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/MATH60630A_2024A/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: Jianan Zhao, Xinyu Yuan
  - Ph.D at Mila and University of Montreal
  - Email: jianan.zhao@mila.quebec, xinyu.yuan@mila.quebec
- Discord: https://discord.gg/JVcJs2Ww
  - 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

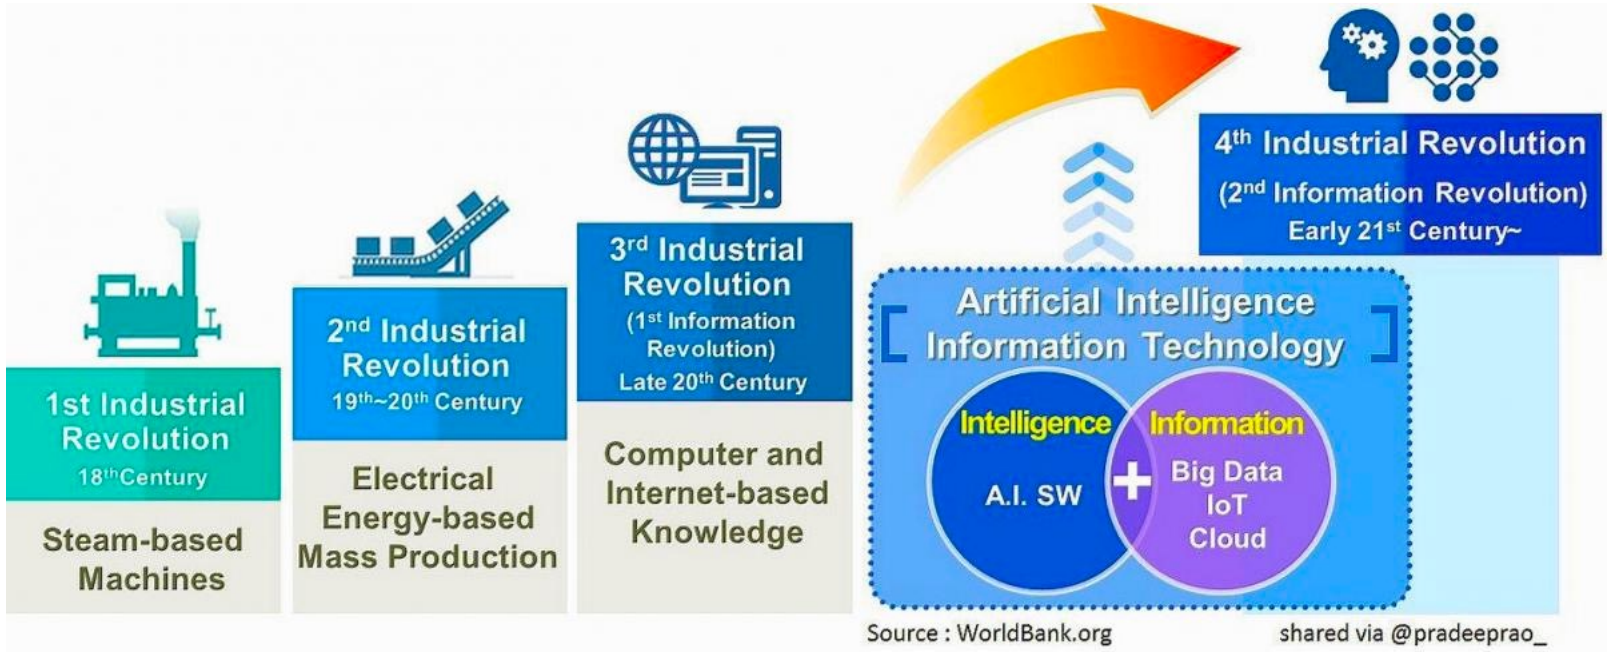

-image from Internet

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

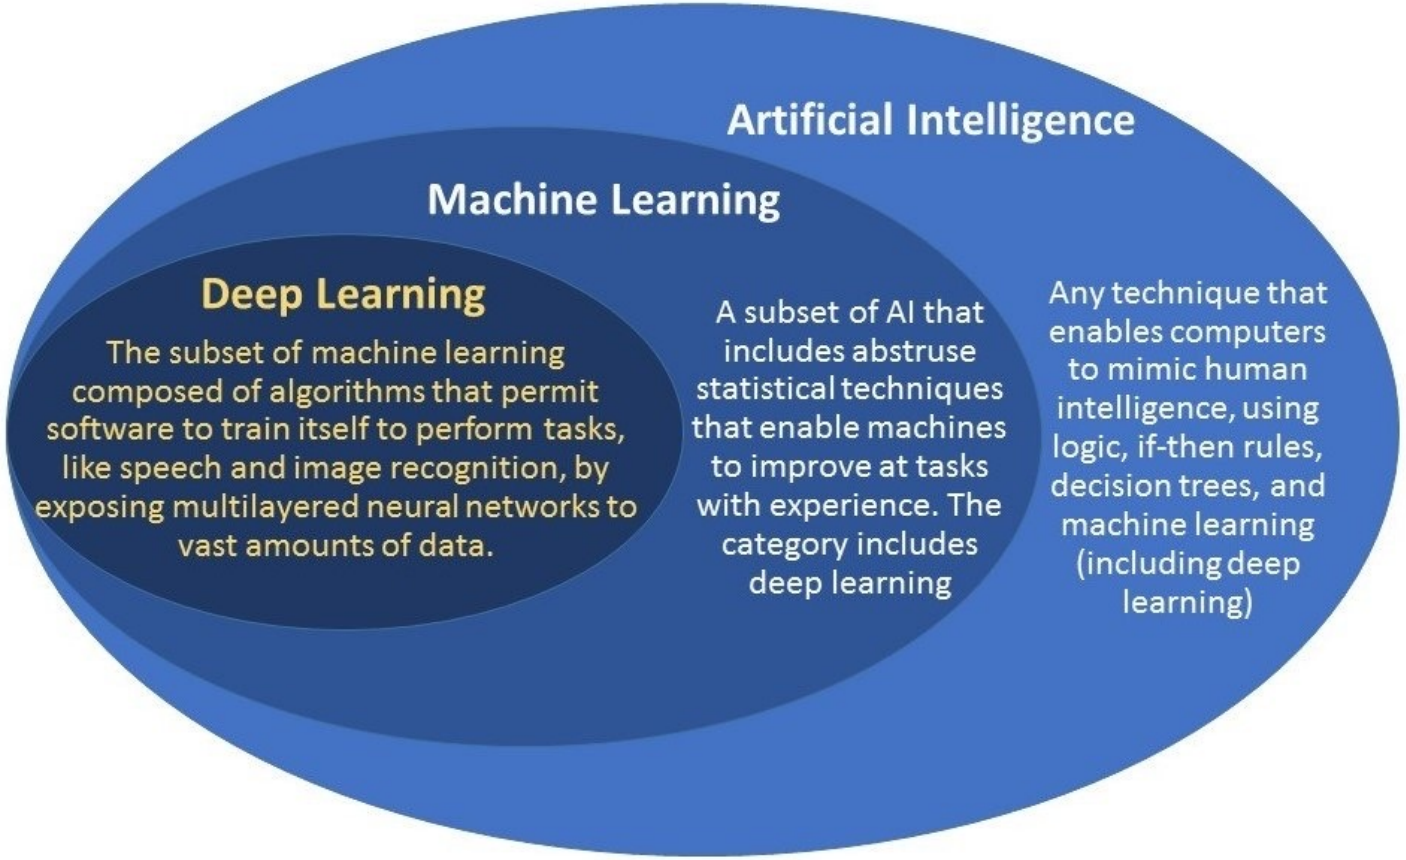

Machine Learning

- “**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

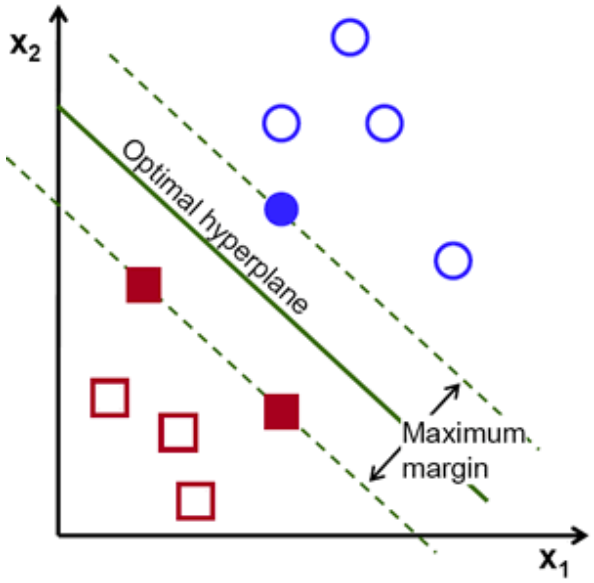

Support vector machines

Domain experts

Deep Learning = Feature Representation Learning

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

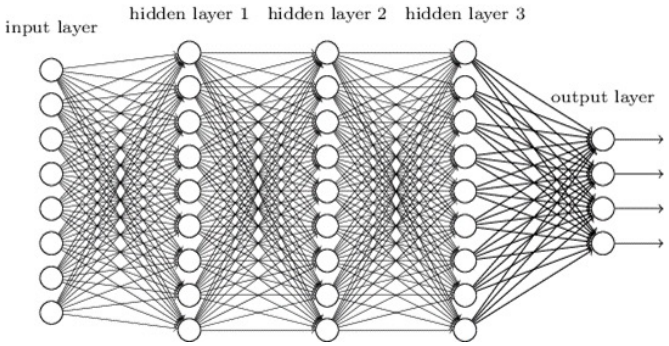

Deep Neural Networks

Trainable
Feature Extractor

Simple Trainable Classifier
e.g., SVM, LR

Domain experts

Why Deep Learning Now?

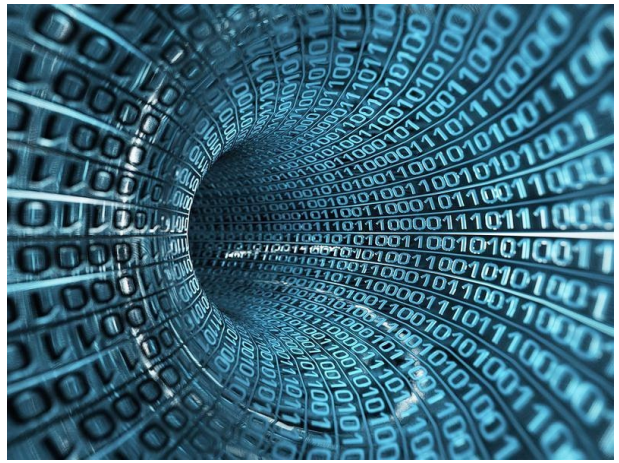

Big Data

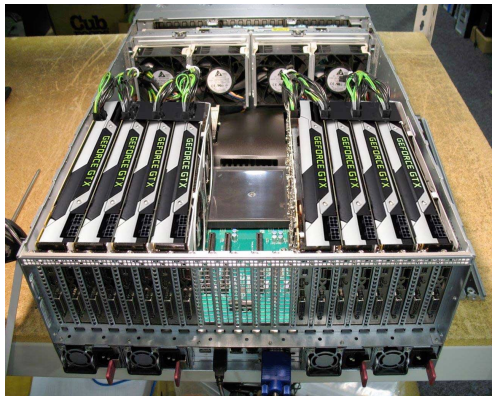

Big Computation

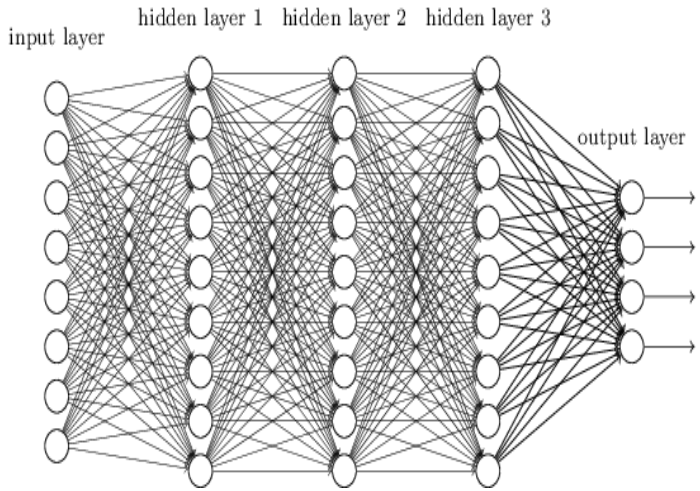

Big Model


FULL SELF DRIVING
[V12.3]
10
MPH
Now
Alum Rock Ave

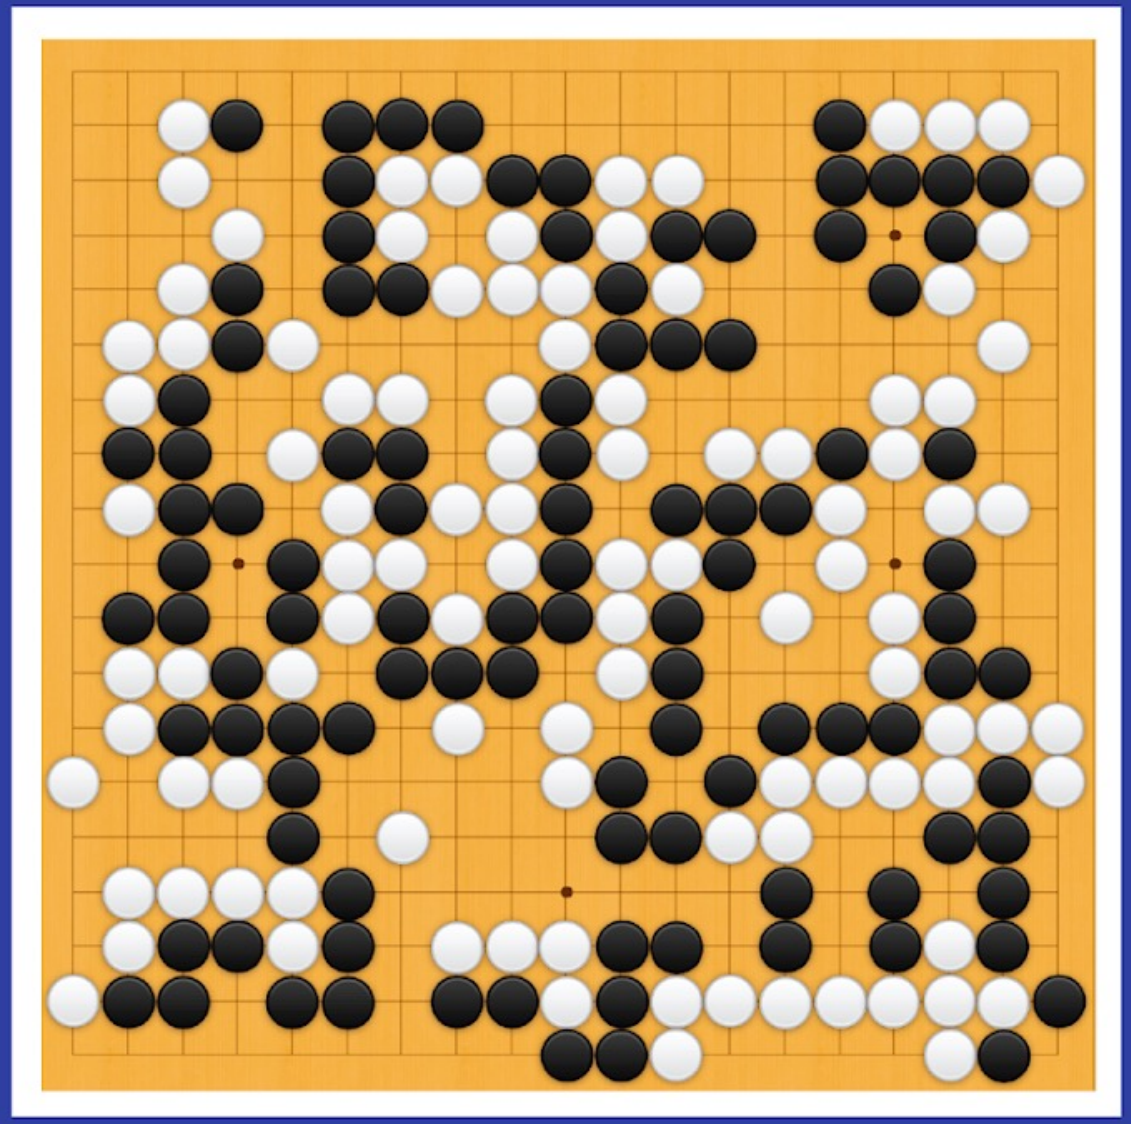

THE ULTIMATE GO CHALLENGE

GAME 3 OF 3

27 MAY 2017

AlphaGo vs Ke Jie

Winner of Match 3

RESULT B + Res

DEEP LEARNING ENABLES RAPID IDENTIFICATION OF POTENT DDR1 KINASE INHIBITORS

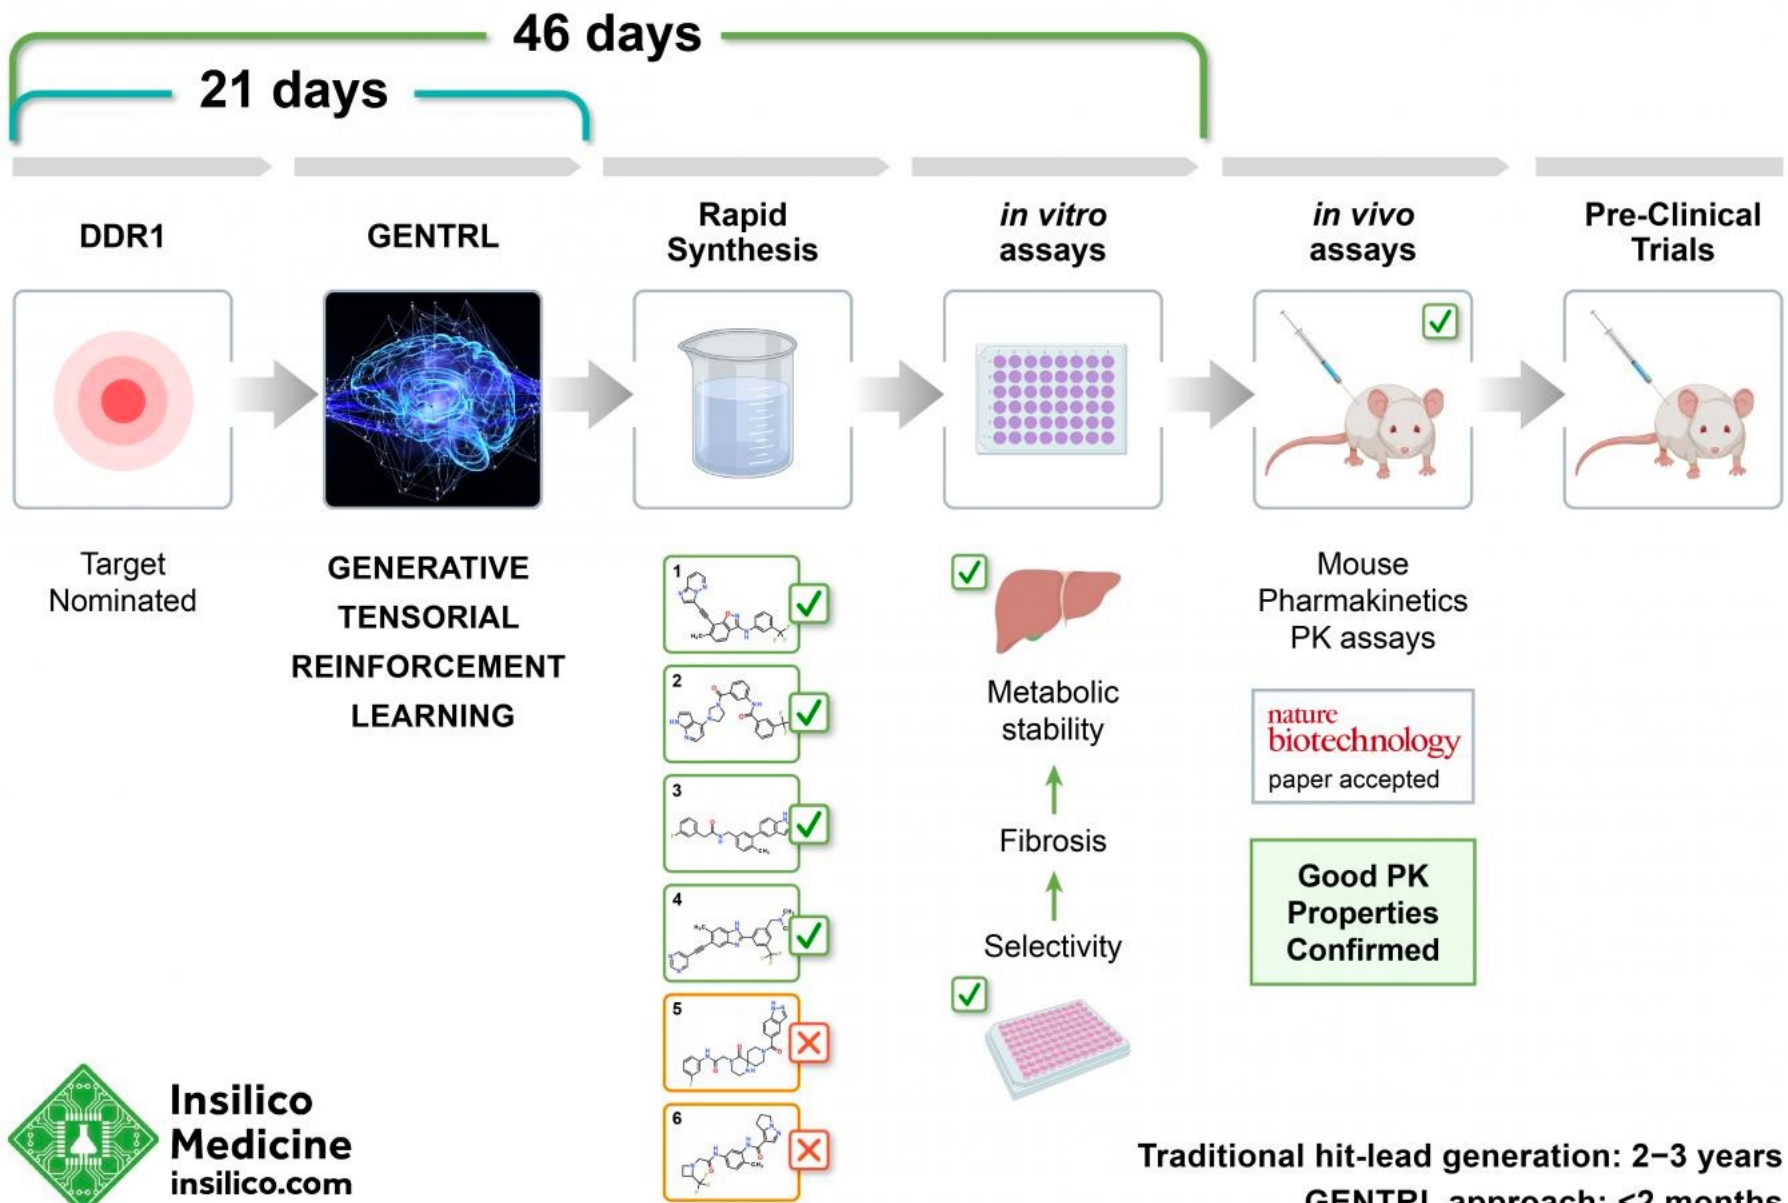

Speech Recognition

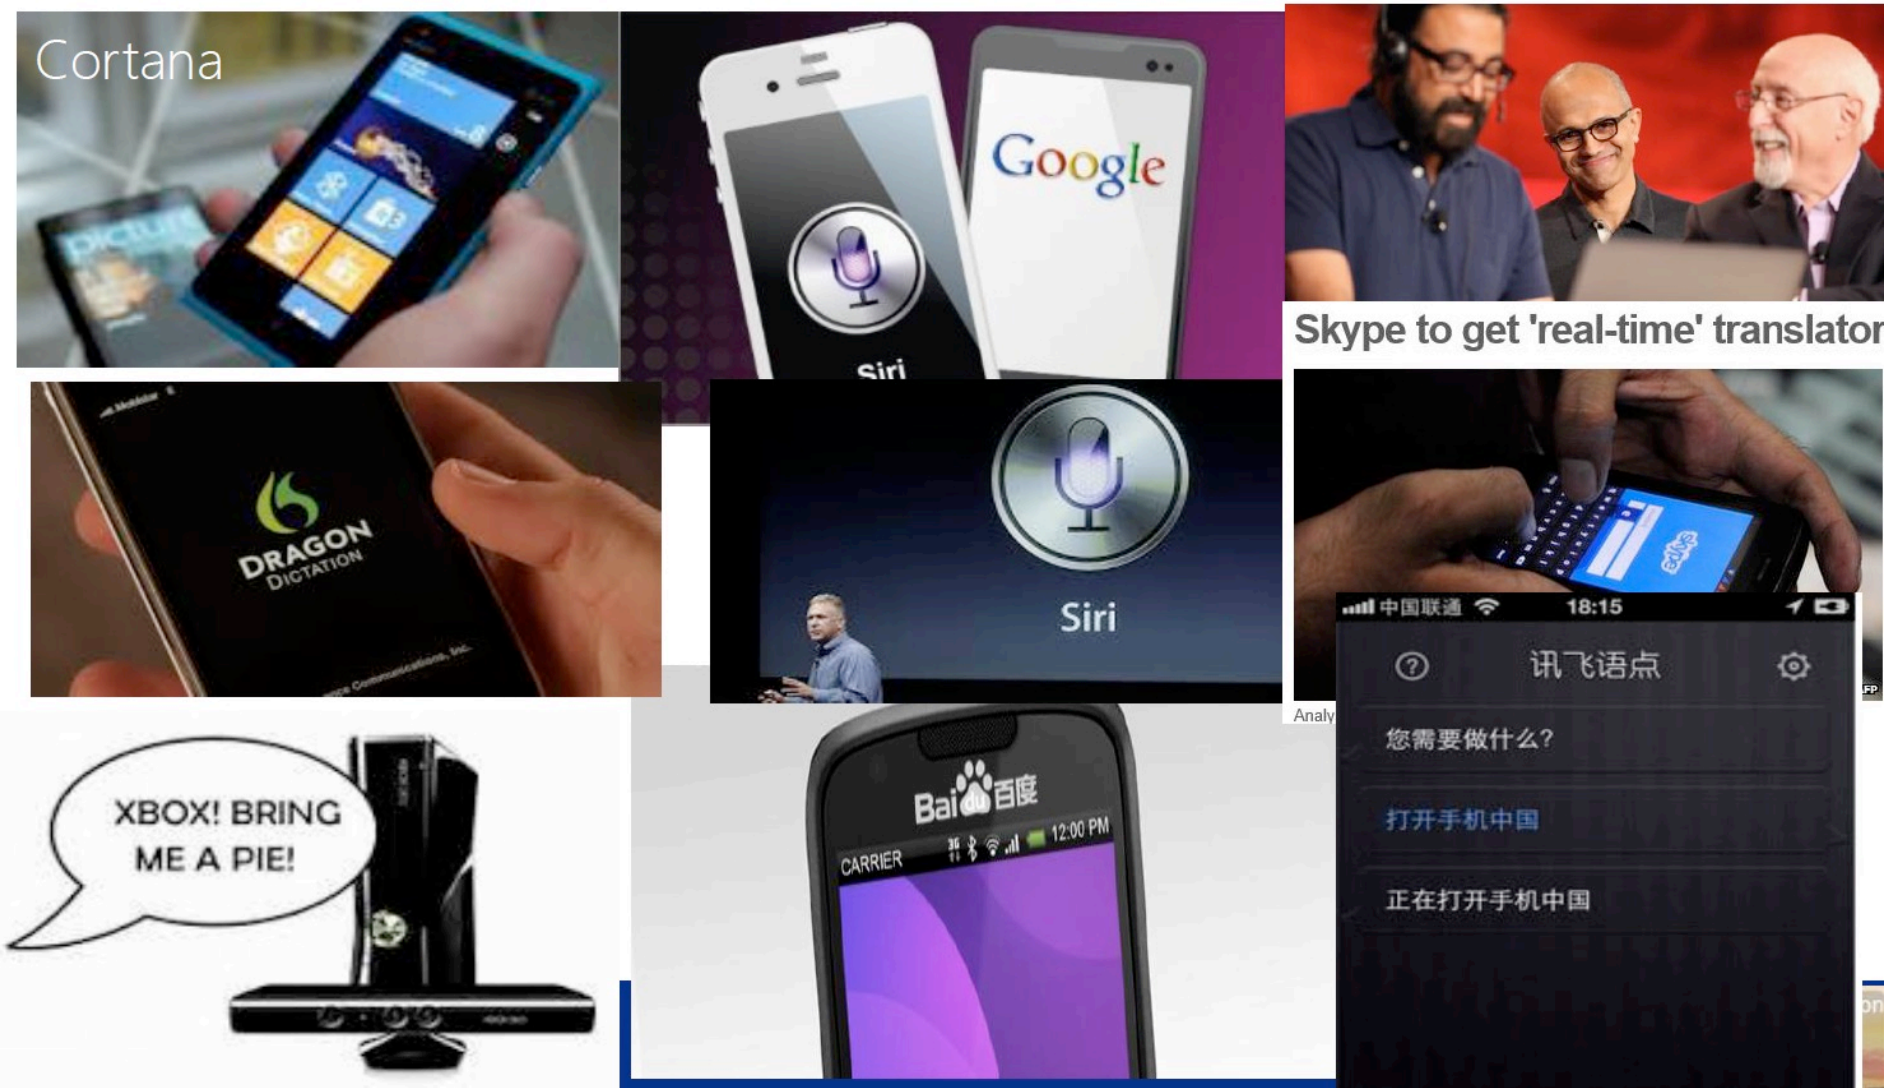

Speech Recognition Results

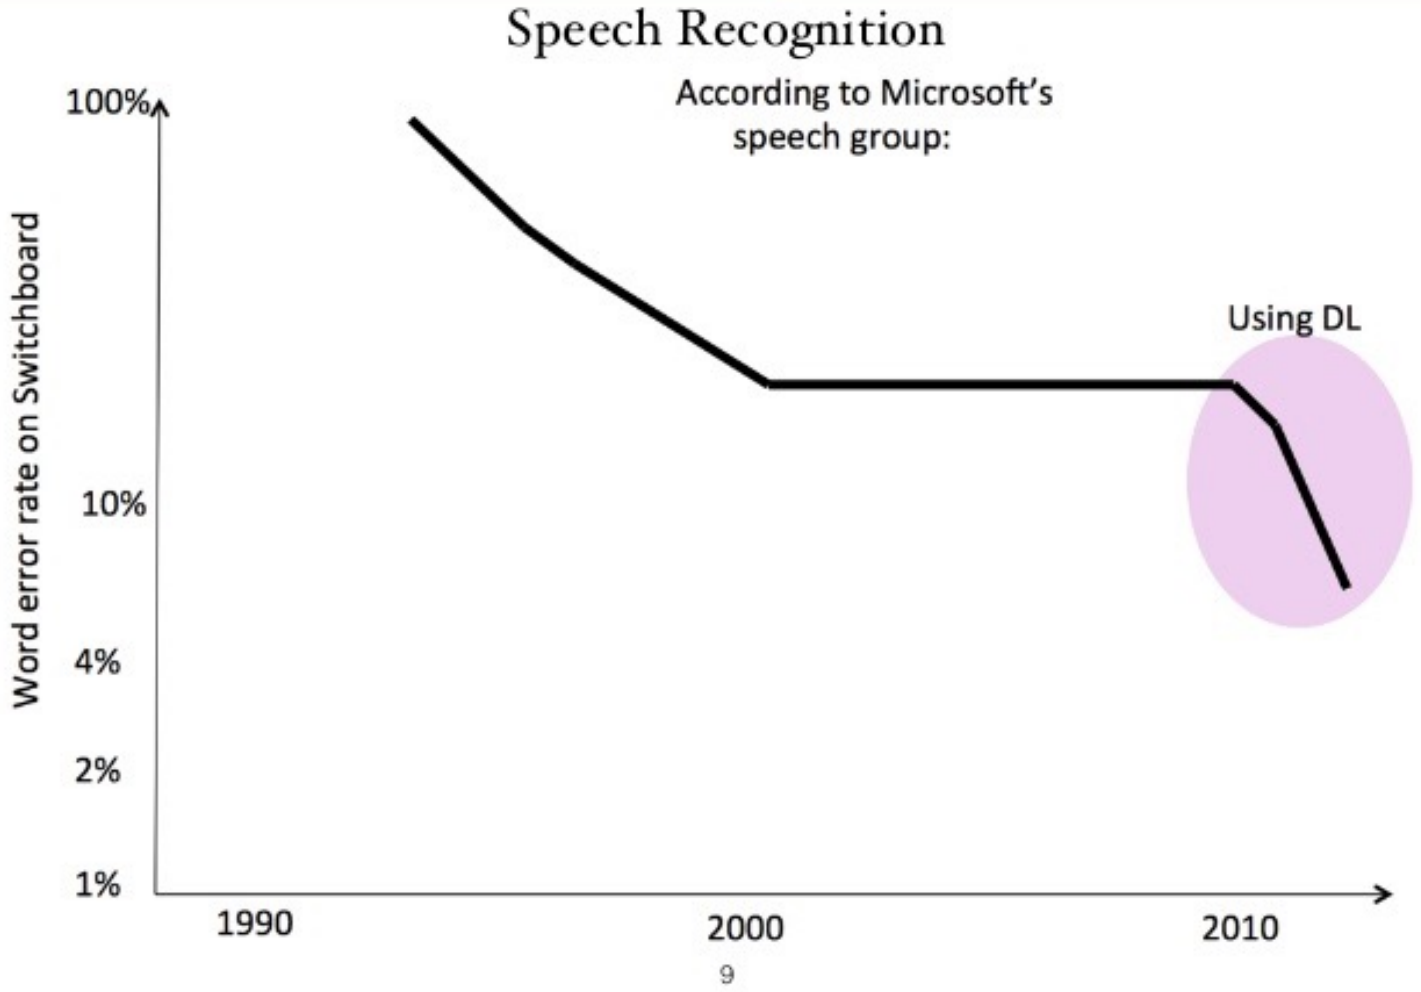

(Figure from Microsoft's speech Group)

Image Recognition

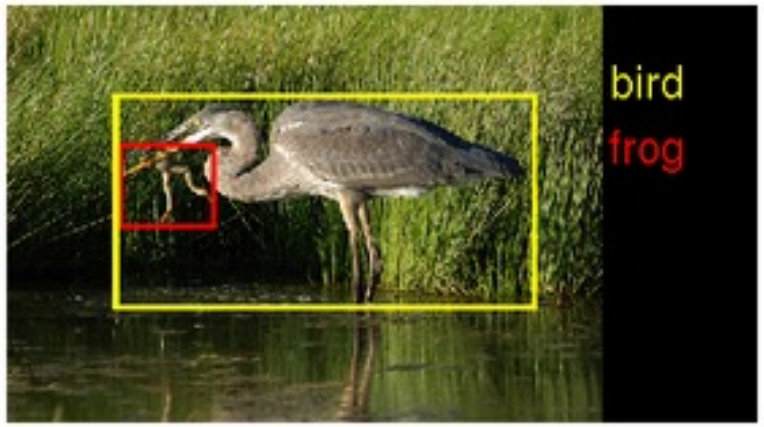

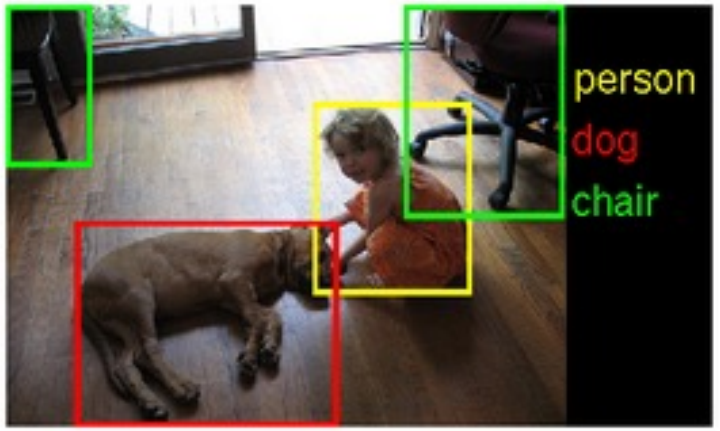

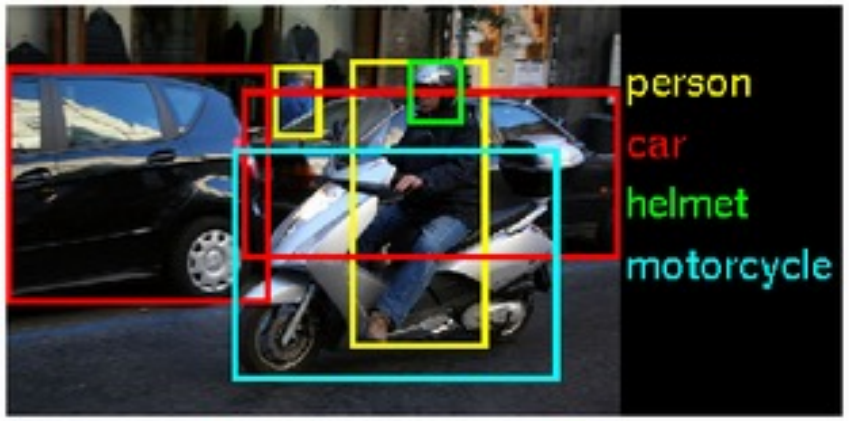

Results on ImageNet

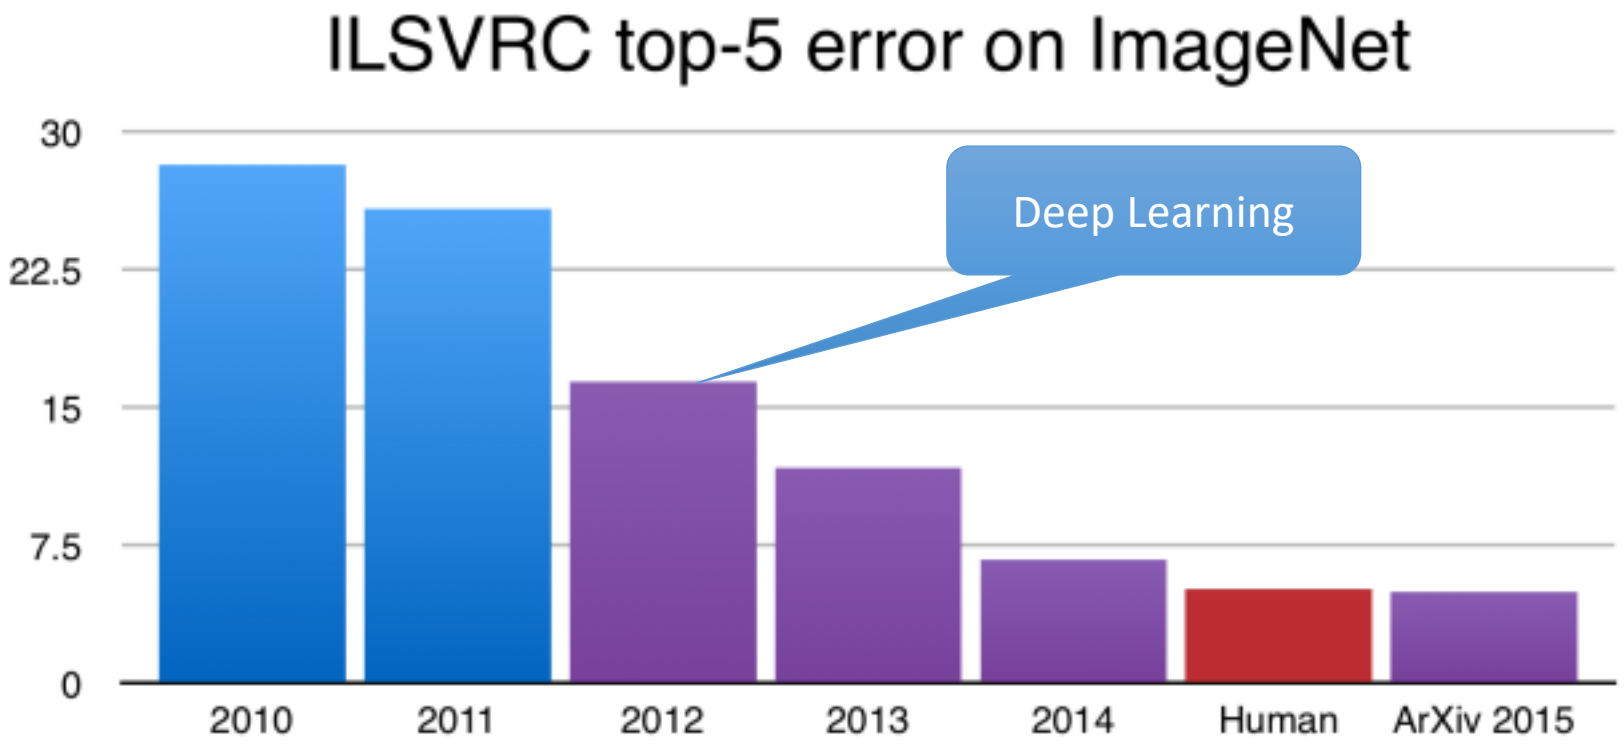

Image Generation

Image generation
(by StyleGAN, From Internet)


你好
Hi
Hola

Machine Translation

- 2016.9, Google announce its *neural machine translation* system.
- 2018.3 , Microsoft claimed its NMT achieved “human parity” on automatic Chinese to English news translation.

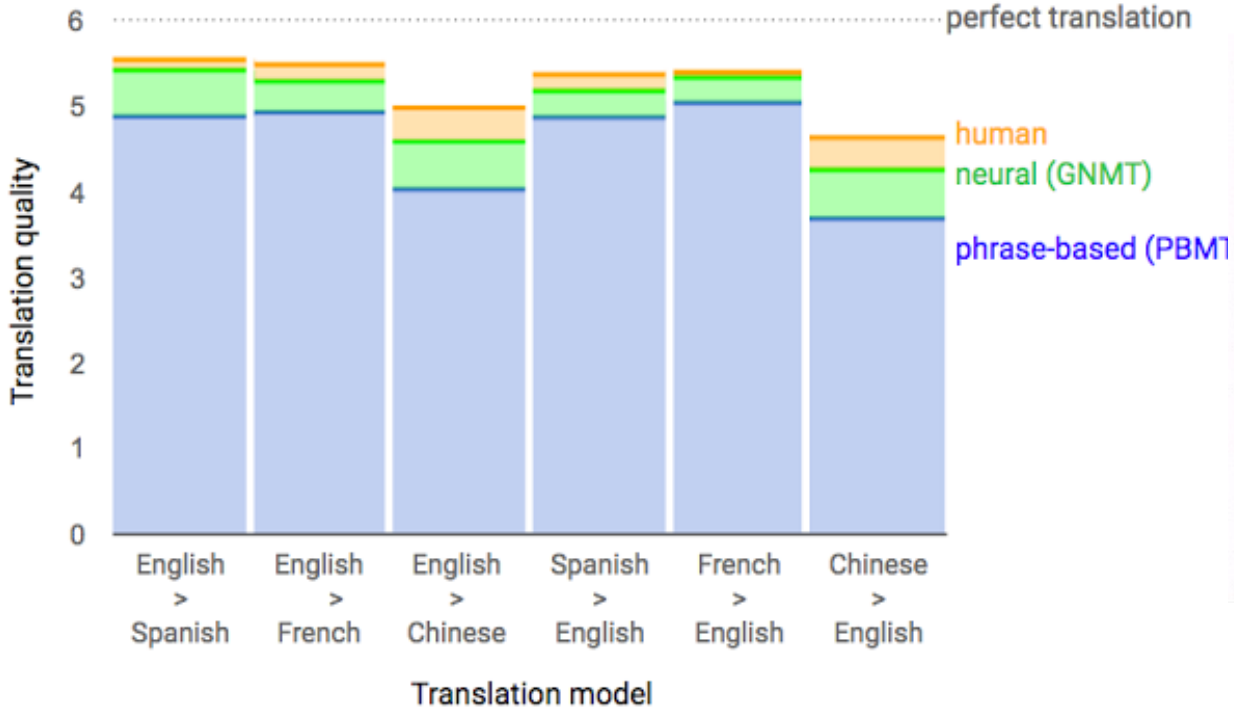

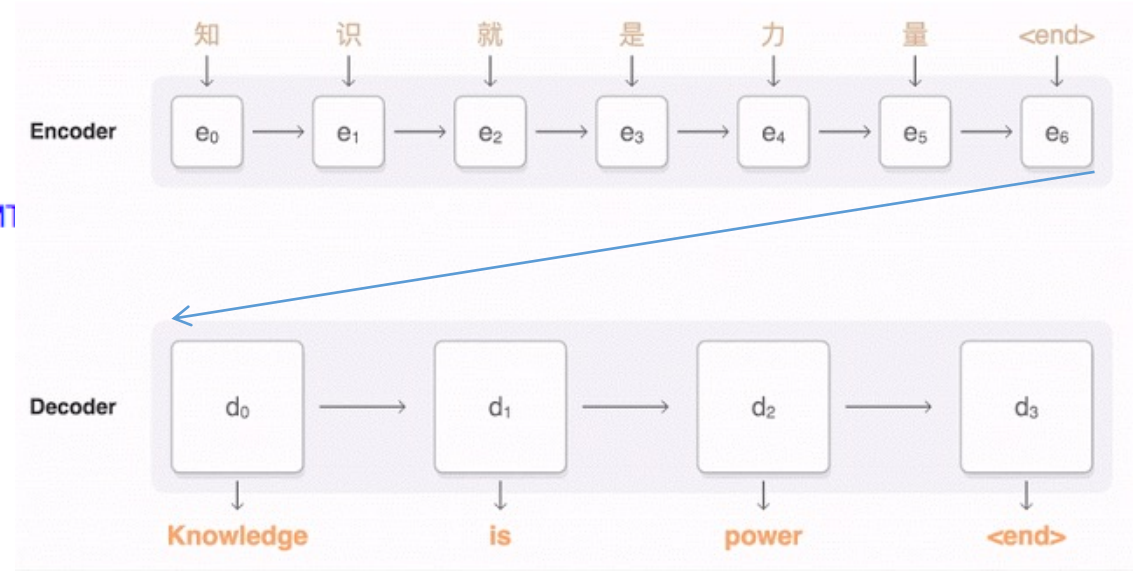

(Seq2Seq, Sutskever et al. 2014)

Machine Reading Comprehension

Passage: Tesla later approached Morgan to ask for more funds to build a more powerful transmitter. **When asked where all the money had gone, Tesla responded by saying that he was affected by the Panic of 1901**, which he (Morgan) had caused. Morgan was shocked by the reminder of his part in the stock market crash and by Tesla's breach of contract by asking for more funds. Tesla wrote another plea to Morgan, but it was also fruitless. Morgan still owed Tesla money on the original agreement, and Tesla had been facing foreclosure even before construction of the tower began.

Question: On what did Tesla blame for the loss of the initial money?

Answer: Panic of 1901

Rank	Model	EM	F1
	Human Performance *Stanford University* (Rajpurkar et al. '16)	82.304	91.221
1 Mar 19, 2018	QANet (ensemble) *Google Brain & CMU*	**83.877**	**89.737**
2 Jan 22, 2018	Hybrid AoA Reader (ensemble) *Joint Laboratory of HIT and iFLYTEK Research*	82.482	89.281

Large Language Models

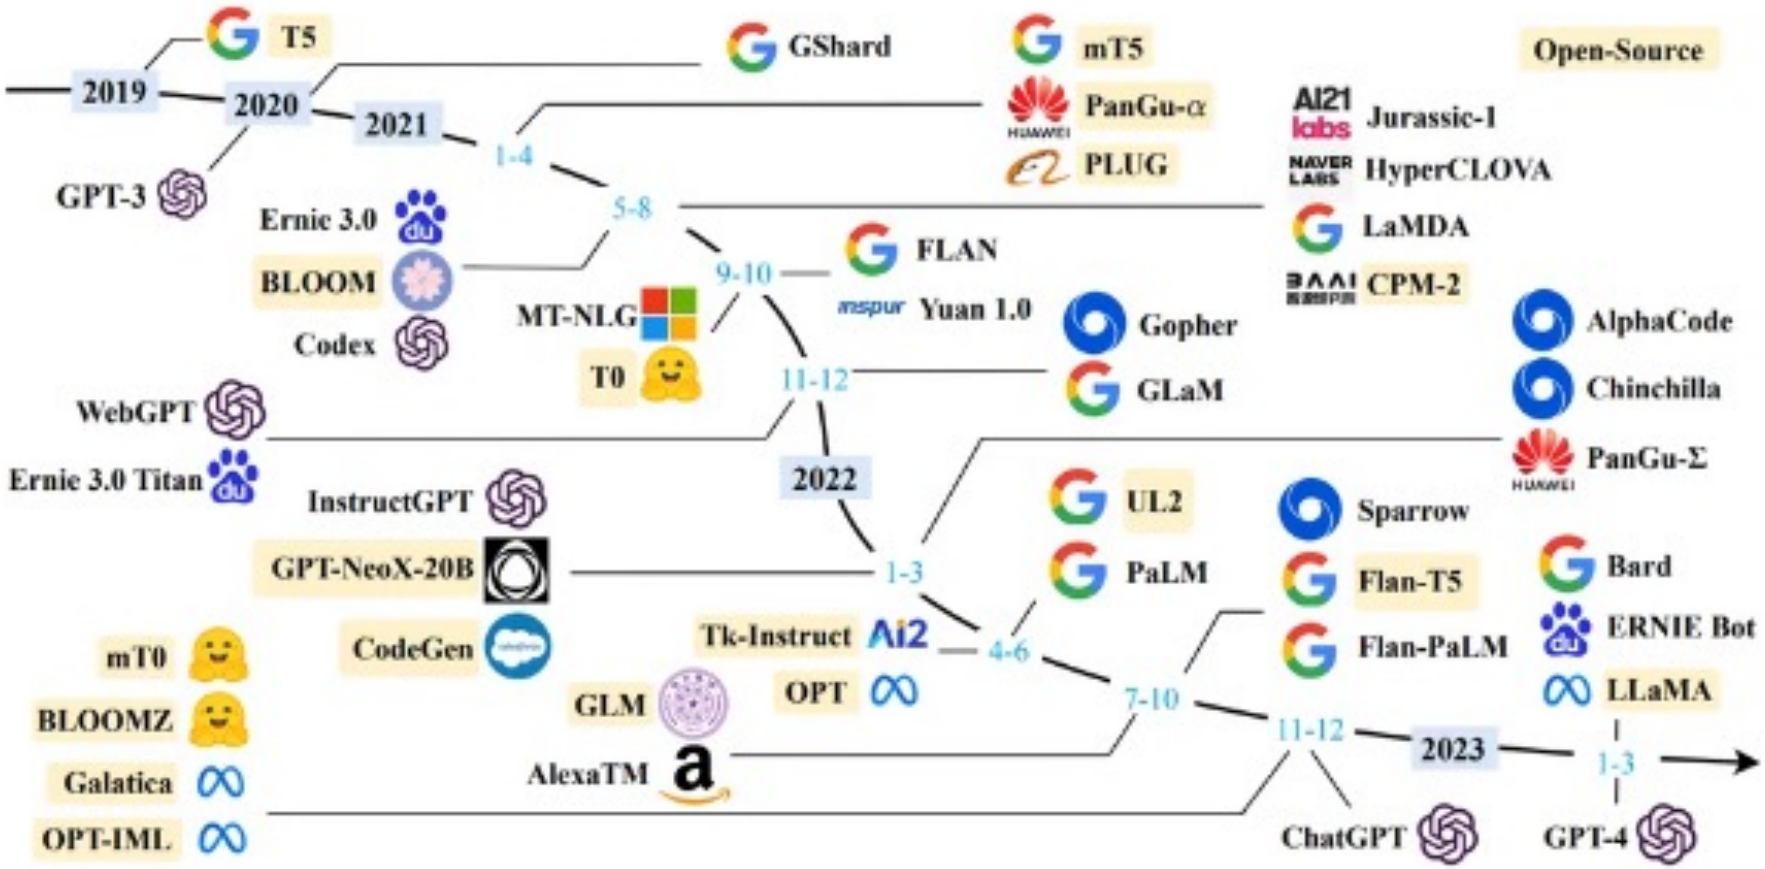

ChatGPT/GPT-4

"**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

Graph/Networks Analysis

- Social networks, World Wide Web
- Molecules

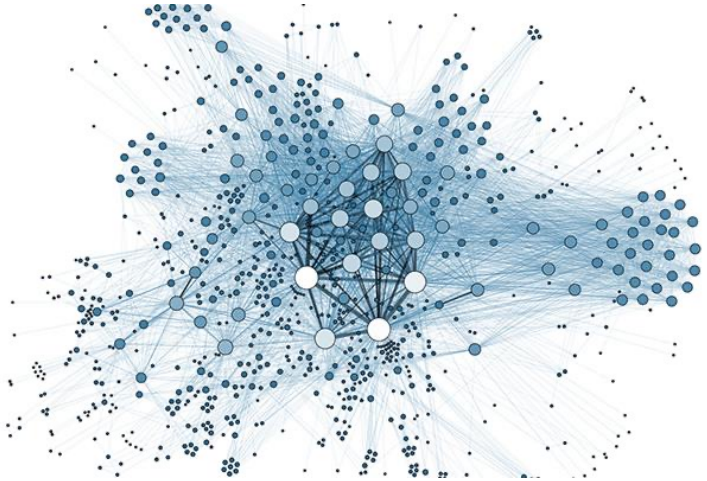

Social Networks

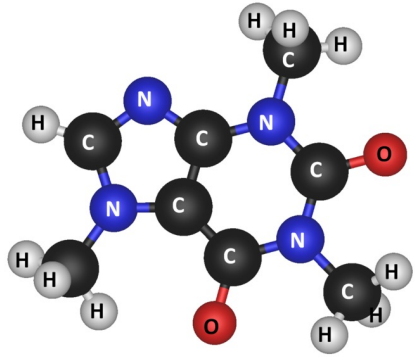

Molecules

Recommender Systems

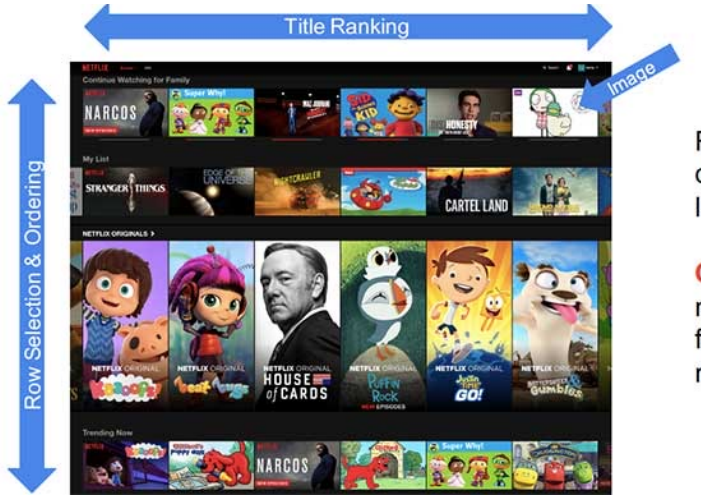

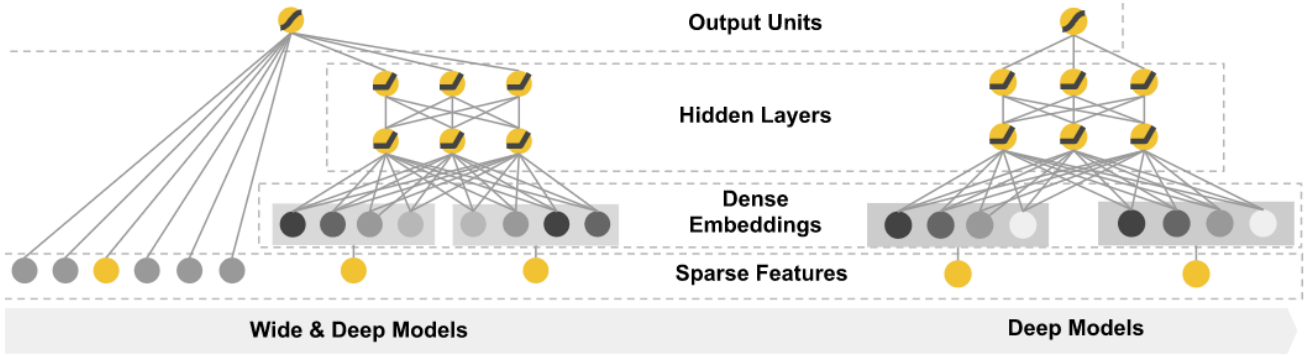

Wide & deep learning for recommender systems (Google 2016)

Workshops on Deep Learning for Recommender Systems

AI for Science

- 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)

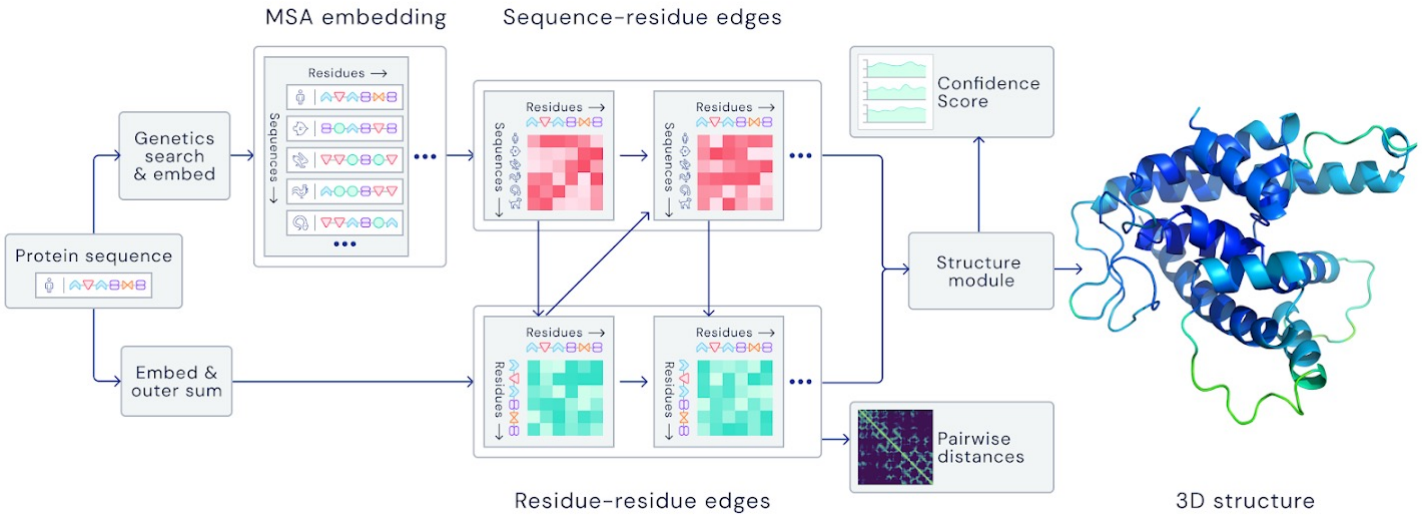

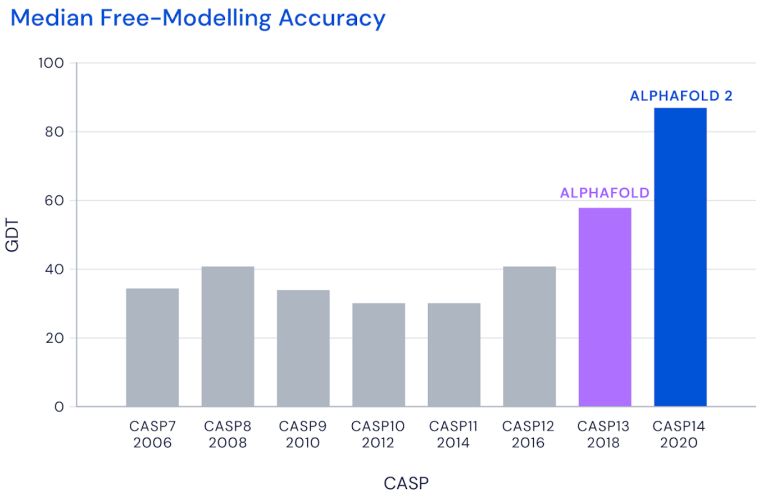

AlphaFold3

- 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,

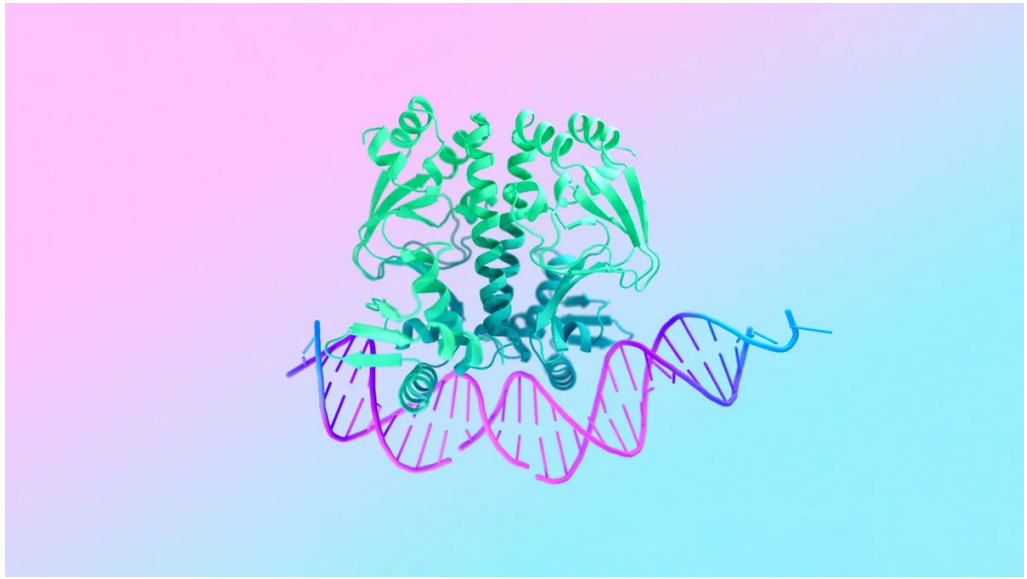

AI for Protein Design/Generation

- RFDiffusion

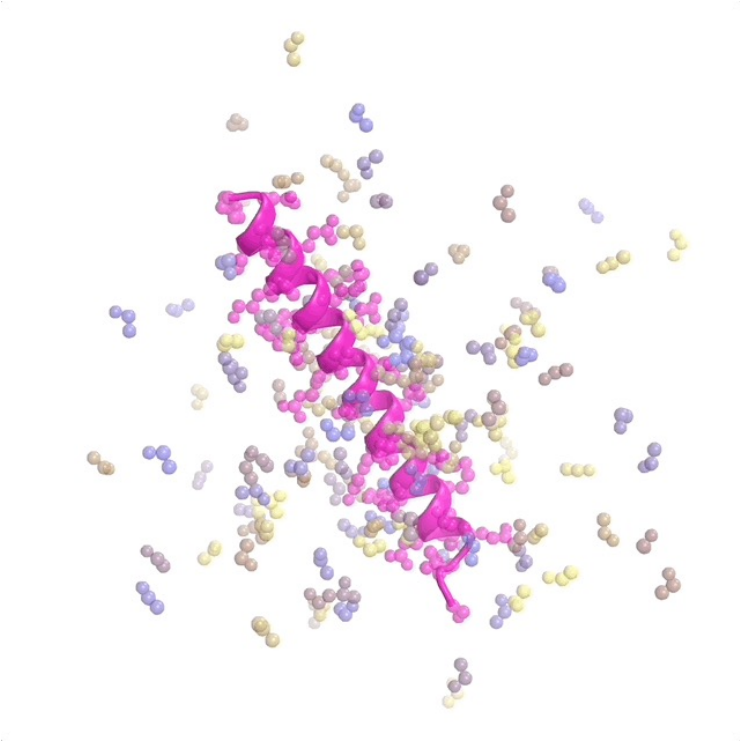

Watson et al. **De novo design of protein structure and function with Rfdiffusion**. Nature, 2023.

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

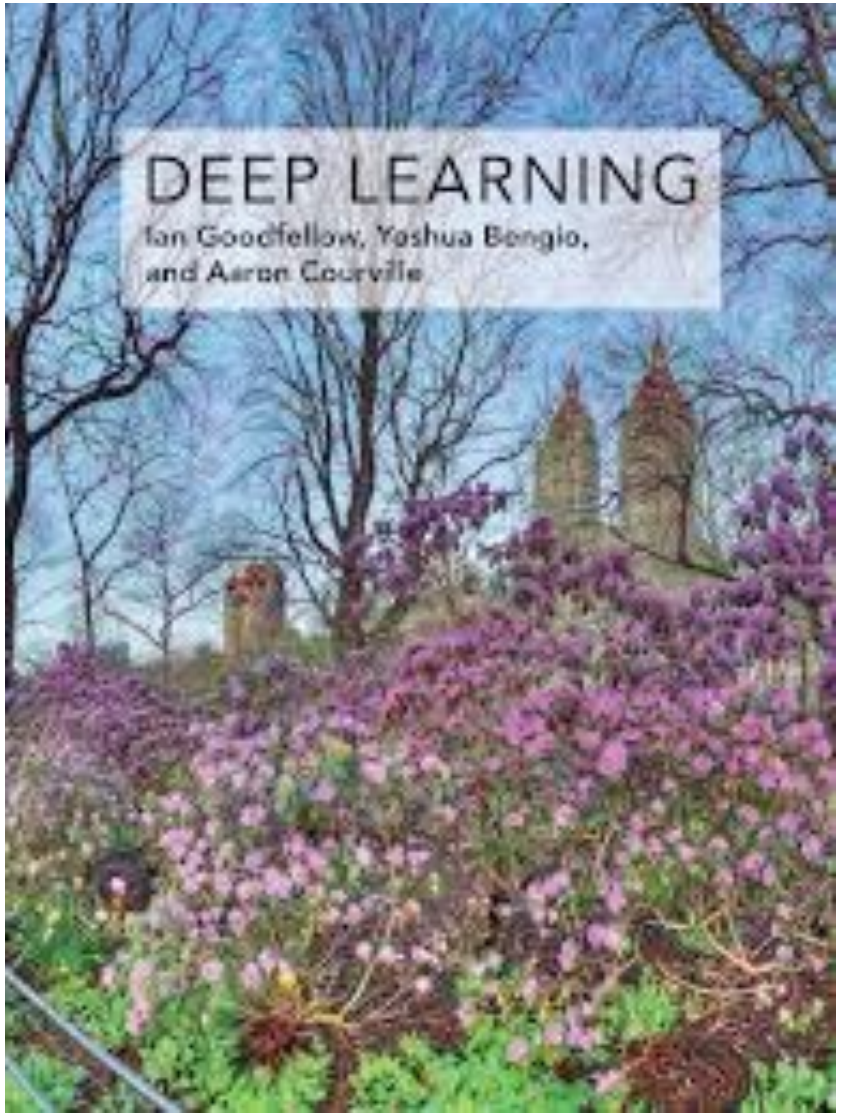

Ian Goodfellow, Yoshua Bengio and Aaron Courville. “Deep Learning”. MIT, 2016.

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/MATH60630A_2024A/index.html)
  - First come, first serve

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: Image Generation, Multi-modality Training
  - Week 12: Graph Representation Learning, Graph Neural Networks
  - Week 13: Poster Session

Questions?