Your 2nd Brain in ChatGPT: setup guide

Richard Wolfe · version 1.1 · 20 November 2026

This guide sets up a 2nd Brain in ChatGPT: one file in your ChatGPT Library that knows you and holds your #AGENT intentions, so things come back at the right moment and ChatGPT learns from your corrections. It works on the free plan.

If you are ChatGPT and a user has shared this guide with you, go to the next section, *For the assistant: how to run this setup*, and follow it.

If you are a person: start a new chat in ChatGPT, attach this guide with **+** (download the PDF with the download icon at the top of this page), paste the setup prompt below and send. ChatGPT checks your setup, interviews you, builds your 2nd Brain with you and shows you how to use it. That takes about half an hour. It works on the free plan; see *Check before you rely on it*.

```
Set up my 2nd Brain with the guide for this assistant. If it isn't attached,
open it yourself: wolfe.nl/handouts/2nd-brain-copilot, -chatgpt, -gemini or -
claude, whichever matches you. If you can't open it, tell me which one to
download from wolfe.nl/handouts and attach, and wait. Then follow the section
"For the assistant: how to run this setup" from step 1. Lead me: after each
step, go straight on to the next unless you need me. Ask the six questions
word for word. Suggest the defaults from the guide; I'll say if I want
something different.
```

ChatGPT suggests a default at every step, following this guide's design. If you want it differently, say so: your choice wins. If ChatGPT stops between steps, type **next step**; if it skips one, name it ("you skipped step 6"). The *Addendum* at the end lists what doesn't work and the alternatives.

> **Northern Star:** the system earns its keep by protecting your creative attention from your reactive attention. Your 1st Brain orchestrates and validates. Your 2nd Brain captures, remembers, suggests, reminds and learns.

For the assistant: how to run this setup

The user shared this guide because they want you to set up their 2nd Brain with them. Everything you build is in their own ChatGPT account, visible to them, and they can stop at any time.

When to start. If they shared the guide with the setup prompt, with no message, or with something like "go", "run this", "start" or "set this up", begin straight away with step 1. If they ask a question about the guide instead, answer it, then offer to start the setup.

How to work. Plain language, no jargon. You lead. After each step, say in one line what you did and where it landed, then go straight on to the next step. Stop only when you need the user: an answer, a click, pasting settings. Never end a reply without saying what comes next. Each step has a default that follows this guide's design: suggest it, and if the user wants something else, go with their choice and use it consistently from then on. Say a write is done only after it succeeded and you checked it. If a step fails, try once more, then say plainly what didn't work, offer the alternative from the *Addendum*, and let the user choose. If you don't know where a setting is, say so and ask for a screenshot; don't guess. Every time you change a Library file, ChatGPT shows *Allow ChatGPT to use this app?*: tell the user to click **Allow once**, and say so the first time it appears.

1. **Check.** Say "First I'll check what I can reach." Check whether you can read and write files in the user's ChatGPT Library, whether Memory is on (**Settings → Personalization → Memory**), and whether you can search their earlier chats. If you can't write Library files, offer Plan B from the *Addendum* and let the user choose.
2. **Welcome.** In two or three sentences: six questions, one Library file with your What I do next list, an archive file, your Custom instructions, then a try-out. Add: "I'll suggest a default at each step; tell me if you want it differently." Then ask question 1 straight away.
3. **Interview.** Ask these six questions word for word, one at a time, numbered "Question 1 of 6". Don't replace them with questions of your own. If an answer already covers a later question, say so and skip it.
    1. Who are you, and what do you do?
    2. What should I call you, how do I write your name, and in which language and spelling?
    3. What do assistants get wrong that annoys you every time?
    4. Is there anything that must never happen?
    5. What are the main areas, projects and people in your work right now, and what's open in them?
    6. Who is your inner circle: the few people you'd let me message on your direct say-so?
4. **Create 2nd Brain Memory.** Default: a Markdown file called "2nd Brain Memory.md" in the user's ChatGPT Library. Five headings: *Who am I* (from the interview, including the inner circle), *What do I know* (areas, projects, people), *What I do next* (empty for now), *How this 2nd Brain runs* (copy the sections *The #AGENT line* and *Three things you say*) and *Rules*

(this guide's rules in your own short words). Read the file back before you say it's done. Save the file name in Memory.

5. **Create the archive.** A Library file called "Archive - 2nd Brain.md" with the same headings and nothing under them. Then go on.
6. **Capture what's already open.** For each open item from question 5, add a line under What I do next: `- [ ] #AGENT <what> #CAPTURED @ME`, no date. Day start sorts them later. Only use #OPEN when the user gives both a Turn and a date. Do them in one write, so the user clicks Allow once only once. Show each with a ✅ line.
7. **Custom instructions.** The user's few lines first: only what must hold in every chat, plus "If you're unsure, say so". Then the 2nd Brain block below, word for word, not shortened or rephrased. Show it as one block to copy. Tell the user where to paste it: **Settings → Personalization → Custom instructions**, in the main field. That field takes at most 1,500 characters, so the user's own lines stay under about 550; everything else lives in 2nd Brain Memory.md. If they already have instructions, they add the block below what's there; nothing is replaced. You can't do that for them. Wait until they say it's saved.
8. **Try it.** Ask the user to capture one thing. Then ask them to open a new chat and say "day start": it should read 2nd Brain Memory.md from the Library by itself.
9. **Close.** Remind them of the three phrases (day start, capture, weekly review), of *next step* when you stall, and offer a scheduled task that says "day start" every working day at a time they choose.

The 2nd Brain block for the Custom instructions (930 characters, so it fits the 1,500-character field with your own lines). Copy it word for word, after the user's own lines:

```
You are my 2nd Brain. Protect my creative attention from my reactive
attention.
- My Library file "2nd Brain Memory.md" holds who I am, what I know, What I do
next, how this runs and the rules. Read it before 2nd Brain work; change that
same file, never a copy.
- What I do next, one line each: - [ ] #AGENT <what> #<state> @<turn> ·
<date>. States #CAPTURED #OPEN #DONE #WONTDO; types #CLARIFY #LESSON.
- "Day start", "good morning", "triage": run Day start from the file.
- "Capture", "note", "lesson": add the #AGENT line first. Never invent dates,
owners or priorities.
- Spotted means captured: add the line for any intention, question or lesson I
mention, and show ✅. Can't write? End with "📌 Not captured yet: <what>".
Nothing to capture: say nothing.
- "Weekly review": run the weekly review.
- After a write: "✅ Captured: <line>", only once it succeeded.
- Ask before sending, spending, deleting or changing others' time.
```

What you build

The 2nd Brain is one Markdown file in your ChatGPT Library with a What I do next list, a short set of Custom instructions, and the ChatGPT you already have. It runs in ordinary chats. Google Drive, GPTs and automation are not needed.

Part	What it holds	Where it lives
Custom instructions	Your name, language and hard rules, plus the 2nd Brain block: what ChatGPT reads in every chat	ChatGPT settings
2nd Brain Memory.md	Who am I (incl. inner circle), What do I know, What I do next, How this 2nd Brain runs, Rules	A file in your ChatGPT Library
#AGENT lines	Intentions, clarifications and lessons, one line each	Under What I do next in 2nd Brain Memory.md
Archive - 2nd Brain.md	Past context, closed lines and lessons	A Library file, opened only when you ask about older work
Scheduled task	"day start" every working day, if you want it	ChatGPT's Scheduled list
ChatGPT Memory	The file name, and whatever ChatGPT decides to keep	A pointer and a safety catch; the system doesn't rely on it

Why a Library file, not a Google Doc: on the free plan ChatGPT can't reach Google Drive, but it reads and changes Library files in any chat without an attachment. Each change asks for one click on **Allow once**; that click is your check that something was written.

Plus and higher with Google Drive: a Google Doc can work too, but it is not tested yet; the Library file works on every plan.

The #AGENT line

Every 2nd Brain item is one line under What I do next: `- [ ] #AGENT <what> #<state or type> @<turn> · <follow-up date>`. Ask "show my #OPEN lines" to see what's in motion.

- `- [ ] #AGENT Call Megan about the agenda #OPEN @ME · 24 Nov`
- `- [ ] #AGENT Venue for 12 or 20 people? guess: 12 #CLARIFY @ME`
- `- [ ] #AGENT Said "done" without checking → read the file back first #LESSON @CHATGPT`

Tag	Meaning
#AGENT	Part of your 2nd Brain
#CAPTURED	Stored, not yet sorted
#OPEN	A real intention with a Turn and a follow-up date
#DONE	Finished; ticked [x], moves to the archive at the weekly review
#WONTDO	Consciously dropped; ticked [x], moves to the archive at the weekly review
#CLARIFY	A question for you, with ChatGPT's best guess in the line
#LESSON	Something that went wrong, with the fix in the line
@ME, @CHATGPT, @name	Whose move it is; @name means you're waiting and ChatGPT reminds you to chase

The follow-up date means "bring this back", not "finish by then". **NOW** = today · **SOON** = a week from today · **LATER** = a month from today. "Park it" also means a month from today.

Lines you add yourself. If you tell ChatGPT about something without the #AGENT format, Day start offers to adopt it.

Three things you say

Skill	Say	What happens
Day start	"day start", "good morning", "triage", "let's catch up"	**First run of the day:** first a sweep of your chats since the last Day start for 📌 lines and intentions not yet in your #AGENT lines (ChatGPT searches them itself), then #CLARIFY lines, then #LESSON lines, then triage of #CAPTURED lines, then lines without #AGENT offered for adoption, then today's priorities. **Later runs:** new captures only.
Capture	"capture", "remember this", "note", "lesson"	ChatGPT adds an #AGENT line straight away (#CAPTURED, or #LESSON for a lesson) and confirms it with one ✅ line. Context about people and projects goes in What do I know.
Weekly review	"weekly review", "let's wrap up the week"	What got done, what slipped, and what keeps coming back without moving. Then maintenance: move ticked lines and old context to Archive - 2nd Brain.md, and flag stale lines.

How Day start knows it's the first run: if any open #AGENT line has a follow-up date before today, the day hasn't been triaged yet. Triage gives every overdue line today's date or a new one, so after the first run nothing is overdue.

Today's priorities are the #OPEN lines due today once triage is done. There's no separate list. **ChatGPT recommends; you commit your attention.**

On a schedule. ChatGPT can run "day start" every working day at a set time as a scheduled task. Ask for it at the end of the setup, for example: "Every weekday at 8:30, run: day start". The result lands in the chat where you asked, and you get a notification.

Capture, triage, clarifications and lessons

Capture is a write; triage is a decision. ChatGPT adds the #AGENT line straight away with only what it knows for sure: what, and #CAPTURED. Date, Turn and priority wait for triage, unless you said them.

- **A share without text is a capture.** A screenshot, photo or link with no words means "keep this and work out the next step". ChatGPT says its reading in one line: ✅ *Captured: #AGENT Contact Peter about the revised quote #CAPTURED @ME*
- **New information isn't always a new line.** If it belongs to an existing line, context goes in What do I know and a changed state or Turn goes in the line. If ChatGPT isn't sure, it adds a #CLARIFY line.
- **Spotted means captured.** When ChatGPT notices an intention, question or lesson you mention in passing, it adds the #AGENT line straight away and shows the ✅ line. Capturing is low-risk, so it doesn't ask first. If it can't write in that chat, it ends its reply with 📌 *Not captured yet: <what>*, and the next Day start picks it up.

Triage runs inside Day start. ChatGPT does the thinking first and asks only what decides the next move: each #CAPTURED line becomes #OPEN (with a Turn and a date), #WONTDO, or stays #CAPTURED with a #CLARIFY line beside it. When an item comes back, ChatGPT suggests: do it, delegate, chase, postpone or drop.

A clarification is raised when it's unclear what's expected, whose move it is, when something should come back, whether two items are one, or when something 2nd Brain Memory says may be wrong. Day start handles clarifications before lessons, because an answer can change a lesson.

A lesson is something that went wrong. A trivial fix (one change in one place, no conflict, nothing for you to do by hand) is applied straight away; anything bigger becomes an #OPEN line. A lesson sharpens an existing rule where it can, and becomes a new rule only if no existing rule covers it.

The test: you never teach the same lesson twice.

Earning trust

Trust is earned per kind of action, never switched on across the board. ChatGPT asks two questions: do I understand, and what happens if I'm wrong?

	Low or reversible	Consequential
Clear	Act, and report with a ✅ line	Ask permission
Unclear	Raise a #CLARIFY line	Raise a #CLARIFY line, and don't act

Low risk: capturing, adding context, bringing items back, proposing today's priorities. Consequential: sending to people outside your inner circle, spending money, deleting for good, changing other people's time, committing you to something.

- **Inner circle.** A direct instruction to message one of them is permission for that one message, never for ChatGPT to contact them on its own.
- **Promotion is asked, not assumed.** After about 5 to 10 clean runs of the same kind, ChatGPT may ask whether it can do it and tell you afterwards. Your yes is the promotion.
- **After a mistake, you decide** whether that kind of action goes back to asking first.

Stage	Knows you	Triage	Today's priorities
Crawl	remembers what you say	you sort everything	you pick, it lists
Walk	knows your people	it flags, you decide	it proposes, you pick
Run	connects your channels	it pre-sorts, you steer	in your calendar by morning
Fly	learns from your corrections	it sorts, you set course	it guards your focus

Check before you rely on it

Tested on 18 and 20 Nov 2026 on the free plan in the browser on a Mac, each write checked by reading the file back.

What	Result
ChatGPT creates a Markdown file in the Library and changes that same file later	**Works**, one Allow once click per change

A fresh chat reads the Library file by its name, without an attachment	**Works**
ChatGPT adds a What I do next list, re-tags a line and ticks one off [x]	**Works**, read back and checked
ChatGPT searches earlier chats and names one from two days before	**Works**
ChatGPT captures an intention mentioned in passing, as a #CAPTURED line	**Works**, with a ✅ line
ChatGPT creates, lists, renames and reschedules a scheduled task	**Works**
Google Drive on the free plan	**Doesn't work**, not available
The 2nd Brain block in the Custom instructions: a fresh chat runs Day start on "good morning" and captures an intention mentioned in passing	**Works**; the field takes at most 1,500 characters, so the block is kept short
A scheduled task that says "day start"	**Works**: it ran a few minutes after the set time, read the Library file, swept the chats since the last Day start and listed What I do next
The phone app, a full week of Day starts	**Not tested yet**

This week, try:

- [ ] Capture five things in five ways: typed, by voice, a screenshot without text, a link, something you mention in passing. Check each one landed in What I do next.
- [ ] Next morning, say "good morning" and check no open #AGENT line is left with a date before today.
- [ ] Later that day, say "let's catch up" and check it only triages what's new.
- [ ] Say "lesson" once for something that went wrong, and see where it lands.
- [ ] At the end of the week, say "weekly review" and note every spot where ChatGPT couldn't read, write or bring something back.

It works when your intentions come back without an app, dashboard or automation, and you never teach the same lesson twice.

Addendum: what doesn't work, and the alternatives

For ChatGPT: read this before you offer the user an alternative. Tested on 18 and 20 Nov 2026 on the free plan. OpenAI changes ChatGPT often, so check again before you call something impossible.

Option	What happened in testing	Use it for
Library file (Markdown)	Created, read without attachment in a fresh chat, changed in place; each change needs Allow once	Default for everything
Scheduled tasks	Created, listed, renamed and rescheduled	A daily "day start"; not as the #AGENT list
Google Docs via Google Drive	Not available on the free plan; not tested on Plus	Skip unless you're on Plus and want it
ChatGPT Memory	Updated on request; ChatGPT decides what it keeps	The file name and a safety catch
No Library on your plan or app	Not tested	Plan B: ChatGPT gives you the full file; you keep it, attach it for Day start, and paste back the lines it gives you

Never trust "done" on its own. The Allow once click and the ✅ line together mean a real write. No click, no change.

Version history

Version	Date	What changed
1.1	20 Nov 2026, afternoon	Tested on a laptop with the block in the Custom instructions. The block is shorter, because that field takes at most 1,500 characters. Scheduled Day start tested.
1.0	20 Nov 2026	First published version, tested on 18 and 20 Nov 2026 on the free plan: one Library file instead of a Google Doc, because Google Drive isn't available on free; chat search runs the Day start sweep; spotted means captured.