A HYPERAUTOMATION LABS WORKING COPY

The One-File Agent

A complete, working AI agent in a single file you can read. It plans, searches the live web, writes, attacks its own draft, and revises — in front of you, step by step. Free key, nothing to install, and the whole "agent" is a page of commented logic you can open in any text editor.

» Run it in one click — or save the file and it's yours forever

» The 5-step loop inside every agent, explained on one page

» The test log — including the draft our own critic caught fabricating news

» Missions worth giving it (and the kind it will fail)

tested live: 4 logged missions · August 27, 2026

hyperautomationlabs.co

GET IT RUNNING — 3 STEPS, ~10 MINUTES (MOST OF IT IS GOOGLE'S SIGN-UP PAGE)

You have watched people talk about AI agents for a year. Today you run one.

STEP 1 Open the agent

It lives here: hyperautomationlabs.co/one-file-agent.html — that page IS the agent, running in your browser. Want to own it? Save the page (Ctrl+S / Cmd+S, or right-click → Save As) and you get the single HTML file: double-click it any time, works from your desktop, edit it freely. There is no server behind it — everything happens between your browser and the model API.

STEP 2 Get your free key

The agent needs a brain. Google AI Studio hands out free API keys: aistudio.google.com/apikey — sign in with any Google account, click "Create API key", copy it. No card. Paste it into the agent's Step-1 box: it is saved only in your browser and sent only to the model API. Don't take our word for it — the file is right there, read it.

STEP 3 Give it a mission

Type what you want and press Run. Good missions have a deliverable in them: *"Compare X and Y for a beginner — table + recommendation"*, *"Research whether Z is worth it — verdict + evidence"*, *"What changed in [your field] this month — 5 bullets with sources"*. Then watch the five steps light up: it plans, runs real web searches, drafts, critiques its own draft, revises, and hands you the deliverable with the source list it actually read.

> **What this is, honestly:** a research agent. It reads the live web and writes evidence-backed deliverables — briefs, comparisons, verdicts. It does not book flights or send emails; giving a downloaded file those permissions is exactly what you should NOT want. It is the "read and reason" half of agents, done properly — sources shown, claims checked.

THE LOOP — WHAT "AN AI AGENT" ACTUALLY IS, ON ONE PAGE

Courses spend three hours diagramming this. It is five model calls in a loop. Here they are.

1 PLAN

The mission goes to the model with one job: turn it into a goal, 2–3 web-search queries, and a deliverable format — returned as strict JSON so the code (not the model) stays in charge of what happens next. This is the difference between an agent and a chatbot: the output of this call becomes the *input of a program*, not a message to a human.

2 RESEARCH — the tool call

One call per query with the web-search tool attached. The model decides what to search, the search actually executes, and the response comes back with the list of pages it read. The prompt demands concrete findings — names, numbers, dates — and permits "NOTHING SOLID" as an answer, because an agent that must always find something will invent something.

3 DRAFT

The writing call gets the mission plus the research notes and one hard rule: every specific claim must be traceable to the notes. If the notes are thin, deliver less and say so. (This rule was added after testing — page 4 shows exactly what happened without it.)

4 CRITIQUE — the step that makes it an agent

A second, adversarial call gets the mission, the notes, and the draft: attack it, claim by claim — is each fact actually in the evidence? It returns JSON: passes yes/no, problems, fixes. Instructed NOT to invent problems when the draft is clean — a critic that always complains is as useless as one that never does.

5 REVISE

Only runs if the critic failed the draft. Applies the named fixes under the same evidence rule — if a fix can't be satisfied from the notes, the claim gets cut, not re-invented. Then the deliverable ships, with every source the agent read linked underneath.

> **The homework that's actually fun:** open the file in any text editor. The whole loop is ~150 commented lines under the heading "THE AGENT". Change a prompt, add a step, make the critic meaner. You now know how agents work — not from a diagram, from the source.

THE TEST LOG — WE TRIED IT ON REAL MISSIONS BEFORE SENDING IT

4 logged missions, every call recorded. Including the one where it fabricated the news — and caught itself.

THE SETUP

Three missions a learner would actually give it: compare free AI coding assistants (table + recommendation), "what changed in AI this week" (5 bullets, sources), and "is prompt engineering still worth learning" (verdict + evidence). Full pipeline, live model, live web search. Every run finished in **34–44 seconds**, reading **12–28 web sources** per mission — 25 model calls total, zero rate-limit hits, all on the free tier.

FINDING 1 — THE DRAFT FABRICATED NEWS. THE CRITIC CAUGHT IT.

On the news mission, the first draft confidently included a multi-billion-dollar acquisition that never happened and a security advisory with a made-up ID. Then step 4 ran — and the agent's own critic rejected the draft, naming 4 fabrications. Its first finding, verbatim from the log:

```
"Fabricated headline: Nvidia has not acquired Hugging Face for $13B. This is a completely hallucinated event presented as factual news."
```

This is why the critique step exists, and why it gets the research notes: a model writing from memory sounds exactly as confident as a model writing from evidence.

FINDING 2 — CATCHING LIES ISN'T ENOUGH. THE FIX IS IN YOUR COPY.

The revision after that catch was honest but stale — it swapped the fabrications for real events from years past. Real, sourced… wrong week. The fix we baked into the file: the draft and revision steps now run under a hard evidence rule (every specific claim traceable to the research notes, "deliver less" allowed), and the critic cross-checks claim by claim against the notes. Re-ran the same mission: every bullet came back dated from that actual week, with sources — and the critic passed it first try.

FINDING 3 — WHERE IT'S GENUINELY GOOD

The comparison mission returned a table with real free-tier limits (completions per month, per-tool trade-offs) and a defensible recommendation on the first pass — the critic found nothing to fix. The verdict mission came back with market signals and a nuanced call ("yes, but not as phrasing hacks — as a baseline skill") backed by 28 sources. Decision-shaped research with a deliverable at the end is exactly its lane.

Honest limits: its sources are whatever web search surfaces — blogs included; the source list is shown so YOU can judge. A wrong-but-sourced claim can still get through: verify anything load-bearing before you act on it. Free-tier limits mean a handful of missions per day, not hundreds. And these are single logged runs from one week (receipts saved), not a benchmark.

READY TO GO PAST ONE FILE? — ALL 4 GUIDES UPDATED AUGUST 2026

Step-by-step beginner playbooks.

The One-File Agent shows you what an agent is. These paid beginner guides put full-strength agents to work on real projects — code, files, deliverables — click by click. All four freshly updated for August 2026.

$14.97 **The Complete Claude Code Guide**
Step-by-step beginner playbook for Claude Code — setup to real projects.
hyperautomationlabs.gumroad.com/l/claude-code-guide

$14.97 **The OpenAI Codex Guide**
The beginner path to Codex — delegate real work to an AI agent.
hyperautomationlabs.gumroad.com/l/codex-guide

$14.97 **Claude Cowork Sales Guide**
Turn Claude into a sales co-worker — pipelines, outreach, follow-ups.
hyperautomationlabs.gumroad.com/l/claude-cowork-sales

$29 **Claude Certified Architect Prep Kit**
Prep kit for the CCA-F certification — drills, mock questions, study plan.
hyperautomationlabs.gumroad.com/l/claude-certified-architect-prep

> This working copy shipped in **The Hyperautomation AI Report**. Daily AI moves on **Instagram** & **Facebook** · full breakdowns on **YouTube** — **Hyperautomation Labs.**