

AI Prompt Writing Mastery — Workbook

This workbook turns the course into reps. Each section mirrors one course module with hands-on exercises, fill-in worksheets, and checklists you run against real prompts. Work through it with a live AI tool open, and finish with a personal prompt library you will actually reuse.

How AI Models Read Your Prompt

Build accurate intuition for next-token prediction, temperature, and the context window before you write a single advanced prompt.

Exercise: Temperature in Action

Open one AI tool. Run each prompt below three times in fresh chats and note how much the answers vary. Then add a constraint that forces a tighter answer and observe the difference.

- Give me five names for a neighbourhood bakery. (Run three times, note how different each set is.)

-
- What is the boiling point of water at sea level? Answer in one short sentence only. (Run three times, note how stable it is.)

-
- Rewrite the bakery-names prompt so it returns exactly one name and nothing else, then run it three times and compare consistency.

Worksheet: Diagnose a Weak Result

Take a real prompt that gave you a disappointing answer recently. Fill in each field to locate the missing ingredient using the five failure modes from the course.

The original prompt I used

What was wrong with the answer (too generic / wrong shape / too long / wrong tone / invented facts)

Which ingredient was missing (role / audience / format / constraints / examples)

The single addition that would most likely fix it

My rewritten prompt with that ingredient added

Checklist: Tool-Selection Quick Check

I matched the task to a tool: long-document work to Claude, web-current facts to Gemini, mixed or plugin tasks to ChatGPT.

I am on the strongest model my plan allows for any research or analysis task.

I confirmed whether browsing or live search is on when the task needs current information.

For an important task, I have a second tool ready to cross-check the first.

The Core Prompt Frameworks

Drill RTCF, few-shot, and chain-of-thought until structuring a strong prompt becomes automatic.

Worksheet: Build an RTCF Prompt

Pick one real task you want AI to help with this week. Fill in all four parts, then paste the assembled prompt into a tool and refine it.

Role (who the model should act as)

Task (one clear action verb)

Context (audience, goal, background, source material)

Format (exact length, structure, and shape of the answer)

Assembled full prompt (the four parts written as one message)

Exercise: Teach by Example (Few-Shot)

Choose a repeating task with a consistent pattern, such as product blurbs or reply emails. Hand-write two perfect examples, then prompt the model to continue the pattern on a new input.

- Write two finished examples of your output exactly how you want them, labelled Example 1 and Example 2.

• Add: Now produce the same for: [new input], matching the structure, length, and tone of the examples above.

- Run it, then improve one example and re-run to see how the output shifts. Note what changed.
-

Exercise: Force the Reasoning (Chain-of-Thought)

Take a small decision or multi-step problem you face. Compare a plain prompt against a step-by-step prompt and judge which answer you trust more.

- Plain: Which of these two options is better: [option A] or [option B]?

• Chain-of-thought: Compare [option A] and [option B] for [situation]. First list the factors that matter, then weigh each, then recommend one and name the deciding factor.

- Note which answer was more useful and why, and whether the shown reasoning revealed anything you had missed.
-

Checklist: Framework Coverage Check

- My prompt names a role and the intended audience.
- It states one clear task, not five stacked requests.
- It specifies length and output format precisely, not vaguely.
- For style or pattern work, I included two or three examples.
- For multi-step reasoning, I asked the model to think step by step.

Prompting for Real Work: Writing, Research, Analysis

Apply the frameworks to your three highest-frequency task types with ready-made patterns.

Exercise: Write in Your Own Voice

Capture a sample of your writing and use it to tune a draft so it sounds like you, not like default AI.

- Paste 150 to 300 words you wrote, then say: Match this tone, rhythm, and vocabulary in everything below.

- Draft a 250-word piece on [topic] for [audience]. Short sentences, one concrete example, no clichés, and never use the phrase in conclusion.

- Editing pass: Tighten my draft without changing my voice or conclusions. Cut filler, fix flow, and flag any claim that sounds unsupported.

Worksheet: Grounded Research Brief

Use this to digest a real document safely. Fill it in, paste the source text, and run the grounded prompts so the model stays anchored to your material.

Source material I am pasting (title and what it is)

My research question

Summary length I want (word count)

Claims flagged by the model as worth verifying

Verification result for each flagged claim (confirmed / wrong / could not confirm)

Exercise: Structured Decision Analysis

Turn a real choice into an auditable comparison. List your options and weighted priorities, then have the model score, recommend, and challenge itself.

- You are a decision analyst. Score these options from 1 to 5 on my priorities [list priorities and any weighting], then show the table.

- Compute a weighted total, recommend one option, and name the single biggest risk of that choice.

- Now argue the opposite of your recommendation and identify the weakest assumption in my framing.

Checklist: Real-Work Output Check

Writing reads in my voice, not the default AI tone, and I removed banned phrases.

Research summaries are grounded in source text I provided, with claims traced to sentences.

Every open-question fact has been verified against a primary source.

Analysis includes explicit reasoning and a self-challenge, not just a verdict.

Refine, Verify, and Build Your System

Convert one-off wins into a repeatable practice through iteration, verification, and a saved prompt library.

Worksheet: Run the Iteration Loop

Take an okay answer and improve it deliberately, changing one variable per round so you learn what works. Starting prompt and what was wrong with the first answer

Round 1: the one variable I changed and the result

Round 2: the one variable I changed and the result

Round 3: the one variable I changed and the result

Final winning prompt to save

Checklist: Pre-Trust Verification Checklist

- Every specific number, date, name, and quote is confirmed against a primary source.
- I clicked through every citation the model gave and confirmed it actually exists.
- I re-read any reasoning chain and the steps genuinely connect.
- I asked whose perspective is missing and checked for one-sided framing.
- I confirmed the information is current given the model's knowledge cutoff.
- For medical, legal, or financial content, a human expert or primary source has the final say.

Exercise: Seed Your Prompt Library

Move five proven prompts out of chat history and into the library template. Make each reusable with bracketed placeholders.

- Pick five prompts that worked well for you in this course and copy each into the library template.
 - Replace specifics with bracketed placeholders like [INSERT TOPIC] so each prompt is reusable.
 - Add a category, the best tool, and a one-line note on what made each output good.
-

Worksheet: Library Maintenance Plan

Set up the light habits that keep a prompt library sharp over time.
Where my library lives (spreadsheet / Notion / in-tool saved prompts)

How I version a prompt when I improve it (dated note in a column)

Monthly review date to prune and tighten prompts

Two prompts I most want to build next

Your Action Plan

1. Pick one primary AI tool and, if research or analysis matters to you, upgrade to its paid tier this week.
2. Rebuild three weak prompts from your recent history using the full RTCF structure.
3. Create one few-shot prompt with two hand-written examples for a task you repeat often.
4. Add think step by step to your next multi-factor decision and compare it against a plain ask.
5. Capture a 200-word writing sample and save a voice-matching prompt you can reuse.
6. Run the four-step iteration loop on one okay answer, changing a single variable per round.
7. Apply the verification checklist to one important AI output before you act on it.
8. Set up your prompt-library template and seed it with five proven prompts.
9. Schedule a recurring 10-minute monthly review to prune and tighten your library.
10. Teach the RTCF framework to one colleague to lock in your own understanding.

