

Name _____ Class _____ Date _____

DAY 11

Where did AI start? — 1958: the neuron that learns

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Weights and threshold

Try it — move the weights to clear the threshold

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What does one perceptron (artificial neuron) do, in one line?

- ① Memorizes whole sentences and repeats them
- ② Multiplies by weights, adds up, 1 past threshold
- ③ Searches the internet for the answer
- ④ Follows rules people wrote, one by one

Q2 What made the 1958 perceptron most different from earlier machines?

- ① It computed much faster
- ② It didn't use electricity
- ③ It fixed its own weights when wrong
- ④ It could make pictures and music

Q3 What else do we call the perceptron's 'weights' in today's AI?

- ① Tokens
- ② Parameters — from DAY 05
- ③ Prompts
- ④ Temperature

The takeaway in my own words

"Today's AI is really just that one 1958 neuron, stacked up many, many times."

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 12

Why did AI freeze over twice? — Two AI winters

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Put it in order

Try it — put the winters in order

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What did Minsky and Papert's 1969 book prove?

- ① Computers calculate faster than people
- ② Some problems a one-layer perceptron can't solve
- ③ Bigger neural nets are always better
- ④ You should never give AI written rules

Q2 What was the biggest reason 1980s expert systems collapsed?

- ① They used too much electricity
- ② Rules were hand-written, so new cases broke them
- ③ There was no internet yet
- ④ Nobody was interested

Q3 What do both AI winters teach us?

- ① AI just doesn't work
- ② Overpromising brings winter
- ③ Money alone always wins
- ④ New theories bring winter

The takeaway in my own words

“Winter didn't come because AI was bad — it came because the promises got ahead of the skills.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 13

One photo contest flipped the whole game — 2012: deep learning comes back

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Before & after

Try it — before and after 2012

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 Why was the 2012 ImageNet result such a shock?

- ① The winner hit 0% error, a first ever
- ② Winner 15.3% vs. 2nd 26.2%: a 10+ point gap
- ③ People entered the contest for the first time
- ④ Only one team entered that year

Q2 What's the biggest reason deep learning only took off in 2012?

- ① Neural networks were invented that year
- ② Photos, graphics cards, deep nets finally met
- ③ The internet was invented that year
- ④ The government made AI research mandatory

Q3 Why were gaming graphics cards (GPUs) used for deep learning?

- ① They make screens look nice
- ② They do tons of math at once
- ③ The internet got faster
- ④ They have lots of storage

The takeaway in my own words

“What ended the winter wasn't a new idea — it was the day data, computing and depth came together.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 14

The words in a sentence started looking at each other — 2017: the Transformer

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Split it up

Try it — What does 'it' point to? Watch the words look at each other

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What's the biggest difference between the Transformer and the way before it?

- ① Words in a sentence look at each other at once
- ② It memorizes words faster
- ③ It searches the internet to answer
- ④ People type in rules by hand

Q2 Why did 'looking at once' matter so much?

- ① It made sentences prettier
- ② All-at-once math is perfect for graphics cards
- ③ It used less memory
- ④ It was easier for people to read

Q3 In 'it was too tired,' how does attention know 'it' means the animal?

- ① It's the first word
- ② A dictionary says so
- ③ Scores how well words fit
- ④ It picks at random

The takeaway in my own words

“Once words could look at each other, AI started understanding whole sentences.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 15

Why did just making it bigger give it new skills? — 2020: GPT-3

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Number sense

Try it — A feel for digits — how big did GPT get?

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What made GPT-3 most different from the models before it?

- ① It searched the internet in real time
- ② Show it a few examples, and it did a new task
- ③ People put in the rules one by one
- ④ It switched its own power on and off

Q2 About how many numbers (parameters) does GPT-3 have?

- ① About 170 million
- ② About 1.7 billion
- ③ About 175 billion
- ④ About 17.5 billion

Q3 What do you call what you give the model, examples tucked in and all?

- ① An update
- ② A download
- ③ A refresh
- ④ A prompt

The takeaway in my own words

“Same design, just more size and more reading, and it learned to copy from a few examples.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 16

A chat window, not the model, changed the world? — 2022: ChatGPT

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Before & after

Try it — 2020 vs. 2022: what changed?

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What's the biggest reason ChatGPT surprised the world?

- ① It invented a whole new design
- ② It was faster than the internet
- ③ Anyone could use it just by talking to it
- ④ It grew to 1 trillion numbers

Q2 What did polishing with human feedback show?

- ① The bigger model always won
- ② Answers always got longer
- ③ Results were the same at any size
- ④ People preferred a 1.3B model over a 175B one

Q3 What year did ChatGPT first come out?

- ① 2022
- ② 2020
- ③ 2017
- ④ 2025

The takeaway in my own words

“Not a bigger model, but a window you could talk to, made AI something for everyone.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 17

Why write out the work instead of answering right away? — Thinking before answering

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Put it in order

Try it — Put the steps of the work in order

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What makes a reasoning model different from earlier models?

- ① It memorizes whole answers
- ② It makes its answers shorter
- ③ It always searches the internet
- ④ It works through the steps before answering

Q2 What happened on the same math competition problems?

- ① It went up from 12% to 74%
- ② It dropped from 74% to 12%
- ③ Both got 100%
- ④ There was no difference

Q3 What does thinking first cost?

- ① Accuracy
- ② Time and words written
- ③ Internet speed
- ④ Screen brightness

The takeaway in my own words

“Once it wrote out its work instead of answering right away, it started getting right what it used to get wrong.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 18

Is grabbing a calculator when you don't know the smarter move? — AI that uses tools

1 • Think first

Before the lesson, write down what you think the answer is.

2 • Experiment • Pick one

Try it — Answer directly, or grab a tool?

What did you change in the experiment, and what happened?

3 • Quiz

Circle the right answer.

Q1 What does it mean for an AI to 'use tools'?

- ① It memorizes whole answers
- ② It calls a calculator, search or code on its own
- ③ It switches to a bigger model
- ④ It makes its answers shorter

Q2 Which of these questions is better handled with a tool?

- ① What's 'cat' in Spanish?
- ② What's the opposite of 'big'?
- ③ What's 847,293 times 6,184?
- ④ What's 3 plus 4?

Q3 Who decides whether to call a tool?

- ① The power button
- ② The user, every time
- ③ The internet line
- ④ The model itself

The takeaway in my own words

“Noticing you don't know and grabbing a tool is a kind of smart, too.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 19

How do photos and sounds get to stand next to words? — AI that sees and hears

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Split it up

Try it — Split text, photos and sound into pieces

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What actually happens when you put a photo into an AI?

- ① The photo gets translated into text first
- ② It's turned black and white
- ③ It's cut into a grid and the pieces lined up
- ④ It's made twice as big

Q2 Why can text, photos and sound all go through one design?

- ① The model was given eyes and ears
- ② Photos were text all along
- ③ Sound has no meaning
- ④ All three become pieces in the same row

Q3 How does sound get split up?

- ① Short time slices
- ② One letter at a time
- ③ By color
- ④ Sorted by length

The takeaway in my own words

“Chop photos and sounds into pieces and line them up, and they can stand right where words do.”

Rewrite this line in your own words.

Name _____ Class _____ Date _____

DAY 20

Why are there so many AIs you can use now? — Everyone's model race

1 · Think first

Before the lesson, write down what you think the answer is.

2 · Experiment · Before & after

Try it — End of 2022 vs. now: what changed?

What did you change in the experiment, and what happened?

3 · Quiz

Circle the right answer.

Q1 What's the biggest difference between the end of 2022 and now?

- | | |
|--------------------------------|-----------------------------------|
| ① Prices to use it went way up | ② All the models merged into one |
| ③ Nobody makes them anymore | ④ Way more places make models now |

Q2 What does 'open model' mean?

- | | |
|--|--|
| ① Anyone can download its numbers and run it | ② A model with no ads |
| ③ A model that always gives the right answer | ④ A model that searches without the internet |

Q3 When more places make models, what happens for users?

- | | |
|----------------------|------------------------|
| ① Answers disappear | ② You get to choose |
| ③ The power goes out | ④ The text gets bigger |

The takeaway in my own words

“As the places making it went from a few to dozens, AI went from something you use to something you choose.”

Rewrite this line in your own words.

Answer key · Part 2: How We Got Here

DAY 11 · Where did AI start? — 1958: the neuron that learns

Q1 · Answer ② Multiplies by weights, adds up, 1 past threshold

Multiply, add, compare with the threshold. That's all three — today's AI just does this math an enormous number of times.

Q2 · Answer ③ It fixed its own weights when wrong

Instead of people writing every rule, it nudged its weights by however much it was wrong. That was the start of machines that learn from examples — the root of the 'training' you saw on DAY 04.

Q3 · Answer ② Parameters — from DAY 05

Parameters. On DAY 05, when we counted how many parameters a model has, these weights are exactly what we were counting. A big model just has hundreds of billions of them.

Today in one line · Today's AI is really just that one 1958 neuron, stacked up many, many times.

DAY 12 · Why did AI freeze over twice? — Two AI winters

Q1 · Answer ② Some problems a one-layer perceptron can't solve

It nailed down, with math, that some problems — like XOR — can't be split with a single line. This one book set off **the first winter**.

Q2 · Answer ② Rules were hand-written, so new cases broke them

It hadn't learned anything — **it had been dictated to**. Without the 'learning that fixes itself' from DAY 04, it can only do exactly what it was told.

Q3 · Answer ② Overpromising brings winter

Both times, the tech didn't go backward — people left because **promises that had grown too fast** weren't kept. It's a yardstick you can use on today's AI news too.

Today in one line · Winter didn't come because AI was bad — it came because the promises got ahead of the skills.

DAY 13 · One photo contest flipped the whole game — 2012: deep learning comes back

Q1 · Answer ② Winner 15.3% vs. 2nd 26.2%: a 10+ point gap

In a contest that had been improving by 1–2 percentage points a year, one team pulled ahead on its own by more than 10 points. That's why everyone sees that year as the turning point.

Q2 · Answer ② Photos, graphics cards, deep nets finally met

Neural networks, and the way to stack and fix layers, already existed by 1986 (DAY 12). What was missing was data and computing power, and in 2012 both were filled in.

Q3 · Answer ② They do tons of math at once

A graphics card is built to calculate tons of similar screen dots at the same time. Neural network math has the same shape, so it could be used as-is.

Today in one line · What ended the winter wasn't a new idea — it was the day data, computing and depth came together.

DAY 14 · The words in a sentence started looking at each other — 2017: the Transformer

Q1 · Answer ① Words in a sentence look at each other at once

Instead of reading in order from the start, every word looks at every other word at once. That's 'attention' — and it's why the paper was called 'Attention Is All You Need.'

Q2 · Answer ② All-at-once math is perfect for graphics cards

Reading in order, you have to finish one word before you can look at the next. Looking at once means one big calculation, and the graphics cards you met yesterday are great at exactly that.

Q3 · Answer ③ Scores how well words fit

Each word scores 'which words should I look at?' What goes with 'tired' is the animal, not the street, so the scores pile up there.

Today in one line · Once words could look at each other, AI started understanding whole sentences.

DAY 15 · Why did just making it bigger give it new skills? — 2020: GPT-3

Q1 · Answer ② Show it a few examples, and it did a new task

With no retraining, it just looked at the examples in the question and followed them. That's 'few-shot' learning — which is where the paper's title came from.

Q2 · Answer ③ About 175 billion

It's 175 billion. GPT-2 was about 1.5 billion, so it grew more than 100 times.

Q3 · Answer ④ A prompt

Everything you hand to a model is called the prompt. Starting with GPT-3, how you wrote it could change the results a lot.

Today in one line · Same design, just more size and more reading, and it learned to copy from a few examples.

DAY 16 · A chat window, not the model, changed the world? — 2022: ChatGPT

Q1 · Answer ③ Anyone could use it just by talking to it

The model family wasn't that different from before. What changed was a window where you could just talk to it, even if you didn't know how to add examples.

Q2 · Answer ④ People preferred a 1.3B model over a 175B one

When people's picks were collected and used for polishing, people liked a much smaller model's answers better. A sign that size isn't everything.

Q3 · Answer ① 2022

November 30, 2022. 2017 was the Transformer, and 2020 was GPT-3.

Today in one line · Not a bigger model, but a window you could talk to, made AI something for everyone.

DAY 17 · Why write out the work instead of answering right away? — Thinking before answering

Q1 · Answer ④ It works through the steps before answering

Instead of answering straight away, it writes out its work and checks the in-between answers. That's why it makes fewer mistakes on hard problems.

Q2 · Answer ① It went up from 12% to 74%

The model that answered right away got 12%, and the one that thinks first got 74%. It was specially trained to write out its work before answering, and that's why the gap got so big.

Q3 · Answer ② Time and words written

The longer it thinks, the slower the answer and the more text it writes. So for easy questions, it doesn't bother with this.

Today in one line · Once it wrote out its work instead of answering right away, it started getting right what it used to get wrong.

DAY 18 · Is grabbing a calculator when you don't know the smarter move? — AI that uses tools

Q1 · Answer ② It calls a calculator, search or code on its own

When the model says 'I need this tool now,' the tool runs outside, and the model takes the result and finishes the answer.

Q2 · Answer ③ What's 847,293 times 6,184?

Big multiplications almost always come out wrong when you're just stringing words together. The other three are things it already knows.

Q3 · Answer ④ The model itself

It used to take a person telling it, but now the model looks at the question and decides for itself. That's today's key idea.

Today in one line · Noticing you don't know and grabbing a tool is a kind of smart, too.

DAY 19 · How do photos and sounds get to stand next to words? — AI that sees and hears

Q1 · Answer ③ It's cut into a grid and the pieces lined up

It's cut up like a checkerboard, and each piece gets a number for which square it is. Then it looks just like word pieces.

Q2 · Answer ④ All three become pieces in the same row

Once cut, all three are 'pieces in a row.' So the Transformer, which lets pieces look at each other, works as is.

Q3 · Answer ① Short time slices

Sound flows along in time, so it gets chopped into very short slices of time. Each slice is one piece.

Today in one line · Chop photos and sounds into pieces and line them up, and they can stand right where words do.

DAY 20 · Why are there so many AIs you can use now? — Everyone's model race

Q1 · Answer ④ Way more places make models now

From a handful of places to dozens. With everyone chasing each other, prices went down and performance went up.

Q2 · Answer ① Anyone can download its numbers and run it

It's a way of releasing the whole bundle of numbers inside a model. You can download it, run it on your own machine and modify it.

Q3 · Answer ② You get to choose

Price, speed and strengths all differ, so you pick one to fit the job. But the job of judging what to trust comes with it.

Today in one line · As the places making it went from a few to dozens, AI went from something you use to something you choose.
