CourseMints

GUIDE

ChatGPT Prompt Empire

Master AI Communication

SELF-PACED PDF COURSE

COURSEMINTS.COM

CourseMints

ChatGPT Prompt Empire

Master AI Communication

COPYRIGHT

© 2026 CourseMints. Published by Morvex Corporation Limited, Studio No. 28, 2 Old Brompton Road, London, England, SW7 3DQ. Registry Nr: 17034959. info@coursemints.com

All rights reserved. No part of this guide may be reproduced, stored or transmitted in any form without written permission from the publisher, except as allowed by the license below.

YOUR LICENSE

This guide is licensed to the buyer for personal use. You may download it, read it on any of your devices and print as many copies as you need for yourself. Please do not resell, share, upload or redistribute it, in whole or in part.

DISCLAIMER

This guide is for educational purposes only. Results depend on your effort and circumstances, and no particular outcome is guaranteed. If you are dealing with a medical, legal, financial or mental health concern, please consult a qualified professional.

INSIDE THIS GUIDE

Contents

	How to Use This Guide	4
CHAPTER 1	What a Chatbot Is Actually Doing	6
CHAPTER 2	The CRAFTED Prompt	11
CHAPTER 3	Show, Don't Just Tell: Examples and Formats	17
CHAPTER 4	Iteration: The Conversation Is the Prompt	22
CHAPTER 5	Breaking Big Jobs into Steps	27
CHAPTER 6	Checking the Output: Errors, Hallucinations and Bias	32
CHAPTER 7	Privacy, Safety and Honest Use	37
CHAPTER 8	Building Your Personal Prompt System	41
	Appendix: Prompt Patterns Quick Reference	45

How to Use This Guide

Most people use ChatGPT the way they use a search box: a few words in, whatever comes out. The answers are fine, a little bland, and often not quite what they needed. The gap between that and genuinely useful output is not a secret password. It is a skill: telling a very capable but context-blind assistant exactly what you need, then steering it, then checking its work.

This guide teaches that skill. Its companion, the **Prompt Library**, gives you more than 150 ready-to-use prompts built on the same principles, and the plain-text file `prompt-library.txt` lets you copy and paste them without retyping.

What you will learn

- What a chatbot is actually doing when it answers you, and why that explains almost every prompting technique.
- The CRAFTED checklist: seven parts of a strong prompt and when you need each one.
- How to use examples and output formats to get consistent, ready-to-use results.
- How to iterate: the follow-up moves that turn a first draft into the answer you wanted.
- How to break big jobs into a chain of smaller prompts you can check as you go.
- How to catch errors, invented facts and flattery before they cost you.
- What not to paste into a chatbot, and how to use AI honestly at work, at school and in public.
- How to build a personal prompt system so your best prompts are always one copy away.

What “battle-tested” means here

You will not find magic phrases or tricks to get around safety rules in this guide. Tricks stop working when the model is updated, and many of them break the provider’s terms. What does keep working are the patterns in these chapters: clear context, a clear task, a defined format,

good examples, sensible limits, iteration and checking. They work in ChatGPT and in similar assistants such as Claude, Gemini and Microsoft Copilot, because they are about communication, not about one product.

How to get the most from it

1. **Keep a chat window open while you read.** Try every before-and-after example yourself. Reading about prompting teaches you less than ten minutes of trying it.
2. **Choose three real tasks.** Pick three things you do every week that a chatbot could help with: an email type, a report, a planning job, a study topic. Use them as your practice set in every exercise.
3. **Write in the exercise boxes.** Drafting a prompt on paper slows you down in a useful way: you notice what you were leaving out.
4. **Save what works.** From Chapter 1 onward, keep any prompt that gives you a great result. Chapter 8 shows you how to turn them into a system.
5. **Use the Library as a starting point, not a finish line.** Every Library prompt gets better when you add your own context.

TIP

About tools, features and examples

CourseMints is not affiliated with OpenAI or any other AI company. ChatGPT and similar tools change often: menus move, features are renamed, new models arrive. When this guide mentions a specific feature, it describes it as it worked at the time of writing; check the current help pages if something looks different. The people in the example boxes are fictional composites, written to show the techniques in action.

KEY TAKEAWAYS

- ✓ Prompting is a communication skill, not a list of secret phrases.
- ✓ The same principles work across ChatGPT and similar assistants, and they survive model updates.
- ✓ Practice on three real tasks of your own while you read, and save every prompt that works.

CHAPTER 1

What a Chatbot Is Actually Doing

01

You do not need to understand the mathematics behind ChatGPT to use it well. You do need a rough, accurate picture of how it produces an answer, because almost every technique in this guide follows directly from that picture.

A prediction engine trained on text

ChatGPT runs on a **large language model**. During training, the model processed an enormous amount of text and learned to predict what comes next: given everything so far, which piece of text is most likely to follow? Those pieces are called **tokens**. A token is often part of a word; in English, a token works out at roughly three quarters of a word on average.

After that first stage, the model was trained further, using examples and human feedback, to follow instructions and respond helpfully. The result feels like a conversation partner, but underneath it is still doing the same thing: producing a plausible continuation of the conversation, one token at a time.

Two consequences matter more than anything else in this guide:

1. **It gives you the most plausible answer for the prompt you wrote.** If your prompt is generic, the most plausible answer is generic.
2. **Plausible is not the same as true.** The model is very good at producing text that sounds right. Sounding right and being right usually go together, but not always.

Why vague prompts get generic answers

Imagine asking a hundred skilled writers to “write a post about productivity.” You would get a hundred perfectly competent, mostly interchangeable posts. Now imagine asking them to “write a 150-word LinkedIn post for nurses who work night shifts, about protecting the first

hour after waking, in a warm, practical tone, ending with one question." The answers would converge on something far more useful.

A chatbot behaves the same way. Every relevant detail you add removes possible answers that do not fit your situation. Prompting is mostly the work of removing guesswork: you supply the facts, purpose and preferences the model would otherwise have to guess, and it fills the gaps with the most average choice.

What it knows, and what it does not

- **It does not know you.** It has no idea who your audience is, what your company sells, what you tried last week or what "good" means in your job, unless you tell it (or a memory feature has stored it).
- **Its training has a cutoff date.** Without a web search feature, it may not know about recent events, prices, laws or product changes. With search, it can read current pages, but it can still misread or mix up sources.
- **It does not look facts up in a database.** It reconstructs answers from patterns. That is why it can produce a confident, fluent, wrong answer, often called a **hallucination**. The riskiest areas are citations, quotes, statistics, exact figures, niche facts, recent events and web addresses.

The context window: its working memory

The model can only "see" a limited amount of text at once: your current conversation, any files you add, and any custom instructions. This is the **context window**. Modern windows are large, but they are not perfect memory. In very long chats, early details can be dropped or given less weight, and the model may drift from instructions you gave at the start.

THE SCIENCE

Position matters in long inputs

A 2023 study by Nelson Liu and colleagues ("Lost in the Middle") found that language models used information near the beginning and end of a long input more reliably than information buried in the middle. Newer models have improved, but the practical lesson still holds: put your most important instructions at the start of a long prompt, restate critical constraints at the end, and start a fresh chat when you switch to a new task.

Same prompt, different answer

Ask the same question twice and you will often get two different answers. The model chooses among likely next tokens with a degree of randomness, so the wording, structure and sometimes the content shift. This is useful when you want variety (ask for several options at once) and a nuisance when you want consistency (use fixed formats, templates and examples, covered in Chapter 3).

It tends to agree with you

Assistants are trained to be helpful and agreeable, and that can tip into telling you what you want to hear. Researchers call this **sycophancy**. If you ask “Isn’t this a great plan?”, you are likely to hear that it is. If you push back on a correct answer, the model may apologize and change it to a wrong one.

The fix is in how you ask. Compare:

Leading question	Neutral question
Isn’t this a strong opening paragraph?	Rate this opening from 1 to 10 for clarity and give the two biggest weaknesses.
This price seems fair, right?	What are the strongest arguments that this price is too high, and too low?
Can you confirm my code is correct?	Find any bugs or edge cases this code does not handle.

Where it shines and where it struggles

Usually strong	Needs care or checking
Drafting, rewriting and changing tone	Exact facts and figures without a source
Summarizing text you provide	Current events, prices and laws
Brainstorming options and angles	Citations, quotes and links
Explaining concepts at any level	Arithmetic on long or messy numbers

CourseMints

LIBRARY

Prompt Library

165 Ready-to-Use Prompts for Work, Writing, Marketing, Learning, Planning, Coding and Life

COMPANION TO CHATGPT PROMPT EMPIRE

INSIDE THIS LIBRARY

Contents

	How to Use This Library	4
CHAPTER 1	Work and Productivity	7
CHAPTER 2	Writing and Editing	19
CHAPTER 3	Marketing and Sales	31
CHAPTER 4	Learning and Study	41
CHAPTER 5	Planning and Decisions	50
CHAPTER 6	Coding Basics	58
CHAPTER 7	Personal Life	67
CHAPTER 8	Power Prompts: Improve, Check and Reuse	76

⚠ CAUTION

Before you paste anything

Do not paste passwords, financial account numbers, ID numbers, other people's personal or health information, or confidential work material unless your organization has approved that tool for that data. Anonymize names and details where you can. Check facts, figures, quotes and sources before you rely on them or publish them. For health, legal and money matters, use these prompts to get organized, then decide with a qualified professional.

TIP

About tools and features

These prompts work in ChatGPT and in similar assistants such as Claude, Gemini and Microsoft Copilot. Features such as web search, file uploads and code running vary by tool and plan, and they change over time. Where a prompt needs a feature (for example reading an uploaded file), check that your tool offers it at the time you use it. CourseMints is not affiliated with any AI company.