DAY 11–20 · Part 2 · How We Got Here

Part 2 Review — How We Got Here

All ten days of Part 2 at a glance · short, with analogies

factnotes

No test — a review to read through

A neuron machine that learned on its own

1. In 1958, Frank Rosenblatt built the perceptron.

2. It's one artificial neuron: multiply each signal by a weight, add them up, output 1 if the sum clears a threshold, else 0.

3. The real breakthrough: when wrong, it fixed its own weights — the first machine to learn from examples, not written rules.

When hopes outrun skill, winter comes

1. In 1969, Minsky and Papert proved mathematically that one layer of neurons can't solve some problems. Money and people left — the first winter.

2. In the 1980s, expert systems with hand-written rules boomed but failed just off-script; in 1987 their special-computer market crashed — the second winter.

3. Both times, it wasn't that AI was bad — the promises ran ahead of its skill.

The day three things came together

1. In 2012, a University of Toronto team's neural network won the ImageNet photo-guessing contest with a 15.3% error rate.

2. Second place had 26.2%, so the gap was more than 10 percentage points.

3. Hardly anything was new: millions of photos, gaming graphics cards and deep neural networks came together for the first time that year.

Words started looking at each other

1. In 2017, a Google team's paper, 'Attention Is All You Need', set out the Transformer — the backbone of today's AI.

2. Before, sentences were read front to back, one word at a time; the Transformer let every word look at all the others at once.

3. Looking all at once means computing all at once, so it was a perfect match for graphics cards.

Scaled up, it copied from examples alone

1. Same Transformer backbone, just made bigger: in 2020 that gave us GPT-3, with 175 billion numbers inside.

2. The surprise: no retraining needed — put two or three examples in the question and it followed along right away.

3. From then on, 'what you tell it to do' — the prompt — started to really matter.

A chat window is what changed the world

1. ChatGPT, opened to everyone on November 30, 2022, was a GPT-family model put in a back-and-forth chat window.

2. A big part was polishing it with people's picks of the answers they liked better, so it actually got what you asked.

3. So what changed that day was less the tech than who got to use it.

Work it out before you answer

1. AI got much better at hard problems not by getting bigger, but by working through steps before answering.

2. At first, people had to say 'think step by step'; later came reasoning models trained to do it on their own.

3. The catch: longer thinking means slower answers and more text, so it's not worth it for easy questions.

Grabbing a calculator when you don't know

1. AI's main job is adding the next word, so it's weak at exact things like big multiplications or today's weather.

2. So it got tools — a calculator, search, running code — and from 2023, models chose by themselves when to grab one.

3. A good AI isn't one that knows everything — it's one that notices it doesn't know and grabs a tool.

Chop anything into pieces, line them up

1. AI didn't learn to see photos and hear sounds by getting eyes and ears.

2. Photos are cut into square grid pieces and sound into tiny time slices, then lined up in one row with the text pieces.

3. Then no new backbone is needed — just feed them into the Transformer, where pieces look at each other.

From using one to picking one

1. In late 2022 you could count good chat AIs on one hand; now dozens of groups in the US, Europe, China and Korea make their own.

2. The biggest change: open models, whose whole bundle of numbers is released for anyone to download and run on their own gear.

3. So now you don't just use one AI — you pick, weighing price, speed and what each one is good at.

Ten days, one road.
From one neuron to everyone's models.

In Part 3: what you can do with all this — starting with summarizing and translation, hands-on, one at a time.