DAY 51 · Part 5 · Looking Inside

Can you teach a fully trained AI one extra job? — Extra training (fine-tuning)

Five-Minute AI Lessons · Oct 27, 2026

You don't raise a grown-up model all over again, you just polish it a little, as much as you need.

In a nutshell

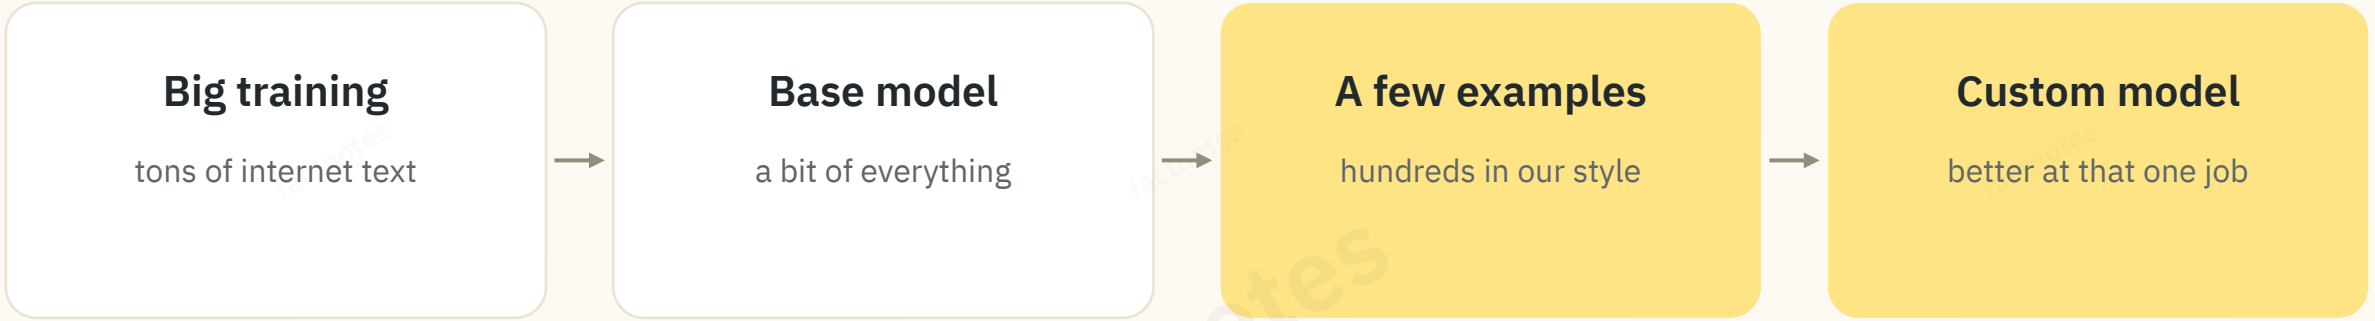
1. You can. You take a **base model** that already went through its big round of training and train it a little more on just the job you need. That's called fine-tuning.

2. It's not learning everything again from scratch. It leaves almost all the numbers inside as they are and **changes only a tiny part** — a method called LoRA cut the numbers it needs to change by 10,000 times.

3. But it's good at some things and not others. Extra training is strong at changing **tone and format**, while yesterday's open-book answers are what you use to add new facts.

In one picture

Raising one AI in two stages



The first two boxes take thousands of graphics cards and months. The last two finish much faster and cheaper by comparison.

Key idea: you don't raise it again, you polish a model that's already grown up

Experiment · Before & after

Same model, before and after a little extra training

Matching a tone

Before

Ask 'What's photosynthesis?' and it answers long and stiff, like a textbook. It's not wrong, bu...

After

We showed it a few hundred more example answers written for middle schoolers. Now it answers in...

What it knows stayed the same, but its whole tone changed. That's what extra tr...

Sticking to a format

Before

Ask for a math solution and sometimes it starts with the equation, sometimes with an explanatio...

After

We gathered only solutions written in the order 'equation → working → answer' and trained it on...

Just like tone, fitting a fixed template is something it picks up pretty well f...

Adding new facts

Before

Ask 'What's on my school's lunch menu this week?' and it makes up something that sounds right....

After

Even after teaching it the menu a few more times, it doesn't really stick. And when next week's...

Extra training is weak at making it memorize new facts. For that, it's better t...

Pick each of the three jobs and tap 'Before' and 'After' back and forth. You'll see what changes and what doesn...

Try it yourself in the web lesson · next-token.net/en/learn/day-51/

Quick check

Just three questions

Q1

What does extra training (fine-tuning) do?

① Rebuilds the model from scratch
② Searches the web and adds it to the answer
③ Teaches an already-trained model a bit more
④ A person fixes the answers by hand

Q2

Why is a method like LoRA such a big deal?

① It makes every answer correct
② It cuts the model's size in half
③ It runs without the internet
④ It needs 10,000 times fewer numbers changed

Q3

How do you get it to tell you a lunch menu that changes every week?

① Find it and attach it
② Teach it more
③ Make the model bigger
④ Write a longer question

Know the why, too

Q1

③

It keeps the big, already-trained base model as it is and just shows it more examples of the job you need. It's not raising it again, it's polishing it.

Q2

④

The 2021 LoRA paper froze almost all the numbers inside and trained only a tiny part. That meant 10,000 times fewer numbers to train.

Q3

①

You saw it in the third case of the experiment. For new facts, especially ones that keep changing, showing it the open book like yesterday works better than making it memorize.

Today in one line

You don't raise a grown-up model all over again, you just polish it a little, as much as you need.

Tomorrow: Humans scoring answers — given two answers, which one would a person pick, and how do those picks shape AI?