

Prompting 101

AI Consultant

GUIDES · 9 MIN READ · AUG 6, 2026

This is the full breakdown behind the Prompting 101 carousel, not just a link list. It covers what prompting actually is, the building blocks every prompt is made of, the techniques that separate a vague ask from a precise one, and where most prompts quietly break. Copy paste templates are included so you can use this immediately, not just read it. The resource list is at the end for when you want to go deeper.

WHAT IS PROMPTING

Prompting is writing instructions that get an AI model to produce the output you actually want. It is the input side of every AI interaction, and it is a skill, not luck.

Before getting good at it, two mindset shifts matter more than any technique. Treat the model as a fast, literal collaborator, not a mind reader, it only knows what is in front of it. And expect to iterate, the first output is a draft, not a final answer.

One more thing worth knowing early: the model only has access to what is in the current conversation, called the context window. Past a certain length, older details get deprioritized, which is why restating what matters in a long conversation helps.

THE TWO BUILDING BLOCKS

Every prompt is built from two things. Get both right and the technique layer gets much easier.

Context, what the model has to work with: background information (what it knows going in), examples (what good looks like), and constraints (what to avoid).

Instruction, what you are actually asking for: the task (the specific action), the format (how you want it delivered), and the role (the lens it should answer from).

Vague context in, vague output out. Specific in, specific out. Most disappointing outputs trace back to one of these six pieces being missing, not the model being incapable.

ZERO SHOT VS FEW SHOT

Zero shot is one instruction with no examples. Fast, and it works fine for simple, familiar tasks.

Few shot is an instruction plus 2 to 3 examples of the exact output you want. The model matches your pattern instead of guessing at it, and it is the single fastest upgrade for any task where tone or format matters.

Specificity changes everything between the two. Vague: "write me a caption." Specific: "write 3 captions, casual tone, under 150 characters, one call to action each." Same task, completely different output quality.

COPY PASTE TEMPLATES

Three templates that cover most of what the techniques above are trying to teach. Replace the brackets and use them as is.

ROLE, TASK, AND FORMAT

You are [a specific role, for example a senior copywriter, a blunt editor, a patient tutor].

Context:

[background the model needs: audience, goal, constraints]

Task:

[the specific action you want done]

Output format:

[exact shape you want back: length, structure, tone]

CHAIN OF THOUGHT

Think through this step by step before answering.

Problem:

[the question or task]

Work through your reasoning first, then give your final answer clearly labeled as Final Answer.

SELF CHECK

Here is my draft answer to [the task]:

[paste your draft]

Review this against the following criteria before finalizing:

1. [criterion one]
2. [criterion two]
3. [criterion three]

Point out anything that fails a criterion, then give the corrected version.

WHERE PROMPTS BREAK

Too vague: "make this better" gives the model nothing to aim at. Say exactly what "better" means, shorter, punchier, more formal, whatever it actually is.

Too much in one prompt: stacking 5 tasks into 1 message drops output quality across all 5. Split it into steps instead.

No context: asking for a caption with no info on brand, audience, or goal gets generic guesses filled in for you. The gaps you leave get filled with the model's best guess, not your intent.

Assuming it remembers: in a long conversation, restate what matters. Older context gets deprioritized, not deleted, but deprioritized enough to quietly stop shaping the output.

ADVANCED: CHAINING AND STRUCTURED OUTPUT

Multi step prompting (chaining): break a big task into a sequence, where the output of step 1 becomes the input to step 2. Each step does one job well instead of one prompt trying to do everything, which is usually where quality drops.

Structured output: ask for a specific format the output has to follow, JSON, a table, a fixed set of fields. This forces consistency and makes the output easy to hand off to another system or person without reformatting it by hand.

BEGINNER PRACTICE

- ☐ Rewrite one of your own captions in 3 different tones
- ☐ Turn a messy paragraph into a clean bullet list
- ☐ Write a prompt that always returns exactly 3 options

INTERMEDIATE PRACTICE

- ☐ Build a reusable template for a task you do weekly
 - ☐ Write a role plus task plus format prompt for a real deliverable
 - ☐ Get the model to catch its own mistakes before you do, using the Self Check template above
-

ADVANCED PRACTICE

- ☐ Chain 3 prompts into one workflow, each step feeding the next
 - ☐ Build a system prompt for a recurring assistant persona
 - ☐ Design a prompt whose output another tool can read directly, like JSON another script parses
-

WEBSITES

Guide: [Learn Prompting](#)

Guide: [Prompting Guide](#)

FREE COURSE

Course: [ChatGPT Prompt Engineering for Developers](#)

Taught by Andrew Ng and Isa Fulford, free and self paced.

OFFICIAL DOCS

Docs: [Anthropic's Prompt Engineering Guide](#)

YOUTUBE

Channel: [Matt Wolfe](#)

Channel: [AI Explained](#)

COMMUNITY

Community: [Prompt Engineering Subreddit](#)



WORK WITH ME

Business owner looking to automate?

Book a free 30-minute discovery call. We'll pressure-test fit, scope, and what to build first.

hexgarcia.com/free-call

Questions? Reply to the newsletter or reach out via email.

Hector Garcia

@hex.gar · contact@hexgarcia.com