

AI Workshop Hirst

Prompt Engineering Beginner to Intermediate Study Material

What this report gives you

A complete learning path for people who want to move from basic prompting to structured, repeatable prompt design. It keeps the presentation focus on Prompt Anatomy, Prompting Techniques, Live Exercises, and Prompt Debugging, while expanding each idea into practical notes, templates, and answer keys.

Skill level

Beginner + Intermediate.

The guide avoids highly technical API topics and focuses on everyday prompting for writing, analysis, content creation, image prompts, business review, and structured problem solving.

Use this as a study guide, workshop handout, or self-practice workbook.

1. How to Use This Guide

This updated report turns the source presentation into a detailed study material. The original deck introduces prompt engineering, explains why prompting matters, presents live activities, teaches prompt anatomy, demonstrates prompting techniques, and closes with prompt debugging. This guide follows the same teaching flow, but expands each slide into practical lessons and reusable frameworks.

Learning outcomes

- Understand what a prompt is and how it influences an AI model response.
- Use the five-part prompt anatomy: Role, Context, Task, Format, and Examples.
- Apply beginner techniques such as zero-shot, role prompting, and clear formatting.
- Apply intermediate techniques such as few-shot prompting, constraint-based prompting, structured reasoning, and iterative refinement.
- Diagnose bad outputs by identifying the bug in the prompt instead of blaming the model immediately.
- Convert weak prompts into better and best prompts using practical exercise patterns.

Suggested study plan

Session	Focus	Practice
1	Definition, mental model, and why prompting matters.	Rewrite 3 vague prompts into specific prompts.
2	Prompt anatomy: Role, Context, Task, Format, Examples.	Build a reusable prompt template for your own work.
3	Core techniques: zero-shot, few-shot, role, constraints.	Complete the coffee ad and beach image exercises.
4	Intermediate workflow: structured reasoning and refinement.	Improve an output through 3 rounds of revision.
5	Debugging prompts and the skincare tagline case.	Diagnose 5 flawed prompts and patch only the cause.

Core mindset

Prompt engineering is not about memorizing magic phrases. It is about giving the model the right job, the right context, the right boundaries, and the right output shape. Good prompting is structured communication.

2. What Prompt Engineering Means

Prompt engineering is the practice of designing effective inputs to guide AI models toward useful outputs. A prompt can be a question, instruction, role description, document, example, constraint, desired format, or a combination of these. In simple terms: you are telling the model what job to do, what information to use, and what the result should look like.

The basic flow

Prompt -> AI Model -> Response

The prompt is the input. The AI model interprets the input using the conversation context and its learned patterns. The response is the generated output. When the prompt is vague, the model fills gaps with assumptions. When the prompt is clear, the model has fewer gaps to fill.

What a prompt can include

- **Instruction:** the action you want: summarize, compare, draft, critique, generate, classify, translate, explain.
- **Context:** background information: audience, situation, goal, product, data, constraints, risks.
- **Role:** the viewpoint or expertise: teacher, copywriter, sales manager, reviewer, interviewer.
- **Format:** the output shape: table, checklist, bullets, email, JSON, 3 variations, 100 words.
- **Examples:** samples that show tone, structure, quality level, and expected pattern.
- **Rules:** things to include, avoid, prioritize, verify, or flag as uncertain.

Prompting is not mind reading

A common beginner mistake is assuming the model understands the unstated goal. For example, "write an ad" does not define the platform, audience, offer, brand voice, word limit, or call to action. The model will still answer, but the answer will reflect generic assumptions. Strong prompts reduce those assumptions.

Weak prompt

Write an ad for a new coffee brand.

Improved direction

Write 3 Facebook ad variations for a new coffee brand aimed at busy professionals aged 28-40. Each variation must include Hook, Body, and CTA. Keep each under 60 words. Tone: confident, warm, never pushy. Avoid the words premium and perfect.

3. Why Prompting Matters

The presentation highlights four reasons prompting matters: better results, efficiency, control, and innovation. These are not abstract benefits. They directly affect the quality of daily work, especially when using AI for writing, analysis, planning, learning, and creative tasks.

Benefit	What it means	Example
Better results	Clear prompts produce outputs that are more accurate, relevant, and usable.	A prompt that defines audience and format creates a better LinkedIn post than "write something about AI."
Efficiency	Good prompts reduce the number of back-and-forth corrections.	Instead of asking for a report, then asking for bullets, then asking for shorter text, specify the final shape upfront.
Control	Prompts steer tone, structure, detail level, and boundaries.	A legal update, social media caption, and student explanation require very different registers.
Innovation	Prompting helps explore ideas, alternatives, simulations, and combinations quickly.	Generate 20 campaign angles, cluster them, score them, and turn the best 3 into drafts.

The cost of vague prompting

- You spend more time editing because the first answer is too generic.
- The model may choose the wrong audience, tone, format, or level of detail.
- Important constraints may be ignored because they were never provided.
- Outputs become inconsistent across repeated tasks because each prompt is shaped differently.
- You may accept a confident answer without checking assumptions or missing information.

The practical goal

The goal is not to write long prompts every time. The goal is to write prompts that are complete enough for the task. A simple request may need one sentence. A high-value business analysis may need role, context, data, output sections, assumptions, and a quality checklist.

Prompt quality test

After writing a prompt, ask: If someone else read only this prompt, would they know exactly what to produce, for whom, in what style, in what format, and with what restrictions? If not, add the missing part.

4. Prompt Anatomy: The Five Building Blocks

The presentation uses a simple and powerful anatomy: Role, Context, Task, Format, and Examples. This is the central framework for the entire guide. You do not need all five in every prompt, but you should check all five before important tasks.

Part	Purpose	Beginner question	Example
Role	Sets the model viewpoint, vocabulary, and expertise level.	Who should the model act as?	You are a senior copywriter.
Context	Provides situation, audience, background, constraints, and goals.	What does the model need to know?	Audience: busy professionals aged 28-40. Platform: Facebook.
Task	Defines the exact deliverable.	What should be created or solved?	Write 3 ad variations.
Format	Controls output shape, sections, length, and tone.	What should the answer look like?	Use Hook, Body, CTA. Keep each under 60 words.
Examples	Shows the pattern, tone, or quality standard to imitate.	Can I show what good looks like?	Example 1: Hook... Body... CTA...

A complete prompt formula

Formula
You are [ROLE]. Context: [AUDIENCE / SITUATION / BACKGROUND / CONSTRAINTS]. Task: [SPECIFIC DELIVERABLE]. Format: [SECTIONS / LENGTH / STRUCTURE / TONE]. Examples: [OPTIONAL FEW-SHOT EXAMPLES]. Rules: [MUST INCLUDE / MUST AVOID / QUALITY BAR].

When to use each part

- Use **Role** when expertise, voice, or professional judgment matters.
- Use **Context** whenever the model could otherwise guess the wrong background.
- Use **Task** always. This is the one part you should never leave vague.
- Use **Format** when you want a table, sections, word limit, tone, or number of outputs.
- Use **Examples** when tone, structure, or style is difficult to describe in abstract terms.

5. Anatomy Deep Dive: Role and Context

Role: tell the model who it is

A role changes the lens through which the model approaches the task. "You are a teacher" creates a different explanation than "You are a senior financial analyst" or "You are a skeptical editor." The role should match the expertise required by the task. A generic role like "helpful assistant" is often too weak for specialized work.

Weak role	Better role	Why better
You are helpful.	You are a senior copywriter for consumer brands.	The second role implies audience awareness, persuasion, and concise messaging.
You are smart.	You are a sales manager reviewing quarterly performance.	The second role matches business analysis and action planning.
You are creative.	You are an art director creating image prompts for travel blogs.	The second role guides visual language, style, and composition.

Context: give the model the situation

Context is the information the model needs before it can make good choices. It includes the client, audience, platform, product, constraints, brand voice, available data, risks, and success criteria. In the presentation notes, the context section is emphasized because the model does not automatically know the client, audience, or constraints unless you include them.

Context brief template

```
Project: [What is being created or analyzed]
Audience: [Who will read/use it]
Goal: [What success looks like]
Background: [Important facts or data]
Constraints: [Length, tone, platform, legal, brand, budget, time]
Avoid: [Words, claims, formats, assumptions, or risks]
```

Common mistakes

- Assigning a role that does not match the job, such as asking a "friendly assistant" to write a technical audit.
- Providing context after the model has already produced the wrong output instead of upfront.
- Giving irrelevant context that distracts from the task.
- Not defining the audience, which leads to the wrong vocabulary and depth.
- Using brand words like "premium" or "innovative" without explaining what they mean in practice.

6. Anatomy Deep Dive: Task, Format, and Examples

Task: define the deliverable

The task is the action the model must perform. A strong task uses a clear verb and defines the expected result. Weak tasks are broad: "help me with sales" or "make this better." Strong tasks are specific: "Analyze this Q3 sales report and identify the top 5 revenue drivers, 3 risks, and 4 actions for Q4."

Format: anchor the output shape

Without a format anchor, the model chooses a structure. That may be fine for casual questions, but it is risky for business, study, and workshop use. Format can specify sections, headings, number of variations, table columns, word count, tone, and order of information.

Useful format instructions

Return the answer as:

1. Executive summary - 80 words maximum
2. Key insights - 5 bullet points
3. Risks - table with Risk / Evidence / Impact / Action
4. Next steps - prioritized list

Tone: clear, direct, no hype.

Examples: show what good looks like

Examples are optional but powerful. In few-shot prompting, you provide 1-3 examples before the real task. The model uses them as anchors for tone, structure, and expected quality. This is especially useful for emails, taglines, summaries, image prompts, product descriptions, and any task where style is hard to explain.

Use examples when...	Why examples help
You need a specific tone.	Examples show the rhythm and register better than adjectives like "professional" or "friendly."
You need consistent structure.	The model can mirror a pattern like Hook + Body + CTA or Problem + Insight + Action.
You need to avoid generic output.	Strong examples raise the quality bar and reduce default wording.
You are training a repeatable workflow.	Examples make team outputs more consistent across users.

Anatomy checklist

Before sending a high-value prompt, check: Did I define the role? Did I provide enough context? Is the task specific? Did I anchor the format? Would examples improve the result?

7. Technique: Zero-Shot Prompting

Zero-shot prompting means asking the model to complete a task without giving examples. It is the simplest technique and is useful when the task is common, the desired style is easy to describe, or you need a quick first draft. Zero-shot does not mean vague. The best zero-shot prompts are often highly detailed instructions with no examples.

When zero-shot works well

- Simple explanations: "Explain prompt engineering to a beginner in 5 bullets."
- Common writing tasks: "Draft a polite follow-up email after a meeting."
- Structured outputs where the format is clear: "Compare option A and B in a table."
- First drafts where you plan to refine later.

Weak -> better -> best pattern

Level	Prompt	What changed
Weak	Write an ad for a new coffee brand.	Too open: no audience, platform, tone, length, or structure.
Better	Write a Facebook ad for a new coffee brand for busy professionals aged 28-40. Focus on morning routine. Keep it under 60 words.	Adds platform, audience, angle, and length.
Best	You are a senior copywriter. Write 3 Facebook ad variations for a coffee brand called Rise & Roast. Audience: busy professionals aged 28-40. Tone: confident, warm, never salesy. Each variation must include Hook, Body, CTA. Avoid the words premium and perfect. Do not use "Shop Now" as the CTA.	Adds role, brand, variations, structure, tone, banned words, and CTA constraint.

Zero-shot quality checklist

- Give a specific deliverable, not a broad topic.
- Add audience and purpose if the output will be read by someone else.
- Add format instructions so the model does not choose the structure for you.
- Add constraints when generic defaults are likely.
- Ask for multiple variations when you want options rather than one average answer.

Key idea

Zero-shot is not the beginner version of prompting; vague zero-shot is beginner. Detailed zero-shot is often enough for many professional tasks.

8. Technique: Few-Shot Prompting

Few-shot prompting means giving the model a few examples before asking for the real output. The source presentation describes examples as optional but powerful because they anchor tone, structure, and register. Few-shot prompting is especially useful when "make it professional" is not specific enough.

How few-shot prompting works

- Example 1 shows the desired pattern.
- Example 2 confirms that the pattern is intentional, not accidental.
- The actual task follows the examples.
- The model mirrors structure, tone, level of detail, and sometimes vocabulary style.

Few-shot structure

Instruction: Rewrite the input in a measured, professional tone.

Example 1

Input: We cannot finish this today.

Output: We will need additional time to complete this accurately.

Example 2

Input: You forgot to send the file.

Output: Please share the file when convenient so we can proceed.

Now rewrite:

Input: This draft is confusing and needs work.

Beach image exercise: few-shot for visual prompts

For image generation tasks, examples help define visual style. The presentation activity asks participants to create a beach image for a travel blog. A weak prompt only names the subject. A strong prompt describes theme, subject, details, style, composition, and example references.

Best-style beach image prompt

Create a feature image for a travel blog article about slow beach vacations.

Theme: calm, spacious, restorative.

Subject: one quiet beach scene at sunrise with a small wooden boat near the shore.

Supporting details: soft waves, warm early light, clean sand, distant palm silhouettes, no crowds.

Style: editorial travel photography, natural colors, realistic, wide composition.

Example style anchors:

1. Minimal travel magazine cover with open negative space.
2. Warm documentary beach photo with soft morning haze.

Few-shot rule

Use examples when words like professional, beautiful, premium, fun, emotional, or modern could be interpreted in many ways.

9. Technique: Role-Based Prompting

Role prompting assigns the model a specific professional perspective. The presentation uses a sales manager exercise: analyze a third-quarter sales report. This is a good role-based task because the desired answer should not be a generic summary; it should sound like business judgment from someone responsible for performance.

Why role matters in analysis

- A sales manager looks for revenue drivers, pipeline issues, team performance, regional differences, and next-quarter actions.
- A finance analyst may focus on margins, forecasts, and variance.
- A marketing strategist may focus on lead quality, campaigns, and conversion rates.
- The same data can produce different insights depending on the role.

Sales manager analysis prompt

You are a Sales Manager reviewing the Q3 sales report.
Goal: identify what happened, why it matters, and what the team should do in Q4.
Context: [paste Q3 sales data, targets, regions, products, pipeline notes].
Task: analyze the report and produce practical business insights.
Output format:
1. Executive summary - 5 lines maximum
2. Key insights - 5 bullets with evidence
3. Risks or gaps - table with Risk / Evidence / Impact / Action
4. Q4 priorities - top 5 actions in order
5. Questions to ask the team
Negative instruction: do not invent numbers not present in the report; mark missing data clearly.

Role prompt quality checks

- Does the role match the task expertise level?
- Does the role imply useful judgment, not just style?
- Have you provided the data or context the role would need?
- Have you asked for evidence-based insights rather than broad opinions?
- Have you included negative instructions such as "do not invent numbers"?

Intermediate move

Combine role prompting with format and constraints. Role alone can improve tone, but role + context + output sections creates a much stronger result.

10. Technique: Constraint-Based Prompting

Constraint-based prompting tells the model what boundaries to follow. Constraints can be positive, such as required sections, or negative, such as banned words. The coffee ad activity uses banned words like "premium" and "perfect" to prevent generic advertising language. This is a powerful way to force fresher outputs.

Types of constraints

Constraint type	Example	Why it helps
Length	Keep each ad under 60 words.	Prevents rambling and makes output platform-ready.
Structure	Use Hook, Body, CTA.	Makes multiple outputs easy to compare.
Tone	Confident, warm, never salesy.	Controls voice and audience perception.
Banned words	Avoid premium, perfect, best, innovative.	Stops default clichés and pushes more original wording.
Required content	Mention morning routine and low-acid blend.	Ensures key selling points appear.
Evidence rules	Use only the provided report data.	Reduces unsupported claims in analysis tasks.

Positive and negative instructions

Positive instructions say what to do. Negative instructions say what not to do. Both are useful. A common mistake is using only negative instructions, which can make the model focus on avoiding errors without knowing what good looks like. Pair negative constraints with a clear positive target.

Balanced constraints

- Do:
- Write in plain language for first-time buyers.
 - Use one concrete benefit in every variation.
 - End with a calm CTA.
- Do not:
- Use the words premium, perfect, best, or revolutionary.
 - Make health claims.
 - Use the CTA "Shop Now".

Constraint debugging question

Ask this after a weak output

What default behavior did the model fall into? Then write a constraint that blocks that default. If the output is too long, add length. If it is generic, ban generic words. If it is unstructured, add sections. If it invents, restrict evidence.

11. Technique: Structured Reasoning

The presentation lists Chain of Thought as a prompting technique and notes that step-by-step reasoning can improve problem solving. For practical use, ask the model for structured reasoning, assumptions, checks, or a concise rationale. This makes the output easier to audit without requiring a long hidden reasoning transcript.

Use structured reasoning when the task needs judgment

- Comparing options or making recommendations.
- Analyzing a report and deciding priorities.
- Explaining why a conclusion follows from given evidence.
- Checking a draft for risks, gaps, or contradictions.
- Teaching a concept step by step to a learner.

Structured reasoning prompt

Analyze the situation using this structure:

1. Relevant facts from the input
2. Assumptions you are making
3. Possible interpretations
4. Recommendation
5. Confidence level and what information would change your answer

Keep the reasoning concise and evidence-based.

High-stakes caution

For medical, legal, financial, or safety-related tasks, AI output should not replace professional judgment. Use prompts to organize information, identify questions, draft summaries, or surface red flags, but verify with a qualified professional or trusted source before acting.

Safer triage-style prompt for learning

You are helping organize a patient note for discussion with a clinician.

Input: [symptoms, duration, history, medications].

Task: summarize the information, list potential red flags that should be escalated, and identify missing questions.

Rules: do not diagnose; do not recommend treatment; advise urgent care if severe or concerning symptoms are present.

Output: Facts / Red flags / Missing information / Suggested next step.

Intermediate move

Ask the model to show evidence and assumptions, not just conclusions. This makes the answer easier to review and improves trust.

12. Technique: Iterative Refinement

Iterative refinement is the meta-skill of prompting. Instead of expecting one perfect answer, you use the first response as feedback. The presentation describes it as starting rough and improving each round. This is how professionals often use AI: generate, diagnose, refine, and finalize.

The refinement loop

Step	Action	Example refinement instruction
1. Generate	Ask for the first version using a clear prompt.	Draft 3 ad variations using Hook, Body, CTA.
2. Evaluate	Compare the output to the goal.	Which version sounds least generic and why?
3. Diagnose	Identify the specific defect.	The tone is too salesy and the CTA is predictable.
4. Patch	Change only the part that caused the defect.	Rewrite with a calmer CTA and remove hype words.
5. Finalize	Ask for polish and quality control.	Create final version; check word count and banned words.

Useful refinement commands

- **Make it more specific:** Add concrete details, examples, or evidence.
- **Make it tighter:** Reduce length without losing meaning.
- **Change the tone:** More direct, warmer, more executive, more beginner-friendly.
- **Change the structure:** Convert paragraphs to table, checklist, or sections.
- **Increase originality:** Avoid clichés and generate less obvious alternatives.
- **Audit the output:** Check against the constraints and list any violations.

Refinement prompt template

```
Review your previous answer against these criteria:
- Audience fit
- Tone
- Specificity
- Format compliance
- Originality
- Missing information
Then rewrite it, changing only the parts that fail the criteria.
```

Key lesson

Do not restart from scratch after every weak output. Diagnose the reason it failed, patch that part of the prompt, and run the next version.

13. Live Exercise 1: Coffee Ad

This exercise demonstrates zero-shot prompting and the weak -> better -> best upgrade pattern. It is simple enough for beginners but rich enough to teach role, context, format, tone, and constraints.

Task

Prompt challenge

Write an ad for a new coffee brand. First write the weak version, then improve it twice. The goal is to see how each added instruction changes the quality of output.

Three prompt levels

Weak

Write an ad for a new coffee brand.

Better

Write a Facebook ad for a new coffee brand for busy professionals aged 28-40. The angle is that coffee is part of a focused morning routine. Keep the ad under 60 words.

Best

You are a senior copywriter. Write 3 Facebook ad variations for a new coffee brand called Rise & Roast. Audience: busy professionals aged 28-40. Goal: make the brand feel like a confident morning ritual, not a luxury product. Format for each variation: Hook / Body / CTA. Tone: confident, warm, never salesy. Constraints: each variation under 60 words; avoid the words premium and perfect; do not use "Shop Now" as the CTA.

Sample output from the best prompt

Variation	Hook	Body	CTA
1	Start before the noise.	Rise & Roast is made for the quiet minutes before the day asks for everything.	Build your morning ritual.
2	Coffee that keeps pace with ambition.	A smooth daily cup for focused mornings, clear plans, and work that matters.	Make tomorrow easier.
3	Your first decision should be simple.	Rise & Roast brings steady flavor to busy mornings without the drama.	Meet your new morning habit.

Debrief questions

- What did the weak prompt leave open?
- Which instruction most improved the answer: role, audience, format, tone, or banned words?
- What did the banned-word rule force the model to avoid?
- How would the prompt change for Instagram, email, or a billboard?

14. Live Exercise 2: Beach Image Prompt

This exercise demonstrates few-shot prompting and visual specificity. Image prompts need different details than text prompts. Instead of only describing the subject, you describe scene, mood, composition, style, lighting, and exclusions.

Weak prompt

Weak

Create beach image for travel blog.

This prompt leaves too much open. The model can choose any beach, any style, any time of day, any subject, and any mood. The result may be attractive but not useful for the blog objective.

Better prompt

Better

Create a realistic beach image for a travel blog article about relaxing weekend getaways. Show a calm beach at sunrise with soft waves, clean sand, and warm light. Use a wide composition suitable for a blog header.

Best prompt with few-shot anchors

Best

Create a feature image for a travel blog article titled "The Slow Beach Weekend."
 Theme: calm, spacious, restorative.
 Subject: a quiet beach at sunrise with one small wooden boat near the shoreline.
 Supporting details: soft waves, clean sand, distant palm silhouettes, warm early light, no crowds.
 Style: realistic editorial travel photography, natural colors, wide composition, open space for title text.
 Examples to mirror:
 1. A minimal travel magazine cover with one clear subject and lots of negative space.
 2. A warm documentary-style beach photo with soft morning haze.
 Avoid: crowded resort scene, neon colors, cartoon style, heavy filters.

Visual prompt checklist

Element	Question to answer
Subject	What should be the main object or scene?
Setting	Where and when is it happening?
Mood	What should the image feel like?
Style	Photo, illustration, cinematic, editorial, minimal, realistic?
Composition	Close-up, wide, centered, negative space, header-friendly?
Details to include	Objects, lighting, colors, environment, textures.
Details to avoid	Crowds, text, logos, unrealistic elements, wrong style.

15. Live Exercise 3: Sales Manager Analysis

This exercise demonstrates role-based prompting, structured analysis, evidence rules, and negative instructions. The goal is to convert a vague analysis request into a business-ready management prompt.

Weak prompt

Weak
Analyze sales report for 3rd quarter.

Why the weak prompt fails

- No role is defined, so the answer may be a generic summary.
- No goal is defined, so the model does not know whether to explain, diagnose, forecast, or recommend.
- No data boundaries are defined, so the model may infer or invent missing numbers.
- No output format is defined, so the answer may not be usable in a meeting.
- No negative instruction is defined, so assumptions may appear as facts.

Best prompt

Best
<p>You are a Sales Manager preparing for the Q3 performance review.</p> <p>Goal: explain what happened in Q3, identify the biggest opportunities and risks, and recommend Q4 actions.</p> <p>Context: [paste report data: total revenue, target, region performance, product performance, pipeline, conversion rate, team notes].</p> <p>Task: analyze the report for leadership.</p> <p>Output format:</p> <ol style="list-style-type: none"> 1. Executive summary - 5 bullets 2. Key insights - each with evidence from the data 3. Underperforming areas - table with Cause / Evidence / Action 4. Q4 priorities - ranked top 5 5. Missing data or questions <p>Negative instruction: do not invent data. If evidence is missing, say "data not provided."</p>

Sample insight format

Insight	Evidence	Action
Revenue concentration risk	Two regions contribute most of total revenue while one region trails target.	Review lead quality and rep capacity in the weaker region.
Pipeline gap	Late-stage pipeline appears thin for next quarter.	Add weekly pipeline review and recovery plan.
Product mix issue	One product line is outperforming while another is flat.	Shift sales enablement and campaign focus toward high-fit segments.

Manager mindset
A good analysis prompt should not ask only "what happened?" It should ask "what happened, why does it matter, and what should we do next?"

16. Prompt Debugging: The Five Bugs

The presentation teaches a crucial debugging idea: read the bad output as a mirror. The bug is usually in the prompt. A weak output reveals what the prompt failed to define, constrain, or prioritize.

Bug	What it looks like	Fix
1. Conflicting instructions	The prompt asks for "premium yet casual," "scientific but fun," or "short but descriptive." The model averages incompatible directions.	Pick the main register and rule out the competing one. Define the hierarchy of priorities.
2. Underspecified task	The request says "summarize this" or "write an ad" without saying for whom, how long, or why.	Add audience, goal, length, platform, and emphasis.
3. Missing format anchor	The answer is useful but in the wrong structure.	Specify sections, bullet count, table columns, examples, or template.
4. Role mismatch	The output sounds like a generic chatbot, not a domain expert.	Assign a role that matches the task expertise level.
5. Ambiguous scope	The prompt says "make it better" and the model changes the wrong thing.	Define the improvement goal: clearer, shorter, more persuasive, more evidence-based, less formal.

Debugging process

- Step 1: Identify what is wrong with the output. Do not rewrite immediately.
- Step 2: Map the defect to one of the five bugs.
- Step 3: Patch only the cause, not the entire prompt.
- Step 4: Run again and compare the change.
- Step 5: Save the improved prompt as a reusable template.

Debugging prompt

The output below is not good enough. Diagnose the likely prompt problems using these categories: conflicting instructions, underspecified task, missing format, role mismatch, ambiguous scope. Then propose a revised prompt that fixes only the identified issues.

Output to diagnose: [paste output]

Key takeaway

Most people respond to a bad output by asking for "better." Skilled prompters name the defect and patch the missing instruction.

17. Case Study: Skincare Tagline Debugging

The skincare tagline case from the presentation is a strong example because the weak outputs sound polished but generic. The goal is not only to improve the tagline; it is to learn how to diagnose why the original prompt produced a cliché.

Original prompt

Original

Write a tagline for our skincare brand. We want something emotional and memorable. The brand is called Bare. We sell minimalist, clean skincare for people who are tired of complicated routines.

Bad output

Bad output

Bare: Pure. Simple. Effective. Clean skincare for a cleaner you. Because you deserve the best, naturally.

Diagnosis

Problem	Why it happened	Prompt patch
Generic wellness language	The prompt used broad words like emotional, memorable, clean, and minimalist.	Define the specific feeling: relief from performance, not luxury aspiration.
Cliché words	No banned-word list blocked default skincare terms.	Ban pure, clean, simple, natural, effective, deserve, best.
Weak use of brand name	The prompt did not ask for a double meaning of Bare.	Ask for taglines that use Bare as both brand name and emotional idea.
Voice unclear	Minimalist describes the product, not the voice.	Choose voice: calm, direct, anti-overcomplication, not poetic hype.

Revised prompt

Fixed prompt

You are a brand copywriter for a minimalist skincare brand called Bare.
 Audience: people tired of complicated routines and beauty-performance pressure.
 Desired feeling: relief, self-acceptance, and permission to stop overdoing it.
 Task: write 12 short taglines.
 Voice: calm, direct, modern, slightly anti-BS; not luxury, not clinical, not poetic hype.
 Use Bare as a possible double meaning: bare skin, bare minimum, nothing to hide, nothing to prove.
 Avoid these words: pure, clean, simple, natural, effective, deserve, best, glow.
 Format: tagline only, maximum 6 words each.

Improved outputs

- Bare. Nothing to prove.
- Stop performing. Start Bare.
- Skin care without the performance.
- Less routine. More you.
- Bare minimum. Maximum calm.

18. Practice Workbook and Answer Key

Use these exercises to test beginner and intermediate prompting skills. Try writing your own answer first, then compare with the suggested prompt structure.

Exercise A: Improve a vague learning prompt

Weak prompt

Teach me Excel.

Suggested answer

You are a patient Excel tutor. Teach me the basics of Excel for office work. I am a beginner. Create a 7-day learning plan with daily topics, 3 practice tasks per day, and a simple checklist. Focus on formulas, formatting, sorting/filtering, pivot tables, and charts. Avoid advanced VBA.

Exercise B: Make a content prompt more specific

Weak prompt

Write a LinkedIn post about AI.

Suggested answer

You are a LinkedIn ghostwriter for early-career professionals. Write 3 LinkedIn post drafts about how beginners can use AI to learn faster. Audience: college students and freshers. Tone: practical, encouraging, not exaggerated. Format each draft with Hook, Body, Takeaway, and CTA. Keep each under 130 words. Avoid fear-based language and generic claims like "AI is the future."

Exercise C: Add a format anchor

Weak prompt

Summarize this meeting transcript.

Suggested answer

Summarize the meeting transcript for a project manager. Output sections: 1. Decisions made, 2. Open questions, 3. Action items with owner and deadline, 4. Risks, 5. One-paragraph executive summary. Do not invent owners or dates. If missing, write "not specified."

Exercise D: Debug a bad output

Bad prompt

Rewrite this to sound professional but also casual, emotional, short, detailed, and funny.

Suggested diagnosis

Bug: conflicting instructions and ambiguous scope. Fix by choosing one primary tone and defining what should improve. Revised task: Rewrite this message to sound professional and approachable. Keep it under 80 words. Preserve the key information. Avoid jokes and exaggeration.

19. Final Cheat Sheet and Templates

One-line prompt upgrade

From vague to useful

For [AUDIENCE], create [DELIVERABLE] about [TOPIC/CONTEXT]. Use [FORMAT]. Tone: [TONE]. Include [MUST-HAVE]. Avoid [MUST-AVOID].

Full prompt template

Reusable template

Role: You are [specific role].
 Context: [audience, situation, goal, constraints, data].
 Task: [exact deliverable].
 Format: [sections, table columns, word count, number of options].
 Examples: [optional examples of desired style or structure].
 Rules: [must include, must avoid, evidence rules, tone rules].
 Quality check: before finalizing, verify the answer follows every constraint.

Technique selection guide

Need	Use this technique
Quick first draft	Zero-shot with clear task and format.
Specific tone or style	Few-shot examples.
Professional judgment	Role-based prompting.
Avoid generic output	Constraint-based prompting and banned words.
Complex decision or analysis	Structured reasoning with assumptions and evidence.
Output is close but not right	Iterative refinement and debugging.

Final prompt quality checklist

- Role: Did I assign the right expertise level?
- Context: Did I include audience, situation, and constraints?
- Task: Is the deliverable specific and action-oriented?
- Format: Did I define the output shape?
- Examples: Would samples improve tone or consistency?
- Rules: Did I state what to include and avoid?
- Verification: Did I ask the model to flag missing data or uncertainty?
- Refinement: If the first output is weak, did I diagnose the exact bug?

Closing idea

Beginner prompting asks questions. Intermediate prompting designs instructions. The difference is structure: you move from hoping the model understands to giving it a clear job, context, format, and quality bar.

Prepared by Ashutosh Singh