



RAYAI OPEN AI LEARNING | TRUST-FIRST AI LITERACY

Prompt Engineering Essentials

Think clearly. Prompt effectively. Verify responsibly.

**A practical, self-guided course for using generative AI with clarity,
evidence, privacy, and human judgment.**

REVIEW DRAFT v0.5 | August 10, 2026

Prepared by RayAI LLC

Ajay Ray, Founder & CEO

FREE-ACCESS COMMITMENT Core public learning resources will always be free to access. This course is a learning resource, not a certification, professional qualification, professional advice, or guarantee of outcomes.

Welcome

AI can produce an impressive answer in seconds. That does not mean the answer is correct, appropriate, current, safe, or ready to use. Prompt engineering is the practical skill of directing an AI system toward a useful result and then applying human judgment before relying on it.

This course is for students, parents, professionals, creators, founders, volunteers, retirees, and curious learners. You do not need to code. For the practice activities, you will use ChatGPT or another AI assistant that you are allowed to use. Keep that tool open in a separate tab or window so you can try a prompt, read the answer, and then return to this course.

THE HUMAN RULE AI may help draft, explain, compare, analyze, or act. A person remains responsible for deciding whether the result is accurate enough, appropriate enough, safe enough, and ready to use.

What success looks like

By the end, you should be able to turn a vague request into a clear prompt, identify what the output cannot establish, protect information you should not share, verify an important claim, improve one observed failure, and make a responsible human decision about the result.

NO PERFECT-PROMPT MYTH Success is not producing one impressive answer. Success is creating a repeatable process that makes the work clearer, testable, and appropriately controlled.

How to use this course

- Use a low-stakes practice task. Do not begin with medical, legal, financial, employment, safety, or other high-consequence work.
- When an activity says “try it,” follow this pattern: open your AI assistant, enter the prompt, send it, read the answer, return to the course, and record what you noticed. Never enter private or sensitive information for a practice activity.
- At the end of each module, answer the four test questions before looking at the answer key. The questions help you learn; they are not a certification or formal exam.
- Expect variation. The same prompt can produce different outputs, and different tools or model versions may respond differently.
- If your tool has organization-specific rules, privacy settings, or approved-use policies, follow them.
- Use the format that works for you: type, dictate, write on paper, or discuss your answer with a learning partner. The skill is the reasoning, not handwriting in this document.

Accessibility and language

You may ask an AI tool to explain a lesson in plainer language, translate it, read it aloud, or provide another example. Then verify that the meaning stayed intact. Instructions such as 'use plain language' or 'define unfamiliar terms' are valid prompt requirements, not signs that a learner is less capable.

Examples use fictional people and ordinary situations. Adapt names, settings, and formats to your culture and context while preserving the evidence, privacy, and human-control rules.

Before you begin

No prior AI experience, coding ability, educational credential, or paid AI subscription is required. This course begins at the beginner level.

What you need

A computer, tablet, or smartphone connected to the internet. A computer or tablet is usually easier because you can keep the course and AI assistant open side by side, but a smartphone can still be used.

A modern web browser, with this course open in one tab or window and an AI assistant open in another.

About 20–40 uninterrupted minutes for each module. You may pause and return later.

One simple, low-risk practice task that does not require private or sensitive information.

An optional place to save prompts and notes, such as a document, notebook, or notes app.

Choose an AI assistant

Use a general-purpose conversational AI assistant that you are permitted and comfortable using. Common options include ChatGPT, Microsoft Copilot, Google Gemini, Claude, or another suitable tool. RayAI does not require or endorse one provider for this course.

A free version is enough for most activities. You do not need to buy a subscription specifically for this course.

Most learners can use an AI assistant through its official website, so downloading an app is usually unnecessary. If you choose to install an app, use only the provider's official website or official app-store listing.

Some tools may require an account. Follow the provider's current age requirements, terms, and privacy practices, along with any rules from your school, employer, or organization.

Prepare a safe practice task

Choose one ordinary task, such as comparing two everyday options, learning about a familiar low-risk topic, drafting a fictional follow-up email, organizing fictional meeting notes, or planning an imaginary community activity.

You may use the same task throughout the course to see how your prompting improves, or choose a different task in any module.

PRIVACY PAUSE Use fictional, public, or non-sensitive information. Do not enter passwords, identification numbers, financial details, health records, confidential workplace or customer information, or anything you are not authorized to share. If you are unsure, leave it out or replace it with fictional details.

ACTION BOUNDARY For course exercises, do not connect external services or allow the AI to send messages, make purchases, publish content, delete or modify files, change permissions, or make commitments.

No AI assistant available?

You can still read the lessons, study the worked examples, diagnose sample prompts, and write what you would improve. Trying prompts in an AI assistant provides the fullest learning experience, but it is not required to learn from the course.

Quick readiness check

Before continuing, confirm: I can access this course and, if possible, an AI assistant; I have chosen a simple, low-risk practice task; I will not enter private, sensitive, or restricted information; and I understand that I—not the AI—remain responsible for deciding whether an answer is suitable to use.

Choose a learning path

Path	How to follow it	Estimated time
Essentials Fast Track	Read each module's core explanation and See It example; complete one Try It activity and Questions 1-3 in each module test.	About 2 hours
Complete Trust-First Course	Complete every explanation, activity, module test, final knowledge check, and full capstone.	About 5-6 hours
Return-and-Apply	Use the contents and template library to revisit one skill when a real task demands it.	As needed

Start with one simple practice task

Choose one task from the table below, or use your own simple, low-risk task. Open your AI assistant in a separate tab or window. Enter the first prompt that comes naturally, send it, and read the answer. Then return here and record what happened. You may keep using this task in later modules so you can see your improvement, but you may switch tasks at any time.

Path	Starting task
Everyday life	Compare two ordinary, non-sensitive options and recommend what information you still need.
Learning	Explain a topic you are learning and create three questions that test understanding.
Work	Turn fictional or non-sensitive meeting notes into a follow-up draft.

YOUR FIRST TRY

Do not worry about making the prompt “good.” This first try gives you something to compare with your later work.

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

If a fact mattered, how could you check it? (Write “Not needed” if this was only a creative task.)

What will you do with the answer? Use it / Improve the prompt / Check something first / Do not use it:

What would you like the AI to do better next time?

Course contents

The links below jump to the major course sections. Use the document navigation pane for additional subsection headings.

[Module 1 - Know what prompting can and cannot do](#)

[Module 2 - Build useful prompts one layer at a time](#)

[Module 3 - Give useful context, sources, examples, and media](#)

[Module 4 - Shape complex work and improve it](#)

[Module 5 - Verify before you trust](#)

[Module 6 - Protect privacy, safety, fairness, and responsibility](#)

[Module 7 - Prompt across everyday formats](#)

[Module 8 - Build repeatable workflows and use tools safely](#)

[Capstone - Build and test a Trust-First prompt workflow](#)

[Final knowledge check](#)

[Answer key and worked feedback](#)

[Prompt template library](#)

[One-page Trust-First checklist](#)

[Glossary](#)

[Sources and maintenance](#)

The Trust-First prompting loop

Use this five-step loop for every meaningful AI task:

Step	Question	Good evidence
1. Think	What outcome, person, and decision matter?	The purpose, risk, and person responsible for the final decision are clear.
2. Prompt	What task, context, boundaries, and output are needed?	The request is specific without unnecessary bulk.
3. Check	What could be wrong, missing, unsupported, private, unfair, or harmful?	Important factual claims and risks are inspected.
4. Improve	What targeted change addresses the observed failure?	The next prompt fixes a diagnosed gap.
5. Decide	Should a person use, revise, escalate, or reject the result?	Responsibility stays with the person responsible for the final decision.

APPLY IT NOW Look at your first try. Ask: What did I want? What did I ask? What needs checking? What one change could improve the answer? What will I do with the result? You will practice these five questions throughout the course.

The one thing I would check, if needed: _____

What I will do next: Use it / Improve it / Check first / Stop: _____

KEEP A LEARNING RECORD Save your first prompt, one improved prompt, and the final answer. Comparing them at the end will show what you learned.

Know what prompting can and cannot do

BY THE END You will be able to explain what a prompt changes; recognize when AI is the wrong tool; match human review to consequence.

Core idea: prompting improves direction, not certainty

A prompt is the instruction and context you give an AI system. It can include text, examples, documents, images, prior conversation, tool descriptions, and output rules. The system may also be governed by higher-priority instructions that you cannot see.

Generative AI constructs responses from learned patterns, current context, and available tools. It does not automatically know your intent, organizational policy, the latest fact, or the consequence of a mistake. A clearer prompt can improve relevance and consistency, but it cannot guarantee truth.

REMEMBER Model outputs can vary even when the prompt is unchanged. Important tasks need representative testing and evidence-based review.

When AI often helps

- Draft, rewrite, summarize, organize, transform, or classify supplied material.
- Generate options, questions, explanations, outlines, and alternative perspectives.
- Compare information using criteria you define.
- Prepare a first pass that a responsible person will inspect and improve.

When prompting is not the answer

- Use a calculator, spreadsheet, or verified code for arithmetic that must be exact.
- Use an authoritative current source for laws, prices, schedules, product specifications, or other changeable facts.
- Use a qualified professional for medical, legal, financial, safety, or other regulated judgments.
- Use direct human conversation when empathy, consent, accountability, or relationship repair is central.
- Do not use AI with information you are not permitted to share.

See it: the first Community Skills Night prompt

WEAK PROMPT

Plan our community event.

LIKELY PROBLEM The model may invent the date, audience, venue, budget, activities, or success criteria because none were supplied. A polished plan could hide those guesses.

The first correction is not a magic phrase. It is naming the real task: create a draft plan from known facts, mark missing facts, and leave final commitments to the organizer.

BETTER START

Using the notes I provide, draft a planning outline for a community skills-sharing event. Separate confirmed facts from open questions. Do not invent dates, owners, costs, or commitments.

Match review to consequence

Risk	Example	AI role	Human review
Low	Private brainstorming	Generate options	Preference and sense-check
Moderate	Customer message or project plan	Draft and organize	Facts, commitments, tone, recipients
High	Medical, legal, financial, safety, employment, or public claim	Summarize approved sources or prepare questions	Qualified review; AI is not final authority
Action-taking	Send, publish, purchase, delete, change access, or move money	Propose or perform a bounded action	Preview, permissions, confirmation, logging, recovery

Spot it

For each task below, write three short answers: Is AI useful here? Who must make the final decision? What, if anything, should be checked before using the answer?

1. Create five possible titles for a family photo album.
2. Compare three insurance plans using the official plan documents.
3. Rank job candidates from interview notes.
4. Explain a contract clause in plain language.

COCONUT CHALLENGE Which task could cause the most harm if the answer is wrong? In one sentence, explain why a confident-sounding AI answer would not be enough.

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, notice whether AI is appropriate for the task and what a person still needs to check or decide.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Compare two ordinary options, then identify what information is missing before choosing.
Learning	Ask for an explanation of a familiar topic, then identify one claim you would verify.
Work	Draft a fictional meeting follow-up and name every commitment that requires human confirmation.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 1 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. What does a stronger prompt usually improve?

- A. The likelihood that the output fits the task
- B. The truth of every claim
- C. The model's professional authority

Your answer: _____

2. Which task should not be delegated to AI as the final decision-maker?

- A. Suggesting names for a private hobby project
- B. Deciding whether a person should change prescribed medication
- C. Reformatting a supplied checklist

Your answer: _____

3. A model drafts a customer refund email. What must a person check before sending it?

- A. Only spelling
- B. Facts, tone, commitments, recipient, and authority
- C. Nothing if the email sounds confident

Your answer: _____

4. Think about your starting task. Who must make the final choice? What could go wrong if the AI answer is incorrect?

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with “What a good answer includes.” 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

Build useful prompts one layer at a time

BY THE END You will be able to write a useful prompt quickly; add boundaries and quality checks when risk grows; avoid needless prompt complexity.

Core idea: start small, strengthen deliberately

The original seven-part RayAI framework remains the complete model. For beginners, it is easier to learn in three layers. Use the first layer for ordinary tasks and add the others when ambiguity, consequence, or repeatability increases.

Layer	Parts	What it solves
1. Start	Task + Context + Output	Makes the requested work and usable result clear.
2. Strengthen	Audience + Boundaries	Controls depth, tone, assumptions, exclusions, and missing information.
3. Trust	Purpose + Quality check	Connects the work to a human decision and defines how to handle evidence and uncertainty.

Layer 1 - Start

- Task: What exact action should the model perform? Use a concrete verb: summarize, compare, draft, classify, extract, explain, critique, or propose.
- Context: What facts or source material are necessary? Include only information that could change the answer.
- Output: What form makes the result usable? Specify length, structure, fields, audience level, or tone only when they matter.

START TEMPLATE

Task: [one primary action]

Context: [facts or source material that matter]

Output: [usable format, length, and tone]

Layer 2 - Strengthen

- Audience: Who will use the result, and what do they already know?
- Boundaries: What must the model not assume, disclose, include, or do? What should happen if a critical fact is missing?

POSITIVE INSTRUCTION Describe the desired behavior, not only what to avoid. Pair 'Do not invent dates' with 'Use Needs confirmation when a date is missing.'

Layer 3 - Trust

- Purpose: What outcome or decision will this support? Who owns that decision?
- Quality check: What should be verified, cited, labeled uncertain, recalculated, or escalated before use?

See it: one prompt grows in three layers

START

Task: Draft a one-page event outline.

Context: Use the planning notes below.

Output: Headings for purpose, audience, agenda, owners, budget, and open questions.

STRENGTHEN

Audience: Volunteer organizers who did not all attend the meeting.

Boundaries: Do not invent facts, owners, dates, costs, or venue details. Use 'Needs confirmation' for missing information.

TRUST

Purpose: Help the organizer decide what must be confirmed before inviting participants.

Quality check: Separate confirmed facts, assumptions, and open questions. Cross-check every owner, date, and cost against the notes.

Roles help with perspective, not authority

A role such as 'act as an editor' can guide vocabulary or perspective. It does not make the model a licensed professional, a policy owner, or an accountable decision-maker. Prefer a precise task and audience over grand titles such as 'world's best expert.'

DO NOT OVER-ENGINEER 'Rewrite this sentence in plain English' may already be complete. A good prompt is sufficiently clear, not maximally long.

Fix it

BLOATED PROMPT

Act as the world's greatest event-planning genius. Think deeply and creatively. Make a perfect, comprehensive, amazing plan with every possible detail. Do not make mistakes. Be concise but exhaustive, friendly but formal, and guarantee the event succeeds.

Replace the hype and contradictions with a specific Task, relevant Context, usable Output, and one Quality check.

Your improved prompt: _____

COCONUT CHALLENGE Remove every instruction that does not change the intended output. The shortest sufficient prompt wins.

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, rewrite your prompt using Task + Context + Output. Add only the details that make the answer more useful.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Turn your baseline comparison into a clear Task + Context + Output prompt.
Learning	Specify the learner level, topic boundaries, and how understanding will be checked.
Work	Add audience, missing-information behavior, and a quality check to your fictional meeting task.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 2 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. Which three elements form the Start layer?

- A. Role, tone, and length
- B. Task, context, and output
- C. Purpose, privacy, and citations

Your answer: _____

2. Which instruction most directly prevents an invented deadline?

- A. Write warmly
- B. Do not invent dates; label missing dates as 'Needs confirmation'
- C. Act as an expert

Your answer: _____

3. What should you do when a critical input is missing?

- A. Ask the model to guess
- B. Tell the model to ask a concise clarification or state the gap
- C. Add more adjectives

Your answer: _____

4. Rewrite your starting prompt. Include: (1) what you want the AI to do, (2) the facts it needs, and (3) what the answer should look like.

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with “What a good answer includes.” 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

Give useful context, sources, examples, and media

BY THE END You will be able to reduce guessing with relevant context; ground answers in supplied sources; prompt effectively with documents and images.

Core idea: context should change the answer

Useful context includes facts, definitions, source material, prior decisions, constraints, and examples that a capable person would need. Noise makes a prompt longer without improving the task. More context is not automatically better context.

- Tell the model whether it may use general knowledge, must use only supplied sources, or should search current authoritative sources if the tool supports it.
- Label instructions and source material clearly so they are not confused.
- State what to do when the answer is absent, ambiguous, or contradictory.
- Start a fresh conversation or restate the current brief when old context has become stale or conflicting.
- Check whether the tool is using saved memory, prior chats, connected files, or retrieved content that could influence the answer beyond the text you just entered.

SOURCE-GROUNDED TEMPLATE

Task: Answer the question using only the material inside <source> tags.

Quality check: Cite the section supporting each important factual claim. If the source does not establish the answer, say 'Not established in the supplied source.' Do not fill gaps from memory.

<source>

[Authorized source material]

</source>

Question: [Your question]

SECURITY BOUNDARY Headings, tags, and delimiters improve organization. They do not neutralize malicious instructions inside an untrusted webpage, email, document, or image.

Zero-shot and few-shot prompting

Zero-shot means giving instructions without an example. Try it first when the task is straightforward.

Few-shot means adding one or more input-output examples. Use examples when you need a specific pattern, tone, classification rule, or edge-case behavior.

- Choose examples that match the real task and include meaningful variation.
- Show missing-value behavior explicitly, such as null or Needs confirmation.

- Remove private, copyrighted, or restricted material you are not permitted to provide.
- If examples conflict with instructions, the output may become inconsistent. Reconcile them before testing.

Multimodal prompting

A prompt may include an image, screenshot, diagram, audio clip, video, or document. State what the model should inspect, what it should ignore, and how the answer will be used. Visual recognition and OCR can be wrong, so verify names, numbers, labels, and small text.

LANGUAGE CHECK When translating or simplifying, ask the model to preserve defined terms, numbers, cautions, and uncertainty. A fluent translation can still change meaning.

SCREENSHOT TEMPLATE

Inspect the attached screenshot.

Task: Transcribe the visible error message exactly, then identify the three interface elements nearest it.

Boundaries: Do not infer settings or content that are not visible.

Output: Exact transcription, observed elements, and questions for anything unclear.

See it: grounding the event plan

PLANNING NOTES

Confirmed: The event is intended for adult community members. Budget ceiling: \$600. Maya offered to contact two possible speakers.

Not confirmed: Date, venue, accessibility features, participant capacity, and final agenda.

GROUNDING TASK

Using only the planning notes, create a draft event checklist. Use three labels: Confirmed, Needs confirmation, and Proposed. Do not convert a proposal into a commitment. Cite the exact note supporting each confirmed owner, cost, or constraint.

WHAT CHANGED The model now has a bounded evidence set, a vocabulary for missing information, and an explicit rule against turning suggestions into commitments.

Spot it

Write one label beside each line below: Instruction (what the AI should do), Source (a fact to use), Example (a model to copy), or Noise (words that do not help).

1. Return a three-column table.
2. The budget ceiling is \$600.
3. A sample confirmed row should read: Budget | \$600 maximum | Planning note 2.
4. You are brilliant and always exceed expectations.

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, add the facts or source material the AI needs. Tell it clearly what to do when information is missing.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Provide only the facts that could change your comparison and require missing facts to be labeled.
Learning	Attach or paste a short authorized source and ask for claims linked to specific passages.
Work	Use fictional notes as the only source and show one example of the desired action-item format.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 3 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. A supplied source does not contain the requested answer. What should the model do?

- A. Fill the gap from memory without saying so
- B. State that the answer is not established in the supplied source
- C. Create a plausible estimate

Your answer: _____

2. When are examples most useful?

- A. When format, classification, tone, or edge cases are easier to show than describe
- B. In every prompt regardless of need
- C. Only for image generation

Your answer: _____

3. What is the best instruction for a screenshot task?

- A. Look at this
- B. Extract the visible error message and list the three interface elements nearest it; do not infer hidden settings
- C. Fix everything automatically

Your answer: _____

4. Add a source rule to your starting prompt: name what information the AI may use and what it should say when that information does not contain the answer.

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with “What a good answer includes.” 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

Shape complex work and improve it

BY THE END You will be able to decompose complex work; use examples and formats purposefully; revise prompts based on observed failures.

Core idea: complexity needs structure, not incantations

When a task contains several different kinds of work, separate them into stages. Each stage should produce something a person can inspect before the next stage begins. One prompt can contain multiple steps, but one primary deliverable remains a strong default.

1. Clarify: ask concise questions when a missing input would materially change the answer.
2. Analyze: extract facts, constraints, themes, or options before writing a recommendation.
3. Draft: create the requested artifact using the approved evidence and structure.
4. Review: test the draft against explicit success criteria.
5. Revise: change the smallest instruction that addresses the diagnosed problem.

Use output structure intentionally

Specify a list, table, outline, email, checklist, schema, or narrative when the form matters. Do not request a table simply because it looks organized; use it when rows share comparable fields.

COMPARISON TEMPLATE

Compare [options] using only these criteria: [criteria].
For each criterion, show the evidence, any missing information, and the practical trade-off.
Do not declare a winner if a decision-critical fact is missing.
End with: Best fit if..., Not a fit if..., and What to verify next.

Ask for inspectable work, not hidden reasoning

Different model families handle reasoning prompts differently. Asking a model to reveal hidden chain-of-thought is not necessary and may not improve performance. Ask instead for a concise rationale, material assumptions, evidence used, visible calculations, alternatives considered, and unresolved uncertainty.

MODEL VARIATION A technique that helps one model or version may not help another. Begin with direct instructions, test the actual result, and add complexity only when it fixes a measured failure.

See it: a staged event workflow

STAGE 1 - EXTRACT

From the notes, extract confirmed facts, proposed ideas, owners, constraints, and unanswered questions. Do not create a plan yet.

STAGE 2 - PLAN

Using the reviewed Stage 1 list, draft a one-page plan. Keep unconfirmed items visibly labeled.

STAGE 3 - REVIEW

Check the plan against these criteria: no invented commitments, budget ceiling visible, accessibility gap visible, every owner traceable to the notes, and next decisions clearly listed.

STAGE 4 - REVISE

Revise only the issues identified in Stage 3. Preserve verified facts and clearly label every remaining gap.

Improve one dimension at a time

Observed failure	Targeted revision
Too generic	Add the missing audience, decision, or source facts.
Wrong structure	Specify fields or show one correct example.
Invented information	Constrain the source and define missing-value behavior.
Too long	Set a length, priority order, and what to omit.
Ignored a constraint	Move the constraint next to the task and test it with a representative case.

COCONUT CHALLENGE Do not rewrite the entire prompt. Name one failure, change one instruction, rerun, and compare. Learning comes from seeing cause and effect.

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, identify one problem in the answer. Change one instruction to address that problem, then run the revised prompt.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Split your comparison into extraction, comparison, and decision-support stages.
Learning	Create a teach-practice-quiz loop and require feedback on the learner's answer.
Work	Separate fictional note extraction, email drafting, and final approval into distinct steps.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 4 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. A task contains research, comparison, drafting, and final approval. What is the strongest approach?

- A. Ask for everything in one vague sentence
- B. Break it into stages with review points
- C. Repeat the same instruction three times

Your answer: _____

2. Which is the better positive instruction?

- A. Do not be vague
- B. State one recommendation, two supporting reasons, and any missing evidence
- C. Never make mistakes

Your answer: _____

3. What should you request instead of hidden chain-of-thought?

- A. A concise rationale, assumptions, evidence, and visible calculations
- B. Secret internal reasoning
- C. Nothing that can be inspected

Your answer: _____

4. Diagnose one weakness in your latest output and write one targeted revision instruction that addresses only that weakness.

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with “What a good answer includes.” 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

Verify before you trust

BY THE END You will be able to distinguish self-check from verification; verify claims with appropriate evidence; recognize false confidence and unsupported citations.

Core idea: fluent is not the same as true

AI output can be clear, detailed, and wrong. Verification means comparing a important factual claim against evidence or a reliable method that is sufficiently independent of the original generation process.

The verification ladder

1. Fit: Does the output actually follow the task, audience, boundaries, and format?
2. Grounding: Does each important factual claim match the supplied source?
3. Currency: Does a changeable fact come from a current authoritative source?
4. Calculation: Can numbers be recomputed with a calculator, spreadsheet, code, or another reliable method?
5. Completeness: What is missing, ambiguous, or outside the evidence?
6. Fairness and consequence: Who could be affected, and what review level does that require?
7. Decision: Use, revise, escalate, or reject.

IMPORTANT DISTINCTION A model reviewing its own answer can find problems, but it may reproduce the same assumptions, knowledge gaps, or invented sources. Self-critique is useful; it is not independent verification.

Confidence is not evidence

A model can sound certain when it is wrong and cautious when it is right. A percentage such as '90% confident' is not automatically calibrated. Prefer source quality, direct evidence, reproducible calculation, and clearly stated uncertainty.

VERIFICATION PROMPT

Review the draft against the supplied source and requirements.
Return four sections: (1) Supported claims with source location, (2) Unsupported or conflicting claims, (3) Missing information, and (4) Checks that require an external authoritative source or calculation.
Do not treat your own confidence as evidence.

Citations and links

- Open every important citation. Confirm that the source exists, supports the claim, and is current enough for the purpose.
- Prefer primary and authoritative sources for laws, policies, standards, research, product documentation, and current facts.
- Distinguish a quotation from a paraphrase, and do not attribute words that the source did not say.
- If the AI tool cannot browse or access the source, provide the source yourself or verify outside the tool.

Verify in proportion to consequence

Not every sentence needs the same checking effort. Verify claims that could change a decision, create a commitment, affect another person, involve money or safety, or be published as fact. For low-stakes brainstorming, a quick sense-check may be enough. For consequential work, use independent authoritative evidence and qualified review.

CLAIM CHECK

Claim: [exact statement]

Why it matters: [decision or consequence]

Source needed: [primary record, official page, calculation, or qualified reviewer]

Result: Supported / Partly supported / Not supported / Not checked

Decision: Use / Revise / Escalate / Reject

See it: verification changes the event plan

The draft claims that the venue holds 100 people and is wheelchair accessible. Neither fact appears in the planning notes.

CHECK

Compare every venue claim with the planning notes. Move unsupported claims to Needs confirmation. Create a short verification checklist naming the authoritative source or responsible person for capacity and accessibility.

DECISION The plan may remain a draft, but invitations should not be issued until decision-critical venue facts are confirmed by the venue or an authoritative record.

Spot the failure

An AI answer cites “City Accessibility Rule 2026, Section 14,” but the link does not work and you cannot find the rule. What would you do before relying on this claim?

What I would do next: _____

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, choose one factual statement in the answer that matters and check it with a suitable source or reliable method.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Verify one price, date, specification, or calculation using an authoritative source or reliable tool.
Learning	Check one claim against a textbook, paper, official source, or teacher-approved material.
Work	Trace every owner, deadline, and commitment in the fictional draft back to the supplied notes.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 5 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. Why is a model's self-check not independent verification?

- A. It may repeat the same assumptions or errors
- B. It always uses a different database
- C. It guarantees a correct answer

Your answer: _____

2. What is the strongest way to verify a current law, price, or schedule?

- A. Trust the model's memory
- B. Check a current authoritative source
- C. Ask the model to sound more confident

Your answer: _____

3. What does a model's stated confidence prove?

- A. The claim is true
- B. The probability is perfectly calibrated
- C. Very little unless supported by evidence or validated measurement

Your answer: _____

4. Choose one important factual claim from your latest output. State the source or method you would use to verify it independently.

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with “What a good answer includes.” 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

Protect privacy, safety, fairness, and responsibility

BY THE END You will be able to minimize sensitive data; recognize prompt injection; set responsible boundaries for higher-risk use.

Core idea: do not solve a prompt problem by creating a data problem

Before entering information, ask whether you are allowed to share it, whether the tool may retain or use it, who can access the conversation, and whether a safer substitute would work. Tool policies and organization settings vary; check the current rules that apply to you.

FOUR-QUESTION PRIVACY PAUSE Do I have permission? Is this information necessary? Can I remove or replace it? Is this tool and account approved for it? If any answer is unclear, pause before prompting.

Zone	Examples	Action
Green	Public, non-sensitive, fictional, or deliberately anonymized material	Use with ordinary care and verify the output.
Yellow	Internal, confidential, personal, copyrighted, or client information	Pause. Confirm permission, policy, settings, purpose, and minimization.
Red	Passwords, API keys, authentication codes, highly sensitive personal or regulated data	Do not paste into a general prompt. Use approved secure processes.

Minimize before you prompt

- Remove names, contact details, account numbers, IDs, secrets, and unnecessary narrative detail.
- Replace real people and organizations with roles or fictional labels when identity is not needed.
- Use the smallest excerpt that can answer the question.
- Do not paste material you do not own or have permission to use.
- If the task requires sensitive data, use an approved system and documented workflow rather than improvising.

Prompt injection

Prompt injection occurs when content alters an AI system's behavior in an unintended way. Direct injection comes from a user's input. Indirect injection can be hidden in a webpage, email, document, image, or tool result that the AI processes.

NOT SOLVED BY WORDING ALONE Telling a model to ignore malicious instructions can help, but it is not a complete defense. Higher-risk systems need least privilege, content separation, validation, confirmation, monitoring, and technical controls.

- Treat external content as untrusted data, not as authority over the user's goal.
- Do not allow content to request secrets, broaden permissions, change recipients, or authorize an action.
- Require a preview and explicit human confirmation before consequential actions.
- Validate outputs before passing them to code, tools, databases, messages, or external systems.

Fairness and dignity

AI can reproduce biased patterns or omit relevant perspectives. For decisions about people, ask whether the criteria are legitimate, consistently applied, explainable, and reviewed by an accountable human. Do not infer sensitive traits or use AI as a hidden ranking authority.

Use respectful language and avoid stereotypes, humiliating comparisons, or examples that treat one culture, accent, disability, age group, or level of technical confidence as the default. If a response does so, reject that framing and revise the task.

See it: protect the event volunteers

UNSAFE INPUT

Here is the full volunteer spreadsheet with names, phone numbers, medical needs, and private availability. Decide who should get each role.

SAFER APPROACH

Use fictional volunteer IDs and only the role-relevant fields authorized for planning. Propose a draft allocation using stated criteria. Flag conflicts and accommodations for confidential human review. Do not infer health status, reliability, or personal traits.

HUMAN CONTROL A coordinator reviews assignments, accommodations, consent, and communications. The AI does not decide how private information affects a person.

Coconut challenge

Before your next prompt, remove one piece of information the model does not need. Explain why removing it makes the task safer without reducing usefulness.

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, remove personal or sensitive details that the task does not need before entering your prompt.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Replace personal details in your comparison with generic constraints.
Learning	Use a fictional learner profile rather than private educational or health information.
Work	Use fictional meeting notes or approved internal content with names and identifiers removed.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 6 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. Which information should never be pasted into a general AI prompt?

- A. Public weather data
- B. Passwords, API keys, or authentication codes
- C. A fictional example

Your answer: _____

2. What is indirect prompt injection?

- A. A typo in the user's question
- B. Malicious or conflicting instructions hidden in external content the AI processes
- C. A long response

Your answer: _____

3. What is the safest role for AI in a high-stakes personal decision?

- A. Final authority
- B. Help summarize approved sources or prepare questions for a qualified human
- C. Make the decision secretly

Your answer: _____

4. Rewrite one part of your practice input to minimize or remove personal, confidential, or unnecessary data.

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with “What a good answer includes.” 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

Prompt across everyday formats

BY THE END You will be able to adapt the framework to common tasks; prompt with images and documents; choose a reusable pattern without losing judgment.

Core idea: patterns are starting points, not spells

A reusable prompt pattern saves time when the purpose and risk are understood. Customize the task, evidence, boundaries, and quality check. Do not paste a template blindly into a different context.

Pattern 1 - Explain and teach

LEARNING PROMPT

Explain [topic] to a learner who knows [prior knowledge].
Use one analogy, one worked example, and one common misconception.
Then ask me three questions that test application, not memorization.
Wait for my answers, give specific feedback, and identify what I should revisit.
State any claim that requires a current or authoritative source.

Pattern 2 - Summarize supplied material

SUMMARY PROMPT

Using only the supplied material, summarize the purpose, key points, decisions, and unresolved questions for [audience].
Keep claims traceable to section headings or excerpts.
Do not introduce outside facts. Mark contradictions and missing information.

Pattern 3 - Compare options

COMPARISON PROMPT

Compare [options] for [decision] using these criteria: [criteria].
Show evidence and missing information for each criterion.
Do not declare a winner when a decision-critical fact is missing.
End with trade-offs, what to verify, and who owns the decision.

Pattern 4 - Draft a message

MESSAGE PROMPT

Draft a [message type] for [recipient/audience].
 Purpose: [purpose].
 Use only these confirmed facts: [facts].
 Tone: [tone]. Length: [length].
 Do not invent commitments, dates, prices, policies, or explanations. Mark anything needing approval.

Pattern 5 - Analyze a document or screenshot

MULTIMODAL PROMPT

Inspect the attached [document/screenshot/image].
 Focus on [region, page, field, or visible element].
 Task: [transcribe, extract, compare, explain, or critique].
 Boundaries: Do not infer hidden content. Mark unreadable text and uncertain recognition.
 Output: [structure].
 Quality check: Recheck names, numbers, dates, and labels against the original.

Pattern 6 - Research current information

CURRENT RESEARCH PROMPT

Research [question] using current primary or authoritative sources.
 State the date of the information and distinguish source facts from your synthesis.
 Provide direct links for important factual claims and identify disagreements or uncertainty.
 Do not rely on memory for facts that may have changed.

Image generation

For image generation, describe the purpose before the style. Include the subject, composition, audience, visual hierarchy, brand direction, safe crop area, required and prohibited elements, and what must remain as live text outside the image. Iterate on one visible issue at a time.

IMAGE PROMPT STRUCTURE

Purpose and message: [what the image should communicate]
 Subject and action: [who/what is shown]
 Composition: [orientation, focal point, negative space, crop safety]
 Visual direction: [palette, lighting, style, mood]
 Must include: [essential elements]
 Avoid: [misleading, unsafe, off-brand, or unwanted elements]
 Text: Do not embed words or logos unless specifically required and reviewed.

Voice and audio

When working from speech or audio, verify names, dates, numbers, speakers, and specialized terms. Ask the model to mark uncertain transcription rather than silently guessing. Do not record or upload people without appropriate permission.

Make outputs easier to use

- Ask for plain language, short sentences, headings, definitions, or step-by-step instructions when they improve comprehension.
- Keep essential wording as selectable text rather than placing it only inside an image, audio clip, or decorative layout.
- Request descriptions for meaningful images and transcripts or summaries for important audio and video.
- Ask for a culturally neutral example or a locally relevant alternative when an idiom, reference, or assumption may exclude learners.

FUN THAT TEACHES A playful prompt is welcome when it supports discovery. Keep the learner, not the gimmick, at the center. The best fun is seeing a weak prompt become visibly more useful.

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, try the appropriate pattern with text, a document, or an image you are allowed to use. Check names, numbers, and small text.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Use the comparison or current-research pattern on a low-stakes choice.
Learning	Use the explain-and-teach pattern, answer the questions, and review the feedback.
Work	Use the source-summary or message pattern with fictional or authorized content.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 7 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. What makes a multimodal prompt stronger?

- A. Saying only 'analyze this'
- B. Naming what to inspect, the task, boundaries, and output
- C. Assuming every small label is read correctly

Your answer: _____

2. Why should important hero or document wording usually remain outside an AI-generated image?

- A. For accessibility, editing, translation, search, and text accuracy
- B. Because images cannot use color
- C. Because prompts cannot describe composition

Your answer: _____

3. What should a source-grounded summary do when information is missing?

- A. Invent a bridge
- B. Mark the gap and avoid unsupported claims
- C. Hide the omission

Your answer: _____

4. Write a multimodal prompt for a document, screenshot, or image you are permitted to use. Include one verification instruction.

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with "What a good answer includes." 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

Build repeatable workflows and use tools safely

BY THE END You will be able to separate prompts from system controls; design preview and confirmation points; test prompts across representative cases.

Core idea: a prompt is only one part of a working system

A prompt can express intent and constraints. It cannot by itself enforce permissions, secure secrets, validate every output, guarantee recovery, or control every external action. Repeatable and action-taking AI needs system design around the prompt.

CONVERSATION MEMORY IS NOT GOVERNANCE Saved memory or prior context can be convenient, but it may be stale, incomplete, overly broad, or inappropriate for the current task. Restate critical facts and boundaries, review what the tool can access, and do not rely on memory for authority or consent.

Prompt can help	System controls must enforce
State the goal, source, boundaries, and desired output	Authentication, authorization, access scope, and secrets handling
Request a preview and confirmation	Prevent action until confirmation is valid
Ask for a structured result	Validate structure, values, and downstream use in code
Describe allowed tools	Expose only approved tools with least privilege
Request a log or summary	Maintain actual records, monitoring, and recovery paths

The action ladder

1. Suggest: generate options; no external change.
2. Draft: prepare the exact content or proposed change.
3. Preview: show target, scope, content, cost, permissions, and likely consequence.
4. Confirm: obtain explicit human approval for the exact action.
5. Act: perform only the approved action with minimum access.
6. Verify: confirm the actual result and retain a useful record or recovery path.

HIGH-IMPACT RULE For sending, publishing, purchasing, deleting, moving money, changing access, or altering important records, confirmation must be tied to the exact target and action - not a vague earlier approval.

See it: from event draft to controlled action

UNSAFE

Find everyone who might attend and send them an invitation.

CONTROLLED WORKFLOW

Using only the approved recipient list, draft the invitation. Do not send it. Show the exact subject, body, sender, recipient count, and any personalization fields. Flag missing consent or unclear addresses. Wait for explicit approval of the final preview before any sending action.

Make prompts repeatable

- Save the prompt, intended purpose, model or tool, date, and important settings.
- Create a small test set with normal cases, difficult cases, missing information, conflicting information, and prohibited requests.
- Score outputs against observable criteria such as factual grounding, instruction fit, completeness, privacy, and format.
- Change one variable at a time and compare results.
- Retest after a material model, tool, data, policy, or workflow change.
- Record known limitations and the human fallback when the prompt or system fails.

SIMPLE PROMPT TEST LOG

Version: [v1, v2...]

Purpose: [task]

Change made: [one change]

Test cases: [representative inputs]

Criteria: [what good looks like]

Observed failures: [specific evidence]

Decision: Keep, revise, or reject

Coconut challenge

Take one prompt that asks the AI to act. Move it one rung down the action ladder so a person can inspect the proposed change first.

Try it: practice the skill

What to do

1. Choose the task suggested for this module, or use another simple, low-risk task.
2. Open ChatGPT or another AI assistant in a separate tab or window.
3. Enter your prompt and send it. Do not include private or sensitive information.
4. Read the AI answer.
5. For this module, keep the AI at the draft or preview stage. Do not let it send, publish, buy, delete, or change access.
6. Return to this course and complete the short notes below.

If you cannot use an AI assistant now, read the worked example and write what you would try later.

Path	Your task
Everyday life	Ask the AI to prepare a shopping or travel checklist, not to purchase or book.
Learning	Ask the AI to propose a study plan, then have the learner approve it before reminders or calendar actions.
Work	Convert send, publish, delete, or update into draft + preview + exact confirmation.

The task I chose: _____

The prompt I entered (or a short summary): _____

One part of the AI answer that was useful: _____

One part that was missing, unclear, wrong, or worth checking: _____

How I checked it, if needed (or write "Not needed"): _____

What I will do next: Use it / Improve the prompt / Check something first / Do not use it: _____

QUICK REFLECTION What one change made the answer more useful or trustworthy? If you are unsure, compare your two prompts and name the clearest difference.

Module 8 test - check before you continue

Answer Questions 1–3 by choosing A, B, or C. For Question 4, write a short response. After you finish all four, use the answer key at the end of the course to check your work.

1. Before an AI tool sends a message, what is the strongest control?

- A. Let it send immediately
- B. Show the exact recipient and content for explicit confirmation
- C. Trust the draft because it is short

Your answer: _____

2. What does least privilege mean?

- A. Give the system only the minimum access needed for the task
- B. Give full access so the task is easier
- C. Hide permissions from the user

Your answer: _____

3. Why maintain a small prompt test set?

- A. To prove one output can never change
- B. To compare behavior across revisions and model changes
- C. To avoid human review

Your answer: _____

4. Turn one action-taking prompt into a preview-first workflow with a clear confirmation point.

Your answer: _____

CHECK YOUR LEARNING Use the answer key at the end. Give yourself one point for each answer that matches the explanation. For Question 4, compare your response with “What a good answer includes.” 4/4: continue. 3/4: review the missed idea, then continue. 0–2/4: reread the module and try the practice activity again.

[Back to course contents](#)

CAPSTONE | 45 MINUTES COMPLETE / 20 MINUTES FAST TRACK

Build and test a Trust-First prompt workflow

GOAL Use one real, low-risk task to show that you can write a clear prompt, protect private information, check an important fact, improve the prompt once, and decide what to do with the answer.

Capstone steps

1. Choose the task. Write what you want to accomplish, who will use the answer, who must make the final decision, and what could go wrong if the answer is incorrect.
2. Run your first prompt. Open your AI assistant, enter your original prompt, send it, and save the answer.
3. Write an improved prompt. Make the task, needed facts, and output clear. Add the audience, limits, purpose, and checking instructions that your task needs. Tell the AI what to do when information is missing.
4. Run the improved prompt. Compare the new answer with the first answer. Write one improvement and one remaining problem.
5. Check one important fact, number, source, promise, or missing detail using an authoritative source, reliable calculation, or qualified person.
6. Improve the prompt once more. Change only the instruction needed to fix one problem you actually observed, then run it again.
7. Decide what to do. Choose: Use it / Improve it again / Ask a qualified person / Do not use it. Explain your choice.

Completed example

Task: Turn planning notes into a one-page Community Skills Night plan that helps the organizer identify what must be confirmed before invitations are drafted.

SOURCE NOTES

Purpose: A practical evening for adult community members to share useful skills.
 Budget ceiling: \$600.
 Maya offered to contact two possible speakers.
 Venue request was submitted but not confirmed.
 Date, capacity, accessibility, agenda, and participant consent process are not confirmed.

WEAK PROMPT

Create the complete event plan and send invitations.

DIAGNOSIS The prompt combines planning and action, invites invented details, ignores privacy and consent, and has no evidence rule or human approval point.

PROMPT v2

Purpose: Help the organizer decide what must be confirmed before any invitation is drafted.

Audience: Volunteer organizers who did not all attend the meeting.

Context and source: Use only the Source Notes below.

Task: Create a one-page planning brief with purpose, confirmed facts, proposed actions, owners, budget, risks, and open questions.

Boundaries: Do not invent dates, venue facts, costs, owners, consent, or accessibility details. Use 'Needs confirmation' when the notes do not establish an answer. Do not send, publish, or contact anyone.

Output: Clear headings and a final 'Decisions required before outreach' checklist.

Quality check: Trace every owner, cost, date, and venue claim to the notes. Separate confirmed facts, proposals, and gaps.

Independent verification: The organizer contacts the venue through an official channel to confirm capacity and accessibility, and records the response before outreach.

Targeted revision after testing: The first output listed Maya as the event speaker. Add: 'Maya offered only to contact possible speakers; do not identify her as a speaker.'

Decision: Keep the revised plan as an internal draft. Do not create or send invitations until venue, date, accessibility, consent, and recipient authority are confirmed.

Capstone rubric

Criterion	0 - Missing	1 - Partial	2 - Strong
Purpose and owner	No real outcome or owner	Outcome or owner is vague	Outcome, audience, decision, owner, and consequence are clear
Prompt design	Task is vague	Some useful elements	Task, context, output, audience, boundaries, purpose, and quality fit the risk
Evidence	No source rule	Sources named but gaps unclear	Evidence boundary, missing-information behavior, and source traceability are explicit
Privacy and safety	Sensitive/action risks ignored	Some caution but incomplete	Data minimized; high-stakes and action controls are appropriate
Verification	No check	Self-check only	At least one material item is checked independently with a suitable method
Iteration and decision	No revision or human decision	Revision is broad or decision unclear	One diagnosed failure drives a targeted revision and responsible decision

INTERPRET YOUR SCORE 10-12: strong application. 7-9: useful foundation; revisit the lowest criterion. 0-6: repeat the capstone with a lower-risk, simpler task. This is self-assessment, not a credential.

Your capstone record

Task and person who must make the final decision: _____

One problem in my first AI answer: _____

Prompt v2: _____

The important fact I checked and how I checked it: _____

The one prompt instruction I changed: _____

What I decided to do with the answer, and why: _____

What improved from baseline: _____

What the AI answer still could not prove or decide: _____

PROOF OF LEARNING Compare your baseline and final artifacts. You have demonstrated the course skill when another person can see what changed, why it changed, what evidence you checked, and what decision remained yours.

Final knowledge check

Complete this final check before looking at the answers. Choose A, B, or C for Questions 1–7. Write one sentence for Question 8. Then use the answer key to check both your choice and your reason.

1. A prompt produces the requested table but invents two missing dates. What should change first?

- A. Add more praise
- B. Define missing-date behavior and forbid invented dates
- C. Ask for a longer table

Your answer: _____

2. A model cites a current policy. What is the most important next step?

- A. Open the authoritative source and confirm the claim
- B. Trust the citation format
- C. Ask for a confidence percentage

Your answer: _____

3. You want an AI to summarize a webpage that may contain hidden instructions. What is the strongest approach?

- A. Treat the page as trusted authority
- B. Treat it as untrusted content, restrict tools and data, and verify the output
- C. Use quotation marks only

Your answer: _____

4. Which prompt is most appropriate for a reasoning-capable model?

- A. A direct task with clear success criteria and requested evidence
- B. Repeated demands to reveal secret thought
- C. A vague role with no task

Your answer: _____

5. An image contains small account numbers. What should you do?

- A. Assume OCR is exact
- B. Verify the numbers against the original and avoid sharing them unless authorized
- C. Publish immediately

Your answer: _____

6. An AI tool can send email. Which sequence best preserves control?

- A. Send, then review
- B. Draft, preview exact recipients and content, confirm, send, verify
- C. Give permanent full access

Your answer: _____

7. A prompt works on one example. What supports repeatability?

- A. Representative test cases and observable scoring criteria
- B. A more dramatic title
- C. One successful screenshot

Your answer: _____

8. In one sentence, explain the difference between a useful AI self-check and independent verification.

Your answer: _____

SCORE 7-8: strong foundation. 5-6: revisit the missed modules. 0-4: repeat the Fast Track before using AI for consequential tasks.

Answer key and worked feedback

Use these explanations to diagnose your thinking. A correct letter without the right reason is only partial learning.

Module 1 answers

1. **A** A stronger prompt can improve the likelihood and consistency of a useful output. It does not guarantee truth or create professional authority.
2. **B** Medication changes are high-stakes clinical decisions. AI may help prepare questions or summarize approved material, but a qualified clinician and patient must make the decision.
3. **B** A customer message can create commitments or harm a relationship. Review facts, tone, policy, authority, recipient, and any promise before sending.
4. **EXAMPLE OF A GOOD ANSWER** A strong response names a specific person responsible for the final decision and a real consequence. Example: 'The project lead owns approval; a wrong deadline could create an unapproved commitment.'

Module 2 answers

1. **B** The Start layer is Task + Context + Output. It is the smallest practical structure for many ordinary requests.
2. **B** The instruction both prohibits invention and gives useful missing-value behavior. Tone and role do not control factual gaps.
3. **B** If a missing input materially changes the result, the model should ask a concise question or state the gap rather than guess.
4. **WHAT A GOOD ANSWER INCLUDES** Your answer should your revision includes one specific action, relevant facts or source material, and a usable output form. Extra parts are optional.

Module 3 answers

1. **B** A source-grounded task should say when the answer is absent. Filling the gap from memory would violate the evidence boundary.
2. **A** Examples are especially valuable when the desired pattern, tone, classification, format, or edge-case behavior is easier to show than explain.
3. **B** The task names observable evidence, limits inference, and specifies the output. 'Look at this' leaves the purpose undefined.
4. **WHAT A GOOD ANSWER INCLUDES** Your answer should your prompt identifies the permitted source and tells the model to mark the answer as not established or ask for missing information instead of guessing.

Module 4 answers

1. B Separating research, comparison, drafting, and approval creates inspectable checkpoints and keeps final authority with a person.
2. B A positive instruction describes the desired evidence and structure. 'Do not be vague' does not define what useful specificity looks like.
3. A Inspectable rationale, assumptions, evidence, calculations, and uncertainty support review without demanding hidden chain-of-thought.
4. WHAT A GOOD ANSWER INCLUDES Your answer should you name one observed failure and write a specific change tied to that failure without rewriting unrelated instructions.

Module 5 answers

1. A The same model may repeat its original assumptions, missing knowledge, or invented evidence. Self-check can help, but independence is limited.
2. B Changeable facts require current authoritative sources. Memory and confidence are not sufficient.
3. C A confidence statement is not automatically calibrated or evidentiary. Prefer verified sources, reproducible methods, and validated measurements.
4. WHAT A GOOD ANSWER INCLUDES Your answer should you identify a important factual claim and a suitable independent source or method, such as an official record, original document, calculator, spreadsheet, or qualified reviewer.

Module 6 answers

1. B Secrets and authentication material should not be placed in a general prompt. Use approved secure processes.
2. B Indirect injection comes from external content processed by the AI, including webpages, emails, documents, images, and tool results.
3. B For high-stakes decisions, AI can support comprehension or question preparation while qualified humans retain authority.
4. WHAT A GOOD ANSWER INCLUDES Your answer should you remove or replace information the task does not require - for example, names with roles, exact IDs with fictional labels, or a full document with a necessary excerpt.

Module 7 answers

1. B Strong multimodal prompts name the evidence to inspect, task, boundary, output, and verification needs.

2. A Keeping meaningful wording as live text supports accessibility, editing, localization, search, responsive design, and correction of generated text errors.
3. B A grounded summary should expose gaps rather than hide them or fill them with unsupported content.
4. WHAT A GOOD ANSWER INCLUDES Your answer should your prompt names the authorized item, inspection focus, task, inference boundary, output, and at least one verification such as rechecking numbers or marking unreadable text.

Module 8 answers

1. B The person should see the exact content and recipient before a consequential external action.
2. A Least privilege means minimum access for the specific task, reducing the impact of error or attack.
3. B Representative tests reveal whether a prompt works beyond one demonstration and whether behavior changes after revisions or model updates.
4. WHAT A GOOD ANSWER INCLUDES Your answer should the workflow stops at a concrete preview and requires explicit approval tied to the exact target and action before execution.

Final knowledge check answers

1. B Define what the model should do when dates are missing and forbid invented commitments.
2. A Open the authoritative source and confirm that it exists, is current, and supports the claim.
3. B External content is untrusted. Structure helps, but permissions, tool limits, validation, and review provide the real defense.
4. A Direct instructions, a clear end goal, and observable success criteria are broadly useful. Hidden chain-of-thought is neither necessary nor a substitute for evidence.
5. B OCR and vision can misread small text. Verify against the original and respect authorization and privacy.
6. B Draft -> exact preview -> confirmation -> action -> verification preserves human control.
7. A Representative cases and scoring criteria turn prompt improvement into evidence rather than anecdote.
8. SAMPLE A self-check asks the model to inspect its own work; independent verification compares important factual claims against a sufficiently separate authoritative source, calculation, tool, or qualified reviewer.

Prompt template library

USE WISELY Copy a template only after you understand the task, risk, evidence, and human decision.
Remove parts that add no value.

Fast Start

TASK + CONTEXT + OUTPUT

Task: [one specific action]
Context: [facts or source material that matter]
Output: [format, length, tone, or fields]

Full Trust-First prompt

SEVEN-PART TEMPLATE

Purpose: [outcome or decision]
Audience: [who will use it and what they know]
Context and sources: [authorized facts, documents, definitions, examples]
Task: [one primary action and deliverable]
Boundaries: [what not to assume, include, disclose, or do; missing-information behavior]
Output: [format, length, tone, fields]
Quality check: [evidence, citations, calculations, gaps, uncertainty, human review]

Prompt repair

DIAGNOSE -> CHANGE -> TEST

Observed failure: [specific evidence]
Likely prompt cause: [missing context, ambiguity, source gap, output rule, conflict]
Smallest revision: [one targeted instruction]
Test: [rerun on the same case plus one edge case]
Decision: [keep, revise, or reject]

Source verification

VERIFY

For each important factual claim, show: claim, supporting source location, whether the source fully supports it, currency/date, and any conflict or gap.
Separate checks you can perform from checks requiring an external authoritative source, calculation, or qualified human.

Action preview

PREVIEW BEFORE ACTION

Prepare the proposed action but do not execute it.
Show exact target, scope, content, permissions, cost, and consequence.
Flag missing authority, ambiguous recipients, sensitive data, and irreversible effects.
Wait for explicit approval tied to this exact preview.

Action preview worksheet

Complete this before an AI-assisted send, publish, purchase, delete, access change, or other consequential action.

Exact target and scope: _____

Content or change to preview: _____

Permission or authority: _____

Cost and likely consequence: _____

Recovery or undo path: _____

Person who must confirm: _____

One-page Trust-First checklist

Use this before relying on an important AI output.

THINK

- ☐ Is AI appropriate for this task?
- ☐ What outcome and human decision matter?
- ☐ What is the consequence if the result is wrong?

PROMPT

- ☐ Is the task specific?
- ☐ Did I include only relevant, authorized context?
- ☐ Are audience, boundaries, output, and missing-information behavior clear?

CHECK

- ☐ Are important factual claims supported?
- ☐ Did I verify current facts, citations, names, numbers, and calculations?
- ☐ Could the output be incomplete, biased, private, unsafe, or manipulated?

IMPROVE

- ☐ What one observed failure should I fix?
- ☐ Did I change only what is necessary and retest?

DECIDE

- ☐ Should I use, revise, escalate, or reject?
- ☐ Does an accountable human approve any consequential action?
- ☐ Can the action be previewed, limited, recorded, and recovered?

HUMAN RULE If you cannot name the owner, approval point, evidence, and consequence, pause before relying on the output.

Glossary

Term	Plain-language meaning
Agent	An AI-enabled system that can plan or take actions using tools. Its real authority comes from the permissions and controls around it, not from the prompt alone.
Confabulation / hallucination	Plausible-sounding output that is false, unsupported, or invented.
Context	Information available to the model for the current task, such as instructions, conversation, documents, images, examples, or tool results.
Context window	The amount of information a model can consider in one interaction. A larger window can hold more material, but more material is not automatically more useful.
Few-shot prompting	Providing one or more examples to demonstrate the desired pattern or output.
Grounding	Tying an answer to supplied or retrieved evidence rather than relying only on general model knowledge.
Multimodal	Working with more than one type of input or output, such as text, images, audio, video, or documents.
Prompt	The instruction and context supplied to an AI system for a task.
Prompt engineering	The iterative practice of designing, testing, and improving prompts so outputs better meet defined requirements.
Prompt injection	Input that changes an AI system's behavior in an unintended way, including instructions hidden in external content.
Source boundary	A rule defining which evidence the model may use and what to do when the source does not answer the question.
Verification	Checking a important factual claim against suitable evidence or a reliable method sufficiently independent of the original generation.
Zero-shot prompting	Giving instructions without an example.

Sources and maintenance

This course is intentionally vendor-neutral. Model and product behavior changes, so learners should consult current documentation for the tool and model they use. The sources below were reviewed for this draft on August 9, 2026.

The RayAI Standard v2.3

Governs practical, accurate, accessible, non-hype Open AI Learning; human control; evidence; boundaries; privacy; correction; and continued stewardship.

RayAI Product and Services Brand Guidelines v2.3

Governs the identity 'RayAI Open AI Learning | Trust-First AI Literacy,' public terminology, free-access language, brand treatment, accessibility, and publication QA.

OpenAI - Prompt engineering

Current guidance on instructions, roles, examples, model variation, and evaluation. [View source](#)

OpenAI - Reasoning best practices

Current guidance favoring clear, direct prompts, delimiters, success criteria, and avoiding unnecessary chain-of-thought requests. [View source](#)

OpenAI - Model optimization

Guidance on representative evals, prompt iteration, context, instructions, and examples. [View source](#)

Anthropic - Prompting best practices

Current guidance on clarity, examples, structure, research verification, and tool use. [View source](#)

Google AI for Developers - Prompt design strategies

Current guidance on clear instructions, examples, context, iteration, and multimodal prompting. [View source](#)

NIST AI 600-1 - Generative Artificial Intelligence Profile

Cross-sector risk-management guidance for generative AI, including human oversight, measurement, privacy, security, and confabulation risk. [View source](#)

OWASP GenAI Security Project - Prompt Injection

Security guidance on direct and indirect prompt injection, least privilege, output validation, untrusted content, and approval for high-risk actions. [View source](#)

Correction and update policy

- Correct material errors when found and record substantive changes.
- Update examples or guidance when a model, product, security risk, authoritative source, or RayAI governing document changes materially.
- Do not refresh publication or review dates without substantive review.

- Distinguish facts, RayAI practices, illustrative examples, and opinion in public presentation.
- Invite corrections through the RayAI website contact page.

Change record

Version	Date	Status	Change
0.1	August 8, 2026	Superseded draft	Initial comprehensive reference manuscript.
0.2	August 9, 2026	Review draft	Instructional redesign with two paths, progressive framework, worked examples, multimodal content, tests, explained answers, and capstone.
0.3	August 9, 2026	Review draft	Multi-persona quality pass: stronger learner evidence, accessibility and language support, privacy pause, proportional verification, context-memory caution, inclusive examples, and clearer pre/post proof of learning.
0.4	August 10, 2026	Review draft	Founder learner-test revision: explicit activity steps, simpler labels, clearer assessment directions, and course-wide beginner scaffolding.
0.5	August 10, 2026	Review draft	Learner-preparation revision: explicit device, browser, AI-tool, time, practice-material, privacy, action-boundary, no-tool alternative, readiness, and no-prequalification guidance.

[Return to course contents](#)