DAY 21–35 · Part 3 · What It Can Do

Part 3 Review — What It Can Do

From shrinking text to handing over a job · with the real terms · DAY 21-35

No test — a review to read through

Summarizing is choosing what to throw out

1. Summarizing isn't cutting a text off from the front.

2. AI skims the whole thing first, then asks of every sentence, **how important is this?**

3. Then it keeps only the high scorers and rewrites them in its own words.

Pulling out vs. writing anew

1. A summary that lifts sentences straight out of the original is called extractive summarizing.

2. What AI does now is **abstractive summarizing** — it writes new sentences.

3. That reads more smoothly, and that is exactly where words that weren't there get room to slip in.

Summaries taught by human ratings

1. A way to pre-train a model for summarizing alone came out in 2019.

2. In 2020 came a way to learn from summaries people had picked as good, and **quality jumped**.

3. People set the standard there — which is why you have to say your standard too.

Translation carries meaning, not words

1. Translation isn't swapping one word for another.

2. You catch the meaning of the whole sentence first, then **build it again in the other language**.

3. That's why a word-by-word 'You're taking my hair' says nothing to an English ear.

From swapping a dictionary to catching meaning

1. Older translators swapped words and short chunks using statistics.

2. Now, as you saw on DAY 14, the words in a sentence **look at each other** so the meaning gets caught first.

3. A word like bank has several meanings, and you can only pick one by reading what comes before and after.

2016, the year the service changed

1. The design that holds meaning in a middle slot and builds a new sentence from it was proposed in 2014.

2. Google put it into its live translation service in 2016, **Chinese to English first**.

3. That's why feeding in a whole passage gives steadier results — the model picks words with the context in view.

Pieces to hand over, pieces to keep

1. AI isn't a machine that writes for you — it's a tool for turning your own writing over and looking at it.

2. Your ideas, your experience and checking the facts stay **yours**; tightening and shortening you can hand over.

3. Whatever AI writes still goes out with your name on it.

Why the choosing can't be handed over

1. AI is good at making plausible sentences, but what matters comes out of your purpose.

2. So hand over the boxes where judgment lives and **the writing stays while your writing disappears**.

3. UNESCO's 2023 guidance said the same thing: students should keep hold of the judgment.

The last gate is reading and nodding

1. The last gate is keeping only sentences you can read and nod at.
2. Numbers and quotes whose source you couldn't find are **better deleted**.
3. A report on how students actually use AI came out in 2025 too.

Getting the code is only the start

1. It's true that code comes out when you describe what you want.

2. But the real half of coding sits **after the code appears**.

3. Running it, reading the errors and fixing them — that's where it lives.

Code that looks right vs. code that runs

1. The judge isn't your eyes, it's the result of running it.

2. In a 2021 evaluation, a code model got a problem right on the first try **about three times in ten**.

3. So the habit of running a test alongside is the skill itself.

The year a helper became everyday

1. A code-completion tool opened to every developer in 2022.

2. More tools, but **the checking box** stayed with people, exactly where it was.

3. Being able to read an error turns out to be what lets you use the tool at all.

Drawing by clearing away the speckle

1. A picture AI doesn't fetch a picture from somewhere.

2. It starts from a screen full of dots and draws by **clearing the speckle away bit by bit**.

3. Your description is the handle that sets which way the clearing goes.

Diffusion — from noise to a picture

1. Training practices covering a picture in noise; drawing practices undoing it the other way.

2. The name for this approach is a **diffusion model**.

3. In 2021 came a way to do the clearing in a compressed space instead of on the full-size picture, and speed jumped.

The year it took to open up

1. In 2021 DALL·E showed this kind of picture AI to a wide audience for the first time.
2. In 2022 **a model was released openly** so anyone could download and use one.
3. Adding three or four elements at a time and fixing as you go beats cramming everything in at once.

Thirty pictures per second

1. Video isn't a moving picture — it's still pictures stitched together fast.
2. One second usually holds **twenty-four to thirty** of them.
3. So making video is tens of times heavier than yesterday's picture-making.

The weight in numbers

1. One picture is called a frame, and one dot inside it is a pixel.

2. A ten-minute video comes to **about 37.3 billion pixels**.

3. The volume is large, but the harder part is that you can't draw each frame on its own.

Toward imitating a world

1. Research stretching the diffusion approach along the time axis came out in 2022.

2. In 2024 came a piece describing video generation models as **simulators of a world**.

3. The longer it runs, the harder remembering what came before gets — that's the homework still left.

It listens by chopping sound up small

1. AI doesn't take in your speech whole.
2. It writes the sound down as numbers 16,000 times a second and sweeps it with **a 25-millisecond window**.
3. Each piece on its own is just a pattern of highs and lows.

Speaking is the same road backwards

1. From letters it builds a pattern of highs and lows, then stamps out vibration values and sends them to the speaker.

2. That's **more than 16,000 stamps** a second.

3. In 2016 came a way to generate those vibration values directly, and machine voices moved much closer to human ones.

Noise and dictation

1. The less noise mixed into each piece, the clearer the pattern.

2. In 2022 a dictation model trained at scale on many languages and noisy audio was **released**.

3. It's the same story behind why a dictation app works better in a quiet room.

A list of places, or one solid block

1. Search doesn't give you an answer, it gives you a list of places to go.

2. AI takes what it has already read and builds **one solid block of sentences**.

3. There's only one thing to read, so it's fast — but where the words came from isn't written down.

Answering from memory vs. answering after fetching

1. If a model answers only from what it memorized, no source is left behind.

2. A way to have it **fetch the material first** and then answer came out in 2020.

3. Even so, the fetched material and the answer can drift apart, so the checking stays exactly where it was.

If it's going into homework

1. The trick is splitting jobs by whether a source is needed or not.

2. If the sentence is going into homework, in the end **you click through to the original** once.

3. What that 2023 study measured was precisely that degree of support.

What speeds up is the front end

1. Researching with AI makes starting faster.

2. It throws out search terms and outline options quickly, so **the finding-and-sifting front end** is what speeds up.

3. Numbers, years, people's names and quotes you have to click through and check yourself.

Breaking the checking into steps

1. Instead of believing the answer as given, a method of splitting it into questions to check and asking again came out in 2023.

2. Statistics from the body that published them, quotes **straight from the book**.

3. A number copied by someone else sometimes changes in the copying.

Tools that go and find things themselves

1. In 2025 came tools that roam the web and bring research back on their own.
2. The faster it gets, **the more there is to check**.
3. Decide up front which pieces need your hands and you won't get knocked around.

Just the answer, or the steps?

1. Ask about the same problem and what comes back changes completely depending on whether you asked for the answer or the steps.

2. There's one test for choosing — **will I have to do this on my own later?**

3. If you'll be on your own, take the steps; if you only need to know whether it's right, take the answer.

Why steps help

1. Answer only, all at once, and the middle of the calculation gets squashed out of sight.

2. Write the steps out and you can **check each box on its own**.

3. In 2023 came a method that teaches by grading whether each step is right.

The one step that makes it yours

1. Ask again about only the box you got wrong and the explanation comes back short and precise.

2. Taking the steps for something you only need graded is **a waste of time**.

3. Knowing where the hole is makes the question you ask shorter — that's the trick of this lesson.

Collect → table → fix → chart

1. You can't chart scattered material straight away.
2. A table has rules — **one row per case, one value per cell**.
3. Line things up like that and both you and AI can calculate right away.

Tidy rows as a standard

1. A piece laying out 'one row per case, one value per cell' as a standard came out in 2014.
2. Blanks have to be written in too, as **no answer**, or they can't be counted later.
3. Line the rows up in a table and the answers you get from AI are far more accurate too.

Leave the arithmetic to code

1. Instead of having it add numbers in its head, have it write calculation code and the mistakes drop.

2. This method was **written up in 2022**.

3. Which is why making the table well comes before the chart.

A talk breaks into pieces

1. Preparing a talk isn't one lump.
2. Narrowing the topic, gathering material, building slides, rehearsing — it **breaks into pieces**.
3. The pieces AI is good for are the rough draft and the labor-heavy work.

Where you spend the time the draft saved

1. Listing outline options and shortening long material are pieces you can fix yourself if they come out wrong.

2. Spend the time you save there **on rehearsing**.

3. UNESCO's 2023 guidance says the same: students keep hold of the judgment.

The questions come to you

1. In the room, you are standing there alone.

2. So check numbers where they were first published and quotes **in the original**, and write the sources down.

3. In 2025 AI tools made for schools arrived as well, but the checking box is unchanged.

The knob between boring and off-the-wall

1. AI picks the next word by probability.

2. Grab only the high-probability ones and it turns boring; **reach down to the low ones** and it turns odd.

3. That handle for widening and narrowing really exists, and it's called temperature.

How a word gets drawn from probability

1. Grabbing only the single highest option makes text repeat itself.

2. So a way of drawing from only part of the high-probability side was **written up in 2019**.

3. Real temperature values usually run between 0 and 2.

Draw many times and see what overlaps

1. A method of drawing the same question many times and taking the most-overlapping answer came out in 2022.

2. Pulling wide and **picking well** are a pair.

3. Making a pile of candidates is AI's strength; picking is yours.

Goal → break up → order → check

1. When you're stuck, don't ask AI for the answer — ask it to break the job up.

2. Planning isn't knowing the right answer, it's deciding **what to do first**.

3. Goal, break up, order, check — fill those four boxes and you have a plan.

Break it small and it solves better

1. Split a hard problem into small ones and have it solve them in turn, and it does better than being asked whole.

2. This method was **written up in 2022**.

3. A style of planning first and solving afterwards came out in 2023 as well.

Mistakes show up box by box

1. Ask for the four boxes one at a time and where it went wrong shows up box by box.

2. Find out on the last day and **there's no time left to fix it**.

3. Writing the thinking out in steps spread widely from 2022 for the same reason.

If you can undo it, hand it off

1. An AI that takes only a goal and then searches, opens files and retries on failure by itself is called an agent.

2. There's really one test for what you can hand over — **can you undo it when it's wrong?**

3. A file name you can undo; money sent and a report handed in you cannot.

Why many steps are hard

1. One wrong step in the middle throws the whole thing off.
2. So a style that makes it **alternate** thinking and acting came out in 2022.
3. In 2023 evaluations built to measure this kind of work arrived too.

Checks are what let you hand over more

1. Slot in a box where a person looks and you can safely hand over more.
2. Scores have risen a lot since, but it still doesn't mean **handing it over is free**.
3. Anything with someone's feelings in it, don't hand over — write it in your own words.

Pick by probability. read the neighbors. split into steps.
The last box always stayed with a person.



Part 4 looks at what to watch out for.