

Prompt Engineering for Business — Workbook

This workbook turns the course into action. Each section pairs with a course module and gives you exercises, fill-in worksheets, and checklists to build real prompts, a reusable library, and a quality-control habit. Work through it with ChatGPT, Claude, or Gemini open in another window, and keep your filled-in sheets as the first entries in your own prompt system.

How Prompting Actually Works

Build the mental model and choose the right tool before you write a single business prompt.

Exercise: Vague Versus Specific Showdown

Open your chosen AI tool. Run the vague prompt, then the specific one, and compare the two outputs side by side. Note in writing how much editing each would need before you could use it.

- Vague: Write a follow-up email to a client who has not replied.
- Specific: You are an account manager at a B2B firm. Write a warm follow-up email under 120 words to a client named Dana who requested a proposal two weeks ago and went quiet. Reference the proposal, offer a 15-minute call this week, end with one clear question, friendly tone.
- In two sentences, explain which output was better and exactly why.

Worksheet: Pick Your Primary Model

Fill in each field to decide which AI tool you will learn first and why. Base it on the tasks you do most and the apps you already use.

My three most common AI tasks

Typical input size (short text / long documents)

Tools my work already lives in (Google, Microsoft, neither)

Free or paid plan I will use

My chosen primary model

One reason this model fits my work

Checklist: Strong Prompt Self-Audit

- I named a role or perspective for the model
- I stated the task in one clear, verb-first sentence
- I included the names, numbers, or source text it needs
- I specified the output length and format

- I added tone or any do-not constraints
- I read my prompt back to confirm it says what I meant

Core Frameworks for Business Prompts

Practice RTCF, role prompting, and few-shot examples until they become automatic.

Worksheet: Build an RTCF Prompt

Choose a real task from your work this week. Fill in each part, then combine the four into a single prompt and run it. Save the result.

Role (who the model should be)

Task (the single action, verb-first)

Context (names, numbers, audience, source text)

Format (length, structure, table or list)

Assembled prompt (all four combined)

Quality of the result (1-5) and what to fix

Exercise: Few-Shot Style Match

Pick a recurring piece of writing you produce (product blurb, outreach message, post). Gather two examples you consider excellent, then run a few-shot prompt to generate a new one in the same style.

- Here are two examples of how we write [TYPE OF CONTENT]. Example 1: [paste]. Example 2: [paste].
- Now write a new [TYPE OF CONTENT] in exactly the same style for the following: [new details].
- Compare the new output to your examples: did it match tone, length, and format? Note one adjustment.

Checklist: Persona Quality Check

- My role names a profession and seniority level
- It includes a specialty or industry
- It states who the persona is speaking to
- It sets a standard (reading level, no jargon, cite sources) where useful
- I considered a review persona (skeptical CFO, hostile customer) for critique
- I saved any persona that worked well for reuse

Advanced Techniques and Reasoning

Apply chain-of-thought, decomposition, and grounding to higher-stakes business tasks.

Exercise: Reasoning on a Real Decision

Take an actual decision you are weighing (a price, a vendor choice, a budget line). Run it first without reasoning, then with chain-of-thought and step-back prompting. Compare the depth of the two answers.

- Plain: Recommend a price for [product/service] given these details: [details].
- Step-back: Before recommending a price, first list the factors that determine pricing in this category, then apply them to [details], thinking step by step.
- After the reasoning, ask: Now give a clean final recommendation I could share, with the key justification

only.

Worksheet: Decompose a Big Task

Choose one project too large for a single prompt. Break it into an ordered chain of smaller prompts, listing what each step produces and feeds into.

End goal (the final deliverable)

Step 1 prompt and its output

Step 2 prompt (uses Step 1 output)

Step 3 prompt (uses Step 2 output)

Step 4 prompt (final assembly)

Where I will review and edit between steps

Checklist: Anti-Hallucination Checklist

- I pasted the real source text the model should rely on
- I instructed it to use only that source and say I do not know otherwise
- I required figures and quotes to trace to a specific sentence
- I treated every unverified number, date, and citation as a draft
- I ran a follow-up asking the model to flag any unsupported claim
- A human will review anything legal, financial, medical, or client-facing

Building Libraries and Quality Systems

Turn proven prompts into a versioned library and a repeatable scoring and rollout workflow.

Worksheet: Start Your Prompt Library

Capture your five most-used prompts as proper library entries. Use bracketed variables for anything that changes each time.

Prompt name

Task group (sales, summaries, social, etc.)

The prompt text with [VARIABLES]

One-line note: what it does and when to use it

Version date and last improvement

Average quality score (1-5)

Exercise: Score and Iterate

Take one AI output from this week and score it on the five criteria. Then run targeted follow-ups to lift the lowest score, without rewriting from scratch.

- Score this output 1-5 on accuracy, completeness, tone, format, and usability, then identify the weakest one.

- Refine with a specific instruction, e.g. Make it 30 percent shorter, or That second point is wrong because [reason], fix it and keep the rest.

- Re-score after the fix. If it now succeeds, save this version to your library.

Worksheet: Draft a One-Page AI Policy

Write a short team guideline so everyone uses AI consistently and safely. Keep each answer to one or two lines.

Approved tools

Data that must never be pasted in

How to anonymize sensitive inputs

Tasks that require human review before sending

Where the shared prompt library lives

Who owns accountability for AI-assisted work

Checklist: Team Rollout Checklist

- The one-page AI policy is written and shared
- The shared prompt library is set up and accessible
- I demoed three high-value prompts to the team
- Data privacy rules and opt-out settings are confirmed
- Everyone knows which tasks need human review
- We have a place to suggest and improve prompts over time

Your Action Plan

1. Choose one primary model (ChatGPT, Claude, or Gemini) and commit to learning it deeply.
2. Memorize RTCF and run it on every prompt for one full week.
3. Create three reusable expert personas for your most common tasks.
4. Build a few-shot example set for one recurring piece of writing you produce.
5. Start a prompt library and capture your five most-used prompts with variables.
6. Apply chain-of-thought and grounding to one high-stakes task this week.
7. Adopt the five-criteria scoring rubric and use it before sending any AI output.
8. Set up an iteration habit: refine with follow-ups instead of rewriting from zero.
9. Write a one-page AI policy covering tools, data privacy, and human review.
10. Roll the library and policy out to your team with a short demo of your best prompts.

