DS4440
practical neural networks

Instructor Byron Wallace
TA Sanjana Ramprasad

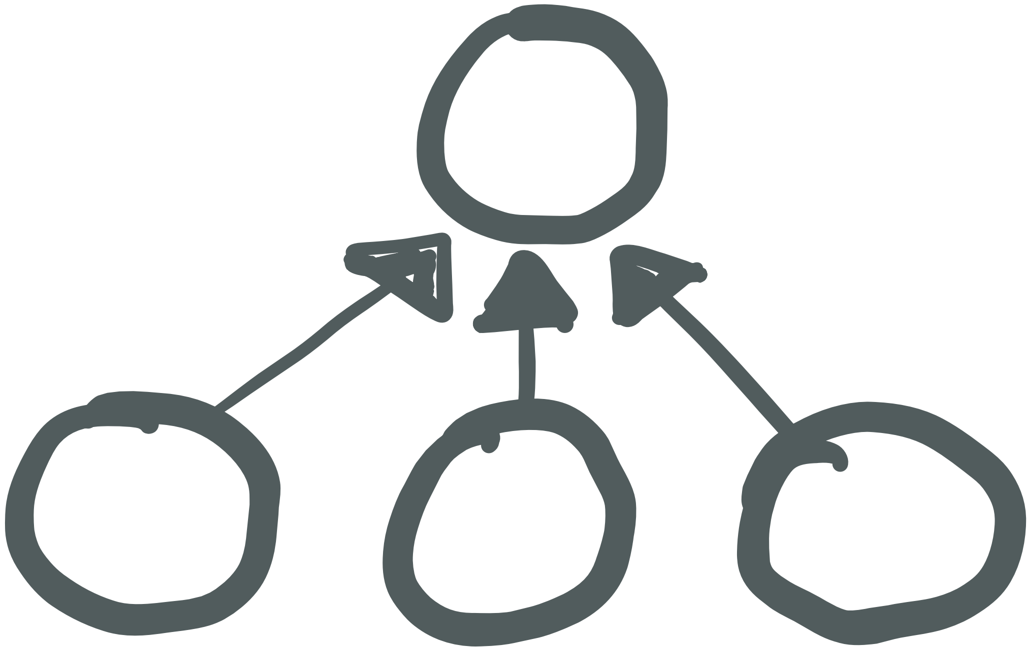

Northeastern University
Khoury College of Computer and Information Sciences

About me

I work mostly in NLP

It’s been an interesting ~5 years

Language modeling, circa 2017

Language modeling, circa now

T

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.

Generative AI: All the Hype

The New York Times

The Brilliance and Weirdness of ChatGPT

A new chatbot from OpenAI is inspiring awe, fear, stunts and attempts to circumvent its guardrails.

npr

We made ChatGPT write a song for us

THE WALL STREET JOURNAL.

The Jobs Most Exposed to ChatGPT

New study finds that AI tools could more quickly handle at least half of the tasks that auditors, interpreters and writers do now

The Economist

Is GPT-4 the dawn of true artificial intelligence?

What this course is about

How do these things work?
How do we train them?
How do we evaluate them?

The McCullough-Pitts neuron

brain stuff

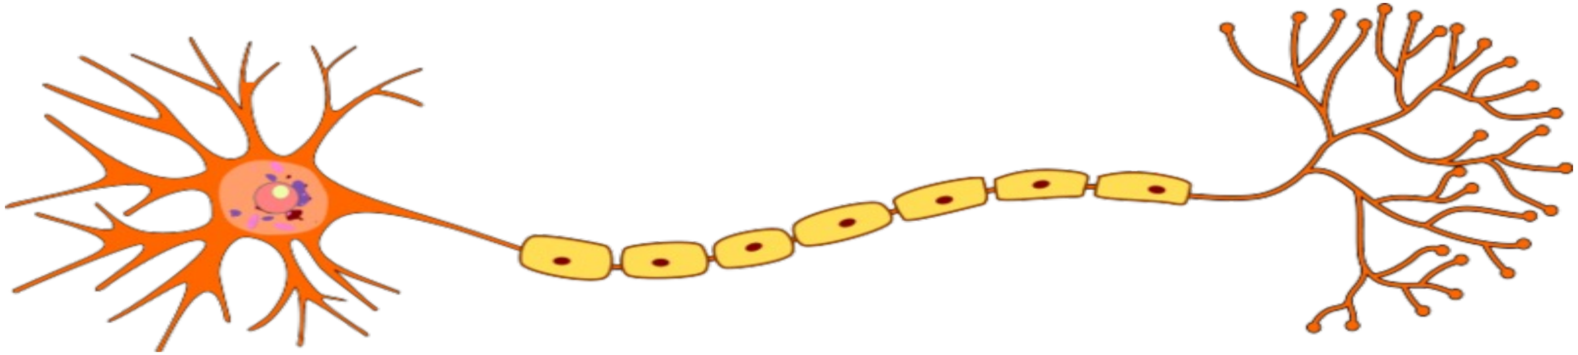

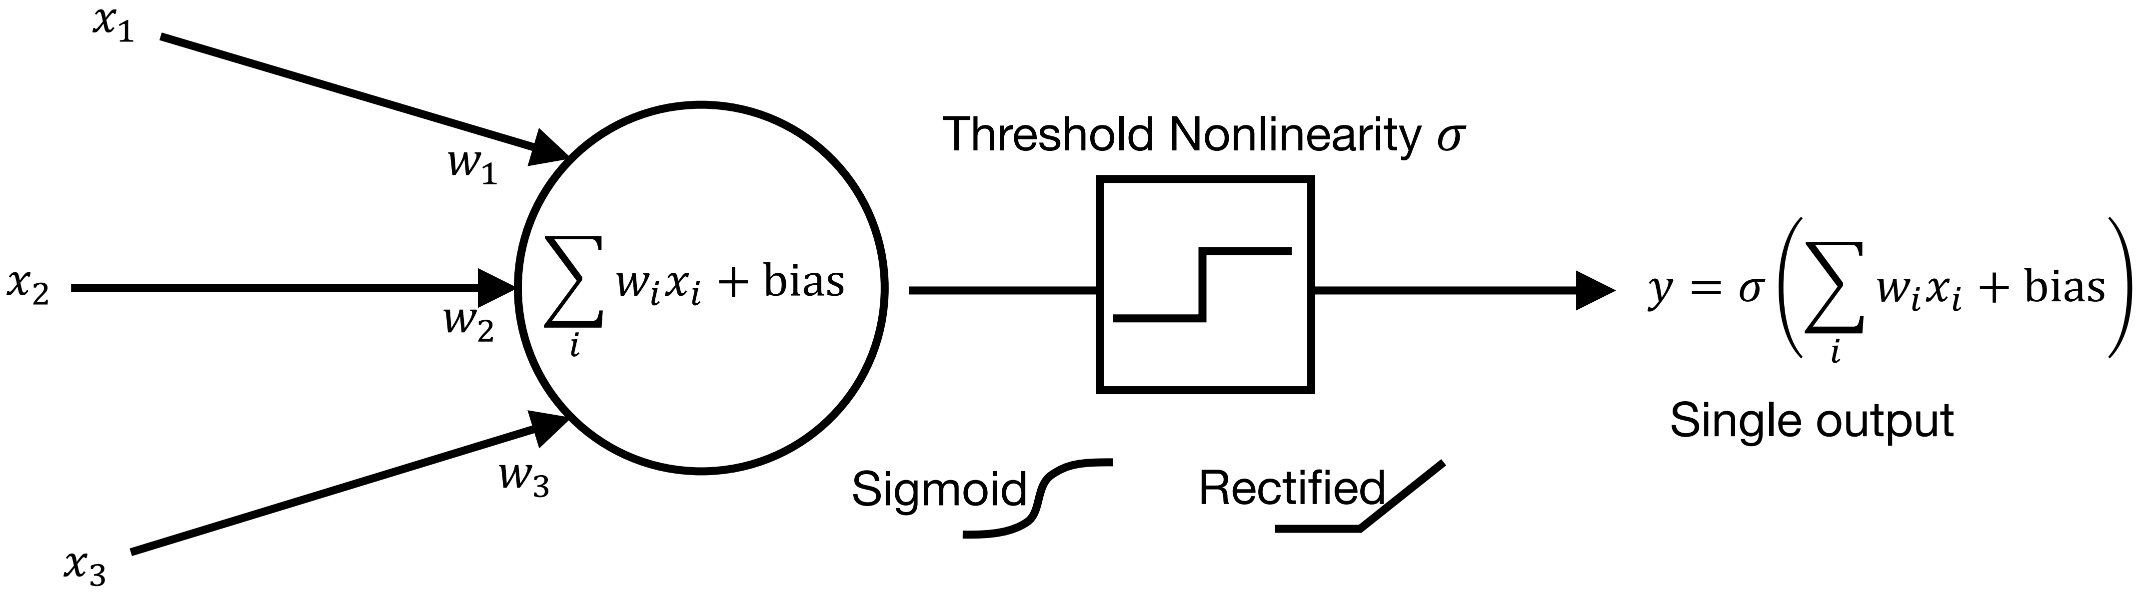

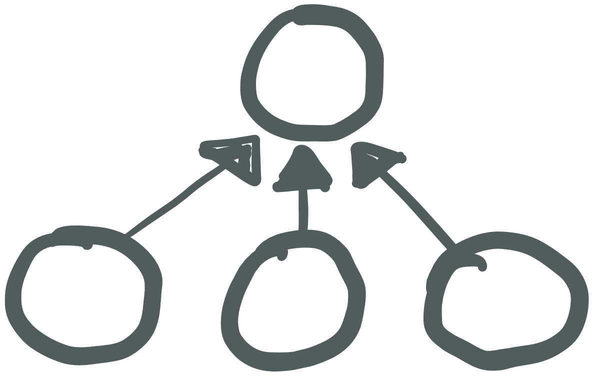

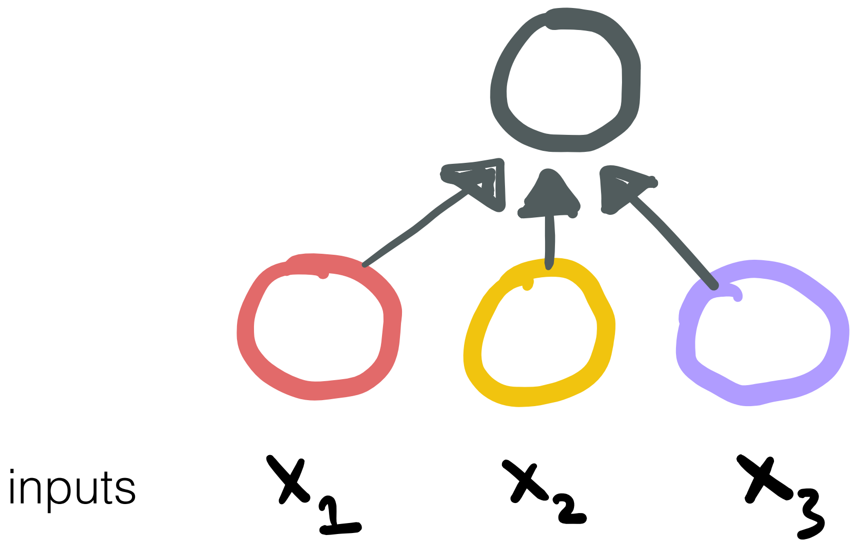
inputs
x_1
x_2
x_3

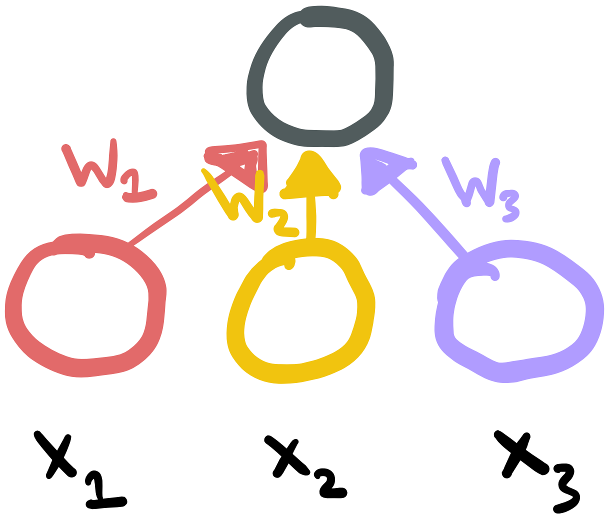

inputs

$$y = \sigma(w_1 x_1 + w_2 x_2 + w_3 x_3)$$ output

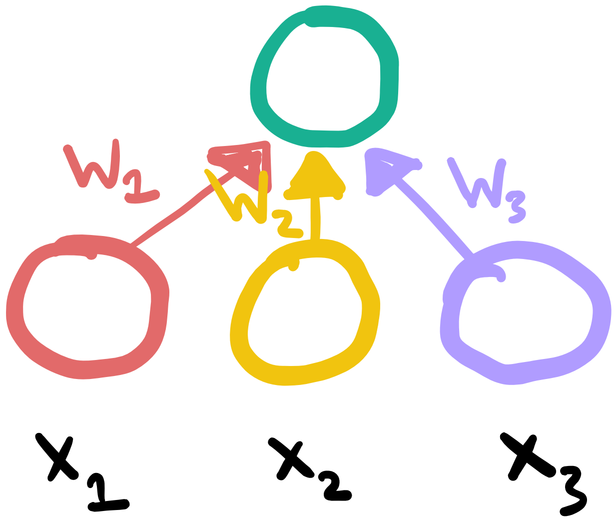

inputs

Learning

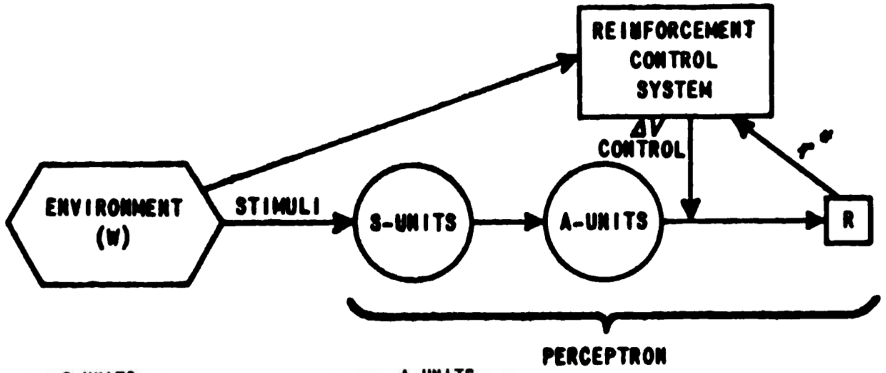

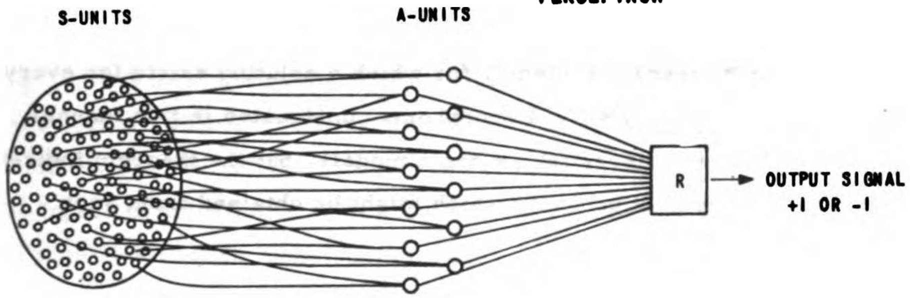

An error-corrective reinforcement system (error correction system) is a training procedure in which the magnitude of η is 0 unless the current response of the perceptron is wrong, in which case, the sign of η is determined by the sign of the error. In this system, reinforcement is 0 for a correct response, and negative (see Definition 34) for an incorrect response, or, more generally, $\eta = f(R^* - r^*)$ where R^* is the required response, r^* is the obtained response, and f is a sign-preserving monotonic function, such that $f(0) = 0$.

Rosenblatt’s 1958 Perceptron AI Boom

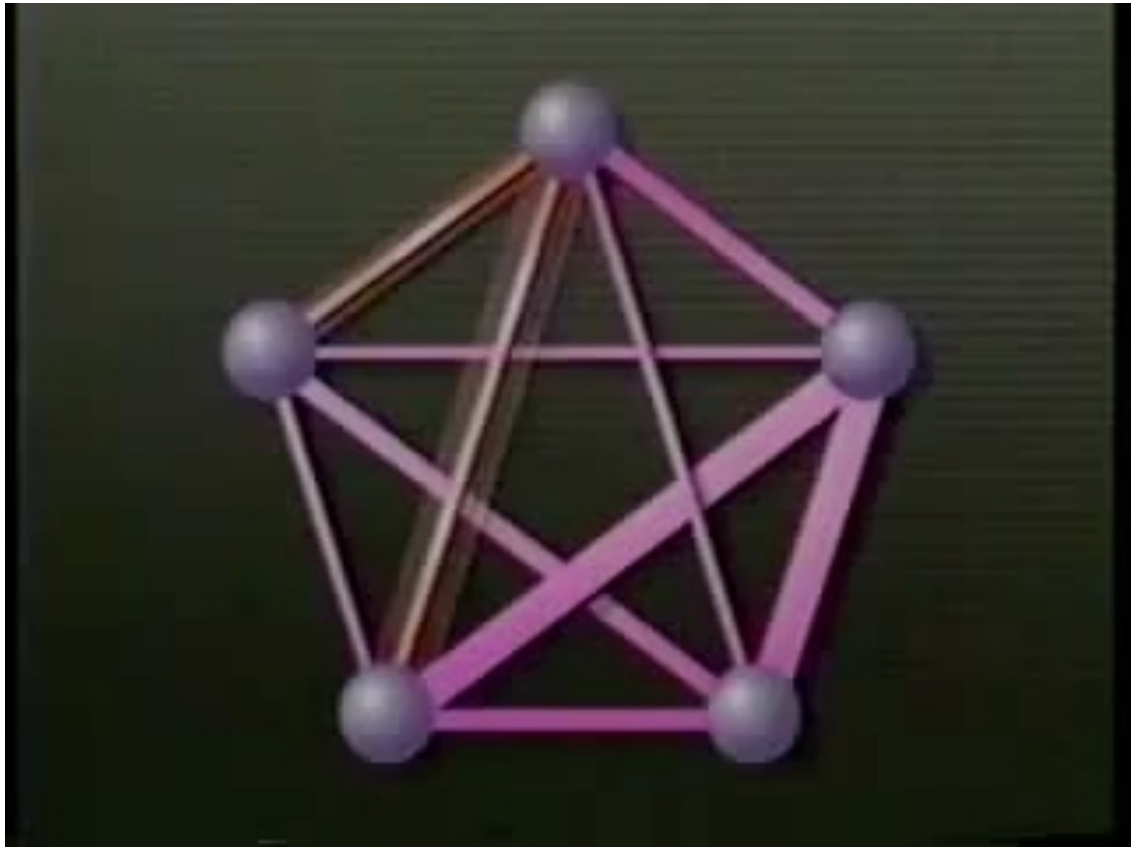

NEW NAVY DEVICE LEARNS BY DOING

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 "704" 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, de-

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 prin-

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 "704" 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.

The New York Times

Published: July 8, 1958

Dr. Rosenblatt was passing through town recently, on his way to a consultation with Dr. Yovits, and we conned him, over a cup of coffee, into a brief exegesis of their brilliant offspring. "Our success in developing the perceptron means that for the first time a non-biological object will achieve an organization of its external environment in a meaningful way," Dr. Rosenblatt said. "That's a safe definition of what the perceptron can do. My colleague disapproves of all the loose talk one hears nowadays about mechanical brains. He prefers to call our machine a self-organizing system, but, between you and me, that's precisely what any brain is."

Of what practical use, we asked, would the perceptron be? "At the moment, none whatever," Dr. Rosenblatt said cheerfully. "Someday we may find it useful to send one out into space to take in impressions for us. In these matters, you know, use follows invention.

about perceptrons." What, we asked, *wasn't* the perceptron capable of? Dr.

Rosenblatt threw up his hands. "Love," he said. "Hope. Despair. Human nature, in short. If we don't understand the human sex drive, why should we expect a machine to?"

THE NEW YORKER DECEMBER 6, 1958

Perceptron Learning Algorithm

Input

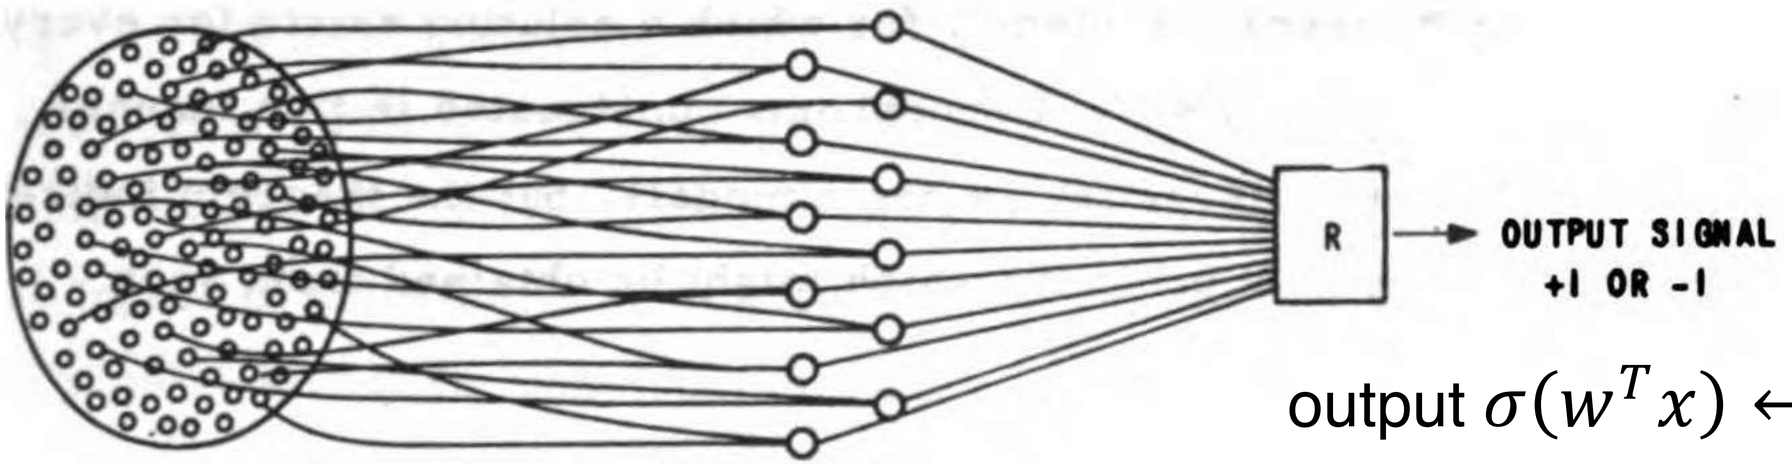

Label

output $\sigma(w^T x) \leftrightarrow$ $y = \pm 1$

feature x

learned weights w

Run on random training data

A review of supervised learning & intro to the Perceptron

(whiteboard / see notes)

Let’s see this in action …
(Perceptron notebook)

Why does this work?

Limitations of the Perceptron

A simple failure of perceptron learning

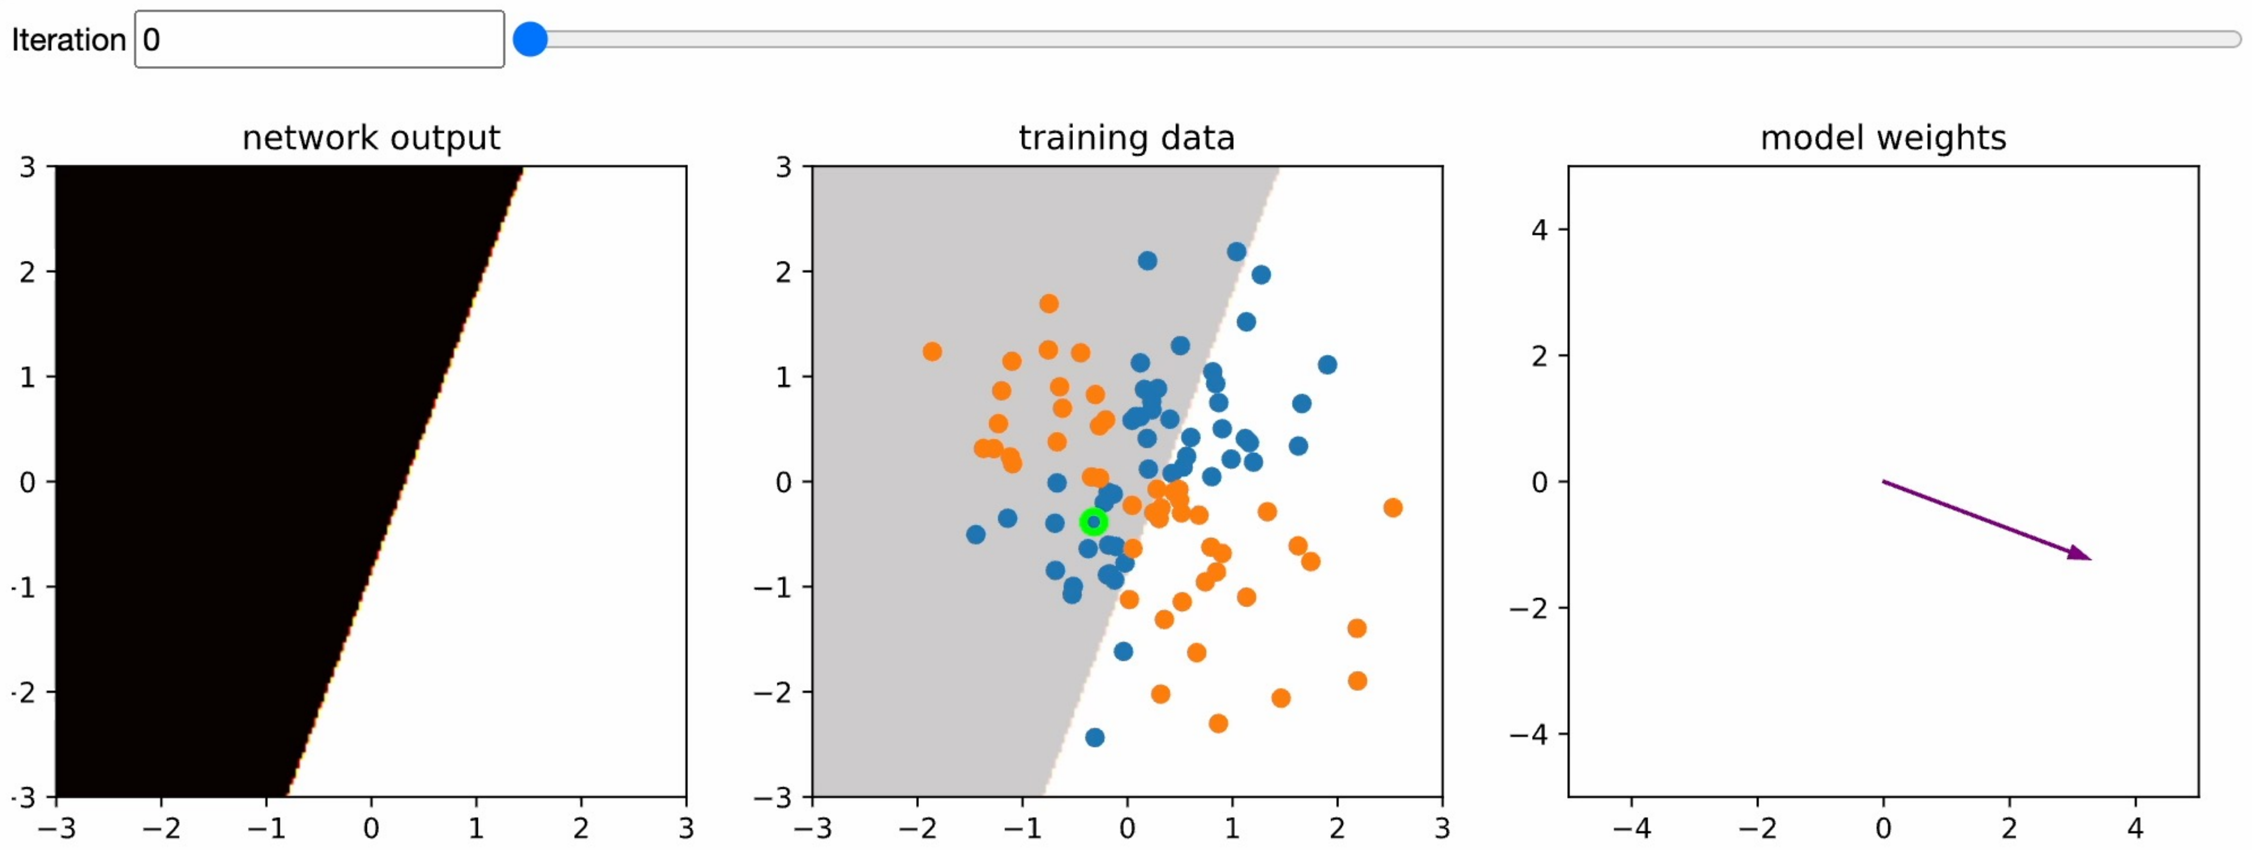

Non-linearly separable data cannot be classified by the perceptron algorithm.

Minsky and Papert prove limitations [1969]

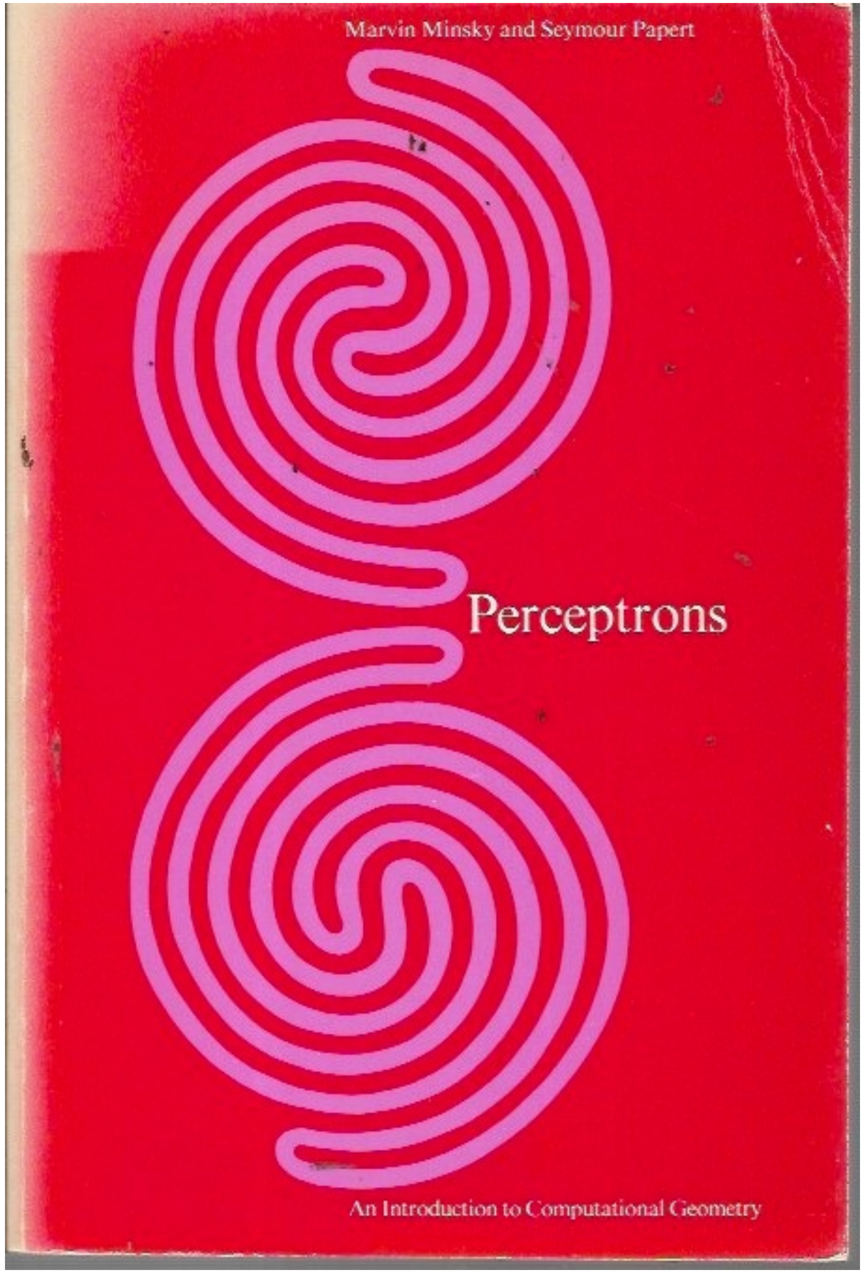

We needed **deeper** networks with non-linearities. This is what we will explore in this course.

OK now some logistics &etc.

Course aims & scope

Basics of neural networks

- Backpropagation and optimization via gradient descent
- “Classic” architectures (MLPs, ConvNets, RNNs); emphasis on Transformers and emerging models

Course aims & scope

Basics of neural networks

- Backpropagation and optimization via gradient descent
- “Classic” architectures (MLPs, ConvNets, RNNs); emphasis on Transformers and emerging models

Hands-on experience in modern toolkits

- And final projects related to **model interpretability** using the NDIF API, a Northeastern-based research infrastructure initiative!

NDIF

Northeastern University
Khoury College of Computer Sciences

NDIF

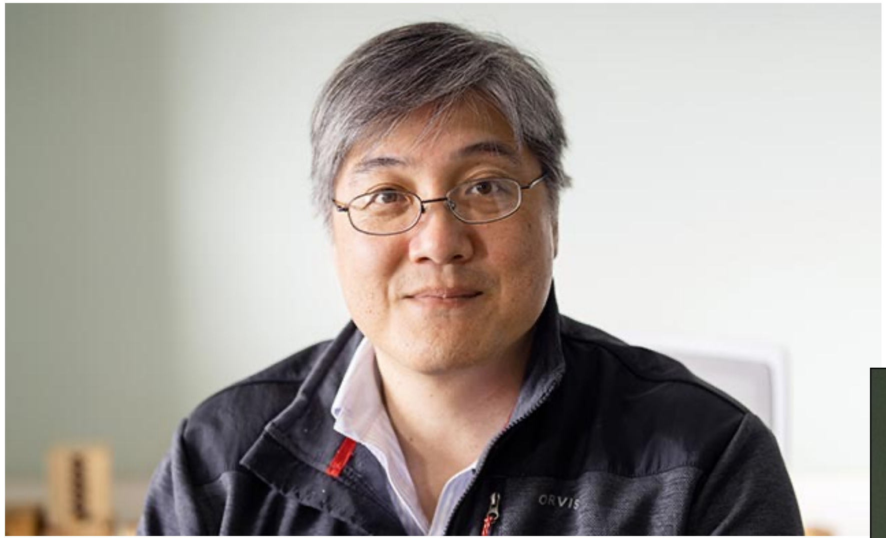

Northeastern University
Khoury College of Computer Sciences

NDIF.us

National Deep Inference Fabric

Home | Fabric | About Us | Citing | Get Involved

National Deep Inference Fabric

The National Deep Inference Fabric is a research computing project that will enable us to crack open the mysteries inside large-scale Artificial Intelligence systems.

Get Started | Our Mission

NDIF

Expectations

I will assume…

- Python knowledge
- Comfort with matrix/vector notation (brush up if rusty!)
- Basic calculus and probability (but we will introduce/review most as we go)

Resources

Homeworks submitted via Canvas
Website: https://course.ccs.neu.edu/ds4440f24/
Join piazza! https://piazza.com/northeastern/fall2024/ds4440

Grading

- Homeworks (written exercises + coding): 30%
- In-class participation (mostly exercises): 5%
- Mid-term: 25%
- Final project (interpretability work; you will use NDIF!): 40%

Questions for me?