



Stop Typing Prompts That Get You Mush

The 20-minute reality check: the 5-part formula behind every great AI prompt, plus a taste of the 200-prompt library — for ChatGPT, Claude & Gemini.

WHAT'S INSIDE

Everything that matters, in pointers

This mini gives you the honest map in 20 minutes: why prompting is now a measurable skill, the 5-part anatomy behind every great prompt, and a working sample from the 200-prompt library. Read it once and you'll write better prompts than most people who use AI daily.

01	Why prompts matter — the adoption data	p.3
02	The anatomy of a great prompt	p.5
03	The Work Library — email, reports, meetings	p.7
04	The Career Library — resumes, interviews	p.9
05	Marketing, sales & analysis, teased	p.11
06	Advanced patterns — chaining & self-critique	p.12

THE ONE IDEA TO HOLD THE WHOLE WAY

Every great prompt — and every framework with a fancy acronym — is the same five parts wearing different clothes: Task, Context, Format, Examples, Constraints. Learn the skeleton once and you can build any prompt yourself, forever.

Why prompts matter — the adoption data

- **Adoption crossed 50%.** Gallup's Q1 2026 survey of 23,717 U.S. employees found 50% now use AI at work — first time past halfway, up from 46% one quarter earlier.
- **The stricter count says 21%.** Pew Research counts more conservatively. Both surveys have more than doubled in a year — adoption is no longer the edge. Skill is.
- **The Harvard/BCG proof:** 758 consultants with AI access completed 12.2% more tasks, 25.1% faster, at 40% higher quality — but only inside AI's "capability frontier."
- **Outside that frontier:** the same consultants scored 19 percentage points worse than colleagues working unaided. The tool that made them brilliant on one task made them confidently wrong on the next.
- **The missing 12 words:** successful workplace prompts average ~21 words with context; most people type ~9. Chapter 2 turns that gap into a formula.

REALITY CHECK

One experimental arm in the Harvard/BCG study received prompt-engineering training before the tasks. This is a learnable skill — and the learning takes hours, not years.

CHAPTER 01 · CONTINUED

The adoption gap, in numbers

50%

of U.S. employees use AI at work (Gallup, Q1 2026)

28%

use it daily or weekly — an all-time high

+40%

higher quality output for skilled prompters (Harvard/BCG)

Inside firms that push adoption it's further along still — 72% of McKinsey's 45,000 employees actively use its internal AI tool. The gap between "uses AI" and "prompts well" is now the real productivity divide.

- **Three habits past their sell-by date:** "think step by step" (redundant on reasoning models), heavy personas (theatrical costumes over substance), and lists of "don't do X" (models handle negation poorly).
- **The pink elephant problem:** "Don't think of a pink elephant" — you just did. Every word in a prompt, even a banned one, pulls output toward it. Phrase behavior positively.

REMEMBER

Tens of millions of people now type prompts every week — and the tools reward the minority who ask well.

The anatomy of a great prompt

OpenAI, Anthropic and Google publish separate prompting guides. Strip the branding and they agree on one skeleton:

- **Task:** a verb + a deliverable — "Write a 150-word email that..." Anthropic's guidance is blunt: give a direct action verb, not "can you help me with..."
- **Context:** who you are, what this is for, and the pasted source material. Fence pasted text in `"""triple quotes"""` so instructions and data never blur.
- **Format:** the container — table, 5 bullets, email, one paragraph. No format spec = a wall of text.
- **Examples:** the single highest-leverage part. Start with one; add more only if results stay inconsistent.
- **Constraints:** tone, audience, length, exclusions — your quality bar stated up front.

BEFORE → AFTER

Before: "Write an article about remote work." **After:** "Write a 700-word opinion piece arguing hybrid beats full-remote for junior staff. I run a 40-person design agency... Blog post: strong hook, 3 arguments... Audience: startup founders. Plain English." Same ask, all five parts filled in.

CHAPTER 02 · CONTINUED

Frameworks & the honest verdict

FRAMEWORK	BEST FOR	WEAKNESS
RTF	Quick jobs: emails, posts	Too thin for complex work
CO-STAR	Audience-facing content	Six blanks for a two-line reply
CRISPE	Analysis & creative work	Clunkiest acronym of the three

- The honest take:** frameworks are memory aids, not magic. They all work on ChatGPT, Claude and Gemini alike because what matters is the five components — not the acronym.
- System prompts for civilians:** a standing instruction the model reads before every message. Set once in ChatGPT's Custom Instructions or a Claude Project, and every future prompt gets shorter.
- The 8 mistakes that ruin most prompts:** vague asks, no format spec, single-shot expectations, no context, kitchen-sink prompts, negative-only instructions, trusting confident output blindly, and not saving your winners.

DEBUGGING MINDSET

"Treat prompts like code: test edge cases, fix the failing part, save the version that works."

The Work Library — a working sample

Thirty copy-paste prompts live in the full book for email, reports, meetings and productivity. Here are three, ready to use today:

- **The Difficult-Email Reply:** "Here is an email I received: [PASTE]. Write a reply that [DECLINES / PUSHES BACK / ASKS FOR X] while keeping the relationship warm. Under 120 words. Apologize at most once."
- **The CFO Summary:** "Summarize this for a CFO with 60 seconds: [PASTE REPORT]. Exactly 3 bullets of findings — each with its number — then 1 risk, then 1 recommendation."
- **The Action Extractor:** "From this transcript, extract every action item as a table: Owner | Task | Deadline | Source quote. Only include actions actually stated. Mark inferred owners as [unconfirmed] — never guess."

THE 3 EMAIL RULES

Every email prompt gets a word cap, no exceptions. Ban exclamation marks explicitly. Higher stakes + colder audience = more structure.

Reports execs actually read

The pattern for every summary prompt: set the audience altitude, set the length budget, and say what to preserve — otherwise the model compresses the wrong things.

- **The Preserve-the-Numbers Summarizer:** preserve every figure, owner and date exactly as written; decisions and numbers first, narrative second — left alone, models compress numbers first.
- **The Sunday Planner:** time-block your week, put the hardest work in your peak hours, and name the ONE task to protect if everything goes wrong. The highest-leverage prompting ritual in the book.

PRO INSIGHT

Chain prompts 12 → 13 → 11 in the full book: raw notes become a status report, a "so-what" pass sharpens it, then a CFO summary tops it. Big deliverables are chains, not single prompts.

The Career Library — the honest warning

- **The most dangerous prompt in the book:** "write my resume." Used raw, AI costs people interviews — ~49% of hiring managers auto-reject resumes they suspect are AI-generated.
- **The #1 tell:** generic filler — "results-driven," "proven track record," "dynamic professional." Harvard's career service line: AI for tailoring and editing is endorsed; wholesale generation is the fast lane to rejection.
- **The JD Keyword Extractor:** "Extract the 10 skills and keywords this JD screens for, ranked by prominence. Split into must-haves and nice-to-haves, quote the exact JD phrase for each." A JD is a scoring rubric wearing prose.
- **The Bullet Upgrader rule:** rewrite as verb + what I did + measurable result — using ONLY facts you give it. Where a result needs a number you haven't provided, insert [METRIC?] — never estimate or invent one.
- **The Mock Interviewer:** the model plays hiring manager, asks one question at a time, pushes back on vague answers, then scores you against the JD's top requirements after 6 questions.

CHAPTER 04 · CONTINUED

Edit-and-match, never generate

	THE GENERATION TRAP	EDIT-AND-MATCH
Prompt	"Write me a resume for this job"	Decode the JD first
Numbers	Invented metrics	[METRIC?] placeholders
Letters	Same letter, every company	Your voice, cloned from writing

- The Specific-Hook Letter:** match the voice of 2–3 of your real writing samples, open with a true hook — never "I am writing to apply" — and map two proof points to the JD's top requirement.
- The Salary Negotiation Script:** anchors your ask at your target without an ultimatum, gives calm responses to "that's above our budget," and a polite walk-away line.

Marketing & sales — teased

- **Ads live or die on constraints:** "Write 10 Google Ads headlines under 30 characters... 5 urgency-led and 5 benefit-led, with character counts next to each." Hard limits force genuine variety.
- **Voice lives in examples:** "Here are 2 product descriptions in our brand voice: [PASTE BOTH]. Write one for this new product in exactly that voice." Two pasted samples beat two hundred adjectives.
- **Cold outreach math:** trigger-event hooks reply at ~2.3× the rate of generic problem statements; personalization roughly doubles reply rates; follow-ups carry 42% of all replies.
- **The Objection Dojo:** name what the prospect really means, a 2-sentence acknowledge-and-reframe, one proof point, then roleplay it with the model pushing back twice.

THE CHAPTER'S ONE RULE

Never publish marketing copy without reading it aloud once. AI drafts fast; your ear catches what no prompt can.

Analysis & the business library

- **The Excel Formula Builder:** describe your columns precisely, name your Excel version, and demand the formula, a plain-English explanation, and 2 failure cases — the model's #1 failure is guessing your layout.
- **The [VERIFY] convention:** "Mark every specific claim — numbers, dates, names — with [VERIFY]." Converts hidden hallucination risk into a visible checklist, used across research AND business prompts.
- **The Actionable SWOT:** a SWOT where every item is company-specific — no "strong brand" generics — then 2 rated initiatives per quadrant, each with effort and impact scores.
- **The support reply skeleton:** validation → explanation (no excuses) → action with timeline → warm close. "Kind but professional" measurably outperforms plain "professional."
- **Legal-ish drafts, boundaries in bold:** AI is for understanding contracts and preparing drafts — never final legal documents. A licensed attorney always reviews before anyone signs.

GOLDEN RULE

Test any AI formula on 10 rows before applying it to 10,000. Ten rows cost a minute; ten thousand wrong rows cost a weekend.

Advanced patterns

- **Prompt chaining:** break big deliverables into research → outline → draft → edit, where each prompt consumes the last output. Quality holds at every step.
- **Persona + audience:** who the AI pretends to be matters less than who the output is for. "Explain this" and "explain this to a first-time manager" produce different documents.
- **Self-critique loops:** generate → critique → revise reportedly improves output quality 10–25%. "Do not rewrite yet" during critique keeps the feedback honest.
- **Style mimicry:** 3–5 samples + a written style breakdown reach 80–90% perceived similarity in blind tests; adjective-only instructions ("friendly, casual") reach ~40%.

META-PROMPTING — THE CLOSER

"Ask me 5 clarifying questions — one at a time — about context, audience, format, constraints and success criteria. Then write the ideal prompt." The book that extends itself.

Long-context rules & the whole book

- **Lost-in-the-middle:** models attend best to the start and end of what you paste — performance drops 13.9–85% as input grows, even within claimed limits (Stanford study).
- **The 3 rules:** instructions at top and bottom, paste the 5 relevant pages — never all 100 — and trim boilerplate before pasting.
- **Model notes:** ChatGPT — state the outcome, don't force "think step by step." Claude — strongest instruction adherence, great with pasted source material. Gemini — largest context, native web grounding.
- **Same prompt, different answers:** normal, not a bug — models sample from probabilities. When output matters, run it in two models and keep the better draft.

WHOLE-BOOK TAKEAWAY

Every great prompt = task + context + format + examples + constraints. No prompt for your task? Meta-prompting writes it for you.

YOU'VE GOT THE FORMULA. NOW GET THE LIBRARY.

This was the 20-minute version. The full guide is 200 prompts deep.

You now write better prompts than most daily AI users. The complete **Prompt Engineering Cheat Book** takes every chapter above and gives you the full 200-prompt library, organized by the work you actually do.

- **All 8 chapters** — the full anatomy, plus Work, Career, Marketing, Analysis, Business and Advanced Pattern libraries.
- **200 battle-tested prompts** across 12 domains — email, resumes, ads, Excel, contracts, strategy decks and more.
- **The complete debugging guide** — all 8 mistakes that ruin prompts, with the fix for each.
- **Model notes & FAQ** — ChatGPT vs Claude vs Gemini conventions, tested as of 2026.
- **The advanced-pattern toolkit** — chaining, self-critique loops, voice cloning and meta-prompting, fully worked.

UPGRADE TO THE FULL EDITION

~~₹199~~ **₹99**

Complete 200-prompt library · instant download

SCAN / TAP
to upgrade

GO FURTHER, PAY LESS

Two ways to keep going

THE AI & TECH STACK BUNDLE

Prompt Engineering + AI Engineer Roadmap + Python + Data Science

Everything you need to work faster with AI, in one download.

~~₹396~~ **₹199**

THE COMPLETE LIBRARY

All 30 Skill Series books

Spirituality, money, career, health, tech — the entire collection, one price.

₹499

GET EVERY NEW BOOK FREE

Follow [@weknowwhatyourelookingfor](#) and join our WhatsApp list — buyers get first access and reader-only pricing on everything we publish next.

weknowwhatyourelookingfor.com

COPYRIGHT & LICENSE

© 2026 weknowwhatyourelookingfor.com — All rights reserved.

This copy is licensed for **personal use only**. Reproduction, resale, sharing, or redistribution of this book in whole or in part, by any means, is prohibited without written permission.

NOTICE

This guide is for **general educational purposes only**. Individual results vary with effort and circumstances, and no specific outcome is warranted.

Published by weknowwhatyourelookingfor.com · For queries, corrections, or licensing, visit weknowwhatyourelookingfor.com.