The Complete ChatGPT Guide

Every model, every plan, every feature: GPT-6 and the Thinking slider, Work, Codex, the desktop app, Space, dots, browser tools, plugins, skills and 25 power tips.

By Maverick Maltin · Last verified October 16, 2026 · 28 min read

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 takes a goal and works on it for hours until there is a finished spreadsheet or deck sitting in front of you.

The one-line version: ChatGPT is 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.

In regular Chat, the model menu now just says GPT-6, with a Thinking slider for how hard it thinks. In Work and Codex you pick from models like GPT-6 Astra for demanding work, GPT-6.1 Sol for everyday work and coding, and GPT-6 Luna for focused tasks. Here is how to make sense of it.

Getting Started on Every Platform

Setup takes about two minutes. Where you use ChatGPT matters, because the three platforms are not equivalent.

Web (chatgpt.com)

Sign up with email, Google, Microsoft or Apple. Free includes unlimited everyday text chats on GPT-5.6 Luna, subject to abuse guardrails. Uploads, images, research and agent work have separate limits. Start at chatgpt.com, then use Work when you want a finished result.

Desktop app: the one that changed everything

Download from chatgpt.com/download. The desktop app brings Chat, Work and Codex together with a built-in browser, Computer Use, Appshots and access to local files. It is available on every plan, including Free, on macOS and Windows, with Linux in public preview. Synced Work can also reach an enabled, online computer for local steps.

Mobile (iOS and Android)

Full voice, camera, and 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.

PROMPT 1 · COPY & PASTE

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.

Then connect the apps you actually use through Plugins. Select the service, sign in, and grant the access you want ChatGPT to have. In your next task, name the app with @ or ask it to use that connected source. Connecting an account and enabling access in a task are separate steps.

What changed at Dev Day on September 29?

The big everyday additions are dots, synced Work and Space. Their sections below explain where they fit. GPT-6.1 Sol and Astra Ultrafast also expand the available model choices.

A few more additions matter when you start building or working with a team:

- **Codex Cloud:** Reusable development environments let coding tasks continue from desktop, web or mobile while your computer sleeps.
- **Security Cloud:** A research-preview plugin can inspect connected repositories and propose fixes for review.
- **Connected Sites:** Supported workspace-private Sites can use each visitor's approved app connections.
- **Plugin interfaces:** Plugins can include panels, editors and forms. MCP Events support work triggered by updates from connected tools.
- **Website annotations:** Supported desktop browsers can attach selected page content and preview controls to feedback.
- **For teams:** Slack and Teams deployments, Team Tasks, shared workspace connections and Agent Security give organizations more ways to manage work and access.
- **For developers:** GPT-6.1 Sol and Astra Ultrafast have separate API options, billing and availability.

Use the sections relevant to your work. The full Dev Day announcement links to setup and rollout details for each addition.

The Models, Explained Properly

Start with the job. The most expensive model is not necessary for every question.

Chat, Work and Codex use different controls

In regular Chat, the model menu shows **GPT-6** with a **Thinking** slider: Instant, Medium, High, Extra High and Pro. Plus gets Medium and High. Pro, Business and Enterprise get all five, and Pro uses GPT-6 Pro, powered by GPT-6 Astra. Free and Go use a lighter model (GPT-5.6 Luna) with a **Think** button for harder questions. In Work and Codex, the model and reasoning control beneath the composer lets you choose an available model and effort level. A model appearing in Work does not mean it is also available in regular Chat.

GPT-6.1 Sol, Astra and Luna

- **GPT-6.1 Sol:** A starting point for everyday work, code fixes and tasks that need several steps. Use it when it appears in your picker.

- **GPT-6 Astra:** The choice for the hardest work across research, apps, design and code. I like Astra for design.
- **GPT-6 Luna:** Quick summaries, extraction and smaller tasks where speed matters.
- **GPT-6 Sol:** Still available during the 6.1 rollout. A useful choice when 6.1 has not reached your account.

GPT-6.1 Sol, GPT-6 Sol and GPT-6 Luna are Work and Codex models. In Chat, GPT-5.6 Sol can still be picked as a separate option if you prefer it. Free and Go have limited desktop agent access, while paid plans expand access and usage.

Reasoning effort

Start with the default effort. Raise it for a difficult decision, a large code change or a task that needs careful planning. Higher effort takes longer and uses more of your allowance. Ultra can split work across subagents, which makes more sense for a large project than a quick question.

The older models

GPT-6 replaced "Latest" in the Chat model menu on October 7, 2026, and GPT-5.5 retired from ChatGPT, Work and Codex on **October 14, 2026**. If an older scheduled task or saved default pointed at GPT-5.5, open it and check it still runs the way you want. The OpenAI API was not part of that retirement.

How much can it read?

Context is the material a model can work with at once. In regular Chat it depends on your plan and Thinking level: paid plans get more than Free, Pro gets the most, and higher Thinking levels can hold more. The usable space also includes the conversation and the answer, so a very long chat eventually crowds out your files.

Work, Codex and API models have their own limits. Do not assume an API context size is your ChatGPT upload allowance. For a pile of documents, give Work a folder or project and ask it to work through the relevant files.

Pricing: What You Really Get

These are the current US web subscription prices. Your checkout shows local pricing and applicable tax.

- **Free ($0):** Everyday chat, with limited uploads, images, research and desktop agent access.
- **Go ($8/month):** Higher everyday limits than Free. It can include ads.
- **Plus ($20/month):** Expanded reasoning, research and agent work. A good starting point if you use ChatGPT for work.
- **Pro ($100, $200 or $500/month):** Three levels of included usage. The $500 plan includes Astra Ultrafast.
- **Business (Standard seats $25/user/month, or $20 with annual billing; Premium seats $125, or $100 annual):** A shared workspace, admin controls and business data excluded from model training by default. Premium seats get about 5x the usage plus dots. Two users minimum.
- **Enterprise and Edu:** Organization plans with additional controls and their own agreements.

The Pro change to know about

Pro $200 is open to new subscribers again. New subscriptions receive a lower included allowance than the previous version. Eligible existing subscribers retain the earlier allowance through **October 29, 2026**, then move to the new allowance at the same price. Pro $200 does not include Ultrafast.

Work and Codex share usage. A long task with a powerful model can use much more than a short chat. Check Settings > Usage or Usage & Billing for your account's remaining allowance and available credits.

Which plan should you choose?

Try Free first. Choose Plus when you regularly use research, files and Work. Move to Pro when the actual limits interrupt your day. Pick the tier for the work you do, then check the current plan comparison before upgrading.

The subscription and API are billed separately. Buying ChatGPT Pro does not give you a free API balance.

Ads, and How to Kill Them for Free

Yes, ChatGPT has ads now. Testing started in the US on February 9, 2026, and ads are rolling out to Free and Go users in select regions, including Canada, the UK, Australia and New Zealand. They appear below the end of a response, labeled as sponsored.

Free and Go see ads. Plus, Pro, Business, Enterprise and Edu are completely ad-free. Ads are never shown to under-18 accounts 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. 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. Where ads are running, OpenAI offers a free no-ads option with extra limits. Look for it in Settings > Ad Controls. The trade: fewer messages per day and no access to tools like images. Availability varies by region.

Settings, Ad Controls also lets you turn off ad personalization, dismiss ads and share feedback, see why you got a specific ad, and clear your ads data. If personalization is on, ads can draw on your past chats and memory. If that bothers you, that toggle is the one to find. Temporary Chats never show ads.

Navigating the Interface

The sidebar

Use the sidebar for your chats, Projects, Plugins and Scheduled tasks. Library holds your files, and on Pro, Business and Enterprise it becomes Space, with Pages and Slides. On desktop, the product and mode controls let you move between Chat, Work and Codex. The items you see depend on your account and the selected view.

The model picker

In regular Chat, open the model menu to pick GPT-6 and move the Thinking slider. In Work and Codex, look beneath the composer for the model and reasoning control. Choose the model for the task, then raise effort only when the job needs it.

The composer

- **+ or 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.** Choose the available Search tool or ask explicitly for a web search with sources.
- **Image.** Native image generation and editing.
- **Mic or voice control.** Dictate a prompt or start a live conversation. Camera and screen sharing depend on the voice experience and platform.
- **Writing and code blocks.** Long drafts and code open as editable blocks inside the chat, so you change one paragraph or one function without regenerating the whole thing.
- **Tools menu.** Deep research, study mode, record, and (on business plans) company knowledge.

Keyboard shortcuts

In the current desktop app:

- Enter sends. Shift + Enter adds a line.
- Cmd/Ctrl + N starts a chat.
- Cmd/Ctrl + B toggles the sidebar.

- Cmd/Ctrl + F finds text in the current chat.
- Cmd/Ctrl + / shows the shortcuts for your platform.
- On Mac, pressing both Command keys takes an Appshot where enabled.

Check Settings > Keyboard Shortcuts for your app's current assignments. Older ChatGPT Classic shortcuts and browser shortcuts can differ.

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.

ChatGPT Work

This is the headline feature of ChatGPT in 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. On the biggest jobs it can now run on GPT-6 Astra, the model built exactly for this.

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

1. Open the Work tab in the desktop app sidebar, or on web and mobile if you are Plus and above.
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.
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.
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

PROMPT 2 · COPY & PASTE

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.

PROMPT 3 · COPY & PASTE

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.

PROMPT 4 · COPY & PASTE

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.

PROMPT 5 · COPY & PASTE

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.

PROMPT 6 · COPY & PASTE

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 and limits

Plus and Pro include Work on desktop, web and mobile. Free and Go have limited desktop access. Business, Enterprise and Edu access also follows workspace settings. Your available models and tools depend on the account, client and rollout.

Work and Codex share usage. Start ambitious tasks with a clear goal, relevant files and a sensible model. You can review the direction before a large job consumes your allowance.

Continue across devices

New synced Work tasks can continue between desktop, web and mobile where available. Cloud work uses uploads and connected apps. Steps that need your local files or apps require your connected computer to be online and available. Older tasks keep the mode they started with.

What companies are actually doing with it

- Zapier uses Work for lead triage across HubSpot, Gong and email. A lead used to take 35 to 45 minutes manually. Now it feeds a weekly exec dashboard.
- NVIDIA uses it for GTC event prep, replacing a spreadsheet process that used to eat around 40% of the team's time. It now runs twice weekly.
- Virgin Atlantic used it for competitor benchmarking for a 5-year plan. Weeks of work compressed into hours.
- RingCentral uses it for launch-readiness checks across Jira and Salesforce. Scaled from supporting one 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.

The Desktop App

The desktop app is one application for Chat, Work and Codex. If the app on your computer is called ChatGPT Classic, download the current app from chatgpt.com/download instead.

The desktop app is available on every plan, including Free, on macOS and Windows, with Linux in public preview. And it is meaningfully more capable than the website. For two years the web app was the full product and desktop was a convenience wrapper. That is over.

What the desktop app adds

- **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 macOS. On Windows it takes over the foreground during a task.
- **Local file access.** Point it at a folder. "Analyze every file in here." Synced cloud tasks can use a connected computer for local steps when that computer is online and access is enabled.
- **Appshots.** Capture an app as context. On macOS, pressing both Command keys captures the front window where Appshots is enabled.
- **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.
- **Work with apps.** ChatGPT can read from your open apps and, in supported IDEs, edit code directly: it shows you a diff, you review, you apply.

Set the desktop shortcuts you use most in Settings > Keyboard Shortcuts. A shortcut that opens a chat or starts dictation makes it much easier to use ChatGPT in the middle of your normal work.

Remote: your phone as a remote control

Pair the mobile app with a Mac or Windows machine running the desktop app (QR code pairing). 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.

The Built-In Browser and Computer Use

ChatGPT comes with a browser built directly into the desktop app, plus a browser extension that runs in the sidebar of Chrome, Edge, Brave, Opera or Vivaldi.

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 operate inside the login-required SaaS tools that have no API and no plugin. All the tedious web-based admin work that lives in tools with no integration becomes 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. 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, and it got about twice as fast with the GPT-6 Astra release.

PROMPT 7 · COPY & PASTE

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.

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.

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.

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.

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 makes the current version strong:

- **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.

PROMPT 8 · COPY & PASTE

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.

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.

To start one, type /Deepresearch in your prompt, pick Deep research from the + menu, or open it from the sidebar. You can also pull research into a ChatGPT Work task.

Free and Go get limited deep research. Plus gets expanded. Pro gets maximum. OpenAI does not publish 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.

Voice: GPT-Live

ChatGPT Voice lets you talk through a task, interrupt, change direction and check progress. In **Settings > Voice** you can choose **Live** (the newest, powered by GPT-Live-1.1, and it can use web search, memory and plugins), **Advanced** (use this one for video or screen sharing, which Live doesn't do) or **Standard**. In the desktop app it

works with Chat, Work and Codex on supported paid plans. Free gets limited Live access; Go and Plus get 3 hours a day.

Open a task and choose **Start voice chat**, or start a new voice chat. Allow microphone access and pick a voice. If the option is missing from an existing task, update the app and check your plan and rollout.

On macOS, **Settings > Voice > Screen context** lets you show the current app during the conversation. Through Remote on iOS, you can talk to work running on a paired desktop. Voice dictation is separate and simply turns speech into a typed prompt.

Use the app's current limits and available voice options for your account. Desktop voice, mobile voice and camera sharing do not all have the same controls.

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:

PROMPT 9 · COPY & PASTE

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.

Image Generation

ChatGPT image generation creates and edits images from a prompt, now on ChatGPT Images 2.5. Open **More > Images** to start from a template or a sketch, or just ask in chat. Attach the image you want to change and say exactly what to preserve.

Open a generated image to review it. On mobile you can open it full-screen to edit it or leave comments on what to change. For a small edit, describe the region and the change in one clear sentence.

What matters:

- **Reference images.** Show the style, subject or layout you want.
- **Specific edits.** Name what changes and what stays fixed.
- **Purpose and size.** Explain where the image will appear and the shape you need.
- **Review.** Check faces, lettering, products and factual details before using it.

PROMPT 10 · COPY & PASTE

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.

PROMPT 11 · COPY & PASTE

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.

PROMPT 12 · COPY & PASTE

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.

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.

Sites

Long-form writing and code do not need a separate editor. They open as editable blocks inside the chat: highlight a paragraph, ask for a change, only that paragraph changes, and the version history is right there. Sites goes one step further.

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. Sites is in public beta on Plus, Pro, Business, Enterprise and Edu (not Free or Go, and public publishing isn't in the EEA, UK or Switzerland yet). There's no longer a cap on how many Sites you can make; usage and storage limits still apply.

PROMPT 13 · COPY & PASTE

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.

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.

Codex

Codex is the autonomous coding agent. It writes, tests, debugs and ships code, handles multi-file and multi-repository projects, and keeps running until the task is done. It can run on GPT-6 Astra now, which was built with exactly this kind of long-horizon work in mind.

One number worth knowing: OpenAI says over ten million people use Codex and ChatGPT Work weekly, and about one in five Codex users is not a developer. That second stat 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:

- 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.

PROMPT 14 · COPY & PASTE

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.

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.

Plugins

The naming here has changed twice, which is why everyone is confused. Connectors became Apps in December 2025. In July 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 well over a thousand of them, at chatgpt.com/plugins, and the library keeps growing: recent additions include Apple Messages on Mac, restaurant reservations through OpenTable, Resy and Yelp, and Zendesk.

What's connectable:

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

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.

If ChatGPT cannot see a file, check the connected account and the plugin permissions. Then name the app and the file in your request. Installing a plugin does not automatically grant access to every account or document you own.

Interactive apps inside the chat

Some plugins 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 @. Availability varies by region, and some apps are not offered in the EEA, Switzerland or the UK.

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.

Skills

Skills are on every plan, and they are the most underrated feature in ChatGPT right now. Find them in the Plugin Directory's Skills tab.

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.

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.

PROMPT 15 · COPY & PASTE

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:

PROMPT 16 · COPY & PASTE

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

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.

Still using a custom GPT? Custom GPTs retire on **December 11, 2026**. Open **My GPTs** and choose **Migrate to plugin** for the ones you want to keep (custom actions don't carry over), or rebuild them as a skill or a project.

Memory and Personalization

Memory was rebuilt in mid-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.

You can 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 find out why and fix it.

Settings that matter:

- **Reference saved memories: on.** This is the core of it.
- **Reference chat history: on.** Lets it draw on past conversations, not just explicit memories.
- **Memory sources.** Tap the book icon under any answer to see and correct what it drew on.
- **Temporary Chat.** For anything you do not want remembered.

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. The version of that I actually run is a folder of plain text files that works in ChatGPT or Claude, in my file system setup guide. That is the asset that compounds.

PROMPT 17 · COPY & PASTE

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]

Library and Space

Library keeps uploaded and generated files available beyond a single conversation. On Pro, Business and Enterprise, Library becomes **Space** (chatgpt.com/space or the sidebar), which brings files, editable Pages and, from October 16, Slides together. Use a Page when you want to keep revising a document and work with other people on it, and Slides when you need a deck your team can edit together. Space works on web and desktop and is read-only on mobile.

Start with one project folder and a few useful files. Ask ChatGPT to create a Page from them, review the result, then use the sharing controls if you want someone else to view, comment or edit. Availability follows your account and workspace rollout.

Your Dot

A dot is an always-on agent that keeps working toward an ongoing goal. It runs in the cloud with its own computer and browser, powered by GPT-6 Astra, so cloud work can continue while your laptop is closed. Dots are rolling out to Pro and Business Premium users aged 18 and up (Pro access excludes the EEA, UK and Switzerland). My five features guide walks through setting one up: mavgpt.ai/resources/five-chatgpt-features-setup-guide-2026

Give it a real responsibility, like keeping a project plan current or preparing a weekly research brief. Explain what it can do independently and what needs your approval. Keep the first job specific so you can judge the results before expanding it.

Give your dot its first job

Create your dot in the ChatGPT mobile app, desktop app or a desktop browser (mobile web isn't supported). Give it one responsibility, provide the relevant background and connect the sources it needs. Explain when it should bring you a result and which decisions need your approval.

Try an ongoing job you already understand, so you can judge whether it is helping:

PROMPT 18 · COPY & PASTE

Help me keep one ongoing project moving: [project].

Here is the outcome I want: [outcome]. Here are the notes, files and sources you can use: [sources]. My deadline and constraints are [details].

Start by reading that material and showing me the three most useful next actions. Ask for any missing information that would change your plan.

Keep a short project page with the current goal, decisions, open questions and next actions. When new information arrives, update that page and tell me what changed. Bring me decisions I need to make, with a recommendation and the reason for it.

You can research and draft within those sources. Ask before sending messages, publishing, spending money or changing another account. Tell me what access or schedule you need before setting it up. Do not claim ongoing monitoring is active until it is configured.

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, 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:

1. Create the project and name it after an outcome. Not "Marketing." Try "Q3 launch, target 5,000 signups."
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.
3. Write project instructions. This is the step everyone skips, and it is the one that makes it work.

PROMPT 19 · COPY & PASTE

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.

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.

Scheduled Tasks

Scheduled tasks are how ChatGPT stops being something you visit and starts being something that shows up. Create and manage them in Scheduled. Free gets 3 active tasks, and paid plans get more. My scheduled tasks guide has the full setup.

Four kinds of schedule:

- **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 (Plus and Pro):** "When a new message lands in #launches, update the launch doc." Triggers can come from Gmail, Slack or GitHub.
- **Monitor:** "Watch this page. Tell me if the pricing changes."

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

PROMPT 20 · COPY & PASTE

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.

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

Web scheduled tasks can use uploaded context, connected tools, skills and plugins. Local scheduled tasks need the computer on and the app running. Manage them under Scheduled, test the prompt first, and check its first runs. Eligible web and mobile accounts can also use supported app events as triggers.

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 newer is that it plugs directly into the tools you already use.

ChatGPT has extensions for Word, 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.

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.

PROMPT 21 · COPY & PASTE

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.

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.

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.

ChatGPT for Business and Enterprise

Business Standard seats start at $20 per user per month billed annually, or $25 monthly, with two users minimum. Premium seats ($100 annual, $125 monthly) get about 5x the usage and dots. Enterprise uses a separate agreement. Business, Enterprise and Edu data is excluded from training by default. Check your organization's current seat and usage options before choosing a plan.

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.

Use the Company knowledge option where your workspace provides it. In Work, installed plugins can also bring connected company sources into a task. A missing option can reflect your surface or workspace permissions, so check the available tools before changing plans.

Workspace agents

Workspace agents are shared, Codex-powered agents. OpenAI describes them as the evolution of Custom GPTs, and with Custom GPTs retiring on December 11, 2026, they're where that kind of team tool lives now. 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 vs Enterprise: what you do and do not get

Both plans include SAML SSO, admin console, SOC 2 Type 2, company knowledge, workspace agents, and no training on your data.

Enterprise adds SCIM provisioning, IP allowlisting, an analytics dashboard, compliance API logs, data residency across ten regions (US, EU, UK, JP, CA, KR, SG, IN, AU, UAE), ISO certifications, and enterprise key management.

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 of the Enterprise-only rows above, you need Enterprise.

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.

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.

Settings, Privacy and Safety

The settings worth changing today

- **Custom instructions** (Personalization). The highest-leverage five minutes in the product.
- **Connected apps and permissions.** Check that the right account is connected and that the plugin can access the source you want to use. Name that source in your request.
- **Model and reasoning effort.** Start with the default effort, then raise it for tasks that need more careful work. Watch your usage when running long tasks.

- **Memory sources** (Memory). See and correct what it's drawing on.
- **Ads controls.** Where ads are available, manage personalization and feedback here. Free accounts in the test can also choose fewer messages in exchange for no ads.
- **Data controls** (Data controls). Opt out of training on your content.
- **Temporary Chat** (Composer). No memory, no history.

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 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.

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.

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. Today's models can think. Stop doing their thinking for them.

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

The old way (stop doing this):

PROMPT 22 · COPY & PASTE

```
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):

PROMPT 23 · COPY & PASTE

```
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:

1. **State the outcome, not the method.** The model plans better than your scaffold does.

2. **Plan in one chat, execute in another.** Long chats degrade. Get the plan, then open a fresh chat and say "implement step one."
3. **Ask for its uncertainty.** "Tell me what would change your answer" surfaces the weak link.
4. **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.
5. **Don't max the effort by default.** High reasoning on a simple task produces a worse answer, slower.
6. **Kill the chat when it gets dumb.** It's not broken. The context window is full of your own conversation. Start fresh.

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.

8 Prompt Frameworks

The difference between a mediocre AI response and a brilliant one usually comes down to how you structure the prompt. These eight frameworks give you a repeatable system.

1. R-T-F (Role, Task, Format)

The simplest framework. Specify the role, task, and output format.

PROMPT 24 · COPY & PASTE

Act as a brand strategist. Write a messaging hierarchy for a B2B SaaS founder targeting finance teams. Use bullet points with core message, value props, proof, and CTAs.

2. S-O-L-V-E (Situation, Objective, Limitations, Vision, Execution)

Great for strategic planning. Define where you are, what you want, what constraints exist, where you want to end up, and how to get there.

PROMPT 25 · COPY & PASTE

Situation: At a startup that just launched a new B2B product.
Objective: Generate qualified inbound leads through educational content on LinkedIn.
Limitations: No ad spend. Small team. Must publish 3x/week.
Vision: Become the go-to thought leader for mid-market SaaS operations leaders.
Execution: Review performance biweekly and double down on top-performing content.

3. T-A-G (Task, Action, Goal)

Cut to the chase. Define the task, state the action, clarify the goal.

PROMPT 26 · COPY & PASTE

Task: Reduce customer churn in a SaaS subscription business.
Action: Analyse churn data and launch a targeted retention programme.
Goal: Improve customer retention rate by 15% over the next 6 months.

4. R-A-C-E (Role, Action, Context, Expectation)

Ideal for complex business prompts.

PROMPT 27 · COPY & PASTE

Role: Senior marketing strategist.
Action: Develop a content marketing plan.
Context: B2B SaaS company targeting mid-market operations teams.
Expectation: Tailored to buyer roles (Ops Lead, CTO, Procurement).

5. D-R-E-A-M (Define, Research, Execute, Analyse, Measure)

A full project lifecycle in one prompt.

PROMPT 28 · COPY & PASTE

Define: Declining retention in a B2B SaaS product.
Research: Through user interviews and CRM data.
Execute: Launch a 3-month retention programme with onboarding improvements.
Analyse: Compare behaviour between control and test cohorts.
Measure: Track churn rate, product usage frequency, and NPS.

6. P-A-C-T (Problem, Approach, Compromise, Test)

Perfect for decision-making with trade-offs.

7. C-A-R-E (Context, Action, Result, Example)

Provide context, describe needs, clarify results, give examples.

PROMPT 29 · COPY & PASTE

Context: Engagement with its onboarding flow.
Action: Redesign the onboarding journey to focus on quick wins and key features.
Result: Higher product adoption within the first 7 days.
Example: Show how the revised onboarding helped increase day-7 activation from 25% to 42%.

8. R-I-S-E (Role, Input, Steps, Expectation)

Great for data-driven tasks.

PROMPT 30 · COPY & PASTE

Role: Commercial director.
Input: Quarterly sales performance data and pipeline forecasts.
Steps: Identify weak points, prioritise high-conversion leads, and optimise resource allocation.
Expectation: Deliver a plan to exceed revenue targets in the next quarter.

Quick Decision Guide

> **Simple tasks** → R-T-F or T-A-G. **Strategic work** → S-O-L-V-E or D-R-E-A-M. **Business prompts** → R-A-C-E.

Getting a Finished File From One Prompt

Give ChatGPT the outcome, the audience, the format and the files it needs. Upload the spreadsheet, document or reference material, then describe the finished deliverable you want. Tell it the constraints, let it use the relevant tools, and save prompts that work so you can reuse them.

For a large job, ask for a plan first. Then have it carry out the work, check the result and give you the finished file. Browsing, analysis and file creation can be parts of the same task.

Two More Framework Examples

P-A-C-T: Choose a Pricing Model

PROMPT 31 · COPY & PASTE

Problem: Deciding between per-seat pricing and usage-based pricing for our SaaS product. Approach: Evaluate both models against our 200 accounts. Compromise: We accept some revenue unpredictability if expansion revenue is significantly higher. Test: Run a 60-day pilot with 20 new customers on the alternative model.

C-A-R-E: Improve Onboarding

PROMPT 32 · COPY & PASTE

Context: A SaaS company struggling with low engagement in its onboarding flow. Action: Redesign the onboarding journey to focus on quick wins and key features. Result: Higher product adoption within the first 7 days.
Example: Show how the revised onboarding helped increase day-7 activation from 25% to 42%.

6 Prompts From the Original Model Guide

Analyze Sales Data

PROMPT 33 · COPY & PASTE

Here's my sales data for Q1. Analyze the trends, identify the top 5 performing products, find any concerning patterns, and build me a presentation with charts and 3 specific recommendations for Q2. Keep the presentation under 10 slides.

Build a Business Plan

PROMPT 34 · COPY & PASTE

I run a DTC skincare brand doing $50K/month. Growth has stalled for the last 3 months. Give me a 90-day plan to get back to 20% month-over-month growth. Focus on acquisition channels, retention tactics, and 3 quick wins I can implement this week.

Research AI Workflows

PROMPT 35 · COPY & PASTE

Research how AI agents are changing developer workflows in 2026. Separate your findings into: what's actually working, what's overhyped, and what this means for junior vs senior developers. Cite your sources.

Debug a Large CSV Script

PROMPT 36 · COPY & PASTE

Here's my Python script. It's supposed to process CSV files and output a cleaned version, but it keeps crashing on files larger than 1GB. Find the issue, fix it, optimize it for large files, and add error handling. Then write tests to verify the fix.

Write a Detailed Article

PROMPT 37 · COPY & PASTE

Write a 2,000-word blog post about the future of remote work for a tech-savvy audience. Avoid generic advice. Use specific examples, data points, and contrarian takes. Conversational tone, not corporate. Include a strong hook in the first paragraph.

Review a Vendor Contract

PROMPT 38 · COPY & PASTE

I'm uploading a 50-page vendor contract. Find every clause related to termination, liability, auto-renewal, and price escalation. Summarize each one in plain English and flag anything I should negotiate before signing. Organize by risk level: high, medium, low.

25 Power Tips

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

1. **Connect the right account.** Install the plugin, connect the service and name the source in your request. Check permissions if a file is missing.
2. **Start with the default effort.** Give it a clear task and raise effort when the first result shows the problem needs more thought.
3. **Stop maxing out the reasoning level.** On a simple task, high effort makes the answer worse, not better. It overthinks and invents problems.
4. **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.
5. **Set a useful shortcut.** Open Settings > Keyboard Shortcuts and choose a shortcut for the action you use most.
6. **Use desktop for work on your computer.** Local files, Computer Use and Appshots fit naturally there. Use synced Work when you want to continue across devices.
7. **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.
8. **Run the automation audit.** Same trick for scheduled tasks. Ask it what you keep asking it, then schedule those.
9. **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.** "Analyze every file in this folder." Use local Work or an enabled, online computer connection.
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 see exactly which chats and files it drew on, and mark them not relevant.
21. **Free plan users, get ad-free.** Settings, Ad Controls. Trades message limits and some tools for zero ads, where it's offered.
22. **Use Temporary Chat for anything sensitive.** No memory, no history.
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.

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, and GPT-6 Astra just widened that lead.
- 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.
- Naming and menus. Features get renamed and moved often, and things you rely on can get cut.

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: Add custom instructions and connect one app you use every day.
- Day 2: Install the desktop app and set your favorite shortcut.
- 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.
- Days 6 and 7: Seed your memory properly. That is the part that compounds.

Updated October 16, 2026. Models, plans and menus change, so use your account's current controls when they differ. The habits are what carry forward: give it context, ask for a useful result, and review what it makes.

Download the full PDF below to keep it on hand as ChatGPT keeps changing.

MORE FREE GUIDES

100+ free AI guides, prompts and systems at mavgpt.ai/resources

New guides every week. Newsletter every Wednesday and Sunday at mavgpt.ai/newsletter.