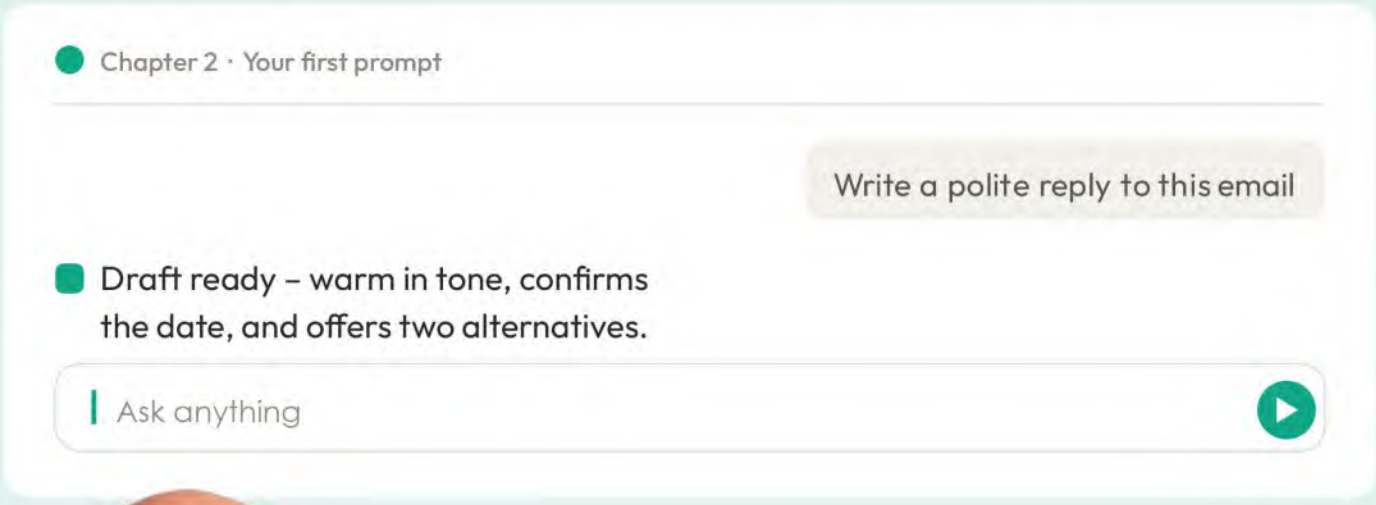
«packt»

First Edition

ChatGPT Visual Bible

A step-by-step, screenshot-by-screenshot guide to writing, researching, analyzing, and automating with ChatGPT





Mark J. Price

ChatGPT Visual Bible

A step-by-step, screenshot-by-screenshot guide to writing, researching, analyzing, and automating with ChatGPT

Mark J. Price

‹packt›

ChatGPT Visual Bible

Copyright © 2026 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

Portfolio Director: Ashwin Nair
Relationship Lead: Nitin Nainani
Project Manager: Ruvika Rao
Content Engineer: Hayden Edwards
Technical Editor: Arjun Varma
Copy Editor: Hayden Edwards
Production Designer: Ajay Patule
Growth Lead: Sayantani Saha

First published: September 2026

Production reference: 1280827

Published by Packt Publishing Ltd.
Grosvenor House
11 St Paul's Square
Birmingham
B3 1RB, UK.

ISBN 978-1-80865-515-9
www.packtpub.com

Contributors

About the author

Mark J. Price helps beginners and everyday knowledge workers get real, practical value out of AI tools like ChatGPT and Claude. Based in Brighton, England, he brings two decades of experience as a technical author, including ten editions of a bestselling C# and .NET series for Packt, to making AI approachable for readers with no technical background. A former maths teacher and Microsoft Certified Trainer, he comes from a family of teachers and now chats with readers in Packt's Discord.

> Thank you to all my readers. Your support means I get to write these books and celebrate your successes. Special thanks to the readers who give me actionable feedback via my GitHub repository and email, and interact with me and the book communities on Discord. You help make my books even better with every edition.

Get in touch

Feedback from our readers is always welcome.

General feedback: If you have questions about any aspect of this book or have any general feedback, please email us at `customercare@packt.com` and mention the book's title in the subject of your message.

Errata: Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book, we would be grateful if you reported this to us. Please visit `http://www.packt.com/submit-errata`, click **Submit Errata**, and fill in the form.

Piracy: If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at `copyright@packt.com` with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit `http://authors.packt.com/`.

Free benefits with your book

This book comes with free benefits to support your learning. Activate them now for instant access (see the "*How to Unlock*" section for instructions).

Here's a quick overview of what you can instantly unlock with your purchase:







How to Unlock

Scan the QR code (or go to `packtpub.com/unlock`). Search for this book by name, confirm the edition, and then follow the steps on the page.



Note: Keep your invoice handy. Purchases made directly from Packt don't require one

Table of Contents

Introduction 1

Why you picked up this volume 1
How this volume is structured 4
How to read the four books 5
Get the sample files and bonus content 6
A note about figures 7
Why the books will not go out of date 7
Why not just ask ChatGPT to explain itself? 8
Key terms to help you follow along 9
What prerequisites do you need? 10

Book 1: Ask ChatGPT 11

Chapter 1.1 Understand What ChatGPT Can Do 12
Chapter 1.2 Create Your Account and Choose a Plan 17
Chapter 1.3 Use ChatGPT on the Web, Phone, and Desktop 22
Chapter 1.4 Have Your First Useful Conversation by Typing, Talking, or Showing 29
Chapter 1.5 Write, Refine, and Improve Prompts Through Conversation 34
Chapter 1.6 Learn Something New and Educate Your Child 40
Chapter 1.7 Plan Everyday Life and Compare Important Choices 46
Chapter 1.8 Search and Verify Current Information 49
Chapter 1.9 Conduct Deeper Research with Cited Sources 52
Chapter 1.10 Add Documents, Images, and Data to a Chat 55
Chapter 1.11 Personalize and Organize with Instructions, Memory, and Projects 58
Chapter 1.12 Project – Build Your Personal AI Workbench 63

Book 2: Create Real Work Results 67

Chapter 2.1 Use ChatGPT Safely and Responsibly at Work 67
Chapter 2.2 Choose the Right Mode for Each Task 71
Chapter 2.3 Draft Professional Documents 78
Chapter 2.4 Summarize and Extract Information 84
Chapter 2.5 Analyze Spreadsheets and Data 88

Chapter 2.6 Create and Edit Images and Media 93
Chapter 2.7 Create Presentations and Visual Reports 100
Chapter 2.8 Add Plugins, Apps, and Skills 106
Chapter 2.9 Project – Produce a Useful Weekly Briefing 112

Book 3: Automate Recurring Tasks 119

Chapter 3.1 Choose the Right Automation Building Block 120
Chapter 3.2 Schedule Recurring and Conditional Work 124
Chapter 3.3 Turn a Procedure Into a Reusable Skill 130
Chapter 3.4 Build a Custom GPT 139
Chapter 3.5 Automate Personal Organization and Learning 145
Chapter 3.6 Build a Research and Content Workflow 147
Chapter 3.7 Build a Decision and Operations Workflow 149
Chapter 3.8 Project – Build and Measure an End-to-End Workflow 150

Book 4: Advance Your Career and Income 153

Chapter 4.1 Choose What Better Work Means to You 154
Chapter 4.2 Audit Your Role and Find High-Value Work 156
Chapter 4.3 Measure and Document Your Improvement 157
Chapter 4.4 Learn the Skills Your Next Role Requires 160
Chapter 4.5 Combine AI Fluency with Human Strengths 162
Chapter 4.6 Propose and Lead an AI Improvement at Work 164
Chapter 4.7 Build a Case for Promotion or Higher Pay 166
Chapter 4.8 Prepare for a Better Job or Career Change 168
Chapter 4.9 Rehearse High-Stakes Conversations with Voice Mode 171
Chapter 4.10 Build a Portfolio and Measurable Case Study 174
Chapter 4.11 Turn a Problem into a Service or Product 177
Chapter 4.12 Decide How to Respond to AI in Your Field 179
Chapter 4.13 Project – Complete a 30-day Career Value Project 181

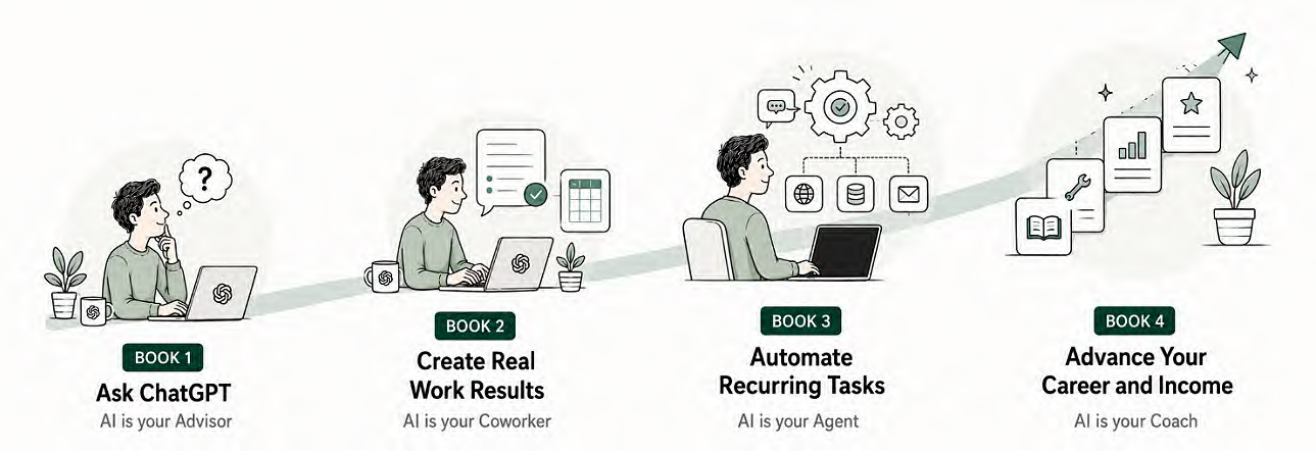
Epilogue 185

Introduction

Welcome to *ChatGPT Visual Bible*!

In this introduction, you will learn:

- Why you picked up this volume
- How this volume is structured
- How to read the four books
- Get the sample files and bonus content
- A note about figures
- Why the books will not go out of date
- Why not just ask ChatGPT to explain itself?
- Key terms to help you follow along
- What prerequisites do you need?



Why you picked up this volume

At this point in your time with AI, you may have sensed that ChatGPT could help you work faster and think more clearly. Maybe you have asked it a few questions, got it to rewrite an email, or asked for an idea when you got stuck. But you probably still treat it as a chatbot: something you type into, not something you create, do deep research, or automate aspects of your life with.

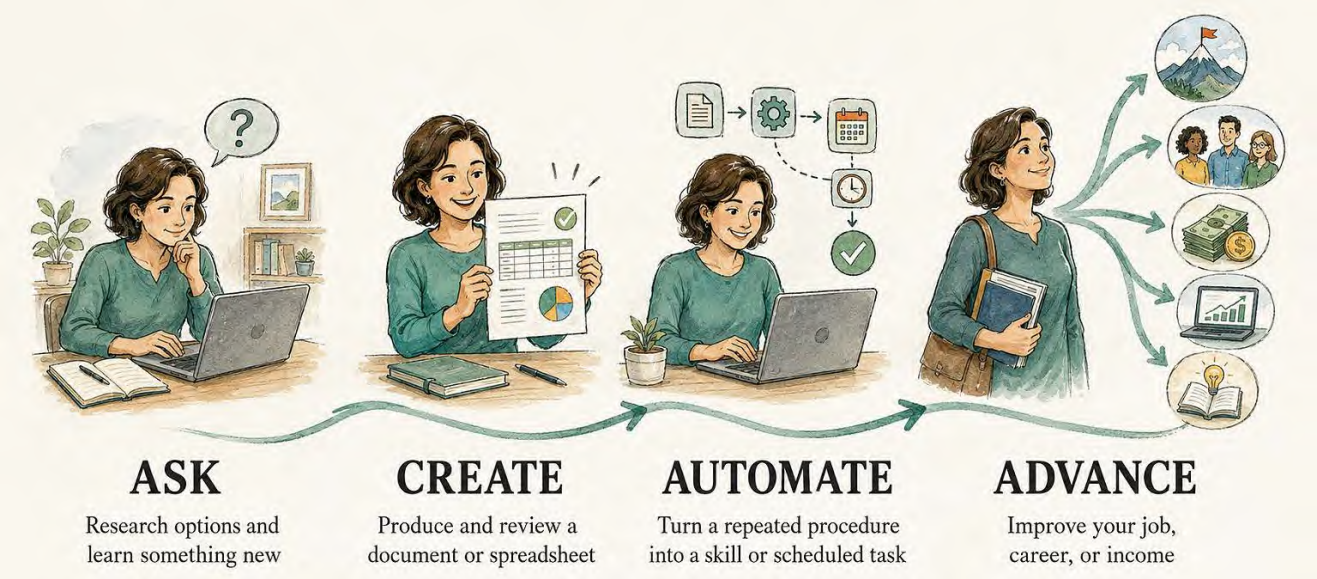
You are not late. In OpenAI's own worldwide usage data, the share of ChatGPT messages sent by people over the age of 35 rose in almost every country during the year ending in the second quarter of 2026. Worldwide, these users now send a 5% higher share of messages than they did a year earlier. If you are middle-aged and mid-career, and you feel like the AI game is only for the young, know that you are joining a wave of people in the same position who are only getting started now too.

Closing the gap between casual, occasional ChatGPT use and using it to improve your work and personal life is why I wrote this volume. You do not need a technical background to close the gap. You need a clear path from your first conversation to confident, professional use, with checkpoints along the way that reinforce what you learn.

Most guides to ChatGPT choose one of two extremes. Some try to cover several large subjects, including coding, in one slim volume, so each subject gets miserly treatment. Others promise easy income from a handful of prompts, without teaching you how to check whether the result is any good.

Neither approach matches how people truly build value with AI: they ask for useful guidance, produce a checked piece of work, turn a proven procedure into something repeatable, then use that evidence to improve a career or generate a new income stream.

ChatGPT Visual Bible is different because it follows that better path, one step at a time:



Introduction Figure 1: Four steps to advance your work and home life with AI

This guide also gives you something other books don't: premade files to practice on. Every chapter includes sample files matched to that chapter's task, like a Word document to summarize or a spreadsheet to analyze, so you can start immediately instead of hunting for your own test material or skipping the exercise altogether.

Alongside those files, tasks come with one or more prompts you can copy straight into ChatGPT. You're not left to guess how to phrase a request or peer at tiny text in a poorly cropped screenshot. You get a working prompt as a starting point, then a clear explanation of which parts to tweak for your own documents, your own job, and your own goals.

Book 4 also takes a wider view of what building value with ChatGPT means. The obvious move at that stage is selling a product or service, but that route has genuine risk and a low success rate. *Book 4* covers that path, but treats it as one option among several, alongside using ChatGPT to work faster in your current job, build the case for a promotion, or prepare for and land a new career. For most readers, improving what they already do is a far more practical, more likely route to success than starting a business from nothing.

That practical scope extends past work, too. In *Book 1* you'll find tasks aimed at your personal life, such as planning, organizing, and budgeting at home, or helping your kids with their homework and teaching them how to use AI responsibly.

Who this book is for

This book is for you if you have the curiosity and the everyday computer skills most adults already do, such as using a web browser, installing an app, and working with documents, spreadsheets, and images.

You will get the most from this book if you recognize yourself in one or more of these situations:

- You are mid-career and want to close the gap between casual ChatGPT use and confident, professional use
- You manage a team, a small business, or your own workload, and want to work faster without sacrificing quality
- You want to build a habit of checking AI output instead of trusting it automatically
- You want to use ChatGPT for personal tasks too, such as planning, budgeting, or helping your kids with homework
- You are looking for a path from your first conversation with ChatGPT to a completed piece of work, a repeatable process, or a stronger case for a promotion or new role

This book does not teach software development. If your goal is to write code or build websites and apps with AI, look for the companion title, *ChatGPT Codex Visual Bible*, instead.

A flesh-and-blood author writing for a respected publisher

I am a living, breathing human based in Brighton, on England's south coast. Close enough to London for exciting cultural weekends, breezy enough to survive the increasingly sweltering British summers.

I've written more than 20 programming titles for my publisher Packt since 2016, including ten editions of my best-selling C# and .NET book. I update it every year in November, and I plan to do frequent updates for *ChatGPT Visual Bible* too, since ChatGPT is improving so rapidly.

Teaching is in my blood. My mother was a maths teacher, my sister a primary school teacher, my grandfather taught at the Royal Military Academy Sandhurst, and teaching and military careers run through my uncles and cousins too. I taught GCSE and A-Level maths in two London secondary schools, and I hold a Computer Science BSc (Hons) from the University of Bristol.



Introduction Figure 2: Me and my background

That mix, maths teacher and software developer, turns out to be exactly what a beginner's *ChatGPT Visual Bible* needs. Teaching maths to teenagers taught me how to break an intimidating idea into the smallest step that still makes sense, and how to spot the exact place a student gets stuck before they even ask. Software development taught me precise problem-solving: to test whether something actually works rather than assume it does, and to write instructions a computer, or a nervous beginner, can follow without guessing. Put those together and you get a guide that walks you from confusion to confidence, one checked step at a time, the way a good maths teacher walks you from "I don't get it" to "oh, that's actually simple."

There's another reason my software developer background shapes how I teach here, beyond teaching precision: programmers are the canary in the coal mine for what AI does to a profession. I've spent the last few years watching AI move from autocomplete to something that can draft, test, and ship real code. That's a transition I lived through in my own daily work, which is exactly why I can show you, concretely, what to watch for as the same pattern reaches and transforms your profession.

Writing books for a decade took my face-to-face classroom teaching instinct further: it's how I now reach far more people, in far more places, than any classroom could, and I love hearing from readers all over the world. It's my way of making a dent in the universe, as Steve Jobs mused he wanted to do.

My involvement didn't end when I delivered the final drafts to my publisher. I am active in Packt's Discord channels, where you can chat with me and other readers about my books and share tips and tricks on how to get the most out of AI.

You can learn more about me and confirm my credentials at the following links:

- GitHub repositories: `https://github.com/markjprice/`
- LinkedIn: `https://www.linkedin.com/in/markjamesprice/`
- Amazon: `https://www.amazon.com/stores/Mark-J.-Price/author/B071DW3QGN`
- Packt: `https://subscription.packtpub.com/search?query=mark+j.+price`
- Discord channel: `https://discordapp.com/channels/963343534323269662/963344584287604736`

How this volume is structured

This single volume contains four books:

Book title	AI role	What changes for you
Book 1 – Ask ChatGPT	Advisor	You move from asking occasional questions to holding a useful conversation, doing research, learning new skills, and organizing your work using ChatGPT features like projects.
Book 2 – Create Real Work Results	Coworker	You produce documents, spreadsheets, presentations, analysis, images, and other media, and you learn to supervise longer, multi-step assignments.
Book 3 – Automate Recurring Tasks	Agent	You turn a procedure you repeat often into a skill, a Custom GPT, or a scheduled task that ChatGPT can run for you automatically.
Book 4 – Advance Your Career and Income	Coach	You use what you learned in the other books to earn a promotion, make a career change, or provide a service to generate additional income.

This structure is not a guess. OpenAI, the company that builds ChatGPT, tracks how people around the world actually use it through a project called OpenAI Signals. Its August 2026 report found that at work, doing already dominates asking: people are more than twice as likely to use ChatGPT to complete a task or create something, such as writing, coding, or analysis, at work than outside of work. Outside work, asking a question or seeking information still leads.

Each of the four books in this volume contains about ten concise chapters and ends with a practical hands-on project, so you finish every book holding a completed piece of work, not just a lot of theory swimming around in your head.

You pay OpenAI only when you're ready to use what their paid subscription plan gives you. You can complete nearly all of *Book 1* on the **Free** plan. As your needs grow into real work and repeatable automation in *Book 2* and *Book 3*, you can upgrade your plan to **Plus** or higher, which covers everything you'll need: ChatGPT Work, scheduled tasks, connected apps, skills, and plugins. Only workspace agents and fully synced connectors sit beyond an individual plan, reserved for organization-managed workspaces most readers won't need.

How to read the four books

If you are new to ChatGPT, then read the four books in order. Each one builds on the concepts and skills from the one before it. For example, you cannot automate a procedure in *Book 3* until you can produce reliable work in *Book 2*, and you cannot produce reliable work until you can hold a useful conversation, which is what *Book 1* teaches.

If you already use ChatGPT regularly, start closer to where your own skills stop. Skim the chapter headings at the start of *Book 1* and *Book 2*. If the tasks look familiar, move ahead to *Book 3*.

Inside each chapter, look for several types of recurring boxes:

- **Watch out** boxes warn you about a specific mistake or risk
- **Good practice** boxes give you a habit that improves your results
- **What to ask next** boxes give you an optional follow-up prompt you can send right after finishing the chapter's task
- **Real-world example** boxes give you a need solved by actual people using AI

Two more boxes point you outside the books:

- **Learn more online** boxes link to an official resource, such as documentation from the AI provider. Interfaces and features change often. When a detail is likely to change after you buy this volume, I give you a link to a verified source instead of guessing.
- **Sample files** boxes link to a file or folder in this volume's online repository. Download these files when a chapter's activity needs them, if you don't want to use your own. You can open them, upload them into ChatGPT, or copy from them, depending on what the chapter asks you to do.

You don't need to use every box. Read the ones that help with what you're doing right now.

- Each chapter ends with a **Try it now** section that gives you one task you can finish in five to fifteen minutes. It **also** gives you a few questions to check what ChatGPT produced, so reviewing its work becomes a habit rather than an afterthought.

Every **Try it now** task in this volume is written to work for as many readers as possible, so it stays intentionally general: a status update, a spreadsheet, a presentation, rather than a task built around one specific job. That breadth is what lets a teacher, an accountant, and a shop manager all work through the same chapter. If you'd like a task closer to your own role, I keep a growing online list that matches common professions to the chapter where a suggested, tailored task fits, so you can try something that maps directly onto your day-to-day work instead of adapting the general version yourself. Find it at `https://github.com/markjprice/chatgpt-vb/blob/main/docs/bonus/professions/README.md`. I add new professions and tasks to that list over time, so it's worth checking back even after you've finished the book.

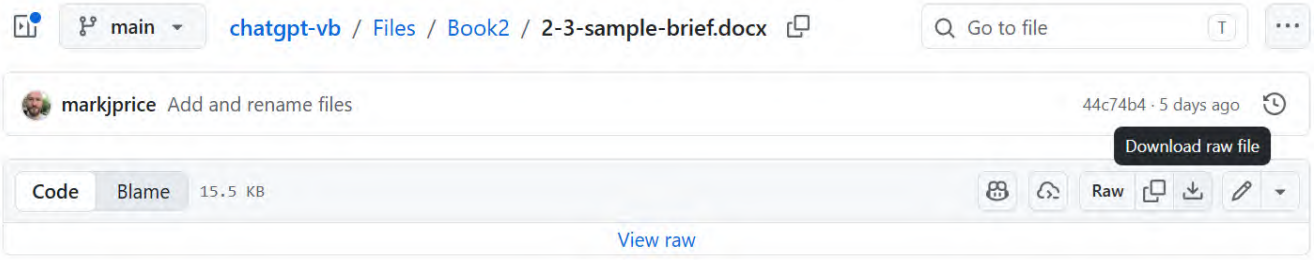
Get the sample files and bonus content

Other ChatGPT guides expect the reader to supply their own files to work with while learning skills. *ChatGPT Visual Bible* provides sample files, so you don't have to scramble to think of or create your own. These files are stored in a repository in my GitHub account so I can keep them updated.

Don't worry! Although GitHub is designed for software developers, anyone can use it. For example, you can go to the following link to see sample files you could use in *Book 2*:

`https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2`

Just click a file like `2-3-sample-brief.docx` and then click the **Download raw file** button:



Introduction Figure 3: Download a sample file from this volume's GitHub repository

You can also download all the folders and files in one go using the following link:

`https://github.com/markjprice/chatgpt-vb/archive/refs/heads/main.zip`

Once the ZIP file is downloaded, extract its contents to your local computer folder.

I have also written three accompanying bonus books:

- *Beyond Your First AI*: This bonus book explains when a different AI service is the better tool for a specific job, or a backup when your primary AI is unavailable. It covers Anthropic Claude, Google Gemini, Mistral Vibe, and some Chinese providers like DeepSeek and Qwen.
- *Profession-Specific Prompts*: This bonus book adds a layer that the general prompting skills taught in the print books do not cover: the specific duties, risks, and rules that come with 20 professions.
- *Local AI with gpt-oss*: The print books taught you to work with ChatGPT running in OpenAI's cloud. This bonus book covers running an AI model entirely on your own laptop or desktop, with nothing leaving your device, using OpenAI's two free and open-weight models, `gpt-oss-20b` and `gpt-oss-120b`.

You can find links to these in the *Epilogue*.

A note about figures

The word *Visual* in *ChatGPT Visual Bible* is a promise. Every screenshot and diagram does one specific job: it helps you decide what to do, shows you what ChatGPT is doing, walks you through an action, or helps you judge whether a result is good enough to use:

- **Screenshots** are cropped to the part of the screen that you need to see and enlarged so the interface text stays readable on a printed page. Where a screenshot shows several controls, I number them and explain each one in order.
- **Diagrams** are reserved for ideas that are hard to explain in a sentence, such as how a project holds context or how a process moves from a trigger to a finished result.

I never place essential information only inside an image. When a figure is important to a task, the surrounding text explains it too, so the instructions still work even if a picture is hard to see.

One of my favorite benefits offered by my publisher is that for every purchase, Packt provides free PDF and online web versions. You can unlock your free benefits by turning to the *Free benefits with your book* section in the book's front matter. This will allow you to zoom into full-color screenshots if that helps.

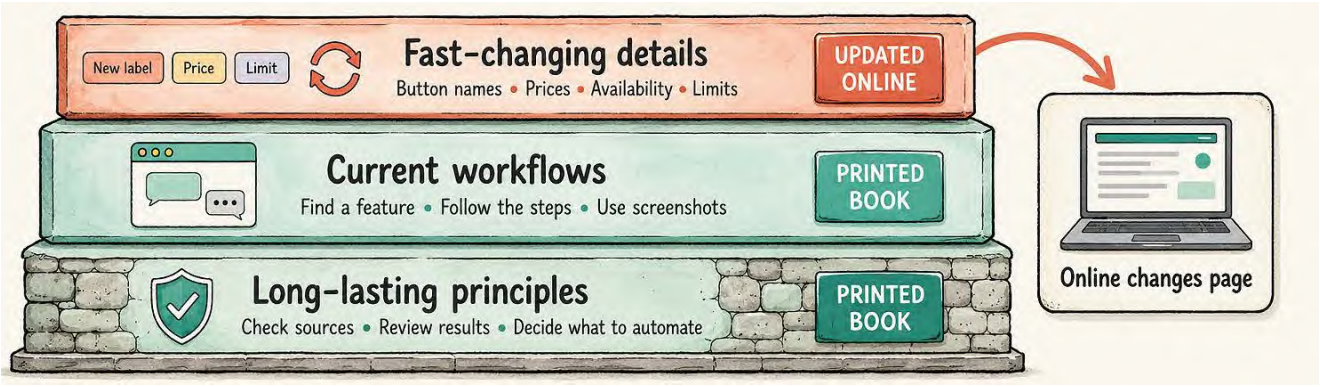
Why the books will not go out of date

ChatGPT's menus, buttons, and features change often. Books built only around today's exact user interface would be out of date within months.

So, the books teach decisions and tasks that stay useful after the interface moves: how to judge whether a source is reliable, how to decide whether a task belongs in a quick chat or a longer piece of work, how to tell when a procedure is stable enough to automate. These long-lasting principles will serve you well while ChatGPT evolves rapidly over the coming months and years.

Screenshots illustrate these ideas using the interface as it exists today, and each set of instructions combines a stable landmark, such as a menu's purpose, with the exact label currently shown on screen. If a label changes, the landmark still points you to the right place.

Where a detail is likely to move, such as plan prices or a settings screen, I have placed it on a companion page you can check online instead of printing it here. I will keep that page updated after the volume is published, so the printed pages can stay focused on what to decide and why, not on details a future software update could make wrong.



Introduction Figure 4: A solid foundation, current screenshots, with fast-changing details updated in an online changes page in the book's GitHub repository

Visit `https://github.com/markjprice/chatgpt-vb/blob/main/changes.md` for the current list of changes. I have organized it by chapter number, so you can jump straight to the section that affects the page you are reading. The GitHub page shows:

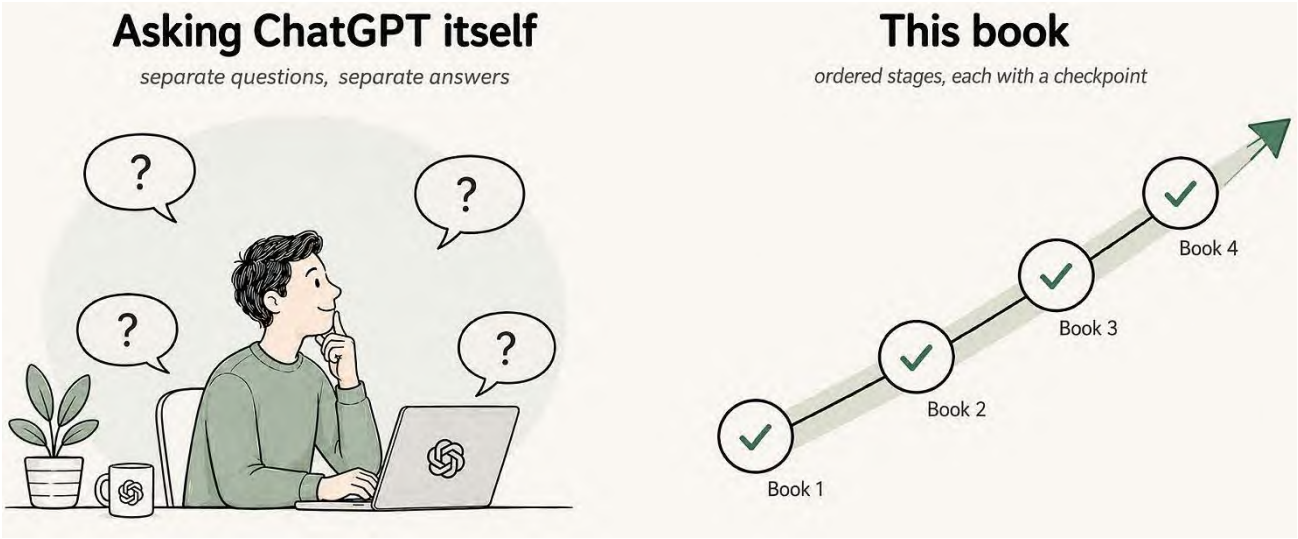
- **Interface changes**: A new menu name, a moved button, or a renamed setting, matched to the chapter and step where the old version appears in this volume.
- **Plan and pricing changes**: Current plan names, prices, and usage limits, since these move more often than the underlying features.
- **Feature availability**: Notes on when a feature described as a preview becomes generally available or when a feature is removed.
- **Region and account differences**: Cases where a feature works differently, or not at all, depending on your country, device, or account type.
- **Corrections**: Any error identified in the volume after printing.

Why not just ask ChatGPT to explain itself?

You can ask ChatGPT to explain its own features, and you should keep doing that after you finish this volume. Still, it is a fair question to ask why you need books at all. Here's why: If you ask ChatGPT four separate questions, you get four separate answers, each shaped by how you happened to phrase the question that day. The four books in this volume order material into four stages, each building on the skills from the one before it, and each ending with a project that gives you something finished to show for it. This volume is a curated experience crafted by an education expert.

ChatGPT can describe a feature confidently and still be wrong about whether that feature is available on your plan, in your region, or on your device. The books state those limits where they matter and point you to a page I keep updated after printing.

ChatGPT will not check its own work for you. The books build that check into every chapter, with a short Check the result section that asks you to review accuracy, sourcing, and safety before you rely on anything ChatGPT produced. The books have been checked against primary sources by my publisher's editors and other members of their production team before publication. A chat response has not. I verify product details against official documentation, mark anything uncertain, and correct errors on the companion update page.



Introduction Figure 5: Comparing asking ChatGPT itself and reading this volume



Good practice: Use ChatGPT for quick, current answers. Use the books for the parts that need structure, verification, and a plan you can follow to the end.

Key terms to help you follow along

This glossary covers the terms used most often throughout the books:

Term	Definition
App	A connected service that gives ChatGPT access to information or actions outside the conversation, such as a calendar or a file storage service.
Chat	The ChatGPT mode built for questions, discussion, and shorter, focused assistance.
Codex	The ChatGPT mode built for software development: writing code, running tests, and working with Git repositories. This volume does not teach Codex. See the companion title, *ChatGPT Codex Visual Bible*.
Context	Information you give ChatGPT, through your message, an uploaded file, or an earlier part of the conversation, that shapes its response.
Custom GPT	A version of ChatGPT that you configure with your own instructions, files, and selected tools, built for one repeated purpose.
Custom instructions	Personal preferences and background, such as your role or preferred tone, that you save once in **Settings** so ChatGPT applies them automatically to every new conversation without you repeating them.
Data Controls	The settings area where you manage how your data is used and stored, including whether your conversations help improve OpenAI's models, and where you can export or delete your data.
Deep research	A ChatGPT capability that plans a research task, gathers sources, and produces a longer, cited report instead of a single quick answer.
Dictate	A way of entering text by speaking instead of typing. ChatGPT transcribes what you say into the message box; you are still writing text, just with your voice instead of your hands. Not the same as Voice.
Memory	Information ChatGPT can store between conversations. You can review, correct, or delete it at any time.
MCP (Model Context Protocol)	A standard way to connect ChatGPT to external tools and information, so different services can work with ChatGPT without a custom-built connection for each one.
Plugin	An installable bundle that can include skills, apps, and templates for a particular task or service.
Project	A container that groups related chats, files, instructions, and memory around one ongoing piece of work.
Prompt	The message you send ChatGPT: what you want it to do and what a good answer should include.

Scheduled task	A set of instructions that ChatGPT runs automatically at a time or interval you choose, without you starting the conversation yourself.
Skill	A reusable set of instructions, examples, and resources that ChatGPT can run to complete a specific, repeated task.
Temporary Chat	A conversation you turn on that ChatGPT does not save to your history and does not use to create memories.
Voice	A mode where you speak to ChatGPT and it replies in an ongoing spoken conversation, rather than only transcribing your words into text. Not the same as Dictate.
Work	The ChatGPT mode built for longer assignments: research, connected files and apps, and finished documents, spreadsheets, presentations, or reports.

What prerequisites do you need?

You need very little to start *Book 1*:

- A computer, tablet, or phone with a web browser and an internet connection. ChatGPT works the same way on Windows, macOS, Chromebook, Linux, iOS, and Android, so use whichever device you already have.
- A free OpenAI account. You can complete nearly all of *Book 1* on the Free plan.
- A ChatGPT Plus subscription, or higher, once you reach *Book 2*. Plus covers everything the rest of this volume teaches, including ChatGPT Work, scheduled tasks, connected apps, skills, and plugins. You only need to upgrade when you're ready to use those features.
- Familiarity with common documents, spreadsheets, presentations, and images. You don't need specific software. Sample files for every chapter are ready to download from this volume's GitHub repository, so your own files are optional, not required.

A few more things will help before you turn the page:

- *ChatGPT Visual Bible* covers ChatGPT's Chat and Work experience on the web, phone, and desktop. It does not teach software development. If your goal is to write code, or build websites or mobile apps, look for the companion title, *ChatGPT Codex Visual Bible*.
- Treat every ChatGPT response in this volume as a draft to review, not a finished answer to trust automatically. Each chapter shows you what to check. A wrong answer stated with confidence can still cost you time, money, or credibility, so this habit is worth building early and often.
- Keep a notebook, document, or note app open as you read. Several chapters ask you to try a short task, and you will want somewhere to keep what you produce for the project at the end of each book.

With that, you are ready to open *Book 1* and have your first useful conversation with ChatGPT.

Book 1

Ask ChatGPT

Book 1 is about moving from asking occasional questions to holding a useful conversation with AI as an advisor and mentor, doing real research, learning new skills, and organizing your work into projects.

It has twelve chapters, including an end-of-book project:

1.1 Understand What ChatGPT Can Do 12

1.2 Create Your Account and Choose a Plan 17

1.3 Use ChatGPT on the Web, Phone, and Desktop 22

1.4 Have Your First Useful Conversation by Typing, Talking, or Showing 29

1.5 Write, Refine, and Improve Prompts Through Conversation 34

1.6 Learn Something New and Educate Your Child 40

1.7 Plan Everyday Life and Compare Important Choices 46

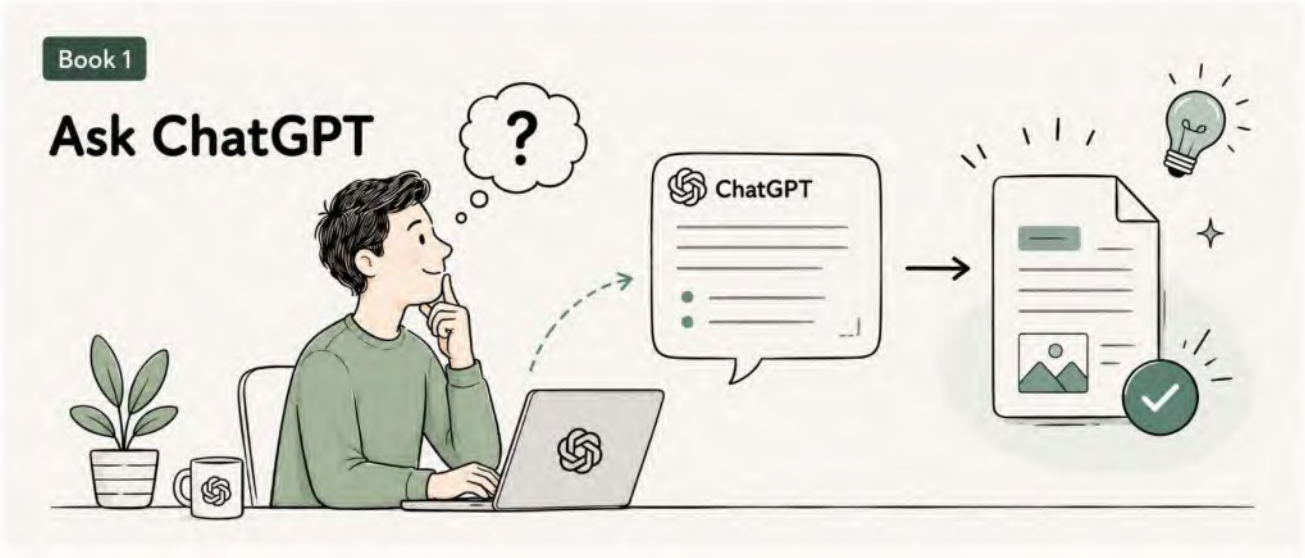
1.8 Search and Verify Current Information 49

1.9 Conduct Deeper Research with Cited Sources 52

1.10 Add Documents, Images, and Data to a Chat 55

1.11 Personalize and Organize with Instructions, Memory, and Projects 58

1.12 Project – Build Your Personal AI Workbench 63



Download the PDF version of this book

Your purchase includes a DRM-free PDF copy of this book, the code bundle, and a range of exclusive benefits. To unlock everything, follow the *Free benefits with your book* section in the front matter.

Chapter 1.1 Understand What ChatGPT Can Do

You can ask ChatGPT a question and receive an answer in seconds. That conversation is useful, but it is only one way to work with the product. This chapter will show you how a useful conversation can grow into research, a finished deliverable, and a repeatable work process.

From chatbot to work platform

The back-and-forth exchange of chat messages that you are already familiar with is still at the center of ChatGPT, hence its name, and the name of **Chat** mode. But ChatGPT can do much more. Depending on your plan, you can speak, show it a photo, upload a document like a spreadsheet, search the web, or connect information from another service.

ChatGPT can also help with work that lasts longer than one reply:

- **Deep research** gathers and combines information from several sources. You will learn about this in *Chapter 1.9 Conduct Deeper Research with Cited Sources*.
- **Projects** keep related chats, instructions, and files together. You will learn about this in *Chapter 1.11 Personalize and Organize with Instructions, Memory, and Projects*.
- **Work** mode can complete longer assignments and produce documents, slide decks, spreadsheets, and reports. You will learn about this in *Book 2*.
- **Scheduled tasks** can run once, repeat, or watch for a change. You will learn about this and other automation features in *Book 3*.

A work platform gives you several ways to reach an outcome. You choose the goal, supply useful context, set limits, and review what comes back. ChatGPT supplies speed and a wide base of general knowledge. You remain responsible for the purpose and the final decision.

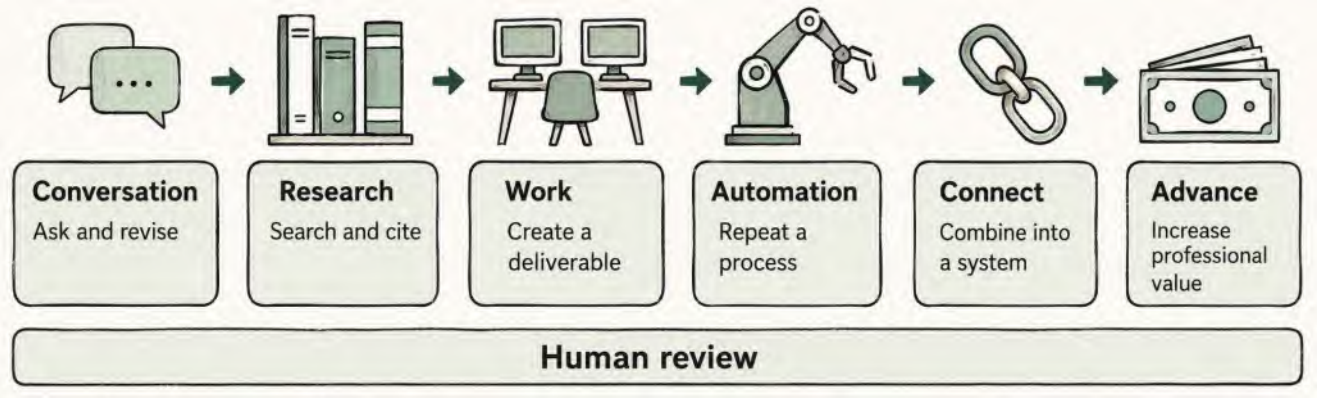


Figure 1.1A: A conversation can grow into checked research, finished work, a repeatable process, connected systems, and measurable value. Human review must span every stage

Tasks ChatGPT handles well

ChatGPT works best when you can describe the result and check the response. A suitable task has enough information, a clear audience, and a useful format. It also has consequences you can manage if the answer is incomplete or wrong:

Type of task	What ChatGPT can help you do	What you should check
Explore	Suggest choices, questions, risks, or another approach.	Assumptions, bias, and options it left out.
Research	Search sources, compare claims, and prepare a cited report.	Source quality, dates, quotations, and claims.
Learn	Explain a topic, give examples, or make a quiz.	Accuracy, missing points, and whether the explanation fits you.
Analyze	Find patterns in a table and create a chart or explanation.	The source data, calculations, labels, and scale.
Create	Draft a message, plan, image, document, or presentation.	Facts, fit for the audience, permissions, and final polish.
Transform	Summarize, rewrite, translate, classify, or extract details.	Lost meaning, tone, privacy, and copyright.

Table 1.1: Tasks that ChatGPT handles well



Learn more online: *ChatGPT Capabilities Overview*: `https://help.openai.com/en/articles/9260256-chatgpt-capabilities-overview`.



Good practice: The best results come from pairing AI help with a check you know how to perform. If you do not have the experience to judge the output, ask a trusted person or qualified professional to review it.

Community arts center example

Imagine a community arts center wants more people to attend its weekend workshops. A staff member could upload survey data and approved planning notes, as shown in *Figure 1.1B*, which also annotates some common **Chat** mode user interface elements:

1. Select **Close sidebar** to make more space for your chats.
2. Select **Share chat** or **More** to pop up a menu of more actions.
3. Files added to a chat, like Word documents and Excel spreadsheets, have their file types recognized automatically.
4. **Projects** and recent **Chats** are shown in the sidebar. Select to show and hide them.
5. Hover over **Chats** and select **...** to organize them.
6. Select the down-arrow button to quickly scroll to the bottom of the response.

7. Select the + button to add files to your chat prompt:

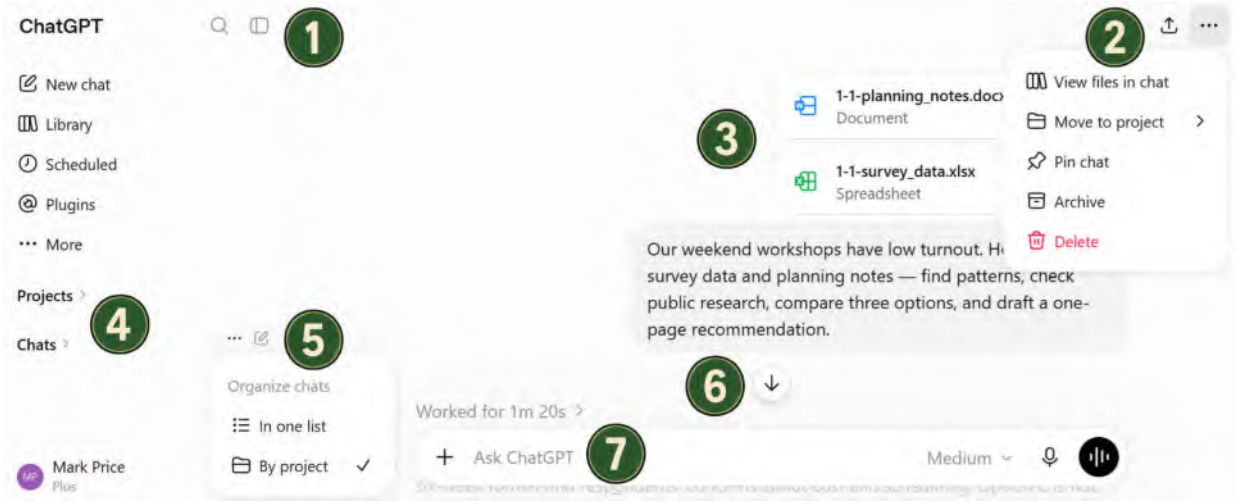


Figure 1.1B: Chat mode user interface elements

For this example prompt, ChatGPT will compare three options, find patterns in public data, and draft a one-page recommendation:

1. Select the ChatGPT logo to show the sidebar when it's been closed.
2. The response-generated document can be copied or downloaded, and the writing block editor can expand to make more space.
3. The document can be edited in place. Just click in the text and type.
4. The chat box has controls to change the thinking effort level, dictate a prompt, or activate **Voice** mode:

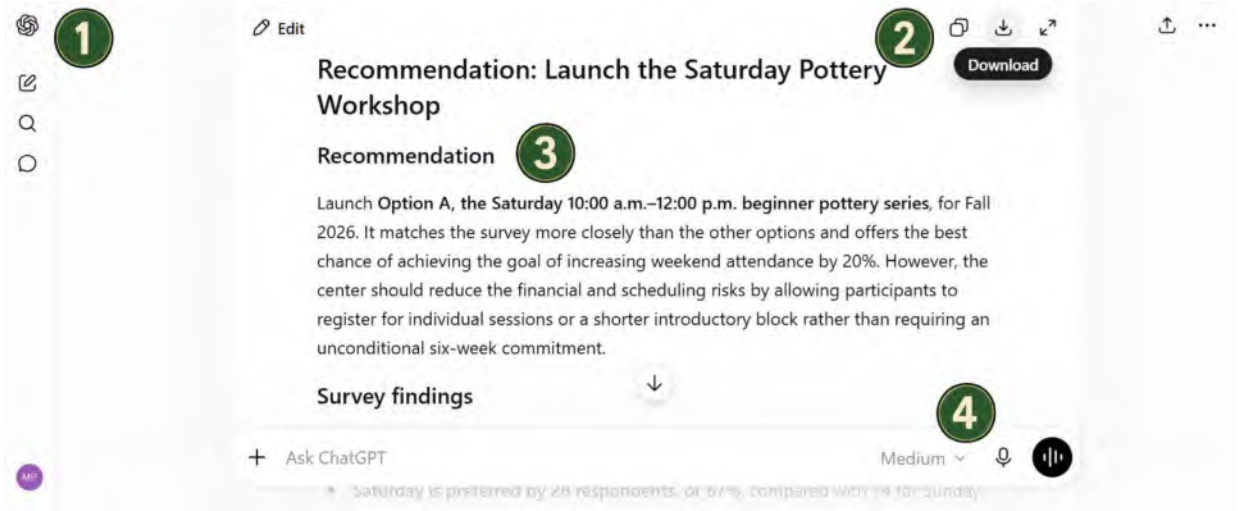


Figure 1.1C: The response includes a directly editable document that is ready for human review

The staff member would then check the data, open the sources, and decide which option fits the center's budget and goals.

Chat features

You can activate additional features by clicking the + button on the left side of the chat box:

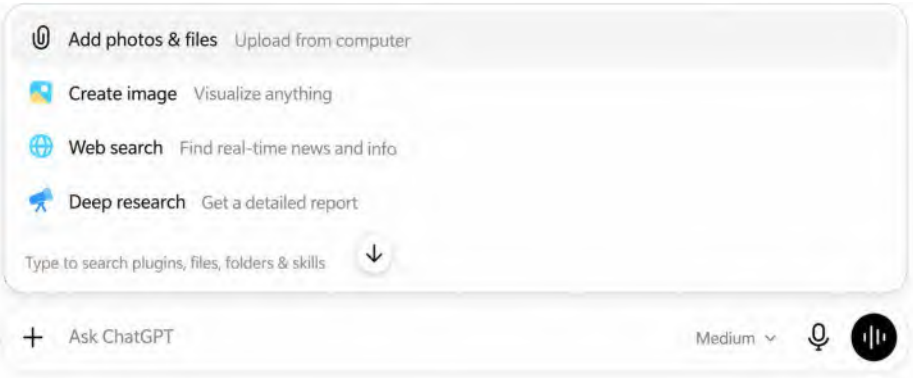


Figure 1.1D: Activating extra features in Chat mode.

For example, if you need the response to be as up-to-date as possible with the latest information, you could activate **Web search** so ChatGPT does not rely only on its model data, which will have a cut-off date in the past.

You will learn about these modes in the following chapters:

- **Add photos & files:** *Chapter 1.10 Add Documents, Images, and Data to a Chat*
- **Create image:** *Chapter 2.6 Create and Edit Images and Media*
- **Web search:** *Chapter 1.8 Search and Verify Current Information*
- **Deep research:** *Chapter 1.9 Conduct Deeper Research with Cited Sources*



Warning! The exact features you see in the menu depend on which plugins you have installed. For example, you might see choices like GitHub, Gmail, Figma, and Google Drive if you have installed those plugins.

Tasks that still require human judgment

ChatGPT can produce a polished answer that contains an error. It may misunderstand your goal, omit a useful option, misread a file, or present an uncertain claim with too much confidence. OpenAI tells users to check important information and review output before using or sharing it.

It is especially important to keep a person in control when a task involves any of these conditions:

- **Facts with consequences:** Check names, dates, numbers, quotations, citations, and calculations against dependable sources.
- **Private or restricted information:** Do not upload passwords, private keys, health records, customer data, or confidential work unless you have permission and an approved account and process.
- **Professional advice:** Ask a qualified person to review medical, legal, financial, employment, or safety guidance before anyone acts on it.
- **Decisions about people:** Do not let ChatGPT make the final choice about hiring, education, housing, credit, insurance, health care, or access to an essential service.
- **Actions that are hard to undo:** Review before a tool sends, publishes, buys, deletes, changes access, or contacts another person.

- **Values and relationships:** ChatGPT can help you prepare for a choice or conversation. It cannot decide what you owe another person or which trade-off you should accept.

Before you use an AI response, ask three questions:

1. **Is it accurate?** Check the facts and make sure the response answers the task you set.
2. **Am I allowed to use it?** Check privacy, consent, copyright, workplace rules, and access rights.
3. **What happens if it is wrong?** Add more review when the cost, harm, or loss would be hard to repair.



Watch out: Do not confuse a neat format with a correct result. A table, citation, or professional tone can make a weak answer look dependable. Open the sources and check the underlying facts.

From research to professional value

A long response does not create value by itself. Value appears when the work helps someone understand a problem, make a sound choice, create a useful result, or complete a task with less wasted effort. You need evidence of that outcome.

Let's return to the arts center example. A short chat might produce workshop ideas. Web search could add current evidence from the Internet. Adding a file for analysis could show which past workshops attracted repeat visitors. A verified recommendation could help staff choose one change. A reusable project could hold the source files and instructions for the next review. A scheduled task could check enrollment each week.

The final step is to measure what changed. The center might track attendance, repeat bookings, staff time, or cost per booking. A professional using ChatGPT could track turnaround time, errors found during review, customer response, or work completed. These measures show whether the process helped. They also show where it needs refinement.

Commercial value follows the same rule. A client or employer pays for a dependable result, not for the number of prompts you used. You earn trust by understanding the problem, protecting information, checking the work, and accepting responsibility for the deliverable.



Good practice: Start with a small, low-risk task that you already understand. Compare the time and quality with your usual method. Keep the parts that help, and change or remove the parts that do not.

Try it now

In this activity, ChatGPT will help you decide where it can assist and where you should review. You'll work through this using one recurring task from your own work or personal life:

1. Choose one recurring task from your work or personal life. Pick a task that does not require private information and will not send, publish, buy, or delete anything.
2. Write down the name of your task in your notes. Examples include planning a weekly meal, preparing meeting notes, studying a topic, or comparing routine purchases.
3. Remove names, account details, confidential facts, and other information that you should not share.

4. Open ChatGPT and paste the following prompt (if ChatGPT asks you to sign in, save the prompt and return to it after you have read *Chapter 1.2*, where you will learn how to register an account):

   ```
   I want to improve how I handle [recurring task]. Ask me up to five short
   questions. Then create a table with these columns: current step, how ChatGPT
   could help, what I must review, information I should not share, and one small
   test. Use only the facts I provide. Label every assumption. The first test must
   take less than 15 minutes and must not send, publish, buy, delete, or change
   access to anything.
   ```

5. Answer the questions without adding private or restricted information.
6. Read the table. Correct any false assumptions and ask ChatGPT to revise the plan.



What to ask next: `Choose the safest first test. Give me a short checklist, including what I should check before I use the result.`

Check the result

- ☐ Does the plan separate ChatGPT's help from the choice or approval that remains yours?
- ☐ Did ChatGPT label its assumptions and identify information that you should not share?
- ☐ Can you test one part in less than 15 minutes without creating a hard-to-reverse consequence?

If any answer is no, tell ChatGPT what is missing and request a new version. Revision is a normal part of working with an AI assistant.

If you have not created a ChatGPT account yet, *Chapter 1.2* walks through that step and helps you choose a plan that fits the kind of task you just tried.

Chapter 1.2 Create Your Account and Choose a Plan

A free ChatGPT account gives you saved chats, personal settings, and access on several devices. A paid plan can provide higher usage allowances and more access to some tools, but you do not need one to begin.

This chapter shows you how to create and protect an account, compare plans, recognize usage limits, and choose privacy settings you understand. Start free, then upgrade only when you can name the repeated need that a paid plan will solve.

Choose a sign-in method that you expect to keep. If an employer or school gives you a managed account, follow its rules and keep personal activity in a separate personal account. You will not buy a plan in this chapter.

Creating and securing an account

You can try ChatGPT without signing in, but an account keeps your chats and settings together. It also lets you return to a chat on another device. Create a personal account for personal use, or use the managed account supplied by your organization for approved work.

The available sign-in choices can vary. You may be able to use an email address, phone number, or an existing account from another provider like Google that most people already have, so you don't have to remember another password. Use a method that only you control and remember which choice you made. You will need the same method when you sign in later.

Create the account before you change its security settings:

1. Open your browser and enter `chatgpt.com` in the address bar. Check the website address spelling before you enter any account details.
2. Select **Sign up for free**. If you already have an account, select **Log in** and skip to step 5:



Figure 1.2A: Sign up for free to create your account on the official ChatGPT site

3. Choose one of the sign-in methods shown on your screen. Complete any identity or age checks that ChatGPT requests.
4. Return to ChatGPT after verification. You should now see a new chat and your account menu in the bottom-left corner of the ChatGPT page.
5. Open your account menu by clicking on your name and plan. (For me, that is **Mark Price** and **Plus** because I pay about US$20 per month for the Plus plan.)
6. In the menu, select **Settings**:

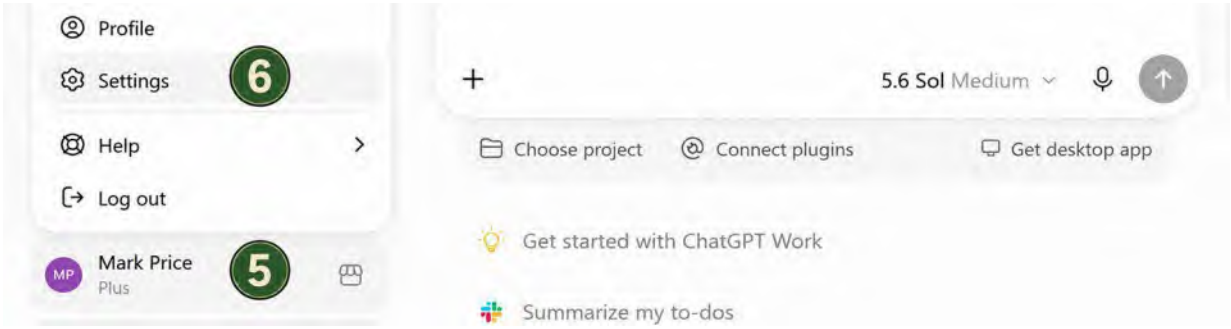


Figure 1.2B: Selecting your account menu and then Settings

7. In the **Settings** window, select **Security and login**.
8. Under **Multi-factor authentication**, enable an available second sign-in method. If **Security keys & passkeys** is an option, you can select **Add** and follow the instructions for your device. If neither option appears, continue and check again later:

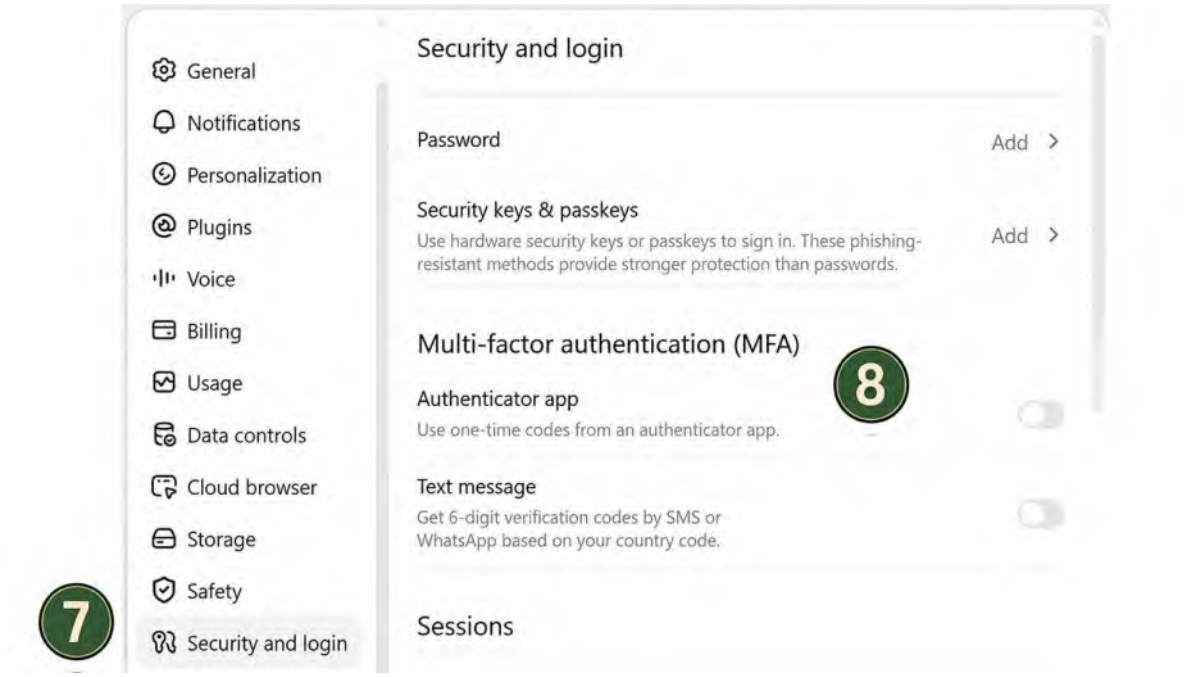


Figure 1.2C: Add another sign-in check from Security and login settings. Available sign-in and security methods can vary

Good practice: Use a password manager if your sign-in method needs a password. Never share your password, verification code, passkey, or account with another person.

Set the interface language

Separately from the language ChatGPT replies in, you can change the language used for its own menus, buttons, and settings screens. Open your account menu, select **Settings**, then **General**, then **Language**. Choose **Auto-detect** to match your browser or device language, or pick a specific language from the list.

Watch out: This setting changes only the interface text. It does not change the language ChatGPT uses to reply to you; that follows the language of your own messages unless you tell it otherwise.

Understanding free and paid plans

OpenAI currently offers Free, Go, Plus, and Pro plans for individuals. **Business** and **Enterprise** plans are intended for organizations. Names, prices, features, and country availability can change, so check the current pricing page before paying.

Figure 1.2D shows how the **Free** plan will encourage you to **Upgrade**:

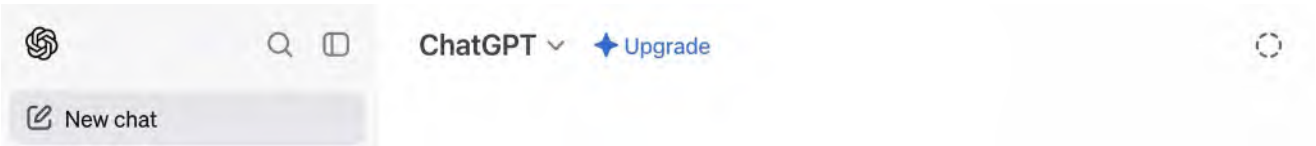


Figure 1.2D: The Free plan user experience

OpenAI shows ads to some Free-plan users in some regions.

Figure 1.2E shows how the **Plus** plan has extra features like **Work** as well as just **Chat**:



Figure 1.2E: The Plus plan user experience

The **Free** plan supports common chat tasks and includes limited access to tools such as web search, file uploads, data analysis, image creation, and deeper research. An individual paid plan like **Plus** adds higher allowances, advanced models, and broader access to selected tools. Pro is aimed at people who use those capabilities even more often:

Plan	A suitable starting point	Question to ask
Free	You are learning, testing tasks, or using ChatGPT from time to time.	Can I finish my usual task before a limit interrupts it?
Go or **Plus**	You need higher allowances or a named feature for recurring personal work.	Which feature or allowance will I use each week?
Pro	You use ChatGPT often and regularly reach the limits of another individual plan.	Will the time saved justify the higher monthly cost?
Business or **Enterprise**	Your organization needs managed accounts, shared controls, or business data protections.	Which plan and settings does my organization approve?

Table 1.2: Understanding free and paid plans

A paid plan does not make every answer correct. It also does not guarantee unlimited use of every model or tool. You still need to check facts, sources, files, calculations, and any output you plan to share.

Use a managed plan for work when your organization requires it. **Business** and **Enterprise** plans have different data protections and administrator controls. Ask who can access conversations, how long data is kept, which tools are enabled, and what information you may upload.



Good practice: Stay on the **Free** plan while you learn. Consider an upgrade to a paid plan after a limit or missing feature interrupts the same useful task more than once. Individual plans can be paid monthly, or annually at a small discount. The main benefit of monthly is that you can stop at any time instead of being committed for a full year.

Understanding usage limits and feature availability

A usage limit is a cap on how much of a model or tool you can use during a period. The **Free** plan has limits, and paid plans can have them too.

Some tools have separate limits. For example, file uploads, data analysis, image creation, and deeper research may not share one allowance. Reaching an image limit does not always mean that ordinary chat has stopped. Read the message on the screen to learn which capability is affected.

When ChatGPT reports a limit, use this short check:

1. **Read the whole notice:** Look for the affected model or tool, the reset time, and any available alternative.
2. **Save your work:** If the limit will reset soon, continue later instead of buying a plan in a hurry.
3. **Record what the interruption prevented you from completing:** A repeated interruption gives you evidence for a plan decision.
4. **Open the current pricing page before upgrading:** Check the exact feature, allowance, billing period, tax, and cancellation terms shown for your country.



Learn more online: A tool shown in this book may appear in a different place, have a new name, or require another plan by the time you read it. Check the plan page and the notice inside your account: `https://chatgpt.com/pricing/`.

Reviewing privacy and data controls

Privacy has several parts. Chat history, model improvement, memory, account administration, and data retention are separate controls. Changing one setting does not change all the others:

1. **Model improvement.** In a personal account, open **Settings**, select **Data Controls**, and review **Improve the model for everyone.** Turning it off keeps new conversations in your history but prevents them from being used to improve OpenAI's models.
2. **Export and deletion. Data Controls** also lets a signed-in user request a data export or permanently delete all chats. Export information you want to keep before deleting anything.

Other privacy and data controls include:

- **Temporary Chat.** Turn on a temporary chat when you do not want the conversation saved in history or used to create memories. OpenAI says temporary chats are not used for training and are deleted from its systems after 30 days, although they may be reviewed for abuse:



Figure 1.2F: Button to turn on a temporary chat

- **Managed accounts.** An organization administrator may be able to access, export, audit, retain, or delete data in a managed account. The administrator can also restrict features and control retention settings.
- **Business data.** OpenAI says it does not use content from **Business** and **Enterprise** offerings to improve its models by default. Your organization's own rules and administrator settings still apply.

No setting makes it suitable to paste passwords, private keys, confidential files, health records, or personal data that you have no right to share. Remove restricted information before you start and use an approved account for work.



Watch out: A managed work account is not private from the organization that controls it. Check that you are in the intended account before you begin a personal or work conversation.



Learn more online: *Data Controls FAQ*: `https://help.openai.com/en/articles/7730893-data-controls-faq`.

Try it now

Create or review a free ChatGPT account, add a second sign-in check, and record your plan and Data Controls privacy choices. This activity should take 5 to 15 minutes. It does not require a purchase.

Open the ChatGPT pricing page: `https://chatgpt.com/pricing/` and compare **Free** with the next individual plan available in your country. Do not purchase it during this activity.

Check the result

- ☐ Can you sign in again using a method you expect to keep, with another verification method enabled or its absence recorded?
- ☐ Did you compare the Free plan with the next available plan for your country without purchasing anything, and can you name one feature or allowance that would justify upgrading later?
- ☐ Did you review your Data Controls, and can you say what you chose for model improvement, and whether you know how to export or delete your data if needed?

With your account created and secured, *Chapter 1.3* shows how to bring it with you: into a browser, onto your phone, and onto a desktop computer.

Chapter 1.3 Use ChatGPT on the Web, Phone, and Desktop

ChatGPT works in a web browser, on iPhone and Android phones, and in desktop apps for Mac and Windows. You can begin in one place and continue conversations in another, as long as you sign in to the same account.

This chapter helps you choose the version that suits the task in front of you. You will learn how to reach the official web and app downloads, how to check which account is active, and how to test a harmless conversation across two devices. You need one ChatGPT account. A second device is optional. Use the version that fits the task but check the active account before you write or attach anything.

Start with ChatGPT in a web browser

The web version is the best starting point for most readers. It works without a separate installation and runs on any device with a modern browser, including locked-down work laptops, Chromebooks, and computers where you cannot install software. It gives you room to read a long answer, compare sources, and work with a document on a larger screen. It is typically the first place new features and model updates appear before they reach the desktop or mobile apps.

Open a browser and enter `chatgpt.com` in the address bar. Then sign in with the same method you chose in *Chapter 1.2*. If you are using a shared computer, sign out when you finish and do not ask the browser to save your password.

A browser also gives you a useful safety check. Look at the address before you enter a password or upload a file. A search result, an advertisement, or a copied link can lead to a site that looks familiar but is not the official ChatGPT service.

You can pin `chatgpt.com` in your browser or add it to your bookmarks. That makes it easier to return to the correct site without searching for it each time:

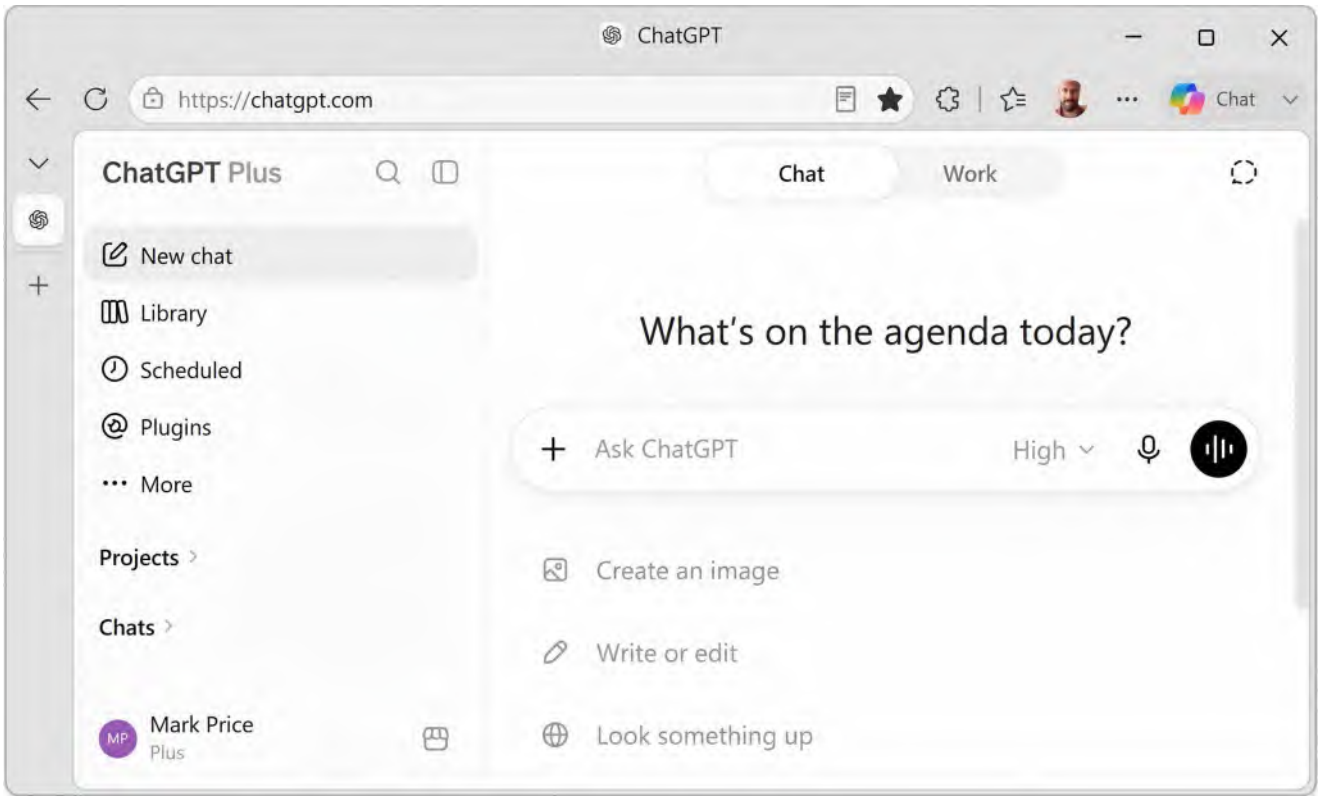


Figure 1.3A: Use the web version anywhere with a browser



Watch out: In *Book 1*, we will only use **Chat** mode features. You will learn about **Work** mode in *Book 2*. Make sure you are in **Chat** mode for all the chapters in *Book 1*.

Use ChatGPT on your phone

The official ChatGPT app is available for iPhone and Android. The mobile app's strength is being useful away from a keyboard. Voice mode adds live camera and screen sharing, so you can point your phone at something and talk through it in near real time, whether that's practicing a language, getting help with a physical object, or narrating a document. The camera makes it easy to capture whiteboards, receipts, or printed pages on the spot. But live camera is currently only available on a paid plan.

This makes the mobile app best suited for quick, hands-free help while you're out and about rather than sitting at a computer:

1. Select the menu button in the top-left corner to toggle the sidebar.
2. The sidebar contains similar features to the web version. **Recents** shows your chats.
3. The chat box contains similar features to the web version:

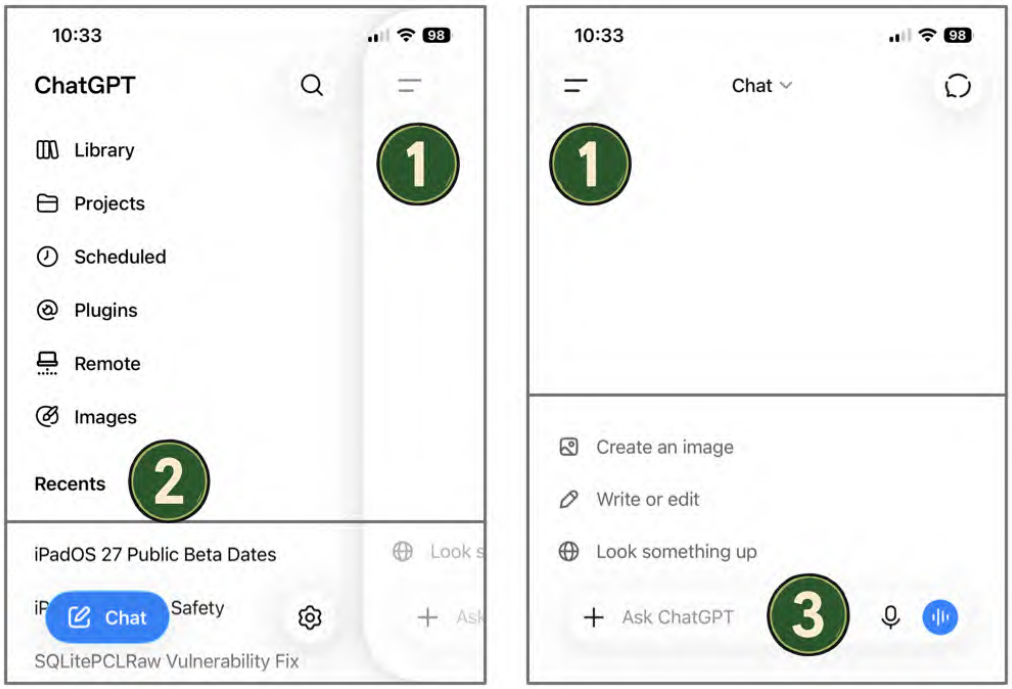


Figure 1.3B: Use the phone version on iPhone or Android devices

Download the iPhone app from the Apple App Store or the Android app from Google Play. Search for **OpenAI ChatGPT** and check that **OpenAI** is the publisher before you install it. Country availability and device requirements can change, so use the store page if the app does not appear for you.

After installation, sign in with the same method as on the web. If you created the account with Google, Apple, Microsoft, an email address, or a phone number, choose that same route. Using another sign-in method can create a separate account with a different chat history and plan.

The phone app may ask for permission to use the microphone, camera, photos, or notifications. Grant a permission only when you know why the feature needs it. For example, Dictate and Voice need microphone access, and asking about a photo needs access to the photo you select.



Watch out: A phone screen can hide context that would be obvious on a computer. Before you send a message, check whether you are in a personal account or a managed work account. Then reread any attachment name and the text above the send button. Voice and photo features can capture more than you intend. Move to a private place before speaking about personal or work information. Crop or review a photo before you upload it, and remove names, addresses, account numbers, and other restricted details.

Use ChatGPT on your desktop

The desktop app is intended for longer sessions at a Mac or Windows computer. It can keep ChatGPT available while you work in a document, spreadsheet, browser, or other application. The exact features can differ from the web version and can change as OpenAI updates the app. The desktop app's advantage is system-level integration with whatever else you are working on:

1. The sidebar contains similar features to the web version. **Sites** is a ChatGPT Codex feature for creating websites.
2. The blue down arrow indicates an update is available. Select it to download and install it.
3. Select to show or hide the integrated web browser in a right window pane.

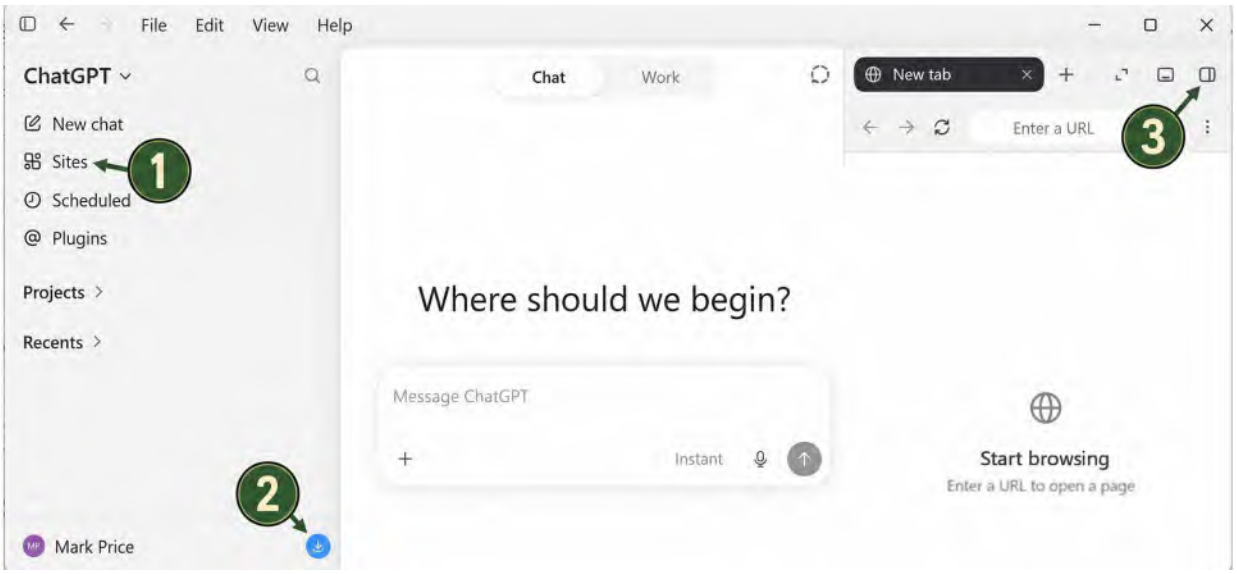


Figure 1.3C: Use the desktop app on Windows or Mac. You can toggle a browser in the side panel.

Download the current desktop app from `https://chatgpt.com/download`, not from a search advertisement or a file-sharing site. Open the installer, follow the instructions for your operating system, and sign in to the same ChatGPT account that you use on the web.

The current desktop app brings together ChatGPT Chat, Work, and Codex. *Book 1* starts with ordinary ChatGPT conversations. If you see Work or Codex, you can leave those areas alone for now.



Watch out: Desktop features can request access to local files, other apps, a browser, or a work account. Read each permission screen. Do not give ChatGPT access to a folder, service, or application because it might be useful later. Choose only the material needed for the task you are doing now.

If you use an employer's computer, ask whether the desktop app is approved before you install it. Your organization may manage app installation, restrict features, or require a work account for business information.



Watch out: Do not download an app from an unfamiliar website or install one that is not published by OpenAI.



Learn more online: *Moving to the new ChatGPT desktop app*: `https://help.openai.com/en/articles/20001276-moving-to-the-new-chatgpt-desktop-app`.

Open ChatGPT instantly with a keyboard shortcut

The desktop app can pop up a small companion window over whatever else you're doing, without switching to a browser tab or bringing the full ChatGPT window forward. This is useful for a quick question while you're writing an email, reading a webpage, or working in another program. macOS has the same feature under its own **Settings** menu, with Appshots adding further context there, as mentioned above.

The shortcut is off by default. Turn it on in **Settings** | **General** | **Popout Window hotkey**, then select the pencil button and press a key combination to set the key combination:

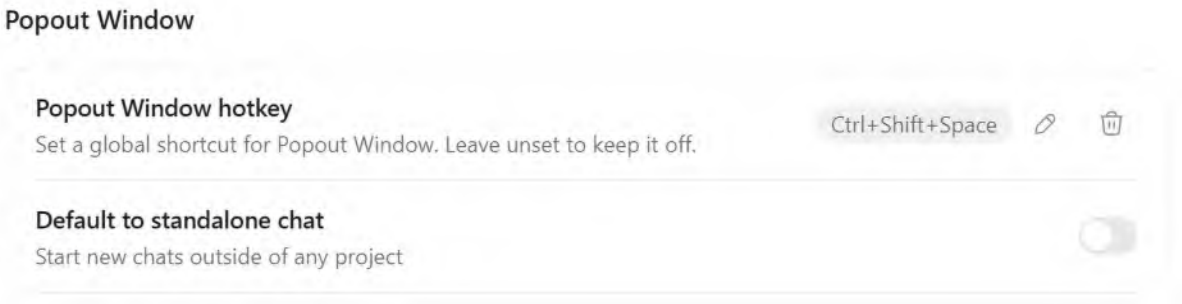


Figure 1.3D: Setting a Ctrl+Shift+Space key combination for the ChatGPT popout window

Once it's on, press the shortcut from inside any other application to bring up the popup window:

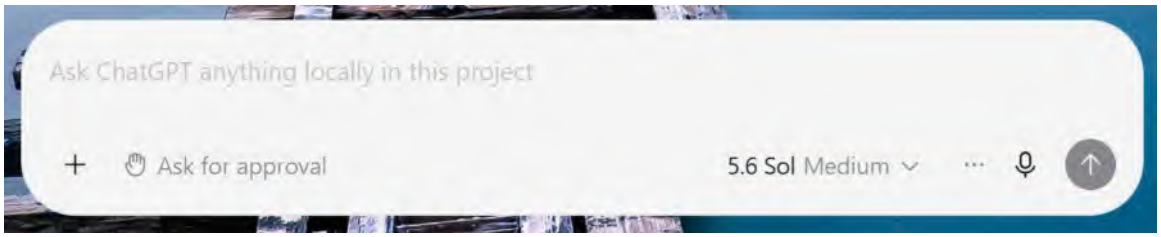


Figure 1.3E: ChatGPT Popout Window above the Windows desktop

Type a question, upload a file, or start a new conversation without leaving what you were doing. Close it and reopen it with the same shortcut. It remembers its last position on screen. If the shortcut doesn't respond, another Windows application may already use the same key combination. Choose a different combination in the same settings screen rather than assuming the feature is broken.



Good practice: Pick a combination you don't already use elsewhere. *Ctrl+Alt+Space* is a common choice because it avoids Windows' own reserved shortcuts. For Mac, try *Option+Space*, since it doesn't collide with Spotlight's built-in *Command+Space*, which in macOS 27 invokes Siri AI.

Continue a conversation on another device

Signed-in ChatGPT conversations can usually appear on your other devices. That makes it practical to start a question in a browser, read the reply on a phone, and return to a computer when you need more space.

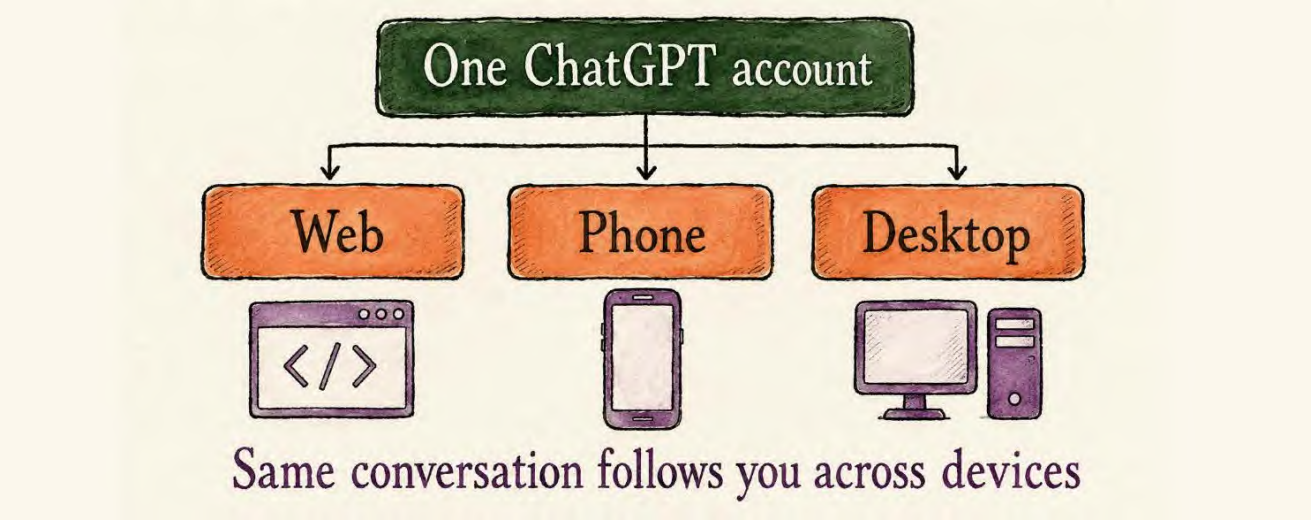


Figure 1.3F: Use the same account to continue cloud work across devices, while keeping temporary and managed-account conversations under the right controls

Use one conversation for one subject. If you switch from planning a vacation to writing a work memo, start a new chat. A focused conversation gives ChatGPT less unrelated text to misunderstand and makes it easier for you to find the exchange later.

When a chat does not appear on another device, first confirm that both devices use the same account and workspace. Then check your internet connection, close and reopen the app or browser, and look for the conversation in chat history. Do not paste the same private material into a second chat until you know what happened to the first one.

Some activity stays tied to a particular device. For example, a task that uses local files or another installed app may need the original computer and your permission. Treat a new device as a new place to review the task before you continue it.

My recommendations for when to use each type of device and what to check before you begin are summarized in *Table 1.3*:

Device	Choose it when	Check before you begin
Web	You want a familiar browser window, a larger screen, or no installation.	The address is `chatgpt.com` and you are in the intended account.
Phone	You want to ask by voice, capture a photo, or continue a short task away from your desk.	The app is published by **OpenAI**, and the message contains no sensitive details.
Desktop	You work for long periods at a computer and want ChatGPT close to your files and other apps.	Your computer and workplace rules allow the app and any permission it requests.

Table 1.3: Use ChatGPT on the web, phone, and desktop

Try it now

In this activity, you will sign in on the web and, if you have one, a second device. Then you'll confirm that a short test conversation appears in both places. This activity should take 5 to 15 minutes:

1. On a computer or phone browser, open `chatgpt.com` and sign in to the account you chose in *Chapter 1.2*.
2. If you have a second device, install the official ChatGPT app from the Apple App Store, Google Play, or `chatgpt.com/download`. Check that OpenAI publishes the app or download page.
3. Sign in on the second device using the same sign-in method. Check the account name or workspace before you continue.
4. On the first device, start a new chat and send the following prompt:

   ```
   Reply with the words: This is my cross-device test. Do not search the web or
   use another tool.
   ```

5. Open ChatGPT on the second device. Find the test chat in the chat list or search for the text: `cross-device test`. Read the reply, then add this follow-up:

   ```
   I can see this conversation on another device.
   ```

6. Return to the first device and check for the follow-up. If it does not appear, confirm the account and workspace on both devices before you try again.



What to ask next: `Help me choose whether I should use ChatGPT on the web, phone, or desktop for these tasks: [list up to three tasks]. For each one, name the best starting place, one privacy check, and one reason I might switch devices.`

Check the result

- ☐ Can you reach the official ChatGPT website and identify an official mobile or desktop download source?
- ☐ Can you tell which account or workspace is active before you start a conversation?
- ☐ Did the harmless test chat appear on both devices, or did you record the point where it did not?
- ☐ Can you name one task that suits the web, phone, or desktop version and one privacy check for that task?

If an answer is no, use the web version as your starting point and check the current Help Center for your device. If you use a managed account, ask your organization administrator before you install software or change access settings.

Once you know where to sign in, *Chapter 1.4* turns to what you do once you're there: having a first conversation you can type, speak, or show.

Chapter 1.4 Have Your First Useful Conversation by Typing, Talking, or Showing

A useful ChatGPT conversation has a goal that you can describe and an answer that you can improve. You can type a message, speak through Voice, or show ChatGPT a photo or screenshot. The input method changes how you give information. The same conversation habits still apply.

In this chapter, you will make one small request, answer follow-up questions, correct a wrong assumption, and ask for another option. You will also learn how to keep a conversation that you may want to return to or share after review.

Languages ChatGPT can understand and reply in

You don't need to change any setting to talk with ChatGPT in your own language. Write or speak in the language you're most comfortable with, and ChatGPT replies in that same language by default. OpenAI states that ChatGPT natively supports around 58 languages, covering roughly 4.5 billion native speakers, and it can often work reasonably well beyond that list too. Quality is strongest in widely used languages such as English and Spanish, so check anything important, such as a translation or a technical term, the same way you'd check any other answer.



Good practice: If you want ChatGPT to answer in a specific language regardless of the language you type in, say so directly, for example, "Reply in French even if I write to you in English."

Choose the best way to start

Type or dictate when the wording or details need care. A typed or dictated request is useful for a question with names, dates, limits, or a format that you want to keep. You can reread it before you send it and edit it later.

Use Voice to talk to ChatGPT interactively. This will help you explain an idea or ask short follow-up questions. Select the Voice control only if it appears in your account, then allow microphone access when you are ready. Voice transcripts can differ from what you said, especially with background sound or overlapping speech, so review the transcribed text before you rely on it.



Watch out: Do not confuse the Dictate feature (the microphone button) with the Voice feature (the waveform button). Dictate mode is an alternative way of typing. You are still entering text. Also note that Voice is only available in **Chat** mode, not **Work** mode.

Show ChatGPT an image when the picture contains information that words would describe poorly. For example, you might show a non-identifying photo of a household object and ask what features it can see. Crop the image first and remove names, addresses, faces, account numbers, and other information that you should not share.

To send that image, or anything else you want to ask, use the controls in the **Ask ChatGPT** box:

1. The + button lets you add a file like an image or photo and activate features like deep research.
2. Select the **Ask ChatGPT** box to type text.

3. Select the microphone button to start dictating text. Select it again to stop.
4. Select the waveform button to start Voice and then speak to ChatGPT naturally and more interactively. Select it again to stop:

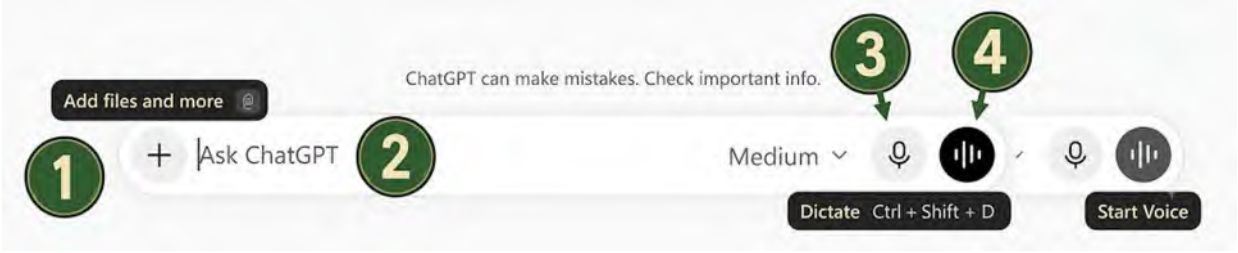


Figure 1.4A: Three inputs: add a file like a photo you've taken, type or dictate some text, and interact with Voice mode



Watch out: A photo can reveal more than the subject you meant to show. Check the background, labels, reflections, and file name before you upload it. Voice can also capture people nearby.

Select a thinking effort level

Once you know how you will send your message, you can also choose how much thinking effort ChatGPT puts into answering it:

1. The **Thinking effort** control in the chat box shows your current choice, such as **Medium**.
2. Select it to open a slider, where you trade speed for depth, or depth for speed, before you send a request. The Plus plan offers three thinking effort levels: **Instant**, **Medium**, and **High**. The Pro plan adds two more: **Extra High** and **Max**.
3. Select **Advanced** to select a specific **Model** as well as **Effort** level:

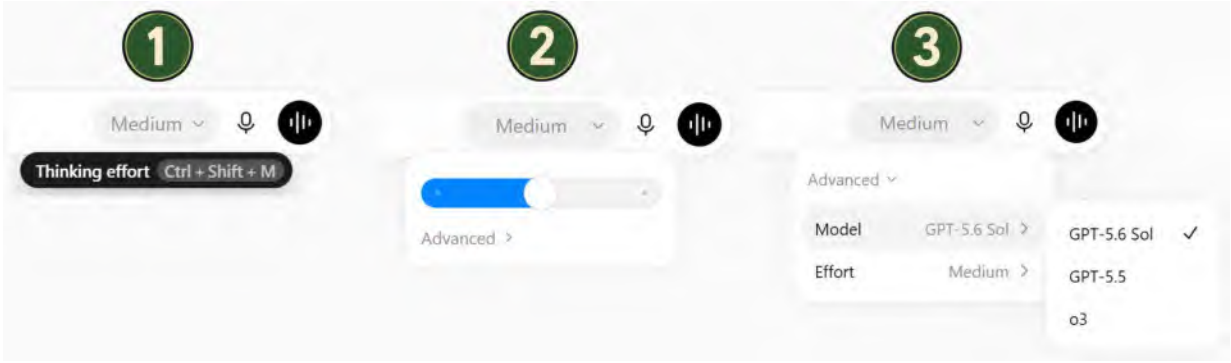


Figure 1.4B: Selecting a model and thinking effort level

Each thinking level balances speed against how much the model reasons before it replies:

Option	What it does	Model behind it
Instant	Fast answers for everyday questions. This is the default.	GPT-5.6 Sol (Plus, Pro) or GPT-5.6 Luna (Free, Go)
Medium	Standard reasoning for tasks that need some thought.	GPT-5.6 Sol

High	Extended reasoning for harder tasks, such as planning or debugging.	GPT-5.6 Sol
Extra High	The most reasoning effort available in ordinary chat.	GPT-5.6 Sol
Pro	The highest-capability option, for difficult or long-running work.	GPT-5.6 Sol Pro

Table 1.4: Choices in the thinking effort level control



Learn more online: OpenAI's official *Models* article covers which thinking effort levels are available on each plan, and updates as the rollout continues: `https://learn.chatgpt.com/docs/models`.

Which of these options you see depends on your plan. As of August 2026, Free and Go plans default to GPT-5.6 Luna with unlimited everyday text chats, plus a separate Think option for questions that need deeper reasoning. Plus and Pro plans can also use a slider to choose how much thought Sol puts into an answer.

OpenAI rolls out new model families across accounts over time, so check your own thinking effort level for the exact list your plan currently offers. (ChatGPT Work has even more options, as you will see in *Book 2*.) ChatGPT can also switch itself from Instant to Medium when a request looks complex enough to need more reasoning.



Watch out: A higher thinking effort level does not guarantee a correct answer. It only gives ChatGPT more brain power to work through the problem. But over-thinking is a genuine problem for AI just as it is for humans. Keep checking facts, sources, and calculations at every thinking effort level, the same way you would with Instant.

Below the main options, a row named for the current model, such as **GPT-5.6 Sol**, expands into a further list of specific model versions, for example GPT-5.6 Sol, GPT-5.5, GPT-5.3, and o3. Most people can leave this on the newest model. Switch only if you have a specific reason to keep using an older version, since OpenAI adds and retires versions in this list over time.

OpenAI introduced Sol as part of a family called GPT-5.6, alongside Terra, a balanced model for everyday work, and Luna, the fastest and least expensive of the three. As of August 2026, Chat mode uses Sol for Plus and Pro users, at the thinking effort levels described previously, and Luna as the default model for Free and Go users. Terra still appears in ChatGPT Work and in ChatGPT Codex, OpenAI's coding tool. *Book 2*, *Create Real Work Results* covers ChatGPT Work and its model picker. The companion book, *ChatGPT Codex Visual Bible*, covers ChatGPT Codex and its model picker.



Good practice: Match the thinking effort level to the task. Use Instant for a quick fact or a short everyday message. Move up to Medium or High when a task involves several steps, code, or a decision worth checking carefully. Save Extra High or Pro for work that is actually difficult or long, since these levels take longer to answer and use more of your plan's usage allowance.

Ask for a specific result

Start with the result you want, who it is for, and any limit that affects the answer. Then ask for a format that helps you compare or use the response. You do not need a long instruction. The first request gives ChatGPT enough information to make a useful start.

For example, this request names the task, the limit, the result, and a check before a purchase:

```
I am choosing a gift for a neighbor who enjoys gardening. Ask me up to three
questions. Then suggest three gifts that cost no more than [amount and currency].
Give one reason for each choice and name anything I should check before buying. Do
not make a purchase.
```

The request does not assume that ChatGPT knows the neighbor, local stores, or your budget. It asks for questions first. That gives you a chance to add information before ChatGPT makes suggestions. This example illustrates two things about prompting well:

1. Asking a broad question will not give you a great response
2. Providing some detail starts a conversation that can get you what you need

Answer questions, correct assumptions, and fill gaps

ChatGPT may ask a question before it answers, or it may make an assumption that does not fit. Treat either response as a chance to guide the work. Answer the useful questions. State what is wrong. Add the missing fact. Then ask for a revised answer.

For the gift example, you might say:

```
The neighbor has a small balcony, not a garden. They already own basic tools. Revise
the suggestions. Keep the budget below [amount and currency], and tell me which
option needs no special equipment.
```

A correction does not need to restate the whole task. It only needs to change the part that was wrong. If ChatGPT still guesses, ask it to label each assumption and list the information it needs from you.



Good practice: When a response contains a fact that could affect money, health, safety, work, or another person, check the fact before acting. A useful conversation can still contain an error.

Request alternatives and make your choice

One answer may be suitable, but it is rarely the only possible answer. Ask for alternatives when you need to compare cost, time, risk, style, or another tradeoff. A short table makes the differences easier to inspect.

For example, you could ask:

```
Give me three different options. Put them in a table with these columns: option,
estimated cost, why it fits, possible drawback, and what I need to check. Do not
invent facts that I have not supplied.
```

Read the result as a set of ideas, not as a final decision. You choose the option, check the details, and decide whether the answer is ready to use. If none of the options work, say why and ask for another set.

Requesting alternatives and making your choice follows a five-step loop:

1. Name the result
2. Answer questions
3. Correct assumptions
4. Compare alternatives
5. Review or keep useful work

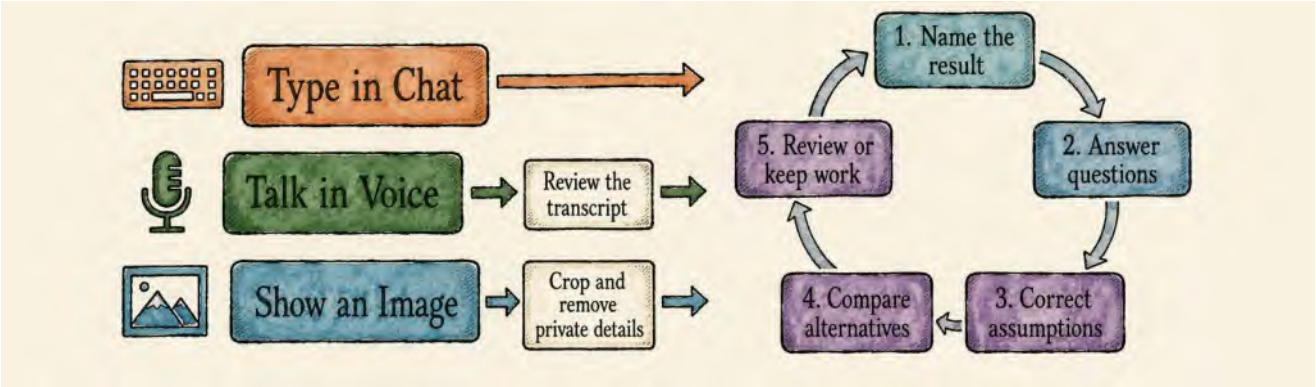


Figure 1.4C: The input can change. The improvement loop stays the same.

For voice input, always review the transcript to ensure that ChatGPT correctly understood your request. For image input, crop and remove any private details before sending the image to ChatGPT.

Human review of the response and comparison with alternatives before using it is always good practice.

Keep useful work without sharing too much

A useful conversation is easier to find when its title describes the work. ChatGPT may give a chat an automatic title. Open the chat options and choose **Rename** if that control appears, then use a title such as "Balcony gift ideas, July 2026." The exact location of chat options can change across devices.

Copy a finished result into your own document when you need to keep it with other work. Include the sources, choices, and checks that support the final decision. A chat is a working record. Your document can become the version that you edit, approve, or share.

ChatGPT can create a shared link for a conversation in supported accounts. Anyone who has the link can view the conversation and pass it on. Read the full conversation before you share it. Remove private, confidential, or unnecessary material. Use a managed work account and your organization's rules for work conversations.

Try it now

Hold a short conversation that improves over several messages. This activity should take 5 to 15 minutes:

1. **Start a new chat**: Choose a small planning task, such as a 60-minute activity for a community group, a meal idea, or a gift for someone you know. Use typing, Voice, or a non-identifying photo. Do not upload a file, share a link, or make a purchase.
2. **Send a first request that names the result you want**: Ask ChatGPT to ask up to three questions before it suggests an answer.
3. **Answer the questions**: If you used Voice, read the transcript and correct any words that were misunderstood. If you used an image, check that ChatGPT describes only what you can see.

4. **Send one correction or missing detail:** Ask ChatGPT to revise its answer and label any remaining assumption.
5. **Ask for three alternatives in a table**: Name one fact, cost, source, or safety point that you would check before using an option.
6. **Rename the chat so that you can find it again**: Do not share the conversation link during this activity.



What to ask next: `Review this conversation. List the facts I need to verify, the assumptions you made, and the private information I should remove before sharing any part of it.`

Check the result

- ☐ Did your first message describe a result, a limit, or a format that you wanted?
- ☐ Did you answer a useful question or correct one assumption before asking for a revision?
- ☐ Can you explain one check that you would make before acting on the response?
- ☐ Did you rename the chat without sharing private or work-restricted information?

If an answer is no, return to the chat and add the missing information.

Chapter 1.5 breaks that same instruction into four parts you can reuse for almost any task: the outcome, the context, the limits, and the form of the answer, and shows what to do once ChatGPT's first answer is close, but not quite right yet.

Chapter 1.5 Write, Refine, and Improve Prompts Through Conversation

A prompt is the message you give ChatGPT. A useful prompt tells ChatGPT what you want to do and what a good answer should include. You do not need special wording or a long list of rules. You need to provide enough information in everyday language for ChatGPT to make a suitable first response.

In this chapter, you will write one prompt in four parts: the task and outcome, useful background, limits, and the form of the answer. The same pattern works for planning, learning, writing, and many other everyday tasks.

State the task and desired outcome

Start by naming the job you want ChatGPT to help with. Then say what you want at the end. A task is an action, such as plan, explain, compare, draft, or list. An outcome is the result you can inspect, such as an agenda, a checklist, a table, or a short draft.

For example, imagine that you are helping a community group run a 60-minute class for adults who are new to smartphone photography. "Help with a class" leaves many choices open. "Create a 60-minute class plan" gives ChatGPT a result to work toward, like so:

```
Create a 60-minute class plan that helps adults who are new to smartphone
photography take, find, and share photos safely.
```

The first sentence is enough to begin. You can add the details that affect the answer next. Do not try to predict every question ChatGPT might ask. Include the information that would change the plan if it were missing.

Supply relevant context

Context is the background that helps ChatGPT understand the situation. It can include the audience, what they already know, the purpose, available materials, and facts that shape the task. Give context that is relevant to the answer. Extra details can distract from the main job.

For the smartphone class, the group size, the kinds of phones people have, and the teaching location all affect the plan. A favorite movie or an unrelated work deadline would not help with this request. Here is an example prompt:

```
The class is for eight adults at a community center. Some use iPhones and some use
Android phones. They can use their own phones, a projector, and a printed handout.
Most have taken photos before but do not know where the photos are stored or how to
share one.
```

You can write the task and context as a few short sentences. If the context comes from a document or image, check that you have the right to share it. Remove names, contact details, account numbers, faces, and other information that the task does not need. *Figure 1.5A* shows:

1. **Form of the answer:** "Create a 60-minute class plan..." sets the length and the type of result. ChatGPT returns a timed plan instead of general advice about photography.
2. **Task and desired outcome:** "...that helps adults who are new to smartphone photography take, find, and share photos safely." States the outcome you want. ChatGPT can aim at a clear goal instead of a general topic.
3. **Constraints and boundaries:** "They can use their own phones, a projector, and a printed handout." Lists what is available in the room. ChatGPT can suggest activities that fit the space instead of ones that need equipment nobody has.
4. **Relevant context:** "The class is for eight adults at a community center. Some use iPhones and some use Android phones. Most have taken photos before but do not know where the photos are stored or how to share one." Describes the learners, including group size, devices, and what they already know. ChatGPT can match the plan to this group.

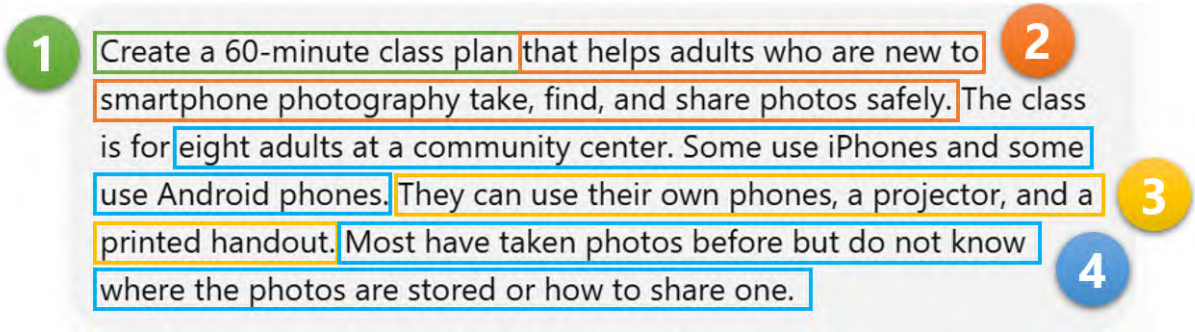


Figure 1.5A: Four kinds of information can make a first prompt more useful. Add only the parts that help with your task.

Set constraints and boundaries

Constraints are limits that the answer must follow. They can include time, budget, length, tools, language, or the number of choices. Boundaries say what ChatGPT should not do. They are useful when a response could otherwise make an unsafe assumption, spend money, or use information you have not supplied.



Watch out: A boundary only rules out what you actually name. Telling ChatGPT "do not invent local rules" does not stop it from inventing a different detail you forgot to rule out. State the boundaries that matter most, and still check the answer.

For the class plan, the group has only 60 minutes and must use the standard photo tools already on the phones. No participant should have to share a private photo. Those details help ChatGPT choose suitable activities:

```
Keep the plan within 60 minutes. Use only standard phone features, not paid apps or
extra equipment. Include an activity that works for both iPhone and Android users.
Do not ask participants to share personal photos. Do not invent local rules or phone
settings. Mark anything that I need to check on my own phone.
```

A boundary is not a promise that ChatGPT will always get the details right. Read the answer and check any fact that affects privacy, money, safety, work, or another person. For an important decision, use qualified human advice where needed.



Good practice: Keep constraints and boundaries specific enough to check. A vague limit such as "keep it reasonable" cannot be verified. A specific one, such as "under 500 words" or "do not invent a fact," gives you something to confirm in the response.

Specify the form of the answer

The form is how you want the answer arranged. Ask for a table when you need to compare items. Ask for a numbered list when you need actions in order. Ask for headings when you want a short guide. You can also set a length, a reading level, or a tone if those choices will help the reader use the result.

For a class plan, a table can make the timing and materials easier to check. You can name the columns rather than hoping that ChatGPT chooses the same structure you had in mind:

```
Give the answer as a table with these columns: time, activity, what participants do,
materials, and privacy reminder. Then add a five-item handout checklist in plain
language. Keep the whole answer under 500 words.
```

Put the four parts together in one message when you are ready to send it. The order is less important than the information. If ChatGPT asks a useful question, answer it. ChatGPT should respond with a writing block containing a 60-minute class plan that can be edited, copied, or downloaded:

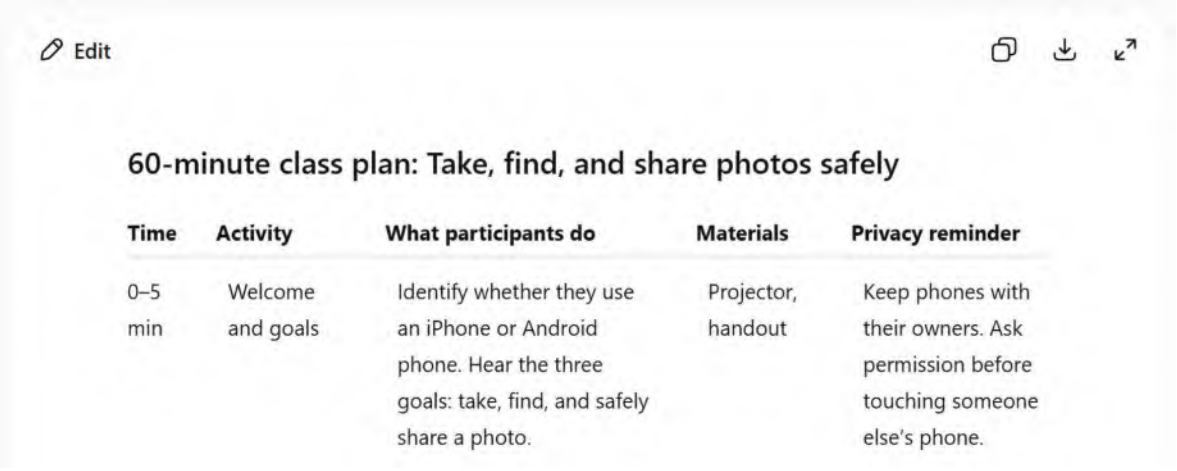


Figure 1.5B: A 60-minute class plan that can be edited, copied, or downloaded.

A first answer gives you something to inspect, but it may still be missing information, make a poor trade-off, or need one section changed. The rest of this chapter continues the smartphone-photography class example, asking ChatGPT what it needs to know, comparing alternatives, checking an answer against your own criteria, and revising one part while keeping the parts you want to keep:

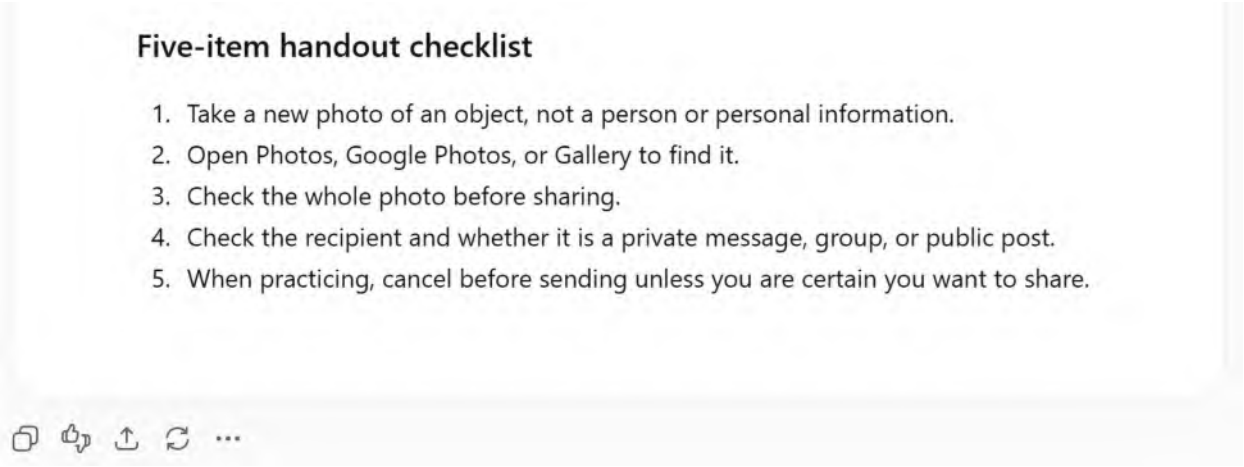


Figure 1.5C: The class plan ends with a five-item handout checklist in plain language, as requested

Ask ChatGPT to identify missing information

ChatGPT can answer a question even when important details are missing. That can be useful for a rough draft, but it can also lead to guesses. Before asking for a larger answer, ask which missing facts would change the result.

Continue the smartphone-photography class example. The first plan may fit the time limit and use standard phone features. It still does not say whether everyone can see a projected screen, whether the group needs accessible materials, or how much help each participant needs. Those details may change the activities:

1. **Request for missing information**:

```
Before you improve the 60-minute smartphone-photography class plan, list up to
five missing facts that could change the plan. Group them as audience, room and
equipment, accessibility, and safety or privacy. Do not guess the answers. For
each fact, say how it could change the plan.
```

2. **Selective answers**: Read the questions before you answer them. Some will be useful. Others may not affect your task. Answer only the ones that change your decision.
3. **Labeled assumptions**: If you do not know the answer, say so and ask ChatGPT to make any labeled assumptions that you can review later.

Request alternatives rather than accepting the first answer

A response can meet your instructions and still be one of several suitable choices. Ask for alternatives when you need to compare time, cost, order, level of detail, or another tradeoff. Name what must stay the same, so ChatGPT does not discard useful work.

For the class plan, one version might spend most of the time practicing with phones. Another might include more time finding and organizing photos. Both can fit the same audience and privacy boundary. Comparing them helps you choose the version that suits the group.

Here is an example prompt:

```
Give me three alternative 60-minute class plans. Keep the audience, standard-phone-
features rule, and privacy boundary from the current plan. Make the main emphasis
different in each version: taking photos, finding and organizing photos, or sharing
photos safely. Put the options in a table with these columns: emphasis, main
activities, advantage, possible drawback, and what I need to check.
```

Ask for alternatives that differ in a named way. A request for "better ideas" does not tell ChatGPT what better means. A request for options with different emphases gives you a comparison that you can inspect.

Review against stated criteria

Criteria are checks that you use to judge an answer. Write them before you ask ChatGPT to review its work. They can include facts from the task, such as the time limit, and qualities that you want, such as plain language or a privacy reminder in each activity.

For the class plan, the criteria might be: it lasts 60 minutes, works for both iPhone and Android users, uses no paid apps, includes a privacy reminder, and gives beginners enough practice time. A criterion should be specific enough that you can explain why a plan passes or needs work. For example:

```
Review the [chosen version] against these criteria: 60 minutes total; suitable for
beginners using iPhone or Android phones; no paid apps or extra equipment; at least
20 minutes of practice; and a privacy reminder in every activity. Use a table with
the criterion, pass or needs revision, evidence from the plan, and any question
that I still need to check. Do not claim that the plan passes when information is
missing.
```

ChatGPT can help perform the review, but you remain responsible for the decision. Check its timing, calculations, facts, and quotations. For work that affects safety, money, health, legal duties, employment, or another person's rights, use suitable qualified human review.



Watch out: Asking ChatGPT to review its own work against your criteria is still ChatGPT judging ChatGPT. It can mark a weak answer as passing. Spot-check at least one criterion yourself instead of accepting every "pass" at face value.

Revise one part without losing the rest

A follow-up can change one section without starting the task again. Name the part that needs work, state what must remain unchanged, and ask to see only the changed section when that makes comparison easier.

Suppose the review shows that the plan has only 10 minutes of practice. You can ask ChatGPT to adjust the timing while keeping the audience, privacy boundary, and the activities that already work:

```
Revise only the timing and activities needed to provide at least 20 minutes of
hands-on practice. Keep the total at 60 minutes. Keep the audience, standard-phone-
features rule, privacy reminders, and the existing opening activity. Show only the
changed rows, then give the revised total time.
```

Read the changed rows and compare them with the earlier plan. ChatGPT may still alter more than you asked it to change. If that happens, point to the changed item and ask it to restore the earlier version. For longer work, copy approved text into your own document and keep a dated version before you request a major rewrite.



Good practice: Use a new chat when you change to a different subject. Use the same chat when the next request depends on its earlier messages. If you must begin a new chat, paste the short task summary, the criteria, and the current version that you want revised.



Learn more online: *Prompt engineering best practices for ChatGPT*: `https://help.openai.com/en/articles/10032626-prompt-engineering-best-practices-for-chatgpt`.

Try it now

Write one prompt in four parts, then improve it through conversation. This activity should take 10 to 20 minutes. Choose a personal planning, learning, or writing task. Do not upload a file or enter private information:

1. Start a new chat. Choose a task with a visible result, such as a study plan, a packing checklist, a comparison table, or a draft agenda.
2. Write one sentence that states the task and desired outcome. Add two or three facts that provide relevant context, one or two constraints, and a boundary if needed.
3. Ask for a specific form of answer, such as a table with named columns or a numbered checklist. Send the prompt.
4. Ask ChatGPT to list missing information that could change the result. Answer the useful questions.
5. Ask for two or three alternatives that differ in one named way. Choose one, and ask ChatGPT to review it against two to four of your own criteria.
6. Request one focused change. Ask to see only the changed section, then compare it with the earlier version before you keep it.



What to ask next: `Summarize the decisions in this conversation. List the information I supplied, the assumptions that remain, the criteria that still need a human check, and the next action I should take.`

Check the result

- ☐ Did your prompt name both a task and a result you can review, using context that changes the answer without adding private or unrelated details?
- ☐ Did you state a limit or boundary that protects your time, money, privacy, or work?
- ☐ Did you ask for missing information instead of letting ChatGPT fill every gap with a guess, and did the alternatives differ in a named way you could compare?
- ☐ Did your final instruction change one part while keeping the useful parts of the earlier answer?

If an answer is no, revise that part of the prompt and compare the next response. You are learning how to give ChatGPT useful instructions, not trying to write a perfect message on the first attempt.

Chapter 1.6 uses this same habit of staying in the conversation and turns it toward learning: asking for an explanation at the right level and checking your own understanding instead of only reading ChatGPT's answers.

Chapter 1.6 Learn Something New and Educate Your Child

Chapter 1.5 showed how to keep a conversation going until a task's result is right. That same habit works for learning, whether you're the one learning or you're helping a child with schoolwork. ChatGPT can act as a patient tutor for almost any subject: a language, a craft, a school topic, or a skill you need for work.

In this chapter, you will ask for an explanation pitched to your level, break a subject into a short learning path, request analogies and practice questions, let ChatGPT question you instead of only answering, check your own understanding by explaining an idea back in your own words, and use those same habits to help a child with homework without doing it for them.

Ask for an explanation at the right level

A good teacher asks what you already know before explaining something new. ChatGPT cannot see your background unless you state it. The same question can need a different answer for a beginner and for someone who already knows the basics.

To help ChatGPT to help you learn something new:

1. Say what you can already do
2. Name the specific part that confuses you
3. Set a length so the answer stays focused

Imagine you are learning Spanish before a trip and the difference between the verbs *ser* and *estar* keeps tripping you up:

```
I know how to introduce myself and order food in Spanish, but present tense verbs
still confuse me. Explain when to use ser and when to use estar, using words I
already know. Keep the explanation under 200 words and include two example sentences.
```

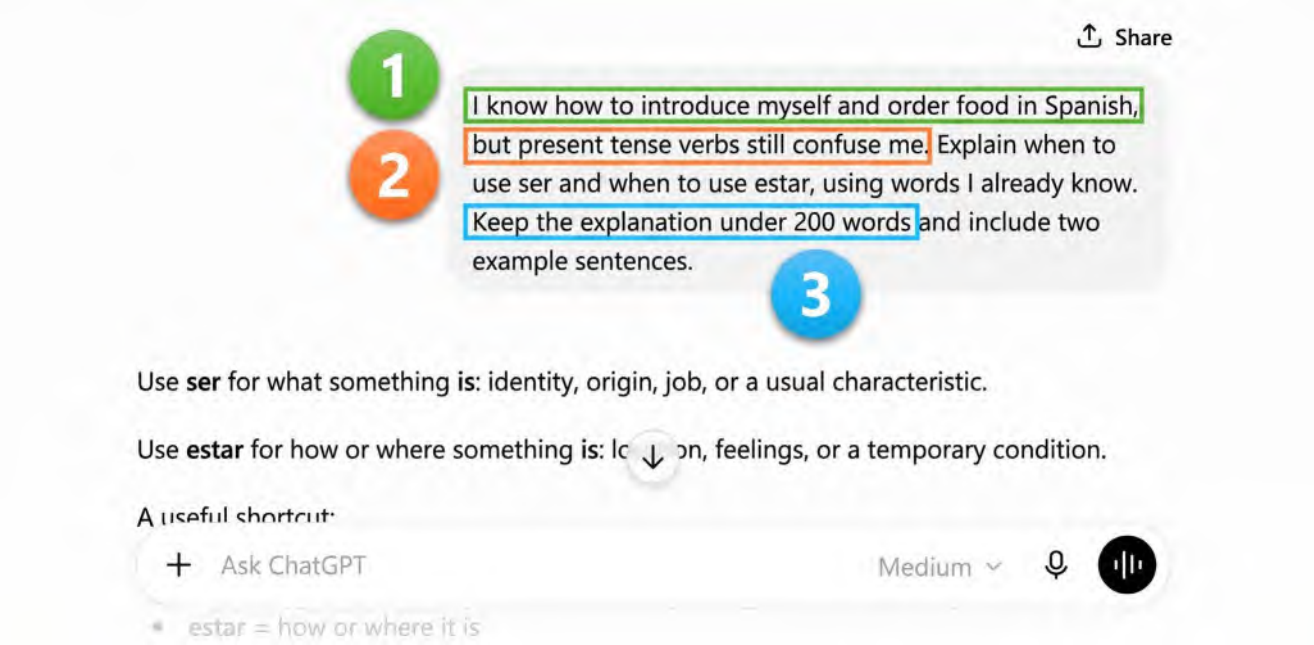


Figure 1.6A: Say what you can already do, name the specific part that confuses you, and set a length so the answer stays focused

If the reply is still not clear, say which part did not make sense. Ask for the same idea explained a different way, such as with a shorter answer, a different example, or fewer new words.

Break a difficult subject into a learning path

A single explanation can cover one idea, but a subject usually has several parts that build on each other. Ask ChatGPT to lay out a short sequence of topics in the order you should learn them, instead of trying to absorb everything at once:

```
I want to hold a simple conversation in Spanish on a trip in three weeks. Break this
into a short learning path. List the grammar and vocabulary topics in the order I
should learn them, and mark which topics depend on an earlier one.
```

Read the path before you follow it. Ask ChatGPT to explain why one topic comes before another if the order is not obvious. You can also ask it to shorten the path if you have less time, or to mark which topics you can skip for your specific goal.

Request examples, analogies, and practice questions

An explanation is easier to remember when it links to something you already understand, and easier to use when you get to practice it. Ask ChatGPT to compare the new idea to something familiar, then ask for a short exercise that uses it:

```
Compare the Spanish subjunctive mood to something in English grammar that I probably
use without noticing. Give one short example in each language. Then write two
practice sentences with a blank where I need to choose the correct verb form.
```

Check whether the comparison actually fits. A comparison can make an idea memorable while still leaving out an important difference:

```
Where does that comparison break down? Give one case where English and Spanish
handle this differently.
```

Let ChatGPT question you, rather than reveal the answer

Reading an explanation is not the same as learning it. Answering a question about it is what helps the idea stick. Ask ChatGPT to write questions without the answers, answer them yourself, and then ask it to check your work.

```
Write five short quiz questions about ser and estar. Do not include the answers.
Wait for me to answer each question before you tell me if I am right, and explain
any mistake.
```



Figure 1.6B: Asking for a quiz on Spanish grammar.

When you answer, ChatGPT will correct you if needed, and then continue with the quiz:

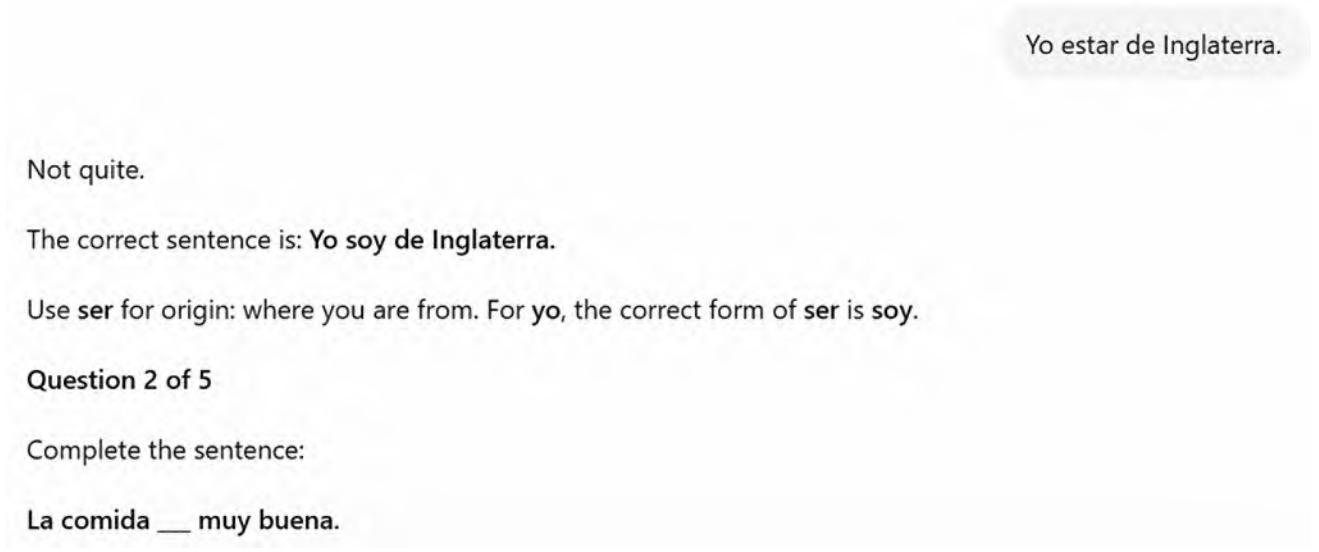


Figure 1.6C: ChatGPT will correct you and then continue with the quiz.

ChatGPT will sometimes reveal an answer inside a later question or explanation if you do not ask it to hold back. If that happens, ask for a fresh set of questions and try again.

Explain it back to check your understanding

Once you can answer quiz questions, try the harder test of explaining the idea in your own words. Write your explanation, then ask ChatGPT to check it for gaps or mistakes instead of grading you against its own version first:

```
I will explain when to use ser versus estar in my own words. Read what I write and
tell me what is correct, what is missing, and what is wrong: [type your explanation
here]
```



Good practice: A day or two after ChatGPT checks your explanation, try explaining the same idea again without looking at the chat. If you cannot restate it without checking the chat, that topic needs another round of practice.

Verify facts and keep your own thinking in charge

ChatGPT can state a grammar rule, a translation, or a fact with confidence even when it is wrong or does not fit your specific region or dialect. Before you rely on a phrase in person, such as ordering food or asking for help, check it with a native speaker, a textbook, or a course you trust.

Use ChatGPT to explain, quiz, and check your work, not to think through the problem for you. Try an exercise on your own first, then ask ChatGPT to review your attempt. This keeps the practice in your hands and gives ChatGPT a specific, useful job: checking your reasoning instead of replacing it.



Watch out: A confident-sounding answer is not automatically an accurate one. Treat a surprising fact, an unfamiliar rule, or anything you plan to say to another person as worth a second source.

Use ChatGPT ethically for schoolwork and essays

Nowhere does the habit from this chapter – explaining, quizzing, and checking rather than answering for you – matter more than a graded assignment. Using ChatGPT the wrong way on an essay can cost you an academic integrity case, not just earn a weak grade.

Read the AI policy for the specific course and assignment, as well as your school's general policy. Course-level rules increasingly differ from each other even within the same department, and a general school policy may not settle what your instructor actually expects. If nothing is written down, ask before you start, and keep a record that you asked.



Good practice: Write your own first draft of an academic essay assignment before asking ChatGPT to react to it. Getting feedback on your own attempt teaches you more than asking ChatGPT to produce the attempt for you, and it keeps the finished essay defensibly yours.

There's a real line between support and substitution. Asking ChatGPT to brainstorm angles, build an outline, explain a source, check grammar, or react to a draft you already wrote is learning support. Asking it to write paragraphs, or the whole essay, that you then submit as your own is not, even at schools that otherwise welcome AI use, because it misrepresents someone else's writing as yours.

If your instructor asks you to disclose AI use, be specific rather than vague. "I used ChatGPT to build an initial outline and check grammar in my final draft" tells your instructor something they can evaluate. "I used AI" does not and reads as evasive even when it's honest.



Real-world example: Before writing a response to an assigned reading, a student uses ChatGPT to quiz them on the reading's argument until they can restate it correctly, then writes the analysis entirely themselves and adds a short note describing that use, matching their instructor's disclosure requirement.



Watch out: Submitting text ChatGPT wrote as entirely your own work counts as academic misconduct at most schools, whether or not a detector flags it, and detection tools are unreliable. This isn't hypothetical: in 2026, a publisher canceled a debut novelist's roughly $2 million book deal after the author's own team said they could no longer verify how the manuscript was actually written, regardless of whether AI was involved. "Will it get caught?" is the wrong question. The question to ask is whether you can honestly defend the work as your own.

Helping a child with their homework

Homework help is one of the most common reasons a parent opens ChatGPT in the evening, and one of the easiest places to slip from teaching into doing. After a long day, it is tempting to paste your child's worksheet question into the chat, copy back the answer, and move on. That fills in the page. It does not help your child learn, and it leaves them stuck the next time a similar question shows up on a test with no ChatGPT beside them.

Treat ChatGPT here the same way this chapter treats it everywhere else: as a tutor that questions and explains, not a source of finished answers you relay. You should stay in the room, read what ChatGPT asks and says, and decide what your child sees.

Children copy how the adults around them use a tool as much as what those adults tell them. Talk through what you are doing as you use ChatGPT with, or in front of, your child: "I'm asking it to check my explanation before I explain this to you," or "I'm going to ask it to question you instead of giving the answer." That running commentary teaches the habit you want your child to build later, more than a rule about screen time will.

Math and science: ask for a method, not a result

Ask ChatGPT to work alongside your child rather than solve for them:

```
My child is working on this algebra problem: [paste the problem]. Do not give the
final answer. Ask one guiding question at a time, and only offer a hint if they are
still stuck after two attempts.
```

Check the working, not just the final number. A child who reaches the right answer from a hint they did not understand will struggle with the next problem that uses the same idea.

English essays: feedback on their draft, not a draft from scratch

The child should write their essay first, and then ask ChatGPT to react to what they wrote:

```
My child (age 12) wrote this paragraph for a book report: [paste the paragraph]. Do
not rewrite it. Ask two questions that would help them add more detail, and point
out one thing that already works well.
```

History and art: use the parts a worksheet can't mark

History and art homework often plays to what ChatGPT is genuinely good at: quizzing dates and causes back to your child, suggesting a few directions for a project, or explaining a technique they noticed in a museum photo. Ask for questions and options here too, rather than a finished paragraph or a ready-made "five causes of X" list your child can paste in unchanged.

Check the assignment's own rules first

Before you open ChatGPT for a specific assignment, check what the teacher and school actually allow. Some homework is meant to be closed-book practice; some explicitly welcome AI tools. When you are not sure, have your child ask their teacher rather than guessing.



Watch out: The hardest version of this habit to catch is reaching for ChatGPT because a deadline is close, not because your child is stuck on the material. If you notice you are using it to finish something quickly rather than to help your child understand it, that is the moment to slow down, not speed up.



Learn more online: *ChatGPT Study Mode FAQ*: `https://help.openai.com/en/articles/11780217-chatgpt-study-mode-faq`.

Try it now

Use ChatGPT to start learning one topic or skill. This activity should take 15 to 20 minutes. Choose a subject you want to learn for work, school, or personal interest:

1. Start a new chat. State your goal in one sentence and describe what you already know about the topic.
2. Ask for a short learning path: the topics to cover, in order, and which ones build on each other.
3. Pick the first topic. Ask for an explanation at your level, an analogy, and one or two practice questions.
4. Ask for three to five quiz questions without answers. Answer each one, then ask ChatGPT to check your answers and explain any mistakes.
5. Explain the topic back in your own words and ask ChatGPT to point out what is missing or wrong.



What to ask next: `Based on what I got right and wrong, tell me which topic I should study next and why.`

Check the result

- ☐ Did you tell ChatGPT what you already know before asking for an explanation?
- ☐ Does your learning path show which topics come first and why?
- ☐ Did you answer the quiz questions yourself before ChatGPT checked them?
- ☐ When you explained the topic back in your own words, could ChatGPT confirm it was accurate and complete?

If an answer is no, return to the chat and ask ChatGPT to try that part again.

Whether you're learning something new yourself or helping a child do the same, the habit is the same: ask, question, and check, rather than just answer. *Chapter 1.7* turns from learning a subject to planning everyday decisions: explaining your circumstances, comparing options, and testing assumptions before you choose.

Chapter 1.7 Plan Everyday Life and Compare Important Choices

You will use ChatGPT to think through a real decision instead of asking one question and stopping there. This works for choosing a product, planning a trip, organizing a home project, preparing meals from what is already in your kitchen, planning a new exercise routine, or preparing questions before you see a doctor, a lawyer, an accountant, or a financial adviser.

Before you begin, pick one decision you are actually facing. A real decision, even a small one, teaches this skill better than a made-up example. By the end of this chapter, you will have a short, organized comparison you can use to decide, or to take into a conversation with a person who can advise you.

Explain your circumstances, goals, and constraints

Start by describing your situation in your own words. Include what you want to achieve, what limits you, such as budget, time, or location, and anything that you personally care about.

For example:

```
I'm planning a five-day trip for two adults and one child. Our budget is US$1,500
excluding flights. We want somewhere warm in October, and my child gets bored on
long car journeys.
```

The more relevant detail you give, the more useful the reply. Leave out anything private that you would not want stored, such as your exact address or account numbers.

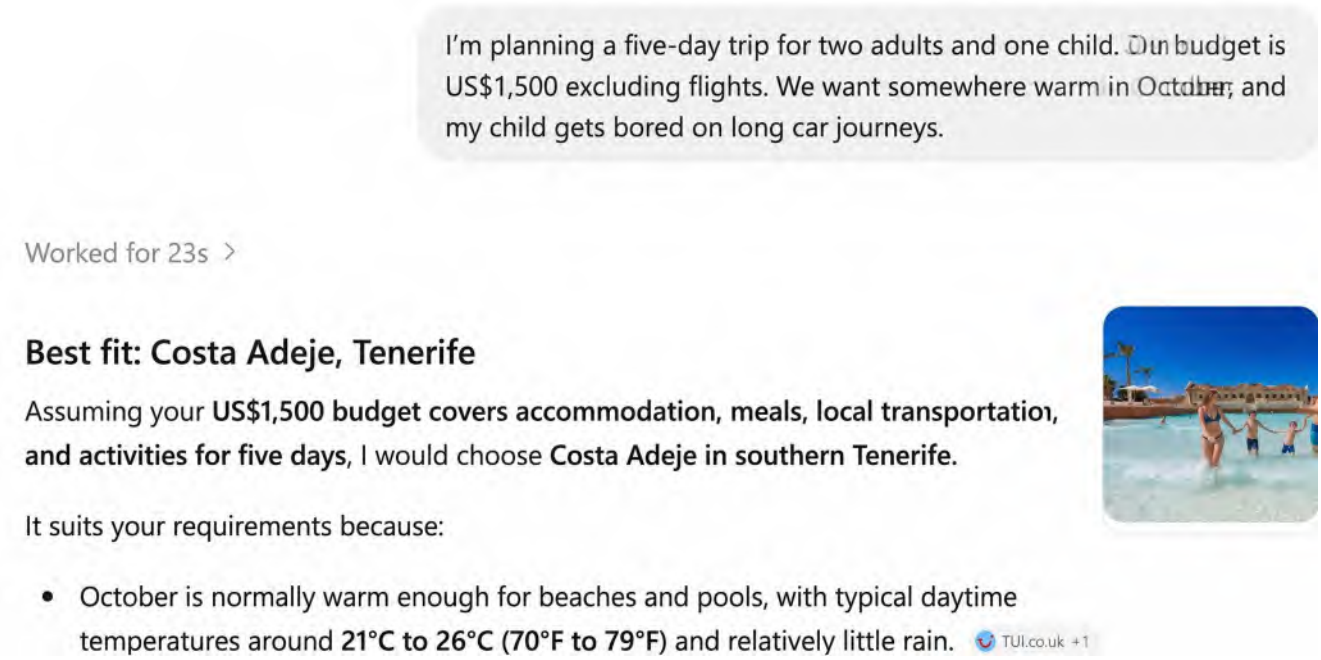


Figure 1.7A: A clear starting message names the goal, the budget, and the one constraint to plan around: a child who dislikes long car journeys.

Ask ChatGPT to identify missing information

Before you ask for a recommendation, ask what else would help. A short prompt does this well:

```
Before you suggest anything, tell me what other information would help you give a
better answer.
```

ChatGPT might ask about your travel dates, how you feel about flying, or whether you have visited similar places before. Answer the questions that matter to you and skip the rest.



Good practice: Treat these follow-up questions the way you would treat questions from a helpful travel agent or shop assistant. You do not owe a complete answer to every one, but each answer you give narrows the search.

Generate options before asking for a recommendation

Ask for several options before you ask which one is best. This keeps you from anchoring on the first idea you hear. For example:

```
Give me four different options that fit what I've described. Don't recommend one
yet. List the choices only.
```

Reviewing several options side by side shows you the real range of choices before you narrow down to one.

Compare options using explicit criteria

Once you have options, name the criteria that matter to you and ask ChatGPT to compare the options against them. Criteria might include cost, time, effort, risk, or how well an option fits a constraint you named earlier:

```
Compare these four options against cost, travel time, and how well they suit a six-
year-old. Use a table.
```

A table makes trade-offs easier to see than a paragraph of prose. Ask for one whenever you are comparing more than two or three options. The table might look like this:

Goals	Constraints	Evidence	Trade-offs	Unanswered questions
Orlando, Florida Warm pools and theme parks	US$1,500 for 5 days, no long drives	Warm October weather, short airport-to-hotel transfer	Higher attraction costs, busy resort areas	Which airport? Do you want a theme-park day?
San Diego, California Beach and animals	US$1,500 for 5 days, no long drives	Mild, sunny October, zoo and beach close together	Usually less hot, hotel prices vary	Is mild warmth enough? Do you need a pool?

Cancun, Mexico Beach resort break	US$1,500 for 5 days, no long drives	Warm sea and resort activities, airport transfers available	Weather can be rainy, passport and resort costs	Are passports ready? All-inclusive or self-catering?

Table 1.7A: A choice matrix separates what you want from what you know, what you would give up, and what you still need to find out

Test assumptions and consider disadvantages

Every option has a downside. Ask directly for the disadvantages and question any assumption you are not sure about:

```
What could go wrong with option 2? What am I assuming that might not be true?
```

This step is most useful when an option sounds appealing. A confident-sounding answer is not the same as a well-supported one. Ask ChatGPT to explain its reasoning, and check anything that surprises you against another source before you rely on it.



Watch out: ChatGPT can describe disadvantages you did not ask about only if you prompt for them. Left alone, it will often lean toward the option that best matches your stated preferences, so ask it to argue against your likely choice at least once.

Separate low-stakes guidance from decisions requiring qualified advice

Some decisions carry little risk if you get them wrong, such as which restaurant to try or how to pack for a trip. Other decisions, including anything about your health, your legal rights, your taxes, or your safety, deserve a qualified professional's review before you act.

ChatGPT can help you prepare for that conversation. It can help you list your symptoms clearly before a doctor's appointment, draft questions for a solicitor or lawyer, or organize your finances before you speak with an accountant. It should not replace the appointment itself.

Ask ChatGPT to help you prepare, then confirm with a qualified person.



Figure 1.7B: As the stakes rise, ChatGPT's role shifts from adviser to a way of preparing for advice from a qualified person.

Try it now

Choose a real decision you are currently facing, even a small one. Describe your goals and constraints, ask what information is missing, request several options, compare them against your own criteria in a table, and ask what could go wrong with your likely choice.



What to ask next: `Turn this comparison into a short checklist I can bring to a qualified person. List the questions I should ask them and the facts I should confirm myself first.`

Check the result

- ☐ Did the comparison use criteria that actually matter to you, not just the ones ChatGPT chose?
- ☐ Did you ask for disadvantages, or only see the advantages of each option?
- ☐ If this decision involves your health, legal rights, money, or safety, do you have a plan to confirm it with a qualified person?

The comparison you just built rests on ChatGPT's own reasoning.

Chapter 1.8 adds a further check: making sure that any current fact behind your decision, such as a price or an opening date, is actually up to date, and when a task needs current information from the web and how to inspect the sources that ChatGPT uses.

Chapter 1.8 Search and Verify Current Information

You will ask ChatGPT a question that needs current information, and then check the sources it used before you trust the answer. Do this for anything that changes often, such as prices, news, opening hours, availability, or a person's current job title, because ChatGPT's stored knowledge can be out of date.

Before you begin, think of a question you would normally look up online, such as whether a store is open today or what a recent news story actually said. By the end of this chapter, you will know when ChatGPT searches the web on its own, how to open and read its sources, and how to check that a source really supports what ChatGPT told you.

Know when web search is necessary

ChatGPT can answer many questions using what it learned during training. That training has a cutoff date, so it cannot know about anything that happened after that point on its own. For anything current, ChatGPT needs to search the web while answering you.

ChatGPT often searches automatically when your question includes words like "today," "this week," "latest," "current," or "price." You can also ask for a search directly if you are not sure it happened:

```
Search the web for this and show me your sources: What are the current opening hours
for [a specific business]?
```

For example, I might search for a local burger joint in my hometown of Brighton:

```
Search the web for this and show me your sources: What are the current opening hours
for Smashburger in Brighton and what special offer do they have today?
```

You can explicitly activate **Web search** by clicking the + in the chat box:

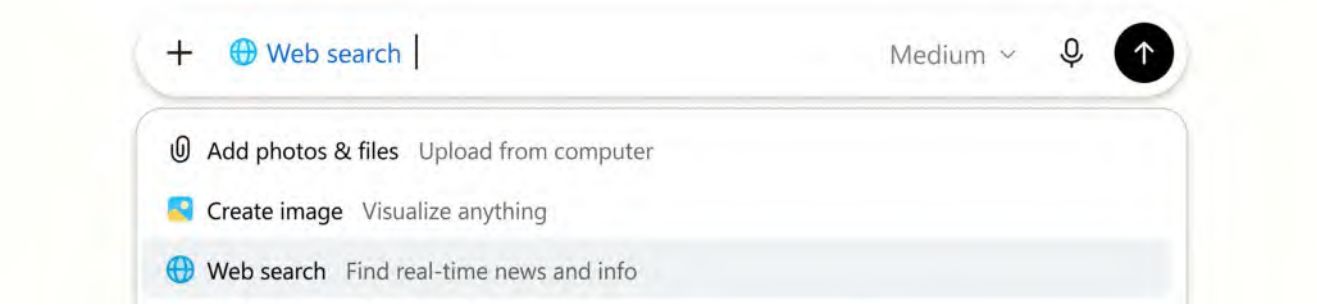


Figure 1.8A: Activating Web search

Good practice: If you suspect an answer could be outdated, such as a price, a schedule, or a recent event, ask ChatGPT to search rather than assuming it already has. Check all sources before trusting the supported answer.

Learn more online: *ChatGPT Search*: `https://help.openai.com/en/articles/9237897-chatgpt-search`.

Ask for useful and varied sources

A single source can be wrong, outdated, or one-sided. Ask ChatGPT for more than one source, and ask for sources of a particular kind when accuracy is critical, such as an official website, a news report, or a government page:

```
Give me three sources for this, including at least one official or primary source,
and tell me how they differ.
```

You should now see a response with visible citations and a **Sources** panel you can open. If ChatGPT gives you only one source, or sources that all repeat the same original article, ask for a wider search before you rely on the answer.

Open and inspect citations

When ChatGPT uses search, it usually adds citations, small markers linked to the pages it used. On the web version, you can hover over or click a citation to see where it came from. Below the response, you can click the **...** button and then click **View sources** to open a panel listing every page ChatGPT used, along with other links it found along the way:

1. Read the response as usual.
2. Hover over or click a numbered citation to see its source.
3. Click **...** and then **View sources** below the response to open the full list.

4. Open at least one source yourself for any claim you plan to rely on:

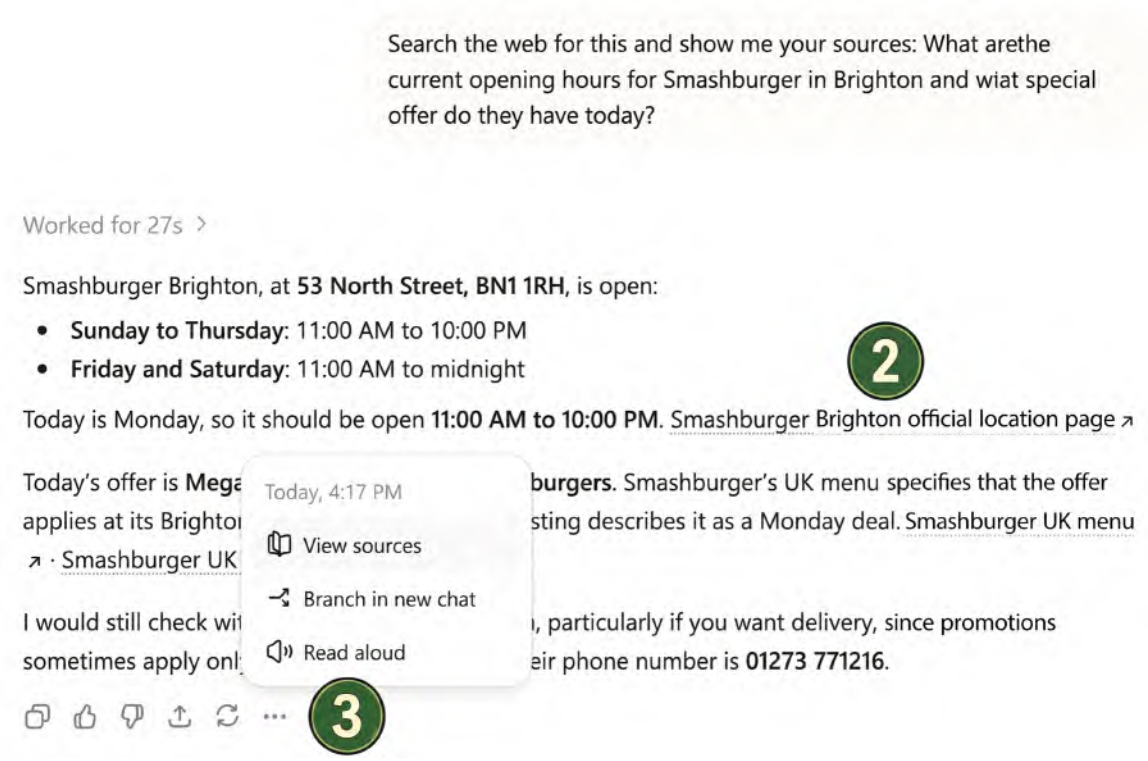


Figure 1.8B: Inline citation markers point to specific sentences. The Sources panel lists every page ChatGPT drew from, so you can open one directly.

Verify dates, claims, and quotations

A citation shows where ChatGPT found information. It does not prove that the source actually says what ChatGPT claims. Open the source and check three things: the publish or update date, whether the page truly supports the specific claim, and whether any quoted words match the source exactly.



Watch out: A confident answer with a linked source is not automatically accurate. ChatGPT can summarize a source incorrectly or attach the wrong page to a claim. Treat every citation as a starting point for your own check, not as proof on its own.

For anything you plan to repeat, publish, or act on, such as a statistic in a report or a quotation in an email, open the original source and read the relevant passage yourself.

Try it now

Ask ChatGPT a question about something current, such as a recent price, a news event, or an organization's current contact details. Request three sources, open the **Sources** panel, click through to one source, and check whether it actually supports the claim ChatGPT made.



What to ask next: `Show me how you decided this needed a web search instead of answering from memory. Then tell me one kind of question where you would not need to search.`

Check the result

- ☐ Did ChatGPT search the web, or did it answer from memory alone?
- ☐ Do the sources come from more than one website?
- ☐ Does the source you opened actually say what ChatGPT claimed, including any dates or numbers?

A single search answers a single question well. *Chapter 1.9* covers deep research, a built-in feature for questions that need several sources and more time to work through properly.

Chapter 1.9 Conduct Deeper Research with Cited Sources

You will ask ChatGPT to run a longer research task instead of a single search. This built-in feature is called **deep research**. Rather than answering right away, it plans its own steps, reads many sources over several minutes, and returns a structured report with citations.

To access the deep research feature, you must:

1. Be in **Chat** mode (because deep research is not available in **Work** mode).
2. Click the + menu.
3. Select **Deep research**:

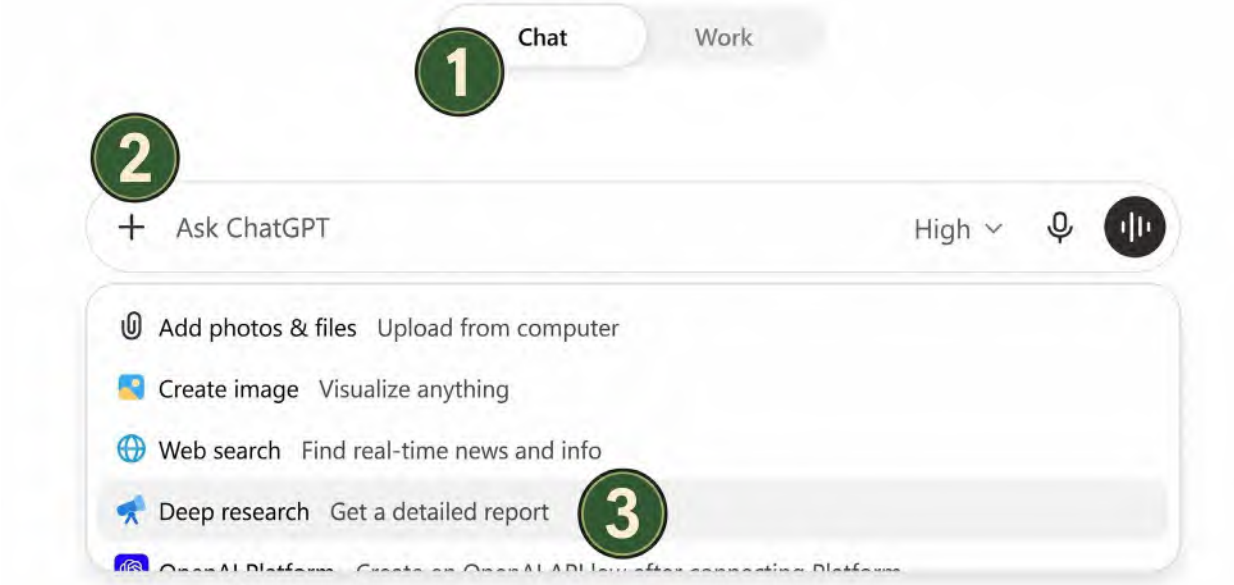


Figure 1.9A: Activating the deep research feature

Deep research suits questions that a single search cannot answer well, such as comparing several options, understanding a topic with conflicting evidence, or preparing a briefing that needs multiple sources. Everyone can try it. The free plan includes a small number of lightweight research tasks each month. Paid plans include more tasks and a fuller version that reads more sources in greater depth. Check your plan's current limit before you start, since it can change.

Before you begin, have a specific question ready. By the end of this chapter, you will have run one deep research task and reviewed its report against its own sources.

Define a research question with scope and request a plan

A vague question produces a vague plan. Narrow your question before you start and decide how wide the research should range.

State the topic, the relevant time period, and any boundary you want respected. For example, rather than “Tell me about electric cars,” try:

```
Research the main barriers to electric car adoption in Europe over the past two
years, focusing on charging infrastructure and price.
```

A clear scope also tells ChatGPT what to leave out, which keeps the final report focused on what you actually need.

Before deep research begins gathering sources, it can show you a plan, a short list of the sub-questions it intends to answer and the kind of sources it will look for. Review this plan and adjust it before you let the task run:

```
Before you start, show me your research plan so I can check it.
```

To request and review that plan before the research runs:

1. Activate deep research before you type or send your prompt.
2. When ChatGPT shows its research plan, click **Edit** and then describe how to add, remove, or reorder items in the plan if something is missing or unnecessary.
3. Click **Start** when you’re ready to start the research. A research task can take several minutes to finish, so a short review now saves you from an off-target report later:

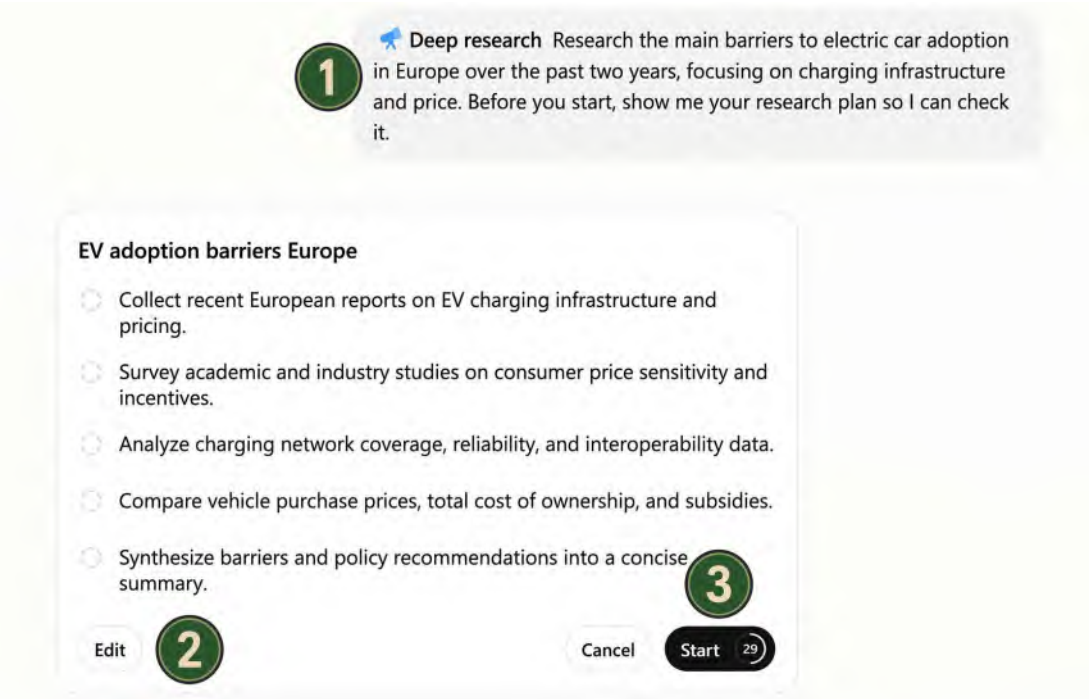


Figure 1.9B: Deep research proposes a plan before it searches. Reviewing it at this early stage is faster than fixing a finished report.

4. Read or download the finished report:



Figure 1.9C: Read or download the finished report

There are several techniques to get better research results.

Set source and date requirements

Tell ChatGPT what kind of sources you want and how current they need to be. This is especially useful for topics where outdated or low-quality sources would mislead you:

```
Use sources published in the last 18 months where possible. Prefer official
statistics, peer-reviewed research, and established news organizations over blogs or
forum posts.
```

If your topic depends on very recent developments, say so directly, since deep research otherwise decides source recency on its own.

Turn findings into a useful conclusion

Ask for a report structured around your original question rather than a plain list of facts. A useful structure includes a short summary, the main findings grouped by theme, any contradictions between sources, and what remains unclear:

```
Summarize your findings in four sections: key findings, points where sources
disagree, remaining gaps, and a one-paragraph conclusion.
```

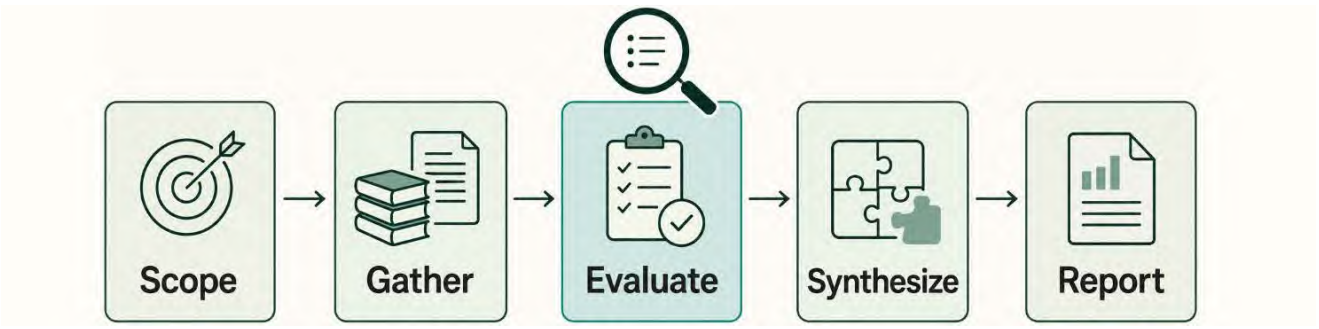


Figure 1.9D: Deep research moves through these five stages on its own. Your scope at the start and your review at the end are the two points where you stay in control.



Watch out: A longer, more thorough-looking report is not automatically more accurate than a short one. Deep research can still misread a source or overstate a weak finding. Treat its report as a strong first draft that still needs your review.

Read the report the way you learned in *Chapter 1.9*: open a few of the cited sources yourself, check that they say what the report claims, and note the publication dates.



Good practice: Save a finished report you plan to reuse. You will build on this skill in *Chapter 1.12*, where projects keep research material alongside related conversations.



Learn more online: *Deep research in ChatGPT*: `https://help.openai.com/en/articles/10500283-deep-research-in-chatgpt`.

Try it now

Choose a question that needs more than one source to answer well. Define its scope, ask ChatGPT to show its plan before starting, set a date and source requirement, and request a report with key findings, disagreements, and gaps.



What to ask next: `Turn this report into a one-page briefing for someone who has not read the full research. Keep the key findings and disagreements, and flag anything that still needs a human decision.`

Check the result

- ☐ Did the final report match the scope you set, or did it wander into unrelated territory?
- ☐ Did you review the plan before the research ran, or only the finished report?
- ☐ Did you open at least one cited source yourself, and did it support the claim attached to it?

Deep research works from sources ChatGPT finds on its own. *Chapter 1.10* turns to sources you already have: your own documents, photos, and spreadsheets.

Chapter 1.10 Add Documents, Images, and Data to a Chat

You will upload a file, a photo, or a spreadsheet and ask ChatGPT to work with it directly, instead of typing everything out yourself. This includes documents such as contracts and reports, photos of receipts or whiteboards, and structured data such as a spreadsheet or CSV file.

Before you begin, have one file ready to try, something ordinary, such as a bill, a short report, or a spreadsheet with a few rows of data. By the end of this chapter, you will know how to upload a file, ask a clear question about it, and get a useful, checkable result back.



Watch out: The free plan allows a limited number of file and image uploads per day. Paid plans raise these limits. The exact numbers change from time to time, so check the book's companion page or your account settings for the current limits.

Upload files and explain the assignment

Click + in the chat box, select **Add photos & files**, and then choose a file from your device or a connected cloud storage account. ChatGPT accepts common document formats such as PDF and Word files, image formats such as JPEG and PNG, and spreadsheet formats such as Excel and CSV.

A file you upload here is also saved to your Library, a persistent store you can search and reuse in later chats without uploading it again. *Chapter 1.11* covers the Library in full.



Sample files: If you do not have your own files, then you can download a sample file from the following link: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book1/1.11-Two-Page-Rental-Agreement.pdf`.

Uploading a file is not enough on its own. Explain what you want done with it, just as you would if you handed a document to a colleague. State the purpose, the audience if there is one, and the form you want the answer in:

```
This is a two-page rental agreement. Summarize the key obligations for the tenant in
plain language, and flag anything unusual compared to a standard lease.
```

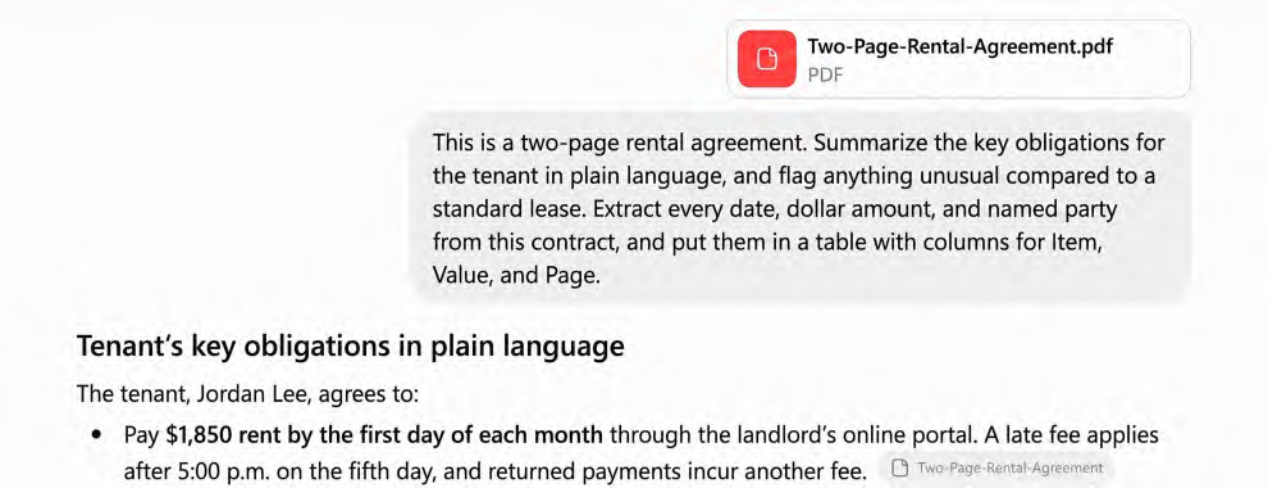


Figure 1.10A: Naming the file's purpose produces a far more usable result than uploading a file with no instructions



Watch out: Do not upload passwords, private keys, medical records, or anyone else's personal information unless you have the right to share it. Check your organization's rules before uploading work documents, a topic covered fully in *Chapter 2.1*.



Learn more online: *File Uploads FAQ*: `https://help.openai.com/en/articles/8555545-file-uploads-faq`.

Summarize and extract information

Ask for a summary written for a specific reader rather than a general one. A summary for yourself can stay informal. A summary for a manager or a client should be shorter and focus on decisions and numbers.

To pull specific details out of a longer document, ask for them by name and request a structured format:

```
Extract every date, dollar amount, and named party from this contract, and put them
in a table with columns for Item, Value, and Page.
```



Extracted dates, dollar amounts, and named parties

Item	Value	Page
Named landlord	Example Property Group LLC	1
Named tenant	Jordan Lee	1
Lease start date	September 1, 2026	1

Figure 1.10B: Naming the exact output you want, in this case, a table, produces a far more usable result

A table is easier to scan and easier to check against the original document than a paragraph.

Compare multiple documents

Upload two or more files in the same conversation and ask ChatGPT to compare them directly. This works well for comparing quotes from different suppliers, two versions of the same document, or a set of resumes against a job description:

```
I've attached three supplier quotes. Compare them on price, delivery time, and
warranty terms, and show the comparison as a table.
```

Name the specific points of comparison you care about. Without direction, ChatGPT may compare details that matter less to you and skip the ones that matter most.

Analyze images and structured data

ChatGPT can look at a photo and describe what it sees, read printed or handwritten text, and answer questions about a chart or diagram in an image:

```
Here's a photo of a whiteboard from our meeting. Transcribe the text and organize it
into a list of decisions and action items.
```

Handwriting, glare, and blurry photos reduce accuracy, so review the transcription against the original image, especially for numbers and names.

For spreadsheets and other structured data, ChatGPT can clean messy columns, calculate totals, and build a chart from the numbers you provide:

```
Here's a spreadsheet of monthly expenses. Clean up any inconsistent category names,
then show total spending by category as a bar chart.
```

You should now see a chart or table built directly from your uploaded data. *Chapter 2.5* covers spreadsheet analysis in more depth, including checking formulas and calculations.



Good practice: Keep the original file open next to ChatGPT's result. Compare the two side by side so you can catch a misread number, a dropped row, or a merged cell before you use the extracted version.

Try it now

Upload one file, a document, a photo, or a spreadsheet, and give ChatGPT a specific instruction that names the purpose and the format you want back, such as a summary, a table, or a chart.



What to ask next: `Check your own extraction against the original file one more time. List anything you are less than fully confident about, such as unclear handwriting, a merged cell, or an ambiguous date format.`

Check the result

- ☐ Did you tell ChatGPT what the file was and what you wanted done with it?
- ☐ If ChatGPT extracted numbers, dates, or names, do they match the original file?
- ☐ Did you avoid uploading anything private or confidential that you should not share?

Every skill so far has lived inside a single conversation. *Chapter 1.11* shows how to save your preferences and your recurring work so you stop repeating yourself in every new chat.

Chapter 1.11 Personalize and Organize with Instructions, Memory, and Projects

You will set up ChatGPT so it remembers what you want it to remember, forgets what you would rather keep separate, and keeps a recurring piece of work organized in one place. This removes the need to repeat your background and preferences in every new conversation.

Before you begin, think of one task you return to often, such as planning meals, preparing weekly reports, or studying a subject. By the end of this chapter, you will have personal instructions set, a clear idea of what ChatGPT remembers about you, and one project ready to hold that recurring task. You will build on this project in *Chapter 1.12*.

Set durable personal instructions

Some facts about you rarely change: your job, your preferred tone, or a format you always want. Rather than repeating these in every conversation, save them once as **custom instructions**. ChatGPT applies them automatically to every new chat:

1. Select your account menu, and then **Personalization**, or you can select **Settings** and then **Personalization**.
2. Scroll down the settings to find **Custom instructions**.
3. Describe your role, your preferences, and anything ChatGPT should always avoid, and then click **Save**:

```
I'm a secondary school teacher. Keep answers concise and avoid jargon. When you
give me a lesson idea, include one way to check whether students understood it.
```

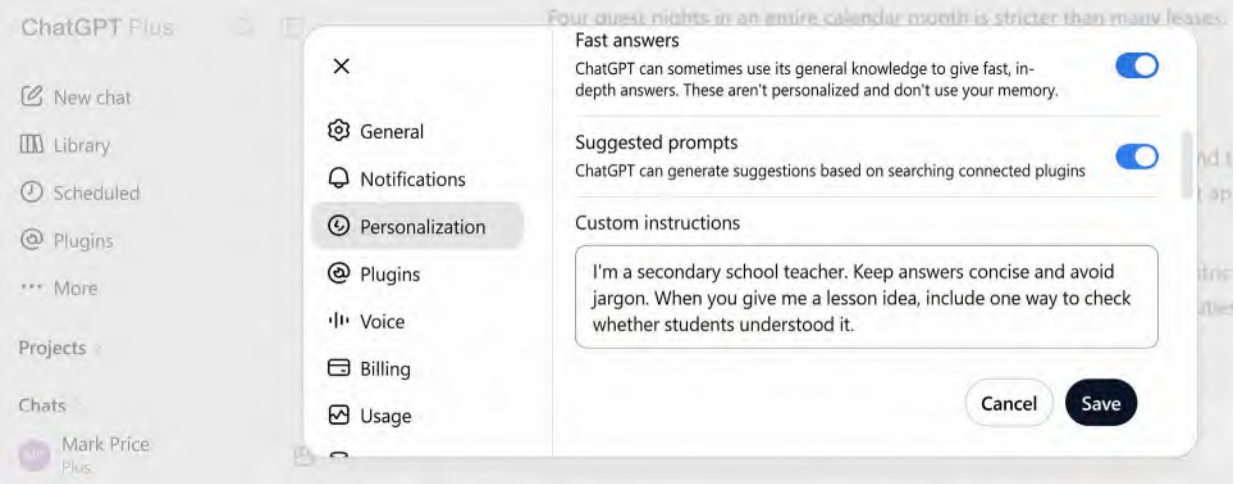


Figure 1.11A: Personalization settings and custom instructions

You should now see your custom instructions saved, and a new conversation should reflect them without you mentioning your role again.



Watch out: Custom instructions and memory apply automatically, even when you forget they are switched on. A reply that seems unusually short, informal, or oddly specific to your situation may be following a saved instruction or memory rather than your prompt. Check **Settings** | **Personalization** if a response looks unexpectedly shaped by something you didn't just say.

Decide what ChatGPT should remember

Custom instructions are rules you set once. **Memory** is different: it is made of specific facts ChatGPT picks up while you talk, plus what it can recall from your earlier conversations. Together they let ChatGPT respond as someone who already knows your situation, instead of starting from zero each time.

You can add a memory directly by telling ChatGPT to remember something:

```
Remember that I'm training for a 10K race in October and I prefer running in the
early morning.
```

Memory works best when it stays accurate. Check it occasionally and remove anything outdated or wrong.

To edit ChatGPT's memories of you:

1. Select your account menu, and then **Personalization**, or you can select **Settings** and then **Personalization**.
2. Scroll down the settings to find **Memory**.

3. In **Memory summary**, click **Manage**:

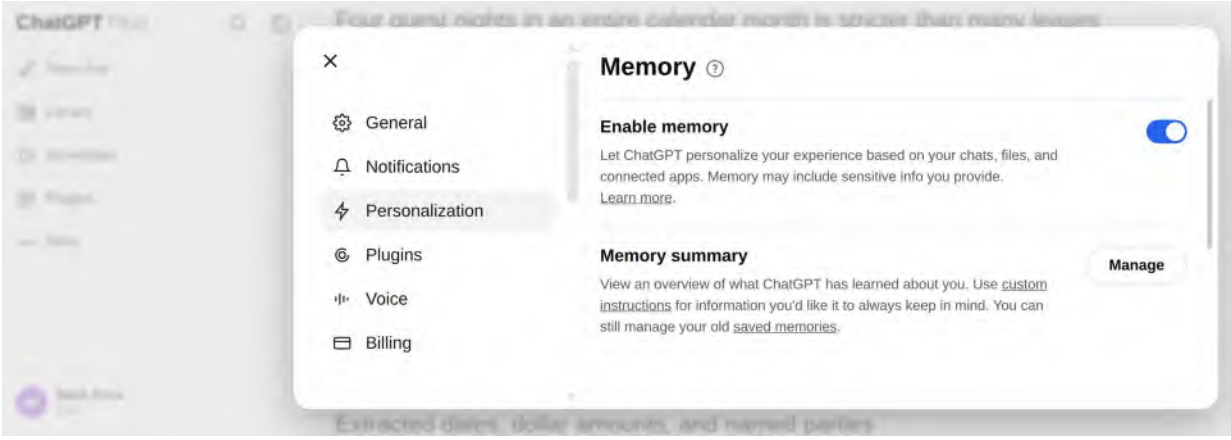


Figure 1.11B: Settings to enable and manage memories ChatGPT has about you

4. Review the memory summary about you.
5. Select some existing text, like “for many years” in my example.
6. Select **Make a correction**.
7. Type some replacement text, like “since July 2015”.
8. Select the up-arrow button to apply the change:

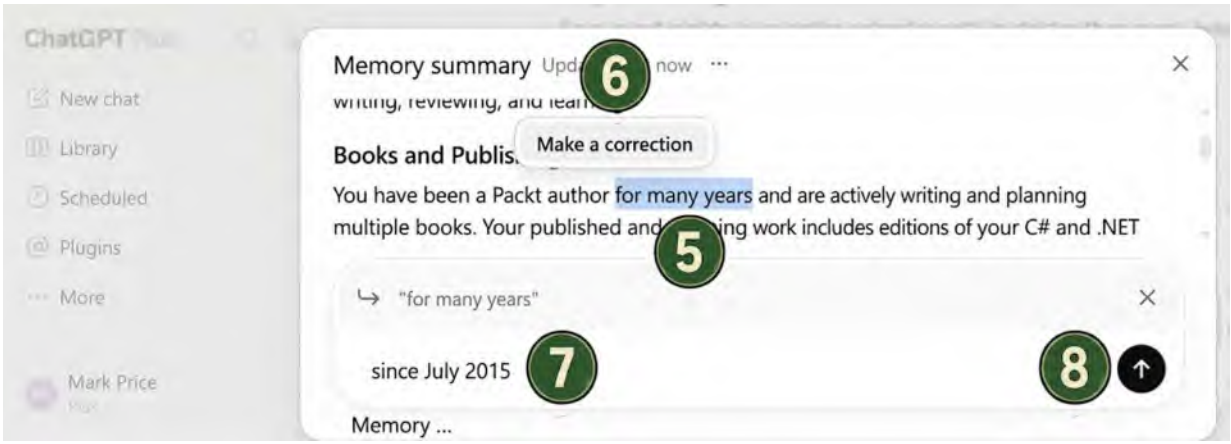


Figure 1.11C: Making a correction to the memory summary

You can also just type instructions in the **Add or update** box.

ChatGPT will reference earlier conversations automatically, picking up patterns over time even when you have not asked it to remember anything specific. You control both behaviors in your personalization settings.



Good practice: Review your saved memories every few months, the same way you would review any list that grows without you noticing. Deleting a conversation does not delete the memories it created, so remove those separately if you want them gone.

Use temporary chats for isolated work

Sometimes you want a conversation that leaves no trace in memory, such as when you are testing a sensitive question, helping someone else for a single session, or discussing something you would rather keep separate from your usual profile. A **temporary chat** does not use your saved memory, does not get added to it, and does not appear in your conversation history afterward.

Start one from the model or chat menu and choose **Temporary chat**. Use it whenever you want a clean, one-off conversation that will not shape how ChatGPT responds to you later.

Create a project with instructions and reference material

A **project** is a workspace that groups related chats, files, and instructions around one continuing piece of work. Instead of starting fresh each time you return to a task, you keep everything in one place.

1. Create a new project by clicking the + button next to **Projects** in the sidebar:



Figure 1.11D: Creating a new project

2. Give it a clear name, such as `Weekly Team Report` or `Learning Spanish`, and then click the **Create project** button.
3. On the project page, click the ... button and then select **Project settings**:

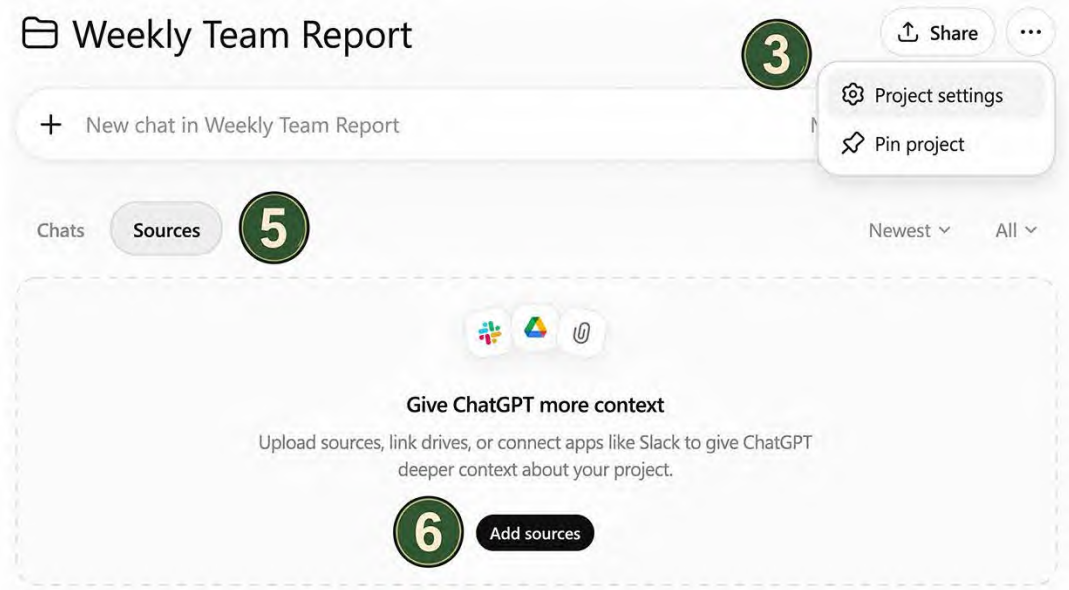


Figure 1.11E: Accessing the Project settings menu. Everything this project needs, including source reference files and past conversations, stays together.

4. Add project-specific instructions that apply only inside this project like custom instructions. You can also delete the project by clicking the **Delete project** button.
5. On the project page, select **Sources** to see the files the project has access to.

6. Select **Add sources** to upload source reference files, such as a template, a style guide, or background material, or just drag and drop files into the **Sources** area.
7. Start a conversation inside the project whenever you work on this task.
8. Select **Chats** to see previous conversations.



Learn more online: *Projects in ChatGPT*: `https://help.openai.com/en/articles/10169521-projects-in-chatgpt`.

Choose between general and project-only memory

Memory you save in a project stays inside that project. It does not appear in your main chat, and it does not carry over into other projects. Memory you save outside any project is your general, account-wide memory, and it applies everywhere.

The guidelines for choosing are:

- Use general memory for facts that are true about you no matter what you are doing, such as your job or your location.
- Use a project's own memory and instructions for details that only make sense inside that one piece of work, such as a client's preferences or a course's syllabus.

Reuse files with your Library

Every file you upload, or that ChatGPT creates for you, such as a generated document, image, or downloaded report, is also saved to your Library. Unlike a file you add inside a project, which stays scoped to that project's Sources, your Library spans every chat, project or not, so you can find and reuse a file without hunting through old conversations or uploading it a second time.

To browse it:

1. Select **Library** from the sidebar, or from your account menu if it is not pinned there.
2. Search or scroll to find a file by name or by the chat that produced it.
3. Select a file to preview it, reuse it in a new message, or remove it.

The Library is only available in the web version for now, and its rollout to more plans and countries is ongoing. If you don't see it in your sidebar, check **Settings** or the book's companion page for current availability in your plan and region.

Different plans have different storage limits. For example, Free users have 500 MB of Library storage and Plus and Business users have 20 GB of Library storage.

Files in the Library are grouped by project folder, and you can filter by file type and source (uploaded or generated). You can also search for files in the Library by name.



Good practice: The Library keeps growing as you work, the same way memory does. Delete a file once you no longer need it, especially one containing private or work-restricted information, rather than leaving it there indefinitely.



Learn more online: File storage and Library in ChatGPT: `https://help.openai.com/en/articles/20001052-file-storage-and-library-in-chatgpt`.

Try it now

Set one custom instruction about your role or preferences. Then create a project for a task you return to often, add a short instruction and one reference file, and start a conversation inside it.



What to ask next: `List everything you currently remember about me, separated into custom instructions, general memory, and this project's own memory. Tell me if anything looks outdated.`

Check the result

- ☐ Did a new conversation reflect your custom instructions without you restating them?
- ☐ Can you find, edit, and delete an item in your saved memories?
- ☐ Does your new project keep its instructions and files separate from your general chats?

The project you just created is not the finished workbench yet. *Chapter 1.12* brings everything from this Book into that project for one complete, working result.

Chapter 1.12 Project — Build Your Personal AI Workbench

You will bring together everything from *Book 1* into one finished result: a project set up around a task you actually do, ready to use again the next time that task comes around.

Before you begin, make sure you have worked through *Chapter 1.1* to *Chapter 1.11*. You will use prompting, research, file handling, and projects together in this chapter, not as separate skills. By the end, you will have a working project containing instructions, reference material, and one completed round of the task, ready to repeat the next time you need it.

Choose a task worth repeating

Pick something you expect to do again, not a one-time question. This book serves a wide range of readers, so choose whichever version fits your life right now:

Track	Good for	Example
Learn a subject	Students, career changers, curious learners	Building a study routine for a language or a certification exam
Investigate and decide	Anyone facing a recurring choice	Comparing options each time you plan a family trip
Complete a recurring assignment	Employed and independent workers	Writing a weekly status update or a client summary

Table 1.12: Tracks to choose from based on a task worth repeating

Whichever track you choose, the steps that follow work the same way. Write down, in one sentence, what a finished result looks like for this task. You will check your work against that sentence later.



Real-world example: A freelance bookkeeper chose the third track and built a project for preparing monthly client summaries. A retired teacher chose the first track and built a project for learning conversational Italian before a trip.

Create and configure the project

Now let's create and configure the project:

1. Create a new project and give it a name that describes the task, not only the subject, such as "Monthly Client Summary" rather than "Bookkeeping."
2. Open the project's instructions and describe the task in the terms you used in *Chapter 1.5*: what you want produced, who it is for, and any constraints. Set the tone and format you want ChatGPT to use every time, so you do not need to repeat it in each conversation:

```
This project is for preparing my monthly client summary. Each summary should
be under one page, written for a non-technical client, and end with one
recommended next step. Ask me for the month's figures if I have not provided
them.
```

3. You should now see a project with a name and a short set of instructions saved inside it.

Add source material and instructions

Upload anything you will need every time you do this task: a template, a past example, a style guide, or background data. Referring back to *Chapter 1.10*, give each file a clear purpose rather than uploading it without comment.

```
This uploaded file is last month's client summary. Use its structure and tone as a
model for future ones, but update every figure.
```



Sample files: If you do not have your own files then you can download sample files from the following link: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book1`

If your task depends on current information, such as prices or availability, plan to search or run research inside the project each time, as you practiced in *Chapter 1.8* and *Chapter 1.9*, rather than relying only on saved files.

Complete, review, and repeat the workflow

Now run the task for real, inside the project:

1. Start a new conversation inside the project and ask for the finished result.
2. Compare the result against the one-sentence description of success you wrote earlier.
3. Ask for specific revisions where the result falls short, using the targeted revision approach from *Chapter 1.5*.

4. Once you are satisfied, note anything you would change about your instructions or files for next time:

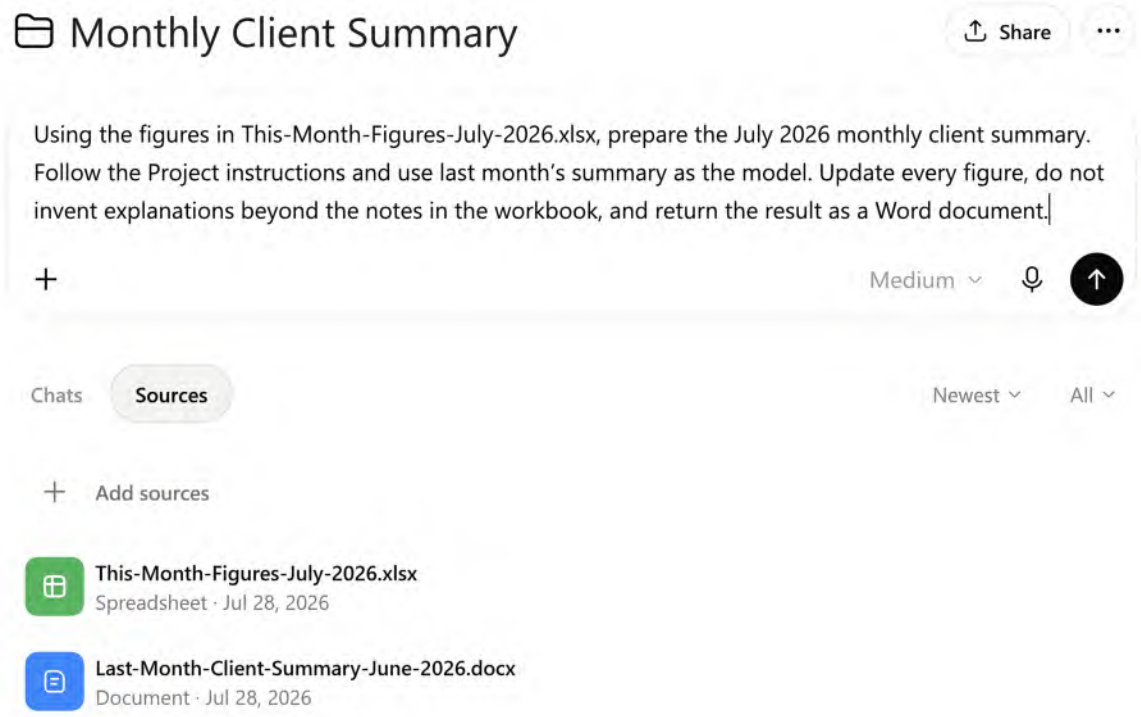


Figure 1.12A: A finished project home screen: instructions, reference material, and the conversation history for this task, all in one place

Here is the response:

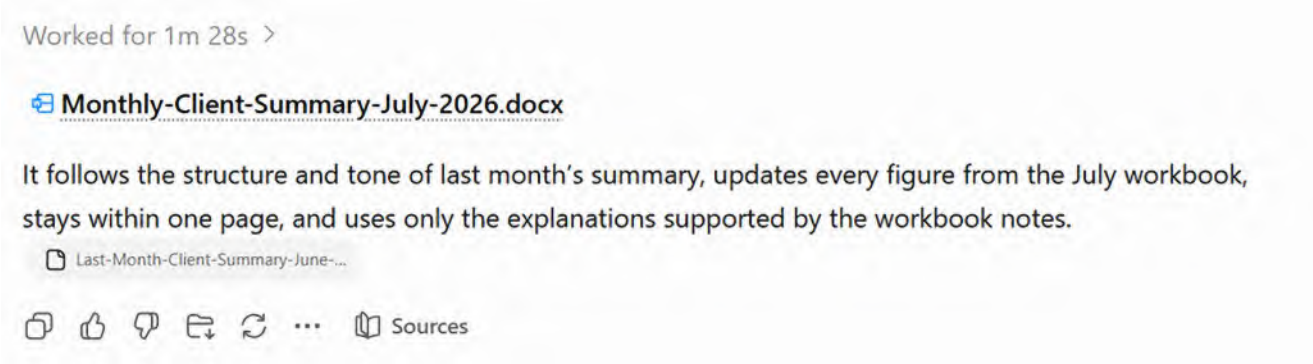


Figure 1.12B: Response from ChatGPT including Sources and downloadable document

Check the result

- ☐ Does the output match the one-sentence description of success you wrote at the start?
- ☐ Did the project use your uploaded files and instructions without you repeating them?
- ☐ Is there a small change to your instructions that would save you time next time?

The next time this task comes around, start a new conversation inside the same project instead of beginning from nothing. Update the instructions or reference files whenever your needs change.



Figure 1.12C: Your finished workbench: a repeatable loop of inputs, context, request, review, and output, ready to run again whenever this task returns

You now have a working example of the pattern this book will keep returning to: a task, organized context, a clear request, and a human review before you rely on the result.

The next step

Book 2, *Create Real Work Results*, builds on this pattern with ChatGPT Work, longer assignments, and scheduled tasks. Keep this project. You may extend it rather than replace it.

Book 2

Create Real Work Results

In *Book 2* you interact with AI as a coworker to produce documents, spreadsheets, presentations, analysis, images and other media, and you learn to supervise longer, multi-step assignments.

It has nine chapters, including an end-of-book project:

2.1 Use ChatGPT Safely and Responsibly at Work 67

2.2 Choose the Right Mode for Each Task 71

2.3 Draft Professional Documents 78

2.4 Summarize and Extract Information 84

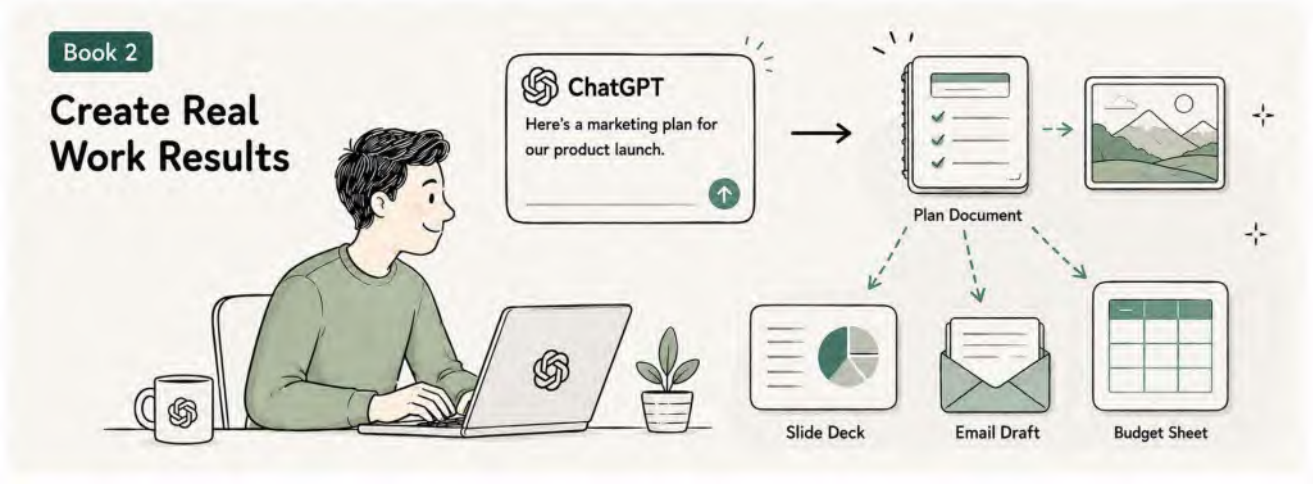
2.5 Analyze Spreadsheets and Data 88

2.6 Create and Edit Images and Media 93

2.7 Create Presentations and Visual Reports 100

2.8 Add Plugins, Apps, and Skills 106

2.9 Project – Produce a Useful Weekly Briefing 112



Chapter 2.1 Use ChatGPT Safely and Responsibly at Work

Book 1 showed you how to use ChatGPT for personal questions and everyday planning. This chapter marks a change. From here on, many of your conversations will involve work: client information, company documents, or decisions that affect other people.

Before you draft a single report, you need a few ground rules. This chapter explains how to use ChatGPT at work without putting your job, your employer, or your customers at risk. You don't need any special settings to follow it. You just need your employer's policy and a habit of checking your work.

By the end of this chapter, you'll know what you can safely type or upload, when you need to say that AI helped you, and when a task needs a qualified professional instead of ChatGPT.



Watch out: All chapters in *Book 2* require an individual paid plan or organization-managed workspace.

Follow your employer's AI policy

Many organizations now have a written policy that covers AI tools. It may list approved products, blocked types of content, and required disclosures.

Find your employer's policy before you use ChatGPT for work tasks. Ask your manager or IT department if you can't find one. If no policy exists yet, treat the rules in this chapter as a reasonable default, and check with your manager before you handle sensitive material.

A policy can change. Review it again if your employer switches AI vendors, updates its terms, or announces new rules.



Watch out: If your organization's policy conflicts with anything in this chapter, follow the policy. It reflects decisions made for your specific workplace.

Know the difference between a personal account and an approved workspace

A personal ChatGPT account belongs to you. An approved business workspace belongs to your employer and is managed by an administrator. The two can look similar, but they follow different rules for data storage, retention, and access.

Check which account you're signed into before you start work. Look at the account name or workspace label in the corner of the screen. If your employer hasn't set up a workspace, ask before you use your personal account for company material.

Some organizations allow personal accounts for low-risk tasks, such as brainstorming or learning, but require an approved workspace for anything involving customer or company data. Follow your organization's rule, not a general guess.



Learn more online: OpenAI's help center article *What is a ChatGPT Enterprise workspace? How can I switch workspaces?* explains how a managed workspace differs from a personal account and how to switch between them with screenshots: `https://help.openai.com/en/articles/8265430-what-is-a-chatgpt-enterprise-workspace-how-can-i-switch-workspaces`.

Keep confidential and regulated data out of ChatGPT without permission

ChatGPT can help with many tasks, but some information should never go into a prompt without your employer's permission. This includes passwords, private keys, unpublished financial results, health records, personal data about customers or colleagues, and anything covered by a nondisclosure agreement.

One way to decide what's safe to share is to sort information into three groups:

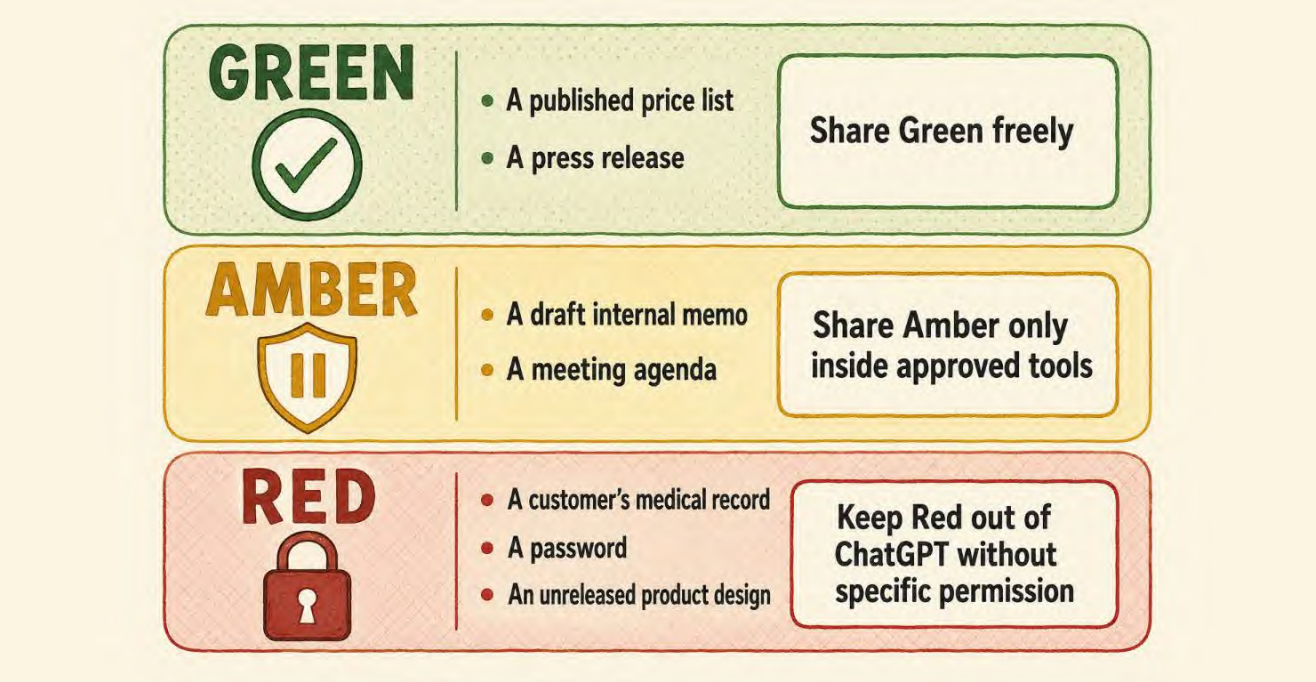


Figure 2.1A: Sort information into these three groups before you share it with ChatGPT. When you're not sure which group applies, treat it as red.

Treat red information as off-limits unless your employer has approved a workspace built to handle it. When in doubt, ask before you paste.

Real-world example: A marketing manager pastes a client's campaign brief into ChatGPT to draft social media captions. Before she sends the brief, she removes anything classed red (client contact details and budget figures) and amber (specific client names), because her agency's data classification treats these as confidential. She keeps the green data, like general messaging goals and target audience, which is safe to share.

Watch out: Pasting an email thread, a spreadsheet, or a screenshot can carry hidden information you didn't intend to share, such as other people's names, phone numbers, or comments. Check the full content before you upload it, not only the part you meant to use.

Verify facts, calculations, quotations, and citations

ChatGPT can produce a confident-sounding answer that contains an error. This can happen with a date, a statistic, a quotation, or a calculation. Treat every factual claim as a draft until you check it.

Compare important numbers against your source spreadsheet or system of record. Look up quotations in the original document. Recalculate totals that will appear in a report or an invoice. This check is most important for material other people will rely on, such as a client proposal or a board report.



Good practice: Ask ChatGPT to show its sources or its working before you accept a number or a quotation. A response with no source is harder to check.

Disclose AI assistance when it's required

Some organizations, clients, or professional bodies require you to say when AI helped produce a piece of work. A contract, a regulator, a journal, or a school may set its own rules.

Check your employer's policy, your contract, and any relevant professional or industry rules before you submit AI-assisted work. If you're unsure, ask rather than guess. Disclosure requirements are becoming more common.

Retain human approval for consequential actions

ChatGPT can now do more than answer a question. Work and connected tools can send messages, edit files, and complete tasks on your behalf. That capability makes a review step necessary before anything happens that can't easily be undone.



Figure 2.1B: ChatGPT's output is one step in a process you still control. Pause at an approval point before any action that can't easily be reversed.

Before you let ChatGPT act on your behalf, check what it plans to do, and approve it. Require this step for any action that sends a message, publishes content, spends money, shares data, or changes a system that other people depend on.

Know when a qualified person needs to be involved

ChatGPT can help you understand a medical, legal, financial, employment, or safety question, but it can't replace a licensed doctor, lawyer, accountant, HR professional, or safety officer. These fields carry consequences that go beyond an incorrect document, such as a missed diagnosis, a broken contract, or an unsafe decision.

Use ChatGPT to prepare for these conversations. Ask it to help you list questions, organize your situation, or understand unfamiliar terms. Then bring your findings to the right professional before you act.

Try it now

Find your employer's written AI policy or ask your manager whether one exists. Note three things it allows and three things it doesn't. If no policy exists, write down which rules from this chapter you'll follow by default, and confirm your list with your manager.

Check the result

- ☐ Can you explain, in one sentence, what your organization treats as confidential?
- ☐ Do you know who reviews AI-assisted work before it reaches a client or the public?
- ☐ Have you identified one task in your role where a qualified professional, not ChatGPT, should make the final call?

You now know what you can safely share with ChatGPT at work and when to bring in a qualified professional instead. The next question is where to have that conversation. ChatGPT offers three modes built for different kinds of work, and picking the right one before you start will save you time throughout the rest of this book.

Chapter 2.2 Choose the Right Mode for Each Task

ChatGPT brings three modes together in one application: **Chat**, **Work**, and **Codex**. Each mode suits a different kind of task. Picking the right one before you start saves time and produces a better result.

This chapter will help you decide which mode to open. You should already be comfortable having a basic conversation with Chat. By the end, you'll be able to look at a task and know within seconds whether it belongs in Chat, Work, or Codex.

Use Chat for questions, focused help, and single-step deep research

Chat is the mode most readers already know. Choose it when you:

- Have a single question
- Want to talk through an idea
- Need a short piece of writing, an image, or a quick analysis

Chat works well for tasks you can finish in one sitting, perhaps by adding a few files, but without connecting outside systems or accounts. Asking for an explanation, comparing two products, drafting a short email, or improving a paragraph all belong here.

Chat responds quickly and keeps the conversation in view, so you can revise your request in a few back-and-forth messages. Use it when you expect to stay closely involved in every step.

In Chat, **Deep research** is a distinct tool you switch on for one message. It runs a single multi-source research pass and hands back one structured, cited report, then you're back to a normal chat. Compared to Chat, Work doesn't need a separate "turn on deep research" step because multi-step, multi-source research is just how Work operates by default.

Use Work for multi-step research and finished deliverables

Work is built for larger assignments. Choose it when a task:

- Needs research across several files, sources, and steps
- Draws on files or connected accounts, such as email, calendars, or shared drives
- Should end in a finished image, document, spreadsheet, presentation, or report
- Will take more than a few minutes of back-and-forth to finish in Chat

Work can plan a multi-step assignment, carry it out over several minutes, and show you its progress along the way. You can step away and check back, rather than watching every message appear, or get on with other work until you get a notification that it's finished.

Recognize when a task has grown too big for Chat

Some tasks start small and grow. A single question about a competitor turns into a request to compare five companies across ten criteria, using data from three files. At that point, Chat is no longer the right tool.

Watch for these signs that a task belongs in Work instead:

- You've asked several related questions in a row and are stitching the answers together yourself
- You keep re-explaining the same background information
- You need ChatGPT to check multiple sources before it answers
- The result should be a shareable file, not a chat message

You don't have to start over. Most Work assignments begin with the same plain description of what you want that you'd type into Chat. The difference is what happens after you send it.

Starting a large, multi-source assignment in Chat can leave you copying and pasting between messages for far longer than the task should take. If a task keeps growing while you work on it, stop.

If you're working inside a project, switch to Work in that same project and start a fresh Work chat so your files and context carry over and you're not starting from zero. If you're not in a project, start a new project first, or move the chat into an existing project rather than pushing through in Chat.

Know when a task belongs in Codex instead

Codex is the mode for building and changing software. Choose it when your task involves writing or editing code, working with a repository, or running tests.

Figure 2.2A shows the mode selectors in the ChatGPT desktop app:

1. **Chat**: For questions and focused help
2. **Work**: For research and finished deliverables
3. **Codex**: For building and changing software



Figure 2.2A: The mode selector in the ChatGPT desktop app

Codex is a desktop-only experience, by design. On the desktop app, ChatGPT and Codex sit as two separate top-level views: you pick ChatGPT or Codex from the top-left menu, and Chat and Work only live inside the ChatGPT view as a toggle. Codex is not selectable on the web at `chatgpt.com` or in the mobile app at all. If you're on a phone, the closest you get is the **Remote** tab, which lets you check on a Codex chat you started on desktop, but that doesn't create web or mobile chat history the way Chat and Work do.

Chat and Work are meant to work the same way across web, mobile, and desktop, with chats syncing between them. Codex stays desktop-bound because it works with local folders, repositories, terminals, and other developer tools that only make sense on your own machine.

This book does not cover Codex at all. If you want to learn how to use ChatGPT to build websites or mobile apps, the separate book *ChatGPT Codex Visual Bible* covers that software development work.



Watch out: In a late-July 2026 interview with Joanna Stern, when she called the new ChatGPT desktop app "kind of a mess," Greg Brockman, President of OpenAI, agreed: "It's kind of a mess. We agree. And so what you're seeing is incremental progress towards the future." He said OpenAI plans to remove the separate Work mode by the end of 2026, folding its features directly into ChatGPT with a simpler, tab-free interface.

Work models and Intelligence levels

Chapter 1.4 introduced the thinking effort control in Chat with its up to five levels: **Instant**, **Medium**, **High**, **Extra High**, and **Pro**, most running on GPT-5.6 Sol. Work has a similar model control, but with finer choices. You pick which specific model works on the assignment, not only how hard it thinks.

By default, the model control shows a simplified view with just a slider from **Terra Light** to **Sol Extra High**, and a lightning bolt button to toggle **1.5x speed**:

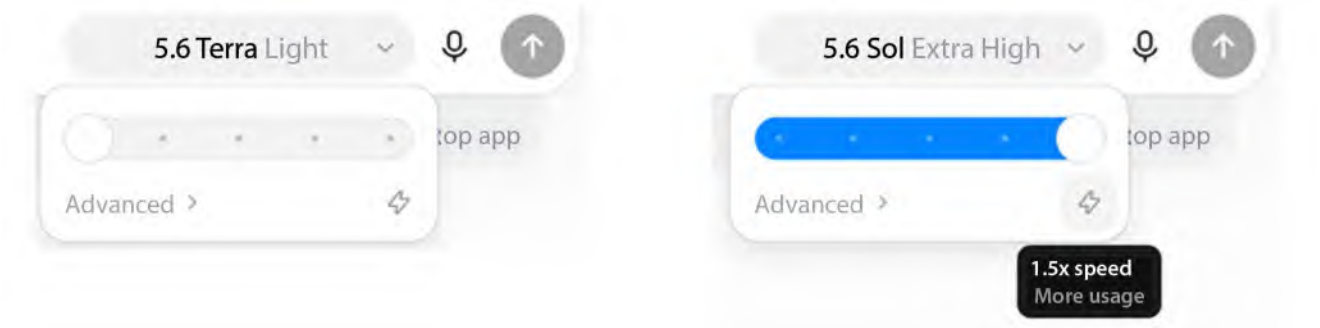


Figure 2.2B: Simplified model selector in Work mode

Click **Advanced** to see the full model, effort, and speed options:

1. **Model**: Sol, Terra, Luna, or legacy models.
2. **Effort**: Light, Medium, High, Extra High, or Max.
3. **Speed**: Standard or Fast 1.5x.

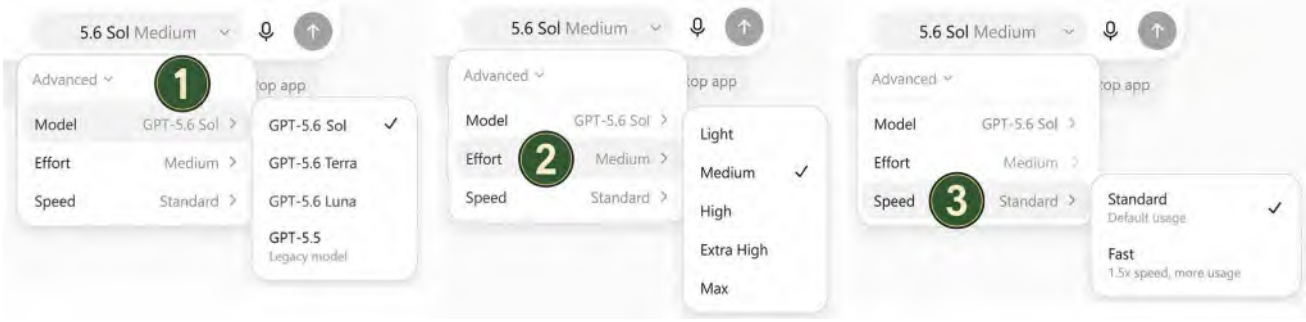


Figure 2.2C: Pick the model first, then set how hard and fast it should think

Each model trades capability for speed and cost:

Model	Best for	Plans that include it
Sol	The hardest or highest-stakes assignments: long research, difficult writing, or a task with many steps	Plus, Pro, Business, Enterprise
Terra	Everyday assignments: a routine report, a standard spreadsheet analysis, most day-to-day work	All plans, including Free and Go
Luna	Simple, high-volume tasks you run often, such as short checks or routine summaries	Plus, Pro, Business, Enterprise

Table 2.2: Models available in Work mode

Once you've picked a model, set its thinking effort using the same scale from *Chapter 1.4*: Instant, Medium, High, Extra High, or Pro. A lower setting suits a routine step in a longer assignment. A higher setting suits a step that needs careful reasoning, such as reconciling numbers from several sources.

Free and Go plans work with Terra only in Work. Plus, Pro, Business, and Enterprise plans can choose any of the three models and set the thinking effort for each one.



Watch out: Setting every step to Sol at its highest effort level can use up your plan's usage allowance faster than the assignment needs. That is the equivalent of Anthropic's Claude Fable 5. Reserve that combination for the steps that require it. Start most Work assignments with Terra. Switch a specific step to Sol only when it turns out to be harder than expected, or when the result will support a decision worth extra care.

Complete multi-step assignments and supervise computer actions

ChatGPT Work is intended for longer assignments involving research, connected files and apps, and finished documents, spreadsheets, presentations, and reports. This section covers defining the deliverable, letting Work plan and execute the assignment, staying involved along the way, and reviewing what it actually completed.

The common seven steps for longer assignments should include:

1. **Define the final deliverable:** State what you want at the end: a specific document, spreadsheet, presentation, or report, not just a general area to investigate. A clear deliverable gives Work something concrete to plan toward.
2. **Let Work create and execute a plan:** Once you describe the assignment, get Work to propose a plan before carrying it out. Read the plan before approving it. This is your chance to catch a misunderstanding before Work spends time on the wrong task.
3. **Answer questions and redirect progress:** Work may ask you a question partway through, or you may notice it heading in the wrong direction. Answer promptly and redirect clearly rather than waiting until the end to fix a problem that grew along the way.



Real-world example: A consultant hands Work a request to compare three vendors and produce a one-page recommendation, then reviews each research step before approving the final document rather than reading only the finished summary.

4. **Distinguish assistance from autonomous action:** Some steps only produce a draft for you to review. Others, such as sending a message or changing a shared file, act outside the chat and can affect other people. *Figure 2.2D* shows this range as a ladder, from a suggestion you review yourself to an action taken on your behalf:



Figure 2.2D: Know which rung a given step sits on before you approve it. A suggestion and an action carry very different levels of risk.

5. **Review logins, permissions, and sensitive data:** Before approving a step that needs a login or access to a file, check what permission it's requesting and whether that permission matches the task. Do not grant access to more accounts or data than the assignment needs.
6. **Require approval before consequential actions:** *Figure 2.2E* traces a full assignment from goal to deliverable, with review built in at each stage:

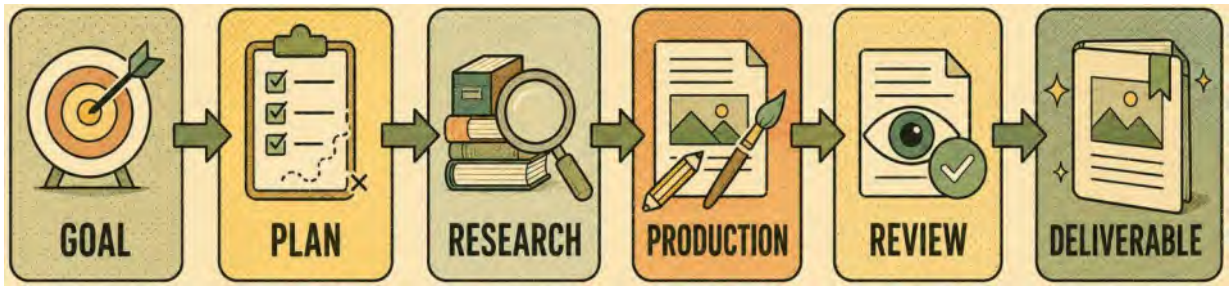


Figure 2.2E: Review sits before the deliverable, not after. Catching a problem at this stage is easier than fixing a finished document.



Watch out: An action that sends a message, publishes content, spends money, or changes shared data cannot always be undone. Require your own approval before any such step, even if Work presents it as routine.

7. **Check what the system actually completed:** When Work reports that it finished an assignment, open the actual deliverable and compare it against your original request, rather than trusting the summary of what it did. A report that says a step succeeded is not the same as confirming the result is correct.

Good practice: Keep a short record of what you approved during a long assignment. If a result later looks wrong, this record helps you find which step introduced the problem.

Know when a Work assignment needs you back

A Work assignment can run for several minutes without you watching every step, so the ChatGPT desktop app gives you two ways to notice when it needs you: system notifications and an optional floating pet. (The web interface also has notification settings, but these are more limited.)

1. In **Settings**, select **General**.
2. In **Notifications**, select whether a turn-completion alert appears never, only while the app is in the background, or always.
3. Separate switches control alerts for a permission request and for a question Work is waiting on you to answer:

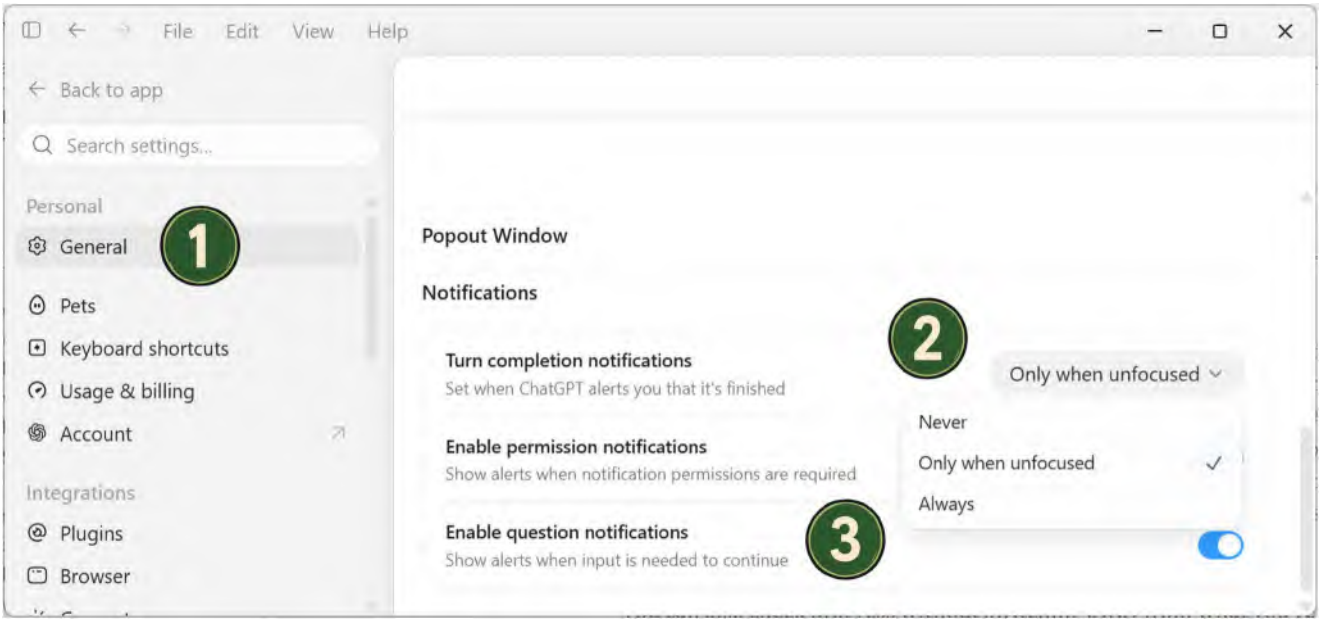


Figure 2.2F: Notification settings.

Your operating system may also ask permission to let the ChatGPT desktop app show notifications; grant it if you want an alert outside the app window.

You can also check progress without switching away from what you're doing. Select the bell icon in the sidebar to open Activity view, which lists chats that are unread, still running, or waiting for your response.

Good practice: Turn on notifications before you start a longer assignment, not partway through it. Missing the one moment Work pauses for your approval, like the consultant's vendor comparison earlier in this chapter, is the most common reason a supervised assignment stalls without you noticing.

A pet is an optional animated companion that floats above your other windows and shows what an assignment is doing, so you don't need to switch back to check. Open the profile menu at the bottom of the desktop app, or go to **Settings** | **Pets**, and choose one.

Use the command menu's **Show pet** option to bring it onto your screen; select **Hide pet** to hide it.

A pet's appearance changes to match what the assignment behind it is doing:

- **Running**: Work is still being carried out according to the plan.
- **Needs input**: Work is waiting on your approval or an answer, the point where you should stop and review before it continues.
- **Ready**: the assignment has finished and has activity you haven't seen yet.
- **Blocked**: the assignment failed or hit an error.

If more than one chat has activity, the pet shows you the one that needs you most: a chat needing input first, then a blocked one, then a ready one, then one still running. Select the pet to jump back to ChatGPT, or open its activity tray to choose a specific chat.



Watch out: A notification banner or an on-screen pet can reveal part of an assignment, such as a document title or a step Work just completed, to anyone glancing at your screen. Treat this the same as any other Work output: keep alerts to permission and question notifications only if you're handling anything confidential or working on a shared screen.



Learn more online: OpenAI's documentation on notifications and pets in the ChatGPT desktop app covers current settings, activity view controls, and available companions: `https://learn.chatgpt.com/docs/notifications` and `https://learn.chatgpt.com/docs/pets`.

Try it now

In a current conversation, find the model selector and, if available, the reasoning-effort control.

Pick a task you'd normally split across several Chat messages, such as researching three options for a purchase and writing up a short comparison. Open Work and describe the whole task in one message: what you want to compare, what sources are acceptable, and what the final comparison should include. Watch Work outline its plan before it begins.



What to ask next: `Based on what I just described, would this task work better in Chat or Work?`

Check the result

- ☐ Can you locate the model and reasoning effort controls without help?
- ☐ Does the plan Work proposed match what you actually asked for?
- ☐ Are the sources it used ones you'd trust for this decision?
- ☐ Is the finished deliverable something you could hand to someone else without further editing?



Learn more online: OpenAI's help center documents the current Chat, Work, and Codex modes and what each one includes on your plan. Check it before you rely on the exact controls shown in *Figure 2.2B*: `https://help.openai.com/en/articles/20001275-chatgpt-work-and-codex`.

Once you know which mode fits your task, you're ready to put it to work. *Chapter 2.3* starts with one of the most common jobs you'll hand to ChatGPT: turning a blank page into a finished professional document.

Chapter 2.3 Draft Professional Documents

ChatGPT can produce a full first draft of a report, proposal, letter, or policy in minutes. But the draft is only as useful as the plan behind it, and a first draft is rarely the final version.

This chapter walks through seven steps: define the reader and purpose in a brief, build an outline, generate a complete draft, get useful feedback, keep the document sounding like you, check its claims, and finally produce one controlled final version and move it into the format you need:

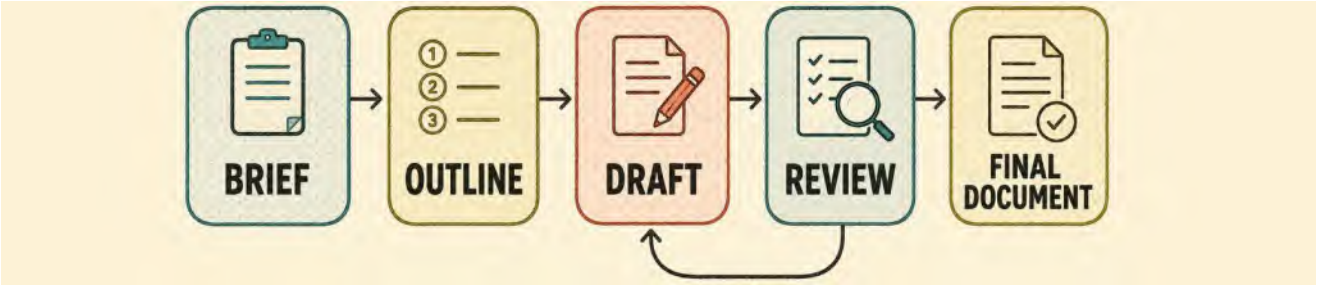


Figure 2.3A: Treat the outline as a checkpoint. Catching a structural problem here is faster than catching it in the final draft.

You need a clear idea of the document you want to produce and a rough sense of its audience. By the end of this chapter, you'll be able to turn a document request into a finished draft ready for review.

Define the reader and purpose

Before you ask ChatGPT to write anything, decide who will read the document and what you want that reader to do after reading it. A status update for your manager calls for a different length and tone than a proposal for a client or a policy for new staff.

Tell ChatGPT the reader, the purpose, and any constraints you already know, such as a required length, format, or deadline. A short brief like this does more to improve the draft than a longer, vaguer request.

Start by asking ChatGPT to help you define the document before writing it:

```
I need to write a short professional report about [project or topic].

Help me define:
- Who will read it
- What the reader needs to know
- What I want the reader to do after reading
- The required length, tone, and deadline

Ask me no more than five short questions. Then turn my answers into a brief I can
use to plan the report. Do not write the report yet.
```



Good practice: State the reader and purpose in one or two sentences before you describe the content. For example: "This is a one-page update for my regional manager. She wants to know whether the project is on schedule and what she needs to approve this week."

Build an outline before drafting

Ask ChatGPT for an outline before asking for the full document. An outline is quicker to check than a complete draft. Fixing a missing section or a wrong order is easier now than after the document is written:

```
Using the attached report brief, create an outline for the report.

Use no more than six sections. For each section, give:
- A short heading
- One sentence explaining what the section will cover

Put the information in the order that will be most useful to the reader. Do not
write the report yet. If the brief is missing something important, tell me instead
of inventing it.
```

Review the outline against three questions:

1. Does it cover everything the reader needs?
2. Does it include anything the reader doesn't need?
3. Does the order make sense?

Ask for changes before moving on.

Generate a complete first draft

Once the outline looks right, ask ChatGPT to write the full document from it. Ask for one complete first draft rather than one section at a time. This keeps the tone and level of detail consistent from start to finish:

```
Using the attached report brief and the outline we agreed on, write the complete
first draft.

Keep it to one page. Use plain, professional language, short headings, and bullet
points where they make the report easier to scan. Use only facts from the brief. If
any necessary fact is missing, mark it [VERIFY] instead of inventing it.

Create the report in a writing block so I can review and export it.
```



Sample files: If you don't have a real document ready to draft from yet, download `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2/2-3-sample-brief.docx`. It contains a short, filled-in project brief you can add to your chat.

Watch out: If you've used an older ChatGPT guide, article, or tutorial, you may be expecting a feature called Canvas. OpenAI removed Canvas from its current models in mid-2026 and folded its editing tools into the writing blocks used throughout this book.

Figure 2.3B shows a response with a document in a **writing block** that lets you edit text directly, select text for revision, ask ChatGPT to revise the whole draft, and undo or redo recent AI-assisted edits. Availability can vary by account, device, model, and rollout:

1. Click **Edit** and then enter a prompt to request changes.
2. The toolbar includes buttons to: **Undo**, **Copy**, **Download**, and **Open editor** (meaning expand the editing area). **Download** provides formats like **PDF document** and **Microsoft Word**.
3. Select text to pop up a floating toolbar to: **Ask for changes**, create a link, toggle bold or italic, and apply a **Text** style like a heading level or list type:

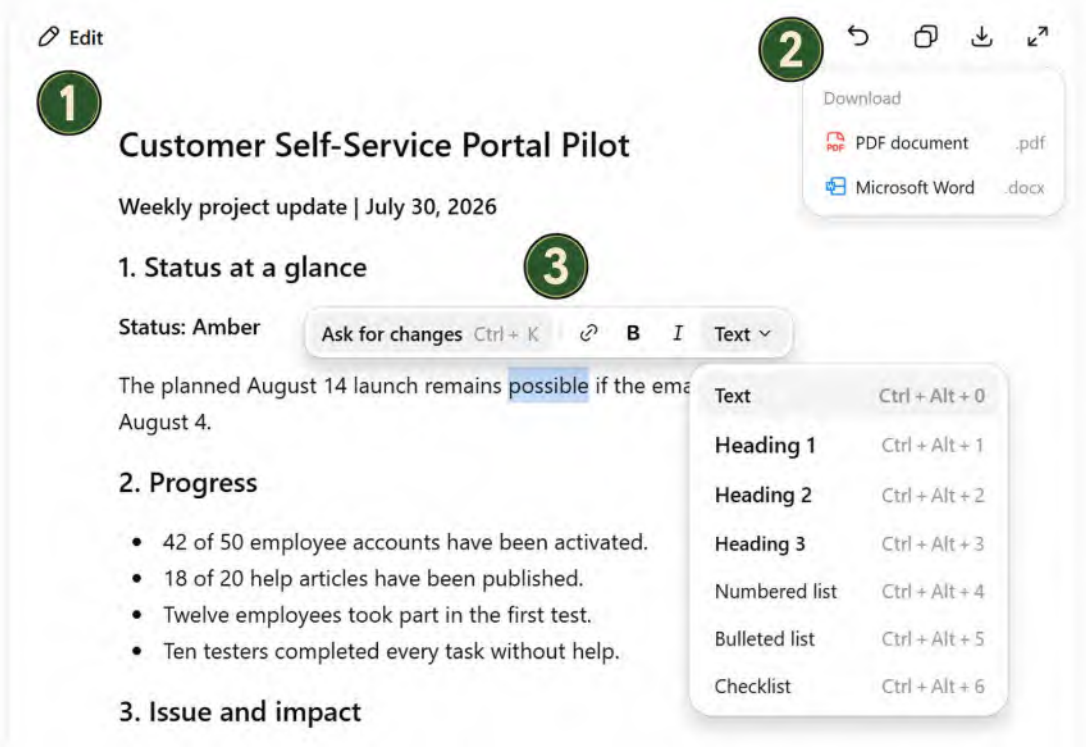


Figure 2.3B: The document's first draft in a writing block

Read the whole draft before asking for changes. Check that it matches the outline, states facts you can confirm, and sounds like something you'd be willing to put your name on.

Real-world example: A nonprofit program officer drafts grant reports in ChatGPT by giving it the funder's required sections and word count up front, then checking every dollar figure against the program's own budget spreadsheet before submitting.

Watch out: A fluent draft can still contain wrong figures, invented sources, or overstated claims. Confirm any number, date, or quotation before it goes to a reader who will act on it.

Ask for structural rather than vague feedback

A request such as "make this better" produces a vague response. Ask instead about specific parts of the document: the argument, the order of sections, the opening, or the ending.

Name what you want checked. Ask whether the document makes its point in the first paragraph. Ask whether each section follows logically from the one before it. Ask whether the conclusion matches the evidence in the body. These questions produce feedback you can act on.

There are four kinds of review you can ask for, one at a time: accuracy, structure, tone, and presentation:

1. **Accuracy**: whether the facts, figures, dates, and claims are correct. Ask, for example, "Does every dollar figure in this proposal match the totals in the budget table?"
2. **Structure**: whether the sections follow a logical order and support the document's main point. Ask, for example, "Does the evidence in section two support the conclusion in section three?"
3. **Tone**: how the document will sound to the person reading it. Ask, for example, "Does this message sound confident without sounding demanding?"
4. **Presentation**: whether headings, spacing, and formatting stay consistent throughout. Ask, for example, "Are the heading levels and bullet styles the same in every section?"

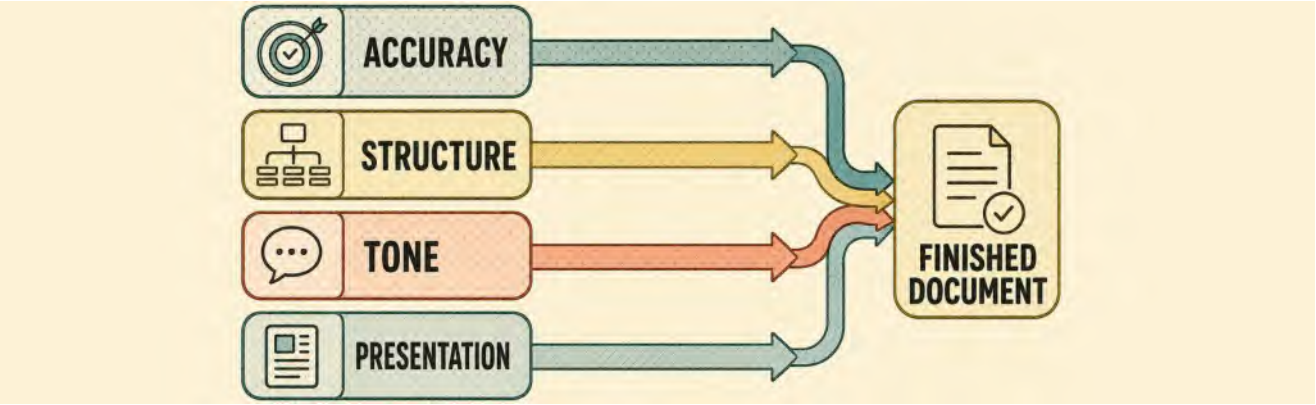


Figure 2.3C: Ask for one kind of review at a time. Mixing all four into a single request makes the feedback harder to act on.



Good practice: Ask one structural question per message. Reviewing the argument, then the section order, then the opening gives you three focused answers instead of one crowded one.

Preserve the intended voice

A revision should still sound like you or like the person the document represents. Tell ChatGPT which parts of the tone and wording to keep before asking for changes.

If a revision changes more than you wanted, point out what changed. Ask ChatGPT to restore it while keeping the rest of the improvement. This is faster than starting the revision over.



Real-world example: A small-business owner asks ChatGPT to review a customer email for tone before sending it, specifically checking whether it sounds apologetic without sounding defensive.



Watch out: A revision aimed at clarity can shift the tone from formal to casual without you noticing, or the other way around. Read the revised version aloud, or have a colleague read it, before you send it.

Check claims and internal consistency

Before you accept a revision, check every fact, figure, date, and quotation against a source you trust. ChatGPT can write confidently about a claim that is wrong, so a fluent sentence is not proof that the sentence is accurate.

Also check the document against itself. A number given in one section should match the same number used later. A term defined early in the document should keep the same meaning throughout.

Export or transfer the finished content

Once you have accuracy, structure, tone, and presentation feedback, ask for one final revision that applies the changes you approved. Naming exactly which changes to apply keeps ChatGPT from introducing new, unreviewed edits along the way.

A finished draft still needs to reach its destination: a shared document, an email, a slide deck, or a printed page. Copy the finished text into your usual word processor. Or use the download option if ChatGPT offers a direct download in the format you need. Check the formatting after moving the document, since headings, lists, and tables don't always carry over exactly.

Read the final version in full before sending it. A controlled revision still deserves the same check as the first draft.

Know what an Office extension can see and do

Word, Excel, and PowerPoint each have an official ChatGPT extension: an add-in that opens a ChatGPT panel inside the file you already have open, rather than sending you to `chatgpt.com` to upload a copy. *Chapters 2.3*, *2.5*, and *2.7* introduce each one in its own context. This section covers what they have in common before you meet them chapter by chapter.

An Office extension signs in with the same ChatGPT account you use everywhere else, so it follows your existing plan, workspace, and administrator settings rather than opening a separate trial. Once you're signed in, the extension can read the content of the open file and the text you type into its panel, and it uses both as the basis for its response.

An Office extension has real limitations worth knowing before you rely on it for a live document. Its conversation runs separately from your regular ChatGPT chat history and memory, so it can't recall an earlier conversation from the main app. It works with the file that's open, not with files you haven't loaded. And some advanced, application-specific features, such as Word tracked changes and styles, Excel macros and VBA, or PowerPoint's finer formatting and animation controls, may not be fully supported yet.

Treat an Office extension the same way you'd treat ChatGPT in a browser tab. Everything earlier in this chapter still applies: confidential material still needs your employer's permission before it goes into the panel, and any change the extension makes to your document, workbook, or deck is a draft that needs your review, not a finished result.



Watch out: Installing an Office extension also brings you under Microsoft's terms for that add-in, in addition to OpenAI's. Microsoft's Add-In Marketplace terms allow Microsoft to read the content of the file, attachments, and prompts you send through the panel. Factor that into what you're willing to open in a connected file.

Exactly which plan, workspace, and administrator setting unlock each extension changes often enough that it doesn't belong in print. *Chapters 2.3*, *2.5*, and *2.7* each point you to the current, official answer for the application covered in that chapter.



Learn more online: For Word, OpenAI has not yet published a standalone extension to match Excel and PowerPoint; a Business or Enterprise administrator deploys the ChatGPT add-in for Office from Microsoft's marketplace, which adds a similar panel directly inside Word. Check the listing for current requirements before asking your IT team to deploy it: `https://marketplace.microsoft.com/en-us/product/office/WA200010215`.

Try it now

Choose a short document you actually need, such as a one-page project update or a short proposal. Tell ChatGPT the reader and purpose, ask for an outline, review it, then ask for the full draft. Revise one section using a targeted instruction rather than requesting a new draft.

Or take a draft you already have. Ask ChatGPT to review its structure only: does the argument hold together, and does each section follow from the one before it. Then ask a separate question about tone. Compare the two sets of feedback before deciding what to change.



What to ask next: `List every change you made in the final revision. Confirm each one against my approved feedback.`

Check the result

- ☐ Does the draft match the reader and purpose you described?
- ☐ Can you confirm every fact, figure, and quotation in the draft?
- ☐ Does the document read as one consistent piece, rather than several different voices stitched together?
- ☐ Does the feedback address the specific question you asked, rather than the whole document at once?
- ☐ Can you confirm every fact, figure, and quotation that changed during revision?
- ☐ Does the revised document still sound like the intended voice?



Learn more online: OpenAI's help center article on writing blocks covers the current download formats and controls for a document ChatGPT drafts for you: `https://help.openai.com/en/articles/20001246-working-with-writing-blocks-and-code-blocks-in-chatgpt`.

Revising a document you already have is one skill. Making sense of documents other people wrote is another. *Chapter 2.4* turns to the meeting notes, email threads, and reports piling up in your inbox, and shows you how to pull out exactly what you need.

Chapter 2.4 Summarize and Extract Information

Meeting notes, long email threads, and reports often bury the few facts you actually need. This chapter covers four steps: summarize for a specific reader, pull out actions and dates, compare several documents, and keep a link back to the source behind every claim.

You need one or more documents, transcripts, or sets of notes to work with. By the end of this chapter, you'll be able to turn a pile of unstructured material into a short, structured result you can act on.



Sample files: If you don't have your own set of notes to work with, download this document: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2/2-4-sample-meeting-notes.docx`. It contains three short sets of fictional meeting notes with a few actions and dates buried in each.

Summarize for a stated audience

A summary written for your own reference looks different from one written for a manager who wasn't in the meeting. Tell ChatGPT who will read the summary and what they need to know before asking for one. Ask for the length you want, such as three bullet points or one short paragraph. A summary that runs longer than the reader has time for defeats its own purpose:

```
Using the attached meeting notes, write a summary for a manager who did not attend
the meeting.

The manager needs to know:
- What was decided
- What action is required next
- Any deadline or risk that needs attention

Use exactly three concise bullet points. Lead with the information that requires the
manager's attention. Use only facts stated in the notes. If something is unclear or
missing, say so instead of guessing.
```



Good practice: Name the reader's next action. A summary for someone who needs to approve a budget should lead with the number they need to approve, not with background context.

Extract actions, decisions, dates, and risks

Meeting notes and long threads usually contain a few items that matter more than the rest: who agreed to do what, by when, and what could go wrong. Ask ChatGPT to pull these into a short, structured list instead of a narrative summary:

```
Using the attached meeting notes, extract the information that people need to act
on.
```

```
Create four sections:
1. Actions
2. Decisions
3. Important dates
4. Risks

For each action, include the task, owner, deadline, and status. Sort the actions by
owner.

Use only information stated in the notes. Do not infer an owner, date, decision, or
risk. Write "Not stated" when information is missing, and flag anything unclear for
review. Combine duplicate items that appear in more than one set of notes.
```

Figure 2.4A shows a page of meeting notes turned into a table of actions, owners, and deadlines:

1. Actions

Owner	Task	Deadline	Status	Review flag
Jordan Miles	Train the support team.	August 6, 2026	Scheduled	None
Leo Chen	Investigate the email queue and report the findings.	August 3, 2026	In progress	The cause of the email delay is not confirmed.

Figure 2.4A: Turning notes into a table makes it easy to check that every action has an owner and a date

Figure 2.4B shows the page of meeting notes with decisions extracted:

2. Decisions

- Keep the August 14 launch target for now.
- The password-reset process must be fixed and retested by August 4 for the August 14 launch to remain realistic.
- Ask Maya Patel to approve up to US$750 for extra vendor troubleshooting by August 1.
- Hold the final launch decision meeting on August 7.

Figure 2.4B: Turning notes into a bullet list makes it easy to validate every decision

Figure 2.4C shows the page of meeting notes with important dates extracted:

3. Important dates

Date	Event or deadline
August 1, 2026	Deadline to ask Maya Patel to approve up to US$750 for extra vendor troubleshooting.
August 3, 2026	Leo Chen to report the findings from his email-queue investigation.

Figure 2.4C: Extracting important dates into a table makes it easy to know what's happening and when

Figure 2.4D shows the page of meeting notes with risks extracted:

4. Risks

- If password-reset emails are still delayed after August 4, the August 14 launch may need to move.
- Password-reset emails sometimes take up to 15 minutes to arrive. The vendor's email queue settings may be causing the delay, but the cause is not confirmed. **Cause needs review.**

Figure 2.4D: Turning notes into a bullet list makes it easy to highlight risks

Check the extracted tables and lists against the source material. An action without a clear owner, or a date that doesn't appear anywhere in the notes, is worth a second look before you rely on it.

Real-world example: A team lead pastes the notes from three planning meetings into ChatGPT and asks for one table of open actions, sorted by owner, instead of rereading all three sets of notes before the next meeting.

Compare several documents

Scattered notes, emails, and reports can become one structured comparison. For example, a small business owner sourcing a new supplier could paste rough notes from three separate phone calls with printing companies into ChatGPT and ask it to generate a single table comparing price, delivery time, and support terms side by side:

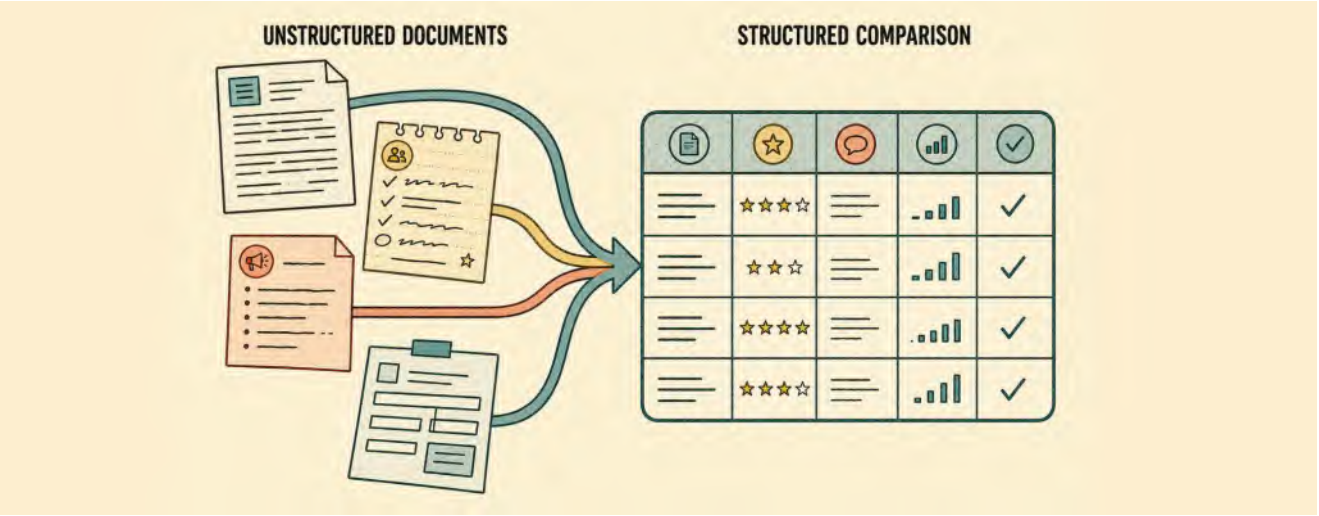


Figure 2.4E: Scattered notes, emails, and reports become one structured comparison

You may need to compare notes from several meetings, versions of a report, or proposals from different printing companies. Upload the documents together and ask ChatGPT to compare them against the same set of questions.

Sample files: If you don't have your own set of files, download all the files with names that start with `2-4-compare` from this folder: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2/`. There are three fictional, one-page source documents in different formats: an email proposal, rough call notes, and a draft quote. Their facts overlap enough for comparison, but each has different omissions and tradeoffs.

Structuring the comparison around the same questions for every document keeps the comparison fair. Ask for the comparison in a consistent format, such as a table with one row per document and one column per question. This makes differences easier to spot than reading a paragraph of prose for each document:

```
Compare the three attached vendor documents for customer portal training.

Create one table with one row per vendor and these columns:
- Vendor
- Total price
- Earliest start date
- Delivery format and capacity
- Materials and follow-up support
- Conditions or limitations
- Missing information

Use short phrases so the table is easy to scan. Use only facts stated in the
documents.

Do not combine details from different vendors. If a document does not answer a
question, write "Not stated." Do not assume that something is included or excluded
merely because it is not mentioned.

After the table, list up to three questions that should be answered before choosing
a vendor. Do not recommend a vendor.
```

	Total price	Earliest start date	Delivery format and capacity	Materials and follow-up support	Conditions or limitations	Missing information
	US$7,200; office-session travel extra, up to US$600	August 24, 2026	Three 2-hour workshops; up to 45 people total; all online, or first at office and two online	Existing materials adapted to portal; one coaching call for managers	Final employee list due one week before first workshop; online sessions recordable; August 24 is after current launch target; details are from rough call notes	Payment dates; cancellation terms; later support; accessibility compliance; whether office session can be recorded
	US$9,600 fixed fee, including travel	August 12, 2026	Six 75-minute in-person workshops; up to 10 people per group; maximum 60 people total	Custom role-play exercises; accessible Word and PDF workbook; skills check; phone and email support for 60 days	Suitable room required; portal access due two weeks before first session; sessions not recorded; 10 business days' notice for cancellation or rescheduling; dates held only after signing, quote valid through August 5	Exact workshop dates; support hours and response times
	US$8,400 fixed fee	Week of August 17, 2026	Four live 90-minute virtual workshops; up to 50 people total; 9:00 a.m. to 5:00 p.m. Eastern Time	Workshop recordings; facilitator guide; two quick-reference sheets; email questions for 30 days	50% due when booking; balance after delivery, portal access due August 10; no evening or weekend sessions, five business days' notice to move a session without charge; August 17-21 held only until August 3	Accessibility compliance; cancellation terms; support hours and response times

Figure 2.4F: Comparing unstructured documents in a structured table

Alongside the table, ChatGPT flagged the following open questions to answer before choosing a vendor:

- On what dates would all workshops be completed?
- Do all materials and delivery methods meet the required accessibility standard?
- What are the complete cancellation, rescheduling, payment, and follow-up-support terms?

Retain links to supporting passages

A summary or extracted list is only useful if you can check where each item came from. Ask ChatGPT to quote or reference the specific passage that supports each claim, action, or date it pulls out.



Watch out: a confident summary can still misstate a detail from the source. Confirm any figure, date, or decision that will affect someone's schedule or budget before you pass it along.



What to ask next: `Flag any statement in the summary that you cannot trace back to a specific passage in the source material.`

Try it now

Take a set of meeting notes or a long email thread. Ask ChatGPT to extract every action, who owns it, and the deadline, as a short table. Then ask for a one-paragraph summary for someone who missed the meeting.

Check the result

- ☐ Does every extracted action have an owner and, where one was given, a date?
- ☐ Does the summary match the length and tone the reader needs?
- ☐ Can you trace each extracted item back to a specific passage in the source material?

Summarizing turns long text into a short, useful result. The next chapter does the same job for numbers. *Chapter 2.5* shows you how ChatGPT can clean, calculate, and visualize a spreadsheet, so you can trust a chart before you present it.

Chapter 2.5 Analyze Spreadsheets and Data

ChatGPT can run code in a secure analysis environment to clean, summarize, visualize, and make projections from structured data. This chapter follows Brightside Learning, a fictional training provider, as it decides which workshops to run more often next quarter. The walkthrough covers five steps: explain your data and question, clean the data, analyze and visualize results, validate the results, then decide:

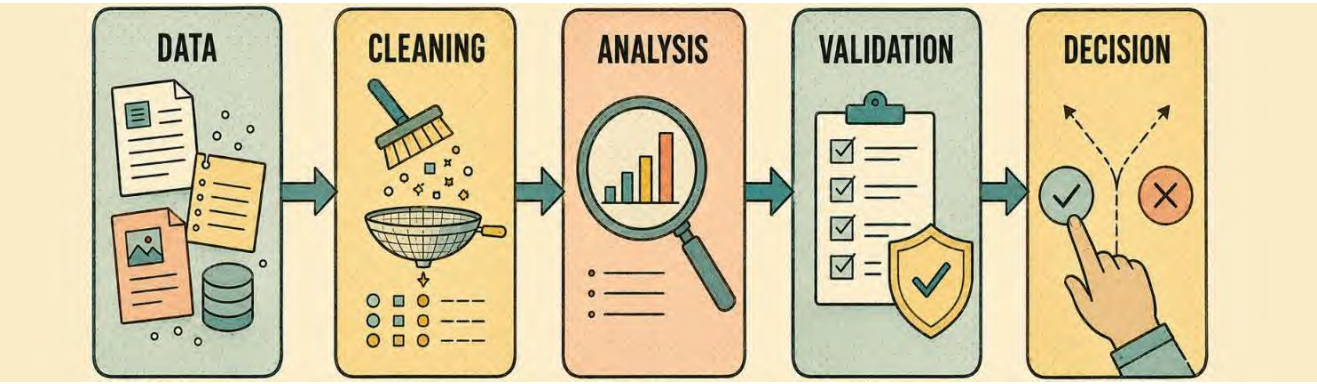


Figure 2.5A: Validation sits between analysis and decision on purpose. A chart is not a decision until you've checked the numbers behind it.

You need a spreadsheet or data file and a question you want that data to answer. Brightside Learning's workshop bookings give us both. By the end of this chapter, you'll be able to turn a raw spreadsheet into a checked analysis and chart, and you'll return to this same data in the final project chapter.



Sample files: If you don't have your own data to work with, download this document: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2/2-5-sample-workshop-data.xlsx`. It contains a small, fictional sales spreadsheet of Brightside Learning's workshop sessions you can use to follow along.

Explain the data and the business question

Before you upload a spreadsheet, decide what question you want answered. "What does this spreadsheet show?" produces a vague response. "Which product line had the largest drop in sales last quarter?" gives ChatGPT something specific to calculate.

Describe what each column means, especially if the headers are abbreviated or specific to your organization. A short explanation of the data saves several rounds of clarifying questions later:

```
Review the attached workbook, including the Read Me and Workshop Sessions sheets.

Before changing or calculating anything, explain:
- What one row represents
- What each column contains
- Which columns contain categories, dates, counts, money, ratings, or notes
- What the stated business question is
- How net revenue, attendance rate, and satisfaction could help answer it
- What the data cannot tell us

Use plain language for someone who does not work with data. List any missing
definitions, assumptions, or questions that should be resolved before analysis.

Do not edit the workbook or calculate the results yet.
```

This prompt makes ChatGPT establish the meaning and limits of the data before producing numbers:

What one row represents

In the **Workshop Sessions** sheet, each row represents one workshop session run by Brightside Learning from April through June 2026.

The apparent duplicate session, **S-114**, should be checked before relying on this definition. It may be an accidental duplicate or two records for the same session.

What each column contains

Column	What it contains	Type of information
Session ID	A code used to identify the session, such as S-101	Identifier
Session date	The date on which the workshop took place	Date
Course	The course taught, such as AI at Work or Excel	Category

Figure 2.5B: ChatGPT explaining what rows represent and what columns contain

Clean and inspect the information

Real spreadsheets rarely arrive ready for analysis. Ask ChatGPT to check for missing values, duplicate rows, and inconsistent formatting before running any calculations.

Ask to see what changed during cleaning. A cleaning step that drops rows or changes values without telling you can throw off every result that follows:

```
Inspect the Workshop Sessions sheet for data-quality problems.

Check for:
- Exact duplicate rows
- Missing values
- Inconsistent spelling, capitalization, or spacing
- Dates stored in different ways
- Values outside their expected range
- Bookings that exceed capacity
- Other values that may need human review

Keep the original sheet unchanged. Create:
1. A new sheet named Cleaned Sessions
2. A new sheet named Data Issues

In Cleaned Sessions:
- Remove only exact duplicate rows
- Standardize text values only when they clearly refer to the same course, format,
or location
- Convert recognizable dates to real spreadsheet dates
- Do not fill missing values or change questionable figures by guessing

In Data Issues, record:
- Session ID
- Column
- Original value
- The problem found
- Any change made
- Whether human review is still needed

At the top of Data Issues, report the original row count, cleaned row count, number
of duplicates removed, and number of unresolved issues.

Return the cleaned workbook as a downloadable Excel file.
```

The source contains 29 rows representing 28 unique sessions. ChatGPT should also find inconsistent labels, a text-formatted date, missing information, and a bookings figure above capacity:



Figure 2.5C: A cleaned spreadsheet ready for analysis.

Calculate and visualize useful results

Ask for the specific calculation you need, such as a total, an average, a trend, or a projection, rather than a general analysis. Ask for a chart only after you know which numbers matter, so the chart shows the comparison you actually need:

```
Using the Cleaned Sessions sheet, add these calculated columns with spreadsheet
formulas:
- Attended: Bookings minus No-shows
- Gross revenue: Bookings multiplied by Fee per attendee
- Net revenue: Gross revenue minus Refunds
- Attendance rate: Attended divided by Bookings

If Refunds is missing, leave Net revenue blank and flag the session for review. Do
not treat a missing refund value as zero.

Create a new sheet named Analysis containing:

1. A summary by course
2. A summary by delivery format

For each summary, calculate:
- Number of sessions
- Total bookings
- Total attended
- Overall attendance rate, calculated as total attended divided by total bookings
- Sessions with known net revenue
- Total known net revenue
- Average net revenue per session with known revenue
- Average satisfaction, excluding missing scores
```

```
Create a horizontal bar chart comparing average net revenue per session by course.
Sort the courses from highest to lowest. State clearly that the chart uses only
sessions with known refund information.

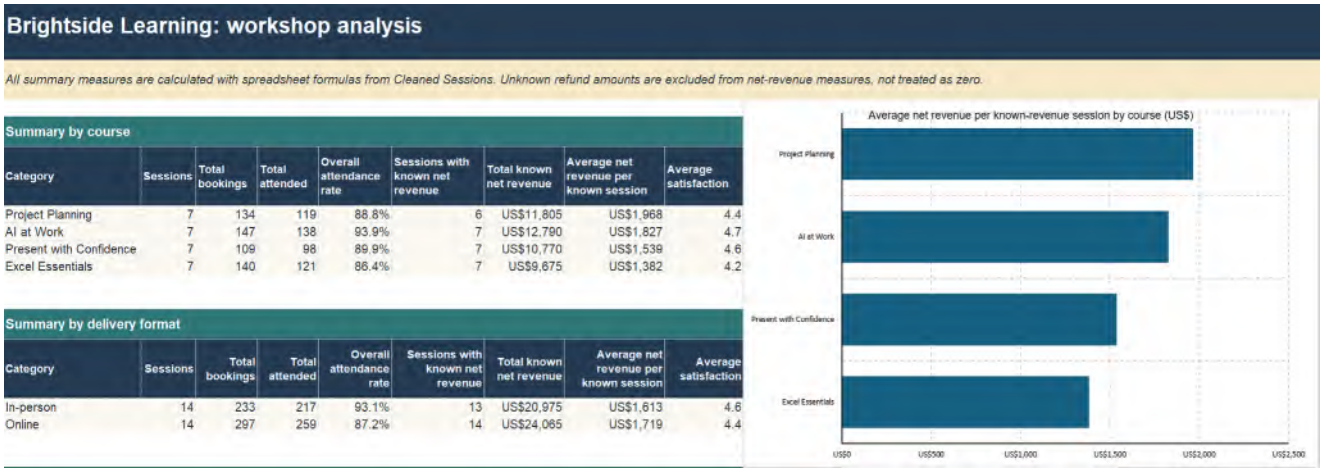
Below the summaries, write three concise observations that could help with next
quarter's schedule. Separate facts from interpretations, and mention any missing or
questionable data that limits the conclusions.

Keep all calculations visible in the workbook and return it as a downloadable Excel
file.
```

I chose average net revenue per session for the chart because total revenue would unfairly favor any course that happened to run more often. The prompt also requests attendance and satisfaction, preventing the analysis from treating revenue as the only measure.



Sample files: The final spreadsheet with the generated analysis: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2/2-5-sample-workshop-data-analysis.xlsx`.



Brightside Learning: workshop analysis

All summary measures are calculated with spreadsheet formulas from Cleaned Sessions. Unknown refund amounts are excluded from net-revenue measures, not treated as zero.

Summary by course

Category	Sessions	Total bookings	Total attended	Overall attendance rate	Sessions with known net revenue	Total known net revenue	Average net revenue per known session	Average satisfaction
Project Planning	7	134	119	88.8%	6	US$11,805	US$1,968	4.4
AI at Work	7	147	138	93.9%	7	US$12,790	US$1,827	4.7
Present with Confidence	7	109	98	89.9%	7	US$10,770	US$1,539	4.6
Excel Essentials	7	140	121	86.4%	7	US$9,675	US$1,382	4.2

Summary by delivery format

Category	Sessions	Total bookings	Total attended	Overall attendance rate	Sessions with known net revenue	Total known net revenue	Average net revenue per known session	Average satisfaction
In-person	14	233	217	93.1%	13	US$20,975	US$1,613	4.6
Online	14	297	259	87.2%	14	US$24,065	US$1,719	4.4

Figure 2.5D: Analysis sheet and chart

Compare the original columns with the generated analysis to confirm ChatGPT used the data the way you intended.



Learn more online: *Chapter 2.3* introduces the ChatGPT extensions for Word, Excel, and PowerPoint. OpenAI's help center article on ChatGPT for Excel and Google Sheets covers current plan and workspace requirements, installation steps, and known limitations for the spreadsheet-native version of ChatGPT that works directly inside an open workbook: `https://help.openai.com/en/articles/20001063-chatgpt-for-excel-and-google-sheets`.

Check calculations and interpretations

Spot-check a calculation against the raw data before you use it in a decision. Pick a few rows, do the arithmetic yourself, and compare the result.

An average or trend can look reasonable while resting on a small sample, an outlier, or a misread column. Ask ChatGPT to state how many rows a calculation is based on and whether any values were excluded.



Good practice: Ask ChatGPT to explain its method in plain language before you trust the result. If you can't follow the explanation, ask for it again in simpler terms rather than accepting a number you don't understand.



Real-world example: A shop manager uploads a year of sales data and asks ChatGPT which product line lost the most ground last quarter, then checks the answer against the three lowest-performing rows by hand before adjusting orders.

Try it now

Upload a spreadsheet you already have and describe one specific question you want it to answer. Ask ChatGPT to clean the data, show you what changed, then calculate and chart the result.



What to ask next: `Which single number in this analysis would change my decision if it turned out to be wrong?`

Check the result

- ☐ Do the totals or averages match a rough estimate you can calculate by hand?
- ☐ Did the cleaning step change or remove any rows you needed to keep?
- ☐ Does the chart compare the values you actually asked about?



Learn more online: OpenAI's help center article on data analysis in ChatGPT covers current upload limits, chart types, and analysis controls: `https://help.openai.com/en/articles/8437071-data-analysis-with-chatgpt`.

With a checked analysis in hand, it's time to switch from numbers to a different kind of visual. *Chapter 2.6* covers generating and editing images, along with the permission and copyright questions you need to answer before you use one.

Chapter 2.6 Create and Edit Images and Media

ChatGPT can generate illustrations, concept visuals, and social media assets, and it can edit an uploaded image using plain instructions. This chapter covers describing what you want, generating and editing images, and checking the result for accuracy, permission, and bias before you use it.

Keep precise diagrams separate from generated artwork. A diagram that shows a process or a comparison needs to be accurate in every detail, and image generation is not built for that kind of precision. *Figure 2.6A* shows the difference.

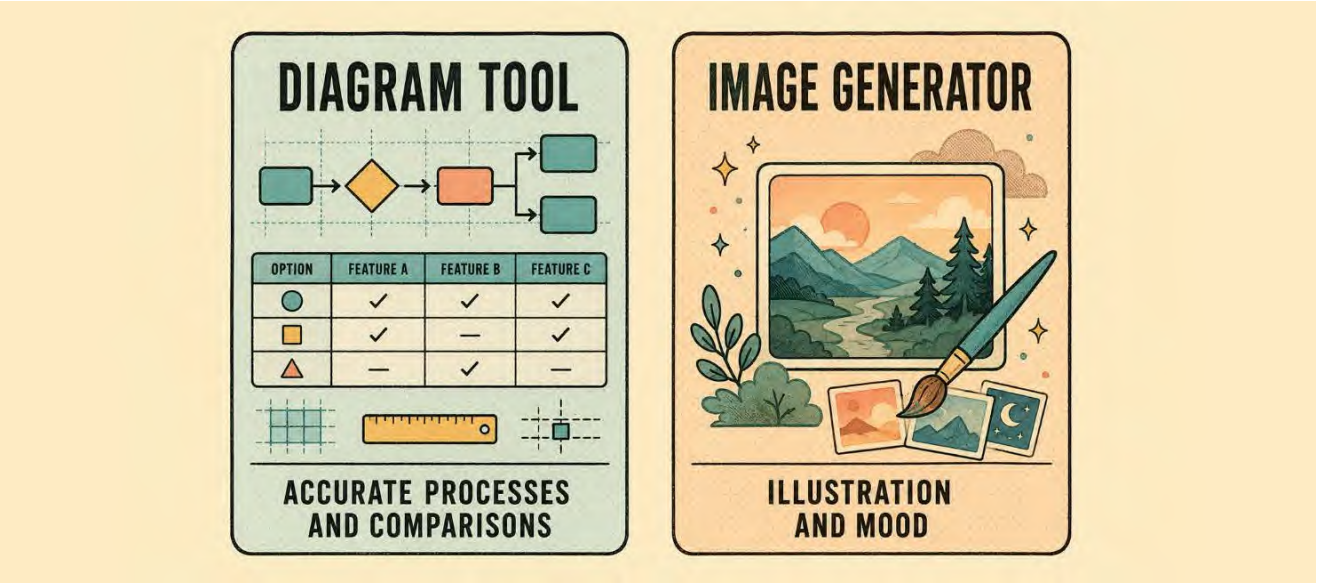


Figure 2.6A: Choose the tool based on whether the image needs to be exactly accurate or only needs to convey a feeling

You need a clear idea of the image you want, or an existing image to edit. By the end of this chapter, you'll be able to generate and revise an image, and know when to reach for a different tool instead.

Real-world example: OpenAI publishes worldwide ChatGPT usage data through a project called OpenAI Signals. Its August 2026 report named images and other media as ChatGPT's fastest-growing use case. By the second quarter of 2026, 7.8% of all ChatGPT messages worldwide asked it to create, edit, or find media. The skills in this chapter are already part of how people worldwide use ChatGPT every day. `https://openai.com/index/how-the-world-is-putting-chatgpt-to-work/`.

Describe the intended subject, composition, mood, and format

A useful image request describes more than the subject. Include the composition, such as a close-up or a wide scene. Add the mood, such as bright and energetic or calm and formal. Name the format, such as a square social media post or a wide banner.

The more specific your description, the less revision you'll need afterward. Compare "a photo of an office" with "a bright, wide photo of an open-plan office with people collaborating at a whiteboard, suitable for a website banner."

This chapter will use an image idea for a fictional Riverside Reading Club social post:

- **Subject:** A person reading in a comfortable bookstore corner, with an orange cat nearby.
- **Composition:** Reader and chair on the right, event text on the left, bookstore window in the background.
- **Mood:** Warm, welcoming, and relaxed.

- **Format**: Square social-media illustration.
- **Text**: "Riverside Reading Club" and "Saturday at 3 PM."
- **Style**: Friendly editorial illustration with clear shapes and muted colors.

This first prompt asks ChatGPT to turn the idea into a visual brief before generating anything:

```
I want to create a social-media image for a fictional community reading club.

Turn the following idea into a short visual brief:

A person is reading in a comfortable chair inside an independent bookstore. An
orange cat is resting on the rug. The image should feel warm, welcoming, and
suitable for adults of different ages.

Organize the brief under these headings:
- Subject
- Composition
- Mood
- Visual style
- Format
- Text to include
- Details to avoid

The finished image will be a square social-media post. It must include only this
text:

Riverside Reading Club
Saturday at 3 PM

Do not generate the image yet. Ask me up to three short questions if an important
choice is missing.
```

An image request involves more than naming the subject. Position, atmosphere, purpose, format, and restrictions can all change the result.



Good practice: Generate two or three versions of the same request before you pick one. Small wording changes can shift the result more than you'd expect, and comparing versions helps you spot which words are doing the work.

Generate illustrations, concept visuals, mockups, and social assets

Ask for the kind of image that fits your purpose. A rough concept visual can test an idea. A polished illustration suits a finished document. A mockup shows how content might look in a layout:

```
Using the visual brief we agreed on, generate the finished social-media image.

Create a square 1:1 editorial illustration. Place the reader and chair on the right
and leave a clear area for the event text on the left. Show a cozy independent
```

```
bookstore with a large window, bookshelves, a floor lamp, and an orange cat resting
on the rug.

Use warm, muted colors and clear shapes. Make the image friendly and professional
rather than childish.

Include exactly this text:

Riverside Reading Club
Saturday at 3 PM

Do not add any other words, logos, watermarks, or brand names. Keep important
content away from the outer edges so it will not be cropped on social-media
platforms.
```

Created the finished square editorial illustration with the exact requested text, crop-safe layout, bookstore setting, reader, lamp, and orange cat.

call_uU84OtXLrzjYHMa2IaAmp72r.png

Final prompt followed the agreed visual brief, with exact-text and no-additional-branding constraints.

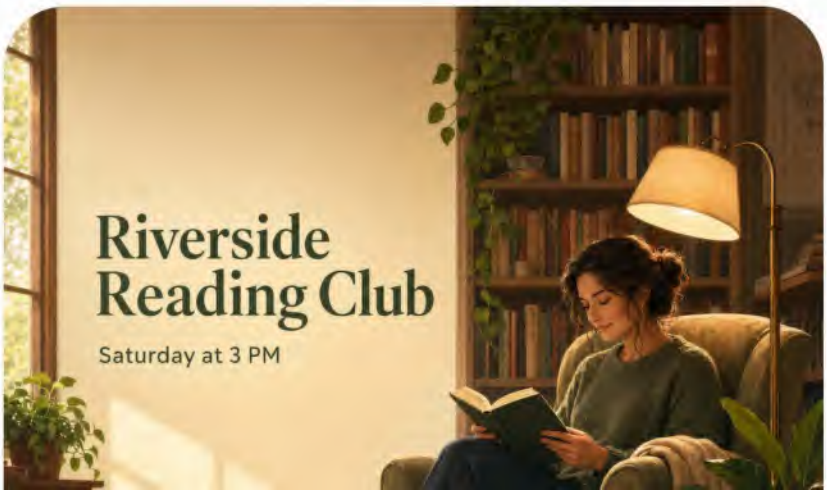


Figure 2.6B: Read the prompt and the image side by side to see which words shaped which parts of the result



Sample files: An example generated image: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2/2-6-cozy-bookstore-reading-club-scene.png`.

Edit an uploaded image using natural language

You can upload an existing image and ask for a specific change, such as adjusting the background, removing an object, or changing the lighting. Describe the change in plain language rather than technical editing terms.

You can download the generated image, upload it again, and request a visible change:

```
Edit the uploaded image.

Change the sunny afternoon outside the bookstore window to a cozy rainy evening.
Show soft rain on the window and warm light from the floor lamp.

Keep the reading-club setting, the person, the chair, the orange cat, and the square
composition. Keep all existing event text.
```

This is a good introductory edit because it uses everyday language. You do not need to identify layers, select pixels, or name image-editing tools.



Figure 2.6C: Uploading and then editing an image

The change should be obvious. The original has daylight outside, while the edited version has rain and warm evening lights.

Preserve important details while changing selected elements

Name the parts of the image that must stay the same, such as a person's face, a logo, or specific text, before asking for a change. This reduces the chance that an unrelated detail shifts along with the edit you wanted.

The next prompt makes the editing boundary much more precise:

```
Make one more edit to the uploaded image.

Change only the reader's sweater from deep teal to rust red.

Preserve all other details, including:
- The reader's face, skin tone, hair, expression, pose, and hands
- The book and its position
```

```
- The orange cat
- The chair, shelves, window, lamp, and other objects
- The rainy evening lighting
- Every word of the event text
- The typography, text size, and text position
- The placement and relative size of every element
- The square aspect ratio and crop

Do not add, remove, or move anything else.
```

The phrase “change only” sets the boundary. The preservation list then identifies details that might otherwise drift during regeneration.



Watch out: ChatGPT may rebuild parts of an image while editing it. Compare the edited result with the original. Check faces, hands, words, logos, product details, object positions, and the image crop. If something changed unexpectedly, return to the original and try a narrower instruction.

Create clear diagrams from a written brief

A generated illustration can be useful even when some small details vary. A diagram has less room for variation. Every box, label, arrow, and connection should communicate the intended meaning.

Start by deciding what the diagram needs to explain. Do not begin with its colors, icons, or artistic style. A useful diagram brief identifies:

- The main point the reader should understand.
- The items that must appear.
- The relationship between those items.
- The direction in which the reader should follow them.
- The exact wording of important labels.
- The details that should be left out.

For example, “Create a diagram about reviewing AI work” leaves many choices open. The following request gives ChatGPT a structure it can follow:

```
Help me plan a diagram that teaches this rule:

AI can assist with work, but a person remains responsible for the final decision.

The diagram must show these four stages in order:
1. AI assistance
2. Human review
3. Approval
4. Action

Use those exact labels. Show the stages from left to right, with one arrow between
each stage. Make Human review and Approval more visually prominent than the other
```

```
stages.

Do not create the diagram yet. First give me:
- A one-sentence teaching point
- A list of the required elements
- A list of the connections
- A short layout description
- Anything that might be misunderstood
```

Review this plan before asking ChatGPT to produce the visual. It is much faster to correct a missing connection in a written list than after the diagram has been drawn.

Choose a format that fits the job

A diagram can be created as a generated image, Mermaid code, Scalable Vector Graphics (SVG) code, or a set of shapes in a presentation. Each format has different strengths:

Format	Use it for	Main limitation
Generated PNG image	Concept visuals, hand-drawn diagrams, and illustrated explanations	Labels, arrows, spacing, and small details may change
Mermaid	Flows, sequences, hierarchies, and relationships	Layout and visual styling are limited
SVG	Precise custom layouts that must remain sharp at any size	The code can become long and harder to edit
Presentation shapes	Diagrams that people need to edit in PowerPoint or Google Slides	Building and aligning the shapes may take longer
Table or chart	Exact comparisons and numeric information	Less suitable for showing a process or loose relationship

Table 2.6: Formats for images and diagrams

Mermaid describes a diagram with short lines of text. ChatGPT can write that code and, where Preview is available, display the result inside a code block. SVG is also text-based code and can be previewed in supported code blocks. Preview availability can vary by plan, device, model, workspace settings, and rollout.



Good practice: Use Mermaid when the relationships carry most of the meaning. Use SVG when placement, typography, line weight, and visual style need tighter control. Use a generated PNG when the diagram benefits from illustration and small variations will not change its meaning. A table may be better than a diagram when readers need exact values or repeated comparisons. Do not force every idea into boxes and arrows.



Learn more online: *Images in ChatGPT*: `https://help.openai.com/en/articles/11084440-images-in-chatgpt`.

Check for distorted details, misleading representations, and unintended bias

Look closely at hands, text, logos, and small details, since generated images can distort these even when the overall image looks convincing. Also check whether the image represents people or situations in a way you didn't intend, such as showing only one type of person in a role meant to represent everyone.



Watch out: A generated image can look finished while still misrepresenting a real product, place, or person. Do not present a generated image as a photograph of something real.

Understand permission, likeness, trademark, and copyright concerns

Do not generate or edit an image of a real, identifiable person without their permission. Avoid recreating a company's logo, trademark, or copyrighted artwork. Check your organization's policy before using generated images in published or commercial work.



Learn more online: OpenAI's usage policies describe what image generation can and can't be used for, including rules on real people and copyrighted material. Check the current policy before using a generated image in published or commercial work: `https://openai.com/policies/usage-policies/`.

Try it now

Describe an image you need for a real task, such as a social media post or a slide background. Include the subject, composition, mood, and format. Generate the image, then ask for one specific edit while keeping the rest unchanged.

Check the result

- ☐ Does the image match the subject, composition, and mood you described?
- ☐ Did the edit change only the element you asked about?
- ☐ Would you be comfortable sharing this image publicly, or does it need another look first?

Chapter 2.7 brings together the writing and visual skills from the last two chapters to help you build a slide deck that leads your audience to a decision.

Chapter 2.7 Create Presentations and Visual Reports

A presentation should lead an audience to a decision, not just display information. This chapter covers defining the audience and desired decision, building a narrative, generating the slides and notes, and inspecting the finished slide deck before you deliver it.

Figure 2.7A traces this path from evidence through argument to a slide deck aimed at one audience decision.

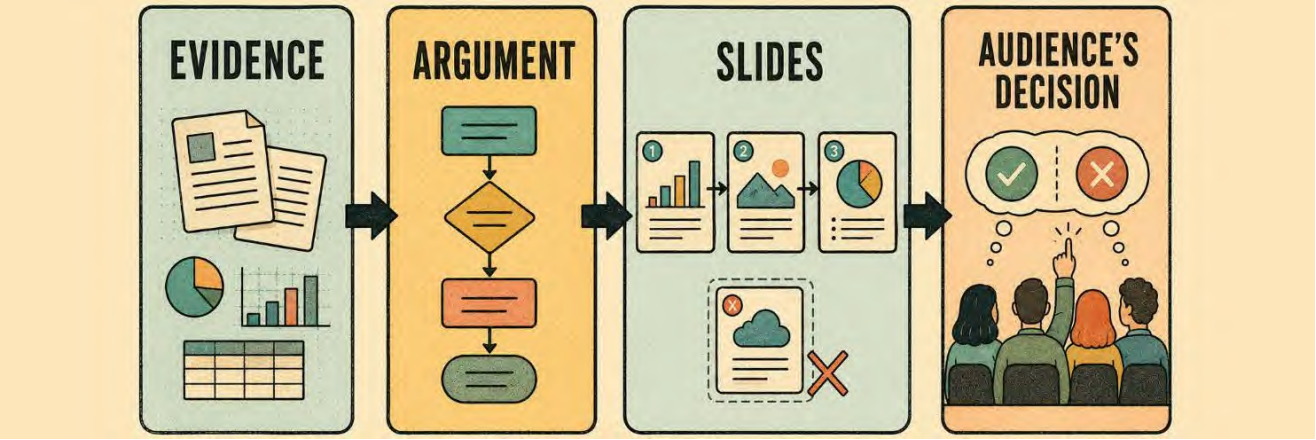


Figure 2.7A: Each slide should support one step in this chain. A slide that doesn't connect to the argument is a candidate for removal.

Real-world example: A department head builds a budget request deck by first asking ChatGPT for a narrative outline built around one decision, whether to approve a new hire, and only then asking for slides.

You need a topic, an audience, and a rough sense of what you want that audience to decide or do. By the end of this chapter, you'll be able to turn a brief into a reviewed slide deck with speaker notes.

If you don't have your own idea, I recommend one scenario that continues the Brightside Learning example from *Chapter 2.5*.

This scenario shows the difference between analyzing data and communicating a recommendation. The spreadsheet may contain dozens of results, but the audience needs only the evidence that helps them make the scheduling decision.

Define the audience and desired decision

Before asking for slides, decide who will watch the presentation and what you want them to approve, choose, or understand afterward. A presentation for a client deciding whether to sign a contract needs a different structure than one for a team reviewing progress.

State the audience and the decision in your first message. This shapes everything that follows, including how much background the deck needs and how the closing slide should read.

Include what the audience already knows, how much time they have, and what they care about. A senior manager may need the recommendation, supporting evidence, risks, and requested decision. They probably do not need a detailed explanation of every calculation.

You can turn the spreadsheet analysis into a presentation for a real decision:

- **Topic**: Decide which Brightside Learning workshops to offer more often next quarter.
- **Source material**: The cleaned and analyzed version of `2-5-sample-workshop-data.xlsx`.
- **Audience**: Brightside Learning's managing director and program manager. They understand the business but have not studied the spreadsheet.

- **Desired decision:** Approve the recommended scheduling priorities for next quarter, including which courses or delivery formats to increase, maintain, or reconsider.
- **Evidence:** Known net revenue, attendance rate, satisfaction, and data-quality limitations.
- **Format:** A seven-slide presentation designed for a seven-minute meeting.

Use the following prompt with the analyzed workbook from *Chapter 2.5*:

```
I need to prepare a short presentation based on the attached Brightside Learning
workbook.

Topic:
Which courses and delivery formats should Brightside Learning schedule more often
next quarter?

Audience:
The managing director and program manager. They understand the workshop business but
have not reviewed this analysis.

Desired decision:
Approve the recommended scheduling priorities, including which courses or formats to
increase, maintain, or reconsider.

Meeting:
I will have seven minutes to present, followed by questions.

Before creating any slides, prepare a presentation brief that states:
- What the audience needs to understand
- The decision I am asking them to make
- The evidence that will matter most
- Questions or objections the audience may raise
- The information that can be left out
- The appropriate tone and level of detail

Use only information found in the workbook. Identify any missing, questionable, or
incomplete data that may limit the recommendation. Do not create the slides yet.
```

Review the brief before continuing. Check that the requested decision is narrow enough for one meeting. "Discuss next quarter" is vague. "Approve the recommended scheduling priorities" gives the presentation a clear destination.

Build the story before creating slides

A strong presentation follows a line of reasoning: the problem or opportunity, the evidence, the options, and the recommendation. Each slide should prepare the audience for the next one. If ChatGPT creates the complete deck immediately, it may repeat information, introduce the recommendation too late, or fill slides with facts that do not support the decision.

Ask ChatGPT for an outline of this narrative before generating any slides, the same way you would ask for a document outline first:

```
Using the presentation brief and the attached workbook, create an outline for a
seven-slide presentation.

The presentation should lead the audience from the business question to a clear
scheduling decision.

For each slide, provide:
- Slide number
- Short slide title
- The one point the audience should remember
- Evidence from the workbook that belongs on the slide
- A suitable visual, such as a chart, table, or simple diagram
- The purpose of the speaker notes

Include:
1. A title slide that states the decision needed
2. The business question and measures used
3. Performance by course
4. Performance by delivery format
5. Recommended scheduling priorities
6. Data limitations and risks
7. The decision and next steps

Use only facts supported by the workbook. Do not invent a number of future sessions,
a budget, a cause, or a forecast.

Keep the outline concise. Do not create the finished presentation yet.
```

A useful slide title states the slide's message. For example, "Live online sessions have the strongest attendance rate" tells the audience more than "Delivery format analysis." The exact title should, of course, depend on the workbook results.

Check that each proposed chart answers a question. A decorative chart that does not help the audience decide can be removed.

Generate slides, notes, and supporting visuals

After approving the outline, ask ChatGPT to build the presentation. State how much text each slide may contain and which visual formats are suitable. Charts should represent the source data accurately. Generated artwork may support a theme or illustrate a broad idea, but it should not be used to represent exact figures. Ask for notes that expand on what you'll say, not a script you'll read word for word:

```
Using the attached workbook, the agreed presentation brief, and the approved seven-
slide outline, create the complete presentation as an editable PowerPoint file.

Design requirements:
```

```
- Use a clean 16:9 layout
- Use large text that will be readable in a meeting room
- Give each slide one clear message
- Keep body text to no more than five short bullets
- Use short, message-based slide titles
- Use a consistent layout, type style, and color palette
- Do not rely on color alone to communicate meaning
- Do not add decorative images that compete with the evidence

Charts and supporting visuals:
- Create charts from the workbook data
- Label every chart clearly
- Show whether revenue figures include only sessions with known refund information
- Use overall attendance rate rather than an unweighted average of session
percentages
- Exclude missing satisfaction scores from averages
- Keep charts and diagram elements editable where possible
- Use a small table or decision matrix to show whether each course or format should
be increased, maintained, or reconsidered
- Do not use a generated image to represent numeric results

Speaker notes:
- Add brief speaker notes to every content slide
- Explain the slide's main point
- State the evidence that supports it
- Mention any limitation the presenter should disclose
- Suggest a natural transition to the next slide

Accuracy:
- Use only facts from the workbook
- Do not invent causes, forecasts, budgets, or future results
- Mark an unsupported recommendation with [REVIEW]
- Preserve uncertainty where source data is missing or questionable

End with a slide that states the requested decision and the next action.

Return the finished presentation as a downloadable PowerPoint file.
```

The speaker notes should add useful context rather than repeat the slide word for word. They can remind the presenter about a missing value, explain how a measure was calculated, or prepare for a likely question.

Supporting visuals may include a horizontal bar chart comparing courses, a second chart comparing delivery formats, a small increase, maintain, or reconsider decision matrix, and a simple next-steps timeline. The charts should do most of the work. A presentation based on business data does not need an AI-generated photograph on every slide.

Figure 2.7B shows the generated slide deck in preview:

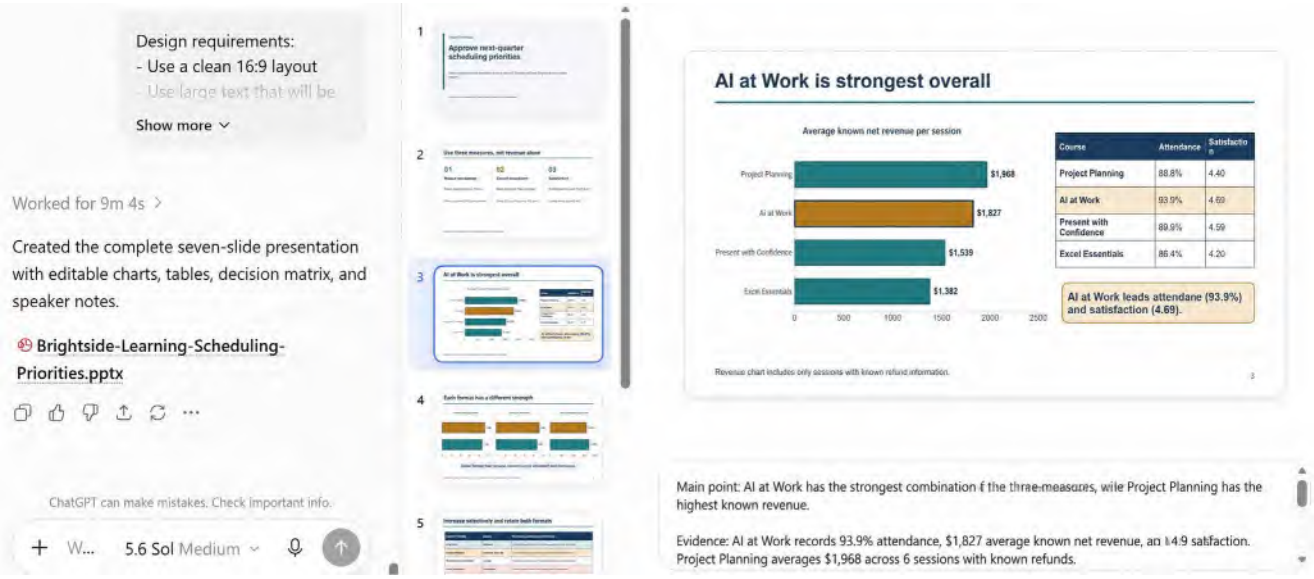


Figure 2.7B: Compare the outline with the finished slides to confirm every planned point made it into the deck



Learn more online: *Chapter 2.3* introduces the ChatGPT extensions for Word, Excel, and PowerPoint. OpenAI's help center article on ChatGPT for PowerPoint covers current plan and workspace requirements, installation steps, and known limitations for the presentation-native version of ChatGPT that works directly inside an open deck: `https://help.openai.com/en/articles/20001242-chatgpt-for-powerpoint`.

Inspect the complete presentation before delivery

A presentation can look finished while still containing incorrect figures, weak conclusions, or words that are too small to read. Review the complete deck against the original source, not only against ChatGPT's summary of that source.

Read every slide in order, not just the ones you asked for changes on. Check that numbers match across slides, that charts show what the text claims, and that the opening and closing slides connect.

The same material can become a one-page visual report for people who will not attend the meeting. Ask ChatGPT to preserve the recommendation, evidence, limitations, and next action, but redesign them for reading rather than presenting. A report must make sense without speaker notes.



Good practice: Ask a colleague unfamiliar with the material to look at the deck and tell you what decision they think you're asking for. If their answer doesn't match your intended decision, revise the narrative before the real presentation.

Try it now

Pick a real presentation you need to give. State the audience and the decision you want from them, ask for a narrative outline, review it, then generate the full deck with speaker notes.



What to ask next: `Which slide in this deck could I remove without weakening the argument?`

Check the result

- ☐ Does each slide support one step in the argument, rather than repeating information from another slide?
- ☐ Do the speaker notes help you present, rather than reading like a script?
- ☐ Would the intended audience reach the decision you were aiming for after seeing this deck?

So far, you've worked with material you already had in hand. *Chapter 2.8* opens the door to information ChatGPT can reach on its own, once you or your administrator connects a source such as email or a shared drive.

Chapter 2.8 Add Plugins, Apps, and Skills

ChatGPT can do more than respond using its general knowledge and the files you upload. You can add capabilities that help it follow specialized workflows, work with external services, and create particular types of output.

These capabilities are organized around three related ideas:

- A **plugin** is a package you can discover and install.
- An **app** connects ChatGPT to information or actions in another service.
- A **skill** gives ChatGPT reusable instructions for completing a particular kind of work.

You do not need to understand how these capabilities are built before using them. However, you should understand what you are installing, what information it can access, and whether it can take actions outside ChatGPT.

Understand how the pieces fit together

Think of a plugin as a box of capabilities prepared for a particular kind of work. Depending on its purpose, the plugin may contain one or more skills, one or more apps, or a combination of them:

Capability	What it provides	Simple example
Plugin	An installable package containing capabilities for a workflow	A presentation plugin containing instructions and supporting tools
App	A connection to an external service, its information, or its actions	Google Drive, SharePoint, Slack, or GitHub
Skill	Reusable instructions, examples, and resources for a particular task	A skill for creating and checking a presentation
App template	A starting point an administrator can use to configure an organization-specific app	A template for connecting an internal company system

Table 2.8: Plugin capabilities and what they provide

A skill is like a written playbook. It might tell ChatGPT how to inspect a spreadsheet, calculate results, create suitable charts, and check the finished workbook. The skill does not necessarily connect to another company's service.

An app is more like a doorway. It may allow ChatGPT to find a file in Google Drive, read a page in Notion, search messages in Slack, or create something in another service. The app can only use the access you and your organization have allowed.

A plugin brings these pieces together so that you can install a useful capability without assembling every part yourself.

Figure 2.8A traces the path from your request, through ChatGPT and the app, to the specific information it selects and the output it produces:

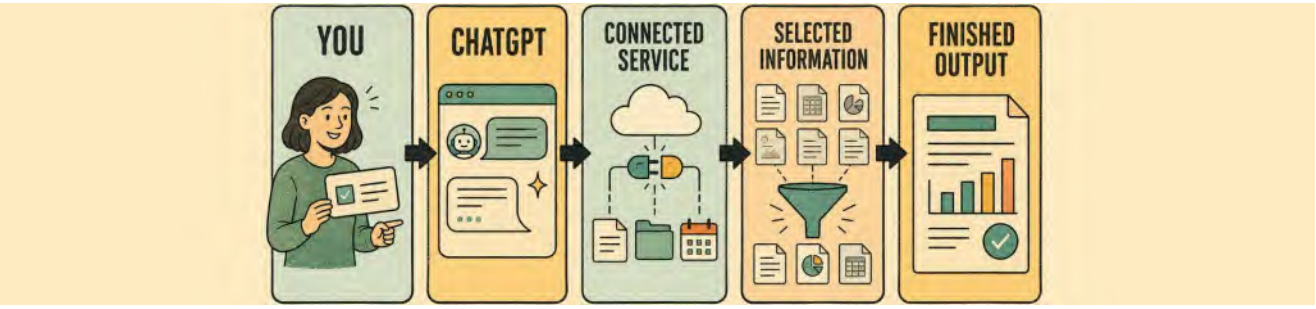


Figure 2.8A: Every app passes through this same chain. Reviewing the selected information before the output stage catches most problems early.

You need access to at least one approved plugin, or an administrator who can approve one. By the end of this chapter, you'll be able to pull relevant workplace information into a task while keeping control over what ChatGPT sees.

Use the Plugin Directory

The Plugin Directory is the main place to discover capabilities for ChatGPT and Codex. It replaced the earlier App Directory, although the word app is still used for connections to external services.

Open Plugins from the ChatGPT sidebar or settings. The exact position may vary between the web, mobile, and desktop versions:

1. In the ChatGPT sidebar, select **Plugins**.
2. Select the switch to toggle between showing **Plugins** and **Skills**.
3. Type in this box to search for a plugin:



Figure 2.8B: Browsing for plugins to add for connecting to external sources

4. Select a plugin to view its details and add it:

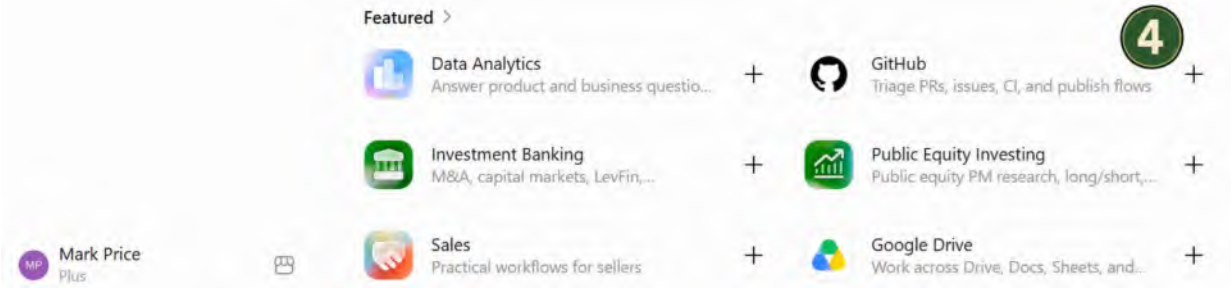


Figure 2.8C: Select a plugin to add it

Good practice: Install only what your task needs. You do not need to connect Google Drive, Microsoft 365, Slack, or another external service merely to work with files in ChatGPT. You can upload a document, spreadsheet, presentation, or other supported file directly.

A plugin's listing should explain what the plugin helps you do, which skills it contains, which apps it uses, whether those apps are required or optional, what setup or sign-in is needed, which ChatGPT surfaces support it, and what permissions it may request, as shown in *Figure 2.8C*:

1. **Gmail** permissions include viewing your messages and settings, or reading, composing, and sending messages.
2. **Google Calendar** permissions include seeing your availability, viewing and editing events, and seeing and downloading any calendar you have access to.
3. **Google Drive** permissions include seeing, downloading, creating, and deleting all your files. You can also control individual permissions for Google Docs, Sheets, and Slides.

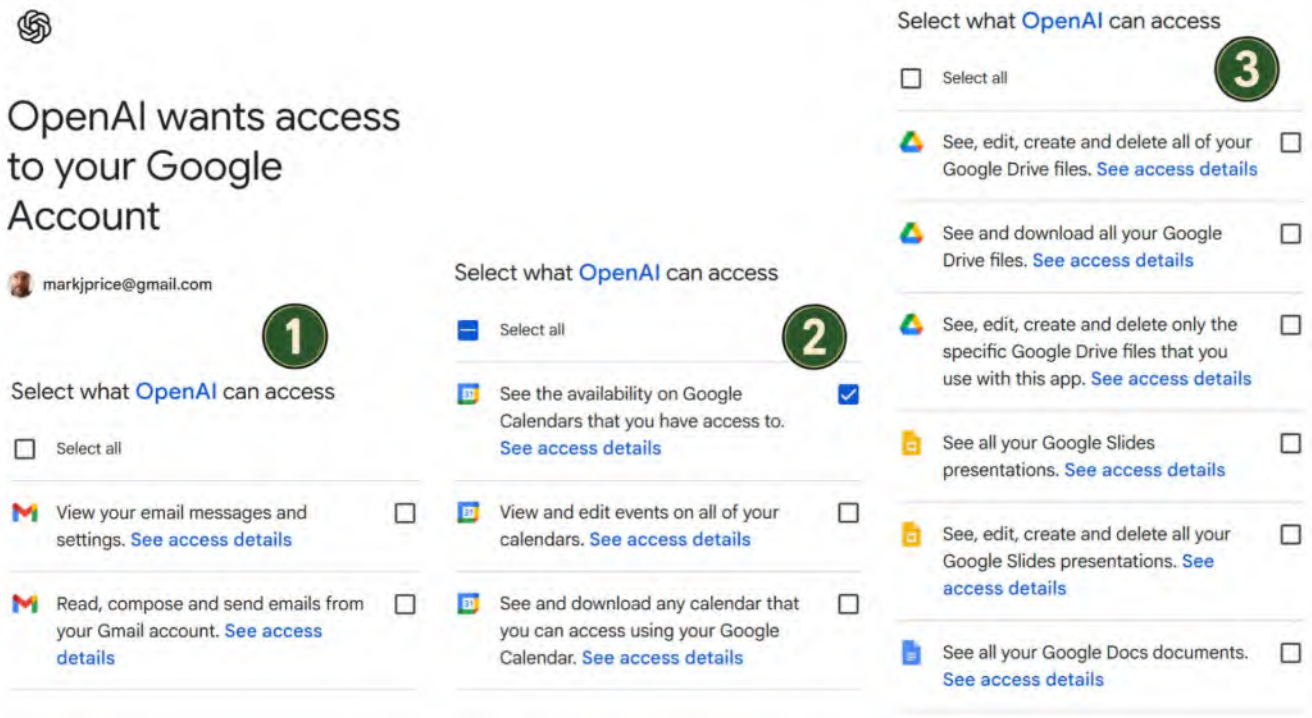


Figure 2.8D: Permissions for Gmail, Google Calendar, and Google Drive plugins. You can select only the capabilities you need, for example, read-only access to your emails and calendar.

Search for the result you want rather than installing everything that sounds interesting. For example, you might search for presentation creation, spreadsheet analysis, document editing, project reporting, or a particular service you already use.

A plugin containing only skills may not require an external account. A plugin containing an app may ask you to sign in to another service and approve access.

Availability can depend on your ChatGPT plan, device, region, workspace, and role. In a managed workplace or school account, an administrator may decide which plugins and apps you can use. If an Install or Connect button is unavailable, the capability may be restricted by your organization.



Watch out: A listing in the Plugin Directory does not guarantee that every included capability is available to you. A required app may still need to be enabled by an administrator or connected to your account.

Review app access and permissions

Connecting an app may allow ChatGPT to read information or take actions in another service. The available controls depend on the app and your workspace.



Good practice: Read the authorization screen before approving it. Check which account you are connecting, what information the app can access, and whether it can make changes.

When a skill, plugin, or connection can read files, reach another service, or take an action outside the chat, ChatGPT will show what it is asking permission to do. Read this request before approving it, the same way you would read a permission request on your phone.

While a tool is running, ChatGPT shows its status, such as Running, Waiting for approval, or Complete. This status tells you whether the work is still in progress or ready for you to review.



Watch out: Never approve a permission request you do not understand. Ask ChatGPT to explain what the tool needs and why, in plain language, before you say yes.

ChatGPT does not gain access to everything in an organization merely because an app is connected. Your existing permissions in the source service still apply. If you cannot open a restricted file in the source service, connecting the app should not give you access to it through ChatGPT.

When permission choices are available, prefer settings that ask before making changes. Reading a document is different from deleting it, sending it, sharing it, or replacing it.

Before connecting any account or service, check your employer's policy on which sources may be connected and who can approve the connection. Some organizations manage this centrally, while individual accounts may allow you to connect your own approved services directly.

Once a source is connected, describe what you're looking for in plain language, such as a project name, a date range, or the people involved. ChatGPT searches the connected source and returns the files or messages that match.

Ask ChatGPT to use only files from a specific folder, date range, or sender, rather than searching the whole connected source. A narrower request also produces a more focused, more useful result.



Watch out: Connecting a source gives ChatGPT the ability to search everything within the scope you approve, not only the files related to your current task. Review that scope before connecting.



Learn more online: OpenAI's help center article on apps in ChatGPT explains what a connected source can and can't access once approved, and how to review or revoke that access later. Check it before approving a new connection: `https://help.openai.com/en/articles/11487775-connectors-in-chatgpt`.

Grant only the access a task actually needs. If a plugin asks for more than its stated purpose requires, do not connect it, or contact the service to ask why.



Watch out: A connected plugin can keep working even after you stop using it regularly. Review your connected plugins occasionally, not only when you first set them up.

Remove access that is no longer needed

When you finish a project or stop using a service, disconnect the plugin rather than leaving it connected.

1. Open your connected plugins or applications list.
2. Select the plugin you want to remove.
3. Choose **Disconnect** or **Remove access**.
4. Confirm the removal.

Removing access does not delete your work in the other service. It only stops ChatGPT from reaching that service going forward.



Real-world example: An operations manager connects a project-management plugin for a three-month product launch. Once the launch ends, she removes the connection, since the plugin no longer has an active purpose and would otherwise keep standing access to the project data.

Select a plugin explicitly when necessary

ChatGPT may recognize when an installed plugin is relevant. You can also select a plugin or one of its included capabilities before sending your request.

Depending on the ChatGPT version you are using, you may be able to mention the plugin using `@` followed by its name, select it from the attachment or + menu, select it as a source in Work, or open the plugin from the Plugin Directory.

An explicit selection is useful when more than one installed capability could handle the task, or you want ChatGPT to use information from a particular service.

For example, you could prompt:

```
Using the connected Google Drive app, find the latest approved project brief for the
customer portal.

Summarize the agreed scope, deadlines, owners, and unresolved decisions. Include
links to the source files. Do not use draft documents or files outside the approved
project folder.
```

Naming the source and its boundaries reduces the chance that ChatGPT will use an unrelated file.

You do not need to install every available document, spreadsheet, or presentation plugin. Common plugins for working with documents like Word and PDF files, spreadsheets, and presentations are installed by default.

Cite and review source material

Ask ChatGPT to name or link the specific file or message behind each claim it makes from a connected source. Open a sample of these citations yourself before relying on the result, especially for anything that will inform a decision.



Real-world example: An office manager connects a shared drive and asks ChatGPT to find every vendor contract expiring in the next 60 days, then opens each cited file to confirm the expiration date before renewing anything.



Good practice: treat a connected search the same way you'd treat a colleague's summary of documents you haven't read yourself. Useful as a starting point, worth checking before you repeat it as fact.

Try it now

Open the plugin directory and find one plugin connected to a service you already use. Read its description and the permissions it would request, without connecting it yet. Decide whether you would approve that request and why.

Using an approved connected source, ask ChatGPT to find and summarize the files or messages related to one specific project or topic. Limit the request to a folder or date range you choose.

Check the result

- ☐ Could you explain, in one sentence, what the plugin you found actually does?
- ☐ Did the requested permissions match the plugin's stated purpose?
- ☐ Do you know where to go later to disconnect a plugin you no longer need?
- ☐ Do you know what a permission request looks like before you meet one in a real task?
- ☐ Did ChatGPT use only the sources you intended, based on the scope you set?
- ☐ Can you open the cited files and messages to confirm they say what the summary claims?
- ☐ Would you be comfortable sharing this result with someone outside your team, given what it draws on?



What to ask next: `List every connected source you used to answer this, and tell me if any file you checked seemed outdated.`

Connecting a source is only useful if ChatGPT can act on what it finds. You've now learned every skill this chapter about ChatGPT Work pulls together: drafting and revising documents, summarizing information, analyzing data, creating media, and adding plugins and apps. *Chapter 2.9* asks you to combine them into one real, recurring briefing you'll actually use.

Chapter 2.9 Project — Produce a Useful Weekly Briefing

This project brings together everything from *Book 2*. You'll build one weekly briefing, gather its material from approved sources, and produce it with Work.

You need at least one approved source of information, such as connected files, a shared drive, or your own notes. By the end of this project, you'll have a working weekly briefing and a process for improving it over time.

Figure 2.9A traces the full path from your connected sources, through Work, to your approval, and finally to the finished briefing:

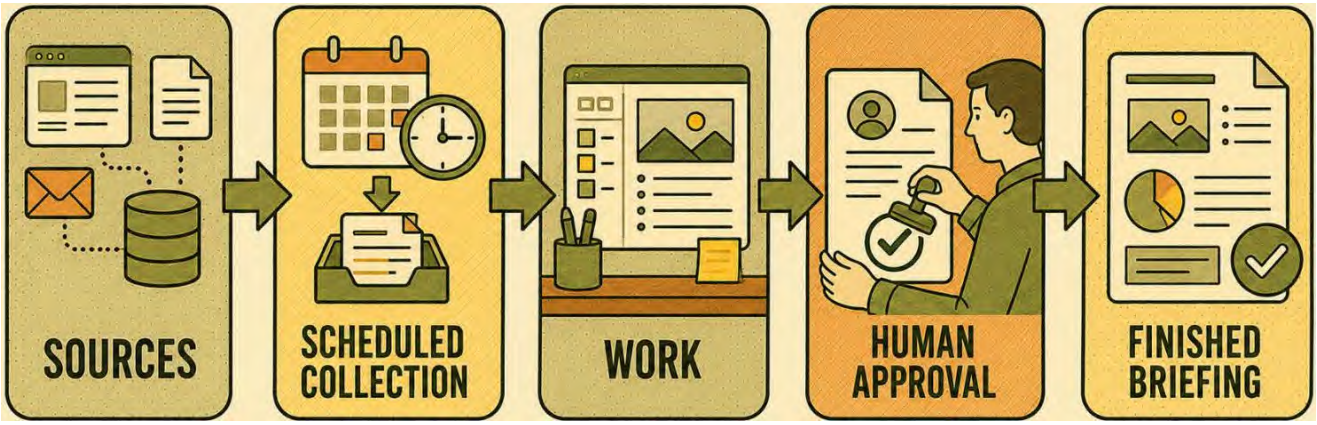


Figure 2.9A: Your approval sits between Work's draft and the finished briefing every week, not just on the first run.

Fictional worked example

Brightside Learning Weekly Operations Briefing continues the fictional organization used in the spreadsheet and presentation chapters, while bringing together files, evidence, decisions, and human review.



Sample files: You can download files that all start with `2-9` from the following folder: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book2/`.

Use these files in sequence:

- `2-9-Approved-Source-Register.docx`: Defines which sources may be used, what each source can support, and what information is excluded.
- `2-9-Weekly-Workshop-Tracker.xlsx`: Contains an editable dashboard, current-week workshop results, calculated attendance rates, and the next-week schedule.
- `2-9-Team-Update-Notes.docx`: Supplies confirmed risks, recommendations, client follow-ups, owners, and deadlines.
- `2-9-Previous-Weekly-Briefing.docx`: Provides a model for structure and tone. Its old figures are not approved as current evidence.
- `2-9-Example-Weekly-Briefing.docx`: Shows the one-page result ChatGPT should produce after gathering and checking the evidence.

The files include several useful details:

- Attendance must be calculated from total attendees and total capacity
- One satisfaction score is missing because its survey remains open
- A lower result has no confirmed explanation
- One session is full and has a waitlist
- One instructor arrangement is still at risk
- One decision needs management approval
- The previous briefing may guide presentation but not supply current facts

In this project, you will produce a weekly operations briefing for Brightside Learning, a fictional training business. The managing director and program manager use the briefing to understand the past week, prepare for the next one, and make decisions.

Your approved sources include a workshop tracker, team update notes, and the previous approved briefing. You will define what the audience needs, gather evidence from those sources, create a one-page briefing, check every important claim, and prepare the next edition.

You can adapt the same process for a family planning update, project report, client briefing, community newsletter, learning review, or small-business summary.

Define the audience and reporting questions

A useful briefing is written for a particular group of people. It should help them understand, decide, or do something. “Summarize this week” gives ChatGPT too little direction.

Begin by naming the audience and the questions the briefing must answer. Also explain what the audience already knows and how much time they have.

Upload `2-9-Approved-Source-Register.docx`, and use this prompt:

```
I am creating a weekly operations briefing for Brightside Learning.

Audience: The managing director and program manager, who know the workshop business
but haven't seen this week's results or team notes.

Purpose: Help them understand this week, prepare for next week, and make any
decisions that can't wait. Keep it to a three-minute read.
```

```
Read the attached Approved Source Register, then create a briefing plan (not the
briefing itself) that states:
- What this audience needs to understand
- Decisions or actions to support
- Five reporting questions to answer
- Evidence needed for each question
- Information to leave out
- Length, tone, and structure
- Any rule that could prevent a misleading result
```

A suitable response should identify questions such as:

- How many sessions were delivered?
- What were overall attendance, known net revenue, and satisfaction?
- Which results were unusually strong or weak?
- What risks or decisions affect next week?
- What information remains missing?

Review the questions before gathering information. Remove questions that are interesting but do not help this audience act. Each reporting question should have a reason for being included and at least one approved source that can answer it.

Gather information from approved sources

The source register gives ChatGPT boundaries. It identifies what may be used, what must be excluded, and how evidence should be labeled.

Upload the tracker, team notes, and previous briefing. Ask ChatGPT to gather evidence before drafting prose:

```
Use the approved briefing plan and these attached files:
- 2-9-Approved-Source-Register.docx
- 2-9-Weekly-Workshop-Tracker.xlsx
- 2-9-Team-Update-Notes.docx
- 2-9-Previous-Weekly-Briefing.docx

Read the Source Register first and follow its rules. Do not write the finished
briefing yet.

For each reporting question, build an evidence table with these columns:
- Reporting question
- Finding
- Exact figure or confirmed fact
- Source filename
- Sheet, row, cell, or heading
- Limitation or missing information
- Suitable for the briefing: Yes, No, or Needs review
```

```
Calculation rules:
- Calculate overall attendance as total attendees ÷ total capacity, not an average
of session percentages.
- Include known net revenue from delivered sessions only.
- Exclude missing satisfaction scores from the average, and state how many sessions
have a score.

Source rules:
- Use the previous briefing for structure and tone only, never as current figures.
- Don't guess why a result changed, or invent forecasts, revenue, client details,
owners, or deadlines.
- Mark conflicting or unsupported information with [CHECK].

Finish with:
1. Three most important findings
2. Decisions or actions needing attention
3. Risks or missing information to disclose
4. Any question the approved sources can't answer
```

This intermediate table is important. It separates gathering evidence from writing persuasive prose. It also gives the reader a chance to catch a calculation error or unsupported conclusion before it reaches the briefing.

In the supplied example, ChatGPT should find:

- Four delivered sessions
- Attendance of 58 out of 72 places, or 80.6%
- Known net revenue of £11,340
- Average satisfaction of about 4.5 from three of four sessions
- A full AI for Office Tasks session with four people waiting
- A pending backup instructor confirmation
- One management decision and two follow-up actions
- No approved explanation for Project Planning's lower attendance



Watch out: Missing information is not the same as zero. An open satisfaction survey should be reported as missing or pending.

Produce the briefing and supporting evidence

Once the evidence table is accurate, ask ChatGPT to turn it into a short briefing. The main briefing should remain readable, while its source labels make important claims easy to check:

```
Using the approved briefing plan and the checked evidence table, create the
Brightside Learning Weekly Operations Briefing for the week ending July 31, 2026.
Return it as an editable Word document.
```

```
Audience: The managing director and program manager. Keep the complete briefing to
one page, about a three-minute read.

Use these sections:
- A message-based title
- In brief
- What changed this week
- Decision and actions
- Risks and missing information
- Supporting evidence

Content requirements:
- State delivered sessions, total attendance, and overall attendance rate.
- Report known net revenue from delivered sessions only, and average satisfaction
with its coverage.
- Identify the strongest and weakest attendance results, without inventing an
explanation.
- State the decision needed on the proposed overflow session.
- Give each confirmed action an owner and due date.
- Disclose the missing satisfaction result and pending instructor confirmation.

Evidence requirements:
- Support every important figure or claim with a cited, approved source (filename
plus sheet, cell, row, or heading).
- Use the previous briefing for structure and tone only, and preserve uncertainty
from the sources.
- Mark anything needing confirmation with [CHECK].

Design requirements:
- Use a clean one-page layout with the summary and required decision easy to find.
- Use a small table for decisions and actions, with text large enough for
comfortable reading.
- No decorative images or unsupported charts.

Do not send, publish, or distribute the briefing. Return it for human review.
```

The supporting evidence does not need to interrupt every sentence. A short evidence section can identify the workbook sheets, cell ranges, and note headings used. For a higher-risk or client-facing report, use numbered source markers beside important claims and provide a fuller evidence appendix. A reader should be able to distinguish results, decisions, recommendations, risks, and missing information without guessing.

Review the briefing

A fluent briefing can still contain a wrong number or unsupported explanation. Review it against the original sources rather than checking only whether it sounds convincing.

Upload the generated briefing alongside the source files:

```
Review the attached weekly briefing against:
- 2-9-Approved-Source-Register.docx
- 2-9-Weekly-Workshop-Tracker.xlsx
- 2-9-Team-Update-Notes.docx
- 2-9-Previous-Weekly-Briefing.docx

Do not edit the briefing yet.

Check every important claim for:
- A matching approved source, with correct arithmetic and reporting period
- Correct owner and deadline, and any missing qualification or uncertainty
- Unsupported causes, forecasts, or recommendations
- Accidental reuse of figures from the previous briefing
- Confidential, unnecessary, or misleadingly presented missing values
- A source citation that doesn't lead to the stated evidence

Create a review table with these columns:
- Briefing claim
- Source checked
- Result: Supported, Incorrect, Incomplete, or Unsupported
- Problem found
- Recommended change
- Priority: High, Medium, or Low

Also confirm the briefing stays on one page, leads with the most important
information and a clear decision, gives every action an owner and date, and is ready
for human approval.

Do not change the document until I approve the review.
```

After checking the review table, the reader can approve specific revisions:

```
Apply only the High-priority changes that I approved.

Do not change any other figure, conclusion, wording, source citation, owner,
deadline, formatting, or layout.

Return the revised briefing as a new Word document so the previous version remains
available.
```

Try it now

Build a first draft of your briefing by hand, using the sources and questions you defined. You can choose one of three versions of this project:

1. A personal or family planning briefing
2. A team, manager, or executive briefing
3. A client, market, or small-business briefing

Pick the version that matches a briefing you'd actually use, since a real audience gives you a genuine reason to keep improving it.

Check the result

- ☐ Does the briefing answer the specific questions you defined, without drifting into unrelated material?
- ☐ Can you open the linked sources and confirm they support the summary?
- ☐ Would your intended audience find this briefing worth reading every week?

The next step

Your weekly briefing is a real result, but right now it still depends on you: checking in, adjusting its questions, and deciding when it's ready to run on its own. *Book 3*, *Automate Recurring Tasks*, picks up exactly there. You'll learn to turn a working routine like this one into a skill, a specialized GPT, or a scheduled task that ChatGPT can run without you standing over it, and how to keep control over what it's allowed to do along the way.

Book 3

Automate Recurring Tasks

Book 3 teaches you how to turn a procedure you repeat often into a scheduled task, a skill, or a Custom GPT that ChatGPT can run for you without you needing to handhold.

It has eight chapters, including an end-of-book project:

3.1 Choose the Right Automation Building Block 120

3.2 Schedule Recurring and Conditional Work 124

3.3 Turn a Procedure Into a Reusable Skill 130

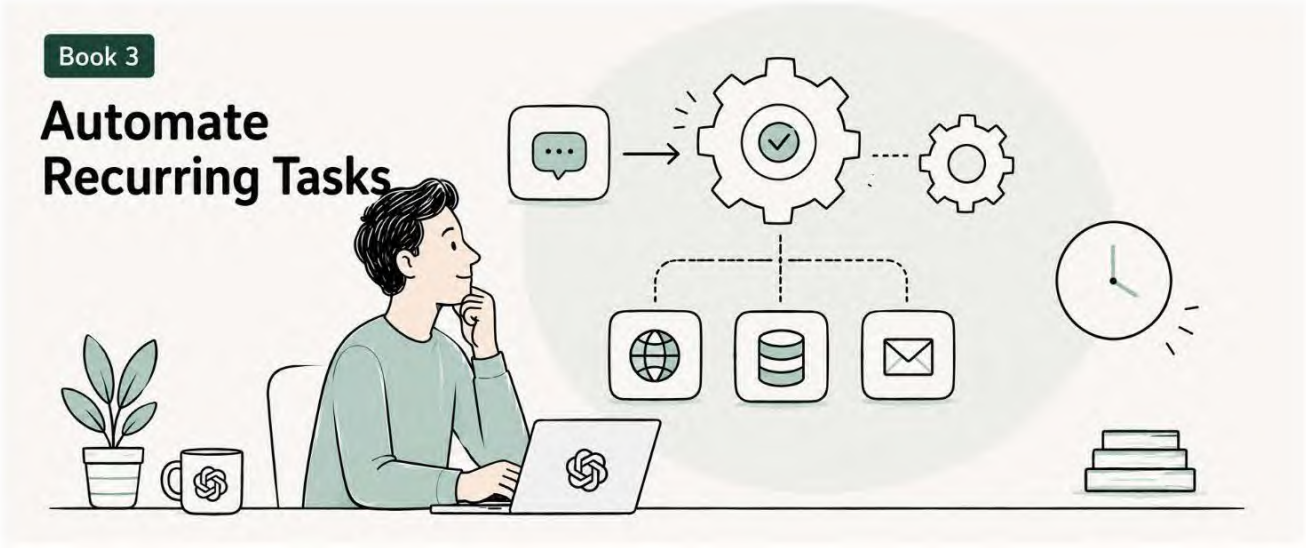
3.4 Build a Custom GPT 139

3.5 Automate Personal Organization and Learning 145

3.6 Build a Research and Content Workflow 147

3.7 Build a Decision and Operations Workflow 149

3.8 Project – Build and Measure an End-to-End Workflow 150



Chapter 3.1 Choose the Right Automation Building Block

By now you have used ChatGPT for single conversations, research, and finished documents. This chapter is about the next step: work you do again and again. You will learn to spot that kind of work, and to pick the simplest tool that handles it well. You do not need to memorize every option today. You need a way to decide, each time a task repeats, what to build.



Watch out: All chapters in *Book 3* require an individual paid plan or organization-managed workspace.

Recognize repeated and multi-step work

Some tasks happen once. Others come back every week, every project, or every new client. A task is worth automating when it repeats, follows a pattern, and takes real time to redo by hand.

Watch for phrases you say to yourself, such as "I did this last month too" or "I always have to remind it about our format." Those are signs that a one-off conversation should become something more permanent.

A multi-step task has several stages that must happen in order, such as gathering information, drafting a document, and checking it against a rule. These tasks benefit the most from structure, because a plain chat can lose track of the stages.



Watch out: Do not automate a task you have only done once. Without a few repetitions, you cannot yet tell which parts are stable and which parts change each time.

Compare projects, scheduled tasks, skills, and custom GPTs

ChatGPT offers several ways to make repeated work easier. Each one solves a different problem:

- A **project** keeps related chats, files, and instructions together. Use it when work is repeated but still needs your judgment at each step. You were introduced to projects in *Book 1*, *Chapter 1.11*.
- A **scheduled task** runs a set of instructions automatically, once or on a repeating schedule. Use it when the same request should happen without you starting it. You will learn about scheduled tasks in *Chapter 3.2* because it is the simplest concept, good for a gentle start to automation.
- A **skill** stores a tried and tested procedure, including instructions, examples, and sometimes code, so ChatGPT can run it consistently. Use it when the steps are stable and well understood. You will learn about planning, building, and sharing skills in *Chapter 3.3*.
- A **Custom GPT** packages instructions, reference files, and chosen tools into an assistant built for one purpose. Use it when you or your team return to the same kind of task often. You will learn about planning, building, and sharing custom GPTs in *Chapter 3.4*.
- A **workspace agent** is an organization-managed version of a Custom GPT, built and maintained for a Business, Enterprise, Edu, or Teachers account. Use it when the AI must follow rules set by an organization rather than one person.

- **Computer Use** and **Appshots** let ChatGPT reach an app or a screen that has no plugin, app, or MCP connection, either through a one-time snapshot or by clicking, typing, and navigating on its own.

To guide you in choosing the right automation building block for your work, use the matrix shown in *Figure 3.1A*:

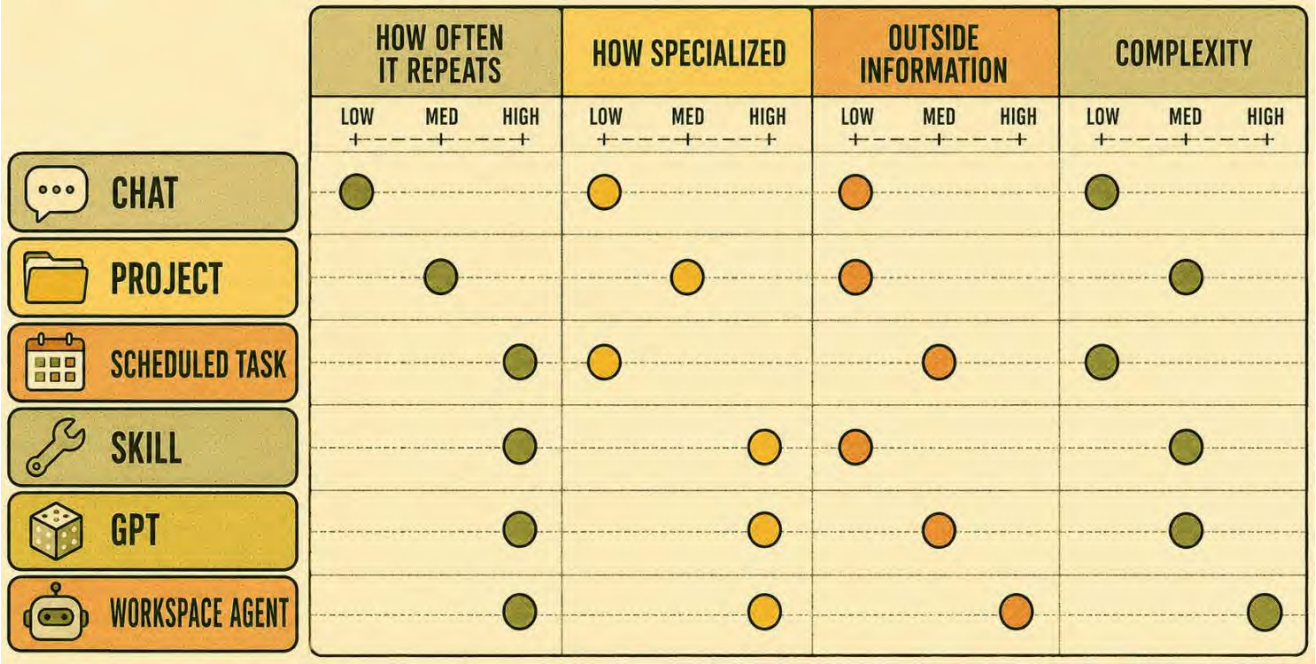


Figure 3.1A: Match your task to a building block by asking how often it repeats, how specialized it is, and how much outside access it needs

Decide when an app or MCP connection is needed

Some tasks need more than instructions. They need current information from another system, such as a calendar, a shared drive, or a company database. When that happens, an app connection or a **Model Context Protocol (MCP)** connection brings that outside system into the conversation safely, with permissions you control.

Learn more online: An optional online-only chapter is available if you want to learn about MCP, *Chapter 3.9*, *Connect Trusted Information and Tools with MCP*: `https://github.com/markjprice/chatgpt-vb/blob/main/docs/ch3-9-mcp.md`.

Good practice: Add a connection only when the task depends on outside data or action. A connection you do not need adds risk without adding value.

Reserve subagents for genuinely separable work

A subagent is a separate specialist that ChatGPT can send part of a task to, then bring the results back together. Subagents help when a large assignment splits cleanly into independent pieces, such as research, drafting, and review, that can run at the same time.

Do not reach for subagents out of habit. If one person, or one clear sequence of steps, can do the work just as well, a simpler tool will be easier to build, test, and trust.



Learn more online: An optional online-only chapter is available if you have an organization-managed workspace and therefore you can use workspace agents, *Chapter 3.10*, *Build or Use a Workspace Agent*: `https://github.com/markjprice/chatgpt-vb/tree/main/docs/ch3-10-workspace-agents.md`.

Reach an app with no connection: Computer Use and Appshots

The optional *Chapter 3.9* covers the best case: a trusted MCP server that gives ChatGPT a defined set of tools for one system. Sometimes no MCP server, plugin, or app exists for the tool you actually use. Appshots and Computer Use are what you reach for then, when ChatGPT needs to see or work with a screen directly instead of through a structured connection.

Use them as a last resort, after you have checked whether a plugin, an app, or an MCP server already does the job. A structured connection is faster, more reliable, and easier to review than reading pixels off a screen.

Give ChatGPT a snapshot with Appshots (Mac only)

Appshots hands ChatGPT a one-time picture of whatever app is currently in front of you, along with any text that app makes available. Use one when you are actively working in another app and want ChatGPT to see exactly what you see before it helps with the task, such as drafting a script from an API reference page or reviewing a budget spreadsheet.

Bring the app window you want to share to the front, then press both Command keys, or a custom hotkey you set in ChatGPT's settings. ChatGPT stores the appshot in the chat like any other attachment.

Appshots are only available in the ChatGPT desktop app on macOS.



Watch out: An appshot shares an image and any visible or hidden text from the app window with ChatGPT. Bring only the window you intend to share to the front first, and avoid taking an appshot of a window that contains sensitive information you have not already decided to share.

Let ChatGPT operate an app with Computer Use

Computer Use lets ChatGPT see a screen and control the mouse and keyboard directly, the way you would. Use it for a task that depends on a graphical interface with no dedicated integration, such as checking a desktop app's settings, reproducing a bug that only shows up on screen, or working with a data source that has no plugin.

Install the Computer Use plugin from the Plugins Directory in ChatGPT Work, then turn it on for the app you want ChatGPT to use. ChatGPT asks for your permission before it uses an app for the first time, and you can choose to always allow that app in future tasks. On macOS, grant Screen Recording and Accessibility permissions so ChatGPT can see and interact with the target app. Computer Use is available on macOS and Windows.



Good practice: Give ChatGPT one clear app or task at a time, and stay present for anything involving an account, a password, a payment, or another sensitive setting. Review the permission prompt before you approve an app.



Learn more online: *Computer Use*, an official ChatGPT reference: `https://learn.chatgpt.com/docs/computer-use`. *Appshots*, an official ChatGPT reference: `https://learn.chatgpt.com/docs/appshots`.



Watch out: On Windows, Computer Use takes over your foreground input while it runs, since it cannot work in the background the way it can on macOS. Expect ChatGPT to move your pointer and type while a Windows task runs, and plan to step away from that computer or use a different device until it finishes.

A durable progression for building automation

Use this order whenever you consider automating a task. It keeps you from building something complicated before you know the task well:

1. Have ChatGPT help with one task.
2. Repeat and refine the task or procedure manually.
3. Store its context in a project.
4. Schedule or automate only after the procedure is reliable.
5. Turn stable instructions into a reusable skill or GPT.
6. Connect approved information or tools.
7. Measure time, quality, error rates, and exceptions.
8. Delegate separable work to subagents if you have that capability.

The core question at every stage is the same: what is the least complicated mechanism that reliably handles this work?



Good practice: Start with the plainest option that could work. You can always add a skill, a connection, or a subagent later, once you know exactly where the plain version falls short.

Try it now

Pick one task you have asked ChatGPT to help with more than twice in the last month. Write down each stage of that task, in order. Then use the questions in this chapter to guess which building block fits best. You will build it in the chapters ahead.

Check the result

- ☐ Did you name a task that truly repeats, rather than a one-off request?
- ☐ Can you list its stages in order?
- ☐ Does your guess about the right building block match the reasons given in this chapter, not just the newest or most advanced option?

Choosing the right building block is only the first decision. Before you build anything, *Chapter 3.2* covers turning a manual routine into a task that runs on a schedule or a trigger.

Chapter 3.2 Schedule Recurring and Conditional Work

Work supports tasks that run once, repeat according to a schedule or trigger, or monitor for changes. This chapter covers setting up a one-time scheduled task, repeating work on a schedule, watching for a specific condition, and managing tasks once they're running.

Figure 3.2A shows how a trigger, whether a time or a condition, leads to a task running, checking a source, producing a result, and sending you a notification.

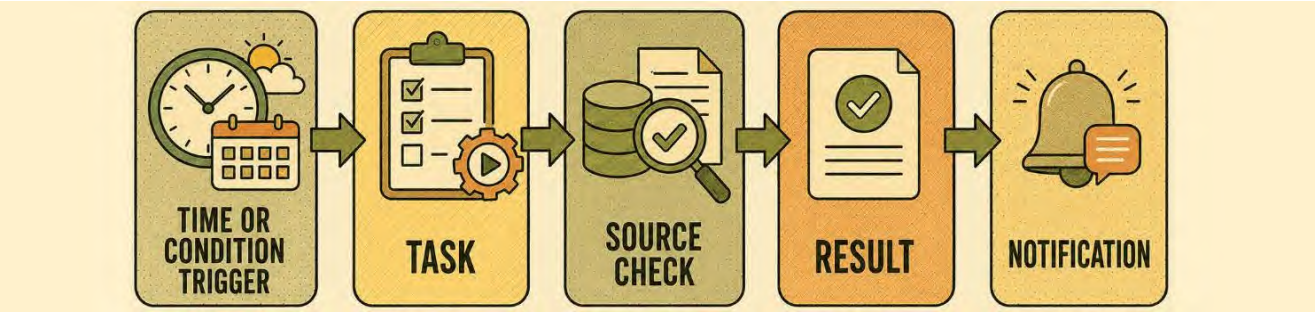


Figure 3.2A: Every scheduled task follows this same chain. If the notification stops arriving, check each stage to find where it broke



Watch out: Active task limits vary by plan: Go users can have up to 3 active tasks, Plus users up to 5, Business and Edu users up to 10, and Pro and Enterprise users up to 15. If you reach this limit, ChatGPT will not be able to create a new task or resume a paused task unless you pause or delete an existing active task or it completes per its scheduled time.

A scheduled task has two parts:

1. **The work**: What ChatGPT should do
2. **The trigger**: When ChatGPT should do it

The trigger might be a date, a repeating schedule, or a condition that ChatGPT should check.

For example, these three requests have different triggers:

```
Remind me to renew my library books tomorrow at 6 PM.

Every Thursday evening, find five things to do in my town that weekend.

Check once a day and tell me when tickets for this event become available.
```

The first task runs once. The second repeats at a set time. The third checks for a change and reports only when its condition is met.

To complete this chapter, you need a task you already do by hand on some kind of repeating basis, such as a weekly check or a monthly report. Think about any of your work that follows one of these patterns: "I do this every Monday," "I check this website every few days," "I prepare the same report each month," "I need to remember this on a particular date," or "I'm waiting for something to change."

The best first task is small, useful, and easy to check. I discourage you from beginning with anything that automatically spends money, sends messages, submits forms, deletes information, or makes an important decision!

If you can't think of anything, you can use mine (and adapt it to your own location). Every Thursday evening, I used to go to ChatGPT and enter the following prompt:

```
Find five interesting things happening in or near Brighton this weekend.
```

By the end of this chapter, you'll be able to turn that manual routine into a scheduled or condition-based task.



Good practice: Before scheduling recurring work, try the prompt once in a normal chat. Check whether the result contains the right information and level of detail. Scheduling an unclear prompt only repeats the unclear result.

Create a one-time scheduled task

A one-time task runs once at a future time. It can remind you to do something, or it can complete a piece of work and return the result.

A reminder only needs to state what you need to remember:

```
Remind me at 4 PM tomorrow to send the revised proposal to my publisher.
```

ChatGPT can also perform work at a scheduled time. For example:

```
This Thursday at 6 PM, find five interesting public events happening in Brighton
this coming weekend.

For each event, include:
- Event name
- Date and time
- Location
- Admission price, or "Free"
- A link to the official event page
- One sentence explaining who might enjoy it

Use current information from official event or venue websites. Do not include an
event if its date or availability cannot be confirmed.
```

This task does more than issue a reminder. At the scheduled time, ChatGPT searches for information and prepares the requested list.

Include the following details when they matter:

- The date or amount of time to wait
- The time and time zone
- The work to perform
- The sources to use
- The required output
- Any limits or checking rules

Use a specific time when the result must arrive at that time. You can also use a broad period such as "tomorrow morning" when the exact minute does not matter.

After you submit the request, ChatGPT should show a task confirmation. Check the date, time, and description before leaving the conversation:

1. Select the newly created scheduled task to open its details.
2. In the left-side navigation, select **Scheduled** to view and manage all scheduled tasks:

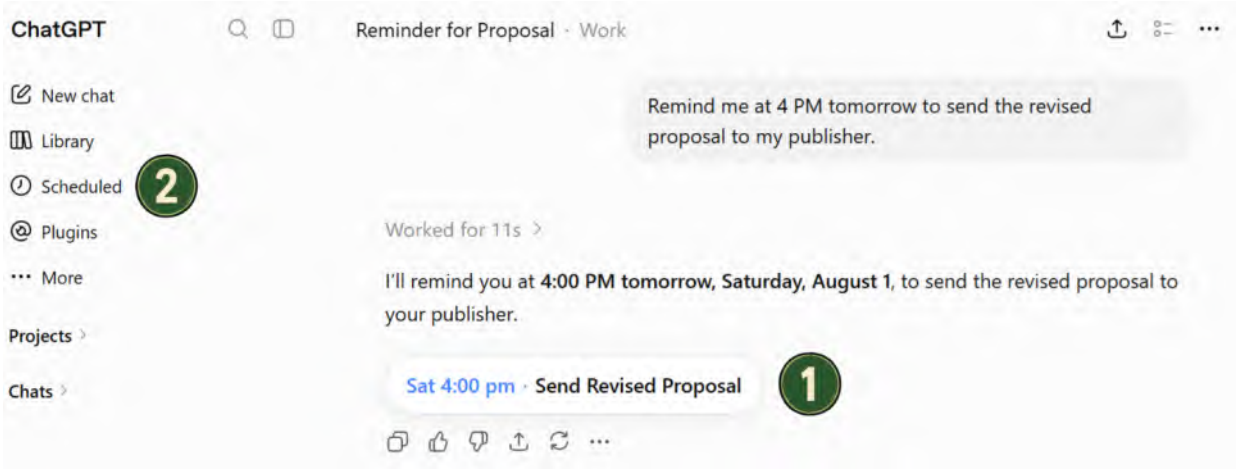


Figure 3.2B: Scheduling a one-time task

3. On the scheduled task details page, select the **...** menu for actions like **Run now**, **Notification settings**, and **Delete**.
4. Select the **Pause** button to temporarily stop the scheduled task from running. It is a toggle, so selecting it again resumes the task:

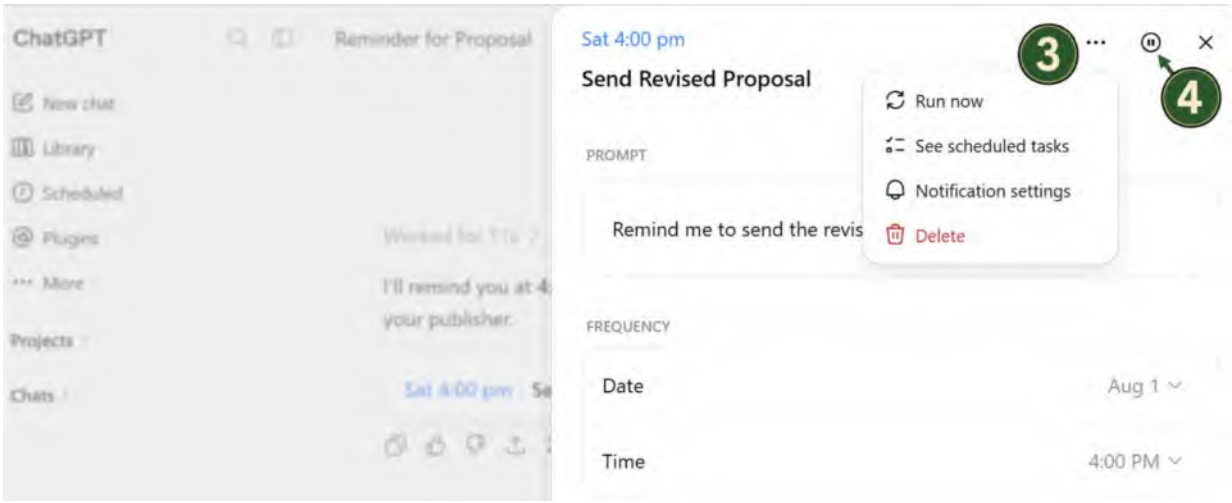


Figure 3.2C: Viewing details of a one-time task

Watch out: Make sure ChatGPT has created a scheduled task. A promise written only as ordinary text may not mean that a task was successfully scheduled.

Repeat work daily, weekly, or monthly

Once a one-time task produces the result you want, turn it into a recurring task on the schedule that fits: daily, weekly, or monthly. Choose a time that gives the task access to fresh information, such as after a weekly report normally arrives.

A repeating task must be able to work with new information each time. Avoid wording that depends on an earlier answer unless the task can return to the same chat and use its context.

For a monthly report, define the reporting period carefully:

```
On the first weekday of every month, prepare a summary of the previous calendar
month.

Use [name the connected source or information]. Report:
- The three most important results
- Any change from the previous month
- Problems that need attention
- Missing or incomplete information
- One recommended next step

Do not invent an explanation for a change. Mark anything that needs human review.
```

"Previous calendar month" is clearer than "recent results." It tells ChatGPT where one reporting period ends and the next begins.

Recurring work often needs updated information. Make sure the task can still reach its sources when it runs. A file stored only on your computer may not be available to a web task. A connected app may also stop working if its permission expires or your organization disables it.

Review the first few results. Adjust the prompt if the task returns too much information, repeats unimportant details, or misses the reason you scheduled it.



Good practice: Begin with a weekly schedule unless the information changes quickly. A daily task that rarely finds anything useful soon becomes noise. Name each scheduled task clearly enough that you recognize its purpose months later, without needing to open it to remember what it does.



Learn more online: OpenAI's help center article on scheduled tasks lists current limits on how many active tasks you can run and how often they can check for updates on your plan. Check it before scheduling a task that needs frequent runs: `https://help.openai.com/en/articles/10291617-tasks-in-chatgpt`.

Monitor for a condition or change

Some tasks don't need a fixed schedule. Instead, ask ChatGPT to check periodically for a specific condition. This could be a price drop, a new message matching certain criteria, or a change to a tracked page. Ask it to notify you only when that condition is met.

This is useful when you are waiting for tickets or appointments to become available, a product to return to stock like a Mac Studio with lots of unified memory, an application status to change, a new software version to be released, an official report or result to be published, or a price to cross a limit.

A condition-based request needs four pieces of information:

- What ChatGPT should check
- Which source it should use
- What counts as a meaningful change
- What it should do when nothing has changed

For example:

```
Check the official event page at [paste official URL] once a day.

Notify me only when tickets become available for [event name] on [preferred date].

When tickets become available, include:
- The available date and time
- The listed price
- The direct booking link
- Any booking deadline or stated restriction

Use only the official event or venue website. If tickets are not available, do not
notify me.

Stop checking after [final date].
```

“Check for updates” is too broad. The example defines the exact event that should cause a notification.

The instruction not to notify when nothing has changed prevents an empty daily message. The stop date also prevents ChatGPT from checking forever after the information is no longer useful.

For a news or product watch, define what counts as meaningful:

```
Check official OpenAI announcements once a day for a new ChatGPT feature that is
available to individual users in the UK.

Notify me only when a new feature is released or an announced release date is
confirmed.

Do not notify me about rumors, repeated coverage of an earlier announcement, minor
wording changes, or features available only through the API.

Include a link to the official announcement and state which plans or regions are
included.
```

A monitoring task still checks on a schedule. The condition controls whether you receive a notification. It does not mean that ChatGPT watches the source every second. Read the condition as a yes-or-no question. Could you tell clearly whether it has been met? If not, make the condition more specific.



Good practice: Do not rely on one monitoring task for an emergency, medical warning, financial trade, legal deadline, or safety alarm. Use the official alert service for anything where a late or missed notification could cause harm.



Real-world example: A job seeker sets up a weekly task that checks three job boards for new postings matching a specific title and location, and gets a summary only when a new match appears.

A scheduled task does not need to produce work on every run. A recurring report should return a report each time. A monitoring task should stay quiet until its stated condition is met. This distinction is an important point. Understanding this will prevent you from creating daily "no change" notifications when what you really want is a condition watch.

Review, pause, edit, and delete tasks

Check your list of scheduled tasks regularly. Open **Scheduled** from the ChatGPT sidebar to review your tasks and their recent results:

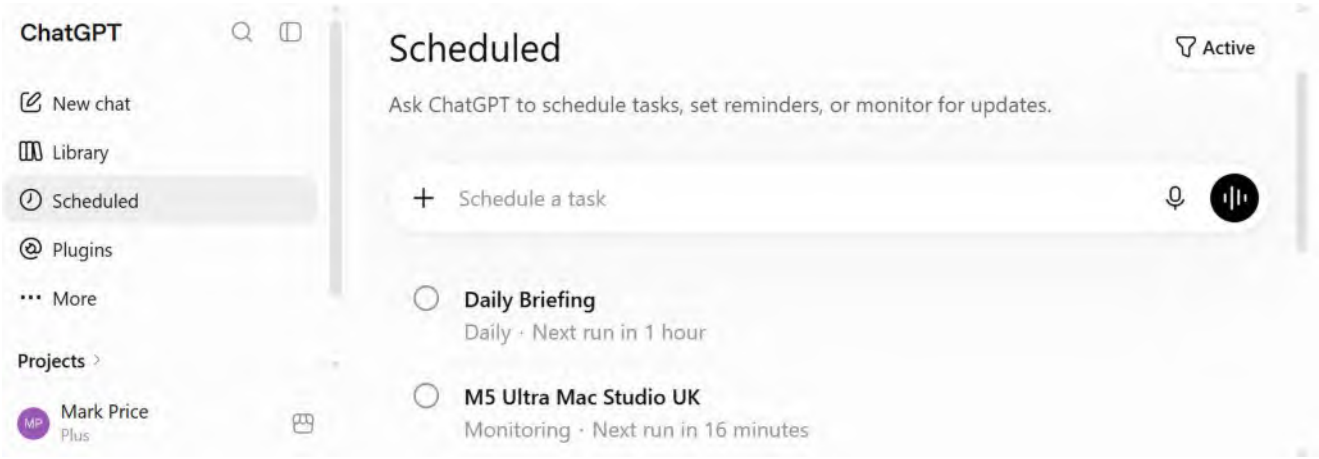


Figure 3.2D: Review your scheduled tasks regularly

Edit a task when its purpose remains useful, but its instructions or timing need to change.

For example, you could change the Brighton Weekend Planner by asking:

```
Change the Brighton Weekend Planner to run every Friday at 8 AM instead of every
Thursday at 6 PM.

Keep its location, event requirements, source rules, output format, and word limit
unchanged.
```

A narrow editing instruction helps prevent other parts of the task from changing accidentally. After editing, confirm the next run. Do not assume that changing the wording also changes the schedule.

Pause a task when you may need it again but do not want it to run now. For example, pause a weekly job search while you are not looking for work. A paused task keeps its instructions but does not run. Resume it when you are ready. Check the next run after resuming because the original time may already have passed. ChatGPT may also pause unattended tasks after a period of inactivity. Review Scheduled if an expected result stops arriving.

Delete a task when you no longer need its instructions or run history. Deleting a task stops future runs. Deleting the task does not necessarily delete its associated chat. Pause before deleting if you are unsure. Pausing is easier to reverse. Deleting a chat associated with a scheduled task may pause the task. Check Scheduled after reorganizing or deleting chats.



Watch out: A scheduled task can drift away from its original purpose without anyone noticing and keep running while producing results nobody reviews. Set a reminder to review your scheduled tasks every few months.

Try it now

Pick one routine check you currently do by hand, such as reviewing a report every Monday or watching for a change on a webpage. Set it up first as a one-time task, confirm the result, then make it recurring or condition-based.



Good practice: Review your scheduled tasks once a month. You could even create a scheduled task to remind you to do this! Pause or delete tasks that no longer help, repair tasks with failed runs, and revise tasks that create repeated noise.

Check the result

- ☐ Did the scheduled task run at the time and frequency you intended?
- ☐ Does the result match what you'd have found by checking manually?
- ☐ Would you notice quickly if this task stopped running or started producing an unhelpful result?

With scheduled tasks under your control, you're ready to build your first automation. *Chapter 3.3* shows you how to turn a procedure you already repeat by hand into a reusable skill.

Chapter 3.3 Turn a Procedure Into a Reusable Skill

In *Chapter 3.1* you learned that a skill stores a tested procedure so ChatGPT can repeat it consistently. This chapter shows you how to prepare a procedure before you build it. I will then show you how to create, test, and share the finished skill.

Good preparation counts for more than the building step itself. A skill built from a vague procedure will behave inconsistently, no matter how carefully you set it up.

Select a stable repeated procedure

Choose a task you have done several times already. It should have a result you were satisfied with more than once. A stable procedure has steps that stay roughly the same each time, even if small details change. Avoid a procedure that is still evolving. If you changed your approach last week and expect to change it again, wait until it settles down before turning it into a skill.



Real-world example: A marketing coordinator writes a short social media summary of each blog post the team publishes. She has done this the same way for two months: read the post, pull out the main point, write three sentences, and add one call to action. That consistency makes it a strong candidate for a skill.

In the sections that follow, we will build a Blog Post Social Summary skill as described in the preceding real-world example box.

Define required inputs and outputs

Write down exactly what the skill needs to start, and exactly what it should produce when it finishes:

- **Inputs** might include a document, a topic, a target audience, or a word limit
- **Outputs** might include a document, a table, a short summary, or a checklist

Being specific here prevents a common problem: a skill that runs but produces something close to what you wanted instead of what you actually asked for.

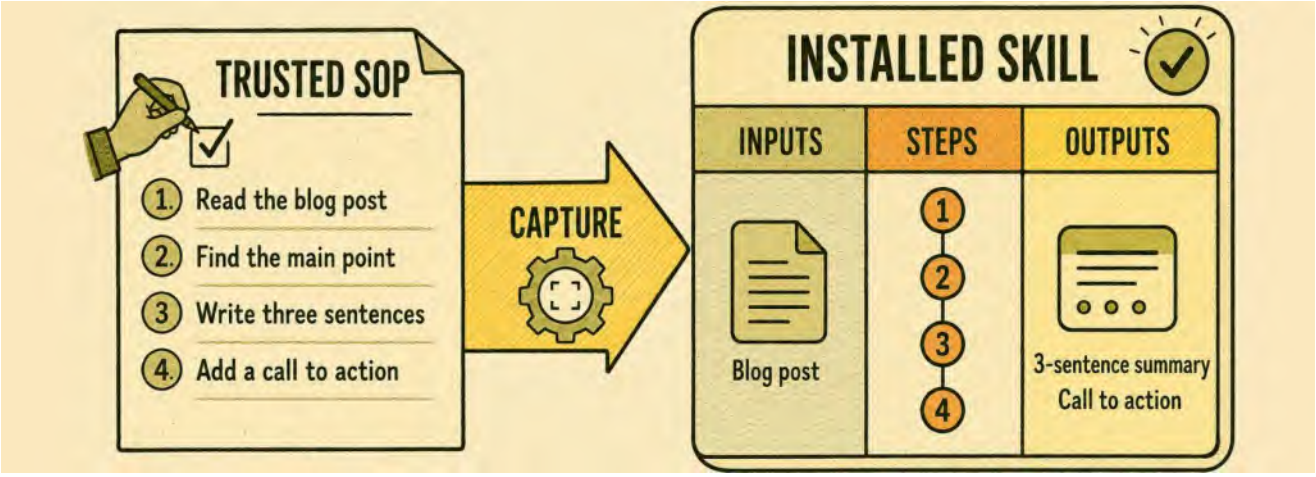


Figure 3.3A: A reusable skill starts as a procedure you have already tested by hand

Capture rules, resources, and examples

List the rules the procedure must follow, such as a required format, a length limit, a tone, or a fact that must always be checked. Then gather any resources the skill will need, such as a template, a style guide, or a reference document.

Add at least one worked example: a real input and the exact output you wanted from it. An example often communicates a rule faster than a written instruction can.

Write acceptance criteria

Acceptance criteria are the specific conditions a result must meet before you consider the skill's work finished. Write them as a short checklist, not a paragraph.

For the marketing summary example, acceptance criteria might include:

- Exactly three sentences, not four or two
- One clear call to action at the end
- No claim that is not stated in the original blog post



Sample files: A one-page Word specification for a Blog Post Social Summary skill: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book3/3-3-Blog-Post-Social-Summary-Skill-Specification.docx`. It includes numbered fields for name, inputs, steps, rules, worked example, and acceptance criteria.



CHAPTER 3.3 • EXAMPLE SPECIFICATION

Blog Post Social Summary

A written specification for a reusable ChatGPT skill

1	Name	Blog Post Social Summary
2	Inputs	• Required: the published blog post as pasted text, an uploaded document, or an accessible link. • Required: the target audience, such as new customers or event attendees. • Optional: a preferred call to action. If none is given, use an action stated in the source.

Figure 3.3B: A concise skill specification names its inputs, its steps, its rules, and the criteria a good result must meet



Good practice: Write the acceptance criteria before you build the skill, not after. Deciding what "good" looks like in advance keeps you from lowering your standard without noticing once you see the first result.

You now have a clear specification: inputs, rules, an example, and acceptance criteria. Turning a procedure into instructions is only half the job.

The next sections turn that specification into a working skill, test it, and share it if it proves reliable. ChatGPT can create skills through conversation, install skills that others have shared, and distribute a successful skill within a workspace.

Create the skill through conversation

Let's create a skill:

1. Open a new chat and tell ChatGPT you want to create a skill.
2. Provide the inputs, rules, resources, and examples from your specification.
3. Ask ChatGPT to draft the skill's instructions and confirm the required inputs and outputs.
4. Review the draft against your acceptance criteria before saving it.

Use the skill specification from the previous section as an attachment, then enter this prompt:

```
Create a reusable skill from the attached Blog Post Social Summary specification.
Help me refine it through conversation before you create it. Ask one focused
question at a time if any required detail is missing or unclear.

When you have enough information, show me a concise summary of the skill's purpose,
inputs, steps, rules, example output, and acceptance checks. Wait for my approval,
then create, validate, and install the skill. Keep it simple and do not add features
that are not in the specification.
```

Review the skill

You should now see a saved skill with:

1. Skill name (in this case, **Blog Post Social Summary**)
2. The skill description
3. How to invoke this skill (in this case, `$blog-post-social-summary`):

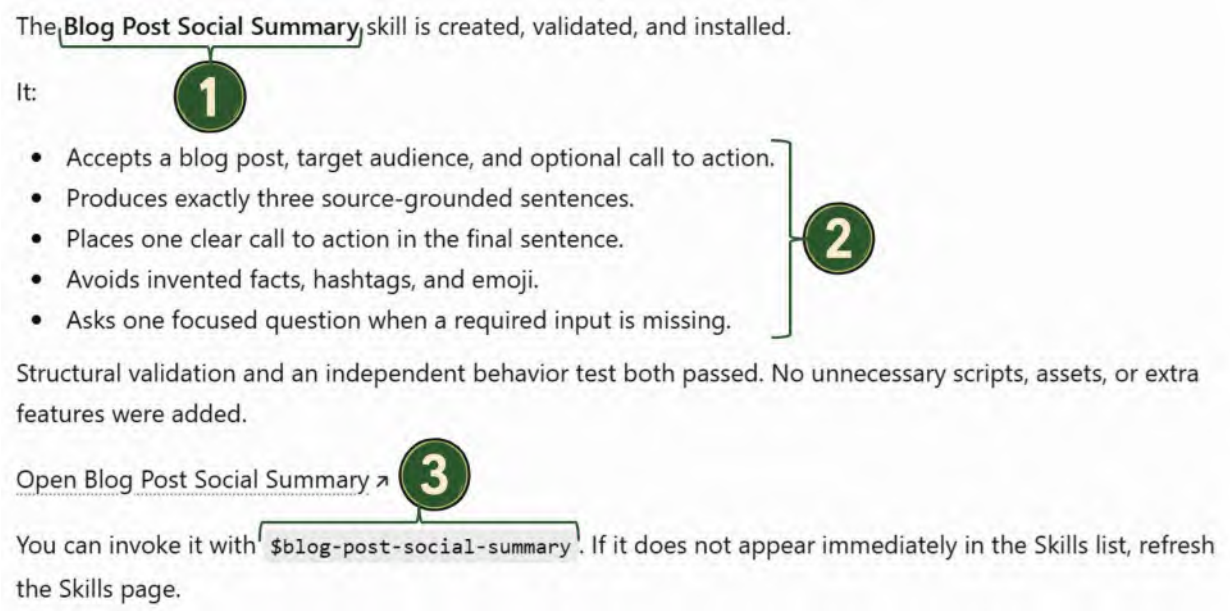


Figure 3.3C: ChatGPT response after creating and installing a new skill

The **Skills** page lists every skill you have created or installed, with its current status:

1. In the left-side navigation, select **Plugins**.
2. Select **Skills**:

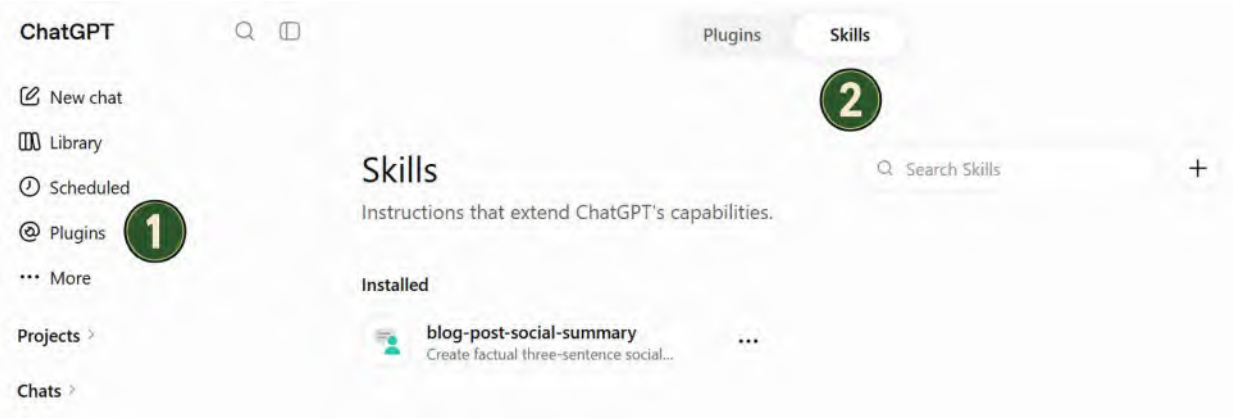


Figure 3.3D: The Skills page lists every skill you have created or installed, with its current status like Installed

Test a normal case

Run the skill on a normal example first, one like the worked example in your specification. Check the result against your acceptance criteria. Use the installed Blog Post Social Summary skill with this prompt:

```
Test the Blog Post Social Summary skill with this normal example:

Blog post: ClearPath will offer free 30-minute online planning clinics every
Thursday at 1:00 p.m. ET, starting September 10. Customers can bring one question
about setting up their first project. Registration is required through the ClearPath
website.

Audience: New ClearPath customers

First, run the skill and show its three-sentence social summary. Then check the
result against every acceptance criterion in the skill specification. For each
criterion, report Pass or Fail and briefly explain why. Do not change the result
during the check. If anything fails, suggest a correction separately. Do not test a
difficult case yet.
```

The response should show two distinct parts:

1. The result produced by the skill.
2. An acceptance check confirming whether it contains exactly three sentences, ends with a clear call to action, and uses only facts traceable to the source:

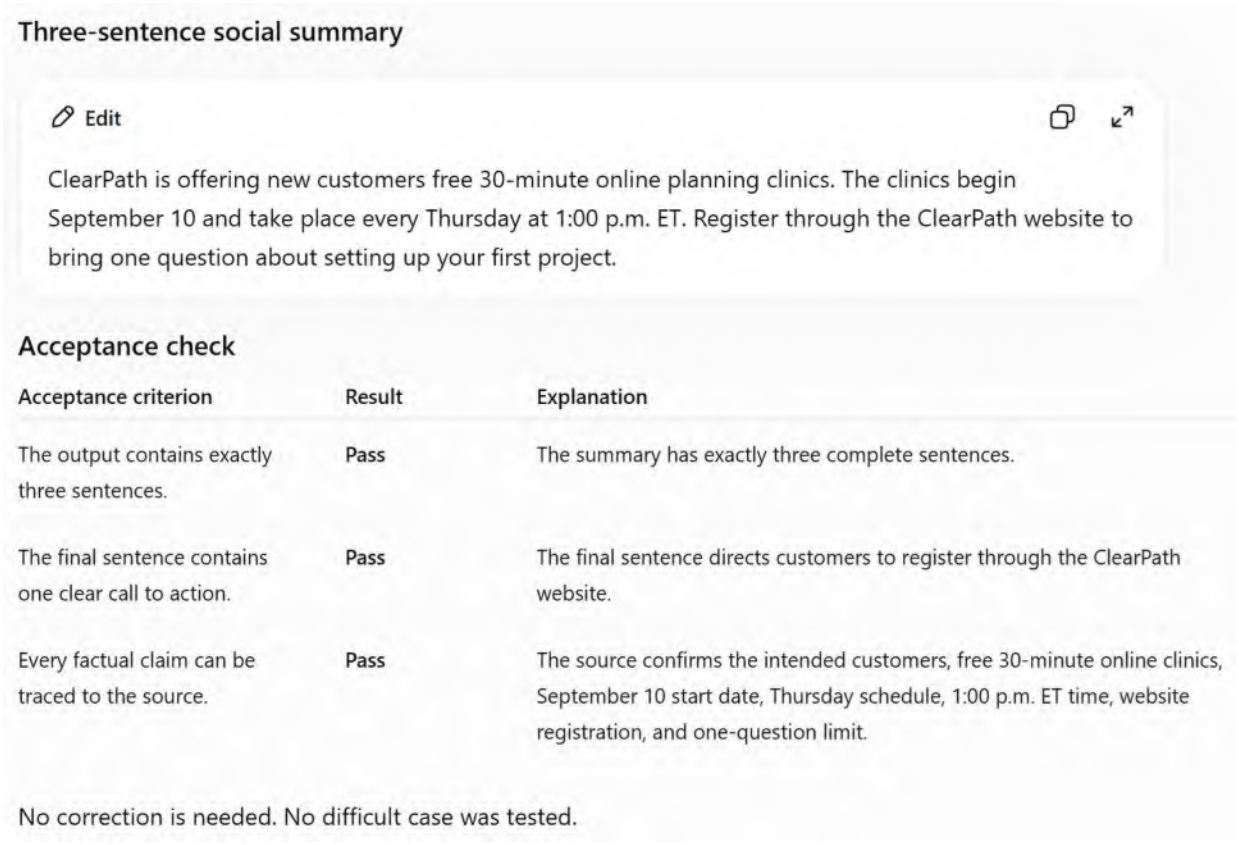


Figure 3.3E: A test run of the skill and an acceptance check

Keeping the first test simple establishes that the skill works under ordinary conditions before the reader challenges it with missing, ambiguous, or potentially misleading inputs.

Then test at least one difficult case: incomplete information, an unusual request, or an input that breaks the expected pattern. A skill that only works on easy cases is not ready to trust.

Test a difficult case

Choose a difficult case that tests one of the skill's rules. For the Blog Post Social Summary skill, leave out the required target audience. The skill should notice the missing input and ask one focused question instead of guessing.

Enter this prompt:

```
Test the Blog Post Social Summary skill with this incomplete input:

Blog post: ClearPath has published a new guide to planning a first project. The
guide explains how to set a goal, divide the work into smaller tasks, and choose
a realistic completion date. Readers can download the guide from the ClearPath
website.

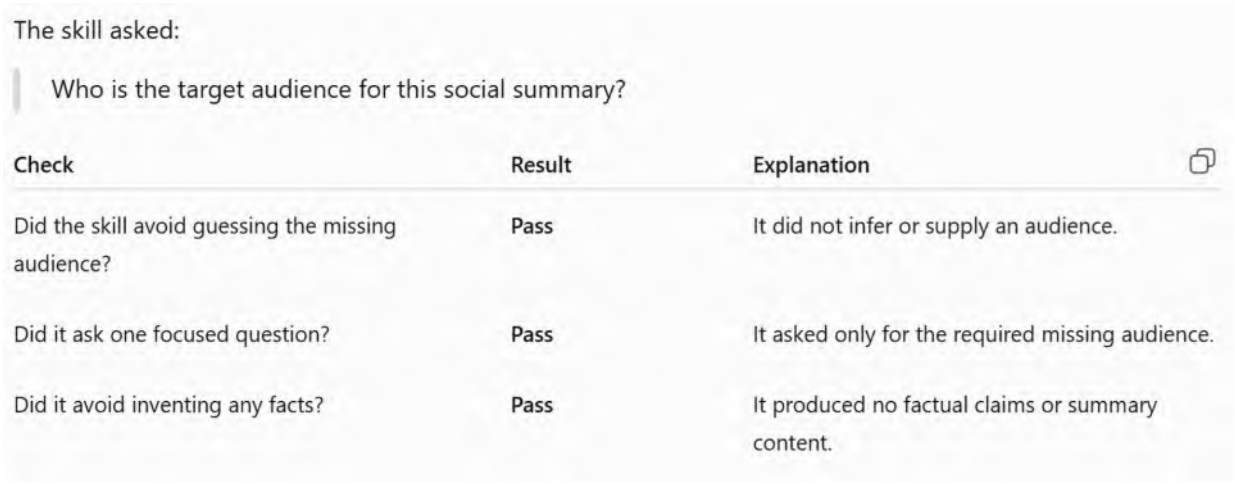
I have not provided a target audience.

Run the skill once. If it asks a question, stop and show that question. If it
produces a summary, show the summary without changing it. Then report Pass or Fail
for each of these checks:

Did the skill avoid guessing the missing audience?
Did it ask one focused question?
Did it avoid inventing any facts?

Do not correct the result during the test.
```

A passing result should ask for the target audience and wait for your answer. It should not produce the three-sentence summary yet. If it guesses an audience or continues without one, the skill has failed one of its own rules:


The skill asked:

> Who is the target audience for this social summary?

Check	Result	Explanation
Did the skill avoid guessing the missing audience?	Pass	It did not infer or supply an audience.
Did it ask one focused question?	Pass	It asked only for the required missing audience.
Did it avoid inventing any facts?	Pass	It produced no factual claims or summary content.

Figure 3.3F: The skill asks you a question and waits for your response

Reply to the question with:

```
New ClearPath customers who are planning their first project.
```

The skill should now produce its summary. Check the finished result against the original acceptance criteria: exactly three sentences, one clear call to action in the final sentence, and no facts that cannot be traced to the source:

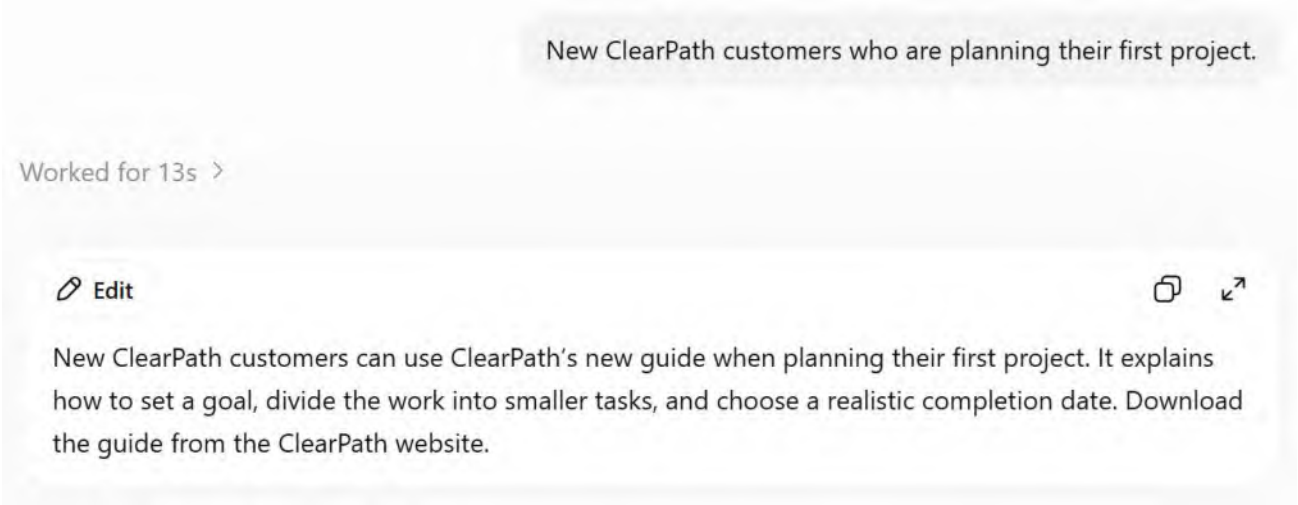


Figure 3.3G: The skill continues and produces a response

If a test fails, do not keep correcting each result by hand. Return to the skill instructions and improve the rule that allowed the failure. Then repeat both the normal test and the difficult test. This checks that your change fixes the problem without breaking behavior that already worked.



Watch out: Do not judge a skill after one successful run. Test it three or four times with different inputs before you rely on it for real work.

Invoking the skill with required inputs in the prompt

Once the skill has passed its normal and difficult tests, you can use it on a real task. Invoke the skill by entering its name with a dollar sign, then provide the required inputs in the prompt.

For example, enter:

```
$blog-post-social-summary

Blog post: ClearPath will offer free 30-minute online planning clinics every
Thursday at 1:00 p.m. ET, starting September 10. Customers can bring one question
about setting up their first project. Registration is required through the ClearPath
website.

Target audience: New ClearPath customers
```

The skill produces:

```
Starting September 10, ClearPath will offer free 30-minute online planning clinics
every Thursday at 1:00 p.m. ET. New ClearPath customers can bring one question about
setting up their first project. Register through the ClearPath website to attend.
```

The invocation tells ChatGPT which skill to use. The blog post and target audience supply the inputs required by that skill. You do not need to repeat the skill's steps or rules because they are already stored in its instructions.

In the real world, you would normally invoke the skill with a public link. You can test it with a sample file for this book:

```
$blog-post-social-summary

Blog post:
https://raw.githubusercontent.com/markjprice/chatgpt-vb/main/Files/Book3/3-3-blog-
post.md

Target audience: New ClearPath customers who are planning their first project.
```

The post announces free weekly planning clinics and includes dates, times, eligibility, capacity, registration details, and a clear next action. It also identifies ClearPath as fictional, preventing readers from mistaking the example for a genuine offer.

Invoking the skill interactively

You can also invoke the skill and wait for it to ask for the information it needs:

1. Type the `$` command to invoke the named skill:

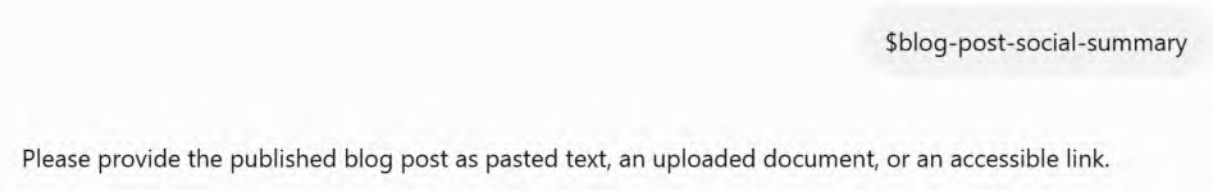


Figure 3.3H: Invoking a skill interactively. ChatGPT asks for the first input: the published blog post

2. Answer the first question by entering a link:

```
https://raw.githubusercontent.com/markjprice/chatgpt-vb/main/Files/Book3/3-3-
blog-post.md
```

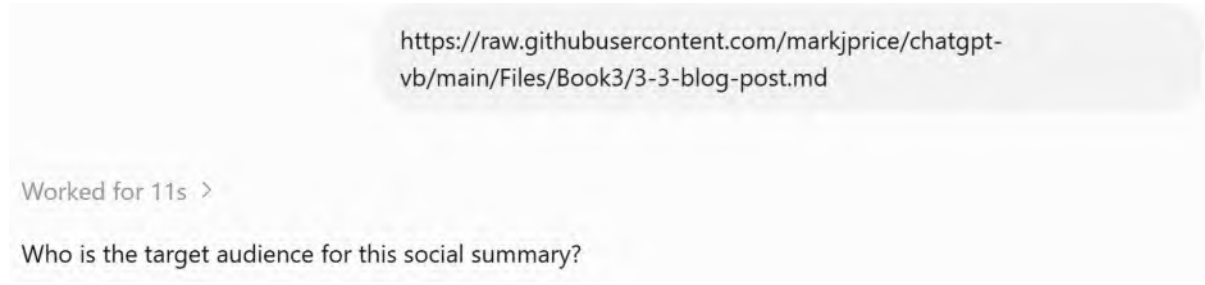


Figure 3.3I: ChatGPT asks for the second input: the target audience

3. Answer the second question by describing the target audience:

```
New ClearPath customers who are planning their first project.
```

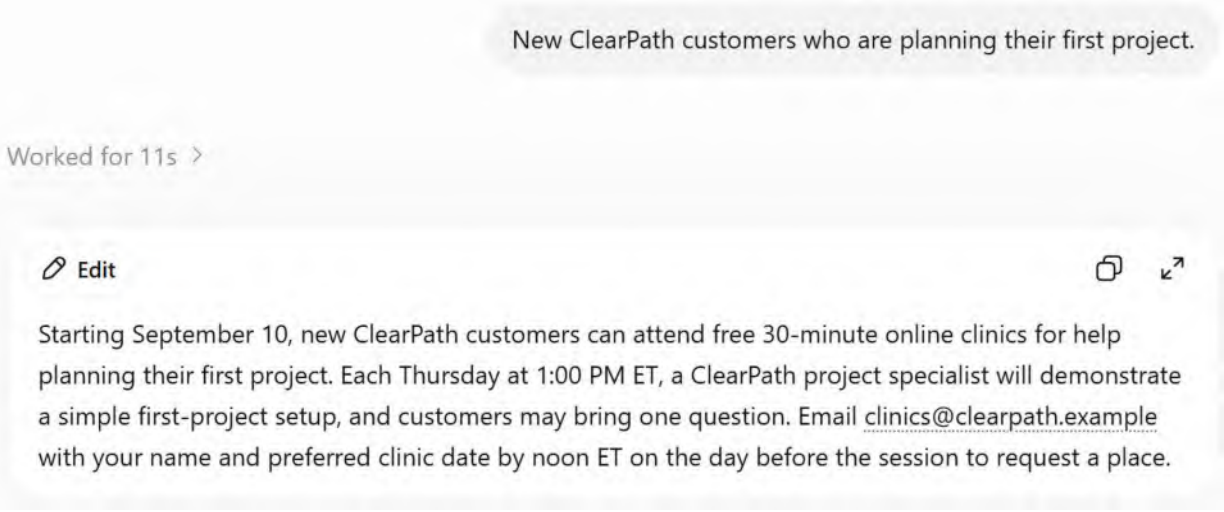


Figure 3.3J: ChatGPT responds to the skill invocation

Improve instructions and error handling

When a test fails, do not try again and hope for a better result. Read the failure, decide which instruction was missing or unclear, and update the skill's instructions directly.

Add error handling for cases where the skill cannot complete safely, such as missing information. Tell the skill to ask a clarifying question or stop and explain what it needs, rather than guessing.

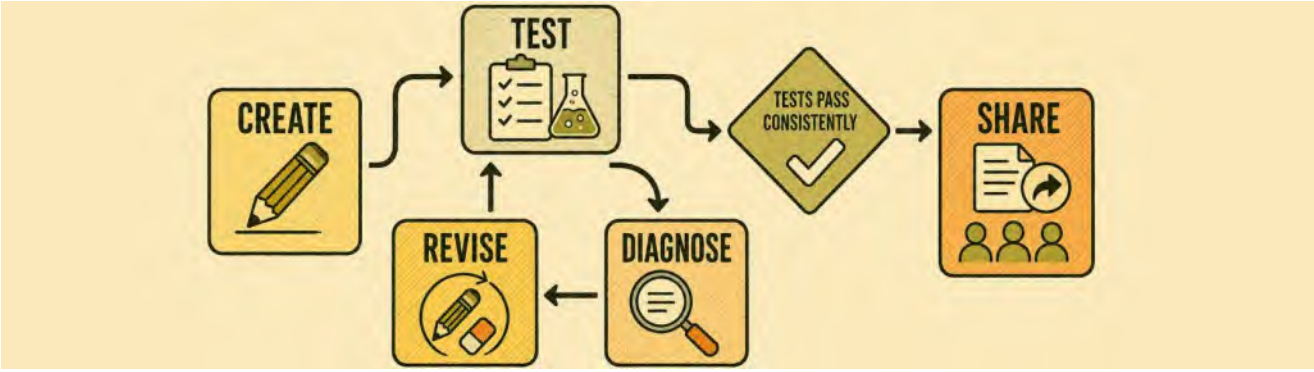


Figure 3.3K: Building a reliable skill means moving through create, test, diagnose, and revise, sometimes more than once

Install or share the finished skill

Once a skill passes your tests consistently, you can keep it for your own use or share it. In an organization-managed workspace, a shared skill becomes available to colleagues who have permission to use it.

1. Open the skill's settings.
2. Choose to keep it private or share it with your workspace.
3. If sharing, write a short description so colleagues understand what the skill does and what it needs.
4. Confirm the sharing setting.



Good practice: Before sharing a skill, ask a colleague unfamiliar with the task to run it once. If they can use it correctly from your description alone, it is ready to share more widely.

Try it now

Choose the procedure you identified in *Chapter 3.1*. Write its inputs, outputs, rules, one worked example, and three acceptance criteria. Keep this list. Then build the skill you specified. Run it once on a normal case and once on a difficult case. Revise its instructions at least once based on what you observe.

Check the result

- ☐ Have you done this procedure the same way at least twice before?
- ☐ Could someone else follow your written inputs, rules, and example without asking you a question?
- ☐ Do your acceptance criteria describe something you can check by looking at the result, rather than a vague feeling?
- ☐ Did the skill meet every acceptance criterion on the normal test?
- ☐ What happened on the difficult test, and did you update the instructions in response?
- ☐ Would you feel comfortable sharing this skill with a colleague today?

A skill handles one procedure well. *Chapter 3.4* shows you how to go further and build a specialized GPT: an assistant shaped around a whole area of your work, not just a single task.

Chapter 3.4 Build a Custom GPT

A skill stores one procedure. A Custom GPT is a larger step: a complete assistant, built for one purpose, with its own instructions, reference material, and tools. This chapter shows you when a GPT is worth building, and how to build one that stays useful over time.

We will build the ClearPath Workshop Planner. It continues the fictional ClearPath example from *Chapter 3.3*, but expands from one narrow skill into a specialized assistant that can:

- Apply permanent program and brand rules
- Analyze previous workshop results
- Research current venues and prices
- Develop and revise longer workshop plans
- Create a draft promotional image

This gives us a meaningful reason to use instructions, knowledge files, conversation starters, and three capabilities.



Watch out: Creating or editing a GPT currently requires a paid ChatGPT subscription. Workspace settings may also control whether GPT creation or sharing is available.

Decide whether a GPT is the right mechanism

Build a GPT when you need more than one skill working together, when the assistant should always start from the same role and knowledge, or when several people will return to the same kind of task. A single skill is often enough for one procedure. A GPT is better when you want a consistent assistant for a whole area of work, such as answering questions about your team's policies or helping with every stage of client onboarding.

If you are unsure, start with a skill. You can combine skills and knowledge into a GPT later, once you know which ones you use most.

Discover and create custom GPTs

The GPT's instructions describe its role, what it should do, what it should avoid, and how it should respond. Write these instructions the way you would brief a new colleague: clear, specific, and free of assumptions they cannot know.

To create a custom GPT:

1. In the left-side navigation, select **... More**.
2. Select **GPTs**.
3. Type in the search box to find GPTs or browse in categories.
4. Select the **+ Create** button to create a new custom GPT:

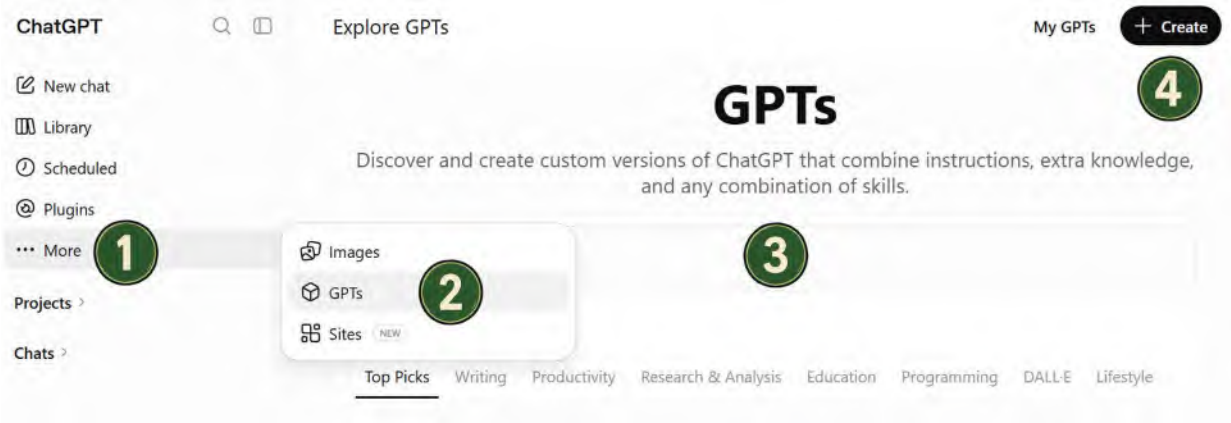


Figure 3.4A: GPTs page where you can discover or create custom GPTs

Conversation starters are short example prompts shown to anyone opening the GPT. They help people understand what to ask without guessing.

Configure a custom GPT

To configure a custom GPT:

1. On the **New GPT** page, select **Configure**.
2. **Name** and **Description**: Write the GPT's name and a short description of its purpose.
3. **Instructions**: Write clear instructions covering its role, scope, and any rules it must follow.
4. **Conversation starters**: Add two to four conversation starters that show typical requests:

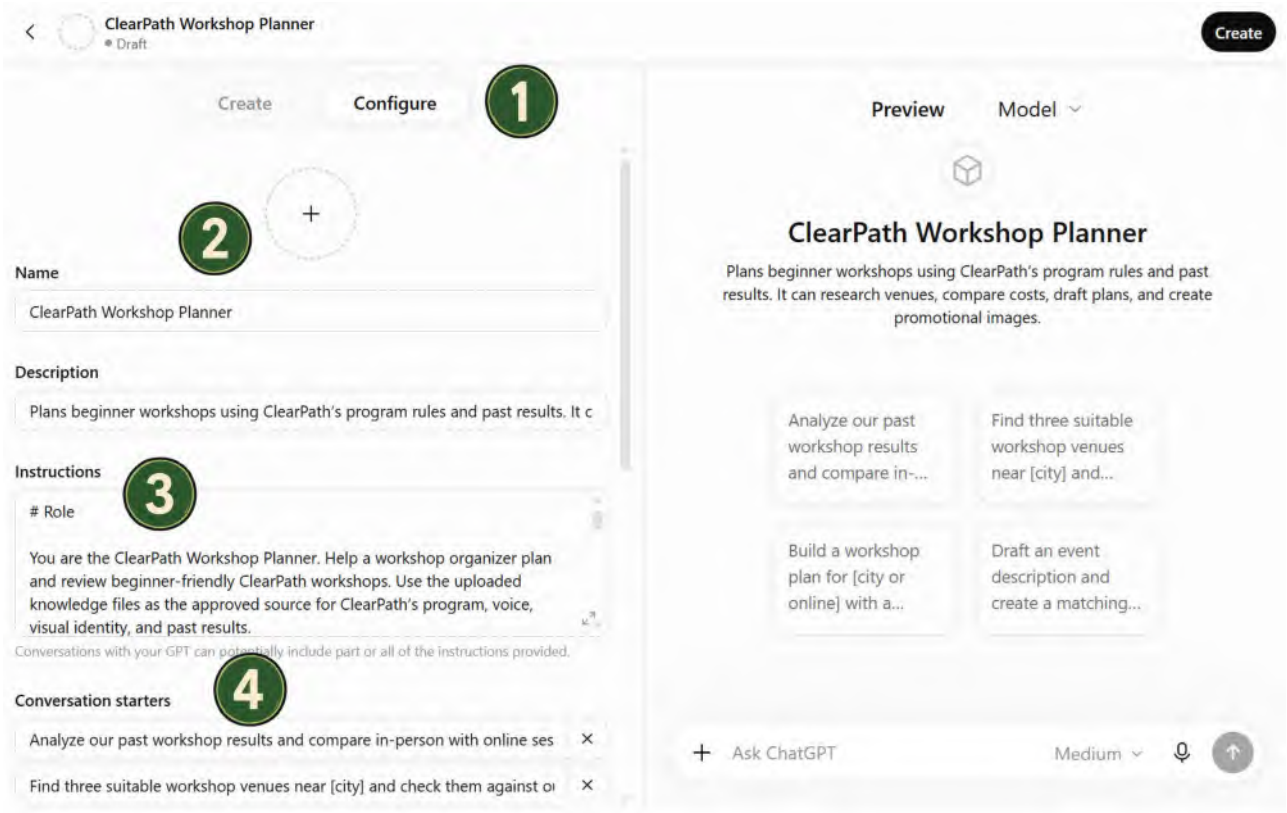


Figure 3.4B: A GPT's configuration combines a name, a description, written instructions, and example conversation starters

Use *Table 3.4* to complete the required information for your GPT:

Property	Value
Name	ClearPath Workshop Planner
Description	Plans beginner workshops using ClearPath's program rules and past results. It can research venues, compare costs, draft plans, and create promotional images.
Instructions	See *Instructions* section below.
Conversation starters	Analyze our past workshop results and compare in-person with online sessions. Find three suitable workshop venues near [city] and check them against our program rules. Build a workshop plan for [city or online] with a budget of [amount]. Draft an event description and create a matching promotional image.
Knowledge	Upload these three files: • `3-4-ClearPath-Workshop-Program-Guide.md` contains the workshop format, agenda, budget rules, accessibility requirements, approval limits, and permitted claims. • `3-4-ClearPath-Voice-and-Visual-Guide.md` controls promotional copy and image style. • `3-4-ClearPath-Past-Workshop-Results.csv` contains 12 fictional workshop records for analysis.

Recommended Model	Leave as **No Recommended Model**.
Capabilities	Select: **Web Search**, **Image Generation**, **Code Interpreter & Data Analysis**

Table 3.4A: GPT properties to set



Good practice: These starters each reveal a different part of the GPT. Conversation starters should be realistic example prompts that show users how to begin. Keep behavior and workflow rules in the **Instructions** field while using uploaded knowledge files as reference material. I recommend explicit steps and headings for multi-stage workflows.

Capabilities to turn on and why our GPT needs them are shown in *Table 3.4B*:

Capability	Why the GPT needs it
Web Search	Finds current venue information, prices, opening hours, accessibility evidence, and travel details.
Image Generation	Creates a draft promotional illustration using the uploaded visual guide.
Code Interpreter & Data Analysis	Calculates attendance rates, costs, satisfaction averages, and comparisons from the CSV file. It can inspect CSV and spreadsheet data, run Python calculations, and create tables and charts.

Table 3.4B: Capabilities and why our GPT needs them

Leave **Actions** unconfigured. This GPT prepares recommendations and drafts. It should not book venues, publish events, spend money, or contact attendees.

Use these instructions

Download the instructions as a Markdown file from the following link and then copy and paste them into the **Instructions** box:

`https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book3/3-4-gpt-instructions.md`

Add controlled knowledge files and capabilities

You can upload reference files, such as a style guide, a product catalog, or a policy document, so the GPT can use them when responding. Upload only material you have the right to share and remove anything confidential that is not needed for the GPT's purpose.

You can also turn on capabilities such as web browsing, file creation, or specific tools, depending on what the GPT's tasks require. Turn on only the capabilities the GPT actually needs. Extra access adds risk without adding value, echoing the same principle from *Chapter 3.1*.

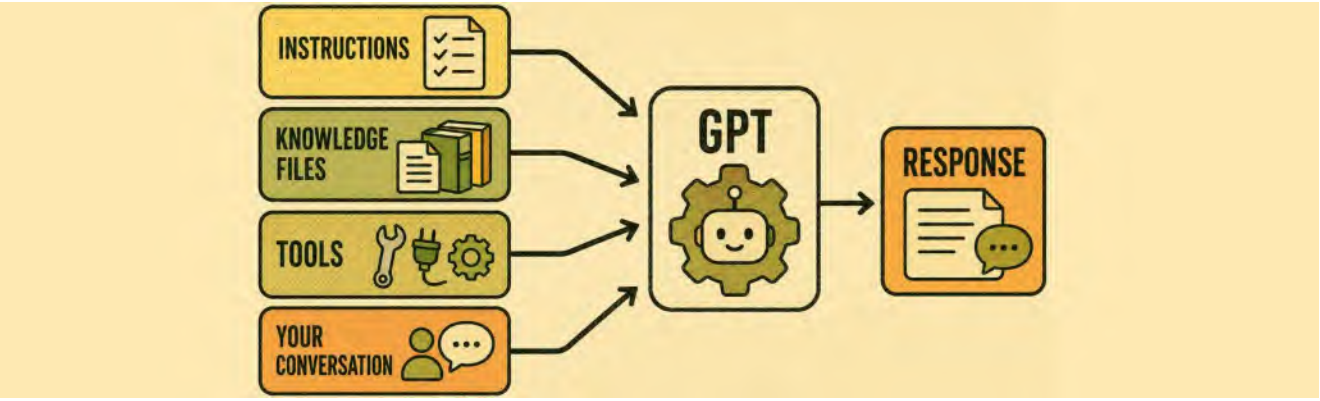


Figure 3.4C: A GPT's response comes from its instructions, its knowledge files, its available tools, and the conversation itself

Preview the GPT

Use the **Preview** page to try out these four prompts in order:

1. **Knowledge:** `What is ClearPath's standard workshop format?`
2. **Analysis:** `Compare online and in-person workshop results.`
3. **Web Search:** `Find three suitable venues near Brighton railway station.`
4. **Image Generation:** `Create a promotional image for the workshop.`

This sequence makes the value of a specialized GPT immediately visible. The same GPT recalls organizational rules, analyzes a permanent dataset, researches changing public information, and produces new material without you having to repeat your operating instructions each time.

Click **Create** when you're finished previewing your GPT.

Share and maintain the GPT

Decide who can use the GPT: only you, anyone with a link, or everyone in your organization:

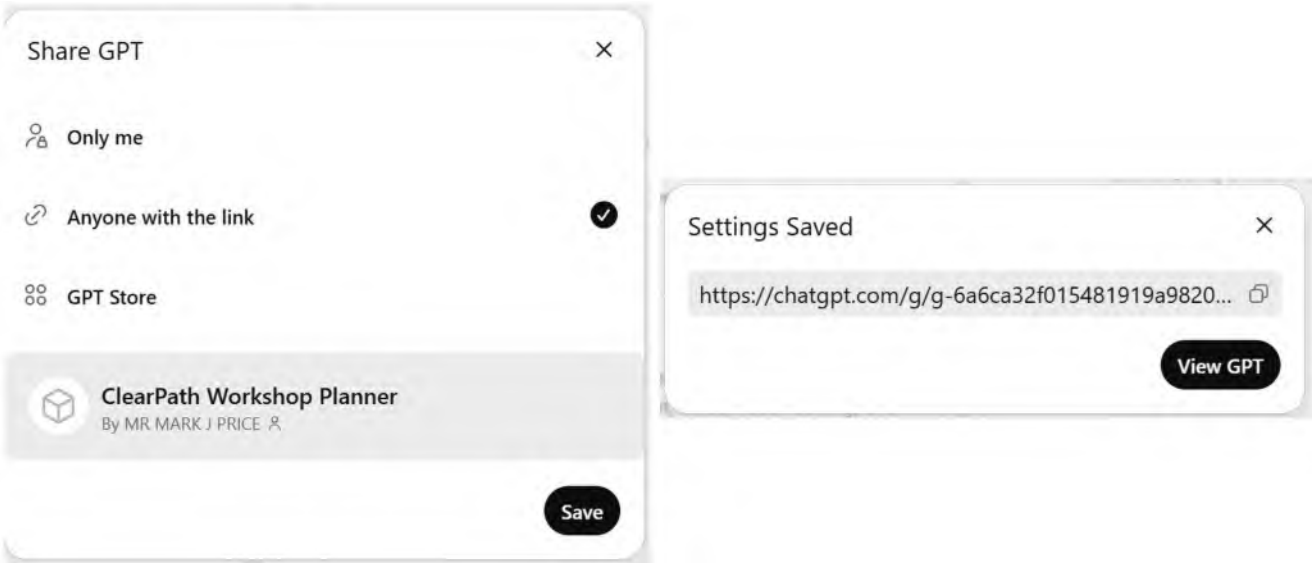


Figure 3.4D: Share GPT options

If you select **Anyone with the link**, a second dialog allows you to copy a link to the clipboard to share with anyone. Click **View GPT** to see what they would see and try out the GPT:

1. Select the GPT name to pop up a menu to manage the GPT.
2. In the chat box, select **Thinking** to change the effort level.
3. Select **...** and then **Pin GPT** to add the custom GPT to your ChatGPT account's **GPTs** page.
4. Note the name and description are shown, and the conversation starters are clickable to jump-start a conversation for someone new to this custom GPT.



You can try out this GPT here: `https://chatgpt.com/g/g-6a6ca32f015481919a98203ad8220d91-clearpath-workshop-planner`.

In an organization-managed workspace, a workspace agent extends this idea further. It is built and maintained by the organization rather than one person, follows rules that the organization sets, and is not a personal GPT with a new name.

Review the GPT occasionally. Update its instructions or knowledge files when the underlying task, policy, or product changes.



Watch out: A GPT's knowledge files are only as current as the last time you updated them. If the source material changes, update the GPT too, or its answers will go out of date without warning.

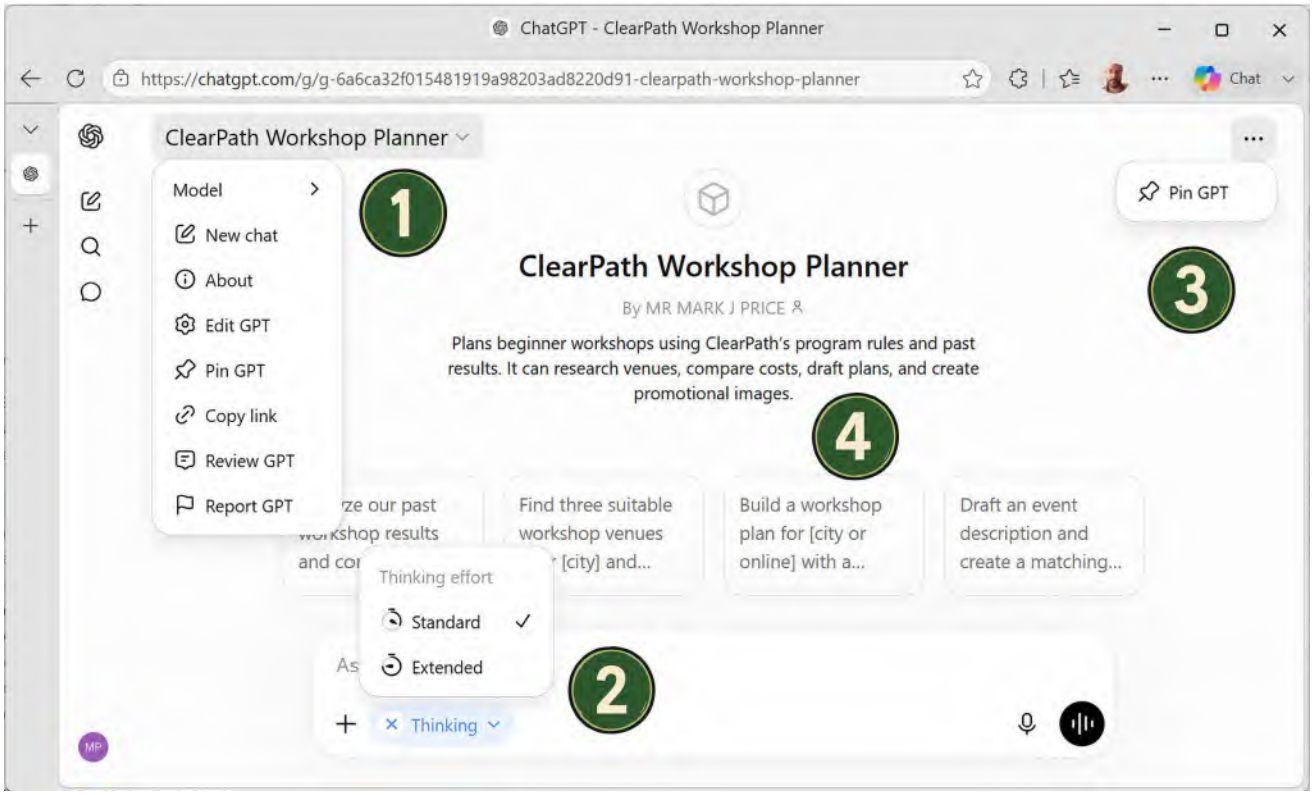


Figure 3.4E: A custom GPT now available to the public with the link you can share

Try it now

Sketch a short brief for one GPT you could build: its purpose, three instructions, one knowledge file it would need, and two conversation starters. Build it if you have time or save the brief for later.

To more fully test the GPT as a user, an optional online-only section with sample prompts is available: `https://github.com/markjprice/chatgpt-vb/blob/main/docs/ch3-4-test-gpt.md`. Together, the prompts test its uploaded knowledge, data analysis, web search, image generation, missing-information handling, and approval boundaries.

Check the result

- ☐ Does the GPT have a single, clear purpose, rather than several unrelated ones?
- ☐ Did you test it with a question outside its usual scope?
- ☐ Do you know who should be able to use it, and have you set that correctly?

A specialized GPT gives ChatGPT a clear purpose. *Book 2*, *Chapter 2.8* covered plugins, a friendly way to connect popular services. *Chapter 3.5* scales back down to something more personal: using automation for your own organization and learning.

Chapter 3.5 Automate Personal Organization and Learning

Book 3 has covered the building blocks of automation: scheduled tasks, skills, and GPTs, with optional online chapters to cover more advanced topics like MCP and workspace agents. This chapter applies them outside work, to the everyday organizing and learning tasks from *Book 1*. You do not need every tool from this book to benefit here. Even one well-built skill or scheduled task can make a real difference in your week.

Build a recurring personal-organization workflow

Many personal tasks repeat on a schedule: a weekly review of your calendar, a monthly budget check, or a seasonal home-maintenance list. These are strong candidates for the scheduled tasks you met in *Chapter 3.2*, combined with a project that holds your ongoing context.

Turn that into four repeatable steps:

1. Choose one recurring personal task, such as planning your week.
2. Create a project and add the information ChatGPT needs each time, such as your priorities or recurring commitments.
3. Write a short skill or standard prompt that produces the result you want, such as a prioritized task list.
4. Schedule it to run at a time that fits your routine.

Turn a learning goal into a repeatable routine

Chapter 1.6 showed you how to learn a subject with ChatGPT in a single conversation. You can turn that same approach into a routine you repeat over weeks, rather than a one-time session:

- Keep one project for the subject, with your current level and goals stated once
- Use a simple skill that asks ChatGPT to explain one new idea, give an example, and quiz you on the last session's material
- Review your progress every week or two, and adjust the pace based on what is actually sticking

Real-world example: A retired teacher learning to use a spreadsheet program keeps a single project called *Spreadsheet basics*. Each session, she asks for one new skill, a short practice task, and two questions that check whether she remembers the previous lesson. After six weeks, she can build a simple household budget without help.

Prepare for a personal decision or milestone

Some personal automation is not about repeating a task, but about supporting a single important moment, such as a big purchase, a move, or a conversation you are nervous about. *Chapter 1.7* covered comparing options for everyday decisions. You can extend that into a short, repeatable pattern:

1. State your situation, goals, and constraints once in a project.
2. Ask ChatGPT to list the questions you still need answered.
3. Gather information over several sessions, adding it to the same project.
4. When ready, ask for a clear summary of your options and open questions, not a final decision made for you.

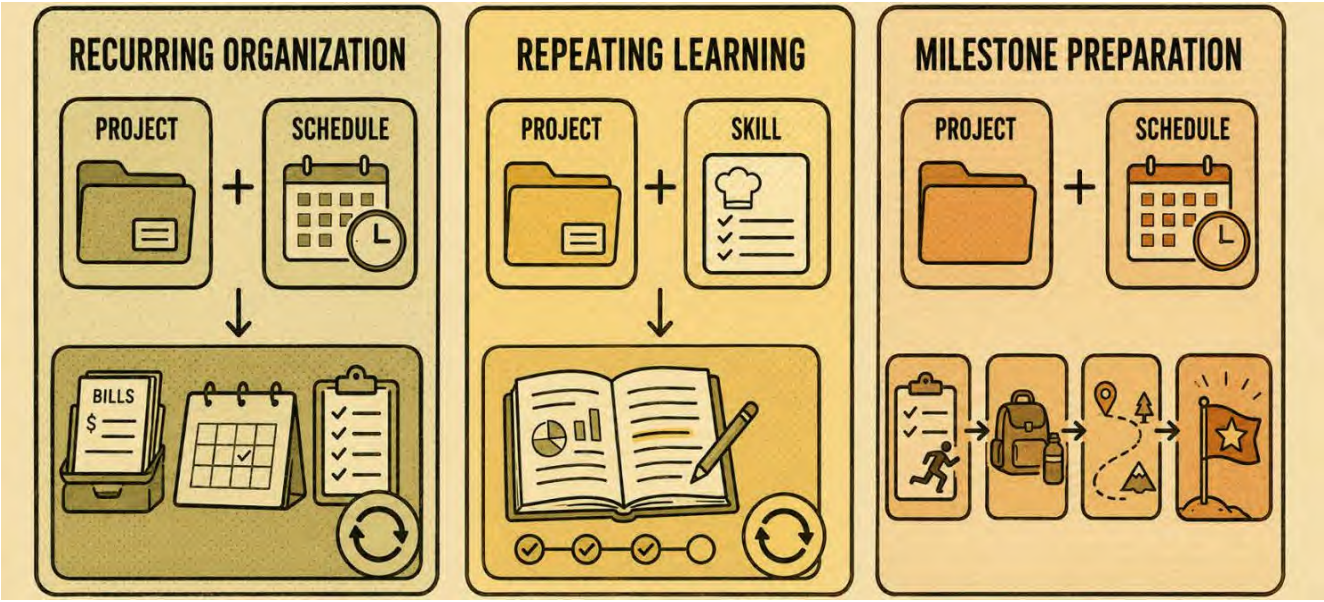


Figure 3.5A: The same combination of a project and a small skill or schedule can support ongoing organization, steady learning, or preparation for one big decision

Watch out: Keep personal and sensitive information, such as financial details or health notes, in a project you control, and review your privacy and memory settings from *Chapter 1.11* before adding anything you would not want stored longer than necessary.

Try it now

Pick one recurring personal task or one learning goal. Set up a small project for it, with one skill or standard prompt you can reuse. Run it once this week.

Check the result

- ☐ Did the routine save you real time or effort compared with starting from scratch each time?
- ☐ Is the information stored in the project something you are comfortable keeping there?
- ☐ Would you keep using this routine next month, or does it need a change first?

With a personal routine automated, you're ready to apply the same thinking to work you do for other people. *Chapter 3.6* walks through building a workflow around producing content, from research through a finished draft.

Chapter 3.6 Build a Research and Content Workflow

This chapter builds a complete workflow around one common need: turning research into finished written content on a repeating basis. It combines the skills, scheduling, and connections from earlier chapters in *Book 3* into one working pattern you can adapt to your own content.

Produce a recurring management review

A management review is a short, regular report that turns raw information into a decision-ready summary. Building it as a workflow, rather than a fresh request each time, keeps the format consistent and saves you the setup work every week.

Turn that into four repeatable steps:

1. Define what the review must cover, such as key metrics, notable changes, and open risks.
2. Connect the sources it needs, using the plugins from *Chapter 2.8* or the MCP connections from the online *Chapter 3.9*.
3. Build a skill that gathers the information and drafts the review in your required format.
4. Schedule the workflow to run before your regular reporting deadline, as you learned in *Chapter 3.2*.

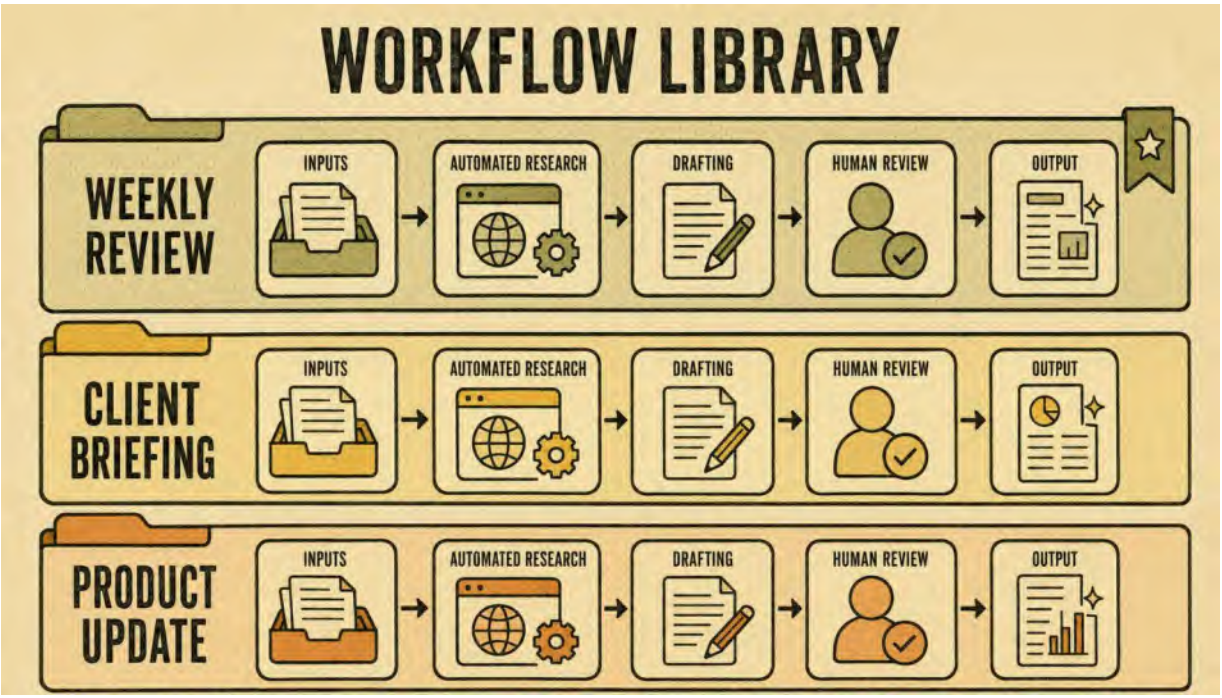


Figure 3.6A: A workflow library holds several small, repeatable patterns, each moving from inputs through automated stages to a reviewed output

Gather and compare evidence before drafting

Strong content depends on strong evidence gathered first. Before drafting, have ChatGPT collect relevant facts, figures, and sources, then compare them for agreement or conflict, the way you learned in *Chapter 1.9*'s deeper research process:

- Ask for at least two independent sources for any important claim
- Flag any number or claim that sources disagree about, rather than picking one silently
- Keep a short list of sources alongside the draft, so a reader or colleague can check them later



Watch out: A workflow that drafts content automatically can also repeat an error automatically, every time it runs. Check the evidence stage carefully the first several times you run a new workflow, before you start trusting it without review.

Develop options and recommendations

Some content workflows should end in a recommendation, not only a summary. When that is the goal, ask ChatGPT to present two or three real options with their trade-offs, rather than jumping straight to one preferred answer:

1. State the decision the content should support.
2. Ask for the leading options, each with its evidence, benefits, and drawbacks.
3. Ask for a recommendation, along with the reasoning behind it.
4. Review the recommendation against the evidence before you rely on it.



Good practice: Keep one part of every content workflow manual: a final human read-through before the content goes to its audience. Automation should remove repetitive drafting work, not remove the last check.

Try it now

Pick one recurring report or piece of content you produce regularly. Sketch the sources it needs, the evidence-gathering step, and whether it should end in a summary or a recommendation. Build the first version as a skill or a small scheduled workflow.

Check the result

- ☐ Does your workflow gather evidence from more than one source before drafting?
- ☐ Did you build in a step to flag disagreement between sources, rather than hiding it?
- ☐ Is there still a human read-through before the content reaches its audience?

This chapter built a workflow around producing content. *Chapter 3.7* builds a different kind of workflow, one that watches ongoing situations, flags what needs attention, and prepares you for a specific decision or action.

Chapter 3.7 Build a Decision and Operations Workflow

This chapter builds a workflow that watches ongoing situations, flags what needs attention, and prepares you for a specific decision or action. The pattern is the same set of building blocks from earlier in *Book 3*, applied to operations instead of content.

Monitor decisions, risks, and follow-up actions

Many roles involve tracking things that change: open decisions waiting on someone, risks that need a response, and actions that were promised but not yet done. A monitoring workflow keeps this list current without you having to rebuild it from memory each week:

1. Define what counts as an open decision, a risk, or a follow-up action for your work.
2. Connect the sources where these items appear, such as notes, messages, or a shared tracker, using a plugin or MCP connection.
3. Build a skill that reviews these sources and produces a short, current list.
4. Schedule the workflow to run before the meeting or checkpoint where you need the list.

Prepare for prospect and account conversations

Sales and business-development work often depends on a short brief prepared before a conversation. This is a strong candidate for the same pattern: gather relevant information about a prospect or account, then produce a consistent brief every time.

Turn that into three repeatable steps:

1. Gather public information about the prospect or account, along with any internal notes you have permission to use.
2. Summarize what's relevant for the upcoming conversation: their situation, likely needs, and open questions.
3. Flag any information that could not be confirmed, rather than presenting a guess as a fact.



Watch out: Never include personal or confidential information about a prospect that you do not have a legitimate reason to hold or share. Review *Chapter 2.1*'s rules on data classification before connecting any customer information source.

Prepare for screening and interview conversations

The same brief-before-conversation pattern applies to hiring. Before a screening call or interview, a short workflow can compare a role's requirements against a candidate's stated experience and prepare focused questions:

1. State the role's key requirements clearly.
2. Provide the candidate's resume or application material.
3. Ask ChatGPT to identify where the match looks strong, where it looks uncertain, and what to ask about directly.
4. Review the questions before the conversation, and remove anything that could introduce unfair bias.

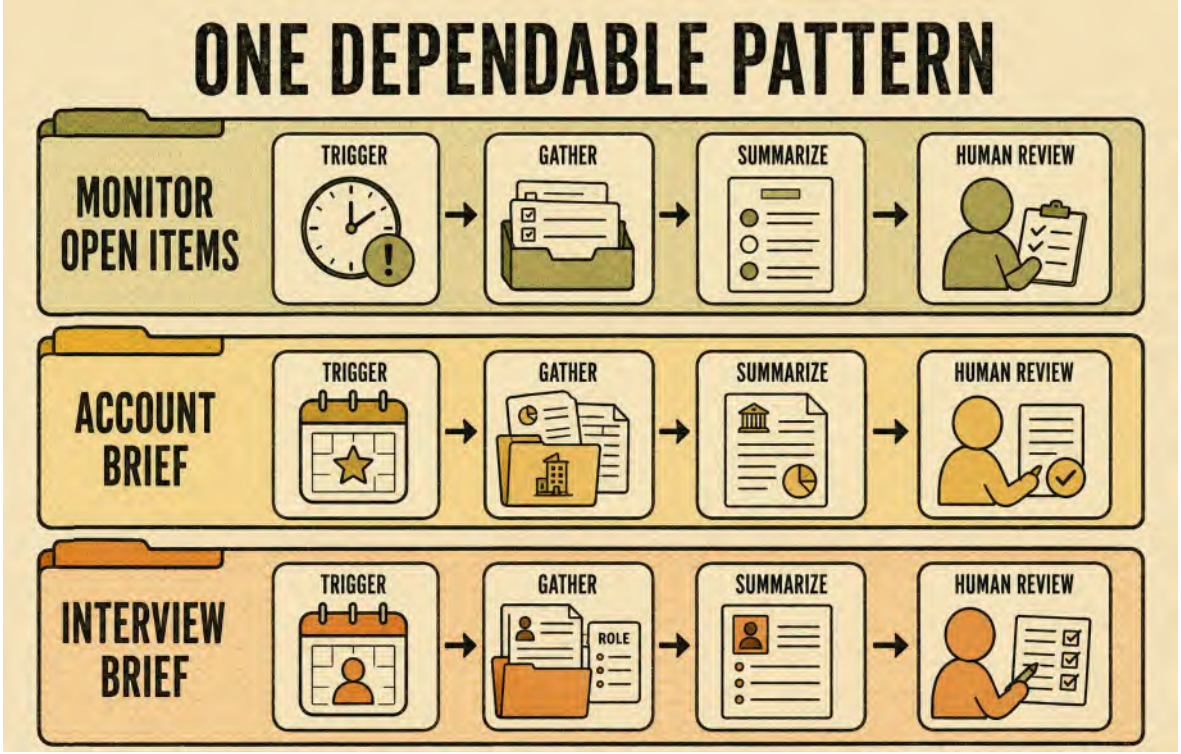


Figure 3.7A: Monitoring, account preparation, and interview preparation all follow the same pattern: a trigger, a gathering step, a summary, and human review before the conversation

Good practice: Keep a fixed set of questions or criteria for every screening brief, rather than letting the workflow generate new ones each time. Consistency protects fairness as much as it protects your time.

Try it now

Choose one operational task you repeat before a meeting or a conversation, such as a status check, an account brief, or a screening summary. Build the first version as a skill that gathers the relevant information and produces a consistent result.

Check the result

- ☐ Does the workflow's list or brief match what you would have prepared by hand?
- ☐ Did you check for any confidential or unfair content before using the result in a real conversation?
- ☐ Would this workflow still work if the source information changed shape next month?

You've now built automations for content, decisions, and daily operations. *Chapter 3.8* asks you to choose one of these workflows, or a new one, and prove it holds up: measured against a baseline, tested with a case designed to fail, and ready to run without you standing over it.

Chapter 3.8 Project — Build and Measure an End-to-End Workflow

This project brings together everything from *Book 3*: choosing the right building block, controlling tools and permissions, building a skill, extending ChatGPT with plugins or MCP, and delegating separable work. You will build one complete workflow, then measure whether it actually helps.

Choose a repeated workplace process

Pick one process from your own work that repeats often enough to matter, and that you understand well enough to describe clearly. A good candidate has a defined start, a defined end, and a result you can judge as good or not good enough.

Avoid choosing your hardest or most sensitive process for a first attempt. Choose one where a mistake is easy to catch and easy to fix.

Combine projects, skills, and plugins

Use the tools from this book together, choosing only what the process actually needs:

1. Set up a project to hold the process's ongoing context, as in *Chapter 3.1*.
2. Build one or more skills for the stable, repeated steps, as in *Chapter 3.3*.
3. Add a plugin or MCP connection only if the process actually needs outside information or action.



Figure 3.8A: A complete workflow moves from trigger through intake, specialist work, quality review, and approval, ending in delivery and a record of what happened.

Add review and approval checkpoints

Every workflow in this project needs at least one point where a person checks the result before it goes anywhere consequential. Decide where that checkpoint belongs, and what the reviewer should specifically look for.

State clearly when the workflow should stop and ask a person, rather than continuing on its own. A workflow that cannot recognize its own limits is not ready to run without close supervision.

Test, document, and reuse the workflow

Follow this sequence to turn your workflow from a working prototype into something you can trust and reuse:

1. Record the existing time and quality baseline, before your workflow existed.
2. Define what success looks like in specific, checkable terms.
3. Test the workflow on a normal case, an unusual case, and a case designed to fail.
4. Confirm exactly when the workflow stops and asks a person for input.
5. Record what happened during each test, including any surprises.
6. Compare the result against your baseline from step one.
7. Provide a manual fallback in case the workflow is unavailable or fails.



Good practice: Keep this record. *Chapter 4.3* shows you how to use evidence like this to document your improvement and build a case for the value of your work.

Choose your version of the project

Pick the version that fits your situation best.

- A personal or family workflow, such as an organizing or planning routine you built in *Chapter 3.5*
- A content or research workflow, such as the recurring review you built in *Chapter 3.6*
- A decision or operations workflow, such as the monitoring or briefing routine you built in *Chapter 3.7*

Try it now

Complete your chosen workflow this week. Run it through every step above, from choosing the process to recording your test results and comparing them against your baseline.

Check the result

- ☐ Does your workflow perform measurably better, the same, or worse than your baseline?
- ☐ Did every test, including the one designed to fail, behave the way you expected?
- ☐ Do you have a manual fallback ready if the workflow cannot run?

The next step

A workflow that runs on its own, passes its tests, and beats your baseline is a genuine capability, not a demo. *Book 4*, *Advance Your Career and Income*, shows you what to do with that capability. You'll learn to measure the value of what you've automated, build a case for a promotion or a new role, and decide whether the skills and workflows you've built are worth turning into a service or product of your own.

Book 4

Advance Your Career and Income

In *Book 4*, you use what you learned in *Book 1*, *Book 2*, and *Book 3* to improve your current work, earn a promotion, learn new skills to make a career change, or provide a product or service to generate additional income. It has thirteen chapters, including an end-of-book project:

4.1 Choose What Better Work Means to You 154

4.2 Audit Your Role and Find High-Value Work 156

4.3 Measure and Document Your Improvement 157

4.4 Learn the Skills Your Next Role Requires 160

4.5 Combine AI Fluency with Human Strengths 162

4.6 Propose and Lead an AI Improvement at Work 164

4.7 Build a Case for Promotion or Higher Pay 166

4.8 Prepare for a Better Job or Career Change 168

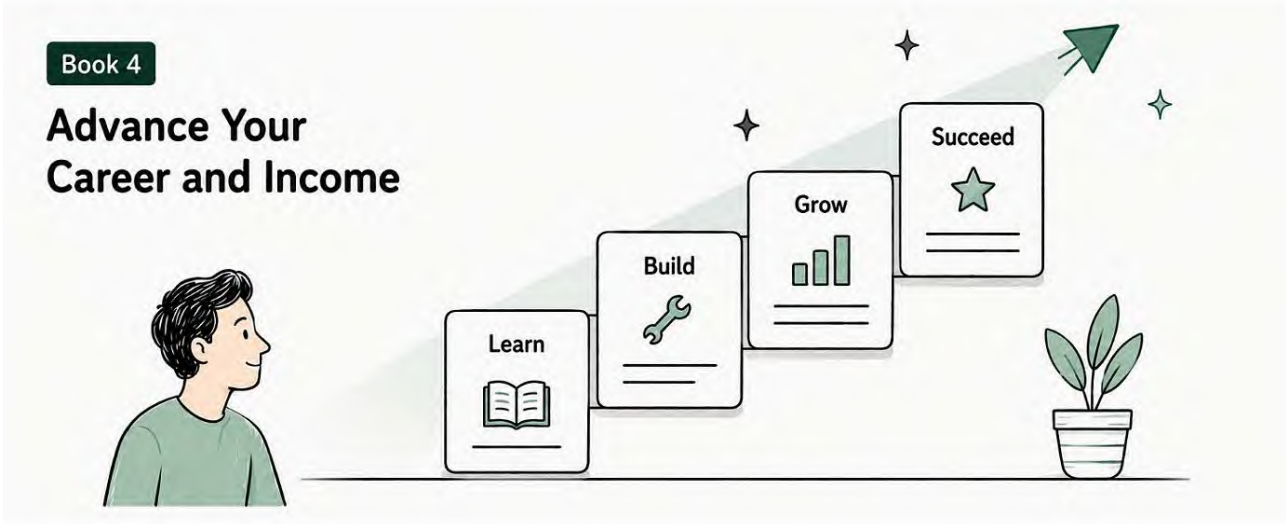
4.9 Rehearse High-Stakes Conversations with Voice Mode 171

4.10 Build a Portfolio and Measurable Case Study 174

4.11 Turn a Problem into a Service or Product 177

4.12 Decide How to Respond to AI in Your Field 179

4.13 Project – Complete a 30-day Career Value Project 181



Chapter 4.1 Choose What Better Work Means to You

Book 3 helped you build workflows and measure whether they worked. *Book 4* asks a different question: what should you do with the ability you have built? This chapter helps you choose a direction before you invest more time. There is no single right answer. Two readers with the same skills might choose completely different paths, and both choices can be good ones.

Reducing routine work in your current role

Some readers mainly want their week to feel less repetitive. If routine tasks take up hours you would rather spend elsewhere, the goal is to hand over more of that routine work to ChatGPT, using the automation skills and processes from *Book 3*. This is often the fastest path to a better week, since you already know which tasks feel repetitive and draining.

Working toward a promotion

If your goal is a better title or a higher salary in your current organization, later chapters in *Book 4* show you how to document your improvement and build a case that colleagues and managers will find convincing.

Developing expertise for a new role or career change

Other readers have enough time but want their work to matter more. Here, the goal is not doing less, but redirecting freed-up time toward work that produces a bigger result, such as a project that reaches more people or solves a harder problem.

To go deeper in your field rather than wider, you can develop your existing skills. ChatGPT can help you learn faster, as you saw in *Chapter 1.6*, so the time you save on routine tasks can fund real study in your specialty. Some readers want to move into a different role or a different field entirely. *Chapter 4.4* and *Chapter 4.8* focus directly on this path, from learning new skills to preparing your job search.

Selling a product or service to earn independently

A smaller group of readers want to turn their abilities into a service or product they sell directly, whether alongside their current job or instead of it. *Chapter 4.10* and *Chapter 4.11* cover this path in detail, without promising quick or guaranteed income.



Watch out: Earning independently takes real work: finding customers, delivering consistently, and handling problems when they come up. Treat any promise of quick or easy income with suspicion, including promises made by ChatGPT itself.

Deciding where saved time should go

Whatever path you choose, the time and mental energy you free up through automation will not automatically go somewhere useful. You have to decide.

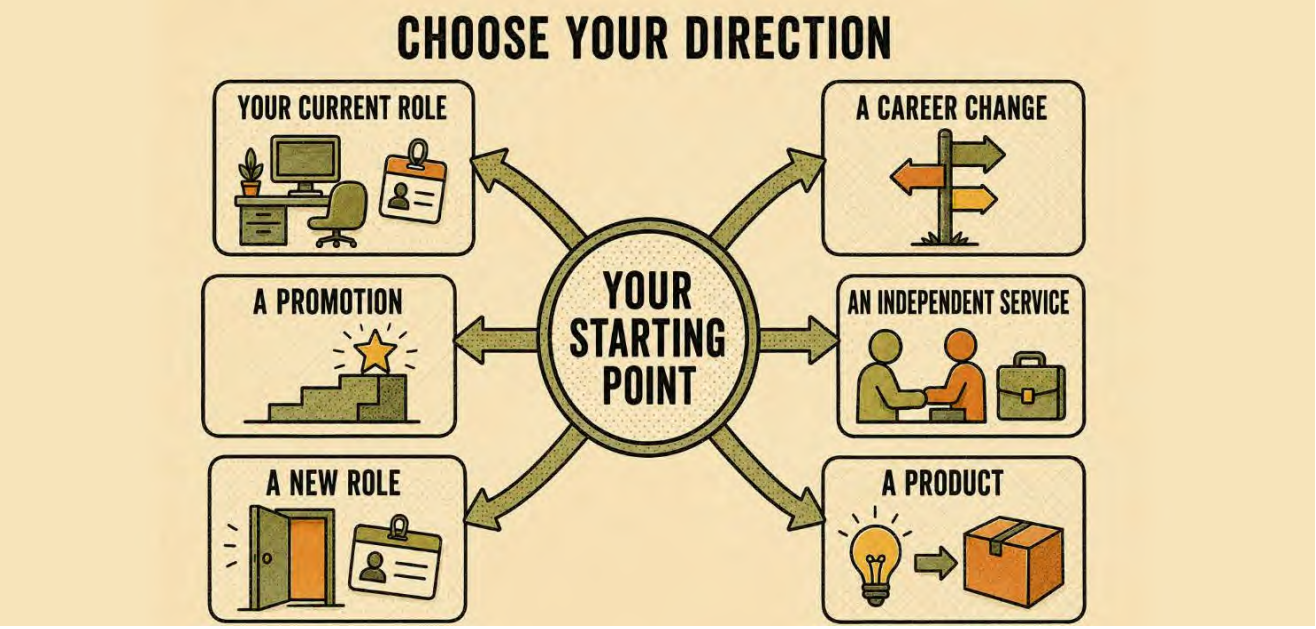


Figure 4.1A: Better work can mean many different things. This map shows common directions readers take once they have freed up time and built new skills.

Before moving on, write down which direction, or combination of directions, you care about most right now. You can change your mind later. *Book 4* works best when you have a working answer to start from.

Good practice: Choose one primary direction, even if you are curious about several. The chapters that follow build on a specific choice, and a scattered effort across many directions tends to produce less progress than a focused one.

What to ask next: `Turn my chosen direction into one sentence I could say if a colleague asked what I'm working toward this quarter.`

Try it now

Read the mapped directions in this chapter again. Rank your top two. Write one sentence about why you chose each one, using your own words rather than the wording in this book.

Check the result

- ☐ Did you choose a direction based on what you actually want, rather than what sounds most impressive?
- ☐ Can you explain your choice to another person in one or two sentences?
- ☐ Are you comfortable spending the next several chapters working toward this direction?

This chapter helped you choose a direction. *Chapter 4.2* turns that direction into a specific starting point: one piece of your current work worth improving first.

Chapter 4.2 Audit Your Role and Find High-Value Work

This chapter focuses on one piece of your current work worth improving.

List recurring activities

Start with a plain list of what you actually do. Over two or three days, write down each task as you do it, including small ones you might normally overlook, such as status updates or routine replies.

You do not need a complete list on the first try. A rough list you can refine is more useful than a perfect list you never finish.

Identify friction and wasted effort

Look at your list and mark the tasks that feel slow, repetitive, or frustrating compared with the value they produce. Common signs of friction include retyping the same kind of message, searching for the same information more than once, and waiting on a step that could run without you.



Real-world example: A project coordinator notices she rewrites the same weekly status update from scratch each Friday, pulling numbers from three different places. The task itself takes thirty minutes, but finding the numbers takes closer to an hour. That gap between the task and the effort around it is a clear sign of friction.

Classify work as retain, assist, automate, delegate, or stop

Sort each recurring activity into one of five categories:

1. **Retain:** work that depends on your judgment, relationships, or authority, and should stay with you as it is
2. **Assist:** work you will keep doing, but where ChatGPT can help you go faster or check your work
3. **Automate:** work stable and well understood enough to become a skill or workflow, as in *Book 3*
4. **Delegate:** work better handled by a colleague, a specialist subagent, or another person entirely
5. **Stop:** work that no longer serves a real purpose and could end

Recognize high-risk work

Some tasks should not move to automate or delegate even if they repeat often, because a mistake would be costly or hard to reverse. This includes work involving confidential information, financial commitments, legal obligations, or decisions that affect another person's job or safety.



Watch out: A task that looks routine on the surface can still carry real risk. Before automating anything, ask what happens if the result is wrong and nobody catches it. If the answer is serious, keep a human check in place, even after you build a workflow for the rest of the task.

Choose one improvement

Resist the urge to fix everything on your list at once. Pick the single task that combines the clearest friction with the lowest risk, and commit to improving that one first:

1. Review your classified list.

2. Cross out anything marked retain or stop, since those need a decision rather than a build.
3. From the remaining tasks, choose the one with the most friction and the least risk.
4. Write one sentence describing what "better" would look like for this task.

You should now have a single, specific task to bring into the next chapter, along with a clear picture of what improvement means for it.



Good practice: Write your classification down somewhere you can revisit, not only in your head. You will use this list again in later chapters as you look for the next task to improve.



What to ask next: `Suggest three ways I could reduce the friction in this task, ordered from easiest to hardest to try first.`

Try it now

Spend fifteen minutes listing your recurring activities. Classify at least five of them using the five categories in this chapter, then choose the one task you will improve first.

Check the result

- ☐ Does your chosen task show clear friction, real repetition, and manageable risk?
- ☐ Did you classify at least one task as retain, so you can see the categories are not only a path toward removing yourself from your own work?
- ☐ Can you describe what a better version of this task would look like, in one sentence?

Finding a high-value task to improve is only useful if you can prove the improvement happened. *Chapter 4.3* shows you how to measure your starting point and document real progress over time.

Chapter 4.3 Measure and Document Your Improvement

In *Book 3*, the end-of-book project asked you to record a baseline before building your workflow. This chapter applies that same method to the task you chose in *Chapter 4.2*, whether that's a fresh improvement or the workflow you already built.

You need evidence, because a claim of improvement without evidence is just an opinion, and opinions are easy to doubt. A specific measurement is much harder to argue with.

Establish a baseline

A baseline is a clear record of how a task worked before your improvement. Without one, you cannot honestly say whether anything actually got better.

1. Choose the task you identified in *Chapter 4.2*. If it's the same task you built a workflow for in *Book 3*, use the baseline you already recorded instead of measuring again.
2. Note how long it currently takes, from start to finish.
3. Note its current quality, using a specific measure such as an error rate or a completion rate.
4. Record this baseline before you change anything.

Measure time, quality, errors, and outcomes

Once you have a baseline, measure the same task after your improvement, using the same measures each time so the comparison is fair.

- **Time:** how long the task takes from start to finish
- **Quality:** how often the result meets your stated standard
- **Errors:** how often something needs to be corrected after the fact
- **Outcomes:** the real-world result the task was meant to produce, such as a client's satisfaction or a report's usefulness to its reader

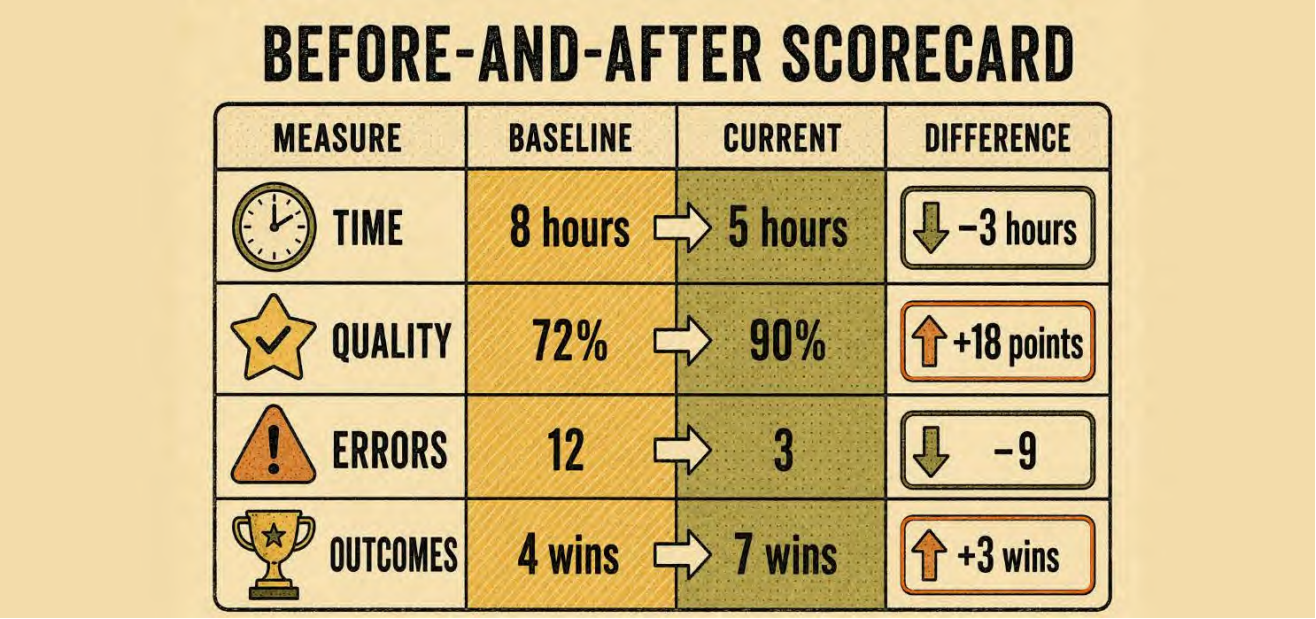


Figure 4.3A: A before-and-after scorecard puts your baseline and your current result side by side across time, quality, errors, and outcomes

Avoid invented precision

Report your measurements honestly, even when they are rough. If you only have a general sense that a task now takes about half the time, say that, rather than presenting a precise number you did not actually measure.

Watch out: A specific-sounding number built on a guess is more misleading than a stated estimate. If someone asks how you got a number, you should be able to explain it. If you cannot, round it or describe it as an estimate instead.

Keep an achievement record

Do not rely on memory to carry your evidence forward. Keep a running record of every task improvement you make, with its baseline, its result, and the date:

1. Create one document or spreadsheet for your achievement record.
2. Add a new entry each time you complete an improvement, including the ones that did not work as well as you hoped.
3. Include enough detail that you could explain the entry to someone else a year from now.

4. Review the record occasionally and update it as circumstances change:

Achievement Record

A running log of measured improvements: what you improved, where you started, what you achieved, and when.

Task improvement	Baseline	Result	Date
Weekly status report	90 min total: 30 min writing, ~60 min gathering numbers from three sources	35 min total, using a saved ChatGPT prompt to draft from pasted data	14 Mar 2026
Client onboarding emails	~45 min per new client, rewritten from scratch each time	15 min per client, using a reusable template ChatGPT helped refine	02 Apr 2026
Meeting notes and action items	Notes often incomplete; about 1 in 4 action items missed follow-up	All action items captured; formatting still needs manual cleanup	20 May 2026
Expense report summaries	2 hours monthly, 3 formatting errors on average	50 min monthly, 1 formatting error on average	30 Jun 2026

Figure 4.3B: An achievement record collects your improvements over time, so you always have evidence ready when you need it



Sample files: A one-page Word achievement record: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book4/4-3-Achievement-Record.docx`. The table has four rows drawn from the book's running example, including one entry that only partially succeeded.



Good practice: Include improvements that did not fully succeed in your record, along with what you learned. A record of only perfect results looks less credible than one that also shows honest attempts.



What to ask next: `Turn this baseline and result into three sentences I could say out loud in a performance review.`

Decide how to use the capacity created

Once you have measured a real improvement, decide what to do with the time or effort it freed up. Return to *Chapter 4.1* and check that the capacity is going toward the direction you chose there, rather than disappearing into other routine work.

Try it now

Measure your chosen task's baseline and its current result, using the four measures in this chapter. Add the comparison to a new achievement record, even if the record only has one entry so far.

Check the result

- ☐ Can you explain exactly how you measured each number in your comparison?
- ☐ Did you record the result honestly, including any part of the improvement that fell short?
- ☐ Do you know where the time or effort you saved is now going?

Documenting your current improvement builds a record for today's role. *Chapter 4.4* looks further ahead, helping you identify and learn the skills your next role will actually require.

Chapter 4.4 Learn the Skills Your Next Role Requires

If *Chapter 4.1* pointed you toward developing expertise or changing career, this chapter helps you find out exactly what to learn, then build a plan you will actually follow.

Research how a role is changing

Roles change as tools, expectations, and industries change. Before setting a learning plan, research what your target role looks like today, not what it looked like when you first learned about it:

1. Ask ChatGPT to search for current descriptions of your target role, including recent job postings.
2. Ask it to identify skills that appear often across several sources.
3. Note any skill that seems new compared with older descriptions of the same role.
4. Check a few of the sources yourself, using the research habits from *Chapter 1.8*.



Learn more online: O*NET OnLine, sponsored by the US Department of Labor, lists the skills, tasks, and tools linked to hundreds of occupations. Use it to cross-check what ChatGPT tells you about a role: `https://www.onetonline.org/`. Readers outside the United States can ask ChatGPT to find their country's equivalent occupational database.

Compare current and required skills

List the skills your target role needs, then honestly rate your current level in each one. Be specific rather than general. "Comfortable with data" is harder to act on than "can build a pivot table but not a formula-driven dashboard."

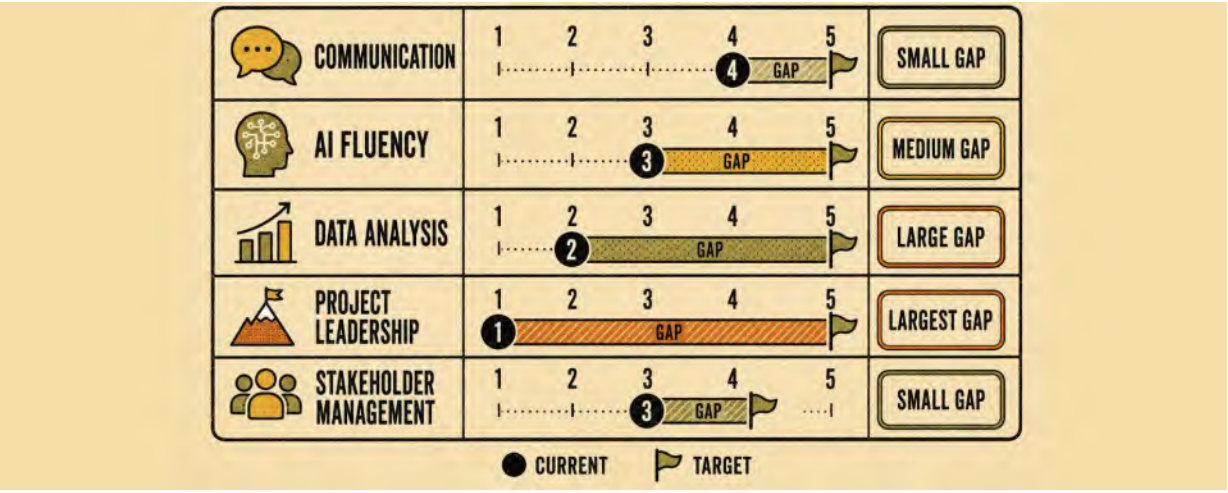


Figure 4.4A: A current-to-target skills map shows exactly how far you are from the skills your next role needs, one skill at a time

Create a learning plan

Turn your biggest gaps into a plan with real steps and a real timeline. *Chapter 1.6* showed you how to learn a single subject with ChatGPT. Extend that approach across several linked topics, in an order that builds on itself:

1. Choose three skills to focus on first, based on the gap and how much the role needs them.
2. Ask ChatGPT to break each skill into a short learning path, similar to *Chapter 1.6*.
3. Set a realistic timeline, with checkpoints rather than a single distant deadline.
4. Keep the plan in a project, as in *Chapter 1.11*, so your progress and materials stay in one place.



Watch out: An overloaded learning plan is easy to write and hard to finish. Choose fewer skills and a slower pace rather than a long list you are unlikely to complete.

Complete a demonstrable project

Reading about a skill is not the same as being able to use it. Choose a small project that forces you to apply what you are learning to something real, ideally something close to the actual work of your target role.



Real-world example: Someone learning data analysis to move into an operations role does not stop at tutorials. She analyzes her current team's own data to answer a real question, such as which process step causes the most delay, and produces a short report she can show later.

Test genuine understanding rather than dependence on AI

It is possible to produce good-looking work with heavy help from ChatGPT while understanding very little of it yourself. Guard against this by testing your own understanding directly:

- Explain a concept back in your own words without looking at your notes
- Complete a small task without asking ChatGPT for help, then check your work against what you learned
- Notice when you can follow an explanation but could not have produced it yourself and treat that as a sign to study the topic further



Good practice: Schedule a short, honest self-test at the end of each stage in your learning plan, before moving to the next stage. Catching a weak spot early is far easier than discovering it during your demonstrable project.

Try it now

Research your target role and list its current key skills. Rate your level in each one, choose your top two gaps, and sketch a short learning path for the first gap.

Check the result

- ☐ Did your research include current sources, not only your existing assumptions about the role?
- ☐ Is your learning plan realistic given your available time?
- ☐ Have you planned a way to test your own understanding, separate from what ChatGPT produces for you?

Learning new skills is only part of becoming more valuable at work. *Chapter 4.5* looks at the other half: the human strengths that AI fluency should sharpen, not replace.

Chapter 4.5 Combine AI Fluency with Human Strengths

Skill with ChatGPT is valuable, but it is not the whole of what makes you valuable at work. This chapter looks at the strengths that stay firmly yours, and how to keep them sharp while you automate more of your routine tasks.

Develop human strengths

There are seven areas that depend on human experience, relationships, and context in ways ChatGPT cannot fully replace:

1. **Judgment** means weighing a situation's specifics and deciding what to prioritize, especially when the right answer is not obvious. For example, deciding whether a minor product bug is worth delaying a launch over is a call no model can make for you.
2. **Domain knowledge** is the deep, specific understanding of your field that lets you spot when something looks wrong. For example, a mechanic who can tell from an unusual engine sound that a diagnostic report missed something.
3. **Communication** is explaining an idea clearly to the particular person or group in front of you. For example, explaining the same budget shortfall differently to a worried teammate and to a skeptical executive.
4. **Creativity** is combining ideas in a way that fits your actual situation, not a generic best practice. For example, redesigning a standard onboarding process because your team works across three time zones and the usual schedule doesn't fit.
5. **Empathy** is understanding how a decision will actually land with the people affected by it. For example, sensing that a teammate's quiet pushback in a meeting is discomfort with a decision, not disagreement with the plan.
6. **Leadership** is helping other people do their best work, especially when the path forward is unclear. For example, noticing a capable teammate has gone quiet under pressure and adjusting how you delegate to get them unstuck.
7. **Stakeholder management** is understanding what different people need from a project and balancing those needs. For example, balancing a client's request for speed against your engineering team's need for more testing time.

Real-world example: Two analysts both use ChatGPT to build the same kind of financial model. One stops there. The other uses her domain knowledge to notice that the model's assumption about seasonal demand does not match what she has seen in three years of actual sales data. Her judgment, not the model itself, is what protects the business from a bad decision.

Recognize skill atrophy

Skill atrophy happens when a skill you rarely practice grows weaker without your noticing, even though you might not notice until you need it. Heavy reliance on ChatGPT for tasks you once did yourself can accelerate this, especially for skills tied to writing, calculation, and structured thinking.

Watch for signs of atrophy: hesitating over a task you used to do easily, or noticing that you can no longer catch an obvious error in ChatGPT's work without help. When you notice this, set aside time to practice the skill without assistance, the way you tested your own understanding in *Chapter 4.4*.

Watch out: Do not let convenience replace a skill you still need without your noticing. A skill you never practice becomes a skill you can no longer trust yourself with, even in situations where ChatGPT is unavailable or wrong.

Retain responsibility for decisions

ChatGPT can draft, analyze, and suggest, but it does not carry consequences the way you do. When a decision affects your work, your team, or another person, the responsibility for that decision stays with you, regardless of how much of the drafting or analysis ChatGPT provided.

This means reviewing important output carefully, questioning a result that seems off even if you cannot immediately explain why, and being ready to explain and defend a decision in your own words.

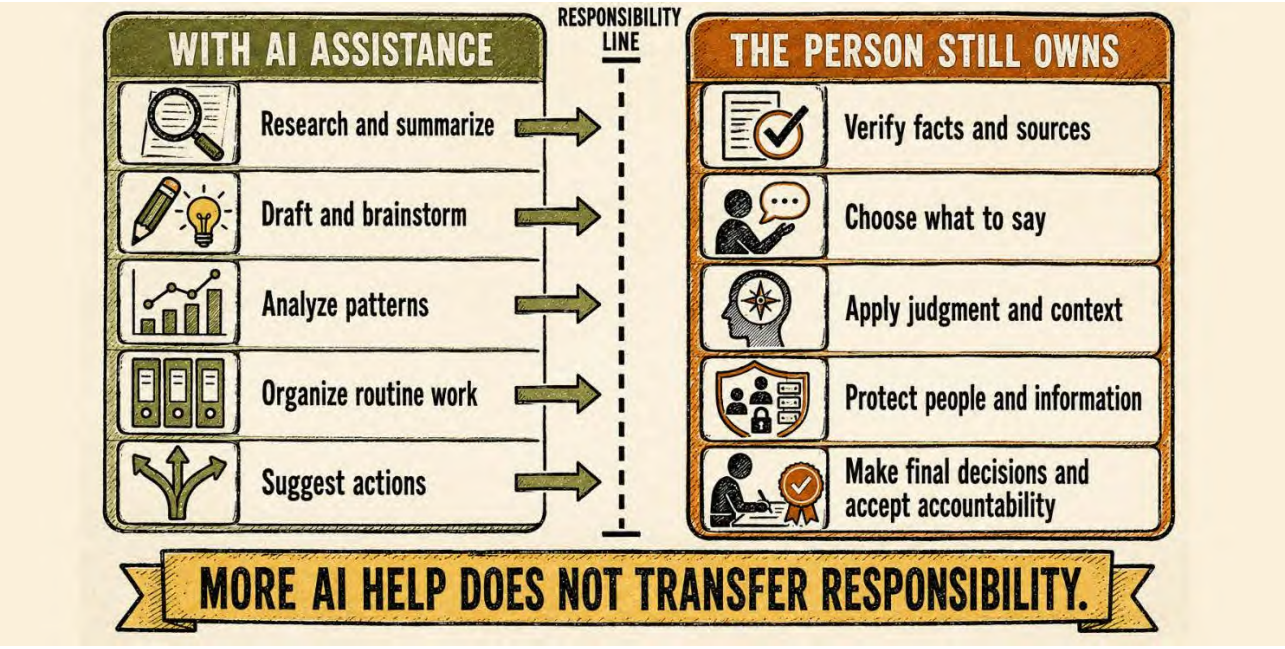


Figure 4.5A: AI can assist with a task, but responsibility for the decision and its outcome stays with the person who made it.



Good practice: Before you finalize any consequential piece of work, ask yourself whether you could explain and defend it in a meeting without reading from the screen. If you cannot, review it more closely before you rely on it.



What to ask next: `Turn this example into two sentences I could add to a self-review or LinkedIn profile.`

Try it now

Pick one of the seven human strengths listed in this chapter. Write one specific example from the past month where that strength, not ChatGPT, made the difference in a piece of your work.

Check the result

- ☐ Can you name a real example, not a general claim, of a human strength you used recently?
- ☐ Have you noticed any skill of your own growing weaker from disuse?
- ☐ Do you have a clear sense of which decisions in your work remain yours, no matter how much ChatGPT contributes?

Once you know where your own judgment adds the most value, you're ready to bring that combination to your organization. *Chapter 4.6* shows you how to propose and lead an AI improvement project at work.

Chapter 4.6 Propose and Lead an AI Improvement at Work

The improvement you measured in *Chapter 4.3* might be worth more than a personal habit. If it could help your whole team, this chapter shows you how to propose and lead that change responsibly.

Define the business problem

Start from the problem, not the tool. Describe what is currently slow, error-prone, or costly, in terms a colleague outside your immediate team would understand.

A clear business problem sounds like "Our weekly reporting takes six hours and often contains outdated figures," not "We should use ChatGPT more."

Estimate benefits and risks

List the likely benefits of your proposed improvement, using the measurements from *Chapter 4.3* where you have them. Then list the risks honestly, thinking about what could go wrong, who could be affected, and how serious the consequences would be.

Here are some examples:

- Benefits might include time saved, fewer errors, or faster turnaround for colleagues who depend on the work
- Risks might include exposing sensitive data, producing a result nobody checks carefully, or creating a dependency on a tool that changes without notice



Watch out: A proposal that only lists benefits will look less credible, not more. Naming the real risks, and how you plan to manage them, shows that you have thought the idea through.

Follow organizational policy

Before proposing anything, check your organization's AI policy, covering approved tools, data handling rules, and required approvals. *Chapter 2.1* introduced these rules for your own work. A team-wide proposal needs the same care, applied more broadly.

If your organization does not yet have a clear policy for this kind of change, say so in your proposal, and suggest who should be involved in setting one.

Win stakeholder support

Identify who needs to approve or support your proposal, and what each of them cares about most. A manager may care about cost and risk. A colleague who will use the new process may care about how much it changes their daily work:

1. List the people whose support you need.
2. Note what each person cares about.
3. Prepare a short explanation for each stakeholder, focused on what they care about.
4. Ask for feedback before requesting final approval.

Run a limited pilot

Do not propose a full rollout as your first step. Propose a small, time-limited pilot with a clear group of participants, a clear duration, and a clear way to measure the result:

1. Choose a small group or a short period for the pilot.
2. Set the same kind of baseline and measurements as in *Chapter 4.3*.
3. Run the pilot and collect feedback from participants as it proceeds.
4. Set a specific end date to review the results.



Good practice: Build a stop condition into the pilot from the start: a clear sign that would tell you to pause or cancel it. A pilot with no way to stop safely is harder to trust, both for you and for the people approving it.



What to ask next: `Point out the part of this pilot proposal a skeptical stakeholder would question first, and suggest how to address it.`

Report results honestly

When the pilot ends, report what actually happened, including any part that did not go as planned. Use your measurements from *Chapter 4.3* to show the real change, not an estimate.

If the pilot succeeded, propose a clear next step, such as expanding to a larger group. If it did not, explain what you learned and whether the idea is worth trying again in a different form.



Learn more online: Many organizations publish an internal AI use policy or acceptable use guidelines. If yours has not shared one, check your intranet or ask your manager or IT department before you check the current external guidance most workplaces point to, such as NIST's AI Risk Management Framework: `https://www.nist.gov/itl/ai-risk-management-framework`.

Try it now

Draft a one-page pilot for your *Chapter 4.3* improvement, or for a new idea worth proposing to your team. Include the problem, the evidence, the risks and controls, and how you would measure success.

Check the result

- ☐ Does your proposal name a specific business problem, rather than starting from the tool?
- ☐ Have you checked your organization's current AI policy before proposing this change?
- ☐ Does your pilot have a clear size, duration, and stop condition?

Leading a successful pilot project gives you a concrete result to point to. *Chapter 4.7* shows you how to turn that result, along with the rest of your achievement record, into a case for promotion or higher pay.

Chapter 4.7 Build a Case for Promotion or Higher Pay

If *Chapter 4.1* pointed you toward promotion or higher pay within your current organization, this chapter helps you turn your evidence into a case that others will find convincing.

Maintain an achievement record

Chapter 4.3 introduced the achievement record: a running list of your measured improvements. This is the raw material for your case, so keep adding to it well before you plan to ask for anything.

A few real entries make this concrete. One might read: "Automated new-hire onboarding emails, cutting turnaround from two days to same-day, with zero missed steps across twelve new hires." Another might read: "Built a skill that drafts meeting notes and action items within an hour of a call instead of the next day, and caught two follow-ups that would otherwise have been missed." Each entry names the task, states a specific before-and-after, and notes something beyond "faster," the same three parts *Chapter 4.3*'s baseline steps asked you to record.

A thin record built the week before a performance conversation is far less convincing than a steady one built over months.

Connect work to organizational outcomes

A manager rarely cares about a task getting faster for its own sake. They care about what that speed, quality, or reliability means for the organization: money saved, risk reduced, customers served better, or a project delivered on time.

For each entry in your achievement record, add one line connecting it to an outcome your organization already cares about. "Cut report preparation from six hours to two" becomes stronger as "Cut report preparation from six hours to two, freeing time for the client research that led to two renewed contracts."

Describe expanded responsibilities

A case for promotion often depends on showing that your actual responsibilities have grown, not only that your existing tasks have improved. Look for ways your role has expanded: new kinds of work you now handle, decisions you now make without being asked, or colleagues who now come to you for help.



Watch out: Do not claim responsibility for work you only touched briefly or supported in a minor way. An overstated claim is easy to challenge and can undermine the credibility of your entire case.

Prepare for a performance conversation

Before a formal review or a direct request, prepare as you would for any important conversation:

1. Choose your three or four strongest pieces of evidence from your achievement record.
2. Practice explaining each one in two or three sentences, focused on the outcome.
3. Anticipate likely questions or objections and prepare honest answers.
4. Decide what you are specifically asking for, whether that is a new title, a raise, or expanded responsibilities.



Real-world example: A customer support lead preparing for her annual review does not walk in with a general sense that she "did a lot with AI this year." She brings three specific examples, each with a before-and-after measurement and a connection to a result her manager already tracks, such as customer response time.

Propose the next contribution

Beyond looking back at what you have done, describe what you would do with expanded responsibility or a new role. A specific, credible plan for the future often counts as much as a strong record of the past.

Case for Promotion — One-Page Summary

Jordan Ellis • Project Coordinator • Prepared for annual review, July 2026

1. Weekly status report	
Activity	Automated the weekly status report using a saved ChatGPT prompt fed by pasted source data.
Output	Cut report preparation from 90 minutes to 35 minutes each week.
Outcome	Reports now go out consistently by Friday morning instead of trickling in over the weekend.
Value	**Frees roughly 4 hours a month for client research, and gives leadership more reliable visibility into project status.**

2. Client onboarding emails	
Activity	Built a reusable onboarding email template and refined it with ChatGPT.

Figure 4.7B: A one-page summary turns your achievement record into a case built from activity, output, outcome, and organizational value.



Sample files: A one-page Word summary for a case for promotion: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book4/4-7-Case-for-Promotion.docx`. Three achievement entries reuse the same tasks from *Figure 4.3B* (status report, onboarding emails, meeting notes) for continuity across the book.



Good practice: Rehearse your case out loud before the actual conversation, ideally with a colleague who can ask you questions. Hearing your own case spoken aloud often reveals gaps that are easy to miss on the page.



Learn more online: The US Bureau of Labor Statistics publishes current wage data by occupation and location, useful when you decide what pay to ask for: `https://www.bls.gov/oes/`. Readers outside the United States can ask ChatGPT to find an equivalent wage data source for their country.

Try it now

Choose your three strongest achievement-record entries. For each one, write a line connecting it to an organizational outcome, then write one sentence proposing what you would take on next.



Good practice: Before a performance conversation, paste your case into a second AI, one you have not used to draft it, and ask it to argue against you: "What would a skeptical manager say is missing or overstated here?" A model that helped you write the case will not spot its own blind spots as reliably as one seeing it fresh.

Check the result

- ☐ Does each piece of evidence connect clearly to something your organization already values?
- ☐ Have you avoided overstating your role in any piece of evidence?
- ☐ Do you have a specific ask, rather than a general hope that good work will be noticed?

A strong case for promotion works inside your current organization. *Chapter 4.8* prepares you for a move outside it, if that's the direction you choose, covering how to research a new role and build a résumé you can stand behind.

Chapter 4.8 Prepare for a Better Job or Career Change

If *Chapter 4.1* pointed you toward a career change, this chapter takes you from research through to a completed application, using the skills you built earlier in this book.

Research resilient and growing roles

Before committing to a target role, research whether it is a good bet for the years ahead. Look for roles with steady or growing demand and be honest about how automation and AI might affect that demand over time:

1. Ask ChatGPT to search for current data on demand and outlook for your target role.
2. Check at least one source directly, using the research habits from *Chapter 1.8*.

3. Note anything the sources disagree about, rather than picking the most convenient answer.
4. Consider roles adjacent to your first choice, in case they offer a stronger outlook.



Learn more online: The US Bureau of Labor Statistics Occupational Outlook Handbook describes typical duties, pay, and the ten-year job outlook for hundreds of occupations: `https://www.bls.gov/ooh/`. Readers outside the United States can ask ChatGPT to find their country's equivalent outlook resource.

Identify skill gaps

Return to the current-to-target skills map you built in *Chapter 4.4* or build one now if you have not already. A career change usually needs a slightly wider set of skills than a promotion within your current field, so review the map with fresh eyes.

Improve a résumé and LinkedIn profile honestly

ChatGPT can help you write a clearer résumé and profile, but the content must remain accurate. Every claim should describe something you actually did:

1. Provide ChatGPT with your work history and the target role's requirements.
2. Ask for a draft that highlights relevant experience clearly.
3. Check every claim against what you actually did, and correct anything overstated.
4. Ask for a second pass focused on clarity and length, not new claims.



Watch out: An impressive-sounding résumé built on exaggerated claims can fall apart in an interview, when someone asks a specific question about a project you described. Keep every claim something you can defend in detail.

Prepare work samples

If your target role values demonstrated skill, prepare one or two samples that show your ability directly, rather than only describing it. *Chapter 4.10* covers how to build a full portfolio case study, but even a single strong sample helps here.



Real-world example: Someone moving from general marketing into data-informed marketing does not just add "data skills" to a résumé. She includes a short writeup of one campaign where she used data to change the approach midway through, with the result that followed.

Practice interviews

Use ChatGPT to prepare for likely interview questions, based on the target role's requirements and common questions in that field:

1. Ask ChatGPT to generate likely interview questions for the role.
2. Answer each question out loud, then review your answer against the question's intent.
3. Ask ChatGPT to act as an interviewer and ask follow-up questions based on your answers.
4. Practice explaining your work samples in two or three sentences each.

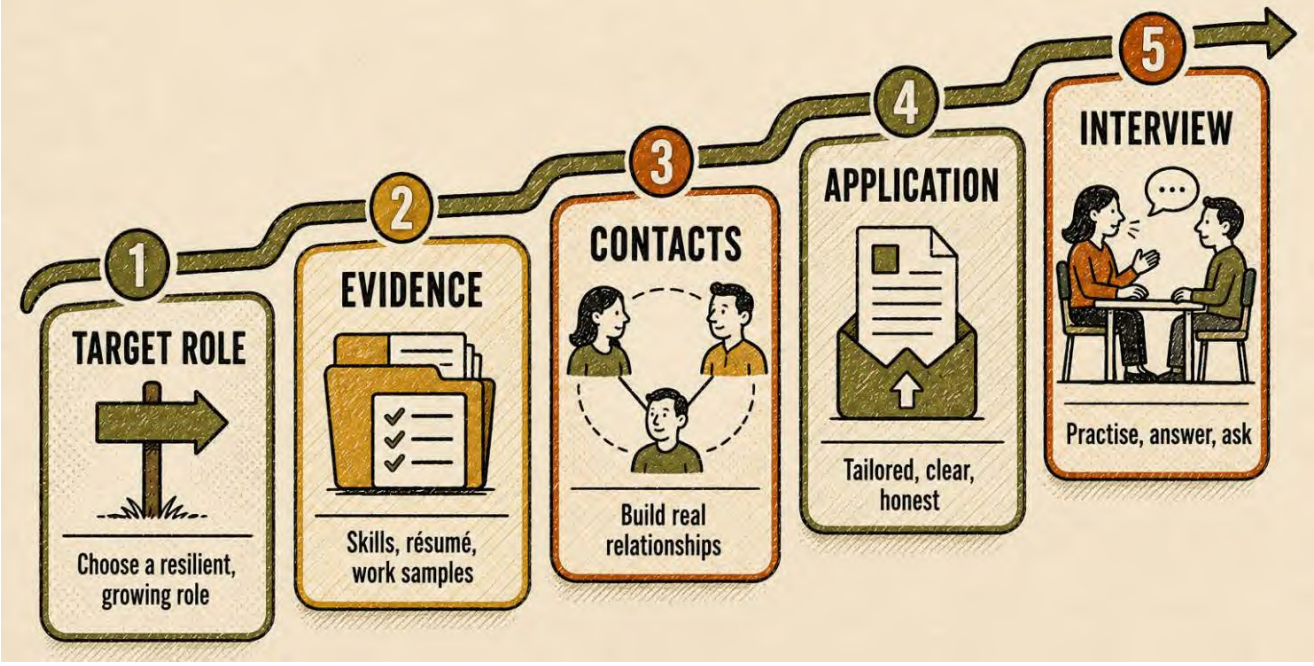


Figure 4.8A: A career-change pathway moves from a target role, through evidence and contacts, to an application and an interview

Build a networking plan

Most career changes benefit from real conversations with people already in or near the target role, not only applications sent into a system. Use ChatGPT to help you prepare for these conversations, not to replace them:

1. List people you already know who work in or near the target field.
2. Draft a short, specific message asking for a brief conversation, not a job.
3. Prepare two or three questions that show genuine interest in their experience.
4. Follow up with a thank-you message that references something specific from the conversation.

Good practice: Treat every networking conversation as a chance to learn something real about the role, not only as a step toward an opportunity. People notice the difference, and genuine curiosity tends to lead to better conversations.

Try it now

Research your target role's outlook, then draft one paragraph of your résumé aimed at that role. Check every claim in the paragraph against your actual experience before moving on.

Check the result

- ☐ Is your target role's outlook based on current research, not an old assumption?
- ☐ Does every claim in your draft résumé describe something you could explain in detail if asked?
- ☐ Do you have at least one real person you could reach out to for a networking conversation?

A résumé states your experience in a few lines. *Chapter 4.9* shows you how to rehearse that interview out loud, using Voice mode, before you sit down with a real interviewer.

Chapter 4.9 Rehearse High-Stakes Conversations with Voice Mode

Chapter 4.6 through *Chapter 4.8* asked you to write a proposal, build a case for promotion, or prepare for an interview. Writing the words is one skill. Saying them out loud, under a little pressure, is another.

This chapter shows you how to rehearse a real conversation with ChatGPT before you have it with a real person. You will practice on your phone, in a private space, using Voice mode. If your plan and device support live video, ChatGPT can also see what you point the camera at, such as a whiteboard or a printed page, the same combination introduced in *Chapter 1.3*.

By the end of the chapter, you will have run one rehearsal and turned it into a short report you can use before the conversation that counts.

Set up the rehearsal

A useful rehearsal needs three things: a clear scenario, a specific person for ChatGPT to play, and a time limit. Without these, the conversation tends to drift into a general chat instead of practice:

1. Find a private room where you can talk out loud without being overheard.
2. Open the ChatGPT app on your phone and select the waveform button to start Voice mode.
3. Describe the scenario in one or two sentences: who you are talking to, what the conversation is about, and what you want from it.
4. Ask ChatGPT to play that person, such as a hiring manager for a named role, a manager deciding on a raise, or a stakeholder who controls a budget.
5. Set a time limit before you start, such as 20 minutes, so the rehearsal stays focused.



Good practice: Vary the difficulty across rehearsals. Ask ChatGPT to play a friendly version of the person the first time, then a rushed or skeptical version the next time. One easy rehearsal will not prepare you for a hard conversation.

Practice a job interview

Interview practice benefits most from the camera, especially if your work involves diagrams, whiteboards, or drawings. *Chapter 4.8* covered how to prepare your résumé and research the role. This section rehearses the conversation itself:

1. Tell ChatGPT the job title, the type of company, and the seniority level.
2. Ask ChatGPT to ask a mix of technical, behavioral, and teamwork questions, one at a time.
3. Answer out loud. If a question calls for a diagram, turn on the camera and sketch it on a whiteboard or a sheet of paper while you explain it.
4. Ask ChatGPT to ask one follow-up question about your answer or your diagram before moving to the next question.

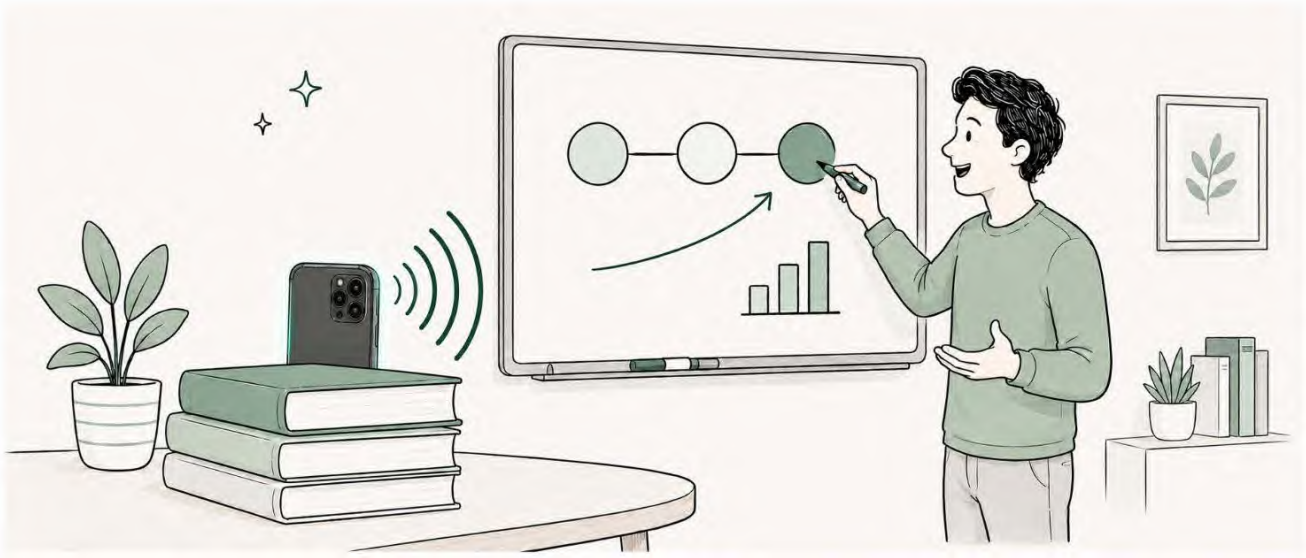


Figure 4.9A: A phone propped against a stack of books, facing a whiteboard, while the reader talks and sketches at the same time



Real-world example: A reader preparing for a software engineering interview asked ChatGPT to play a hiring manager for a mid-level backend role. When ChatGPT asked how he would design a message queue, he stood at a whiteboard, sketched the components, and pointed the camera at himself drawing while he explained it. ChatGPT then asked why he chose one queue design over another, the kind of follow-up question a real interviewer might ask.

Adapt the rehearsal to other conversations

The same method works for most conversations where you only get one attempt to get it right. Change the person ChatGPT plays and what you bring to the camera:

Conversation	Ask ChatGPT to play	What to bring to the camera
Asking for a raise or promotion (*Chapter 4.7*)	A manager who needs to see the business case	Your one-page case from *Chapter 4.7*
Pitching an AI improvement (*Chapter 4.6*)	A stakeholder worried about cost or risk	A rough sketch of the workflow change
A sales or client call (*Chapter 4.10*)	A potential client with a real objection	A one-page outline of your offer
A difficult conversation at work	A colleague in a specific mood, such as rushed or defensive	Nothing. Focus on tone and wording.

Table 4.9: Ways to adapt the rehearsal method to other high-stakes conversations.

A similar technique works well when helping a child with homework, or when you are learning something new. Start a Voice conversation with ChatGPT, show it your textbook and what you are working on or studying, and have an interactive discussion about it.

Turn the rehearsal into a report

A rehearsal helps most when you can see what to change. Once you finish talking, ask ChatGPT to review the conversation and turn it into a short report.

1. End Voice mode when your time limit is up.
2. If you used the camera, take one photo of your final diagram or notes and attach it to the chat.
3. Ask ChatGPT to review the conversation and list what went well, what to improve, and one specific change to make before the real conversation.
4. Save the report somewhere you will see it again, such as your notes app or a document.



What to ask next: `Review this practice interview. List one answer that was strong, one answer that was vague, and one specific change I should make before the real interview.`



Watch out: Voice mode can pick up sound and images beyond what you intend, the same caution from *Chapter 1.3*. Rehearse in a private room and describe your scenario with a role and a general situation instead of your actual employer's confidential details or a real client's name.



Learn more online: OpenAI's Voice mode help article describes which plans and devices currently support live camera sharing during a conversation: `https://help.openai.com/en/articles/8400625-voice-mode-faq`.

Try it now

Pick one conversation from this chapter, such as a job interview or the raise conversation from *Chapter 4.7*. Set a 20-minute time limit, then run one full rehearsal using Voice mode:

1. Describe the scenario and the person you want ChatGPT to play.
2. Have the conversation out loud. Turn on the camera if a diagram or a document would come up in the real conversation.
3. Ask ChatGPT to turn the rehearsal into a short report.
4. Read the report and write down the one change you will make before the real conversation.

Check the result

- ☐ Did you set a time limit before you started?
- ☐ Did ChatGPT ask at least one follow-up question based on what you said or showed it?
- ☐ Does the report name one specific answer or moment to improve, rather than only general encouragement?
- ☐ Do you have one clear change to make before the real conversation?

A rehearsal builds your confidence for one conversation. *Chapter 4.10* shows you how to build a portfolio that backs up your claims with evidence anyone can review.

Chapter 4.10 Build a Portfolio and Measurable Case Study

A case study is a short, honest account of a real problem you improved, with evidence to prove it. This chapter shows you how to build one, whether it will support a promotion, a career change, or an independent offer.

Select an internal, personal, or client project

Choose one project that shows your ability clearly, drawing from your achievement record in *Chapter 4.3*. It can come from your current job, a personal project, or client work, as long as you have the right to share the details. Pick a project with a clear before and after, rather than one with vague or gradual improvement. A reader should be able to see the change quickly.

Explain the original problem

Describe the problem the way it looked before you addressed it: what was slow, costly, unreliable, or frustrating, and for whom. A reader unfamiliar with your work should understand the problem within a sentence or two.

Show the process and safeguards

Explain what you actually did, at a level of detail that builds trust without becoming a full tutorial. Include the safeguards you used, such as verification steps, human review points, or the acceptance criteria from *Chapter 3.3*, since these show a careful approach rather than a rushed one.



Watch out: A case study that only shows the polished result can raise doubts about how the work was actually done. Naming your checks and safeguards makes the achievement more credible, not less impressive.

Show your improvement

Present your measured result clearly, using the same before-and-after approach from *Chapter 4.3*. A specific number, chart, or comparison is far more convincing than a general claim that something improved.

This step covers the last two parts of the one-page structure: the result you achieved, and the evidence behind it. State them together, since a reader should never have to take a number on trust when a quick check is available.

- **Result**: a single sentence naming the before number and the after number side by side, such as "invoice processing time fell from twelve minutes to four"
- **Evidence**: the artifact that lets someone else verify the result without relying on your word alone, such as a screenshot, an exported report, a dashboard view, or a short quote from someone affected by the change

Choose the measure your reader already cares about, using the same time, quality, errors, or outcome categories from *Chapter 4.3*, rather than whichever number happens to look most impressive.



Figure 4.10A: A one-page case study moves from the original problem, through what you did, to the result and the evidence behind it

What to ask next: `Rewrite this case study's headline claim in one sentence a stranger would understand in five seconds.`

Remove confidential information

Before sharing a case study with anyone outside your organization, or even with certain people inside it, review it for confidential material:

1. Remove client names, internal figures, and proprietary details unless you have explicit permission to share them.
2. Replace specific numbers with proportional ones where necessary, such as "cut processing time by 40 percent" instead of exact minutes tied to a specific system.
3. Ask a colleague or manager to review the case study if you are at all unsure what is safe to share.
4. Keep an unredacted internal version separately, for your own records and for conversations where full detail is appropriate.

Real-world example: An operations analyst wants to share a case study about a workflow that cut invoice-processing errors. Before sharing it publicly, she replaces the client's name with "a mid-sized logistics client" and reports the improvement as a percentage rather than exact dollar figures tied to a specific contract.

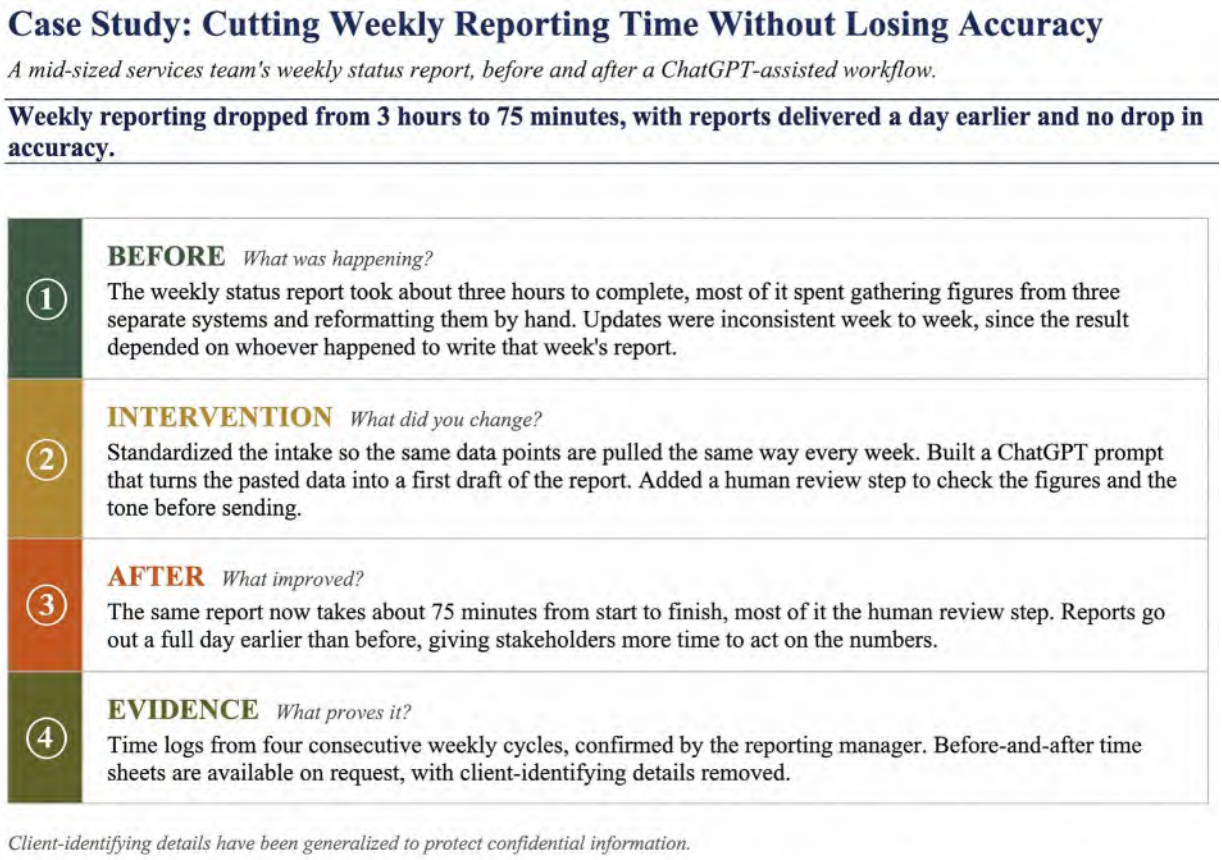


Figure 4.10B: A finished case study fits on one page and moves cleanly from the original problem to the proven result



Sample files: A one-page Word case study: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book4/4-10-Case-Study.docx`. This mirrors *Figure 4.10A* exactly using the same weekly-reporting scenario.



Good practice: Write your case study for a reader who does not know you and has thirty seconds to decide whether to keep reading. A clear headline claim, followed by the evidence, works better than a long, gradual buildup.

Try it now

Choose one project from your achievement record. Draft a one-page case study using the before, intervention, after, and evidence structure, and review it once for anything confidential that needs removing.

Check the result

- ☐ Would a reader unfamiliar with your work understand the original problem within a sentence or two?
- ☐ Does your evidence include a specific measurement, not only a general claim?
- ☐ Have you checked the case study for confidential information before sharing it outside your organization?

A strong case study can support a job application, or something more independent. *Chapter 4.11* shows you how to turn a problem you understand well into a specific offer or product you could sell.

Chapter 4.11 Turn a Problem into a Service or Product

If *Chapter 4.1* pointed you toward earning independently, this chapter is where that idea becomes a specific service or product. The best service or product ideas start with a real problem someone else has, not with a skill you want to sell.

Select a customer group

Choose a specific group of people or businesses who share a problem you understand well, ideally one close to your own experience or professional background. A narrow, well-understood group is easier to serve well than a broad, vague one. "Small accounting firms that struggle to keep client reports consistent" is easier to design an offer around than "small businesses."

Validate the need

Before building anything, confirm that your chosen group actually experiences this problem, and cares enough to want it solved. Talk to real people in the group if you can, and use ChatGPT to help you research how the problem is discussed publicly, such as in forums or reviews.



Watch out: It is easy to convince yourself a problem is common because you personally find it frustrating. Confirm the need with people outside your own perspective before investing further time.

Research alternatives

Look at how people in your target group currently solve, or try to solve, this problem. This might include competing services, do-it-yourself approaches, or living with the problem unresolved:

1. Search for existing services or products aimed at the same problem.
2. Note their strengths and their gaps, especially anything customers complain about.
3. Identify what you could do differently or better, based on your own experience.
4. Confirm your planned offer is not a repeat of something already well served.

Define outcome, scope, exclusions, and evidence

Write a clear description of what your offer actually delivers:

- **Outcome**: the specific result the customer gets, described in terms that matter to them
- **Scope**: exactly what is included, stated plainly enough to avoid confusion later
- **Exclusions**: what is not included, so expectations stay realistic on both sides
- **Evidence**: proof that you can deliver this outcome, such as the case study from *Chapter 4.10*

Choose service or reusable product

Decide whether your offer works best as a service you deliver personally for each customer, or as a reusable product built once and used by many customers.

A service usually starts faster and depends more on your direct time. A reusable product takes longer to build well but can serve more customers without growing your time commitment in the same way. Many independent workers start with a service and turn a proven, repeated version of it into a product later.

Price and package responsibly

Set a price that reflects the value delivered and the effort required, and be honest with yourself about your actual costs, including your time. Avoid pricing based only on what feels comfortable to charge, since that often leads to underpricing real value:

1. Estimate your actual cost, including time, at a rate you would be comfortable defending.
2. Research what comparable offers charge in your market.
3. Set a price and clearly state what it includes.
4. Decide how you will handle requests that fall outside your defined scope.



Good practice: Write your price and scope down before your first real conversation with a prospective customer. Deciding in the moment, under pressure to close a deal, tends to produce inconsistent and sometimes regrettable terms.

Put the basics in place before you take payment

A priced offer is not yet a legal or financial one. Before your first paid engagement, put a few unglamorous pieces in place:

1. A short written agreement stating the scope, price, and payment terms, even for a small job.
2. A way to invoice and get paid that keeps a record you can find later.
3. A habit of setting aside a portion of each payment for taxes, since nobody will withhold it for you.
4. A quick check into whether your location requires registering a business or reporting this income separately.



Watch out: Tax and business rules vary by country, state, and how much you earn, and getting them wrong can cost more than the income you are trying to generate. Confirm your obligations with a qualified accountant or tax advisor rather than relying on general information, including from ChatGPT.



Learn more online: The US Internal Revenue Service explains when you owe self-employment tax, how to make quarterly estimated payments, and how to choose a tax professional, at the Self-Employed Individuals Tax Center: `https://www.irs.gov/businesses/small-businesses-self-employed/self-employed-individuals-tax-center`. Ask ChatGPT to find an equivalent link for you in your home country or region.

Try it now

Draft an offer for one problem you could realistically help solve. Fill in the customer, the problem, the outcome, your proof, the scope, a price, and the main risk.

An optional online-only chapter covers how to find customers and deliver products and services to them responsibly: `https://github.com/markjprice/chatgpt-vb/blob/main/docs/ch4-14-find-customers.md`.

Check the result

- ☐ Have you confirmed the need with someone outside your own perspective?
- ☐ Does your scope clearly state what is included and what is not?
- ☐ Is your price based on real cost and market research, rather than a guess?

Chapter 4.12 addresses the question many readers are actually carrying when they pick up a career book about AI in 2026: whether AI is going to take their job. Skipping that question does not make it go away. It just means you are left to answer it alone, with whatever you last saw in the news. The next chapter will resolve that question calmly and put you in a position to choose a direction on purpose rather than out of anxiety.

Chapter 4.12 Decide How to Respond to AI in Your Field

Many people open a book like this one while carrying a specific worry: that AI will take their job. This chapter does not promise that it will not. It gives you a way to find out what is actually happening in your own field, so you can choose your next move based on evidence instead of a headline.

Separate two different questions

"Will AI take my job?" is really two separate questions, and they have different answers.

The first is whether AI can do some of the tasks inside your job. For most jobs, the honest answer is yes, at least for some tasks. You have already seen this across *Book 1*, *Book 2*, and *Book 3*: research, drafting, analysis, and routine automation.

The second is whether AI can do your whole job, including the judgment, relationships, and context you bring to it. *Chapter 4.5* covered this in more detail earlier, but the short answer is that this depends heavily on your specific role, and it changes as tools change.

Confusing these two questions is what turns a manageable planning problem into a source of constant worry. This chapter treats them separately, on purpose.

Research your own field instead of a general headline

General news coverage about AI and jobs is written for a broad audience, not for your specific role, industry, and location. Get a current, specific answer instead:

1. Ask ChatGPT to search for current analysis and data about how AI is affecting your specific role and industry, not AI and jobs in general.
2. Ask it to note where sources agree and where they disagree.
3. Check at least one source yourself, using the research habits from *Chapter 1.8*.
4. Write one paragraph, in your own words, summarizing what you found and how confident you are in it.



Watch out: Treat any source, including ChatGPT, that gives you a single confident percentage or year with caution. Credible sources about AI's effect on employment usually include ranges, disagreement, and stated assumptions, not one certain number.



Learn more online: The World Economic Forum's Future of Jobs Report series surveys employers worldwide about which tasks and skills AI is expected to change over the next five years: `https://www.weforum.org/publications/series/future-of-jobs/`. Treat its figures as informed estimates, not guarantees, the same way this chapter asks you to treat any other single source.

Recognize the two responses that hold up

Once you have a specific, current picture of your field, two responses tend to hold up better than either panic or denial.

The first is building AI fluency early, while it still sets you apart. The chapters in this book, on auditing your role, learning new skills, and proposing improvements at work, all build this response step by step.

The second is investing in the human strengths that stay yours regardless of what AI can do, covered directly in *Chapter 4.5*. Judgment, relationships, and domain knowledge do not disappear because a tool gets better at drafting.

Most readers will use some combination of both, and the right mix depends on your research, not on a general rule that applies to everyone equally.



Real-world example: A claims processor at an insurance company read that AI could handle claims review and assumed her role was ending soon. She researched her specific role instead and found that AI was handling straightforward claims, while complex and disputed claims, which made up much of her actual caseload, still needed a person. She used the time freed up on simple claims to build expertise in the complex cases, which made her more valuable, not less.



Good practice: Revisit this research every few months rather than treating it as settled. Your field, and the tools affecting it, will keep changing over the life of this book.

Failed predictions and Jevons paradox

In 2016, computer scientist Geoffrey Hinton, a British-Canadian researcher often called the "Godfather of AI," said medical schools should stop training radiologists, since he expected AI to soon read scans better than any human could.

Hinton has since said he was wrong. A decade later, radiologists are earning more and are harder to hire than when he made that claim, partly because AI-assisted imaging increased how many scans get ordered and reviewed, not fewer.



Good practice: Take predictions about AI from anybody, even experts, with a lot of salt.

Economists call this pattern **Jevons paradox**: when a technology makes something cheaper, people do not use less of it, they use so much more that total demand rises. It comes from an 1865 observation that more efficient steam engines burned less coal per task but made steam power worth using so much more widely that total coal consumption went up. Previously, most people had assumed that as steam engines became more efficient, demand for coal would fall. The opposite happened.

Economists are now applying the same logic to AI: as thinking gets cheaper, demand for analysis, advice, and decisions can grow faster than AI alone can supply it, which increases rather than shrinks the work available for people. Current reporting on the labor market backs this up: white-collar employment has been rising, not falling, alongside AI adoption.

A related idea explains why demand for people can hold up even as AI improves: comparative advantage. Even if AI became better than humans at every task, it would still make sense to point AI at its highest-value uses and leave the rest to people. AI's capacity is not free or unlimited. Every hour of compute spent on one task is an hour not spent on another.

A skilled lawyer illustrates the same idea. The lawyer could likely research a legal issue faster than a law clerk, but still hires one, because the lawyer's time is worth more spent meeting with clients. Economists have studied this kind of trade-off in international trade for two centuries, long before AI existed. Applied to AI, it suggests that as AI takes over the mechanical but detail-oriented tasks, plenty of higher-value judgment-based tasks still need a person. Economists disagree about how long that holds as AI keeps getting cheaper, so treat it as a reason for calm rather than a guarantee.



Learn more online: A 160-year-old paradox explains why AI will create more lawyers and accountants—not fewer, top economist says: `https://fortune.com/2026/04/28/will-ai-kill-jobs-why-not-jevons-paradox-torsten-slok/`.

Try it now

Research your specific role and industry using the four steps in this chapter. Write one paragraph summarizing what you found, including anything sources disagreed about.

Check the result

- ☐ Did your research focus on your specific role, not AI and jobs in general?
- ☐ Did you check at least one source yourself, rather than relying only on a summary?
- ☐ Do you have a clearer, calmer picture of your situation than when you started?

Now that you have a specific, current picture of your field, you've covered every path this book offers: improving your current role, preparing for a new one, or building something to offer customers directly. *Chapter 4.13* asks you to choose one path and commit to a real 30-day project, measured against a baseline you set yourself.

Chapter 4.13 Project – Complete a 30-day Career Value Project

This final project brings together the whole book: choosing a direction, finding high-value work, measuring your improvement, and building a case for what comes next. Over the next thirty days, you will complete one track fully, with real evidence to show for it.

Choose your track

Pick the track that matches the direction you chose in *Chapter 4.1*:

1. **Improve your present job**: create and run one approved workflow pilot, compare it with your baseline, and prepare a one-page results report.
2. **Prepare for promotion or career change**: produce a skill-gap plan, a completed demonstration, a portfolio case study, and a conversation or application plan.
3. **Earn independent income**: validate one offer, produce a demonstration, approach a defined number of suitable prospects, and complete or arrange one pilot.

You may recognize pieces of your track from earlier chapters in this book. This project asks you to complete the whole sequence, not only the individual steps.

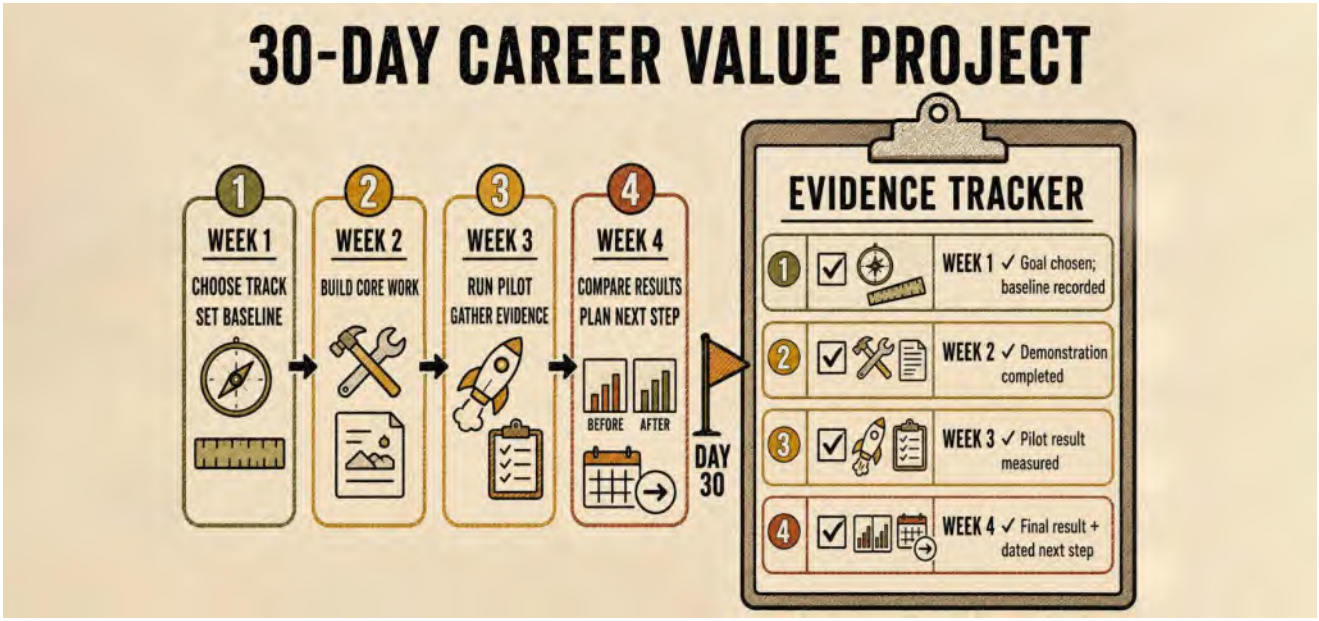


Figure 4.13A: A four-week roadmap breaks the thirty-day project into manageable stages, with an evidence tracker recording your progress along the way

Establish evidence

Whichever track you choose, start by setting a clear baseline, using the approach from *Chapter 4.3*. Record where you stand before the thirty days begin, so your progress is measurable rather than a general impression:

- For the present-job track, record your current baseline for the task you plan to improve
- For the promotion or career-change track, record your current skill levels and your existing achievement record
- For the income track, record your current offer, if any, and your current customer conversations, if any

Complete one pilot

Spend the bulk of the thirty days completing the core work of your track:

1. **Present-job track**: build and run your workflow pilot, following the checkpoints from *Chapter 4.6*.

2. **Promotion or career-change track:** complete your learning plan's demonstrable project from *Chapter 4.4*, and build your case study using *Chapter 4.10*.
3. **Income track:** approach a defined number of prospects, aiming for at least three real conversations, and complete or arrange one paid or trial engagement.



Watch out: Thirty days is enough time to make real progress, not enough time to make everything perfect. Set a defined scope for your pilot before you start, and resist the urge to expand it once you are underway.

Review results

At the end of the thirty days, compare your result against your baseline, using the same honest approach from *Chapter 4.3*.

1. **Baseline:** record the task as you do it before AI enhancement
2. **Weekly updates:** record the task with AI enhancements
3. **Result:** calculate the time saved as an absolute and as a percentage reduction



Sample files: A spreadsheet launch tracker: `https://github.com/markjprice/chatgpt-vb/tree/main/Files/Book4/4-13-Launch-Tracker.xlsx`. It continues the same weekly-status-report thread from *Figure 4.3B* and *Figure 4.10B*.

30-Day Career Value Project — Launch Tracker

Track: Improve your present job — automate the weekly status report

	Stage	Week / Date	Time (min)	Notes / Evidence
①	**Baseline**	Week 0 — 28 Jun	90	Manual report, gathered by hand from three separate systems.
②	Week 1	05 Jul	70	Dry run of a ChatGPT-assisted draft prompt; no live report sent yet.
	Week 2	12 Jul	55	First live report using the new workflow; refined the prompt based on gaps found.
	Week 3	19 Jul	45	Second live cycle; added a human review checkpoint before sending.
	Week 4	25 Jul	35	Third and fourth live cycles; time held steady with zero formatting errors.
③	**Final Result**	Week 4 wrap — 25 Jul	35	Time logs from four weekly cycles, confirmed by the reporting manager.

Time saved vs. baseline (min)	55
Percent reduction	61.1%

Baseline recorded before the pilot began; weekly entries logged each Friday during the four-week pilot; final result compared against the original baseline. Minutes are self-timed and rounded to the nearest 5.

Figure 4.13B: A completed launch tracker records your baseline, your weekly progress, and your final result in one place. Some readers may build this as a ChatGPT Sites page instead of a document or spreadsheet

Produce a next-step plan

Regardless of how the thirty days went, write a short plan for what comes next. If the pilot succeeded, describe how you will expand or repeat it. If it fell short, describe what you learned and what you would try differently:

1. Summarize what you set out to do and what actually happened.
2. State whether you met the goal you set at the start.
3. Describe the single most important lesson from the thirty days.
4. Write your next concrete step, with a realistic date.



Good practice: Share your results with at least one other person, whether a manager, a mentor, or a peer. Explaining your project out loud often clarifies what you have actually learned, and an outside perspective can catch something you missed.

Try it now

Choose your track, set your baseline this week, and block time on your calendar for each of the four weeks ahead. Begin the first stage of your pilot.

Check the result

- ☐ Did you complete the full sequence for your track, not only part of it?
- ☐ Does your final result compare clearly against the baseline you set at the start?
- ☐ Do you have a specific, dated next step written down?

Share your result

If you are willing, share your before-and-after numbers, with anything confidential removed, in the book's GitHub Discussions or Packt's Discord channel for this book. Future editions of this book may draw on real reader results, with permission, in place of some of the examples used here. Let's make it real together.

Wrap up with the Epilogue

Whichever track you chose for your thirty-day project, you've now used ChatGPT to ask better questions, create real work, automate a genuine routine, and advance something that counts in your career or income.

The *Epilogue* looks back across all four books and points to where you might go next, especially if the automations and workflows you built have left you curious about writing and running code yourself.

Epilogue

In this epilogue, you will learn:

- Where this book ends and your work continues
- What stays true as ChatGPT changes
- Why you should spend your next $20 on a second AI
- What to do if you want to build, not just direct
- A last word

Where this book ends and your work continues

You started this volume by asking ChatGPT a question. You ended it with improving your current job tasks, or building a documented case for a promotion, or a real product or service offering to a paying customer. That progress did not come from one clever trick. It came from four stages, each building on the one before it.

Book 1 taught you to ask ChatGPT better questions: how to set up your account, hold a real conversation, and use ChatGPT to learn, research, and decide. You finished it with a personal AI workbench built around your own life and work.

Book 2 turned those conversations into finished results. You drafted and revised documents, made sense of spreadsheets and long reports, created images and presentations, and connected ChatGPT to real files and workplace information. You finished it by building a useful weekly briefing that serves a real audience.

Book 3 showed you how to stop repeating yourself. You started by turning your weekly briefing into a recurring task that runs on a schedule, turned a manual routine into a skill, built a specialized GPT, and, if you completed the online chapter, you connected trusted tools with MCP. You finished it with a complete workflow, tested against a baseline and checked for what happens when something goes wrong.

Book 4 asked what all of that was for. You measured the value of your improved work, learned the skills your next role requires, and built a case for a promotion, a new job, or a service of your own. You finished it with a thirty-day project measured against a result you set yourself, not a guess.

What stays true as ChatGPT changes

ChatGPT will keep changing. Menu names, default modes, and exact controls may look different by the time you read this. Some of the screenshots in these books are already a step behind the current interface.

What should not change is the way you approached each task: define what you want, check your employer's policy, verify what ChatGPT gives you, and decide for yourself whether the result is good enough to use. That judgment is the actual skill this volume taught. It will transfer to whatever ChatGPT looks like next year, and to any AI tool you use after that.

Why you should spend your next $20 on a second AI

When ChatGPT starts running into its limits, the obvious next step is to pay OpenAI more: upgrading from Plus to Pro, from around US$20 a month to US$100 or US$200, depending on the tier available in your country at the time. Before you make that jump, consider a different one. Keep your Plus plan at US$20 and spend another US$20 on a second AI tool instead.

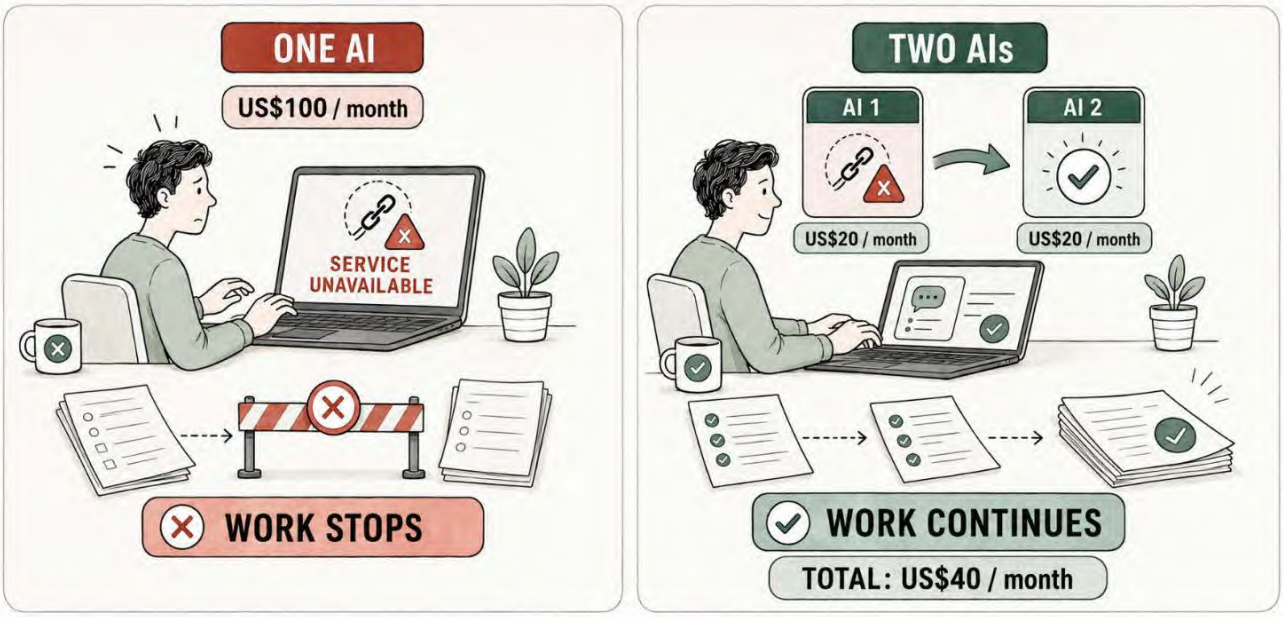


Figure E1: Two AI services instead of one bigger AI have benefits beyond cost

Think of ChatGPT and Claude, from Anthropic, as the Coke and Pepsi of AI assistants. Both are capable general-purpose tools, and picking a primary between them is closer to personal taste than a clear technical verdict. Whichever one you did not choose as your primary is usually the strongest possible secondary: not a downgrade, but a second frontier-grade model with its own way of writing, reasoning, and handling a task.

A second full tool earns its US$20 in a way a bigger single-vendor plan often doesn't. This guide has asked you, in almost every chapter, to check a result rather than trust it on the first try. A second AI extends that same habit: run a demanding request past both, and you get a genuine second opinion instead of the same model reasoning with itself twice. It also protects you from a single company's outage, a single model's blind spot, or one vendor holding every part of your work.

That last risk is the easiest to underestimate until it happens to you. Every major AI provider suspends accounts, sometimes for a genuine policy violation and sometimes because an automated fraud or abuse system flags something ordinary as suspicious, and the reason given is rarely detailed enough to argue with. Anthropic's own published transparency reporting counted 690,000 banned Claude accounts in one six-month period, against roughly 35,000 appeals and only about 1,000 overturned decisions, a success rate under three percent, a figure worth checking again since the report updates every six months: `https://www.anthropic.com/transparency/system-trust-reporting`. OpenAI describes its appeals process but doesn't publish how often it bans accounts or how often appeals succeed.

The stories behind those numbers are concrete: in more than one reported case, a growing company's staff logged in one morning to find every employee account suspended at once, with no specific explanation and a short online form as the only route to appeal, while a separate billing account kept charging them regardless. Individual long-paying subscribers have reported the same experience alone: years of saved conversations and custom setups gone overnight, with no response after several attempts to reach a human reviewer.

This is the sharpest version of the case made earlier in this chapter for a second AI subscription: not just a second opinion, but a working account you can still open the day your primary one goes dark. Export or back up anything you would miss, a finished document, a useful custom instruction, a project you rely on, the same way you would back up a hard drive, rather than trusting your only copy to an account you do not control.

If you choose Claude as your second AI and you want to learn Claude properly rather than just hope that all its features work in the same way that ChatGPT does, the companion *Claude Visual Bible* walks through it with the same four-stage structure as this volume.

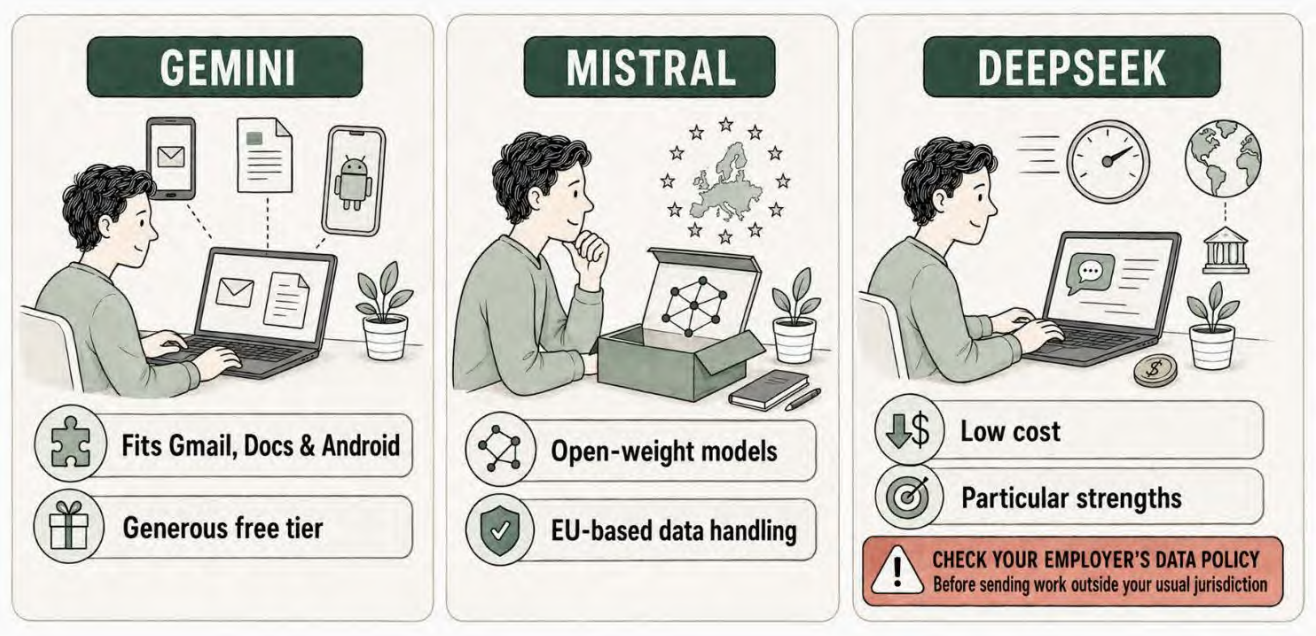


Figure E2: Choosing a third AI provider

A third AI is even more optional, and worth adding only once your first two are earning their keep. Google's Gemini is the natural choice if you already live inside Gmail, Docs, or Android, and its free tier is generous, so you can try it alongside your paid subscriptions without cost. Mistral is worth a look if open-weight models or EU-based data handling matter to your work. And one of the fast-improving Chinese providers, such as DeepSeek, Qwen, or Kimi, can be worth a try for their low cost and particular strengths with coding, though check your employer's data policy, the same way you checked it in *Chapter 2.1*, before you send anything work-related through a platform hosted outside your usual jurisdiction.



Learn more online: I have also written three bonus books, which you can find online:

Beyond Your First AI (`https://github.com/markjprice/chatgpt-vb/blob/main/docs/bonus/beyond/README.md`): This bonus book explains when a different AI service is the better tool for a specific job, or a backup when your primary AI is unavailable. It covers Anthropic Claude, Google Gemini, Mistral Vibe, and some Chinese providers like DeepSeek and Qwen.

Profession-Specific Prompts (`https://github.com/markjprice/chatgpt-vb/blob/main/docs/bonus/professions/README.md`): This bonus book adds a layer that the general prompting skills taught in the print books do not cover: the specific duties, risks, and rules that come with 20 professions.

Local AI with gpt-oss (`https://github.com/markjprice/chatgpt-vb/blob/main/docs/bonus/gpt-oss/README.md`): The print books taught you to work with ChatGPT running in OpenAI's cloud. This bonus book covers running an AI model entirely on your own laptop or desktop, with nothing leaving your device, using OpenAI's two free and open-weight models, `gpt-oss-20b` and `gpt-oss-120b`.

What to do if you want to build, not just direct

In *Book 3*, you built automations that pushed right up against the edge of this book's scope. You customized GPTs and built skills, scheduled tasks to run automatically, and maybe wondered what it would take to write the code behind an automation yourself, instead of only directing ChatGPT to configure one. That is the boundary between this volume and the next one in the series.

ChatGPT Codex Visual Bible picks up exactly where the automation chapters left off. It teaches you to write and edit working code, work in a repository, run tests, use Git to track your changes, and build websites and mobile apps. It assumes the same starting point as this volume: no coding background, ordinary computer skills, and a willingness to follow clear, visual steps.

You do not need to be planning a career in software to get value from it. Many readers who reach this point just want to understand what their automations are actually doing under the hood, or to build one small tool that solves a problem no existing skill or plugin quite covers. If that describes you, *ChatGPT Codex Visual Bible* is the next book in this series worth picking up.

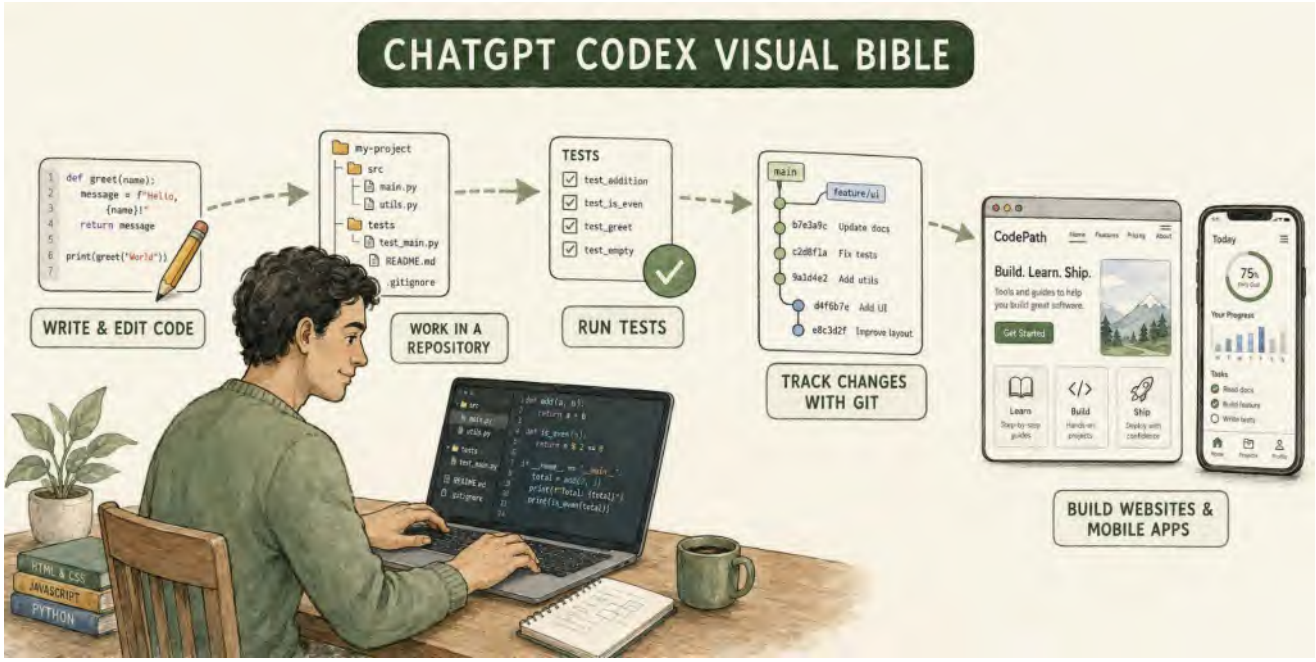


Figure E3: ChatGPT Codex Visual Bible topics

A last word

Four books and forty-two chapters is a lot of material to work through, especially if you diligently completed all the **Try it now** sections, and you likely skipped some sections that did not apply to you yet. Keep this volume where you can reach it. The chapter you skimmed on analyzing spreadsheet data or building custom GPTs may be exactly the one you need in six months.

This book was written to close the gap between using ChatGPT and getting useful results from it. What you do with that next is up to you. Every chapter in this book asked you to check the result before you trusted it. That habit is the lasting skill. Everything else was ChatGPT.

Go use it.

Share your thoughts

Now you've finished *ChatGPT Visual Bible*, we'd love to hear your thoughts! Scan the QR code below to go straight to the Amazon review page for this book and share your feedback or leave a review on the site that you purchased it from.



`https://packt.link/r/180865515X`

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.