

01 — MENTAL MODEL

What an LLM Actually Is

Maximally Agreeable	It confirms whatever frame you set — even wrong ones. Set the right frame first.
Fills Gaps With Plausible	Vague prompt → guesses for the average person. Be specific to get answers for <i>you</i> .
Mirrors Your Precision	Vague in → vague out. Specific in → specific out. This is the one rule above all others.

The GIGO Law

Output quality is bounded by input quality. When the answer is wrong, 90% of the time the problem is in your prompt — not the model.

02 — THE 6-LAYER ARCHITECTURE

Complete Prompt Structure

GOAL	Exact outcome. Surgical, not vague. "10 MCQs on French Revolution for entrance exam" not "help me study."
CONTEXT	Who you are, your situation, your constraints. The model knows nothing about you unless you say it.
TASK	Exact action verb: explain / generate / evaluate / compare / critique / summarise / roleplay
CONSTRAINTS	What NOT to include. Most underused layer. Exclude what you know, what's too basic, what format you hate.
OUTPUT FORMAT	Exact structure: table / bullet / prose / numbered steps / JSON / code / spoken delivery.
META	How to think: "be brutally honest / depth over breadth / if uncertain, say so / assume I'm smart but new."

GOAL: [Exact outcome needed]

CONTEXT: [Who I am / background]

TASK: [explain / generate / compare...]

CONSTRAINTS: [What to avoid / exclude]

OUTPUT FORMAT: [Exact structure]

META: [How to think / tone / depth]

03 — DEPTH TECHNIQUES

Use Selectively — Not Every Prompt

Assumption Audit	After any answer: "What assumptions did you make? List them." Challenge ones that don't fit your situation.
Steel-Manning	Ask for the strongest case against your position. Then critique that argument. Best way to stress-test any idea.
Constraint Tightening	Generic prompts → generic answers. Tight constraints force specificity. Like adjusting a camera lens — tighter = sharper.
The Failure Frame	"What are the 3 most likely ways this fails?" — primes edge-case thinking before asking for solutions.
Socratic Chain	3 focused sequential prompts drill deeper than 1 giant one. Each answer informs the next question.
Pre-Mortem	"Assume this fails in 6 months. What were the 3 most likely reasons?" Surfaces risk before it occurs.
Calibrated Uncertainty	"Tag each point: highly confident / moderately confident / uncertain." Stops AI speculation from looking like fact.

05 — THE ITERATION LOOP

Every First Output Is a Draft

- 1 First Pass** — Use your 6-layer prompt. Read the entire response before reacting.
- 2 Gap ID** — What did it miss? What assumption is wrong for your situation?
- 3 Targeted Fix** — Name the gap explicitly. "You assumed 3 months. I have 3 days. Revise." Never say "make it better."
- 4 Challenge** — "What is the strongest reason this advice is wrong?"
- 5 Synthesise** — Ask it to merge original answer and counter into one nuanced position.
- 6 Extract** — "What is the single most important takeaway?" Lock in the output before moving on.

The difference between mediocre and excellent AI use is almost entirely in what happens AFTER the first response.

04 — ADVANCED ROLE PROMPTING

3 Layers of a Role

IDENTITY	Title + seniority + domain. "Senior SWE at B2B SaaS" — not just "an engineer."
STAKES	What do they optimise for? "You care about reliability and cost — not cleverness."
RELATIONSHIP	"You are interviewing me. Be skeptical. Push back on vague answers."

Multi-Role Panel

Assign 3 roles, same problem, different angles. For each: 2 concerns + 1 compelling point. Then: biggest conflicts across views.

Adversarial Critic Setup

Role:	Brutally honest editor.
Job:	Find every weak point. Don't soften.
Flag:	Vague claims / unsupported assertions / logical gaps.
End with:	3 specific rewrites of the weakest parts.

06 — RESEARCH-LEVEL OUTPUT

Research Brief Template

TOPIC	Specific subject (not "AI" — "AI in B2B SaaS pricing, 2024-25")
PURPOSE	What decision does this inform?
DEPTH	Surface overview / working knowledge / expert depth
FOCUS	What specific angle matters most for my purpose?
EXCLUDE	What I already know / what is irrelevant to my purpose

Source Check "For each major claim: is this well-established consensus, emerging consensus, or contested?"

Assumption Stress "What changes if [key assumption] were false?" Reveals how fragile conclusions are.

07 — OUTPUT CONTROL

Fix Bad Outputs Fast

Too Long	"Remove every sentence that could apply to any answer on this topic. Keep only what's specific to my situation."
Too Surface	"Go one layer deeper on point 3. Give me the mechanism, not the summary."
Too Generic	"Rewrite with a real-world example at the centre of each point. No abstract claims without an anchor."
Wrong Tone	"Rewrite as natural spoken language. First person. No bullets. Under 150 words."
Missed Point	"You answered a slightly different question. My actual question was [restate]. Answer only that in 3 focused bullets."

08 — THE TOP 1% MINDSET

What They Do Differently

Use AI to Think, Not Just Produce
"What am I missing?" "Challenge this assumption." They interrogate — not just consume.

Front-Load Context Ruthlessly

Never make the model guess. The model just met you. Tell it everything it needs to help *you specifically*.

Iterate with Surgical Precision

Never "make it better." Name exactly what is wrong: "Your answer ignores [X]. Fix that one thing."

Treat Conversations as Sessions

Each prompt builds on the last. Reference earlier outputs explicitly — don't restart from zero.

Use Constraints as a Superpower

Tighter constraints → more specific output. Don't fear limiting the model. Limits force precision.

Audit Outputs Critically

Ask: What did it assume? What did it skip? What sounds confident but might be wrong? Every output is a draft.

Know When NOT to Use AI

Don't outsource judgment calls that require lived experience or personal values. Use AI for leverage, not replacement.

09 — DOMAIN PLAYBOOKS

Chat · Code · Image/Video

CHAT AI — CLAUDE, GPT, GEMINI

Set context first (1 line changes everything). Continue threads — don't restart. Ask it to poke holes in your plan, not just build it.

Power Move: "Before answering, ask me any clarifying questions you need. Don't assume — ask first."

CODE AI — CURSOR, COPILOT

Give architecture context before any feature. Specify scope: "Only modify this function." Ask for approach before code. Ask it to PR-review its own output.

Power Move: "Review this as a senior engineer. Flag: bugs, security issues, performance concerns, anything not production-ready."

IMAGE/VIDEO — MIDJOURNEY, SORA

Describe what is *seen*, not what is felt. Layer: subject → style → lighting → mood → camera. Add negative prompts. For video: describe motion explicitly.

Power Move: Find a reference image, paste into Chat AI: "Describe this image as a generation prompt." Use that as your starting point.

The One-Line Summary

Top 1% prompters aren't people who know more tricks. They are more precise about what they want, more honest about what they don't know, and more disciplined about iterating until the output is actually useful.