

FROM USER TO ORCHESTRATOR

Free extract
One full chapter.
Plus four ready prompts.

This is a real chapter from the course, not a summary.
You will find it exactly like this inside the full
200+ page manual.

Read it. Apply it to something on your desk today.
Then judge for yourself whether the method works.

If it helped, the full course is at ai.evolbot.com

In the previous chapter we established that context is the difference between a generic output and a useful one. The next step is understanding what kind of context really matters, and how to organize it so the AI receives exactly the information it needs to do a good job.

Over time, among professional AI users, a standard structure for prompts has consolidated. It's not a rigid protocol to memorize, it's more of a checklist: four questions to ask yourself before hitting send. Most poor outputs come from ignoring at least one of these four questions.

The four components are context, role, objective and format. Let's see what each one does, why it matters, and what happens when you skip it.

Context is the starting point. It's everything the AI can't know on its own about your specific situation: who you are, what industry you work in, who is the recipient of the output, what history or circumstance led to this request. Without context, the AI responds to a generic version of your question. With the right context, it responds to your question. The distinction seems obvious but has enormous practical implications.

A practical test to understand if your context is sufficient: imagine sending this request to someone who has never seen your project, doesn't know your company, and has no access to any information beyond what you wrote in the prompt. Could they complete the job? If the answer is no, more context is still missing. AI doesn't fill gaps reliably: it fills them with statistical plausibility, which is a fancy way of saying it guesses in slightly wrong ways. "Write an email to handle a dissatisfied client" is a generic request. "Write an email to handle a dissatisfied client who bought our warehouse management software three months ago, experienced a four-hour outage yesterday due to a known bug, and with whom we have a renewal contract being discussed next week" is your situation.

Role is the instruction that tells the AI from which perspective to respond. It's not a matter of fancy, it's of practical utility: the kind of response you get from a "lawyer expert in commercial contracts" is structurally different from what you get from a "promotional copy writer" or from an "HR manager with twenty years of experience in multinational companies." The AI doesn't change models when you assign it a role, but it changes how it weighs information, the language it uses, the level of technical detail in the response, and the angle from which it approaches the problem. Assigning the correct role means you're talking to the right version of the same tool for the specific task you have in front of you.

Objective is the part that most often gets defined too vaguely. It's not enough to say what you want in abstract. You have to specify what you want in concrete as output. One thing is "help me with this presentation," another is: "write the texts for five slides presenting Q3 results to our investors, with focus on operating margin and customer base growth." The first formulation doesn't give the AI an endpoint. The second does. The objective must answer the question: when I read the output, how will I know if it succeeded?

Format is the element most often overlooked by beginners and most appreciated by advanced professionals. You tell the AI how to present the response: length, structure, tone, style, language, level of formality. Do you want continuous text or bullet points? Do you want 200 words or 800? Do you want formal or conversational tone? Do you want it to conclude with an explicit recommendation or just with the facts? Without format guidance, the AI will choose for you, and the choice might not match what you have in mind. With the right guidance, the output arrives already in the form you need, without needing to rework it.

To see the principle in action, take a real request and observe how it changes as you add the components one by one.

Objective only (the minimum): "Write a risk analysis of the project."

Objective + context: "Write a risk analysis for an ERP implementation project in a medium-sized manufacturing company of 300 employees. The project is six months from go-live, with a budget of 800,000 euros. The main risks concern the migration of historical data and staff training."

Objective + context + role: "You are a project management consultant with experience in ERP implementations in the manufacturing sector. Write a risk analysis for an implementation project in a company of 300 employees, six months from go-live, budget 800,000 euros. Focus on data migration and staff training."

All four components: "You are a project management consultant experienced with ERP in manufacturing. Write a risk analysis for a board of directors (non-technical) on a SAP S/4HANA implementation in a company of 300 employees, six months from go-live, budget 800,000 euros. Focus on data migration and training risks. Format: max 400 words, direct tone, three main risks with estimated probability and impact, one final recommendation."

Each addition produces a qualitatively different output. Not just longer: more relevant, more calibrated to the actual recipient, more usable without reworking.

There's a secondary effect of this discipline worth highlighting. Shopify CEO Tobi Lütke observed that forcing himself to provide complete context to the AI made him a better communicator in general: more precise emails, clearer memos, more solid decision frameworks. The logic is the same as always: clarity of thought is the same whether the request goes to an AI system or to a person. Those who learn to build a precise prompt tend to become more precise in everything else too.

The results of this approach are measurable. Those who use structured methods register improvements in perceived output quality between 40% and 60%, and develop effective prompts 65% faster than those who write by instinct. An analysis of over 2,000 real corporate prompt templates conducted by TU Munich in 2025 confirmed that the order of components makes a difference to the result: the framework that achieves the best results in production corresponds exactly to the sequence we're using in this course.

REALITY CHECK

Not every prompt needs to use all four components. The framework is a map, not a form to fill out. On simple requests, context and objective are enough. On short and generic requests, even just the objective is fine. The value of the framework is not in mechanical application, but in the mental habit it creates: before you write, mentally run through the four questions and consciously decide which to include and which to skip. This awareness is worth more than any template.

Case: A lawyer running a small practice started using AI to prepare draft letters to clients. His prompts were good on context, but completely lacking in format guidance. The result was texts that were too long, with inconsistent structure, requiring almost as much rewriting time as a manual draft would have. When he added a final line along the lines of "professional tone, maximum 200 words, structure: situation, legal position, next step," the review time cut in half. The content was already good. It was the form that was creating extra work.

Counterexample: A common mistake is assigning an overly specific role on topics where the AI doesn't have enough granularity to really differentiate. "You are an expert in B2B software sales in the retail sector in northeastern Italy" rarely produces something very different from "you are an expert in B2B software sales." Role works better for defining general disciplines (lawyer, doctor, copywriter, engineer, university professor) than for hyper-specific niches. On the latter, descriptive context is worth more than the role label.

EXERCISE

Exercise 2.2: Take a task you need to complete this week, something real, not an invented exercise. Write the prompt using all four components: context, role, objective, format. It doesn't have to be perfect. Then execute the prompt and, before reading the output critically, ask yourself: which of the four components could I have defined better? Modify that component and re-execute. Compare the two outputs. The goal isn't to get the perfect prompt on the first try, but to learn to diagnose what was missing.

2.3 – Iterative prompting: how to refine the request until you have what you want

There's a way of using AI that's almost universal among beginners, and that's almost always wrong. It works like this: you write a prompt, read the output, it doesn't quite satisfy you, you close the window. Or you accept a mediocre result because "at least I saved some time." The problem isn't the AI. It's that you're using a conversational tool like it's a vending machine: insert the coin, press the button, accept whatever comes out.

AI is not a vending machine. It's a conversation.

You are an expert professional communication consultant.
I need to respond to this email: [paste the email].
My goal is [e.g. decline the request while maintaining the relationship / ask for clarification without sounding passive-aggressive / respond to criticism without over-apologizing].
The tone should be [e.g. direct but respectful / warm but firm].
Write three versions of the response, from softest to most direct.

First draft of an email from scratch

I need to write an email to [recipient's role, e.g. a client who hasn't responded in three weeks].
Context: [what happened, what I want to achieve].
Tone: [e.g. professional and patient / urgent without being aggressive].
Length: maximum [e.g. 5 lines / one paragraph].
Write the draft.

Summary of a long email thread

Here's an email thread between me and [who]: [paste the thread].
Summarize it in three points: (1) what it's about, (2) where things stand now, (3) what they expect from me as the next action.

Responding to a sales objection

A potential client said: "[exact objection]".
Our product/service is [brief description].
Write a response that: acknowledges the objection without dismissing it, reframes it from a different angle, and proposes a concrete next step. Tone: consultative, not defensive.

Documents and Reports

First draft of a report from notes or bullet points

I have these raw notes on [topic]: [paste notes].
Turn them into a structured professional report with: introduction (context and objective), body (data and analysis), conclusion (implications and recommendations).
Audience: [e.g. board of directors / project team / external client].
Tone: [e.g. formal and concise / analytical with data highlighted].

FROM USER TO ORCHESTRATOR

This was one chapter out of eight modules.

You just read how to structure a request. The full course also covers how to refine it by iterating, how to choose the right tool, how to measure AI ROI in your work, how to bring it into a resistant team, and what changes when AI stops answering and starts acting.

200+ pages. Eight modules. Four appendices, including the full prompt library with 9 categories.

Download once. Yours forever. Updated as tools change.

Get the full course: ai.evolbot.com - EUR 67

Secure payment via Stripe. PDF available immediately.