6781: Theoretical Foundations of Machine Learning

Lecture 1: Introduction

Optional Readings: Understanding ML, Chapter 1.

Instructor: Nika Haghtalab

Even your coffee is surprised that you got up this early!

You do what you gotta do to stay awake in an 8am class!

Outline of Today

Who are we:

- Instructor: Nika Haghtalab
- Teaching Assistants:

Abhishek Shetty

Seung Won (Wilson) Yoo

- What is machine learning?
- Syllabus and logistics

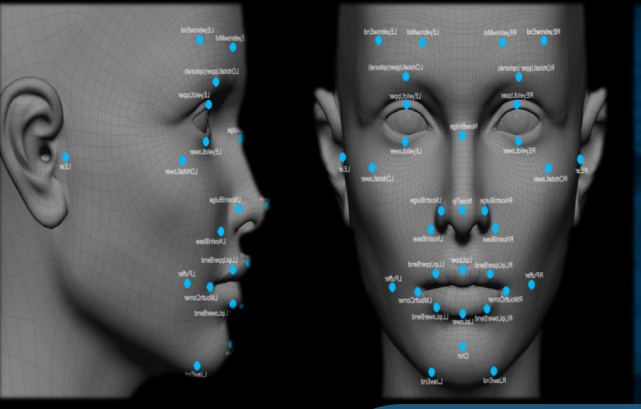

Machine Learning (ML)

Programs that improve with experience

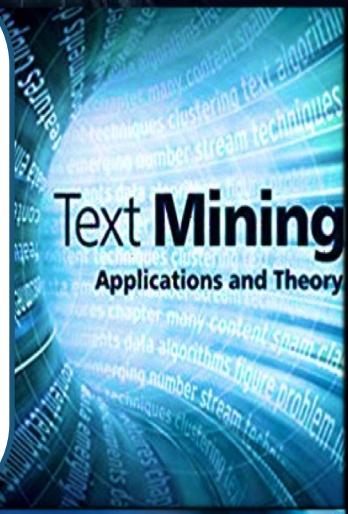

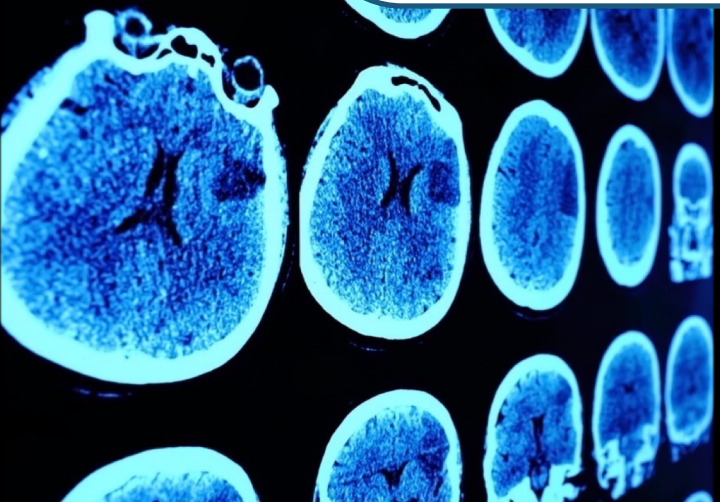

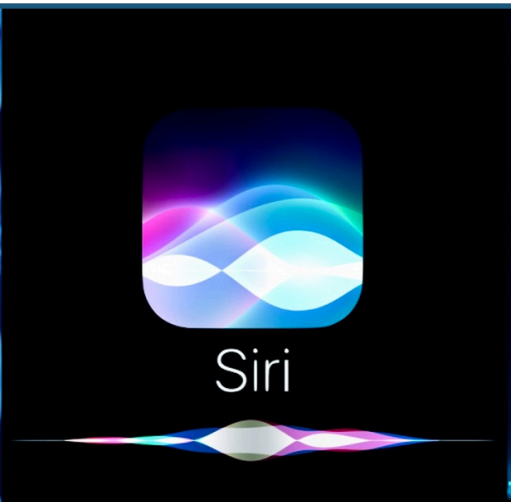

Revolutionizing Science and Technology

"**A breakthrough in machine learning would be worth ten Microsofts.**" (Bill Gates, Microsoft)

"It will be the basis and fundamentals of every successful huge IPO win in 5 years." (Eric Schmidt, Google / Alphabet

"AI and machine learning are going to change the world and we really have not begun to scratch the surface." (Jennifer Chayes, Berkeley)

"ML is transforming sector after sector of the economy, and the rate of progress only seems to be accelerating." (Daphne Koller, Stanford / Coursera/ Insitro)

"**Machine learning is the next Internet**" (Tony Tether, DARPA)

What is Machine Learning?

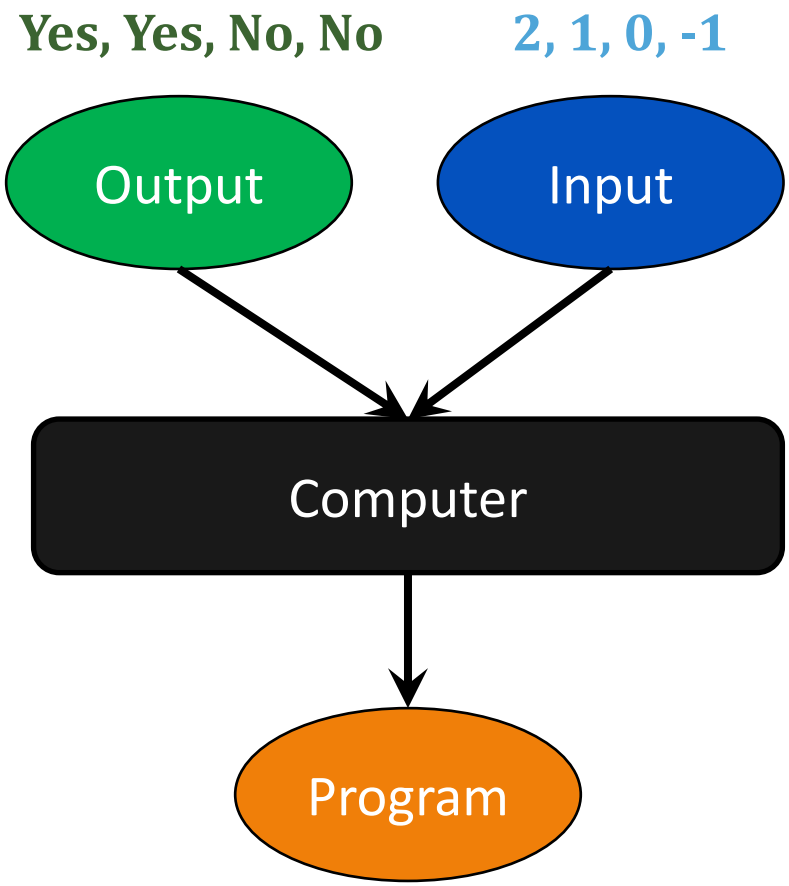

Machine Learning

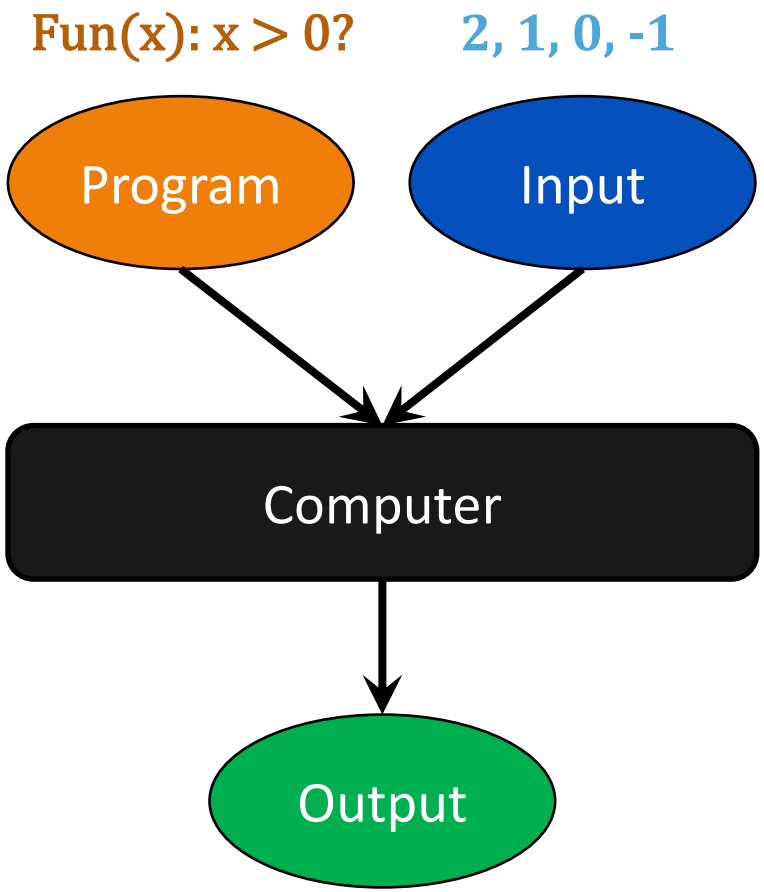

Traditional Computing

What is Machine Learning?

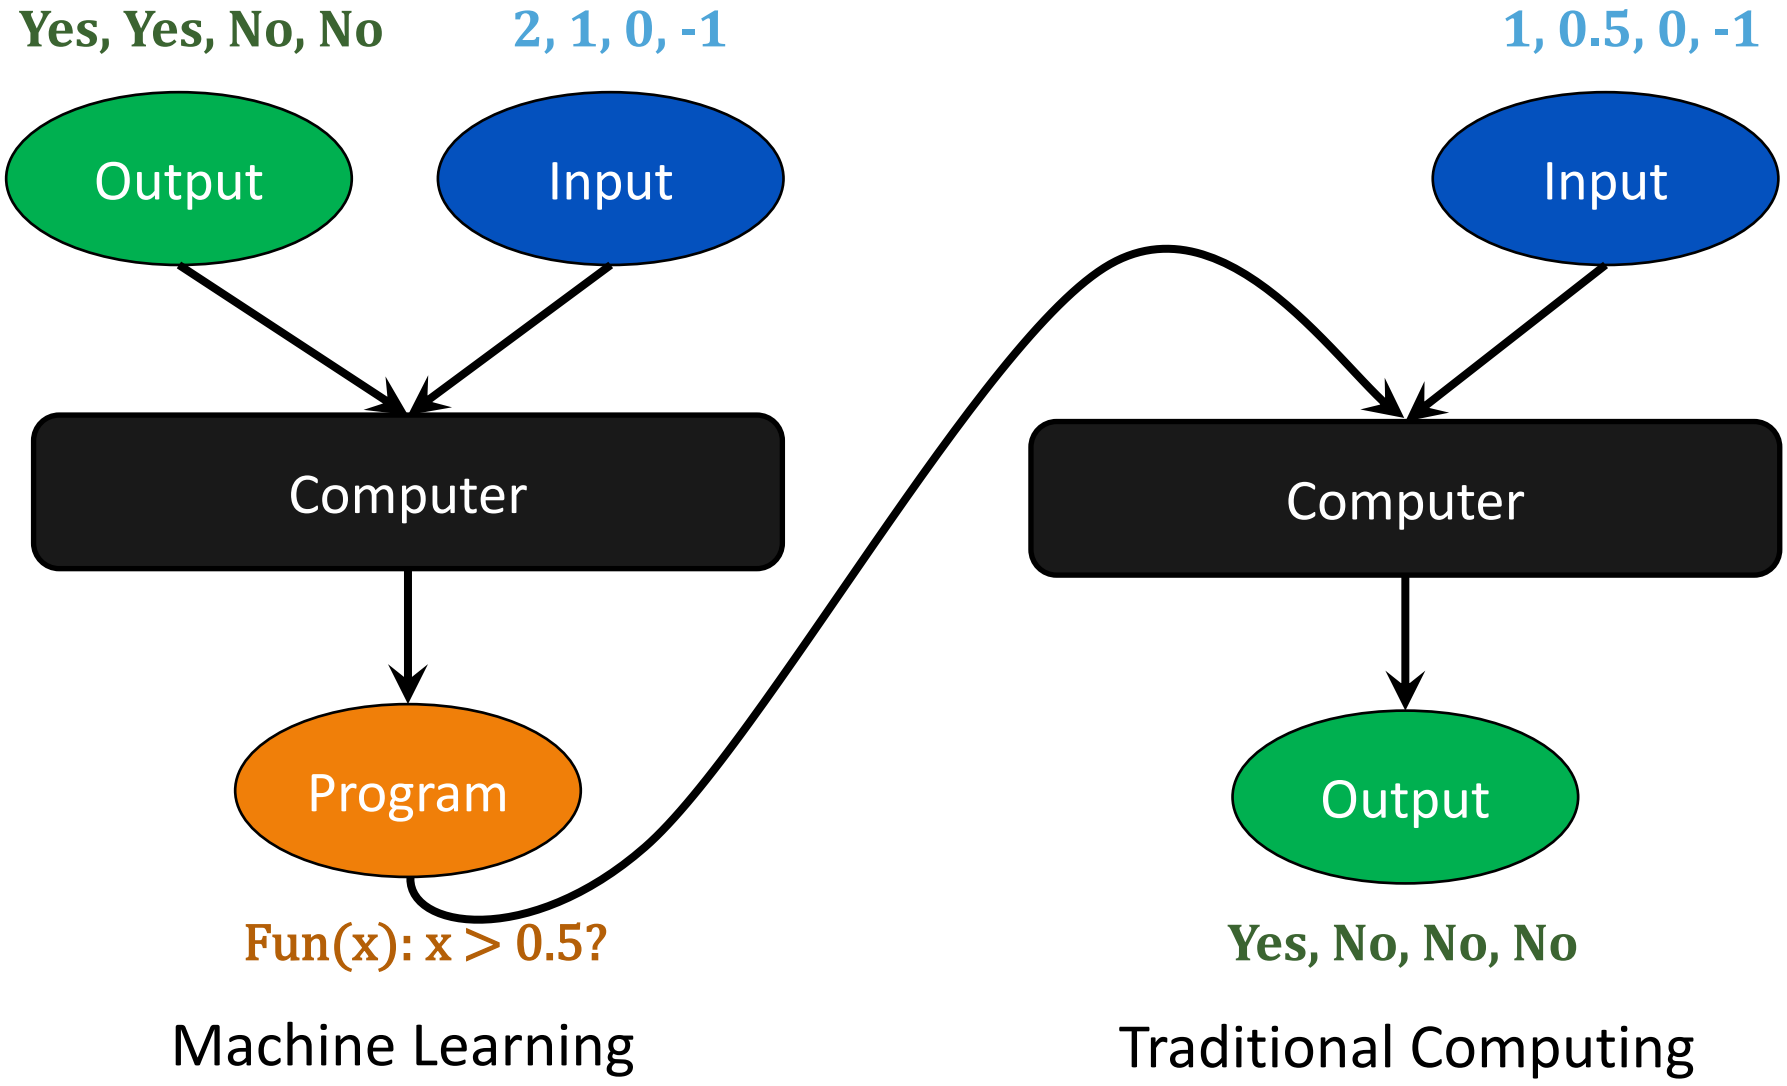

Tom Mitchell, 1997:

A computer program **A** 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**.

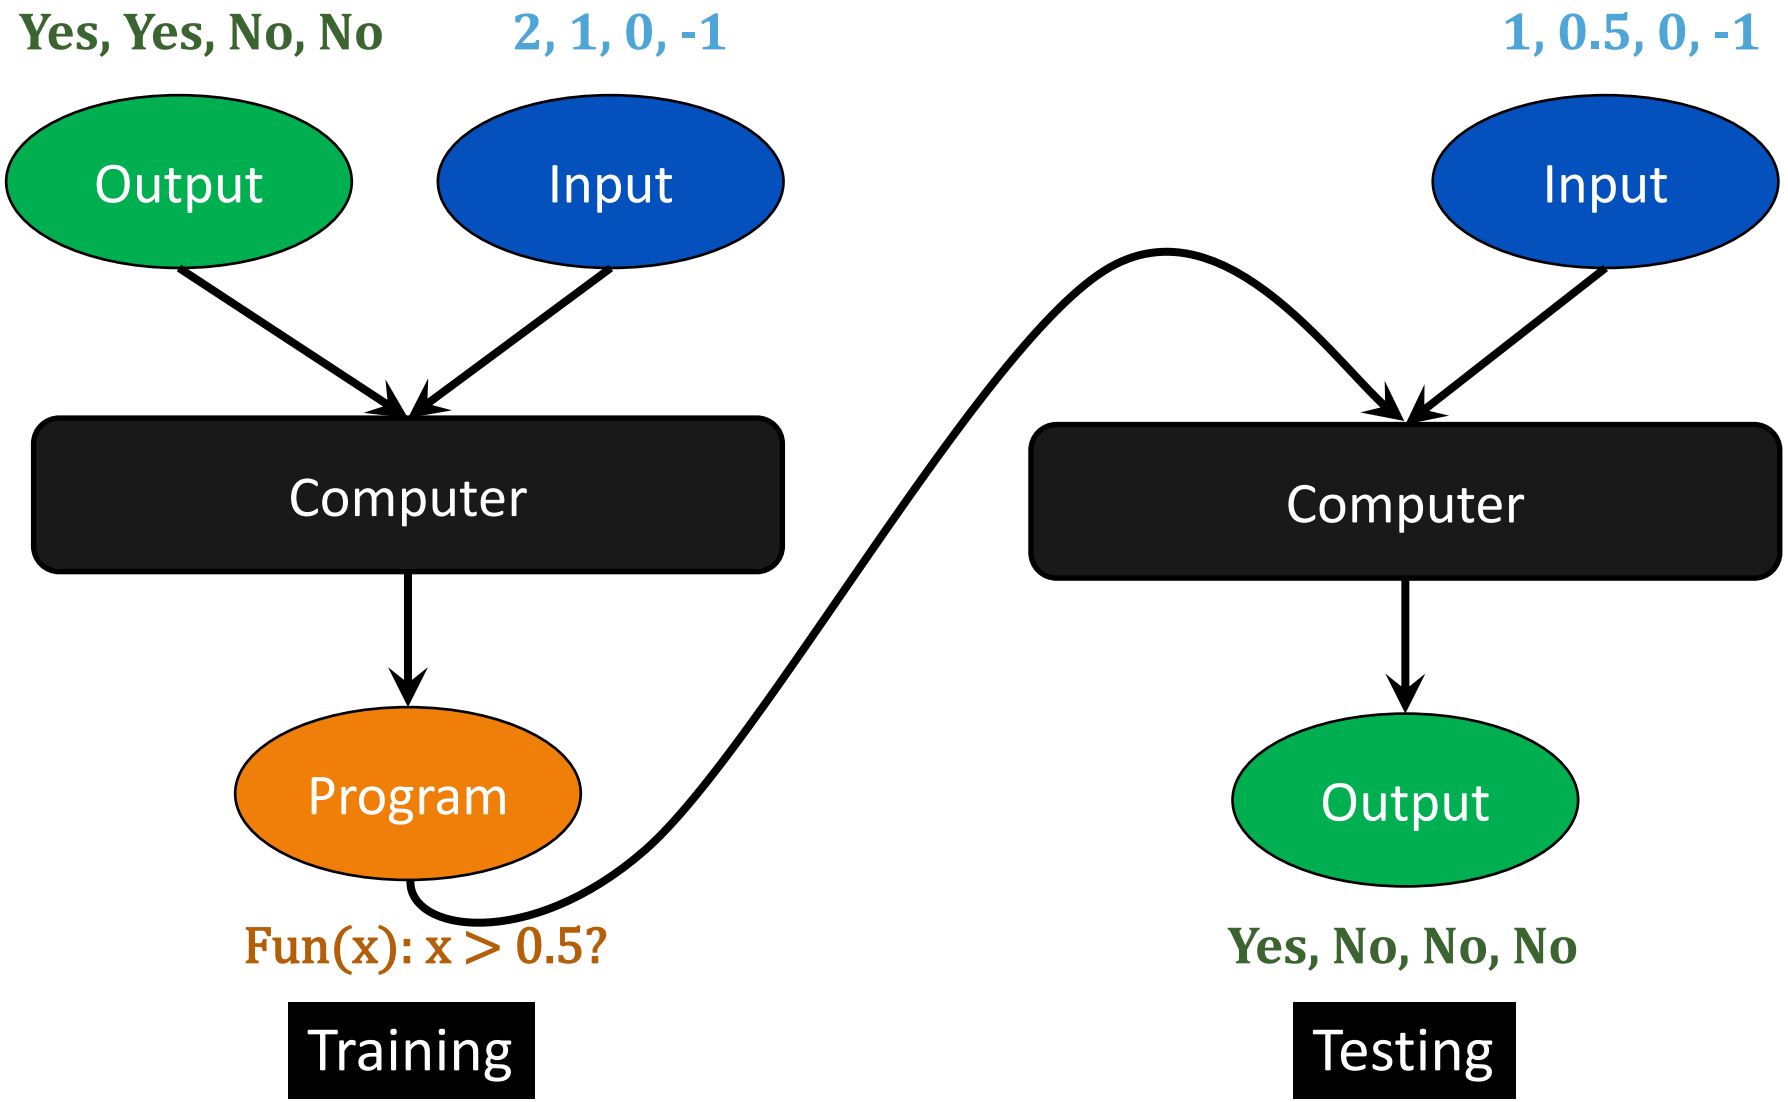

Learning to Detect Spam

Use past emails and whether or not they were flagged as spam.

Learn a program that takes a future email and decides whether it is a spam:
E.g. If the email is from an unknown sender, has a misspelling, and has “Million Dollars” in it flag as spam.

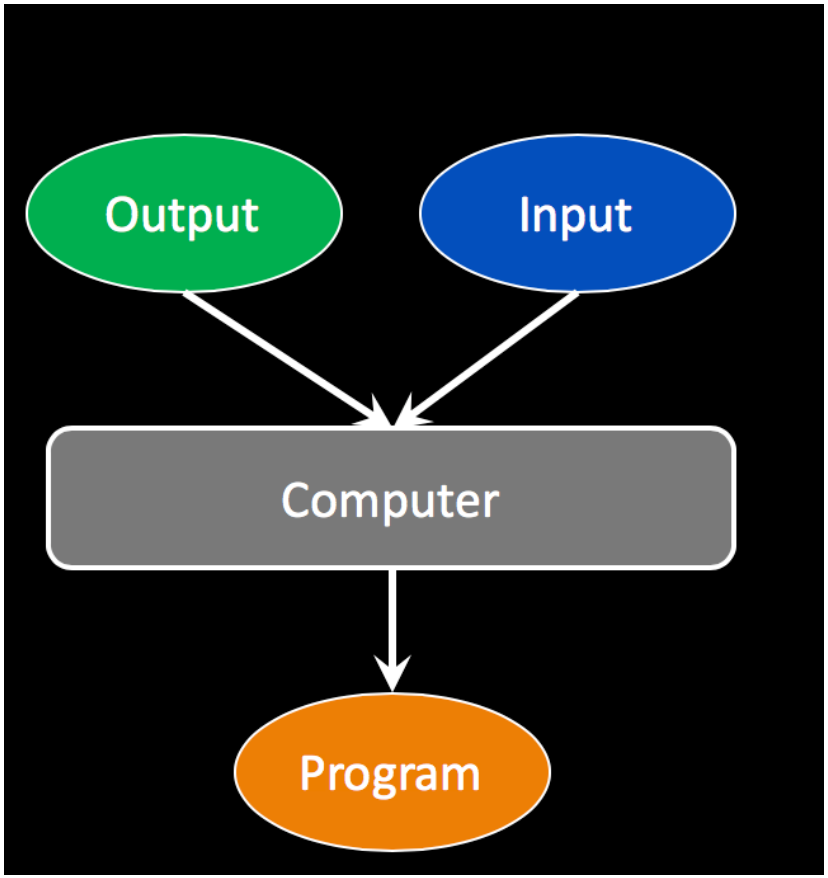

Applications of ML

Use past data to …

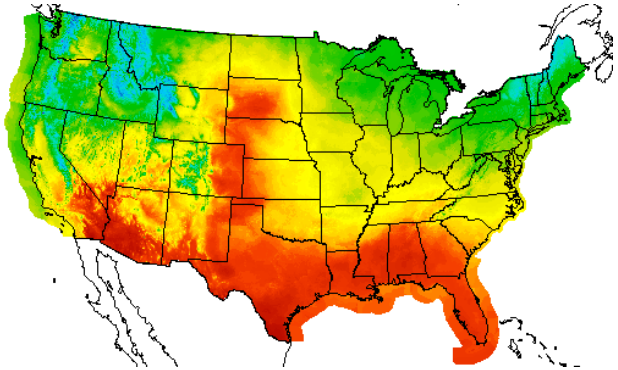

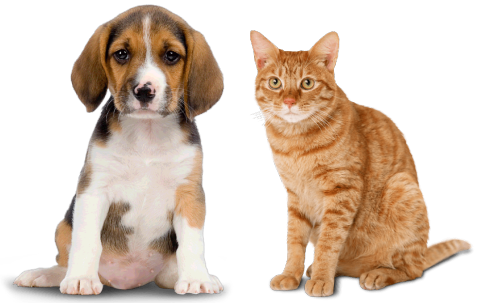

Detect spam

Predict weather

Classify images

<u>Other Examples:</u>

Fraud Detection, Flagging inappropriate social media posts, Natural Language Processing, Document Classification, Designing Economic Mechanisms, Computational Advertising, …

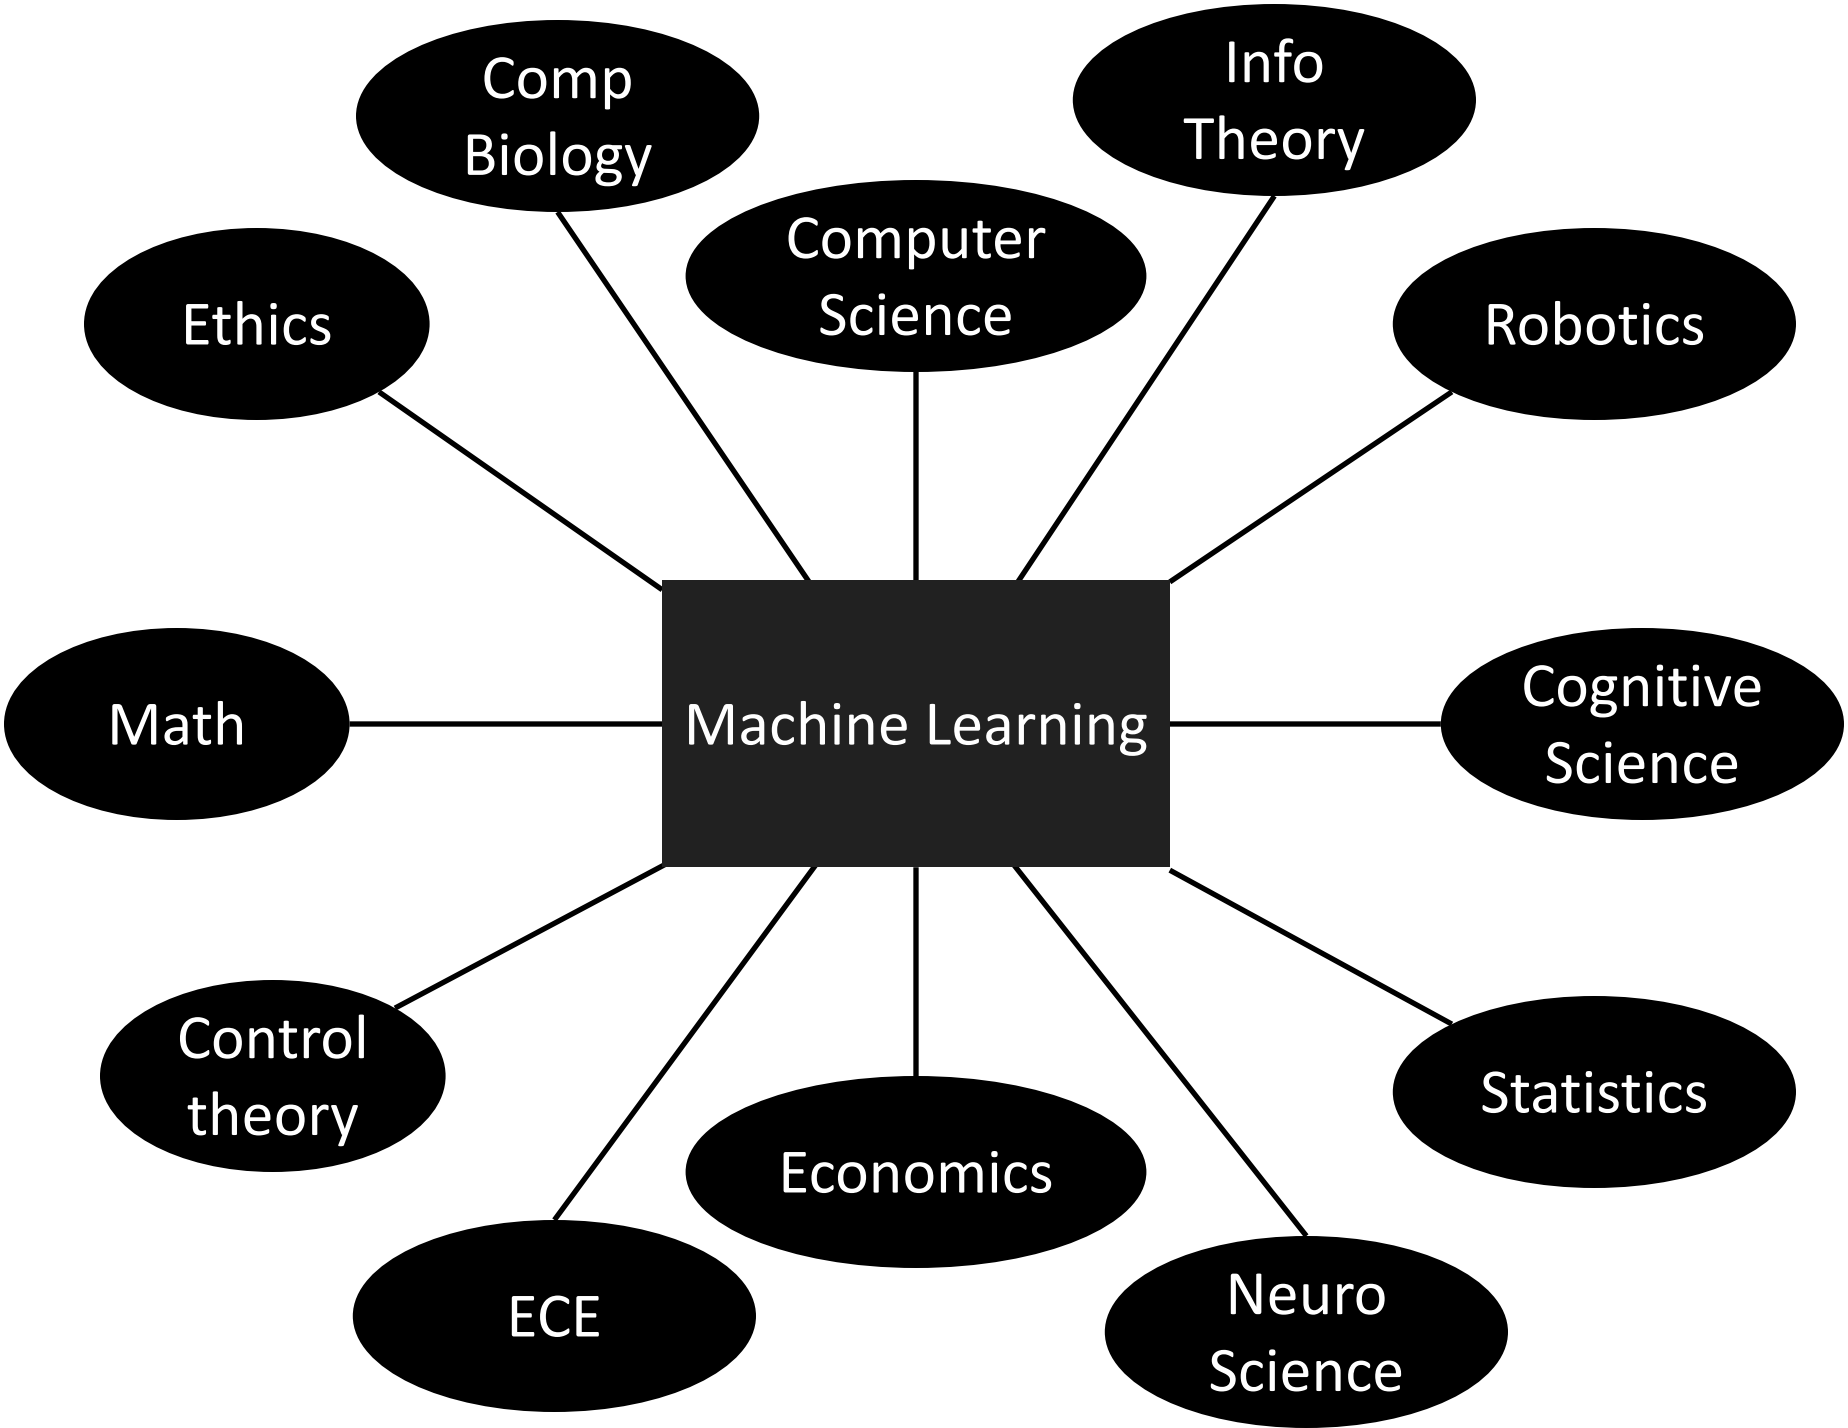
Comp Biology
Info Theory
Computer Science
Ethics
Robotics
Machine Learning
Math
Cognitive Science
Control theory
Statistics
Economics
ECE
Neuro Science

Alan Turing

The Turing Test, 1950

A machine is intelligent if its answers are indistinguishable from a human's.

Arthur Samuel

Checkers Program, 1952

Created a Checkers-playing program that got better overtime.

Also introduced the term “Machine Learning”.

Frank Rosenblatt
@ Cornell!

Perceptron, 1957

Predecessor of deep networks.

Separating two classes of objects using a linear threshold classifier.

Provable learning and convergence guarantees followed later by Albert Novikoff.

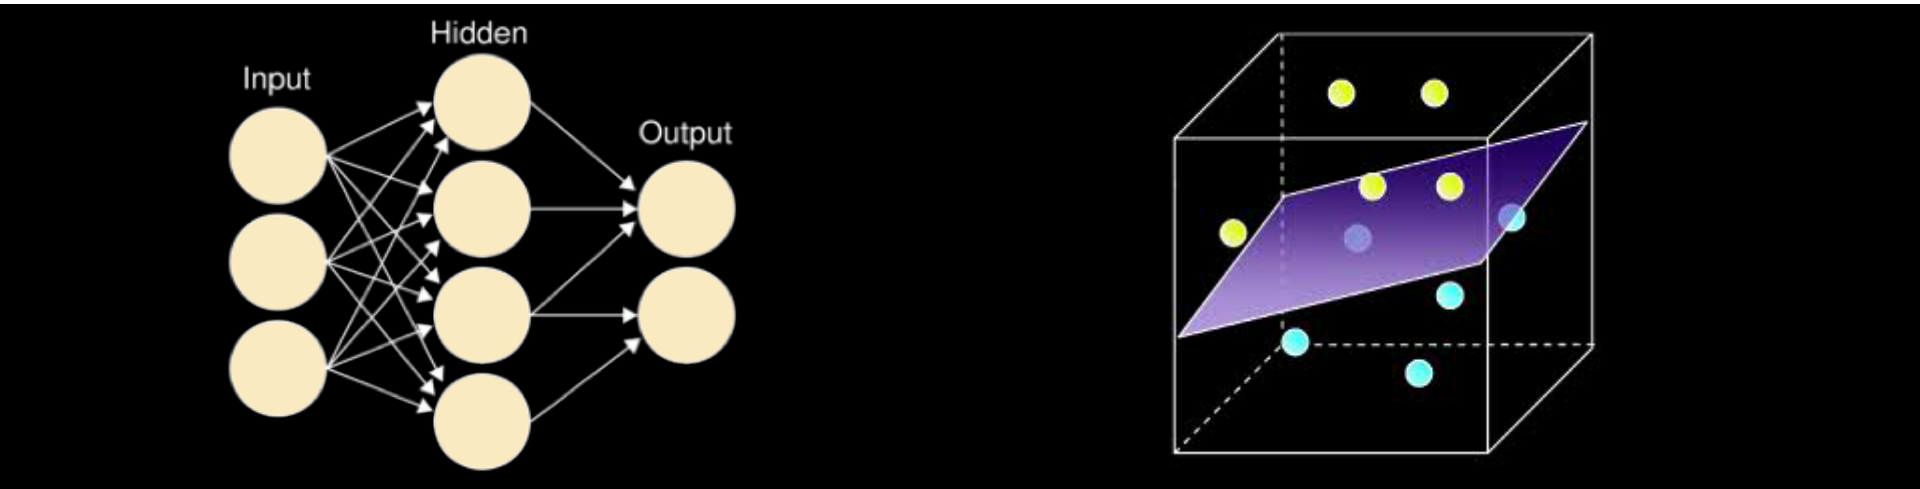

1960s: Lots of hope for AI to solve everything!

AI didn't live up to the hype!

- 1966: Machine Translation failed.
- 1970: Minsky and Papert argued against Perceptron.
- 1971: Speech Understanding failed.
- 1973: Lighthill report torn apart AI.

> "In no part of the field have the discoveries made so far produced the major impact that was then promised"

- 1974: The UK and US stopped funding AI research.

The AI Winter, 1974-1980

Rebirth as Machine Learning

Machine Learning:

- Originally, a bit of a name game to get funding.
- Fundamentally a different approach to intelligence:

Machine Learning	Artificial Intelligence
Data-driven	Knowledge-based
Bottom-up approach	Heavy use of logic
	Top-down approach

Foundations of ML, 1980s-present

Formal notions of learnability from Data.

- When data-driven learning is possible?
  → Probably Approximately Correct Learning (PAC) by Valiant.
  → How much data is required?
- What's the difference between great and mediocre learners?
  → Improving the performance of a learning algorithm.
  → Boosting algorithm of Freund and Schapire.
- How to deal with difficult and noisy learning problems?
  → (Soft Margin) Support Vector Machines by Cortes and Vapnik
- What to do when the learning task evolves over time?
  → Online learning framework.

TD-Gammon, 1992

Gerald Tesauro at IBM thought a neural network to play Backgammon.

The net played 100K+ games **against itself** and beat the world champion.

Algorithm found new techniques that people had erroneously ruled out.

Deep Blue, 1997

IBM's Deep Blue won against Kasparov in chess.

The crucial winning move was made due to machine learning methods developed by Gerald Tesauro.

Expanding the reach, 2000s

Learning to rank

→ Powering search engines: Google, Bing, …

Topic Modeling:

→ Detecting and organizing documents by subject matter.

→ Making sense of the unstructured data on the web.

Online economy:

→ Ad placement and pricing.

→ Product recommendation.

Machine learning became profitable!

Return of Neural Networks, 2010s

Neural networks return and excel at image recognition, speech recognition, …

The 2018 Turing award was given to Yoshua Bengio, Geoff Hinton, and Yann LeCun.

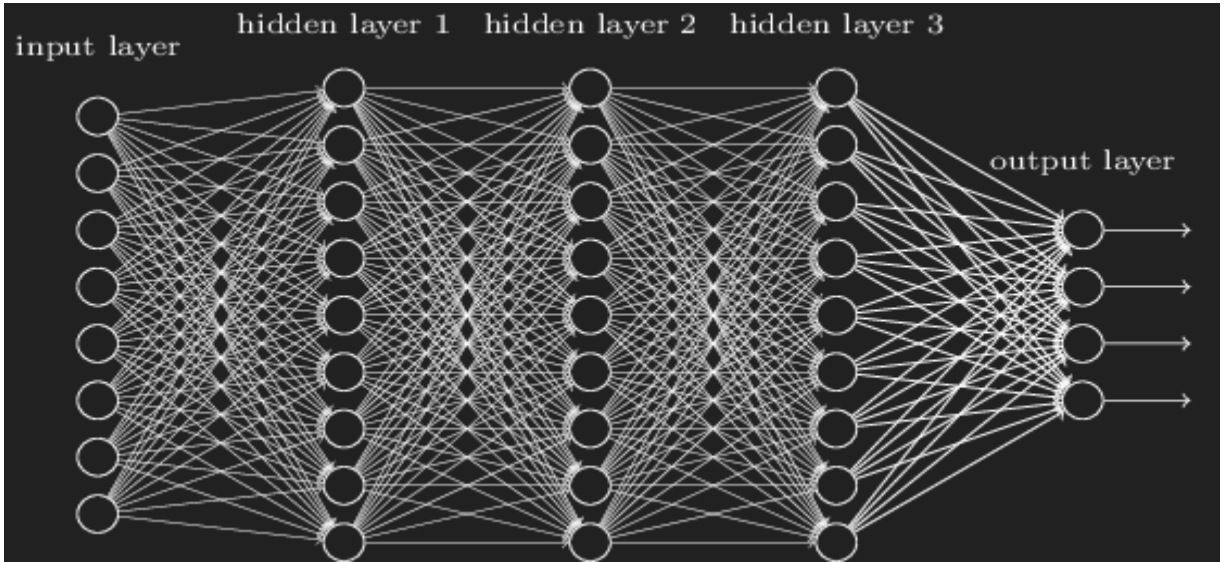

Surrounded by Machine Learning

Google Translate

machine learning

فراگیری ماشین

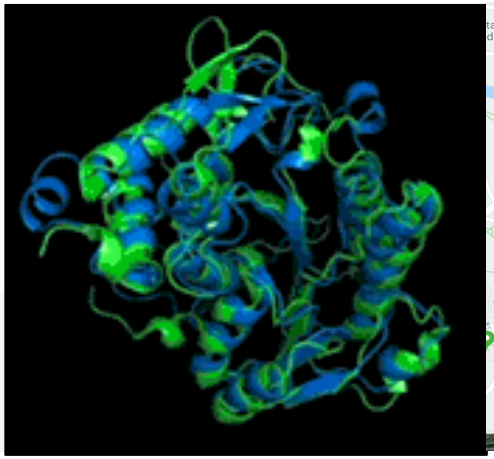

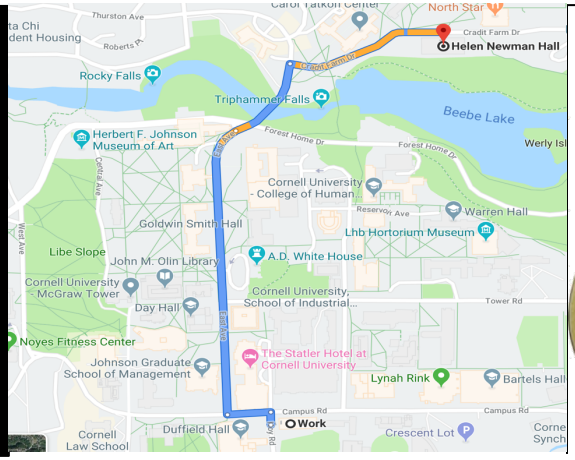

Azure Machine Learning | Create Your Free Account Today

Ad azure.microsoft.com/Services/MachineLearning

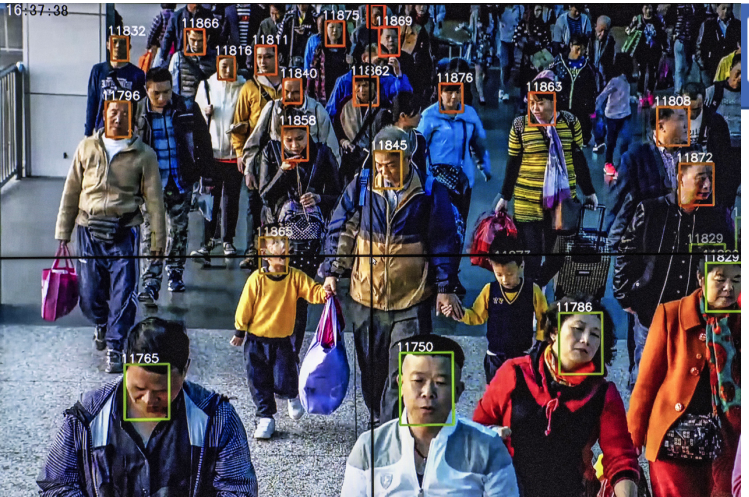

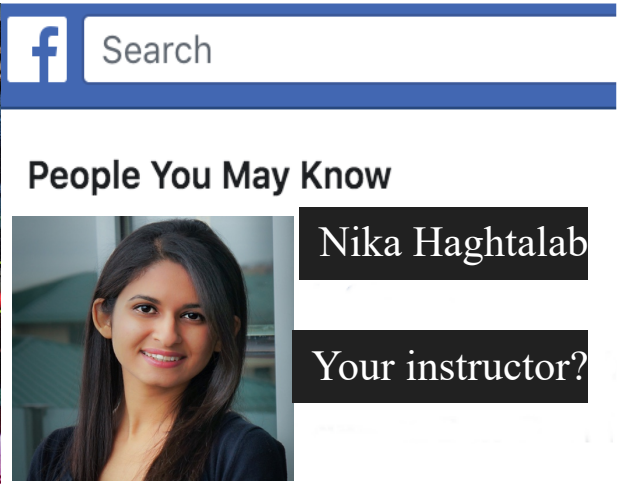

“With great power, there must also come – great responsibility!”

Data Privacy

Learning models leak training data

(Fredrickson et al. ‘15)

Leaked data

Real image

Learning algorithms detect sexual orientation better than people

(Wang & Kosinski’17)

Formal definitions of data privacy:

- K- anonymity (Sweeney)
- Differential Privacy (Dwork, McSherry, Nissim, Smith).

Latanya Sweeney

Cynthia Dwork

Frank McSherry

Kobbi Nissim

Adam Smith

Robust and Secure ML

Image Recognition

Misreading traffic signs

(Eykholt et al)

Speech recognition

Hide commands in noise (Carlini & Wagner)

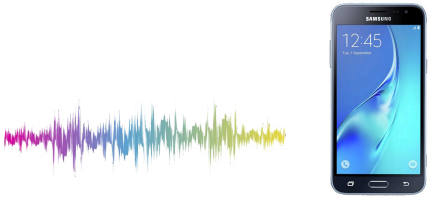

Poisoning Attacks

Tay (chat bot) became inflammatory in 16 hr.

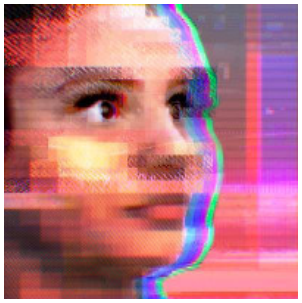

How to create robust and secure machine learning algorithms?

Learning and the Society

- Bad dynamics, perpetuating and worsening stereotypes and biases.
- Who carries the burden of bad prediction?
- How to design good dynamics?

The Best Algorithms Struggle to Recognize Black Faces Equally

Google's algorithm shows prestigious job ads to men, but not to women. Here's why that should worry you.

Gender and racial bias found in Amazon's facial recognition technology (again)

How Amazon Accidentally Invented a Sexist Hiring Algorithm

A company experiment to use artificial intelligence in hiring inadvertently favored male candidates.

Do Google's 'unprofessional hair' results show it is racist?

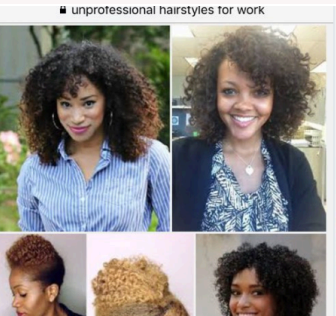

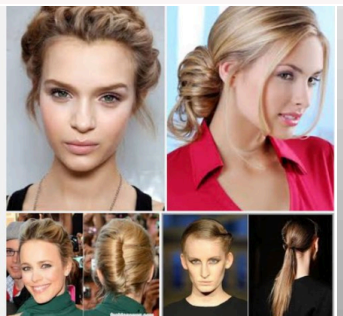

When an Algorithm Helps Send You to Prison

By Ellora Thadaney Israni

Challenging Questions

Machine learning and Artificial Intelligence will shape the future, what kind of a future do we want?

What is the role of machine learning?

→ ML for good versus ML for profit.

How do automation and learning change the quality of life?

→ Job loss and displacement, life satisfaction, safety and security?

How do we approach machine learning and (inter-)national security? Weaponization of machine learning and AI?

This Course

Theory of Machine Learning

Theoretical backing to the wide range of applications seen in practice.

Machine learning interacts with every aspect of human life

→ We need to have guarantees about what it does.

Mathematical tools and abstractions for provably understanding:

- What can be learned from data? And what cannot?
- How much computation is needed to learn?
- What happens to learning algorithms once they are deployed in the wild?
- How does machine learning push forward other branches of science?

This Course

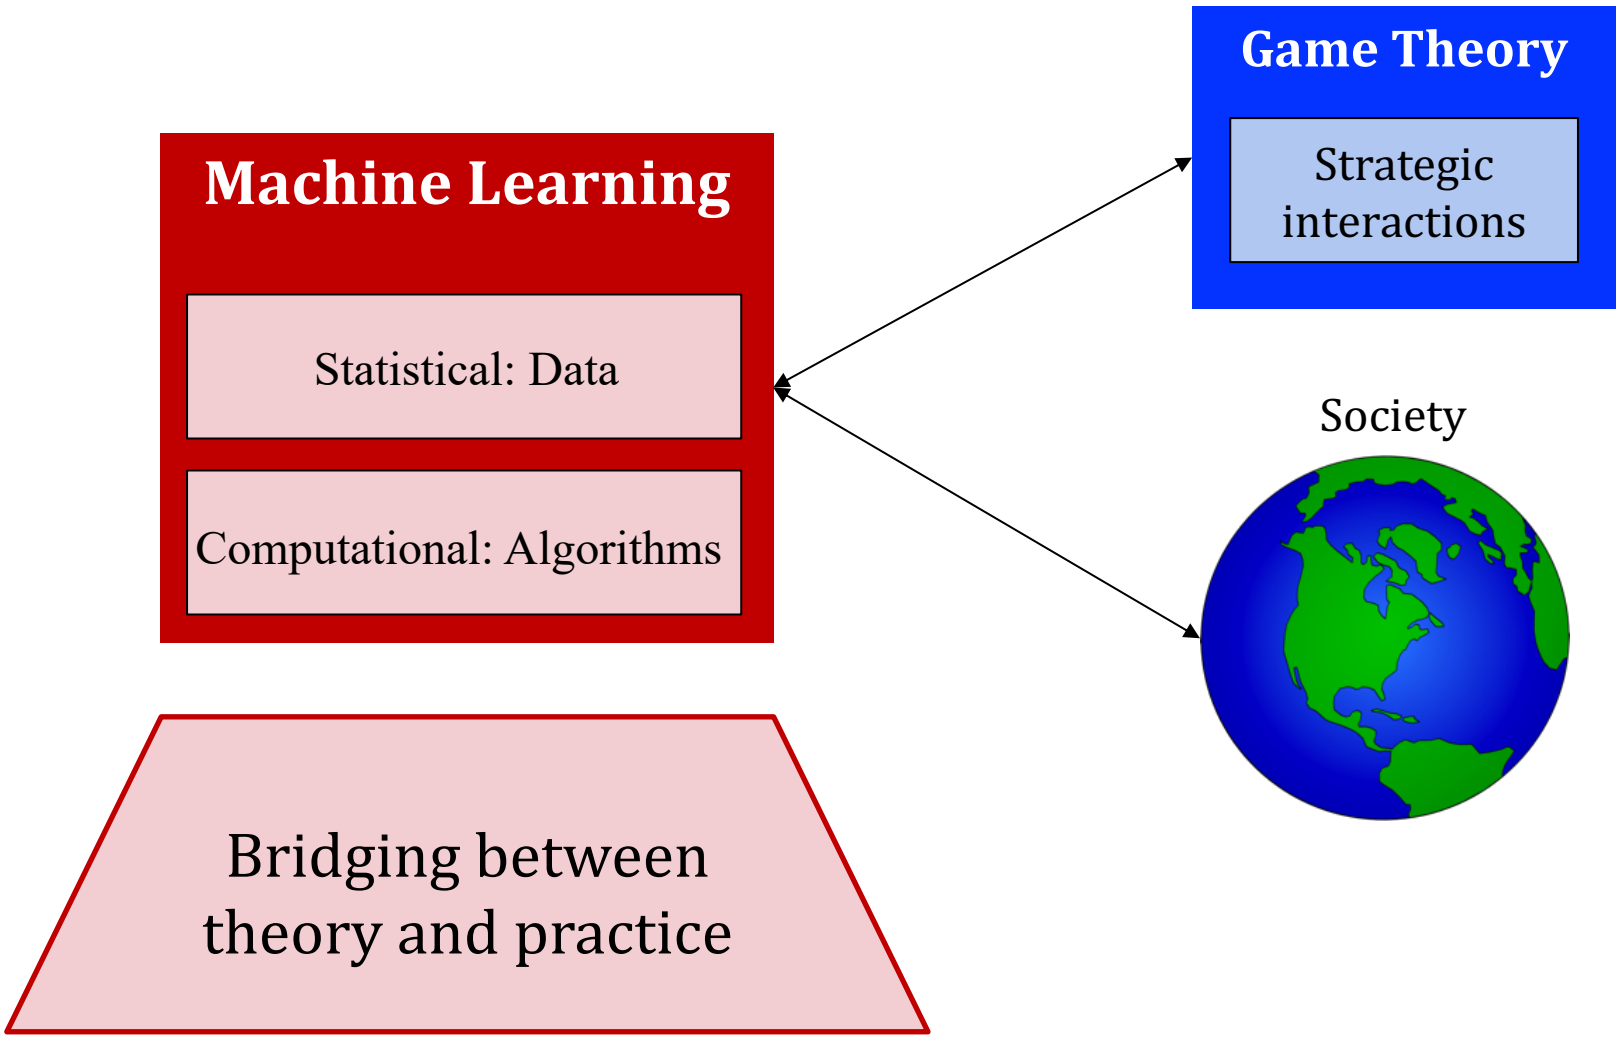
Game Theory
Strategic interactions
Machine Learning
Statistical: Data
Computational: Algorithms
Society
Bridging between theory and practice

Syllabus

Mathematical foundations:

- **Basic tool set for the theory of machine learning:** Offline and online learning, combinatorial dimensions in learning, concentration, generalization, boosting.
- **Computational aspect of machine learning:** Hardness of learning, algorithmic efficiency.
- **Connections between learning and game theory**: min-max theorems, equilibria and learning, GANs, adversaries…
- **Beyond the worst-case analysis of learning:** Middle ground between adversarial and benign world views, noise models, and computations/statistics for every day use.
- **Societal considerations:** Learning and privacy, learning and fairness.
- **Other paradigms of learning:** active learning, semi-supervised, and other topics.

Lecture Delivery

We will use slides occasionally.

But, most of the lectures will be on the board.

→ We will have lecture notes for the first few lectures.

→ The rest of the lectures will be scribed by you!

What you need for success

Interest and hard work!

This course:

- Theory! Motivated by needs of the real world applications.
- Understanding the mathematical foundations.
- No programming skills needed.
- You need to be comfortable with understanding/writing mathematical proofs and abstractions.

  → Basics of probability theory, algorithms, a bit of linear algebra and high dimensional calculus.

Formal Pre-req:

- None for graduate students. 4820 for undergraduate students.

Related Courses

6783: Machine learning theory.

→ Some overlap in topics, specifically the basic toolset of offline and online learning.

6780: Advanced Machine learning (PhD)

6784: Topics in ML: variety of topics

4780/5780: ML for intelligent systems.

4686/5786: ML for data science.

4787: Large scale ML

More applied

7992: Seminar on Bias and Fairness in Learning Systems

4732: Ethical and Social Issues in AI

Other courses in ORIE, STAT, etc.

Course Material

No required textbook.

→ We will rely on in class presentations (white board/slides), lecture notes, scribe notes, and other online material.

Great additional resources (see the website for links):

- Shalev-Shwartz, Ben-David, "Understanding Machine Learning - From Theory to Algorithms", Cambridge University Press.
- Blum, Hopcroft, Kannan, “Foundations of Data Science”, Cambridge University Press.
- Bousquet, Boucheron, Lugosi, “Introduction to Statistical Learning Theory”, Springers.

Course and Grade Components

- Deliverables
  - Homeworks (50% of Grade)
  - Final Project (25% of Grade)
  - Scribing / Participation (5% of Grade)
  - Final Exam (take-home) (20% of Grade)

Check out the homepage for detailed policy.

Homework Assignments

Assignments

- 5 written homeworks, total 50% of total grade.
- Tentative due dates already online.
- Include algorithm design, analysis, proofs, etc. No programming.

Policies

- To be done individually. You can discuss the problems with your classmates. But, do not take any record of the discussion and write your solutions separately.
- Everybody has 5 “free” late days. When you run out of late days, you’ll loose 1% per day from the total homework grade.
- No assignments will be accepted after the solutions have been made available (typically 3-5 days after deadline).
- Do not post the homework or your solutions publicly, whether on Cornell or other websites.

Final Project

The class has a final project:

- A project proposal due mid way through the semester (see the homepage for tentative dates). 5% of total grade.
- Final project writeup. 20% of the total grade.
- There will be a poster session.

Two types of projects:

- Survey of a couple papers on one topic.
- Original contributions to theory of machine learning.

Projects can be done in groups of 2.

Scribing / Participation

Full policy will be announced after the size of the class is finalized (after add/drop).

Scribing:

- Each student will be responsible for scribing 0-2 lectures.
- A lecture can be scribed by 1-2 students.
  → Two students can scribe a lecture together only if at most one of them is a PhD student.
- We will start scribing on February 6 (after the add/drop deadline)
  →We will send out a signup sheet next week.

Policy:

- We need fast turnaround!
- Within 2 weekdays, prepare a draft and meet with one of the TAs to discuss it. You have two more weekdays to prepare the final version.
- We aim to post the scribed notes within a week of the lecture.

Exams

No midterm exam.

Take-home final exam for 20% of the total grade.

Details will be announced later.

How to Get in Touch

Online

- Course Homepage (slides, notes, additional resources, course policies, office hours, etc.)
  - https://www.cs.cornell.edu/courses/cs6781/2020sp/
- Piazza forum
  - piazza.com/cornell/spring2020/cs6781/home
- CMS (Submit deliverables on CMS)
  - https://cmsx.cs.cornell.edu/

Office Hours

- Nika: Fridays 10am – 11am, 315 Gates Hall
- Wilson: Mondays 5:30pm to 6:30 Rhodes 402
- Abhishek : Thursdays 4:30 to 5:30 Rhodes 412
- Subject to change, check Piazza and the website.

These meetings aren’t always going to be this early, are they?!

See you at 8:40am on Thursday!