

M A V E R I C K A I

THE COMPLETE CHATGPT GUIDE

**EVERYTHING CHATGPT DOES IN 2026,
FROM FIRST LOGIN TO POWER USER**

*Fully rewritten after the July 9 update. Every model, every plan, every feature, plus the stuff
OpenAI quietly killed that most guides are still telling you to use.*

GPT-5.6 Sol, Sol Pro, Terra, Luna and GPT-5.5 Instant

ChatGPT Work: the agent that does the whole job

The new desktop app: Chat + Work + Codex in one

Built-in browser, Computer Use, Record and Replay

Plugins, Skills, Memory, Projects, Scheduled Tasks

Business, company knowledge and workspace agents

Real pricing, real limits, and 25 power tips

2026 Edition | Updated July 12, 2026

What's In Here

- 01 What ChatGPT Actually Is Now
- 02 What Got Killed (Read This First)
- 03 Getting Started on Every Platform
- 04 The Models, Explained Properly
- 05 Pricing: What You Really Get
- 06 Ads, and How to Kill Them for Free
- 07 Navigating the Interface
- 08 ChatGPT Work
- 09 The New Desktop App
- 10 The Built-In Browser and Computer Use
- 11 Deep Research
- 12 Voice: GPT-Live
- 13 Image Generation
- 14 Canvas and Sites
- 15 Codex
- 16 Plugins (Formerly Connectors and Apps)
- 17 Skills
- 18 Memory and Personalization
- 19 Projects
- 20 Scheduled Tasks
- 21 Data Analysis and the Office Extensions
- 22 ChatGPT for Business and Enterprise
- 23 Settings, Privacy and Safety
- 24 Prompting for the Reasoning Era
- 25 25 Power Tips
- 26 The 90-Day Reality Check

HOW THIS GUIDE WAS BUILT

Every fact here was checked against OpenAI's own pricing page, help centre articles and release notes as of July 12, 2026, then cross-checked against tutorials published in the last 90 days. ChatGPT is changing fast enough that anything written before July 9 is already wrong in at least three places.

01. What ChatGPT Actually Is Now

ChatGPT stopped being a chatbot a while ago. What you are logging into in 2026 is closer to an operating system with a text box on the front of it.

It writes, reasons, codes, researches the web, generates and edits images, holds a real-time voice conversation, reads your files, connects to over a thousand third-party tools, drives a browser, clicks around your actual desktop, and now takes a goal and works on it for hours until there is a finished spreadsheet or deck sitting in front of you.

The biggest shift landed on **July 9, 2026**. OpenAI shipped the GPT-5.6 model family, launched ChatGPT Work, merged the Codex app and the ChatGPT desktop app into a single application, and renamed the whole connectors/apps system to Plugins. If you learned ChatGPT before that date, a good chunk of what you know is now out of date.

The platform at a glance

Layer	What's in it
Where you use it	Web (chatgpt.com), desktop app (macOS + Windows), mobile (iOS + Android), API
Models	GPT-5.5 Instant (default), GPT-5.6 Sol, GPT-5.6 Sol Pro, GPT-5.6 Terra, GPT-5.6 Luna, GPT-5 Thinking Mini
Core features	Deep Research, ChatGPT Work, Voice (GPT-Live), Images, Canvas, Sites, Codex, Memory, Projects, Skills, Scheduled Tasks
Connections	Plugins (1,400+), interactive in-chat apps, built-in browser, Computer Use, Excel / PowerPoint / Sheets extensions
For teams	Company knowledge, workspace agents, admin console, SAML SSO, SOC 2

THE ONE-LINE VERSION

ChatGPT is now three products in one window: **Chat** for answers, **Work** for finished deliverables, and **Codex** for code. Most people only ever use the first one, which is why they think it peaked in 2024.

02. What Got Killed (Read This First)

Most ChatGPT guides floating around right now are teaching features that no longer exist. Start here so you do not waste a week learning a dead product.

Thing	Status	What to do instead
Sora (video)	Discontinued. App and web shut down April 26, 2026. The API shuts down September 24, 2026.	Export anything you care about. There is no in-ChatGPT video generation right now.
ChatGPT Atlas (the standalone browser)	Shutting down August 9, 2026.	Use the browser now built into the ChatGPT desktop app, or the ChatGPT Chrome extension.
GPT-4o, GPT-4.1, o4-mini, original GPT-5	Retired from ChatGPT on February 13, 2026.	GPT-5.5 Instant is the everyday replacement. Legacy models remain for Plus and Pro only.
"Connectors" and "Apps"	Renamed. The App Directory became the Plugin Directory on July 9, 2026.	Same idea, new name and a bigger library. See chapter 16.
Advanced Voice Mode	Replaced by GPT-Live (July 8, 2026).	It is better. See chapter 12.
The old desktop app	Renamed "ChatGPT Classic."	Download the new unified desktop app. See chapter 9.

HEADS UP

If a guide, video or course published before July 9, 2026 tells you to generate video with Sora, install Atlas, or pick GPT-4o from the model picker, it is out of date. That includes the previous version of this guide.

WHY THIS KEEPS HAPPENING

OpenAI ships and cuts aggressively. The lesson is not to memorise features. It is to build habits (memory, projects, skills, plugins) that survive model and product churn. More on that in chapter 25.

03. Getting Started on Every Platform

Setup takes about two minutes. Where you use ChatGPT now matters more than it used to, because the three platforms are no longer equivalent.

STEP 1: Web: chatgpt.com

Sign up with email, Google, Microsoft or Apple. Free gives you GPT-5.5 Instant with limited messages in a rolling 5-hour window. The web app is the only place company knowledge works for business accounts, and it is where most people live.

STEP 2: Desktop: the app that changed everything

Download from chatgpt.com/download. This is the important one now. The desktop app is the only place you get the built-in browser, Computer Use (ChatGPT clicking around your actual machine), Record mode, Appshots, and local file access. It is available on every plan, including Free.

STEP 3: Mobile: iOS and Android

Full voice, camera, and now Remote, which pairs your phone with a Mac or Windows machine running the desktop app so you can kick off and check on tasks from your pocket. ChatGPT Work on mobile requires Plus or above.

The first thing to do after signing up

Before you ask a single question, go to **Settings > Personalization > Custom Instructions** and tell it who you are. This context is applied to every conversation and it is the single highest-leverage five minutes you will spend.

CUSTOM INSTRUCTIONS TEMPLATE

WHAT I DO

I run [business/role]. My audience is [who]. My main goals right now are [goal 1] and [goal 2].

HOW TO ANSWER ME

Be direct. No preamble, no filler, no 'great question'. If I'm wrong, say so and tell me why. Give me the answer first, then the reasoning if I ask. Use plain language. Never use em dashes.

WHAT I ALREADY KNOW

Assume I understand [X]. Don't explain

the basics of [Y] unless I ask.

TIP

Then go to **Settings > Personalization > Advanced** and switch on **web search** and **connector search**. Connector search is off by default for most people, which is why they think ChatGPT cannot see their Drive or Gmail. It can. You just never turned it on.

04. The Models, Explained Properly

OpenAI changed how model naming works on July 9. The **number** is the generation (5.5, 5.6). The **name** is a capability tier that stays consistent across generations. Sol is the flagship, Terra is the balanced one, Luna is the fast cheap one.

But here is the part that trips everyone up: in normal chat, **you do not pick a model. You pick a reasoning level.** The model picker now shows effort settings, and the effort you choose determines which model runs.

The reasoning levels

You pick	What actually runs	Use it for	Who gets it
Instant	GPT-5.5 Instant	Everyday questions, writing, lookups, translation. 80% of your work.	Everyone (default)
Medium	GPT-5.6 Sol	Standard reasoning. Analysis, planning, moderately hard problems.	Plus, Pro, Business, Enterprise
High	GPT-5.6 Sol	Extended reasoning. Hard maths, legal analysis, multi-step code.	Plus, Pro, Business, Enterprise
Extra High	GPT-5.6 Sol	Maximum Sol effort. Slow, expensive, occasionally worth it.	Pro, Business, Enterprise
Pro	GPT-5.6 Sol Pro	The hardest tasks and long-running workflows.	Pro, Business, Enterprise

HEADS UP

Plus does not get Extra High or Pro. Plus gets Instant, Medium and High. Free and Go do not get GPT-5.6 Sol at all. Go has a separate, weaker 'Thinking' option (press + then Thinking, capped at 10 messages per 5 hours).

The other models

Model	What it is	Where you meet it
GPT-5.5 Instant	The default. Fast, and much better than it sounds.	Every plan, every chat
GPT-5.6 Terra	Balanced capability, speed and cost.	Not selectable in chat. Available in Work and Codex.
GPT-5.6 Luna	Fastest and cheapest of the 5.6 family.	Not selectable in chat. Available in Work and Codex (Plus and up).
GPT-5 Thinking Mini	Lightweight reasoning fallback.	All plans. Kicks in when you exhaust reasoning limits.
GPT-5.5 Instant mini	What you get after burning your Instant allowance.	Plus and Go, automatically
Legacy models	Older picks kept around for people who need them.	Plus and Pro only

TIP

Turn on **automatic switching**: model picker > Configure > let Instant escalate to Medium on its own.
Auto-switches do **not** count against your manual reasoning allowance, so this is free extra intelligence.

Which one should you actually use?

Start on Instant. Always. The single most common mistake right now is people cranking every task to maximum reasoning because they assume higher is better. It is not. On simple tasks a high-effort model over-thinks, invents problems that do not exist, and produces a worse answer more slowly.

Escalate when Instant gives you a surface-level answer, when the task has multiple dependent steps, or when being wrong is expensive.

THE CHEAPEST TRICK IN THIS GUIDE

Before running a big job, paste the prompt into a normal chat and ask: **"Which reasoning level should I run this at, and why?"** It will tell you. This alone will save you a lot of wasted time and usage on Pro-tier work that Medium would have nailed.

Context windows by plan

	Free	Go	Plus	Pro
Instant context	27K	54K	54K	128K
Instant input max	~12 pages	~40 pages	~40 pages	~250 pages
Reasoning context	Varies	256K	256K	400K
Reasoning input max	Varies	~320 pages	~320 pages	~680 pages

That number is the **total shared** window. System instructions, tools, personality and your memories all eat into it, so the space actually available for your input is smaller than the headline figure. This is why long chats get dumb: you are slowly filling the window with your own conversation.

05. Pricing: What You Really Get

Six plans. Here is the honest breakdown, including the bits OpenAI does not put in the marketing copy.

Plan	Price	The real story
Free	\$0	GPT-5.5 Instant with limited messages in a 5-hour window. No GPT-5.6. No scheduled tasks. Shows ads. ChatGPT Work only on desktop, limited.
Go	\$8/mo	More messages, more uploads, longer memory, scheduled tasks, you can create GPTs. Still no GPT-5.6 Sol. Still shows ads.
Plus	\$20/mo	The one most people should be on. GPT-5.6 Sol (Medium + High), 160 Instant messages per 3 hours, expanded deep research, Sites, Record mode, ChatGPT Work everywhere, ad-free.
Pro	\$100/mo	5x Plus usage. Adds GPT-5.6 Sol Pro, Extra High reasoning, maximum deep research and Codex.
Pro	\$200/mo	Identical features to \$100 Pro. You are buying 20x Plus usage instead of 5x. Nothing else changes.
Business	\$20/user/mo annual \$25/user/mo monthly	2-seat minimum. Adds company knowledge, workspace agents, admin console, SAML SSO, SOC 2. Your data is not used for training.
Enterprise	Custom	Adds SCIM, RBAC, IP allowlisting, data residency, compliance API logs, analytics dashboard, ISO certifications.

THE TWO PRO TIERS ARE THE SAME PRODUCT

This confuses everyone. \$100 Pro and \$200 Pro have **identical capabilities**. The only difference is how much you can use them (5x versus 20x Plus). Start at \$100. You can switch tiers anytime in Settings > My Plan.

Which plan is right for you

You just want to try it

Free. It is genuinely capable now. GPT-5.5 Instant is not a crippled model, it is the same default everyone else uses, you just get fewer messages.

You use it a few times a day for simple stuff

Go, at \$8. More messages, longer memory, scheduled tasks. But be clear about what you are **not** getting: no GPT-5.6 Sol, no Sites, no Record mode, and ads.

You use it for work

Plus, at \$20. This is the real product. The jump from Go to Plus is \$12 and it buys you the actual reasoning models, deep research, Sites, Record mode, full ChatGPT Work on every platform, and no ads. If you are choosing between Go and Plus, choose Plus. It is not close.

You hit limits constantly, or you run agents all day

Pro at \$100. Then \$200 only if you are actually running out at \$100. ChatGPT Work and Codex are usage-metered, and a single ambitious Work task at high effort can chew through a lot of allowance in one go.

You have a team

Business. Two seats minimum. The thing you are really buying is **company knowledge** (ChatGPT searching across your Slack, Drive, SharePoint and Jira with citations) and the guarantee that your data is not used for training.

TIP

Nonprofits get up to 75% off Business or Enterprise. Verified US K-12 teachers get **ChatGPT for Teachers free through June 2027**. Neither of these is advertised well.

HEADS UP

There is **no annual billing** for Go, Plus or Pro. Only Business and Enterprise have annual plans. If someone offers you a discounted annual ChatGPT Plus subscription, it is a scam.

06. Ads, and How to Kill Them for Free

Yes, ChatGPT has ads now. Testing started in the US in February 2026. They appear below the end of a response, labelled as sponsored, one per response.

Who sees them

Plan	Ads?
Free	Yes
Go	Yes
Plus, Pro, Business, Enterprise, Edu	No. Completely ad-free.

Ads are also never shown to under-18 accounts, never in Temporary Chats, and never near sensitive topics like personal health, mental health or politics. Political advertising is not allowed at all. Advertisers see aggregate clicks, not your chats, memories, email or location. And ads do not influence the answer itself, those are separate systems.

THE HIDDEN AD-FREE OPTION ON THE FREE PLAN

Almost nobody knows this exists. Go to **Settings > Ads controls > Change plan to go ad-free** and pick **"Reduce message limits."** You get zero ads on the free plan. The trade: fewer messages, and you lose image generation and deep research. You can switch back any time. **This is Free-plan only.** Go users cannot do it.

Ad controls worth knowing

Settings > Ad Controls lets you turn off ad personalisation, hide or report individual ads, see why you got a specific ad, and clear your ads data. Note that since February 2026, if personalisation is on, ads can draw on your past chats and memory. If that bothers you, that toggle is the one to find.

TIP

If ads bother you enough to think about them, the \$20 Plus plan removes them permanently and gives you the reasoning models. That is a better trade than crippling the free plan.

07. Navigating the Interface

The sidebar

Conversations grouped by date, with search, pin, rename, archive and delete. Above them sit Projects, Skills, Scheduled tasks, GPTs and the Plugin directory. On the desktop app you also get the mode switcher: **Chat, Work, Codex**.

The model picker

Top of the conversation. Shows your reasoning level, not a model name (see chapter 4). Click **Configure** inside it to turn automatic switching on or off.

The composer

The message box has more in it than people realise:

Control	What it does
+ / paperclip	Attach files. Images, PDFs, DOCX, XLSX, CSV, code. Drag and drop works.
@	Call a specific plugin or app directly. Type @ then the app name. Works inside scheduled tasks too.
Search toggle	Force a web search with citations.
Image	Native image generation and editing.
Mic	Voice. On mobile, with live camera.
Canvas	Side-by-side editor for writing and code.
Tools menu	Deep research, study mode, record, and (on business plans) company knowledge.

Keyboard shortcuts

Shortcut	Action
Enter	Send
Shift + Enter	New line without sending
Cmd/Ctrl + Shift + O	New chat
Cmd/Ctrl + Shift + C	Copy last response
Cmd/Ctrl + Shift + S	Toggle sidebar
Option + Space (Mac) / Alt + Space (Win)	Summon ChatGPT from any app (desktop app)
Both Command keys (Mac)	Appshot: send the front window to ChatGPT as context
Cmd/Ctrl + G	Search tasks (desktop app)
Cmd/Ctrl + /	Show all shortcuts

Response controls people never touch

Hover any response and hit the three-dot menu. You can retry **more concisely** or **with web search** in one click, with no re-prompting. You can also **branch** the conversation into a new chat, which forks it from that point without polluting the original thread. That last one is genuinely useful and almost nobody uses it.

08. ChatGPT Work

This is the headline feature of 2026 and the reason this guide needed rewriting. ChatGPT Work launched July 9, 2026.

Chat gives you an **answer**. Work gives you a **finished deliverable**. You give it a goal. It gathers context from your files, your connected apps and the browser, breaks the job into steps, shows you a plan, and then goes away and works on it, sometimes for hours, and comes back with a spreadsheet, a deck, a document, a report or a live website.

Under the hood it is Codex technology pointed at knowledge work instead of code. Which, as one YouTuber put it, makes it "Codex with the scary word removed." That is a fair jab, but it undersells what it does for a non-coder.

How to use it

STEP 1: Open the Work tab

In the desktop app sidebar, or on web and mobile if you are Plus and above.

STEP 2: Give it a goal, not a task

"Summarise this file" is a chat request. "Build me a competitor benchmark deck covering these five companies, using our Q2 numbers from Drive" is a Work request.

STEP 3: Use Plan mode

It will gather context, ask you clarifying questions, and produce a step-by-step plan. **Read the plan.** This is your one cheap chance to correct it before it spends an hour going the wrong direction.

STEP 4: Approve and walk away

It works in the background. On the desktop app it can keep working while your Mac is locked.

Five prompts that show what it can actually do

THE DAILY COMMAND CENTER

Review my inbox, my calendar and my Slack from the last 24 hours. Build me a one-page command center with: my top 3 priorities today, anything genuinely urgent, decisions that are waiting on me, and anything I said I'd do and haven't. Output as a PDF I can read in 60 seconds.

THE FOLDER ANALYSIS (DESKTOP ONLY)

Analyse every file in this folder. Build me a spreadsheet with one row per file: what it is, when it was last touched, what it's for, and whether it's still relevant. Flag duplicates and anything that looks out of date.

THE COMPETITOR BENCHMARK

Research these 5 competitors: [list]. For each: pricing, positioning, target customer, main weakness, and what they've shipped in the last 6 months. Cross-reference against our positioning in the attached doc. Deliver as a slide deck with a one-slide 'so what' at the end.

THE RECURRING REPORT

Every Monday at 8am, pull last week's numbers from [connected tool], compare to the previous 4 weeks, and send me a one-page summary. Call out anything that moved more than 15% in either direction and give me your best guess as to why.

THE MESSY-DATA CLEANUP

Take the attached spreadsheet. It's a mess. Standardise the columns, dedupe the

rows, fix the date formats, flag anything that looks like a data-entry error, and give me back a clean .xlsx plus a short note on what you changed and why.

Availability

Plan	Where Work runs
Free	Desktop app only, limited
Go	Desktop app only, limited
Plus	Desktop, web and mobile
Pro	Desktop, web and mobile, expanded
Business / Enterprise	Everywhere, with admin spend controls

HEADS UP

Work is usage-metered, not unlimited. It follows the same credit structure as Codex. A big task at high reasoning effort burns through allowance fast. One creator reported a five-part marketing plan at Extra High taking 33 minutes and a serious chunk of his quota. Budget it. Do not point it at trivial jobs.

SCHEDULED WORK IS THE REAL UNLOCK

Work tasks can run once, on a schedule, on an event trigger, or as a monitor that watches for changes. That is the difference between a clever assistant and an employee. Set up three recurring Work tasks and you will feel it immediately.

What companies are actually doing with it

Company	The workflow	The result
Zapier	Lead triage across HubSpot, Gong and email	A lead used to take 35-45 minutes manually. Now it feeds a weekly exec dashboard.
NVIDIA	GTC event prep, replacing a spreadsheet process	Pre-event prep used to eat ~40% of the team's time. Now runs twice weekly.
Virgin Atlantic	Competitor benchmarking for a 5-year plan	Weeks of work compressed into hours.
RingCentral	Launch-readiness checks across Jira and Salesforce	Scaled from supporting 1 PM to around 50.

The pattern in all four: a repetitive, cross-tool, context-gathering job that a smart human was doing badly because it was boring. That is the shape of a good Work task.

09. The New Desktop App

On July 9, 2026 OpenAI merged the Codex app and the ChatGPT desktop app into one application. The old app was renamed **ChatGPT Classic**. If you are still using it, you are using the wrong app.

The new desktop app is available **on every plan, including Free**, on macOS and Windows. And it is now meaningfully more capable than the website. This is new. For two years the web app was the full product and desktop was a convenience wrapper. That is over.

What only the desktop app can do

Feature	What it means
Built-in browser	A real browser inside ChatGPT. Multiple tabs, downloads, logins, saved passwords, cookies, Chrome profile import, live previews.
Computer Use	ChatGPT clicks, types and moves files around your actual machine. Runs in the background. On Windows it takes over the foreground during a task.
Local file access	Point it at a folder. "Analyse every file in here." The web app cannot do this.
Appshots (macOS)	Press both Command keys to send the front window (screenshot plus text) straight into ChatGPT as context.
Record and Replay (macOS)	Demonstrate a workflow once. It turns it into a reusable skill.
Record mode	Meeting transcription with speaker labels. Plus and above.
Three modes in one window	Chat, Work and Codex in the sidebar, plus Scheduled and Sites.
Code edits on macOS	Direct edits in supported IDEs.

WORK WITH APPS: THE LOCAL APP LIST

On macOS, ChatGPT can read from (and in IDEs, edit) these directly: **Apple Notes, Notion, TextEdit, Quip, Xcode, Script Editor, VS Code** (plus Cursor, Windsurf, VSCodium), **the JetBrains suite, Terminal, iTerm, Warp and Prompt**. Editing is IDE-only: it shows you a diff, you review, you apply, and Cmd+Z reverts. It reads the last 200 lines of an open pane, and focuses on your selection if you highlight text first.

TIP

Option + Space (Mac) or **Alt + Space** (Windows) summons ChatGPT from anywhere on your machine without switching apps. If you learn one shortcut from this guide, learn that one. It is the difference between using ChatGPT and living in it.

Remote: your phone as a remote control

Pair the mobile app with a Mac or Windows machine running the desktop app (QR code pairing, available on all plans). You can start tasks, check on running ones, browse workspace files, review diffs and approve actions from your phone while the heavy lifting happens on your desktop.

10. The Built-In Browser and Computer Use

Atlas, OpenAI's standalone browser, is being shut down on August 9, 2026. What replaced it is better: a browser built directly into the desktop app, plus a Chrome extension that runs in Chrome's sidebar.

Why the built-in browser matters more than it sounds

It supports multiple tabs, saved passwords, cookies, and importing your Chrome profile. Read that again. It means ChatGPT can now operate inside the login-required SaaS tools that have no API and no plugin. All the tedious web-based admin work that was previously impossible to automate because the tool had no integration is now automatable.

One creator's example: a podcast host who spent 45 minutes a day moving information between five tools, none of which had an API. Previously unautomatable. Now a scheduled Work task.

Computer Use

This goes a step further. ChatGPT clicks, types and moves files across your local apps, tools and browser. On macOS it runs in the background, and can even continue approved work after your machine locks. On Windows it takes over the foreground while a task runs. It can be triggered as part of a Scheduled Task.

A REAL COMPUTER USE TASK

```
Open the invoices folder. For each PDF,
pull out the vendor, date, amount and
invoice number. Put them into a new
spreadsheet, one row each. Rename each
PDF to YYYY-MM-DD_Vendor_Amount.pdf.
Move anything older than 12 months into
an Archive subfolder.
```

Safety, and why you should care

Agentic browsing is the single biggest new attack surface in AI. A malicious webpage can contain hidden instructions aimed at your agent. OpenAI runs an auto-review system where advanced models check important actions on connected tools before they execute, and sensitive actions (purchases, form submissions, account changes) require your approval.

HEADS UP

Do not let an agent loose in a browser session that is logged into your bank, your payroll system or anything that moves money. The safeguards are good. They are not perfect, and neither is any competing product. Treat agent permissions like you would treat handing someone your unlocked laptop.

11. Deep Research

Deep Research is an autonomous research agent. Instead of answering from what it knows, it writes a research plan, browses dozens of sources, cross-references them, and produces a structured report with inline citations. Typically 3 to 10 minutes per query.

How it works

You ask. It plans and shows you the plan. It browses. It synthesises. You can **interrupt it mid-run** to add constraints, which is the feature most people never discover.

What changed in 2026

Upgrade	What it lets you do
Plugin sources	Point deep research at your connected apps and internal data, not just the open web.
Trusted-site restriction	Restrict the search to specific domains you trust. Kills most of the slop problem.
Mid-run steering	Interrupt and refine while it is still working.

A GOOD DEEP RESEARCH PROMPT

```
Research the top 10 project management
tools for small teams as of 2026.
For each: pricing, key features,
integrations, and what real users
complain about.

Restrict sources to vendor sites, G2,
and Reddit. Ignore affiliate blogs and
listicles.

Give me a comparison table and a single
recommendation for a team of 5-10 with a
budget under $20/user/month. Tell me what
would make you change your recommendation.
```

TIP

That last line, **"tell me what would make you change your recommendation,"** is worth adding to almost every research prompt. It forces the model to expose its own uncertainty instead of presenting a confident guess.

Limits

Free and Go get limited deep research. Plus gets expanded. Pro gets maximum. OpenAI no longer publishes exact per-plan numbers, so ignore any guide that quotes a specific figure. Check the usage counter in the product. For plans with a fixed monthly allowance, it resets 30 days from first use.

12. Voice: GPT-Live

Advanced Voice Mode was replaced by **GPT-Live** on July 8, 2026. The difference is architectural, not cosmetic.

GPT-Live is **full-duplex**: it listens and speaks at the same time. It backchannels ("mhmm", "right") while you talk. It waits when you pause mid-thought instead of jumping in. It ignores background noise. And when you ask it something hard, it hands the question off to a frontier model in the background while keeping the conversation going, so you do not sit in silence.

What you get

Feature	Detail
Full-duplex conversation	It can listen while speaking. Interrupting it finally feels natural.
Visual answer cards	Weather, stocks, sports, maps rendered on screen mid-conversation.
Reasoning levels in voice	Instant, Medium and High.
Nine voices	All remastered.
Voice with video	Camera and screen share. Go and above (not Free).

Free users get GPT-Live-1 mini. Go, Plus and Pro get the full GPT-Live-1.

TIP

Two settings people miss: turn on **auto-detect language** in voice settings, and check which voice model you are actually on. The default can quietly be the weaker one. Also switch on **dictation** so you can talk straight into the composer without starting a full voice session.

The best use of voice that nobody does

Stop treating it as a novelty. Use it as a thinking partner while you walk. Talk through a problem for ten minutes with no structure, then finish with:

THE WALK-AND-TALK CLOSER

Okay, I've been rambling. Take everything
I just said and give me back: the actual
decision I'm trying to make, the three
options I circled around, what I clearly
already want to do, and the thing I'm
avoiding saying out loud.

13. Image Generation

ChatGPT Images 2.0 landed in April 2026. It is the first image model with native reasoning built in, it is roughly four times faster than the previous generation, and it added a selection tool for editing specific regions of an image instead of re-rolling the whole thing.

What's new that matters

Feature	What it does	Plans
Region editing	Select an area of an image and change only that.	All
Aspect ratio control	Set it in the UI instead of begging in the prompt.	All
Image generation with Thinking	Searches the web while generating, batches up to 8 images, verifies its own output.	Plus and Pro only

Prompts that actually work

THE THUMBNAIL RECREATION

Here's a screenshot of a thumbnail whose style I like. Recreate that exact style, but the headline should read [YOUR TEXT] and replace the central image with the photo I've attached. Keep the colour grading and the text weight identical.

THE INFOGRAPHIC

Turn the attached document into a single vertical infographic. 5 sections. Clean sans-serif type, lots of white space, one accent colour (#E8834A). No stock-photo people. No fake charts. Every number on the graphic must come from the document.

THE PRODUCT SHOT FIX

Use the selection tool on the background only. Replace it with a soft neutral studio backdrop. Do not touch the product itself, the lighting on it, or its shadow.

HEADS UP

Text inside generated images is still the weak spot. If the headline matters, generate the image without text and add the type yourself. Do not spend forty minutes re-rolling because it keeps spelling your brand name wrong.

14. Canvas and Sites

Canvas

Canvas is a side-by-side editor. Your document or code sits in a panel next to the chat, and you edit it collaboratively rather than regenerating the whole thing every time. Highlight a paragraph, ask for a change, only that paragraph changes.

It is available on **every plan, including Free**. Despite what you may have read, it is not dead.

Where it earns its keep:

- Long-form writing where you want to keep 90% and fix 10%
- Code you are iterating on rather than generating once
- Anything where the previous version mattered and you do not want it rewritten from scratch

Sites

Sites lets you build and publish a real, hosted website or web app with a live URL and built-in analytics, straight out of a conversation. No deployment, no hosting setup, no code required.

Plus and Pro only. Not on Free or Go.

BUILD A SITE FROM A CONVERSATION

```
Build me a single-page site for [thing].  
Sections: hero with one clear promise, 3  
benefit blocks, social proof, FAQ, and one  
email capture form.  
  
Design: dark navy and gold, Helvetica,  
lots of white space, mobile-first. No  
stock photos, no gradients, no emoji.  
  
Publish it and give me the URL.
```

TIP

Sites plus a Work task is a genuinely strong combo: have Work research and write the content, then have it publish the result as a live page. Internal dashboards, one-pagers and client reports go from a document nobody opens to a link people actually click.

15. Codex

Codex is the autonomous coding agent. It writes, tests, debugs and ships code, handles multi-file and now multi-repository projects, and keeps running until the task is done.

Two numbers worth knowing: over five million people use Codex weekly, and more than a million of them use it for **non-coding** work. That second number tells you what it has actually become. It is a general-purpose task agent that happens to be very good at code.

Where Codex lives

Surface	What it is
ChatGPT desktop app	Codex mode in the sidebar. The main way most people use it now.
ChatGPT web	Cloud-based tasks.
Codex CLI	Terminal.
Codex IDE extension	Inline in VS Code and friends.
Codex cloud	Long-running sandboxed tasks.
Integrations	GitHub, Slack, Linear. Plus an SDK and a GitHub Action.

If you are not a developer, read this bit

You do not need to write code to get value out of Codex. What you need is a task that is repetitive, file-based, and currently done by hand. Renaming and sorting a thousand files. Converting a folder of documents. Cleaning up a data export. Building a small internal tool. Codex will do all of that and you never see the code.

CODEX FOR NON-CODERS

```
I have a folder of 400 client PDFs with
inconsistent names. Build me a small tool
that reads each one, pulls out the client
name and date, renames the file properly,
and sorts it into a folder by year.
```

```
Show me it working on 5 files first. Don't
touch the rest until I say go.
```

TIP

That last instruction, **"show me it working on 5 files first, don't touch the rest until I say go,"** should be in every destructive task you ever give an agent. Test on a sample. Always.

16. Plugins (Formerly Connectors and Apps)

The naming here has changed twice in seven months, which is why everyone is confused. Connectors became Apps in December 2025. On July 9, 2026, the App Directory became the **Plugin Directory**.

A **plugin** is a packaged workflow. It can bundle **skills** (reusable instructions), **apps** (the actual integrations and connections), and templates. There are over 1,400 of them, at chatgpt.com/plugins.

What's connectable

Category	Examples
Files and docs	Google Drive, SharePoint, Box, Dropbox, Adobe Acrobat
Comms	Gmail, Outlook, Slack, Microsoft Teams
Calendar	Google Calendar, Outlook Calendar
Work tracking	Jira / Atlassian, GitHub, GitLab, Azure Boards
Data	Snowflake, Databricks, BigQuery
Business	Salesforce, Stripe
Custom	Anything, via developer mode and the Apps SDK (built on MCP)

How to actually use them

Type @ in the composer and name the app. ChatGPT will also auto-select plugins based on what you asked, and suggest them mid-conversation. The @ trick also works **inside a scheduled task**, which is how you build a real automation.

HEADS UP

The setting that ruins this for most people: Settings > Personalization > Advanced > **connector search**. If it is off, ChatGPT will not proactively look in your Drive, Gmail or Slack to answer you. Half the people complaining that "ChatGPT can't see my files" simply never flipped this switch.

Interactive apps inside the chat

Separate from plugins, there are apps that render an actual interactive interface inside the conversation: Canva, Spotify, Figma, Zillow, Booking.com, Expedia, Coursera, DoorDash, Uber, Instacart, Target, Wix, Quizlet. Start the prompt with the app name ("Spotify, build me a playlist for a 3-hour drive") or use @. Available on all logged-in plans, though not currently in the EEA, Switzerland or the UK.

PERMISSIONS ARE GRANULAR ON PURPOSE

Each plugin can be set to **always ask**, **allow read actions**, or **allow low-risk actions**. There is deliberately no "allow everything" button. That is a feature, not an oversight. Read actions are safe. Write actions are where things go wrong.

17. Skills

Skills are new, they are in beta, and they are on **every plan including Free**. They are also the most underrated feature in ChatGPT right now.

A skill is a reusable set of instructions. Think of it as a saved standard operating procedure that ChatGPT loads when it is relevant. Instead of re-explaining how you want your emails written, or how your reports should be structured, or what your brand rules are, you write it down once as a skill and it is applied every time.

Three obvious ways to make one, and one clever one

You can build a skill in the chat, in the skills editor, or by uploading a document. All fine. But the best method is the one nobody talks about: **let ChatGPT tell you which skills you need**, based on how you already use it.

THE SKILL AUDIT (RUN THIS TODAY)

```
Based on how I've interacted with you in
the past, and everything in your memory
about me, what skills should we create so
you give me better answers with less
explaining from me every time?
```

```
Be specific. Name each skill, say what it
would contain, and tell me how often it
would actually get used.
```

Then follow up with:

TURN IT INTO SKILLS

```
Good. Now write the actual prompts and
instructions we need to create those
skills. Full text, ready to paste.
```

TIP

Any time ChatGPT produces something you love, say: **"Turn this into a skill I can reuse."** That single sentence is how you stop rebuilding the same output every month.

DO THIS NOW

Run the skill audit prompt above. It takes 60 seconds and it will show you three or four things you have been re-explaining to ChatGPT for months without realising it.

18. Memory and Personalization

Memory was rebuilt in June 2026 and is now considerably smarter. It updates itself: when something you told it becomes stale, it revises the memory instead of holding two contradictory facts. Plus and Pro get double the capacity.

The transparency upgrade nobody noticed

You can now see **exactly which chats, memories and files** ChatGPT drew on for any given response, and mark those sources as relevant or not relevant. This is a quiet but major feature. If ChatGPT keeps giving you answers coloured by an old project you have moved on from, you can now find out why and fix it.

The settings that matter

Setting	What to do with it
Reference saved memories	On. This is the core of it.
Reference chat history	On. Lets it draw on past conversations, not just explicit memories.
Reference browser memories	On if you use the built-in browser.
Reference record history	On if you use Record mode for meetings.
Memory sources	Use this to see and correct what it is drawing from.
Temporary Chat	For anything you do not want remembered. Also: no ads in Temporary Chats.

THE MOST IMPORTANT STRATEGIC POINT IN THIS GUIDE

Models change every three months. Your memory and context layer does not. Whoever wins the next benchmark is irrelevant to you if ChatGPT knows your business, your voice, your customers and your standards. **Invest in the memory layer, not in model loyalty.** That is the asset that compounds.

SEED YOUR MEMORY DELIBERATELY

Remember the following, permanently:

MY BUSINESS

[what you sell, to whom, at what price]

MY VOICE

[3 rules about how you write]

MY STANDARDS

[what a good output looks like to you,
and what an unacceptable one looks like]

WHAT I'M WORKING ON THIS QUARTER

[2-3 goals]

19. Projects

Projects are folders with a brain. Each project holds its own chats, its own uploaded files, and its own custom instructions. Everything inside inherits that context.

They are on every plan now, including Free, and there are **shared projects** too, so a team can work in the same context.

How to set one up so it is actually useful

STEP 1: Create the project and name it after an outcome

Not "Marketing." Try "Q3 launch, target 5,000 signups."

STEP 2: Upload the context files, not everything

Brand guidelines, past examples of good work, the strategy doc, key data. Resist the urge to dump your entire Drive in. Context is a budget, and junk files spend it.

STEP 3: Write project instructions

This is the step everyone skips, and it is the one that makes it work.

PROJECT INSTRUCTIONS TEMPLATE

WHAT THIS PROJECT IS

[one sentence]

THE GOAL

[what success looks like, with a number]

THE FILES IN HERE

[what each one is and when to use it]

HOW TO WORK IN THIS PROJECT

Always check the brand doc before writing anything public-facing. Never invent a statistic. If you're not sure, say so. Default to short.

TIP

One project per outcome, not per topic. A project that runs for six months slowly fills with dead context. When the outcome is achieved, archive it and start a fresh one. Projects are cheap. Confusion is not.

20. Scheduled Tasks

Scheduled tasks are how ChatGPT stops being something you visit and starts being something that shows up. They are available on **Go and above** (not Free), and there is now a Scheduled hub in the sidebar.

Four kinds of schedule

Type	Example
Once	"Tomorrow at 7am, remind me to send the invoice and draft it for me."
Recurring	"Every Monday at 8am, give me last week's numbers with commentary."
Event-triggered	"When a new message lands in #launches, update the launch doc."
Monitor	"Watch this page. Tell me if the pricing changes."

The prompt that builds your automations for you

Do not sit and try to think of automations. Ask it what you already do repeatedly:

THE AUTOMATION AUDIT

Based on your memory of me and the things I ask you over and over, give me a list of tasks I should schedule so you can just do them automatically.

Be as granular as possible. For each one: the exact prompt, the cadence, and which apps it needs. Rank them by how much time they'd save me per month.

TIP

Inside a scheduled task, type @ to pull in a plugin. That is the difference between a scheduled reminder and a scheduled **action**.

HEADS UP

Scheduled tasks have real limits: roughly one run per hour maximum, plain-text prompts only, and no voice or file uploads inside the task itself. Build accordingly.

21. Data Analysis and the Office Extensions

Upload a CSV or spreadsheet and ChatGPT will write and run code to analyse it, then hand you back charts and findings. That has been true for a while. What is new is that it now plugs directly into the tools you already use.

The Excel, PowerPoint and Google Sheets extensions

ChatGPT now has extensions for Excel, PowerPoint and Google Sheets. Limited on Free and Go, full on Plus, expanded on Pro. It can open and refine Google Workspace and Microsoft 365 files directly. This closes the most annoying gap in the product: output that lived in a chat window instead of in the file you actually work in.

Interactive tables and charts

Plus and Pro get interactive tables and charts in the response itself, so you can sort and explore rather than staring at a static image.

THE ANALYSIS PROMPT THAT AVOIDS GARBAGE

```
Here's the data. Before you analyse
anything:

1. Tell me what's wrong with this dataset.
Missing values, duplicates, weird
outliers, columns that don't mean what
they're named.

2. Tell me what questions this data CAN'T
answer, so I don't ask them.

Then, and only then, answer this: [Q].
Show your working. If a finding is based
on fewer than 30 rows, say so explicitly.
```

TIP

Asking what the data **cannot** tell you, before you ask it anything, is the single best defence against a confident, beautifully-formatted, completely wrong analysis.

HEADS UP

ChatGPT will happily find a pattern in noise and present it with total confidence. It does not know your business context. Every number it gives you that you plan to act on, check.

22. ChatGPT for Business and Enterprise

Business is \$20 per user per month billed annually, or \$25 monthly, with a two-seat minimum. Enterprise is custom priced. Both guarantee that **your content is not used to train OpenAI's models**.

Company knowledge: the actual reason to buy Business

Company knowledge lets ChatGPT search across your connected apps (Slack, Google Drive, SharePoint, GitHub, Jira and so on) and answer with citations and links back to the source. It respects existing permissions, so people only ever see what they already had access to.

Turn it on per chat from the composer, then choose which apps to include.

HEADS UP

Company knowledge only works on ChatGPT Web. Not on desktop. Not on mobile. This is the single most common thing people get wrong about Business, and it is the kind of detail that makes a rollout look broken when it is not.

Workspace agents

Workspace agents are shared, Codex-powered agents. OpenAI describes them as the evolution of Custom GPTs, and that is accurate. They run in the cloud, they have memory, they can write and run code, use your connected apps, run on a schedule, and be deployed straight into Slack.

You build one by describing the workflow in the sidebar. There are templates for finance, sales and marketing. Real examples from OpenAI's own customers: a lead outreach agent, a product feedback router, a weekly metrics reporter, a third-party risk manager, and a month-end close agent that produces accounting workpapers.

Currently a research preview on Business, Enterprise, Edu and Teachers. Not available on any consumer plan.

Business versus Enterprise: what you do and do not get

Capability	Business	Enterprise
SAML SSO, admin console, SOC 2 Type 2	Yes	Yes
Company knowledge, workspace agents	Yes	Yes
No training on your data	Yes	Yes
SCIM provisioning	No	Yes
Role-based access controls	No	Yes
IP allowlisting	No	Yes
Analytics dashboard	No	Yes
Compliance API logs	No	Yes
Data residency (US, EU, UK, JP, CA, KR, SG, IN, AU, UAE)	No	Yes
ISO 27001 / 27017 / 27018 / 27701	No	Yes
Enterprise key management	No	Yes

A lot of blogs get this wrong and tell you Business includes SCIM and data residency. It does not. If your security team needs any row in bold above, you need Enterprise.

SPEND CONTROLS

ChatGPT Work and Codex are usage-metered. Enterprise and Edu admins can set workspace defaults, per-group limits and individual overrides in the Admin Console, with a request-more-credits flow. If you are rolling this out to a team, set the controls before you announce the feature, not after.

TIP

The highest-ROI first deployment for most teams is not a fancy agent. It is turning on company knowledge and letting people ask "where is the doc about X" and get a real answer with a link. Do the boring thing first.

23. Settings, Privacy and Safety

The settings worth changing today

Setting	Where	Why
Custom instructions	Personalization	The highest-leverage five minutes in the product.
Web search + connector search	Personalization > Advanced	Off by default for many. This is why it 'can't see' your files.
Automatic reasoning switching	Model picker > Configure	Free extra intelligence. Doesn't count against your allowance.
Memory sources	Memory	See and correct what it's drawing on.
Ad controls	Advertising	Free and Go only. Includes the ad-free option.
Data controls	Data controls	Opt out of training on your content.
Temporary Chat	Composer	No memory, no history, no ads.

Privacy, plainly

On Free, Go, Plus and Pro, your content **can** be used to train OpenAI's models, and there is an opt-out. On Business, Enterprise and Edu, it is not used for training, full stop. That difference is worth more than any feature on the Business plan for a lot of companies.

Data is encrypted in transit (TLS 1.2) and at rest (AES-256). Enterprise adds data residency across ten regions.

Parental controls and age prediction

ChatGPT now predicts whether an account belongs to someone under 18 and defaults to restrictive settings when it is unsure. Parents can link a teen's account by email invite, disable memory and history, and receive alerts in cases of acute distress. Parents do not get access to the conversations themselves. Ads are never shown to under-18 accounts.

HEADS UP

The genuine risk in 2026 is not the model saying something silly. It is **prompt injection**: a webpage or document containing hidden instructions aimed at an agent that has access to your accounts. Be deliberate about what you connect and what write permissions you grant. Read-only is safe. Write is where the damage happens.

24. Prompting for the Reasoning Era

Everything you learned about prompting in 2023 is now actively hurting you. The long, elaborate, role-play-heavy, step-by-step-scaffolded prompt was a workaround for a model that could not think. GPT-5.6 can think. Stop doing its thinking for it.

The single biggest change: outcome over process

OpenAI's own guidance is that short outcome-focused prompts beat step-stacked ones.

THE OLD WAY (STOP DOING THIS)

```
You are a world-class expert marketing
strategist with 20 years of experience.

Take a deep breath and think step by step.
First, analyse the market. Then identify
the segments. Then consider the
competitors. Then think about positioning.
Then, and only then, write the plan...
```

THE NEW WAY (DO THIS)

```
Market analysis for a sustainable hiking
gear brand.

I need:
- Top 3 customer segments, ranked
- Top 3 competitors and each one's real
weakness
- A market size estimate with your working

Tell me what you're least confident about.
```

Six rules that hold up in 2026

Rule	Why
State the outcome, not the method	The model plans better than your scaffold does.
Plan in one chat, execute in another	Long chats degrade. Get the plan, then open a fresh chat and say 'implement step one.'
Ask for its uncertainty	'Tell me what would change your answer' surfaces the weak link.
Give it your standards, not just your task	Show it one example of good and one of bad. Worth more than 500 words of description.
Don't max the effort by default	High reasoning on a simple task produces a worse answer, slower.
Kill the chat when it gets dumb	It's not broken. The context window is full of your own conversation. Start fresh.

TIP

The 'plan in one chat, execute in another' rule is the one that separates people who get great output from people who get mush. Reasoning models degrade over a long session. Treat the plan as an artefact you carry to a clean room.

25. 25 Power Tips

The things that actually move the needle, roughly in order of how much time they will save you.

- 01. Turn on connector search.** Settings > Personalization > Advanced. Off by default. This is the number one reason people think ChatGPT cannot see their Drive or Gmail.
- 02. Turn on automatic reasoning switching.** Model picker > Configure. Auto-switches do not count against your manual reasoning allowance. Free intelligence.
- 03. Stop maxing out the reasoning level.** On a simple task, high effort makes the answer worse, not better. It overthinks and invents problems.
- 04. Ask it which model to use.** Paste your prompt into a normal chat and ask which reasoning level to run it at. Then go run it there.
- 05. Learn Option + Space.** Or Alt + Space on Windows. Summons ChatGPT from any app on your desktop. The single highest-frequency shortcut in the product.
- 06. Use the desktop app, not the website.** Only the desktop app gets local file access, the built-in browser, Computer Use, Appshots and Record and Replay.
- 07. Run the skill audit.** "Based on how I've interacted with you, what skills should we create?" It knows your patterns better than you do.
- 08. Run the automation audit.** Same trick for scheduled tasks. Ask it what you keep asking it, then schedule those.
- 09. Type @ inside a scheduled task.** That pulls a plugin into the automation. It is the difference between a reminder and an action.

- 10. Plan in one chat, execute in another.** Reasoning models degrade in long sessions. Carry the plan to a clean chat.
- 11. Branch, don't restart.** Three-dot menu on any response > branch into a new chat. Forks the conversation without polluting the original.
- 12. Use the one-click retry options.** Three-dot menu lets you retry more concisely or with web search, no re-prompting needed.
- 13. Point Work at a folder.** "Analyse every file in this folder." Desktop only. Genuinely feels like magic the first time.
- 14. Test destructive tasks on 5 items first.** "Show me it working on 5 files, don't touch the rest until I say go." Every single time.
- 15. Restrict deep research to trusted sites.** Kills the affiliate-blog slop problem entirely.
- 16. Interrupt deep research mid-run.** You can steer it while it works. Almost nobody knows this.
- 17. Add "tell me what would change your answer".** Forces it to expose uncertainty instead of guessing confidently.
- 18. Ask what the data can't tell you.** Before any analysis. Best defence against a beautiful, confident, wrong chart.
- 19. Seed your memory deliberately.** Do not wait for it to learn you. Tell it who you are, what good looks like, and what you are working on.

20. Check your memory sources. You can now see exactly which chats and files it drew on, and mark them not relevant.

21. Free plan users: get ad-free. Settings > Ads controls > change plan to go ad-free. Trades message limits for zero ads.

22. Use Temporary Chat for anything sensitive. No memory, no history, no ads.

23. Voice-and-walk, then ask it to summarise your own rambling. "What decision am I actually trying to make, and what am I avoiding saying?"

24. Turn a great output into a skill immediately. "Turn this into a skill I can reuse." One sentence, permanent leverage.

25. Invest in the memory layer, not the model. Models change every 90 days. Your context does not. That is the asset that compounds.

26. The 90-Day Reality Check

A guide that only tells you how good something is, is an advert. So here is the honest part.

What ChatGPT is genuinely best at right now

- Agentic work. Nothing else touches the desktop app's combination of a real browser, local file access and Computer Use.
- Research at depth, with citations.
- Coding, via Codex.
- Voice. GPT-Live is not close to anything else.
- Being the default everyone already has, with the biggest plugin library by a distance.

Where it is still weak

- Long-form writing with genuine voice. It still writes a bit academically. Claude is generally better here and it is not particularly close.
- Text inside generated images.
- No video generation at all, now that Sora is gone.
- Naming and menus. Connectors became apps became plugins in seven months. Features move. Things you rely on get cut.
- Sync bugs between the Work tab and the Chats panel on desktop, at the time of writing.

The fair criticism of ChatGPT Work

One well-regarded AI YouTuber called ChatGPT Work "Codex with the scary word removed," and argued that on the web it is largely a repackaging of things that already existed. He has a point. If you are technical and you were already using Codex, Work is a rebrand with better defaults.

But that is exactly why it matters. Codex was sitting there the whole time and normal people would not touch it, because it looked like a developer tool. Removing the scary word is not nothing. It is the difference between a capability existing and a capability being used.

YOUR FIRST WEEK

Day 1: Custom instructions, and turn on connector search. **Day 2:** Install the new desktop app. Learn Option + Space. **Day 3:** Run the skill audit and the automation audit. **Day 4:** Set up your three most valuable scheduled tasks. **Day 5:** Give ChatGPT Work one real, ugly, cross-tool job you have been putting off. **Day 6 and 7:** Seed your memory properly. That is the part that compounds.

A NOTE ON DATES

Everything in this guide was accurate on July 12, 2026. Given how fast OpenAI ships and cuts, assume something in here will be wrong within a quarter. When it is, the principles in chapters 24 and 25 will still hold. The features are rented. The habits are owned.

Want More AI Tips Like This?

I post daily AI content, tool breakdowns, and prompts that actually help you get more done. Follow me for the latest.

@MavGPT

TikTok | Instagram | Facebook

Stop asking ChatGPT questions. Start giving it jobs.