

Prompt Engineering for Business School Students

This handout gives business school students a general introduction to prompt engineering. It is not about magic words. It is about **task design**: telling an LLM what job it should do, what information it may use, what good output looks like, and how the result should be checked. The techniques here are deliberately broad and durable. They work for case analysis, research synthesis, writing, spreadsheet support, interview preparation, and many other business-school tasks.

Core claim

Better prompts are clearer work orders. The best prompt is the one that most clearly defines the task, the context, the constraints, the output format, and the quality bar.

1. Define the job

What decision, output, or question matters?

2. Supply context

What materials, facts, or assumptions should the model use?

3. Set constraints

What should the model avoid, compare, or verify?

4. Shape the answer

What structure, tone, or table do you want back?

5. Review and iterate

What is weak, missing, or uncertain?

1. Why prompt engineering matters

Many students first meet LLMs through casual chat: ask a question, get an answer, move on. That can work for light tasks. It works much less well for high-value academic and professional work. In business school, many tasks are not single questions. They are **multi-step jobs**: summarize a case packet, compare alternatives, design an interview plan, critique a draft memo, explain a spreadsheet, or prepare for class discussion. Those tasks benefit from better prompting because they require structure, context, and judgment.

A useful mental model

Treat the LLM less like a search bar and more like a **trainable junior collaborator**. A junior collaborator can be fast, organized, and helpful. But that collaborator still needs a clear assignment, source boundaries, examples, and review.

Question	LLMs are usually good at	LLMs are weaker at
What can the tool do well?	Drafting first passes, organizing notes, summarizing long material, comparing alternatives, reformatting content, generating questions, and tightening prose.	Checking the truth of facts by itself, knowing which source is authoritative, doing long calculations reliably, or understanding hidden context that you never supplied.
How should a student respond?	Use the model to accelerate low-level work and broaden the first pass.	Verify material facts, own the final judgment, and be explicit about what the model may and may not assume.

Two habits to avoid

Slot-machine prompting: typing vague one-line prompts and hoping for brilliance.

Polish bias: trusting an answer because it sounds fluent, professional, or confident.

2. A simple example

A fast way to understand prompt engineering is to compare a weak prompt with a stronger one. Suppose a student is preparing a case memo for a hypothetical company, **Northstar Foods**, which is considering whether to launch a direct-to-consumer subscription business.

Weak prompt

Analyze whether Northstar Foods should launch a subscription business and write a recommendation.

Stronger prompt

You are helping me prepare a one-page class memo on Northstar Foods. Use only the case notes I provide below. First extract the facts that matter for the launch decision. Then compare the strongest arguments for and against the subscription model. Evaluate the decision on four criteria: customer value, economics, execution risk, and strategic fit. Return: (1) a 6-row fact table, (2) a decision table with pros, cons, and open questions, and (3) a recommendation in no more than 180 words. If the evidence is insufficient, say what is missing rather than inventing facts.

Why the second prompt works better

It narrows the task, specifies the source boundary, separates **fact extraction** from **interpretation**, defines the decision criteria, constrains the output format, and tells the model what to do when the evidence is incomplete.

What improved	What the stronger prompt added	Why that matters
Scope	One concrete memo for one decision	Reduces generic filler and keeps the answer relevant.
Source boundary	“Use only the case notes I provide”	Lowers the chance of fabricated outside facts.
Process	Extract facts, then compare, then recommend	Encourages cleaner reasoning and easier checking.
Decision rule	Four named criteria	Produces a more balanced recommendation.
Output shape	Tables + short recommendation	Makes the answer easier to review and reuse.
Uncertainty rule	“Say what is missing”	Rewards honesty instead of hallucination.

3. The five building blocks of a strong prompt

A useful prompt does not need to be long. It does need to be specific enough to remove avoidable ambiguity. Five elements do most of the work.

Element	What to specify	Quick example	Typical failure if omitted
Objective	The exact job to be done	“Compare two market-entry options and recommend one.”	The answer becomes broad, generic, or off-target.
Context and inputs	The materials, facts, audience, and setting	“Use only the interview notes and class case packet.”	The model guesses missing facts or uses the wrong frame.
Constraints	Limits, exclusions, tone, length, assumptions, or steps	“Do not use outside facts. Keep the memo under 250 words.”	The response becomes too long, too confident, or too open-ended.
Output format	The structure of the answer	“Return a table, then three bullets, then a recommendation.”	The answer may be hard to reuse or compare.
Quality control	How the model should handle uncertainty or weak evidence	“Label assumptions and open questions explicitly.”	The model fills gaps with plausible but unsupported claims.

A simple fill-in template

Role or job: You are helping me as a [research assistant / writing coach / case analyst / spreadsheet explainer].

Task: I need [specific deliverable or decision].
Inputs: Use [documents, notes, data, or facts].
Constraints: Do not [invent facts / use outside sources / exceed word limit / ignore counterarguments].
Output: Return [table / memo / bullets / outline / action list].
Quality bar: Flag uncertainty, label assumptions, and say what additional information would most improve the answer.

A useful rule of thumb

As the task becomes more expensive, more visible, or more consequential, your prompt should become more explicit. A short prompt may be fine for a quick rewrite. A board memo, case recommendation, or research brief needs a much tighter work order.

4. Seven durable prompt-engineering techniques

The techniques below are simple, widely useful, and effective across most modern LLM systems. None depends on one special platform or one secret phrase.

4.1 Technique 1. Ask for clarification before the answer

When the task is underspecified, the model should not race to completion. It should first identify the missing information that would change the result.

When to use it

Use this when the task depends on audience, goal, format, or decision criteria that the model cannot infer safely.

Prompt pattern

Before answering, ask up to five clarifying questions that would materially change the output. If the task is already clear enough, state your assumptions briefly and proceed.

4.2 Technique 2. Separate extraction from interpretation

Many bad answers mix evidence gathering, analysis, and recommendation too early. A better workflow asks the model to extract facts first, then interpret them, then decide.

Why this works


Fact extraction is easier to verify than interpretation. Once the factual base is clean, later reasoning is usually stronger and less error-prone.

Prompt pattern

Step 1: extract the relevant facts into a table. Step 2: identify the major themes or trade-offs. Step 3: provide the recommendation. Do not skip directly to Step 3.

4.3 Technique 3. Break large tasks into stages

One enormous prompt often produces one enormous answer. A staged workflow usually produces better work. This is sometimes called **prompt chaining**. Instead of asking for everything at once, ask for the next useful artifact.



Prompt pattern

For this project, do not produce the final memo yet. First return the 10 most important facts in a table with categories and source labels. After I confirm the table, we will move to option analysis.

4.4 Technique 4. Ground the answer in provided material

LLMs are much more reliable when the relevant material is in the prompt or attached context. If you care about factual accuracy, tell the model what sources it may use and how it should handle missing evidence.

Good grounding habits

Name the material. Delimit it clearly. Ask for quoted evidence when wording matters. Tell the model to mark unknowns instead of guessing.

Prompt pattern

Use only the meeting transcript and the market research notes below. When a claim matters, quote the relevant sentence or number. If the packet does not support a conclusion, label it as an open question.

4.5 Technique 5. Request structured outputs

A strong answer is easier to inspect when it comes back in a useful shape. Tables, labeled bullets, issue trees, action lists, and memo outlines are often better than dense paragraphs.

Why structure matters

Structure improves review, comparison, and transfer into slides, spreadsheets, and written assignments.

Prompt pattern

Return the answer in this order: (1) a table with issue, evidence, implication, and uncertainty; (2) three prioritized recommendations; (3) two risks that could change the recommendation.

4.6 Technique 6. Use examples, rubrics, or success criteria

If you know what good output looks like, tell the model. A short example or rubric often improves results dramatically, especially for writing or classification tasks.

What to provide

A good sample, a grading rubric, a scoring scheme, or a simple “good vs weak” contrast can all help.

Prompt pattern

Write the response in the style of a concise consulting note. A strong answer should be specific, source-grounded, balanced, and no longer than 300 words. Avoid slogans, repetition, and unsupported claims.

4.7 Technique 7. Ask for critique, not just completion

A useful LLM is not only a drafter. It can also act as a reviewer. After a first pass, ask the model to identify weak logic, missing evidence, risks, or counterarguments.

A powerful habit

Use one prompt for production and a separate prompt for criticism. The second prompt should behave like a skeptical reviewer, not a cheerleader.

Prompt pattern

Act as a skeptical reviewer. What are the five weakest claims in this memo? For each one, explain why it is weak, what evidence is missing, and how the claim should be revised.

5. How these techniques show up in common business-school tasks

Prompt engineering is most useful when tied to real assignments. The table below shows how the same core techniques can travel across typical business-school work.

Task	Best prompt shape	Helpful output	Main watch-out
Case discussion prep	Extract facts first, then compare options, then draft a position	Fact table, issue tree, recommendation bullets, likely cold-call questions	Jumping straight to a recommendation without using the case packet carefully
Research brief	Ground to named sources and force citation or quotation	Source table, theme summary, open questions, contradictions	Mixing reliable sources with weak web summaries or uncited claims
Spreadsheet or model support	Ask for formula logic, assumptions, and interpretation separately	Line-by-line explanation, error checklist, sanity checks, scenario ideas	Treating the model as a calculator instead of a reasoning and explanation tool
Writing and slides	Provide audience, tone, structure, and a quality bar	Executive summary, slide headlines, tightened prose, speaker notes	Producing smooth but generic text with no real point of view
Interview preparation	Specify role, company context, and likely question types	Question bank, STAR story refinement, follow-up questions, feedback	Memorizing generic answers that do not sound like you

5.1 A sample workflow for case analysis

One workable chain for a case packet

Prompt 1 - Extract: Pull the 10-12 most decision-relevant facts from the packet.

Prompt 2 - Compare: Build a table comparing the main alternatives on named criteria.

Prompt 3 - Draft: Write a short memo using the fact table and comparison table.

Prompt 4 - Review: Critique the memo, identify weak claims, and suggest the strongest revision.

6. Common failure modes and how to fix them

Failure mode	What it looks like	Better move
Vague objective	“Help me with this case”	Name the exact deliverable: memo, table, slide outline, recommendation, or question list.
Too many jobs at once	One prompt asks for research, analysis, writing, and formatting all together	Split the work into stages and confirm the early outputs before moving on.
No source boundary	The model invents facts or blends outside material with your packet	Tell the model which sources it may use and what to do when the evidence is thin.
No output specification	The answer is long, hard to scan, and hard to reuse	Ask for a specific structure such as a table, memo, or checklist.
No verification step	The work sounds polished but contains quiet errors	Add a reviewer pass that challenges claims, assumptions, and missing evidence.
Treating fluency as truth	The answer feels professional, so the student stops checking	Verify the facts, question the logic, and compare against the underlying material.
Over-engineering simple tasks	A tiny job gets a giant prompt that is slow and awkward	Match prompt complexity to task complexity. Keep small tasks small.

One important limit

Prompt engineering does not remove the need for judgment. It improves the quality of the first pass and the quality of the interaction. It does not turn an LLM into an authority on facts, a substitute for domain knowledge, or a replacement for checking your work.

7. A compact prompt library

The templates below are short on purpose. They are starting points, not magic scripts.

7.1 Template 1. Decision memo prompt

Reusable template

You are helping me prepare a decision memo for [audience]. Use only the material I provide. First extract the facts that matter most. Then compare the leading options using these criteria: [criteria]. Return (1) a fact table, (2) an option-comparison table, and (3) a recommendation of no more than [length]. Label any assumption or open question explicitly.

7.2 Template 2. Source-grounded research prompt

Reusable template

Use only the sources below. Build a table with four columns: source, key point, why it matters, and uncertainty or conflict. After the table, give a brief synthesis in plain language. Do not add outside facts. If the sources conflict, say so clearly.

7.3 Template 3. Spreadsheet or technical explainer

Reusable template

Explain this spreadsheet section to a business-school student. Start with the purpose of the block, then explain each major line item, then list the three most common modeling mistakes. Keep the explanation practical rather than mathematical unless I ask for more detail.

7.4 Template 4. Writing improver

Reusable template

Revise the text below for a [CEO / professor / recruiting manager / client] audience. Keep my main argument, remove repetition, tighten the logic, and improve clarity. Preserve any numbers or factual claims exactly. After the rewrite, give a short note on the three biggest changes you made.

7.5 Template 5. Red-team reviewer

Reusable template

Act as a skeptical reviewer. Identify the weakest claims, hidden assumptions, and missing evidence in the material below. Return a table with claim, concern, why it matters, and suggested fix. Focus on substance rather than grammar.

8. A practical workflow students can reuse

A five-step operating rhythm

1. **Frame the job.** What is the deliverable?
2. **Supply the packet.** What material may the model use?
3. **Get the first artifact.** Ask for the next useful output, not the whole project.
4. **Review critically.** Use a second prompt to challenge the first pass.
5. **Capture what worked.** Save strong prompts and revise weak ones.

A practical student habit is to keep a tiny prompt log. After an important task, note the prompt you used, what worked, what failed, and what you changed. Good prompting skill compounds because students stop rewriting from zero every time.

Mini prompt-log template

Task: What job was I trying to do?

Prompt version: What did I actually ask?

What worked: Which part of the answer was useful?

What failed: What was vague, wrong, or too generic?

Revision: What would I change next time?

Verification note: What facts or claims still required checking?

9. What not to obsess over at the start

Business-school students sometimes jump immediately to advanced settings, model variants, or internet debates about secret prompting tricks. Those things can matter later. For most students, the highest-return improvements come earlier.

Focus on first	Usually later or optional
Clear objective, source boundaries, output structure, critique loop, and verification habit	Exotic prompt tricks, highly specialized syntax, or tool settings you do not fully understand

A strong beginner prompt is usually not glamorous. It is simply clear, scoped, source-aware, and easy to review.

10. Conclusion

Prompt engineering is best understood as a practical management skill. You are assigning work to a fast but imperfect collaborator. The more clearly you define the job, the evidence, the output, and the review standard, the more useful that collaborator becomes. For business-school students, this matters because many class and career tasks are not about generating text alone. They are about extracting signal, comparing alternatives, presenting judgment, and doing so with enough discipline that the final work deserves trust.

Final takeaway

Do not ask only for answers. Ask for **better work processes**: fact extraction before interpretation, structure before prose, critique before submission, and verification before trust.

Appendix A. A one-page prompt canvas

Field	Fill-in question
Task	What exactly do I need the model to produce?
Audience	Who is this for, and what tone or level is appropriate?
Inputs	What documents, notes, or facts may the model use?
Source boundary	What may the model <i>not</i> use?
Criteria	How should the model judge alternatives or prioritize ideas?
Output shape	What structure should the answer have?
Constraints	What limits on length, time, scope, or style matter?
Verification	What assumptions, uncertainties, or open questions must be labeled?
Next step	What is the next artifact after this one?

Appendix B. Practice exercises

1. Take a weak prompt you have actually used. Rewrite it using the five building blocks from Section 3.
2. Choose a case or article from class. Write one extraction prompt and one red-team prompt for it.
3. Ask an LLM to improve a short paragraph of your own writing. Then ask a second prompt to explain what changed and why.
4. Build a two-step prompt chain for a business-school task you face this week: first artifact, then review artifact.

Appendix C. A short self-check before submission

Before relying on an LLM-assisted output, ask:

- Did I give the model enough context to do the job well?
- Did I specify what good output looks like?
- Did I separate evidence gathering from judgment where that matters?
- Did I ask for assumptions, uncertainties, or missing information to be labeled?
- Did I verify the facts that actually matter?