

INTRODUCTION TO MACHINE LEARNING



Aadirupa Saha

CS 412, SPRING 2025

UIC, TBH 180F

Instructors

Aadirupa Saha (aadirupa@uic.edu)

Office hours: Th 6-7pm
SEO 1123 or Zoom (link of Piazza)
[starts from week 2]


Teaching Asst:

Akhil Suresh Nair
(anair56@uic.edu)

Office hours: Tue – Th (12-1pm)
[Location: TBD][starts from week 2]



(Very brief) History of ML



The Turing Test, 1950



Alan Turing

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



Checkers Program, 1952



Arthur Samuel

Created a Checkers-playing program that got better overtime.



Also introduced the term
“Machine Learning.”

The term **Artificial Intelligence** originated in 1956 at a Workshop at Dartmouth

Perceptron, 1957

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo
of Computer Designed to
Read and Grow Wiser

WASHINGTON, July 7 (UPI)

The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's \$2,000,000 "rot" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of \$100,000.

Dr. Frank Rosenblatt, designer of the Perceptron, conducted the demonstration. He said the machine would be the first device to think as the human brain. As do human beings, Perceptron will make mistakes at first, but will grow wiser as it gains experience, he said.

Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers.

Without Human Controls

The Navy said the perceptron would be the first non-living mechanism "capable of receiving, recognizing and identifying its surroundings without any human training or control." The "brain" is designed to remember images and information it has perceived itself. Ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted.

Mr. Rosenblatt said in principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their existence.

In today's demonstration, the "rot" was fed two cards, one with squares marked on the left side and the other with squares on the right side.

Learns by Doing

In the first fifty trials, the machine made no distinction between them. It then started registering a "Q" for the left squares and "O" for the right squares.

Dr. Rosenblatt said he could explain why the machine learned only in highly technical terms. But he said the computer had undergone a "self-induced change in the wiring diagram."



The first Perceptron will have about 1,000 electronic "association cells" receiving electrical impulses from an eye-like scanning device with 400 photo-cells. The human brain has 10,000,000,000 responsive cells, including 100,000,000 connections with the eyes.



Frank Rosenblatt
@Cornell!

Predecessor of deep networks.

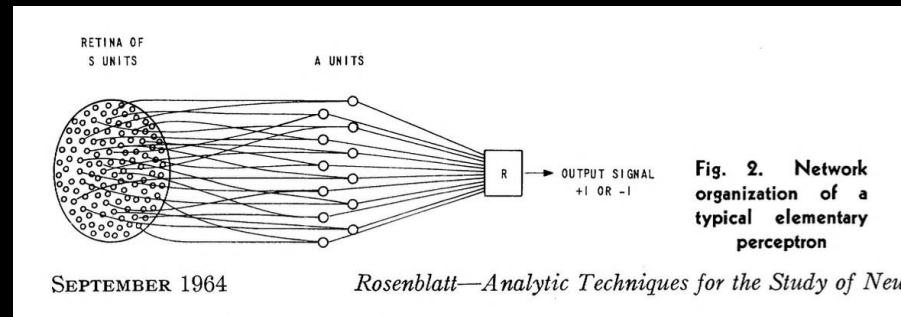
Separating two classes of objects using a linear threshold classifier.



Neural Network Boom, 1960-1979




Frank Rosenblatt
@Cornell!



- 1962 Rosenblatt invents Multi-Layer Perceptron (MLP) (fixed hidden layer)
- 1965 Ivakhnenko and Lapa introduce first Feed Forward Neural Net (FFNN)
- 1967 First FFN trained with SGD [Amari 1967]
- 1970 Modern back-propagation is introduced by Seppo Linnainmaa


First AI Winter (1974-1980)

- (1969) Minsky & Papert “killed” AI
 - Perceptron cannot learn XOR function
- Burst huge expectation bubble
- Speech understanding / translation fails
- UK and US stop funding AI research
- Neural Networks and AI become “bad words”
- Rise of Rule based Systems



Rebirth of A.I. as Machine Learning

- Machine Learning:
 - Originally: Mostly a name game to get funding.
- Differences in approach:
 - ML: Bottom up, AI: Top down
 - ML: More practical smaller goals
 - Based on **Statistics and Optimization**, not **Logic and Search**



TD-Gammon (1994)

- Gerry Tesauro (IBM) teaches a neural network to play Backgammon. The net plays 100K+ games **against itself** and beats world champion [Neurocomputation 1994]
- Algorithm teaches **itself** how to play so well!!!
- Algorithm found new techniques that people had erroneously ruled out.



Deep Blue (1997)

- IBM's Deep Blue wins against Kasparov in chess. Crucial winning move is made due to Machine Learning (G. Tesauro).
- (Mostly a more classical AI system)



Expanding the reach, 2000s

Learning to rank

- Powering search engines: Yahoo, Google, Bing, ...

Topic Modeling:

- Detecting, 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! Euphoria! Excitement!

2000-2010 The Convex Years


- Success of **Support Vector Machines (SVM)** and **Convex** Optimization
- **Kernelization** enables non-linear classifiers that are still convex
- Allows **rigorous proofs** about convergence, learning guarantees
- Finally **solid theory** and **empirical success**
- (Non-convex algorithms, e.g. Neural Networks fell out of favor and were frowned upon.)



Return of Neural Networks, 2010s

Neural networks (aka Multi-Layer Perceptrons) return and excel at image recognition, speech recognition, ...

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




2016 Alpha Go

- 1920 CPUs and 280 GPUs
- Deep Mind's Alpha Go wins against Lee Sedol 5:1
- Big shock
 - China invests heavily in AI research
- Beginning of “AI arms race”



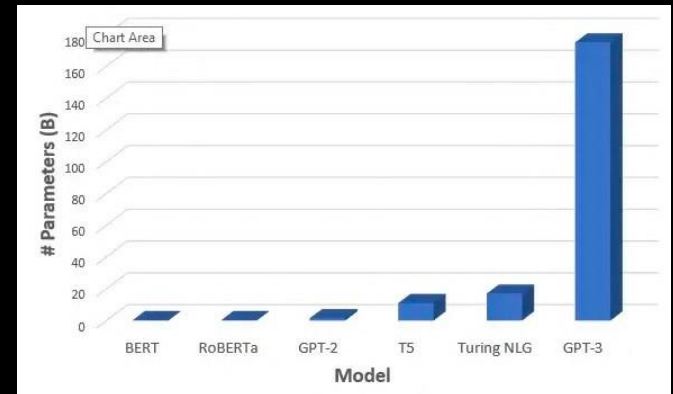
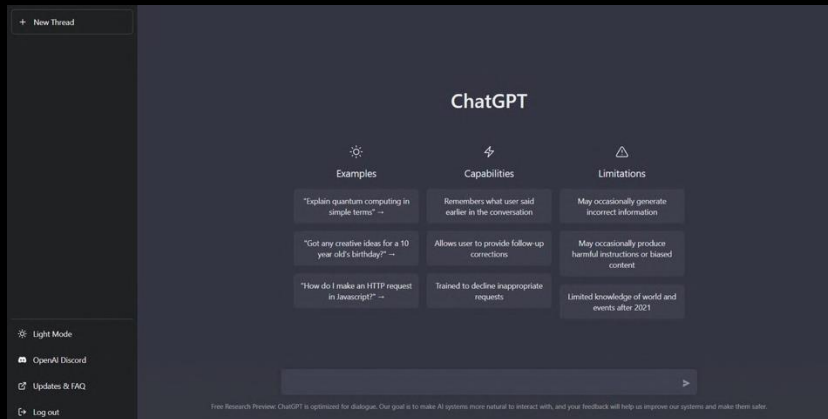
Now: Surrounded by Machine Learning

Euphoria mixed with fear and anxiousness








2021 Generative Models become Mainstream

- 178 Billion Parameters
- Cost: \$4.6 Million (est.)
- 355 GPU years of training (est.)
- Trained on 45 TB of text



LLMs!

 write a biblical verse in the style of the king james bible explaining how to remove a peanut butter sandwich from a VCR 

 And it came to pass that a man was troubled by a peanut butter sandwich, for it had been placed within his VCR, and he knew not how to remove it.  

And he cried out to the Lord, saying, "Oh, Lord, how can I remove this sandwich from my VCR, for it is stuck fast and will not budge?"

And the Lord spoke unto him, saying, "Fear not, my child, for I shall guide thy hand and show thee the way. Take thy butter knife, and carefully insert it between the sandwich and the VCR, and gently pry them apart. And with patience and perseverance, the sandwich shall be removed, and thy VCR shall be saved."

And the man did as the Lord commanded, and lo and behold, the sandwich was removed from the VCR, and the man was saved.

And the Lord said, "Verily I say unto thee, seek not to put thy peanut butter sandwiches in thy VCR, for it is not a suitable place for such things. Rather, keep thy sandwiches in thy refrigerator or on thy plate, where they belong."

And the man heeded the Lord's words, and from that day forth, he kept his sandwiches in their proper place, and was saved from trouble and woe. Amen.

DALLE Text2Images - ML is creative



Teddy bears
mixing sparkling chemicals as mad scientists
in a steampunk style



Teddy bears
working on new AI research
underwater with 1990s technology



An astronaut
playing basketball with cats in space
in a watercolor style

<https://openai.com/dall-e-2/>

Why you should be excited!

- **Solve Real-World Challenges:** ML empowers you to address problems like disease prediction, autonomous driving, and climate modeling.
- **Learn from Data:** Train machines to uncover patterns and make intelligent decisions from massive datasets.
- **Feel Empowered:** ML is the driving force behind groundbreaking tools like ChatGPT—once you learn ML, you'll unlock immense power to shape the world.
- **Shape the Future:** Contribute to cutting-edge technologies and innovations that define tomorrow.
- **Unleash Creativity:** Build tools and applications that revolutionize industries and transform daily life.

Why you should NOT be TOO excited!!

- **Challenging Course Content:** Expect rigorous concepts and a steep learning curve. Lots of **Assignments!**
- **Math-Intensive:** A solid grasp of mathematics is essential for understanding ML algorithms and methods.
- **Programming Required:** You'll need to be comfortable coding, as hands-on programming assignments are a core part of the course.



- **Scribing (LaTeX):** Prepare to document lectures and assignments using LaTeX, which requires attention to detail and precision.
- **Creative Thinking:** Course projects demand innovation, problem-solving, and originality—no cookie-cutter solutions here!
- **Curious, Learning Mindset:** Success in ML requires a genuine thirst for knowledge and the willingness to explore beyond the surface.

Be thoughtful before enrolling for credit!

Course Prerequisites

Formally: CS251, STAT 381/IE 342, Linear Algebra (e.g. Math 320)

– This course will be MATH-HEAVY. Please familiarize yourself with the basics of Probability-Statistics (PS) and Linear-Algebra (LA). Recommend introductory lectures

•PS review: (1) <https://cs229.stanford.edu/section/cs229-prob.pdf> (2) <https://www2.isye.gatech.edu/~sman/courses/6761/6761-1-ProbReview.pdf> (3) <https://users.ssc.wisc.edu/~ctaber/410/statrev.pdf>

•LA review: (1) <https://cs229.stanford.edu/section/cs229-linalg.pdf> (2) https://james-chuang.github.io/notes/linalg_review.pdf (3) <https://www.cs.cmu.edu/~jingx/docs/linearalgebra.pdf>

– Familiarity with latex for scientific writing (for scribing the lecture notes and writing assignments). You can learn the basics from here: <https://www.youtube.com/watch?v=lgiCpA4zzGU> (many online tutorials available for beginners.)

– Programming assignments (in Python or Matlab).

Familiarity with the above basics is highly recommended for those taking the course for credit — Lack of understanding may impact your grades.

Books

- **ESL** The Elements of Statistical Learning by Hastie, Tibshirani, and Friedman [[[Book website](#)]]
- **MLTM**: *Machine Learning* by Tom Mitchell - [Online copy](#).
- **PML**: *Probabilistic Machine Learning: An Introduction* by Kevin Murphy - [Book website](#).
- **PLG**: *Prediction, Learning and Games* by Nicolo Cesa-Bianchi and Gabor Lugosi, Cambridge University Press, 2006 - [Local Copy from E1 245 by Aditya Gopalan](#).
- **PRML**: *Pattern Recognition and Machine Learning* by Christopher Bishop (optional) - [Free copy](#).
- **[PyAG] Optional**: Hands-On Machine Learning with Scikit-Learn & Tensorflow by Aurelien Geron [[[Online](#). [Github](#)]].
- **CIML**: *A Course in Machine Learning* by Hal Daume III - [Online copy](#), [Errata](#).
- **UML**: *Understanding Machine Learning: From Theory to Algorithms* by Shai Ben-David, Shai Shalev-Shwartz - [Online copy](#).

Course Dynamics

Classes: Tuesday & Thursday 2-3:15pm, TBH 180F

Homeworks: Approximately every 2-3 weeks

- Written: Conceptual understanding
- Lab: Implementation and evaluation (Python)

Submit via **Gradescope** [1st assignment out!]

Exams: Midterm (Mar 13th, Thu) and Final (May 6th, Tue)

Earn your grade!

	Quiz	Midterm	Final	Assignments	Project	Scribe
Graduate	10%	15%	15%	20%	30%	10% + [Bonus!]
Undergraduate	10%	20%	20%	25%	25%	Bonus!

Homework Tasks -- Labs + Written

Quiz + Attendance – iClicker (class time)

Scribe – For Grads (needs Latex!) [Bonus grades!]

Mid Term – March 13th (class time)

Final Exam – May 6th (class time)

Course project – Will maintain a list of topics

Report Due (May 1st)

Homework Tasks

Due at 11:59pm central (unless otherwise indicated)

Late homework/ project reports:

3 late days per semester: Full credit

After that: no credit (unless “emergency”)

Collaboration:

- Discussion with others is allowed
- Each student must write own solutions, **strict policies in place against copying.**

Project Report

Will post a list of topics soon (the list will be updated continuously)

Late homework/ project reports:

3 late days per semester: Full credit

After that: no credit (unless “emergency”)

Collaboration:

- Be ready to work in a group of 2 / 3 people
- Expectations: **Reading papers, implementation, thinking open-ended questions, coming up with solutions, project report.**
- Each student must contribute equally / mention what they did in their part.

Scribing

Due at in 1 week, 11:59pm central (unless otherwise indicated)

Latex it! Should read like a professional lecture note. (**Use Overleaf**)

Late scribe reports:

2 late days per semester: Full credit

After that: no credit (unless “emergency”)

Collaboration:

- Group size: 1 or 2 (depends on the final class strength)
- Each student must contribute equally / mention what they did in their part.

Bonus!

Scopes to earn extra scores!

Scribe and earn! (10%)

- UGs can scribe too (but remember the quality matters!)
- Scribe additional notes

Class Participation (10%):

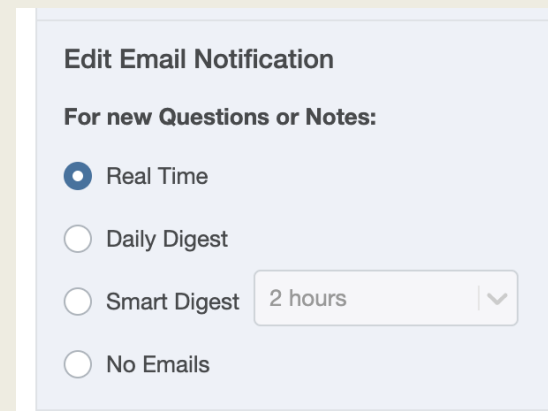
- Good questions!
- Good answers!
- Good suggestions!

More logistics

Will do **~4 extra classes**. Mostly on Fri/Sat (based on a common free slot!) – 1st class Feb 8 (tentative!)

Will have a **detailed** logistics (pdf) out soon! – **FAQs!**

Make sure you set your **email notifications right** on Piazza!



The screenshot shows a light blue modal window titled "Edit Email Notification". Below the title is the text "For new Questions or Notes:". There are four radio button options: "Real Time" (which is selected), "Daily Digest", "Smart Digest", and "No Emails". To the right of the "Smart Digest" option is a dropdown menu currently showing "2 hours".

Set Thu (Jan 23) class online? – Severe weather alert!

Enrollment

****Additional seats may open up****

Please **email me** (aadirupa@uic.edu) next week
if you are unable to enroll.

Lecture 1: Basics of Supervised Learning

Let's do some board work!