DécodAI

PRACTICAL GUIDE

The Guide to Context-Driven AI

Giving AI the context it needs to be useful at work, whatever the tool.

Bruno Poirier
AI Adoption Specialist, DécodAI

2026 Edition · Version 2.0 · November 2026

Foreword

I've spent more than twenty years in data and IT. In my last years in a large company, I led data and AI literacy programs for more than 10,000 employees. I founded DécodAI in 2026 to help Quebec SMBs make the same shift.

What I take away from it fits in one sentence: AI is all about people. Tools change every few months. What lasts is how you think, how you phrase a request and how you check the result.

When I started working seriously with language models, I looked for the right formulas and tricks like everyone else. That doesn't get you far. What makes the difference is the context you provide: who you are, what you're looking for, for whom, and under what constraints. Exactly what you'd explain to a competent new colleague on their first day at the office.

This guide brings together the techniques that hold up, whether you use Copilot, ChatGPT, Claude, Gemini or another assistant. The prompts are ready to copy. Replace whatever is in square brackets with your own situation.

One piece of advice: don't try everything at once. Pick one technique, use it for a week on your real work, then move on to the next. Chapter 10 explains how.

Bruno Poirier

AI Adoption Specialist, DécodAI

Contents

01	Understand the model to instruct it better	4
02	The Socratic method applied to AI	6
03	Guided self-critique	8
04	Context engineering	10
05	Personalize your assistant	12
06	Professional writing and editing	14
07	Research and analysis	16
08	Clients and business development	18
09	Security, limits and good practices	20
10	Make AI a habit	22
11	Ready-to-use prompt library	24
A	Appendix: Before you publish or send	27

01

CHAPTER 01

Understand the model to instruct it better

It predicts rather than knows. Here's how to make the most of that.

How a language model produces an answer

Copilot, ChatGPT, Claude and Gemini all rely on large language models. These models don't look things up in a verified knowledge base. They produce, word by word, the most likely continuation given what they've learned and what you've given them.

The direct consequence: a model can state something false with great confidence. It isn't lying. It produces a plausible answer, and plausible doesn't mean true. This is called a hallucination.

Recent tools reduce the problem without eliminating it. Several now search the web, read your files or take time to reason before answering. They get some tasks wrong less often. They still get things wrong.

The other consequence is more encouraging: the more precise the context you provide, the more closely the answer fits your reality. This whole guide builds on that.

Three ways to ask

In training sessions, I keep seeing the same pattern: most people stay at the first level.

LEVEL 1 · VAGUE QUESTION

A generic answer that could be meant for anyone.

LEVEL 2 · STRUCTURED REQUEST

A goal, a format, a few constraints. Usable, with some editing.

LEVEL 3 · FULL CONTEXT

Who you are, what's at stake, for whom, with which examples and which constraints. This is where the tool becomes truly useful.

The same need at all three levels, for a manufacturing SMB:

LEVEL 1 · VAGUE QUESTION

```
Do a risk analysis for my plant.
```

LEVEL 2 · STRUCTURED REQUEST

```
Identify the 5 main operational risks for a 45-employee manufacturing SMB that makes metal
components in Quebec. Present them as a list, with a priority level for each.
```

LEVEL 3 · FULL CONTEXT

```
Context: I'm the operations director of a 45-employee machining SMB. We mainly serve the
aerospace sector. Our main client accounts for 60% of our revenue. We've had 3 quality
incidents in 6 months.

Goal: an analysis of the 5 priority operational risks, with their criticality level and one
concrete action for each.

Audience: our management committee, who will have 15 minutes to read it.

Constraints: real risks and realistic actions for an SMB, not large-company theory. If
you're missing information to properly assess a risk, say so instead of assuming.
```

KEY TAKEAWAYS

- AI predicts, it doesn't know. Check what matters.
- The quality of the answer follows the quality of the context you give.
- Writing a level 3 request takes two extra minutes and often saves several attempts.

02

CHAPTER 02

The Socratic method applied to AI

Let the model question you before it answers

Reverse the roles

A simple technique: instead of explaining everything at once, ask the model to ask you the questions it needs. You don't have to guess what it's missing, it tells you. This is especially useful when the deliverable is complex, or when you don't yet know exactly what you want.

A word on numbers: spectacular results attributed to this technique circulate online. They don't come from any published study. Trust what you observe in your own work instead.

SOCRATIC PROMPT · BASIC VERSION

```
I want you to help me with the following task: [describe your task in 1 or 2 sentences].

Before producing anything, ask me the questions you need to fully understand my goal, my
audience, the tone I want and the format constraints. Ask them one at a time.

Only start writing once you have what you need.
```

Have it name its assumptions

This variant forces the model to put its assumptions on the table before answering. It's mainly useful for analyses and decisions.

SOCRATIC PROMPT · ADVANCED VARIANT

```
Before answering my question, first do this:

1. What assumptions would you normally make when answering this type of question?
2. What contrary information should you take into account?
3. What is the most likely scenario, and what is the riskiest but plausible scenario?

Then answer my question taking this work into account.

My question: [your question]
```

A concrete example

You're the HR manager of an 80-employee company, many of whom work in the field. You need to announce a new management software. Rather than asking AI to write the announcement right away:

EXAMPLE · INTERNAL COMMUNICATION

```
I need to prepare the announcement of new operations management software for our 80 field
employees, several of whom aren't comfortable with technology.

Before proposing a plan, ask me the questions that will help you understand: the level of
resistance I expect, the concerns of seasonal employees, the best way to reach them and the
message I want everyone to remember.
```

KEY TAKEAWAYS

- Don't start with the answer you expect. Start with the questions the model should ask you.
- The more complex the deliverable, the more useful the technique.
- Having the model name its assumptions is a good reflex for important decisions.

03

CHAPTER 03

Guided self-critique

Generate, critique, improve, while knowing the method's limits

What self-critique does well

Asking the model to reread and critique its own answer often improves the result: clearer text, a better-built argument, blind spots named. The method works in three steps: generate, critique, improve.

What it does less well

Left to itself, a model is poor at correcting its own reasoning errors. Google DeepMind researchers even observed that it sometimes turns a good answer into a bad one when simply asked to review itself (Huang et al., 2023).

What works better for facts is checking against outside information: a source, a document, a figure you know. Meta researchers showed that hallucinations are reduced by having the model draw up a list of verification questions, answer them separately, then correct the answer (Dhuliawala et al., 2024). The prompt below is based on that.

So-called reasoning models already do part of this work before answering. A critique you ask for is still useful, especially if you give it precise criteria.

UNIVERSAL TEMPLATE · GUIDED SELF-CRITIQUE

```
Step 1: give a complete answer to the following question: [your question or task].

Step 2: reread your answer as a skeptical expert. Name 3 to 5 possible problems: logical
flaws, unsupported claims, blind spots, oversimplifications.

Step 3: list the factual claims in your answer. For each one, say whether you're certain of
it, whether it needs to be checked or whether you assumed it.

Step 4: produce an improved final version and say what you changed and why.
```

The variant: a simulated expert debate

For a decision where several readings can be defended, ask the model to present different points of view in turn. Useful for an investment, an organizational change or a regulatory question.

PROMPT · SIMULATED EXPERT DEBATE

```
I need to make a decision about [decision topic].

First analyze the situation from the point of view of a [point of view 1, e.g.: cautious
CFO]. Then from the point of view of a [point of view 2, e.g.: growth-oriented sales
director]. Finally, from the point of view of a [point of view 3, e.g.: change management
specialist].

For each one: their reading of the situation, their main reservations, their recommendation.
End with what they agree on and what they still disagree on.
```

In practice

The owner of a service company is considering automating part of his customer service with a chatbot. Before deciding, he has three points of view debate: the effect on his employees, the risk to his reputation and the real return. He isn't looking for an answer, he's looking for what he hadn't seen.

KEY TAKEAWAYS

- Generate, critique, improve: effective for form, structure and blind spots.
- For facts, self-critique isn't enough. Check against an outside source.
- Ask the model to separate what it knows from what it assumes.

04

CHAPTER 04

Context engineering

Set the frame before asking the question

From prompt to context

The term context engineering spread in 2025. Anthropic presents it as the natural extension of prompt engineering: instead of focusing only on how the request is worded, you take care of everything the model has in front of it when it answers, meaning the instructions, the documents provided and the conversation history.

For you, the idea translates simply. Before your question, give the frame: who you are, the situation, what you expect, for whom, under what constraints. Attach the relevant documents rather than summarizing them from memory. Part of this context can even be saved once and for all: that's the subject of chapter 5.

FULL CONTEXT FRAME · TEMPLATE

```
═══ CONTEXT ═══
My role: [your position, your industry]
Situation: [the problem or issue]
Key information: [data, facts, attached documents]

═══ EXPECTED ANGLE ═══
Point of view: [e.g.: demanding editor, plain-language explainer, devil's advocate]

═══ GOAL ═══
Expected result: [specific deliverable]
Audience: [who will read or use the result]
Success criteria: [how you'll recognize a good result]

═══ CONSTRAINTS ═══
Format: [list, table, paragraphs]
Length: [number of words or pages]
Tone: [professional, direct, plain-language]
Avoid: [what you don't want]

═══ TASK ═══
[Your request, now that the frame is set]
```

A word about roles

Giving the model a role ("act as a demanding editor") helps set the tone, the level of detail and the angle of the answer. However, several studies show that telling it it's an expert doesn't make it more accurate on facts (Zheng et al., 2024; Wharton Generative AI Labs, 2025). Use the role for the how, not

for the truth.

Show rather than describe

Giving one or two examples of what you want is one of the best-established techniques. Two paragraphs in your own hand say more about your style than a page of instructions.

MATCH THE TONE TO YOUR STYLE

```
Here are two examples of the way I write. Analyze my style: sentence length, vocabulary, way of arguing, transitions. Summarize it in 5 points. Then apply this style to write [your task].

Example 1: [a representative paragraph]
Example 2: [a second paragraph]
```

Separate the instructions from the content

When a request mixes instructions with a text to process, separate them clearly. Tags like the ones in the example below are recommended by Anthropic and work well with other assistants. They help the model tell your instructions apart from the content. On their own, they don't protect against a malicious instruction hidden in a document (see chapter 9).

STRUCTURED EXAMPLE WITH TAGS

```
<instructions>
Edit the text provided to make it clearer without changing the message. Shorten it by about 20%. Replace passive constructions with active ones.
</instructions>

<reference_style>
Direct tone. Short sentences. Professional but accessible vocabulary.
</reference_style>

<text_to_edit>
[Paste the text to edit here]
</text_to_edit>
```

KEY TAKEAWAYS

- Before the question, the frame: who, what, for whom, under what constraints.
- The role sets the tone. It doesn't make the model more accurate.
- One or two examples beat a long description.

05

CHAPTER 05

Personalize your assistant

Say who you are once, rather than in every conversation

Two different settings

Custom instructions are what you write once and for all: who you are, how you want to be answered. The assistant applies them to every new conversation.

Memory is what the assistant retains over your exchanges: an ongoing file, a preference mentioned in passing. Depending on the tool, it fills up on its own or when you ask ("remember that I prefer short answers"). You can review it, correct it and erase it.

Several tools also offer **projects** or workspaces: a folder that groups documents and instructions specific to a mandate, a client or a topic. Handy so you don't have to re-explain everything each time.

Where to find them

TOOL	CUSTOM INSTRUCTIONS	MEMORY
Copilot (organization account)	Copilot Chat, "..." menu at the top right, Settings, Personalization, Custom instructions	Same screen: saved memories
ChatGPT	Settings, Personalization, Custom instructions	In the personalization settings
Claude	In the profile settings, personal preferences section	Settings, Memory
Gemini	Settings & help, Personal Intelligence, Instructions for Gemini (personal Google account; organization accounts have a separate setting)	In the same personalization settings

Menus change often and their names vary with the interface language. In an organization account, the administrator can also turn these features off. If you can't find them, ask your IT lead.

A template to adapt

CUSTOM INSTRUCTIONS · TEMPLATE

```
I am [your position] at [type of organization, e.g.: a 20-person accounting firm]. I mainly work on [your main tasks].

Answer in English, in a [professional and direct] tone.
Format: [the answer first, the explanation after / short lists / paragraphs].
If you're not sure about a piece of information, say so. Never make up a figure, a rate, a date or a reference.
If my request isn't clear, ask me a question before answering.
```

Let the assistant help you write them

If the blank page stops you, reverse the roles, as in chapter 2:

WRITE YOUR INSTRUCTIONS WITH THE ASSISTANT

```
Ask me 10 questions to better understand my work, one at a time.

[After your answers]
Based on my answers, suggest custom instructions of 8 to 10 lines that I can paste into the settings.
```

What not to put in them

- Personal information about your clients or colleagues.
- Passwords, account numbers or other identifiers.
- Confidential information about your organization that has no business being in a permanent setting.

Clean house

A few times a year, reread what the assistant has retained. Erase anything that's wrong, outdated or too personal. A memory full of outdated information does more harm than good.

KEY TAKEAWAYS

- You write the instructions. The assistant builds the memory. Both can be reviewed and corrected.
- Ten minutes of setup means every conversation starts further ahead.
- Nothing confidential in a permanent setting, and a clean-up a few times a year.

06

CHAPTER 06

Professional writing and editing

Make AI an editor, not a generic text machine

Write first, then have it edited

Most people ask AI to **write** a text. Those who get the most out of it mainly use it to **improve** their own.

A text written from scratch by a model is quickly recognizable: generic phrasing, put-on enthusiasm, no voice. A text that **you wrote**, then reworked with AI following your instructions, stays yours.

> *Write first, then have it edited. The result sounds like you, and your readers can tell.*
>
> Bruno Poirier, DécodAI

The demanding editor

PROMPT · THE DEMANDING EDITOR

```
Act as a demanding business communication editor. For the text below:

1. Remove redundancies and filler.
2. Replace passive constructions with active ones where possible.
3. Cut sentences longer than 25 words.
4. Replace vague adjectives (important, significant, interesting) with precise wording.
5. Start each paragraph with its main idea.
6. Keep the author's tone and voice.

Then, in 3 to 5 lines, name the main structural problems.

Text to edit: [your text]
```

Long documents, section by section

Today's models handle long texts much better than before. Yet a document produced in one go is often flat and repetitive, and you lose control of the content. Validate the structure first, then work one section at a time, giving the table of contents again at each step.

LONG DOCUMENT · STEP 1: THE STRUCTURE

```
Based on the notes below, propose the table of contents of a [X]-section document on
[topic]. For each section, give: (1) the title, (2) the question it answers, (3) 3 or 4 sub-
themes, (4) the reader it's mainly intended for.

My notes: [your notes]
```

LONG DOCUMENT · STEP 2: THE WRITING

```
Here is the document's table of contents: [structure]

We're now working on section [number]: [title]. It answers the question: [question]. Here are my notes: [notes]

Write this section in about [X] words. Start with a concrete example, develop it, then link to the next section.
```

KEY TAKEAWAYS

- Having your text edited gives a better result than having it written from scratch.
- Give precise editing instructions, not just "improve this text".
- For a long document, have the structure validated before writing.

07

CHAPTER 07

Research and analysis

Search, analyze, verify

What has changed

Not long ago, most assistants answered only from what they had learned, with no web access. That's no longer the case. Copilot, ChatGPT, Claude, Gemini and Perplexity can all search the web and cite their sources. Most also offer a deep research mode, which consults many sources over several minutes and produces a report. Depending on your licence, some of these features may not be available.

QUICK SEARCH

A specific question, an answer with links. To check a date, a rule, a definition.

DEEP RESEARCH

A broad topic, a structured report with its sources. For monitoring, a market overview, a literature review.

For analysis, give your documents directly to the assistant: reports, spreadsheets, notes. A model is generally more reliable when it works on what you provide than on what it finds on its own.

A cited source is not a verified source. Open the link. A model sometimes cites a page that doesn't say what it claims. Check the primary source before publishing or making a decision.

LITERATURE SUMMARY (WEB SEARCH ON)

```
Research [specific topic]. Find the 5 to 8 most relevant sources from the last 24 months.
For each one: (1) the reference with a direct link, (2) the main thesis in two sentences,
(3) the method used, if applicable. If you can't find a direct link, say so rather than
making one up.

Then summarize the points of consensus, the disagreements and the questions that remain
open.
```

Strategic analysis

STRATEGIC ANALYSIS

```
Help me analyze the following situation, in the [your industry] sector in Quebec.

Context: [the situation, the issues, the available data]

Produce: (1) the industry trends that will weigh most on us over the next 18 months, with
your sources, (2) a SWOT analysis specific to our situation, (3) three scenarios
(optimistic, realistic, pessimistic) with priority actions for each, (4) a recommendation
based on effort versus impact.

Format: a summary of 600 words maximum, jargon-free.
```

Verify a claim

Before quoting a statistic or a claim obtained from a model, especially in front of a client or in public, take two minutes for this:

VERIFYING A CLAIM

```
The following claim seems important to me, but I want to make sure it's reliable: [specific
claim].

Using web search: (1) which sources support it? (2) Which sources dispute or qualify it? (3)
Is it widely accepted, debated or marginal? Tell me whether I can quote it as is, with
nuance, or not at all without the original source.
```

KEY TAKEAWAYS

- Assistants now search the web. Turn search on when the answer depends on current events.
- Open the cited sources before relying on them.
- A SWOT analysis produced by AI is a starting point, not a deliverable.

08

CHAPTER 08

Clients and business development

Preparing a proposal, reading a document, anticipating a reaction

Simulate a client's reaction

Models tend to agree with the person talking to them. Asking "is my proposal good?" often leads to a flattering answer. Instead, ask the model to **play a specific client**, with their doubts, their resistance and their real priorities.

SIMULATING A CLIENT PROFILE

```
Play the following client: [profile, e.g.: CEO of a 30-employee service company, careful with spending, already let down by a supplier in the past].

I'm going to present an offer to you. React exactly as this person would, with their doubts and priorities. Don't be polite to please me.

The offer: [description of your offer]

In your answer: (1) your first reactions, (2) the questions you'd ask before accepting, (3) what would make you say no for sure, (4) what would have convinced you to go further.
```

Preparing a service proposal or a quote

After a meeting with a potential client, your notes are often all over the place. The assistant can put them in order and spot what's missing. The commercial content, however, remains your business.

STRUCTURING A SERVICE PROPOSAL

```
Here are my notes from a meeting with a potential client: [your notes]

Help me structure a service proposal with these sections: (1) our understanding of their need, in their words, (2) the proposed approach, (3) the deliverables, (4) the timeline, (5) the terms.

Don't make up any price, rate, deadline or commitment. Leave a [TO BE COMPLETED] placeholder wherever one is needed.

Then reread the proposal from this client's point of view: what would make them hesitate, what's missing, what's unclear?
```

Prices, rates, deadlines and commitments come from you, never from the model. A proposal that's been sent commits your organization.

Analyzing a document you've received

A call for tenders, a contract or a supplier quote can run dozens of pages. The assistant can pull out the essentials quickly, provided you work with the enterprise version of your tool if the document is confidential (chapter 9).

ANALYZING A CALL FOR TENDERS, A CONTRACT OR A QUOTE

```
Here is a document I received: [attach the document]. Bring out:

1. The requirements and expected deliverables.
2. All dates and deadlines.
3. Our obligations and those of the other party.
4. Unusual or risky clauses.
5. The questions to ask before signing or bidding.

For each point, cite the section of the document. If a piece of information isn't there, say
so.
```

This reading helps you prepare your questions. It doesn't replace a contract review by a legal professional.

KEY TAKEAWAYS

- Don't ask whether your proposal is good. Have a simulated client critique it.
- The assistant structures the proposal. Prices and commitments stay yours.
- For a document you've received, require the section reference for every point.

09

CHAPTER 09

Security, limits and good practices

What you need to know before using AI with real data

Where does your data go?

Two very different situations:

- **A consumer tool** (free version or personal account): your organization has no control over what's kept, and your exchanges may be used to train models depending on the settings. No identifiable client, employee or financial data.
- **An enterprise version** (for example, Copilot connected to your organization's account, or other vendors' enterprise offerings): your data is covered by your organization's agreement and, in general, isn't used to train models. Check the terms with your IT lead. It's the right place for work, but still send the minimum of personal information.

Prompt injection

Prompt injection is an instruction hidden in content you give the model: a web page, an email, a document. The model may follow it as if it came from you. The risk grows with tools that act on your behalf, for example by sending emails or modifying files.

According to OWASP, a reference in IT security, no method fully prevents this type of attack. You limit its impact by restricting what a tool can do without your approval and by keeping a human in the loop for important actions.

The risks and how to reduce them

RISK	RECOMMENDED MEASURE
Data leaks	Use the enterprise version provided by your organization. Send the minimum of personal information. Never use a personal tool for client or employee data.
Hallucinations	Check every factual claim in a primary source before using it. Ask the model to separate what it knows from what it assumes (chapter 3).
Prompt injection	Be wary of documents and pages from unknown sources. Don't let a tool send, publish or modify anything without your approval.
Outdated practices	Tools change often. Retest your important prompts after a major update and review your practices a few times a year.
Copyright	Treat generated content as a first draft. Don't reproduce other people's work as is. Check the tool's terms for commercial use.

Guideposts for Quebec organizations

- **Law 25:** it governs the protection of personal information in Quebec. Before entrusting personal information to an AI tool, especially if it processes it outside Quebec, your organization may need to carry out a privacy impact assessment. This guide is not legal advice: talk to the person in charge of protecting personal information in your organization and consult the Commission d'accès à l'information (cai.gouv.qc.ca).
- **Systematic verification:** every specific fact obtained from a model (statistic, date, quote, reference) is checked in a primary source. The appendix offers a checklist to keep near your desk.
- **Method before formula:** "secret" tricks circulate constantly on social media, often with numbers no one can trace. What comes up in serious research is less spectacular: clear context, examples, a task broken into steps and a check of the result. That's what this guide teaches, and it holds from one tool to the next.

KEY TAKEAWAYS

- Enterprise version for work, never a personal tool for client or employee data.
- Every factual claim from a model is checked before use.
- Be wary of spectacular numbers with no source, including those about AI itself.

10

CHAPTER 10

Make AI a habit

What makes the difference isn't the tool, it's what you do with it every week

The early dip

The first few times, using AI takes longer than doing the task the old way. You have to phrase the request, reread, correct. That's normal, and it's exactly where most people drop off: "I'd have been faster on my own."

The gain comes after a few weeks, when requests become natural and you have your own templates. Knowing this dip exists helps you get through it.

Pick a trigger

A habit sticks better when it's tied to a specific moment. Rather than "I'll try to use AI more often", pick a concrete trigger. A few examples:

- **Professional:** when an email thread goes past three messages, I have it summarized before replying.
- **Professional:** before sending an important text to a client, I run it through the demanding editor (chapter 6).
- **Manager:** after every meeting, I give my rough notes to the assistant to get the list of follow-ups.
- **Manager:** before an important decision, I launch an expert debate (chapter 3).
- **Administrative support:** at the end of the day, I ask which emails are still waiting for a reply.

Move forward in small steps

One technique a week, not everything at once. Keep a "My requests that work" file (Word, OneNote or your notes) and paste in every request that saved you time. You'll reuse it as is next time.

Block 15 minutes a week in your calendar to try something new, on a real task rather than an example.

As a team, it goes further

Adoption rarely happens alone. Share what works: an effective request sent to a colleague, five minutes at a team meeting, a shared folder of prompts. What works for one person often ends up serving the whole office.

My 4-week plan

WEEK	TECHNIQUE TO TRY	MY TRIGGER	WHAT I TAKE AWAY
1	Full context (chapters 1 and 4)		
2	Personalize my assistant (chapter 5)		
3	Have my texts edited (chapter 6)		
4	The Socratic method or self-critique (chapters 2 and 3)		

KEY TAKEAWAYS

- The first weeks are slower. That's normal, and it passes.
- A specific trigger beats a good intention.
- Keep your requests that work and share them with your colleagues.

11

CHAPTER 11

Ready-to-use prompt library

Copy, paste, adapt

Replace whatever is in square brackets with your content. These prompts work with Copilot, ChatGPT, Claude, Gemini or any other assistant.

WRITING AND COMMUNICATION

EXECUTIVE SUMMARY

```
Read the following document and summarize it in 200 words maximum. Structure: one sentence on the main issue, the three key points, then the recommendation or conclusion. The summary must be understandable to someone who hasn't read the document.

[Document]
```

DELICATE PROFESSIONAL EMAIL

```
I need to send an email about [delicate situation]. The recipient is [profile]. Our relationship: [description].

Write two versions: one direct and concise (under 150 words) and one more diplomatic (200 to 250 words). For each, briefly explain your choices.
```

SOCIAL MEDIA · SEVERAL PLATFORMS

```
Based on the following text, create a version for each of these platforms:

1. LinkedIn: professional tone, about 150 words, expertise angle.
2. Facebook: conversational tone, about 100 words, a question to spark discussion.
3. X: under 280 characters, strong hook.

Source text: [your text]
```

ANALYSIS AND DECISIONS

COMPARE TWO OPTIONS

```
I need to choose between two options.
Option A: [description]
Option B: [description]
Context: [issues, constraints, priorities]

Compare them on: (1) pros and cons, (2) short- and long-term risks, (3) resources required,
(4) reversibility. End with a well-argued recommendation.
```

PREPARE A NEGOTIATION

```
I need to negotiate [topic] with [profile of the other party]. My ideal outcome: [outcome].
My walk-away point: [limit].

Prepare me: (1) the other party's likely interests, (2) my best arguments, (3) the likely
objections and how to answer them, (4) a way to open the discussion, (5) the signs that it's
going well or badly.
```

MONITORING AND BUSINESS INTELLIGENCE

ANALYZE A COMPETITOR (WEB SEARCH ON)

```
Analyze the following competitor based on public information: [name and description]. Cite
your sources.

Analysis: (1) their positioning, (2) their target clients, (3) their strengths and
weaknesses, (4) the way they communicate, (5) what they do that we don't, and vice versa.

End with three implications for our own positioning.
```

INDUSTRY MONITORING (WEB SEARCH ON)

```
I want to keep track of what's happening in [your industry in Quebec]. Summarize the
developments of the last 7 days in the following areas: [3 to 5 areas].

For each development: (1) what happened, with the source, (2) the likely effect on [your
type of organization], (3) the level of urgency (monitor, analyze, act). Keep only the 5
most important.
```

TRAINING AND CHANGE MANAGEMENT

TRAINING PLAN · ADOPTING A NEW TOOL

I need to design a training plan for [number] employees who will use [tool] in their daily work. Their comfort with technology is [low, medium, high].

Propose a 4-phase plan over [X weeks] with: the goals of each phase, the training methods, the adoption indicators to track and ways to support the most reluctant people.

ANNOUNCE A CHANGE

I need to announce [nature of the change] to my team of [size and profile]. This change will affect [effect on daily work]. The resistance I anticipate: [your assumptions].

Propose a 3-phase communication plan (announcement, support, anchoring) with the key messages, the communication channels and the moments not to miss, for a Quebec SMB.

OPERATIONS MANAGEMENT

MAP A PROCESS

I want to map and improve the following process: [description of the process].

Step 1: ask me the questions you need to understand the current steps, the people involved, the tools used and the friction points. One question at a time. Step 2: once I've answered, describe the current process and its inefficiencies. Step 3: propose 3 improvement scenarios with an effort and impact assessment for each.

MEETING NOTES INTO FOLLOW-UPS

Here are my rough notes from a meeting: [your notes or the transcript]

Turn them into a summary of 5 points at most, then the list of follow-ups as a table: who, what, by when. If a deadline or an owner isn't mentioned, write "to be confirmed".

A

APPENDIX

Before you publish or send

Print it and keep it near your desk

Before sending, publishing or submitting a document prepared with the help of AI, check every box.

- [ ] Facts and figures are checked in a primary source.
- [ ] Calculations are redone or cross-checked.
- [ ] Cited sources have been opened and really say what we claim they say.
- [ ] No personal or confidential information is exposed unnecessarily.
- [ ] The tone and voice are mine, reread out loud.
- [ ] I can defend every sentence if questioned.
- [ ] My organization's AI use policy is followed (AI disclosure if required).

DécodAI

Technology moves fast. Strategy must stay human.

This guide suggests treating AI like a colleague: you give it context, clear instructions, and you check its work. Not like a search engine.

To go further, whether it's personalized support or a workshop for your team, book a meeting.

Let's take it further

Pick a time in my calendar from the Contact page of my website.

Book a meeting

DécodAI

Bruno Poirier
AI Adoption Specialist
bruno.poirier@decodai.ca
decodai.ca