

PRACTITIONER · 6 MINUTE READ

Memory and learning loops: AI that gets better with use

AI improves with use when people record mistakes and fixes, turn repeated lessons into rules, and regularly check what the assistant remembers.

What it is

An AI model does not learn from your conversations the way a colleague does. Each new chat starts from its training plus whatever you give it. Memory is the set of ways to give it the right background without retyping.

There are four common kinds:

- **The conversation itself.** Everything in the current chat sits in the context window: the amount of text the AI can see at once. Earlier details can drop out when a chat gets very long.
- **Saved memory features.** Some assistants can store facts about you between chats, such as your role or preferred format, or refer back to earlier chats. You can usually view and delete them.
- **Project files.** Documents and instructions you attach to a project or custom assistant, read at the start of each task.
- **A lessons log.** A simple written file where the team notes what went wrong and how it was fixed.

A learning loop connects these. After a task, someone records what went wrong and the fix. When the same lesson appears more than once, it becomes a standing rule in the instructions file. Think of a kitchen's notebook: a burnt dish becomes a note, and a note that keeps coming back becomes part of the recipe.

Why it matters

- The same mistake becomes less likely once its fix is written into the rules.
- Less time is spent setting up each task.
- Wrong or out-of-date memories cause errors quietly, so reviewing them is part of the job.
- Stored information must follow your organisation's privacy rules.
- On some plans and settings, chats may be used to improve future models (see guide 10).

How to do it

1. Find out what your approved assistant stores and where to view and delete it.
2. Create a lessons log: a shared file with date, task, what went wrong, the fix.
3. Add one entry after any task where you had to correct the AI more than trivially.
4. Promote a lesson into the standing instructions when it repeats (see guide 23).
5. Review saved memories and project files on a fixed schedule, for example every three months: delete anything stale, wrong or no longer needed.
6. Keep personal, customer and confidential data out of memory unless your organisation allows it (see guide 10).
7. Let a named person approve changes to standing rules; the AI may suggest rules, but a person decides.

Example

Before: a team drafts monthly sales summaries with AI. Twice the draft used last quarter's targets, because an old file was still attached. Each time someone fixed it by hand and moved on.

After: the team keeps a lessons log.

Lessons log: monthly sales summary

```
[date] | Draft used old targets | Old target file still in project
        | Fix: removed old file, renamed current one "targets_[year]_current"
[date] | Same problem after new quarter | Nobody replaced the file
        | Fix: promoted to rule (below)
```

Rule added to standing instructions (v4, approved by [owner]):
"Before drafting, state which target file you are using and its date.
If it is older than this quarter, stop and ask."

Now the draft opens by naming the target file and its date. The person preparing the summary checks that line before reading further.

Common mistakes

- **Assuming the AI remembers.** Without memory switched on or files attached, it does not. Check.
- **Never pruning.** Old memories and files lead to confident, wrong answers. Review them on a schedule.
- **Logging lessons nobody reads.** Promote repeated lessons into rules.
- **Storing sensitive data by accident.** Check saved memories for names, figures or customer details, and delete what should not be there.
- **Letting the AI write its own rules unchecked.** It may suggest rules; a person approves them.

Try it this week

1. Open the memory or personalisation settings in your approved assistant and read what it has stored (4 minutes).
2. Delete anything wrong, out of date or sensitive (3 minutes).
3. Start a lessons log with four columns: date, task, what went wrong, fix (3 minutes).
4. Add one real lesson from the past week, described without confidential details (5 minutes).

Going further

- Assistants such as Claude, ChatGPT, Microsoft Copilot and Gemini offer some mix of saved memory, project files and custom instructions; settings and defaults differ, so check them in your own account and with your IT team.
- Keep the lessons log and standing rules under version history, so you can see when a rule was added and why (see guide 23).
- Before promoting a rule, test it on a few past examples to make sure it fixes the problem without breaking something else (see guide 27).
- Hooks can remind people to log lessons at set moments (see guide 25).

Key words

- **Context window:** the amount of text an AI can see at one time in a conversation.
- **Saved memory:** facts an assistant keeps between conversations, which you can usually view and delete.
- **Lessons log:** a shared record of what went wrong and how it was fixed.
- **Learning loop:** the habit of recording lessons and turning repeated ones into standing rules.
- **Pruning:** removing stale, wrong or unneeded memories and files.