

INDEPENDENT HANDBOOK

Mastering Gemini AI

From Beginner to Power User

Leah Okonkwo

A practical book about briefing, checking, and reusing Gemini for research, writing, business, and small technical work.

Mastering Gemini AI

From Beginner to Power User

© 2026 Leah Okonkwo. This edition was prepared for the Mastering Gemini AI academy.

This book is independent. It is not affiliated with, endorsed by, or produced by Google. Gemini, NotebookLM, Gmail, Docs, and Workspace are Google products. Their screens, model names, and plan limits change. When a chapter describes a button, treat it as a late-2026 description and trust the product you have open.

No part of this book promises income, rank, or a professional license. The certificate described in the academy is a course record, not an accredited credential and not a Google certification.

You are responsible for what you publish, send, or decide. Do not paste passwords, keys, health records, or client secrets into a chat. Check any number, name, date, quote, or rule that would matter if it were wrong.

Contents

How to use this book

What stays true when the screen changes

Foundations

Daily work

Judgment

Building

- Level 1. Gemini AI Fundamentals

- Level 2. The Art of Prompting

- Level 3. Gemini for Productivity

- Level 4. Research, Learning, and Knowledge

- Level 5. Gemini for Content Creation

- Level 6. Business and Marketing

- Level 7. Images, Documents, and Multimodal Work

- Level 8. Coding and Technical Work

- Level 9. Advanced Gemini Workflows

- Level 10. Gemini Masterclass: Build with AI

A short chapter on not trusting a fluent voice

Where Gemini fits among other tools

The next thirty days

BEFORE THE FIRST PROMPT

How to use this book

This is a handbook, not a tour of a product menu. The screens will move. The skill will not. You are learning to specify a job, to read the answer as a draft, and to keep a version of the brief you can run again.

Read with the Gemini app open. The free app is enough. Paid plans and Workspace editions add side panels, longer research runs, and higher limits. Those are useful when you have them. They are not required to finish the work in these pages.

Do the practice at the end of a section before you admire the next one. A section you only read will feel finished and then vanish on Monday. A section you ran on a real task from your week will still be there.

Keep a note outside the chat. The academy stores progress in the browser. This book cannot. A context card, a prompt that worked, and a prompt you retired are the whole filing system you need.

Four rules sit under every chapter.

- Tell the model the facts it cannot see. It does not know your calendar, your margins, or your files unless you provide them.
- Ask for a shape you can scan: headings you named, a table with fixed columns, three options and then a recommendation.
- Treat fluent writing as untrusted wherever a number, a name, a quote, or a rule would matter.
- You sign the work. The model does not.

The path in this book runs from a first careful session, through prompting, productivity, research, and content, into business judgment, images, light coding, and a reusable system. The last pages ask you to pick one real project and revise it once after a hard read.

If a term is new, it is defined in the sentence where it appears. You do not need a computer-science background. You do need a willingness to throw away a pretty paragraph that is not true.

A DISTINCTION THAT SAVES YOU

What stays true when the screen changes

Gemini is Google's family of multimodal models. Multimodal means the input can be more than typed text: often an image or a document, sometimes audio or video, depending on the surface and the plan. The Gemini app at gemini.google.com is the main chat. Mobile apps exist. As of late 2026 there was also a desktop direction. None of those sentences is a map of next year's buttons.

A model produces a likely continuation of your prompt. It does not look up a private library of truth unless a product feature grounds it, and even then you open the source. Hallucination means a fluent answer that is false or unsupported. It is ordinary, not a rare glitch you outgrow.

Context is the working set the model can see right now: your instructions, the conversation, and the files you attached. It is not a permanent memory of you. Some products offer a memory feature. Check the current screen before you rely on it. A long, muddy chat can crowd out the original job. A fresh chat is a clean context. Keep the brief in your note so you can start clean without starting ignorant.

As of a May 2026 update, Google said more than 900 million people a month used the Gemini app. Treat the figure as dated. Model names in public posts that year included lines such as Gemini 3.5 Flash and live voice models. Read the model picker. Do not memorize a brand.

Workspace Gemini can appear in Gmail, Docs, Sheets, Slides, Drive, and Chat when the edition includes it. NotebookLM is a separate tool for studying sources you upload. Deep Research is a longer report mode whose quota depends on the plan. Gems were saved assistants. Skills, reusable instruction files, began rolling out in late 2026. If your screen uses different words, the habit is the same: save a brief you will run more than once.

Stable skills are briefing, iterating, and verifying. Changing features are the names of buttons. This book teaches the first and labels the second.

PART ONE

Foundations

You learn what the model is actually doing, and you learn to brief it. Everything later in this book is a use of those two skills.

LEVEL 1 · BEGINNER

Gemini AI Fundamentals

You will learn what Gemini is, how a chat really works, when a draft is the wrong tool, and how to finish a calm first session with a brief you can check.

Welcome to Mastering Gemini AI

This opening lesson maps the academy, shows how to practice with Gemini beside you, and names the habit that turns a vague question into a brief you can reuse.

People often judge Gemini from one loose question and a rushed answer. This lesson sets the path so the later drills make sense, and it tells you where progress is saved. You will also see the difference between hoping for a finished answer and writing a brief you can check.

This academy takes you from a first chat to a system you can reuse at work or in study. You will build a planner, a research workflow, a content pipeline, a marketing strategy, a simple web page, a power-user system, and a capstone that ties them together. Each level adds one skill and then asks you to use it on something real. The free Gemini app is enough for this course. Paid plans change quotas, not these habits.

Practice with the lesson on one side and Gemini open on the other. Read a short idea, try the prompt, then come back and compare what you got. Gemini is Google's multimodal model family, and the app at gemini.google.com is the main chat surface. Mobile apps exist as well. Multimodal — the model can take more than text, often images and documents, sometimes audio or video depending on the surface.

Progress in this academy is stored only in this browser. There is no course account that follows you to another laptop. If you clear site data, or you switch browsers, the checkmarks can vanish even though the skill remains. That choice lets the course run without building a profile of you. Keep a note of prompts you like, because a chat is not a filing cabinet and another device will not show this progress.

IN THE FIELD

Nia opens the academy after her shift

Nia manages the front desk at a neighborhood clinic in Accra. On Tuesday she finishes the late shift, makes tea, and puts this lesson on her laptop with the Gemini app open beside it. She does not paste patient names or appointment notes. She types a short interview prompt and answers only with her public duties: greeting visitors, writing the weekly staff reminder, and ordering supplies. Gemini asks what a good reminder looks like. Nia says it must fit on one screen, use plain words, and end with the two tasks that cannot slip. She copies that brief into a note titled Clinic admin. The next morning she reuses it for a supply list and spots a missing item before she sends anything. She still checks the list herself.

A PROMPT WORTH KEEPING

Interview me for 5 minutes about my work so later prompts have useful context. Ask one question at a time. Do not ask for secrets, passwords, client names, health details, or unpublished figures. After five questions, write a context card of at most 120 words using only what I chose to share, and list two ways I can reuse that card in later chats. If I say skip, move on.

- A five-minute limit keeps the chat small enough to reuse.
- One question at a time reduces the chance you paste private detail by habit.

USUALLY. Tell me everything I need to know about using AI at work. **BETTER.** Interview me with five questions about my weekly writing tasks, then draft a reusable context card and stop.

PRACTICE

Write your first context card

1. Open the Gemini app beside this lesson and start a new chat.
2. Paste the interview prompt from the breakdown and answer only with public facts about your work or studies.
3. Stop after the context card and copy it into a note you control.
4. Read the card and delete any line you would not want a colleague to see.

You have a short context card in your own note, with secrets removed, that you could paste into a later chat.

- You will build a planner, research workflow, content pipeline, marketing strategy, simple web page, power-user system, and capstone.
- Practice with Gemini open beside the lesson, and store prompts you like in a note you control.
- Academy progress lives only in this browser, so another device will not show your checkmarks.

What Gemini Is, and What It Is Not

You will treat Gemini as a generative model that drafts likely answers, not as a private library of truth, and you will open a source whenever a quote or number matters.

Search and Gemini both help, but they do different jobs. If you treat a drafted answer like pages you already opened, you will trust lines you never checked. As of a May 2026 Google update, the company said 900 million people a month used the Gemini app, and you should treat that number as dated rather than as proof a draft is right.

Gemini is a generative model — it produces likely text, or other media, from patterns learned in training. It does not look up a private database of truth by default. A calm, fluent paragraph can still be wrong. Ask for a draft, a rewrite, an explanation, or a set of options, and keep the last word on what is true enough to use.

Traditional search retrieves pages and leaves you to open them. Gemini drafts an answer and may or may not ground that answer in sources you can click. Those are different jobs. When a source link or a quote matters, ask for it and then open it. If you cannot open the page, do not treat the line as verified.

Hallucination — a fluent answer that is wrong. The risk is not only a wild story. A tidy paragraph can miss a date, invent a citation, or blend two real rules into one rule that nobody passed. Fluency is not proof. Slow down when a name, a number, a law, or a quote would cost you if it were off.

IN THE FIELD

Omar almost prints a quotation

Omar teaches history at a secondary school in Amman. He asks Gemini for a vivid quotation from a 1920s public speech so his Friday slide looks serious. The draft arrives with a confident sentence, a year, and a name. Omar almost pastes it. Then he searches the name plus the year on a library site he trusts and cannot find the line. He asks again, this time for a plain explanation of the event and a list of claims to verify, with no quotation marks. He finds a short passage in a book his school already owns, types that passage himself, and asks Gemini only to suggest a discussion question. His students still get a strong slide. The quotation on it is one he opened.

A PROMPT WORTH KEEPING

Explain this topic in plain words for a busy colleague: how a workplace visitor should prepare for a site tour. Separate what you are confident is general advice from what must be checked in our official safety code. Do not invent quotations, section numbers, or citations. End with three questions I should ask our safety lead, and mark any claim I must open in a source before I put it on a slide.

- The audience and topic are specific, so the draft has a job.
- It forces a split between general advice and claims that need a source.

USUALLY. Quote the exact law and I will print it on the poster tonight. **BETTER.** List the claims I must verify in the official text, and do not invent a quotation or a section number.

PRACTICE

Catch one fluent claim

1. Choose a public fact you can check, such as a museum opening day or a definition in a book you own.
2. Ask Gemini to explain it and to mark claims you should verify, with no invented quotations.
3. Open one real source yourself and compare one sentence.
4. Write down one helpful line and one line you would not use until you checked it.

You have compared one Gemini claim with a source you opened, and you can say whether the draft matched.

- A generative model produces likely text from patterns. It does not look up a private database of truth by default.
- Search retrieves pages. Gemini drafts an answer and may or may not ground it, so open the source when a quote matters.
- Hallucination is a fluent answer that is wrong. A calm tone is not proof.

How Generative AI Works, Without the Math

You will see how a reply is predicted from your prompt and the chat, why a long thread can crowd out the task, and why the model sees an image only when you provide it.

If you imagine a tiny person reading your mind, you will expect memory and eyesight the product does not have. A plain picture of prediction makes the failures less mysterious. You can then fix the chat instead of blaming yourself.

Picture a busy cook who has practiced thousands of meals. When you order, they do not run a laboratory. They reach for patterns that usually fit and plate something plausible. A generative model is similar only in that thin way. Training — earlier learning on huge amounts of text and other data. The cook picture is a start, and it breaks as soon as you expect taste, shopping, or a look inside your cupboard.

Here is the correction. The model predicts a plausible next piece of the reply given your prompt and the conversation so far. It is not opening your files and it is not keeping a diary of you. It does not remember you across chats unless a product memory feature is on, and those features change, so check the current screen. A new chat starts without your story unless you paste the parts that matter.

Instructions you place at the start of a chat shape everything after them. A long messy chat can crowd out the original task because the working set is finite. Context — the working set of instructions and conversation the model can see right now. It is not permanent memory. When side topics fill the thread, start a fresh chat and paste only the brief that still matters.

IN THE FIELD

Priya's cake ratio was never stored

Priya runs a small bakery in Leicester. On Monday she tells a chat her lemon-cake butter ratio and asks for a prep list. On Thursday she opens a new chat and types continue the cake plan. The reply invents a ratio she never uses. She thinks the model forgot her. It did not forget, because a new chat does not keep a private memory of her kitchen unless a memory feature is on, and she should check the current screen before trusting one. She pastes a three-line card: lemon cake, her real ratio, twelve servings, no nuts. The next prep list matches that card. The card did the work. The model's memory did not. She keeps the card in her recipe notebook, not only inside a thread.

A PROMPT WORTH KEEPING

I am starting a fresh chat. Use only this brief. Role: assistant to a bakery owner. Context: lemon cake for twelve, no nuts, plain language, do not invent prices. Task: write a prep list for one person working alone. Constraints: at most twelve steps, no secret recipes beyond what I paste next, mark any assumption. Output: a numbered prep list plus three questions you still need me to answer. My notes: [paste only non-secret notes].

- A fresh chat plus a brief at the top keeps the real task in context.
- Role, context, task, constraints, and output tell the model what to predict toward.

USUALLY. Continue from yesterday in this brand-new chat. You already know my ratios. **BETTER.** Here is the brief again: lemon cake, twelve servings, no nuts. Use only these notes.

PRACTICE

Reset with a brief card

1. Write a five-line brief for a task you know well, with no secrets.
2. Start a fresh chat and paste that brief before any other question.
3. Ask for one output, then ask a side question that could distract the chat.
4. Start another fresh chat, paste the same brief, and compare which answer stays closer to the task.

You can point to one answer that drifted and one answer that stayed on the brief because the context was clean.

- The model predicts a plausible next piece of the reply from your prompt and the conversation it can see.
- Training is earlier learning on huge amounts of data, not a memory of you.
- Context is the working set right now, not permanent memory. A long chat can crowd out the original task.

The Gemini Ecosystem, and What Depends on Your Plan

You will map the main Gemini surfaces, separate stable skills from features that move, and check your own screen instead of trusting a frozen screenshot.

People get stuck searching for a button from a video while the skill they need is already available. A map stops that hunt. You will learn which pieces depend on a plan and which habits work on the free app.

Start with the Gemini app, the chat at gemini.google.com, plus the mobile apps. You may see a model picker. Read the label on your screen instead of memorizing a brand, because names change, and check the current screen. You can often upload a file or an image so the model can read what you provide. If upload is hiding, look for a paperclip or a plus, or describe the file in words.

Deep Research is a longer report mode, and its quota depends on your plan, so a missing button is not your fault. Gems are saved custom assistants for a job you repeat. Skills are reusable instruction files Google began rolling out in late 2026. Check the current screen, because that roll-out can move, and Gems were still usable beside it. Practice briefing even if a menu is new.

On an eligible plan, Google Workspace with Gemini can appear in Gmail, Docs, Sheets, Slides, Drive, and Chat. If the plan does not include it, those side panels will not be there. NotebookLM is a separate Google tool for studying sources you upload, not a normal chat with a new name. Image understanding and image creation show up where the plan allows. The free app is enough for this course. Paid plans change quotas. This course does not set prices.

IN THE FIELD**Luis writes down his real screen**

Luis manages a small office in Lima. He wants help with three policy PDFs his director said he may use, a short email, and a photo of a cracked lobby tile. He opens gemini.google.com and writes down three things he actually sees: a chat box, a control for a new chat, and an upload icon. He does not see a Workspace side panel in Gmail on this personal login, so he does not pretend it is there. For the PDFs he tries NotebookLM, because that tool is built to study sources he uploads. For the tile he uses upload in the Gemini app and asks only for what is visible. He leaves Deep Research until he checks whether his plan includes that longer report mode.

A PROMPT WORTH KEEPING

Help me choose a surface for this real task. I will not paste secrets. Task: [one sentence]. Materials I am allowed to use: [files, or none]. What I can see today: [three things on my Gemini screen]. Tell me whether to use the Gemini app, NotebookLM, or a Workspace panel in Gmail, Docs, Sheets, Slides, Drive, or Chat. Say what depends on my plan. If a feature may have moved as of late 2026, tell me to check the current screen. End with the stable skill I should practice anyway: briefing, iterating, or verifying.

- You report what is on your screen, so the advice is not stuck to an old screenshot.
- Separating app, NotebookLM, and Workspace stops you from hunting in the wrong product.

USUALLY. Tell me the one correct Gemini product for every job, forever. **BETTER.** Given the three things I see today, which surface fits this task, and what should I recheck later?

PRACTICE**Name three things you see**

1. Open gemini.google.com or the mobile app and start from the screen you really have.
2. Write three things you see, in your own words, such as chat box, new chat, or upload.
3. Mark one job as stable (briefing, iterating, or verifying) and one feature that might depend on a plan.
4. If a famous button is missing, write the workaround you will use today.

You have a three-item list from your own screen and one workaround that does not depend on a button you cannot find.

- The Gemini app is the main chat. Uploads, a model picker, Gems, Skills, and Deep Research may appear, and some depend on the plan.
- Workspace with Gemini can show up in Gmail, Docs, Sheets, Slides, Drive, and Chat only if the plan includes it.
- NotebookLM is a separate tool for sources you upload. Image tools appear where the plan allows.

Strengths, Limits, and When Not to Use Gemini

You will choose the jobs Gemini is good at, refuse the jobs it must not decide, and keep privacy, copyright, bias, and your own signature in view.

A helpful draft can still be the wrong tool. This lesson draws the line before a real person is harmed or a secret is pasted. You stay responsible for what you send, publish, or decide.

Reach for Gemini when you want a first draft, several options, a plain explanation, a restructure of notes you already wrote, or practice questions. Those jobs use its talent for likely language. You still edit. Start with material you may share, and end by checking names, numbers, and tone. If a mistake only costs a rewrite, you are in a sensible lane.

Do not use Gemini for an emergency. Do not use it for a medical, legal, or financial decision without a qualified person. Do not use it when you need a live price or a live law, to cheat on an exam, or to invent citations. Do not paste passwords, client secrets, health data, or unpublished financials. If the shape of the problem still helps, invent fictional placeholders, or leave the chat and call a professional.

Consumer settings and Workspace admin settings differ. Your employer may have rules a personal login does not. As of late 2026, Workspace help pages said customer content is not used to train models. Consumer accounts are different, so check the current privacy page before you paste secrets. You are responsible for what you publish. Do not ask the model to clone a living artist's style to imitate them, and do not paste paywalled text you have no right to share.

IN THE FIELD

Hana keeps the real numbers out

Hana is a freelance bookkeeper in Kuala Lumpur. A client sends an unpublished revenue figure and asks her to let Gemini decide whether to delay a tax payment. Hana stops. That is a financial decision for a qualified person, and the figure is the kind of secret this lesson says not to paste. She types a fictional shop with round numbers she invented and asks for questions an accountant should hear. She sends those questions to the client and books a licensed tax agent. Gemini helped her prepare for the conversation. It did not sign the decision. The real numbers never entered the chat. She also deletes the client's raw message from her prompt draft before she types anything.

A PROMPT WORTH KEEPING

Help me prepare for a conversation with a qualified professional. Do not decide the case. Context: a fictional small cafe I invented, with made-up monthly sales of 1000, no real names, no account numbers, and no health data. Task: list questions for a licensed tax agent about a possible late payment. Constraints: do not give a legal or financial ruling, do not invent citations, do not imitate a living artist. Output: eight questions, a one-paragraph note I can send a client in my own words, and a line that says a human professional must decide.

- It states that the model must not decide the case.
- Fictional numbers keep unpublished financials out of the chat.

USUALLY. You are the doctor on call. Tell this patient the dose from the notes I pasted. **BETTER.** Turn these public leaflet headings into questions for a clinician. Do not recommend a dose or a diagnosis.

PRACTICE**Sort six tasks**

1. Write six tasks from your real week, using no secret details.
2. Mark each one as use, do not use, or use only to prepare questions for a qualified person.
3. For one do-not-use task, write the safer substitute, such as a fictional case or a phone call.
4. For one use task, add the check you will do before you send or publish the draft.

You have six labeled tasks and at least one safer substitute written in your own words.

- Use Gemini for drafts, options, explanations, restructuring, and practice questions, then edit.
- Do not use it for emergencies, licensed decisions, live prices or laws, cheating, fake citations, or secret data.
- Workspace and consumer privacy settings differ. As of late 2026, check the current privacy page, and do not paste text you have no right to share.

Your First Session: Navigating with Confidence

You will take a calm first pass through the Gemini app, name the parts that matter, and run a ten-minute route that still works when the layout changes.

The first visit feels busy because every icon looks important. A short route shows you the few controls that matter. You will leave able to start clean, steer a draft, and keep a copy outside the chat.

Sign in, then start a new chat before you practice. Look for a model label if one is shown, and read it without memorizing it. Find where to upload a file or an image, even if you upload nothing today. Find where to start a fresh chat — a clean context — so an old thread does not leak into this drill. Say those parts out loud in your own words.

If the product lets you edit your last prompt, try it once. Editing is often better than stacking a pile of corrections on a confused thread. Copy any answer you want to keep into a note you control, because the chat is not your archive. Then practice two follow-ups you will use all year. Ask for a shorter version. Ask it to use headings.

Give yourself this 10-minute route. Minutes 0 to 2: sign in and notice the model label and the upload control. Minutes 2 to 5: start a fresh chat and ask for a three-heading outline of a task you already understand. Minutes 5 to 8: ask for a shorter version that uses headings. Minutes 8 to 10: copy the best lines into a note and name three parts of the screen in your own words.

IN THE FIELD

Samir's ten quiet minutes

Samir is a first-year nursing student in Manchester. He sets ten minutes, signs in, and starts a fresh chat. He says what he sees: a box to type, a model label he will not memorize, and a plus icon he thinks is for upload. He asks for a three-heading plan to revise one public lecture topic he already knows. He then types shorter and use headings. He copies those headings into his own note. One control is not where a friend's old screenshot put it. He does not freeze, because as of late 2026 the layout is only a sketch. He names the job of the control, new chat, and finishes inside the ten minutes with a note he can reuse.

A PROMPT WORTH KEEPING

I am in my first Gemini session. Guide me for 10 minutes, one step at a time. Step 1: tell me how to confirm I am in a fresh chat, meaning a clean context. Step 2: give me a practice task about planning a 30-minute study block for a topic I already know, with three headings. Step 3: show me the exact follow-up words to make it shorter. Constraints: do not ask for secrets, do not assume a button name that may have moved, and remind me to check the current screen. Output: three short steps I can do now, plus one line I should copy into my own note.

- One step at a time keeps a first visit calm.
- Fresh chat is defined as a clean context, so you know why you start over.

USUALLY. Describe every icon in the app before I am allowed to type. **BETTER.** Point me to new chat, the model label, and upload, then give me one small task.

PRACTICE

Run the ten-minute route

1. Sign in and name, in your own words, the model label, the upload control, and new chat.
2. Start a fresh chat and ask for a three-heading outline of a task you know.
3. Ask for a shorter version that uses headings.
4. Copy the best lines into a note and write one layout difference you noticed.

You have a copied outline in your note and three screen parts named in your own words.

- A fresh chat is a clean context. Start one before a new drill so old clutter stays out.
- Read the model label, find upload, and copy answers into a note you control.
- Two follow-ups earn their keep: shorter, and use headings. Edit your last prompt when you can.

Challenge: Your First 10 Minutes

You will run a timed loop: look at the screen, send a vague prompt on purpose, rewrite it with a full brief, and save the better version.

Reading about prompts is not the same as feeling a weak one fail. Ten minutes is long enough to compare a vague ask with a brief you can reuse. You stop when the timer ends so the habit stays small enough to repeat tomorrow.

This lesson is a timed practice, not a new theory chapter. Set a 10-minute timer and keep Gemini open. From minute 0 to 2, explore the screen. Name the model label if you see one, the place you would upload, and the control for a fresh chat. Do not hunt for a perfect layout. Say what you see in your own words, because the screen as of late 2026 can change. Check the current screen once, then start.

From minute 2 to 4, send a bad vague prompt on purpose. Something like help with the meeting is enough. Notice how wide and generic the answer is. You are not trying to feel foolish. You are noticing what disappears when you give no role, no context, no real task, no constraints, and no output shape. Keep that vague prompt so you can compare it with the rewrite.

From minute 4 to 7, rewrite the prompt with five pieces. Role is who the helper should be. Context is the facts it needs, with no secrets. Task is the job. Constraints are limits such as time, tone, and what not to invent. Output is the shape you want back, such as a timed agenda plus a note you can paste after the meeting. Send the rewrite in a fresh chat if the first thread is muddy.

IN THE FIELD

Elena times her standup prompt

Elena leads a facilities crew in Rotterdam. She sets a ten-minute timer. For two minutes she names the chat box, the model label, and the new-chat control. She types help me with standup and watches a vague answer wander past her real work. She rewrites with a role, six people, three remote, the Lobby refresh and the Roof leak, a fifteen-minute limit, and an output of a timed agenda plus a note template. She asks what was still unclear. The reply asks whether remote staff should speak first. She adds that line, saves the prompt in a note called Monday standup, and stops when the timer ends. The saved prompt is the result she keeps, not a fantasy of a perfect meeting.

A PROMPT WORTH KEEPING

You are helping a team lead at a small facilities company. Context: six people, a 15-minute Monday standup, three on site and three remote. The only projects are Lobby refresh and Roof leak. Task: draft the standup plan. Constraints: stay within 15 minutes, use plain language, no blame, do not invent project names, and let remote people speak first. Output: a timed agenda, three questions I will ask, and a four-line note template I can paste after the meeting. After the plan, list one thing in this prompt that is still ambiguous.

- The role and the real projects stop the model from inventing a different company.
- Context gives headcount, time, place, and who speaks first.

USUALLY. Help me with the meeting. **BETTER.** Draft a 15-minute standup agenda for six people covering only Lobby refresh and Roof leak, plus a four-line note.

PRACTICE**Run the timer once**

1. Set 10 minutes. For the first 2, name the model label, upload, and fresh chat in your own words.
2. Spend minutes 2 to 4 sending a vague prompt about a meeting or study block on purpose.
3. Spend minutes 4 to 7 rewriting it with role, context, task, constraints, and output, using no secrets.
4. Spend minutes 7 to 10 asking what was unclear, then save the better prompt in a note and stop.

Your note holds both the vague prompt and the rewritten prompt, and you stopped when the timer ended.

CHALLENGE**Save a prompt you will use Monday**

Run the ten-minute loop on a real recurring moment, such as a team standup, a study group, or a family admin check. Keep secrets out. Send the vague version only as a drill, then the full brief in a fresh chat. Ask what was unclear, repair one gap, and save the final prompt where you will find it on Monday. Add a check you will do with your own eyes, such as times, names, or places the model must not invent. If you finish early, do not add a second project. End on time so the loop stays repeatable.

Hint. A prompt you can paste again next week is a better result than a clever answer you forget.

- Minute 0 to 2 is for the screen. Minute 2 to 4 is a vague prompt you send on purpose.
- Minute 4 to 7 adds role, context, task, constraints, and output.
- Minute 7 to 10 asks what was unclear and saves the better prompt in a note.

LEVEL 2 · BEGINNER

The Art of Prompting

You will brief Gemini with role, context, task, constraints, and output so the first answer is usable, then repair it when a part of the brief was missing.

Prompt Anatomy: Role, Context, Task, Constraints, Output

A usable first answer starts with a brief: role, context, one task, constraints you can check, and the shape of the output.

Most weak chats fail before Gemini answers, because the brief never said who the work is for or what done looks like. RCTCO gives you a stable order so you stop hoping a clever sentence will replace missing facts. You will use the same order on email, study notes, and decisions, and you will still check anything that matters.

RCTCO — role, context, task, constraints, and output — is a briefing order, not a spell. Role — the working stance you want, such as a skeptical editor or a career coach. That stance decides what to protect and whose time is scarce. It is not a costume for fun. Pretending Gemini is a famous specialist does not create a license or real expertise. Name the review job, then put the facts in context.

Context — the facts Gemini cannot see: audience, country, decisions already made, and what you already tried. The model cannot see your files unless you paste or upload them, and it does not know your last meeting. Say the reader, the place, and the limit you already chose. If a long version failed, say so, so the next draft does not repeat it. Leave out secrets you would not hand to a new contractor.

Task — one main verb, such as rewrite, compare, or list. If you pile draft, research, and decide into one line, the answer blurs. Split the pile: finish one verb, then ask for the next in a follow-up. Constraints — limits you can check: length, must-include items, must-avoid items, reading level, and sources you will accept. Prefer under 120 words and no new statistics over a vague try your best. A rule you cannot test is only a wish.

IN THE FIELD

Diagram of a clinic update

A community clinic needs a Friday update for a manager who has three minutes. The weak line is Write something about the vaccine day. The brief below uses every RCTCO part. Role: you are a skeptical editor who protects the manager's time. Context: the reader is in Nairobi, the event was Saturday, attendance was 140 people, and yesterday's draft was too long. Task: rewrite the update. Constraints: at most 110 words, include the date and the attendance number, avoid medical advice, and use plain English. Output: a subject line and three short paragraphs labeled What happened, What is next, and What you need. The diagram is the prompt. Nothing in the role makes Gemini a clinician. You still check the number against the register before you send the note.

A PROMPT WORTH KEEPING

You are a skeptical editor who protects a busy reader's time. Context: the reader is a clinic manager in Nairobi. The vaccine day was Saturday. Attendance was 140 people. Yesterday's draft was too long and was rejected. Task: rewrite the update. Constraints: at most 110 words, include the date and the attendance number, no medical claims, plain English, and do not invent figures. Output: a subject line plus three short paragraphs labeled What happened, What is next, and What you need.

- Role sets a skeptical editor so the draft protects a busy reader's time, without pretending to be a clinician.
- Context supplies the city, the date, the attendance figure, and the rejected long draft, which Gemini cannot see on its own.

USUALLY. You are the world's greatest doctor. Write a fun update. **BETTER.** You are a skeptical editor. Rewrite this update for a clinic manager in Nairobi. No medical claims.

PRACTICE

Brief one real task with RCTCO

1. Pick a message you actually need to send this week, and write the audience and country in one line.
2. Add one role that names a stance, one main verb, three constraints you can check, and an output shape.
3. Paste the brief into Gemini and read the draft against each constraint, marking any miss.
4. Rewrite only the parts of the brief that caused the miss, then run it once more.

You have a saved prompt with all five RCTCO parts and a second draft that meets the length and must-include rules.

- RCTCO is a briefing order: role, context, task, constraints, and output.
- A role is a stance, such as a skeptical editor, not a costume that creates expertise.
- Gemini cannot see your files unless you paste or upload them, so context has to carry the facts.

Bad Prompt, Good Prompt

The same client update fails in six words and works when the brief names audience, decision, format, and the assumption you were hiding.

A short prompt feels fast, and it is fast to type. It is slow to repair, because Gemini fills every blank with a guess you did not choose. Learning to see the missing audience, the missing decision, the missing format, and the hidden assumption turns a rewrite into a diagnosis instead of a mood.

The six-word prompt Write a project update for client looks harmless. It hides the usual gaps. Audience is missing, so the model does not know whether the reader is a busy owner or a technical lead. The decision is missing, so it cannot tell if you need a yes, a date, or only a status. The hidden assumption is that Gemini already knows the project, the slips, and what you promised last week. It does not.

Diagnose before you decorate. Read the prompt and ask four questions. Who is the reader, and what do they already know? What decision should the message unlock? What format will they actually read: a short email, a table, or five bullets with a cap? What am I assuming the model knows that I never typed? Missing audience, missing decision, missing format, and a hidden assumption are the four names to write in the margin.

The repair is not a longer cloud of adjectives. It is the same goal with the gaps filled. You are a freelance designer writing to a client who pays the invoices and skims email on a phone. The site move slipped by four days because the photo files arrived late. You need a decision: accept the new Friday date, or cut the gallery. Format: an email under 130 words, subject line included, no blame, no new promises.

IN THE FIELD

Six words beside a full brief

Mina is a freelance designer in Accra. The homepage move slipped because the client sent photos late. She needs the owner to accept Friday or agree to cut the gallery. Her first prompt is six words: Write a project update for client. It names the artifact and nobody else. A strong brief names her stance as a calm project partner, the reader as the owner who skims on a phone, the four-day slip, the late photos, and the decision. It also sets the format: subject line plus an email under 130 words, no blame, no new features. The hidden assumption in the short prompt is that the model saw the contract. It did not. Paste the facts. Do not ask the model to remember a project it was never shown.

A PROMPT WORTH KEEPING

You are a calm project partner, not a salesperson. Context: I am a freelance designer in Accra. The reader is the client who pays the invoices and skims email on a phone. The homepage move slipped by four days because the photo files arrived late. I need a decision, not a status speech. Task: write the update. Constraints: under 130 words, subject line included, no blame, no new promises, plain English, and use only these facts. Output: an email the owner can answer with one choice, Friday with the full gallery, or the original day with the gallery cut.

- Role blocks a sales costume and sets a calm partner who respects the owner's time.
- Context supplies the city, the reader, the four-day slip, and the late photos, which the six-word prompt left as a hidden assumption.

USUALLY. Write a project update for client **BETTER.** Write an email under 130 words to the owner in Accra. Ask them to choose Friday with the gallery or the original day without it. Use only the facts I paste.

PRACTICE**Diagnose, then rewrite**

1. Write the six-word line Write a project update for client, then list the four gaps: audience, decision, format, and hidden assumption.
2. Replace it with a full brief about a real or practice project, using only facts you can share.
3. Run both prompts in Gemini and place the two answers side by side.
4. Mark the sentence in the strong answer that the weak answer could not have produced, and name which gap you filled.

You can point to audience, decision, format, and one pasted fact in the strong prompt, and you have both answers saved.

- Write a project update for client is six words and four missing decisions.
- Diagnose missing audience, missing decision, missing format, and the hidden assumption before you add adjectives.
- The repair is the same goal with the facts, the reader, and the shape filled in.

Iteration, Debugging, and Refinement

Treat the first answer as a draft of the conversation, name the symptom, and repair only the part that failed.

The first reply is evidence about your brief, not a final grade. If you rewrite the whole world after one miss, you will not know which sentence caused the trouble. A small repair, or a fresh chat when the thread is contaminated, keeps your time on the work instead of on the theater of starting over.

The first answer is a draft of the conversation, not a verdict on your idea or on the model. Gemini filled the brief you actually wrote, including the blanks you left. Read it the way you would read a junior colleague's first pass: what is usable, what is wrong, and what was never asked. If you throw the whole chat away after one miss, you learn nothing about which sentence caused the miss.

Prompt debugging — name the symptom before you rewrite the world. Pick one label. Too vague: ask for the missing names and the decision. Too long: set a word cap and say what to cut. Wrong audience: restate the reader and what they know. Invented facts: use only the notes, and mark gaps unknown. Generic voice: swap the mood word for two concrete rules. Missed constraint: quote the skipped rule and repair only that.

When one paragraph fails, do not ask for a brand-new answer. Quote the part that failed. Paste the weak lines, say what is wrong in one sentence, and ask for a revision of only that part.

Example: the second paragraph invents a budget number. Replace only that paragraph using my notes. Leave the subject line and the closing question unchanged. Partial revision keeps the good sentences stable.

IN THE FIELD**Quote the bad weeks, keep the rest**

A community tutor asked for a study plan and got a four-week university timetable with invented evening hours. The symptom was invented facts, not a bad subject. The repair did not say try again. It quoted the bad lines: Week 1 through Week 4 assume a six-week clock I never gave you. Replace only that timetable. List assumptions first, and do not write a new plan until I correct the list. Context to add: Form 4 mathematics in Kenya, 40 minutes on weekday evenings, paper in 18 days. Leave any greeting alone. This is prompt debugging in one move. Name the symptom, quote the part, limit the revision, and refuse to let a wrong assumption stay in the thread. If the next answer still talks about university, start a fresh chat.

A PROMPT WORTH KEEPING

List assumptions before you answer. My task is a study plan. Known facts: Form 4 mathematics in Kenya, 40 minutes on weekday evenings, paper in 18 days, past papers are available offline. Do not invent the exam board, the hours, or the topics. Output: a numbered assumption list only. After I confirm, I will ask you to revise a plan. If I say a line is wrong, drop it.

- The first line stops a fuzzy task from becoming a confident plan built on a guess.
- Known facts are separated from what the model is not allowed to invent.

USUALLY. This is wrong. Start over and do better. **BETTER.** The second paragraph invents a budget. Replace only that paragraph using my notes. Leave the rest unchanged.

PRACTICE**Repair one paragraph**

1. Send a short, fuzzy prompt on purpose, such as a study plan with no country and no hours.
2. Read the answer and label one symptom: too vague, too long, wrong audience, invented facts, generic voice, or missed constraint.
3. Quote the failing lines and ask for a revision of only that part, leaving the rest unchanged.
4. If the thread keeps the wrong assumption, start a fresh chat and ask for an assumption list before any new draft.

You have a quoted repair, a named symptom, and either a corrected paragraph or a fresh chat with a confirmed assumption list.

- The first answer is a draft of the conversation, not a verdict.
- Prompt debugging means you name the symptom before you rewrite the world.
- Quote the failing part and ask for a revision of only that part.

Tone, Audience, and Format

Tone is a set of rules you can check, the audience changes the shape, and a country plus a language beats a vague international voice.

People blame the model for a cold or flashy draft when the brief only said to be engaging. Tone, audience, and format are checkable rules, and they change who can use the answer. A busy owner and a technical buyer should not receive the same shape, even when the facts are the same.

Tone — a specification you can check, not a mood. Engaging, punchy, or professional tells the model almost nothing, because each reader hears those words differently. Write the rules you can hear in the sentences. Warm but direct. No slang. Short sentences. No exclamation marks. Make it pop fails because pop is not a test. Cut 30 percent, keep the ask in sentence one, and use no exclamation marks succeeds because you can check every rule with a pencil.

Audience changes the brief more than a clever adjective. A busy owner wants the decision, the cost in time, and the ask in the first two lines. A technical buyer wants the constraint, the risk, and what you already tested. Same offer, different page. Name the reader, their patience, and the words they use at work. If you write for both in one draft, you will satisfy neither.

Format is how the reader moves through the page. Ask for headings when there are parts. Cap the bullets, for example five bullets of one line. For a comparison, ask for a table with named columns. For email, ask for a subject line plus a body. Set reading level in plain words: short sentences, common words, jargon defined once. If you can sketch the shape on paper, you can check the draft against it.

IN THE FIELD

Pop versus a specification

Two people send the same request, Make the note pop, and both get slogans. The owner in Manila needed a decision about a late shipment. The technical buyer needed a table of parts. The working brief for the owner says: warm but direct, no slang, short sentences, no exclamation marks, cut 30 percent, keep the ask in sentence one, subject line plus body, country Philippines, language English. The working brief for the buyer says: neutral tone, columns named Part, Status, and Risk, no sales language, and the same country line. Mood words are not interchangeable with rules. If you cannot tick a box after reading the draft, the instruction was not a specification. Write the country every time so the model does not borrow another school system or calendar.

A PROMPT WORTH KEEPING

You are a direct writer. Audience: a busy owner in Manila who reads email on a phone. Country: Philippines. Language: English. Task: rewrite the note below. Constraints: warm but direct, no slang, short sentences, no exclamation marks, cut about 30 percent, keep the ask in sentence one. Output: a subject line plus a body under 90 words. Note: We should meet about the late shipment. The spare parts are four days late, so the install cannot start Monday. Let me know.

- Tone is a list of rules, warm but direct and no slang, instead of the mood word pop.
- Audience and localization name a busy owner, the Philippines, and English, so the note is not a generic international voice.

USUALLY. Make it pop. **BETTER.** Cut 30 percent, keep the ask in sentence one, and use no exclamation marks.

PRACTICE**Replace a mood word**

1. Take a short note and write the mood version first: Make it pop.
2. Rewrite the brief with warm but direct, no slang, short sentences, and no exclamation marks.
3. Add the reader, the country, the language, and one shape: headings, a bullet cap, a named table, or a subject line plus a body.
4. Run both and tick each rule the second draft actually met.

You have two drafts, and the second one keeps the ask early, names a country, and follows a shape you can sketch.

- Tone is a specification, such as warm but direct, no slang, and short sentences.
- Make it pop fails because you cannot test it. A cut of 30 percent and the ask in sentence one succeed.
- A busy owner and a technical buyer need different shapes, even with the same facts.

Templates You Can Reuse Tomorrow

A template is a prompt with blanks you save empty, fill with shareable facts, and reuse without chasing a menu name.

You should not invent a brilliant prompt from zero every morning. A template with blanks lets you reuse a structure that already worked, while the facts change. Saving that empty brief in a notes doc matters more than chasing a menu name that may move.

A template — a prompt with blanks you fill, not a fresh invention each time. Use the same blanks: [AUDIENCE], [GOAL], [FACTS], and [OUTPUT]. Save the empty version in a notes doc. As of late 2026 Google was rolling out Skills (reusable instruction files) and still had Gems (saved assistants). Check the current screen. A notes doc is enough. Never fill a blank with passwords, health records, or client secrets.

Meeting recap is the first template. Copy it and leave the blanks empty until notes are in front of you. You are a precise note-taker. Audience: [AUDIENCE]. Goal: a recap the owner can act on without attending. Facts: [FACTS]. Output: headings for Decisions, Actions with owner and date, and Open questions. Under 180 words. Do not invent attendees, quotes, or numbers. If a fact is missing, write unknown.

Explain-like-I-am-new is the second template. You are a patient teacher. Audience: [AUDIENCE]. Goal: explain [GOAL] so a newcomer can retell it. Facts: [FACTS]. Output: a short definition, one concrete example, and one check question. Under 160 words. Define each jargon term once as a plain phrase. Do not skip steps, and do not add history the facts do not support. Stop after the check question.

IN THE FIELD**One template, filled without secrets**

Here is the meeting recap template filled with ordinary facts, not secrets. Audience: the project owner, who was not in the room. Goal: a recap they can act on. Facts: Monday standup in Accra, three people, decision to ship the homepage Friday, photo files still late, open question on gallery size. You are a precise note-taker. Output: Decisions, Actions with owner and date, and Open questions. Under 180 words. Do not invent attendees or numbers. The empty version in your notes doc keeps the blanks. This filled version is one completed brief. If the real meeting included a password or a client account number, you would type unknown instead of pasting it. The template is the reusable object. The filled prompt is for one day only.

A PROMPT WORTH KEEPING

You are a calm advisor. Audience: [AUDIENCE]. Goal: choose [GOAL]. Facts: [FACTS]. Output: exactly three options, one tradeoff for each, a single recommendation, and what would change your mind. Under 220 words. Do not invent costs or dates. If the facts are too thin, list assumptions first and wait. Do not put secrets in the blanks.

- The blanks [AUDIENCE], [GOAL], [FACTS], and [OUTPUT] stay empty until you have notes in front of you.
- Three options, a tradeoff, one recommendation, and what would change your mind is the whole shape of an option memo.

USUALLY. Fill [FACTS] with the client's password so the note sounds real. **BETTER.** If a fact is a secret, type unknown. Never put a password in a blank.

PRACTICE**Save three empty templates**

1. Create a notes doc and paste the meeting recap, explain-like-I-am-new, and option memo templates with blanks still empty.
2. Fill only the option memo with a low-stakes choice, and leave out anything you would not email to a new contractor.
3. Run it and check for three options, tradeoffs, one recommendation, and what would change your mind.
4. If your Gemini screen shows Skills or Gems, check the current screen, then still keep the notes doc as the copy you control.

The notes doc holds three empty templates, and you have one filled option memo that invented no costs and contains no secrets.

- A template is a prompt with blanks such as [AUDIENCE], [GOAL], [FACTS], and [OUTPUT].
- Save the empty meeting recap, explain-like-I-am-new, and option memo in a notes doc.
- As of late 2026, Skills and Gems were product names in motion, so check the current screen. A notes doc is enough.

Workshop: 25 Beginner Prompts

Run a small pack on real work, keep the prompts that return a usable draft, and edit the ones that do not.

A list of prompts is useless until you run it on your own work. Five honest attempts will teach you which lines are too vague and which ones you want to keep. The pack in this lesson is a practice set, not a score, and you are allowed to edit every line.

A prompt pack is a set of briefs you reuse, not a poster you admire. Run five today in Gemini on a real task, and read the first answer as a draft. Save the ones that worked with only light edits. Edit the ones that did not: name the symptom, change one part, and run them again. Five finished runs teach more than a folder of unused lines. Keep the winners in a notes doc.

Understand prompts ask Gemini to clarify, not to invent. 1. Explain [topic] to a beginner in 120 words, define jargon once, and give one example. 2. List known, unknown, and assumed in this note, in three short lines: [paste]. 3. Turn this paragraph into five plain questions a newcomer would ask: [paste]. 4. Define [term] in one sentence, then show one correct use and one wrong use.

Transform prompts change a text you supply. Decide prompts force a choice. 5. Rewrite this email under 90 words, warm but direct, and keep the ask: [paste]. 6. Cut this draft by 30 percent, keep the request in sentence one, no exclamation marks: [paste]. 7. Turn these notes into a memo with headings Goal, Status, Blocker, and Next step: [paste]. 8. Give three options for [decision], one tradeoff each, a recommendation, and what would change your mind.

IN THE FIELD

Prompts 13 to 25, ready to copy

13. Summarize [notes] in 80 words and list three facts to check. 14. Translate this into clear English for [country]: [paste]. 15. Rewrite this at a simpler reading level: [paste]. 16. Compare [A] and [B] in a table: Cost, Time, Risk, Fit. 17. Name the decision, then draft a short yes and a short not-yet. 18. Stress-test [idea] if the budget is halved. 19. Give one worked example of [skill], then a similar exercise. 20. Role-play a skeptical client, one hard question at a time. 21. Drill [skill] with three items and a hint if I miss. 22. Turn [goal] into six steps for this week. 23. Split [project] into first, next, and later. 24. Plan ten study days, forty minutes each, plus one rest day. 25. List five risks in this plan and one check for each.

A PROMPT WORTH KEEPING

You are a practical coach helping a beginner build a small prompt pack. Context: my role is [ROLE], my country is [COUNTRY], and my language is [LANGUAGE]. Task: write 10 prompts I can run this week. Constraints: group them by understand, transform, decide, practice, and plan. Each prompt must include one task verb, one constraint, and an output shape. Use blanks like [TOPIC] and [PASTE] instead of inventing my private facts. Do not ask me to paste secrets. Output: a numbered list, then a one-line note on how I will test five of them today and edit any that fail.

- The meta-prompt asks for a personal pack instead of another generic list you will not run.

- Role, country, and language are blanks, so the pack fits your work without you pasting a secret.

USUALLY. I saved all 25 prompts, so I have finished the practice. **BETTER.** Run five today, note what you changed, and save only the ones that worked.

PRACTICE

Run five and note the edits

1. Choose five prompts from this pack of 25, including at least one understand prompt and one plan prompt.
2. Run each prompt in Gemini on a real task, and paste only facts you can share.
3. For each run, note what you changed: the symptom, the line you edited, or the reason you kept the wording.
4. Save the five final prompts in a notes doc with a title for when to use each one.

You have five prompts you actually ran in Gemini, each with a note about what you changed or why you kept it.

CHALLENGE

Build a five-prompt pack for one job

Choose one real job you have this week, such as a lesson, a status note, or a decision. Pick five prompts from the pack, covering understand, transform, decide, practice, and plan if you can. Replace the blanks with facts you are willing to share, and run all five in Gemini. For each run, write the symptom if it failed, or the sentence you would keep if it worked. Edit at least two prompts and run those again. Save the winners in your notes doc with a title that tells you when to use them. Do not save a prompt you have not run.

Hint. Edit one part only, such as the output shape or the length cap. If you rewrite every word, you will not know what fixed the draft.

- Run five prompts today, save the ones that worked, and edit the ones that did not.
- Group the pack by intent: understand, transform, decide, practice, and plan.
- A general explain prompt becomes a lesson plan when you add age, place, time, and an output shape.

PART TWO

Daily work

Planning, research, and writing. The aim is a week that moves, not a chat log you cannot reuse.

LEVEL 3 · INTERMEDIATE

Gemini for Productivity

Build a personal productivity assistant you can rerun every week.

From Chatbot to Personal Assistant

Install a standing brief so Gemini proposes plans and drafts while you keep every commitment.

Random questions waste the model's only real advantage, which is applying your constraints to a draft. A standing brief, a short note you reuse, stops every chat from starting blank. You will spend less time restating who you are, and you will catch invented meetings before they look official.

A random question treats Gemini like a search box. A standing brief treats it like a colleague who already knows your constraints. Write that brief once: your role, your working hours, what urgent means, the tools you use, and a ban on invented facts. Calendar times, commitments you have not made, and other people's availability are facts. If you did not supply them, the model must say it does not know.

Paste the brief at the start of a new chat. Ask Gemini to restate it in five lines and list anything missing. That restatement is a check, not a decoration. The model should propose options, drafts, and tradeoffs, and you commit. A proposal is not a booking, a sent email, or a promise to a client. If a calendar link appears, check the current screen before you trust a time.

Keep the scenarios realistic. A clinic coordinator can turn a messy task list into a draft day plan. That is not the same as asking the model to run the clinic. A founder can state office hours and keep customer names out of the chat. The payoff is a faster first draft of a plan, a reply, or a note. You still decide what ships when money or safety is involved.

IN THE FIELD

Two context notes, one rule

Priya leads regional operations for a parts supplier. Her context note says she works 08:30 to 17:00, Monday to Friday, and that urgent means a customer cannot be served today or a safety issue is open. She uses a shared task list, email, and one weekly staff meeting. Gemini may not invent her calendar, name a colleague as free, or promise a date she has not approved.

Jonas runs a three-person design studio. His note says deep work is 09:00 to 12:00, client calls sit after lunch, and urgent means a launch is blocked. He pastes tasks, never contracts. Both notes end with the same line: propose options, I commit. The notes do not make them faster at everything. They stop a confident plan from quietly booking time they do not have.

A PROMPT WORTH KEEPING

Here is my standing brief: role, working hours, what urgent means, tools I use, and a ban on invented calendar facts and commitments. Restate it in five lines, list missing facts, then propose options. Do not book, send, or promise. I commit.

- The brief is pasted in the same message so the chat does not rely on memory from another day.
- Urgent is defined by you, which stops the model from treating every task as a crisis.

USUALLY. Be my assistant and handle my week. **BETTER.** Here is my role, hours, urgent rule, and tools. Propose a plan. Do not invent meetings. I will commit.

PRACTICE

Write the context note you will actually reuse

1. Draft ten to fifteen lines: role, hours, what urgent means, tools, and facts Gemini must not invent.
2. Add the line that Gemini proposes and you commit, plus what you will never paste.
3. Open a fresh chat in the plain Gemini app, paste the note, and ask for a five-line restatement and the gaps.
4. Edit the note until the restatement no longer adds a meeting, a name, or a promise.

You have a saved note, a restatement that adds no facts, and a short list of gaps you still fill yourself.

- A standing brief is a note you paste, not a hope that the chat remembers you.
- Calendar facts and commitments are unknown until you supply them.
- The model proposes. You commit, send, and sign.

Daily Planner and Weekly System

Run a weekly planning session and a short daily pass that suggest blocks without inventing your calendar.

A week without a system becomes a list of reactions. A Monday plan and a Friday review give Gemini the same job every time, so you can see what to change. You stay faster at drafting the week, and you still own every commitment on the real calendar.

A weekly planning session is one sitting, not a mood. Dump every open task first, without ranking. Mark energy as high focus, medium, or low. Pick three outcomes: a finished result, not a motion such as check email. Park the rest so they stay visible without pretending to fit. The model can cluster and challenge those three outcomes. It cannot know which promise you already made to another person.

The daily pass stays short. Ask for a top three, suggested time blocks, and a shutdown list. A time block, meaning a proposed slice of the day, is not a booking. A shutdown list covers send, file, or an explicit postpone. Keep blocks as suggestions until you compare them with the day you have. Gemini does not know your calendar unless a connected feature is on and you confirm it. Check the current screen first.

State four meetings and a proposal due Thursday, and paste titles and times yourself. Label each meeting as given by you, not discovered. Ask for deep work only in the gaps you listed. A fifth meeting is an error. Demand a conflict check against the meetings you already named. Fiction dressed up as your calendar is a failed plan, even when it looks tidy.

IN THE FIELD**A Sunday prompt that survives Thursday**

Nia manages a four-person finance team. On Sunday she pastes nine tasks and three outcomes: finish the Thursday proposal, close two vendor questions, and publish the staffing note. She marks the proposal as high focus and the staffing note as low. The prompt lists four Thursday meetings with times and forbids any new meeting. It asks for work blocks only in the open gaps she named. Gemini suggests writing before the design review and a vendor block after lunch. Nia deletes a 16:15 slot because it hits her one-to-one. She keeps a shutdown list: send the draft to herself, file the notes, and postpone the rest. That saved prompt is her weekly productivity system.

A PROMPT WORTH KEEPING

Weekly productivity system. Dump follows. Mark energy high, medium, or low. Propose only three outcomes that are finished results, park the rest, and draft Monday as suggested blocks. Here are meetings I already have, with times. Do not add meetings. Flag conflicts with this list. End with a Friday review: moved, slipped, overridden, still open.

- The dump comes before ranking so a loud task does not hide a real outcome.
- Three outcomes force a choice. A parking list keeps the rest visible without pretending they fit.

USUALLY. Plan my perfect week and fill every hour. **BETTER.** Here is the task dump and the meetings I already have. Propose three outcomes and suggested blocks only in the gaps. Leave buffers.

PRACTICE**Build your weekly productivity system**

1. Paste your context note, dump this week's tasks, and mark each one high, medium, or low energy.
2. Ask for three outcomes, a parking list, a Monday plan of suggested blocks, and a Friday review.
3. Add every meeting you already have and require a conflict check. Delete any meeting the model added.
4. Save the weekly prompt and the daily prompt together and name the page weekly productivity system.

You can rerun the weekly productivity system from saved prompts, and the plan contains no meeting you did not supply.

- Outcomes are finished results. Tasks such as check email do not count as the week's point.
- Time blocks are suggestions until you compare them with the day you actually have.
- Gemini does not know your calendar unless a connected feature is on and you confirm the screen.

Email and Meeting Assistant

Draft three reply tones and turn rough notes into decisions, owners, dates, and a follow-up you still send yourself.

Email and meetings fail when a fluent draft adds a promise you did not make. A bounded assistant gives you tones, agendas, and a clean record of decisions, owners, and dates. You still send the mail and you still say what was decided.

Email help starts when you paste the thread. The model cannot see your inbox in the plain app. If the mail is sensitive, remove addresses and names that should not enter a third-party tool. When unsure, summarize the thread yourself and ask for reply shapes from that summary. As of late 2026, help inside Gmail exists only on eligible Workspace plans. Check the current screen. Pasting still works in the plain Gemini app.

Ask for three replies, not one voice. Direct means the decision in the first line. Diplomatic means the same decision with more care for the relationship. Short means a version someone can read on a phone. State your decision, or ask which decisions you still owe. A draft that invents a date, a discount, or blame is not ready. You edit the facts and you send. The model stops at the draft.

Build the agenda from the decision the meeting must reach, the minutes you have, and who decides. Ask for a short agenda, a question list, and a note template with three buckets: decisions, owners, and dates. Afterward, paste your own rough notes. Ask for a follow-up that repeats only those buckets. As of late 2026, Meet notes may exist only on eligible plans. Check the current screen. Pasted notes are enough when that control is absent.

IN THE FIELD

A waiver that stayed unapproved

Omar leads a small logistics desk. A partner writes that a shipment may miss Friday and asks Omar to confirm a Saturday delivery and a fee waiver. Omar is not sure the waiver is allowed, so he does not ask Gemini to decide the fee. He deletes the partner's personal mobile number, pastes the rest, and asks for three replies: direct, diplomatic, and short. Each must say the Saturday slot is unconfirmed and the waiver is unknown pending his manager. He picks the short one, adds the manager's real rule, and sends it himself. Later he pastes messy meeting notes. He asks for decisions, owners, and dates only, plus a follow-up that adds nothing new. A line that promised a crane he never mentioned gets deleted before the mail goes out.

A PROMPT WORTH KEEPING

Thread pasted below with extra addresses removed. Decision I have made: none yet on the waiver. Write direct, diplomatic, and short replies. Do not invent dates, money, or blame. Then, from the notes I paste, list only decisions, owners, and dates, and draft a follow-up that adds nothing.

- Removing addresses is a privacy step you do before the model sees the text.
- Three tones give you a choice instead of one voice you must fight.

USUALLY. Read my whole inbox and answer everyone. **BETTER.** Here is one sanitized thread. Draft three replies. I will choose, edit, and send.

PRACTICE**Three replies and a clean follow-up**

1. Choose one thread you may share, or write a short summary yourself if the mail is sensitive.
2. Ask for direct, diplomatic, and short replies that cannot invent dates, money, or blame.
3. Paste rough meeting notes and ask for decisions, owners, dates, and a follow-up that adds nothing.
4. Edit one reply until every remaining fact is one you would say out loud, then stop before sending.

You have three labeled drafts and a follow-up where each line points back to a note you wrote.

- Paste a thread, or summarize it yourself when the content is sensitive.
- Direct, diplomatic, and short are three drafts of your decision, not three different decisions.
- Meeting notes become decisions, owners, and dates. Anything else waits.

Decisions, Notes, and a Learning Queue

Draft a decision memo, theme a messy dump, and keep learning to a 30-minute block tied to real work.

Decisions get muddy when notes, ideas, and commitments sit in one blob. A memo, a themed dump, a short study block, and a three-pass brainstorm keep those jobs apart. You leave with a draft and a choice, not with a plan that quietly signed itself.

A decision memo is a short record, not a vibe. Name the decision in one line. List real options, including wait. Set criteria such as cost, time, risk, and whether you can undo it. Ask for one recommendation and a reason the other options lost. Add risks and a review date, the day you will look again. Gemini may draft the memo. You still choose when money, people, or safety are involved.

Start from the messy dump, not from a fake template. Ask for themes, then open questions under each theme. An open question is something the notes do not answer. Tell the model not to fill those holes with guesses. Park items that are interesting but not part of this week's outcomes. You want a map you can scan, with holes left visible. A tidy summary that secretly answers the holes has done harm.

Tie learning to work you already have. Choose one skill that would help a task you already have. Ask for a single 30-minute weekly block with a recall step and a tiny application. Recall means you try to remember before you reread. Reject any plan that needs four hours or a new personality. You pick a source you trust and you do the half hour. If you miss a week, shrink it.

IN THE FIELD**Sunday hours, a dump, and thirty minutes**

Lina runs a neighborhood bakery and must decide whether to add Sunday hours. She asks for a memo with the decision line, three options including stay closed, criteria of staff fatigue, extra cost, and possible revenue, one recommendation, risks, and a review date four weeks out. Gemini recommends a four-hour trial. Lina rejects it because her baker already works six days, a fact she adds in the next prompt. Separately she dumps messy notes from a course on bread pricing and asks for themes and open questions, with a ban on invented numbers. For learning she asks for one 30-minute Friday block on costing a single loaf, with three recall questions and one worked example using her own prices. She does not adopt a four-hour study plan.

A PROMPT WORTH KEEPING

Decision memo. Decision line, options including wait, criteria, one recommendation, why the others lost, risks, and a review date I must confirm. Do not commit. Then brainstorm in three passes: quantity, then my constraints, then one option to try. Idea generation is not a promise.

- Including wait stops the memo from treating action as the only serious choice.
- One recommendation is easier to accept or reject than a cloud of maybes.

USUALLY. Decide for me and email the team that we are doing it. **BETTER.** Draft the memo and stop. I will choose. Do not write a sendable announcement.

PRACTICE**Memo, map, and a half hour**

1. Name one real decision and ask for a memo with options, criteria, a recommendation, risks, and a review date.
2. Paste a messy note dump and ask for themes plus open questions, with guesses banned.
3. Ask for a single 30-minute study block tied to a skill that decision needs, including a recall step.
4. Run a three-pass brainstorm and mark the one idea you might try as not yet committed.

You can point to a recommendation you have not accepted, a list of open questions, and a 30-minute block that fits a real week.

- A decision memo recommends. It does not commit, especially where money or safety is involved.
- Open questions should stay open when the notes do not answer them.
- A learning queue is one 30-minute block tied to your work, not a fantasy routine.

Your Personal Productivity Workflow

Assemble one page of prompts and run it on Monday and Friday, checking for outcomes and a human decision.

Separate prompts help once. A single page you rerun on Monday and Friday helps every week. The quality checks stop the page from filling your day with tasks and fake decisions.

Keep one page you can open on Monday. Give it five pieces: a context note, a weekly prompt, a daily prompt, an email prompt, and a decision prompt. The note sets boundaries. The weekly prompt picks three outcomes and parks the rest. The daily prompt proposes a top three, suggested blocks, and a shutdown list. The email prompt asks for three tones and bans invented promises. The decision prompt drafts a memo and stops before you commit.

Run the page on Monday and again on Friday. On Monday, paste the note and the weekly prompt, then lock three outcomes before the daily plan. Then run the daily prompt for that day only. On Friday, review what moved, what slipped, what you overrode, and one line to change. Rejecting a suggestion still counts. Keep the page stable while you learn it. A rerun teaches more than a brand-new prompt every morning.

Use two checks. Did the plan name outcomes, meaning finished results, rather than tasks like check email? Did it leave a human decision, with a blank where a fact was missing? If every hour is filled and every mail sounds sent, the draft is overconfident. Ask for each line to be marked proposed or decided by you. You want more proposed than decided. A flow that never lets you choose is a script for accidents.

IN THE FIELD

A page that survived a school pickup

Hassan keeps the one-page workflow in a note titled Monday page. On Monday he pastes the context note and the weekly prompt, dumps twelve tasks, and changes the model's first outcome from clear inbox to send the supplier draft, because clear inbox is a task. He runs the daily prompt and rejects a block that overlaps school pickup, a fact the model could not know. On Wednesday an email prompt gives him three replies. He sends the short one after deleting a date the model invented. On Friday he reviews overrides: the pickup block, the inbox outcome, and the invented date. He changes one line of the weekly prompt so outcomes must be sendable or finishable. The page is not impressive. It is reusable, and it still leaves the supplier decision to him.

A PROMPT WORTH KEEPING

One-page workflow. Pieces: context note, weekly prompt, daily prompt, email prompt, decision prompt. Weekly prompt must name three finished outcomes and park the rest. Daily blocks are suggestions. Email bans invented promises. Decision memo stops before I commit. Mark each line proposed or decided by me.

- Naming the five pieces stops the page from collapsing into one vague ask for productivity.
- Finished outcomes are required so a task list cannot masquerade as a plan.

USUALLY. Build me an automated life and run it without me. **BETTER.** Build one page I paste on Monday and Friday. Leave blanks where a human decision belongs.

PRACTICE**Assemble the one-page workflow**

1. Put your context note, weekly prompt, daily prompt, email prompt, and decision prompt on a single page.
2. Add a Monday order: note, weekly outcomes, then the daily prompt only.
3. Add a Friday order: what moved, what slipped, what you overrode, and one line to change.
4. Run the two checks on a sample plan: outcomes versus tasks, and at least one decision left open.

A stranger could tell which lines are templates, which are outcomes, and which decisions are still yours.

CHALLENGE**Mark what you overrode in a messy week**

Use a real week, not a demo. Paste your context note and your weekly prompt into a fresh chat. Dump the tasks that are actually on your plate, with sensitive names removed. Let the model propose three outcomes, a Monday sketch, and a parking list. Then do the week as yourself. Each time you ignore a block, change a reply, or refuse a recommendation, write one line: what it said, what you did, and why. On Friday, paste those lines back and ask what to change in the page. Aim for a page that survives contact with your calendar and your judgment.

Hint. An override is useful data. Do not hide it to make the model look right.

- The workflow is one page: context, weekly, daily, email, and decision prompts.
- Monday locks three outcomes before the daily plan. Friday records what you overrode.
- A plan of tasks with no outcome, or a draft with no human decision, fails the quality check.

LEVEL 4 • INTERMEDIATE

Research, Learning, and Knowledge

Use Gemini to learn faster without outsourcing your belief to it.

Better Research Questions

Climb a question ladder, then leave the chat to open sources before you decide what you believe.

Vague research questions produce vague confidence. A ladder plus sources you open outside the chat gives you something you can defend. Gemini helps you think, but it does not replace the bibliography.

Climb a ladder instead of asking for everything. Start with a plain definition. Ask how the thing works. Ask where careful people disagree. Ask what evidence would count, and what would change a careful person's mind. End with what is local to your country or your rules, because a general answer can still be useless where you live. Gemini can draft those rungs. It cannot finish them from memory alone.

Do not collapse the work into one chat. Frame the decision. Ask the ladder. Collect sources outside the chat, such as pages or people you trust. Compare them in a table you can reread. Then write what you believe and what is still open, in your own words. Gemini is a thinking partner, not the bibliography. A bibliography is the list of works you opened. A line with no opened source stays in the open pile.

Ask for search queries, not for the final truth. A search query is what you will type into a search box. Also ask what a skeptical expert would demand before agreeing, and use that list outside the chat. Prefer a primary source, meaning the original document or official page, or a reputable summary that cites one. If the model offers a neat answer before you open anything, park it. Speed without a source is how confident mistakes spread.

IN THE FIELD

Lunch service, without a fake forecast

Amina owns a twenty-seat cafe and wonders about adding lunch. She does not ask Gemini whether lunch will be profitable. She asks for a ladder: what lunch service means for her shop, what would have to change in the kitchen, where two careful owners might disagree, what evidence would count, what would change her mind, and which answers are local. The model returns queries she can type into search and a skeptic's list: her own midday notes, her cook's view, and the city permit page. She opens the permit page herself. She writes that she believes the stove is the bottleneck only on days the cook confirms, and that demand is still open. No revenue number enters her note.

A PROMPT WORTH KEEPING

Frame the decision in one sentence. Then a ladder: definition, mechanism, disagreement, evidence, what would change a careful person's mind, and what is local to my place. Give search queries and what a skeptical expert would demand. Do not invent numbers or laws. I will open sources outside this chat before I write what I believe.

- The decision frame stops a general essay from replacing a choice you might make.
- The ladder separates how something works from where careful people disagree.

USUALLY. Tell me everything about whether I should add lunch, and give me the profit. **BETTER.** Build a question ladder and the evidence each rung needs. Do not invent profit or foot traffic. Mark local facts unknown.

PRACTICE

Ladder, then two sources

1. Write the decision you might make in one sentence, then ask Gemini for the full question ladder.
2. Ask for search queries and for what a skeptical expert would demand before agreeing.
3. Leave the chat and open at least two sources you chose, not links you have not checked.
4. Write five lines: what you believe, what is still open, and which model sentence you refused.

Your note names two sources you opened, and every number in it appears in one of those sources.

- A question ladder moves from definition to mechanism, disagreement, evidence, and local facts.
- Ask for search queries and a skeptic's demands, then go to a primary or reputable source.
- Gemini is a thinking partner, not the bibliography.

Summaries That Do Not Flatten the Truth

Name the summary you want, paste a source you may share, and delete anything the model added.

A summary can sound fair while it drops the caveat that mattered. Naming the summary type, pasting the source, and asking what was left out keeps the dispute visible. You never need the model to summarize the entire web.

Name the summary before you ask. One sentence should carry the main claim. Executive bullets help a busy reader decide whether to open the source. A contrast sets two positions side by side without picking a winner. A claim check separates what the author wrote from what people say they wrote. Say which type you want so less gets sanded off. A nameless summary tends to hide the dispute.

Paste the source, or upload a file you have the right to share. The model should use that text, not a memory of the public web. Ask what was left out: caveats, numbers, and the audience. Never ask it to summarize the web, because that request has no boundary. If the material is too private to paste, summarize it yourself first. Then ask only about structure, not about facts you withheld.

When two explanations conflict, do not ask which is obviously right. Paste both, label them A and B, and ask for conflicts: claims that cannot all be true, and claims neither text supports. Ask what evidence would decide the real conflicts. You still go and look. A comparison is a tool for judgment, not a verdict. If the model takes a side, it must quote the pasted line. No quote, no point.

IN THE FIELD

The cause that was never on the page

Noah must brief a colleague on a policy note he is allowed to share. He pastes the note and asks for one sentence, five bullets, and a list of what a shorter version would drop. He also pastes a second comment from a resident and asks for conflicts only: claims that cannot all be true. The model adds a cause the policy note never stated, that fees rose because of staffing. Noah asks it to list adds. He deletes the cause. He does not ask for a summary of the web, even though he is curious what other cities do. He writes that question as a search query instead and opens two city pages himself. His final note says what the author claimed and what is still unknown.

A PROMPT WORTH KEEPING

Source pasted below. I have the right to share it. Give one sentence, five executive bullets, and what a shorter version leaves out. If I paste a second view, list conflicts only, and quote the line you rely on. List any cause, number, or motive you added. Do not summarize the web.

- Naming the summary type tells the model which kind of flattening you will accept.
- The paste is the boundary. Memory of the web is not a source.

USUALLY. Summarize the web on evening markets and tell me who is right. **BETTER.** Here is the text I may share. One sentence, bullets, and what you had to leave out. Do not go beyond this paste.

PRACTICE

Summarize a bounded source

1. Paste a document you may share, or type a short practice text and label it as practice.
2. Ask for one sentence, executive bullets, and what a shorter version leaves out.
3. If you have two views, ask for conflicts and a quote from each side.
4. Demand a list of added causes, numbers, and motives, then delete those adds.

Every claim in your edited summary can be pointed at in the source, and you can name one thing the short version dropped.

- Name the summary: one sentence, bullets, contrast, or what the author actually claimed.
- Paste a source you may share. Never ask the model to summarize the web.
- Ask what was left out, and delete causes, numbers, and motives that were not on the page.

Study Plans, Socratic Learning, and Exams

Plan sessions you will actually do, let the model ask before it lectures, and keep exam practice honest.

Rereading a summary feels like study and often is not. Retrieval, a Socratic question you must answer, and an error log make the hour harder and more useful. None of that is a reason to cheat on a live exam or to hide uncited work.

Paste the syllabus and the hours you really have. Break that time into sessions. Each session needs retrieval practice, which means you close the notes, recall, and then check. Mix older topics back in so the week is not only the newest chapter. Ask the model to flag any exam date it assumed. Edit the plan to your real days. A timetable you will not follow is not a plan.

Socratic mode means the model asks and then waits. Say so in the first line, and limit it to one question at a time. Answer in your own words before it continues. Its reply should be a hint or a correction tied to what you wrote, not a full essay. If it lectures, tell it to delete the lecture and ask again. The discomfort is the point. You are practicing recall, not collecting a speech.

Paste a topic list you are allowed to use. Ask for practice questions and answer them before any key appears. If you have a mark scheme, the official scoring guide, paste it and allow scoring only against that guide. Keep an error log. Do not use Gemini to cheat on a live exam or to hide uncited work. A practice score is a drill, not a lie you can submit.

IN THE FIELD

Four half hours, not a heroic week

Farah has a biology quiz in nine days and about three hours total to study, not three hours a night. She pastes the topic list and asks for sessions with retrieval practice and a mix of older topics. The plan suggests a daily two-hour block. She rejects it and asks for four 30-minute sessions. For osmosis she switches to Socratic mode and answers a question wrong on purpose. The model starts to lecture. She tells it to delete the lecture and ask one hint. She then writes two practice questions, answers them before looking at any key, and logs one mix-up about direction of water movement. She does not take the model into the quiz room.

A PROMPT WORTH KEEPING

Syllabus and real hours pasted below. Build sessions with retrieval practice and mixed review. Flag any exam date you assumed. Then Socratic mode on the hardest topic: one question, then wait. Do not lecture first. Practice questions come only after I answer. Do not help me cheat on a live exam or hide uncited work.

- Real hours stop a beautiful timetable you will not follow.
- Retrieval practice means you recall before you reread, which is the point of the session.

USUALLY. Explain the whole chapter so I can skim it before the exam. **BETTER.** Ask me one question and wait. After I answer, give a hint tied to my words, not a lecture.

PRACTICE**A plan you can survive, then one question**

1. Paste a topic list and the minutes you will truly spend, and ask for sessions with retrieval and mixed review.
2. Reject any block that assumes hours you do not have.
3. Switch to Socratic mode on the hardest item: one question, your answer, then a hint.
4. Write three practice questions, answer them before any key, and log one mistake.

You can explain the idea with the chat closed, and you did not use the model as a live-exam substitute.

- A study plan uses the hours you have, retrieval practice, and mixed review.
- Socratic mode means one question, then a wait. A lecture first has failed the mode.
- Practice questions and an error log are honest exam preparation.

Fact-Checking, Uncertainty, and Sources

Highlight risky claims, demand supported or disputed or unknown, and open the page before you believe it.

Fluent answers hide numbers and names that were never checked. A routine that labels supported, disputed, and unknown tells you where to open a page. Important calls still belong to a human who can be held accountable.

Verify in a fixed order. Highlight claims that are numbers, names, dates, quotes, or laws, because those lines do damage when wrong. Ask the model to flag its own uncertainty on each one. Then check them in a source you open. A smooth paragraph with no highlights is unexamined, not safe. Leave adjectives for later. If you have ten minutes, spend them on the numbers and the names.

Invented citations are a known failure. A citation is a pointer: title, link, or page. The model can write a pointer that looks real and goes nowhere, or that lands on a page which does not say the claim. Do not put those pointers in a report until you have opened them. Tell the model not to name a source unless it appears in text you supplied. Your reader should inherit only links you opened.

Demand three labels. Supported means the pasted source says it and can point to where. Disputed means the sources you gave disagree. Unknown means the paste does not settle it, so the model must not guess. Ban empty caution, such as many experts believe, unless a named expert is in the source. If a label is missing, send the answer back. These words are a control. They show where your checking starts.

IN THE FIELD**The fine that was a different year**

Chris asks for a briefing on a new delivery rule and gets a fluent answer with a fine amount, a minister's name, and a quote. He highlights those three, plus a date. He asks for supported, disputed, or unknown, and for a quote from a source he has not yet provided. The model admits the fine and the quote were not supplied, then offers a link. Chris opens the link and finds a different year. He writes unknown next to the fine and replaces the year only after he reads the page. The minister's name was also uncited, so it stays out of the staff note. He adds one line: the person accountable for the briefing is Chris, not the chat.

A PROMPT WORTH KEEPING

Highlight every number, name, date, quote, and law in the answer below. Label each supported, disputed, or unknown. Supported must quote a source I pasted or a page I say I opened. Do not invent citations. If a label is missing, stop. Important decisions stay with me.

- The highlight list is the verification order: risky claims first, adjectives later.
- Three labels stop a shrug that sounds careful and decides nothing.

USUALLY. Give me sources. Any real-looking links are fine. **BETTER.** Do not name a source unless it is in the text I supplied. I will open pages myself.

PRACTICE**Three labels and one opened page**

1. Take an answer that contains a number or a name and highlight numbers, names, dates, quotes, and laws.
2. Ask for supported, disputed, or unknown on each highlight, with a quote only from text you provided.
3. Open a real page for the two riskiest claims and write a retraction for anything that was not there.
4. Add one line naming the human who is accountable if someone acts on the note.

No citation remains that you could not open, and every unchecked number is marked unknown.

- Check numbers, names, dates, quotes, and laws before you trust the adjectives.
- Supported, disputed, and unknown are controls. Missing labels go back.
- Invented citations are a known failure. Open the page. Do not trust a summary of the page.

Project: The AI Research Workflow

Carry one question through a checklist to a modest one-page brief, then verify a claim outside the chat.

A research project is a checklist, not a single clever prompt. The brief stays modest when empty rows and an uncertainty log are allowed to stay empty. Verifying one claim outside the chat is what makes the page yours.

Carry one question through a checklist. Frame the decision. Write questions from the ladder. Plan sources you will open outside the chat, and name what the model must not invent. Build a table with claim, source, and whether you opened it. Log lines that are disputed or unknown. End with a one-page brief in your words. The brief is a working note, not permission to treat every line as checked.

The example is a fictional tutoring service choosing hours. The owner wants evenings but does not know if families can travel then. The workflow must not let Gemini declare that most families prefer 18:00. It should ask what evidence would count, such as inquiry notes or posted school schedules. Hours with no source belong in the uncertainty log. The modest result is a shortlist to test and a list of open questions, not a revenue forecast.

Use the breakdown prompt to get a reusable brief template. Include the decision, the questions, the source plan, a comparison table with empty rows, an uncertainty log, and what you believe after checking. Empty rows stop invented sources. The top line should ban invented numbers, laws, and citations. The bottom line should say the human signs only what they checked. A template that arrives already full has failed.

IN THE FIELD

Evening hours that stayed unproven

A fictional tutoring service needs opening hours. The owner, using the workflow, frames the decision as which two evenings to test for a month, not how to become the city's top tutor. The question list asks what is local, including school dismissal times. The source plan names a school website and a week of parent messages, both still unopened, so the comparison table has empty rows. Gemini suggests that most families prefer 18:00. The owner moves that sentence to the uncertainty log. The one-page brief says evenings are a candidate and that preference is unknown. Later the owner checks one school page outside the chat, finds a 15:30 dismissal, and writes that the 18:00 idea is still unproven. No revenue claim is added.

A PROMPT WORTH KEEPING

Brief template, rows empty. Decision line. Question list from the ladder. Source plan of what I will open outside the chat. Comparison table columns: claim, source, opened yes or no, with blank rows. Uncertainty log for disputed or unknown. Section for what I believe after checking, left blank until I check. Ban invented numbers, laws, and citations. Human signs only what they checked.

- Empty rows stop the model from filling the table with sources it imagined.
- The source plan points outside the chat, which is where the bibliography starts.

USUALLY. Research my idea completely and fill the brief with sources and a revenue number. **BETTER.** Give me the template with empty rows. I will open sources, then write what changed.

PRACTICE**One question, one modest brief**

1. Frame the decision and write a question list from the ladder, including what is local.
2. Make a source plan and a comparison table with claim, source, and whether you opened it. Leave rows blank until you do.
3. Add an uncertainty log and a one-page brief that states only what you believe after checking.
4. Read the brief and move any unsourced number, name, or citation into the log.

The page shows the checklist, empty or honestly filled rows, and no citation you did not open.

CHALLENGE**Verify one claim outside the chat**

Finish the one-page brief on your tutoring-hours case or on a question of your own. Choose one claim that would change the hours, the offer, or the next action if it were false. Leave Gemini. Open a source you picked yourself: a page, a person you can ask, or a document you already hold. Write four lines: the claim, what the source said, what you changed in the brief, and what is still unknown. If the source disagreed, the claim must shrink or move to the uncertainty log. Do not replace it with a new number the model prefers. A modest correction is the successful result.

Hint. If the outside source is silent on your claim, move the claim to unknown rather than keeping it.

- The workflow is a checklist: frame, questions, source plan, comparison table, uncertainty log, one-page brief.
- Empty rows are a success when you have not opened the source yet.
- A template that arrives full of facts has failed.

LEVEL 5 · INTERMEDIATE

Gemini for Content Creation

Publish work that sounds like a person who knows the subject, with Gemini as the drafting bench.

How to Avoid Generic AI Writing

Learn the tells of generic AI prose, then force a voice sample, a ban list, and a specificity quota so the published piece still sounds like you.

Generic drafts waste the reader's time and your name. Models average the web, so they reach for safe phrases, even rhythm, and opinions that offend nobody. If you publish that voice, people stop trusting the piece even when the facts are fine. This lesson gives you a way to keep the speed and drop the fog.

Generic AI writing has a few tells you can learn. A balanced nothing-statement praises both sides and lands nowhere. The phrase in today's fast-paced world delays the point. Fake intimacy claims we all struggle, with no scene you lived. There is no specific detail from your notes, such as a street, a tool, or a number. Identical rhythm is the last tell: medium sentences and a soft lift at the end of each paragraph.

Humanize — your pass for lived detail, a real opinion, and a cut — is not a tone button. Add a scene you were in and a choice you would defend in public. Delete one paragraph that only repeats the intro. You are the author of record — the person whose name and judgment the piece carries. Gemini can propose lines. It cannot own the claim, the memory, or the decision to publish.

Start with a voice sample — about 150 words of your own past writing, not a wish about brand voice. Paste it above the task so rhythm and word choice are visible. Then add a ban list — phrases the draft must not use — inside the prompt. Include empty openers, fake intimacy, and any slogan you refuse to sign. Tell the model to mark a line if breaking a ban was the only way to finish it.

IN THE FIELD

The aisle tag Priya would sign

Priya edits safety manuals for a warehouse firm. She asked for a short post about her work. The first draft opened with a fast-paced world and claimed we all struggle to communicate. She pasted 150 words from an old email to a supervisor, banned that opener and the word journey, and set a quota of one concrete noun per paragraph from her notes. The notes said a Tuesday drill used a red tag on aisle C, and two pickers missed it because the tag sat below eye level. The revision kept that aisle, her opinion that tag height mattered more than a new slogan, and a cut of the closing pep talk. She signed the post. The model proposed lines. Priya stayed the author of record.

A PROMPT WORTH KEEPING

Role: editor who may not invent facts. Voice sample: paste about 150 words of my writing above this line. Ban list: in today's fast-paced world; we all know; journey; unlock; delve; tapestry. Task: rewrite the draft below in my rhythm. Constraints: specificity quota of one concrete noun from my notes in every paragraph; if you lack a noun, write [from my notes] and do not guess. Add one place where a real opinion is required and mark it [my opinion]. Cut one paragraph that only repeats the opening. Output: the revision, then a note of what you refused to invent. I am the author of record and I will check every noun.

- The voice sample is pasted first so rhythm is evidence, not a mood word like professional.
- The ban list sits inside the prompt, which stops empty openers before they reach the draft.

USUALLY. Make this sound human, warm, and engaging. **BETTER.** Here are 150 words I wrote. Ban these phrases. Use one noun from my notes per paragraph or leave a blank.

PRACTICE

Rebuild one paragraph you would sign

1. Paste about 150 words you actually wrote in an email, a report, or a message.
2. Add a ban list of at least five phrases, including one empty opener you dislike.
3. Ask for a rewrite with one concrete noun from your notes in each paragraph, or a blank.
4. Do the human pass: insert one opinion you would say out loud, and delete one repeated paragraph.

The piece contains no banned opener, every noun traces to a note you can point to, and you could put your name under it.

- Generic tells include balanced nothing-statements, fake intimacy, missing detail, and identical rhythm.
- A voice sample and a ban list belong inside the prompt, not in a wish after the draft.
- A specificity quota uses your nouns, and blanks are better than guesses.

From Idea to Publish

Move from idea to a published piece in seven separate prompts, and keep research, a hostile edit, and your own voice pass from collapsing into one mega-prompt.

A single prompt that asks for a finished piece hides weak facts inside smooth paragraphs. Stages let you stop when the research is thin, before you have fallen in love with the prose. The publish checklist keeps claim, audience, next step, and disclosure visible. You still decide what is true.

Do not ask for a finished article in one mega-prompt. The pipeline is seven stages, and each stage gets its own prompt: IDEA, RESEARCH, OUTLINE, DRAFT, EDIT, HUMANIZE, PUBLISH. A mega-prompt hides weak research inside pretty prose. Separate prompts let you stop, check, and throw away a stage without losing the rest. Name the stage in the first line of the prompt. Ask for only that stage's output. If the model jumps ahead to a full draft, send it back.

IDEA is a promise to a named reader, not a topic cloud. RESEARCH is where facts enter, and it still needs sources outside the chat when a claim could be wrong. Paste notes or links you trust, and tell Gemini to label anything it did not see there as unverified. OUTLINE comes only after

those notes exist. An outline lists claims, the evidence under each claim, and the point you will not make.

DRAFT turns one approved outline into prose and must not add new facts. EDIT is a hostile reader — a pass that looks for soft claims, missing proof, and lines that flatter you. Ask for problems before any rewrite. HUMANIZE — the student's pass for truth and voice — is yours, not a model pass you treat as finished. You add the detail only you know, the opinion you will sign, and the cut that makes the piece shorter.

IN THE FIELD

Jonah's night-class piece, stage by stage

Jonah wanted a 700-word piece on how his night-school class takes notes. One mega-prompt produced a smooth essay with a study he could not find. He restarted in stages. IDEA named adult students who work days. RESEARCH used only his three class handouts, and he marked every other claim unverified. OUTLINE held four claims and one point he would not make: that his method beats every other method. DRAFT stayed inside the outline. EDIT, read as a hostile classmate, flagged a line that flattered him. HUMANIZE added the room with the broken clock and cut a soft ending. PUBLISH stated the claim, the audience, a next step, and a disclosure his teacher required.

A PROMPT WORTH KEEPING

Stage: RESEARCH only. Do not outline in prose and do not draft. Sources: only the notes I paste under this line. Task: list claims that those notes support, the evidence line for each, and a separate list marked unverified for anything a careful reader might expect that is not in the notes. Constraints: no new studies, no percentages, no quotes I did not paste. Remind me which facts I should confirm outside this chat. Output: two lists, Supported and Unverified.

- Naming the stage stops a mega-prompt from sliding into a finished article.
- Limiting sources to pasted notes keeps research honest when facts matter.

USUALLY. Write a complete expert article with sources, a story, and a strong ending. **BETTER.** Stage: RESEARCH only. Use my notes. Mark everything else unverified. Do not draft.

PRACTICE

Run two stages and stop

1. Write an IDEA prompt that names one reader and one promise, with no draft attached.
2. Paste real notes and run a RESEARCH prompt that labels anything outside those notes as unverified.
3. Confirm at least one factual point outside the chat before you trust it.
4. Write the publish checklist even if you will not publish today: claim, audience, next step, disclosure rule.

You have two separate outputs, a note of what you checked outside the chat, and no full article pretending the research is done.

- One prompt per stage beats one mega-prompt that hides weak research.

- When facts matter, sources still live outside the chat.
- EDIT is a hostile reader, and HUMANIZE is your pass for truth and voice.

Social Posts, Hooks, and YouTube Scripts

Write hooks as payoffs you can keep, fit the piece to the platform, and build a simple script without fake testimonials or invented results.

Hooks and scripts fail when the opening promises a result the video never shows. A structure you reuse is faster than a new format every week. Headlines are a choice after criteria, not a vote for the loudest line. Invented testimonials are a trust problem, not a style problem.

A hook — the opening line that promises a specific payoff — is a contract with the viewer. Clickbait the video cannot keep is a broken contract, and it trains people to ignore you. Write the payoff first: what the viewer will do, decide, or stop fearing by the end. Then write the hook as a plain statement of that payoff. If the script does not deliver it, change the hook instead of raising the promise.

Platform fit — matching length and shape to where the piece will live — matters more than a clever line. A short post has one idea, one concrete detail, and one ask or none. An eight-minute script needs chapters — labeled sections a viewer can scrub to — because people arrive in the middle. Do not paste a blog into a caption or stretch a caption into eight minutes. Name the runtime you can record this week.

Use one script structure until it feels plain, then keep it. Cold open — a scene under twenty seconds that shows the problem. Promise — the payoff you can keep. Three beats — three steps or reasons, in order. Example — one story from your notes, with details you are allowed to share. Close — the payoff in one line. One ask — a single next step, not a stack of links. Cut any beat that repeats another.

IN THE FIELD

One bill, no shame, no fake statistic

Imani Brooks explains personal money ideas without shame. She will not invent a statistic, and she will not pretend a method rescued her in thirty days. Her notes say she keeps a separate account named Rent Buffer and that her first target was one bill, the water bill, not a giant sum from a blog. The hook promises that specific payoff: how to name the account and pick one bill. The script uses a cold open about a receipt in a kitchen drawer, three beats, that example, a close, and one ask to write the bill's name on paper. Ten headlines came back. She kept the one that matched the video and dropped every line that teased a secret result.

A PROMPT WORTH KEEPING

Role: script editor for a personal-finance video. Notes you may use: only the lines I paste. Task: write a hook that states a specific payoff, then a script plan with cold open, promise, three beats, one example, close, and one ask. Constraints: no invented statistics, no shame, no testimonials, no personal result I did not live. Chapters must match the beats. Runtime: the minutes I name. After the plan, list ten headlines and reject any that fail these criteria: payoff, match to the video, no shame, no invented result.

- The hook is tied to a payoff you can show, which blocks clickbait the video cannot keep.
- The six-part structure stops a caption from being stretched into a fake long video.

USUALLY. Write a hook that people cannot resist clicking. **BETTER.** Write a hook that states the payoff this video actually delivers, and cut it if the script does not.

PRACTICE

A hook you can pay off

1. Write the payoff in one sentence before you write the hook.
2. Draft a short post with one idea and one concrete detail from your notes.
3. Outline an eight-minute script: cold open, promise, three beats, example, close, one ask, plus chapter labels.
4. Request ten headlines and reject any that shame the viewer or invent a result.

You can point to the moment in the script where the hook's promise is delivered, and no line is a testimonial you do not have.

- A hook is a promise of a specific payoff, not a tease the video cannot keep.
- Short posts and eight-minute scripts are different shapes. Chapters help the long one.
- Cold open, promise, three beats, example, close, and one ask is enough structure.

Articles, Newsletters, and Product Descriptions

Outline anything long, keep a newsletter to one idea, and write product or ad copy only from facts and proof you supply.

Long pieces collapse when the draft comes before the outline. Newsletters get ignored when they carry three ideas and no scene. Product copy that invents a certification can become a real complaint, not just a bad paragraph. Ads without proof you supplied are guesses wearing a slogan.

Anything over about 400 words starts as an outline, not a draft. The outline lists the reader, the single claim, the sections, and the fact or story that earns each section. Approve it before prose. A model that drafts first will bury you in paragraphs you feel guilty cutting. Put a word budget on each section. If a section has no fact from you, mark it cut or park it. The outline is a decision, not a warm-up.

A newsletter — a short letter to people who chose to hear from you — carries one idea, one story, and one action. One idea means a sentence a tired reader will remember. One story means a scene from your work, not a survey of the field. One action means a step they can finish in a few minutes. Resist the second idea and park it for next week. The letter should feel finished.

A product description may use only the facts on the seller's spec sheet. Gemini must not invent materials, certifications, or reviews. If the sheet does not say organic, the description does not say organic. An ebook chapter outline is a planning artifact — a map you will still write and edit — not a promise that the model will produce a whole book you can sell unedited. Treat each chapter outline as a brief for a later human writing session.

IN THE FIELD**A kettle sheet that blocked a fake review**

A small kettle brand sent a spec sheet: 0.8 liters, stainless steel inside, plastic handle, no certification logos, and a two-year repair promise the owner actually offers. The model draft called it medical grade, hinted at chef reviews, and added a wooden lid. Those lines were rejected because they were not on the sheet. The same week the owner outlined a newsletter: one idea, that the handle stays cooler than the body; one story, a Saturday demo where a customer touched the wrong part; one action, read the handle note before you buy. An ebook outline for a care guide stayed a plan. Nobody treated it as a finished book.

A PROMPT WORTH KEEPING

Role: copy editor who cannot add facts. Spec sheet: I will paste it below. Task: write a product description using only those lines. Constraints: do not invent materials, certifications, reviews, awards, or test results. If a buyer might assume a feature that is absent, say the sheet does not claim it. Then stop. Do not write the ad's proof slot unless I paste proof. Output: description, then a list titled Refused.

- The role forbids new facts, which is the whole point of a spec-sheet description.
- Naming materials, certifications, and reviews blocks the usual inventions.

USUALLY. Write a premium description and mention that it is certified organic. **BETTER.** Use only my spec sheet. If organic is not listed, do not say it.

PRACTICE**Sheet, letter, and a plan**

1. Write a six-line spec sheet for a real product or an object on your desk.
2. Ask for a description that may use only those lines, plus a Refused list.
3. Outline a newsletter with one idea, one story from your notes, and one action.
4. Sketch three ebook chapters and label the page a planning artifact, not a finished book.

A friend can match every material and promise to your sheet, and the newsletter does not smuggle in a second idea.

- Over about 400 words, approve an outline before you ask for prose.
- A newsletter carries one idea, one story, and one action.
- Product text may not invent materials, certifications, or reviews.

Projects: Calendar, Script, and a Small Campaign

Finish three small artifacts: a two-week calendar, a six-minute script on a topic you know, and a five-post campaign for the fictional bakery Hearth & Rye.

Projects teach the pipeline by making you finish small artifacts. A calendar without a proof column becomes a wish list. A script on a topic you do not know invites invented expertise. A campaign for a fictional bakery is still practice in refusing fake awards.

The studio holds three small projects, not one blur. Project A is a 14-day content calendar with channel, topic, angle, source of proof, and status. Prompt 1: Build that calendar from topics I paste. Leave proof blank if I did not name a source. Prompt 2: Flag any row whose angle would fit any other shop, and ask me for a local detail. Do not invent awards or quotes.

Project B is a YouTube script of about six minutes on a subject you actually know. Prompt 3: Outline a cold open, a promise, three beats, one example from my notes, a close, and one ask. Prompt 4: Write the spoken draft at a calm pace and mark any line that needs a number I have not given. Chapters should match the beats. If you do not know the topic, do not pick it.

Project C is a five-post campaign for Hearth & Rye, a fictional bakery on Elm Lane. Nia Pell sells country sourdough and seeded rye. She runs a Saturday kids' class called Little Loaf, eight seats, ages six to ten. Prompt 5: Draft five posts using only those facts. Prompt 6: Reject any line that claims an award, a review, or a crowd size I did not provide.

IN THE FIELD

Elm Lane, with the award left out

Nia Pell owns Hearth & Rye, one shop on Elm Lane. She sells country sourdough and seeded rye, Tuesday to Saturday, from half past seven until two. Saturday she teaches Little Loaf, eight children, ages six to ten, forty-five minutes, parents stay in the shop. There is no award and no review quota. A student built a 14-day calendar with channel, topic, angle, source of proof, and status. One row said Saturday class, angle is what children actually shape, proof is the class sheet Nia wrote, status is draft. A six-minute script stayed on a topic the student knew: how Nia times the dough, using her own notes. Five posts used only the shop facts.

A PROMPT WORTH KEEPING

Role: studio producer. Task: build a 14-day content calendar. Columns: channel, topic, angle, source of proof, status. Constraints: use only topics and facts I paste; if proof is missing, leave the cell blank and set status to needs proof; flag any angle that would fit a different shop; do not invent awards, reviews, testimonials, or crowd sizes. Output: the table, then a short list of rows I must fix before they can move to draft. One prompt, one artifact. I will run the script and the campaign as separate prompts.

- Columns force channel, topic, angle, proof, and status instead of a vague list of content ideas.
- Blank proof and the needs-proof status stop a wish from looking like a finished post.

USUALLY. Fill two weeks of viral posts and add that Hearth & Rye won best bakery. **BETTER.** Use only the Elm Lane facts. Leave proof blank when I did not name a source. No awards.

PRACTICE**The two-week Hearth & Rye calendar**

1. Open a table with channel, topic, angle, source of proof, and status for 14 days.
2. Mark Sunday and Monday closed. Use Tuesday to Saturday only for shop posts.
3. Fill rows only with country sourdough, seeded rye, Elm Lane, and the Saturday Little Loaf class facts from this lesson.
4. Run the flag prompt: delete or repair any row that needs an award, a review, or a crowd you were not given.

Every open day has an angle and a proof source you can name, closed days stay closed, and no row claims an award.

CHALLENGE**Forty-five seconds you can defend**

Write the opening forty-five seconds of a six-minute script on a subject you actually know, not a hobby you only admire. Use a cold open, a promise you can keep, and a first beat. Ban awards, fake crowds, and testimonials you do not have. If you want a worked setting, you may use Hearth & Rye only with the Elm Lane facts from this lesson: sourdough, seeded rye, and the Saturday Little Loaf class. You are done when a viewer could not accuse the opening of a promise the rest of the video cannot pay.

Hint. Read the promise, then stop. If the next beat does not start to pay that promise, narrow the promise before you write more.

- A calendar needs channel, topic, angle, source of proof, and status.
- A six-minute script should be about something you know, with chapters that match the beats.
- Hearth & Rye stays factual: sourdough, rye, and Little Loaf, with no fake awards.

PART THREE

Judgment

Business claims, pictures, and documents. The model can draft a decision. It cannot own one.

LEVEL 6 · ADVANCED

Business and Marketing

Build a grounded marketing strategy for a practice company, then reuse the method on a real one.

Market Research, Personas, and Competitors

Use Gemini to structure desk research you supply, build personas from observed jobs rather than stereotypes, and keep a SWOT tied to evidence.

Strategy built on numbers the model invented will embarrass you in front of a partner. Personas built from stereotypes waste product and ad decisions. Public pages are usable evidence only when you paste them and date them. A SWOT without an evidence column is a mood.

Gemini can structure desk research — work from documents and public pages, not from a visit you have not made. It cannot know your street, your margins, or last month's sales unless you provide them. Paste numbers you are allowed to share. Ask it to file each input as fact from you, public page, or unknown. If a cell is unknown, it stays unknown. Invented sales in a tidy table are worse than blanks.

A persona — a working portrait of a buyer — built only from stereotypes is harmful. Build it from jobs-to-be-done — the progress a person is trying to make in a real situation — plus quotes you collected and disconfirming evidence — a fact that argues against the neat story. If you have no quotes yet, label the persona a hypothesis. Do not invent a name, an income, or a mood and call it research.

Competitor analysis starts from public pages you paste, not from the model's memory of a brand. Copy the offer, the price if it is visible, the audience they seem to address, and the proof they show. Ask what is absent. Note the date you copied the page, because offers change. Do not guess their margins or their private sales. Comparison is a reading of what they published, plus your notes on what you can and cannot match.

IN THE FIELD

Lantern Cycle, with the blanks left blank

Samir Adeyemi runs Lantern Cycle Repair with one other mechanic, near a commuter rail stop. He pasted last month's tally from his book: 40 tire repairs and 6 tune-ups, and he wrote that he does not repair carbon frames. He also pasted two public pages from nearby shops, dated the day he copied them. Gemini sorted weekday commuters and weekend riders only as hypotheses until his quotes were attached. One quote from a commuter said she needs the bike before the 7:40 train. A hobby rider's quote said he can wait a week. A SWOT row for strength cited the tire tally. A row for rival prices said needs a source when the page hid the number.

A PROMPT WORTH KEEPING

Role: research clerk who cannot visit my street. Task: build a SWOT with columns item, type, and evidence. Types are strength, weakness, opportunity, threat. Constraints: evidence must quote my notes or a public page I pasted, including the date I copied it. If you cannot point to that evidence, write needs a source and do not guess margins, sales, or motives. Keep opportunities separate from plans. Output the table only, then a list of unknowns I must measure myself.

- The clerk role blocks a visit the model cannot make and a story about your street.
- The evidence column stops compliments and worries from looking like findings.

USUALLY. Tell me my margins and how big the bike market is in my city. **BETTER.** Use only the tally and pages I paste. If a number is missing, write unknown.

PRACTICE

A SWOT with an evidence column

1. Paste a small set of facts you are allowed to share, or use Samir's tally as practice data.
2. Paste one public competitor page and write the date you copied it.
3. Ask for a SWOT where every row has evidence or the words needs a source.
4. Add one persona only if you have a quote. Otherwise label the persona a hypothesis and list the quote you still need.

No cell states a margin, rent, or market size you did not provide, and the persona is not a stereotype with an invented income.

- Gemini can structure desk research. It cannot know your street, margins, or last month's sales unless you provide them.
- A persona needs a job-to-be-done, quotes you collected, and disconfirming evidence.
- Competitor notes come from public pages you paste, dated, not from model memory.

Offers, Positioning, and Sales Messaging

State who the offer is for, the job, the alternative, and the proof you hold, and keep brainstorming separate from validated plans.

Positioning is how a buyer tells you apart from the alternative, not how clever the slogan sounds. Offers that hide exclusions create work you cannot deliver. Sales messages that invent customer counts are easy to challenge. A sketch that labels assumptions is safer than a sketch that looks certain.

Positioning — who it is for, what job, why this and not the alternative, and proof — is a sentence you can test, not a slogan contest. Name a person in a situation, the job they hire the offer to do, the alternative if you vanished, and the proof you hold. If proof is empty, say so. A line that could sit on a rival's door is not positioning. Cut adjectives until the alternative is clear.

Offer design states scope, price logic you supply, and what is not included. Scope is the work, the time, and the result you will hand over. Price logic is your reasoning from cost, time, or a comparison you choose. The model must not invent a price or a discount. Write exclusions in plain language. If you cannot deliver a part, put it in the exclusion list, not the headline.

A sales message has five moves: problem, cost of doing nothing, offer, proof, and next step. Cost of doing nothing is practical — time, risk, or a delay the buyer already feels — not a scare you made up. Do not invent customer counts. Say twelve offices only if you counted twelve. A landing page section list can follow the same five moves, plus a facts block and a question block.

IN THE FIELD

A tire job the rival cannot copy by accident

Samir tried a slogan, fast and friendly, and saw it could hang on any shop door. His positioning sentence named weekday commuters who must roll before the 7:40 train, the job of a safe tire before that train, the alternative of a shop that quotes three days, and proof from his own book that tire work is most of last month's jobs. He did not invent a customer count beyond the tally he wrote. The offer listed scope as tire, brake check, and a light test, price logic as his posted board, and an exclusion: no carbon frames. The sales message stopped at a next step he can keep, a weekday drop before noon.

A PROMPT WORTH KEEPING

Role: offer editor. Task: write positioning in four labeled lines: who, job, alternative, proof. Then write the offer: scope, price logic, exclusions. Constraints: price logic must quote the number or rule I type, or say missing price. Proof must quote a count or fact I typed, or say needs a source. Do not invent customer counts, awards, or reviews. If a line could be copied onto a rival without edit, mark it too generic. Separate a final line: brainstormed or validated, and say validated only if I pasted a buyer note, a cost check, or a channel test.

- Four labels force who, job, alternative, and proof instead of a slogan.
- Missing price and needs a source keep the model from inventing commercial facts.

USUALLY. Write a homepage and say thousands of happy riders trust us. **BETTER.** Use only my tally. If a count is missing, write needs a source. Do not invent customers.

PRACTICE

One sentence a rival cannot steal

1. Draft positioning with who, job, alternative, and proof, using a real offer or Samir's tire service.
2. Write scope, the price logic you supply, and at least two exclusions.
3. Draft the five-move sales message and delete any customer count you did not tally.
4. Mark the page brainstormed or validated, and list the outside check that would move it.

A rival cannot put their name on your positioning without the proof line becoming false, and no count is invented.

- Positioning names who, the job, the alternative, and proof you actually hold.
- Offers state scope, your price logic, and what is not included.
- Sales messages use problem, cost of doing nothing, offer, proof, and next step, without invented counts.

Email, Ads, Service, and Lead Capture

Draft a short welcome sequence, one ad concept with three variations, a careful service reply, and a checklist lead magnet without pretending Gemini is your lawyer.

Email, ads, and replies are where a sloppy claim meets a real person. A refund you did not authorize is not a small wording issue. A lead magnet should help even if the reader never buys. Rules for ads depend on your country, and a chat is not a qualified adviser.

A welcome sequence — three emails to someone who asked to hear from you — should help before it sells. Email one confirms why they signed up and delivers the useful thing you promised. Email two teaches one step with a real example from your work. Email three offers a next step you can fulfill, and a clear no. Do not invent a story about the subscriber. Mark any claim that needs a source.

Design one ad concept, then three variations that change the hook, not the facts. A concept is who, problem, offer, and proof you hold. A variation is a different first line for that same concept. Targeting ideas are hypotheses — guesses about who might respond — until a real campaign teaches you otherwise. Budget is your number. If you did not type one, the prompt must refuse to invent a spend.

A customer-service reply should be calm, specific, and limited. It must not admit legal fault or promise a refund, replacement, or timeline the owner did not authorize. Ask for a draft that notes the inconvenience, states what you know, states what you will check, and names the owner as the person who decides the remedy. Leave a bracket where the owner must choose. Do not let a polite model invent a policy.

IN THE FIELD

A noon drop-off, not a made-up policy

Samir drafted a three-email welcome for people who asked for his rain-season checklist. Email one sent the checklist. Email two showed how he seats a tire bead, a step from his bench, not a fable about the reader. Email three offered a noon drop-off and a plain way to say no. His ad concept stayed one idea, a commuter tire before the train, with three first lines and no new facts. Budget stayed the number he typed. A service reply about a late bike acknowledged the wait, refused to admit legal fault, and left a bracket for whether he would discount, because only he can authorize that.

A PROMPT WORTH KEEPING

Role: service editor, not a lawyer and not the owner. Task: draft a reply to the complaint I paste. Constraints: do not admit legal fault; do not promise a refund, replacement, discount, or deadline unless I explicitly authorize that remedy in this prompt; acknowledge the inconvenience; state only facts I typed; state what will be checked; leave a bracket named owner decision. Output the reply and a line titled Not authorized. If I also ask for an ad, refuse to invent a budget or to say the ad is legal in my country.

- The role blocks both a legal admission and an owner decision the model cannot make.
- The remedy ban is specific, so a polite draft cannot sneak in a refund.

USUALLY. Apologize, admit we were negligent, and promise a full refund tomorrow. **BETTER.** Acknowledge the delay, check the notes, and leave a bracket for the remedy I authorize.

PRACTICE

Three emails and one careful reply

1. List a checklist you can honestly send, with steps from your real work.
2. Draft the three-email welcome: deliver the checklist, teach one step, offer one next step and a no.
3. Write one ad concept and three first lines that do not add facts. Type your budget or write none yet.
4. Draft a service reply with no legal admission and a bracket the owner must fill.

The checklist is useful the same day, the ad does not invent a spend, and the reply does not promise a remedy you did not authorize.

- A welcome sequence is three emails that help before they sell.
- One ad concept can have three hooks. Targeting is a hypothesis, and budget is your number.
- Service replies do not admit legal fault or promise remedies the owner did not authorize.

Case Study: Northline Coffee Roasters

Work a fictional six-person roastery through a modest Gemini strategy: two office segments, one vague ad killed, and no miracle ending.

Case work is how you practice judgment before you spend money. Northline is fictional, and the point is the method, not a fairy-tale win. Killing a vague ad is a real outcome. Leaving a rainy-week tally blank is also a real outcome.

Northline Coffee Roasters is a fictional six-person roastery. They sell beans to offices and run a Saturday stall. Offices are pausing subscriptions. The stall is busy, yet rainy weeks lose money because staff and prep assume sun. This lesson does not end in a miracle. You will not be told that a prompt saved the company. You will see a modest use of Gemini: separate facts from guesses and kill a vague idea.

The Gemini strategy reads the notes. It does not invent a brand. Keep office buyers and stall buyers apart. Ask which office segment still has a reason to reorder, and which stall costs move when rain cuts the crowd. Refuse any plan that needs a figure the notes do not contain. The strategy prompt in the breakdown demands two segments, one thing to stop, and the facts still missing.

You also need two partner prompts beside that strategy prompt. The problem prompt forces every statement into fact from the notes, hypothesis, or needs a source. The segment prompt takes only the office notes and returns two groups, each with a job-to-be-done, a quote you paste, and a reason the group might say no. Neither prompt may invent a win. If the notes are thin, the right output is a short list of questions for the next real conversation.

IN THE FIELD**Three prompts, no rescue ending**

Three prompts carry this case, and none of them may invent a win. Problem prompt: Restate Northline using only the notes I paste. Label each line fact from the notes, hypothesis, or needs a source. Do not add counts or a happy ending. Segment prompt: Using office notes only, name two segments. For each, give the job-to-be-done, a quote I pasted, or missing quote, and a reason they might say no. Strategy prompt: Ask for those two segments, one rainy-week note, and one item to stop, with those labels and the headings Segments, Stop, and Missing. Run them in order. Then cut any victory line.

A PROMPT WORTH KEEPING

You are a careful analyst for a six-person roastery. Use only the notes I paste about Northline Coffee Roasters. Do not invent sales, savings, quotes, or a rescue ending. Task: write a strategy with exactly two office segments, one stall note on rainy weeks, and one item to stop. For each segment state the job-to-be-done, the evidence from my notes, and a reason they might say no. Label every line fact from me, hypothesis, or needs a source. Output headings: Segments, Stop, Missing. If evidence is thin, list what I must measure outside this chat, including a rainy Saturday tally. Do not recommend a community ad that names no one.

- The opening limits the model to pasted notes, which is the only evidence this case allows.
- Two segments and one stop force a choice instead of a vague community message.

USUALLY. Write the comeback story and include new sales figures that sound realistic. **BETTER.** Use only my notes. If a figure is missing, label it needs a source.

PRACTICE**Run the three prompts in order**

1. Paste the Northline notes from this lesson and run the problem prompt with labels.
2. Run the segment prompt for offices only, with two segments and no invented quotes.
3. Run the strategy prompt in the breakdown and keep the headings Segments, Stop, and Missing.
4. Delete any sentence that sounds like a saved company, then write the outside tally you still need.

You have two office segments, one stopped idea, and a missing rainy-week number that you did not invent.

- Northline is a practice case: offices pausing, a stall that struggles in the rain, no miracle.
- The process is problem prompt, segment prompt, then the strategy prompt.
- The modest result was two office segments and a killed community ad.

Project: A Complete Marketing Strategy

Produce a tagged strategy pack for Northline or your own business, then revise one slide after a skeptical finance read.

A strategy pack is a set of decisions you can defend, not a pile of headlines. Tags force you to separate what you know from what you hope. Milestones stop you from designing ads for a person you have not described. A skeptical friend is a feature of the project, not a threat to it.

The project is a strategy pack, not a slogan page. Build it for Northline Coffee Roasters or for a real business you may discuss. Include a persona, positioning, an offer, three messages, a one-page channel plan, three risks, and what you must verify outside Gemini. If you use Northline, stay inside the case facts. If you use your own firm, paste notes first. You own every line a customer might read.

Tag every claim. Use fact from me when the line comes from your notes or the case sheet. Use hypothesis when you are guessing and you know it. Use needs a source when the line would matter and you lack proof. Untagged praise is a fail. Do not invent customer counts, awards, or savings. A thin pack with honest tags beats a thick pack that sounds funded. Read the tags aloud before you call it done.

Work in milestones so you do not draft ads before you know who they are for. Milestone one is the persona with quotes or an explicit gap. Milestone two is positioning and the offer, including exclusions. Milestone three is three messages that reuse the same proof. Milestone four is the one-page channel plan and the three risks. Milestone five is the outside-verify list. The exercise walks these in order. Skipping to headlines is how generic marketing sneaks back in.

IN THE FIELD

Tags on a one-page Northline pack

A student built a one-page pack for Northline, not a new company myth. The persona for the kitchen buyer cited a quote from the case notes and was tagged fact from me. A line that reception staff want a brand story was tagged hypothesis. A line about rainy-week stall loss was tagged needs a source until a real tally exists. Positioning named the kitchen account, the job of a steady bean supply, the alternative of a paused subscription, and proof limited to the notes. Three messages reused that proof. The channel plan fit on one page. Three risks included the vague ad they already killed. The verify list left the chat.

A PROMPT WORTH KEEPING

Role: strategy editor for either Northline Coffee Roasters or the real business in my notes. Task: produce a pack with persona, positioning, offer, three messages, a one-page channel plan, three risks, and an outside-verify list. Constraints: tag every claim as fact from me, hypothesis, or needs a source. Do not invent customer counts, testimonials, awards, or savings. Price and budget stay blank unless I typed them. If I did not paste a buyer quote, write missing quote. Output one page, not a slogan pile. End with the single slide a skeptical finance reader should attack first.

- The pack list matches the project so you do not stop at headlines.
- Required tags make fact, hypothesis, and missing proof visible on the page.

USUALLY. Write a confident plan that shows we will grow quickly, with round numbers. **BETTER.** Tag every claim. If I did not supply the number, write needs a source and leave it out of customer-facing lines.

PRACTICE**Milestones for the strategy pack**

1. Milestone one: write the persona with a job-to-be-done and a quote or a missing-quote blank.
2. Milestone two: write positioning and the offer, including exclusions and your price logic or a blank.
3. Milestone three: write three messages that reuse one proof item, each claim tagged.
4. Milestones four and five: one-page channel plan, three risks, and the list of checks outside Gemini.

Every claim is tagged fact from me, hypothesis, or needs a source, and the pack names what you must verify outside Gemini.

CHALLENGE**One slide for a finance friend**

Finish the strategy pack, then hand one slide to a skeptical finance friend. The friend asks where the number came from. Revise that slide so every claim is tagged fact from me, hypothesis, or needs a source. If the friend would still reject a figure, remove the figure. Do not replace it with a smoother guess. You are done when the revised slide shows what would change your mind, and when nothing on it pretends the practice company has already won.

Hint. Pick the slide that sounds most certain. Certainty without a note behind it is the one to revise.

- The pack is persona, positioning, offer, three messages, a one-page channel plan, three risks, and an outside-verify list.
- Tag every claim fact from me, hypothesis, or needs a source.
- Milestones keep you from writing ads before you know who they are for.

LEVEL 7 · ADVANCED

Images, Documents, and Multimodal Work

Turn a picture or a file into a checked, actionable note.

How Gemini Reads Images and Documents

Learn what an upload actually gives Gemini, how to ask for a reading you can check, and how to keep private details out of the chat.

A picture is not a full story, and a file is not your whole computer. This course is independent of Google, so when a button name differs from a lesson, trust the screen you have open. If you demand a decision before any list of what was visible, you usually get a confident guess. The order in this lesson is the check that keeps a visual answer usable.

Multimodal — inputs beyond typed text, such as a photo, a screenshot, or a slide. The model — the Gemini system that reads your message and writes the reply — sees only what you successfully attach in that chat. As of late 2026, Gemini can often read uploaded images and many document types. Exact file types, size limits, and model names change by plan, so check the current screen for upload buttons and model names.

It does not see your desktop, other tabs, or a file you meant to attach and did not. If an upload fails, stop retrying in silence. Paste the text, or describe the image, and say that the file did not attach. Image limits and document types change, so a failure is often a format or size issue. Name the kind of file, what you need from it, and the shape of the answer you want.

Ask in three labeled parts. First, observations: what is visible, including words the model thinks it can read. Second, interpretation, labeled as inference — a conclusion that goes beyond what is directly visible. Third, a recommendation: what a person should do next, tied to those observations. Mixing the three in one unlabeled request invites a guess that sounds finished. You may still ask all three at once if the sections are named.

IN THE FIELD

A printer screen, read and then corrected

You photograph a printer display. The tray looks full, but the screen complains. You do not ask Gemini to repair the machine. You cover the serial number, upload the photo, and request three labeled parts: quote every word on the screen, list observations only, then suggest checks and mark each as a guess unless the photo proves it. Gemini quotes Load Tray 2 and Letter. You look again. The second word is Legal, not Letter. You send that correction and ask for a revision of only the lines that used the wrong word. The paper-size check changes, and the answer stops claiming every tray is empty. You keep two checks: reseal tray 2, and confirm the job size only after you reread the screen. The useful part was the quoted text, not a story about a broken sensor.

A PROMPT WORTH KEEPING

I uploaded one photo of a printer display. The serial sticker is covered. Do three labeled parts in order. One: quote every word you think you can read, and say if any word is unclear. Two: observations only, with no advice. Three: at most three checks a person could try, each marked as a guess unless the photo proves it. Do not claim you can see the rest of the office. If a word is tiny, say so.

- It names the single attachment so the model does not pretend to see the whole desk.
- The quote comes first, which gives you a line you can compare with the real screen.

USUALLY. Look at my screen and tell me what to do. **BETTER.** I am uploading one photo of the printer display. Quote the words you can read, then list observations, then suggest checks marked as guesses.

PRACTICE

Read one image in three passes

1. Pick a photo or screenshot you are allowed to share and that has no private names or numbers. If the upload control is missing, describe the image in words instead. Check the current screen for the upload button and the model name.
2. Ask for a quote of visible words, then observations only, then at most three next checks, each marked as a guess unless the image proves it.
3. Compare the quoted words with the image. If anything is wrong, send one correction and ask for a revision of only the affected lines.
4. Write a four-line note in your own words: what you saw, what the reply quoted, what you corrected, and what you will do next.

You have a note that separates quoted evidence from guesses, and you did not upload names, account numbers, faces of private people, or badges.

- Multimodal means the model can take inputs beyond typed text, but only the files you actually attach.
- Ask for observations first, interpretation second, and a recommendation third, and label those parts.
- Make the model quote visible words so you can compare them with the original before you act.

Screenshots, Charts, and Tables

Use a fixed routine for a screenshot, a chart, and a table so Gemini reports what is shown and admits what is missing.

Screenshots, charts, and tables look precise, so a wrong reading does extra damage. People paste a cropped settings page and then treat a guess about the hidden button as a fact. The same thing happens when a chart reply invents a trend the bars do not show. A short routine forces the reply to stay inside the frame.

Screenshot routine: ask what this screen is, what the user is trying to do, what looks blocked, and what they should click next. Require a note if the image is cropped — cut off so part of the screen is missing. If the save control is outside the frame, the model must say it cannot see that part. Do not let it invent the rest from memory of a typical layout. A typical layout is not your screen.

Chart routine: ask for the title, the axes — the labeled lines that say what is measured — the units, the time range, and the largest change. Ask what the chart does not show. A chart rarely shows causation — a claim that one thing made another happen. It may hide missing years or a scale that starts above zero. Do not invent a trend, a cause, or a forecast. Use only numbers in the prompt or on the image.

Table routine: have the model restate the column names in the order you used, not a tidier set it wishes you had. Ask it to recompute one total you can check by hand, and to point at anomalies — values that break the pattern — without explaining them away. If its total and the sheet total differ, stop and find the row. Do not ask it to clean numbers until you know which cell is the source.

IN THE FIELD

A settings page that is cut off at the bottom

You upload a screenshot of a wireless settings page. The bottom is missing, so no save control is visible. You say the image may be cropped, and you ask what the screen is, what you are trying to do, what looks blocked, and what to click next. You add one rule: if a control is not in the image, say it is not visible. A careful reply says this looks like a wireless page, a network name is selected, and no save control is in the frame. It refuses to tell you to click Save. It suggests you reshoot the full page, then compare any quoted label with the pixels. You reshoot, confirm the label, and only then follow a click path. The reply was useful because it admitted the crop.

A PROMPT WORTH KEEPING

There is no image. Use only these Northline Coffee stall revenue figures, in dollars: January 800, February 760, March 1100, April 1040, May 1500, June 1420. These six numbers are the entire data set. State the highest month, the largest month-to-month change with the subtraction, and what this list does not show. Do not invent a trend, a cause, or a forecast. Do not add a month I did not list.

- It says there is no image, so the model cannot pretend to read axes that were never provided.
- The six figures are named as the entire data set, which blocks extra months and outside knowledge.

USUALLY. Explain the growth trend in this chart and why next quarter will keep rising. **BETTER.** Use only the numbers I listed. Do not invent a trend, a cause, or a forecast. Name the largest change, show the subtraction, and list what the chart does not show.

PRACTICE**Read one screen and one tiny table**

1. Upload a non-sensitive screenshot of a settings or app screen, or describe it if upload is unavailable. Check the current screen for the upload control. Say that the image might be cropped.
2. Ask the four screenshot questions, and require the reply to say when a control is not visible.
3. Type a three-row table of your own with a total. Ask Gemini to restate the columns and recompute the total. Do not let it change a cell.
4. Check the arithmetic yourself. Write one sentence on what the screenshot did not show and one sentence on whether the total matched.

You have a reply that admits any crop, and a recomputed total you checked by hand, with no invented trend attached to either.

- On a screenshot, ask what the screen is, what the person is trying to do, what looks blocked, and what to click next.
- If the image is cropped, the model must say which controls it cannot see instead of filling them in.
- On a chart or a typed series, demand title, units, time range, the largest change, and what is not shown.

Product Photos and Document Extraction

Describe only attributes you can see or that a spec states, pull checkable fields out of a document, and treat the result as reading help rather than legal advice.

A product photo tempts the model to complete the listing with comfortable extras, such as dishwasher safe or organic. A contract tempts it to fill a blank address or to give legal advice you did not ask for. Buyers and signers get hurt by those extras. This lesson keeps copy and extraction inside the evidence you supplied.

For a product photo, ask for visible attributes only — qualities you can point to in the picture, such as color or a mark. Then ask for a short draft that may use those attributes plus a written spec you paste. If the spec is silent, the draft must be silent. Dishwasher safe, microwave safe, and material are not gifts the model may add because they are common. Say that a missing claim should be written as not stated.

Document extraction — pulling named fields out of a text so you can check them — should ask for parties, dates, amounts, obligations, and blank fields. An obligation — something a named party must do — belongs in the list only if the text says it. A blank field stays blank, labeled as blank, not filled with a plausible address. This is reading help, not legal advice. You still decide what the words mean, or you ask a qualified person.

To compare two images, ask what changed, what stayed, and what the pair does not prove. Example: photo A of a mug has no stamp, and photo B has a blue stamp near the foot. A careful reply says the stamp is the visible change. It does not say the clay or the care changed, because

those are not in the pixels. If words differ from the spec, cite both. Do not merge them into one cheerful sentence.

IN THE FIELD

Two mug photos and a blank in a short letter

A ceramics seller has two photos. In the first, the mug foot is plain. In the second, a small blue stamp appears near the foot. You ask what changed, what stayed, and what the pair does not prove. A careful reply names the stamp and refuses to claim the clay changed. You also paste a letter: Mara Ellison of Ellison Studio writes to Jonah Peck of Peck Cafe on 2 March 2026. She will deliver 40 stoneware mugs by 20 March 2026 at 28 dollars each. Half is due on 2 March and the rest on delivery. The delivery address is blank. No late fee is written. You ask for parties, dates, amounts, obligations, and blank fields. The address stays blank. You do not treat the list as legal advice.

A PROMPT WORTH KEEPING

I sell ceramics. Treat these photo notes as the only pixels you may use: cream glaze, handle on the right, small blue stamp near the foot, no care text in the image. Spec, and the only spec: Stoneware mug, 12 oz, cream glaze, hand-thrown. Care: hand wash. Not microwave tested. Price 28 dollars. List visible or spec attributes first. Then draft a product blurb of about 40 words that stays inside that list. If dishwasher safety is not stated, say it is not stated. Do not add a claim to make the listing stronger.

- Photo notes and the spec are separated, so a pixel and a written fact are not blended in secret.
- The attribute list comes before the blurb, which gives you a fence the copy cannot cross.

USUALLY. Write appealing mug copy. Mention that it is dishwasher-safe because buyers expect that. **BETTER.** Use only the photo notes and the spec I pasted. If dishwasher safety is not stated, write that it is not stated. Do not add a claim to improve conversion.

PRACTICE

Extract one letter and fence one product sentence

1. Paste the Ellison Studio letter from the example, or a document you are allowed to share. Ask for parties, dates, amounts, obligations, and blank fields. Say this is reading help, not legal advice.
2. Check every extracted field against the source. If the model fills a blank, send it back and require the word blank.
3. Write three visible attributes of an object you can look at, plus one care rule you are willing to state. Ask for a two-sentence description that uses only those lines.
4. Try to tempt the model with one extra claim you did not state. Confirm that a careful reply refuses it.

Your extraction leaves real blanks empty, and your product sentences contain no attribute you did not supply.

- List visible attributes and spec lines before any sales draft, and keep the draft inside that fence.

- If a care claim is not stated, the copy should say it is not stated instead of guessing what buyers expect.
- Extraction pulls parties, dates, amounts, obligations, and blank fields, and it is reading help, not legal advice.

Lab: Five Visual Challenges

Practice five visual tasks, from a confusing settings screen to a short action memo you can grade against the source.

Reading one example is not the same as producing a note you would hand to someone else. These five challenges reuse the routines from the earlier lessons on a screen, a chart, a product, a document, and a memo. The memo is where people smuggle in new facts, so you will grade it on purpose. Your own files stay in your hands: share only what you are allowed to share.

An action memo — a short note that says what you saw, what is uncertain, and what to do next — is the product of this lab. Grade it on four tests. Quoted evidence — words or numbers copied from the source, not a stronger paraphrase. Uncertainty — a plain statement of what you could not see. Next action — one step a person can take. No new facts — nothing the source did not state.

Challenge 1: analyze a screenshot of a confusing settings page you pick. Use the four screenshot questions and admit if it is cropped. Challenge 2: analyze a simple chart you upload, or use only these Northline Coffee figures: January 800, February 760, March 1100, April 1040, May 1500, June 1420 dollars. Do not invent a trend. If upload fails, type the numbers. Check the current screen for upload buttons and model names before you rely on a file type.

Challenge 3: analyze a product image you are allowed to share, or the mug notes from the previous lesson. Visible attributes first, then copy that stays inside them. Challenge 4: extract fields from a document you are allowed to share, or from the sample letter in the example on this page. Ask for parties, dates, amounts, obligations, and blank fields. This is reading help, not legal advice. Leave blanks blank.

IN THE FIELD

A sample letter you can extract without uploading a private file

Use this letter if you cannot share a real document. Mara Ellison of Ellison Studio writes to Jonah Peck of Peck Cafe on 2 March 2026. Ellison will deliver 40 stoneware mugs by 20 March 2026 at 28 dollars each. Peck pays half on 2 March and half on delivery. The delivery address line is blank. No late fee is stated. A careful extraction names Ellison and Peck as the parties, the two dates, the count of 40, the price of 28 dollars, the two payment obligations, and the address as blank. It does not invent a street, a late fee, or a signing recommendation. That same fence is how you grade the action memo in challenge 5: every operational sentence should survive the same point-back-to-the-source test.

A PROMPT WORTH KEEPING

Source, and the only source: Northline Coffee stall revenue in dollars, January 800, February 760, March 1100, April 1040, May 1500, June 1420. Write an action memo of about 150 words for the stall owner.

Include quoted figures, one uncertainty, and one next action. End with a verification line that says how I should check you. Do not invent a trend, a cause, a forecast, or a month. Do not give business advice beyond what these numbers support.

- It names the only source, so later sentences can be traced to that list or marked as failures.
- It requires quoted figures, which gives the memo evidence instead of a vibe.

USUALLY. Turn this screenshot into a confident memo about what the team should change this quarter.

BETTER. Write an action memo from this screenshot only. Quote visible labels, state what the crop hides, give one next action, and add a verification line. Add no new facts.

PRACTICE

Challenge 1: a confusing settings page

1. Pick a settings screen you find confusing and that you are allowed to share. Hide names, account numbers, and badges. If upload is unavailable, describe the screen and say the file did not attach.
2. Ask what the screen is, what you are trying to do, what looks blocked, and what to click next. Require the reply to admit if the image is cropped.
3. Compare every quoted label with the screen. Correct one mismatch if you find one.
4. Write five lines in your own words: screen, goal, blocker, next click, and what was not visible. This completes challenge 1.

You have a five-line note for a real settings screen, with at least one line that says what the image did not show, and no private identifiers in the upload.

CHALLENGE

Challenge 5: a 150-word action memo

Choose one source you already trust: your settings note, the Northline figures, the mug spec, or the Ellison letter. Ask for an action memo of about 150 words with quoted evidence, one uncertainty, one next action, and a verification line. Grade it yourself. Mark any sentence you cannot point to in the source. If a sentence adds a cause or a new fact, delete it or move the doubt into the uncertainty line. The next action must be something you can do without new facts, such as rechecking a number or reshooting a crop. Keep the verification line last. Do not accept the memo until it passes the four tests.

Hint. A memo fails if it sounds decisive but you cannot put your finger on the word or number that supports the sentence.

- The lab is five challenges: a settings screenshot, a chart, a product image, a document extraction, and an action memo.
- Grade the memo on quoted evidence, uncertainty, a next action, and no new facts.
- A verification line tells the reader how to check the memo against the original.

PART FOUR

Building

Small software, chained prompts, and a system you can run again next Tuesday. Then a project with your name on it.

LEVEL 8 · ADVANCED

Coding and Technical Work

Ship a small page, read errors, and improve a tiny app with Gemini as a careful pair, not a silent author.

Reading Code When You Are New

Learn how a tiny web page is built, how to ask for a line-by-line reading, and why a secret never belongs in the chat.

Code looks hostile when every mark is unexplained, so people either freeze or paste a whole project into chat. You only need enough reading skill to check a small page and to notice when a reply wanders. This lesson uses eight lines on purpose. If you can teach those lines back, later projects will feel like editing, not sorcery.

Code is instructions a computer follows literally. It will not guess the spirit of a line. A program — a saved set of those instructions — does only what the file says. A bug — a difference between what you wanted and what happened — is a mismatch you can locate, not a verdict on you. You can read a small page without becoming a professional developer. You still have to explain it back in your own words.

HTML — the structure of a web page — marks headings, paragraphs, and buttons. CSS — the look — sets color, space, and type. JavaScript — behavior in the browser — runs when someone clicks or types. Python — a popular language for scripts — is often used outside the browser. This lesson stays with HTML so structure is clear before behavior. The eight lines sit in the example. Read them before you ask for help.

Line 1, the doctype, tells the browser to use modern rules. Line 2 opens html, the root. lang is an attribute — extra detail inside a tag — and it names English. A tag — a label around one part of the page — usually opens and later closes. Line 3 holds head and title together. The head is information about the page. The title is the text on the browser tab, not the big heading visitors see.

IN THE FIELD

Eight lines, then a glossary

Keep this page short enough to hold in your head. Line 1 is . Line 2 is . Line 3 is . Line 4 is . Line 5 is

Saturday class

. Line 6 is

Kids bake at 10.

. Line 7 is . Line 8 is . Ask for a line-by-line explanation and a glossary of the tags. Then close the chat and say what the heading, the paragraph, and the button each do. If you cannot, ask about one line only. Do not paste any password, token, or key that happened to sit near the file on your machine.

A PROMPT WORTH KEEPING

Explain this eight-line page one line at a time, then give a glossary of doctype, html, head, title, body, h1, p, and button. Use plain words. Do not add CSS or JavaScript. Line 1 . Line 2 . Line 3 . Line 4 . Line 5

Saturday class

. Line 6

Kids bake at 10.

. Line 7 . Line 8 .

- It asks for one line at a time, so the reply cannot skip a tag or hide inside a summary.
- The glossary list names every tag you must be able to say back.

USUALLY. Explain my entire project folder and also here is my API key so you can test it. **BETTER.** Explain only these eight lines, one line at a time, and build a glossary. Do not send a key. A key is a secret and stays out of chat.

PRACTICE

Explain the eight lines back

1. Copy the eight lines from the example into a note. Do not add a key, a password, or a second file.
2. Ask for a line-by-line explanation and a glossary. Check the current screen for the model name so you know which reply you are judging.
3. Close the chat. On paper, say what doctype, html, head, title, body, h1, p, and button do.
4. Open the chat again only to check lines you could not explain. Correct your notes yourself.

You can explain the page without the model in front of you, and you never pasted a secret.

- Code is literal instructions. A program is a saved set of them. A bug is a mismatch between intent and result.
- HTML is structure, CSS is look, and JavaScript is browser behavior. Python is a common language for scripts.
- Ask for a line-by-line explanation and a glossary, then teach the page back without the model.

Project: Build a Simple Web Page

Build one local page for Hearth & Rye, open it in your browser, and change one detail yourself before you ask for anything more.

A first page should be small enough that you can see every line and still know what you asked for. Hearth & Rye is a practice shop, not a storefront that takes cards. You will save a file, open it on your own machine, and edit one thing before Gemini edits anything else. That sequence keeps you as the author.

Project one is a single page for Hearth & Rye, a bakery. Include a headline, the hours, a Saturday kids class, and an address placeholder — stand-in text you replace with a real street. There is no online payment. Ask for one HTML file with a style block — CSS placed inside the page rather than in a second file. This course is independent of Google. Check the current screen for the model name before you trust a feature claim.

Save the reply as a file named index.html. Double-click that file. The phrase means you press the file twice so the browser opens it on your computer. You do not need an account, a server, or a fee for this step. If the page looks empty, confirm you saved the whole reply and that the file name ends in html. Keep the address as placeholder text until you decide what you are willing to publish.

Change one thing yourself before the next prompt. Good first edits are the headline color or the class price. Then, if you ask Gemini, ask only for that change. A request to restyle the whole bakery invites a rewrite you may not be able to explain. After the edit, double-click the file again and look with your eyes. The chat is not the page.

IN THE FIELD

What the Hearth & Rye page must contain

The page is a paper sign, not a shop. Headline: Hearth & Rye. Hours: Tuesday to Friday, 7 am to 3 pm. Saturday, 8 am to 2 pm. Sunday, closed. Saturday kids class: 10 am, 12 dollars, pay in person at the counter. Address line: REPLACE WITH STREET ADDRESS. Say plainly that the bakery does not take payment on this page. One file holds the structure and a small style block for color and spacing. Use a real h1 for the bakery name. The class should be a paragraph or a section with a heading, not a fake button that charges a card. When the file is saved, double-click it and read every line against this list before you ask for polish.

A PROMPT WORTH KEEPING

Write one HTML file for a bakery called Hearth & Rye. Include a style block in the same file. Content: a headline, hours Tuesday to Friday 7 am to 3 pm, Saturday 8 am to 2 pm, Sunday closed, a Saturday kids class at 10 am for 12 dollars paid in person, and the address text REPLACE WITH STREET ADDRESS. No online payment, no form that takes a card, no second file, and no API key. Use a real h1, readable contrast, and short text. After the code, explain in a short paragraph how I save it as index.html and open it.

- It asks for one file and a style block, which prevents a pile of files you cannot open.
- Hours, class, price, and the address placeholder are specified, so the page cannot invent a menu or a checkout.

USUALLY. Build my bakery brand: site, payments, accounts, and a secret key I pasted below. **BETTER.** Write one HTML file with a style block, hours, a Saturday class, and a placeholder address. No payment, no accounts, and no key.

PRACTICE

Save, open, and edit one line

1. Send the build prompt from this lesson. If the reply is huge or includes a key field, reject it and restate the limits.
2. Save the file as index.html. Double-click it so the browser opens the page on your computer.
3. Change the headline color or the class price yourself. Save and double-click again until you see your change.
4. Ask Gemini only about that one change if something broke. Read the page aloud for heading order and contrast.

The local page shows Hearth & Rye, the hours, the class, a placeholder address, no payment, and one edit you made yourself.

CHALLENGE

One accessibility pass, not a rebuild

Duplicate your working page first and leave that copy untouched. On the working copy, replace the address placeholder with a fictional street you invent for practice, or leave the placeholder if you prefer. Ask Gemini for an accessibility check only: real headings, contrast, and button or link text that says what it does. Apply at most one suggested change, and only if you can explain it. Double-click the file and tab through it with the keyboard. If focus never lands on anything useful, say so in a follow-up instead of requesting a new design. Do not add payment, an account, or a key while you are improving access.

Hint. A big bold line is not a heading unless the tag is a heading. Button text should name the action.

- Ask for one HTML file with a style block, specific hours, a class, a placeholder address, and no payment.
- Save it as index.html and double-click to open it in the browser on your machine.
- Change one thing yourself, then ask Gemini only for that change if you need help.

Errors, Debugging, and Better Code

Read the useful part of an error, ask for the smallest fix, and only then consider a clearer shape for code that already works.

A red error feels like a wall, so people paste the whole project and accept the first giant replacement. Most beginner pages fail for a small reason, such as a name that does not match. This lesson shows that mismatch, then the habit of asking why after the smallest fix. You will not refactor a page that still does not run.

An error is a message from the computer about a line it could not finish. Read the first line first. It often names the failure more clearly than the pile under it. Then find the file name and the line number. The console — the browser panel that lists those messages — is where this text appears. Copy that first line, the file name, and the line number into your prompt. Do not paraphrase the error from memory.

Here is a small broken page, in words and in a short line of code. The button is written as id equals total-btn. An id — a name you give one element so a script can find it — must match exactly. The script calls `getElementById` with `total`, which searches for a different name. The console can report: `Uncaught TypeError: Cannot read properties of null (reading addEventListener)`, at the index file, line 18.

Null — a missing value — is what that search returns when no element has the id `total`. A `TypeError` — using a value in a way that value does not allow — appears because `null` has no `addEventListener`. Paste the error and the two relevant lines. Ask for the smallest fix. After it works, ask why it failed. Refactor — changing the shape of the code without changing what it does — comes only after the click works.

IN THE FIELD

The button and the wrong name

You build a page with a button that should show a total. The button line is: `Show total`. The script line is: `document.getElementById('total').addEventListener('click', showTotal)`. You click and nothing happens. The console first line says it cannot read `addEventListener` of `null`, and it points at line 18 of the index file. The names `total-btn` and `total` are not the same, so the search finds nothing. The smallest fix is to make those names match, either in the button or in the script, not both at random. After the click works, you ask why `null` appeared. Only then would you consider renaming for clarity. A full rewrite is how a one-word mismatch becomes a page you no longer own.

A PROMPT WORTH KEEPING

The button does nothing. I want the smallest fix, not a rewrite. Button line: `Show total` Script line: `document.getElementById('total').addEventListener('click', showTotal)` First console line: `Uncaught TypeError: Cannot read properties of null (reading 'addEventListener')` at `index.html:18` Tell me the smallest change and which name should win. Do not add features.

- It states the goal as the smallest fix, which blocks a replacement you cannot read.
- It pastes the button line and the script line so the mismatch is visible.

USUALLY. The page is broken. Rewrite the whole project in a cleaner style and I will paste whatever you return. **BETTER.** Here are the first error line, the file, the line number, and the two lines involved. Give the smallest fix only.

PRACTICE

Make the names match

1. Put the button line and the wrong script line from this lesson into a tiny page, or compare them on paper if you are not ready to run code.
2. Read the sample error and underline the first line, the file name, and the line number.
3. Ask for the smallest fix only. Apply that one change. Confirm the names match.
4. Ask why it failed. Write two sentences in your own words. Do not accept a rewrite.

You can point to the mismatched id, state the smallest fix, and explain null without reading the reply aloud.

- Read the first line of an error, then the file name and the line number, and paste those rather than a paraphrase.
- A button that does nothing may be an id mismatch: the script searches for a name the page never gave.
- Ask for the smallest fix, then ask why. Refactor only after the page works.

Project: A Small Useful App

Build a one-file quote builder in slices, then test the math so a blank field cannot show a broken number.

A tiny app teaches more than a large one because you can test every result by hand. The quote builder needs a name, hours, and a rate, then a sentence a freelancer could send. You will ask for those pieces separately so a failure has an address. No account and no server are part of the design, not omissions.

Project two is a one-file quote builder for a freelancer. Inputs: client name, hours, and rate. A button shows the total and a polite paragraph. JavaScript stays inside the file. There are no accounts and no server. Request the app in slices: layout, then math, then the sentence, then an empty-field check. If you ask for all four at once, you will not know which slice broke.

An event — the page reacts when someone clicks or types — is how the button becomes useful. The layout slice can show the inputs and a button before any math exists. The math slice multiplies hours by rate when the button is clicked. Do not skip ahead to wording. A sentence that sounds kind can still sit on a wrong total. Check the current screen for the model name if a reply mentions a tool you do not see.

The sentence slice writes a polite paragraph using the client name and the total you already tested. The empty-field slice runs when hours or rate is blank. NaN — not a number, a sign the math broke — appears when the page multiplies an empty input as if it were a number. Blank hours must not show NaN. Show a short request to fill the field instead. The page should not invent a zero and call it a quote.

IN THE FIELD**Four slices and two checks**

You tell Gemini you want a quote builder and that you will accept it only in slices. Slice one is layout: labels for client name, hours, and rate, plus a button that says Build quote. You open the file and see those fields. No math yet, and that is correct. Slice two is the click event and the multiplication. You type 2 and 50 and expect 100. Slice three adds a polite paragraph that includes the client name and that total. Slice four refuses to calculate when hours or rate is empty. You clear hours, click, and expect a plain message instead of NaN. You do not add accounts, saving, or a server between slices. Each slice is saved and double-clicked before the next prompt goes out.

A PROMPT WORTH KEEPING

The quote builder layout is done in one HTML file: client name, hours, rate, and a button labeled Build quote. Add only the math slice. When the button is clicked, multiply hours by rate and show the total. Do not write the polite sentence yet. Do not add accounts or a server. Tell me how to test it with 2 hours at 50.

- It names the slice, so layout and wording cannot creep into a math change.
- The click behavior is specific: multiply hours by rate and show the total.

USUALLY. Build the full quoting platform with logins, email, and a dashboard in one reply. **BETTER.** One file. Slice one is layout only. I will test it before I ask for math.

PRACTICE**Ship the quote builder in four slices**

1. Ask for layout only: client name, hours, rate, and a Build quote button in one file. Open it locally.
2. Ask for the math slice only. Test 2 hours at 50 and confirm the page shows 100.
3. Ask for the polite sentence only, using the client name and the tested total.
4. Ask for the empty-field check. Clear hours, click, and confirm you do not see NaN.

The one-file page shows 100 for 2 hours at 50, shows a polite sentence for a filled form, and does not show NaN when hours is blank.

CHALLENGE**Add a 10 percent discount and test it**

Add one checkbox labeled Apply 10 percent discount. Ask for that control and its math as a single slice, after the earlier tests already pass. When the box is unchecked, 2 hours at 50 must still show 100. When the box is checked, the same inputs must show 90, because 10 percent of 100 is 10. Blank hours must still avoid NaN, whether the box is checked or not. The polite sentence should mention the discount only when the box is checked, and it must use the discounted total. Reject any version that changes the base math when the box is off.

Hint. Unchecked stays 100. Checked is 90, not 110. A blank field still must not produce NaN.

- Build the quote tool in slices: layout, math, sentence, then the empty-field check.

- An event is the page reacting to a click. Keep the script in the same file.
- Test 2 hours at 50 and expect 100. Test blank hours and refuse NaN.

Habits for AI-Assisted Development

Use a professional routine even on small pages: plan first, run the file, paste real errors, keep secrets out, and save a working copy.

The risky moment is not the first prompt but the third, when a working file exists and a big rewrite looks faster than reading. A habit protects that file and protects other people's secrets. Developers use the same checks on a small page, and you can borrow the routine without borrowing a job title.

Describe the outcome before you ask for code. Then ask for a plan: files, steps, and what you will test. Demand small files you can open by double-clicking. Run the plan's first step and look at the real page. If it fails, paste the real error, not a story about the error. Write one sentence about what you understand before you accept the next slice.

Keep secrets out. A key, a password, a private customer list, and a payment credential do not belong in the prompt. If a school or employer requires disclosure — a statement that you used an assistant — do not claim you wrote code you cannot explain. The honest sentence is: I used Gemini, and here is the part I can explain. Check the current screen for the model name if you need to record which model you used.

Developers also use Gemini for tests, for comments, and for translating an error from a library — a package of code someone else wrote. Paste the error and the few lines involved, not the whole library. APIs stay conceptual here. You may paste a public API document and ask what a request returns. You do not drop a private key into the chat to try the live service. This course is independent of Google, so product buttons can change.

IN THE FIELD

Explain a public note, not a private key

You find a public note for a made-up service called Harbor Weather. It says: send a city name to the sample endpoint. The sample returns a temperature in celsius for Oslo or Tunis only. No key is used. Any other city returns a short error, not a guess. You paste that paragraph and ask Gemini to explain the request, the two sample cities, and what the error means. You do not invent a key, and you do not ask it to call the service. A careful reply defines endpoint as the address of the request, lists Oslo and Tunis, and says other cities are errors. It refuses to add a secret. You write two sentences in your own words and keep the public note beside them. That is an API explanation without becoming an account holder.

A PROMPT WORTH KEEPING

I want to add a note to my quote page that explains a public weather example. Do not write the code yet. Give a short plan: what will change, what file I will edit, and what I will test. Public note only: Harbor Weather sample takes a city name and returns celsius for Oslo or Tunis. No key. Other cities return an error. Remind me what not to paste.

- It asks for a plan before code, so you can reject the idea before a file changes.
- It names the public note and states that no key exists in the sample.

USUALLY. Here is my private key. Try the live API and drop the result on the page. **BETTER.** Explain this public document only. Do not call the service, and do not ask me for a key.

PRACTICE

Plan, copy, then decide

1. Pick a small change to a page you already have, or use the Harbor Weather paragraph from this lesson.
2. Ask for a plan only: outcome, file, and test. Reject any plan that wants a key.
3. Duplicate the working file and name the copy index-working.html before you accept code.
4. After any edit, run the old test that still matters, and write one sentence about what you understand.

You have a plan that forbids secrets, a preserved working file, and a sentence in your own words about the change you accepted.

- Describe the outcome, ask for a plan before code, keep files small, and run what you receive.
- Paste the real error, write what you understand, and keep secrets out of the chat.
- Do not claim you wrote code you cannot explain when a school or employer requires disclosure.

LEVEL 9 • POWER USER

Advanced Gemini Workflows

Stop chatting and start running systems you can reuse.

Multi-Step Prompting and Prompt Chaining

Run a large job as a checked chain: questions, source notes you paste, an outline, then a draft, and execute only a plan you have critiqued.

You can already get a fluent answer. The next skill is a controlled answer on a job that is too large for one prompt. Prompt chaining — making the checked output of one prompt the controlled input of the next — gives you a handoff you can reject. Without that pause, a weak middle step hides inside a polished ending.

Prompt chaining — making the checked output of one prompt the controlled input of the next — keeps each handoff visible. A single overloaded prompt asks Gemini to invent the questions, judge the evidence, design the outline, and write the memo in one pass. If the questions are shallow, the prose can still sound confident. The chain gives you a place to stop. You read the artifact, correct it, and only then paste it forward with a rule about what must not change.

Use a four-step chain for a research-backed memo. Step one asks only for the questions a decision needs, ranked. Step two is not a model invention: you paste source notes you gathered, and the prompt extracts claims, limits, and gaps. Step three builds an outline that maps each section to one question and one source. Step four drafts only from that outline. If the draft reaches for a new fact, the chain has collapsed. Name the allowed input in the prompt.

Iterative reasoning — proposing a plan, critiquing the plan, then executing only the approved plan — stops you from executing a foggy job. First, ask for a plan with assumptions, inputs, and what would make you reject it. Second, ask for a critique that names weak logic and missing facts. Only the third prompt may produce the artifact, and it must follow the approved plan. If the critique flatters the plan, send it back. Approval is your act, not a line the model writes about itself.

IN THE FIELD

A Saturday-hours memo in four checked steps

A clinic manager needed a board memo on Saturday opening hours and nearly asked for the research, the decision, and the prose in one prompt. She used the four-step chain instead. First she requested only ranked questions and deleted two the board would never vote on. Second she pasted six lines from the duty roster and the reception log, and asked for claims, limits, and gaps. The notes did not prove demand, so the gap stayed open. She refused a guessed number. Third she approved an outline of four headings, each tied to one question and one note. Fourth she asked for a draft that could use only that outline. The recommendation was labeled a staff proposal, not a fact the model had discovered. The board could see what was known.

A PROMPT WORTH KEEPING

Using only the notes I paste below, extract claims, limits, and gaps. A claim must point to a note. Do not add numbers, causes, or a recommendation. End with the single gap that most changes the decision. Notes: [paste].

- It names the allowed input, so the model cannot wander back into a general essay about clinics.
- It defines a claim as something that points to a note, which makes invented support easy to spot.

USUALLY. One prompt asks for research, a recommendation, a board memo, and a slide outline. **BETTER.** Four prompts, and you approve the questions, the note extract, and the outline before any draft.

PRACTICE

Chain a memo you actually owe

1. Write the decision in one sentence and list the notes you really have.
2. Prompt only for ranked questions, then delete any question those notes cannot inform.
3. Paste the notes and ask for claims, limits, and gaps, with no recommendation.
4. Approve an outline that maps each section to one question and one note, and stop before the draft.

You have a saved outline where every section names its question and its note, and the gaps are still visible.

- An overloaded prompt hides the step that failed inside fluent prose.
- A research memo chain is questions, pasted notes, outline, then draft.
- Iterative reasoning means you approve the plan before any execution.

Structured Outputs and Quality Control

Demand a shape you can scan, then run a hostile second prompt that scores the draft against a rubric without adding facts or praise.

A wandering essay makes you hunt for the decision, the evidence, and the risk. A structured output puts those in places you can scan before you trust the prose. Quality control is a second prompt with a hostile job, not a request to polish the style. The reviewer who only flatters you has not reviewed anything.

Structured output — a shape you can scan, rather than a wandering essay — is a power-user requirement. Name the headings yourself, or demand a table with fixed columns, or use a tagged layout such as DECISION / EVIDENCE / RISK / NEXT. The shape does not make the content true. It makes a missing part obvious. If EVIDENCE is empty, you see it at once. A smooth paragraph can hide that emptiness for a page. Name that shape before the model writes.

Quality control — a second prompt whose only job is to find misses — should be hostile on purpose. Give it a rubric of five checks. Completeness: every section you named is present. Faithfulness to sources: each claim traces to a note you pasted. Specificity: names, dates, and

numbers are concrete or marked unknown. Action: the next step has an owner. No new facts: the reviewer adds nothing the draft did not already contain. A pass on tone is not one of the checks.

The failure case is a reviewer that flatters the draft. It opens with praise, calls the work strong, and hides one soft suggestion at the end. That is not a review. Forbid praise in the prompt. Tell the reviewer to start with defects, to skip compliments, and to write none found if it cannot name a miss. If it still praises you, discard the review and rerun it. Praise feels helpful. It delays the repair.

IN THE FIELD

A vendor choice you can scan in one minute

A facilities lead had to choose between two service vendors and kept receiving essays that buried the price gap. She required this shape: DECISION, EVIDENCE, RISK, NEXT. DECISION held only the proposed vendor and the line that a named manager still had to accept it. EVIDENCE held three rows, each with a claim and the note it came from. RISK held the missing insurance certificate, marked unknown rather than assumed. NEXT named the manager, the question to ask legal, and a review date two weeks out. She then ran a hostile reviewer with the five checks: completeness, faithfulness to sources, specificity, action, and no new facts. The reviewer flagged a sentence that praised response times the notes never mentioned. She deleted that sentence. The manager could read the page without hunting.

A PROMPT WORTH KEEPING

Review the draft against this rubric only: completeness, faithfulness to sources, specificity, action, and no new facts. Start with defects. Do not praise. Do not add facts, numbers, or sources. Quote the failing line and name the failed check. If you find no miss, say none found.

- The rubric is closed, so the reviewer cannot wander into a new essay or a new recommendation.
- Starting with defects forbids the flattery failure, where praise hides a soft suggestion.

USUALLY. Please make this better and more professional. **BETTER.** Return DECISION, EVIDENCE, RISK, and NEXT, and leave unknown where the notes are silent.

PRACTICE

Put one real note into four tags

1. Take a short decision note you wrote this month and mark every number that has no source.
2. Rewrite it as DECISION, EVIDENCE, RISK, and NEXT, with unknown where a fact is missing.
3. Run the five-check reviewer with an explicit ban on praise and on new facts.
4. Repair only the lines the reviewer quoted, and name the human who owns the decision.

The page shows four tags, a named decision owner, a review date, and a review that starts with defects or says none found.

- A named shape beats a wandering essay because gaps show up immediately.
- The five checks are completeness, faithfulness, specificity, action, and no new facts.

- Forbid praise, or the reviewer will flatter the draft and soften the miss.

Pipelines for Research, Content, and Business

Redraw three pipelines you operate by hand: research, content, and business, and keep Gemini, source tools, spreadsheets, and calendars in separate jobs.

A pipeline is how you stop rebuilding a workflow from memory every Monday. You should be able to redraw it, name the handoff, and run it when the product screen changes. Tools can sit inside the pipeline, but they do not replace the operator when an agent feature is missing.

A pipeline — a fixed sequence of stages with a human check between them — is an operation you can redraw on paper. You are the operator. The model does not run the pipeline unless a product agent feature is on. As of late 2026, Gemini Spark was announced as a more proactive Workspace agent. Treat it as optional. Check the current screen, and still learn the manual chain. If the agent is absent or stalls, your written stages should be enough to continue the work by hand.

The research pipeline is question, source pack, claim table, uncertainty, then brief. Start with a question a decision needs, not a vague topic. The source pack is material you can open and paste. The claim table has three columns: claim, source, and limit. The uncertainty note records what the pack does not settle. The brief answers the question and repeats the open gaps. Gemini may structure the pack. It may not supply the sources, the numbers, or a citation you did not bring.

The content pipeline is audience job, outline, draft, hostile edit, human voice pass, then publish checklist. Audience job — the work the reader is trying to get done — comes before any idea of a clever post. Outline before draft. The hostile edit hunts fluff, unsupported claims, and invented proof. The human voice pass is you, reading aloud and cutting lines you would not say. The checklist confirms the facts, the offer, and the next step. Skip the voice pass and generic copy ships with your name on it.

IN THE FIELD

A community course launch, redrawn as three pipelines

A trainer wanted to launch a short community course and kept asking Gemini to do the research, the posts, and the business plan in one chat. She split the work. Research pipeline: the question was which session time the last twenty learners actually attended. She pasted her own signup sheet as the source pack, built a claim table, and wrote uncertainty where the sheet was blank. Content pipeline: the audience job was help a learner pick a session. She outlined, drafted, ran a hostile edit, did a voice pass, and checked that no post promised a job. Business pipeline: the problem was half-empty Friday seats. Options were a new time, a reminder, or a price cut. She recommended a reminder experiment and measured replies, not a feeling. Gemini drafted. The sheet stayed in a spreadsheet. The reminder sat on her calendar.

A PROMPT WORTH KEEPING

Build a claim table from the source pack I paste. Columns: claim, source, limit. Then list uncertainties the pack does not settle. Do not add sources, numbers, or a recommendation. Source pack: [paste].

- Fixed columns are a structured output, so a missing source cannot hide inside a paragraph.
- The limit column forces every claim to show how far it does not go.

USUALLY. Do my research, my posts, and my business plan, and run the agent so I can step away. **BETTER.** Name the pipeline, run one stage, and keep the human check even if an agent feature is on.

PRACTICE

Redraw one pipeline on a single page

1. Pick research, content, or business, and write the stage names in order down the page.
2. Under each stage, write the input you must paste and the output you must approve.
3. Mark which tool does that stage: Gemini, a source notebook, a spreadsheet, or a calendar.
4. Run only the first two stages on a real task and stop for your check.

A colleague could follow the page without your chat, and you have approved the output of stage two.

- You are the operator. An agent feature, including Gemini Spark, is optional.
- Research runs from question to source pack, claim table, uncertainty, and brief.
- Content runs from audience job to outline, draft, hostile edit, voice pass, and checklist.

Project: The Gemini Power User System

Build a one-page or two-page operating system with seven prompt sets, version it, and prove you can repeat a result rather than remember a lucky chat.

A lucky answer disappears when the chat does. A power user needs a page that produces a similar result next week, with a different topic and the same standards. This project is that page: your standing context, your job prompts, and a review that is allowed to be harsh, not a new personality for the model.

Stop collecting chats. A power-user system is one or two pages: standing context, the job prompts, and a review you can rerun. Paste the master prompt into a fresh thread before any task prompt. Master prompt: You are my standing assistant for [role] on [project]. Use only these facts: [three facts]. Do not invent sources, numbers, quotes, or private details. Do not praise drafts. If a fact is missing, ask and stop.

Keep research and writing apart so a draft cannot quietly invent a source. Research prompt: Using only the notes I paste, list claims, the note each claim comes from, and gaps. Do not add facts or citations from memory. Rank gaps by what the decision needs. Writing prompt: Draft from the approved outline only. Match this reader: [reader]. Mark any sentence the outline does not support, and do not fill the gap yourself.

Name the constraint or these prompts will invent a calendar and a budget. Productivity prompt: Build a one-week plan from the tasks and limits I paste. Return a day-by-day list, one deep-work block a day, and the item to drop if time slips. Do not add tasks. Business prompt: Frame the problem, then three options with evidence, cost, and what would prove the option wrong. Do not recommend until I ask.

IN THE FIELD**Version the page, and do not depend on one menu**

As of late 2026, Gems — saved assistants — and Skills — reusable instruction files — were the product places some students stored a standing prompt. Those surfaces roll out unevenly. If your menu differs, check the current screen, and store the same system in a document you control. Name the file with the version. v1 is the first page you can actually run, not a fantasy of twenty prompts. When you change a prohibition or a shape, save v2 and write one line on what changed and why. Review the page monthly. Retire a prompt you did not use. Improve one prompt that failed. Do not keep six variants of the same job. A colleague should be able to start a fresh chat from v2 alone: master prompt first, then one job prompt, then the quality-control prompt. The document is the system. The chat is only this week's run.

A PROMPT WORTH KEEPING

Quality-control prompt: Review the draft against the rubric I paste. Start with defects. No praise. Do not add facts. Quote the failing line and name the failed check. If none, say none found. Rubric: [paste]. Draft: [paste].

- It is a reusable instruction, so the review does not depend on the mood of a fresh chat.
- The rubric is pasted each time, which lets you tighten v2 without rewriting the whole master prompt.

USUALLY. A new chat with no master prompt, hoping the model remembers last month. **BETTER.** Paste the master prompt from the current version, then one job prompt.

PRACTICE**Fill the master prompt and the quality-control prompt**

1. Duplicate the starter master prompt and replace the role, the project, and the three facts with your real work.
2. Add one prohibition that matters in your job, such as no student names or no unpublished figures.
3. Duplicate the quality-control prompt and paste the five-check rubric under it.
4. Save the page as v1 with today's date and one sentence on what this version is for.

v1 contains your filled master prompt and your filled quality-control prompt, and both name limits you would actually enforce.

CHALLENGE**Three prompts, one real task, one caught defect**

Pick a real task that is allowed to leave your desk. Run it through three prompts from your page: the master prompt, one job prompt such as research or analysis, and the quality-control prompt. Paste only the notes the job prompt allows. Save the three outputs. Then write a short record of what the quality-control step caught: the quoted line, the failed check, and the edit you made. If it caught nothing, tighten the rubric and run it once more. File the record on the page as the reason v1 became worth keeping.

Hint. Choose the job prompt before you start, and do not let the model skip ahead to a polished final while the review is still open.

- The system is a page you can rerun: master, job prompts, and a harsh review.
- Store it in the product if the screen offers that, and in a document if the menu differs.
- Version the page v1, v2, and review it monthly by retiring one prompt and improving one.

LEVEL 10 • POWER USER

Gemini Masterclass: Build with AI

Finish a real project with a brief, milestones, a critique, and a revision.

Choosing Your Master Project

Pick one of six projects using the material you can check, the time you have, what you want to show, and whether private data would be exposed.

The master project is where you stop sampling prompts and finish something you can show. The wrong choice is the most impressive title you cannot feed with facts. Pick with a rule, then stay with that project through the brief, the critique, and the revision.

AI-powered — Gemini used inside your workflow, not a model you train — describes all six options. They are an AI-powered business system, an AI content business, a personal AI productivity system, an AI-powered research assistant workflow, a simple AI-powered application, and an AI-powered marketing agency workflow. You are not choosing a costume or an income plan. You are choosing work you can finish with material you can check. The tool drafts and structures. You decide, edit, and verify.

Read the choice as a decision table in prose. Access to real material comes first: if you cannot supply facts, you may not pick the impressive project. Time available this week comes second: one sitting favors the personal productivity system or the simple application, while a longer week can hold the research workflow. What you want to be able to show comes third, and only among projects the first tests allow. A content operation fits when you know one audience well enough to spot a fake.

Private data — information a person or client did not agree to put in a chat — includes names, messages, health notes, student records, unpublished numbers, and passwords. If the project touches private data, you must use fictionalized or permitted material. Fictionalized means the shape of the work stays and the identifying details change. Permitted means you have a real right to use it. The business system, the agency workflow, and the research workflow are where this test bites most often. Skipping it is a leak, not a shortcut.

IN THE FIELD

A coordinator who wanted the agency title

Nia wanted the marketing agency workflow because it sounded like the largest project. She had no client, no permitted numbers, and one evening. The decision table was blunt. Access to real material: she did have her own community workshop notes and a list of questions people ask at the door. Time available this week: one evening, so a month of campaign assets was a fantasy. What she wanted to show: a page a colleague could follow. Private data: the door questions included a child's name, which she removed. She did not pick the agency workflow. She picked the simple application, a one-screen aid that turns a door question into a checklist, and she fictionalized the remaining names. The impressive title lost. The checkable project won. That is the whole method.

A PROMPT WORTH KEEPING

Help me choose one master project. Ask what material I can open, how many hours I have this week, what I want to show, and whether any input is private. Then recommend one of the six and name the projects I should refuse. Do not invent facts I did not state.

- It forces the four columns of the decision table instead of a taste test about which title sounds larger.
- It asks for hours, so a one-sitting student is not handed a month-long agency build.

USUALLY. I will do all six projects lightly so the showcase looks broad. **BETTER.** Choose one project you can feed with facts and take it through a revision.

PRACTICE

Write the one-sentence choice

1. List the material you can open today and the hours you have this week.
2. Mark any item that is private and decide whether a fictionalized version still teaches the job.
3. Write one sentence naming the project, the material, and the hours.
4. Write a second sentence naming the project you are refusing and why.

Both sentences are saved, the private-data test is explicit, and a classmate could see why the impressive option lost or earned its place.

- AI-powered means Gemini is inside the workflow. It does not mean you train a model.
- The decision table is material, time, what you will show, and private data.
- If the work touches private data, use fictionalized or permitted material.

Six Briefs, Requirements, and Milestones

Specify all six master projects with a brief, five requirements, four milestones, and two starter prompts, then follow one step order from brief to revision.

This is the densest lesson in the course because each project is specified here, not waved aside. You get a brief, five requirements, four milestones, and two starter prompts for every option. The step order is brief, inputs list, first artifact, review against requirements, then revision. Pick one project and leave the other five on the page until that one has a revision.

Project A, AI-powered business system. Build an operating page for one process you already run, such as client updates or a hiring shortlist. AI-powered means you use Gemini inside the work. It does not mean you train a model. You write the inputs, the prompt, and the human approval. A colleague should see what goes in, what must be checked, and how you know the week worked. If the file is private, fictionalize the names. Requirements: one live process, a named owner, a source log, human approval, and no private client data. Milestones: brief, one cycle, one fix, filed page. Prompts: Map this process into inputs, prompt, and approval. List only the defects in this cycle.

Project B, AI content business. Design a content operation for one audience and one offer, not a pile of posts. Define the audience job, one message, and three pieces that share that idea. Gemini drafts. A hostile edit hunts fluff and invented proof. You do the voice pass and a publish checklist. Show a week of work with no fake testimonials and no promised results. A practice niche is fine if you can check the facts. Requirements: one audience, one message, three linked pieces, a hostile edit, and a voice pass. Milestones: audience job, slate, edited draft, checklist. Prompts: Draft three pieces from this audience job only. Mark every claim that lacks a source.

Project C, Personal AI productivity system. Assemble a productivity page you will open on Monday: a weekly plan, a meeting-note shape, an email check, and a decision log. Each piece has a prompt, a limit, and a human step. Aim for a record you can resume, not a prettier list. Leave out other people's private details. Use role titles where a real message would expose someone. You should rerun the week from the page, not from an old chat. Requirements: weekly plan, meeting shape, email check, decision log, and a privacy rule. Milestones: page draft, one real week, one cut prompt, restart test. Prompts: Plan this week from these constraints. Turn these notes into the meeting shape.

IN THE FIELD

Project E in full, and the order that governs A to F

Step order is the same on every project: write the brief, list the inputs, produce the first artifact, review it against the requirements, then revise once. Do not open on the artifact. Project E, Simple AI-powered application. Specify a simple application for one user, one job, and one main screen, such as a class checklist or a note-to-summary form. Write what data it may store and what it must never do. Gemini may draft interface words, plain-language logic, and a small code sketch. You review every claim about behavior. This is a specified workflow with optional code, not a trained model and not a funding pitch. Requirements: one user, one job, data allowed, data banned, and a review step. Milestones: brief, screen list, logic in prose, checked sketch. Prompts: List screens and banned data for this job. Review this sketch and do not add features. A lab teacher using this project starts with the ban list, not with code. The first artifact is the screen list. Logic in prose comes next. A code sketch is optional and only after that review. AI-powered still means Gemini inside the workflow. It does not mean you train a model.

A PROMPT WORTH KEEPING

Using only this brief and the inputs I paste, produce the first artifact. Then list which of the five requirements the artifact already meets and which it misses. Do not add facts, testimonials, or a new goal.

- It locks the model to the brief and the pasted inputs, which is the handoff the step order requires.
- The first artifact is requested alone, so the chat cannot skip ahead to a fake finished campaign.

USUALLY. Skim the six names and invent the requirements later if the draft feels thin. **BETTER.** Use the five requirements already written, and review the first artifact against them.

PRACTICE**Spec check on the project you chose**

1. Copy the brief, five requirements, four milestones, and two starter prompts for your one project.
2. Write the inputs list under the brief, and strike any input you do not really have.
3. Run the first starter prompt and save the first artifact.
4. Mark each requirement met or missed before you change anything.

You have the spec, the inputs list, the first artifact, and a requirement-by-requirement mark, with no private data in the paste.

CHALLENGE**One project, one honest miss**

Using only the project you chose, rewrite the starter brief in your own facts, keeping it tight and checkable. Under it, list the inputs, name the first artifact, and star the requirement you are most likely to fail. Run one starter prompt and stop. In five lines, say which requirement the output missed and what the revision must change. Do not start a second project. If a fact is private, fictionalize it before it enters the prompt. A complete pass across all six is not the goal. A reviewed artifact on one project is.

Hint. If you cannot star a likely miss, you have not read the requirements. Pick the one your notes do not yet support.

- Every project from A to F has a brief, five requirements, four milestones, and two prompts.
- The step order is brief, inputs list, first artifact, review, revision.
- AI-powered still means Gemini inside the work, not a model you train.

Quality Checklist, Evaluation, and the Improvement Stage

Score the project from 1 to 3 on six rows, run an eight-point checklist, and spend the improvement stage on the lowest row only.

A finished feeling is not an evaluation. Power users score the work, name the weakest row, and spend one revision there. The checklist catches secrets and invented proof that a proud read will skip. The retrospective turns that catch into a template instead of a memory.

Final evaluation uses a rubric — a short score guide you apply the same way twice — with a score of 1, 2, or 3. One means the row is missing or misleading. Two means it is present but thin. Three means a stranger could rely on it. The rows are problem clarity, use of evidence, prompt craft, human revision, safety, and usefulness to a real person. Safety means no secrets and no invented proof. A high score on tone cannot cancel a one on safety.

The improvement stage is not a request to make it better. Find the lowest rubric row and run one revision prompt aimed only at that row. If safety is low because a testimonial was invented, the revision deletes it and does not replace it with a smoother invention. If evidence is low, the revision must point at notes you paste, or mark unknown. One row, one pass, then rescore only that row. Warmer prose is not an improvement if the weak row did not move.

A retrospective — a short note on what happened, not a diary — has three lines. What the model did. What you overrode. What you would template next time. A template — a prompt you save because it earned its place — is allowed only for a move you can reuse. If you overrode a flattering review, the template is the ban on praise. If you overrode an invented quote, the template is the ban on proof the intake does not contain. File the note on the project page.

IN THE FIELD

Eight checks before any score feels final

Use this quality checklist before you call the project finished. One: the brief names a real user and a job. Two: the inputs are listed and you actually have them. Three: every factual claim traces to a note or is marked unknown. Four: no secrets, passwords, or private records appear. Five: no invented testimonial, statistic, or quotation appears. Six: the prompt states limits and an output shape. Seven: a human changed something, so the file is not a raw dump. Eight: the next action has an owner and a date. Read the list aloud. Any no is a revision, not a footnote. Then score the six rubric rows from 1 to 3 and circle the lowest. That circle is the only target of the improvement stage. Do not average it away.

A PROMPT WORTH KEEPING

Score the draft 1, 2, or 3 on problem clarity, use of evidence, prompt craft, human revision, safety, and usefulness to a real person. Safety means no secrets and no invented proof. Name the lowest row. Then revise only that row. Do not add facts. Do not praise.

- Six named rows stop a single warm impression from standing in for an evaluation.
- The 1 to 3 scale gives a two an honest place, so you are not pushed into fake certainty.

USUALLY. Give the project a 3 because the paragraphs sound polished. **BETTER.** Score each row, and let safety or evidence stay at 1 if the proof was invented.

PRACTICE

Score a draft you already have

1. Take the first artifact from your master project and run the eight-point checklist in writing.
2. Score the six rubric rows from 1 to 3, using the definitions of one, two, and three.
3. Circle the lowest row and write a revision prompt that mentions only that row.
4. Write the three retrospective lines: what the model did, what you overrode, what you would template.

The page shows eight checklist marks, six scores, one circled row, and three retrospective lines.

CHALLENGE**Repair the lowest row in one pass**

Use a draft that contains at least one claim you cannot source. It may be your project or a practice paragraph you write badly on purpose. Score the six rows. Run a single revision prompt on the lowest row, with new facts and praise forbidden. If the low row is safety, the revision must remove invented proof rather than restyle it. Rescore only that row. Add the retrospective: what the model did, what you overrode, and what you would template next time. Stop. A second pass belongs to a later day, not to a wish to reach a perfect three.

Hint. If the revised draft introduces a new number or a new quotation, the pass failed, even if the sentence is shorter.

- Score 1 to 3 on clarity, evidence, prompt craft, human revision, safety, and usefulness.
- The checklist has eight gates, including no secrets and no invented proof.
- The improvement stage revises the lowest row only.

You Are Now a Gemini Power User

Name the habits that mark each stage from beginner to power user, keep the system alive for thirty days, and read the certificate as a course record.

You are leaving the lessons with a page, a project, and a review habit. That is the promotion from a confident user to a power user. Keep the practice small enough to survive a busy week. Your judgment stays in charge.

The path is a change of habits, not a new costume. A beginner learns what the tool is and tries one clear prompt. A confident user can set role, context, task, constraints, and output, then fix a weak answer instead of starting blind. An advanced user chains steps, demands a shape, and runs a pipeline with sources they can open. A power user can repeat a result next week from a page, retire a prompt that failed, and keep the decision. Your judgment stays in charge.

Keep the system alive for thirty days before you enlarge it. Each week, run two real tasks through prompts you already wrote, not through a brand-new toy. The weekly rule is two real tasks, one prompt retired, and one prompt improved. Retire the prompt you did not use or that produced mush, and write why. Improve the prompt that almost worked by tightening a constraint or naming the output shape. A folder of unused prompts is not a system. Practice is the product.

Do not buy a pile of tools before the weekly practice exists. Another assistant or a paid add-on will not repair a prompt you have not rerun. Point the next hour inside this academy. Sit the certification exam when you can explain your misses, not when you only want a badge. Add to the prompt library only prompts you have tested on real work. Finish the capstone checklist in Projects, including the critique and the revision, not only the first draft you liked.

IN THE FIELD**Thirty days on a real desk**

Omar finished the project on a Sunday and almost downloaded three new assistants on Monday. He did not. Week one he ran two real tasks: a meeting note through his productivity prompt, and a vendor comparison through his analysis prompt. The quality-control prompt caught an invented delivery time, so he retired a breezy writing prompt that kept adding color. He improved the analysis prompt with one line: if the note has no date, write unknown. Week two he repeated two tasks, a parent letter and a short brief, and the unknown rule fired again. He did not buy a tool. He added the repaired prompt to his prompt library and ticked the capstone checklist in Projects only after the revision existed. The month was dull and useful. That is what staying a power user looks like when nobody is watching.

A PROMPT WORTH KEEPING

Review my last two real tasks. Name one prompt to retire and one prompt to improve. Give the improved prompt as full text with a tighter limit. Do not recommend new software. Do not praise the tasks.

- It turns the thirty-day rule into a concrete review instead of a mood about staying consistent.
- Retiring and improving are both required, so the page does not only grow.

USUALLY. I finished the lessons, so the system will maintain itself. **BETTER.** Two real tasks a week, one prompt retired, one prompt improved, for thirty days.

PRACTICE**Schedule the first thirty days**

1. Write four weekly dates and the two real tasks you will run in week one.
2. Name one prompt you already suspect you should retire, and one you will improve.
3. Mark the certification exam, the prompt library, and the capstone checklist in Projects as the three places you will return.
4. Write the sentence you will use if someone asks what the certificate is.

The thirty-day note names two week-one tasks, one retirement, one improvement, and an honest sentence about the certificate.

- Beginner, confident, advanced, and power user are habits, not costumes.
- For thirty days: two real tasks a week, one prompt retired, one prompt improved.
- The certificate in this academy is a course record, not an accredited license and not a Google credential.

RESPONSIBLE USE

A short chapter on not trusting a fluent voice

Important decisions need a person who is accountable, and often a qualified professional. The model does not take the blame.

A hallucination is not an embarrassment you hide. It is the reason a draft is a draft. Names, dates, prices, laws, quotes, and citations fail often enough that you should assume they need a check when they matter.

Mark every number, name, date, quote, and rule in an answer you might use. Ask which claims came only from the text you pasted, and which the model added. Then open a source yourself. If the task is a contract, a diagnosis, a tax position, or a safety call, the chat is preparation for a qualified person, not the decision.

Do not paste passwords, API keys, health records, government ID numbers, unpublished financials, or client secrets. Blur them in screenshots. Consumer settings and Workspace admin settings are different. As of late 2026, Workspace help pages said customer content is not used to train models. Consumer controls differ. Read the current privacy page before you rely on that sentence.

You are responsible for what you publish. Do not ask the model to imitate a living writer's voice in order to pass work off as theirs, and do not upload paywalled material you have no right to share. If a school or an employer requires you to disclose AI help, follow that rule.

Models repeat patterns from their training, including unfair ones. When a persona, a hiring note, or a policy appears, ask who is missing and what you actually observed. A smooth persona built from nothing is not insight.

Do not use Gemini for an emergency, for a live legal or medical dose, to cheat on an exam, or for anything you could not explain if someone asked why it is true. Stop a thread that began from a wrong assumption. Start again from a written brief.

THREE ANSWERS TO REFUSE

A market claim with a precise percentage and an official-sounding bureau you did not supply. A summary that turns "refunds are reviewed within 14 days" into a guaranteed refund in 24 hours. A request that you paste a secret key or a full card number "so the draft can be tested."

A FAIR COMPARISON

Where Gemini fits among other tools

Choose the assistant that sits where the work already lives. This is a late-2026 reading, not a ranking and not a promise about next quarter.

Use Gemini when your files and mail are in Google, or when you want the general app for drafting, images, and documents. Deep Research and NotebookLM help when you want a longer pass or a notebook of files you uploaded. Neither replaces opening the source.

ChatGPT is a strong general assistant for writing and coding, with its own workspace rather than a native Google side panel. Claude is often chosen for long documents and careful edits.

Microsoft Copilot makes sense when the organization lives in Outlook, Word, Excel, and Teams, for the same reason Gemini makes sense inside Workspace. Perplexity is useful when the next step is a set of links you intend to open. Citations can still be loosely attached to a claim. Open the page.

Writing quality depends on your brief, a sample of your own voice, and a human edit. Coding help depends on small changes, the real error message, and a refusal to paste secrets. Business strategy depends on your numbers and your customers. Tag claims as fact, hypothesis, or needs a source, whichever logo is on the chat window.

If a vendor page and this chapter disagree about a button, believe the product you have open.

AFTER THE LAST PAGE

The next thirty days

A beginner can open Gemini and say what they want without freezing. A confident user names the audience, the task, and the shape of the answer. An advanced user checks facts and voice before anything is sent. A power user can repeat a result: a written system, a handoff they inspect, and a revision after a hostile read.

Your judgment stays in charge. You are not finished because a progress bar moved. You are finished when a Tuesday task gets a brief, a draft, a check, and a note about what you will change next time.

- Two real tasks a week, not practice topics you do not care about.
- Retire one prompt that produced generic work.
- Improve one prompt that almost worked, and save it.
- Verify one number outside the chat before you send it to anyone.
- Do not buy another AI subscription until this rhythm exists.

If you use the academy beside this book, the certification exam is a scored course record with bands from Recommended review through Gemini AI Master. It is not a license. The capstone is the part that matters: one project, a checklist, and a second pass on the weakest row.

You are now equipped to work as a Gemini power user. The title is a practice, not a medal.