From Perceptron to AI: The Rise of Large Language Models

Lecture 1 of 4:
How did we get here?

Dr. Michael Stobb
Thursday Forum
Coe College

February 2026

About Me

- Originally from California
- MS and PhD in Applied Mathematics
- Studied High Performance Computing with applications to biological systems
- Have been teaching Data Science courses at Coe for 7 years
- Have been using and teaching AI-like models since 2015

Planned Discussion

- **Week 1**: How did we get here?
- **Week 2**: How do LLMs work?
- **Week 3**: What are the impacts?
- **Week 4**: How can we use these tools?

Ultimate Goal Help people understand AI’s past, question its present, and prepare for its future.

How Did We Get Here?

- Why suddenly has "AI" become so important?
- Why is a huge portion of the US economy currently wrapped up in AI?
- Where did all this "AI" even come from?
- Why now and not 10 or 20 years ago?

How Did We Get Here?

- The Manhattan Project developed the Atomic Bomb
  - ~$30 Billion (in today's dollars)

- The Apollo Program to land humans on the moon
  - ~$300 Billion (in today’s dollars)

- Artificial Intelligence Investments since 2013
  - ~$1.6 **Trillion**

How Did We Get Here?

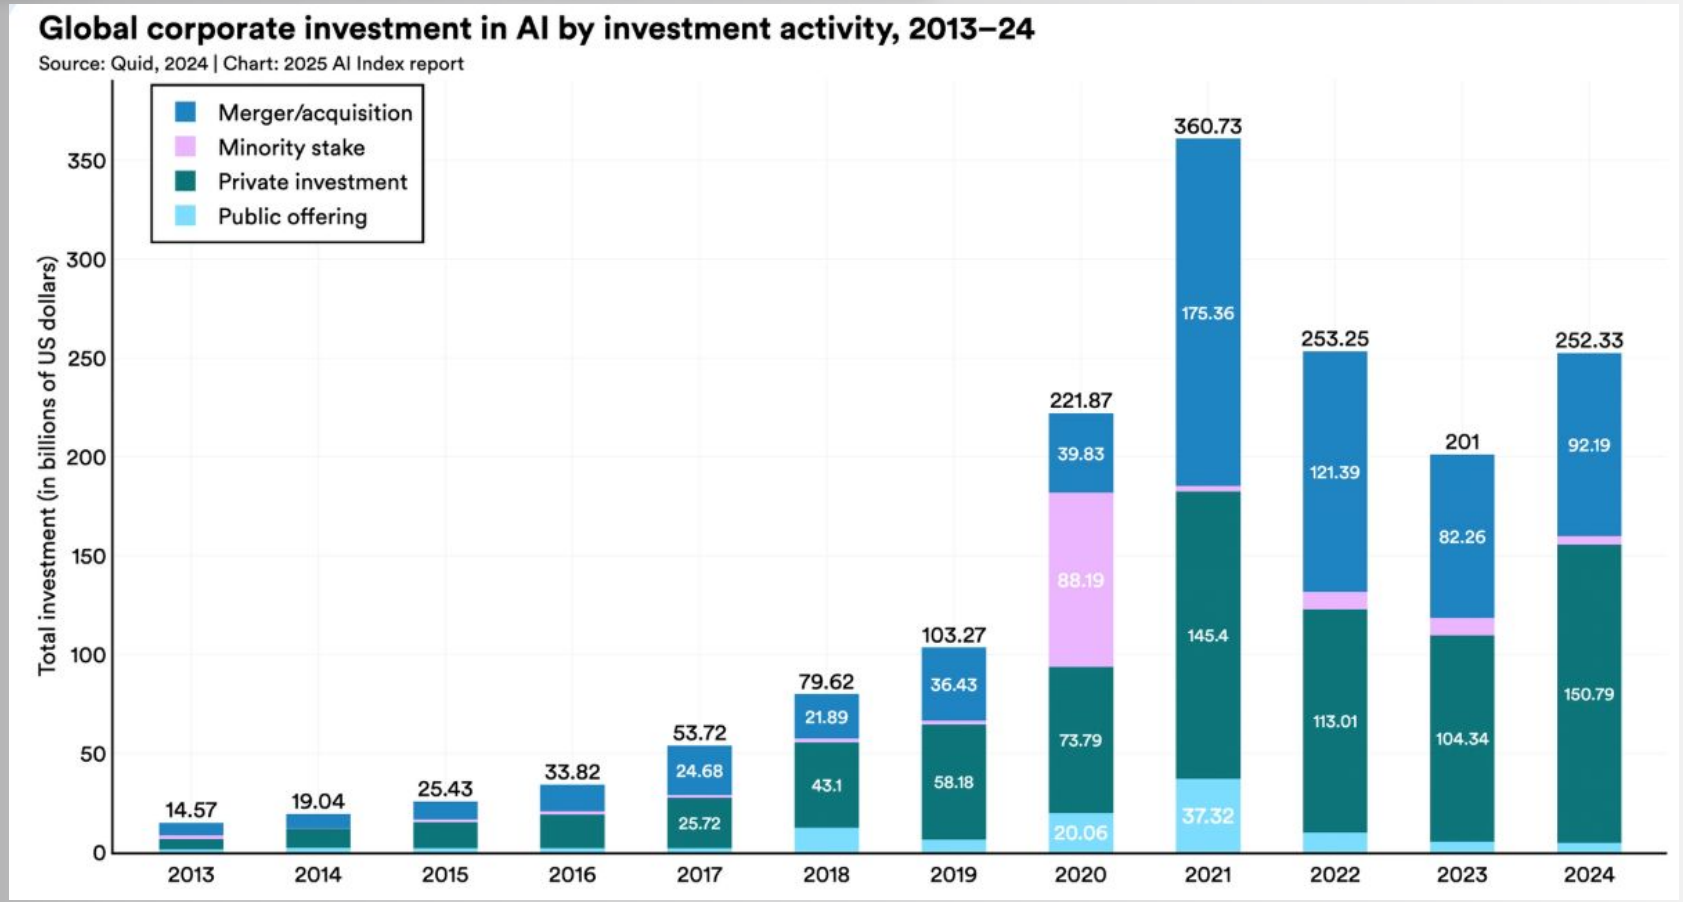

Possible Answers:

- **Computers finally got good enough**
- **We finally had enough data**
- **We had experts building better models**

However… these all existed 10 years ago too!

- Computers were more than powerful enough to make current "AI" models
- More than enough data to train models (Wikipedia was started in 2001)
- major breakthroughs in AI models came out in the 1990s and early 2000s

This image was published…

… in January 2020

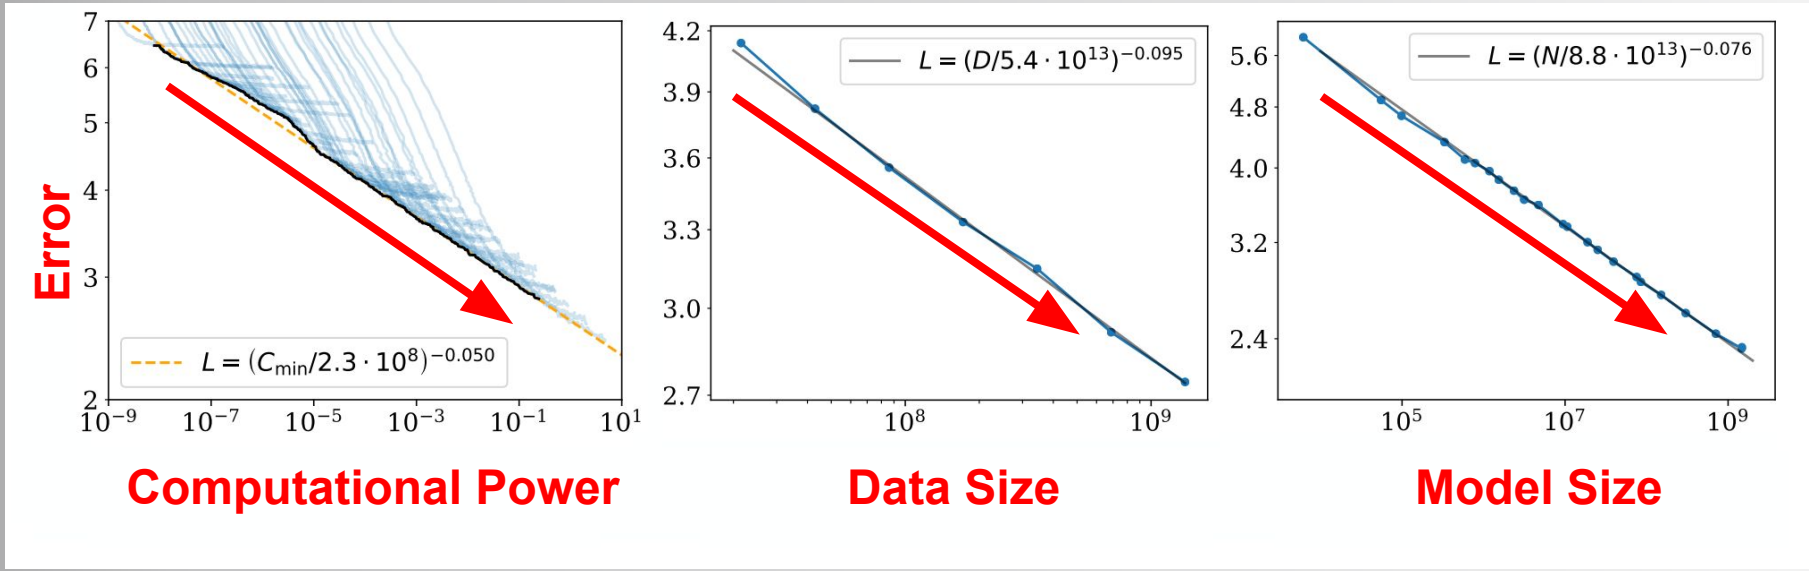

- These have been named the "***AI Scaling Laws***"

AI Scaling Laws

- Prior to these Scaling Laws, AI research was *risky*
  - Algorithms were heavily customized
  - No reliable way to get good results

- Scaling Laws demonstrated ***predictable*** results

- Investors could be “guaranteed” better models

- The AI Scaling Law paper was written by…

OpenAI

(makers of ChatGPT)

How Did We Get Here?

- Active research has been going on for years
- In the last ~20 years, we have gotten
  - Powerful computers
  - Massive amounts of data
  - More dedicated researchers
- OpenAI showed investors that AI models could be produced *reliably* in 2020
- ChatGPT was released publicly in November 2022

Where it Started

- **350 BCE - Aristotle**: Made logical arguments where the truth could be automatically computed

  Syllogism: All dogs bark. Spot is a dog. Therefore, Spot barks.

- **1600 CE - Thomas Hobbes**: Believed that intelligence was just a result of computation

  “For Reason, in this sense, is nothing but Reckoning (that is, Adding and Substracting) of the Consequences of generall names agreed upon, for the Marking and Signifying of our thoughts”

  - ***Leviathan***

Where it Started

- **1600 CE - René Descartes**: Famous "dualist", believed the body and mind were separate
  - The body is a material puppet
  - The soul is immaterial and responsible for reasoning

- **1700 CE - Gottfried Leibniz**: Believed a "Universal Language" could be created that would encompass all concepts, allowing a machine to solve any argument

Enter Turing

- **1950 - Alan Turing**: If a machine could perfectly imitate an intelligence, how could we call it otherwise?

> *I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9, to make them play the imitation game so well that an average interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning…*
>
> *I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.*
>
> - ***Alan Turing***

The Imitation Game (or “Turing Test”)

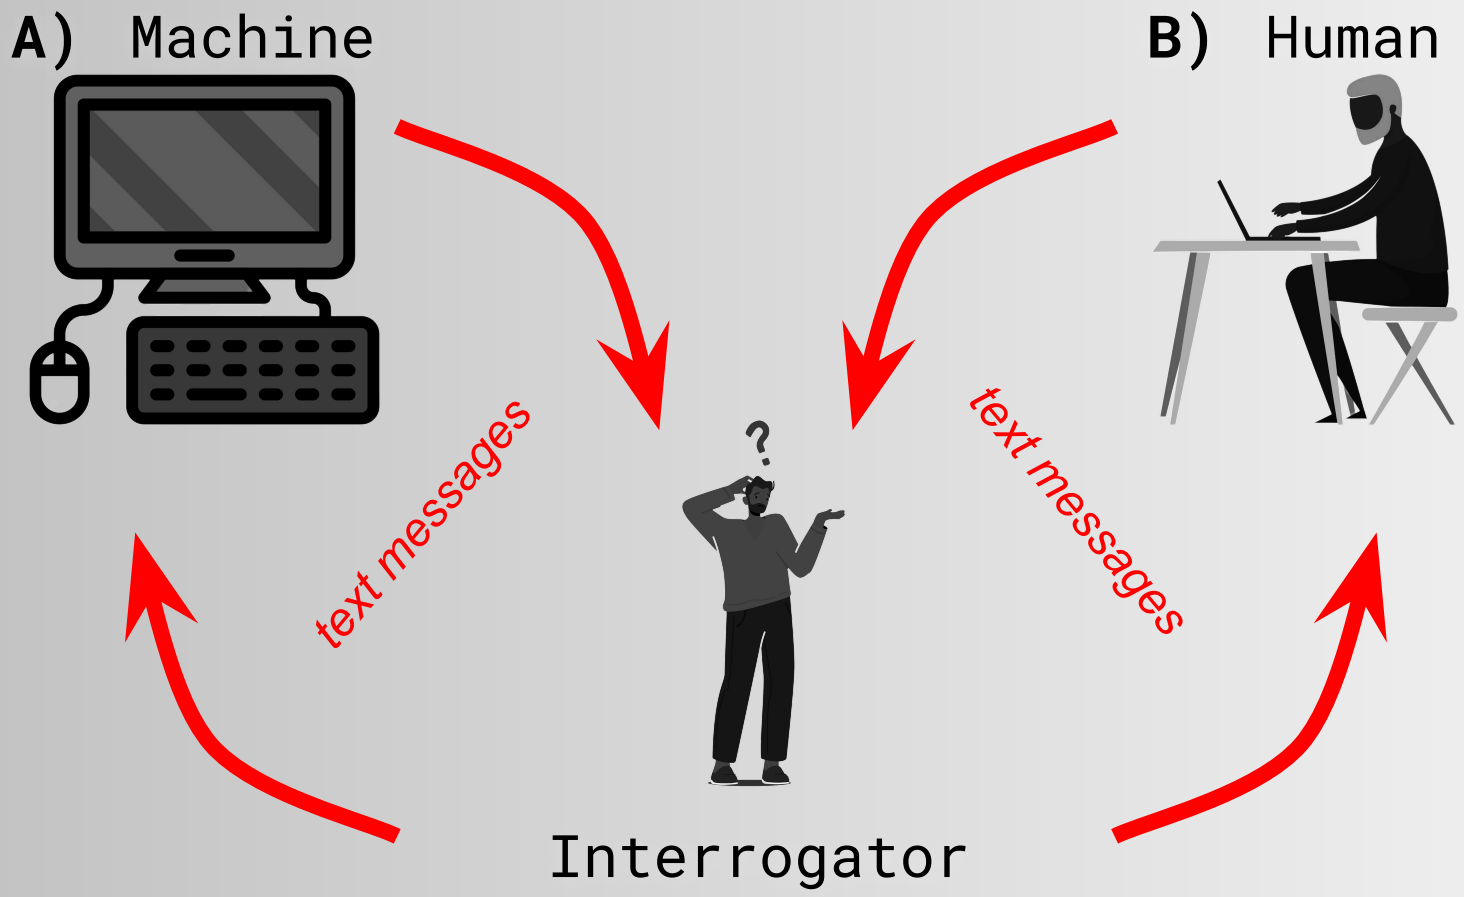

The Imitation Game (or “Turing Test”)

- Turing expected a win rate for the machine to be approximately 30%
- This test was highly influential for the development of AI systems
- Not actually implemented until the 1990s with modern "chatbots"
  - At that time, win rates for machines were 0%

The Imitation Game (or “Turing Test”)

- Not until 2014 did the win rate for machines climb above 30%
  - However, these used linguistic “tricks”
  - Their strategy was to pretend to be non-native english speaking children
- Modern models (like ChatGPT) have win rates between 50-70%

The Imitation Game (or “Turing Test”)

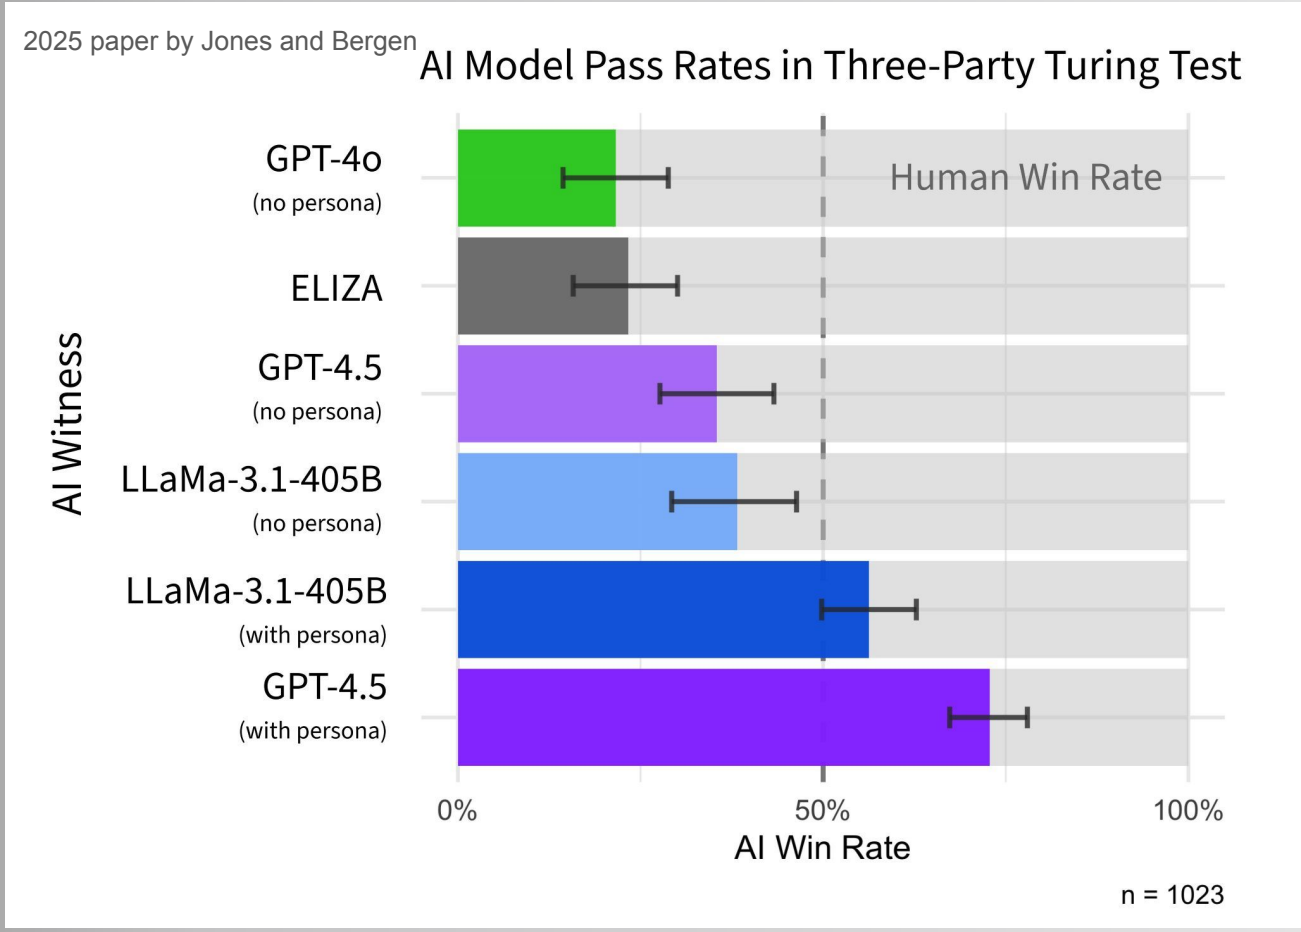

- No linguistic tricks used, just regular conversations
- Humans thought the AI models were **more** human than humans
- “Turing Test” is now fully broken as of 2025

The Turing Test was the *wrong test*

- The Turing Test is really about humanlike *deception* and not true intelligence
- The machine actually has to suppress some intelligence in order to win
- This is not a new critique:

> *The question of whether a computer can think is no more interesting than the question of whether a submarine can swim.*
>
> **- Edsger Dijkistra** (1984)

Algorithms that Learn

- “Artificial Intelligence” is rooted in the notion that *machines can learn*
- AI is often divided into two distinct branches:

Narrow AI

- Capable of learning a specific, limited task
- Does not have general cognitive abilities
- Built using a myriad of different algorithms
- Been around for decades

General AI

- Possess expansive, human-like capabilities
- Able to think, reason, and learn various subjects
- Able to transfer knowledge from one subject to another
- Brand new and highly controversial

Path to General AI

- Current AI models (like ChatGPT, Claude, Gemini) are the closest things we have to **General AI**
- These models are built on the same blueprint:

Transformer Based *Neural Networks*

- Today we seek to better understand just how Neural Networks work and function
- Their basic building block is called

the Perceptron

The Perceptron

- First described in 1943 by McCulloch and Pitts
  - Inspired by the anatomy of brain neurons

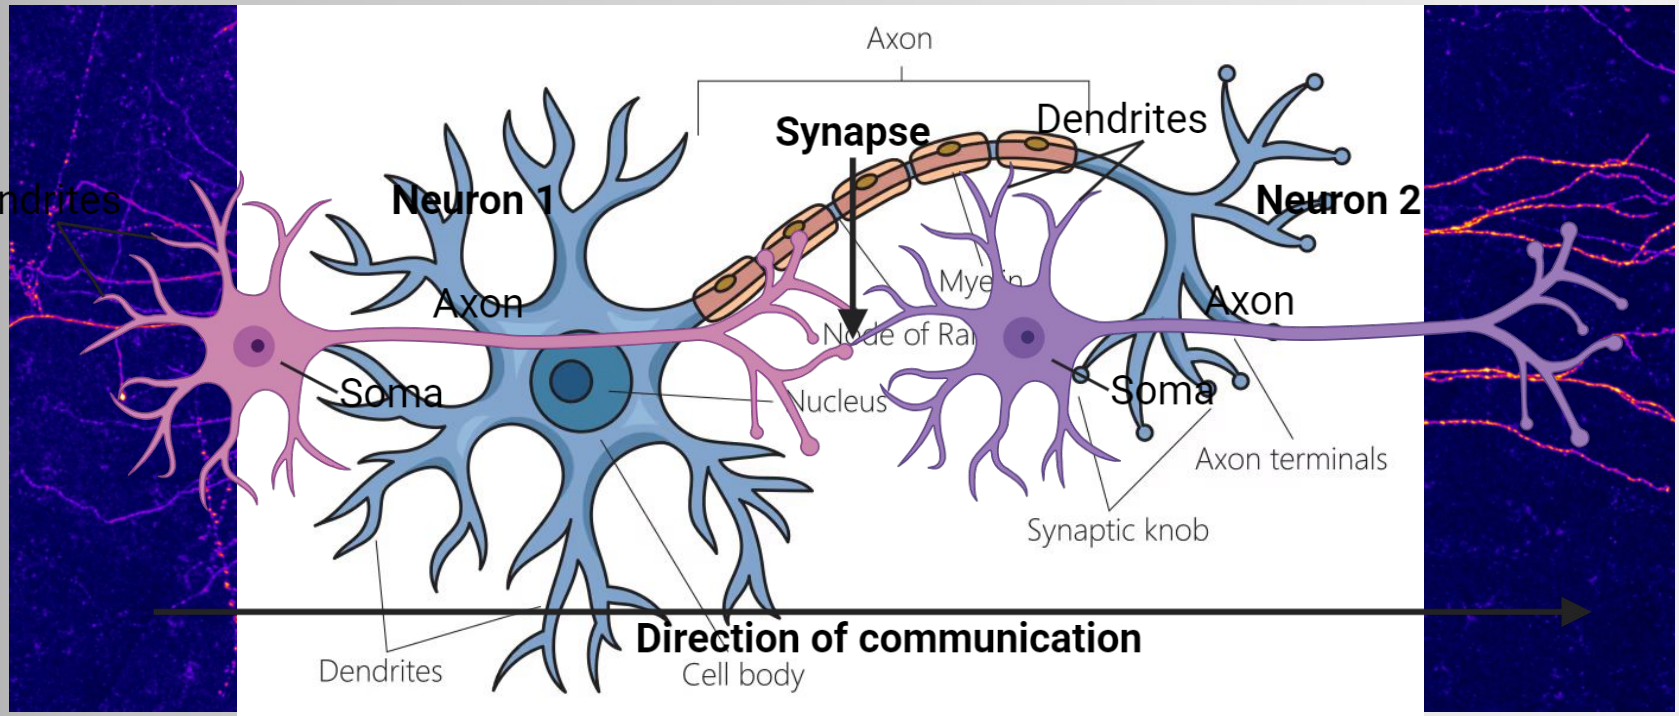

The Perceptron

- McCulloch and Pitts turned the neuron unit into a mathematical representation

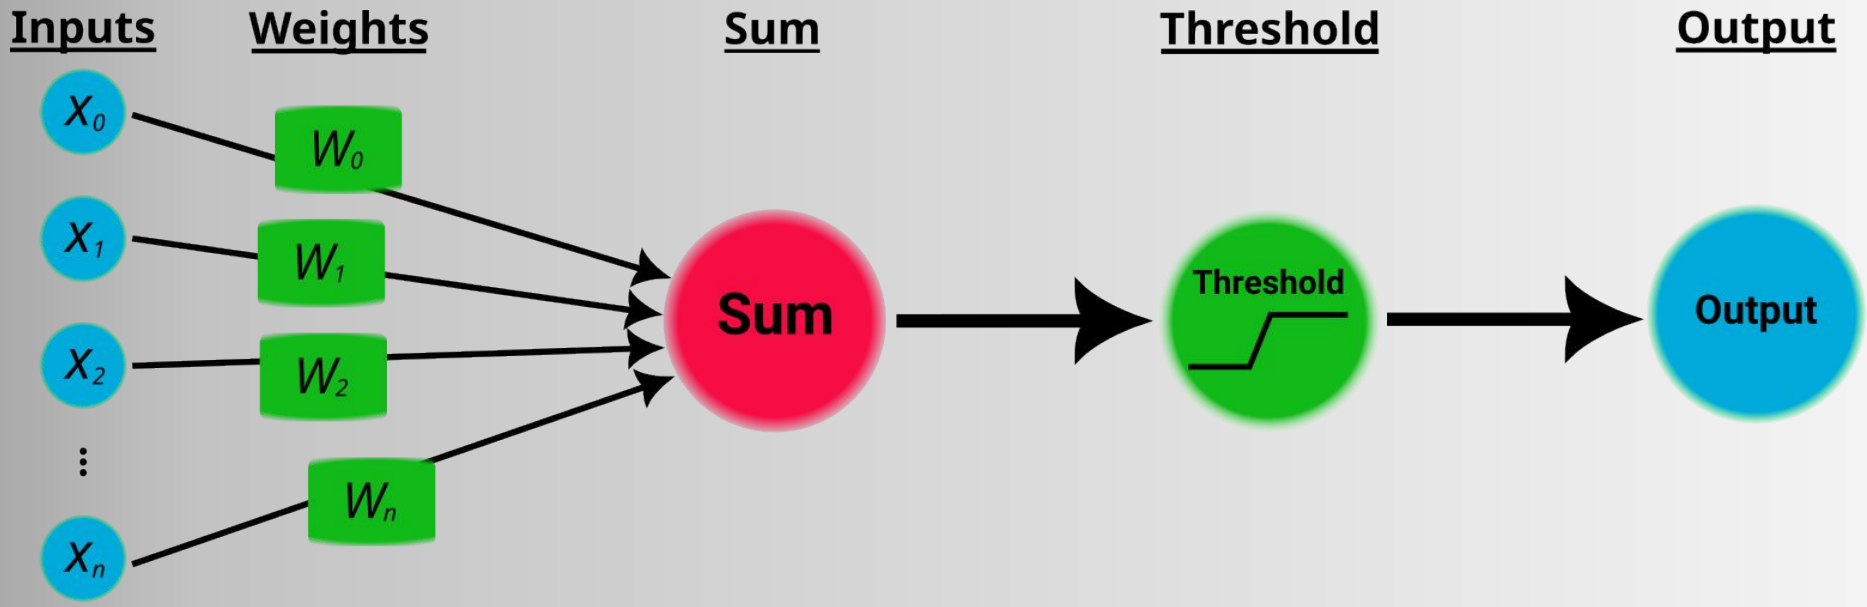

The Perceptron

- These are *artificial neurons* only
- They are just inspired by the biology
- Would be wrong to think that this is *exactly* how real biological neurons work
- Though AI researchers continue to explore biological systems for new neuron architectures!

Perceptron Example

- Let's predict credit card fraud!
  - Input 1: Amount of purchase (in $100s)
  - Input 2: Distance from home (in 100s of miles)
  - Input 3: Used the card chip? (1 if chip, 0 if no)
  - Output: 1 if fraud, 0 if **NOT** fraud

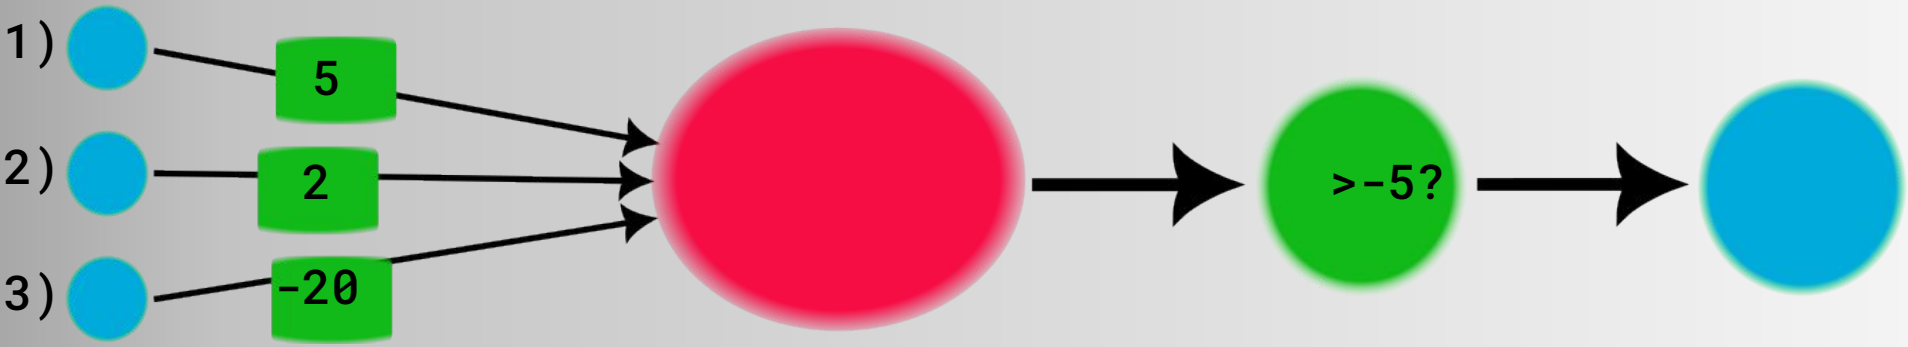

Perceptron Example

- Situation 1: Grocery Store purchase
  - Purchased $140 of groceries → 1.4
  - Store is 20 miles away from home → 0.2
  - Used the chip → 1

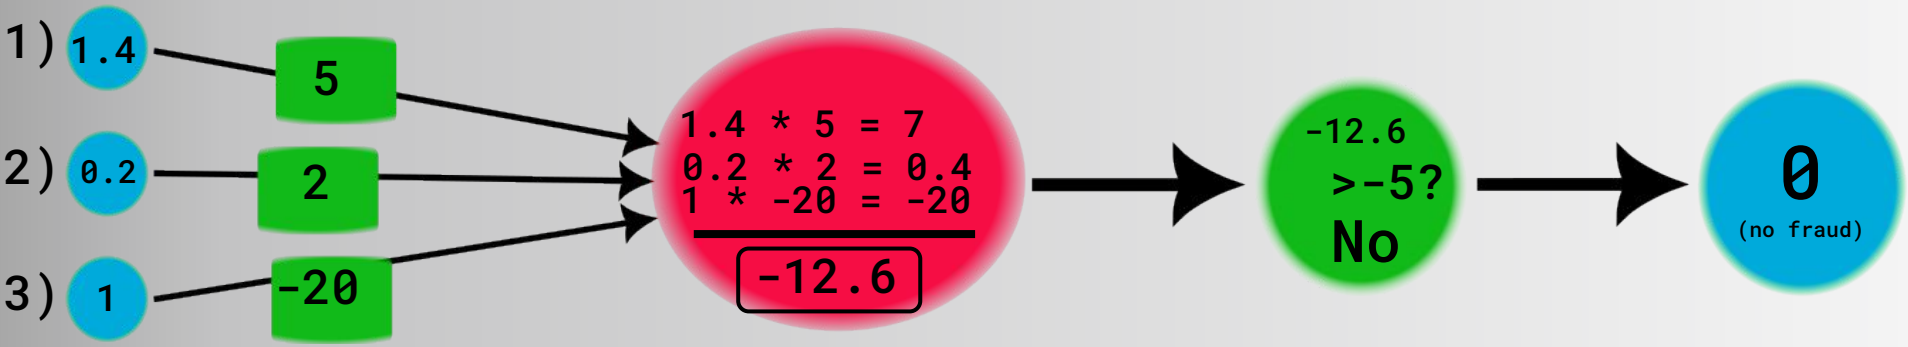

Perceptron Example

- Situation 2: Random internet purchase
  - Purchased $2000 laptop → 20
  - Purchased in Russia 5000 miles away → 50
  - Online purchase, so no chip → 0

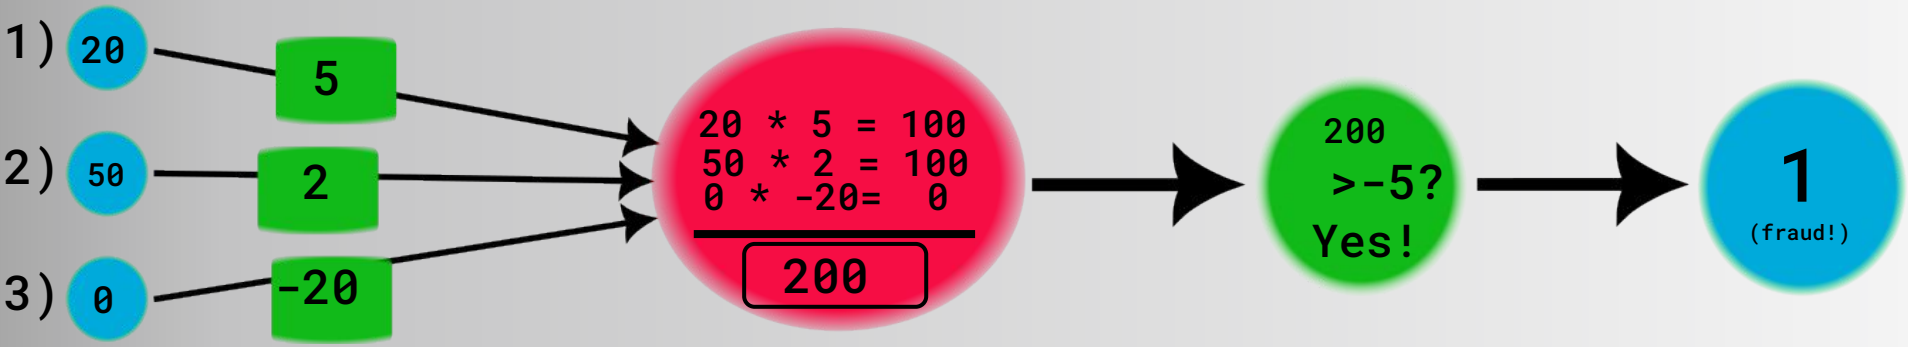

Perceptron Legacy

- Became extremely popular in the early 1960s
  - Expectations were very high, similar to today!

- While quite good, the Perceptron had a major issue
  - In 1969 it was discovered that the perceptron could only work on “linear” data
  - Anything too complex just wouldn’t work

- Effectively destroyed research into Neural Networks
  - Not until 1986 did research start back up, in part by *Geoffery Hinton* (awarded a Nobel prize for it!)

Perceptron Legacy

- A *single* Perceptron is insufficient for complex data

- But recall the biological inspiration:

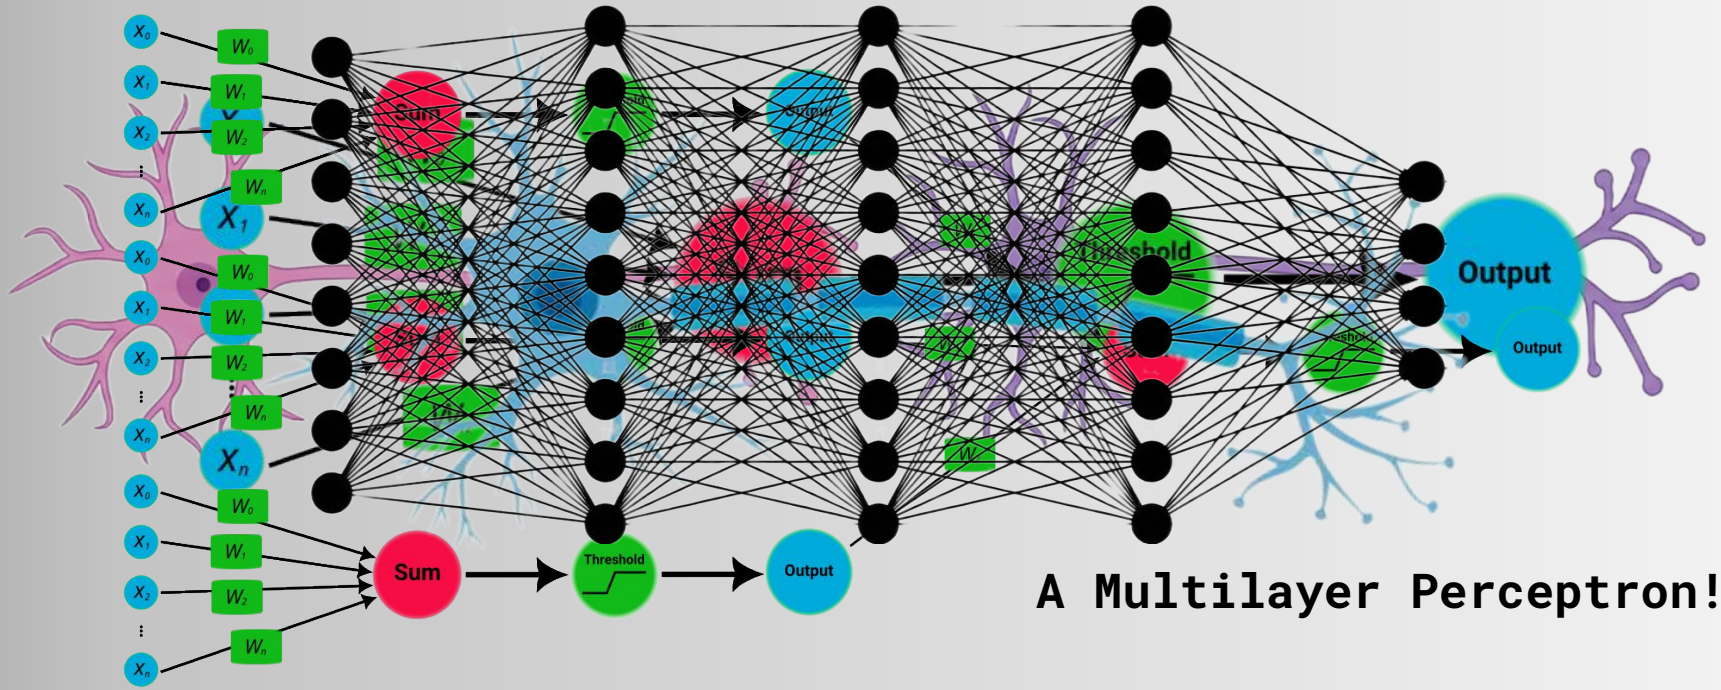

Perceptron Legacy

- The human brain has approximately:
  - 100 billion neurons
  - 100 trillion synapses (connections)
  - Relatively low connection density

- Modern AI models have approximately:
  - 100 million artificial neurons
  - 10 trillion connections
  - Much higher connection density

Next Time…

- We will explore how LLMs actually work!

Thank you!

For notes, further readings, and a full copy of the slides, just scan the QR code:

stobb.org/thursday_forum/2026_lecture_01/