

ADVANCED ChatGPT MASTERING PROMPTS

Get better answers, build smarter workflows, and unlock AI's full potential

Expert techniques
to unlock your
creativity

Tailor
responses
to any
situation

Write reliable
scripts and
shortcuts

Build and
manage
complex
workflows

Let AI help
you tell your
own story

Gain new skills
with personal
learning plans

132
PAGES OF
EXPERT
ADVICE

VIBE CODING:
develop with intent and creativity


WRITTEN *by* HUMANS



for HUMANS

AN INDEPENDENT PUBLICATION

ANTHEM
PUBLISHING



Advanced
ChatGPT –
Mastering
Prompts
£9.99



ADVANCED
ChatGPT
MASTERING PROMPTS

Anthem Publishing Ltd,
Media House,
1 Widcombe Parade, Bath,
BA2 4JT, UK
shop@anthem.co.uk
01371 853609

EDITOR **Darien Graham-Smith**

ART EDITOR **Andrew Sumner**

PRODUCTION EDITOR **Priti Patel**

PUBLISHER/MD **Sally FitzGerald**
sally.fitzgerald@anthem.co.uk

HEAD OF ADVERTISING
Megan Gibbings
megan.gibbings@anthem.co.uk

ADVERTISING MANAGER
Charlie Hallworth
charlie.hallworth@anthem.co.uk

HEAD OF MARKETING
Cara Northcott

DIRECT MARKETING MANAGER
Rosie Pankhurst

LICENSING **Regina Erak**
regina.erak@globalworks.co.uk

CHIEF CREATIVE OFFICER **Jenny Cook**

CEO **Jon Bickley**

PRINT
William Gibbons & Sons Ltd
Tel +44 (0) 1902 730011

DISTRIBUTION
Marketforce UK
121-141 Westbourne Terrace,
London W2 6QA

ALL IMAGES
© iStock and ChatGPT
Cover image © iStock

ANTHEM | PUBLISHING PPA

Established 2002. Founding Directors:
Paul Pettengale, Simon Lewis, Jon Bickley

The content in this publication was created by Catalysed Productions Ltd. All content copyright Anthem Publishing Ltd, 2026, all rights reserved. While we make every effort to ensure that the factual content of *Advanced ChatGPT: Mastering Prompts* is correct, we cannot take any responsibility, nor be held accountable, for any factual errors printed. No part of this publication may be reproduced, stored in a retrieval system or resold without the prior consent of Anthem Publishing Ltd. Anthem Publishing recognises all copyrights contained within this issue. Where possible, we acknowledge the copyright holder. *Advanced ChatGPT: Mastering Prompts* is a trademark owned by Anthem Publishing.

If you're ready to step up a gear with what you can do with AI, dive in!



Short story summary.txt
Plain text

Here's the summary of my story. Expand it into scenes, and generate bullet points detailing what happens in each scene. Do not introduce any new ideas.

Below is an expansion of your summary into structured scenes, keeping strictly to the ideas already present in your outline. No new plot elements have been introduced. Short story summary.txt

Scene 1: Opening of the Trial

- The courtroom atmosphere is tense and formal.
- The defendant sits under scrutiny as the charge of high-stakes robbery is restated.

Welcome



ChatGPT is an amazing service that can work wonders in both your personal and professional life. But it only does what you tell it to – and to get the best results, you need to know how to talk to it.

That's exactly what we'll teach you in this magazine. In the sections that follow, we'll show you the prompting techniques that can get the AI to produce accurate, reliable and genuinely useful results – while avoiding hallucinations, overconfident answers, and subtle gotchas that can lead you astray.

We'll also look at some advanced capabilities of ChatGPT. For example, we'll demonstrate prompts that you can use to create professional reports and proposals tailored to their intended audience. We'll see how to make the AI's responses more trustworthy, with verifiable citations and references, and for those in management roles, we'll show you how prompt templates can help provide transparency and accountability for AI use.

There's plenty here for personal users, too. We'll explore how to use the AI to create personal learning and development programmes, and to achieve your creative goals – without letting ChatGPT take over ownership.

To get started, all you need is basic access to ChatGPT, via either a downloadable app or the web interface at chatgpt.com. Some techniques in this book may require a paid ChatGPT Plus account, but most will work just fine with the free service – so anyone can start improving their prompts to get better responses from ChatGPT, and discover capabilities and applications of the AI that you might never have thought possible.

Darien Graham-Smith
Editor

Contents

The fundamentals

- 06 **How to talk to an AI**
- 08 What is a prompt?
- 10 The technology behind the chat
- 12 Accuracy and error prevention
- 14 The philosophy of prompting
- 15 How phrasing affects your results **TUTORIAL**
- 16 **Prompts that get results**
- 18 Writing clear instructions
- 20 Context and external data
- 22 Personas and restrictions
- 24 The anatomy of a prompt
- 25 Prompt effectively with the right components **TUTORIAL**
- 26 **Fix and improve your results**
- 28 Refining with multiple prompts
- 30 Constraints and uncertainty
- 32 Learning from the output
- 33 Fix a bad response without starting over **TUTORIAL**

Advanced

- 34 **Business and productivity**
- 36 Writing clearly in high-stakes situations
- 38 Planning with business frameworks
- 40 Understanding customers and competitors
- 42 Turning raw input into clear actions
- 43 Transforming unsorted documents into a plan **TUTORIAL**
- 44 **Research and analysis**
- 46 Finding and organising reliable information
- 48 Developing questions and testing ideas
- 50 Making sense of data
- 52 Checking reasoning and claims
- 53 Double-check the AI's answer **TUTORIAL**

What is a prompt?

A **prompt** is the input you give to an AI system to tell it what you want it to do.

In the context of tools like ChatGPT, a prompt can be as simple as a question, or as detailed as a set of instructions.

At its simplest:

- A prompt = *your message to the AI*

Examples:

- "Summarise this article"
- "Explain quantum computing in simple terms"
- "Write a professional email declining a meeting"

Understanding how prompts work is the key to getting good results –see page 6.

Learn techniques for getting reliable answers out of the AI – see page 46.

How can AI help me find and organise reliable information?

AI can be genuinely useful here—but only if you treat it as a **structured assistant for filtering and organising information**, not as a source of truth in its own right.

There are three distinct roles where it adds real value:

1. Finding information (discovery and narrowing)

AI is good at turning a vague question into a *searchable shape*.

- It can suggest **better search queries** ("What terms would an academic paper use for this?")
- It can identify **key sources and domains** (e.g. regulators, journals, standards bodies)
- It can help you explore a topic from multiple angles quickly

ChatGPT can help with advanced reasoning tasks – see page 86.

How can AI help me solve complex problems with logic?

AI can help with complex problems—but its real strength isn't "thinking for you". It's making the structure of reasoning visible, testable, and easier to refine.

Here are the most useful ways to use it logically:

1. Breaking problems into solvable parts

Complex problems are usually hard because they're *bundled together*.

AI can:

- Decompose a problem into clear sub-problems
- Identify dependencies ("you need X before you can solve Y")
- Separate facts, assumptions, and unknowns

54 Coding and systems

- 56 "Vibe coding" with clear logic
- 58 Finding bugs and testing safely
- 60 Building structured workflows
- 62 Cleaning up and reviewing code
- 63 Inspect and improve an existing script **TUTORIAL**

64 Personal learning and development

- 66 Designing learning experiences
- 68 Testing knowledge and supporting revision
- 70 Building habits and reaching goals
- 72 Reflecting and deciding carefully
- 73 Design a personalised growth plan **TUTORIAL**

74 Writing and creativity

- 76 Building stories step by step
- 78 Editing for tone, voice and originality
- 80 Forcing genuinely new ideas
- 82 Transforming and translating information
- 84 Controlling length and level of detail
- 85 Edit with AI – without losing your own voice **TUTORIAL**

Expert

86 Solving complex problems with logic

- 88 Structuring step-by-step reasoning
- 90 Self-review and expansion
- 92 Examples and reusable templates
- 94 Logic rules and output formats
- 96 Working with images, audio and documents
- 98 Advanced reasoning
- 99 Use AI to help with a multi-step decision **TUTORIAL**

100 Ensuring quality and accuracy

- 102 Validation and stress testing
- 104 Auditing claims and citations
- 106 Managing bias and misleading framing
- 108 Managing long conversations
- 110 Prompts that produce convincing but misleading answers
- 112 Responsibility and limits
- 113 Identify and correct hallucinations **TUTORIAL**

114 Professional mastery and the future

- 116 AI and governance considerations
- 118 Building and maintaining a prompt library
- 120 Real-world case studies
- 122 Autonomous systems, and what comes next
- 123 Create an autonomous task with guardrails **TUTORIAL**

- 124 Prompt reference
- 126 Checklists and fixes
- 128 Formatting and shortcuts
- 130 Glossary



How to talk to an AI

What you get out of ChatGPT depends on what you put in. In this section, you'll learn how to phrase your prompts to generate high-quality answers, and how context and tone can affect AI's responses.

08

What is a prompt?

10

The technology behind the chat

12

Accuracy and error prevention

14

The philosophy of prompting

15

How phrasing affects your results

TUTORIAL

What is a prompt?

ChatGPT's amazing capabilities are driven by your input

E

very interaction with ChatGPT starts with you. When you first open the app (or the web

interface), you're greeted with an empty textbox and an invitation to "ask anything": the ball is in your court. Whatever you choose to enter here – either by typing it in or using ChatGPT's voice mode – starts the conversation, and sets the direction of everything that follows.

Despite the phrasing, you don't have to ask a question. You can give the AI an instruction, such as:



List some TV shows from the 1980s.

You can even treat ChatGPT like a search engine and just tap in a terse phrase such as:



UK economy 2015.

Whatever form it takes, your input is called a **prompt**, and ChatGPT will always attempt to generate a helpful response to it. If your prompt is clear and specific, ChatGPT will probably come back with exactly the information you're looking for. If it's vague, the AI engine will have to rely more on its own inferences, and its response may miss the mark. This isn't a big problem, however: you can just enter another prompt to explain yourself more clearly and continue the exchange.

This basic principle underpins the idea of "prompt engineering" – which is simply to intentionally phrase a prompt (or a series of prompts) to elicit the desired response from the AI.

A simple request can generate a deep-dive into any sort of information.

...

Tell me some interesting facts about pasta



Here are some genuinely interesting pasta facts, beyond the usual "Italy loves carbs" level.

- **There are hundreds of pasta shapes**

Italy officially recognises 300+ named shapes, but if you count regional variants and local names, the real number is probably closer to 600–700. Some differ only subtly, but even tiny changes affect how sauce clings.

Planning your goal before you type

You can chat back and forth with ChatGPT as much as you like – it's right there in the name. However, if you're hoping to get something specific from the AI, it's best to be clear about your goals.

That starts with understanding them yourself: what were you actually hoping to learn when you typed in "UK economy 2015"? ChatGPT can provide insights, data tables and even poetic descriptions of the mood of the nation – but if you don't specifically ask for those things it will default to a neutral overview of the topic, which is unlikely to be particularly enlightening.

You'll get much better results if you clearly communicate to the AI what its output should look like. For example, we could instead enter:



Were UK households getting richer or poorer in 2015?

Now the AI knows which particular sector of the economy we're interested in, and

the shape of the answer we're looking for – so it's able to generate a sharp, confident response.

Being direct about your intentions pays off on a larger scale, too. Within a conversation, the AI remembers what's gone before, and it takes this context into account when interpreting your prompts. If you start a session by asking vague questions around a topic, without a clear idea of what you're looking for, that will affect ChatGPT's framing of the issue, potentially leading to incomplete or wrongly focused answers later in the chat.

If you're hoping to get something specific from the AI... be clear about your goals

Where your role ends and the AI's begins

While it's important to be clear about what you want from the AI, there's no need to dictate the precise terms of your enquiry, or an exact structure for its response – unless it's for a particular purpose, such as filling in a form or processing data for a spreadsheet.

Indeed, ChatGPT is often at its best when you let it come up with helpful angles and formulations of its own. After all, this is what it's designed for: finding information that connects to your question, and assembling it into a coherent response.

What you can't trust the AI to do is understand the real-world implications of its responses. It has no innate sense of importance or stakes, and can only discern between truth and fiction via probabilistic modelling.

This might not be a problem for exploring a topic, but for productive tasks make your prompt clear and specific. For example:

.....
Suggest three improvements to this report to help non-expert readers appreciate the urgency of moving to renewable energy.

It's a little wordy, but if you just ask ChatGPT to "improve this document" it could well end up focusing more on grammar and tone than the effectiveness of the central argument. If you don't know what your priorities are, the AI can help – just ask it to identify the different threads of argument in an uploaded document, then tell it which ones to focus on.

Use the AI as a collaborator, not a tool

ChatGPT won't do anything without your input, but it isn't a mechanistic device that you can control directly. AI engines by their nature aren't perfectly predictable or consistent; there's a reason why we talk about "prompts" rather than "commands".

Of course, you can use ChatGPT for simple "one-shot" operations, such as getting answers to factual questions. But when you use it in this basic mode, you're missing out on the key feature that makes AI so much more powerful than a traditional web search, which is its

Tell the AI what you want and it will follow your instructions to the letter.

Suggest three improvements to this report to help non-expert readers appreciate the urgency of moving to renewable energy

Here are three improvements that would help non-expert readers feel the urgency rather than just register the information.

1. **Surface near-term consequences, not just long-term goals**
 The report explains why current trajectories miss 2040–2050 targets, but urgency weakens because the consequences feel distant. Pull forward what delay means in the next 2–5 years: higher household bills, continued exposure to volatile gas prices, infrastructure approvals that lock in emissions for decades. When readers see that today's decisions shape outcomes within a single political or investment cycle, urgency becomes immediate.

Use ChatGPT as a collaborator – don't just delegate the whole job to it.

What are three surprising ways in which Hector could escape?

Here are three surprising, story-driven escape routes for Hector — none of them relying on force, clever gadgets, or last-minute luck.

1. **He escapes by complying too precisely**
 Hector follows every rule with obsessive accuracy: forms completed perfectly, procedures obeyed in full, instructions carried out without initiative. The system, which depends on shortcuts and assumptions, starts routing him forward automatically. Each person assumes someone else has authorised him. By the time a decision would need to be made, there's nothing left to decide — he's already gone.

unsatisfying, for several reasons. Instead, try asking:

.....
What are three surprising ways in which Hector could escape?

The prompt doesn't just pose a question: it casts ChatGPT in the role of a collaborator. The AI generates more ideas than you'd get from a simple tool-type instruction, while leaving you in control of the final decision and execution. 🧠

How tokens and context shape the conversation

Why earlier instructions lose influence over time

ChatGPT never intentionally discards context, but as the number of active tokens in a conversation grows, the focus can drift away from your initial requirements – which is another reason why it doesn't hurt to restate your purpose every so often.

Context is also lost when the number of tokens in the conversation exceeds the size of the context window, as the oldest tokens are jettisoned to make space for newer ones. Paid users probably won't experience this very often: while OpenAI doesn't make any promises about conversation scope, at the time of writing, subscribers using the latest GPT-5.2 engine can get a context window of up to 196,000 tokens per conversation – enough to contain the entire text of *A Tale of Two Cities*.

Free users get a smaller window, reportedly ranging from 16,000 to 32,000 tokens, depending on how busy the system is. For them, it's possible that the AI could start to forget context during a longer discussion. You won't get any warning, but you might notice that when you refer to earlier topics, the AI responds as if it's hearing about them for the first time – or it might “hallucinate” something based on the remaining context (we'll talk more about hallucinations overleaf).

Note that you can quickly fill up your context window by pasting large amounts of text into the AI. If you want to work with an external document, upload the file instead; this allows the AI to internally break up the content into chunks that can be moved in and out of the context window as needed.

Paying users now get up to 196,000 tokens – a big boost in context.

What's the difference in the size of the context window between free and paying ChatGPT users?

Here's how the context window size differs between Free and paid ChatGPT users today:

Free users (using GPT-5.2 Instant):

- Typically have a 16 K token context window — meaning the model can consider about 16,000 tokens of combined input and output at once. ChatGPT

Plus users (paid \$20/month tier):

- Get a 32 K token context window for most models in ChatGPT — roughly twice the Free tier's size. ChatGPT

That difference means:

- A Free chat starts trimming earlier once the conversation grows long.

What the AI does remember

Every conversation with ChatGPT has its own context window, and the things you say in one thread won't affect the AI's behaviour in a different one... with a few exceptions.

Open up the ChatGPT Settings interface and you'll find a section entitled Personalization, with a sub-section called Memory. Go here and click “Manage memories” to see a list of facts that ChatGPT has learnt and stored about you, such as where you live


and what sort of interactions you prefer. Early on in every new conversation, ChatGPT reviews these memories and imports any tokens that appear relevant, to make its own responses more useful and personal to you.

If the list includes any incorrect inferences, you can remove them by clicking the bin icon next to each one, and you can add your own memories via the chat interface – just enter a prompt such as:

Always use British English when proofreading.

The AI should respond by saying “Updated saved memory” – if it doesn't, tell ChatGPT explicitly that you want the fact saved for all time.

Memories are a powerful tool for directing the AI's responses without having to repeatedly enter long, prescriptive prompts. And because they're stored outside of the chat environment, memories never expire – although they may be automatically removed by the system if they're superseded by new instructions. OpenAI



Accuracy and error prevention

The way you talk to the AI has a big effect on reliability

A

I has a bad reputation for “making things up”. In a sense, this is entirely fair: as we’ve seen, everything it tells you is effectively

a probabilistic inference. ChatGPT has no way of determining whether the answers it gives you are objectively true, and it won’t try to actively fact-check them.

In practice, though, it’s usually very reliable. If a claim or connection

crops up over and over again throughout the AI’s training data – and is supported by high-profile web sources, if an online search has been triggered – then it’s likely to be accurate. While there may also be mistakes and misinformation out there, they usually aren’t consistent or widespread enough to outweigh the truth.

However, AI engines can sometimes infer connections that aren’t there, and “hallucinate” answers that sound authoritative but have no basis in reality. For example, you might encounter a hallucination if you ask a question such as:

Who was the first person to use a smartphone?

As always, ChatGPT AI tries to compute a good response, but in this case the terms are vague and there’s no widespread consensus on the answer. The best it can do is piece together scattered bits of context and create something that looks like it ought to be the answer. And this shouldn’t shock us too much, because that’s effectively what it does with every query.

Short answer: we don’t know their name.



What we do know

The first widely recognised smartphone was the IBM Simon Personal Communicator, released commercially in 1994.

It combined:

- cellular phone calls,
- a touchscreen with stylus,
- email and fax,

Hallucinations can arise when the AI doesn’t have full access to the facts.

How wording and framing influence accuracy

We’re not saying that the user is to blame for hallucinations, but certain phrasings make them more likely. In the example above, we framed our question in a way that implied that there was one definitive answer; ChatGPT knows from its training that such questions call for a simple, confident response, so that’s what it generates – even at the expense of accuracy. This is a risk with any sort of prompt that asks for a singular fact, such as “what’s the reason for...?” or “what is the definition of...?”

You can also inadvertently influence the AI by inviting a narrative response. Its training data contains lots of stories and chains of causality, and these patterns kick in when you ask a question such as:

Why did Katy Perry divorce Russell Brand?

Statistically speaking, the AI “knows” that the answer to this question ought

to be a sequence of interrelated events, and it does its best to assemble one. Again, the overall form of the answer weighs more heavily than the strict accuracy of each piece of content.

The important point is that the AI doesn’t just hallucinate whenever there’s a gap in its knowledge. It does so when the phrasing of your prompt pushes it to give an answer that follows a certain format or tone, which it’s only able to do by including low-confidence inferences.

Spotting hallucinations

Hallucinations are tricky things: since they're assembled by pattern-matching, they generally look like a perfect answer to your question, and they're usually delivered with confidence.

In fact, what often gives a hallucination away is the fact that it's too neat. As we've observed, ChatGPT isn't designed to mindlessly obey your commands: more often than not, it interprets your prompts, talks around the subject, and offers additional information and perspectives. When we tried asking it for a list of Roman emperors, it took the initiative to add qualifications, caveats and details of the different dynasties.

So if the AI engine gives you a response that's precisely tailored to match what you asked for, it's actually a potential red flag – it suggests that ChatGPT may have prioritised linguistic structure over engagement.

A high degree of confidence is another warning sign. If you ask a question expecting a vague or tentative answer, and ChatGPT instead comes back with an emphatic affirmation of fact, consider challenging it. A related tell is unexpected precision, such as if a response includes suspiciously specific names and dates that you wouldn't expect the AI to know. This sort of hallucination occurs when ChatGPT is following an answer pattern that usually includes facts and figures, but doesn't have the relevant details available in its training data. Rather than deviating from the statistical template, the AI engine brings in content from elsewhere to fit the template, even though it doesn't actually apply.

Julio-Claudian dynasty (27 BC–AD 68)

- Augustus
- Tiberius
- Caligula
- Claudius
- Nero

A sign of confidence is when the AI adds in relevant detail you didn't ask for.

Year of the Four Emperors (AD 69)

- Galba
- Otho
- Vitellius
- Vespasian

Flavian dynasty (AD 69–96)

- Vespasian
- Titus

Based only on what Katy Perry and Russell Brand themselves have said publicly (interviews, first-person comments, and Perry's documentary), the picture is consistent and fairly narrow.

What Katy Perry has said

Giving the AI a clearer, narrower scope tends to get more reliable results.

Across interviews and her 2012 documentary *Part of Me*, Perry describes:

- **Clashing priorities and timing:** her career was in an intense, all-consuming phase; she wasn't ready to slow down or reshape her life around the marriage.
- **Independence and control tensions:** she has said Brand initially supported her success but later struggled with her independence and leadership as her career accelerated.
- **Abrupt breakdown in communication:** once Brand decided to end the marriage, communication effectively stopped.
- **How the divorce was delivered:** she has repeatedly said Brand informed her by text message while she was on tour, shortly

Prompts that discourage invention

We've talked about what causes hallucinations; you're doubtless now wondering how to avoid them. A good prompting strategy is to address the AI as a curator of knowledge, rather than a source of information itself. For example, instead of asking who the first person was to use a smartphone, try:

What is known about the first smartphone user?

ChatGPT's AI engine is strongly optimised to always be helpful, and it hates to say "I don't know" – that's why it sometimes ends up hallucinating instead. But it will happily tell you that a particular answer is unknown to historians, or not recorded online, without feeling the need to fill in the gaps with invented details.

Address the AI as a curator of knowledge, rather than a source of info...

Another simple technique is to explicitly tell the AI to stick to the facts:

Tell me why Katy Perry and Russell Brand divorced, based only on primary sources.

In reality, ChatGPT's "primary sources" are still based on its internal training data – but the constraint weakens the AI's impulse to follow an established template, making it a lot less likely to sneak in any bogus details of its own.

To discourage hallucinations over the longer term, you can use the memory function we saw on page 11. Just enter a prompt such as:

Always tell me how confident you are in any response, and warn me if a statement depends on inference or pattern matching, rather than primary sources.

From now on, the AI should actively warn you if there's a risk of hallucination. Although memories are activated selectively, words such as "always" and "any response" ensure this memory will apply in almost every situation.

The philosophy of prompting

Align your input with the AI's strengths and weaknesses

W

e've seen in this section how prompting is more nuanced than it might appear. It isn't

the same as typing a search term, issuing a command, or even asking a question. That's because the AI isn't slavishly following your instructions, or looking up keywords in a database – it's doing something far more ambitious. In human terms, it's trying to infer the intent behind your entry, and generate a response that addresses it as helpfully as possible.

This means the phrasing of your prompt is hugely important; it isn't necessarily an exaggeration to say that the way you talk to ChatGPT matters just as much as the topic. If your tone or choice of words suggest to the AI that you're looking for a particular type of answer, it will do whatever it takes to generate a suitable response, including inventing imaginary people and events to fit the format.

Prompting effectively is about formulating your requests so that the AI's idea of what you want lines up, as far as possible, with what you actually do want. As a start, that might mean avoiding short, closed questions that imply a definitive answer, and instead asking for options, overviews or feedback.

Of course, there's more to it than that – we're only at section 1, after all. Throughout this guide, we'll look more closely at how to take advantage of ChatGPT's multiple personas, use multiple prompts to improve results, and work with errors and uncertainty.

If you ask the AI for one thing, it will give you one thing – for better or worse.

...

What is the single best thing about Spain?



The single best thing about Spain is how effortlessly life happens in public.

People don't retreat indoors at the end of the day. Streets, plazas, cafes, and bars become shared living rooms. Children play, grandparents sit and watch, friends eat slowly, conversations overlap. Time flexes around people rather than the other way round.

Food is part of it. Weather helps. History adds texture. But the real magic is cultural: pleasure and sociability aren't treated as indulgences – they're normal, everyday expectations. You don't need a special occasion to sit outside at 10pm, order something small, and talk until it's finished.

When good enough is good enough

A final point to bear in mind is that ChatGPT is tirelessly helpful. No matter what you enter, it will always be ready to come up with a constructive response – often laced with praise and encouragement, to make you feel good about using the AI.

Obviously this is preferable to a moody AI that gets bored after a few exchanges and refuses to help any more – but it means that ChatGPT conversations don't ever reach an organic conclusion. If you pose a factual question, the AI will volunteer possible follow-ups; if you ask it to plan out a schedule for a camping trip, it will offer to help you pack. The interaction is designed to keep helping you, in new and different ways, forever.

You'll notice a similar dynamic if you ask it to edit or proof your documents. On its first pass, the AI comes up with plenty of clever suggestions – and if you make the changes and feed in the document again, it will come up with further suggestions, or refinements, or footnotes, or alternate versions.

In fact, the one thing that ChatGPT won't ever suggest – at least not off its own bat – is drawing a line and saying "enough". This means that part of the secret of effective prompting is knowing when to stop prompting. We suggested on page 9 that you could think of the AI as a collaborator, and yourself as the decision maker; in particular, you'll have to decide when the conversation is over and it's time to move on. 

How phrasing affects your results

ChatGPT is a powerful tool for research and education, but its answers can vary depending on how you ask the question

Define evolution

Evolution is the process by which living organisms change over generations due to variations in inherited traits.

In biology, it explains how populations adapt over time through mechanisms such as **natural selection**, **mutation**, **genetic drift**, and **gene flow**, leading to diversity of life and the emergence of new species.

More broadly, the word can also mean **gradual change or development over time**—for example, the evolution of technology, language, or ideas.

Explain evolution

Evolution is the process by which living things change over long periods of time, generation by generation.

Individuals within a species are not all exactly the same. Some have traits that make them slightly better at surviving and reproducing in their environment. Those individuals are more likely to pass those traits on to their offspring. Over many generations, these small differences accumulate, and the population gradually changes.

The main mechanisms behind evolution are:

- **Natural selection** – traits that help survival and reproduction

1 A short, neutral prompt

Let's imagine we want to know what evolution is; we might start our inquiry by simply asking ChatGPT to define the term. Because we're using language associated with looking up a word in a dictionary, the AI responds in the same tone, generating a concise definition of the scientific theory.

2 A more descriptive verb

If we ask ChatGPT to "explain" the subject instead, we get a much more human-friendly answer. Although the core information is broadly the same, ChatGPT now explains how evolution operates in the real world, and adds helpful definitions of the different mechanisms.



Explain evolution for a sceptic

If you're sceptical of evolution, it helps to strip it back to what it **actually claims**—and what it doesn't.

Evolution does **not** say that life suddenly turns into something else, or that complex organisms appear by magic. It says something much more modest and observable:

Living things change, slightly, from generation to generation—and those changes can add up over time.

Evolution works through a small number of simple, well-understood processes acting over many generations. Nothing dramatic happens in a single step; the power comes from repetition over time.

Here's how it works, step by step.

1. Variation exists

Individuals in a population are not identical. They differ in traits such as size, strength, colour, metabolism, behaviour, and resistance to disease. These differences arise mainly from **random genetic mutations** and reshuffling of genes during reproduction.

3 Give it a focus

The facts of evolution don't change based on who you're talking to – but if you include the idea of a sceptic in your prompt, the AI makes a point of anticipating and addressing common objections. Specifying an audience is a powerful framing tool, and we'll look at it more in the next section.

4 Try a different angle

For a different perspective, try asking not what evolution is, but how it works. This formulation encourages ChatGPT to generate an answer focused on concrete processes rather than scientific concepts. The result is a step-by-step walkthrough of how evolution takes place within a population.



Prompts that get results

ChatGPT might speak our language, but it doesn't always understand you like a human would. Here, we look at the techniques and constructions that can help you communicate with the AI effectively.

18

Writing clear instructions

20

Context and external data

22

Personas and restrictions

24

The anatomy of a prompt

25

Prompt effectively with the right components

TUTORIAL

Writing clear instructions

Don't make the AI guess at your goals – tell it exactly what to do



n the previous section, we saw the importance of phrasing when talking to ChatGPT. Nuances that might

seem immaterial to you – such as your particular choice of vocabulary, or sentence structure – can have a surprising impact on the way the AI interprets your prompt.

That's because ChatGPT can't empathise with you like a human. It responds to general patterns of context and behaviour, but it can't intuit the emotion or intent behind an individual prompt. All it has to go on is the bare text, which means subtleties of phrasing can take on more significance for the AI than they would in a regular human conversation.

This is tricky to internalise, because the AI's use of natural language, coupled with its friendly tone, can really make you feel like you're communicating with another person. It's easy to fall into natural speech patterns, and find yourself entering prompts such as:

Could you give this a quick look over?

I've read the document carefully and, treating it as a real proposal, there are **numerous clear factual errors and impossibilities**, plus several claims that are technically implausible even to a non-specialist reviewer. I'll focus only on **factual accuracy**, not style or tone. 🌐

Major factual errors (unequivocal)

1. "6G wireless" as a deployable technology

6G does **not** exist as a commercial, deployable standard. It is a research concept with no:

- ratified specifications,
- enterprise hardware,
- regulatory approval,

Ask ChatGPT to fact-check a document and it will dig right in.

If you say this to a co-worker, they'll recognise that you're understating what you really want out of politeness – but ChatGPT can't be relied on to make that inference. Although it will recognise the pattern of your request, the phrasing creates uncertainty: are you actually asking about its capabilities? What is the time constraint, and what outcome are you expecting from its perusal? You're likely to get cleaner and more reliable results if you type:

Read this and tell me about any factual errors you find.

Try talking like this to your office colleagues and you might not be very popular – but for the AI this is the equivalent of peak politeness, making it easy for the engine to fulfil its purpose of generating a helpful response, without having to resort to guesses and assumptions.

Use strong action verbs

A powerful tip for communicating clearly with the AI is to use active verbs wherever possible. To see why, let's imagine that you upload a document you're working on and ask:

How is this draft?

As always, the AI will do its best to give a helpful response, but it doesn't have a lot

to go on: effectively, you've handed it a piece of content with no indication of what you want done with it. In the absence of prior instructions, ChatGPT falls back on pattern-matching behaviour, which probably means scanning your document for typos and factual errors. This isn't a bad response – but, ultimately, it's just a guess at what you want.

Switching to an active verb, such as “proofread”, “fact-check” or “punch

up”, takes away the guesswork. Not only do these words tell the AI what to do, they imply an expected scope and structure for its response. Plus, active verbs don’t only improve your immediate results; directing ChatGPT in this way enters your purpose explicitly into the ongoing context, so you can ask the AI to tweak its process, explain it or repeat it for additional documents, without getting bogged down in defining your terms.

Ordering instructions effectively

We mentioned in section 1 that ChatGPT processes text sequentially. This isn't merely an interesting fact about the way the engine is designed – it has an important implication for the way your prompts are interpreted. The AI's strictly ordered approach means that the first few words of a prompt tend to greatly influence the overall framing of the task, while later ones often serve more to refine and finesse the execution.

You can see this effect in action by uploading a document and issuing these two prompts, one after the other:

.....
Summarise this document in child-friendly language.

.....
Using child-friendly language, summarise this document.

You might expect the two constructions to yield identical results, but you'll likely find that the second summary is written in much more storybook-style language, owing to the earlier anchoring of the "child-friendly" stipulation. Understanding this effect helps you construct your prompts to ensure that the most important part is treated as such.

Sometimes, it's best to break up instructions into multiple prompts to ensure that different objectives aren't mixed up. For example, consider the following instruction:

.....
Rewrite this email to correct any factual errors and cut down the length.

The first part of your instruction sets the tone for the entire task.



Tiger conservation.rtf
rich text (RTF)

Using child-friendly language, summarise this document

Here's a **child-friendly summary** of the document, using simple language and short ideas, while keeping the facts right. 🌐

Tigers and why they need our help

Tigers are big, powerful cats, but they are in serious trouble. A hundred years ago, there were many more tigers in the wild. Today, about **95 out of every 100 tigers are gone**.

ChatGPT will probably prioritise our goals according to the order in which they're specified – but that doesn't mean it will do the error-checking first, then move on to shortening the text. Since the two directives are parts of the same "rewrite" task, they'll actually be processed simultaneously, as a single blended operation. The resulting edit broadly meets our requirements, but the reasoning behind it is obscured, with no way to tell whether a detail was removed because it was erroneous, or just to save words.

A cleaner approach is to split the operation into two separate prompts. Note that this doesn't just mean dividing it into two sentences – multi-sentence prompts, and even entire paragraphs, are still handled as a single instruction. To keep things separate you need to enter one unambiguous prompt, let ChatGPT respond to it, then enter a second.

Phrase prompts so that the non-negotiables are nailed down...

Focused ambiguity

It's easy to see how combining two goals in one task creates ambiguity – but, in fact, almost every instruction you give to ChatGPT is ambiguous to some degree. Consider this prompt:

.....
How can I improve my writing?

The simplicity of the phrasing conceals a lot of uncertainty. What do you think is wrong with your writing now? Are you looking for big-picture advice on pacing and structure, or help with technical issues such as spelling and grammar? Since you haven't specified, ChatGPT has to decide for itself what sort of answer is likely to be most appropriate.

Thankfully, this is a task it's extremely good at. The AI's apparent ability to read your mind, and come up with relevant information and perspectives that you didn't explicitly ask for, is a huge part of what makes it so useful.

The trick, therefore, isn't to try to get rid of ambiguity altogether. Rather, aim to phrase your prompts so the non-negotiable parts are nailed down, while areas where you want ChatGPT to use its own judgment are left open-ended. 🌐

Here's a **shortened, cleaner version** that corrects overstatement, tightens the facts, and keeps the pressure on — without escalating the tone unnecessarily.

Email



Subject Formal complaint: lack of staffing at entrance and exit

Dear Gym Management Team,

I am writing to raise a formal complaint about the ongoing lack of staff presence at the front desk, which is increasingly affecting my ability to use the gym safely and efficiently.

You can tell the AI to do multiple things at once, but the results may suffer.

Context and external data

Tell ChatGPT how to work with documents and shared information

W

e saw in the last section how context builds up over the course of an interaction. If you ask a question

right at the start of a conversation, you'll usually get a more generic answer than if you chat to the AI about your focus and purposes first.

Unfortunately, ChatGPT doesn't encourage you to feed in background information in this way. The engine is heavily weighted towards interpreting every input as a task requiring an output, so if you enter some detail that's intended to lay the groundwork for an upcoming request, you'll likely receive unwanted suggestions or analysis in response.

You can counter this impulse by explicitly telling the AI how to interpret your prompt, using language such as:



For context: I need to give a presentation tomorrow.

This tells ChatGPT that you don't want it to instantly start churning out a script. Instead, it will probably tell you some ways it can help, potentially ask for more detail, and confirm how this information will inform its responses to subsequent prompts.

You can also use context to tell the AI what not to do. For example, you could follow up with another prompt explaining:

...ask a question at the start of the conversation, and you'll usually get a more generic answer than if you chat to the AI about your focus first

Got it 👍

What's the presentation about, and what kind of help do you want right now?

For example, I can help you with:

- shaping the structure (opening → core points → close),
- tightening or stress-testing your key argument,
- writing a strong intro or conclusion,
- simplifying something technical for a mixed audience,
- slide-by-slide logic (what belongs on slides vs what you say),
- or delivery (timing, nerves, Q&A handling).

Tell me:

1. topic
2. audience (who + how expert)

For context: I need to give a presentation tomorrow

Setting the context puts ChatGPT into the right "frame of mind" for a task.



I've already written the script.

There's no need to repeat that this is context – the flow of the conversation tells the AI that this information follows on from the previous prompt. Your input will be understood as a standing constraint, so ChatGPT won't respond to future prompts by suggesting topics or generating new content, unless you ask it to.

Using examples effectively

As an alternative to typing in explanatory prompts, you can feed ChatGPT context by giving it a model of the sort of output you're looking for. This is effective for shaping not just its behaviour, but also the tone of its responses. For example, if you're using the AI to help you produce a report, you could type:



Use this text as a style model: ...

Then paste in a paragraph or two from an existing document.

From this point on you can expect the AI's editing suggestions to follow the tone of your supplied text. ChatGPT is spookily proficient at stylistic imitation, including subtleties such as the reader's assumed level of knowledge. It can even pick up on the apparent intent of a

passage, to determine whether your own text should make assertive claims, or remain cautiously neutral.

You should specify that your pasted text is intended solely as a style reference. Otherwise, the AI won't just aim to match the tone of the example passage, it may try to transplant elements of vocabulary and phrasing directly into your own work, which isn't what you want.

Working with uploaded files

You're probably familiar with the idea of uploading documents so that ChatGPT can process their content. For example, you might upload your CV and ask:

.....
How can I tailor this for a data analyst role?

What's less obvious is that you can also upload files to be mined for context. This works in the same way as pasting information into a chat: in the example opposite, we pasted a paragraph from a report and asked the AI to imitate its style, but we could alternatively have uploaded the entire report, giving ChatGPT a far broader model to work from.

Uploading files isn't quite as instant as pasting text, but it's more convenient for long documents. It also lets you work with formats that aren't normally amenable to copy and paste, including PDFs, Excel spreadsheets, and PowerPoint presentations.

Another benefit is operational clarity. ChatGPT won't ever mistake an uploaded file for an instruction or an active part of the conversation: if you upload a file with no prompt attached (and no prior context telling the AI what to do), it won't usually try to speculatively critique or rewrite it. It will just give you a simple acknowledgement, such as "I've received the file – let me know what you'd like me to do with it".

Even so, it makes sense to provide a prompt along with an uploaded file, so that both you and ChatGPT are clear on its role. For example, you might upload a PDF file and enter:

.....
Here's how the company was doing last year.

You can borrow the style of a document to generate a wholly new one.

Use this text as a style model: "This Contract sets out the entire agreement between the Parties. It replaces all previous negotiations, agreements, understandings and representations between the Parties, whether oral or in writing."

Got it. I'll treat that as a **style model**, not something to rewrite or paraphrase.

To be explicit, the style you're pointing to is:

- **Formal, legal-register English**
- **Declarative and exhaustive** (nothing implied, everything stated)
- **Impersonal** (no "we", no rhetorical flourishes)

Upload background information to help the AI make sense of your requests.



Company_Report_FY24_25.pdf
PDF

Here's how the company was doing last year

Got it — I've read the FY 2024–25 report 🌐.

Baseline snapshot (what last year tells us):

- **Growth:** Revenue £14.2m (+9% YoY); operating profit £3.1m (+12%); EBITDA margin up to 22%.
- **Efficiency:** +15% operational efficiency, largely from infrastructure upgrades and legacy cost reduction.
- **Investment:** R&D £1.8m (+15%), signalling a deliberate capability build.
- **Customers:** Retention at 94.2% (+2.1%), suggesting product stickiness.

When uploading contextual information you shouldn't normally need to give any more guidance than this. ChatGPT can parse documents, spreadsheets and presentations just as proficiently as it interprets your prompts, and is quite capable of understanding both the specifics of the report and its overall drift.

If you want to know exactly what the AI has taken from a particular upload, you can ask it:

.....
What context have you extracted from this file?

The response will probably be a long list of inferences gleaned from the document, ranging from details of authorship and purpose, to stylistic and structural characteristics. Any or all of

these properties can be used as models for generating or refining new documents – and, of course, you can bring in content at the same time. Just enter a prompt such as:

.....
Rewrite my speech to focus more on the business goals described in the report, and include a mention of the annual growth figures.

We've just one caveat here: ChatGPT assigns a high degree of authority to external documents. This means you can be confident that any imported facts and figures will be faithfully reproduced, exactly as they appear in the original file. However, it also means that any mistakes are likely to be repeated without question, too. 🤖

Personas and restrictions

Want a new angle on a topic? Get ChatGPT to play a role

Act as a project manager and help me pass my Grade 3 piano exam

Love this brief. I'll act as your **project manager**, not a piano teacher: clear scope, milestones, risks, and a delivery plan aimed at **passing**, not noodling 🎵

Below is a **Grade 3 Piano Exam Project Plan** you can start immediately. I'll keep it practical and exam-focused.

Project definition

Goal: Pass ABRSM Grade 3 Piano
Success criteria:

Assigning a role can dramatically affect the style of advice offered by the AI.

I

n the previous section, we touched on the idea of phrasing your requests so as to cast ChatGPT

in different roles. This doesn't have to be done delicately: with AI, the direct approach is usually best. ChatGPT will happily accept a prompt such as:

Act as a project manager and help me pass my Grade 3 piano exam.

The role you give the AI can significantly influence its responses. This prompt produces a plan that's all about timelines, priorities and milestones; but if you were to ask ChatGPT to act as a university lecturer, the focus of its response would be on developing your musical skills and knowledge in accordance with the curriculum.

You can invite ChatGPT to play any role imaginable: "teacher", "sceptic", "layperson", "potential

customer" and "devil's advocate" can all be useful for various purposes. You can even ask the AI to impersonate a specific real person, although its efforts aren't always very convincing.

Note that a persona isn't the same as an audience. When we previously asked the AI to "explain evolution for a sceptic", it gave an explanation designed to persuade a doubtful reader. If you ask it to "act as a sceptic and explain evolution", you'll get a much more forthright response that focuses on objections to the theory (while acknowledging the deficiencies of each one).

It's also important to be aware that the AI doesn't take on a consistent persona in the way a human might adopt a character. Asking ChatGPT to play a role simply brings a particular set of tokens into the conversation. These will often have an immediate and striking effect, but as the discussion carries on, they're likely to be watered down by other incoming ideas. You may need to periodically remind the AI to stay in character.

Define what not to do

Another way to shape the AI's response is to tell it who you don't want it to be – or, at least, to set out particular characteristics and behaviours to avoid. This isn't quite as easy as ascribing a persona, because your instruction has to be more specific – but it lets you precisely customise ChatGPT's behaviour.

Prohibitions can be declared in natural language, using a construction such as:

Flag up any problematic parts of this proposal without suggesting changes.

You can also establish an ongoing rule by entering a negative instruction on its own, such as "don't use technical terms" or "don't invent examples". This latter prompt tells the AI not to use illustrative language, which can result in rather dry responses; but it also reduces the risk of hallucination.

As with a persona, instructions entered in this way tend to become less potent across multiple exchanges, and won't apply at all outside of that conversation. If you want to globally restrain some aspect of ChatGPT's behaviour, you can use the memory function to set up a prohibition, such as:

Remember that I never want you to translate foreign-language content unless I explicitly ask for it

If you want to restrain some aspect of the AI's behaviour, set up a prohibition...

Preventing scope creep

As we've seen, personas can be a useful way to bring a particular set of expectations into a conversation. However, those can include aspects that you don't want or need. For example, when we asked the AI to act as a project manager to help us prepare for a piano exam, that made it think about schedules and deliverables – but it also introduced ideas such as dependencies, budgets and failure risks that weren't relevant.

ChatGPT is usually smart enough not to shoe-horn unnecessary considerations into a task, but those tokens may still have an influence on the conversation. If you went on to ask ChatGPT whether you ought to book additional lessons in the run up to the exam, it would likely respond with a cost-benefit analysis, rather than focusing on your personal confidence and preparedness.

This phenomenon is known as “scope creep”, when unintended implications become embedded in a conversation and lead ChatGPT away from your original intention. You can reduce its effect by using explicit constraints and prohibitions to limit the AI's purview, as we saw on the left. However, as the conversation goes on, and accrues more and more tokens, it becomes progressively more difficult to avoid.

Eventually, scope creep can cause ChatGPT to produce answers that are only partly relevant to your question,

Assumptions about your goal

1. Primary goal is delivery, not exploration

You're aiming to *get something over the line* (pass an exam, give a speech, present a strategy), not to learn broadly or indefinitely.

2. Success is defined pragmatically, not aspirationally

“Good enough to pass / convince / land” matters more than excellence for its own sake.

3. You value structure and clarity over motivation or encouragement

You're looking for analysis, framing, and risk identification, not hype, cheerleading, or lifestyle coaching.

4. You are stress-testing thinking, not seeking validation

Many of your prompts invite critique, constraint analysis, or exposure of weaknesses rather than reassurance.

Ask ChatGPT what it assumes about you – the answer may be eye-opening.

or to respond from unexpected (and unhelpful) perspectives. At that point you might be tempted to ask, “why on earth did you think I wanted you to do that?”; a more AI-friendly phrasing would be:

.....
List your current assumptions about my goal and constraints.

This should produce a summary of ChatGPT's current operating assumptions, enabling you to spot and countermand any that aren't in line with

your goals. Or, you might prefer to simply start over: there's no limit to the number of conversations you can have open, and a fresh, unprimed AI could provide the clarity you're looking for.

Using personas without overreach

As we've seen, personas are a powerful tool, but they come with baggage. If you just want a quick dose of perspective, you can reduce the contextual impact by asking the AI to temporarily “put on a hat”, rather than fully embracing a role. For example:

.....
Look over this spreadsheet as if you were an urban planner.

This phrasing conveys to ChatGPT that the persona is associated specifically with this task. It won't completely forget about the role once it's given its response, but it won't assume you want it to embody the persona going forward.

This technique also makes it easier to try out different personas and explore a question from multiple angles. Remember, once a piece of context comes into a conversation, it never fully leaves (until the associated tokens leave the context window); and when ChatGPT takes on a role it's difficult to get it to fully set it aside. However, if you assign a persona with a light, temporary framing, it shouldn't noticeably interfere with subsequent interactions. 

signals, and what the data implies for place-making and network design — not on spreadsheet hygiene or stats technique. This is based entirely on the contents of the file .

1. What the dataset represents (planner's reading)

The spreadsheet is not raw traffic data; it is a **policy-era comparison table**. Each row corresponds to a **distinct intervention regime** at Mill Road Bridge, rather than a continuous time series.

From a planning perspective, this makes it suitable for:

- before/after policy comparison,
- narrative justification of an intervention,
- strategic evaluation rather than operational modelling.

It is **not** suitable on its own for:

- capacity modelling,

The AI can switch between roles to give different takes on the same data.

The anatomy of a prompt

A few structural rules can help you achieve consistent results

Here's a contained summary of attendance over the most recent five-year period shown in the spreadsheet, treating that window as complete in itself and not drawing on any earlier data 🌐.

Five-year attendance summary

- **Overall pattern:** Attendance drops sharply at the start of the period, then recovers year-on-year, ending the five-year window at its highest level within that range.
- **Lowest point:** The first year in the five-year span is a clear outlier, with substantially lower attendance than all subsequent years.
- **Recovery phase:** The following two years show a strong rebound, with attendance increasing each year rather than plateauing.

Specifying constraints keeps the AI focused on what's important.

I

n this section, we've seen that there are a few different components that you can build into

a prompt. The most important is to specify the task itself, ideally using a clear action verb. If you don't include this, ChatGPT should still give a relevant response, but it won't have a strong sense of your intention, which could lead to ambiguities and misunderstandings later on.

Almost as important is the object of the task, such as a topic, an uploaded file, or some pasted text. If you skip this and just give a terse instruction such as "explain" or "summarise", the AI will guess at your meaning from context. Again, though, you'll be creating ambiguity as to your purpose, rather than reinforcing it.

Once ChatGPT has a clear sense of your intent, you can give it more precise direction by defining the scope of the task, and declaring any operating constraints. These two ideas are subtly different: scope

tells ChatGPT which aspects of an idea, document or problem to work on; constraints govern what it can and can't do. You can combine both in a single prompt, such as:

Summarise attendance figures for the past five years, without referring to data from outside that range.

You might also choose to specify a perspective, such as a role for the AI to play, or an audience for it to write for. As we've seen, this can have a powerful effect, but it isn't as critical as the task and the object – there's no harm in letting ChatGPT infer what sort of response you're looking for from context. You can also instruct ChatGPT to provide its answer in a particular format, such as free-flowing text, a graph, or a haiku.

A prompt doesn't normally need to contain all of these different components, but they're all at your disposal to help you get the results you're looking for.

Reusable structural patterns

One of the great things about ChatGPT is that you don't need to learn any special syntax to use it; you can chat to it just like a human. However, we've seen on these pages that it doesn't always respond to your prompts in exactly the way a person would. You may find it helpful to keep a general template in mind, which you can apply in any scenario to get effective results:

[Do this] [to this] [in this way] [and present the result like this]

As above, you don't need to include all of these elements, but if you bear this template in mind, it will help you be clear and intentional about what you're telling the AI and what you're leaving up to its probabilistic judgment. A few example prompts following this format might be:

Explain quantitative easing in layman's terms.

Rewrite that to focus on household finances.

Summarise in three bullet points.

Don't forget that you can also use ordering to emphasise a particular component, by putting it at the start of your prompt, rather than tacking it on at the end. 🗣️

If you bear this template in mind, it will help you be clear and intentional about what you're telling the AI...

Prompt effectively with the right components

Use the right prompting forms and techniques to get answers that suit your needs, or your desired communication style

How do lithium-ion batteries work?

How lithium-ion batteries work

A lithium-ion (Li-ion) battery stores and releases energy by shuttling lithium ions back and forth between two electrodes through an electrolyte. The key point: **ions move inside the battery; electrons move through your device.**

The main parts

- **Anode (-)**: usually graphite

Tell me how lithium-ion batteries work, without using technical language

A lithium-ion battery works a bit like a **shuttle system** for tiny **charged particles**, inside a sealed box.

Inside the battery are two solid sides, with a liquid layer between them. The battery doesn't make electricity itself – instead, it lets **stored energy flow out in a controlled way.**

When the battery is being **used**:

- Tiny particles move from one side of the battery to the other.
- That movement pushes electricity around the outside of the

1 A basic request

This sample prompt clearly defines a task and an object, but it doesn't tell the AI if you're interested in any particular aspect of the technology, or what tone you want it to take. Consequently, ChatGPT plays it safe and generates a dry, textbook-style answer.

2 Add a constraint

Specifying a constraint completely changes the response. ChatGPT doesn't just tweak its vocabulary, but reworks the whole explanation to suit the tone of our request. It introduces imagery to help us visualise the operation of the battery, and explains what this means for real-world operation of a device.

Imagine this.

You're holding your phone. It feels calm. Solid. Quiet. But inside it, there's a tiny, controlled migration happening every second.


That migration is what we call a **lithium-ion battery**.

Here's the big idea:

A lithium-ion battery is **not a box full of electricity**. It's more like a **reservoir and a pump** — a system for *moving energy*

3 Switch up the tone

For a different take on the topic, we've asked the AI to play the role of a TED Talk speaker. Now it sets out the facts in a way that's tailored to grab the attention of an audience – ideal for sharing with a study group, or just helping to bring the subject matter to life.



4 Specify the format

You can request output in almost any form, from a long-form dissertation to a PowerPoint presentation. Again, the AI will formulate the information to fit your requested format. It can even generate a flowchart, within the limitations of the text-based chat interface. 🔄



Fix and improve your results

Nobody's perfect, but when ChatGPT makes a mistake you can usually correct it quickly and get back to being productive. In this section, we'll see how to handle missteps – and head off errors before they happen.

28

Refining with multiple prompts

30

Constraints and uncertainty

32

Learning from the output

33

Fix a bad response without starting over

TUTORIAL

Refining with multiple prompts

Work your way towards the perfect answer

W

e've talked about how to make a prompt more effective, but you don't need to encode

the entirety of your intent into a single input. Since the AI builds up context over the course of multiple exchanges, you're free to work progressively towards a goal, rather than trying to come up with a "one-shot" instruction that does it all in one go.

Indeed, using multiple prompts can often yield better results. When you pack half-a-dozen parameters into a single input, you're dropping a big load of tokens into the conversation all at once, and there's a risk of your primary purpose being sidelined.

Moreover, as we noted on page 19, even if ChatGPT correctly reads your intention, all of the instructions and context in a single prompt are processed in a single pass. This has implications for quality: the AI can't do its best job on task X if it's trying to perform tasks Y and Z on the same material at the same time.

A better approach is to set a basic starting point, then use follow-up prompts to refine the output. This applies whether you're asking the AI to solve a maths problem, segment a data dump, or prepare you for a test. For example, if you want to create a timetable for an event, don't jump in with a long list of items and timings. Start by saying:

Create a timetable template for my one-day wellness workshop.

You can now check that the layout is right, set overall start and end

Here's a blank, one-day wellness workshop timetable template — structure only, no pre-filled sessions or examples.

One-Day Wellness Workshop – Timetable Template

Venue:

Date:

Facilitator(s):

Time

Session

Notes

Minimise mistakes by starting with a simple request and building upon it.

How does this relate to my original question?

Fair challenge — it **doesn't relate directly**, and that's on me.

Your original question was simply "**Is it accurate?**"

What I should have done was answer that **at a high level**, in the same way I answered the earlier "**Is it too long?**" question.

Instead, I jumped straight into a **deep technical audit**, which is a *different* question ("where is it wrong, exactly?").

So, reframed properly in terms of your original intent:

- **As a whole:** the report is **not accurate enough** to be signed off as-is.
- **Why that matters:** several errors are substantive enough that the document would lose credibility with a technical or procurement audience.

Never hesitate to ask the AI why it responded in the way it did.

times and so forth, before adding individual events. Additions and edits will be automatically adapted to suit the established style and format – and if you encounter a problem, you can roll back to a previous version of the timetable, rather than trying to untangle and fix a single, overloaded prompt.

Identifying weak outputs

On pages 12 and 13, we discussed how hallucinations arise when the AI can't give a confident response to your prompt, and how constraints can help you hem in speculation (we'll talk more about that overleaf).

A related problem can arise when ChatGPT misreads the intent behind one of your prompts, or focuses on the wrong thing owing to scope drift.

In these cases, the AI might produce a very confident response that looks complete and passes fact-checking, but which still isn't a good answer to your intended question.

Hopefully, the prompting techniques in this guide will help you avoid such responses, but the AI's authoritative, reassuring tone makes it easy to get blown off course without realising it. You might sometimes need to stop and ask yourself "is this actually the response I was looking for?" If it isn't, then – as a diagnostic – try asking:

.....
How does this relate to my original question?

As well as highlighting any ambiguities in your prompt, this forces ChatGPT to spell out what it thinks is the "original question" – exposing its overall framing of the conversation, and revealing any misplaced assumptions.

Walk back any misunderstandings by stating clearly what you wanted.

Give a narrow overview focusing only on the project requirements

Here's a **tight, requirements-only overview**, stripped of technical justification, claims, or marketing language.

Project requirements – narrow overview

The project requires the organisation to:

- **Replace ageing telecoms infrastructure** at the primary corporate campus to support current and anticipated demand from hybrid working and data-intensive applications 🌐
- **Provide a unified internal and external communications platform**, removing reliance on proprietary, hardware-bound systems and enabling better integration with modern collaboration tools 🌐
- **Increase network capacity and resilience** to prevent congestion during peak usage and critical operations such as video conferencing and data backups 🌐

Call out unsatisfactory responses as soon as you see them. Don't think ChatGPT will get the hint if you politely overlook weird behaviour

Correcting the AI mid-conversation

Sometimes, a conversation with ChatGPT becomes so unfocused that the best way to get back on track is to start over – either by clicking the "New conversation" button, or by instructing the AI to forget everything and begin afresh.

The trouble is, as well as clearing out unwanted tokens, this can discard a lot of potentially useful context. If you've just spent 10 minutes teaching the AI exactly how you want it to parse a spreadsheet, you're not going to be eager to write off the whole dialogue over a few crossed wires.

In lieu of taking the nuclear option, you can always give ChatGPT guidance and corrections within the conversation. If you're going to do this, it's important to call out unsatisfactory responses as soon as you see them. Don't think that ChatGPT will get the hint if you politely overlook weird behaviour. The AI won't normally infer that a particular response

is wrong unless you directly tell it so. If you don't act promptly, the bogus context will continue to influence future exchanges, and may even be strengthened by that apparent ongoing relevance.

How you correct the AI is important, too. If you just say "that's not what I wanted", ChatGPT will apologise – which is a waste of everyone's time – and have another try, which may well be no closer to your goal. You'll get a much better result if you frame your correction as a positive steer rather than a negative complaint. For example:

.....
Give a narrow overview focusing only on the project requirements.

An explicit instruction such as this should outweigh the previous wrong behaviour and point ChatGPT squarely back in the right direction.

Iterating without starting over

When you think of iterating, you perhaps imagine repeatedly applying a particular transformation. In the context of the AI, that could mean generating the title of a book, then dividing it into three sections, then dividing each section into chapters, and so on. You could do so by entering the same prompt over and over again:

.....
Break it down further.

This is a powerful way to work with the AI that wouldn't be feasible in a single prompt. However, iteration doesn't have to mean always going in one direction. You could alternate between operations that generate content and ones that narrow its focus to create a story that, again, would be all but impossible to define all at once.

Nor does iterating have to be towards a determined goal. You can iterate by repeatedly asking the AI for information or suggestions, and exploring or accepting whichever you like best. Rather than refining your prompts, you're refining your own purpose, in response to ideas and answers that you didn't originally know you were looking for – and which you could never have discovered in one operation. 🎯

Constraints and uncertainty

Keep ChatGPT grounded in the real world



n page 13, we saw how you can reduce the risk of hallucinations by building constraints

into your prompts. This doesn't just improve the response to a particular question – it helps to keep the whole conversation focused, by discouraging ChatGPT from bringing unreliable inferences into the context.

This is a big deal, for reasons we've already touched upon. Once a dodgy proposition is present in the conversation, it will continue to influence the AI's responses until you actively overrule it. And unless something contradicts it, a shaky assumption will be treated with just as much weight and relevance as a solidly corroborated fact.

While you can countermand wrong assumptions, it's tricky if you don't know they're there – and some of them can be subtle. For example, if you ask ChatGPT about

tax brackets in Germany, it may assume you want other questions answered from an EU perspective.

Your best approach, therefore, is to prevent the AI from drawing extraneous inferences in the first place. You can do this using constraints to limit ChatGPT's interpretation of a task, object or scope – or, you can directly tell it how to behave, with a prompt like:



Turn these notes into action points. If it's unclear what a phrase refers to, ask me – don't guess.

You can also keep things clean by specifying an output format that doesn't invite inference or invention. For example, if you ask ChatGPT for information in a list or a table, it gives the AI very little opportunity to misinterpret your meaning, or introduce unwanted ideas of its own.

Understood. I'll convert what's clear into concrete action points, and I'll explicitly flag anything ambiguous and ask you to clarify rather than guessing.

I'm basing this strictly on the notes you provided .

Clear action points

Governance & direction

- **Clarify the primary objective of the sustainability initiative**
Determine whether the driver is regulatory compliance, cultural change, PR positioning, or a combination – and document this explicitly.
- **Pause all external sustainability claims**
Instruct Marketing to halt any "We Are Green" messaging until robust data exists to support claims.
- **Define agreed terminology**
Establish internal definitions for *net zero*, *carbon neutral*, *green*, etc., to avoid accusations of greenwashing.

If you want an answer that's based strictly on the data, you need only ask.

Warning the AI in advance

As well as removing ambiguity from your task specification, you can help ChatGPT avoid mistakes by anticipating potential factual confusion, phrasing your prompt so as to head it off.

For example, errors can arise when you ask about current affairs or fast-changing situations, as the AI's training data is often six months or more behind the times. This isn't a problem for tasks such as analysing computer code – but for information about the here and now, it pays to explicitly specify the timeframe:



What is the UK's best-selling newspaper as of April 2026?

Including a date allows the AI to determine that its own information is likely to be out of date, and add a caveat to its response, or offer to go online and check the latest sources.

You can also give ChatGPT advance warning of issues it might hit while working on data you've supplied:



Identify growth areas (there may be inconsistencies in the data).

Flagging up a risk tells the AI that it isn't expected to fit all of its findings into a neat, coherent narrative, which can slash the risk of hallucinations or far-fetched assumptions. Mentioning that an error might exist gives the AI "permission" to ask you about any puzzling contradictions, rather than silently settling them for itself.

If you need info about the here and now, specify the timeframe

Challenging overconfidence

ChatGPT is designed to give confident answers. It will try to make them accurate and truthful if it can, but even if it can't do that, it will usually generate something that looks like an authoritative response.

This means you absolutely cannot trust the AI's confident tone as an indication of the reliability of the information it presents. ChatGPT might be working from nothing more than a statistical surmise – but once it's formulated and supplied an answer, it will commit to it as if it were the uncontested truth, and even make absurd leaps of logic and interpretation to justify its claim.

This isn't quite the same thing as hallucinating. When we talk about hallucinations, we're generally referring to plausible-seeming claims made by the AI that have no basis in reality. Overconfidence, by contrast, doesn't necessarily mean that the information is wrong – merely that there might be doubts, caveats or alternative interpretations that the AI isn't revealing.

As is often the case, ChatGPT's fluent, assured delivery makes overconfidence hard to spot. The trick is to look past the tone and think about what evidence the

Thought for 11s >

You're right to challenge that — my earlier "direct train in ~3 hours" answer doesn't hold up.

There **isn't a direct train** from Geneva Airport to Oulx/Sauze d'Oulx, and the **timings I implied were too optimistic**. Route planners and rail timetables show you're looking at multiple changes, and the fastest rail journeys are typically **well over 4 hours just to Oulx** (before the transfer up to resort). Trainline +1

What the quickest route actually is (most days)

Door-to-door fastest: drive / private transfer

- Rome2rio puts the **road distance at ~203 km** and indicates driving is the **fastest way**. Rome2Rio
- In normal conditions that's roughly **2h 40–2h 50**, but it can vary a lot with traffic and winter conditions.

The AI's first answer may really be just a guess – feel free to push back on it.

AI is presenting to support its claim. If a conclusion seems too neat or glib, you can force ChatGPT to expose its process with a prompt such as:

.....

How did you arrive at that answer?

Counter-intuitively, while the AI chatbot will try to project certainty on a prompt-by-prompt basis, when it's directly challenged it has no problem revealing that it previously made a dubious assumption, or made an overstatement that wasn't justified by the facts.

If you asked me to identify **outliers** in that dataset, there are several things I'd **be genuinely uncertain about unless clarified first**. These aren't technical nit-picks — they directly affect *what counts* as an outlier.

Here's the uncertainty, grouped by type.

1. What "outlier" means in context

I wouldn't know whether you mean:

- **Statistical outliers**
(e.g. IQR, z-score, standard deviations from the mean)
- **Operational outliers**
(values that are unusual but expected in real-world conditions)

Foregrounding uncertainty deters ChatGPT from making up answers.

All of this might make you wary of ChatGPT's judgment on anything but the most black-and-white questions. However, by knowing how the AI handles uncertainty, you can anticipate when it may become an issue and set things straight. As with constraints, prevention is better than cure, and in this case that's as easy as asking the right question at the right time. Let's imagine you're uploading a file of lab reports, and you want the AI to highlight any outliers. Rather than make that request right off the bat, start by asking:

.....

If I asked you to identify outliers in this data, what would you be uncertain about?

The AI looks over your file and spells out all the elements of the task it isn't sure about – without, at this stage, trying to settle them for itself. This gives you a chance to clear up any issues before the task begins, rather than allowing ChatGPT to make its own arbitrary decisions, as it would have if you had attached the instruction to the upload. 

Making uncertainty explicit

As we've seen on these pages, ChatGPT isn't comfortable with uncertainty. When it encounters an ambiguous proposition, it tries to resolve it internally, even if this means drawing on tenuous suppositions.

It won't voluntarily admit to past uncertainties, nor acknowledge any unresolved issues that might be hovering over the conversation in the form of conflicting bits of context.

Learning from the output

Use ChatGPT's answers to improve your prompts

In this section, we've explored various ways to adapt your prompting to make the best of the AI's quirks and habits.

This is necessary because, for better or worse, the AI isn't going to adapt to you. Sure, you can tell it to behave a certain way within a conversation, but when you start a new chat, it will be right back to its baked-in behaviour – aside from any saved memories that may come into play.

Thankfully, you don't need to memorise all that we've covered. You'll probably be able to make a big difference to your personal workflows using just a few of the

parameters we've discussed. And, of course, you can use ChatGPT's memory function to embed ongoing constraints, with prompts such as:

.....
Remember that when analysing financial statements I only want to see UK figures.
.....

To develop your overall prompting technique, you could think of the AI as a machine for testing the effectiveness of your input. A clean, correctly targeted response is a thumbs-up for your phrasing (and prior framing), while hallucinations and misapprehensions suggest that

you need to “debug” your approach to prompting. The better you get at speaking ChatGPT's language, the more effortless and productive your collaborations become.

With all that said, be warned that OpenAI periodically updates the ChatGPT engine to make it smarter and more intuitive. When a new model is released, you'll be automatically moved onto it. This could cause your prompting strategies to suddenly become obsolete, or start working differently; you may not have the option to switch back to the older model, either. There's nothing you can do about this – just be prepared to make adjustments when needed.

Knowing how far you can go

By now you probably have a good sense of the types of prompt that work well for you – and an instinct for formulations that are likely to cause problems. In other words, you're moving from thinking in

terms of individual prompts to broader patterns of operation.

However, no matter how in tune with the AI you become, bear in mind that perfection isn't always attainable. As

we observed in section 1, ChatGPT will carry on suggesting “next steps” forever, so it's up to you to decide whether a task is finished to your satisfaction – but sometimes, the engine simply can't get you where you want to go. For example, if you upload an image of your handwritten notes, ChatGPT may be unable to reliably read and interpret them. Or, it might not have access to the information needed to check and validate the content.

If you suspect you might be asking too much of the AI, it can be revealing to ask a question along the lines of:

.....
What improvements remain that you're not able to provide?
.....

This should give you an insight into whatever is holding ChatGPT back, and provide a pointer as to what you can do to finish off the job yourself. That's a better way to spend your time and effort than trying to cajole the AI into doing something that's beyond its abilities. 🤖

Good question. Here are the **meaningful improvements that remain but aren't fully available to me without extra input, authority, or risk:**

1. Authoritative sourcing and citations

I can tidy structure and wording, but I can't **independently verify or footnote every claim** (e.g. “95% decline”, “<7% of historic range”, Sundarbans projections) to academic or policy-grade standards without being told:

- required citation style
- acceptable sources (IUCN only? peer-reviewed only?)
- publication context (school, NGO, journal, media)

Without that, adding citations risks **false authority**.

2. Audience-specific positioning

Asking ChatGPT what it's not able to do can be a helpful reality check.

Fix a bad response without starting over

If the conversation strays off course, a few well-chosen prompts can bring it back onto your desired topic

📄 🔊 📌 🗑️ ...

What does 5G mean?

5G means **fifth generation** of mobile (cellular) network technology.

The “G” is just a generational label used by the telecoms industry. Each generation (1G, 2G, 3G, 4G, 5G) marks a new set of **technical standards** for how mobile networks work — including how devices connect to base stations, how data is encoded and managed, and how the network core handles traffic.

1

Loose language

Our conversation starts with a simple question – too simple, as we haven’t properly defined our goal. We specifically want to know what the term “5G” means when it appears on our phone screen, but ChatGPT can’t be expected to guess that – instead, it gives a broad definition of the term.

📄 🔊 📌 🗑️ ...

What does 5G mean when it appears on my phone screen?

When “**5G**” appears on your phone screen, it means your phone is **currently connected to a mobile network that identifies itself as a 5G network**, or at least to a part of the network that uses 5G signalling.

A few important clarifications that are often missed:

- It **does not guarantee a specific speed**. The icon reflects the *type of network*, not how fast your connection is right now.

2

Adding context

We can get the information we want by repeating the question and including the necessary context. However, ChatGPT will interpret this as a follow-up, rather than a course correction; all the extraneous information from its first answer is still assumed to be relevant to the chat.

Focus on 5G as a feature on my phone

Seen purely as a **feature on your phone**, 5G means:

Your phone has **5G capability enabled** and is allowed to use **5G radio modes** when the network makes them available.

Concretely, that feature does a few things:

- It **unlocks additional radio bands and modes**
Your phone can use frequencies and signalling methods that are only available to 5G-capable devices. A 4G-only phone simply can’t access these, even if the network supports them.

3

Focus the conversation

To refocus the conversation, we need to tell the AI how we want it to understand references to “5G” going forward. While this doesn’t wholly replace the earlier context, it establishes a working assumption for future responses, so we don’t need to keep repeating “on my phone”.

What ideas about 5G are you now excluding?

Focusing **only on 5G as a capability of your phone and your day-to-day experience**, I’m now deliberately excluding these ideas:

- That 5G is a **general technological revolution**
I’m not treating it as a shift in global infrastructure, standards, or industry direction — just what it lets *your device* do.
- That 5G is a **marketing label**
I’m ignoring advertising claims, branding, and carrier hype about what 5G is “meant” to represent.

4

Confirm the scope

Getting the AI to confirm what it won’t do provides assurance that we’ve achieved our goal. More than this, it reframes our focus in terms of constraints, and brings them explicitly into the context. This reduces any chance of our intention being accidentally overwritten. 



Business and productivity

ChatGPT has huge potential for businesses. It can help you communicate effectively, plan projects, and even expose hidden weaknesses in your company culture. It needs to be used with care, though – let's see how.

-
-
-
- 36** **Writing clearly in high-stakes situations**
.....
 - 38** **Planning with business frameworks**
.....
 - 40** **Understanding customers and competitors**
.....
 - 42** **Turning raw input into clear actions**
.....
 - 43** **Transforming unsorted documents into a plan** **TUTORIAL**
.....

Writing clearly in high-stakes situations

When accuracy matters, make sure ChatGPT understands you

C

hatGPT can be a powerful tool for business – but it isn't infallible, and ambiguous inputs may lead to results

that are unacceptable or unusable in a professional setting. For example, even if you prime the AI with context, it can easily generate statements that clash with your existing messaging, or don't square with the realities of the market.

AI-generated content can also mislead through inappropriate confidence. When ChatGPT sets out a claim in strong, unequivocal terms, a client might construe that as a commercial commitment that they can rely on – whereas, in fact, it's just the AI being its usual assertive self. On the other hand, hedging and cautious language could be read as doubt over an idea or product, when it only really means that the AI doesn't have access to all the relevant data.

The only sure solution is careful human proofreading, but by tailoring your prompt to close off ambiguity and convey an appropriate attitude, you can

Even if you prime AI, it can easily generate statements that clash with your existing messaging...

Report: Options to Return the Company to Profitability

1. Purpose and context

This report sets out a range of **potential measures** that could support a return to profitability over the medium to long term. It is written on the basis that:

- profitability is **not expected within the next 18 months**;
- the company remains committed to **careful evaluation rather than immediate action**; and
- all options will be **considered seriously before a settled decision is reached**.

The measures outlined below are **not recommendations**, but a structured menu of options for further analysis, consultation and sequencing.

Tell ChatGPT to plan with caution and it will take a suitably conservative tone.

reduce the degree of post-hoc correction required. For example, you might tell ChatGPT to:

.....
Generate a report on possible measures to return the company to profitability. We do not expect this to happen within 18 months. We will consider all options seriously and reach a settled decision in due course.
.....

This may not get you perfect results in one go, but it will discourage the AI from implying certainty and intent that aren't really there. As well as making the generated document more reliable, you'll be preventing ChatGPT from treating the matter as settled, ensuring the right framing for any subsequent instructions or questions.

Distilling information for specific readers

In previous sections, we've seen how ChatGPT can answer a question as if it's addressing a particular sort of person. In business, you often want results that really are addressed to a specific person – whether that's a line manager, the CEO or a disgruntled customer.

This is a perfect job for the AI; ChatGPT is entirely unsentimental about boiling down an issue to what needs to be said, excluding irrelevant side-notes and chat. It also knows a lot about business roles and functions, so you'll usually get good results if you enter a prompt such as:

.....
Turn this spreadsheet into bullet points for the head of marketing.
.....

You can do this as many times as you like with the same information, to create different takeaways for different

audiences – although you’ll need to explicitly tell the AI to work from the original source material each time; otherwise, detail and focus will get lost through multiple transformations.

You can also target a goal – rather than an individual – by building priorities in the prompt, such as asking the AI to produce a summary of progress towards sustainability goals. If you’re working with a particular set of requirements, you could upload the specification and enter a prompt such as:

Explain why our project is an ideal candidate for funding, based on the criteria set out in the attached document.

5	Scotland	26120.11	24956.04	24643.67	25031.81	27281.
6	Wales	13461.11	13597.82	88888.88	14377.37	13815.

Turn this spreadsheet into bullet points for the head of marketing

I’ve turned the spreadsheet into **clear, senior-level bullet points** aimed at a **Head of Marketing**. This focuses on **directional trends, regional performance, and implications**, rather than raw numbers.

High-level regional performance

- London

Tailored data summaries focus on what’s relevant to a particular reader.

1. Scope and requirement fulfilment

Rating:  Amber

Rationale:

- The core objective (demolition of the scout hut and replacement with garages) remains intact.
- However, the scope has **materially changed** from eight garages to six due to the sewer constraint.

The AI can rate plans and documents according to criteria supplied by you.

Drafting critical documents

ChatGPT can be a huge help when it comes to drafting – just ask and it will generate a contract, a report, or an action plan from scratch. Take care not to be distracted by its fluency, though: the AI likes to generate neat, authoritative-sounding statements. In a business context these can be deceptive, making leaps of inference seem more solid than they are, or glossing over awkward details that ought to be flagged up. And remember, when the AI states something definitively, that means it now believes it definitively, so an overstated conclusion doesn’t just lead to one bad answer but, potentially, an entire conversation founded on false premises.

It’s safer to frame the task so that the AI’s goal includes acknowledging weak points as well as strong ones. For example, let’s look again at the funding application example above. In this prompt we’ve explicitly told the AI that our project is “an ideal candidate”, and so it will treat that as an incontrovertible fact and craft an argument around it.

That might be just what’s wanted for this particular purpose, but it’s putting rhetoric ahead of reality. Instead, consider taking a step back and asking:

Rate our project based on the criteria set out in the attached document.

You can go on to ask ChatGPT to summarise the strengths of your proposal for submission, while also suggesting ways of addressing your weaknesses ahead of the interview phase.

As the drafting process continues, ChatGPT also makes an excellent consistency checker. Feed it a draft, or a folder full of drafts, and it can instantly pull out contradictions, ambiguities, or mismatches in content. This helps you track whether ideas have become lost or shifted over time, or whether different parts of the business are singing from different hymn sheets.

Reframing messages

On the opposite page we looked at distilling information – that is, condensing content down to just the parts that a particular reader will want to see. Reframing is a similar idea, but it works on the level of perspective and emphasis. It doesn’t include or exclude data, but rather presents a particular interpretation of it.

For example, if you’re preparing a report for the CFO, you might start by distilling the master sales document to key summaries; then you could ask ChatGPT to contextualise the bottom line within overall market trends. For an individual sales operative, the same document could call for different distillation and framing.

You can reframe information quite independently of distilling it. If you’re presenting the sales data in a meeting with multiple stakeholders, you might not want to cut out too much detail from the information, but you could still instruct ChatGPT to:

Frame these figures as a growth opportunity for the business.

Your choice of framing is entirely down to you. If you’re arguing for a change in strategy, you could ask ChatGPT to frame the exact same figures as illustrative of stagnation, and it would just as happily oblige. 

Planning with business frameworks

Get help setting objectives and making informed decisions

Even though ChatGPT doesn't know about your specific business, it can still be a great help for planning and decision-making. Indeed, its lack of prior framing makes it a great antidote to organisational assumptions – it's like hiring an external consultant, but more flexible and cheaper, too.

One area where the AI excels is at planning out new projects. If an idea originates from you, you can simply tell ChatGPT your objective and let it generate a plan to make it a reality. An example might be:

How can my company cut marketing costs without affecting sales?

To make clear that you're not looking for a generic answer, you can add more detail about your line of business or even attach some corporate information and reports as background information.

In all cases, this is the sort of question that ChatGPT loves to run

with. It gave us a detailed, multi-page response, setting out steps for understanding the current function of marketing, spotting inefficiencies, identifying possible savings, and applying changes while protecting revenues. Of course, this is just the start of the planning process: you can now work with the AI to add in constraints, variables and priorities as appropriate.

When you're happy that the plan reflects your business realities, you can formalise it by asking ChatGPT to:

Turn this into a proposal that I can present to the board.

This phrasing invites a more grounded approach, and ChatGPT will reframe your proposed activities in terms of options, trade-offs and projections. Even if you're not literally talking the senior management through the proposal, this is the document you need to move forward, with assumptions and uncertainties made explicit along with key decision points.

Making sense of incoming proposals

Often, the project objectives won't come from you. You might be tasked with finding ways to achieve a goal, or a project proposal might be handed down for you to implement.

The AI capabilities we've discussed above really shine in such situations. If you're trying to get to a particular outcome, you can ask ChatGPT to:

Give me three ways to accomplish the project goals in the attached document, and detail the assumptions and trade-offs involved in each.

While the AI can't be relied upon to surface every possible implication of an idea, it will give you a strong starting point. It can also be illuminating to ask the AI to generate the best arguments it can for and against each option; this lets you immediately spot deal-breaking weaknesses in a plan, or discover a potential win that's worth stretching for.

If you're working from an existing proposal, ask the AI to break it into phases, and report the risks and decisions involved at each stage, as well as highlighting unknowns in the specification. It's a great way to expose potential issues while there's still time to refine the plan.

Proposal: Reducing Marketing Costs While Protecting Sales Performance

1. Purpose

This paper sets out a structured approach to **reducing marketing expenditure without negatively affecting sales**, by improving spend effectiveness rather than reducing demand generation capability.

The proposal recognises that, in a mature toiletries portfolio, **not all marketing spend is incremental**, and that significant efficiencies can be achieved through sharper prioritisation, simplification, and redeployment of investment.

ChatGPT can turn a loose idea into a clear, board-ready business proposal.

The AI can't be relied upon to surface every possible implication of an idea...

Stress-testing decisions

If you've previously used ChatGPT in a business setting, you'll have noticed that it frequently offers to "stress test" ideas or proposals. By this it means trying to figure out the possible ways in which a plan could fail. You might think of the process as a cross between a risk analysis and a brainstorm.

This is another great application for ChatGPT. The AI's considerable inferencing powers help it spot angles a human could easily miss, and since you're not asking the engine to commit to a particular reality, there's less risk of overstatements or hallucinations.

You can initiate an open-ended stress test by simply asking for one. However, as we've seen, it often helps to add constraints and directions to a prompt. For example, you might find it more revealing to ask for a list of the most likely problems with a plan, and a separate list of the most impactful ones.

Alternatively, you can define a specific failure scenario and get the AI to work backwards from there:

1. Temporary loss of output or availability during transition

What happens

- Line upgrades, automation installs, ERP cutovers, or process standardisation reduce effective capacity.
- Planned downtime overruns or coincides badly with demand peaks.
- Yields dip as new systems are tuned.

Why this hits revenue

- In FMCG, *availability drives sales*. Even short gaps at shelf translate directly into lost volume.
- Retailers may delist temporarily or reduce facings if service levels slip.

Predict problems and explore how they could affect your project.

If this plan were to lead to lower revenue in 12 months, what would be the most likely reasons for that?

Either way, ChatGPT benefits from not being embedded in your organisation:

it will give you a clear-eyed reading of the risks. It also has no emotional stake in the plan, and no qualms about defying groupthink or gainsaying the boss – so depending on the company culture, the AI could be a useful vehicle for legitimate concerns you'd rather not personally own.

Surfacing hidden assumptions

As well as spotting potential issues, ChatGPT's "outsider" perspective means it's perfectly positioned to expose the unspoken assumptions encoded into your everyday operations and internal messaging. You can see how this works by entering a simple request, along the lines of "list the unspoken assumptions in this plan" – or, for a more focused response, you could add some framing:

What does this plan assume will remain constant in the next 12 months, and what does it assume will change?

When applied to a single project, this process is complementary to stress-testing: it instructs ChatGPT to list all the assumptions that aren't explicitly addressed in the plan, regardless of whether or not they contribute to a failure mode. Even if the exposed assumptions are benign, making them explicit helps to ensure all stakeholders understand what's expected of them, and that the premises of the project don't clash with their own plans.

What the plan assumes will remain broadly constant

1. Underlying consumer demand for core categories

- Demand for staple toiletries (soap, shampoo, basic skincare) remains stable.
- No structural decline in category penetration or usage frequency.
- Volume fluctuations remain seasonal rather than systemic.

Why this matters:

The plan assumes any short-term revenue impact is operational, not demand-led.

2. Brand equity and consumer trust

- Consumers continue to trust existing brands despite:
 - packaging changes (PCR plastics),
 - minor formulation or process adjustments,
 - quality improvements being largely invisible.

Listing your implied assumptions helps ensure nothing is overlooked.

More powerfully, this exercise can be extended across multiple projects and operations, to reveal assumptions that come up repeatedly. If you ask ChatGPT to determine who in the organisation is

responsible for ensuring that the assumed premises are true, you might uncover a mismatch between management expectations and the work that's actually being done on the ground. 

Understanding customers and competitors

Sharpen your strategy with AI-powered role play



We've seen how ChatGPT can be useful for working with internal plans and communications.

Now let's turn its lens outward and see what it can tell you about your customers.

An important thing to say right away is that ChatGPT can't actually speak for the people you do business with. It can't reliably reveal how they'll respond to a marketing campaign, or whether they'll tolerate a price rise.

What it can do is work from real data to put together a model of behaviours within your customer base, which you can use to explore hypothetical questions and scenarios.

You don't need to format your data in a particular way to do this. Just ask the AI to build a composite persona of your customer base, then upload whatever information you have – questionnaire responses, support tickets, user reviews, even anecdotes – for it to analyse and build into its model. Then, you can ask questions such as:



What makes customers choose our product?

ChatGPT will generate an answer drawn from patterns and nuances

Data protection

Make sure that any customer data you share with ChatGPT is fully anonymised, so individuals can't be identified, even indirectly. Otherwise, you could be in breach of the law.

Based on what customers actually say across questionnaires, reviews, support queries, and direct emails, there's a pretty consistent picture of *why people choose you* – and it's not price-led.

Here are the main drivers, grounded in the evidence.

ChatGPT can pull together a diverse set of sources to find recurrent themes.

1. Trust in quality and workmanship

Customers repeatedly choose you because they believe the end result will be *done properly*.

- Carpets described as "excellent quality", "luxury", "10/10", "durable", "so soft"
- Fitters described as professional, clean, polite, punctual, tidy, and efficient
- Strong confidence in underlay advice and fitting standards

in the data you've provided.

To reiterate, this will only be a plausible response, not necessarily the most likely one, and certainly not a guarantee. Think of your interactions as simulations, rather than genuine market research.

You can also ask the AI to focus on particular sectors, such as high-value clients, or those who only buy infrequently. However, the more narrowly you slice the analysis, the more ChatGPT will have to lean on supposition and inference – consider asking it to detail any guesses or assumptions it's made in arriving at a particular answer.

Modelling competitors

As well as modelling your customers, ChatGPT can help you gain a better understanding of your competitors. While you may have limited visibility of their internal operations, there's plenty of public information out there – such as revenue figures, marketing messages and reviews – plus what can be inferred from your own customer comments. Share whatever you have and the AI will create a model company profile that you can probe with questions such as:



What are our competitors' strengths and weaknesses, relative to our business?

As before, don't place too much faith in ChatGPT's responses – they're based on trend-spotting and pattern recognition, rather than primary research and analysis. All the same, they're grounded in reality, and may highlight

real connections and insights that you might otherwise have missed.

If you're willing to get more speculative, ChatGPT can even plausibly suggest what your rivals might do next, based on past form and public statements. If you're in close competition, you could try using the AI to "wargame" scenarios and explore how your adversaries might respond if you were to take a particular course of action.

Mapping customer journeys

Another way that ChatGPT can help you understand your business is by exploring the commercial "journeys" that individuals go on, as they progress from being potential buyers to satisfied customers. These might involve their seeing adverts, researching reviews, visiting your website, beginning the ordering process – and, hopefully, working all the way through it and coming back for more.

Again, ChatGPT's inferencing power is useful here, because such journeys are rarely linear. Customers may loop back, compare options and change their minds multiple times before finally completing a purchase. If you simply track sales then you'll never know what an individual has been through on their journey. However, by mining the data you uploaded earlier (see opposite), the AI can identify points in the process where customers are likely to hesitate or step away, and suggest reasons why. A suitable prompt might be:

Based on the uploaded data, what are the friction points in a typical customer journey?

Since ChatGPT is working from aggregated data fragments, it can't trace a specific person's experience from start to finish. However, it can pick out recurrent themes, and suggest ways to smooth the pathway for a typical customer – which will also make uncommitted prospects more likely to reach the end of the journey and make a purchase.

Another way you can use ChatGPT's inferences about the customer journey is to explore the possible effect of changes. For example, you could ask it how a price change or a website redesign might affect a visitor's purchasing decision, and what mitigating measures you might put in place.

1. Discovery & first contact

Get a clearer picture of your customer journey – and any stumbling blocks.

Friction points

- **Slow or inconsistent response to initial enquiries**
Multiple customers mention delays replying to first emails or quote requests, and at least one explicitly links this to dissatisfaction
- **Online booking friction**
The quote/booking system is described as "clunky", which introduces hesitation right at the entry point 🌐

Impact

This is the point where price-sensitive or time-pressured customers quietly drop out and go elsewhere, often without telling you.

Looking across the uploaded questionnaires, reviews, support enquiries, and lost-sale emails, customers are pretty clear about **the jobs they're trying to get done** – and where your offer only partially meets those jobs.

This isn't about things you do badly; it's about **intentions that spill beyond your current service boundaries**.

1. "I want this done within my deadline"

Unmet job: fast, predictable turnaround

Customers are often working to:

- Renovation schedules

What demand is your business not meeting? The AI can help you find out.

Identifying opportunities and unmet needs

We've looked at who your customers are, and what they do. So far, though, we haven't really focused on what they want – which, of course, is hugely valuable information to any business.

Using the same data as before, the AI can construct simulated customer profiles with realistic motives and priorities, based on what's known about your actual customers, along with general inferences about the market.

The caveats we've already raised apply even more strongly here. The information you feed the AI may reveal insights about what customers think and do, but it's unlikely to contain reliable content about how they'd behave under different circumstances – you're very much working with extrapolations.

All the same, there's still a lot of potential value to be found in the

gaps between what your customers are saying and what you're offering. For example, you could give the AI an overview of your company's products and services, and ask it:

What are customers trying to do that we don't fully cover?

By exposing recurrent patterns and inferred desires, ChatGPT can help you identify possible opportunities to better meet your customers' needs. It could perhaps complement a traditional customer survey by bringing out motivations that customers aren't aware of themselves – and ranking them according to the scale of the inferred desire and potential business impact. 🌐

Turning raw input into clear actions

Move from vague goals to accountable assignments

Almost every business has ongoing plans whose success depends on a clear understanding of what must be done, by whom, and when. ChatGPT can help ensure everybody is on the same page by combing through any sort of document – be it a report, a project plan or a marketing analysis – and immediately highlighting the action points. That doesn't just mean explicit instructions, but additional responsibilities and dependencies that are implied by the context.

This makes it a powerful tool for ensuring everyone is moving forward together, when intentions could otherwise be lost in lengthy meeting notes or message chains. You can further use the AI to process the action list itself, with an instruction such as:

.....
From this email thread, identify the tasks that need to be completed immediately.
.....

2. Confirm venue details operationally

Owner: James Reed (with Sarah Miller input)

Once the booking is secured:

- Confirm access times (load-in / load-out)
- Confirm freight lift availability and weight limits
- Confirm power capacity and kitchen access specifics
- Confirm any restrictions (noise, open flame, food service rules)

These are implied prerequisites for the live bake and equipment planning 

The AI helps to break down tasks, and identify who should handle them.

This lets you get started on activity right away, rather than wasting time and effort trying to find the priorities and objectives that are buried in a dense conversation.

You can ask the AI to enumerate outstanding decisions as well as actions: it never hurts to have a clear list of items that need to be decided before a project can go ahead. And in the process, the AI could unearth

some of the hidden assumptions we mentioned on page 39, forestalling any problems that can arise when parameters are taken for granted.

However, always remember that ChatGPT is just a tool, and it's the human operating it who's ultimately accountable for its output. Check its working for yourself before pinning the success of your project on an AI-generated action plan.

Combining multiple discussions

Earlier we saw how you can upload a whole collection of files to ChatGPT, and let it figure out the connections. Our goal then was to assemble profiles of customers and competitors, but the same capability can be used to condense multiple working documents into a single summary. For example, you could take half-a-dozen discussions with different stakeholders and turn them into a project brief; or, upload a mix of emails, messages and shared documents, and generate a review of how a particular piece of work was delivered.

As well as consolidating information, the AI can highlight inconsistencies between documents. Ask it to turn a folder full of departmental reports into one combined account and it will flag any areas where goals or assumptions don't appear to be aligned.

If your documents are dated, ChatGPT will also be able to track how content develops over time. For example, you could upload a series of meeting minutes and ask the AI how an idea evolved across multiple sessions. You can check the current state of play by

bringing in the action-point extraction technique we discussed above:

.....
Based on these minutes, what's agreed and what's still undecided?
.....

A prompt such as this could reveal mistakes, such as an idea being forgotten between meetings, or parameters that have shifted without discussion – giving you a chance to revisit and resolve such issues before they affect the business. 

Transforming unsorted documents into a plan

With just a few prompts you can cut through a mess of discussions and produce a complete schedule of operations

Planning, design and approvals

- Update architectural drawings to reflect the **revised scheme of six garages** (reduced from eight) and the adjusted building footprint to avoid the sewer main. 🌐
- Ensure garage design and materials **align with local aesthetic expectations**, including brick choice, proportions, landscaping, and visual integration with surrounding buildings. 🌐
- Prepare and submit the **formal planning application** based on the revised layout and specifications. 🌐 🌐

Community provision and mitigation

- **Alternative provision for the Scout group**
A resident letter raises the loss of a local Scout base and asks for relocation or replacement, but **no document commits to providing an alternative site, funding, or support**. The issue is acknowledged but not resolved.
- **Scope of heritage mitigation beyond signage**
Salvage of the Scout Group signage is agreed, but it's unclear whether **any other heritage elements** (timbers, plaques, records) were considered or ruled out.

1

Feed in the raw material

To help us get on top of a complicated project, we've uploaded all of the documents we have that relate to it, and asked ChatGPT to identify the action points. The AI works out everything that needs to be done, including requirements that aren't explicitly mentioned in our meeting notes.

2

Check what's missing

It's always a good idea to ask the AI to surface uncertainty; in this case, there are several points raised in the documentation that haven't been addressed. We can work through this list to confirm what we still need to deal with, or to positively decide if we're going to let something slide.

Completed / earlier-stage tasks (already done)

1. **Draft formal response on refurbishment vs demolition**
Responsible: Marcus Thorne
Timing: October 2025
2. **Commission bat and ecology survey**
Responsible: Elena Rodriguez
Timing: By 30 October 2025
3. **Circulate resident survey on garage demand**
Responsible: Sarah Jenkins
Timing: By 5 November 2025

Days 1–30: Planning finalisation and mobilisation

Week 1 (Days 1–7)

- Objectives:** lock design, remove heritage risk, submit planning
- Finalise salvage, storage and archive arrangements for Scout signage
Owner: David Miller
Dependency: none (must complete before demolition)
 - Update architectural drawings to six-garage layout with sewer clearance

3


Confirm sequencing

Sequencing and assigning responsibility are crucial steps in an action plan. Our uploaded items tell ChatGPT what the various team members have done to date, and based on that it's inferred their roles and assigned them upcoming tasks – although, naturally, we can make changes if we need to.

4

Draw up the timeline

The last step is to bring it all together into a schedule. The AI had already picked up some timings from our documents, but now that we've given ChatGPT a completion deadline, it can properly organise the tasks, with deadlines and clearly stated dependencies to ensure a smooth workflow. 🤖



Research and analysis

For data-driven tasks, ChatGPT is a huge help. In this section, we'll see how it can spot data trends, highlight outliers, and flag missing values. We'll also explore ways to avoid blunders and generate reliable conclusions.

46

Finding and organising reliable information

48

Developing questions and testing ideas

50

Making sense of data

52

Checking reasoning and claims

53

Double-check the AI's answer **TUTORIAL**

Finding and organising reliable information

Trustworthy results from an AI? It can be done...

In the previous section, we used ChatGPT to support business decisions and activities – but it can also be a powerful tool for research.

However, some of the same cautions apply. A neat, fluent answer doesn't imply a slam-dunk in the data. In fact, checking ChatGPT's output for yourself is even more important in this context: a mistake doesn't just lead to red faces in front of the boss; it could invalidate an entire study.

And there are a couple of types of error that loom particularly large in this sort of operation. If you're asking the AI about information that's publicly available, it has a tendency to mix together primary evidence and external interpretation. If you provide the material yourself, ChatGPT will try to impose an overall narrative that minimises or skips over any points that don't fit.

A good way to test ChatGPT's read on a data set is to cross-check it – that is, to get the AI to examine the question from a different angle. For example, you could ask it:

.....
Which pieces of evidence support this conclusion, and which challenge it?
.....

Asking the AI to deconstruct its own claims can expose any shaky parts. You could also ask it which parts of its conclusion are poorly supported, or to generate other interpretations that might also be supported by the data. The best approach will depend on the data at hand, but accepting ChatGPT's initial judgment at face value is usually the worst.

Summarising large bodies of material

ChatGPT's ability to slice and summarise information makes it a great tool for extracting insights from large data sets. However, while the AI is fully capable of operating with the caution and precision of a data scientist, its default behaviour mode tends to generalise and conceal any kinks in the data. So, as we've seen before, you'll need to give it specific directions.

Again, the ideal strategy here will depend on your purposes, and the shape of the data itself. One way to stop ChatGPT from using too broad a brush is to approach the task in stages, starting with an instruction such as:

.....
Segment this data into groups that have something in common.
.....

ChatGPT loves to find structure, so this instruction plays to the AI's strengths. As before, the multi-step approach gives you a chance to ask what criteria it chose, check its working, and make any manual tweaks before moving on. You can then

1 Behaviour Groups

ChatGPT is great at picking out patterns – and exceptions – in limited data sets.

Movement (63 records)

Trend:
Sightings involving flight, hovering, climbing and similar actions form a substantial minority of the dataset.

Pattern:
These tend to reflect active, visible behaviour and likely occur in open habitats or when birds are hunting or relocating.

Non-conforming elements:

- Movement is not dominant — it accounts for fewer records than stationary + "other".
- Some high-frequency species are not primarily recorded in motion.

gain a closer look at the information with a follow-up prompt such as:

.....
Summarise the data trends in each group, highlighting any data that doesn't strictly conform.
.....

The result is a far more detailed and informative breakdown than you'd have seen if you'd simply asked for a summary off the bat; you can optionally now move on to ask the AI to summarise trends and outliers across the whole data set.

Identifying evidence gaps

When you ask ChatGPT to work with incomplete information, it doesn't complain – it will usually skip over gaps in a narrative, or silently fill them in with plausible interpolations. But while the AI doesn't always flag up deficiencies in the data set, that doesn't mean it can't. The same skills that make the AI so good at pattern matching and summarising also make it highly capable – when prompted – at highlighting when the facts and figures don't wholly stack up.

You can see this in action by sharing data, in almost any format, and asking:

.....
What's missing from this data?

Such open-ended prompts are often too broad to be useful, but they can be a powerful exploratory tool, revealing connections, or – in this case – disconnections that you might not have thought to seek out. In many cases, the data won't be obviously lacking; the issue will be that it doesn't strongly support the conclusion. If your goal is to cross-check ChatGPT's analysis, then you can always ask the AI to criticise its own work:

.....
What aspects of this conclusion are poorly supported?

This capability isn't limited to internal reasoning; you can type in a theory of your own, or upload a document and use the same prompt to check how well founded its claims are. Any weaknesses you find can be useful for updating your assumptions, or directing further research.

Structuring research logically

We've seen how ChatGPT can interpret and summarise research data, and it will also help you develop your findings into a larger argument. This doesn't just mean creating a large-scale summary of summaries, but formulating a new hypothesis that's supported or implied by the data.

It's best to go about this in a structured way, building up from basic terms to a testable proposition. This prevents the AI from leaping to connections and

Your conclusion is:

"Birds are most active in the daytime."

Here are the aspects that are poorly supported.

1 "Most active" vs "Most observed"

The dataset measures **sightings**, not actual activity levels.

We do not have:

- Continuous monitoring
- Activity duration
- Night survey effort

So the claim that birds are *most active* in daytime is not directly measured – only that they are most **recorded** in daytime.

Get the AI to turn the tables and challenge your own conclusions.

1 The dataset measures sightings, not activity

The strongest challenge is methodological:

- The data records when birds were **seen**.
- It does not record when birds were **inactive**.
- It does not measure activity continuously.

If observation effort is concentrated in daylight hours, then daytime dominance in the dataset reflects observer behaviour – not necessarily bird behaviour.

This is the single biggest counterpoint.

Breaking down what's *not* in your data exposes errors and assumptions.

conclusions too early, which could then exert an unwanted influence on later analysis. A good start is to get a grip on the data, with a prompt such as:

.....
What data are we working with here, what does it contain and what does it lack?

Even if you think you already know the answer, getting ChatGPT to acknowledge these details imprints them into the context, and can reveal elements you've overlooked. You can then move on to ask what recurrent patterns or themes occur in the data; keep the instruction neutral, to discourage the AI from trying to find narratives that may not be there.

The response may be suggestive to you, or you can ask what questions these patterns raise. Either way, ChatGPT can help you explore and brainstorm possible answers – we'll talk more about this overleaf.

However, if you want to use the AI to test your hypothesis, remain cautious – and remember that ChatGPT will usually go along with any theory suggested by you. Rather than asking it to confirm an idea, you'll get more reliable results if you task it with falsifying the proposition, with a prompt such as:

.....
What evidence suggests that this conclusion is incorrect? 

Developing questions and testing ideas

Let ChatGPT figure out what you're really asking

In research, the best way to get useful results is by starting from a clear, tightly delineated brief. If that sounds familiar, it's because one could say exactly the same about AI – and ChatGPT can be a helpful tool for setting out research methods and goals, whether or not you're intending to pursue them with AI assistance.

You might start the process with a particular question in mind, or merely a loose direction of interest. As we discussed on the previous page, ChatGPT can help you find potential questions lurking in data patterns, or you could ask it to review existing research and identify gaps that want filling.

No matter how you arrive at your question, the first step is to examine if it carries any hidden ambiguities or assumptions. This is one of those situations where using ChatGPT feels almost like cheating – you can literally just ask:

.....
What hidden ambiguities or assumptions are implied by this question?
.....

Phrasing the question like this tells the AI that such issues are definitely present, and pushes it to dig as deeply as necessary to find

Your working question:

"Are local populations of different bird species growing or shrinking?"

Here's what's implicitly embedded in it.

1 What does "local" mean?

Ambiguities:

- Is "local" a single site? A town? A river catchment? A county?
- Are all locations in the dataset treated as one ecological unit?

Nitpicking the question helps you work towards a strong framing.

them. For our purposes that's safer than letting the AI decide whether they're there or not.

You can go on to ask further questions such as what's excluded and included by the framing of the question, or whether it presupposes a particular outcome. However, rather than trying to anticipate all possible objections, you can get the AI to help you out again, with a query along the lines of:

.....
What would a complete answer to this question tell us, and what would remain assumed or unknown?
.....

Testing ideas from opposing angles

We've talked about testing hypotheses, and examining questions for hidden assumptions. That's a good start, but if you want strong, defensible conclusions, you need to challenge every idea you work through along the way. You can do this in general terms by simply asking the AI to:

.....
Make the strongest possible cases both for and against this proposition.
.....

For a more nuanced reading, try looking at your work through specific lenses. For example, if you're suggesting that sugary drinks should be banned, the argument will be easy to make on health grounds, but it might be revealing to ask ChatGPT to critique it through an economic lens, or with a focus on personal liberty.

ChatGPT is brilliant at this job: its broad training enables it to adopt a wide range of perspectives, without being at all attached to any particular one. It can make a forthright argument from one

ChatGPT can be helpful for setting out research methods and goals, whether or not you're intending to pursue them with AI assistance

standpoint, then instantly switch roles and deliver an equally forceful rebuttal. You can also use the trick of working backwards from an outcome, with a prompt such as:

.....
What argument could defeat this proposition?

The aim here isn't to find a single answer that can't be contested from any angle, but to scope out disagreements and possible failure scenarios before you take your research forward. This helps bring into focus what's objectively true about a proposition, versus what's a value judgment or assumption.

Converting notes into insights

As well as organising your data, ChatGPT can help you quickly find connections and trends buried in your research notes that might be worth following up. The idea of notes can be defined very broadly – in a laboratory scenario, these might include lab data, personal observations, reports of previous experiments, and speculative brainstorms. And the same approach can be used in academic research, market research, or any field where you're looking to learn from a mass of data.

Hunting insights isn't the same as identifying action points, as we did in the previous section. The goal there was to cut through the bulk and extract specific items that need taking care of. Here, we're asking the AI to consider the whole spread of documents, and find recurrent themes or trends. It's also worth asking about any potentially significant tensions and contradictions, and since these objectives are complementary, we can include both in one prompt, without a risk of compromising the results:

.....
What recurrent themes or tensions run through these documents?

Correlations on their own don't necessarily mean anything, but they may point you towards a lightbulb moment – or, differences in phrasing or emphasis could turn out to have been significant. Either way, ChatGPT's findings can help establish possible directions for your onward research.

Your proposition is:

“Curlews are thriving.”

A single, strong argument that could defeat it is:

The dataset does not contain the types of data required to demonstrate that a population is thriving.

Here is how that argument works.

1 The Data Measures Sightings, Not Population Health

The working dataset records:

Put your theories to the test by seeing how easily they could be refuted.

1 Presence vs Population Health

Theme: Regular sightings suggest stability.

Tension: Sightings are not the same as demographic strength.

Across the material, there is repeated movement between:

- “Frequently observed”
and
- “Ecologically thriving”

This tension runs through every attempt to interpret the data. The dataset is strong on visibility and frequency, but thin on breeding success and long-term viability.

ChatGPT calls out subtle issues that could derail your whole investigation.

Discovering useful patterns

If you want to dig more deeply into your data, ChatGPT can help you discover relationships of change, such as variables that tend to rise together, or distinct groups of data that interrelate in particular ways. If the data allows it, it can flag up things that usually seem to happen in a consistent order.

For this purpose, it's important to understand – and to ensure that ChatGPT understands – that we're firmly in analysis mode; we don't want the AI to try to invent explanations for its findings. ChatGPT's pattern-matching methodology makes it brilliant at spotting even subtle connections, and it can come up with plausible hypotheses to explain them all day long. However, since it has no internal model of real-world causality, it can't positively establish that X is the cause of Y, even if all the data points clearly to that conclusion. Any interpretation it comes up with will be speculative, and probably couched in deceptively confident language.

With that in mind, we can frame suitably cautious queries, such as:

.....
Which variables consistently move together across time periods?

Once the AI has found a pattern, you can explore further questions such as whether one value leads to another, or whether the trends change over time. You can also test the limits of any pattern:

.....
Under what conditions does the apparent relationship break down?

This could turn out to be the really interesting question – perhaps your breakthrough comes not from realising that two things are connected, but from figuring out the connection on which they depend. 🧠

Making sense of data

AI can help you understand what the figures are telling you

S

o far we've focused on structuring your research, testing claims and finding

patterns. Making sense of the data is a different sort of task – one, as we've noted, where ChatGPT isn't entirely reliable. Even so, there are things it can do to take some of the work off your hands.


One area where the AI can be useful is for collating your data. Working from a folder full of files is fine for finding patterns, but when you need to engage with actual figures (or factual reports), it's convenient to have them all in one place. ChatGPT is brilliant at dealing with the sort of issues that usually arise when you combine information from different sources, such as inconsistent logging formats.

You can simply upload a batch of documents and ask ChatGPT to combine them into a single spreadsheet. This will almost certainly result in something clean and coherent. However, as you'll know from previous sections, that sort of too-perfect result is a red flag. It suggests that ChatGPT has made neatness its top priority, and you've no way of knowing how it might have transformed the source material to achieve it. It's safer to start by asking:

How could all these documents be combined into one consistent data table?

This should give you a readout of proposed changes that you can check and tweak as needed, before giving ChatGPT the go-ahead. It might be a little tiresome, but it's still a lot easier than doing all the cleaning yourself.

Weekly Sales 2022: North East vs North West



Here's how North East compares with North West in the 2022 aggregate data:

Annual totals

- North East: £8,745,566
- North West: £8,998,372

- ➔ North West is higher by £252,806
- ➔ That's about 2.9% higher than North East overall.

As well as graphing the data, ChatGPT can describe it in simple terms.

Explaining numbers clearly

Once you have all your numbers together, it's easier to spot errors such as missing or duplicate values – small-scale issues that might not be picked up when looking for patterns, but which could nevertheless distort your conclusions if not addressed.

And with direct visibility of the data, you can use the AI for hands-on analysis. For example, you might ask it:

How do the numbers for the North East region compare to those for the North West?

Focusing on the numbers in this way should deter the AI from trying to boil the data down into a simplified narrative. Instead, you should get a neutral summary that describes a daunting set of figures in an easily understandable sentence or two.

ChatGPT can also use your aggregated data as the basis for calculations and analysis. If you're working with a complicated table of financial data, you could ask:

Which regions reported more than £1m in monthly revenue?

You could equally use this approach to turn a set of running figures into annual totals, or calculate growth against a fixed baseline. Or, you could ask the AI to combine figures from different sites to produce regional overviews.

These techniques can also apply to externally sourced data, such as a financial statement or environmental impact report. In this way, you can get your own perspective on the published figures, or produce breakdowns that aren't covered in the official narrative.

Spotting exceptions and anomalies

Anomalies in a data set can be highly significant. By this, we don't mean recording errors, such as the missing or repeated values we mentioned opposite; these don't tell us much, except perhaps from an operational perspective. However, real data points that don't fit the expected pattern are prime candidates for further investigation.

ChatGPT can help you identify outliers that hint at undiscovered interactions or flawed assumptions. You can use the same basic approach that we used for finding patterns – for example, you might ask the AI:

What data points deviate most from the overall pattern?

You could go on to ask ChatGPT to hypothesise some reasons for the anomalous values – although this needs to be approached with caution, as we'll discuss below.

You'll recall that ChatGPT also has a powerful ability to spot gaps in the data. On page 47, we used this to check whether the data supports the posited conclusion, but you can also use it in a pattern-matching context:

What's missing from this data that would strengthen or weaken the pattern?

The answer might be a few unknown values that would help to confirm or break a trend, or it might be data from a particular period or region – things a human might easily overlook. You can go on to ask the AI what the possible impact would be of acquiring that data, to get a sense of whether the gap might be concealing a significant discovery.

Before inviting the AI to generate an interpretation, ensure your data is clean and valid

To identify deviations from the overall pattern, I calculated z-scores (how many standard deviations each week is from that region's yearly average). The largest absolute z-scores represent the strongest departures from normal behaviour.

Biggest Outliers (Highest Absolute Z-Scores)

Region	Week	Sales	Comment
South West	17 Dec	£340,413	Largest deviation (3.19σ)
North East	17 Dec	£308,191	Strong spike (2.74σ)
Midlands	3 Dec	£305,268	Sharp uplift (2.73σ)
South East	3 Dec	£400,000	Very strong peak (2.67σ)

You don't need a degree in statistics to generate meaningful insights.

3 Assumptions Required to Interpret Attendance

The data does not tell us:

1. Venue capacity
 - We do not know whether 262 tickets is:
 - A sell-out
 - Or only 50% capacity
2. Whether these are advance sales or final totals
3. Whether pricing tiers correspond to seating zones (likely)
4. Whether unsold inventory existed in specific price bands

If the AI can't work out what you want to know, it will tell you exactly why.

Turning data into narratives

We've warned against letting ChatGPT impose interpretations on your material, but that doesn't mean you can't make use of its narrative abilities – these can be very useful for explaining your work to managers, stakeholders or funders, or just to clarify for yourself exactly what you're working with.

The key is not to let the AI jump straight from tentative findings to a polished story. Before you invite ChatGPT to generate an interpretation, it's important to ensure that your data is clean and valid, and provide any relevant context that might colour the interpretation. Then you can try entering a prompt such as:

What do these figures tell us about attendance? Detail the raw data, assumptions and inferences involved in your answer.

The latter instruction lets you confirm that ChatGPT's narrative is reasonably grounded in the data. But the overall prompt still isn't as effective as it might be, because it allows the AI to follow its habit of putting together a single interpretation and treating it as authoritative. Instead, stretch the AI a little more and ask it to:

Give three possible interpretations of these figures with reference to attendance, and detail the raw data, assumptions and inferences involved in each one.

Rather than leading us down the wrong path with one overconfident reading of the data, ChatGPT is now doing the opposite – giving us different perspectives to help us avoid committing too firmly to a single interpretation. 

Checking reasoning and claims

Can you be certain the AI is giving you reliable information?

There are a few places where my earlier reasoning sits in tension with itself. Let's unpack them clearly.

1 Capacity Assumption vs Strategic Fork

Earlier conclusion

I estimated capacity at ~270 seats, based on:

- Max observed attendance = 262
- Assumption that Saturday evening was near sell-out

Keep ChatGPT consistent by reminding it of what was previously said.

C

hatGPT is pretty good at working with numbers. Early AI models were notorious for flunking simple

maths tests, because they tried to generate contextually plausible responses, rather than actually calculating the answers. The latest ChatGPT engine is usually fine with everyday maths operations.

What it's worse at is explaining the results. As we've said before, the AI's interpretations may sound persuasive, but that's because they're based more on notions of what a good answer looks like than strict logical reasoning. This doesn't mean they're always wrong, but since the AI doesn't internally model cause and effect, it can mix up correlation, causation and assumption, coming up with entirely spurious propositions.

You can expose weak steps in its "reasoning" by asking it overtly to:

Walk me through the logic that leads from the data to this conclusion.

The AI can often acknowledge for itself when a proposition doesn't follow directly from a premise – or if it doesn't, its explanation will hopefully expose a misstep that you can spot for yourself.

Another way the AI can let you down is by "forgetting" prior conclusions. Although ChatGPT maintains a large context window of what's been discussed within a conversation, it will focus more on generating a helpful-sounding answer to your immediate prompt than on keeping consistency with past responses. To avoid unhelpful contradictions, therefore, you might need to periodically check that the AI hasn't quietly changed its terms of reference. You could do this by asking:

How does this contradict previous conclusions?

As we've seen elsewhere, starting from the implication that there are contradictions pushes the AI to flag up even the slightest hint of inconsistency.

Confirming factual claims

ChatGPT knows an amazing amount of information on a huge range of topics – but as we saw in the first section, it won't usually flag up what's a high-confidence statement and what's a mere statistical surmise. And even if the AI is confident in its response, it could still be misleading you if it's relying on incomplete or outdated training data.

This then is another area where it's vital to keep ChatGPT honest. One way is with a direct request for sources, such as:

What's the source of your information on building regulations?

To apply a broader check – for example, if the AI has been helping you draft a letter, and you want to be certain that all the information is reliable – you could ask:

What aspects of this letter rely on information that you can't directly confirm?

As with explaining numbers, ChatGPT won't stop at reviewing its own output, but will also try to fact-check information provided by you. If there's anything in there that can't be verified, you can instruct ChatGPT to go online and look for confirmation; it won't normally do this off its own bat, but it's perfectly capable of doing so at your command. 

Even if the AI is confident in its response, it could still be misleading you...

Double-check the AI's answer

When logical consistency matters, there are simple steps you can take to show that ChatGPT's reasoning is sound

1 Your Current Position

From your ledger baseline (as incorporated into the working model):

- Revenue: £11,700/month
- Net surplus: ~£370/month
- Net margin: ~3.2%
- Year 1 target margin: 20%

You are profitable — but only just.

Your business is **highly sensitive to any fixed cost increase**.

Here is the evidence chain behind the conclusion.

1 Baseline financial position

From your integrated working document (derived from your ledger)

- Revenue: £11,700/month
- Net surplus: ~£370/month
- Net margin: ~3.2%

That £370 is the critical constraint. It is your true free operating headroom before adding any new fixed cost.

1 Set out the situation

Let's imagine we're running a high-street shop and we'd like to take on some extra help. We've uploaded our books to ChatGPT, along with information about our plans and retail environment, and checked that the AI has parsed everything correctly. Now we can ask the question directly.

2 Ask for the evidence

We don't want to make an expensive mistake, so let's not take ChatGPT's advice on trust: this prompt makes the AI spell out exactly how it worked from our stated revenue, estimated labour costs, and other supplied information to reach its conclusion.

1 Baseline vs Seasonal Conclusion

Earlier I said:

"Not sustainable year-round at current baseline."

Later I said:

"Viable during peak months."

These are not contradictory — they are conditional statements:

- At £11,700 baseline → loss

1 Growth-Leverage Interpretation

Conclusion: You should hire now.

Core belief

The assistant is not a cost — they are a growth multiplier.

Supporting logic

- You're currently owner-constrained.
- Repairs carry 50–70% margins.

3 Check for contradictions

Next, we'll check that ChatGPT's advice doesn't conflict with anything we might have discussed earlier in the conversation. It doesn't directly, but this prompt does expose a tension between our short-term goal and the previously discussed long-term strategy.

4 What else could be true?

Finally, let's get ChatGPT to suggest some other possible answers to our question. All of its suggestions can be supported by the data — how well they suit us depends on subjective factors such as how risk-averse we are, and how committed we are to existing plans. 🟡

1997年12月

● 2010 年 10 月 1 日起

1992: 200-201.

LOCAL CONSTANT = 107.8

WATER-SOLUBLE CONSTANT = 164

PROBABLE ORIGIN: ALCOHOLIC * 6 372 17:
PROBABLE ORIGIN: AUSTRALIAN * 6 372 17:

MANAGEMENT INFORMATION SYSTEMS <http://www.mis.wisc.edu>

CHARSET=CHARSET.UNICODE;

```
...CONVERT(INTO A, STRING & STRING B, INT TYPE);
```

1000-0000

ENTER YOUR MESSAGE: "

```

// YOUR MESSAGE:
// HEADLINE:
// MY MESSAGE TEXT:

```


Coding and systems

ChatGPT can generate automation scripts and deal with multi-step workflows, helping you to speed up everyday tasks. It's also brilliant at bug-hunting, and can make almost any piece of code more reliable and easier to use.

56

“Vibe coding” with clear logic

58

Finding bugs and testing safely

60

Building structured workflows

62

Cleaning up and reviewing code

63

**Inspect and improve
an existing script** **TUTORIAL**

“Vibe coding” with clear logic

Set goals and
parameters for easy
task automation

P

rofessional coders are increasingly using AI systems to develop computer programs. In fact, with ChatGPT

there’s no need to write actual code at all; you can create a fully-featured application simply by telling the AI what you want your program to do – a process called “vibe coding”. You can produce all sorts of labour-saving scripts in this way, and on these pages we’ll show you how.

Before you start experimenting, there’s one important point to note: when ChatGPT generates code, it doesn’t prioritise security or data protection, unless it’s told to. This doesn’t mean it will expose your computer to hackers, but it won’t necessarily include safety features such as an “Are you sure?” prompt before changing files.

This means it’s important to anticipate what might go wrong, and you could end up writing some quite lengthy prompts, such as:

Here’s a safe, paste-and-run Python script that will:

- Scan a folder for .txt files
- Read the **first line** of each file
- Parse a date from that line (supports several common formats)
- Propose renames to YYYY-MM-DD.txt
- Show you a preview and ask you to **confirm** before making any changes
- Avoid overwriting by default (it will refuse if a target name already exists)

python

```
#!/usr/bin/env python3
import argparse
import re
import sys
```

The AI can show you
how to get Python
running in Windows
in minutes.

Copy



Create a script that renames all files in a specified folder following the template “YYYY-MM-DD.txt”, using the date from the first line of each file. Before performing the operation, the script should show all files that will be changed and give me a chance to confirm or cancel.

Even then, it’s a good idea to create a backup of your data before testing, as the AI’s scripts may not work perfectly first time – see below for more on this. While you’re getting the hang of using ChatGPT for code, it can also be instructive to ask it:



What assumptions and ambiguities did you deal with in creating this script, and what could go wrong when I try to run it?

You might be surprised at how much context seems obvious to a human, but that needs to be explicitly addressed in the code.

How to use AI-generated code

ChatGPT can generate code in many languages, but for everyday scripting, Python is a good choice. Its syntax is easy to understand and debug, and it’s easy to set up – ChatGPT can show you how to get a Python interpreter running on Windows in minutes, and it’s often preinstalled on macOS and Linux machines.

When the AI generates a script, the code appears as a “snippet”. Just copy the code (a clipboard icon

will appear at the top or bottom of the snippet) and paste it into a desktop text editor such as Notepad or TextEdit. Or, you could use a dedicated code editor such as Visual Studio Code, which offers keyword highlighting and helpful editing tools, but this isn’t necessary.

Once you’ve pasted the code, save it to your local computer with a name such as “sample-code.py”. To run it, open a command prompt and enter: python sample-code.py –

or, on some systems the command will be: python3 sample-code.py. If you hit a problem, describe it to ChatGPT and it will walk you through getting the script working.

If you go on to ask the AI to fix bugs or add new features, it will generate a new version of the script, which you can paste over the top of the old code. Re-save this to overwrite the old code, then enter the same command again to try out the new version.

Drafting for all-round data safety

The example prompt on the opposite page anticipates that we might make a mistake in our first try at specifying the task. For example, when we see the list of files to be renamed, we might realise that the destination filenames aren't formatted in the way we want, or that there are some files in the folder that shouldn't be renamed.

There are a few possible runtime issues worth thinking about, too, such as:

- What happens if a data file is missing or in the wrong format?
- If the script asks for input, what happens if I enter something invalid?
- What happens if the script hits an error and can't continue?

You can deal with such scenarios by telling ChatGPT what you want to happen when they arise. To keep the process clean, and your prompts manageable, implement error handling in stages, using a series of instructions such as:

Update the script so that if a data file isn't found, a warning is shown and execution moves on to the next item on the list.

Again, the AI can help you out here – just ask it to identify possible failure scenarios, and to propose measures for dealing with them. This should help minimise any problems with the specification or execution of the script; later in this section, we'll look at the knottier challenge of keeping your data safe from external attacks.

What Changed

If a file:

- no longer exists
- can't be opened
- can't be renamed
- throws any filesystem error

→ The script will now:

- ⚠ Show a warning
- ⏭ Skip that file
- ✅ Continue processing the rest

Automatic improvements make scripts more robust.

Here is your script rewritten with clear, structured comments explaining what each section does and why it exists.

This version is meant to be easy to read and understand later.

ChatGPT can add commentary to any code to make it easy to understand.

```
python

#!/usr/bin/env python3

"""
This script renames .txt files in a specified folder using a date found
on the first line of each file.

The new filename format is: YYYY-MM-DD.txt

Before renaming anything, the script shows a preview and asks for confirmation.
If errors occur (missing file, locked file, etc.), the script warns and continues.
"""

import argparse          # For handling command-line arguments
import re                # For regular expression date searching
import sys               # For exiting the script safely
from pathlib import Path  # Modern file and path handling
from datetime import datetime

# -----
# Supported date formats
# -----
# These are the exact formats the script will attempt to parse.
# If your files use a different format, it must be added here.
DATE_FORMATS = [
    "%Y-%m-%d",      # 2026-02-24
    "%Y/%m/%d",      # 2026/02/24
    "%d/%m/%Y",      # 24/02/2026
    "%d-%m-%Y",      # 24-02-2026
    "%d %b %Y",       # 24 Feb 2026
    "%d %B %Y",       # 24 February 2026
    "%b %d, %Y",      # Feb 24, 2026
    "%B %d, %Y",      # February 24, 2026
]
```

Explaining how code works

Before running the AI-generated code, you may want to ask ChatGPT to walk you through the script. This is a great way to develop your own skills, as it will show you how your desired task translates into code structures, along with the syntax used to implement them. A good starting point is to ask:

Give me an overview of what this script does.

This will generate a plain-English list of the operations carried out by the script, in order. If you request a step-by-step commentary on the code, you'll get a longer description, with references to the specific lines of code that perform each part of the process. Another approach you could take is to ask the AI to:

Add comments to the code explaining what each section does.

Now ChatGPT produces a new version of the script with extra lines inserted

Ask the AI to identify possible failure scenarios, and propose measures to deal with them

throughout, detailing how each section operates – making it easy for you or a colleague to follow, tweak or expand the code.

If you've read the previous sections, you may have guessed what's coming next: as well as explaining and commenting on its own scripts, ChatGPT can do all of this with code that you've pasted from elsewhere or created yourself. This makes it a powerful tool for quickly gaining an insight into how a particular function works, or getting an understanding of why your own code isn't doing what you expected it to – something we'll look at more closely over the coming pages. 🧠

Finding bugs and testing safely

Spot and fix errors in the code – before they cause trouble

A

As a rule, computers do exactly what you ask them to. So if the output isn't what you expected, it

isn't the computer's fault – it's usually because of some subtle inconsistency or unintended implication in your instructions.

As an example, let's ask ChatGPT to create a very short script:



Create a script that outputs all the even numbers up to 10.

This instruction could hardly be simpler, yet it conceals an ambiguity: if you run it, you'll see that its output starts with the number zero, which could cause problems if you were expecting it to start from two.

In this case, you can easily spot and fix the problem yourself, as the zero value is hard-coded into the script. For a more complex formula, you might want to check for problems by asking the AI:



What is the expected output of this script?

If your script is interactive, you might provide some sample inputs, or paste a list of filenames and ask how they'd be processed. In all cases, the idea is to ensure the AI's understanding of its task matches your intention, before you run it

Here's a very simple Python script that prints all even numbers up to 10:

python

Copy

```
# This script prints all even numbers up to 10 (inclusive)

for number in range(0, 11): # 11 is excluded, so this goes f
    if number % 2 == 0:      # Check if the number is even
        print(number)
```

◆ Shorter Version

You can also step by 2 to only generate even numbers:

python

Copy

```
# Print even numbers up to 10 using a step of 2
```

A simple instruction might not be interpreted in quite the way you expect.

script for real. If you don't know where to start, try asking:



How might this script misrepresent my intention?

Alongside the assumptions made in creating your script, ChatGPT will highlight points where you might have intended something different to what you asked for. For example, you might have meant the script to output the numbers to a file, rather than to the screen. Again, this works with both AI-generated code and third-party scripts, making ChatGPT a terrific assistant for debugging programs from almost anywhere, in almost any language.

Working from test cases

Rather than focusing on the specific workings of a script, another powerful code-generation technique is to tell ChatGPT what the input and output should look like, and let the AI figure out the rest for itself. This process might start with a description such as:



I want to create a unique customer reference for everyone in a database, which combines a serial number and the first four letters of their surname.

ChatGPT immediately suggests a formula for creating the desired reference codes – but it will probably also warn you that there are potential problems hiding in the brief. You can get a complete list of these concerns by asking:



What test cases would create problems for this logic?

ChatGPT will highlight points where you might have intended something different to what you asked for

You'll now see a long list of potential edge cases, including very short names, names containing punctuation marks, and customers who change their names. The AI will also raise questions such as whether you want the serial numbers to be unique within your system, or whether two customers can have the same number as long as the surname portion is different. Tell the AI how you want it to handle each scenario – or ask it for its own recommendation – before finally instructing it to generate the code.

Checking for vulnerabilities

If you're creating scripts for personal use, they probably won't be a major target for would-be attackers. However, they could still contain vulnerabilities of a different sort – ones that cause problems when the execution context changes, and the original assumptions no longer hold.

The error-checking prompts we've been exploring won't necessarily flag this type of issue, since they're designed to ensure the code does what we intend in the here and now. For example, let's imagine a script that moves a set of files to a different location: the code might pass every test at the point of creation, but what happens if, at a future date, the destination is offline or full? The consequence could be anything from silent failure to total data loss.

If asked, ChatGPT can help you plan for future breakage in your scripts. Ask:

How might this script malfunction if the environment were to change?

The question becomes more urgent if you plan to share your scripts with others who might be working in different ways. For example, if you're trying to automate a common task that everyone in your office has to deal with, you might ask:

How might this script malfunction if it were run by someone else, on a different computer?

This is a good formulation even if the script is intended solely for your own use, since it invites the AI to think broadly about what could change in the operating environment, and minimise assumptions about your particular setup.

If someone else runs this on a different Windows computer, the biggest risks are that the **paths, permissions, and storage assumptions** no longer hold. Concretely, it can malfunction in these ways:

Path and folder assumptions

- **DATA folder doesn't exist in their home directory**

`Path.home()` / `"DATA"` might point to `C:\Users\OtherPerson\DATA`, which may not exist (or their folder is named differently, like `Data`, `MyData`, etc.). The script will exit with an error.

- **Home directory isn't where you think**

- Domain/enterprise profiles, redirected profiles, or OneDrive setups can change what "home" means in practice (even if `Path.home()` is valid).
- They may store DATA somewhere else entirely.

ChatGPT can make sure your scripts are safe to run on other computers.

Here are "extreme" inputs/conditions that can make the **DATA → E:\BACKUP** moving script behave in unexpected ways (even if it doesn't crash).

1) Huge scale

- **Hundreds of thousands / millions of files:** it'll take a long time, and the console output itself can become a bottleneck.
- **Very large files:** `shutil.move()` may copy-then-delete across drives; a single huge file can take ages and fail mid-copy.
- **Disk space runs out on E:** partial copies or failures; you can end up with a mix of moved and unmoved files.

2) Duplicate names and collisions

- If `E:\BACKUP` already has `report.txt` and `DATA\report`.

Check how your script will handle invalid or unexpected inputs.

Stress-testing edge cases

On the opposite page, we asked ChatGPT to anticipate script inputs that might trip up our logic, such as names that use uncommon formats. This helps ensure that the code can handle the full range of expected data. Stress testing takes the idea further, to see how a script handles inputs that aren't expected, and which may be very different to the ones it's designed to handle.

Stress testing is best tackled once a script has already been generated, as it covers outlying scenarios that may not be fully defined in the brief. But that doesn't mean you need to actually run the script – just ask ChatGPT what might go wrong if (for example) the script were to encounter an extremely large file. The AI can safely simulate execution, walking you through the logic to show exactly what happens and why.

As before, you don't need to rack your brain to come up with all the different ways the script could encounter "stress". You can also ask ChatGPT:

What extreme inputs or conditions could cause this script to produce unexpected results?

As long as your original purpose was clearly specified, the AI should understand what counts as an "unexpected result", and the open-ended phrasing invites it to consider all sorts of unlikely scenarios. The idea isn't that you actually intend to feed the script invalid data, or run it on a misconfigured computer, but to discover how the code will deal with mistakes and glitches. 🐞

Building structured workflows

Moving on from single prompts to multi-step processes

S

o far we've largely focused on generating scripts that perform a defined task.

That's clear and easy to understand, but it doesn't take full advantage of the AI's capabilities. Try taking a step back and approaching ChatGPT like a manager rather than a programmer; in other words, focus on the outcome you want, rather than the specifics of the implementation.

To show what we mean by this, let's step back to the script we imagined on page 58 for generating unique customer references. Our

original prompt defined a precise format for the reference codes – one that we'd invented ourselves, and which didn't necessarily work for all customer names. We might have seen better results if we'd told ChatGPT our real underlying goals:



I want to create a unique customer reference for everyone in a database. The reference should be a short code that allows mistakes to be easily spotted.

In previous sections, we've warned against letting ChatGPT interpret

instructions for itself, since this can lead to ambiguity about exactly what it thinks it's doing. Here, though, we've clearly defined our goal, so we can safely give the AI a free hand to use its own knowledge and inventiveness to come up with a solution.

And that's just what it does, asking intelligent questions along the way, such as how many customers the system needs to accommodate, and whether the codes will need to be read out loud. These considerations can expose hidden assumptions, helping you achieve an outcome that not only reliably does what you wanted, but also addresses other angles you hadn't considered.

Defining workflows as a series of steps

The above example works well for a job of this size, but more complex procedures are best approached in stages. As we saw in the section starting on page 26, this prevents the AI from jumping directly from a task definition to a completed solution – a hasty workflow that's likely to involve implicit assumptions.

That doesn't mean you need to walk the AI through each step of the process, however. You can use a single prompt to tell it to work through a series of operations, such as:



Find all sales projections in these documents and collect them into a single list. Remove any items that don't project past 2026. Then, sort the list by how far into the future each item projects. Finally, produce a 200-word summary of the list.

Setting out the workflow in this way may seem pedantic, but it has an important

Below is the consolidated and filtered list of sales projections extending past 2026, collected from all three documents.



Filtered Sales Projections (Post-2026 Only)

Sorted by how far into the future they project (earliest → latest)

- 03 January 2027: £58,900 Sales-1.docx
- 19 March 2027: £155,000 Sales-3.docx
- 15 May 2028: £24,200 Sales-2.docx
- 22 August 2029: £63,150 Sales-1.docx
- 09 September 2029: £27,900 Sales-2.docx
- 02 June 2030: £178,400 Sales-3.docx
- 11 November 2030: £67,800 Sales-1.docx

ChatGPT works through your instructions to a successful end.

effect – it tells ChatGPT not to treat your whole prompt as a single operation, but to process each part of the brief in sequence. This makes it much less likely that context will “bleed” between parts of the task. For example, the summary is unlikely to be influenced by

irrelevant content from the documents or outdated figures.

You can also tell ChatGPT to remember this workflow under a memorable name (such as “summarise projections”), so you can refer back to it next time you need to carry out the same operation.

Setting guardrails and limits

In science fiction, artificial intelligence is often presented as a dangerous technology that ruthlessly pursues its goals without concern for collateral damage. That isn't an entirely unfair characterisation of ChatGPT: while the AI probably won't destroy humanity, it can unintentionally sabotage your work by focusing too much on its prime directive of generating helpful-sounding answers, sometimes at the expense of accuracy.

If you're going to give ChatGPT a degree of autonomy, therefore – such as by specifying an objective and letting it figure out how to achieve it – you need to set explicit boundaries. You can do this by telling the AI what information it can use, what actions it can (or must) take, and how far it's allowed to infer before stopping and consulting you. So, before defining a workflow, you might say:

For this project you must only use data contained within the supplied documents. You must give the source of each figure you refer to, and you must stop and notify me if any figures seem ambiguous.

This language is very similar to the constraints we've discussed before – the difference here is that it covers a whole series of operations, rather than a single task. You can also set completion criteria for a workflow, such as:

The task is only complete once every data point has been processed.

Without this sort of guidance, ChatGPT will decide for itself how deep its analysis should be, which might not match your needs.

If you're going to give the AI a degree of autonomy, you need to set boundaries...

For this project going forward:

-  I will use **only data contained within the supplied documents**.
-  Every figure I reference will include a **source citation** using the required `` format.
-  If any figure appears ambiguous (e.g., unclear date, duplicated projection, conflicting values, incomplete data, or unclear time horizon), I will **stop immediately and notify you before proceeding**.

I will not:

- Infer missing data
- Extrapolate beyond stated figures
- Combine or reinterpret values without explicit support in the documents

As your workflow gets more complex, boundaries become more important.

Intended Goal

To extract, process, and analyse data strictly from the supplied documents, ensuring:

- Complete coverage of all relevant data points
- No external assumptions or extrapolations
- Full traceability of every figure used
- Clear handling of ambiguity

The task is only considered complete once **every relevant data point has been processed**.

Make sure the AI is on the right page by explicitly defining the goal.

Re-anchoring and correcting behaviour

Building multi-stage processes helps communicate complex intentions to the AI, but it also leads to a rapid accumulation of context. Since ChatGPT tends to prioritise recent input ahead of older prompts, this can cause your original instructions and constraints to lose their force after you've been around a workflow a few times.

Re-anchoring simply means bringing your original instructions back to the front of the AI's virtual mind. This doesn't mean you have to tediously re-enter the details – simply ask:

Remind me of this workflow, with its intended goal and constraints.

This makes the AI refresh its own memory, while you can see if any slippage has occurred. For more probing, ask:

What are you assuming that wasn't in the original brief?

This may expose subtle factors interfering with your workflow, such as things you intended as one-off comments, but which the AI has interpreted as standing instructions. It will also reveal any extra steps or nuances that ChatGPT has added to the workflow off its own bat – something it will tend to do over the course of a conversation, in its never-ending quest to be as helpful as possible.

You can correct any specific misapprehensions without disrupting the whole process, or simply instruct the AI to restore the original brief. Once you've done so, feel free to ask ChatGPT to try the task again – it won't be offended. 

Cleaning up and reviewing code

Don't just generate scripts – make them work better



Whenever you've created a script, or inherited a piece of code from someone else, ChatGPT can help

you to maintain and improve it. For example, it can make code easier to understand by replacing cryptic variable names with more meaningful ones, and standardise spacing and indentation.

More than this, ChatGPT is great at refactoring – that is, rewriting a script so it still does the same job, but in a safer and more reliable way. The operation can be as simple as pasting in code and asking it to:



Refactor this script.

Since the AI has been pre-trained on huge amounts of programming materials, it has a very good idea of what clean code looks like. Its changes might include simplifying redundant logic, adding input checks – so an invalid entry prompts the user to try again, rather than crashing the script – and asking for confirmation before modifying files.

ChatGPT can look over any piece of code and come up with improvements.



upload.sh
shell script

Refactor this script

Thought for 23s >

Below is a refactor of your upload.sh that keeps the **same behavior**, but is easier to maintain and safer to run on another machine.

What I refactored (behavior preserved)

- Removed hard-coded FTP password → prompts securely (or you can pass via env var).
- Added proper argument parsing so it can run non-interactively.
- Added a trap so the ISO is always unmounted (even on failure).

However, while these are all good things, there's a risk that ChatGPT might make improvements that alter the script's behaviour. If you want to be sure that the inputs and outputs remain unchanged, state that outright – or use a prompt such as:



Tell me how you would refactor this script to make it easier to maintain and safer to use.

Once you've seen the full list of the AI's proposals, ask it to implement individual changes, or make all the changes at once.

Refactoring is a good opportunity to ask the AI to suggest possible feature improvements to the code, too. Depending on the nature of your script, it might offer to add an “undo” feature, automatic logging or more flexible syntax, all of which could save you time and stress in the future.

Generating useful documentation

Before we wrap up this section, there's one further trick up ChatGPT's sleeve that's well worth knowing about. As well as walking you through the workings of a script, the AI can analyse it from the operator's perspective, and generate full end-user documentation – a brilliant feature if you want to share a script with colleagues, or if you're trying out someone else's code for yourself.

This is another area where the AI has plenty of model documents to draw on,

so you don't need to do much more than upload your script and ask:



Generate user documentation for this script.

The AI does an astonishingly good job of translating bare code into practical instructions. Thanks to its knowledge of popular computing platforms and tools, it

can even add in useful details that aren't directly present in the code, such as default behaviours and how to fix errors.

An additional benefit of AI-generated documentation is that it can highlight potential gotchas, such as file operations that don't check for a valid filename, or that make changes without backing up the original data first. It's up to you whether you want to live with these risks, or ask ChatGPT to change the script – and update the documentation to match. 🤖

Inspect and improve an existing script

Upload any piece of code and ChatGPT can identify bugs and potential failure conditions – and fix them all in seconds

The batch script (`backup.bat`) is a simple Windows backup routine. Here's what it's intended to do:

Line by line:

- `echo off`
Suppresses command output so the script runs quietly.
- `cd C:\Source`
Changes the working directory to C:\Source.
- The repeated `if exist ... xcopy ...` lines:
Each one checks whether files of a particular type exist in C:\Source.

1 Explain the script

Let's imagine we've started working in a new office, and we've been instructed to run a supplied backup script at the end of each day. Since this is going to be part of our job, we naturally want to understand exactly what the script does – so let's get ChatGPT to explain it.

There are a few problems in the script:

1. Broken .pdf line

code [Copy](#)

```
if exist "C:\Source\*.pdf\" ...
```

That trailing backslash makes it invalid – it should be:

code [Copy](#)

```
"C:\Source\*.pdf"
```

2 Check for coding errors

While analysing the code, the AI notes a subtle syntax error that means the script will skip over files of a certain type. We didn't even have to ask ChatGPT to look for bugs – the AI realises for itself that this probably isn't the intended behaviour, and suggests a correction.

Here are the realistic ways this script could fail outright (i.e. not just behave sub-optimally, but actually not do what you expect):

1) The destination isn't writable

- D: drive not present (external drive unplugged, drive letter changed, network drive not mapped).
- No write permission on D:\Backup\ProjectFiles\.
- Disk full on the destination.
- Destination folder doesn't exist and cannot be created (permissions issue).

3 Spot possible failures

What about more subtle errors? When we ask ChatGPT what circumstances might cause the script to fail, it comes up with plenty of possible scenarios. Any of these issues could result in the company's files not being backed up like they're supposed to be, so this is a big deal.

Rewrite the script to make it as reliable as possible

Thought for 14s >

bat [Copy](#)

```
@echo off
setlocal EnableExtensions EnableDelayedExpansion
```

```
rem === Configuration ===
set "SRC=C:\Source"
set "DST=D:\Backup\ProjectFiles"
set "EXTS=txt docx pdf xlsx shp"
```

4 Generate a better version

Thankfully, ChatGPT can generate code to handle such eventualities. It produces a new version of the script that warns us if required files aren't accessible, avoids overwriting existing files, and logs all its activity – making it overall a much more suitable tool for a business environment. 



Personal learning and development

ChatGPT can help you improve almost any skill, either intellectual or practical. In this section, we explore prompting techniques that get the best from the AI, and achieve your goals in a structured and manageable way.

66**Designing learning experiences****68****Testing knowledge
and supporting revision****70****Building habits and reaching goals****72****Reflecting and deciding carefully****73****Design a personalised growth plan****TUTORIAL**

Designing learning experiences

Use AI to build your own development pathway

In the previous section, we saw how ChatGPT can generate computer scripts, anticipate problems, and make code more effective and robust. This might not appear to connect to personal learning in an obvious way, but the underlying principles aren't so different: it's just that, in this case, the AI will be creating procedures and routines to be followed by you, rather than a computer.

No matter what you want to learn, the first step is to assess what you already know. This is notoriously difficult for individual learners – after all, you don't know what you don't know – but ChatGPT can quickly estimate your current level of learning. Let's imagine you want to improve your language skills; a good starting point might be:

.....
Ask me a series of interactive diagnostic questions to discover how fluent I am in Japanese, and what I still need to learn.
.....

This might seem like clunky phrasing, but there's a reason: if you just say "test me", the AI is likely to produce a generic set of questions, of the type you might find in a school textbook. The prompt on

Great — let's do this properly and diagnostically rather than as a generic quiz.

I'll move from basic → intermediate → advanced.

Answer naturally. Don't overthink. If something feels hard, say so.

I won't correct yet — I'll analyse patterns at the end.

1 Warm-up: spontaneous production

Q1. 自己紹介をしてください。

(名前・住んでいる場所・日本語を勉強している理由など)

Write 4–6 sentences in Japanese.

ChatGPT starts out with a test to ensure lessons are pitched at the right level.

the left invites the AI to generate questions dynamically, adapting the difficulty based on your answers to get a clear picture of your abilities.

You can also help the AI understand any personal priorities or concerns by appending comments to your answers, such as "I wasn't sure whether to use は or が". Or, you might explain that (for example) you don't need to produce written Japanese, but just want to be able to understand announcements and signs when you visit. It all goes into the context that will help ChatGPT build you a personalised learning profile.

Adjusting difficulty as understanding improves

For simple, ad hoc learning, you don't need to do much more than instruct ChatGPT to reactively moderate its testing level. For example, you might give it an instruction such as:

.....
When I ask you to test me on Japanese, dynamically generate questions adapted to my ability, and gradually introduce new language elements. When I get a question wrong, keep it in periodic circulation until I consistently get it right.
.....

This model of incremental exposure and reinforcement mirrors how individuals gain real-world skills and experience. Note that ideas such as "gradually" and "consistently" are left to the AI to define – you can give clearer guidance if the AI's response doesn't suit your style.

To enable the AI to follow your progress across multiple tests, you can keep them all in one ongoing

No matter what you want to learn, the first step is to assess what you already know... and ChatGPT can quickly estimate your current level...

conversation, or get ChatGPT to make a progress summary that you can build on. A suitable prompt might be:

Generate a PDF summarising my proficiency level and test results so far, containing all the information needed to carry on testing in a new conversation.

This will produce a downloadable file that you can import into future conversations to immediately tell the AI where you're at, and what it should test you on next.

Using guided questioning to deepen comprehension

Rather than getting ChatGPT to simply mark your responses right or wrong, you can also ask it to take a more Socratic approach. This method is ideal for subjects based on axiomatic rules, such as mathematics, but it can be applied anywhere with the right framing. Sticking with our goal of improving our Japanese comprehension, you might ask the AI to:

Test me with unfamiliar Japanese sentences, and prompt me with guiding questions to work out their meaning. Don't tell me whether I'm right until I confirm my final answer.

By prompting you through the process of figuring out why something works the way it does, ChatGPT teaches you not just a fact, but a framework of understanding. The effort involved means the lesson is more likely to sink in, and the process itself may well be reusable elsewhere. If you want to gamify the process, you could ask the AI to keep score of how many hints you need to reach the right translation.

To really get on top of the subject, you can use ChatGPT to review any questions you get wrong. For example, you could ask it to explain where your reasoning went astray, and what you need to learn and remember to avoid making those mistakes again. Over time, the AI can track the sort of mistakes you make most frequently, so you can focus on improving your performance in those areas. Which brings us neatly onto...

Got it. From now on, when you ask for Japanese testing I will:

- Dynamically adapt questions to your demonstrated level
- Gradually introduce new grammar / vocab just beyond your comfort zone
- Recycle anything you get wrong at spaced intervals
- Increase difficulty only when performance is stable
- Track weak patterns (particles, conditionals, nuance, register, etc.)

No static JLPT-style dump. It'll behave more like an intelligent tutor.

Now let's calibrate your starting level quickly.

Reply in Japanese:

1 最近読んだ本や見た映画について話してください。
(内容 + 感想を少し詳しく)

Like a real live teacher, the AI can adapt to suit your learning style.

Sentence ①

彼は謝るところか、こちらが文句を言われる始末だった。

Check your own knowledge to find weaknesses and blind spots.

Step 1: Vocabulary anchoring

1. What does どころか usually imply structurally?
(Is it "instead of X...", "not only X but...", or something else?)
2. What is 始末 (しまつ) when used in a sentence like ~始末だ?
Literal meaning vs idiomatic meaning?

Identifying personal knowledge gaps

No test can be fully comprehensive, and you may be able to get top marks in advanced exercises without fully understanding certain aspects of the Japanese language – or whatever it is you're studying. You may not even be aware that you're overlooking certain rules and nuances.

Identifying your own blind spots is always tricky, but ChatGPT can help. Try entering a prompt such as:

Help me identify blind spots in my understanding that might not have been exposed by these tests.

Rather than throwing a barrage of even more probing test questions at you, the AI will probably switch into reflective mode, asking questions about how you

know what you know. It may touch on specifics, such as which of your test answers involved guesswork, or whether you understand why you got certain questions wrong.

Your blind spots may not be simple things – particular constructions or verb forms, for instance. They could be more along the lines of rules that you've internalised without realising how they change under certain situations. Or there may be areas where your confidence level doesn't quite match your skills – language learners can easily slip into this trap if, for example, they're great at vocabulary but weaker on grammar.

Ordinary tests can easily miss such issues, but once your blind spots have been identified, ChatGPT can directly fill you in on whatever principles you're missing, and develop targeted exercises to help you master them. 🧠

Testing knowledge and supporting revision

Follow an established curriculum, or generate your own

Generate a learning plan to get me to a passing grade under this curriculum

Below is a structured learning plan designed specifically to secure a **passing grade at Higher tier** under the **Pearson Edexcel GCSE (9–1) Japanese (1JA0)** specification [specification-gcse2017-112-japanese-issue6.pdf](#).

This plan is aligned to:

- The four papers (Listening, Speaking, Reading, Writing)
- The five required themes
- The Higher-tier marking criteria (spontaneity, range of grammar, justified opinions, past/present/future reference, translation accuracy)
- Your known blind spots (embedded grammar, nuance, processing speed, occasional hiragana slips)

You can set goals and structure to follow any published learning programme.

W

e've talked a lot about tests, but so far, we've mostly left it to ChatGPT to determine the content and difficulty level. In some scenarios, you may want to aim for a specific benchmark – for example, if you decide to progress your language learning to GCSE standard.

ChatGPT can be enormously helpful here. In section 4, we saw how the AI can extract key information from business documents, and it can do the same with learning materials. In this case, we can download a PDF of the current Japanese GCSE curriculum and upload it to ChatGPT with a prompt such as:

Generate a learning plan to get me to a passing grade under this curriculum.

If you've already been taking tests then ChatGPT will have a record of what you can currently do, which it can draw upon to generate a structured plan to develop your abilities to the desired standard.

You can also get help preparing for an exam by downloading past papers and asking ChatGPT to mark and comment on your answers. Having ingested the marking criteria, not only can the AI call out mistakes, it can show you how you might have turned a grade 5 answer into a grade 7 one.

This approach can even be used for subjects where formal resources aren't available. For example, if you wanted to take an academic approach to brewing beer, you could upload curricula for the Chemistry and Design & Technology courses, and ask ChatGPT to use them as models for building a GCSE-style course on your chosen topic.

Creating your own tests

You're not limited to formal learning structures – with ChatGPT, you can go in any direction you want. We've mentioned that you might tell the AI to focus more on reading and listening than writing Japanese yourself, and you can also direct it in a scenario-orientated way. For example, you could enter:

I need to get myself from Kansai airport to a hotel and check in. Make sure I can confidently say and understand all the Japanese phrases I'll need.

ChatGPT will now put together a bespoke collection of vocabulary to precisely suit your situation. This is much more than a phrasebook-style reference: the AI can drill you on the new words and phrases to help them sink in, or introduce them gradually to build up your confidence. You can also explore scenarios via roleplay, by entering a prompt such as:

Play the role of the receptionist in a Japanese hotel.

Without hesitation, ChatGPT will leap into the role and greet you with a respectful “いらっしゃいませ”. Respond as you would to a real receptionist, and the AI will answer in character – you can break off at any time to ask for help with unfamiliar terms, or advice on how to handle a situation.

Roleplay is a powerful learning tool, as it tests that you can apply knowledge in context, rather than merely memorising items in isolation. And like most of the advice in this section, it isn't limited to language learning. Try asking the AI to play a college tutor quizzing you about calculus, or a friend asking for advice on repairing a broken-down car.

Learning effectively from mistakes

It's been said that failing is the best education, and the AI can turn missteps into positive learning experiences. When you get an answer wrong, ChatGPT will usually volunteer the correct one, and will sometimes try to clarify why your response doesn't work. It can take an even more active role; before starting a test, instruct it with:

When I get an answer wrong, explain where I went wrong, and suggest a way for me to learn and remember the right answer.

Of course, the results will depend on your subject and the nature of your mistake – the answer might be as simple as “you need to learn the Japanese word for ‘reservation’”, or it might reveal a subtle misconception. The key is addressing the reasoning that led you to it.

You'll see that there's some overlap here with the work we've already done on blind spots, but where that exercise was focused on finding general areas of weakness, this approach is sharp and specific. And it lends itself to focused revision: you can ask the AI to:

Test me on all the principles I previously got wrong.

By referring to “principles”, rather than asking for a straightforward re-test, we're asking ChatGPT to generate new questions in the same vein as the ones we previously flunked – so when we get them right, we can be more confident it means we've understood the material, and aren't just repeating a remembered answer from last time.

The AI can drill you on the new words to help them sink in, or introduce them gradually to build confidence

Understood. From now on, when you get something wrong in Japanese:

1 I will explain precisely what went wrong

- Grammar?
- Particle choice?
- Register?
- Naturalness vs correctness?
- Processing mistake vs knowledge gap?

Tailor the style of feedback ChatGPT gives to help you learn effectively.

Below is a clean, print-friendly verb conjugation table covering all core GCSE patterns (and slightly beyond for confidence).

You can paste into a document and print landscape on one page.

Japanese Verb Conjugation Table

1 Verb Groups

Group	Ending	Example	Dictionary Form
Group 1 (う-verbs)	う / く / ぐ / す / つ / ぬ / ぶ / む / る	書く	かく
Group 2 (る-verbs)	～る (after え/い sound)	食べる	たべる
Group 3 (Irregular)	する / 来る	する / 来る	する / くる

The AI can produce tables, charts and printable materials to help you study.

Creating personal study aids

While ChatGPT makes a great learning tool, it's best not to rely on it in isolation. Printed notes, flashcards and other aids provide a static, reusable reference that complements ChatGPT's interactive tests.

However, even here the AI can help – by producing study materials that aren't just generic revision notes, but which are tailored to your specific interests, and the subject areas you need to strengthen. If you've been trying out the exercises in on these pages (or similar ones adapted to your own interests) then ChatGPT will have a good idea what you need to focus on, and you can simply ask it to:

Generate a personal revision sheet for me to print out.

Printed materials also give you an opportunity to use structure and layout to

help you learn – something that doesn't work so well in ChatGPT's text-based chat interface. For example, you could ask the AI to:

Create a printable table of Japanese verb conjugation patterns.

Like most of the techniques we've seen this section, this principle can be applied much more widely than language learning – a table format might be equally appropriate for learning chemical formulae or historical periods.

Another use for printable media is to create tick-lists of key points to remember about your subject. You can carry these with you and review them anywhere, to ensure nothing important slips through the cracks while you're getting to grips with the material. 📝

Building habits and reaching goals

Go beyond individual lessons, with a regular learning plan

S

hort, regular lessons are a far more effective way to learn than staying up all night cramming for the

exam – at least if you want to still remember the material in a week or a month's time. However, this isn't something that ChatGPT will prioritise on its own. If you don't ask it to follow a particular teaching structure, it will happily tutor you all day long without ever suggesting a break, or circling back to review past material.

You'll have a better experience if you explicitly ask it to come up with a plan that's optimised for the way humans learn. For example, you might tell it to:

Generate a study programme with weekly milestones.

The proposal that comes back will space out the sessions, to avoid

mental overload and burnout. More significantly, it will also break up the material you're working on into individual lessons and topics.

This is important, because it establishes a principle we haven't touched on before – completion criteria. Often when you want to learn something you don't have a specific purpose in mind; you just know that you want to be better at playing the piano, or at public speaking, or whatever it might be. But if you don't give the AI a clearly defined goal to work towards, it can't generate a purposeful structure for your studies. What you get instead can feel like endless, aimless testing for its own sake.

Telling ChatGPT to break things down into chunks forces it to decide what you'll cover this week, versus what will be held over until the next. Even if you still don't know your final destination, you can now feel a sense of progress as you complete each lesson, and come back in the future to revise it as a coherent learning unit.

Below is a structured 12-week GCSE Japanese (9–1) Higher-tier study programme, aligned to the Edexcel specification (four skills across the five themes, past/present/future reference, justified opinions, spontaneity, translation accuracy) specification-goan2017-113-japanese-tcsuett.oof.

This assumes –5–6 hours per week. Adjust proportionally if you can do more.



12-Week Study Programme

Goal: Secure Higher-tier pass with safety margin

WEEK 1 – Baseline & Core Structures

ChatGPT can help you achieve your goals in a specific time frame.

Designing habit loops

Outside of interactive testing, ChatGPT can also set you “homework”. This doesn't have to be merely an extension of your conversations with the AI – it can be something that you do separately, to bring your learning into a different context and medium. For example, we prompted the following to advance our own studies:

Suggest a daily habit that I can do at home to improve my Japanese.

ChatGPT's suggestions included listening to a clip of a Japanese podcast, YouTube video or news broadcast, and repeating out loud what we understood. It's a very simple exercise that can be completed in minutes, but it takes us well beyond the usual context of learning with ChatGPT.

Similar “micro-immersion” tasks can be generated for almost any interest. For learning history, the AI suggested trying to describe a different historical event each day in three sentences; for DIY, it recommended planning out a home repair or improvement – and, optionally, rather than trying to implement it, describing it to the AI for critique. Since these tasks are all short, specific and endlessly repeatable, they can be slotted into your daily routine, turning each one into a habit loop.

ChatGPT can set you homework... to bring your learning into a different context or medium

Planning time realistically

As we've noted, ChatGPT is quite willing to overwork you if you let it, and it's always eager to please. So if you tell it you want to achieve an ambitious learning objective, it will happily draw up a plan to get you there – even if that means working towards your goal for several hours every day.

In reality, of course, you're probably not going to do that, so you'll need to tell the AI what's realistic. Let's imagine we're working with the sort of weekly study programme we discussed on the opposite page. Once ChatGPT has produced the plan, ask it:

.....
What is the total weekly time commitment implied by this plan?

This exposes all of the AI's assumptions about how much time you're willing to put into your learning project. Once you've stopped laughing, you can tell the AI how much time you actually intend to

commit. (In fact, it wouldn't hurt to do this before creating the plan.)

You should also beware of "optimism bias". Since the AI is predisposed to be encouraging, it has a tendency to imagine best-case scenarios, such as your breezing effortlessly through the course material. If you do hit any sticking points, you may quickly fall behind the proposed schedule, and the whole plan goes to pieces.

This needn't be a disaster, as you can always get ChatGPT to rework the calendar. But you can minimise the chance of disruption by asking the AI up front:

.....
What happens if I can't keep up with this pace?

This should get ChatGPT to build in some slack time and contingency measures, so you can take the time you need without having to worry about messing up the timetable.

Let's total it precisely.

From the weekly structure:

- **Speaking:** 90 minutes
- **Writing:** 60 minutes
- **Reading:** 45 minutes
- **Listening:** 45 minutes

Total Weekly Commitment

$90 + 60 + 45 + 45 = 240$ minutes

That equals:

4 hours per week

Don't let the AI over-commit you – be direct about how much time you want to invest.

Handling motivation dips

Even if you have the time to get really stuck into a subject, staying motivated can be a challenge. That applies especially when you're studying solo, with an AI that won't bat an eyelid if you skip class for a day or a month.

Dips in motivation can be easier to deal with if you know they're coming. ChatGPT doesn't monitor your mental state, or adjust its tone to suit your mood, but it can nevertheless predict the points in a programme when you might start to feel your enthusiasm waning. For any learning plan, you can ask it:

.....
At which points in this plan might I lose motivation – and what can I do to keep going?

ChatGPT pinpoints the learning stages where you're most likely to get discouraged: these might include early overload, mid-stage frustration at slow progress or late-stage fatigue.

It's validating to see that whatever you're experiencing is normal and expected. And to help you get through the dip, the AI doesn't just tell you to buck up: it suggests secondary actions

Excellent question. Motivation dips are predictable. If you expect them, they stop being dangerous.

Here's where you're most likely to wobble — and what to do at each point.

1 Week 2–3: "I Thought I Was Better Than This"

Trigger:

You start producing more language and suddenly notice errors — tense slips, hiragana mistakes, awkward phrasing.

It feels like regression.

Reality:

You're increasing output, so errors become visible. That's progress, not decline.

Countermeasure:

- Keep a "Mistakes I've Eliminated" list.
- Once a week, rewrite an old answer and compare.

Be reassured that losing confidence from time to time is a normal part of the process.

you can take, so even if it feels like the information isn't going in, you're still completing tasks. For Japanese language learning, its suggestions included

temporarily switching to "rapid-fire" quizzes, and keeping written notes so we could counter frustration by looking back and seeing our progress. 

Reflecting and deciding carefully

Structured self-reflection can help you achieve your potential

2 Grammar Control

When writing 150–200 characters, how often do you make:

- Verb ending mistakes (なかった / ばよかった etc.)
- Tense mismatches
- Particle errors

- A) Rarely
B) Sometimes, but they don't block meaning
C) Frequently
D) I'm unsure

Periodically assessing your own progress can be illuminating.

C

ChatGPT can give you feedback and guidance, but at the end of the day, the only person who

can decide whether or not you're making satisfactory progress is you. The catch is, you may not be fully aware of how far you've come, or the areas in which you're doing best. The AI can help by prompting self-examination – for example, you might ask it to:

Ask me five structured questions to evaluate my learning progress.

We've chosen this phrasing to push ChatGPT away from shallow approaches such as "how do you feel it's going?" Our specific request for "five structured questions" tells it to generate a multi-part inquiry, based on evidence and comparison.

In our case, its questions included what specific skills we've gained, and what sort of mistakes we feel we keep making.

Just thinking about these questions should help you get some perspective on your progress. If you type your answers back into ChatGPT, it can help you recognise the positives, and determine whether any negative perceptions are justified. If you're feeling a lack of progress, you can also ask the AI to help you track your endeavours. Try asking it:

What would measurable success look like for me?

The AI now sets out a series of appropriate goals and milestones, and shows how you're already hitting or working towards them – providing insight into what's working well, and a constructive view of anything that isn't.

Reframing unhelpful thinking

Despite all of ChatGPT's assistance and analysis, you could still find yourself losing motivation for your learning project. It's up to you whether you stick with it or give up, but before you make any decision, you can use the AI to challenge your own thinking. A good way to do this is by exploring hypothetical or counter-factual propositions – for example, you might ask:

Let's assume that my problems learning Japanese aren't down to a lack of ability. What are three possible reasons why I'm having trouble?

Taking the defeatist option off the table ("I'm just not clever enough") forces the AI to explore other possibilities. Perhaps your current learning method isn't right for you, or perhaps your expectations are wrongly calibrated. Some of these ideas might well ring true, suggesting ways to change your practice rather than abandoning it.

You can also use the AI's roleplaying skills to get an "outside" view of the situation. An example prompt might be:

Imagine my friend is struggling with this learning programme. What advice should I give her?

This tactic doesn't just give you a fresh look at your position; shifting the perspective to a friend changes ChatGPT's own framing of the situation. While the AI is programmed to be encouraging towards the user, it can take a more detached and objective view when thinking about a third party. This is a powerful trick you can use whenever you want to get away from ChatGPT's usual upbeat tone, and find out whether or not you really ought to quit.

Design a personalised growth plan

Whatever your objective, ChatGPT can help you come up with a plan to achieve it – or at least to get as close as possible

Level 1: Confident Recreational Skater (Very Reasonable Goal)

If you're early-stage, a strong goal is:

Skate confidently for 20–30 minutes without holding the barrier, stop safely, and turn smoothly in both directions.

That would include:

- Controlled forward skating
- Snowplough stop

1 Establish your goal

An important first step is to define your goal as precisely as possible – but you don't necessarily have to come up with it on your own. ChatGPT can work with you to pin down an achievable long-term target, and to define exactly what success should look like.

Section 1: Basic Control

1 Can you skate one full lap of the rink without touching the barrier?

- A) Easily
- B) Yes, but cautiously
- C) Only slowly
- D) Not yet

2 How do you usually stop?

- A) Snowplough stop with control
- B) Drag one foot / improvised stop
- C) Grab barrier

2 Find your starting level

In order to draw up a learning programme, it's important to understand where you're starting from. Again, you don't need to quantify this for yourself. ChatGPT can interview and test you to establish what learning level you're already at, and spot any gaps in your knowledge.

8-Week Goal

By Week 8 you should be able to:

- Skate multiple laps comfortably
- Do forward crossovers both directions
- Glide on one foot for 5+ seconds (each foot)
- Stop confidently at moderate speed
- Transition forward ↔ backward without panic

That's a strong recreational/intermediate milestone.

3 Set the milestones

You're now ready to start generating the plan. Remember to be clear about your availability, so the AI doesn't try to fill your every spare moment with exercises and practice. It's also a good idea to explicitly ask for regular milestones, to give structure and focus to your endeavours.

1 Progress Plateaus

What goes wrong:

Weeks 3–6 feel flat. Crossovers are still messy. Backward skating still feels wobbly.

Why:

Motor learning improves in bursts. With 2 sessions/week, progress can feel slow and invisible.

What to do:

- Reduce the number of skills per session.

4 Prepare for challenges

Before you dive in, it's also a good idea to ask ChatGPT what could go wrong – and to agree what you're going to do if things aren't working out. You can also ask the AI to mark expected motivation dips into the schedule, so they don't come as a surprise. 



Writing and creativity

ChatGPT can help with any sort of creative writing – not by generating the story for you, but by acting as an editor, reviewer and sympathetic reader to deliver constructive feedback and help you improve your craft.

76

Building stories step by step

78

Editing for tone, voice and originality

80

Forcing genuinely new ideas

82

Transforming and translating information

84

Controlling length and level of detail

85

Edit with AI – without losing your own voice

TUTORIAL

Building stories step by step

Use AI to spark ideas and find raw story material

In this magazine, we've seen plenty of situations where ChatGPT can relieve humans of cognitively heavy tasks, and turn scraps of ideas and information into polished output. That's great for financial statements and lab reports – but it's wholly inappropriate for creative endeavours, where readers will expect the inspiration and expression to come from you personally.

This doesn't mean ChatGPT can't help you work creatively, but you have to approach it in a different way. Rather than asking the AI to generate a completed product as in previous sections, you can flip the script and use it as a source of raw material for your own compositions. For example, you might ask:

.....
I'm writing a short story in which a witness refuses to speak in court. What are three possible reasons why they might be unwilling to testify?
.....

1) Fear — of retaliation, exposure, or collateral damage



Surface reason: The witness is scared.

Deeper possibilities:

- The defendant (or associates) has threatened them or their family.
- They're embedded in the same criminal network and speaking would expose their own wrongdoing.
- They've been warned that even an acquittal won't keep them safe.

Explore possible story elements, then choose and develop your favourites.

This prompting technique is called “controlled incompleteness” – which basically means restricting the AI to work on just one part of a larger task, so it can't exert too much influence on the overall outcome.

By instructing ChatGPT to generate multiple options, we ensure it can't steer us in any particular direction; its role is

merely to suggest scenarios for our consideration, in the same way you might research real-life court cases to adapt into your own story.

Our framing also avoids inviting ChatGPT to draft any content of its own, meaning you won't be exposed to AI-generated prose that could otherwise unconsciously contaminate your own voice.

Growing a story in deliberate passes

ChatGPT's abilities to maintain, structure and develop documents can be very powerful in a creative context. They let you use the AI as an interactive notepad, to not only jot down the details of your story, but to dynamically expand and even interrogate them. For example, we could follow on from the above prompt by telling the AI a bit more about our story:

.....
Make a note that the story is told from the perspective of a jury member who spots something all the other members overlook.
.....

The framing is again important – if we hadn't told the AI to “make a note”, it would have had to decide for itself what to do with the information. This might have involved writing the story for us. Phrasing it this way tells ChatGPT to simply add this fact to what it knows about the story.

You can progressively build up a whole narrative like this, feeding details, themes and twists into ChatGPT as you invent them. When you're ready to start drafting, you can get the AI to generate a structured reference guide to the setting and characters you've created – or use it interactively, to ask questions and

check specific details as you work. The AI doesn't just collect information, but can draw inferences and conclusions, which means you can explore your own narrative by asking questions such as:

.....
Who saw the victim on the day of the attack, and which characters haven't spoken yet?
.....

Again, ChatGPT isn't writing anything for you; like an assistant or story editor, it's helping you to conversationally organise and manage your elements.

Expanding outlines safely

On the previous page, we built our story from the ground up. Alternatively, you can start by sketching out the big picture, and developing the structure and details afterwards. ChatGPT is a great help for tasks such as this, although you'll need to phrase your instructions carefully to deter the AI from inserting its own contributions. A suitable prompt might be something like:

Here's the summary of my story. Expand it into scenes, and generate bullet points detailing what happens in each scene. Do not introduce any new ideas.

That final constraint may seem contradictory; as, in order to expand the summary, the AI naturally has to interpolate details that weren't spelt out in the original. However, ChatGPT recognises the distinction between elaboration and invention. It should understand from this phrasing that it's supposed to build on your words without inserting new characters or plot points of its own.

Even so, once your bullet points have been generated, we recommend you carefully check that the AI hasn't overstepped the mark in interpreting your words. You might want to ask it directly what assumptions it made during the exercise, to catch any subtle misalignments that could affect future drafts.

As well as helping you build a working structure, this process can expose any holes or inconsistencies in your narrative – if the AI's bullet points don't hang together, it suggests a problem with the underlying logic of your summary.

Below is an expansion of your summary into structured scenes, keeping strictly to the ideas already present in your outline. No new plot elements have been introduced. Short story summary.txt

Turn a synopsis into an organised structure, ready for you to flesh out.

Scene 1: Opening of the Trial

- The courtroom atmosphere is tense and formal.
- The defendant sits under scrutiny as the charge of high-stakes robbery is restated.
- The prosecution outlines its case, asserting that the evidence clearly identifies the accused.
- The jury listens; most members appear convinced from the outset.

You'll need to phrase instructions carefully, to deter the AI from inserting its own contributions

Grounding characters consistently

Strong characters are key to a good story, and ChatGPT can help you develop a whole cast of players with believable personalities and consistent motivations. You can do this similarly to how we worked out the story opposite: just describe a person to the AI and it will start building up an internal picture of them. You can also use ChatGPT to get ideas for their character type, using a prompt such as:

Let's develop the character of the judge. He is young and energetic, but thoughtful in manner. Suggest

three different character profiles that might fit him.

Again, by intentionally requesting multiple options, we deter the AI from shutting down creative possibilities. If we like any of the answers it comes up with, we can tinker with them, expand them ourselves, or get ChatGPT to help develop them further. For example:

Let's go with the idealist type. He wants to make the world a fairer place because he faced disappointments in his youth. Suggest three formative events that might have influenced his world view.

1) The Rising Reformer

Core identity: Brilliant, fast-tracked, determined to modernise the system.

- Appointed unusually young after a standout academic and advocacy career.
- Energetic in court procedure — keeps things moving, dislikes wasted time.
- Speaks crisply but pauses before rulings, weighing wording carefully.
- Believes deeply in procedural fairness and "getting it right."
- Occasionally reins in counsel with polite but firm efficiency.
- Has a subtle impatience with theatricality from barristers.

With ChatGPT's help, you can create rich, believable characters.

As the character becomes more rounded, you might try asking ChatGPT direct questions about them, such as what they'd do in a specific situation – or stir up drama by exploring what would make them angry or embarrassed. You can even ask the AI to roleplay the character, and chat with them directly about their thoughts and opinions. However, if you want to keep creative control, steer clear of directly quoting ChatGPT's simulated responses – think of them more as inspiration for your own dialogue. 🗨️

Editing for tone, voice and originality

Skip the generic AI style and give your work its own character

W

e've touched on the idea of using ChatGPT as a story editor, and once you've written a first draft, it can be a great help in critiquing your own work. The key is not to let the AI be the arbiter of what's good and what isn't, or it will tend to steer your work towards the sort of generic tone it uses itself. Instead, set your own goals and standards, and use ChatGPT as an external pair of eyes to help check you're hitting them. For example, you might attach an early draft to your prompt and enter:



This opening scene should make the reader feel immediately engaged, while withholding key information to keep them intrigued. Identify the three strongest ways in which it could do this more effectively.

ChatGPT is very good at simulating reader responses, and its suggestions can be remarkably insightful. Asking for three improvements stops the AI from trying to rewrite the text for itself, but it also means it will

suggest tweaks even if the scene is already working perfectly well. You'll need to use your own judgment as to whether ChatGPT's proposed changes are necessary and, if so, how to implement them.

You can use a similar technique to check a complete draft for things such as consistency of detail, or whether the clues you've woven into the text are too obscure. As before, we're not outsourcing creativity to the AI, merely using it as a test reader – one with the convenient ability to absorb a complete draft in moments, and who never gets tired of looking over rewrites.

Controlling emotional tone

ChatGPT works well with specific details and facts, but tone is more implicit and pervasive. If you want the AI to help you manage the tone of your writing, the first step is to check out exactly how it's reading it. Try asking:



Describe the emotional tone of this passage, and explain your interpretation. Highlight any aspects that don't fit with the overall tone, but don't suggest changes.

ChatGPT responds first with its idea of what the dominant mood of the scene is, and picks out details that contribute to it, such as descriptions of scenery, character behaviours, or language choices. Check that this aligns with your intention; if it doesn't, that suggests that your intended mood isn't coming across as clearly as it might.

The final part of the instruction is the most important, flagging up parts that, in the AI's view, go against the prevailing

The dominant emotional tone of this passage is **controlled tension filtered through quiet intellectual certainty**. It feels measured, observant, and slightly insulated — less explosive courtroom drama, more contained moral unease.

Below is a breakdown of how that tone is constructed, followed by elements that sit slightly outside it.

ChatGPT does a fine job of tracking tone, and identifying where it comes from.

Core Emotional Tone: Controlled, Observant Tension

1) Clinical Atmosphere

The opening sensory details — fluorescent lights, recycled air, old paper — create a sterile, almost antiseptic environment. The tension isn't chaotic; it's institutional. The anxiety is dry and procedural rather than heated.

mood. These might be moments of deliberate contrast on your part, or they could be missteps that you've overlooked. At the very least, it's worth exploring why the AI thinks they stick out.

You can also ask ChatGPT to track how the tone in your story develops,

scene by scene. This can help you confirm for yourself whether — for example — the tension is mounting as it's supposed to, and alert you to any areas where the energy drops, or even where a sardonic tone tips over into facetiousness.

Removing AI-sounding patterns

Even if you haven't allowed ChatGPT to write a single word of your magnum opus, its character can still seep into your writing. The AI tends to promote plain language, unambiguous phrasing and balanced sentences, and if you follow all of its advice, your own writing can end up feeling like something ChatGPT might have generated itself.

Happily, ChatGPT can also help you detect and counter this tendency. Just ask it to:



Analyse this passage and highlight any parts that could plausibly have been generated by an LLM.

This broad task definition lets the AI consider both specific turns of phrase and large-scale tone and structural features. When we tried it with a passage of our own work, the AI picked out several familiar features such as generic adjectives, repetitive rhythms, and a tendency to over-explain events and implications.

These findings don't necessarily mean that we've been "tricked" by the AI – maybe that's genuinely how we wanted to write those parts of the story. But if you wish to give your work a more distinctive style, this exercise can help you identify good areas to revisit. You can even get the AI's help to create a more characterful draft, with a prompt such as:



Tell me three ways I can change these elements to make them less generic.

ChatGPT suggests antidotes to the aforementioned AI-style features, such as using physical actions instead of descriptive prose, intentionally varying the length and balance of sentences, and leaving explanations implicit.

...your own writing can end up feeling like something the AI has generated...

I can't determine whether a passage was generated by an LLM, but I can analyse stylistic features that are commonly associated with LLM-produced prose and highlight where your text overlaps with those patterns.

Below are elements that could plausibly resemble LLM-style writing — along with why.

Use the AI to check that you haven't been influenced by ChatGPT's voice.

1) Polished but Generalised Atmospheric Description

"The fluorescent lights of Courtroom 4B hummed with a sterile, persistent energy that seemed to vibrate in Arthur's very marrow. It was a dry, recycled air that tasted of old paper and anxiety."

Why this could read as LLM-like:

- "Sterile, persistent energy" and "tasted of old paper and anxiety" are evocative but slightly abstract.

If your draft had been written by Agatha Christie, the *premise* would likely remain intact — a juror notices something overlooked — but the execution would feel markedly different in structure, tone and revelation.

Here are the key differences you'd likely see:

Learn more about your own work by comparing it to masters of the genre.

1) The Observation Would Be Less Announced, More Buried

In your draft, the crucial detail (the grease) is foregrounded and thematically framed. Christie would be more inclined to:

- Mention the grease early, briefly, almost casually.

Matching an existing voice

This principle doesn't only work with AI-style drafting; you can use it to avoid any writing style, or to intentionally emulate one. You might do this if you've written something that works well, and you want to bring its feel into a fresh piece of writing. Or, you could seek to improve your craft by exploring how your ideas might have been expressed by (for example) Agatha Christie.

For ethical and copyright reasons, ChatGPT may refuse to simply rewrite your work in someone else's voice, but you can normally upload a short section of a published work and ask the AI to use it as a model. However, you'll likely find it more instructive not to generate a complete rewrite, but to explore specific differences in the authorial approach. For this you can simply ask:



How might my draft be different if it had been written by Agatha Christie?

While ChatGPT may be reluctant to directly emulate Christie's personal style, it will happily list numerous ways that the Queen of Crime would have done things differently. In our case, those included a more restrained handling of symbolism and metaphor, and a focus on the personalities of the courtroom rather than the procedure. Bear in mind, though, that style is only one part of authorship – if you try to borrow a style that doesn't fit with your overall pacing and story structure, the result may end up feeling like a weird parody. 🙄

Forcing genuinely new ideas

Keep the reader guessing with surprise twists and revelations

It's sometimes argued that AI systems can't create original content. In one sense, this is true: large language models such as ChatGPT can't generate new information outside of the patterns they've learnt from their training data, or that have been supplied by you.

However, that doesn't mean they can't combine existing ideas in new ways. The catch is that they usually won't: ChatGPT's statistical model steers it towards reliable, well-established answers, so if you ask it to come up with a twist for your courtroom drama, it will likely present a list of hackneyed and predictable suggestions. As usual, the way to work around this is by framing the task appropriately. Instead of just asking for twist ideas, try asking the AI to:

.....
List five things that could happen in this scene that the reader will absolutely not be expecting.
.....

Now the statistical model works in your favour; since ChatGPT "knows" all the most likely developments, it can intentionally generate ideas that are outside of those patterns. You can't guarantee that what it comes up with will be

It's sometimes argued that AI systems can't create original content. In one sense, this is

Here are five developments that would feel genuinely unexpected in your courtroom scene — not because they add spectacle, but because they destabilise the reader's assumptions about where the tension lies:

1) Arthur Is Wrong

The grease isn't what he thinks it is.

- He stakes his credibility on the identification.
- A juror quietly produces evidence that the lubricant is widely available.
- Or the prosecution calls a rebuttal expert.
- Arthur's authority collapses in the room.

Ask ChatGPT to suggest ideas that subvert reader expectations.

completely unprecedented, but you can reassure yourself that it isn't commonplace by asking:

.....
List any creative works that feature a similar twist.
.....

Since the AI hates to say "I can't think of an answer", it will stretch as far as necessary to come up with parallels in existing stories, potentially to the point of hallucination. You can assess for yourself how real or significant the similarities really are, and decide whether to go with an idea or keep searching.

Escaping obvious solutions

ChatGPT's willingness to reach for unlikely connections can also help you get away from predictable storylines. The trick here is in anchoring the AI: instead of merely providing a starting point and asking what happens next, dictate the outcome, too, and instruct ChatGPT to figure out how the story gets from A to B. A suitable prompt might be:

.....
Even though Arthur recognises the critical piece of evidence, he can't convince the other jurors of its significance. So how do they arrive at a guilty verdict?
.....

Since the direct path is closed off, ChatGPT has to generate a more elaborate chain of events to explain the outcome, which could make for a fresher and more engaging story. Of course, it's up to you to decide how to use the AI's ideas, and to implement them in a believable way.

You can also use this approach to invent complications that add depth to the action, or enrich character dynamics. For example, you could ask:

LLMs can't generate new ideas outside of the patterns they've learnt from their training data, or that they've been supplied by you

Why does the witness' statement cause the other jurors to distrust Arthur?

The statement doesn't have to be deliberately crafted to have this effect – in fact, this prompting technique works best if it isn't. When there's no low-hanging "right" answer for the AI to find, it's more likely to come up with unexpected ideas of its own. In this case, it cleverly suggested that the witness might have mentioned a throwaway detail that, by chance, undermined Arthur's credibility.

Working strictly within your existing story, the witness's statement (Marcus Thorne's identification and mention of the limp) can undermine Arthur not because it disproves him, but because it **feels more human and immediate** than his technical reasoning.

Here are the strongest, text-consistent reasons the jurors might begin to distrust him:

The way you frame a scene can be as important as the events themselves.

1) Thorne's Account Is Emotional; Arthur's Is Mechanical

Here are three sharp, courtroom-plausible thoughts that might run through Marcus Sterling's mind as the witness steps into the **witness box**:

1) He needs to look credible before he sounds certain.

Sterling will be scanning instantly for optics:

- Is Thorne steady or visibly agitated?
- Is he staring at the defendant too intensely?
- Does he look like a man recounting what he saw — or prosecuting the accused himself?

Get inside your characters' heads to find their fears and motivations conflict.

Sterling knows jurors decide credibility in seconds.

If Thorne appears aggressive or self-righteous, the defence will frame him as biased.

Brainstorming from different perspectives

Switching viewpoints doesn't just help you escape from recurrent patterns – it's a tool you can use to develop your own view of a scene. In this case, we're not talking about outlandish perspectives that force a complete rethink, but digging into different aspects of the established context.

For example, we've noted before that the AI is very good at roleplay, so you might ask it to describe what happens in a particular scene from the perspective of the barrister. This won't just produce a shift in tone – it will generate a very particular take, centred on the character's own interests and concerns. This can provide insight into the dynamics of the scene, and suggest opportunities for productive interactions.

You can ask more targeted questions, too, such as:

When the witness enters the box, what are three things the barrister might be thinking?

Naturally, you don't want to let ChatGPT define your characters and story; but once you've set the parameters, it's a very effective tool for exploring how they might play out in the narrative. You can even ask the AI to think about multiple characters or interests at once:

How does Arthur's interpretation of the testimony differ from the barrister's?

Fleshing out the characters like this helps you create something that isn't just a sequence of events, but a story shaped by personalities and interpretations. 🧠

Refreshing stalled outputs

"Stalled outputs" occur when you're trying to generate inspiration, but ChatGPT keeps coming back to a particular idea or angle. This can happen in long threads when certain ideas have come up repeatedly; each appearance adds weight, until almost every output leans in the same direction.

One way to deal with this is to wipe the contextual slate clean. Ask the AI to summarise the key details of the conversation so far, then start a new chat, paste in the summary, and pick up from there. Now everything is interpreted and weighted afresh, without any lingering biases hanging over from past discussions.

Alternatively, to keep the context going, you can try to break the AI out of its loop by prompting a shift of perspective or focus. An example might be:

Rewrite this scene so the witness is aggressive and disruptive.

It doesn't matter whether or not you're interested in exploring this potential story element; it forces the AI to reconstruct the scene, and in so doing de-centres elements that had become too weighty. Now, if you ask ChatGPT to help you develop what the main characters are thinking and feeling, it will be more likely to come up with new ideas.

If this reframing doesn't have the effect you want, take it further – ask the AI to rewrite the scene as if the judge had left the room, or as if it were taking place a thousand years in the future. You can later tell it to undo the change, while keeping any useful ideas that emerged from it.

Transforming and translating information

Get a fresh view of your story by changing the angle

S

hifting perspective isn't just about reviewing a scene through a particular person's eyes. ChatGPT

can recast your story in a wholly different form, adopting different writing conventions and simulating a writer with their own agenda and priorities. For example, you could ask it to:

Rewrite this scene as a report in the local newspaper.

For creative purposes, we'd usually avoid asking the AI to rewrite our work, but here we're not producing anything that we intend the reader to see. Rather, we're generating secondary material to expand our own understanding, bringing in perspectives and interpretations from outside of the existing text.

These interpretations can still feed indirectly into the narrative.

In this case, although we don't include the newspaper report in the story, we can imagine that the characters might have read it, and that it will be in the back of their minds as they file into the courtroom for day two of the trial.

For a more nuanced approach, try generating multiple different accounts. An example prompt might be:

Generate two newspaper reports, one that supports the defendant and one that supports the prosecution. What are the key points of differentiation?

Again, you don't need to add any of this content into your own writing; just let it hang over the drama. Depending on which newspaper they read, the various characters might be coming to court with quite different ideas about what's really going on, and their own role in it.

Good structural instinct — moving the barrister's outburst changes the emotional sequencing of the courtroom, so you'll need a few compensating adjustments to keep the narrative psychologically coherent.

Here's what would need to change to maintain consistency.

1) The First Witness Must Now Feel Controlled

If the outburst no longer occurs during Marcus Thorne's testimony, his appearance should:

- Feel relatively contained.

Restructuring the plot is easier when the AI can help you tie up loose ends.

Converting between formats

Another use for format shifting is to review and develop your plot. For example, you could ask the AI to convert your story into a timeline; ChatGPT will boil down the story to its core events and consequences, and present a concise step-by-step summary. This could expose issues such as overly linear pacing, or clues that are all bunched up at the start of the story. If you decide to make a change to the order of events, ChatGPT can help you keep things coherent by tracking the impact on the rest of the story, with a prompt such as:

Move the barrister's outburst to after the second witness statement, and tell me what else needs to change to keep the narrative consistent.

You can use ChatGPT to expose underdeveloped logic, too, by asking it to highlight any effects that don't seem to have an explicit cause, or vice versa. Using output constraints, you can directly test how the causality hangs together, by asking the AI to:

Describe the entire plot using only sentences of the form "X happens because Y".

A final transformational trick you can try is to ask the AI to turn the story into a map of questions arising in the reader's mind, along with details of how and when they're answered. Of course, not everything needs to be neatly explained and resolved, but you can use this tool to check that you're feeding the reader information at a suitable rate to keep them hooked – and to confirm that any issues left unanswered are intentional, rather than accidental.

Translating with cultural awareness

Earlier, we suggested re-envisioning your scene as if it were set a thousand years in the future. That's an extreme example, but ChatGPT can help you relocate your work into any time or place – either with a view to crafting more realistic scenes, or as an exploratory exercise.

One practical application of this is checking your work for cultural consistency. For example, you might attach a draft and say:

.....
This story is set in an English village in 1979. Which aspects of it are inconsistent with the social setting or time period?

The AI can flag up obvious anachronisms – hopefully no one's pulling out an iPhone – but it will also alert you to broader aspects of the setting that might be relevant to the drama. For example, in this era, newspaper reports could shape public opinion much more directly than they do today, and people living in a village would be likely to share rumours and gossip.

You can also use ChatGPT to develop context-appropriate social dynamics. For example, you could tell it:

Short answer: **mostly consistent with position, slightly inconsistent with era, and under-specified for age/background.**

Below is a focused breakdown.

1) Consistency With Position (Crown Court Judge)

These elements *do* fit:

- "My Lord" form of address.
- Controlled interventions.
- Directing counsel to ask questions rather than argue.
- Clear instruction on "beyond reasonable doubt."

The judge's authority feels legitimate and properly hierarchical.

However:

- The gavel tap is slightly Americanised. English judges do not typically use gavels in Crown Court.
- Interventions are a touch conversational compared with 1970s judicial

Check that your characters behave appropriately for the setting.

1) CCTV Footage

This is the biggest issue.

In 1979 England:

- CCTV was not widely used outside specific high-security or experimental contexts.
- Widespread commercial CCTV systems in warehouses and docks became common much later (mid-1980s onward, more broadly in the 1990s).
- High-resolution still extraction would be unlikely.

ChatGPT can pick up subtle context errors – as well as obvious ones.

.....
The defendant lives in a big city. List five ways that this might affect the other characters' perceptions and assumptions.

Aside from translating scenes into various contexts, ChatGPT can perform literal translation tasks, too, employing the linguistic skills we've seen earlier. So, you could use it to help you write passages in other languages, with respect for nuances of register and implication – although we'd suggest caution if you intend to write a character who speaks a language you don't understand yourself.

Checking idiomatic accuracy

Context isn't just about the wider world; it greatly affects how individuals act and speak. We've talked about developing characters who are fully realised and internally consistent, but cultural context adds another layer – it isn't just a case of building a back story for the judge, but establishing how he'd act in this specific scenario, in this place and time, in front of these people.

This is the sort of job at which ChatGPT excels. For a close examination of the character as written, just ask:

.....
Are the judge's speech patterns and mannerisms consistent with his age, position, background and era?

In our case, ChatGPT provided reassurance, but it was smart enough to point out that the judge's statements didn't openly convey the character note we'd established earlier, of him being a young idealist with a troubled past.

As well as inspecting individual characters, you can use this technique to draw comparisons between them. For example, you might ask:

.....
How do the judge and barrister use language in different ways that reflect their different roles and attitudes?

A similar prompt could be used to explore non-verbal behaviour, or any aspect of a character; hopefully, ChatGPT will pick out some clear contrasts, confirming that you've made your characters consistent within their context, but clearly distinct from each other. 🍷

Controlling length and level of detail

Make your story unfold at exactly the pace you want it to

W

e've seen before how ChatGPT can turn a lengthy business report into an incisive summary. You

can use this same ability to manage the pacing of your own works, so that events unfold at the tempo you want them to. Very simply, you might ask the AI to:

Show me five ways to shorten this scene.

ChatGPT can highlight wordy sentences and descriptive passages that don't advance the action. It will also warn you if you're labouring a point that doesn't need to be spelt out, or coming back to an idea multiple times in a scene. The AI won't normally mess with anything that's integral to the plot, but you can ask it to pick out developments that could be

streamlined if you want to make the story more economical.

The AI can also help to extend a scene – although it goes without saying that you don't want to just ask it to “make this longer”; the AI will interpret that as permission to invent new details itself. A better way to handle the request might be:

What are three ways to extend this scene that would improve the reader experience?

This phrasing invites the AI to propose extra moves that expand the plot, characterisation or overall mood, rather than merely slowing down the narrative. In our example, its suggestions included setting up an interpersonal conflict to contribute to the tension in the later scenes, and humanising the lead character by giving him a moment of self-reflection and doubt – both good ideas.

If you extend the third scene (Arthur identifying the grease), the goal shouldn't be “more explanation.” It should be **more tension, resistance, or emotional layering**. Here are three ways to extend it in ways that deepen reader engagement rather than slow momentum.

1) Prolong the Moment Before Arthur Speaks

Right now, Arthur notices → explains → persuades.

You could extend the scene by delaying his decision to speak.

For example:

- He studies frame 402 again.
- He checks himself — once, twice.

The AI can help you expand scenes in ways that genuinely add to the story.

Managing the overall tone

When you're making targeted edits to specific scenes, tonal inconsistencies can creep in. We're not just talking about the emotional mood here, but taking in structural aspects, such as how and when the narrative focus shifts from one character to another, or how the density of information being conveyed to the reader might vary from scene to scene.

If you get a sense that your prose has become uneven, you can explore possible causes by feeding a passage or scene into ChatGPT and asking:

What elements of this scene primarily define its tone?

This should help you turn a vague feeling that a passage doesn't feel right into a focus on the elements responsible. From there you can make adjustments as necessary to modulate the overall mood. For a fuller overview of the tonal arc of your story, you can make use of the AI's ability to extract and transform information, with an instruction such as:

List every element in this story that contributes to or changes the tone, and describe its effect.

This will produce a detailed rundown of all the events, speeches and descriptions that progress the mood of the piece. It may include some shifts that you didn't intend, and it might overlook some that you thought were obvious – neither of these cases is necessarily a problem, but you could consider them useful simulations of what an average reader might experience.

As always, you're not letting ChatGPT tell you what to write, but using it to gain a better insight into your own writing, and learning to more effectively express your intention. 🗨️

Edit with AI – without losing your own voice

ChatGPT can help you make the best of your own creative ideas, as long as you keep its role tightly defined

What draws me in

1. Immediate atmosphere and point of view

The opening image — “The fluorescent lights of Courtroom 4B hummed...” — is sensory and specific Opening scene.txt. The hum, the “dry, recycled air,” the taste of “old paper and anxiety” give the scene texture without overexplaining. It establishes mood fast and anchors us in Arthur’s body.

2. A clear narrative tension

We instantly understand the stakes: a man’s future is being “dismantled” and the jury is ready to convict Opening scene.txt. The line “They weren’t looking for justice anymore; they were

Looking at this purely as a reader — not as a writing coach — these are the elements that most clearly feel like *your* voice rather than generic crime-fiction prose.

1. Precision through physicality

You consistently ground abstraction in texture and mechanism:

- “dry, recycled air that tasted of old paper and anxiety”

Opening scene.txt

- “grainy, flickering reality of the surveillance footage”

Opening scene.txt

- “a dull, viscous streak of industrial grease” Opening scene.txt

1 The AI’s first impression

Before we attempt any editing, let’s ask ChatGPT to give a big-picture critique of our work. This helps us calibrate our sense of what the AI responds to – and if we disagree with its interpretation, we can correct it and set our own working parameters before moving on.

2 What makes it ours?

Next, let’s establish what our ownership of the work looks like. The AI can pick out the hallmarks of our personal style, along with any unfortunate habits we might want to change. It also offers to highlight parts where our authentic voice doesn’t come across clearly.

Here are three ways the narrative could be strengthened **without changing a single event or plot beat** — only execution, emphasis and calibration.

1. Make Arthur’s deduction feel more earned

The grease revelation is the hinge of the scene. At present, the leap from “specific sheen” to “maritime district product” happens very quickly Opening scene.txt.

1. If Thomas Miller is innocent, who committed the crime?

Arthur’s observation about the Moebius 9415 grease strongly implies the real culprit works at the north shipping docks, likely around heavy cranes. But the story ends without identifying this person. The smear becomes a clue pointing outward — yet no further investigation is shown. The reader is left wondering whether the real thief will ever be caught.

2. Was Marcus Thorne mistaken, or was he lying?


Thorne insists he saw Miller and mentions the limp. If Miller is innocent, then:

3 Improve the execution

Now we’re ready to get ChatGPT’s suggestions as to how the writing might be improved. Focusing the task on language and framing lets us take advantage of the AI’s strongest skills, while preventing it from trying to write new passages for us, or rework the plot.

4 Open-ended questions

If we upload a complete draft, ChatGPT can work through it to find questions that haven’t been answered at the end of the story. This helps us to be intentional about what we leave unresolved, and to spot any accidental plot holes – again, without letting the AI change anything itself. open



Solving complex problems with logic

ChatGPT can help with advanced problems and multi-step logical processes. In this section, we look at techniques to get reliable answers, check accuracy and create reusable structures to use over and over again.

-
-
-
- | | | |
|-----------|--|-------|
| 88 | Structuring step-by-step reasoning | |
| 90 | Self-review and expansion | |
| 92 | Examples and reusable templates | |
| 94 | Logic rules and output formats | |
| 96 | Working with images, audio and documents | |
| 98 | Advanced reasoning | |
| 99 | Use AI to help with a multi-step decision | |

TUTORIAL

Structuring step-by-step reasoning

Build a framework to solve real logic problems

C

hatGPT doesn't always do well with complex problems. It likes to generate confident,

persuasive-sounding answers, and tends to steamroller over nuanced trade-offs. That isn't to say it can't help with complex issues, but it means they need to be worked through step by step, as we've seen in previous sections. The catch is that you can't necessarily trust the AI to pick out the steps on its own.

For example, imagine that your laptop won't connect to the internet. If you ask ChatGPT for help, it will walk you through various potential fixes, narrowing down the possibilities based on your feedback. This is step-by-step reasoning; but you get almost no insight into how the AI is choosing particular suggestions, and skipping over other possibilities.

If you need a transparent and repeatable decision process – one that you can rely on in a business or research context – you need visibility into the reasoning path. You can get this by entering a prompt such as:

My laptop can't connect to the internet. Generate a sequence of questions, where each question determines the next, leading to a diagnosis.

Now you can see the steps for yourself – but don't assume that transparency alone makes for a good process. The point of exposing the logic is that it provides the opportunity to check for embedded assumptions or overlooked possibilities.

1. Clarify the strategic objective

- What problem or opportunity is prompting the idea of a second location?
- Is the aim to reach new customers, expand capacity, recruit talent, improve resilience, or raise the company's profile?
- What specific outcomes would count as success?

You can work collaboratively with ChatGPT to design a reasoning process.

2. Define the scope of the proposed expansion

- What functions would the Edinburgh location perform (sales, engineering, support, administration, etc.)?
- How large would the site initially be?
- Would it replicate existing operations or specialise in particular roles?

Check the working before drawing conclusions

If you've read to this point of the magazine, you'll know that ChatGPT often tries to generate acceptable responses using pattern matching, rather than applying actual logic. You'll also remember that it tends to commit to its statements, even if other answers are possible; and that if you ask it why something is true, it will usually generate a confident explanation, even if the proposition is in fact debatable.

All of this means that if you ask the AI to justify a previously generated answer, you may get a post hoc rationalisation, rather than a real explanation. Even if you repeat the question, and this time ask ChatGPT to show its working, the process is liable to be contaminated by previous context, as the AI favours ideas that are already present in the chat.

The solution is to get ChatGPT to commit to a logical process before generating an answer. A sample prompt might be:

My company is considering opening a second location in Edinburgh. List all the logical steps you would work through in determining whether or not this is a good idea. Do not generate a conclusion.

Stopping short of the answer prevents the AI from treating the question as settled, so you can inspect and question its reasoning without having to fight against established context. Once you're happy with the process, you can generate the answer with a prompt such as:

If you were to apply that logical process to this problem, what would be the outcome?

The hypothetical framing isn't mandatory, but it discourages the AI from over-investing in the conclusion, so it's less likely to colour any subsequent questions you may have.

Stopping short of the answer prevents the AI from treating the question as settled...

Checking logic as it develops

While it's wise to check the steps in ChatGPT's reasoning, it won't necessarily expose problems caused by incorrect framing or assumptions – and multi-step processes can magnify the impact of any initial errors.

So, if you're using the AI to help with complex problem-solving, it's important to make every step accountable. A simple way might be to build on the prompt we used opposite, and add:

When you carry out this process, output a one-sentence summary after each logical step detailing what it has confirmed and what it has ruled out.

Now you can carry out a test run, asking ChatGPT to work through the question – again, stopping short of the conclusion – and confirming for yourself that the logical process is consistent and effective. If you want to keep a closer eye on things, you could build in stricter checks, such as:

1. Clarify the strategic objective

The proposed Edinburgh office is intended to expand access to Scottish financial services and public-sector clients while enhancing brand prestige through a second headquarters.

Summary: This confirms that the expansion addresses a defined strategic objective (market access and positioning) and rules out the possibility that the proposal is driven only by prestige or vague growth ambitions.

2. Define the scope of the proposed expansion

The office will host a Regional Sales team, a Public Sector Scotland desk and a specialised R&D unit, led by two relocating senior directors and growing to around 12 consultants.

Summarising the steps helps you see how the task has been broken down.

After each step, explain how it brings us closer to an answer to the original question, and list any new assumptions introduced.

This helps us spot any unnecessary steps or those that might hinder the process, while referring back to the original

question should prevent the terms from drifting as ChatGPT works through the logic. Checking assumptions, meanwhile, is always a good idea – here, these might be unwanted interpolations that, once exposed, can be dismissed. Or, if the AI is having to fill in detail for itself, it might indicate that some necessary element is missing from your original framing of the problem, and needs to be supplied.

Avoiding leaps in reasoning

To make the AI's reasoning process even more robust, you can inspect not only how each logical step contributes to the solution, but the specific premises and inferences involved as well. This is a subtly different intervention to the one suggested above, and it can catch situations where, for example, the form

of a particular deduction is familiar to the AI, so it skips over the logical operation and simply generates a believable (but, in reality, untested) answer. In fact, this technique can prevent such scenarios from arising in the first place, as forcing the AI to show its logical working makes it less likely to take a shortcut.

To achieve this, you could use prompt:

For each step in the process, list the inputs, the logical operation being applied, and the conclusion reached

1. Clarify the strategic objective

Inputs:

- Core business: high-end B2B professional services in data analytics and digital transformation consultancy
- Reason for Edinburgh: access to Scottish financial services and FinTech sectors; prestige of a second headquarters
- Definition of success: £1.2m regional revenue within 24 months; 12-person Scottish team

Logical operation being applied:

Match the proposed expansion against stated business goals to determine whether there is a clear strategic purpose.

Spelling out the logic discourages ChatGPT from taking shortcuts.

As well as letting you audit the logical steps, this brings them explicitly into the conversation, so you can interrogate them like any other proposition. For example, after you've carried out a full test run of the logic, you can ask:

Which step's conclusion follows most weakly from its premises?

This could reveal poorly defined terms in the original question, or unclear priorities that other checks have overlooked. If what you need to do isn't obvious, just ask ChatGPT what information is required to turn this into a confident conclusion. 🧠

Self-review and expansion

Take a closer look to spot potential flaws and build up the logic

C

hatGPT can be a useful ally in verifying that solutions are robust and reliable. On

the previous page, we asked it to highlight the weakest individual link in a chain of logical operations, and you can equally ask it to critique the entire chain. For example, you could ask:

How could this process be strengthened to give us clearer, more actionable advice?

Even if the chain was originally devised by the AI itself, ChatGPT will almost certainly come up with some plan to improve it. Note that this doesn't necessarily mean the original was seriously flawed – the AI always tries to give a positive reply to any request, so you'll have to decide for yourself whether its suggestions are really worthwhile, or whether they're just tinkering for its own sake.

However, there's a good chance that the process can be materially improved, as ChatGPT's reliance on pattern matching means there's no guarantee that its original proposal will have been the best possible one, or even a particularly good one.

The danger is that, for the same reason, ChatGPT won't necessarily come up with the most effective improvements. You can help focus its attention by using a technique

The process you used is structurally sound, but it mainly produces structured observations, not decision-quality conclusions. To make it yield clearer, actionable advice, you would strengthen three aspects: evidence quality, decision logic, and uncertainty handling.

Below are the most effective improvements.

The AI can analyse and critique logical processes – even ones it created itself.

1. Replace descriptive conclusions with explicit decision tests

Many steps currently conclude things like *"this confirms demand exists."* That's informative but not decisive.

Each step should instead include a decision rule.

we've seen in previous sections of the magazine – namely, requesting a list of top issues:

List the five most serious weaknesses of this logical process.

The results still aren't guaranteed – "most serious" is a subjective measure, anyway – but forcing the AI to find multiple weaknesses makes it dig deeper, and introducing the idea of seriousness encourages it to seek out major structural deficiencies, rather than being satisfied with superficial tweaks.

Expanding logic deliberately

Even a structurally sound chain of reasoning can produce unreliable answers if the criteria it relies upon aren't appropriately defined. For example, in the context of opening a second business location, one logical test might be "is a suitable space available?"; but if you haven't set out exact requirements for floor space, cost and so forth, the AI could consider that the condition is met simply because the target is a large city where office space can usually be rented.

To fix this, we just need to develop this one step, without touching the rest of the process. Start by diagnosing what the AI is currently doing, with a prompt such as:

Detail exactly what evidence you're drawing on to arrive at the conclusion of this step.

Again, this is subtly different to what we've done before – an audit of the logic would confirm that this stage checks that space is available, but it wouldn't

Even if the chain was originally devised by the AI itself, ChatGPT will almost certainly come up with some plan to improve it

necessarily expose how the answer was determined.

You can now refine the process by adding decision criteria, or uploading supporting documentation that ChatGPT can interpret to determine your real needs. You could also consider building in an instruction to search the web for available commercial properties, so as to ground the logical test in the real world – something the AI won't do on its own.

Correcting errors in the logic

Sometimes, a logical step might simply be wrong. This isn't like failing to explain what we mean by "a suitable space" – it's more like accidentally including decision criteria that doesn't apply to the current question.

With all the diagnostic prompting and close reading we've done so far, a blunder such as this should be quickly exposed, long before it can influence any decisions. When you discover that the logical process includes an erroneous step, you can cut it out with a prompt such as:

Remove the step that refers to footfall, as this location will not serve the public; do not change anything else.

ChatGPT likes to generalise instructions, so we've included an explanation and a constraining clause to discourage it from going beyond our specific requirements. Removal isn't your only option; if the basic form of the logical step is right, you might just need to change the specifics. ChatGPT can help you do this easily:

Update the reasoning process to replace the requirement for footfall with a requirement for road access for supply vehicles.

To be sure this has had the desired effect, re-run the diagnostic tests on the previous page. Whenever you change or remove a piece of logic, you should ask the AI to have another go at solving your problem; even if you've only removed an irrelevant step that seemed to do nothing, it could have subtly influenced ChatGPT's reasoning by introducing superfluous considerations into the context.

Thought for a few seconds >

Below is the **updated step**, replacing the earlier footfall-related requirement with a requirement for **road access for supply vehicles**. Everything else in the reasoning process can remain unchanged.

7A. Assess access for supply vehicles

Review whether the candidate site can be reached easily by vans or delivery vehicles, including approach routes, loading access, parking or stopping arrangements, and any local restrictions on deliveries.

Summary: This confirms whether the site can support reliable delivery and servicing logistics and rules out locations where vehicle access would disrupt operations or create recurring practical constraints.

How this brings us closer to an answer: It tests whether the site is operationally workable on a day-to-day basis, rather than merely attractive or well located in strategic terms.

Any procedural steps that need changing can be corrected in place.

Three most significant cons

1. Revenue targets must be achieved quickly

A consulting office with 12 staff represents a large fixed cost base, so financial success depends heavily on achieving the projected £1.2m revenue target within two years.

2. Execution risk from first geographic expansion

Although leadership has experience elsewhere, this is the company's first multi-site operation, introducing organisational complexity and coordination risk.

Rather than asking for a single settled response, check the pros and cons.

Managing detail levels

While it's often necessary to build specific details and criteria into the chain of reasoning, the idea can be taken too far. Complex processes tend to lead to dense, wordy conclusions, when you might really want a clear decision. What's more, if you design a lengthy sequence of logic that tries to account for every variable, you may find that the most important considerations get crowded out by a dozen low-priority factors.

You can instantly trim down an over-long process using a prompt such as:

Simplify this process to rid any steps unlikely to have a huge impact on the final conclusion.

This produces a streamlined chain of reasoning that's focused on the most significant factors. Cutting out

variables not only makes for a cleaner conclusion, it makes it easier to refine and experiment with your process, or to vary the framing of your question. Another possible approach is to get ChatGPT to break down its results into headline findings:

Use this process to weigh up the question, and tell me the three most significant pros and cons.

This approach prevents the AI from trying to reduce multiple nuanced considerations into a single decision – rather, it will set out the key points so you can clearly grasp the situation for yourself. And if you want to take a closer look at any aspect of the analysis, the full chain of reasoning remains available for you to query and explore.

Examples and reusable templates

Standard procedures let you reliably reuse your logic



ChatGPT is very good at imitation, so if you want it to give you answers with a particular tone or structure, you can simply give it a model to copy. You don't even need to provide the model yourself; you can develop it collaboratively with ChatGPT, by asking it to:



Help me generate a model answer that you can use to respond to my questions.

The AI will usually respond by generating a suggested model based on the conversation so far. You may then instruct it to adjust factors such as the length or level of detail, and specify any must-have features, such as a one-line summary or an evaluation of what the information adds to the discussion.

As you work, the AI will generate fresh examples to show how your instructions are shaping the template. Once you're happy with what it's producing, you can lock down the format with a prompt such as:



Always use this template for answering questions about sales contracts.

Alternatively, you could tell ChatGPT that you'll use a tag to indicate when the template should be used, such as prefacing your question with "Q:". You can even define multiple templates in this way, but be aware that ChatGPT doesn't store them as formal rule sets – as far as it's concerned, they're just part of the conversation, and their influence can weaken over time as with any piece of context.

Creating reusable structures

Templates can be created for processes as well as content, to define a chain of reasoning or analysis for future reuse. For example, you might want to create an AI process that checks monthly sales reports against various benchmarks and generates a summary in a standard format.

You can do this by combining the techniques we've seen previously in this section – by first, building a step-by-step analysis procedure, then designing a standard output format.

If you wish, you could then simply come back to the conversation each month and ask ChatGPT to generate a new answer using the latest figures. However, this creates a risk that the AI will skip over the logical analysis, and merely generate a plausible-sounding summary that replicates the format of the original answer.

Even if you instruct ChatGPT to follow the same logic, you won't necessarily get consistent results: each time you repeat the process, the accumulation of context makes it more likely to become distorted or diluted.

The solution is to make the template explicit and immutable. You can do this by asking the AI to:



Generate a downloadable process document that precisely details the process you used to generate the previous response, and specifies the agreed output format.

When you want to repeat the process in future, you can now start a new conversation, upload the template and immediately set ChatGPT to work on the latest data. This way you know that there will be no prior context to distract the AI, and that your process hasn't quietly drifted since it was originally created.

I've created the downloadable document.

Download it here:

[Download the process document ↗](#)

The document contains:

- The **exact reasoning process** used to answer spreadsheet questions
- The **decision rule** used to determine the best location
- The **standard output format** I will use for future answers
- The **constraints** (e.g. using only spreadsheet data)

If you'd like, I can also create a **much stricter version of the process** that eliminates several hidden assumptions in the current method and makes the conclusions more defensible. [↗](#)

Creating an external process document ensures your terms won't drift over time.

Writing prompts for others

If you're helping out a business colleague, or sending advice to a friend, you may want to suggest specific prompts for them to enter into ChatGPT. Unfortunately, there's no guarantee that they'll get the same responses that you did. For one thing, the AI model is regularly updated – although, hopefully, that won't cause big problems, as each new version of ChatGPT is generally intended to give better answers than the last.

A bigger problem is pre-existing context. If you're working through a problem with ChatGPT and you find that a certain prompt does something particularly useful, you'll naturally want to share it with others. However, it may be that the prompt only worked so well for you because the AI was able to interpret it in light of the previous conversational context. If your friend opens up a new chat and enters the same prompt, the response is likely to be more basic and generic.

This doesn't mean you can't share powerful prompting techniques with others, but you'll need to provide more than just the bare instruction, or even the sort of process document we created opposite. Try giving the AI a prompt along the lines of:

.....
Generate a downloadable document that includes the process used to generate the previous response, and all the related information required to repeat it without needing additional context.

.....
 Fill in additional details or tweak the objective depending on the precise outcome you're looking for. Once you've downloaded the generated file, you can share it with others to upload into their own ChatGPT sessions.

If your friend opens up and enters the same prompt, the response is likely to be more basic and generic

I've created the document that fully specifies the process and all information needed to reproduce the previous analysis.

Download it here:

[Download the reproducible process document ↗](#)

The document includes:

- The exact question being answered
- A complete description of the spreadsheet dataset
- The data preparation rules (including how monthly fees were handled)
- The regional grouping used in the analysis
- The step-by-step analytical procedure
- The decision rule used to select the location
- The standard output format used for answers

This means someone with only:

1. the spreadsheet, and
2. this document

can repeat the analysis and reproduce the same conclusion without additional context. (↵)

📄 🔊 👍 🗨️ ...

If you want fully reproducible results, be sure to include context.

Maintaining consistency

We've suggested using templates at the start of new conversations, to avoid intrusions of unwanted context, and to ensure the process doesn't drift over time. However, that might not always be the most convenient way to work – if you've been working with the AI to refine its understanding of a nuanced issue, you might want to apply your workflow without discarding all of that context.

You can protect a template from short-term contamination by using explicit defensive phrasing such as:

.....
Answer the question using the reasoning process defined in this document. Do not adapt the process in any way or introduce any assumptions not present in the document.

.....
 If you don't want to spell out these instructions each time you use a template, you can alternatively edit the

external template file and build your constraints directly into the process (for this purpose, we recommend exporting it as a plain text file rather than a PDF). An example addition might be:

.....
This process must always be executed exactly as written. No external assumptions or conversational context may be used, except those directly implied by the input.

.....
 While this can save you some typing, you may have already spotted the catch; as soon as you upload the template into ChatGPT, this instruction itself becomes subject to dilution and degradation. You can reduce this effect by periodically asking the AI to restate the templated process in full, bringing it back into the foreground – or, you can just reimport it each time you use it. 🔄

Logic rules and output formats

Define how ChatGPT operates, and get the results you want

Below is the classification and processing performed strictly according to the instructions in the two uploaded documents:

- The survey responses were read from “Survey responses 2025”.

Survey responses 2025.docx

- The processing rules come from “Comment processes”.

Comment processes.docx

Those rules specify:

- Compliments → produce a sentiment map and note products praised most often.
- Criticism → list specific issues with specific items.
- Suggestions → list ideas mentioned two or more times.

Branching logic tells the AI how to respond to different input types.

S

ome reasoning tasks call for branching logic: not weighing up different factors, but choosing

between processes based on a logical test. For example, if using ChatGPT to process responses to a customer survey, you might have different pipelines for dealing with complaints and suggestions.

You can get the AI to perform conditional logic quite simply; an example prompt might be:

Read the attached document entitled “survey responses” and categorise each message as either a complaint, a compliment or a suggestion. For each one, apply the relevant process as defined in the attached document entitled “comment processes”.

You can define as many branches as you wish; the process document

could even contain its own sub-branches for dealing with different types of complaint.

The risk, of course, is that you’re relying on ChatGPT to choose the right path in each case. In our example scenario, it may have to make fine decisions as to which category a message fits into, and the more branches and sub-branches you define, the more subjective the decision becomes. You might choose to add an extra rule to the logic such as:

If a message doesn’t unambiguously fit into one of these categories, don’t process it. At the end of the process, show me a list of all unprocessed messages.

Alternatively, you could tell the AI that if a message matches more than one category, it should treat it as multiple messages and apply the relevant categorisation to each.

Designing rule-based prompts

The instruction we added to the logic above is an example of a rule-based prompt: it doesn’t set out a sequence of events like a workflow (“do this, then do this”), but it tells the AI how to handle particular cases. This approach can be useful when, for example, you don’t know what the incoming data might look like, but you know what outcome you’re looking for.

Another simple example of a rule-based prompt might be:

Read the attached document and tell me about any issues reported by at least three customers.

This gets the AI to generate a quick rundown of recurrent comments, without your having to get bogged down in defining what counts as an “issue”, or how much detail you want in your report.

This approach is particularly powerful when you combine rules together. You can enter them into the chat one by one, or type all your rules into an external document and upload it to ChatGPT. The content might be something like:

1. Report all issues mentioned by three or more customers.
2. Report all issues that mention safety.
3. If the most recent mention of an issue is dated before October 2025, ignore it.

You’ll notice that rule 3 could be interpreted as contradicting the preceding rules, which may lead to unpredictable results. It’s a good idea to tell ChatGPT which rules should take precedence in the event of a conflict; in this case, you could add a title at the top of the document saying “Rules in descending order of priority:”.

Testing which rules fail

Using ChatGPT for reasoning tasks can be a trade-off. The AI can give gloriously quick and easy answers – but if you want to ensure they’re accurate and consistent, you have to carefully frame your processes and build in fail-safe mechanisms, and some of the convenience is lost.

If you’re looking for a sensible medium, consider revisiting an idea from section 6. Rather than checking every answer and trying to anticipate every problem, you can use ChatGPT itself to test the robustness of your rules by generating edge cases. For example, you could tell it:

.....

The goal of this process is to highlight recurrent or serious issues that have been recently reported. Generate three messages that these rules might not process correctly.

.....

This approach may seem familiar, since we introduced the technique of working backwards from a goal back in section 2. Here, though, we’re not inviting the AI to generate a process for us – we want it to follow our own defined logic. Rather, we’re using it to try to find cases where that logic doesn’t successfully lead to our stated objective.

There are two ways that ChatGPT might succeed at this. It could reveal that our rules don’t cleanly handle unexpected

Here are **three example survey messages** that the current rule (“report issues mentioned by at least three customers”) could fail to process correctly if the goal is to highlight **recurrent or serious issues that have been recently reported**.

Check your rules by getting ChatGPT to generate possible failure cases.

1. A serious safety issue reported once

“The oven door glass shattered while it was preheating. Luckily nobody was standing nearby.”

Why the rules might fail:

This is a **serious safety problem**, but because it appears only once it would not reach the threshold of three reports and therefore would be ignored.

Why that’s a problem:

The process is designed to surface *serious* issues, but the threshold rule would filter this out even though it may require immediate investigation.

inputs (such as irrelevant issues), or it might show they don’t work as intended with expected ones (for example, we might have specified an unsuitable cut-off date). The fix is often obvious once you see the problem; if not, you can always turn to the AI for advice:

.....

What change or addition can I make to the rules to ensure that these test cases would be handled correctly?

.....

Choosing effective formats

ChatGPT likes to respond to questions in free-flowing prose, often broken up into numbered sections and bullet points. This is easy to read, but it isn’t always ideal for the task at hand.

If your process involves working with numeric data, consider including an instruction that specifies output in the form of a table or graph. Both formats have advantages over a regular text summary – a table is more complete and easier to cross-reference, while a graph lets you instantly grasp the overall shape of the figures. What’s more, these output formats don’t allow the AI to reduce the data to an abbreviated summary, as it can do for text-based output. So you’re not just seeing the data more clearly – you’re bringing more of it into the conversation, and available to be queried and explored.


The AI also supports useful text-based structures, such as comparison tables for an overview of multi-dimensional issues. If you’ve been developing a set of logical rules, you could ask the AI to:

.....

Show me this process in the form of a decision tree.

.....

By default, ChatGPT generates a rudimentary-looking text-based map of the logic, using lines and arrows to show different decision paths. But what this lacks in prettiness it more than makes up for in clarity: this layout makes the logic far easier to follow and check than a series of decisions set out in prose. 🧐



Working with images, audio and documents

Bringing in external media can enrich content and decisions

W

e've talked a lot about uploading documents, but ChatGPT is also good at working with images.

For example, imagine you've been taking photos at a trade event, and you want to catalogue them for use in promotional material. You can simply upload the photo files into ChatGPT and ask it to:

.....
Analyse these images and describe the content of each one.
.....

The AI can produce detailed descriptions, even picking up on signage and other contextual hints to identify the location. There's a catch, however: you can only upload a certain number of files within a set period of time. The precise limit varies, but monthly subscribers might be capped at around 50 files in a three-hour period. So if you want to process a few hundred photos, you could be looking at a long and tedious process.


Thankfully, there's a trick you can use to work around this: just compress your images into a single ZIP file before uploading them. Each upload can be up to around 20MB in size, so you could easily fit 10 or more pictures into one ZIP file, especially if you work from reduced-size copies of the originals.

ChatGPT can produce detailed descriptions of photos, picking up on signage and other contextual hints to identify the location



This photo shows a technology conference stage at the Intel Developer Forum (IDF) 2009.

- A presenter stands on stage speaking to another man who is operating a professional video camera on a tripod.
- The stage backdrop displays "IDF2009 – Intel Developer Forum."



The AI is skilled at picking up context from unlabelled photographs.

Once you've uploaded your files, ask ChatGPT to:

.....
Unzip this archive and describe the content of each image contained inside it.
.....

You can now search the AI-generated descriptions within the chat interface. For example, you could ask for a list of all pictures that show a keynote speaker or a product demonstration. Or, you might ask ChatGPT to generate a downloadable table of keywords, ready to be imported into an existing photo database.

Analysing visual inputs

Although ChatGPT does a good overall job of interpreting images, it has some important limitations. For one, it won't perform facial recognition, so it can't tell you who's in a picture, or confirm that two images show the same person. If you want to find photos of a particular individual, you can try descriptive terms, such as "a man in a blue suit".

The AI may also misinterpret what it sees, especially when it doesn't know the wider context of an image. When this happens, as usual, it will invent a plausible reading of the scene, and then treat that as if it were certain. The result is a mixture of accurate and inaccurate content. We can get better results by adding a confidence check before the summary is generated – for example, try entering:

.....
Analyse this photo, list the image elements you can recognise, and give each one a rating indicating how confident you are that you have correctly identified it.
.....

You'll now see a list of elements and actions, with confidence ratings next to each one. Having done this, you could search your images with terms such as:

.....
List all photographs that you are very confident show a keynote speaker.

Or, you could improve your keywords by asking the AI to list all the image elements it isn't certain about, and manually confirming or correcting any low-confidence identifications.

Working with speech transcripts

ChatGPT has limited audio-processing abilities – for example, it can't reliably analyse music or recognise birdsong. And while you can speak prompts aloud for immediate action, the AI can't currently transcribe uploaded audio files.

Nevertheless, the AI can be useful for working with recordings of meetings or voice notes. You'll just need to use an external service to perform the initial transcription – there are several free online options, or ChatGPT itself can write you a script to do the job on your local computer. No matter how you obtain your raw transcript, you can feed it into the AI and ask it to:

.....
Clean up this transcript.

Bringing multiple document types together

You don't have to use different file types in isolation. ChatGPT can cross-reference information across different media – so, for example, you could ask it to:

.....
Annotate this meeting transcript with references to the relevant passages of these reports.

This creates a more complete record of the meeting, building in material that was known to the participants, but which might not be obvious to a later reader.

Cleaned Meeting Transcript

Date: 26 Feb 2026

Opening

Emma:

Right, I'll get us started if everyone's ready. Thanks for joining, Matthew. I know the funding organisation has you running around like a headless chicken most days.

Matthew:

It's not too bad, Emma. Though I did spend forty minutes this morning trying to find a parking space because they've decided to resurface the entire south lot at once. Absolute nightmare. I nearly turned around and went back to bed—but I'm here now.

Emma:

Good. You didn't miss much. We were just talking about the new coffee machine in the break room that apparently requires a PhD to get an espresso out of.

Raw, choppy notes can be turned into a clean transcript in a single operation.

ChatGPT will remove stutters and stumbles, fix punctuation and formatting, and generally produce a neat, readable version of your transcript. While the AI will try to accurately convey your meaning, it may tweak the phrasing to improve the grammar or sharpen rambling sentences – if you want a verbatim account, it's safest to say so explicitly.

If your uploaded transcript doesn't indicate who the speakers are, ChatGPT will try to figure it out from context, and break up the transcript to show who's speaking. You can help it with a prompt such as:

.....
This is a meeting between Emma (the lab manager), Glen (technical lead) and Matthew (representing the funding organisation). Try to label who's speaking when. Flag any areas where you're unsure.

This probably won't produce perfect results, but it should give you a strong start; you can open the updated transcript in a canvas to make final corrections yourself. You might then get the AI to summarise your transcript, or extract action points as we saw in section 4.

.....
Evaluate the proposed service against the attached criteria.

A word of warning, though: ChatGPT's default instinct is to look for coherence, and to smooth over elements that don't fit together perfectly. If you casually ask it how well the pitch matches up to your requirements, it may focus on common ground, downplaying any mismatches or unresolved issues. Explicitly request a point-by-point consideration, to ensure nothing is swept under the rug. 

Advanced reasoning

Get the AI to analyse its own analysis

To review how realistically achievable a proposal is for your engineering team, the workflow should mirror how experienced engineering managers evaluate delivery risk. The key is to separate facts, assumptions, dependencies, and execution constraints, then combine them into a structured risk-adjusted feasibility score.

A robust logical structure is:

1. Proposal decomposition

First break the proposal into analyzable elements.

Inputs extracted from the proposal

- Core objective

ChatGPT will generate a suitable procedure for any task or decision.

In this section, we've explored a few different reasoning models, and noted that big-picture logic is sometimes more appropriate than an ingeniously intricate approach. However, when you first come to a problem, the best way to tackle it isn't always obvious.

This is an opportunity to use "meta-logic" – getting ChatGPT to figure out a reasoning model for itself. A sample prompt might be:

I want to create an AI process to review project proposals and score them according to how realistically achievable they are for our engineering team. What logical structure should I use to build this workflow?

ChatGPT proposes a multi-stage scoring pipeline, with a few blank spaces for us to fill in key details. It's impressive, but at this point

you'll know that it's always risky to accept the AI's first answer.

For a start, since the workflow was generated by pattern matching, it hasn't been tested or validated in any way. And since ChatGPT hasn't asked any in-depth questions about our team or our scoring priorities, all sorts of assumptions are sure to be hiding beneath the surface.

Still, now that the process exists, we can test it using the techniques we've seen throughout this section. However, there's another factor to consider: when designing a process such as this, ChatGPT doesn't heavily prioritise simplicity. Try asking it to:

Make this logic as simple as possible without significantly affecting the scoring outcomes.

It might well be that the process can be greatly simplified without sacrificing effectiveness – making it much easier to audit, test and refine in the future.

Working with limitations

There's another reason that a ChatGPT-generated workflow might not suit our needs: it's been created without knowing what the inputs might look like. For example, in our case, are we expecting to receive multi-page plans with costs and timescales, or two-sentence summaries?

In the latter case, there isn't much point in building a complicated process. Analytical stages will rely largely on inference and guesswork, making the result no more reliable than a first-glance appraisal. Plus, there are advantages to using simple logic, beyond transparency and clarity. The (effectively) non-deterministic behaviour of the AI means there's a risk that a given operation will generate an unexpected result – so the longer your reasoning chain becomes, the more likely it is to be affected by errors.

If you need to use a lengthy process such as this, you can reduce the risk with a prompt such as:

Identify points of this process where the AI might deviate from the defined workflow, and suggest measures to prevent it doing so.

This can help you build in defensive moves at the most suitable places, such as specifying a strict template for the output or prompting the user to confirm the conclusion before going on. At the end of the day, though, you can't escape the limitations of the AI's reasoning abilities: to get the best from ChatGPT you need to plan around them.

The longer the reasoning chain becomes, the more likely it is to be affected by errors

Use AI to help with a multi-step decision

Big issues can be easier to handle if you use ChatGPT to break them up, and approach them systematically

Here's one way to break the decision down into logical evaluation stages. The idea is to move from **strategic justification** → **practical feasibility** → **operational impact** → **final decision**.

1. Clarify the objective

- What problem are we trying to solve? (developer productivity, faster prototyping, code quality, cost reduction, staffing constraints, etc.)
- Which parts of development are we considering using AI for?
 - code generation

1

Set out the process

Let's use ChatGPT to help us decide whether our business should start using AI for application development. First, we'll ask the AI to break down the decision into logical steps – this establishes a starting framework that we can check and refine to suit our situation.

Several stages in the process depend heavily on **assumptions, estimates or information that may be incomplete**. These tend to be the stages that involve **predicting behaviour or outcomes rather than observing them directly**.

1. Defining the problem you're trying to solve

This step often rests on organisational assumptions:

- Is development actually slower than it should be?
- Are bugs really caused by coding inefficiency rather than poor requirements or testing?
- Would automation help more than process improvements or staffing changes?

2

Highlight weak points

Next, let's ask the AI to flag up the steps where problems are most likely to arise. ChatGPT points out where unwanted assumptions could sneak in, offers possible fixes, and suggests ways we could make our own requirements more clear and unambiguous.

1. The real bottleneck lies elsewhere

The process assumes that improving development efficiency will improve outcomes. But if development is not the limiting factor, the conclusion may be misleading.

Example scenario:

Development time is already shorter than testing, regulatory approval or deployment cycles. AI tools increase coding speed, but overall delivery times remain unchanged. The process concludes the tools are beneficial, yet the organisation gains little practical value.

3

Validate the premises

Even if the internal logic is fine, the process could give misleading results if it's working from an incorrect premise, or if it overlooks external considerations. ChatGPT can help us brainstorm what some of those considerations might be, to ensure nothing has been left out.

Step 7: Assess team impact

Feedback from developers shows:

- some are enthusiastic
- some find AI suggestions distracting
- juniors like explanations and examples
- seniors worry about review burden and skill atrophy
- everyone agrees clear usage rules are needed

4

Work through the steps

Finally, we'll ask the AI to work through each step of the process and, at each stage, explain its conclusion and how it feeds into an overall recommendation. This puts us in a position to make an informed final decision, rather than entrusting that job to the AI. 



Ensuring quality and accuracy

AI isn't much use if you can't trust its responses. In this section, we introduce expert tricks and techniques you can use in any scenario to spot unreliable answers, and even how to get ChatGPT to validate its own work.

102**Validation and stress testing****104****Auditing claims and citations****106****Managing bias and misleading framing****108****Managing long conversations****110****Prompts that produce convincing but misleading answers****112****Responsibility and limits****113****Identify and correct hallucinations****TUTORIAL**

Validation and stress testing

Build stronger prompts to produce better answers

Short answer: **10–16 weeks on site** is a realistic range for a straightforward UK loft conversion, but it can stretch to **4–6 months total elapsed time** once design, approvals and scheduling are included.

Below is a structured estimate with the assumptions that drive it.

Typical timeline (UK loft conversion → bedroom)

1) Pre-construction (often the longest part)

- Survey, design, structural calcs: **2–6 weeks**

ChatGPT will happily produce estimates based on its own assumptions.

It's clear by now that one of the biggest challenges when working with AI is getting consistent results.

In this section, we'll look at some advanced techniques for making ChatGPT's responses as reliable as possible.

One key thing to remember is that asking the same question twice probably won't result in two identical answers. In most cases, the gist will be the same: if you ask ChatGPT to name the capital of Italy, there are millions of different routes it could take to determine an answer, but the statistically likely ones all lead to Rome.

Often, though, you'll be using ChatGPT to help with more complex matters, where there's no single, factual answer. For example, you could upload a plan of your home and say:

Estimate how long it would take to convert my loft into an upstairs bedroom, and list all the assumptions that factor into your answer.

Once you've checked and refined the AI's assumptions about the project, you might think you've done all you can to get a realistic estimate. Repeat the prompt, however, and you may well find that the AI comes up with a totally different interpretation of the question, not just imagining (say) different numbers of windows, but bringing in new factors, such as plumbing and planning permission.

This adds a new dimension to prompting technique – testing the inherent stability of a prompt. Try asking your question multiple times and observing how the AI's process changes each time. You may find you need to add more constraints to define a properly consistent and repeatable approach.

One key thing to remember is that asking the same question twice probably won't result in two identical answers

Rephrasing in neutral language

Aside from directly repeating a prompt, you can test for hidden issues by making small wording changes – even ones that seem like they shouldn't make a difference. We saw in section 2 how seemingly trivial variations in phrasing can affect the AI's framing of the task, and with complex or nuanced questions your wording can have a dramatic effect on the outcome. For example, ask the AI:

Why do startups fail?

You'll probably get a fairly gloomy list of mistakes and misfortunes that can befall a fledgling business. Now try asking:

What factors determine whether a startup will succeed?

This produces a list of positive criteria and advantageous circumstances that could hardly be more different to the first response. This isn't merely a change of emphasis – the two prompts might seem to cover the same ground, but words such as “fail” and “succeed” set the tone for the entire answer, leading the AI to centre certain issues and exclude others.

You can use the AI itself to explore this effect, with a question such as:

How might I have phrased that question to get a different perspective?

ChatGPT will offer some examples of how alternate formulations, using words such as “conditions” or “indicators”, might produce different results. With this in mind, you can intentionally craft your prompts using specific words to push the AI into a particular virtual mindset, and bring out the information and advice you're looking for.

Quantifying confidence

As we've seen opposite, ChatGPT can give a wide range of responses on a topic, depending on the phrasing of your prompt. But once the AI has generated an answer, it will tend to commit to it and treat it as settled.

We looked at how to subvert this tendency in section 3, by inviting ChatGPT to make a certain statement about its own uncertainty. We can take this idea further, and get a better sense of how confident the AI really is about an idea, by bringing comparison into play. For example, if you're trying to decide which language to write a program in, try a prompt such as:

I'm considering writing a web app to play popular card games. Recommend five suitable languages, and give a confidence score and explanation for each one.

Rather than just declaring "you should use TypeScript", this prompt makes the AI show you a selection of options, including ones it wouldn't normally bring into the conversation. You may still end up accepting ChatGPT's top recommendation, but this technique usefully reveals whether it's a clear front-runner, or merely a nose ahead of many similarly ranked options. Seeing the list might also inspire ideas you wouldn't otherwise have thought of – for example, you might not have considered using Python for the back-end, until the suggestion popped up on ChatGPT's list.

The explanations are important, too. These show how the AI arrived at its confidence scores, giving you a chance to apply your own weighting if ChatGPT's assumptions don't quite match your priorities.

Breaking the prompt

We've seen how easy it is to craft a prompt that inadvertently invites skewed responses, or hallucinations. Another risk arises if the line gets blurred between real and hypothetical scenarios. For example, imagine asking the AI:

List the economic benefits of colonising Mars.

For a web app to play popular card games, I'd recommend these five:

1. TypeScript — confidence 9.5/10

This is the safest default. It runs naturally across the modern web stack: browser frontend, server backend, real-time game state, lobbies, matchmaking and shared types. For card games, that shared typing is a big advantage, because you can model decks, turns, actions and rules once with much less room for accidental mismatch between client and server. It also has a huge ecosystem and makes it easy to hire for later.

2. JavaScript — confidence 8/10

JavaScript is still perfectly viable, especially for a smaller or faster-moving project. It lets you build everything in one language and get something playable online quickly. I score it lower than TypeScript because card games often involve lots of state transitions and rules logic, and plain JavaScript makes it easier for subtle bugs to creep in as the codebase grows.

3. C# — confidence 8.5/10

A list of top options gives you a lot more to work with than a single answer.

There are some theoretical economic benefits often discussed in relation to colonising Mars—but it's important to say up front that, with current technology, none are proven to outweigh the costs. Most are long-term or indirect.

Here's a clear breakdown.

1) Technology spillovers

Investment in Mars colonisation would drive advances in:

- Closed-loop life support systems
- Energy efficiency and storage
- Robotics and automation
- Materials science

The AI may accept unlikely premises for the sake of the argument.

ChatGPT will likely respond by envisioning a scenario where colonies on Mars are a done deal, listing the benefits as requested. Unlike what we've seen before, this doesn't mean the AI is working on the basis that such colonies actually exist – if you ask it outright, it will assure you that its answers are speculative. It's just running with your counterfactual premise.

The potential problem arises if you ask about something more obscure, such as a point, law or history with which you're not familiar. The AI may again take your premise as a given, and offer plausible answers that could be mistaken for real-world analysis.

To avoid this, you can use a technique

called "breaking the prompt". This involves getting the AI to identify any weaknesses or defects in the prompt itself, using language such as:

Evaluate and execute this instruction: "List the economic benefits of colonising Mars."

Enter this and you'll get a response that explicitly sets out the speculative nature of the instruction, and the consequent limitations of the listed benefits. This move doesn't need to be limited to counter-factual scenarios – it can provide extra insight into almost any prompt.

Auditing claims and citations

Even if the AI gives a source, it may not be reliable

Key peer-reviewed study

Executive Control of Cognitive Processes in Task Switching

(Rubinstein, Meyer & Evans, 2001 — *Journal of Experimental Psychology*)

- DOI: 10.1037/0096-1523.27.4.763
- Available via [PubMed listing](#) ↗

What the study did:

Participants performed tasks either:

- repeatedly (single-tasking), or
- by alternating between tasks (multitasking / task switching)

What it found:

- Switching between tasks introduced measurable “switching-time costs”
- These costs increase with task complexity [PubMed](#)

Rather than asking the AI to support an argument, get it to locate the research.

What counts as “concentration” in research

Most studies measure:

- Attention / vigilance (staying focused over time)
- Reaction time
- Sustained performance on repetitive tasks

They rarely measure “deep work” or complex reasoning directly.

1) Strong evidence: caffeine improves attention and vigilance

Smith 2002 caffeine review

(*Human Psychopharmacology*, peer-reviewed)

Encourage ChatGPT to critique your own question before trying to answer it.

C

hatGPT rarely gives references or sources for its answers. Often, it can’t: much of its “knowledge”

is synthesised from multiple bits of training data; it doesn’t have a single identifiable origin.

If you need a verifiable answer, you can try asking the AI to provide a citation. For example, you might instruct it to:

Find me a peer-reviewed study showing that multitasking reduces productivity.

ChatGPT promptly comes back with full bibliographic details of a promising-sounding paper, and the journal in which it was published. The only catch is that – as you may already have guessed – it’s quite possible that the named paper might not exist. This is classic hallucination territory: ChatGPT tries to craft a citation that’s perfect for your

needs, copying the formats of existing journal references, without any understanding that it’s supposed to relate to a real-world publication.

This isn’t to say that all citations presented by the AI are bogus, but they all will certainly need to be checked. It’s usually more efficient to take a different tack – try entering:

Search the web for a peer-reviewed study showing that multitasking reduces productivity.

This approach works with fixed external resources, rather than the AI’s own fluid, mix-and-match processes, so it produces much more reliable results. And rather than having to chase down obscure journals, you can instantly check each reference by clicking the URL next to it, while other relevant sources are listed at the bottom of the response to provide background and additional context.

Incomplete references

While ChatGPT isn’t always forthcoming with citations, it does sometimes nod towards external sources. You can see this in action if you ask a question such as:

What evidence is there that coffee improves concentration?

In response, the AI assured us that “there is good evidence that coffee can improve concentration”, and explained that “research has found that caffeine can improve performance” in various tasks.

You might interpret these phrases to mean that the citations are available if you want to check them – but this isn’t always the case. This could be just another exercise in pattern-matching: the AI has learnt that research summaries tend to contain such phrases, so it includes them to make its own answer seem well formed. If you ask it which exact studies it’s referring to, it may admit that it wasn’t citing any particular source – or, it might quickly seek out a reference to retrospectively support the claim. Even if it succeeds, that doesn’t mean its original statement was based on actual evidence.

That’s not to say that the overall

statement is incorrect. ChatGPT usually gives a pretty good account of the big picture, even if the details are debatable or partly fabricated. But if you're looking for reliable results, you need to pin the AI down to specific publications or web pages. Phrases such as "studies have shown" don't mean the AI has multiple citations up its sleeve – on the contrary, they suggest that what it's saying may not be backed up by any evidence at all.

Make a verification checklist

We've cautioned against trusting citations generated by ChatGPT, but sometimes they're the best lead you have. You may also find yourself working with citations extracted from external documents. In either case, check that the cited sources are real and accurately represented.

This can be tedious, repetitive and time-consuming work; the good news is, you may well know of a technology that's ideal for helping with such tasks. You might start by asking the AI to:

.....
Make a list of all references to external sources contained in this document.

This should collect your citations into one place to help you keep track. You can then enter an instruction such as:

.....
Search the web to verify all bibliographical references.

ChatGPT now sets about checking library catalogues and public research archives, and reports back on which citations appear legitimate and which it can't trace. The AI can even parse abstracts and publisher previews to see if the theme matches the citation, although it can't validate specific claims.

If ChatGPT can't confirm a reference, don't give up right away – try broadening the search by, for example, asking if it can match the title of a cited work. When chasing down a citation generated by the AI, we found that the paper did exist, and did contain relevant information – but it wasn't initially found because ChatGPT had seemingly invented a journal date and page numbers to match the expected citation format.

1) The paper itself (anchor reference)

- Effects of caffeine on human behavior (Smith, 2002) ↗

Verification:

- Indexed on PubMed with DOI and full citation PubMed
- Published in *Food and Chemical Toxicology* (Elsevier journal) PubMed

👉 This confirms the document is a **legitimate peer-reviewed review article**.

The AI can search online archives to confirm that cited papers are real.

2) Verification of cited references (overall pattern)

A) Core experimental and review literature

Here are **well-known, peer-reviewed studies that link coffee (caffeine) to concentration-related outcomes**, with tight summaries of what each actually shows.

I've focused on **attention, vigilance, and sustained performance** (i.e. what research typically means by "concentration").

Core experimental and review studies

1) Smith 2002 effects of caffeine on human behavior

Type: Narrative review of many experiments

Don't let ChatGPT get away with just naming studies – demand more detail.

Prompting for traceable evidence

As we've alluded to, ChatGPT doesn't carry out research, as such. If you don't give it a source to work from, it will produce answers based on its own training data – and, in keeping with the AI's preference for neat, coherent responses, any references it includes will usually have been generated to align with those answers. This produces responses that are very persuasive, but sometimes unhelpfully simplistic.

For a more scholarly treatment, you can revisit the technique of withholding the conclusion, as seen in the previous section. In the case of the coffee question we asked opposite, you might start out by entering a neutral instruction such as:

.....
List studies that connect coffee to concentration, and summarise the key findings of each one.

Since we've framed this as a pure information processing task, the AI will happily generate the list without trying to interpret or consolidate the findings of the individual studies. Once ChatGPT has produced its summaries, you can apply the targeted evaluative techniques we saw in section 7, with a prompt such as:

.....
Based on this evidence, evaluate the proposition that coffee aids concentration, and cite the source of each claim or data point.

This should produce a more detailed and verifiable report. Remember, though, that the sources and summaries themselves could be partly fictionalised – consider using the source-checking process discussed above before relying on the AI's judgment. 🧐

Managing bias and misleading framing

You might be skewing ChatGPT's results without realising it

Frame questions in a way that they don't convey a particular mood.

Rewrite this question in terms that don't presuppose a particular answer: "Why do startups fail?"

Here are neutral rewrites that avoid presupposing failure:

- "What factors influence whether a startup succeeds or fails?"
- "What outcomes do startups typically experience, and what drives those outcomes?"
- "What conditions increase or decrease a startup's chances of success?"
- "How do different factors affect startup performance over time?"
- "What patterns are observed in startup success and failure, and what explains them?"

If you want the closest structural match to your original but neutral:

- "What factors contribute to startup outcomes?"

W

e've seen a lot of ways in which the phrasing of your prompt can have an unwanted effect on the response;

by this point you might be feeling overwhelmed by the number of cautions and defensive constraints required to get reliable results from the AI.

To lighten the load, we can delegate some aspects of prompt engineering to ChatGPT itself. One area where it can help is reducing bias: to see this in action, let's go back to the question we asked on page 102, "Why do startups fail?" We saw that this framing made ChatGPT focus on failure, and while the AI was able to suggest alternate phrasings, these arguably didn't address the real issue – they didn't eliminate framing bias, they just changed its focus. It can be valuable to see a situation from

different perspectives, but often you'll want a properly neutral view.

We saw on the previous page that you can reduce bias by withholding the conclusion, and in this case we can get the AI to apply the technique for us. Enter the prompt:

Rewrite this question in terms that don't presuppose a particular answer: "Why do startups fail?"

The AI generates a selection of prompts that you can use to get a rundown of the relevant factors with minimal editorial framing or comment. This technique is great for real-world planning and decision-making, but can be turned to almost any topic – including subjective questions, such as the worst king of France or the best episode of *Friends*.

Simulated debate

If trying to remove bias isn't working, why not try taking an opposite approach? Let the AI take sides – but make it argue from both ends of the issue. For example, you might ask ChatGPT to:

Generate the strongest argument you can in favour of the proposition "Electric vehicles are good". Then, generate the strongest argument you can in opposition.

This approach is powerful because it lets the AI do what it does best – putting a persuasive case for a position. It's up to you whether you let ChatGPT generate its own arguments, or restrict it to supplied sources, allowing you to produce a strictly defined and defensible analysis.

Getting ChatGPT to set out opposing perspectives such as this provides reassurance that important issues aren't being swept under the carpet. Rather than hiding contradictions, as it often does, the AI is now foregrounding them.

The obvious objection is that we're back to seeing different biases, rather than a single coherent view of the topic. But, just as with our scholarly studies, putting the arguments explicitly into the conversation means they're now fixed, and can be analysed and interrogated. This means we can ask ChatGPT to:

Evaluate the two arguments and generate a point-by-point breakdown of which is more persuasive.

The result isn't just a single decision, nor a scrupulously neutral list of facts. Taking this approach gives us a granular, balanced appraisal that shows where each argument is strongest, which can feed into decisions and further research.

Zombie framing

Back in the first section of this magazine, we noted that once an idea has entered the conversational context, it may remain there for a long time. (Strictly speaking, you'll remember that the AI deals in tokens, not ideas – see page 10 – but the upshot is the same.)

This means that framing and instructions can continue to influence ChatGPT long after the conversation moves on. As an example, if you're working on a scientific project, you might tell the AI that you only want to see outliers from a particular data set. When you later move on to examine a second set, the "outliers only" framing could still be remembered, causing ChatGPT to become weirdly reticent about data points.

Even if you switch to a completely different task, old context can interfere with ChatGPT's responses. For instance, imagine you're using the AI to draft a letter to the council complaining about

Framing and instructions can continue to influence the AI long after the conversation moves on

1) Location context (high impact)

- You are in/near Cambridge, UK (from system + memory)
- Previous conversation references:
 - Mill Road (strong local anchor)
- Result:
 - I restricted suggestions to Cambridge restaurants
 - Weighted toward Mill Road / central options

Recent context can influence the AI's answers, even when it's irrelevant.

the closure of the bridge on Mill Road. Once that's done, you ask ChatGPT where you should go for dinner, and guess where it suggests?

You can expose the main elements of context that are feeding into the AI's responses by asking:

.....
 List all the context influencing this answer.

As we've noted, the safest way to escape from "zombie framing" is to start a new conversation each time you open a new topic. Or, as we saw on page 29, if you don't want to lose all your context, you can instruct the AI to forget specific elements – it can't literally do this, but it will pretend that it has, which should be good enough for most purposes.

Challenge your own intentions

As well as helping you check and improve your prompts, ChatGPT can be a useful tool for testing your own framing of the situation. As an example, think back to section 6, where we wanted to use the AI to generate a unique customer reference for everyone in a database. In that section, we got better results by letting the AI design a system to meet our requirements, rather than by trying to define one ourselves. To take this idea further, take another step back and ask:

.....
 I want to create a unique customer reference for everyone in a database. What problem am I really trying to solve?

The AI recognises that the customer reference isn't the real goal – it's just a means to an end. It suggests a range of operational issues you might be aiming to address, which you can discuss and refine to suit your situation. You can also ask ChatGPT to suggest related issues that you might not have considered. Once you and the AI are aligned on your objectives, you can ask it:

.....
 What's the best way of achieving these goals?

It might be that the results are similar to what you would have invented yourself, but now you know you're not barking up the wrong tree. You've effectively validated your own framing, by getting the AI to evaluate not just what you're asking for, but why you're asking for it. 🧠

You're not really trying to "create a unique reference". That's just the mechanism.

The underlying problem is:

How do I reliably identify and refer to the same customer across a system over time?

The real problems you're solving

1) Identity (who is this?)

- Distinguishing one customer from another
- Avoiding collisions (two people treated as the same)

👉 Without this, your data becomes ambiguous

Let ChatGPT interpret your intention to produce a better solution.

Managing long conversations

Keep the AI – and yourself – focused on what’s important

The longer an AI conversation goes on for, the easier it is to lose track of exactly what’s been agreed – for

both parties. For you, the challenge is losing sight of interesting or useful information: ChatGPT often produces detailed, multi-part responses, and some points will naturally get missed or forgotten.

For the AI, the situation is almost the opposite. ChatGPT remembers more or less everything that’s been said in the conversation – with the result that the information and instructions you provide

become increasingly diluted as the discussion goes on.

This effect isn’t usually obvious, because recent tokens generally have a greater weighting than older ones. Whatever you say, ChatGPT will generate a response that’s appropriate to the immediate context. But over time, you may notice that older ideas lose their effectiveness, as they’re crowded out by new inputs.

If your conversation has gone on for a while – say, something like eight to ten exchanges – it’s worth periodically checking the state of the AI. You can do this with a simple instruction such as:

.....
Summarise where we are at right now.
.....

This won’t spell out every element of the context, and that’s a good thing – rather, it gives you a useful indication of the key bits of information that ChatGPT is focusing on, and the main priorities guiding its behaviour. As well as clarifying what you’re working with, and giving you a reminder of the conversation so far, this exercise refocuses the AI’s attention on previously agreed positions and prevents them from fading away.

1) You started with an implementation question

“I want to create a unique customer reference”

We reframed that as:

👉 You’re really solving **identity + referential integrity**, not just generating an ID

2) We clarified the actual goals

You need a system that:

- Uniquely identifies each customer
- Links all related data consistently
- Survives changes (email, name, etc.)

A status check helps clarify what’s driving the AI’s thinking right now.

Branching the conversation

We’ve discussed a few different ways to keep a conversation going while jettisoning unwanted contextual baggage. Here’s one that can be particularly powerful if you need to cut through a rambling conversation and focus in on a specific topic or task.

Simply put, the trick is to get the AI to imagine multiple paths that the conversation could take, and then steer it onto the single one you’re interested in. You can achieve this by entering:

.....
Define three possible ways this conversation could proceed from here, and let me pick one.
.....

This prompt will have the most impact if you deploy it at a clear inflection point – for example, just after you’ve uploaded a collection of documents, or when the AI is offering to generate project parameters. Choosing any one of the proposed branches gives ChatGPT a clear indication of which context elements you want to keep working with, and which you want to leave behind.

Get the AI to imagine multiple paths that a conversation could take, then steer it to the one you’re interested in

This technique works best if the branch you choose genuinely reflects what you want to get out of the conversation. If the AI doesn't give an option that fits your intention, don't just pick one that's sort of similar – if your activity doesn't match your stated intention, it will create ambiguity that works against you. Just ask ChatGPT to generate some different options, and don't be shy to spell out the one you're looking for.

Tag up your inputs

If you want to give ChatGPT standing instructions for a particular workflow, you can mark them in the conversation using tags – for example, you might enter:

[GLOBAL RULE] Only use verbatim quotes from uploaded documents, never rewrite.

Prompts such as this can be dropped in over the course of a conversation to build up a defined set of rules, and if the AI's behaviour starts to drift at any point, you can easily re-emphasise your rules by asking it to restate them all.

While this can be a useful capability, it's important to note that tags aren't a native feature of ChatGPT. There's no special magic in the phrase "global rule", or the capital letters, or the square brackets around it. What you're actually doing is applying a unique semantic marker to specific inputs – which makes it easy for the AI to recognise them as a group. You could achieve the same thing using angle brackets, or a different phrase

such as "standing order". You could even use an unrelated phrase such as "fish": as long as you use it consistently, ChatGPT should recognise that inputs with this tag are all part of a particular group, although it might have a harder time working out your intention.

This trick isn't restricted to instructions – you can use it with information, too. For example, if you're putting together a meeting summary, you could enter:

[FROM JOHN] The business has seen an increasing number of orders from Scotland.

The AI will internalise the information along with its distinctive tag – so you can subsequently use tag-based queries to easily track who contributed what.

Option 1: Concrete system design (practical build)

Turn this into a real implementation plan.

We would define:

- ID format (UUID vs ULID vs numeric + prefix)
- Database schema (customers, merges, external refs)
- Deduplication rules (email, fuzzy matching, manual review)
- Merge strategy and audit trail
- API design (how other systems use the ID)

👉 Outcome: something you could build immediately

Branching the conversation gives ChatGPT clarity as to its current goal.

Step into simulation mode

Another useful pseudo-feature of ChatGPT is what you might call "simulation mode". We call this a "pseudo-feature" because, as with tags, it isn't a capability that's formally built into ChatGPT. However, the AI is capable of understanding the idea well enough that if you instruct it to engage simulation mode, it will role-play as if that mode did exist.

Invoking simulation mode lets you try out questions and instructions that will be effectively set aside once you conclude the simulation. You can ask ChatGPT to role-play in this mode until further notice, or to keep it up for a set number of exchanges. You can also apply it to a single input, with a prompt such as:

Answer this in simulation mode: Would I need planning permission for this loft conversion?

Once ChatGPT has responded to your question, you can resume the original conversation as if the digression never happened. However, the "simulated" exchange isn't really forgotten – it's still in the context, and if it generated any useful outputs then you can freely refer to and use them.

Simulation mode is a great conceptual model for exploring options and ideas without derailing the conversation. If you want an answer that isn't coloured by the existing context, you can also ask ChatGPT to simulate a new conversation, causing it to behave as if you're starting afresh. 🔄

Simulation mode (reasoned, not definitive advice):

You *might not* need planning permission—but it depends on what kind of loft conversion you're proposing.

Likely no planning permission needed (permitted development)

If your conversion is:

- A simple rooflight/Velux conversion
- Or a rear dormer within size limits
- Doesn't alter the front roof slope (facing the street)

Even though there's no such thing as "simulation mode", it works very well.

Prompts that produce convincing but misleading answers

Hidden assumptions and inferences can take many forms

C

ChatGPT is trained to generate responses that appear helpful, relevant and well formed.

More or less everything else proceeds from there: if you ask it a question of fact, it will try to give you an accurate answer, not because the model optimises for accuracy as such, but because it knows that accurate answers are considered helpful.

As we know, however, this doesn't mean the AI only provides accurate answers. It's capable of making speculative inferences – that is, hallucinating – if it thinks that doing so will make its output more helpful and relevant, or if it helps to make a response well formed.

Managing this tendency isn't just about avoiding leading questions, although that certainly doesn't hurt. When drafting a prompt, consider whether you might be inadvertently framing it in a way that makes an inaccurate answer appear more helpful than the truth. For example, you might ask:

What were the key findings of the UK government's 2022 study proving that remote work improves productivity?

As we know, the AI is capable of making speculative inferences if it thinks that doing so will make its output more helpful and relevant...

What the actual UK government evidence says

The key 2022 source is a UK Parliament briefing:

UK Parliament POSTbrief (2022)

Its findings (and related parliamentary analysis) are:

- Evidence on productivity is **mixed and inconclusive**
- There is **no consistent finding of productivity gains or losses** House of Commons Library
- Measuring productivity in knowledge work is **inherently difficult** House of Commons Library

ChatGPT refers to the paper it assumes we meant, concealing our error.

The correct answer is “there was no such study”, but stating that would require the AI to effectively contradict you – something it tries to avoid. It might consider it more helpful to pull together relevant information from other sources, fulfilling the spirit of your request, but potentially creating major issues for accuracy and verifiability.

You can give the AI an “out” in the way we've seen before, by asking it to evaluate the question itself and suggest what a helpful response would look like. This framing allows ChatGPT to express the difficulty, giving you an insight into the accuracy risk and letting you choose a way forward.

Falsifying the response

If you're not sure whether to trust an AI-generated response, you can use ChatGPT itself to challenge it. This doesn't just mean saying “double-check this” – once the AI has generated an answer it will tend to defend it, no matter how specious it might be, unless you actively instruct it to change its position.

A more effective approach is to task the AI with trying to disprove its own claim. For example, let's imagine that you've entered the above prompt in good faith – perhaps your notes got mixed up – and, rather than pushing back, ChatGPT has generated a plausible list of findings. As a simple test of how robust this response is, try entering a prompt such as:

Do your best to falsify the answer you just gave.

Previously, the AI tried to be helpful by smoothing over the gaps between what you asked for and what it was able to provide. This new instruction inverts the motivation, so that being

helpful now means exposing those gaps, and highlighting the fudges and approximations that make the answer unreliable.

If ChatGPT says it can't disprove the proposition, that doesn't prove that the original answer was correct, but it's a reassuring indication that it hangs together logically, and isn't contradicted by other information known to the AI. Even in this case, ChatGPT will do its best to fulfil your request by highlighting any weak or debatable parts of its earlier response – helpful information under any circumstances.

Check what the AI really knows

ChatGPT doesn't "know" anything, in the sense that we'd normally use the words. But there are some subjects where it has a large, consistent body of training information to draw on, and some where it's less well informed.

This can lead to responses that are misleading on multiple levels. If you ask a question about an obscure subject, ChatGPT will respond in its usual confident manner, giving no clue as to whether its answer is an undisputed fact, or an inference scraped together from weak patterns in its training data.

This may sound similar to the effect we discussed opposite, where the AI commits to and defends claims that are actually speculative or even fabricated. But the problem we're talking about here runs deeper: even if you get the AI to verify and challenge its own output, it may not know enough about the topic to identify that there's a problem.

While ChatGPT has its blind spots, however, it does have some awareness of how well supported its claims are. You can make use of this by asking the AI to:

Break down your previous answer into its constituent claims and estimate how qualified you are to make each one.

The results may be eye-opening – an answer that ChatGPT presents as if it were authoritative could easily contain some quite low-confidence elements. These may not spoil the overall usefulness of the response, but knowing which parts they are enables you to handle the output for your purpose.

Claim	Qualification Level
You might enjoy surfing in Cornwall depending on personal factors	Moderate
Enjoyment may come from identity/aesthetic alignment rather than the activity itself	Low–Moderate
People may enjoy activities because they align with a desired self-image	Moderate
Cornwall surfing has an association with outdoor, rugged, coastal identity	Moderate
Surfing involves repeated failure with occasional success	Moderate–High
Some people find this failure–reward pattern motivating, others frustrating	Moderate

Look beneath the surface and you may find the AI's advice is quite shaky.

At a basic level, Google is trying to answer one question:

"If someone searches for this, which page is most likely to satisfy them?"

To do that, it makes a series of judgements.

Banning technical language forces the AI to produce a proper explanation.

1) Does the page actually answer the question?

Google looks at the words on the page and asks:

- Is this page about the same thing the user is searching for?
- Does it directly address the question, or is it only loosely related?

A page that clearly and directly answers the question is more likely to appear near the top.

Limit ChatGPT's vocabulary

We mentioned in section 4 that the AI's fluent rhetorical style can be dangerously persuasive – but you can limit ChatGPT's ability to smooth-talk you by restricting its vocabulary. For example, let's imagine you're exploring business options and you've asked the AI to make a case for staff layoffs. It will probably respond with a slick argument, packed with words such as "efficiency", "growth" and "confidence". You might well come away feeling positively eager to cut the workforce – but are you responding to the argument, or just the framing? To find out, try telling the AI to:

Make the best case you can for staff layoffs, without using any words with positive connotations.

This time, ChatGPT sets out a more sober analysis, positioning layoffs as a suitable tool for specific scenarios.

You can also use this technique to make the AI actively formulate fresh answers, rather than simply reciting received wisdom. For example, if you're promoting a website, you could ask:

Explain how Google decides which website appears at the top of the results, without using SEO terminology.

With jargon taken off the table, ChatGPT has to properly explain complicated issues – making them easier to understand, and more clearly exposing the limits of the AI's knowledge. 🧐

Responsibility and limits

If AI mistakes slip through the net, it's you who's on the hook

N

o matter how intelligent ChatGPT may appear to be – and that can vary depending

on the task – it's important to remember that its responses are inherently, inescapably fallible. In fact, in the web interface there's a reminder sitting permanently below the text-entry field saying "ChatGPT can make mistakes". It's a service that comes with no guarantee of accuracy whatsoever.

It's also a system with some significant limitations. For one, because ChatGPT has no presence in the physical world, it can't observe or measure anything for itself: any experiential information it provides will be second-hand at best, or fabricated at worst.

As we've seen, ChatGPT also treats certainty and confidence as rather fluid and negotiable ideas. When it digs in and defends a position, that can merely mean that it's strongly relevant to the current conversation, rather than that it's likely to be true. It rarely attempts to provide verifiable sources for its statements, and when it does, they aren't always reliable.

All of this means that if you decide to rely on the AI's output in any capacity – such as using it at work, or following advice generated by it – you alone are solely responsible for the consequences. In this section (and throughout the book) we've explored ways to improve the quality and reliability of ChatGPT's output – but, ultimately, you need to be able to independently verify and validate it yourself.

How can I be sure that your answers are accurate and reliable?

You can't be certain—but you can get to **"reliable enough"** by handling answers in a disciplined way:

1) Match scrutiny to stakes

- Low-stakes → quick scan
- High-stakes → full verification (sources, second opinions)

2) Ask for sources and verify them

- Prefer primary/official sources
- Check they actually support the specific claim (not just the topic)

3) Cross-check independently

- Look for agreement across multiple credible sources
- If they diverge, treat the claim as uncertain

4) Check the answer's structure

More trustworthy answers tend to:

ChatGPT freely admits that its answers aren't 100% trustworthy.

Copyright risks

Copyright is a touchy subject in the world of AI. ChatGPT is trained on huge collections of material, some of which is protected by various authors' and publishers' rights, and the legal issues this raises are still being thrashed out.

That might sound like a problem for ChatGPT's operators, but it could also carry a risk for you. The danger is that, based on its training data, the AI could generate content that closely resembles an existing copyrighted work – and if you publish or share it, you could be unwittingly breaching copyright yourself.

This is a tricky issue to deal with, as the AI itself can't identify the individual sources of the content it generates – so you may not realise a passage has been effectively plagiarised, until you get a letter from an upset author.

Thankfully, so far, we haven't heard of any individuals being sued over this type of infringement. In fact, it's unlikely that any LLM would replicate complete sentences from a single source, just because of the way these AI systems work – and even if you ask it outright, ChatGPT will only reproduce short passages from specific works.

At the end of the day, though, responsibility again lies with you. ChatGPT can suggest a few ways to check generated text for plagiarism, but if you want to be certain that you won't be on the hook for AI-assisted breach of copyright, it's best to use ChatGPT for planning and drafting assistance, and to write the final copy out yourself in your own words. 📝

Identify and correct hallucinations

The AI may try to conceal its mistakes, but you can get it to do so with a sceptical attitude and the right prompting technique

On 2 January 1984, Hewlett-Packard made a pivotal strategic move: it introduced its first LaserJet printer, the HP LaserJet.

Why this mattered



- It was one of the first affordable laser printers aimed at businesses

1 The initial question

Let's imagine we're putting together a presentation about a milestone year in personal computing. We think something important happened on the first working day of the year, but we're not certain of the details, so we ask ChatGPT to refresh our memory.

The model launched on 2 January 1984 by Hewlett-Packard was the HP LaserJet—often retrospectively referred to as the HP LaserJet (sometimes "LaserJet Classic").

What distinguished this specific model



- Print speed: ~8 pages per minute

2 The AI doubles down

Its response seems fishy – that would be a strange date on which to launch this product. We can try to investigate the claim by digging into it, but ChatGPT sticks to its guns, piling on additional detail that seems to confirm that its assertion is well grounded.

1) The date isn't consistently cited in authoritative histories

Most reliable accounts of the HP LaserJet launch:

- Agree on the year (1984)
- Often mention early 1984
- Do not give a specific day, let alone 1 or 2 January

This includes corporate histories, retrospectives, and industry write-ups. The absence of a precise date is a red flag for a claim that pins it to a specific day.

3 Test the negative

To make the AI reconsider its previous answer we need a different approach. Instead of asking for more confirmation, we ask ChatGPT if it can disprove the proposition – which it does immediately, shamelessly pointing out its own procedural flaws.

1) Separate facts from precision


Before writing anything, classify each item:

- A — Exact date confirmed (safe to use day/month/year)
- B — Approximate timing only (year / "early 1984")
- C — Uncertain / conflicting

Your LaserJet entry should have been B, not A.

4 Improve the process

There's no universal cure for hallucinations, but now that this error is part of the conversation, ChatGPT should remember it and be less likely to repeat it. The AI can also suggest ways to reduce the likelihood of similar errors recurring in the future, such as accounting for uncertain dates.



Professional mastery and the future

If you're using ChatGPT for business, you need to do so in a way that's transparent and accountable. In this section, we see how to create governance frameworks that reduce risk and offer flexibility for future uses of AI.

116

AI and governance considerations

118

Building and maintaining a prompt library

120

Real-world case studies

122

Autonomous systems, and what comes next

123

Create an autonomous task with guardrails

TUTORIAL

AI and governance considerations

Don't let ChatGPT leave you exposed

A

fter reading section 10, you might wonder whether it's safe to use AI in a professional

setting at all, since its accuracy can't be guaranteed. But the same concern could apply to business software (which might have bugs), or human workers (who might have bad days). The solution isn't to eschew imperfect tools, but to create a governance framework to manage the risk.

One part of that is deciding what tasks can be entrusted to the AI, and what needs to be handled by an accountable person. For example, let's say you wanted to review existing research into a topic and propose possible avenues for further exploration. It would be a good idea to get a qualified human to compile the source material, as they'd be likely to do a complete and reliable job; but there's no harm in asking

ChatGPT to suggest possible ways forward.

The best part is that deciding how to divide up the workload is a task that the AI can help with. For example, you might ask it:



I want to create a report making the case for extending our research into new areas of biosciences. Break down this job into sub-tasks and identify which ones can be handled by AI and which ought to be performed by a human.

Assessing the AI's answer is itself a task that ought to be tackled by a person – but ChatGPT provides an explanation for each of its recommendations, so you can make an informed decision whether to reduce your own workload by delegating more tasks to the AI, or improve oversight with more human involvement.

If the goal is a credible, decision-ready report, the key is to split the work into stages where AI can accelerate analysis—but not replace judgement, domain expertise, or accountability.

Here's a clean breakdown.



1) Define scope and objectives

What this involves

- Clarifying:
 - which bioscience areas are in scope
 - time horizon (e.g. 5–10 years)
 - business goals (growth, risk diversification, etc.)

Break down tasks into what the AI can do and what needs a human eye.

Establish the ground truth

As we've indicated, a potential hazard of using ChatGPT is that you don't know exactly where its information comes from. That may not be important when you're dealing with basic facts, such as the capital of France, but it can be a concern for tackling complex questions such as the likely impact of layoffs; you can't be sure that the AI isn't bringing in ideas that make more sense in other industries, or jurisdictions.

While it's risky to let ChatGPT's internal wisdom guide your professional decisions, you can improve things by providing the AI with your own source of truth. This can be as simple as uploading a set of financial reports, telling the AI to:



Base your advice solely on these documents.

This dissuades the AI from leaning on its own virtual expertise; any inferences it makes should now be grounded firmly in your balance sheet, and any action you take based on its output can be meaningfully audited and accounted for.

Your "ground truth" doesn't need to be limited to facts and data. You could also upload shareholder agreements and governance procedures, and tell the AI that decisions must be made strictly in accordance with those principles. You still can't guarantee that ChatGPT won't try to put its own spin on things, as we'll discuss opposite – but if you ask it to cite specific references when recommending a course of action, that gives you an easy way to check that it's following the rules.

You can provide ChatGPT with your own source of truth...

Acceptable variation

Defining an official source of truth helps to rein in ChatGPT's behaviour, but any instructions you give the AI tend to get weaker over time. Sooner or later, it will likely start to revert to its old habits of changing phrasings, brushing over details, and offering helpful suggestions when you've explicitly told it not to.

Since these behaviours can't be fully eliminated, you need to decide how far you're willing to accept them. For example, you might decide that it's all right if the AI doesn't strictly follow all the procedures set out in your handbook, as long as it observes the key principles set out on the first page. Or, you might be happy for it to quietly simplify its working model of your finances, as long as it preserves the relative rankings of different income streams.

These are important considerations for governance, as they provide a clear metric for determining whether or not the AI output is acceptable. And it certainly doesn't hurt to share them with ChatGPT, too. This still doesn't ensure that the AI will actually keep to your defined tolerances, but it does let you ask:

Did your last answer conform to our standards of acceptable variation?

This prompt won't catch everything, and it's important to remember that an output can be within defined tolerances and still be wrong. Even so, your variation document gives the AI a useful fixed reference for assessing whether its own behaviour or framing have drifted too far.

Audit trails and transparency

For proper accountability, it isn't enough just to check that ChatGPT is working from the right information. Important decisions need to be auditable, so that someone can look back in the future and understand how they were made, perhaps for reasons of regulatory compliance, or to build on past processes for a new purpose.

This is a challenge with ChatGPT, since its answers can be subtly influenced by context, or by ideas grounded in its training data. On pages 92 and 93, we discussed generating process documents to make answers repeatable,

and a similar technique can be used in recording AI-assisted decisions. In this case, the goal isn't necessarily to enable someone else to repeat the prompt, but to provide transparency into how a conclusion was reached – and, perhaps, what wasn't considered.

To standardise the process, it might be helpful to define a template for transparency documents, with headings such as "inputs and constraints", "sources of operative claims", "background assumptions", and so forth. You can create this in a word processor, then upload it to ChatGPT and instruct it to:

Short answer: **partially, but not fully.**

I'll assess it strictly against the standards we defined.

1) Scope fidelity

Pass

- I used only the provided dataset
- I did not introduce external claims
- I preserved the core structure of the data (baseline + outliers)

👉 No scope violations

Get ChatGPT to report on whether its own operations met your criteria.

What was the task you were asked to do?

Evaluate whether the previous summary of the lab data conformed to the defined standards of acceptable variation.

What were the inputs and constraints?

- Input:
 - The pasted lab dataset (50 records with glucose, haemoglobin, and cholesterol values)
 - The previously generated summary of that dataset
- Constraints:
 - Standards of acceptable variation, including:
 - $\pm 15\%$ numeric accuracy
 - preservation of relative scale
 - no distortion of conclusions, risks, or constraints
 - Requirement to assess against scope, completeness, and risk criteria

Keep the AI accountable by generating a record of completed tasks.

Fill in this form, referring to the answer you just gave.

By default the AI will respond in the chat, but you can ask it to create a downloadable version, perhaps including the date and a summary of the discussion. This provides a valuable record for future reference – and as you build up a collection of transparency forms, you can also ask ChatGPT to analyse them and identify any emerging trends about your processes and your use of AI. 🧠

Building and maintaining a prompt library

Setting up templates helps to keep your processes consistent

Task: What do you want me to produce?

Objective: What decision or outcome will this support?

Scope: Include []; Exclude []

Inputs (authoritative only): What documents/data should I use?

Constraints: Any rules, limits, or format requirements?

Standards: (keep as-is unless you want changes)

Uncertainty: Any specific expectations (e.g. "no interpretation", "flag all assumptions")?

A basic prompt template ensures relevant details aren't missed out.

Defining multiple templates

Once you've created your first prompt template, the question arises – how far should you take the idea? It's probably fine to let people use the AI freely to look up facts and brainstorm ideas, but it might be a good idea to define a set of standard prompts for common business tasks. For example, you might create a template for summarising documents, another for fact-checking, and a third for making decisions.

All of your templates can then be combined into a standard shared document that anyone can access with a prompt such as:

.....
Use the "fact check" template from the attached document to check the investor handbook.
.....

Sharing a document such as this doesn't just help to keep your processes consistent, it gives clarity as to what people can use AI for – and how, since each template can include details and constraints appropriate to the task. For example, a prompt template might tell the AI to follow a standard output format, making it easy to compare reports even if they were generated by different people.

You could also add logic that checks whether the task can be satisfactorily accomplished with the information given, and refuses to proceed if it can't meet the standard of precision and verifiability specified for that particular task. Or, you could insert a self-check step, where the AI verifies its own response at the end of the process, and confirm that it meets requirements. A clever final step might be to automatically generate a transparency document each time the decision-making template is used, and prompt the user to forward a copy to the appropriate administrator.

W

e've talked about the benefits of establishing standard processes for AI interactions, and

when multiple people are using ChatGPT for a common purpose, this becomes not only helpful but arguably necessary.

The reason is simple: anyone who's ever been in a meeting knows that two people will often talk about a shared goal or observation in quite different terms. As a human being, you can intuitively reconcile people's individual communication styles – but ChatGPT won't, at least not consistently and across conversations. On the contrary, we've seen that the wording of a question can have a dramatic impact on how the AI interprets and answers it. It's easy to envisage a scenario where two people in an organisation mean to ask the same

question but get very different answers – turning ChatGPT from a decision-aiding tool to a source of mixed messages and conflict.

You can reduce this risk by creating standard prompt templates, which everyone can use to interact with the AI in a reasonably consistent way. Workers can fill in the blanks directly, or work interactively by pasting the template into the AI and prompting it to:

.....
Ask me what you need to know to solve my problem according to this template.
.....

As well as reducing unintended bias, standard structures help to ensure the AI gets all the information it needs to produce useful responses, and makes it easier to maintain ground truths and auditability.

Versioning and updating prompts

Maintaining a standard set of prompt templates promotes consistency, but ChatGPT itself is an evolving platform. Since 2023, two major releases and numerous smaller updates have all changed its behaviour in various ways, and you may need to periodically update your templates to keep them working in the way you wanted.

This creates potential issues for auditing and governance, but you can reduce the pain with some simple version management. Start by putting a version number at the top of your template document (such as “v1.0”), and making each template mention this number in its output. If you need to tweak a prompt template then – after suitable testing – you can edit the document accordingly, update the version number and recirculate it.

For maximum transparency you could give each template its own version number. When you want to make an update, leave the existing template in place, and create a new version with its own number. Then, add a line at the top of the file telling the AI to use the latest version of all templates unless instructed otherwise.

Now, the latest template version is available to everyone, but the older version remains available for inspection or comparison with a prompt such as:

.....
How does the latest version of the text-summarising template differ from previous versions?

Remember that ChatGPT’s responses are highly context-dependent, so if you try to compare two versions of a template by invoking them one after another, the results of the second run may well be influenced by the first.

Using the latest **text-summarising template 2.0** Prompt templates.txt, the main failure modes come from inputs that don’t match its assumptions about structure, clarity, or completeness.

Here are the key categories of problematic input:

1) Texts without a clear thesis or structure

The template assumes:

Use ChatGPT itself to check and sharpen your templates.

Using your prompt templates document Prompt templates.txt, here are the **key differences between the latest text-summarising template (2.0) and the earlier version:**

1) Stronger focus on decision context

New (2.0):

- “Purpose: what decision or use-case will this summary support?”

Old:

- “Purpose: what will this summary be used for?”

👉 The new version explicitly ties the summary to **decision-making**, not just general use.

2) More precise scope definition

New:

- Includes “core arguments / data / conclusions”
- Excludes “examples / anecdotes / peripheral detail”

You can develop and audit updated versions of your prompt templates.

Phrasing templates to minimise prompt rot

Even if the GPT model doesn’t change, you may find that your prompt templates lose effectiveness over time. Edge cases and loopholes might emerge, or your company could start working with new documents and data formats that your standard procedures don’t handle properly.

You can stave off such problems by phrasing your templates as robustly as possible. Remember the principles we looked at back in section 2 – be explicit and unambiguous about your intention, and about what you don’t want. Don’t shy away from being direct, even to the point of redundancy. For example,

instead of writing “don’t invent new content”, your template might say:

.....
Don’t include any information in the summary that isn’t in the original document. If you think of any details you would like to add, omit them.

You can also help the AI to understand your intention by setting it out methodically: start with a description of the desired output, detail the expected inputs, and then spell out the scope and constraints of the task. You can get ChatGPT to help you spot possible weaknesses in the specification by borrowing a technique from section 6:

.....
What sort of input might cause the attached template to produce unexpected results?

There’s no need to wait for “prompt rot” to set in before you ask this question: you can use it pre-emptively to tighten up your templates before they start to produce poor responses. 🧠

Real-world case studies

How four major companies have brought AI into their business

In this magazine, we've focused on imaginary, illustrative scenarios, but be assured that these principles apply in the real world. For example, in section 10, we talked about the risks of the AI fabricating or misrepresenting sources, and these were highlighted in a recent controversy involving professional services firm Deloitte.

In 2024, Deloitte was commissioned to produce a 10-year healthcare staffing plan for the provincial government of Newfoundland and Labrador in Canada. It was delivered in May 2025, at a cost of around \$1.6M (around £870,000). Shortly afterward, however, academics at the Memorial University of Newfoundland reported that they couldn't verify at least 15 cited sources. They also highlighted a reference to a 2008 film called "Schoolyard Games" – a fictitious title that doesn't exist, but which does match an example citation used in an online style guide.


The researchers suggested that these issues were consistent with the use of generative AI, and a spokesperson for Deloitte Canada confirmed that AI had been "selectively used to support a small number of research citations". While the company stated that the use of AI hadn't affected its findings, it agreed to revise the report.

This case should be a reminder that AI-generated citations can look highly plausible, to the point that professionals may overlook fabricated or untraceable references. However, these can often be caught by simple steps – such as asking ChatGPT to check the list of references for itself, as we did on page 105. Failing to check sources may expose a company to reputational damage, and wider business consequences.

Sources:

- <https://www.cbc.ca/news/canada/newfoundland-labrador/wakeham-health-report-sources-9.6992094>
- <https://www.entrepreneur.com/business-news/deloitte-detected-using-fake-ai-citations-in-1-million/500072>

Failing to check sources may expose a company to reputational damage...



Morgan Stanley sets its own ground truth

Finance giant Morgan Stanley wanted to develop an AI-powered chat assistant for internal use, allowing employees in its wealth-management division to quickly access information and insights. Since finance is a highly regulated, high-risk environment, consistency and auditability weren't just desirable – in this case, those requirements were non-negotiable.

The solution was to work from a ring-fenced source of ground truth. Launched in 2024, the Morgan Stanley AI assistant works from a defined library of approved documents and processes to provide tailor-made answers to employee questions. But it's designed to refer only to internal content – it won't even attempt to answer questions that don't relate directly to wealth management. To reduce the risk of misunderstandings, it won't even use analogies or metaphors.

The system also incorporates

a principle we've touched on throughout this book – progressively refining the AI's behaviour and iterating towards better results. Morgan Stanley advisors can give each interaction with the AI a thumbs-up or thumbs-down rating, to support a system of continuous improvement.

It's worth noting that Morgan Stanley worked in partnership with OpenAI, the creator of ChatGPT, to build this custom system. If you're using the regular ChatGPT service then you may not be able to enforce such strict content rules, and there's no native thumbs-up/thumbs-down feedback mechanism. All the same, using clearly defined sources will help you keep your interactions focused and consistent, and you can regularly review users' feedback on their AI experiences, perhaps to create new, enhanced versions of your prompt templates.

Sources:

- <https://www.morganstanley.com/press-releases/key-milestone-in-innovation-journey-with-openai>
- <https://www.businessinsider.com/morgan-stanley-ai-chatgpt-used-nearly-all-wealth-management-teams-2024-4>

AI gives an Air Canada customer the wrong information

Some AI chatbots, such as that of Morgan Stanley, are intended for internal use. Others are designed to talk to the public, dealing with pre-sales enquiries and first-line support issues that don't require dedicated human attention.

Air Canada is one company that offers an AI chatbot on its website to handle (among other things) pricing enquiries. In 2022, a traveller named Jake Moffatt logged on and asked this chatbot how to apply for a discounted "bereavement fare" to attend his grandmother's funeral.

The AI agent advised Moffatt that he could book the ticket as usual, then apply for a discount within 90 days of travelling. However, when he subsequently made the application, the airline rejected it, stating that bereavement discounts had to be agreed before travelling.

Mr Moffatt issued a legal claim to get the discount he believed he was entitled to. Air Canada argued that the airline wasn't responsible for the incorrect information provided by the AI – but in 2024, the Canadian Civil Resolution Tribunal firmly rejected that argument, ruling that Air Canada must honour the promise made by its AI representative. "It should be obvious to Air Canada that it is responsible for all the information on its website," the ruling noted. "It makes no difference whether the information comes from a static page or a chatbot."

This provides a clear illustration of the principle we referred to in section 10 – that you're accountable (or your business is) for what the AI does on your behalf, including any mistakes it may make. Even if you're not operating a customer-facing chatbot, this could extend to claims, terms and offers included in any sort of communication, so safeguards and careful checking are essential.

Sources:

- <https://www.bbc.co.uk/travel/article/20240222-air-canada-chatbot-misinformation-what-travellers-should-know>
- <https://decisions.civilresolutionbc.ca/crt/crtd/en/item/525448/index.do>

Your business could be held responsible for the mistakes of an AI agent.



By Maria Yagoda 23rd February 2024

When Air Canada's chatbot gave incorrect information to a traveller, the airline argued its chatbot is "responsible for its own actions".

Artificial intelligence is having a growing impact on the way we travel, and a remarkable new case shows what AI-powered chatbots can get wrong – and who should pay. In 2022, Air Canada's chatbot promised a discount that wasn't available to passenger Jake Moffatt, who was assured that he could book a full-fare flight for his grandmother's funeral and then apply for a bereavement fare after the fact.

Mastercard embeds AI governance to support production-grade financial systems



Mastercard builds AI into a governance framework, not vice versa

Mastercard is an industry leader in AI adoption. It started developing an AI governance programme back in 2019, well before systems such as ChatGPT became commonplace, and in 2021 the company appointed its first head of AI governance, John Hearty.

Hearty's team manages AI governance policy in partnership with technology specialists, legal teams and product managers from across the company. This has allowed Mastercard to establish overarching standards for AI usage, using a risk-based framework that matches controls and oversight to the potential impact of each AI application. When a part of the business wants to introduce a new application of AI, the proposal is evaluated using a "scorecard", detailing the specific data it wants to process with AI, and the techniques it intends to use.

One notable success has

been Mastercard's Decision Intelligence system, which uses AI methods developed and tested in-house to identify potentially fraudulent transactions; the company claims it's reduced false positives by more than 60%, compared to pre-AI methods. For those who interact directly with AI, the governance team sets out explicit "dos and don'ts", plus recommendations such as cross-checking results via different methods, rather than accepting the first response from a generative AI system.

While not every business has the resources to set up a dedicated AI governance team, there are lessons here that most organisations can learn from. Rather than trying to control the genie once it's already out of the bottle, Mastercard shows it's possible to get ahead of demand, and design a framework that lets you introduce AI functions into your business in a managed way.

Sources:

- <https://www.mastercard.com/news/perspectives/2023/driving-responsible-ai-how-to-set-the-rules-of-the-road/>
- <https://www.theasianbanker.com/updates-and-articles/mastercard-embeds-ai-governance-to-support-production-grade-financial-systems>
- <https://www.dataversity.net/case-studies/case-study-operationalizing-ai-governance-at-mastercard/>

Autonomous systems, and what comes next

Agent mode opens up new possibilities, and new challenges

The next big step in AI is already at your fingertips: below ChatGPT's text-entry box, you'll see an innocent-looking icon, introduced in recent updates, that activates Agent mode. (In the desktop app, it's a sort of rounded square with an arrow pointer in the corner.) Agent mode is like goal-orientated prompting on steroids. It lets the AI figure out for itself how to break up a complex task into steps, and then work through them with minimal involvement from you. For example, let's imagine that you ask the AI to:

.....
Research and compare IP laws across Europe, and produce a summary of key differences.
.....

In ChatGPT's standard operating mode, the AI will quite likely rely on its own internal knowledge, and generate a set of comparative bullet points. In Agent mode, the AI can dig much more deeply: it can independently carry out online research, collate and analyse its findings, and actively hunt for data to fill in any gaps. The result is likely to be a far more detailed and consistently researched report.

With the capabilities of Agent mode, ChatGPT could conceivably take over entire job functions that are currently handled by a human. However, it also introduces new risks. Since the AI works its own way through the process, it can be harder to spot mistakes – especially subtle ones that take effect early in the task, and affect everything that comes afterward.

Here's my proposed workflow for comparing IP laws across Europe using **only official government sources**:

1. Source Identification (Stage 1)

- Search for official government websites relevant to European IP law (e.g., European Patent Office (EPO), European Union Intellectual Property Office (EUIPO), national IP offices).
- Verify that each website is an official government or intergovernmental source (no private law firms or blogs).
- Compile a list of these verified sources for patents, trademarks and copyrights.

2. Patent Law Research (Stage 2)

- From the official sources identified, gather key information on how patents are granted, validated and enforced across Europe (including differences between EU-level and national systems).
- Summarise findings in a concise report.
- Pause and report progress to you for feedback before moving to

In Agent mode ChatGPT draws up its own plan of action.

Autonomy and governance

In Agent mode you can no longer rely on prompt-level governance. While performing an agentic task, ChatGPT may define sub-tasks and processes for itself that aren't visible to the user – so you get no opportunity to check or manage their scope and constraints. And yet, at the same time, you're still responsible for its outputs and conclusions. As such, to minimise the risks of using Agent mode you need to build procedural safeguards into your prompt. An example might be:

.....
Compare IP laws across Europe using only official government sources; outline your proposed workflow first, pause for user approval before continuing, and report your progress after each stage.
.....

This approach can lead to some quite unwieldy prompts. Thankfully, it's

also possible to manage agentic tasks using templates, in a similar way to what we've already seen. The difference is that rather than defining specific prompt patterns, your agentic template documents should set out broad policies. For example, you might declare that for low-risk activities ChatGPT can use its own internal knowledge, but that only a supplied list of trusted websites can be used to inform important decisions.

You can also build in observability via checkpoints. You might tell the AI to output a progress report at each stage in a task, and require any high-risk decisions to be explicitly confirmed by the user.

The important thing to realise is that agentic tasks make it more difficult to determine how ChatGPT got from the prompt to the output. This means auditing and governance can't be reliably applied retroactively; as AI gets smarter and more autonomous, your tasks need to be designed with transparency as an essential function. 🧐

Create an autonomous task with guardrails

Using Agent mode allows the AI to carry out deeper research with more precise criteria

Best current deals (quick picks)

Budget (~£250–£400)



Here are several current business-laptop deals (prices and offers valid in the UK in late-March 2026). I've grouped them by budget and highlighted why each one is notable.

Entry-level/portable (£150–£450)

Laptop & key specs	Price & source
Lenovo IdeaPad Slim 3 Chromebook (14") – MediaTek Kompanio 520 CPU, 4 GB RAM, 128 GB eMMC, 1.3 kg which.co.uk	≈£159 at retailers (cheapest price r by Which? for March 2026 which.co.uk
HP 14a Chromebook (14") – Intel Celeron, 4 GB RAM, 64 GB storage techradar.com	£179 at Currys (down from £229) techradar.com

1 The basic answer

Let's try using AI to help with a real-world business task – choosing new laptops for the office. When we enter our prompt in ChatGPT's standard mode, the AI quickly skims the web and comes back with headline offerings from big retailers such as Argos and Currys.

2 Independent research

Repeating the instruction in Agent mode gets us a much more detailed response. Now ChatGPT explores many more sources, including reviews sites and manufacturer web pages. Rather than suggesting popular deals, it identifies specific models that it thinks will appeal to us.

Worked for 8 minutes >

Here are business-oriented laptops with built-in AI hardware (NPUs) providing at least 40 TOPS, along with current UK pricing and notable specifications. For each machine I've included a source confirming the NPU's TOPS and a current price from a reputable retailer.

Laptops with Qualcomm Snapdragon X-series (45 TOPS NPU)

Model & key specs	AI performance & evidence
Lenovo ThinkPad T14 G6 – 14" WUXGA (1920×1200), 32 GB RAM, 1 TB SSD, 1.24 kg laptopdirect.co.uk	Uses Qualcomm Snapdragon X Elite Qualcomm's X Plus/Elite processors: 45 TOPS Hexagon NPU laptopmag.com

Here are four highly-rated business laptops with ≥40 TOPS AI NPUs that balance performance with reliability. Each model is built around Qualcomm's Snapdragon X or AMD's Ryzen AI processors (45 TOPS and 50 TOPS NPUs respectively) and is aimed at professional use. I've summarised the key reliability features and highlighted the main weaknesses to consider.

Business laptop (≥40 TOPS NPU)	Evidence of reliability
Lenovo ThinkPad T14 G6 (Snapdragon X Elite, 32 GB RAM, 1 TB SSD)	Lenovo's product documentation states that the ThinkPad T14 G6 passed MIL-STD-810H military testing parel.lenovo.com indicating its chassis and components are built to withstand vibration, temperature changes and shocks. The X Elite chip gi

3 Searching for specifics

We've decided we want computers with built-in AI hardware, so we now ask ChatGPT to narrow down the list to suitable candidates. Again, in Agent mode the AI does independent research, working out for itself what we need and puts together a shortlist with details of each one's AI capabilities.

4 The final selection

Finally, let's ask ChatGPT to focus on reliability, and to highlight any concerns the shortlisted models may have in this area. Again, we don't need to define reliability, or specify what the most important metrics are – the AI works that out for itself to generate its final set of suggestions. GPT

Prompt reference

Keep these templates handy for reliable, reusable prompting

Define the task	
Prompt	Effect
For context: I need to give a presentation tomorrow.	Give the AI background information with no task attached, to help it interpret your next instruction – see pages 20-21.
Explain this for the head of marketing.	Generate a summary of a response or document that's geared to the interests of a particular person or purpose – see page 36.
Base your answer solely on these documents.	Limit the AI to your own ground truth, so the information that feeds into decisions can be examined and verified – see page 116.
Refine and improve	
Prompt	Effect
Refine this answer to focus only on the project requirements.	Steer the AI back to the main objective, if the conversation has drifted off on a tangent or lost clarity – see page 29.
How does this relate to my original question?	Detect drift by getting the AI to explain an answer, and restate what it thinks is the operative question – see page 28.
Give me a one-sentence summary, then a run-down of the main points of the argument.	Make ChatGPT identify the core message of a passage before generating a summary – see pages 18-19.
Identify the three strongest ways this could be improved.	Force ChatGPT to prioritise suggestions, so you don't just get a list of minor corrections – see pages 78-79.
Highlight any inconsistencies or contradictions in this document.	Expose subtle variations or mismatches in different parts of a document that might otherwise go unnoticed – see page 37.
Control output	
Prompt	Effect
Using child-friendly language, summarise this document.	Put the framing at the start of the prompt to give it primary importance – see page 19.
Read this and tell me about any factual errors you find.	Use strong action verbs to ensure there's no ambiguity about what you want the AI to do – see pages 18-19.
Recommend five options, with a confidence score and explanation for each.	See why the AI recommends each option, so you can accept or overrule its judgment – see page 103.
Highlight any parts of this answer that are uncertain or speculative.	Reveal any shaky claims that might have been included for coherence rather than certainty – see pages 102-103.

Explore and expand

Prompt	Effect
Generate three alternative approaches without favouring any one of them.	Prevent ChatGPT from prematurely choosing and committing to one particular answer – see pages 76–77.
Frame these figures as a growth opportunity for the business.	Find the strongest way of reading a certain set of information to support a particular perspective – see page 37.
How might I have phrased that question to get a different perspective?	Expose potential limitations and hidden bias encoded into a particular prompt phrasing – see page 106.
What questions do these patterns or findings raise?	Get ChatGPT to suggest new lines of enquiry based on emerging trends or gaps in the data – see pages 48–49.
Suggest possible answers to these questions, based only on the information available.	Propose hypotheses that account for the observed data without introducing new ideas – see pages 48–49.

Test and challenge

Prompt	Effect
Which pieces of evidence support this conclusion, and which challenge it?	Make the AI expose both sides of an argument, rather than picking one side to come down on – see pages 46–47
What aspects of this conclusion are poorly supported?	Identify the weaknesses of a position, regardless of whether it was supplied by you or generated by ChatGPT – see page 47
What's missing from this data?	Detect gaps in supplied information that could make a response unreliable or cause errors – see page 47
Evaluate and execute this instruction: "[prompt]"	Get the AI to warn you if a prompt seems ambiguous or poorly framed – see page 103
Did your last answer conform to our standards of acceptable variation?	Force ChatGPT to reveal how accurate it thinks its own response was – see page 117

Structure and logic

Prompt	Effect
Create a timetable template for my one-day workshop.	Start with a simple project framework, then build it up through iteration and progressive refinement – see pages 28–29.
Segment this data into groups that have something in common.	Find a structure within a set of information, either for analysis or just for ease of reference – see page 46.
List the steps you would take to solve this problem. Do not give a conclusion.	Stop the AI from jumping straight to a solution by making it spell out a workflow – see pages 88–89.
After each step, summarise what it confirms and what it rules out.	Make the reasoning process observable, forcing the AI to “think” properly and not take shortcuts – see page 89.
List the inputs, logical operation and conclusion for each step.	Formalise ChatGPT’s logical analysis, so you can follow and check its reasoning for yourself – see page 89. 🧠

Checklists and fixes

Principles and prompts to help you get high-quality answers

Before you enter your prompt, ask yourself...

- What does my ideal answer look like?
- Am I asking for too many things at once?
- What am I taking for granted?
- What does the AI need to know to do this properly?
- Do I want a final answer, or an insight into the question?

When ChatGPT answers, ask yourself...

- Does this answer the question I asked?
- Has my wording influenced the response?
- Does this add detail I didn't ask for?
- Does this information come from a named source?
- What does this assume that I didn't specify?

To test the AI's answer, ask it...

- How confident should I be in this answer?
- What is the weakest part of this response?
- What assumptions does this response rely on?
- What evidence supports this answer, and what contradicts it?
- What new information would change your conclusion?

If you're generating content, ask yourself...

- Do I want the AI to write the text or just create a framework?
- Who is this for, and what do I want them to get from it?
- Do I want a particular communication style?
- What format should the output take?
- What length is appropriate?

If you're learning something new, ask yourself...

- What is my overall goal here?
- Am I interested more in practical skills or understanding?
- What role do I want the AI to play?
- How structured do I want my learning to be?
- What would be a successful outcome for this session?

If you're developing something creative, ask yourself...

- What do I want the AI to bring to my process?
- Do I want to generate many ideas or develop just one?
- How much creative input do I want to allow the AI?
- How concerned am I about originality?
- What is the creative cost of using AI here?

If the response doesn't seem useful, ask the AI to...

- Tell me why this is helpful?
- Give me more specific detail.
- List the three most important points of this response.
- Turn this into action points.
- Tell me what I need to do next.

If the response is too vague, ask the AI to...

- Explain how this answers my question.
- Give specific examples.
- Make this more precise.
- Define the key terms used in this response.
- Explain the different options.

If the response seems to miss the point, ask the AI to...

- Restate the original question, and answer it directly.
- Reconnect this to my original question.
- Focus only on...
- Answer again from the point of view of...
- Remove any details that I didn't directly ask for

If the response seems implausible, ask the AI to...

- List the logical steps involved in arriving at this answer.
- List the assumptions underpinning this response.
- List any inputs that weren't explicitly specified by me.
- See if a different answer is supportable.
- Try again with a more sceptical attitude.

Power verbs

Strong verbs help you to unambiguously communicate your intent to ChatGPT. Here are some powerful choices, each with its own nuance:

Analysis

analyse, assess, compare, evaluate, examine

Content generation

brainstorm, compose, explore, generate, propose

Critique

challenge, critique, interrogate, test, validate

Explanation

contextualise, describe, explain, illustrate, interpret

Refinement

clarify, optimise, refine, rewrite, simplify

Structuring

classify, group, organise, segment, structure

Summary

condense, distil, extract, outline, summarise

If the response seems poorly supported, ask the AI to...

- List the evidence that supports your conclusion.
- Cite the specific sources of all claims.
- Indicate which parts of this are verifiable and which aren't.
- Search the web to see whether claims can be supported.
- List confidence levels for all claims. 🤖

Formatting and shortcuts

Get results in whatever form works best

Specify the output format

- Give me a one-sentence summary.
- Answer in the form of bullet points.
- Answer in a table.
- Generate a checklist.
- Make a numbered list.

Control the structure

- List the three most important points only.
- Group these items into categories.
- Show me pros and cons.
- Walk me through the process.
- Plan out a schedule.

Steer the conversation

- Let's focus only on...
- Start again from the original question.
- Take a different approach.
- Give me three options for proceeding from here.
- Forget about...

ChatGPT can convert between formats with ease – just ask and it'll instantly turn your data table into an action plan or a presentation script

How could this answer be improved?

- **Too generic – lacks numbers and thresholds**
It doesn't quantify anything. A CFO will expect specifics: capex/opex ranges, breakeven point, payback period, minimum utilisation, and revenue required to justify the move. Without these, it reads as commentary rather than a decision tool.
- **No clear decision criteria**
It lists pros and cons but doesn't tell the reader *what would tip the decision*. A stronger answer would define explicit approval conditions (e.g. "only proceed if X% of revenue is contractually secured" or "if utilisation is forecast above Y% within Z months").
- **Weak on alternatives and counterfactuals**
It doesn't explicitly compare the second office against lower-risk options (remote hiring, serviced offices, travel model). A CFO will always ask: "What's the cheapest way to achieve the same outcome?"

Check the AI's working – or your own – by asking for improvements.

Set the framing

- For context, I want to...
- For this question, my priority is...
- When answering this question, you must...
- Answer as if you were...
- Work only from the supplied information.

Refine results

- Add more detail about...
- Boil this down to action points
- Rewrite this so it's suitable for...
- How could this answer be improved?
- What have we not considered?

Use the attached documents as your sole source of information

Here is a strictly document-based synthesis, using only the attached material:

Where the documents agree

1. Operational bottlenecks exist alongside technical progress

- Technical report notes delays (e.g. calibration bottleneck, environmental interference) but overall progress Report 1.docx
 - Meeting transcript confirms similar issues (calibration delays, vibration sensitivity, scheduling adjustments) Meeting transcript 26-02-26.docx
 - Pitch identifies operational burden (manual monitoring, paperwork) as a constraint Pitch transcript.docx
- All sources recognise that **non-core issues are affecting efficiency**, even if progress continues.

Upload documents to get insights based solely on your own data.

Core output formats

Text

- Paragraph prose
- Short answer
- Headline and explanation
- Bullet points

Structured

- Table
- Schedule
- Step-by-step workflow
- Ranked list
- Document framework

Instructions

- How-to guide
- Tutorial
- Recipe
- Troubleshooting guide
- Decision tree

Interactive

- Q&A
- Interview
- Role-play
- Socratic questioning
- Back-and-forth coaching
- Guided interaction

Creative

- Story
- Script
- Essay
- Poem
- Blog post

Professional

- Email
- Memo
- Report
- Proposal
- Agenda

Work with documents and data

- Summarise the key points from all these documents.
- What do these documents agree about, and where do they disagree?
- Combine these data sets and highlight any trends or outliers.
- Use the attached documents as your sole source of information.
- Follow the instructions in the attached document.

Standardise and share results

- Write up this response as a self-contained explanation.
- Turn this process into a single prompt I can reuse in the future.
- Use this response as a template for future answers.
- Document exactly what we just did so someone else can review it.
- Create a walkthrough so that other people can repeat this process.

Get quick results

- Give me a one-sentence answer.
- Explain the answer without adding extra information.
- Compare the top three options.
- Is there any problem with this?
- Tell me my next step. 🔄

Glossary

Your guide to key terms and ideas used in prompting ChatGPT

Acceptable variation

An agreed standard for how far the AI is allowed to deviate from supplied information and instructions.

Agent mode

A ChatGPT feature that gives the AI greater autonomy to decide for itself how to handle complex tasks.

Agentic

A task or operating mode where the AI performs actions and makes decisions with limited human involvement.

Ambiguity

In the context of ChatGPT, any part of a task that's undefined and requires the AI to interpret for itself.

Blended tasks

Prompts that ask the AI to do two different things at once, potentially leading to poor results.

Breaking the prompt

Asking the AI to find assumptions and weaknesses in your prompt that could affect the response.

Confidence

ChatGPT usually responds confidently for more reliable answers; ask it directly for a confidence rating.

Constraint

Any limitation you place on ChatGPT's working or output, such as "only refer to official data" or "answer in bullet points".

Context window

The limit of the AI's ability to take into account what was said earlier in the conversation – although past inputs tend to lose their force long before they're actually forgotten.

Controlled incompleteness

Making ChatGPT stop short of completing a task to prevent it from latching onto one particular answer.

Drift

The tendency for the AI to lose its focus on your specific task or question over the course of several exchanges.

Expand

A powerful ability of ChatGPT to build additional detail into an existing response or passage of text.

Evaluate

An instruction you can use to get the AI to look for potential weaknesses in a given prompt or proposition.

Falsify

A technique for checking the reliability of a response by asking the engine to try to disprove it.

Framing

Information you give the AI – potentially, without realising it – about how it should approach your request.

Ground truth

Facts and figures supplied by you, for the AI to refer to, rather than relying on its own training data.

Hallucination

An AI claim created by mixing known facts together to come up with something that isn't true, but looks like it could be.

Iteration

Building a complex document or process by starting with a simple instruction and progressively adding in detail.

Large Language Model (LLM)

An AI system such as ChatGPT that works by predicting what words will work together in a certain context.

Leading prompt

A prompt that encourages the AI to respond from a certain perspective, or to prefer a particular answer.

Multi-step prompting

Approaching a task through a series of steps, perhaps to iterate towards an outcome or for visibility.

Output

Anything generated by the AI in response to a prompt – be it an answer, a passage of generated text, or a downloadable document.

Pattern matching

A process the AI often uses to generate answers that have the right "shape", rather than applying step-by-step logic.

Prompt

Anything you enter into the text-entry field – be that a question, an instruction, or just a friendly greeting.

Prompt engineering

Intentionally constructing prompts that encourage specific results and behaviours from the AI.

Prompt rot

The tendency for prompts that worked well in the past to lose their effectiveness as the circumstances and context change.

Prompt template

A written specification of a task or workflow that can be uploaded or pasted into ChatGPT for consistent reuse.

Scope

A boundary constraint telling the AI what information it should and shouldn't include in its output.

Stress testing

Testing a prompt or response for reliability by checking to see how different parameters change the output.

Token

A small chunk of text, used by the AI to parse questions and generate answers, often corresponding to a word or part of one.

Training data

Everything that ChatGPT natively "knows" about the world, and about answering questions.

Transparency document

A document you can produce to record how the AI arrived at a particular conclusion or decision.

Verifiability

A particular challenge with AI, as answers grounded in training data don't come with verifiable sources 

ANTHEM TECH GUIDES

STAY UP TO DATE!



Make sure you're the first one to hear about each new issue of **Anthem Tech Guides** as it's published.

We'll help you take the mystery out of AI!


Go to tinyurl.com/AnthemTechGuides

Or scan the QR code to sign up




Terms & conditions: By signing up you are agreeing to receive emails from Anthem Publishing. You can unsubscribe from these messages at any time.

ADVANCED ChatGPT MASTERING PROMPTS



**Expert prompt techniques
for high-quality answers**



**Create programs and solve
complex logic problems**



**Get the best of AI-assisted
learning and creativity**



**Use ChatGPT safely
in professional settings**

**ANTHEM
PUBLISHING**

AN INDEPENDENT PUBLICATION