

# AI FROM ZERO

## CHAPTER 01

# Your first useful conversation

You do not need to learn a secret language. You need to learn how to ask for help clearly - and how to check what comes back.

### ASK

Describe a real task

### REFINE

Say what needs to change

### CHECK

Decide what you can use

## What you will do in this chapter

Start with an ordinary task, write a clear request, improve the answer in a second message, and check the final result. By the end, you will have a simple method you can reuse for emails, explanations, ideas, and everyday planning.

Reading time: about 15 minutes. Practice time: about 10 minutes. No coding or prior AI experience required.

This is a sample chapter for editorial and design review. Product names are examples, not endorsements. The lesson focuses on skills that remain useful as individual tools change.

# You already know how to begin

Imagine asking someone to help you write an email. If all you say is "Write an email," they have to guess who it is for, what happened, and what you want the reader to do. An AI assistant faces the same missing information. It can produce fluent text anyway, which makes the guessing harder to notice.

A prompt is simply the input you give an AI system. In a conversational assistant, that often means a written request or a spoken question. There is no special ceremony. You can start with the same plain language you would use with a helpful colleague.

## What an AI assistant actually does

A text-generating AI model produces responses based on patterns learned during training and the information available in the current interaction. Depending on the product and its enabled tools, an assistant may also search the web, work with files, or perform other tasks. The conversation itself does not guarantee that those tools are available or that they have been used.

The practical consequence is simple: an answer can sound convincing and still be wrong. A confident explanation is not evidence. Treat the assistant as a useful collaborator whose work you review, especially when the result includes facts, numbers, or consequential advice.

## Three terms without the jargon

Model: the underlying system that generates an answer.

Assistant or app: the product you interact with, which may connect a model to tools and other features.

Prompt: the instruction, question, or other input you supply.

ChatGPT and Gemini are examples of assistant products. Qwen is a model family; SDXL is an image-generation model. Those names describe different kinds of things. Later chapters will explain the differences. For today, begin with a conversational assistant and a small writing task.

### THE FIRST PRINCIPLE

You do not need a perfect prompt. You need a clear starting request and a willingness to improve the result.

# Make your first win small

Do not begin by asking AI to reorganize your entire life. Choose something you can finish and judge in a few minutes. A message you have been putting off is ideal: you already know the situation, and you can tell whether the draft sounds right.

## Good tasks for a first session

Draft a polite scheduling email. Turn a messy list into a short checklist. Explain an unfamiliar term in everyday language. Brainstorm names for a personal project. These tasks give you useful practice without requiring you to trust unfamiliar facts.

For the walkthrough, we will draft a fictional email to a repair company. The company promised an estimate on Tuesday, but it has not arrived. We want an update while keeping the relationship friendly. Nothing needs to be sent; producing a draft is enough.

## Before you type

Open a conversational assistant you have chosen and begin a new conversation. Look for the area where you can enter a message. The exact labels and layout vary between products. You do not need to explore every feature before starting this exercise.

Use fictional details while you learn. Leave out account numbers, passwords, client records, and other information you do not need for the task. If you later want to use confidential material, first understand the service's applicable data controls and your own obligations.

### A WEAK STARTING REQUEST

Write an email asking for an update.

This may produce something usable, but it leaves almost everything to chance. An assistant might invent a project, use a tone you dislike, or add a deadline you never agreed to. The fix is to give it the information you would give a person.

The next page introduces a four-part method. It is a teaching aid, not a technical requirement: you do not have to label the parts or include all four every time.

# Give the answer a destination

A useful request usually tells the assistant what to do, what it needs to know, what a good result looks like, and how to deliver it. We will call those four parts Task, Context, Standards, and Format.

## TASK

What should it do?

Draft an email asking for an estimate update.

## CONTEXT

What does it need to know?

The estimate was promised Tuesday; it is now Friday.

## STANDARDS

What makes the result useful?

Friendly but direct. Do not invent details.

## FORMAT

How should it deliver the answer?

A subject line and a message under 100 words.

The task gives direction. Context reduces guessing. Standards make your preferences visible. Format makes the result easy to use. Together, they turn a vague request into a practical assignment.

You can combine the four parts into one natural paragraph. There is no reward for making your prompt sound technical. "Keep it short and friendly" is often more helpful than "Act as a world-class communications expert."

If you do not know what information is missing, add: "Before drafting, ask me any essential questions." Use that when a wrong assumption would change the answer. For a simple task, it is usually faster to supply the key facts yourself.

# From vague request to usable draft

## COPY AND TRY THIS

Draft an email to a repair company asking for an update on my estimate. They said it would arrive Tuesday, and it is now Friday. Keep the tone friendly but direct. Ask when I can expect it. Use a subject line and keep the message under 100 words. Do not invent names, prices, or additional facts.

Below is an illustrative draft written for this guide. Your assistant may produce different wording. The goal is to evaluate the answer against your request, not to match these sentences exactly.

## ILLUSTRATIVE RESULT

Subject: Update on repair estimate

Hello,

I'm following up on the estimate you said would arrive Tuesday. Since it's now Friday, could you let me know when I can expect it?

Thank you for your help.

## Why this is a useful starting point

The draft states the reason for writing and asks a clear question. It does not accuse the company of bad faith. It stays within the requested length, and it does not make up a price, repair type, or contact name.

It is still your draft to review. Did the company really promise Tuesday? Is Friday the correct day? Would you normally say "Hello" or use someone's name? A good AI response makes the work easier; it does not remove your responsibility for what you send.

## What if the first answer misses?

If it adds a detail you never supplied, ask it to remove that detail. If it sounds stiff, describe the tone you want. If it misunderstands the task, restate the goal. You do not need to abandon the conversation because the first draft is imperfect.

# The second message is part of the work

Many beginners treat the first answer as the final answer. A conversation gives you another option: tell the assistant what worked, what did not, and what to change. Specific feedback is easier to act on than "Make it better."

## A USEFUL FOLLOW-UP

Keep the same facts and request, but make it sound more conversational. Remove "Since it's now Friday" and simply ask for an update. Keep the subject line.

## ILLUSTRATIVE REVISION

Subject: Update on repair estimate

Hello,

Just following up on the estimate you expected to send Tuesday. Could you give me an update and let me know when I can expect it?

Thanks for your help.

Notice what changed: the wording became more relaxed, while the key facts and the request remained. That is a useful way to refine a response. Identify the dimension you want to change rather than asking for a completely different answer without explanation.

## Feedback you can reuse

Too long: "Cut this to 80 words and keep the main request."

Too technical: "Explain this for someone hearing the term for the first time."

Too generic: "Use the details I supplied and remove statements that could apply to anyone."

Wrong assumption: "I did not say the appointment was confirmed. Revise that sentence."

When an important requirement gets lost, restate it. If the conversation becomes tangled with several competing versions, begin a new one with the facts and requirements you still want. Think of this as giving the assistant a clean brief, rather than arguing with an old draft.

# Check before you trust

An AI assistant can make a mistake without sounding uncertain. It may also agree with your premise when the premise deserves a challenge. Your final check should focus on the things that matter for the task.

## For a message, check these four things

Facts: Are names, dates, amounts, and events correct?

Intent: Does it ask for the action you actually want?

Tone: Would you be comfortable saying it yourself?

Additions: Did it introduce promises, claims, or details you never approved?

## Different tasks need different checks

For a brainstorm, usefulness and originality may matter most. For an explanation, look for accuracy and a clear distinction between facts and guesses. For a calculation, verify the inputs and recompute the result. For current information, check a reliable, dated source rather than assuming the assistant has up-to-date knowledge.

### A PROMPT THAT HELPS WITH REVIEW

Review your draft against my original instructions. Point out any assumptions or details I should verify. Then provide a corrected version if needed.

Self-review can help catch omissions, but it is not independent verification. Asking the same assistant "Are you sure?" does not establish that an answer is true. A source link is also not enough by itself: open the source and check whether it supports the claim.

## When the stakes rise, slow down

An email draft and a medication decision do not deserve the same level of trust. For medical, legal, financial, or safety decisions, use appropriate authoritative sources and qualified help. AI can help you organize questions or understand terminology; do not rely on an unchecked answer to make a consequential decision.

For this chapter's exercise, the final test is modest: read the email, confirm the facts, and decide whether you would actually use it. That is already a complete and useful AI workflow.

# Ten minutes. One real result.

Choose a small task from your own day. Use fictional names and remove sensitive details. Write your first request with the template below, then make one focused revision. You are practicing a process, not trying to impress the software.

## YOUR REUSABLE STARTER

Help me [task].

Here is the context: [relevant facts].

The result should [tone, audience, length, and other requirements].

Give it to me as [format]. Do not invent missing facts; flag anything you need me to confirm.

## The practice sequence

1. Ask. Send your request and read the answer.
2. Refine. Identify one thing that would make it more useful and ask for that change.
3. Check. Verify the facts and remove anything you would not stand behind.
4. Keep. Save the useful result and the prompt that helped you get there.

## Try a different kind of task

### AN EXPLANATION PROMPT

Explain the difference between a modem and a router for a beginner. Use one everyday analogy and keep it under 200 words.

Follow up with: "The analogy helped, but why might I need both? Explain that part more simply." The same method works for learning and writing.

## What you should take with you

A useful prompt gives direction and enough context to reduce guessing. A follow-up shapes the result. A final check protects you from fluent mistakes. You can do all three in plain language.

Coming next: choosing an AI tool without getting lost in model names, feature lists, and subscription choices.