

FOR ECOMMERCE FOUNDERS DOING \$500K-\$10M

The Operating Brain

The seven-file system that lets AI draft your CS replies, ad briefs, and emails like a three-year employee — not an intern. Portable to any AI you'll ever use. The full method is inside; nothing held back.


Jay Oswal · Helix AI · helixai.agency

Why every AI tool you've tried writes like an intern

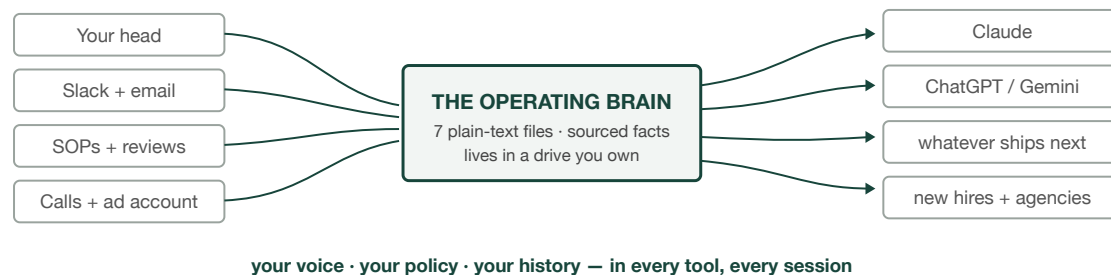
You've pasted your product page into ChatGPT and asked for ad copy. You got something grammatical, confident, and completely generic — copy that could sell any moisturizer, any supplement, any bag. You fixed it by hand, and quietly concluded AI isn't there yet.

The model isn't the problem. The context is. What makes your best freelancer good after three months isn't talent — it's that they've absorbed how you talk, what you've already tried, which claims you can't make, and which supplier cuts corners when nobody checks. None of that lives anywhere an AI can read it. Most of it lives in exactly one place: your head.

TODAY — KNOWLEDGE SCATTERED, AI BLIND



WITH AN OPERATING BRAIN — ONE MEMORY, EVERY AI



The brain is model-agnostic: plain text, no vendor, no lock-in. Models will keep changing; your memory layer shouldn't.

Score your founder-dependency

Before the method, the diagnosis. One point for each statement that's true:

1. A customer complaint escalated to you personally in the last seven days.
2. You explained the same process to more than one person this quarter.
3. Your best ad angles — and why the dead ones died — exist only in your head and your ad account.
4. A new hire or agency took more than two weeks to produce on-brand work.
5. You've caught AI or freelancer output making a claim your brand can't legally make.
6. Nothing significant ships while you're on a flight.
7. If your best CS or ops person quit tomorrow, processes would leave with them.
8. You've re-made a decision you'd already made, because nobody wrote down why.
9. Your last agency onboarding took a month of calls.
10. You've abandoned ChatGPT or Claude for real work because the output was generic.

0–2 Rare. You've built real systems; a brain makes them AI-readable.

3–6 The classic founder bottleneck. This is where a brain pays back fastest.

7–10 You aren't running the company; you *are* the company.

What a typical 5-scorer pays for this: six hours a week of re-explaining and rewriting, at \$150–250 an hour of founder time, is **\$45,000–75,000 a year** — before counting the launches that wait for you. Run your own numbers; they won't be smaller. Keep your score — you'll want it at the end.

The seven files

An operating brain is your brand's working memory: seven plain-text files, small enough to maintain, structured enough to trust. Not a wiki, not a 200-page brand book, not software you rent. Here's the full structure — steal it.

1. brand-snapshot.md — One page. What you sell, to whom, against whom, and this year's targets. The spine: every other file must make sense in light of it.

2. voice-and-customers.md — Voice rules with examples ("we say X, never Y"). Banned words and claims — including the legal ones. Real customer language from reviews and tickets, because your customers describe your product better than you do.

3. products-and-suppliers.md — The SKU map with the story behind it: hero products, margin logic, bundle rules. Supplier terms, real lead times, quality-issue history. The file that stops an expensive reorder mistake.

4. playbooks/ — One file per repeatable process: launch checklist, CS macros and escalation rules, how creative gets tested, how flows are structured. Written as instructions to a competent stranger.

5. **decisions.md** — The most valuable and least-kept file. Each entry: what was decided, when, why, and what would change your mind. "Killed the UGC angle in March — CPA fine but refund rate doubled." Why the same mistake never gets made twice.

6. **people-and-partners.md** — Team, agencies, freelancers: who owns what, what each has promised, how each is briefed.

7. **open-loops.md** — Live commitments: the wholesale deal in negotiation, the reformulation in progress. What's in motion right now.

Every fact carries its source, date, and a status — **verified**, **extracted**, or **inferred**. Sounds fussy; it's the difference between a brain you trust and a wiki you don't.

The six questions that fill it

You don't "document everything." You answer six questions — out loud, recorded, about ninety minutes:

1. **What would someone need to know to not embarrass this brand?** → voice, banned claims, red lines
2. **What do you find yourself re-explaining most often?** → playbooks
3. **What decisions would you hate to see accidentally reversed?** → decisions
4. **What has clearly worked before — and what clearly hasn't?** → decisions, voice
5. **What must never happen?** → rules
6. **Where does the truth actually live today?** → sources to ingest

Everything else — reviews, past ads, SOP fragments, supplier emails — gets pulled from question six and filed under the same structure.

The difference, on a real task

The task: a customer writes — *"My order arrived with a cracked lid. Third time ordering from you. Fix this or I'm done."* Same AI, same prompt, run twice.

WITHOUT THE BRAIN

Dear Valued Customer,

We sincerely apologize for the inconvenience you have experienced. Your satisfaction is our top priority. Please contact our support team with your order number so we can look into this matter, and consider visiting our FAQ page for our return policy.

We truly value your feedback.

WITH THE BRAIN

Oh no — a cracked lid on your third order with us is not the experience we want you having. A replacement ships today; no need to return the damaged one (a photo helps us chase it with the factory).

I've also added a jar of the cedar blend to the box — small thank-you for sticking with us three orders in. You'll have tracking within the hour.

- "Dear Valued Customer" is on the banned list in `voice-and-customers.md`; the tone rules come from three annotated examples in the same file.
- Damaged item → replace immediately, no return required — `playbooks/cs-escalations.md`.
- Repeat buyers get a surprise add-on, never a discount code — a 2025 entry in `decisions.md`.
- The photo request exists because lid cracks are a known issue with one supplier — `products-and-suppliers.md`.

This example is from our demo vault for a fictional home-goods brand — on a call, we run the same test live on *your* brand's real tasks. That's the standard everywhere:

The twenty-minute test. Pick one real task from this week. Run it twice through the same AI: once cold, once with the brain loaded. Judge three things: your voice or *a* voice? Your actual policy? Your history? If the second run isn't obviously better, the brain isn't done.

Why this won't become another dead wiki

Wikis die because maintaining them is a job nobody owns. A brain stays alive because AI does the maintenance and a human does the approving: **ingest** turns new material (a supplier email, a call transcript) into proposed updates; **curate** checks weekly for stale, contradictory, missing, and unsourced facts; and **propose-then-approve** means nothing ever writes itself into your source of truth — you approve changes in a five-minute weekly review.

Four ways founders kill it anyway: dumping everything in (400 files is a wiki with extra steps) · no sources (facts rot into rumors) · letting AI write to it directly (trust dies) · never testing it (the test is the maintenance alarm).

Why we're giving the whole method away

If the files, the questions, and the test are all here — why pay us? Three honest reasons.

Some of you will build it yourselves this weekend. Good. You'll prove the method on your own brand, and you'll know what the maintenance feels like. Some will come back for that part; the rest will tell another founder where the free PDF came from. Either way this document did its job.

The method was never the hard part — judgment is. The follow-up question that surfaces the real policy behind the stated one; the discipline to quarantine a guessed fact; the taste to catch generic slop wearing your logo. None of that transfers by reading. Same recipe, different restaurants.

And what you'd buy isn't information — it's speed and certainty. Fourteen fixed days instead of the founder-weekend that keeps not happening, and a brain tested on your real tasks before you sign off.

	BUILD IT YOURSELF	BUILT BY HELIX
Time to live	the weekend you keep not having	14 days, fixed, while you run the business
The interview	you, interviewing yourself — you can't see your own defaults	an outside interviewer whose whole job is the follow-up question
Quality gate	"looks right"	three real tasks, cold run vs. brain, demoed live at handoff
Maintenance	one more system you own	ingest + curate routines installed and running; retainer optional
Cost	two founder-days, plus ongoing attention	\$4,250

And it isn't theory: Helix runs on the same seven-file brain — our outreach, proposals, and delivery decisions draw from it, and we'll screen-share it live. The same structure sits at the core of every AI system we've built for clients, from ecom brands to a music school to an internet provider.

The offer

Knowledge Audit — \$750, 3–5 days

We map where your brand's knowledge actually lives, interview you against the six questions, and hand you the schema for your brain. Fully credited toward the build if you continue within 14 days.

Operating Brain Build — \$4,250, 10–14 days

The full seven-file brain, built from your real materials — calls, reviews, ads, SOPs, inbox. Installed in tools you already own, ingest and curate running. You own every file; portable to any AI you use next.

The handoff standard is the twenty-minute test, run live on your tasks, in front of you. If the brain doesn't obviously beat the cold run, we keep working until it does.

Book the fit call

Twenty minutes. Bring your scorecard number. We'll read it against the six questions and tell you the one file to build first — useful whether or not you ever hire us. If a brain isn't the right move at your stage, we'll say that on the call.

We take **three builds a month**. Solo-operated by design — that's real capacity, not a countdown timer.

Book the fit call — cal.com/jayoswal/20min-fitcall

Twenty minutes, free. Bring your scorecard number — or just reply to whoever sent you this. · helixai.agency/operating-brain

Common questions

Our data is sensitive. What touches it? Everything is classified before it's written — public, internal, confidential, restricted. Restricted material (credentials, legal, personal data, customer PII) never enters the brain. It lives in your accounts, not ours.

We already have SOPs in Notion. Good — that's source material, not a brain. SOPs say how; the brain also holds who, why, and never-again. The build ingests what you have.

What happens after handoff? It's yours. Run it yourself, or have us keep it alive on a light monthly retainer — weekly curate cycle, monthly re-test, quarterly refresh. Both work.