# CS/ECE 760: Machine Learning
**Course Overview**

University of Wisconsin-Madison

# Outline for today

- Class logistics

- Who should take this class? prerequisites and expectations

- Overview of machine learning

# Logistics

# **Logistics:** Lectures

- Course webpage (also linked from Canvas):
- **Location**: Morgridge Hall 1524
- **Time**: Tuesdays & Thursdays 9:30AM – 10:45AM

# **Logistics**: Enrollment

- Currently at capacity, ~100 students
  - But: if you meet the prereq and are serious about take 760, come to all lectures, and EMAIL ME TODAY (please state your academic background and preparation).

# **Logistics**: Teaching Team

Professor: **Jerry Zhu**

- Office Hours: Thursdays 4-5 PM in MH 5520

TAs: **Michael Yu** **Avi Trost**

- Office Hours: TBD

# **Logistics**: Content

[Info + Schedule] **Course website**: https://pages.cs.wisc.edu/~jerryzhu/cs760.html

[Communication] **Piazza**: https://piazza.com/wisc/spring2026/cs760

- ***Please ask questions on Piazza rather than via emails*!** Sometimes your peers might be able to better answer your questions than the instructor/TAs

[Announcement] **Canvas:** https://canvas.wisc.edu/courses/486141

- ***Do not share materials on Canvas outside of class***

[Homeworks] Gradescope inside Canvas

# **Logistics**: Lecture Format

75 minutes

- You must be present!  We ask randomized quizzes
- You must ask questions!

We will post slides on the course website **before** the class

# **Logistics**: Assignments & Grades

**Class participation: 10%**

- you must attend lectures
- you must ask questions
- we ask randomized quizzes for both attendance and taking pulses on understanding

**Homeworks**:

- 5-6 homeworks, worth 20% total
- Posted after class; due before midnight on the due date (typically two weeks from the release date)
- The lowest grade will be dropped; as a result, it is very unlikely we will grant extensions.

**Exams**:

- Midterm: 30%, date TBD
- Final: 40%, 5/5/2026 12:25-2:25pm, room TBD

Details may be *subject to change*.

# **Logistics**: Drop deadlines

**January 28, 2026:** Final deadline to drop a course so that it does **not** appear on your transcript. After this date, any dropped course will show a **"DR"** notation. (HW0 deadline before this)

**April 20, 2026:** The final deadline to drop a spring course **without** needing approval from an academic dean.

# **Logistics**: Recommended reading

No required textbook, but you should read from the optional books listed on the course webpage:

Machine Learning. Tom Mitchell.

Pattern Recognition and Machine Learning. Chris Bishop.

Understanding Machine Learning: From Theory to Algorithms. Shalev-Shwartz, Ben-David.

Deep Learning. by Goodfellow, Bengio, Courville.

Reinforcement Learning: An Introduction. Sutton, Barto.

# This week's reading assignment

Article by Jordan and Mitchell on course website

REVIEW

## Machine learning: Trends, perspectives, and prospects

M. I. Jordan[1*] and T. M. Mitchell[2*]

Machine learning addresses the question of how to build computers that improve automatically through experience. It is one of today's most rapidly growing technical fields, lying at the intersection of computer science and statistics, and at the core of artificial intelligence and data science. Recent progress in machine learning has been driven both by the development of new learning algorithms and theory and by the ongoing explosion in the availability of online data and low-cost computation. The adoption of data-intensive machine-learning methods can be found throughout science, technology and commerce, leading to more evidence-based decision-making across many walks of life, including health care, manufacturing, education, financial modeling, policing, and marketing.

# Who should take this class?

# **Class Setup**: Goals

Two goals:

- **Understanding** ML
- **Foundation** for future research in ML

(CS760 will lay the foundations of several topics in ML, but will likely not be sufficient on its own to advance a topic.)

This is a technical / mathematical course. If you prefer an “AI appreciation” course, consider CS540 instead.

# Class Goals

Mini-goals:

- **Intuition for each algorithm/model**
- **Big picture/ML ecosystem**

**Examples**:

- When to use what type of ML?
- How hard is it to train?
- What generalizes best?
- Where is the field going?

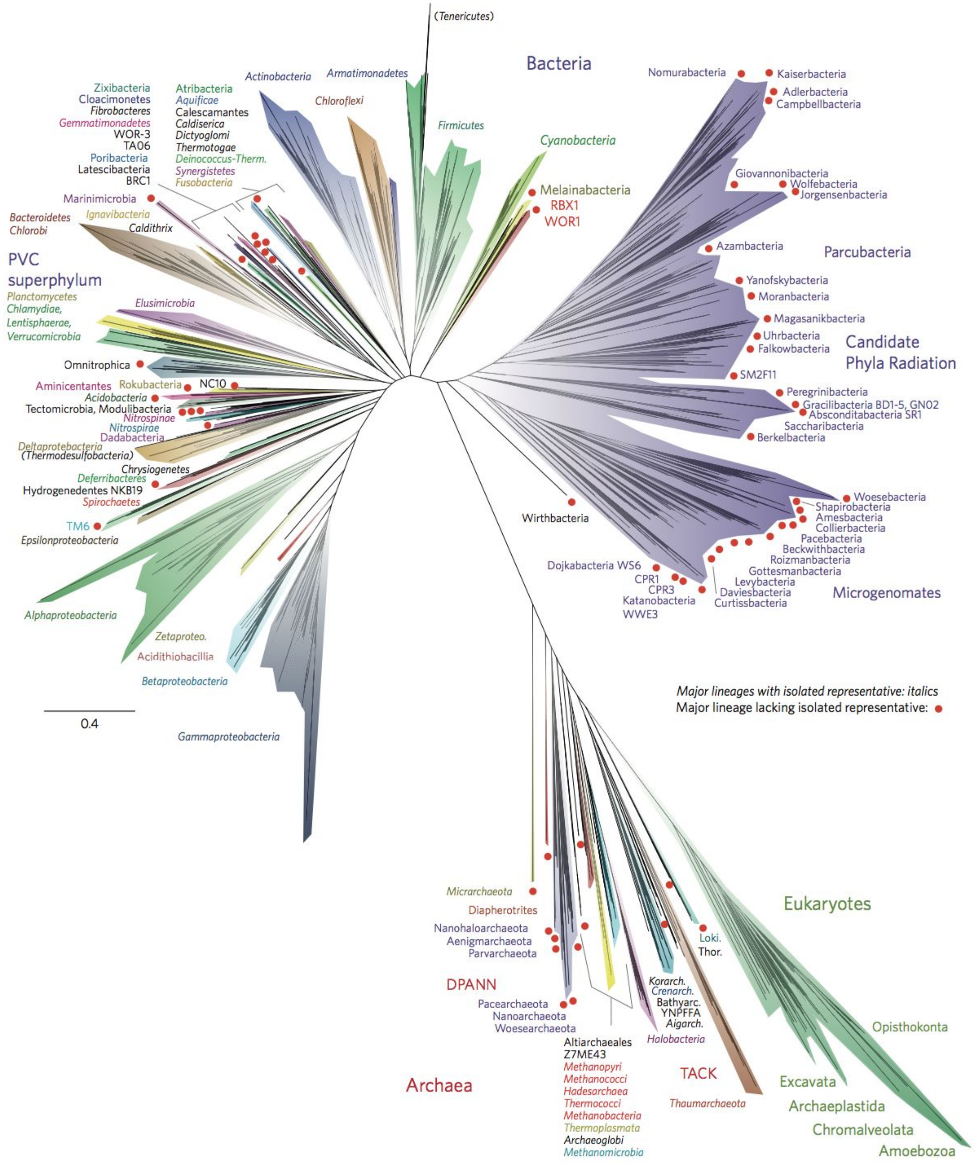

# Required Background

You are expected have (at least) a working understanding of:

- **Linear algebra** (working with data, linear transformations)
- **Calculus** (for optimization, convergence, etc.)
- **Probability & statistics** (dealing with noise, sampling)
- **Programming** (for implementation, mostly python)

Plenty of resources available online

- Just need enough experience/mathematical maturity to pick up missing bits
- Use AI as your assistant, just don’t let it do your homework

# Resources

**Probability**

- Lecture notes: http://www.cs.cmu.edu/~aarti/Class/10701/recitation/prob_review.pdf

**Linear Algebra:**

- Short video lectures by Prof. Zico Kolter: http://www.cs.cmu.edu/~zkolter/course/linalg/outline.html
- Handout associated with above video: http://www.cs.cmu.edu/~zkolter/course/linalg/linalg_notes.pdf
- Book: Gilbert Strang. Linear Algebra and its Applications. HBJ Publishers.

**Big-O notation:**

- http://www.stat.cmu.edu/~cshalizi/uADA/13/lectures/app-b.pdf
- http://www.cs.cmu.edu/~avrim/451f13/recitation/rec0828.pdf

**Using AI for self-study is encouraged (just do your homework “for real”)**

# **Assignment**: Homework

**For HW0,** self-diagnostic on background. Topics:

- Linear Algebra
- Calculus
- Probability
- Big-O notation
- Basic programming skills

**HW0 is available now and due on 1/27 at midnight (one week)**

If these feel very unfamiliar, consider taking relevant courses first and then take CS760 in the future.

# Programming background

**We expect you to be able to**

- Implement simple routines/logic in Python (for/while loops, if/else, break conditions)
  - Familiarity with NumPy would be a plus
- Write simple shell scripts in Linux/Unix
- Install and use ML packages (e.g. scikit-learn, PyTorch)

- Generally, we will **not** help you with these during OHs!
- Usually, you can resolve such issues on Piazza and by asking AI.

# Overview of machine learning

# **ML Overview**: Definition

What is machine learning?

"A computer program is said to learn from experience **E** with respect to some class of tasks **T** and performance measure **P**, if its performance at tasks in **T** as measured by **P**, improves with experience **E**." *Machine Learning*, Tom Mitchell, 1997

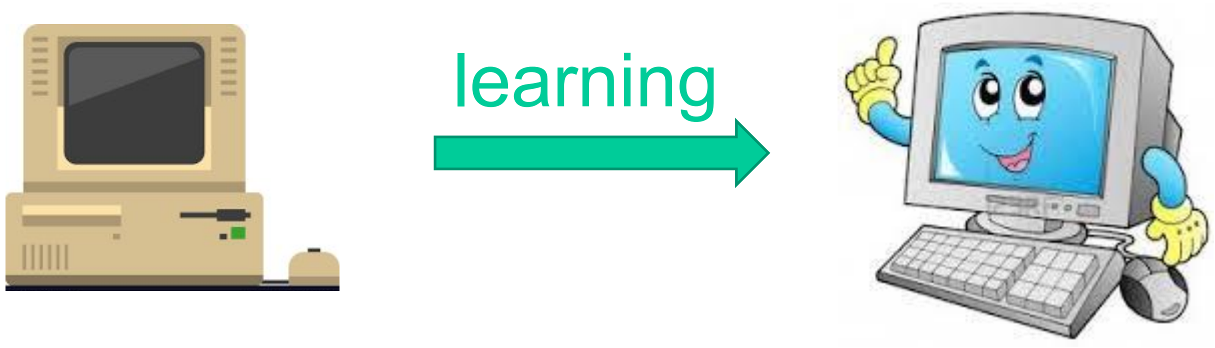

# **ML Overview**: Motivation

Why would we do this?

- We're building a self-driving car. Could just write down rules
  - **Painful**! A lot of cases…
  - **Learn from examples** instead

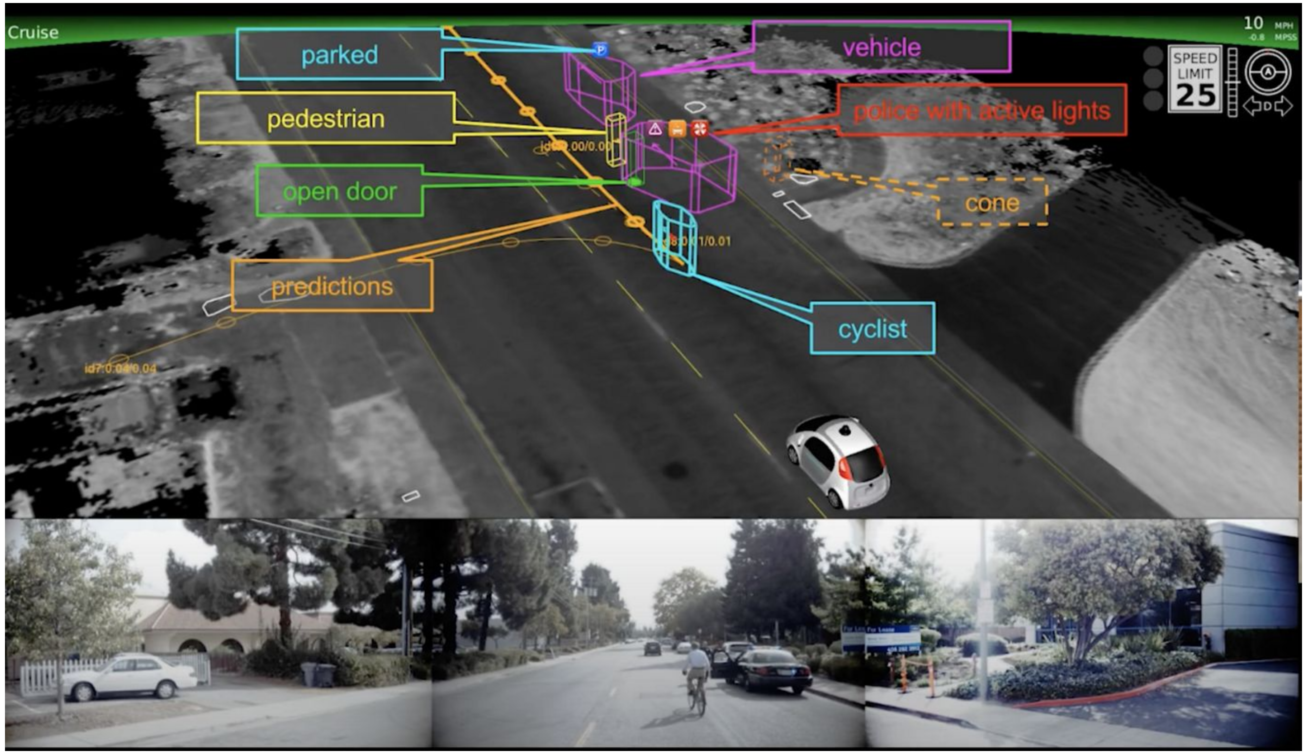


Waymo

# **ML Overview**: Flavors

## Supervised Learning

- Learning from examples ***with* “answers”**
- **Workflow**:
  - Collect a set of examples {data, labels}: **training set**
  - “**Train**” a model to match these examples
  - ”**Test**” it on new data
- **Image classification**:

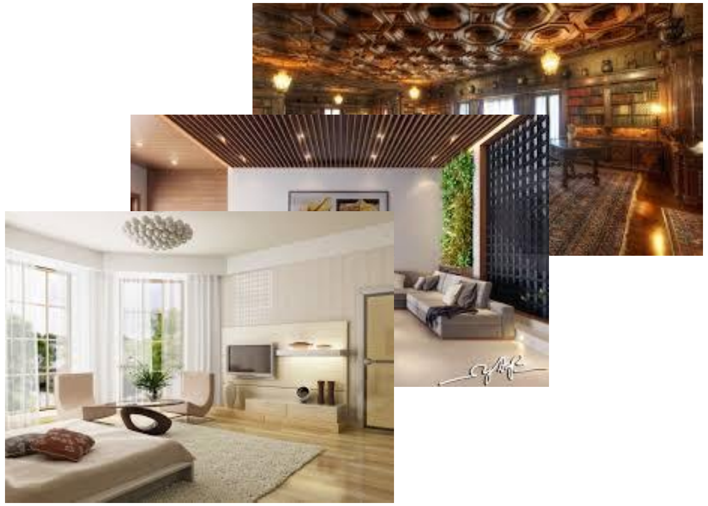

**indoor**

**outdoor**

# **ML Overview**: Flavors

## Supervised Learning

- **Example: Image classification**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: distinguish **indoor** vs **outdoor**
  - **P**erformance measure: probability of misclassifying
  - **E**xperience: labeled examples

**indoor**

**outdoor**

# **ML Overview**: Flavors

## Supervised Learning

- **Example: Spam Filtering**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: distinguish **spam** vs **legitimate**
  - **P**erformance measure: probability of misclassifying
  - **E**xperience: labeled examples of messages/emails

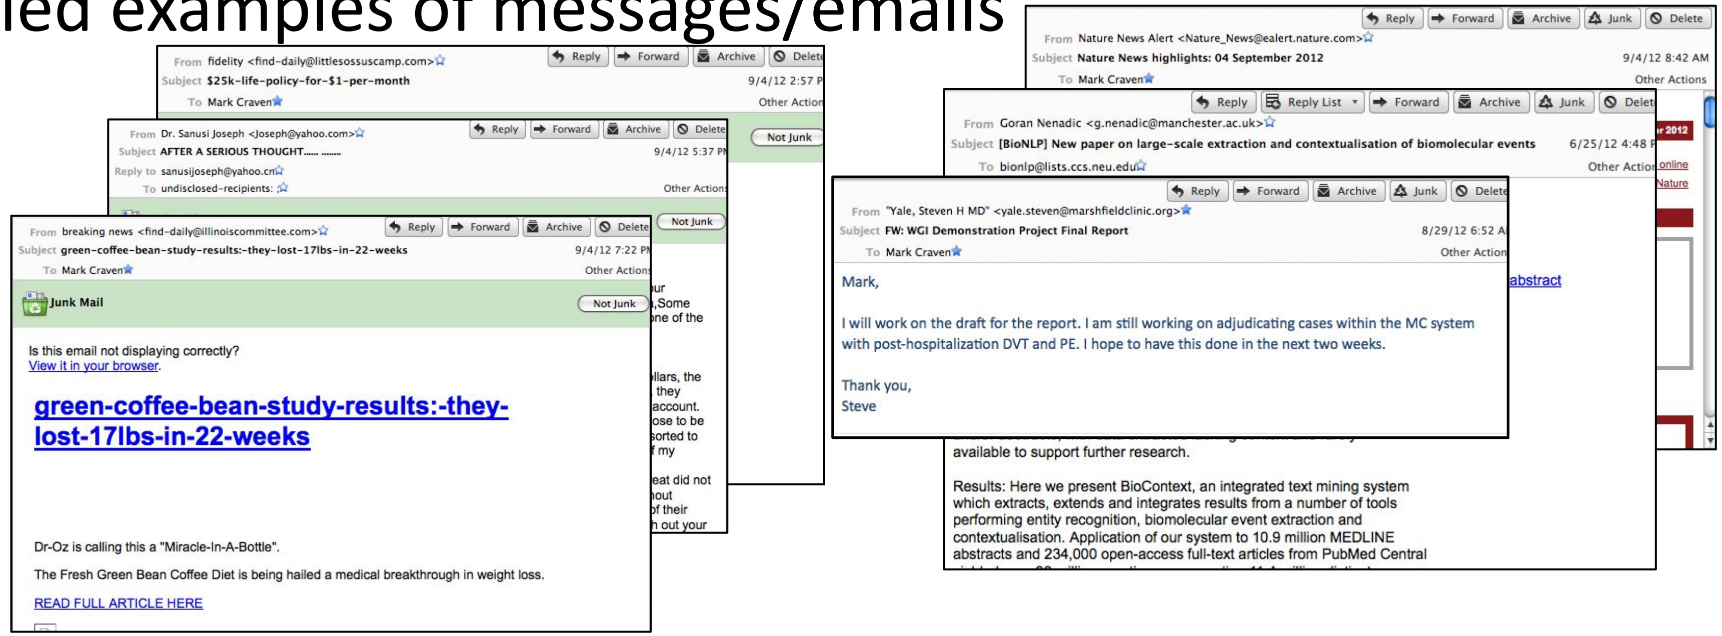

# **ML Overview**: Flavors

## Supervised Learning

- **Example: Ratings/Recommendations**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: predict how much a user will like a film
  - **P**erformance measure: difference between prediction and user's true rating
  - **E**xperience: previous ratings

# High profile success of supervised learning: **Machine translation**

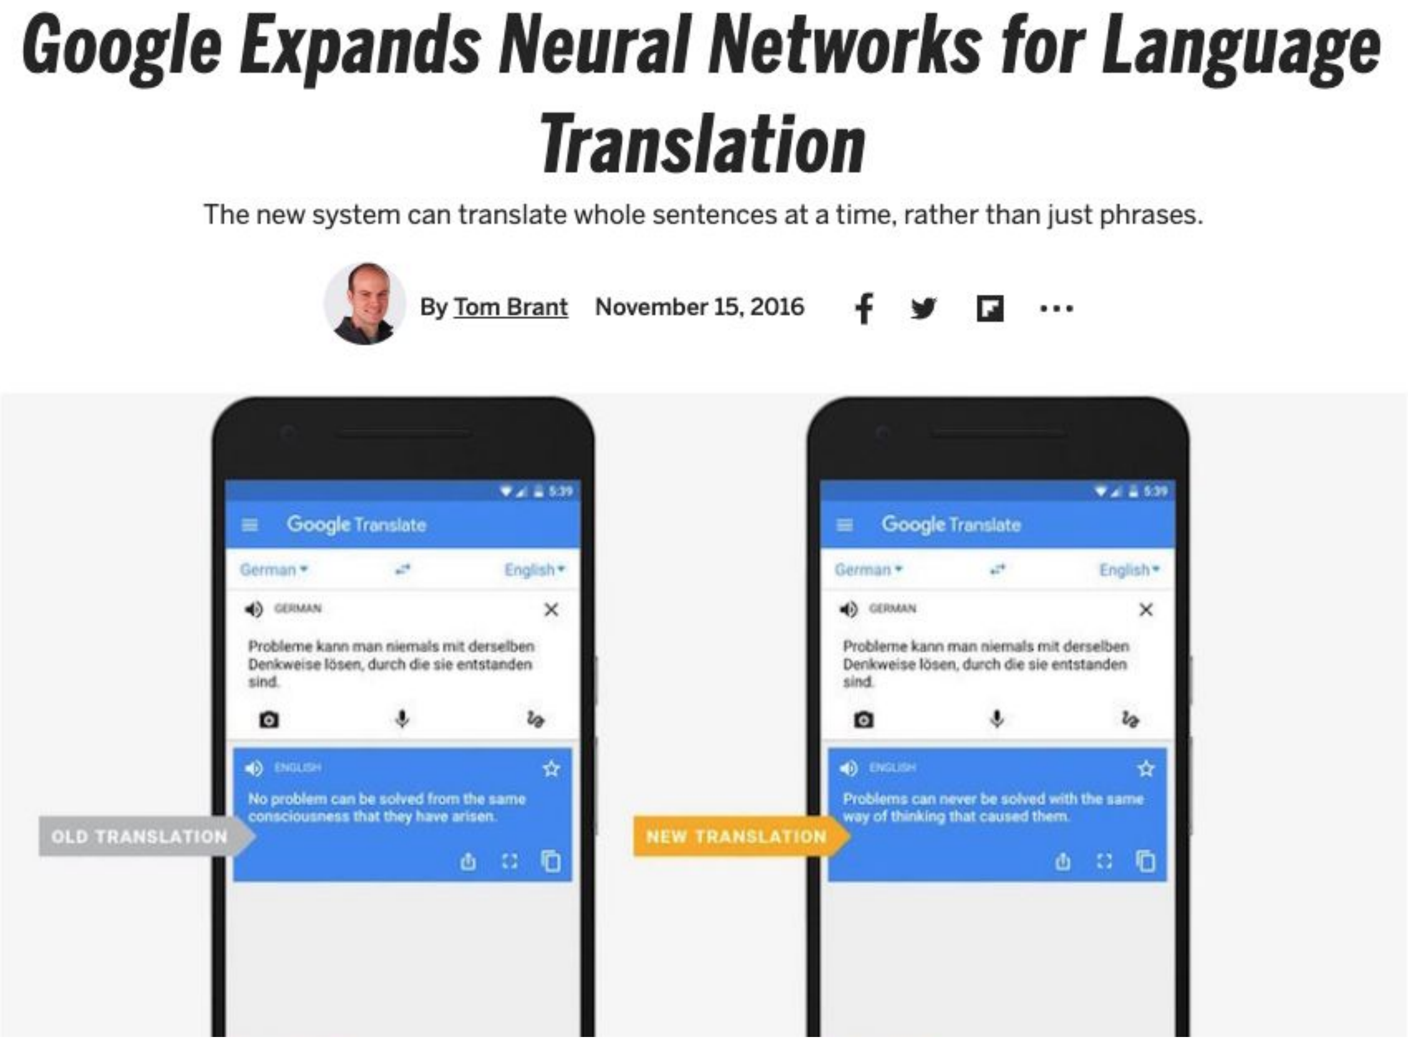

## Google Expands Neural Networks for Language Translation

The new system can translate whole sentences at a time, rather than just phrases.

By Tom Brant November 15, 2016

# High profile success of supervised learning: **Protein-folding**

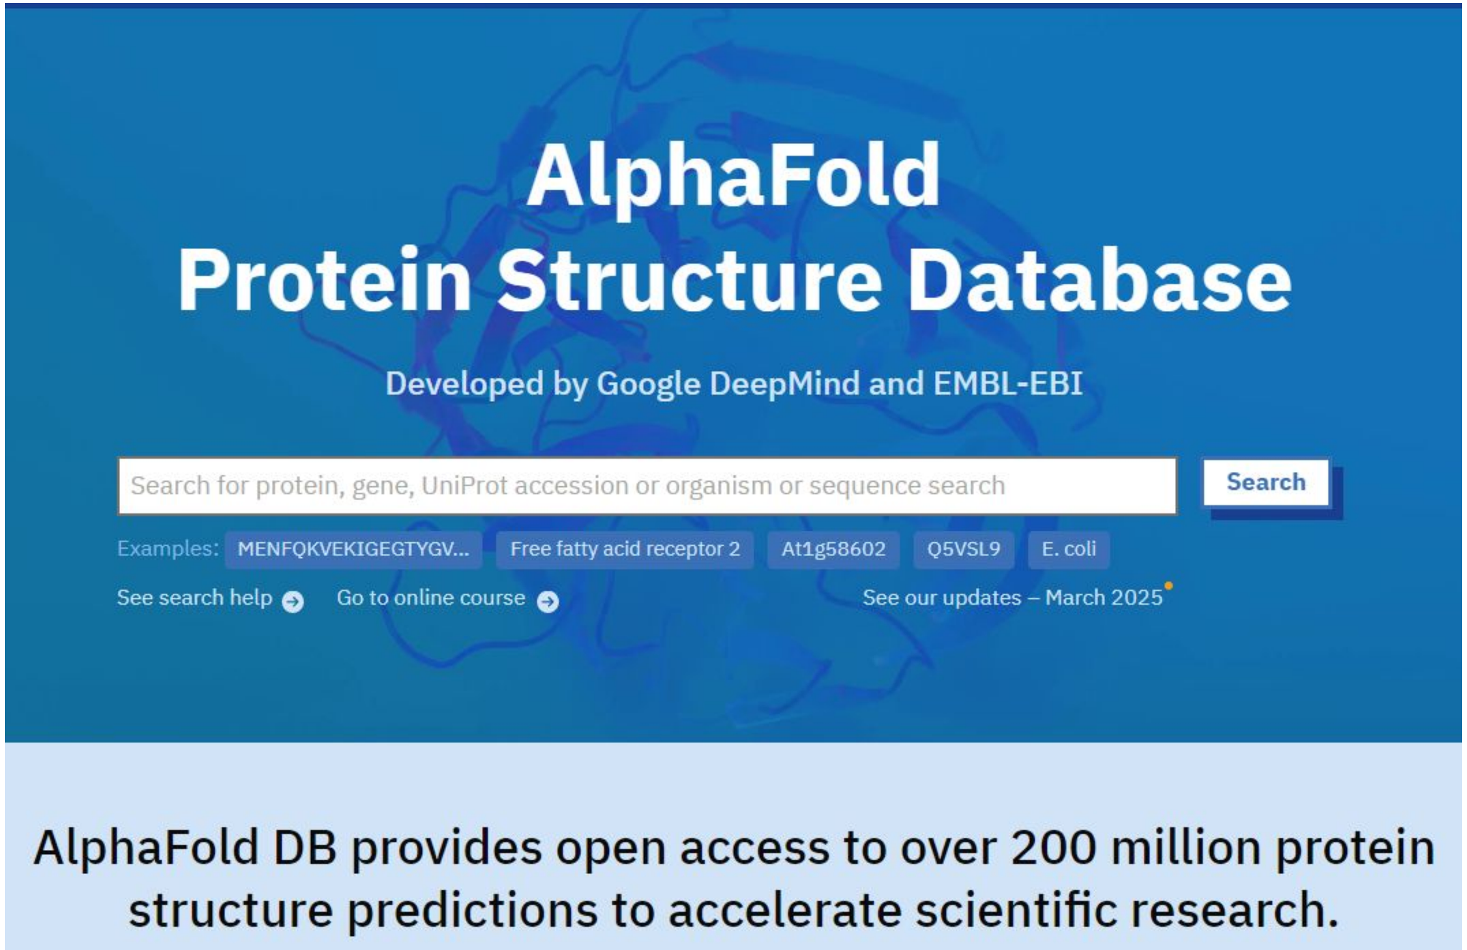

# **ML Overview**: Flavors

## Unsupervised Learning

- Learning from examples ***without* "answers"**
- Goal: "find something": structure, hidden information, etc
- **Workflow**:
  - Collect a set {data}
  - Perform some algorithm on it and draw insights about data
  - Sometimes: test on new data

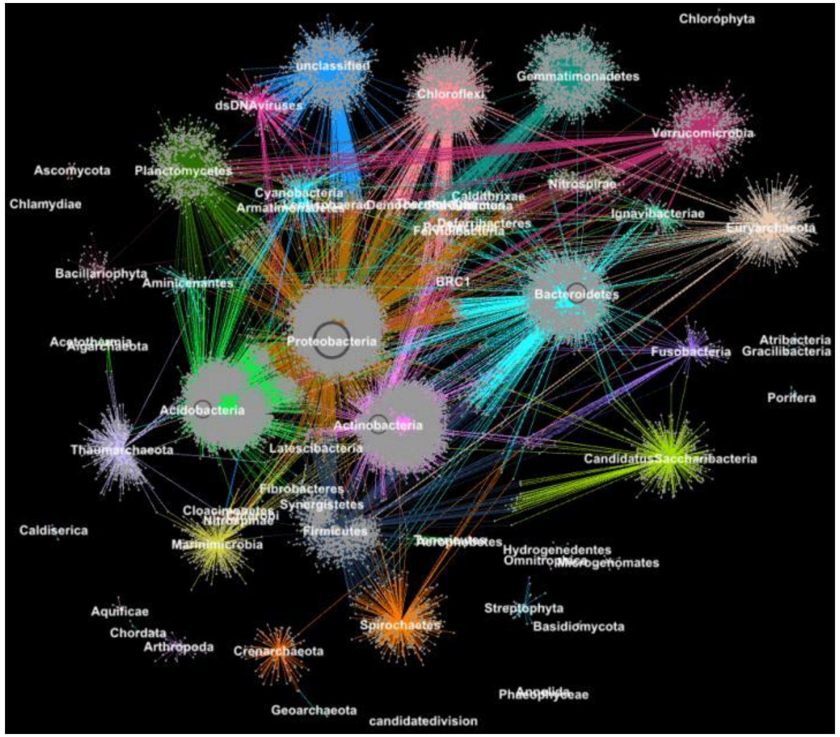

# **ML Overview**: Flavors

## Unsupervised Learning

- **Example: Clustering**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: produce distinct clusters for a set of data
  - **P**erformance measure: closeness to underlying structure
  - **E**xperience: available datapoints

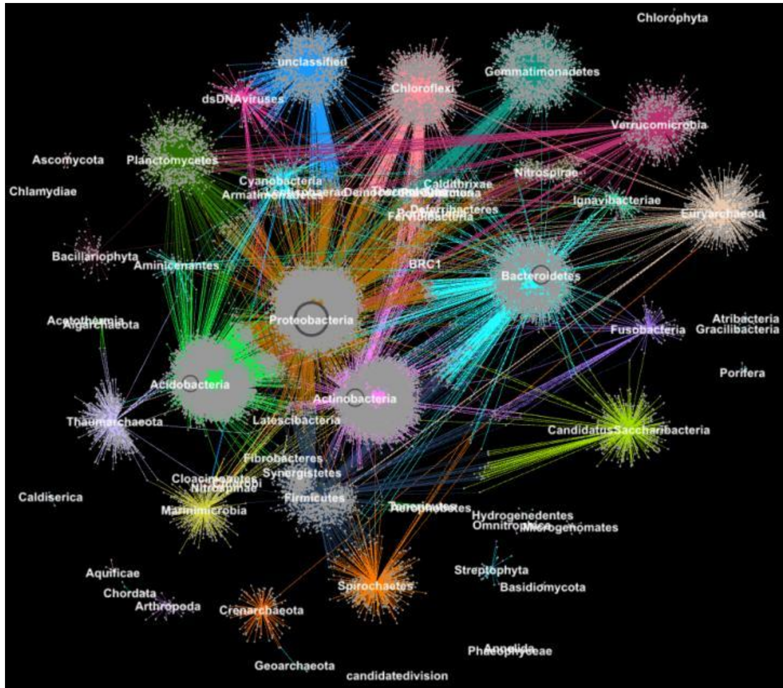

# **ML Overview**: Flavors

## Unsupervised Learning

- **Example: Generative Models (image)**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: produce artificial images of faces
  - **P**erformance measure: photorealism
  - **E**xperience: available images

**StyleGAN2** (Kerras et al ’20)

# **ML Overview**: Flavors

## Unsupervised Learning

- **Example: Language Models**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: next token (word) prediction
  - **P**erformance measure: perplexity (uncertainty or “confusion” of predicting the next word in a sequence)
  - **E**xperience: large and diverse text datasets

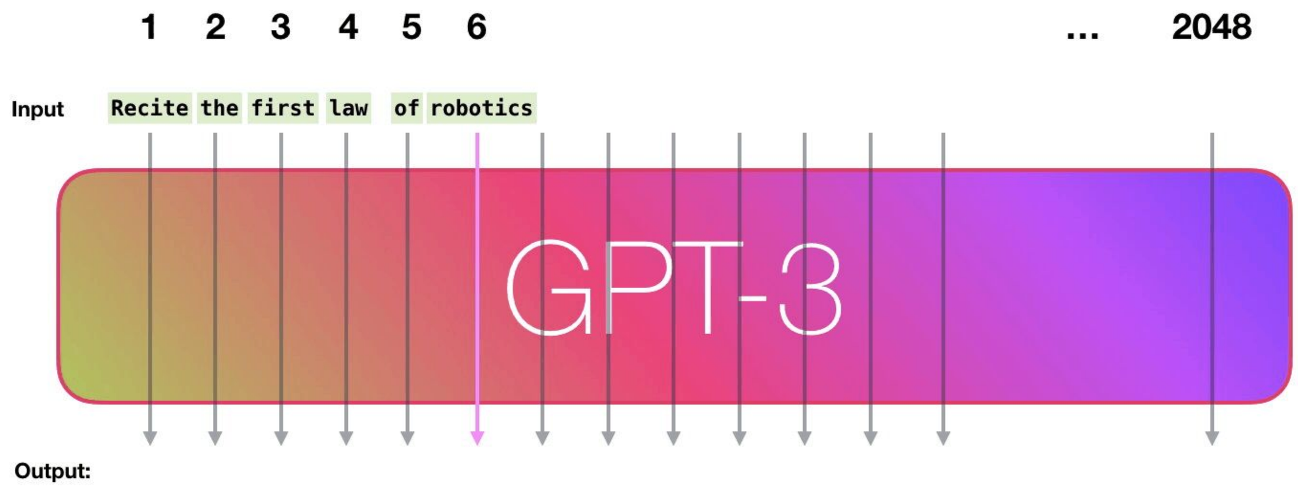

# High-profile success of unsupervised learning: **Contextual image generation**

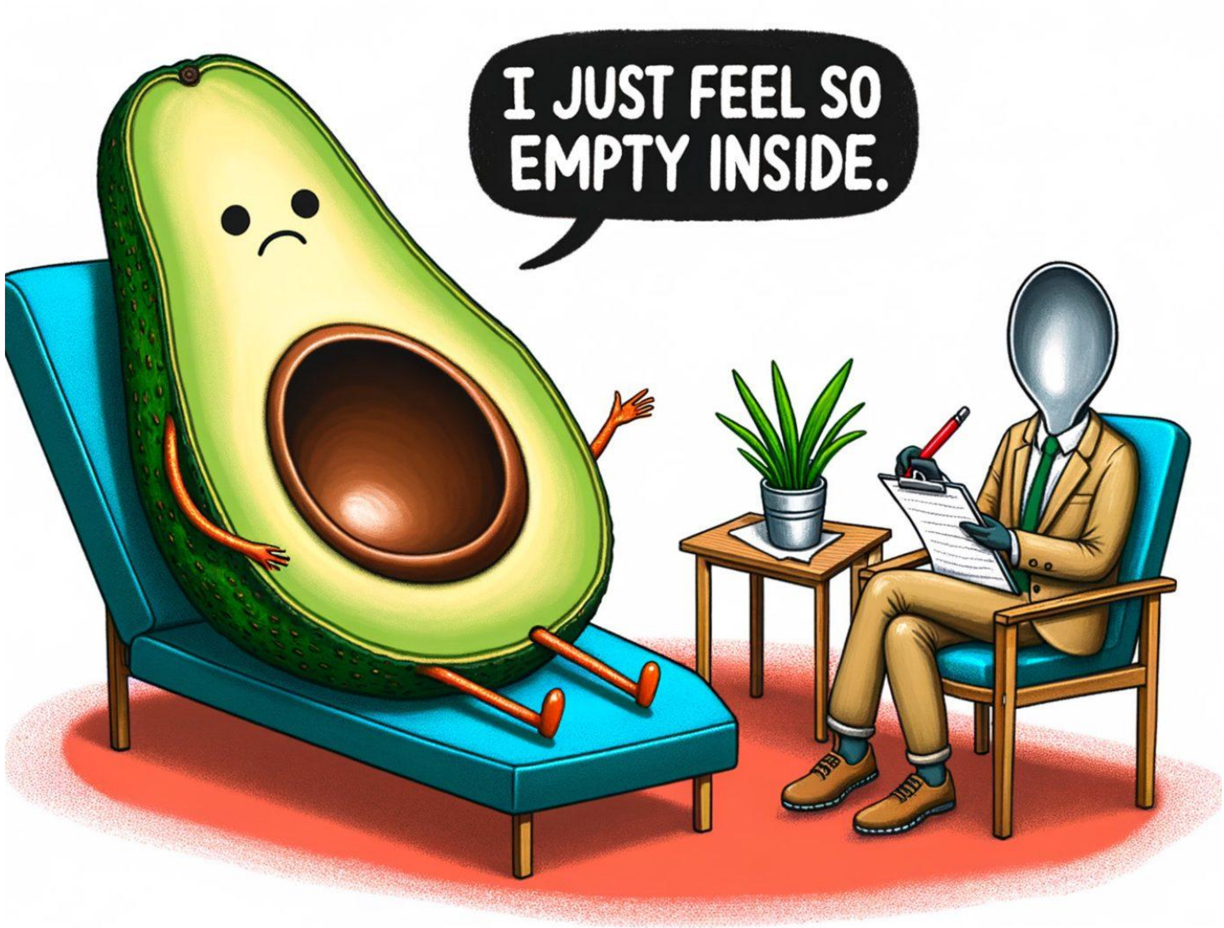


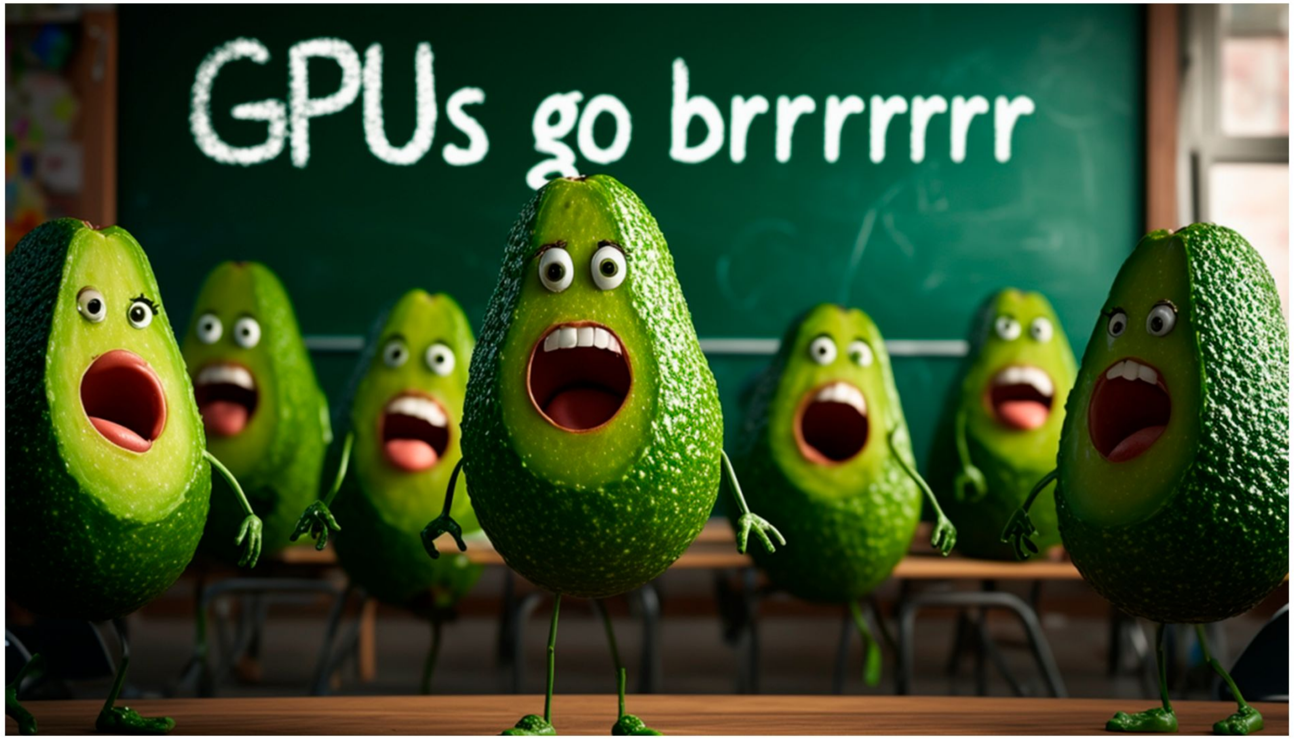


Prompt: A surreal and humorous scene in a classroom with the words 'GPUs go brrrrrr' written in white chalk on a blackboard. In front of the blackboard, a group of students are celebrating. These students are uniquely depicted as avocados, complete with little arms and legs, and faces showing expressions of joy and excitement. The scene captures a playful and imaginative atmosphere, blending the concept of a traditional classroom with the whimsical portrayal of avocado students.

# High-profile success of unsupervised learning: **Large language models (ChatGPT, Claude, ...)**

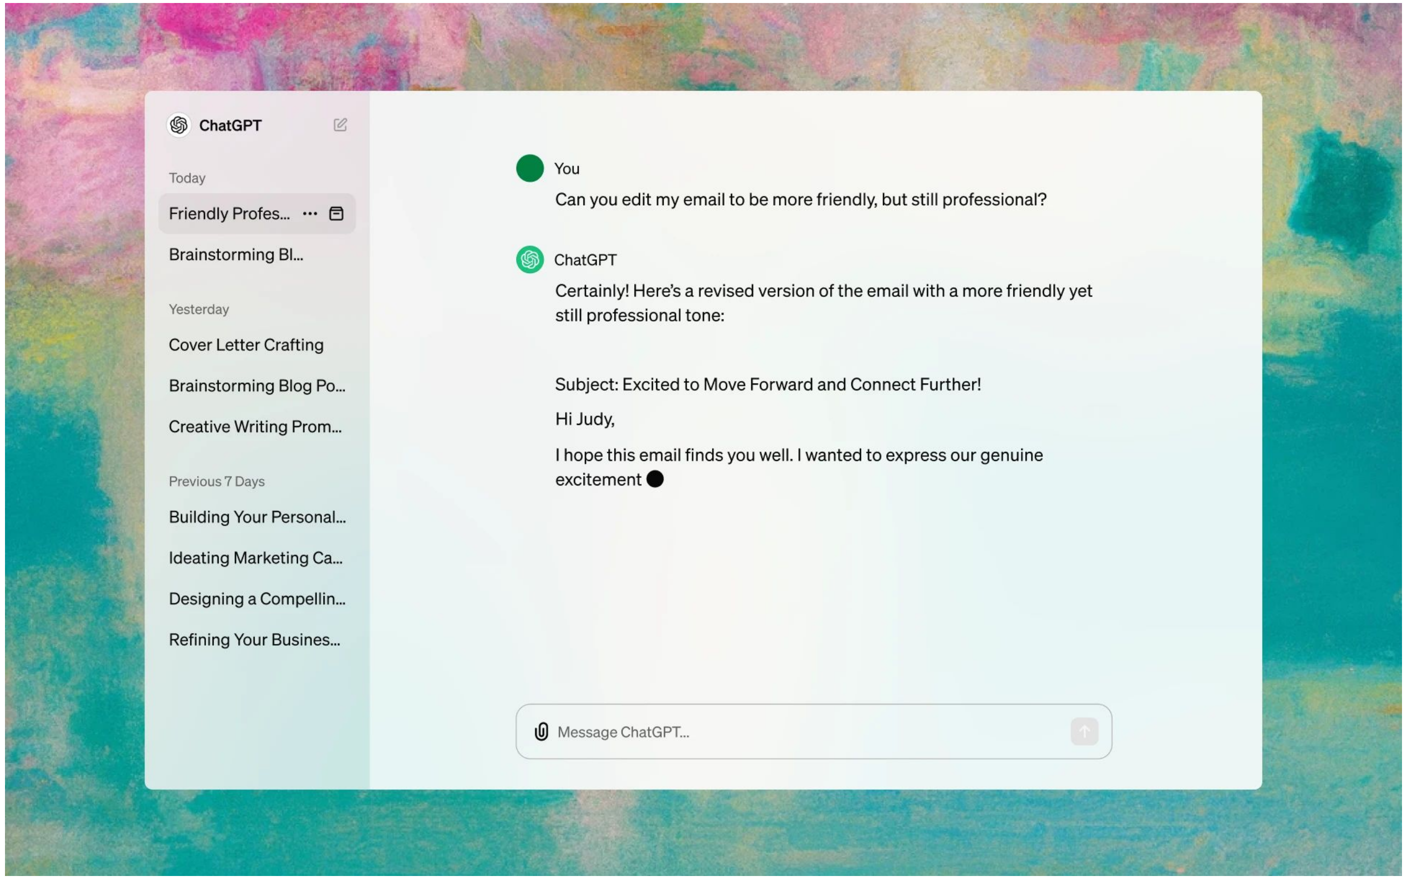

# **ML Overview**: Flavors

## Reinforcement Learning

- Agent interacting with the world; gets rewards for actions
- Goal: learn to perform some activity with high reward
- **Workflow**:
  - Create an environment, reward, agent
  - **Train**: train policy to maximize rewards
  - **Deploy** in new environment

# **ML Overview**: Flavors

## Reinforcement Learning

- **Example: Controlling aircraft**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: keep the aircraft in the air, steer towards a desired goal
  - **P**erformance measure: reward for reaching goal quickly
  - **E**xperience: data (state/action/reward) from previous flights

# **ML Overview**: Flavors

## Reinforcement Learning

- **Example: Playing video games**
- Recall **T**ask/**P**erformance measure/**E**xperience definition
  - **T**ask: play Atari arcade games
  - **P**erformance measure: winning/advancing
  - **E**xperience: state/action/reward from previous gameplay episodes

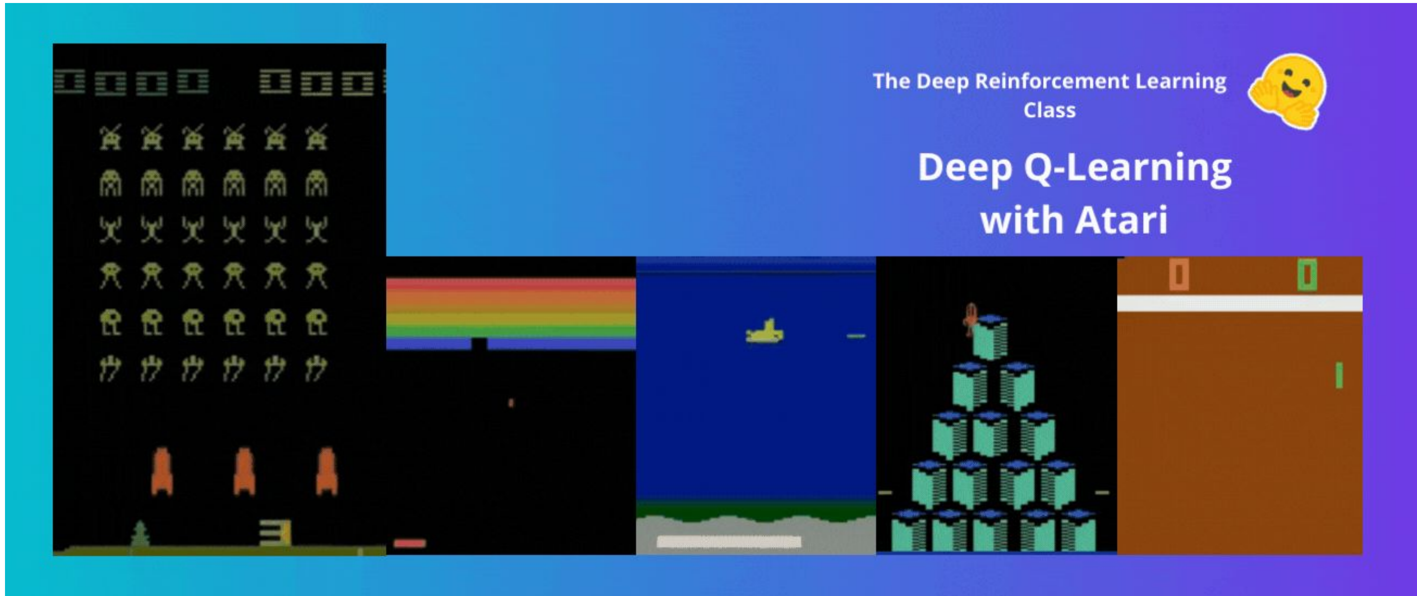

# High-profile success of reinforcement learning: **Game-playing (Go, Starcraft, …)**

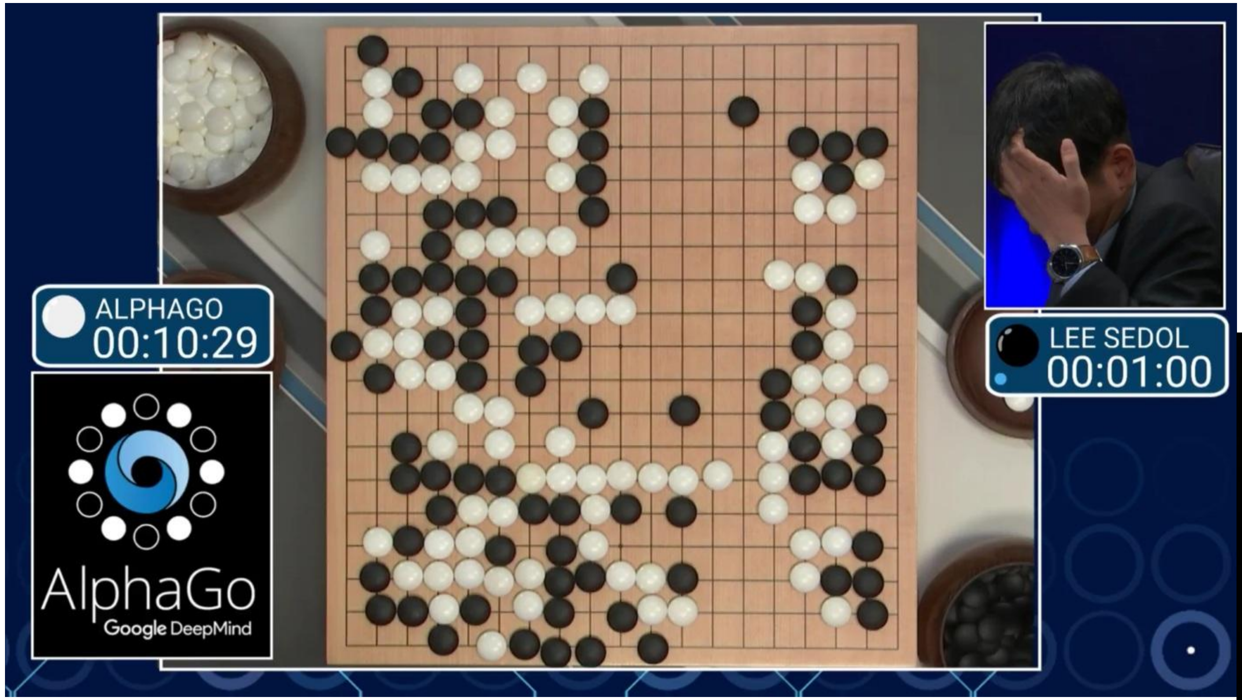


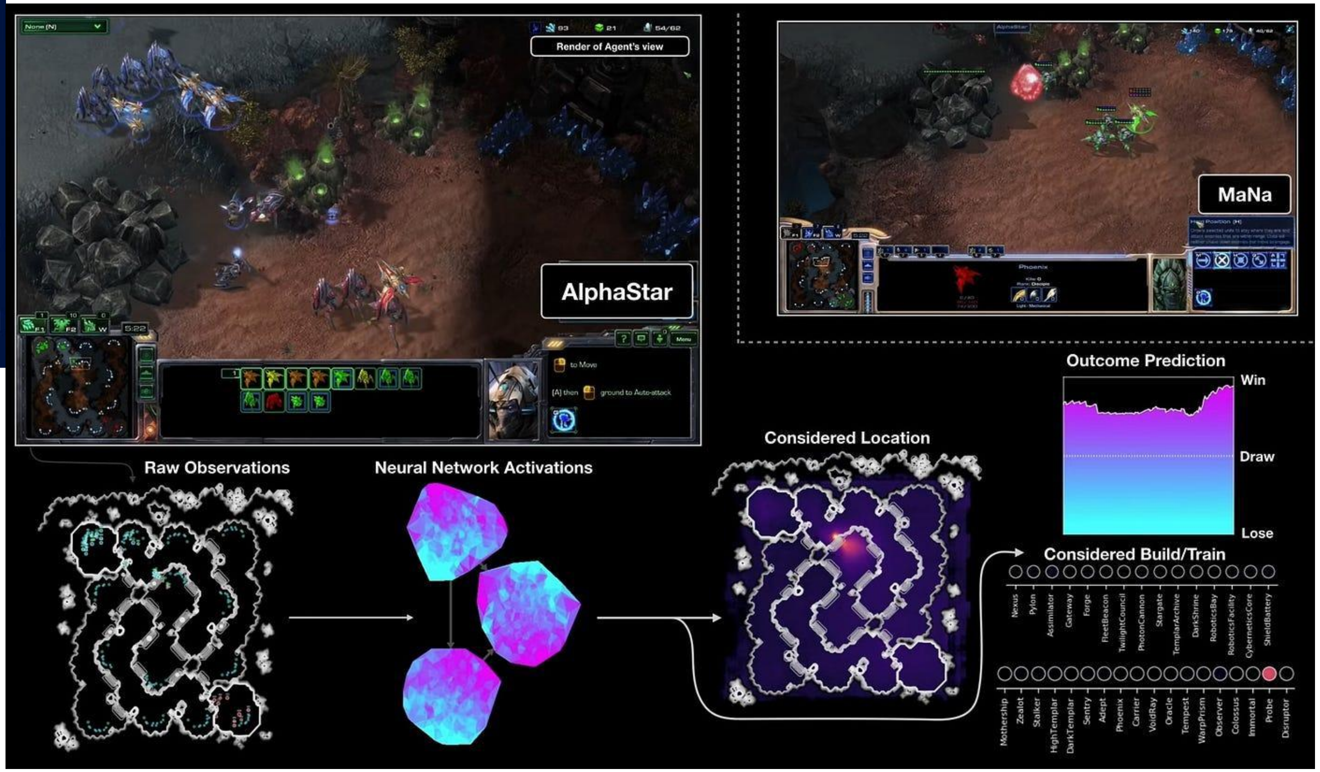

# High-profile success of reinforcement learning: **Advanced competition mathematics**

Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the International Mathematical Olympiad

21 JULY 2025

Thang Luong and Edward Lockhart

Share

# Machine learning used in many other fields

- Other areas of computer science
  - distributed systems, computer architecture, databases, …
- Electrical engineering
- Industrial engineering
- Physics
- Materials science
- Drug discovery
- Finance & economics
- E-commerce
- ….

# Thanks Everyone!

Some of the slides in these lectures have been adapted/borrowed from materials developed by Misha Khodak, Mark Craven, David Page, Jude Shavlik, Tom Mitchell, Nina Balcan, Elad Hazan, Tom Dietterich, Pedro Domingos, Jerry Zhu, Yingyu Liang, Volodymyr Kuleshov, Fred Sala, Kirthi Kandasamy, Josiah Hanna, Tengyang Xie