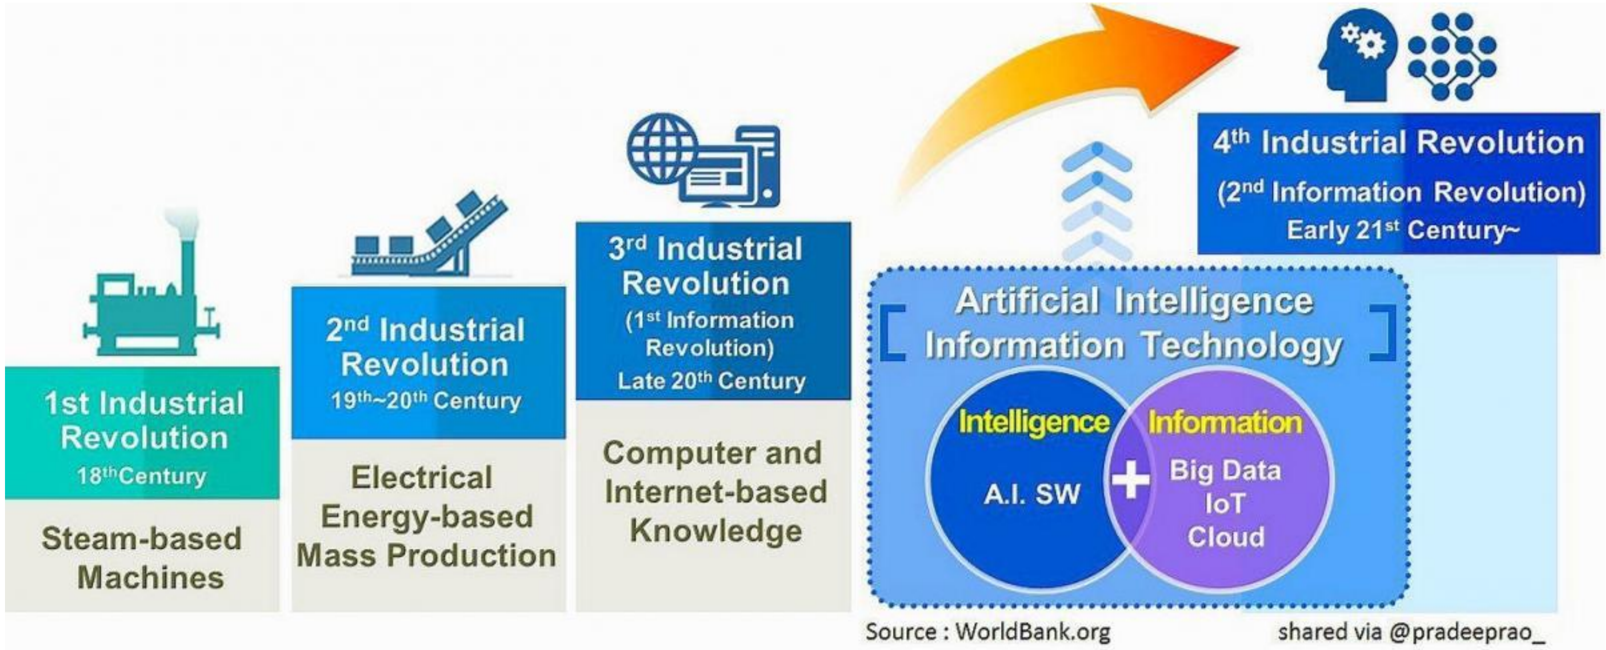

-image from Internet

Artificial Intelligence v.s. Machine Learning v.s. Deep Learning

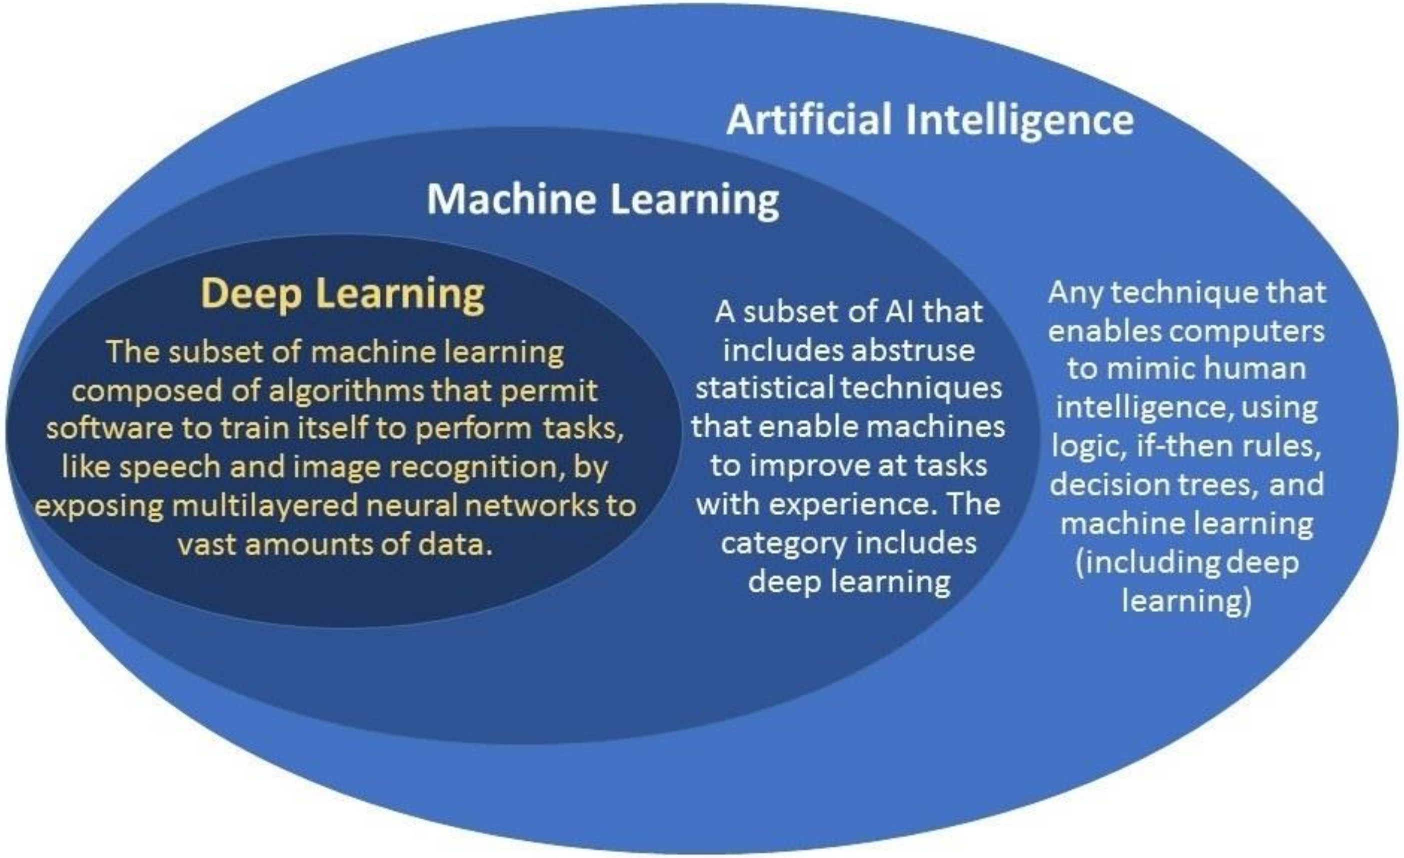

Machine Learning

- "**Machine learning** is a field of computer science that uses statistical techniques to give computer systems the ability to "learn" (i.e., progressively improve performance on a specific task) with data, without being explicitly programmed."

-Wikipedia

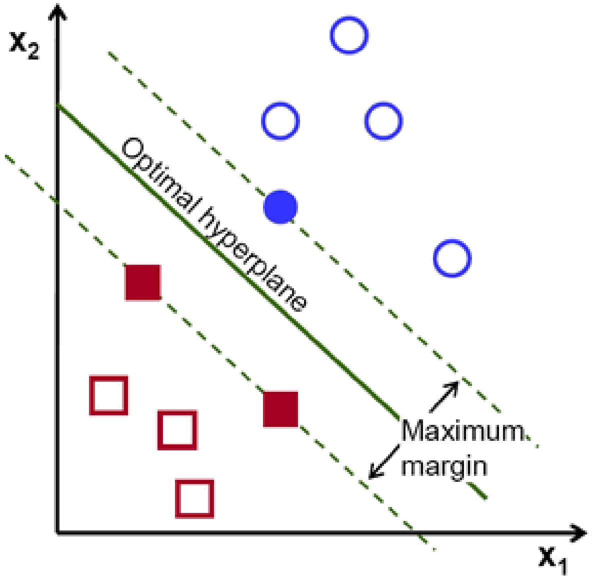

Support vector machines

→ **Hand-crafted** Feature Extractor → Simple Trainable Classifier e.g., SVM, LR

Domain experts

Deep Learning = Feature Representation Learning

- Algorithms that allow to learn features from data (a.k.a, End-to-end learning)

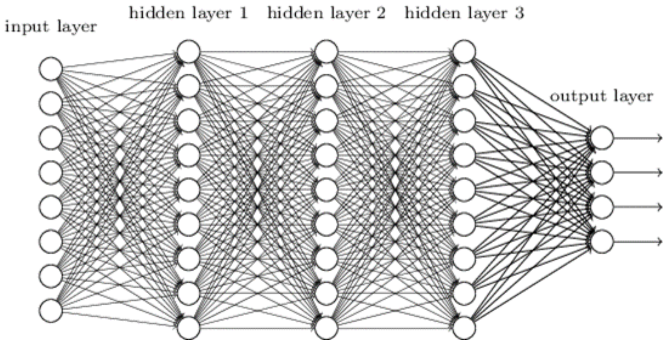

Deep Neural Networks

Trainable
Feature Extractor

Simple Trainable Classifier
e.g., SVM, LR

Domain experts

Why Deep Learning Now?

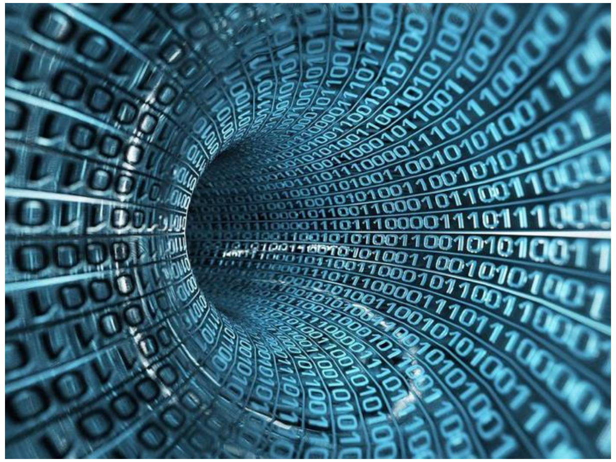

Big Data

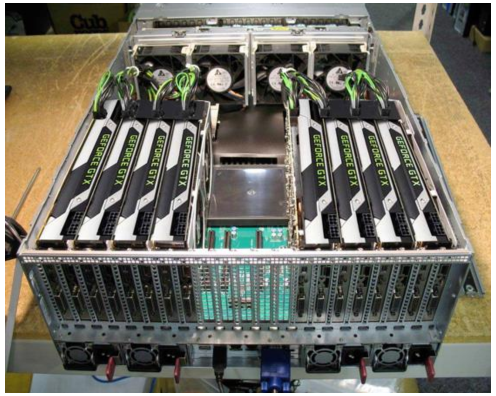

Big Computation

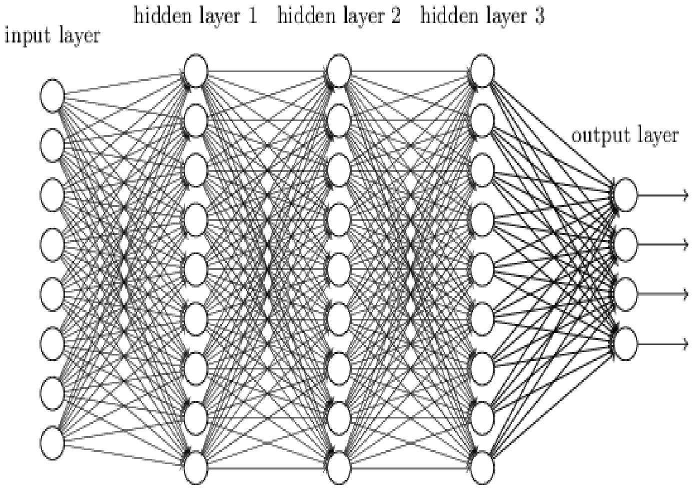

Big Model

AlphaGo/AlphaZero

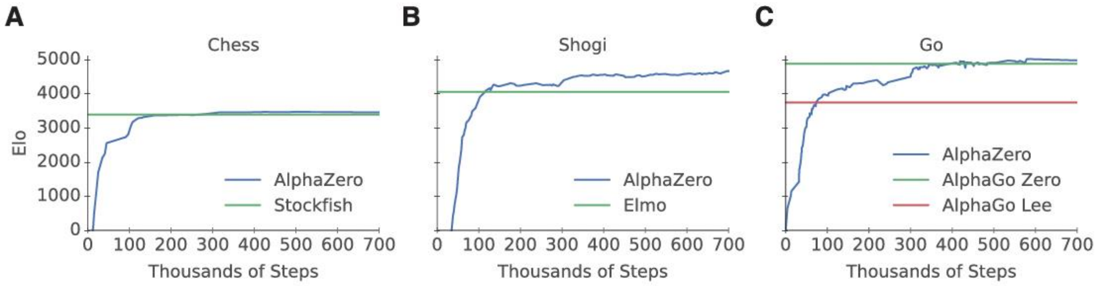

Figure 1: **Training AlphaZero for 700,000 steps.** Elo ratings were computed from games between different players where each player was given one second per move. **(A)** Performance of AlphaZero in chess, compared with the 2016 TCEC world-champion program Stockfish. **(B)** Performance of AlphaZero in shogi, compared with the 2017 CSA world-champion program Elmo. **(C)** Performance of AlphaZero in Go, compared with AlphaGo Lee and AlphaGo Zero (20 blocks over 3 days).

Speech Recognition

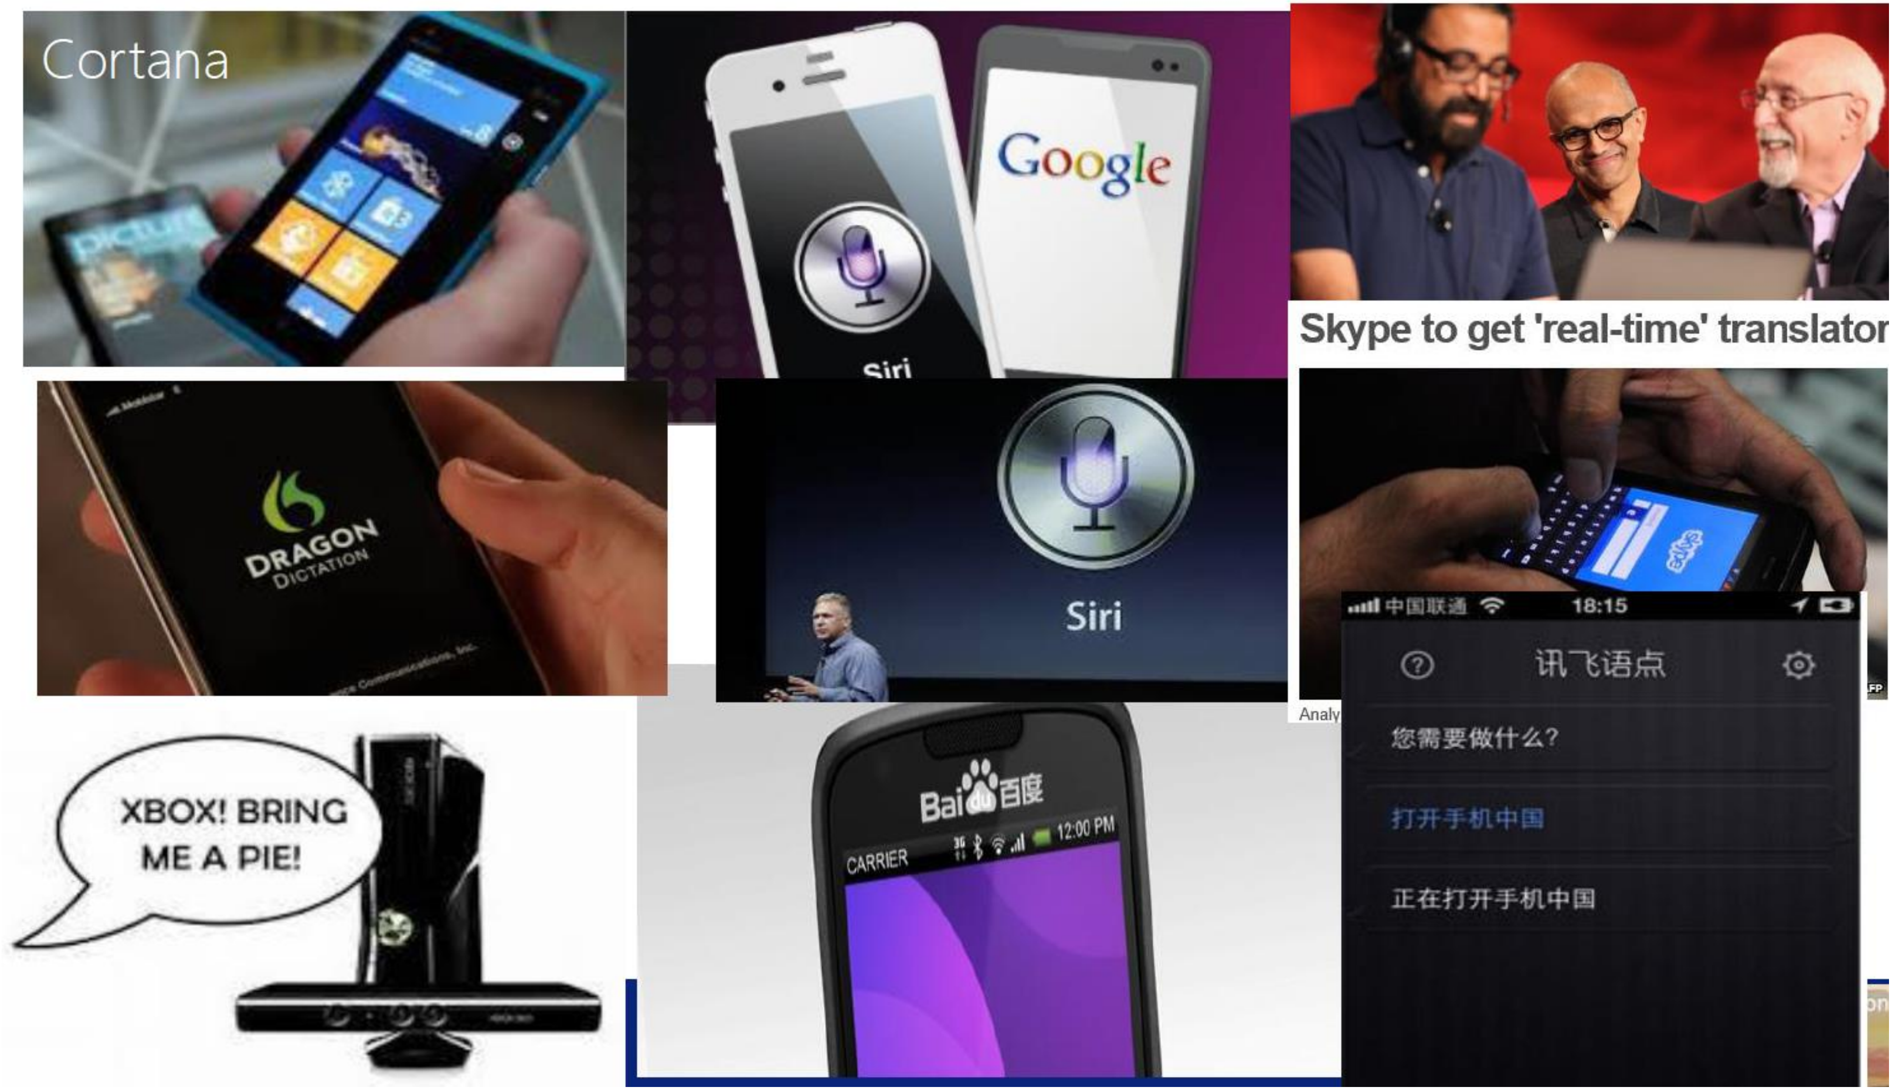

Speech Recognition Results

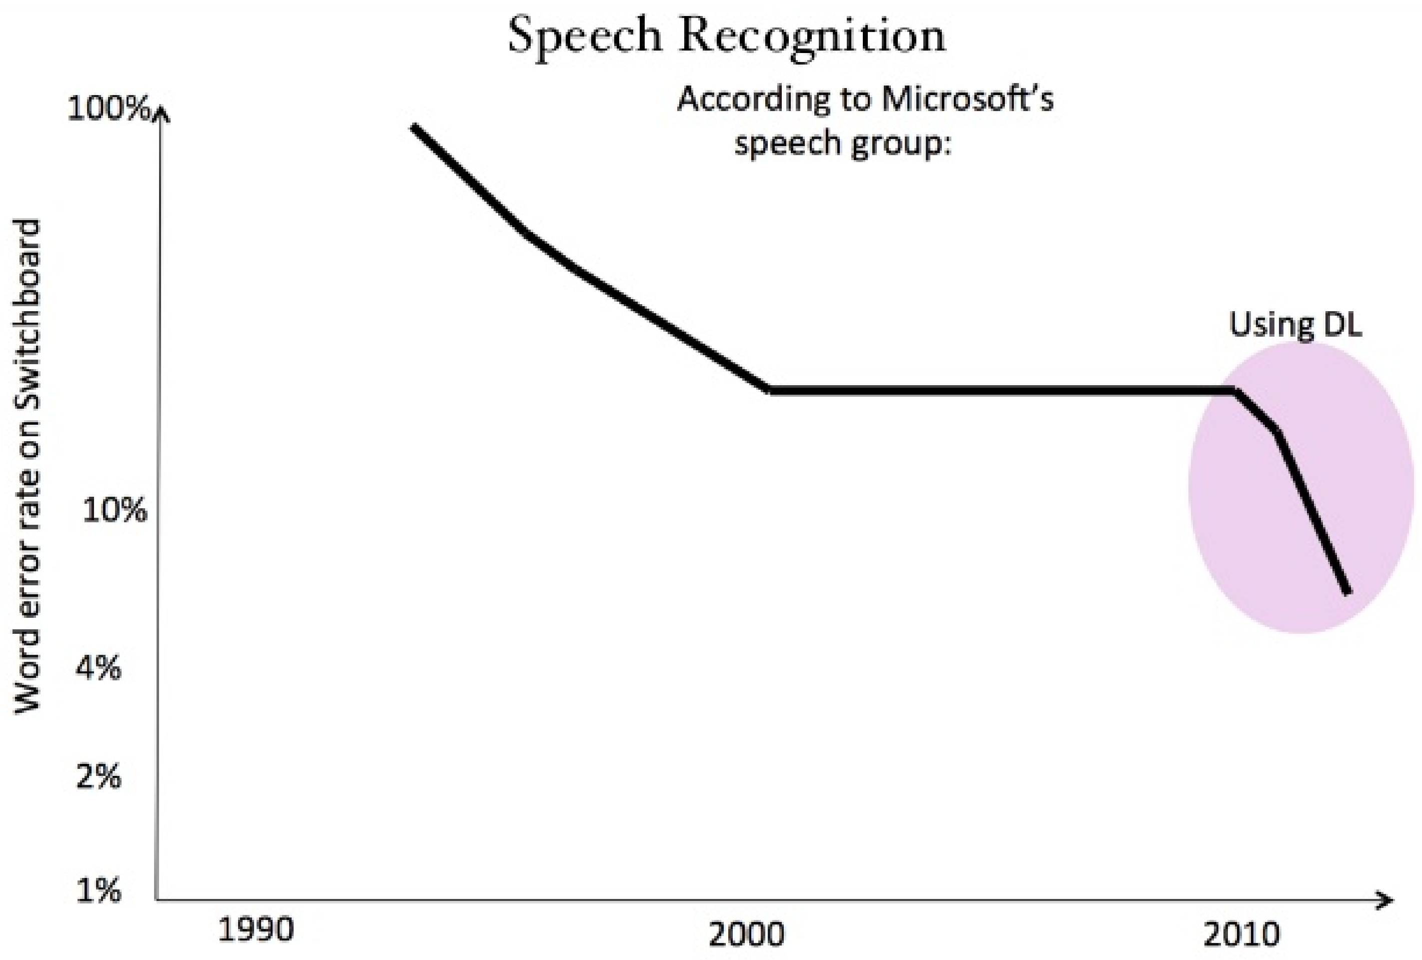

(Figure from Microsoft's speech Group)

Image Recognition

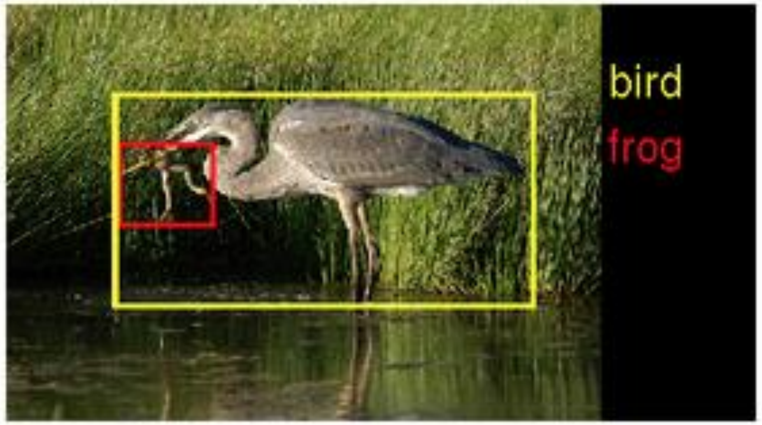

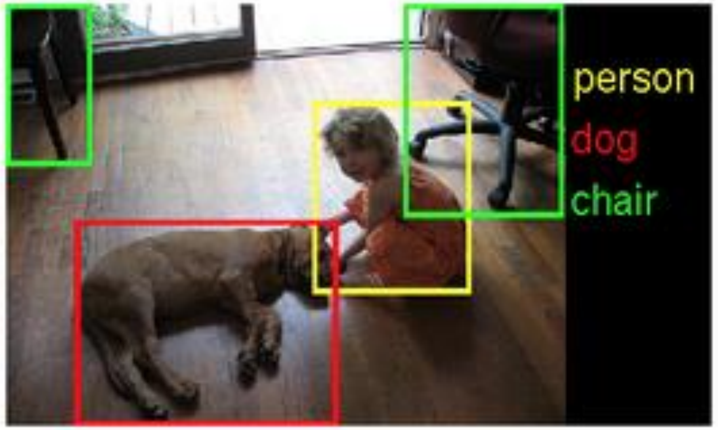

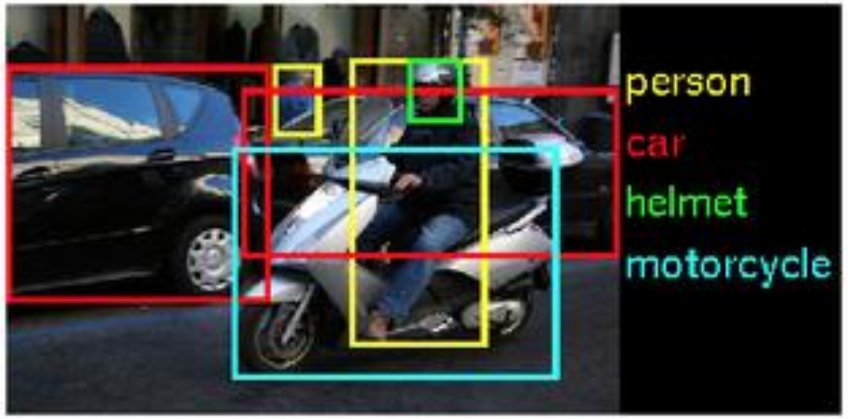

Results on ImageNet

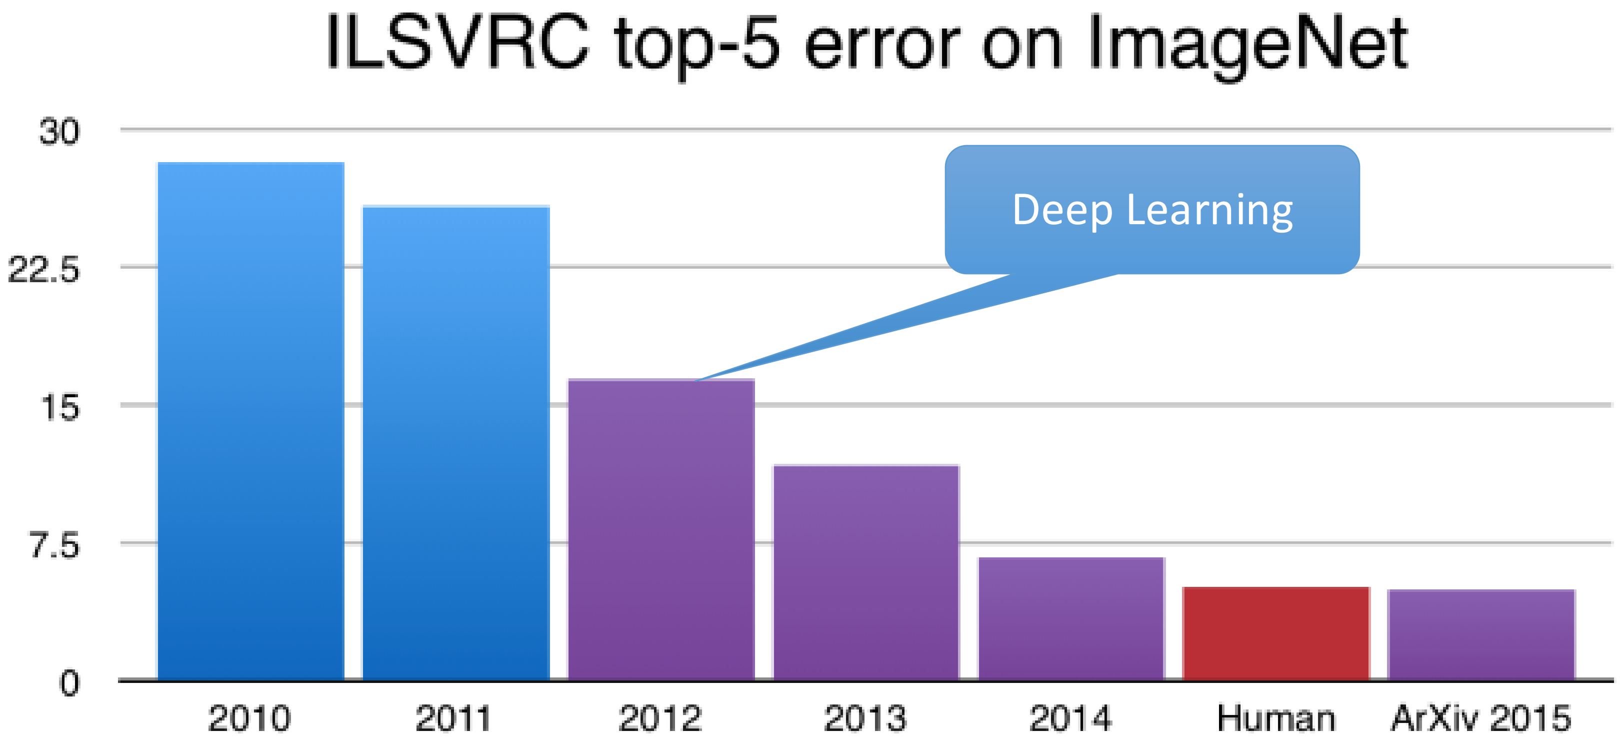

Large Language Models

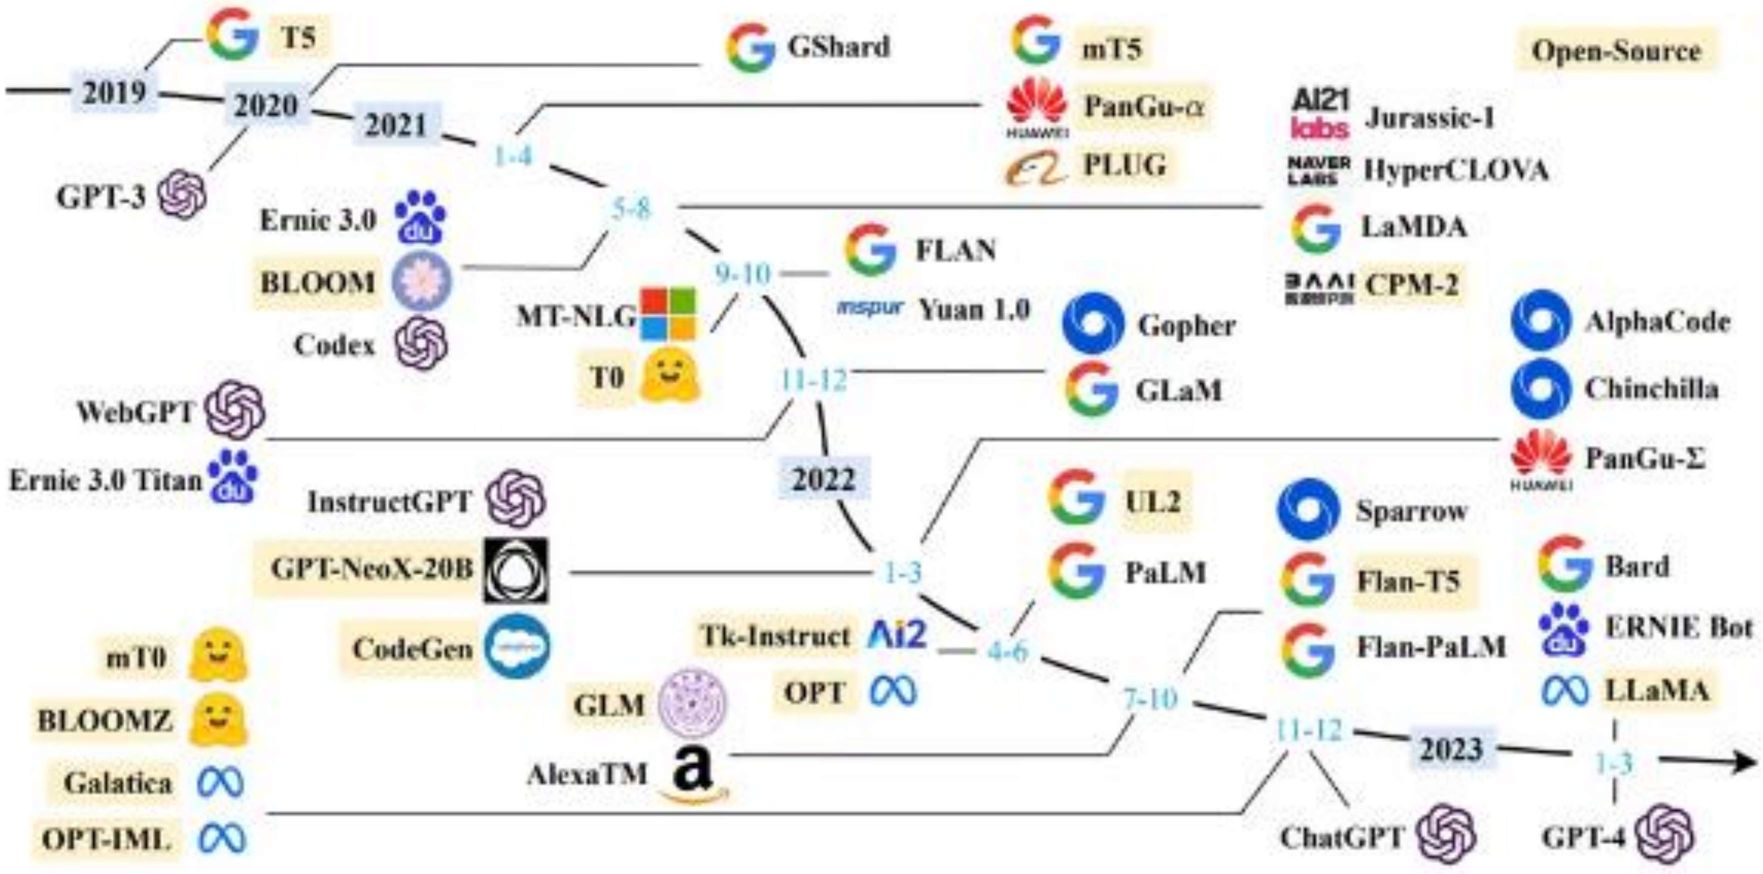

ChatGPT/GPT-4

"**ChatGPT** (Generative Pre-trained Transformer)[1] is a chatbot launched by OpenAI in November 2022. It is built on top of OpenAI's GPT-3.5 family of large language models, and is fine-tuned with both supervised and reinforcement learning techniques." --Wikipedia

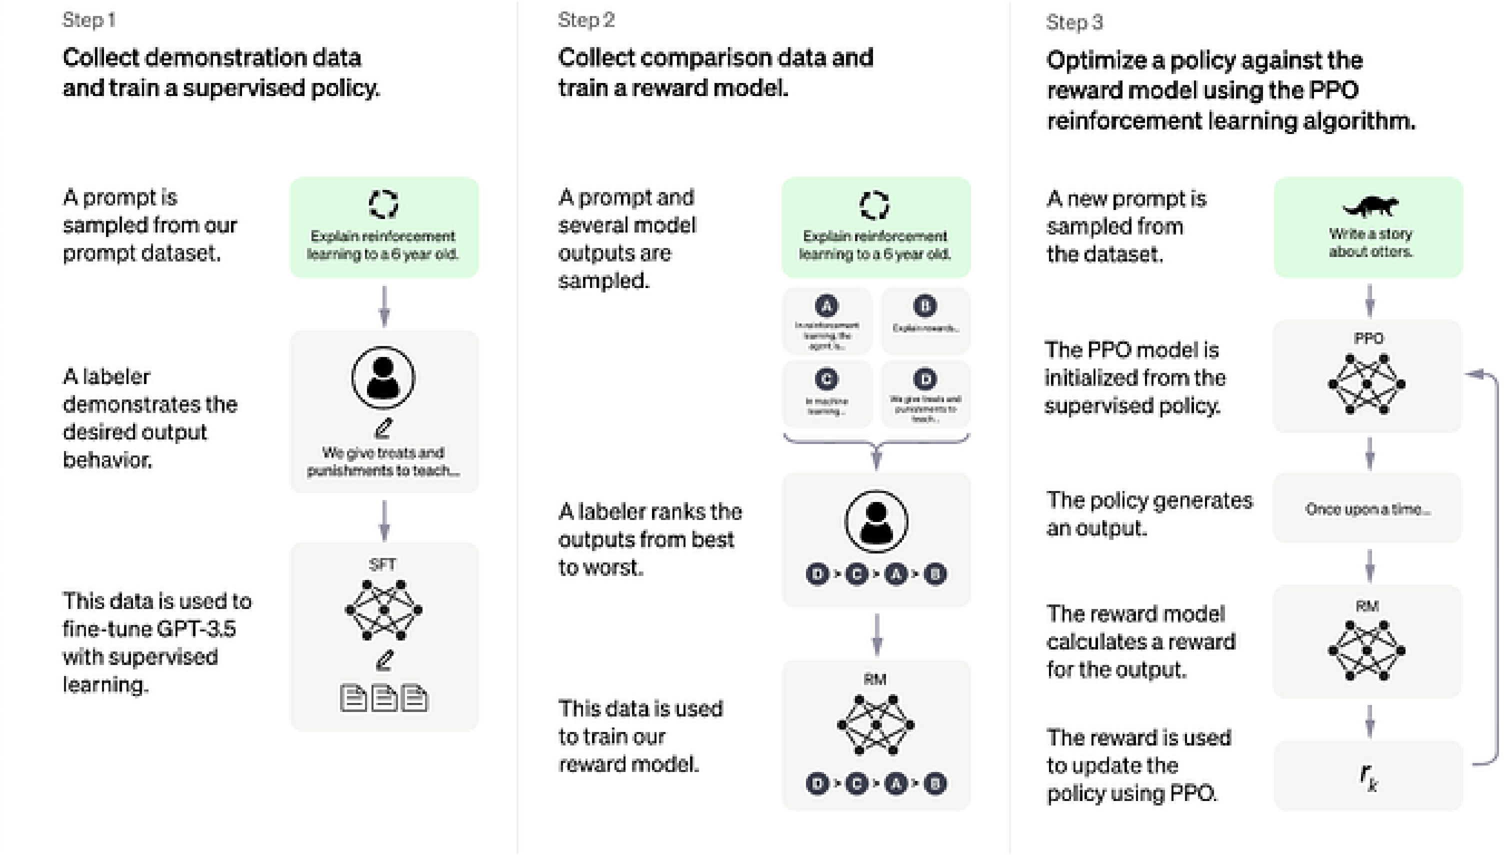

From Text to Image

- DALL·E3: Creating Images from Text by OpenAI
- https://openai.com/index/dall-e-3/

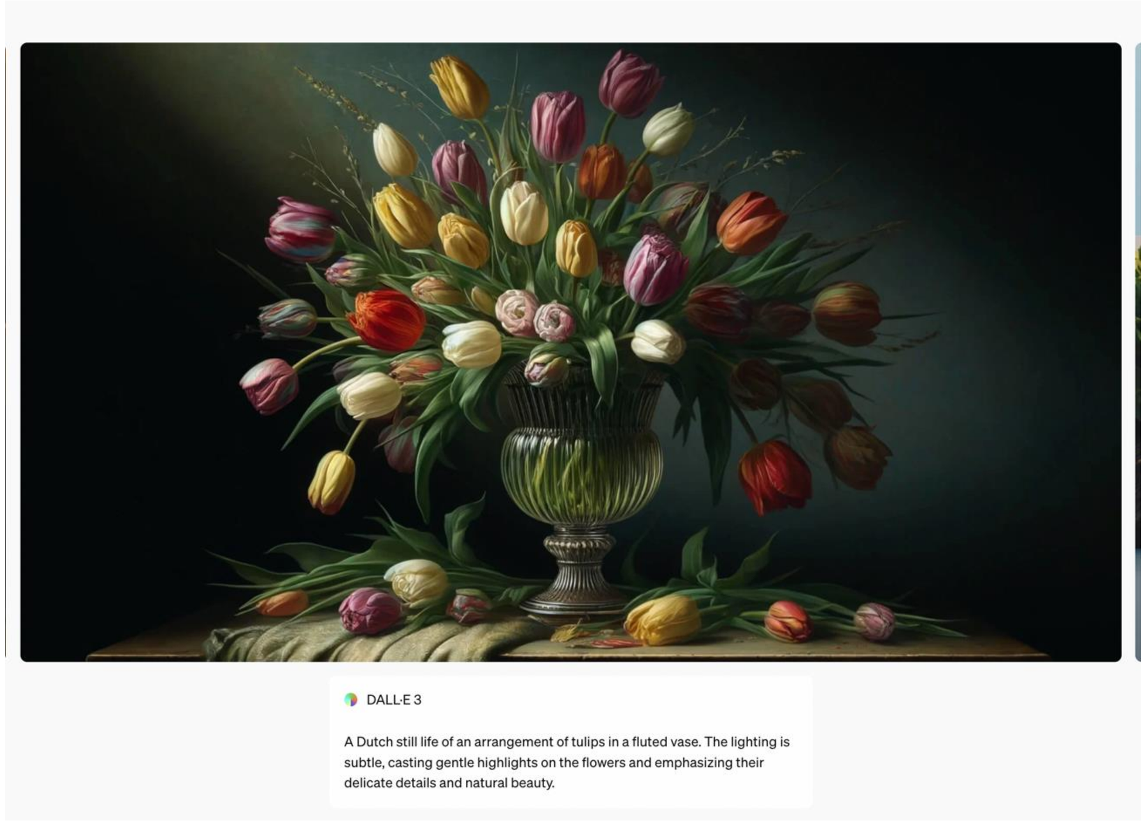

GPT-4o: Multi-modality Learning across Audio, Vision, and Text

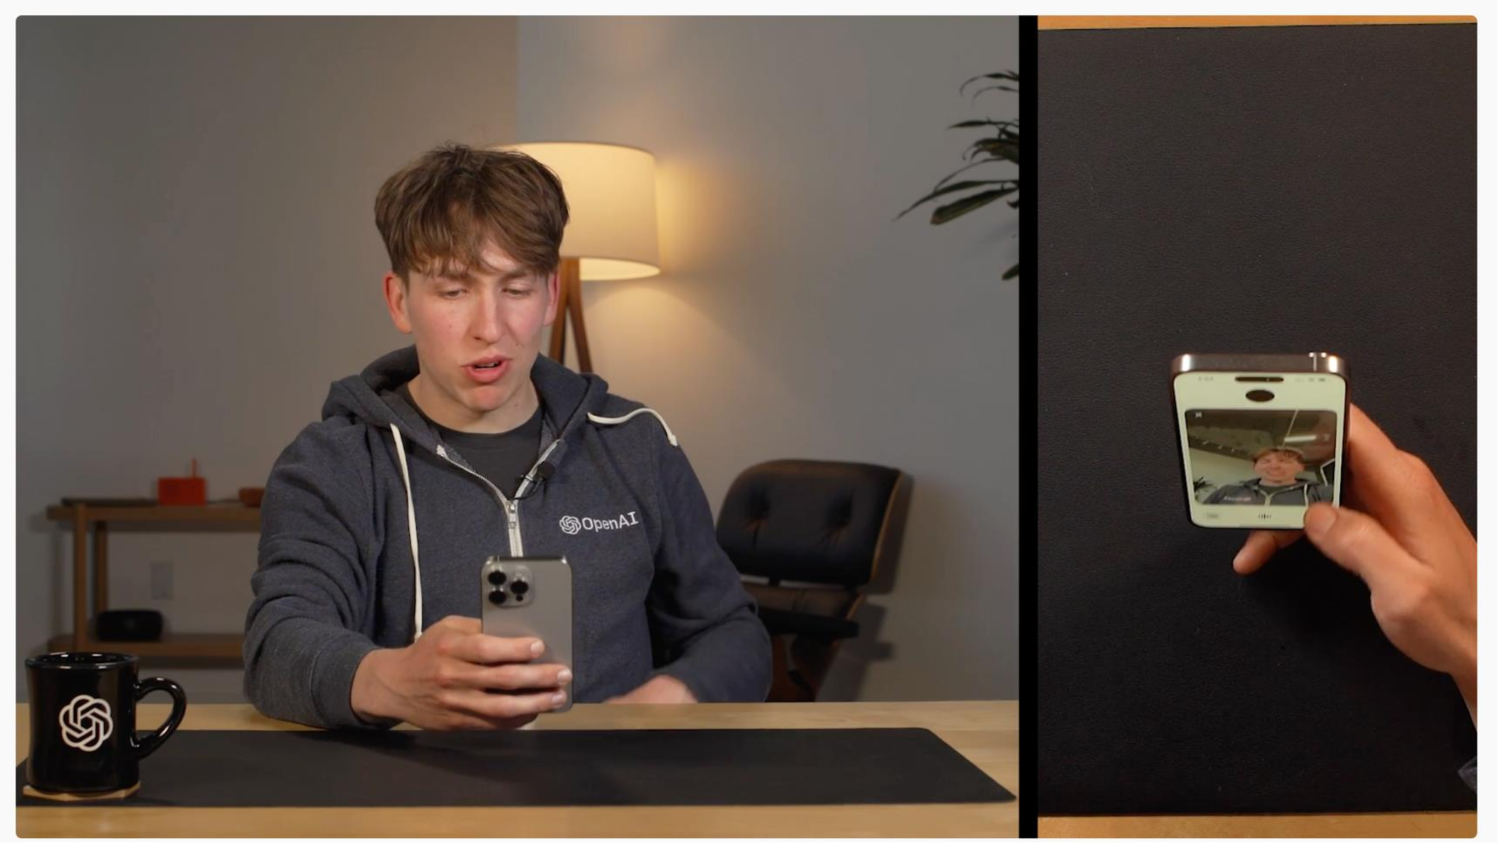

https://vimeo.com/945586717

From Text to Video

- Open AI Sora: https://openai.com/index/sora/

Prompt: A stylish woman walks down a Tokyo street filled with warm glowing neon and animated city signage. She wears a black leather jacket, a long red dress, and black boots, and carries a black purse. She wears sunglasses and red lipstick. She... +

- Kling AI: https://kling.kuaishou.com/en

Robotics

Graph/Networks Analysis

- Social networks, World Wide Web
- Molecules

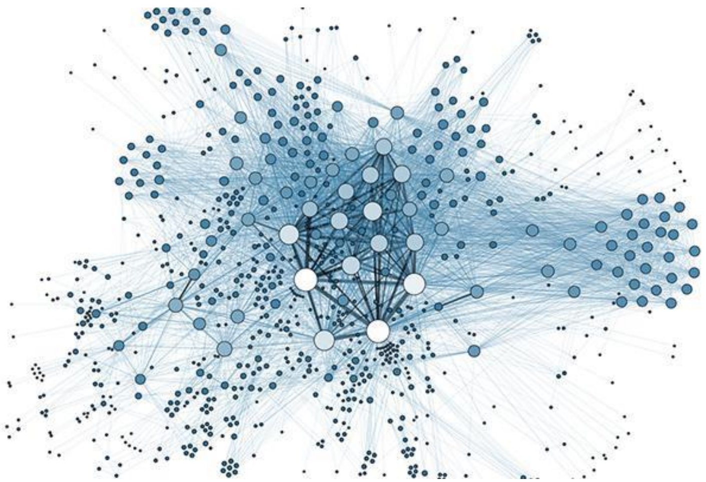

Social Networks

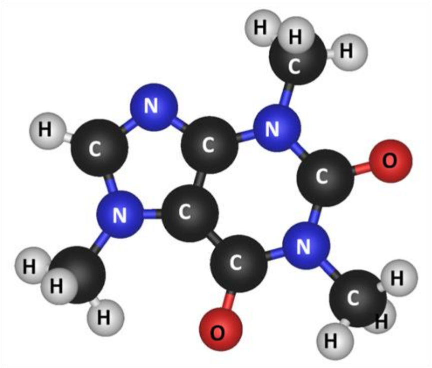

Molecules

Recommender Systems

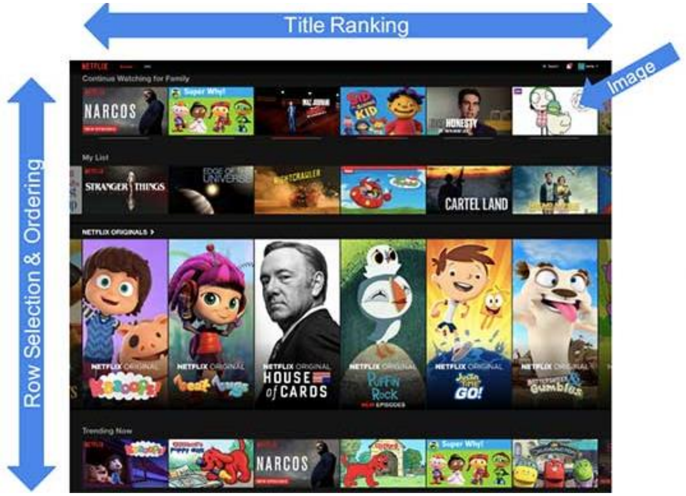

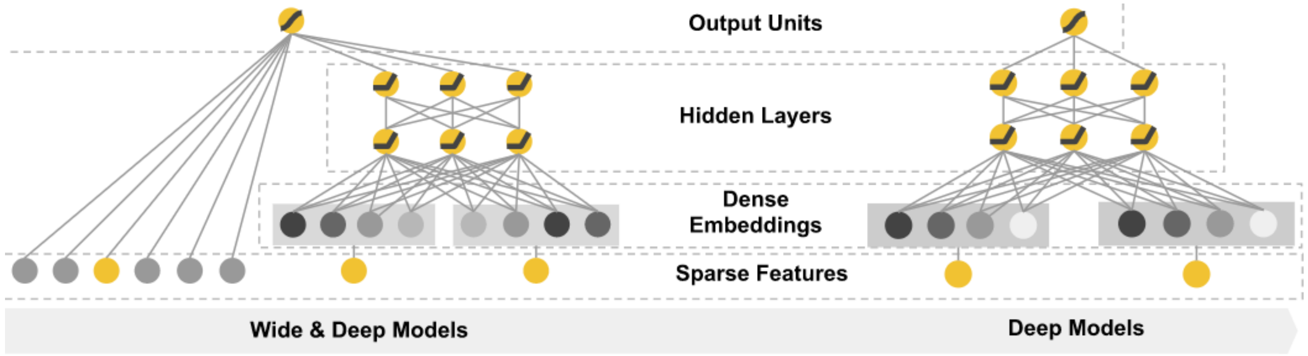

Wide & deep learning for
recommender systems (Google 2016)

Workshops on Deep Learning
for Recommender Systems

AI for Math

Recently, the IMO has also become an aspirational challenge for AI systems as a test of their advanced mathematical problem-solving and reasoning capabilities. Last year, Google DeepMind's combined AlphaProof and AlphaGeometry 2 systems achieved the silver-medal standard, solving four out of the six problems and scoring 28 points. Making use of specialist formal languages, this breakthrough demonstrated that AI was beginning to approach elite human mathematical reasoning.

This year, we were amongst an inaugural cohort to have our model results officially graded and certified by IMO coordinators using the same criteria as for student solutions. Recognizing the significant accomplishments of this year's student-participants, we're now excited to share the news of Gemini's breakthrough performance.

Breakthrough Performance at IMO 2025 with Gemini Deep Think

An advanced version of Gemini Deep Think solved five out of the six IMO problems perfectly, earning 35 total points, and achieving gold-medal level performance. The solutions can be found online here.

Deep Think

Through exploring the frontiers of Gemini's thinking capabilities, we're starting to test an enhanced reasoning mode called Deep Think that uses new research techniques enabling the model to consider multiple hypotheses before responding.

2.5 Pro Deep Think gets an impressive score on 2025 USAMO, currently one of the hardest math benchmarks. It also leads on LiveCodeBench, a difficult benchmark for competition-level coding, and scores 84.0% on MMMU, which tests multimodal reasoning.

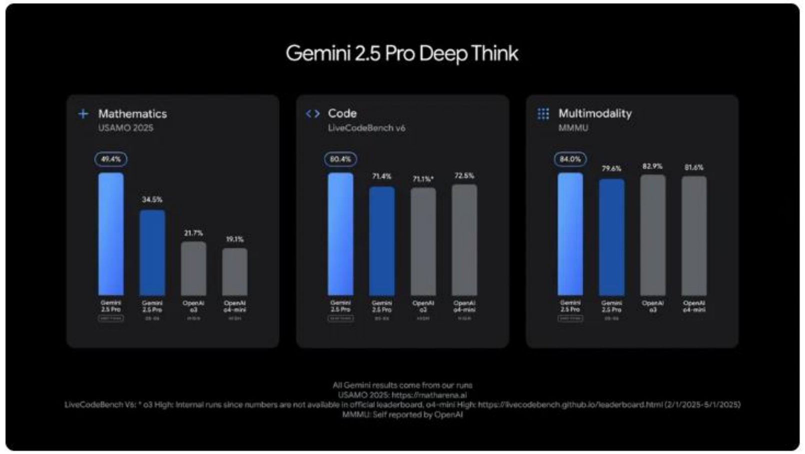

https://deepmind.google/discover/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-international-mathematical-olympiad/

AI for Material Discovery

nature

Article | Open access | Published: 29 November 2023

Scaling deep learning for materials discovery

Amil Merchant ✉, Simon Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon & Ekin Dogus Cubuk ✉

Nature **624**, 80–85 (2023) | Cite this article

Abstract

Novel functional materials enable fundamental breakthroughs across technological applications from clean energy to information processing[1,2,3,4,5,6,7,8,9,10,11]. From microchips to batteries and photovoltaics, discovery of inorganic crystals has been bottlenecked by expensive trial-and-error approaches. Concurrently, deep-learning models for language, vision and biology have showcased emergent predictive capabilities with increasing data and computation[12,13,14]. Here we show that graph networks trained at scale can reach unprecedented levels of generalization, improving the efficiency of materials discovery by an order of magnitude. Building on 48,000 stable crystals identified in continuing studies[15,16,17], improved efficiency enables the discovery of 2.2 million structures below the current convex hull, many of which escaped previous human chemical intuition. Our work represents an order-of-magnitude expansion in stable materials known to humanity. Stable discoveries that are on the final convex hull will be made available to screen for technological applications, as we demonstrate for layered materials and solid-electrolyte candidates. Of the stable structures, 736 have already been independently experimentally realized. The scale and diversity of hundreds of millions of first-principles calculations also unlock modelling capabilities for downstream applications, leading in particular to highly accurate and robust learned interatomic potentials that can be used in condensed-phase molecular-dynamics simulations and high-fidelity zero-shot prediction of ionic conductivity.

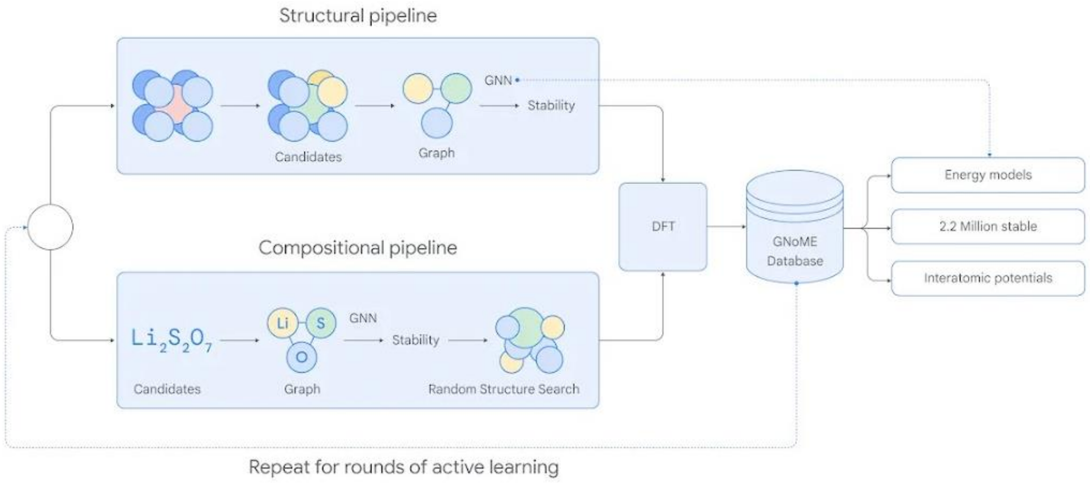

GNoME: Harnessing graph networks for materials exploration

AI for Protein Structure Prediction

- Protein Structure Prediction: a fundamental problem in biology
  - Predict the 3D structures of proteins from their amino acid sequences
  - https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology
- A breakthrough of protein structure prediction by deep learning (December, 2020)

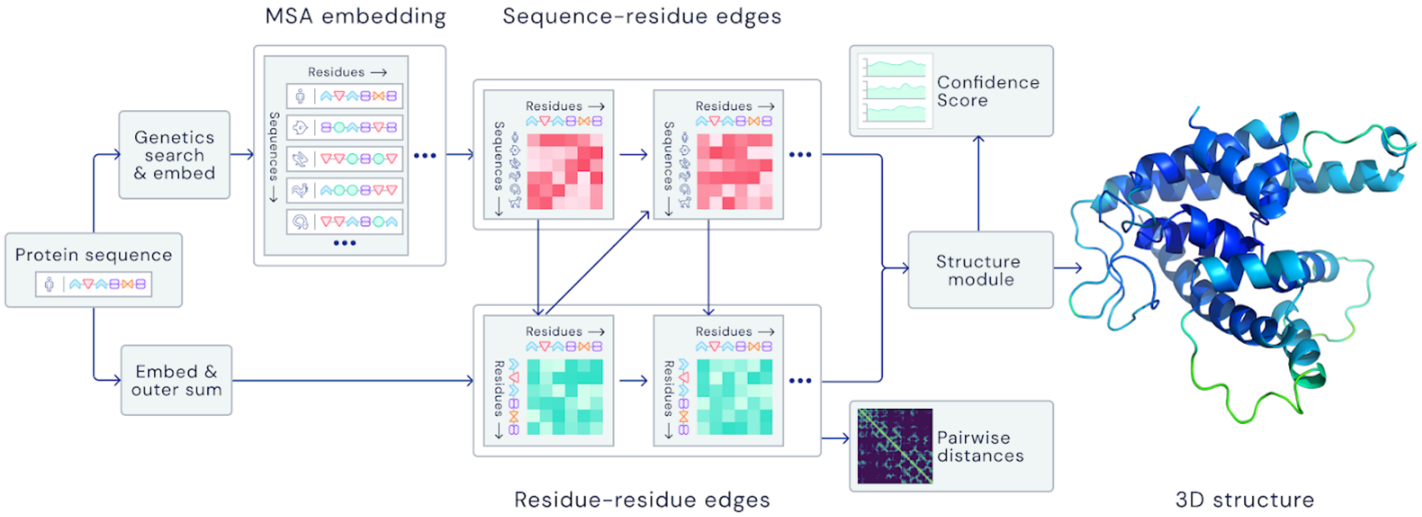

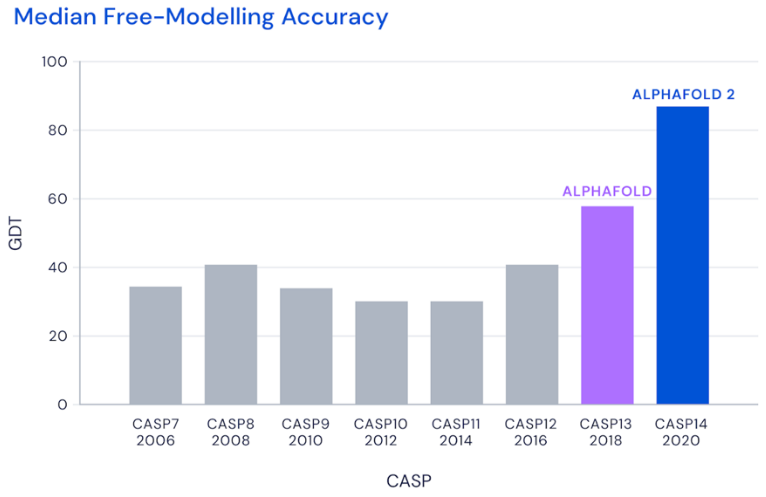

AlphaFold3

- Predicts the structure and interactions of all of life's molecules
  - predicting the structure of proteins, DNA, RNA, ligands and more, and how they interact

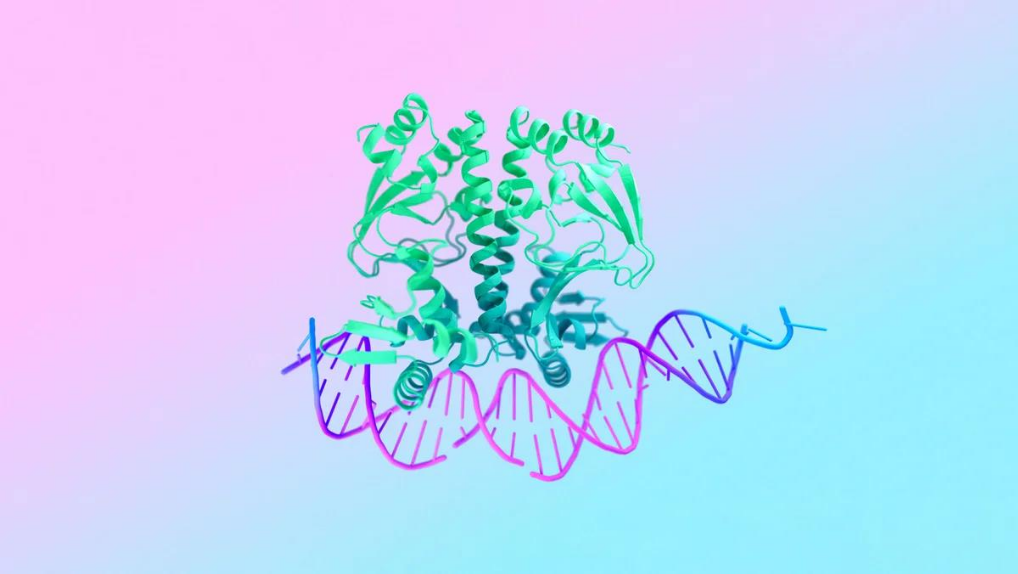

AI for Protein Design/Generation

- RFDiffusion

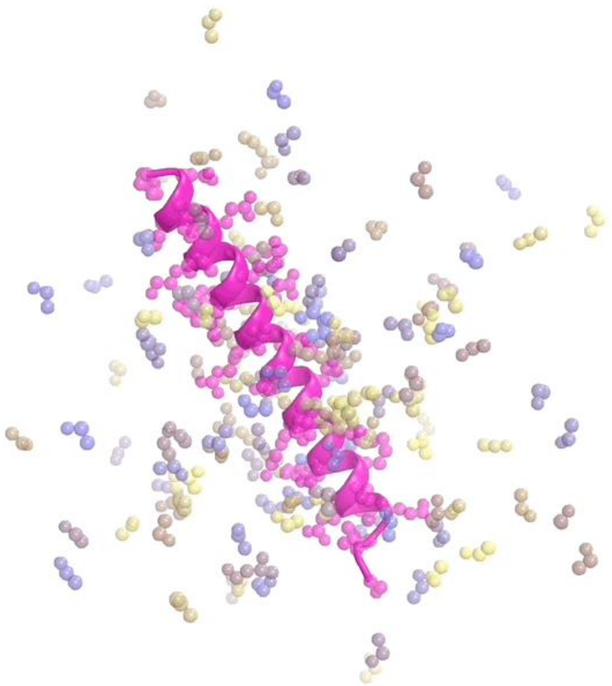

Watson et al. **De novo design of protein structure and function with Rfdiffusion**. Nature, 2023.

This Course

- **Objectives**
  - Introduce the fundamental techniques of deep learning
  - Introduce the applications of deep learning in a few selected domains
  - Learn how to use the Pytorch framework
  - Learn to use deep learning techniques to solve real-world problems
- **Prerequisite**
  - Mathematics **mature** (especially probability, statistics, and linear algebra)
  - Python programming

Textbooks

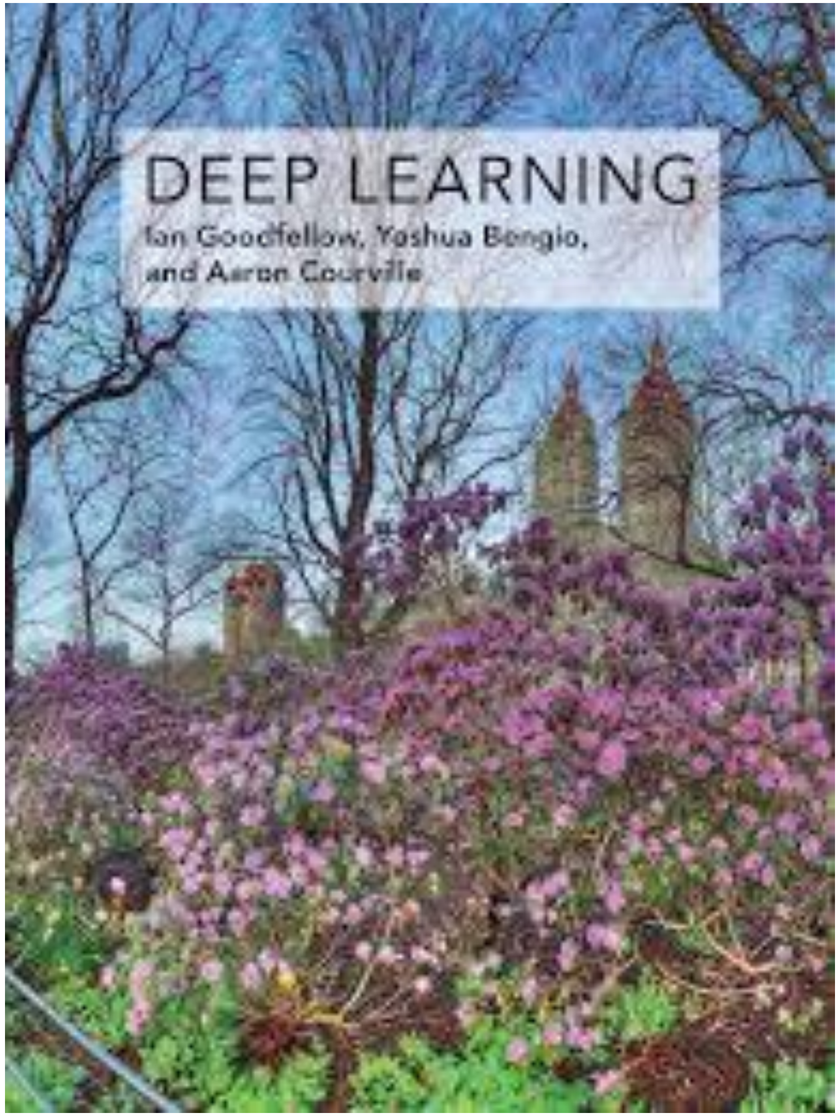

Ian Goodfellow, Yoshua Bengio and Aaron Courville. “Deep Learning”. MIT, 2016.

Evaluation

- Homework (20%)
- Class Presentations (10%)
- Course Project (30%)
  - Research proposal: 5%
  - Poster: 10%
  - Report: 15%
- Final Exam (40%)

Homework (20%)

- Two homework
  - Two Pytorch programming exercises

Class Presentations (10%)

- Each group (<=2 students) will have to present one paper in class
  - 20-25 minutes
- The papers are usually more advanced topics that are not covered during classes
  - You can select the papers listed in the "Reference" column (https://deepgraphlearning.github.io/Math80648A_2025A/index.html)
  - First come, first serve
- Register your group here:
  - https://docs.google.com/spreadsheets/d/1U4Y9uMrM6qQKSRW6KgvhnZLDtpV8nlhVEcF5HVsnQww/edit?usp=sharing

Course Projects (30%)

- Working in groups
  - Each group should have at most 2 students (the same as the presentation group)
- Step 1: Group Registration
- Step 2: Project proposal
- Step 3: Poster session (last class)
- Step 4: Submit a project report
  - One week after the poster session
- More instruction details on this will arrive

Final Exam (40%)

- Open book exam

Homework Late Policy

- Each student should hand in your homework or project report on time
  - 50% grade reduction (within 2 days after the deadline)
  - 0 points (more than 2 days passed after the deadline)

Course Outline

- **Part I: Machine Learning and Deep Learning Basics**
  - Week 1: Introduction, Mathematics, Machine Learning Basics
  - Week 2: Feedforward Neural Networks & Optimization Tricks
  - Week 3: PyTorch
  - Week 4: Convolutional Neural Networks and Recurrent Neural Networks

Course Outline

- **Part II: Applications and More Advanced Topics**
  - Week 5: Word Representation Learning
  - Week 6: Attention, Transformers
  - Week 7: No Class (Project proposal ready)
  - Week 8: Large Language Models
  - Week 9: Large Language Models

Course Outline

- **Part II: Applications and More Advanced Topics**
  - Week 10: Generative Models
  - Week 11: Graph Representation Learning, Graph Neural Networks
  - Week 12: Poster Session

Questions?