

CORTEXA ACADEMY

Getting Started with **ChatGPT**

A complete, step-by-step guide to setting up and mastering ChatGPT — from your first chat to agents and coding tools that finish real work.

1 ChatGPT · Web

2 ChatGPT Work

3 Codex

How to use this guide

This guide takes you from zero to confident across the three ways people actually use ChatGPT. You do not need any technical background. Work through it in order — each part builds on the one before, and the skills you learn in Part 1 carry all the way through.

THE SYMBOLS YOU'LL SEE

- 1** **Numbered steps**
Do these in order, exactly as written.
- >_** **Prompt boxes**
Dark boxes are word-for-word text you can type or paste into ChatGPT.
- ★** **Callouts**
Tips, warnings, and things worth knowing before you click.
- ①** **Screen sketches**
Simplified drawings of what you'll see on screen — not live screenshots.

i A note on the pictures

ChatGPT's interface changes often — button names and colors may look a little different, but the steps stay the same. When in doubt, look for the label we name in the step.

THE THREE-PART JOURNEY

- **ChatGPT on the web** — your daily driver. Chat, files, Projects, memory, research, and personalization in the browser.
- **ChatGPT Work** — workspace agents that carry team knowledge and complete multi-step jobs across your real tools.
- **Codex** — the most powerful surface. Automate work and build things — approachable even if you've never touched a terminal.

WOVEN THROUGHOUT

A running thread on **verification** and **prompt craft** — the two habits that separate good results from great ones. Both get their own sections near the end.

What's inside

01 · Getting oriented — the three surfaces, plans, choosing a model	p. 4
02 · Part 1 — ChatGPT on the web — account, security, first chats, files & Canvas	p. 8
Projects, personalization & power tools — research, data, images, voice	p. 16
03 · Part 2 — ChatGPT Work — workspace agents for real, multi-step work	p. 24
04 · Part 3 — Codex — automation & building, made approachable	p. 32
05 · Prompt engineering — the habits that transform your results	p. 39
06 · Top apps to try — the apps directory & how to connect it	p. 43
07 · Pitfalls & ethical use — hallucinations, privacy, and honest AI	p. 46
08 · Quick reference & next steps — cheat sheet and the 4-week rollout	p. 49

THE ONE-MINUTE VERSION

- Subscribe to **Plus (\$20/mo)**. It unlocks everything in Part 1 — and Codex, at no extra cost.
- Learn the **web app** first. Then add **ChatGPT Work** agents for team jobs, then **Codex** for automation.
- **Verify anything that matters.** ChatGPT is a brilliant colleague, not a source of record.

① Who this is for

Anyone who wants to go from casual questions to using ChatGPT as a genuine work partner — marketers, researchers, analysts, and the teams they work with.

01

Getting oriented

The three surfaces, plans, and choosing a model — the quick lay of the land before you dive in.

IN THIS SECTION

- The three ways to use ChatGPT
- Which plan you need
- Which model to pick

The three ways to use ChatGPT

ChatGPT shows up in three places. They share one account and one brain, but each is built for a different kind of work. You'll learn them in this order — and everything transfers forward.

1 ChatGPT · Web

Type a question, get an answer — plus files, Projects, memory, and research, all in your browser.

BEST FOR

Everyday thinking, writing, analysis, research, and learning. The place you'll spend most of your time.

WHAT IT FEELS LIKE

Like texting the sharpest colleague you know. Nothing to install.

2 ChatGPT Work

Workspace agents that carry your team's knowledge and complete multi-step jobs across your real tools.

BEST FOR

Repeatable team workflows: the Monday pipeline report, meeting prep, research briefs — run end-to-end.

WHAT IT FEELS LIKE

Like handing a project to a capable colleague who works for hours and returns finished files.

3 Codex

A coding agent that runs in your browser, editor, or terminal — automate tasks and build things.

BEST FOR

Repeatable automations, working with lots of files at once, and anything that touches code.

WHAT IT FEELS LIKE

Like having an engineer on call. More power, a small learning curve.

★ The golden rule

Start in the web app and get comfortable. Move to ChatGPT Work when a whole workflow should run without you. Reach for Codex when you want to automate something you do over and over. You do not need all three on day one.

Which plan do you need?

Six tiers now. The short version: Plus at \$20/month is the sweet spot for individuals — every step in Parts 1 and 3 works on it. Part 2's workspace agents live on the team plans.

Plan	Price	What you get	Who it's for
Free	\$0	Current GPT-5–family models with capped usage; small daily reasoning budget; limited images & uploads	Trying it out for light, occasional use
Go	~\$8/mo (US)	Higher limits at a low price — ad-supported	Budget users who accept ads
Plus RECOMMENDED	\$20/mo	Full model picker (the GPT-5.6 family incl.), ~8× usage, Deep Research, Agent mode, Advanced Voice, Projects, Tasks, custom GPTs — and Codex	Almost everyone — the individual sweet spot
Pro	\$100–200/mo	Near-unlimited usage in two tiers, plus Pro-grade research reasoning and max Codex headroom	Power users running research, agents & Codex all day
Business	~\$20–25/seat	Plus for everyone + shared workspace, admin, workspace agents (ChatGPT Work) — and no training on your data, by contract	Teams of 2–200
Enterprise	Custom	SSO/SCIM, compliance logging, residency, agent governance — ChatGPT Work at full scale	Large orgs with IT & compliance needs

Start with Plus

Every step in Parts 1 and 3 works on Plus. Upgrade to Pro only if you repeatedly hit limits; move to Business when a team needs shared agents or the data guarantee.

About “usage limits”

Caps are rolling windows (messages per 3–5 hours + weekly quotas on reasoning/research), not fixed monthly counts — and OpenAI adjusts them with load.

Which model should you pick? (Short answer: let **Auto** decide)

ChatGPT routes every message to one of a few model modes — think of it as choosing horsepower. The picker sits at the top of every chat.

FASTEST · DEFAULT

Instant

The fast everyday workhorse. Drafting, summaries, everyday questions — the large majority of tasks.

REASONING

Thinking

Works step-by-step before answering. Math, logic, tricky analysis, careful writing. Weekly quota on paid plans.

PRO PLAN ONLY

Pro mode

Research-grade reasoning for the hardest problems and the deliverables that really matter.

THE GPT-5.6 FAMILY — LIVE NOW

OpenAI's newest family, **GPT-5.6** (launched July 2026), is in the picker on paid plans in three sizes: **Sol** — the deepest, what Pro mode runs for the hardest work; **Terra** — balanced, the everyday flagship; **Luna** — the fastest and lightest.

★ The beginner rule of thumb

Leave it on **Auto** — it reads your message and picks Instant or Thinking for you. Bonus: Auto-escalations to Thinking **don't** count against your weekly Thinking quota; manual picks do.

i Out with the old

GPT-4o, GPT-4.1, o4-mini and the original GPT-5 were retired from ChatGPT in **February 2026**. If a tutorial mentions them, it's out of date. Names change; the Instant / Thinking / Pro structure is the stable part.

02

PART 1

ChatGPT on the web

Your daily driver: chat, files, Projects, memory, research, and personalization — all in the browser, nothing to install.

IN THIS PART

- Create your account & upgrade
- Lock it down & set privacy
- A tour of the screen & your first chat
- Projects, files & Canvas
- Memory & personalization
- The power tools: research, data, images, voice

Step 1 — Create your account

1 Go to the website

Open your browser and visit **chatgpt.com**. Click **Sign up**.

2 Choose how to sign in

Use **Continue with Google / Microsoft / Apple** to attach ChatGPT to an account you already protect — fastest and usually safest. Or use email + password.

3 Verify


Confirm the code OpenAI sends, then add your name and birth date when asked.

4 You're in — on the Free plan

The account works immediately. Next we'll upgrade so every feature in this guide is available.

No credit card needed to start

Free doesn't ask for payment. You'll only enter a card when you choose to upgrade on the next page.



The screenshot shows the ChatGPT website's sign-up page. At the top, the browser address bar shows 'chatgpt.com'. The page features the ChatGPT logo and the heading 'Create your account'. Below this, there are three large buttons: 'Continue with Google', 'Continue with Microsoft', and 'Continue with Apple'. Under these buttons is an input field labeled 'Email address'. At the bottom of the form, a line of text reads: 'By continuing you agree to the Terms & Privacy Policy.'

SSO buttons = fastest option, and one less password to protect


★ Setting this up for work? Sign in with your **work identity** (Google Workspace / Microsoft 365) so the account can move into a Business workspace cleanly later.

Step 2 — Upgrade to Plus

- 1 Open your account menu**
Click your profile picture (bottom-left on web). A small menu pops up.
- 2 Click “Upgrade plan”**
Choose **Upgrade plan** from the menu (or go straight to chatgpt.com/#pricing).
- 3 Pick Plus**
Select **Plus — \$20/month**. Pro exists for heavy users; upgrade later only on evidence.
- 4 Enter payment & confirm**
Plus activates immediately — you'll see the full model picker appear.

★ What just unlocked

The complete GPT-5.6 model picker, ~8x usage, Deep Research, Agent mode, Advanced Voice, Projects, Tasks, custom GPTs — and **Codex**, which we cover in Part 3. One subscription, all of it.



Two clicks from the profile menu

i Prefer to stay on Free for now?

You can — Part 1's basics mostly work on Free. But the power tools, higher limits, and Codex need Plus, so upgrade before the second half of this part.

Step 3 — Lock the account down

1

Open Security settings

Profile → **Settings** → **Security**.

2

Turn on multi-factor authentication

Use an **authenticator app**, not SMS, if given the choice. This one toggle defeats most account-takeover attempts.

3


Add a passkey (recommended)

Sign-in with your device's fingerprint / face / PIN — phishing-resistant and faster than a password.



Why bother?

Your ChatGPT history becomes a diary of your work: drafts, plans, uploaded files, memory. An attacker with your password reads all of it. Treat this account like your inbox.



Two minutes here protects everything you'll ever type into this account



Going further

OpenAI's **Advanced Account Security** (2026) replaces passwords entirely with passkeys / hardware keys, disables SMS & email recovery, and auto-opts the account out of model training. Consider it for sensitive work.

Step 4 — Decide what OpenAI may do with your data

1 Open Data Controls

Settings → **Data Controls**.

2 Decide on “Improve the model for everyone”


ON = your conversations may train future models (the default on personal plans). OFF = they don't. Turning it off removes **no** features.

3 Know the retention rules

Deleted chats are queued for permanent deletion within ~30 days; even opted-out chats may be held up to ~30 days for safety review.

★ The plain rule

Decide **once, now**, before the account fills with work. If your chats will ever contain client or confidential material on a personal plan, turn training off today — the opt-out is not retroactive.



The toggle most people never find — and the one your clients will ask about

Account type	Training default
Free / Go / Plus / Pro	ON — unless you opt out
Business / Enterprise / API	OFF — contractually

Memory & Temporary Chats

ChatGPT remembers — and it's actually two systems wearing one name. Used well, it's the feature that makes it feel like a colleague. You stay in control.

HOW MEMORY WORKS

- **Saved memories** — facts it explicitly records (“works in agency marketing; prefers concise answers”). View, edit, delete: Settings → Personalization → **Manage memories**.
- **Chat-history reference** — it may also draw on past conversations for context. Separate toggle — people miss it.
- **Ask it directly:** “What do you remember about me?” — or “Forget that.”

★ Good habit

Keep truly sensitive data — ID numbers, health details, credentials, unreleased client material — out of memory-enabled chats entirely. Audit your memory list monthly; it takes 90 seconds.

TEMPORARY CHATS

- Toggle at the top of a new chat: **no history entry, no memory reads or writes, never used for training**. A copy may be held ~30 days for safety, then removed.
- Deleting a normal chat afterward is cleanup; a Temporary Chat was **never stored in your history** in the first place.
- For sensitive topics, make the *first* click the toggle — don't rely on the delete button later.

 The goal isn't memory *off* — it's memory **curated**. The productivity gain is real: by week three it knows your role, your projects, your preferences.

A tour of the screen

Thirty seconds of geography so nothing on this screen is mysterious.

1 New chat

A clean slate. Habit to steal: one chat, one topic.

2 History & search

Every past conversation, searchable. Hover for rename / archive / delete.

3 Projects, GPTs & Agents


Reusable workspaces, purpose-built assistants, and (on team plans) your workspace agents.

4 The composer

+ attaches files; the tools menu holds Search, Deep Research, Create image, Agent mode; the mic starts Voice.

5 Model & mode

The picker (leave it on Auto) and the Temporary Chat toggle, top-right.



The five things you'll use most, labeled

★ Just getting started?

You truly need two things: the message box and Enter. Everything else you'll add as you go.

Your first chat

Click the message box, type, press Enter. Three starter prompts — type any of them word-for-word.

>_ TRY THIS FIRST

[COPY / PASTE](#)

I'm new to ChatGPT. In plain language, explain three ways you could help me with my work this week. Ask me one question about my job first so your answer is specific to me.

>_ LEARN SOMETHING FAST

[COPY / PASTE](#)

Explain **[topic you're curious about]** as if I'm smart but completely new to it. Use a simple analogy, then give me the three things most people get wrong.

>_ TURN A MESS INTO A PLAN

[COPY / PASTE](#)

Here's a rough brain-dump of everything on my plate: **[paste your notes]**. Turn it into a short summary, a prioritized to-do list, and three risks I might be missing.

THREE TINY HABITS THAT HELP IMMEDIATELY

- **Give context.** Say who you are and who the work is for.
- **Say the format.** “In five bullets,” “as a table,” “one paragraph.”
- **Keep going.** Reply “shorter,” “more formal,” “add an example” — don't restart.

- ① **Enter sends, Shift+Enter makes a new line**
Hold Shift while pressing Enter for multi-line prompts without sending.

The anatomy of a great prompt

Five ingredients. You don't need all of them every time — but the more you include, the sharper the output. We go deeper in the prompt-engineering section.

R**Role — tell ChatGPT who to be.**

“You are a senior market researcher...” Sets tone and expertise.

C**Context — situation, audience, goal.**

“This is for a client pitch to a tech brand...”

T**Task — exactly what you want done.**

“Draft a one-page competitive summary...” Specific and direct.

E**Examples — show the style you want.**

“Match the tone of this: [paste]” — the single biggest lever.

F**Format — name the shape.**

“Five bullets, then a one-line recommendation.”

>_ ALL FIVE, COMBINED[COPY / PASTE](#)

You are a senior market researcher. I'm preparing a client pitch to a large PC brand and need to look sharp. Draft a one-page competitive summary of the three leading laptop makers. Match the crisp, factual tone of this example: **[paste]**. Format: a short intro, a comparison table, then three bullets on what it means for us.

**An easy shortcut**

Can't be bothered to structure it? Add: **“Ask me any questions you need before you start.”** ChatGPT fills the gaps itself.

Files in, Canvas out

PUTTING FILES IN

Click the **+** to attach PDFs, Word docs, spreadsheets, slides, and images. ChatGPT reads the contents — summarize, analyze, rewrite, extract, compare. Vision is built in: photograph a whiteboard, a form, a chart — it reads those too.

> TRY WITH A FILE ATTACHED

[COPY / PASTE](#)

Summarize this report in five bullets for a busy executive, then give one clear recommendation. **[attach the file first]**

★ Edit by chatting — or directly

In Canvas you can type in the panel yourself, or select a paragraph and tell ChatGPT what to change — it revises that part, not the whole draft.

GETTING POLISHED WORK OUT: CANVAS

Producing a real deliverable? Say **“draft this in a canvas.”** It opens in a side panel as a living document — the chat is the conversation; the canvas is the work product.

Your chat

WELCOME GUIDE — the Canvas

“Make the intro warmer and add an FAQ” →

Canvas tools: edit directly · adjust length · reading level · suggest edits

A living draft — each change lands in place, versions kept

❶ Export it

Copy or download — Word / PDF for documents. From data-analysis chats, file outputs like .xlsx and .pptx are possible too.

Projects: your reusable workspaces

A Project is a folder that remembers. Instead of re-explaining your context every time, you set it up once — the files, the rules, the tone — and every chat inside inherits it. The single biggest time-saver in the web app.

1 Open Projects

In the left sidebar, click **Projects**, then **+ New project**.

2 Name and describe it


A clear name like “Market Research” and a one-line description of what it's for.

3 Add files

Upload what ChatGPT should always know — brand guides, past reports, personas. It reads them in every chat, no re-uploading.

4 Set instructions

Standing rules: who ChatGPT is here, your audience, tone, do's and don'ts. They apply to every conversation inside.



Set up once — every chat inside starts already knowing it

★ Project vs. plain chat

Use a Project for any recurring stream of work with the same context (a client, a research area, a product). Use a plain chat for quick one-offs.

Teach ChatGPT how you work

Fill this in once and every future chat gets better — the biggest single upgrade to your daily results.

1 Open Custom Instructions

Profile → **Settings** → **Personalization** → **Custom instructions**.

2 Fill in the two boxes


Who you are & what you work on; and how you want responses — as **instructions, not adjectives**.

3 Save

Applies to every new chat automatically, and layers with memory.
Optional: pick a base personality style and a nickname.

Write rules, not personality traits

“I’m detail-oriented” does nothing. “Double-check every number before presenting it” is an instruction it can actually follow.



Steal this outright — then make it yours

THE THREE LAYERS OF PERSONALIZATION

- **Custom instructions** — apply everywhere.
- **Project instructions** — apply inside one Project.
- **The prompt itself** — one message. Most specific wins.

Power tool — Search: current answers, with receipts

Models have a training cutoff — without search they don't know recent events. With Search on, ChatGPT browses the live web and cites sources inline.

- It usually turns search on automatically for time-sensitive questions; force it from the tools menu or by asking it to “search the web for...”.
- **Every claim that matters → click the citation.** Does the source exist? Is it credible? Does it actually say that?
- Learning instead of answering? **Study mode** turns ChatGPT into a tutor that quizzes and scaffolds — great for onboarding yourself to a new domain.

★ Treat ChatGPT-with-search as a **research assistant that drafts**, not a source of record. The citations are there so you can be the editor.

>_ A CLEAN RESEARCH PATTERN

[COPY / PASTE](#)

Search the web for **[topic]** developments in the last 90 days. Cluster what you find into themes, tell me what's accelerating and what's stalling, and cite a source for each claim.

...adoption accelerated in Q2¹²


Sources: ► press release ► analyst note ► filing

Click through before you rely on it.

Citations are the feature, not the decoration

Power tool — Deep Research: the twenty-minute analyst

A different mode of work: give it a brief, and an agent researches autonomously — browsing, reading, and cross-checking dozens of sources — then returns a long, structured, fully cited report.



>_ A BRIEF THAT GETS A GREAT REPORT

COPY / PASTE

Research the current state of **[market/technology]**. I need: key players and their positioning, pricing where public, momentum signals from the last 6 months, risks, and a “what this means for **[my context]**” section. Prioritize primary sources and recent data; flag anything you couldn't verify.

★ Spend runs wisely

Plain search = a question that needs an answer today.
Deep Research = a question that deserves a **document**.
Don't spend a run on what one search answers.

i The rhythm

Runs are quota-limited per month by plan (more on Pro) and take minutes, not seconds. Brief it, do other work, come back to a report.

⚠ What you get is a strong first draft from a fast, tireless junior analyst with occasional overconfidence. A senior human — you — walks the citations on every claim that carries weight.

Power tool — Data analysis: an analyst on call

Upload a spreadsheet or CSV and ChatGPT analyzes it by writing and running real Python code in a sandbox — not by eyeballing the numbers.

- It cleans data, joins tables, computes, charts, and explains what it found — and you can **expand and inspect the code** behind every step.
- Ask for outputs as downloadable files: a cleaned .xlsx, a chart image, a summary doc.



Verify before you present

Spot-check totals against the source file and confirm any figure that will be quoted. Sandboxed code is reliable; **interpretation still needs an owner.**

>_ FIRST PASS ON ANY DATASET

[COPY / PASTE](#)

Here's our raw data **[attach]**. Profile it: what's in each column, what's missing or suspicious, and what are the five most interesting patterns? Charts where helpful. Explain like I'm smart but not a statistician.

>_ THE HONEST SECOND PASS

[COPY / PASTE](#)

Now argue against your own analysis. What could make these patterns misleading? What would you check before presenting this to a CFO?

Power tools — Images & Advanced Voice

IMAGES — IN BOTH DIRECTIONS

- **It can see:** attach photos or screenshots — whiteboards, handwriting, charts, forms. Your camera becomes an input device.
- **It can make:** describe an image, get an image. Iterate conversationally: “**same image, but...**” — it edits rather than starting over.
- **Honest limits:** text in images is now largely reliable; exact logos and fine typography still are not.



Brand & likeness rules still apply to you

No real people's likenesses, no competitors' marks — and label AI-generated imagery wherever your client or platform requires it.

ADVANCED VOICE — THINK OUT LOUD

- ✓ Tap the mic (best on mobile) for a natural spoken conversation. A transcript lands in the chat afterward.
- ✓ Rehearse a pitch or difficult conversation — with a partner that pushes back in character.
- ✓ Brainstorm on a walk; translate live; with the camera on, it can **look at** what you're pointing at and discuss it.



The one habit that unlocks it

Interrupt freely. New users politely wait out every answer like it's a voicemail. Don't — steer in real time, like a call with a sharp colleague.

Make it repeatable: Tasks & custom GPTs

ON A SCHEDULE

Tasks

“Every Monday 8am, search for [competitor] news from the last week and summarize with citations.” It runs and notifies you — your first taste of ChatGPT working without you.

Use when: the same request has a rhythm.
Retire stale ones — active-task limits apply.

PACKAGED ASSISTANTS

Custom GPTs

Instructions + reference files + tools, wrapped under a name (“the Proposal Reviewer”) — built conversationally, no code. Institutional knowledge with an interface.

Use when: a workflow repeats often, or the team should run it too.

❗ GPTs and Projects are shareable — in Business/Enterprise workspaces they become team infrastructure: your best prompter's craft, available to everyone, including the new hire.

★ You've finished the web basics

You can now chat, attach files, build Projects, run research, personalize ChatGPT, and schedule work. Next: agents that carry whole workflows for your team.

★ The repeatability ladder

Great one-off chat → becomes a **Project** (same context every time) → becomes a **Task** (runs on schedule) → becomes a **GPT** (anyone can run it) → becomes a **workspace agent** (Part 2 — runs whole workflows for the team). Climb it deliberately.

03

PART 2

ChatGPT Work

Move from answers to action. Workspace agents carry your team's knowledge and complete multi-step jobs across your real tools — then hand back finished work.

IN THIS PART

- From answering to doing
- Set up your workspace
- Build your first agent
- How agents get work done
- Agent mode — the personal on-ramp
- Real jobs to try · Staying in control

From answering to doing

In the web app, ChatGPT tells you how to do something. In ChatGPT Work, agents do it — planning the job, working across your connected tools for hours if needed, and handing back finished decks, sheets, docs, even small web apps. Launched July 2026 for Business, Enterprise & Edu workspaces, powered by the GPT-5.6 family.

REGULAR CHAT

You ask, it answers

You get text and advice. To act on it, you do the work — open the files, make the edits, send the report. Great for thinking & drafting.

CHATGPT WORK

You describe an outcome

A **workspace agent** makes a plan, works across the apps and files your admin approved, and returns the finished artifact. Great for repeatable, multi-step team jobs.

THE THREE PIECES

- **Workspace agents** — shared, team-owned, run multi-step workflows on schedules or triggers.
- **Codex** — the engineering-side agent (Part 3).
- **Hosted outputs** — agents can ship working artifacts, not just documents.

How it's different from Codex

Work agents and Codex share the same engine family. ChatGPT Work is built for non-technical knowledge work — reports, briefs, decks, pipelines of documents. Codex (Part 3) is aimed at automation and building with code.

- ★ Individual Plus users get the single-player version — **Agent mode** — covered later in this part. Work is the team-scale, shared, administered version.

Set up your workspace

Workspace agents live on the team plans — a Business workspace is the fastest way in; Enterprise adds the full governance stack.

1 Pick the plan

Business (~\$20–25/seat): Plus for everyone, shared workspace, and **no training on your data — by contract**. **Enterprise**: adds SSO/SCIM, audit-log exports, residency, RBAC.

2 Enable workspace agents


An admin turns on Workspace Agents for the org and assigns **who may build** and **who may publish** through role-based controls.

3 Approve the tools

Admins choose which apps and connectors agents may touch — Drive, Slack, SharePoint, CRM. Agents can only use what's approved.

4 Start small

Pilot with one team, one or two low-stakes workflows, then widen. Trust grows with track record.



The banner your compliance team wants to see



Know the ceiling

Business has no SCIM provisioning, limited roles, no compliance-log exports or residency. When IT asks for those, the answer is Enterprise — and 2026's new question, "can we disable agent capabilities we haven't risk-assessed?", is a yes on both.

Build your first workspace agent

No code — the builder is a conversation. Describe the job the way you'd brief a new colleague.

1 Open the builder

Click **Agents** in the sidebar → **New agent**. The conversational builder opens.

2 Describe the job

In plain language: the task, **what success looks like**, the constraints it must follow, and the tone or format of the output.

3 Connect its tools

Pick from the admin-approved apps and connectors, and authenticate the ones it needs. Give it the **minimum** set that does the job.

4 Test, then publish

Run it on a real example, correct it like a new hire's first draft, then publish to the team. Add a **schedule** (“run every Monday 7am”) or a **trigger** (a new message, a data change).

>_ A BRIEF THAT MAKES A GOOD AGENT

[COPY / PASTE](#)


You prepare our weekly pipeline brief. Every Monday at 7am: pull the current pipeline from **[CRM]**, compare to last week, and draft a one-page summary — wins, slips, three risks — in our house style **[attached]**. Save it to the team Drive folder and post the link in #sales. If any number looks off, flag it instead of guessing.

★ The brief is the product

Success criteria + constraints + “flag, don't guess” is what separates an agent you trust from one you babysit. Encode every correction back into the agent's instructions.

How agents get work done

When an agent takes a job, it doesn't fire off one response. It works the way a capable colleague would — and it can stay on the job for hours.



It plugs into everything you set up

Agents use the workspace's approved connectors (pull from the CRM, save to Drive, post to Slack), your shared GPTs and Projects as knowledge, and schedules or triggers to run without being asked.

The economics changed

Workspace agent runs are **metered (token-based)** — **not flat-fee** — since July 2026. An agent that runs for hours is doing real, billable computation. Tell finance before the team gets enthusiastic; watch usage in the admin console.

Set it and (almost) forget it

Scheduled agents are the payoff: a Monday briefing, a weekly cleanup, a daily digest — running before you sit down. Review the output like a colleague's, every time.

Agent mode: the personal on-ramp

No workspace yet? Agent mode (tools menu, Plus and up) is the single-player version: hand ChatGPT a whole task and it operates a browser on your behalf — navigating, gathering, filling, assembling — showing its work as it goes.

GOOD FIRST JOBS

- ✓ Compare offerings across a dozen vendor sites into one table, with links.
- ✓ Gather and structure public data; multi-step web research + document assembly.
- ✓ Draft-and-prepare flows you approve at the end.

❗ It pauses and asks before consequential actions — logins, purchases, sending. You can watch, interrupt, and take over at any point. Those pauses are your seatbelt: if you're reflexively clicking “yes,” slow down.

⚠ Give access like you'd give it to a contractor on day one

No saved payment methods, minimal logins, watch the first runs. Trust is earned by track record — the agent's and the pattern's.

★ Aim agents at tasks that are tedious for humans: **breadth, repetition, and patience** are where agents beat people. Judgment is not. Agent runs are also the most expensive thing you can do with a subscription — capped monthly by plan.

Real jobs to try first

The best way to learn is a small, low-stakes job. Six real-world workflows — each a one-paragraph agent brief away. Start with the ones whose worst case is wasted minutes.

WEEKLY

Pipeline brief

“Every Monday, pull the pipeline, compare to last week, draft the one-page brief, post to #sales.”

DAILY

Morning digest

“Each morning, summarize unread team channels and today's calendar into a one-page brief in Drive.”

PER MEETING

Meeting prep

“Before each external meeting, assemble a prep sheet: who they are, recent news, our history, three talking points.”

MONTHLY

Competitor sweep

“Cluster the last 30 days of [competitor] news into themes with citations; flag what's accelerating.”

ON TRIGGER

Receipts → sheet

“When a receipt lands in this folder, extract date, vendor, category, total into the expense sheet.”

ON DEMAND

Notes → report

“Read this folder of meeting notes and draft the Q1 product-update report with sections and a summary.”

Start small to build trust

Begin with something reversible — a digest, a draft — and watch how the agent plans and asks permission. Once you trust its judgment, hand it bigger jobs.

Staying in control

Agents are powerful because they act across your real tools — so the platform is built to keep humans in charge. Here's the safety model, and the limits to keep in mind.

YOU'RE ALWAYS IN CHARGE

- ✓ **Admins choose the tools.** Agents can only touch the apps and connectors the workspace approved.
- ✓ **Roles gate building and publishing** — who may create agents, who may ship them to the team.
- ✓ **Agents show their plan and check in** before consequential actions.
- ✓ **Every run is attributable** — outputs land where you configured, reviewable like any colleague's work.

GOOD TO KNOW

- **Research preview.** Workspace agents are new (July 2026) and evolving fast — capabilities and limits shift.
- **Metered runs.** Agent work is token-billed; long jobs cost real money. Watch the usage dashboard.
- **An agent inherits the access you grant it.** Scope connectors minimally; widen only with track record.



One firm rule

Don't point agents at regulated workloads (health, financial records under audit rules) until your compliance team has signed off on the logging story — on Enterprise, with audit exports enabled.

04

PART 3

Codex

The most powerful surface, made approachable. Automate the repetitive and build real things — in plain English.

IN THIS PART

- Why it's worth it
- Install Codex
- Your first session
- AGENTS.md & project memory
- Tips, tricks & prompts
- Four ideas to grow into

The most powerful surface

Codex is OpenAI's coding agent: it works across many files at once, automates repetitive jobs, and builds things — from a data-cleaning script to a small website. It has a reputation as a “developer tool,” but a curious beginner can absolutely use it. Two things make it approachable:

1

You can start without a terminal

Codex runs in your browser (chatgpt.com/codex — cloud tasks on your repos), as an **IDE extension** (VS Code and friends), or as the classic **CLI** in a terminal. Pick whichever feels comfortable — they share the same brain and configuration.

2

You talk to it in plain English

You don't write code — you describe what you want. Codex figures out the steps, shows you what it's about to do, and asks permission before changing anything.

INCLUDED WITH PLUS

Codex is included in **Plus, Pro, Business, and Enterprise** — same login, no API key, no extra billing. It runs coding-tuned models from the current GPT-5.x line, with heavier reasoning available for harder tasks.

Two ways in, same power

This part shows the terminal path (worth learning) and notes the no-terminal alternatives at each step. You can switch any time.

★ Is this for you?

If you have a repetitive task involving lots of files (renaming hundreds of photos, reformatting data, generating reports from a template), or you're curious about building small tools — yes. If you just want to chat and make documents, Parts 1 and 2 already have you covered.

Install Codex

1 Install Node.js (one-time)

Codex CLI needs the free Node.js runtime. Download it from **nodejs.org** and run the installer. Skip this if you use the IDE extension or the web.

2 Install Codex

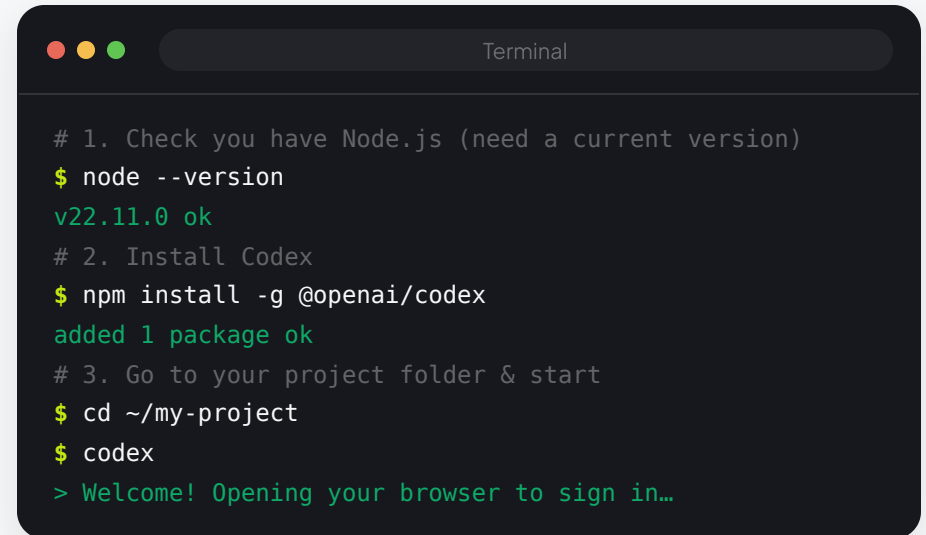
Open your Terminal app, paste the install command shown here, press Enter. A few seconds.

3 Open a project & run it

Navigate to a folder you want to work in (`cd` means “change directory”), type `codex`, press Enter.

4 Sign in

Your browser opens — log in **with your ChatGPT account**. Done. No API key needed.



```
Terminal

# 1. Check you have Node.js (need a current version)
$ node --version
v22.11.0 ok

# 2. Install Codex
$ npm install -g @openai/codex
added 1 package ok

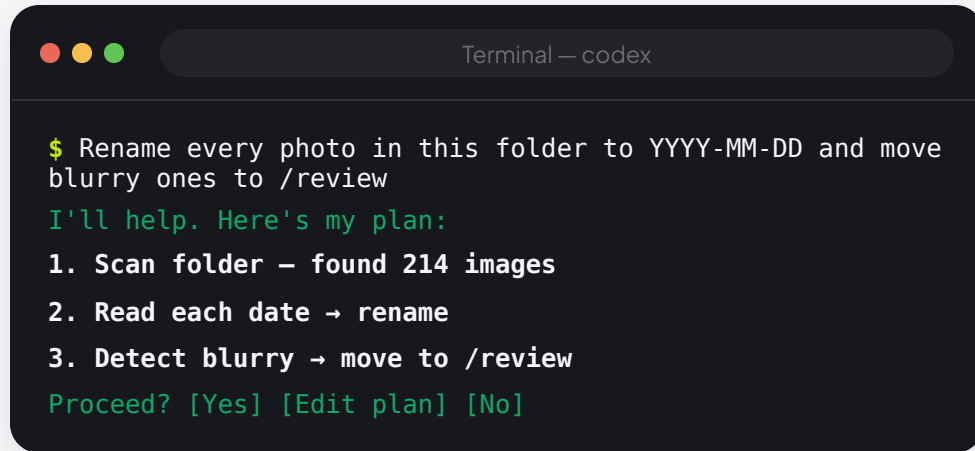
# 3. Go to your project folder & start
$ cd ~/my-project
$ codex
> Welcome! Opening your browser to sign in...
```

Prefer no terminal at all?

Install the **Codex IDE extension** and click the panel, or open **chatgpt.com/codex** to run cloud tasks against a connected repo — same engine, point-and-click. Skip steps 1–3 entirely.

Your first session

Once you're in, it's a conversation — just like the web app, except Codex can act on the files around it. Type what you want in plain English. Codex proposes a plan and waits for your yes.



```
Terminal — codex

$ Rename every photo in this folder to YYYY-MM-DD and move
blurry ones to /review
I'll help. Here's my plan:
1. Scan folder — found 214 images
2. Read each date → rename
3. Detect blurry → move to /review
Proceed? [Yes] [Edit plan] [No]
```

It always shows the plan before touching anything

A FEW COMMANDS WORTH KNOWING

Start in read-only, plan-first mode	<code>/approvals</code>
Create an AGENTS.md memory file	<code>/init</code>
Ask for a review of your changes	<code>/review</code>
Switch model / reasoning depth	<code>/model</code>
See everything you can do	<code>/help</code>

★ Keep a safety net

Approval modes are your friend: start in **read-only**, graduate to auto-with-approvals. Before big changes, copy the folder (or use Git). Codex asks before editing — a backup means you can always undo.

AGENTS.md: your project's memory

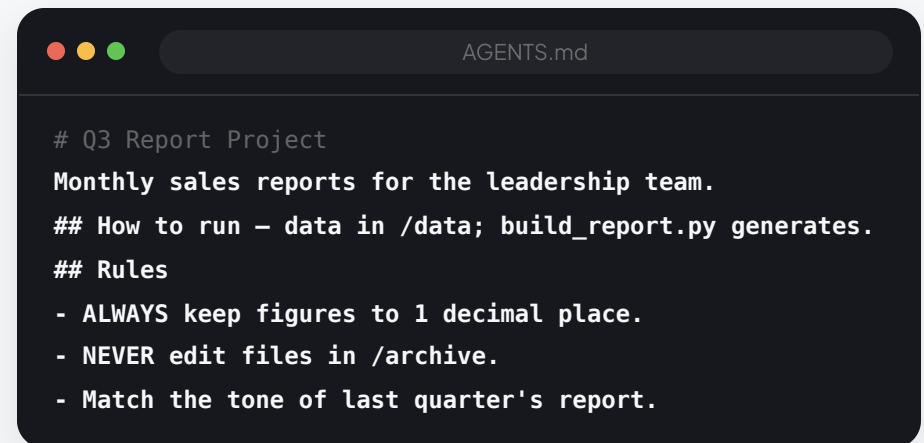
An AGENTS.md file is a plain text note that lives in your project folder. Codex reads it at the start of every session — so it always knows the project's rules without you re-explaining. It's the Codex equivalent of the custom instructions you set in Part 1.

CREATE ONE IN TWO SECONDS

Inside a session, just type `/init` — Codex looks at your folder and drafts a starter AGENTS.md. Then trim it: keep only what matters.

WHAT TO PUT IN IT

- A one-line description of the project
- The commands to run, test, or build it
- Key folders and what they're for
- “Always” and “never” rules Codex can't guess



```
# Q3 Report Project
Monthly sales reports for the leadership team.
## How to run – data in /data; build_report.py generates.
## Rules
- ALWAYS keep figures to 1 decimal place.
- NEVER edit files in /archive.
- Match the tone of last quarter's report.
```

★ Keep it short & high-signal — and alive

Aim for under a page; for each line ask “would removing this cause a mistake?” Every recurring correction you make to Codex’s output is a candidate rule — encode it, and the next session starts from a stronger baseline.

Codex tips, tricks & prompts

The habits that separate a good first week from a frustrating one — plus three prompts worth stealing.

THE HABITS

- ✓ **Small, scoped tasks win.** “Fix the date parsing in report.py” beats “improve my project.” Run three small tasks end-to-end before anything big.
- ✓ **Plan first, then approve.** For anything unfamiliar, ask for the plan and read it. Tiny obvious edits can skip it.
- ✓ **Read the diff, not the vibe.** Review what changed like a colleague's pull request. `/review` makes Codex critique its own changes first.
- ✓ **Ask for tests with the change.** “...and add a test that proves it” turns one-shot fixes into durable ones.
- ✓ **Keep a Backlog.md.** One file of small, actionable open items Codex can pick up across sessions — better than scattered TODOs.

> _ ORIENT ME (GREAT FIRST TASK)

[COPY / PASTE](#)

Explain this codebase to me like I'm a smart newcomer: what it does, the three most important files, how data flows through it, and one thing that looks fragile. Don't change anything.

> _ THE DISPOSABLE TOOL (NO CODE NEEDED)

[COPY / PASTE](#)

Write me a script that merges every CSV in this folder and flags duplicate rows — then walk me through running it step by step, assuming I've never used the terminal.

> _ REFACTOR WITH A SAFETY NET

[COPY / PASTE](#)

Add input validation to `[file]`. Plan first; make the smallest change that works; add a test that proves it; then summarize the diff and anything you're unsure about.

Four ideas to grow into

You don't need these on day one — but as you get comfortable, these four capabilities turn Codex from a helper into a genuine force-multiplier.

PARALLEL

Cloud tasks

Hand Codex a task at chatgpt.com/codex and it runs in the cloud against your connected repo — several tasks in parallel, each in its own sandbox, while your laptop stays free.

REVIEW

GitHub code review

Connect Codex to your repos and it reviews pull requests — a tireless first-pass reviewer that catches the mechanical issues before a human looks.

EXTEND

Skills & add-ons

Packaged abilities you drop into your setup for repeatable work — plus MCP connectors that let Codex reach real tools: a database, your project tracker, a browser.

STRUCTURE

Files as workflow

AGENTS.md for durable rules, Backlog.md for open work, a REQUIREMENTS.md for multi-session projects — plain files that act as structured handoffs between sessions.

The learning path that works

Use Codex for a week with plain requests and plan-first approvals. Add an AGENTS.md once you repeat yourself. Explore cloud tasks, review, and MCP only when you hit a real need. Let the tool grow with you.

05

Prompt engineering

No tricks or magic words — the handful of habits that make every result noticeably sharper. Treat ChatGPT like a brilliant new colleague with zero context on your work.

IN THIS SECTION

- The six habits that transform results
- Structure & the art of iterating
- Prompting reasoning models



The habits that transform results

1

Be clear & direct

- ✗ “Write about our product.”
- ✓ “Write a 100-word product description for first-time buyers; focus on ease of setup.”

2

Give context

- ✗ “Summarize this.”
- ✓ “Summarize this for a busy CEO deciding whether to fund the project.”

3

Show examples

- ✗ “Make it punchy.”
- ✓ “Match the tone of these two headlines: [paste]. Now write five more.” **The single biggest lever.**

4

Ask it to think

- ✗ “Is this a good deal?”
- ✓ “Think through the pros, cons, and risks step by step, then give a verdict.” (Or switch to Thinking.)

5

Assign a role

- ✗ “Give feedback on this.”
- ✓ “You are a veteran editor. Give feedback on this draft...”

6

Name the format

- ✗ “Tell me about options.”
- ✓ “Three options as a table: name, cost, one-line trade-off.”

❗ You won't use all six every time — even one or two lifts results noticeably. For anything important, stack several. And remember the universal shortcut: **“Ask me any questions you need before you start.”**

Structure long prompts; iterate short ones

GIVE LONG PROMPTS A SKELETON

When a request has several parts, **label them** — plain headings remove ambiguity about what's background and what's the instruction.

>_ A STRUCTURED PROMPT

[COPY / PASTE](#)

CONTEXT: We're a B2B agency pitching a retail client. Tone confident, not salesy.

TASK: Write a 3-paragraph intro email requesting a 30-minute call.

FORMAT: Subject line, then body. Under 150 words.

❗ Wrong turn three replies ago?

Edit your earlier message (pencil icon) and re-run from that point — pruning the bad branch beats decorating it.

ITERATE — DON'T START OVER

- “Shorter — half the length.”
- “More formal, drop the exclamation marks.”
- “Add a concrete example to point 2.”
- “Give me three more options in the same style.”
- “What would make this stronger? Then apply it.”

★ That last one — **critique, then apply** — is the cheapest quality upgrade in the product: you're renting the model's critical eye against its own draft.

Reasoning models are briefed, not micromanaged

Thinking models plan their own approach. Your job shifts from steps to brief: goal, constraints, success criteria — then let it work.

	Instant models	Thinking models
Specify	Detailed guidance helps	Goal, constraints, success criteria
Skip	—	“Think step by step” (redundant) · over-scripted procedures (they cap a better plan)
Add	Format & examples	“Flag assumptions rather than guessing” · what a *good* answer contains

★ Ask to **see the reasoning** when it matters: “Summarize the key steps of your thinking before the recommendation.” You're auditing the logic, not just the verdict.

>_ A GOOD THINKING BRIEF

[COPY / PASTE](#)

Evaluate whether we should switch our reporting stack to **[option]**. Constraints: budget under \$X, no downtime during Q4, team of 3. A good answer weighs cost, migration risk, and lock-in, and ends with a recommendation plus the two strongest counterarguments. Flag any assumption you had to make.

i Auto handles the routing day-to-day. These habits are for when **you** choose Thinking deliberately: the analysis, the plan, the decision memo — and they're exactly how you brief a workspace agent or Codex, too.

06

Top apps to try

Apps are the successor to plugins: real products — Spotify, Canva, Zillow — running inside ChatGPT with interactive interfaces. Here are the ones worth knowing, and exactly how to connect them.

IN THIS SECTION

- What an app is & how to connect one
- The apps worth trying right now
- Connectors vs. apps — which is which

What an app is — and how to connect one

An app lets ChatGPT work inside another product — with interactive cards, maps, and designs right in the chat. (If you read about “ChatGPT plugins” years ago: this is what replaced them.)

1 Just name it

Apps respond to natural language — start a message with the app's name: **“Spotify, build me a focus playlist.”** First use, ChatGPT asks you to connect the account.

2 Or let ChatGPT suggest


When an app fits the conversation, ChatGPT offers it at the right moment — “Figma could turn this into a diagram.”

3 Manage connections

Settings → Apps & Connectors: see what's connected, what data each app can see, and disconnect anything, any time.

⚠ A word on data

Connecting an app shares conversation context with that app's maker. Connect what you use, skip what you don't, and review the list quarterly — same discipline as phone-app permissions.



Apps are interactive, not just text answers

📘 Apps vs. connectors

Apps bring another product's *interface* into chat (Spotify, Canva). **Connectors** give ChatGPT quiet access to your *data* (Drive, Gmail, SharePoint, calendar) for search and context. You'll use both.

The apps worth trying right now

The launch partners are still the strongest of the directory — all free to connect, available wherever each service operates.

CREATE

Canva

“Turn these bullet points into a poster / carousel / deck in our brand colors.” Design drafts without leaving the chat.

DIAGRAM

Figma

Turns conversational context into FigJam diagrams — flowcharts and process maps from a description.

LISTEN

Spotify

Playlists and podcast picks from a described mood or task — “deep-work instrumental, 90 minutes.”

LEARN

Coursera

Course recommendations and lesson Q&A while you study — “what should I take to learn SQL?”

HOMES

Zillow

Real-estate search as a conversation — interactive maps, listings, affordability, refined in plain language.

TRAVEL

Expedia + Booking.com

Flights, stays, and itineraries inside the chat — compare, then hand off to book.

YOUR DATA

The connectors

Google Drive, Gmail, Calendar, SharePoint, Teams — quiet context access that makes “find the deck from March” work.

FOR TEAMS

Admin-approved

In Business/Enterprise workspaces, admins choose the approved set — and workspace agents (Part 2) inherit it.

★ How to evaluate any new app

The directory grows monthly (Settings → Apps, or the directory page). Before connecting: who makes it, what data does it see, does it save you real time? Two yeses and a shrug = skip it.

07

Pitfalls & ethical use

Capability without judgment is how people get burned. This is the judgment — five pitfalls with cheap, boring, reliable fixes, and the professional standard for honest AI.

IN THIS SECTION

- The five pitfalls — and the fix for each
- Using AI honestly

The five pitfalls — and the fix for each

1 Hallucination (the big one)

It states false things fluently — fabricated facts, citations, quotes, numbers. Not lying; pattern-completing. **THE FIX** → verify anything with consequence: click citations, check quotes, recompute key numbers.

2 Stale knowledge

Without search, it answers from its training cutoff. **THE FIX** → for anything time-sensitive, Search on + citations present.

3 Over-trust on math & details

Plausible ≠ correct — arithmetic in prose, names, dates, legal/medical specifics. **THE FIX** → route calculations to data analysis (real code); treat specifics as unverified until checked.

4 Privacy leaks

Confidential material in the wrong account or tier. **THE FIX** → client & company data only in workspace tiers (or a training-opted-out account); sensitive one-offs in Temporary Chat.

5 The competence mirage

Polished output switches off your critical reading — usually around week three. **THE FIX** → the new-hire rule: brilliant, fast, and reviewed like a junior colleague's work. Every time, forever.

THE ONE-LINE POLICY

AI drafts; humans decide.

If output leaves your hands with your name on it, you've verified it. If your organization adopts a single sentence of AI governance, make it this one.

- ① Understand hallucination mechanically, not morally: the tone never tells you whether you got truth or plausibility. **Consequence triggers checking** — that's the habit. It applies double to agents, which act on their own conclusions.

Using AI honestly — the professional standard

DISCLOSURE

Say so when it's expected

Know when your client, publication, platform, or regulator expects AI use to be declared — many contracts now have AI clauses. When in doubt, disclose. Nobody ever damaged trust by being honest about their tools.

OWNERSHIP

You own the output

The errors, the biases, the plagiarism risk, the consequences. “The AI wrote it” has never once worked as a defense.

FAIRNESS

Check the framing

Models carry biases from their training data. For work touching people — hiring, screening, targeting — ask: * “What perspectives or groups might this framing overlook or disadvantage?” *

PEOPLE FIRST

Augment, not replace

Use AI to make your team faster and sharper — never to quietly hollow out the expertise your clients are actually buying. Clients buy your thinking; AI accelerates it.

Know your policies

✓ your employer's AI policy · ✓ each client's AI clauses (increasingly in MSAs) · ✓ your industry's rules — legal, medical, financial, and academic contexts have real ones.

★ The standard that never goes out of date

Use AI in a way you'd be comfortable explaining — to the client, in plain language, with the “why.” If the explanation would embarrass you, the usage is wrong.

The one-page cheat sheet

Tear this page out (or bookmark it). Everything you'll reach for most, in one place.

PICK A MODE

Everyday work	Auto
Genuinely hard analysis	Thinking
The deliverables that matter most (Pro)	Pro mode

WHICH SURFACE

Think, write, analyze	ChatGPT · web
Team workflows, end-to-end	ChatGPT Work
Automate / build	Codex

WEB APP SHORTCUTS

Send / new line	Enter / Shift+Enter
Attach a file	+ in the box
Private chat	Temporary Chat toggle
Fix an earlier turn	edit (pencil) & re-run

WHICH TOOL FOR THE JOB

Current events + receipts	Search
A question that deserves a document	Deep Research
Spreadsheet questions	Data analysis
Tedious multi-site legwork	Agent mode
Recurring context	Project
On a schedule (solo / team)	Task / workspace agent
Repeated workflow, shared	Custom GPT
A real deliverable to draft	Canvas
Sensitive one-off	Temporary Chat

THE FIVE-PART PROMPT

- **Role** — who ChatGPT should be
- **Context** — the situation & audience
- **Task** — exactly what to do
- **Examples** — show the style
- **Format** — the shape of the output

CODEX BASICS

Start	type codex
Project memory	/init
Approval modes	/approvals
Review changes	/review

★ The universal shortcut

Stuck? Add: "Ask me any questions you need before you start."

Your 4-week rollout plan

Don't try to do everything at once. Move to the next stage when the current one feels boring — boring means internalized.

STAGE 1

Day 1

Web foundation

- ✓ Sign up & upgrade to Plus
- ✓ Lock down (MFA / passkey) & set Data Controls
- ✓ Fill in custom instructions
- ✓ First chats: the three starter prompts
- ✓ Model on Auto; learn Temporary Chat

STAGE 2

Week 1

Skills & apps

- ✓ R-C-T-E-F on real work daily
- ✓ Files in, Canvas out; build your first Project
- ✓ One Deep Research run, properly briefed
- ✓ Connect one app + one connector
- ✓ Schedule one Task

STAGE 3

Week 2

ChatGPT Work

- ✓ Try Agent mode on a read-only job
- ✓ (Team) enable workspace agents; approve tools
- ✓ Build one agent from a real brief
- ✓ Run it twice; encode every correction
- ✓ Watch the usage dashboard

STAGE 4

Week 3+

Codex

- ✓ Start in the IDE extension or web (no terminal)
- ✓ Or install the CLI
- ✓ Run /init, trim your AGENTS.md
- ✓ Three small tasks end-to-end
- ✓ Keep a backup before big changes

When to upgrade beyond Plus

Repeatedly hitting weekly caps → **Pro**. Two+ people with client data, or you want workspace agents → **Business**. Security review in the room → **Enterprise**.

YOU'RE READY

Start small. Verify everything. Let it grow with you.

You now have the full picture: the web app for thinking, ChatGPT Work for team workflows, and Codex for automation — tied together by the prompting and verification habits that make it all compound. Pick one task this week and try it.

Keep learning

OpenAI ships changes almost monthly. Your own account is always the source of truth for what's live — this guide is the map that keeps the changes legible.

A few honest caveats

Model names change. Limits shift. Workspace agents are a research preview and metered. Verify specifics before you rely on them.

Need a hand?

Cortexa helps teams put AI to work — setup, training, governance, and custom workflows. Come say hello.