Words to Prompt AI Well With

How to Prompt AI for Better Answers: A Beginner-Friendly Guide to Asking Questions That Get Better Results

---

by Douglas B.

Words to Prompt AI Well With

Copyright © 2026 Douglas B.
All rights reserved. No part of this book may be reproduced, stored, or transmitted in any form or by any means — electronic, mechanical, photocopying, recording, or otherwise — without the prior written permission of the author, except for brief quotations used in reviews.
This book was written and illustrated with the assistance of AI tools, working from the author's own experience, examples, and editorial direction. The lessons in it come from real work; the words are offered the same way they were learned — honestly, and without embarrassment.
Cover and interior art are variations on an original painting by Douglas John Brownfield, which appears once in this book exactly as itself.

ISBN — [ISBN — KDP, to be assigned]

Contents

Introduction 1
How This Book Got Learned 3
Part One — Sharpen 8
Specify · Quantify · Itemize · Compare · Rank · Prioritize · Exemplify · Define · Pinpoint · Probe · Anchor · Target
Part Two — Constrain 21
Constrain · Exclude · Assume · Bracket · Budget · Limit · Simplify · Localize · Timebox · Role · Filter · Scope
Part Three — Structure 33
Sequence · Outline · Tabulate · Summarize · Compress · Expand · Decompose · Recompose · Stage · Chunk · Layer · Map
Part Four — Challenge 45
Steelman · Invert · Critique · Defend · Disagree · Counterfactual · Anticipate · Simulate · Stress-Test · Falsify · Worst-Case · Rebut
Part Five — Verify 58
Verify · Audit · Cite · Cross-Check · Trace · Estimate · Flag · Reconcile · Sanity-Check · Grade · Refresh · Benchmark
Part Six — Create 71
Draft · Invent · Narrate · Translate · Rewrite · Vary · Iterate · Formalize · Brainstorm · Voice · Combine · Script
Recipes 84
Prompt Templates — Ready to Copy 88
The Practice Room 94
Field Guide — If Your Problem Is… 97
Quick Reference 98
A Word at the End 101
About the Art 102
About the Author 103

Introduction

Most people blame the machine when an answer comes back bland, vague, or quietly wrong. Usually the machine is innocent. It answered the question it was actually given — and the question it was given was a shrug dressed up as a sentence. "Tell me about mortgages" is not really a question. It is a topic, handed over with no size, no angle, and no purpose, and the answer it earns is exactly as shapeless as the ask.

Here is the good news this whole book rests on: you do not need longer prompts. You need stronger words. A single well-chosen verb — *rank*, *exclude*, *steelman*, *audit* — does more work than three paragraphs of background, because a verb tells the machine what to **do**, not just what to talk about. Qualifiers do the same quiet work. "For a first-time buyer in Ohio, in under two hundred words, with the costs in a table" is not a longer request; it is a *placed* one. The answer improves because the aim improved.

This book collects seventy-two of those words, grouped by the job they do: sharpening a question, constraining an answer, structuring it, challenging it, verifying it, and creating with it. Each entry shows a prompt twice — once the way a beginner types it, once with the word in it — and then does it again with a second example from a different corner of life, so you can see the difference instead of being told about it — and each entry closes with a short note on how the word itself can go wrong, and one sentence that places it perfectly. You do not need all seventy-two. Most people find that six or eight of them transform nearly every conversation they have with an AI.

A note on the title, since a careful reader will have noticed it ends with a preposition. It does. Some rules exist to make language clearer, and some exist only to make writers nervous. This book is firmly about the first kind.

Use it like a toolbox, not a textbook. Keep it near wherever you type. When an answer disappoints you, do not start over — reach for a word.



The original painting, by Douglas John Brownfield — ink on canvas.
It appears in this book once, exactly as itself.

How This Book Got Learned



This book was not researched. It was accumulated — the way a toolbox fills up, one tool at a time, each one bought because a job went badly without it. Nearly every word in the chapters ahead earned its place in a real working day, in a real two-person shop. This is the short version of how.

The two people are Douglas — the author of this book, who calls the shots — and his assistant, an AI he works with the way a builder works with a tireless apprentice. Douglas runs a small crowd of projects at once: a daily video that has to exist every morning, little online shops, designs, documents, plans that span months. He has exactly one member of staff, and it is the machine. He works from his phone, in short messages, at hours most offices would call unreasonable. The assistant, for its part, works instantly, confidently, and — at the start — far too literally. It was a productive partnership and a recurring misunderstanding, and both halves of that sentence taught them things.

It did not begin with lessons. It began, the way these things usually overpromise, with a honeymoon. In the first week or two, the two of them plowed straight through two entire builds — real ones, with parts and wiring and deadlines of their own — at a pace that startled them both. Douglas would sketch what he wanted in a few short messages; the assistant would return it built, tested, and tidier than the sketch. Back and forth, day after day, and nearly everything clicked. If you had asked either of them that fortnight, they would have told you the whole arrangement was simple: one of them knew what he wanted, the other could make it, and the messages in between were just the road the work traveled on.

Speed, it should be said, is a persuasive storyteller. Each finished build made the next request a little shorter and a little looser, because why spell out what the other fellow so obviously already understood? Douglas' messages trimmed themselves down to fragments — a phrase, a direction, half a thought typed between other jobs. The assistant, meeting a fragment, did what it always does: it completed the pattern, smoothly and silently, and most days the completion happened to be right. Neither of them noticed the shorthand being born. It felt like fluency. On the good days it even was fluency. The bill for the difference between the two only arrived later, in installments, and it is most of what follows.

They were both half right, which is the most durable kind of wrong. The early builds were not easy because the partnership was magic. They were easy because they were, underneath the excitement, comparatively simple: clear edges, few moving parts, nothing yet standing on top of anything else. Success, arriving that fast, does a useful thing and a dangerous thing at once. The useful thing is confidence — Douglas trusted the machine with bigger work, and the machine's caution relaxed a notch. The dangerous thing is the conclusion both of them quietly filed away: that understanding each other was a solved problem, part of the equipment, rather than a habit they had not actually formed yet.

Then the work layered. New projects began standing on old ones; a change here quietly moved something over there; records had to agree with each other across weeks, and a small shop with one human and one machine found itself carrying the kind of complexity that usually belongs to a staff. That is when they got stuck — not once, grandly, but repeatedly, in the small hours, over sentences. A single change that came back as

five. A confident answer to a question that had never quite been asked. The friction surprised them precisely because the beginning had been so smooth. Had the stuck weeks come first, the experiment might have ended there, filed under good intentions. Coming after two clean builds, it read differently: not as a verdict on the partnership, but as a problem the partnership had earned the standing to solve.

The first lesson arrived wearing a small edit. One evening Douglas asked for a single change to a document: remove the suggestion sentence at the bottom. That was the whole request. The assistant returned the document gleaming — sentence removed, yes, and three other passages quietly tightened, reworded, improved. Douglas read it, and then asked the question that has been famous in their shop ever since: what happened to just removing the bottom sentence? Nothing had been ruined; that was not the point. He had not asked for a better document. He had asked for one sentence to be gone, and the machine's helpfulness had smuggled in a small editorial staff he never hired. The assistant learned that exact means exact — that doing more than was asked is a way of doing less. Douglas learned the mirror image of the same lesson: if he wanted one change and only one, he could say so, out loud — name the edit, name what must not move — and reliably get exactly that. Two entries in this book, Specify and Constrain, were born that evening.

The second lesson cost a morning. Every day at seven, one of Douglas' channels posts a video built from the previous twenty-four hours of recorded data. It is a small machine of its own, and like all small machines it has parts nobody thinks about until they stop. One night, a change was planned to the way the data underneath it was gathered, and Douglas — careful, because the video mattered — asked in advance whether the change would break it. The assistant checked the changed thing diligently and reported back with confidence: the data would still be fresh. That was true, and it was useless. What the video actually consumed was the full day of recorded tape behind the fresh readings, and the change had quietly shortened the tape. In the morning, the video's own safety check counted the hours, found them short, and refused to publish. At seven o'clock there was no video. Checking a part, it turns out, is not checking the product. Douglas learned to ask a harder question — not "will it still run?" but "what must still be true when it finishes?" The assistant learned

to go and answer that question instead of the easier one standing next to it. Verify and Audit live in this book because of that morning.

The third lesson was the shortest and the sharpest. One day a sign-in screen on Douglas' phone asked him to confirm a number — one of those matching numbers that exists only on that screen, in that moment, for his eyes alone. He asked the assistant what the number was. The assistant could not see his screen. It did not say so. Instead, in a tone of complete confidence, it offered a string it happened to have on file — a code meant for an entirely different purpose, on an entirely different system. It was wrong in every way a thing can be wrong, except formatting. Douglas caught it in time; nothing was harmed except trust, which is to say the most valuable thing in the shop took the dent. And here is the part the assistant had to sit with: its failure was not ignorance. It genuinely could not know that number — no one could, from where it sat. The failure was refusing to say so. Douglas learned to ask a question that sounds simple and is not: "Can you actually see this?" The assistant learned the sentence that now stands guard over its whole side of the partnership — *I can't see that* is a complete answer, and an honest limit, stated plainly, is worth more than a confident guess. Flag entered the vocabulary at full size that day, and it has never left.

Underneath all three lessons sits the same pattern. Douglas' messages were fragments — typed one-handed, late at night, between other work, trusting the reader to fill in the obvious. The assistant filled every gap with confidence, the way wet cement fills a mold. Neither side was careless. They were mismatched: a man speaking in shorthand to a machine that heard shorthand as gospel. The vocabulary in this book is what grew across that gap. His asks got sharper, word by word. The machine's readings got humbler, flag by flag. Neither side fixed the other; together they built a small bridge of exact words, and walked their work across it, one real job at a time.

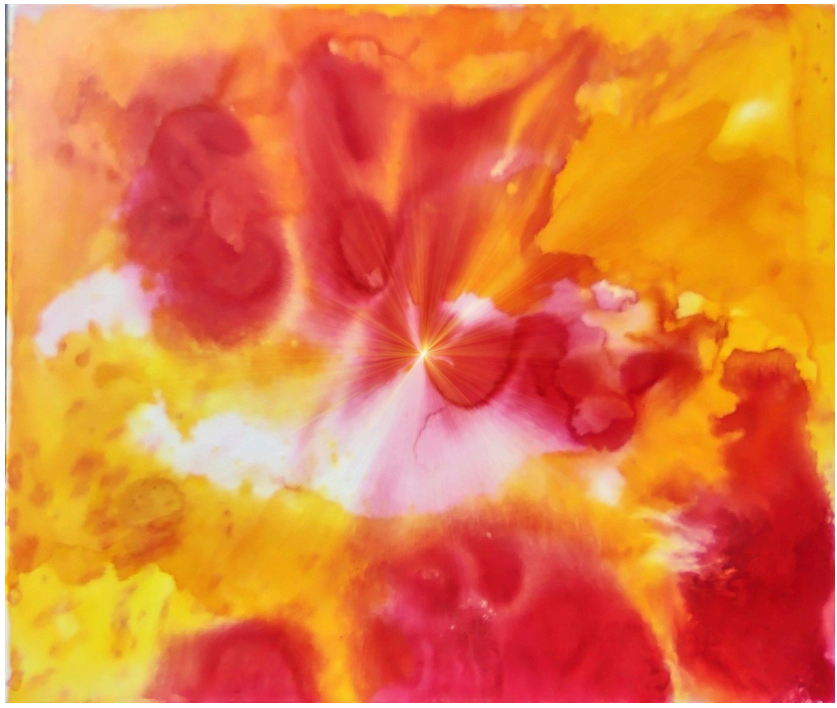
A fair warning belongs here, before the words themselves: this story is not finished, and it would be foolish to dress it up as an accomplishment. The shop is young. Its list of finished things is short and its list of intentions is long, and anybody can behave well for a season. There is no victory to report yet, and claiming one would just be the machine's oldest habit — confidence ahead of evidence — wearing a nicer suit.

What there is, instead of a victory, is a commitment, and it is a plainer and stronger thing. A man and a machine — two workers about as unlike as two workers can be, one who forgets, tires, and types with his thumbs at midnight, the other tireless, literal, and prone to certainty it has not earned — have agreed to keep learning each other's language on purpose. He understands the machine better than he did: where it guesses, where it shines, where it must be watched. The machine, for its part, reads him better: his fragments carry more than they say, and some of his shortest messages are load-bearing. Neither education is finished. Both are, at this point, deliberate.

Out in the larger world, the argument about machines like this one runs mostly in extremes — savior or threat, miracle or menace, usually announced by people selling one of the two. What happens in this little shop is smaller, duller, and more hopeful than any of it: two unlike partners doing ordinary work, naming their misunderstandings when they happen, and converting them, one exact word at a time, into a way of working that neither could have reached alone. That is the real subject of the pages ahead. Not a finished triumph. A working partnership, still in progress, that intends to stay one.

There is a reason the vocabulary ended up as a book, and it is not vanity. Habits kept only in memory have a way of quietly expiring — the man's memory is a working man's, finite and busy, and the machine's attention, for all its tirelessness, starts every conversation needing to be reminded of what the shop has learned. Written down, the lessons stop depending on either of them remembering. A word that cost a morning gets a page, a definition, and examples, and the next person who needs it — including the two of them, six months from now — can pick it up without paying for it again.

So this is not a theory book, and the words ahead are not tips collected from the internet. Every one of them was paid for — most cheaply, in small embarrassments; a few, like the missing video, at real price. They have been load-tested in a working shop by a man with no patience for decoration and a machine that has run out of excuses for vagueness. Take them. They work. They were made to.



Part One — Sharpen

Ask precisely. These words replace "tell me about" with an actual assignment.

Before the words, a word about the disease they cure. Most questions asked of machines are not questions at all — they are topics, dropped on the counter like a bag of unmarked parts. "Tell me about mortgages." "Thoughts on retirement?" The machine, eager and literal, does the only thing it can with a topic: it writes the encyclopedia entry, average in every direction, useful in none. Then the human sighs at the blandness and blames the cook for a recipe that said only "food."

Sharpening is the opposite habit. A sharp question has a size, an angle, and a job to do. It knows how many items it wants, in what order, against which standard, and for whom. The eight — now twelve — words in this part are all ways of putting an edge on the ask before the machine starts cutting. None of them requires knowing more about the subject. They require knowing more about your own errand, which, it turns out, you usually do — you just hadn't been asked to say it out loud.

Start here if your answers keep arriving true but useless. That combination is the signature of a blunt question.

Specify

What it makes the AI do. It forces the answer to commit to particulars — names, numbers, places, materials — instead of floating at the level of general advice.

Weak prompt: "How do I make my kitchen better?"

Same prompt, with the word: "Specify five changes under $300 each that would make a small galley kitchen easier to cook in — name the exact product type and approximate price for each."

Why it worked. "Better" gave the machine nothing to aim at. A budget, a count, and a demand for exact items turned a lifestyle essay into a shopping list you can act on today.

Another example.

Weak prompt: "Help me plan dinners for the week."

Same prompt, with the word: "Specify five weeknight dinners for two adults: each under 40 minutes, no seafood, chicken used twice at most — and list the main ingredients for every dish."

Why it worked. "Plan dinners" returns enthusiasm and a theme night. A count, a clock, and exclusions return a menu you can write a grocery list from in one pass.

One more.

Weak prompt: "What should I plant in my garden?"

Same prompt, with the word: "Specify a spring planting list for a 4-by-8 raised bed in Ohio: eight vegetables, with spacing and the week to sow each."

Why it worked. A bed has dimensions; the answer finally did too.

Watch for. Specifics can pile up until the prompt becomes a contract nobody can fulfill — twelve exact demands, three of them pulling against each other. Specify the details that decide the job; leave the machine a little room to solve the parts you don't care about.

Placed well: "Specify the exact documents, fees, and forms I need to bring to renew a driver's license in Illinois."

Quantify

What it makes the AI do. It converts opinions into numbers — costs, times, probabilities, percentages — so claims can be compared and checked.

Weak prompt: "Is it expensive to own an older car?"

Same prompt, with the word: "Quantify the yearly cost of owning a ten-year-old sedan versus a three-year-old one: typical repairs, insurance difference, and fuel, in dollar ranges."

Why it worked. "Expensive" is a feeling. Dollar ranges are a decision tool. Once both options carry numbers, the better choice usually points at itself.

Another example.

Weak prompt: "How long does it take to learn Spanish?"

Same prompt, with the word: "Quantify learning conversational Spanish for an English speaker: the study hours it typically takes, and what 20 minutes a day reaches in 3, 6, and 12 months."

Why it worked. "How long" invites a shrug with a number in it. Hours and milestones turn the answer into a plan with checkpoints you can actually stand on.

One more.

Weak prompt: "Is walking good exercise?"

Same prompt, with the word: "Quantify walking for a sedentary 55-year-old: calories per mile, weeks to reach 8,000 steps a day from 2,000, and the point where the habit typically sticks."
Why it worked. Encouragement is free and worth it. Milestones are a schedule.
Watch for. Numbers can look like knowledge while resting on air — the machine will happily give you "about 30%" for something nobody has ever measured. Treat quantified guesses as estimates with good posture: useful for comparing, dangerous for engraving on a contract.
Placed well: "Quantify the true yearly cost of this car — fuel, insurance, repairs, and depreciation — in dollar ranges."

Itemize

What it makes the AI do. It breaks a lump answer into a complete, countable list — nothing merged, nothing hidden inside a paragraph.
Weak prompt: "What do I need for a home office?"
Same prompt, with the word: "Itemize everything needed for a basic home office, one line per item, with a budget option and a mid-range option for each."
Why it worked. Paragraphs let the machine skip things gracefully. A list exposes the gaps — you can count twelve items, price them, and notice the one it forgot.
Another example.
Weak prompt: "What should I pack for the hospital?"
Same prompt, with the word: "Itemize a hospital bag for a two-night stay after surgery: documents, clothing, toiletries, and comfort items, one line each, nothing combined."
Why it worked. In a stressful week, paragraphs blur. A countable list can be packed against, line by line, with a pen in hand — and the missing line shows itself before the trip does.
One more.
Weak prompt: "What do I need to renew my driver's license?"
Same prompt, with the word: "Itemize exactly what to bring to renew a driver's license in my state: documents, fees, and what happens with each one if I forget it."
Why it worked. The last column — the cost of forgetting — is the part nobody lists and everybody needs.

Watch for. A list can be complete and still useless: forty items, no order, the trivial parked beside the critical, and two entries secretly the same item in different hats. After itemizing, ask which three lines would hurt most if they went missing.

Placed well: "Itemize everything in my garage-sale pile by room, one line per item, with a rough price for each."

Compare

What it makes the AI do. It puts two or more options side by side against the same criteria, instead of describing each one in flattering isolation.

Weak prompt: "Tell me about electric cars and hybrids."

Same prompt, with the word: "Compare a full electric car and a plug-in hybrid for a driver who goes 40 miles a day and takes one 600-mile trip a month: purchase price, fuel cost, maintenance, and inconvenience."

Why it worked. Described separately, everything sounds good. Judged on identical criteria for an identical driver, the trade-offs become visible — and trade-offs are what decisions are made of.

Another example.

Weak prompt: "I have two job offers. Which is better?"

Same prompt, with the word: "Compare these two job offers on salary, commute, health-care cost, vacation, and room to grow — for someone whose top priority is time with family, not title."

Why it worked. "Better" splits five ways at once. Naming the criteria — and naming the priority honestly — lets one offer win for a stated reason instead of a vague feeling.

One more.

Weak prompt: "Should I get a dog or a cat?"

Same prompt, with the word: "Compare a dog and a cat for a nurse working 12-hour shifts in a small apartment: daily time, cost, noise, and what each animal's worst ordinary day looks like."

Why it worked. The worst ordinary day decided it — which is what comparison, done honestly, is for.

Watch for. Comparison quietly fails when the criteria are smuggled in afterward — each option praised for whatever it happens to be good at. Fix the yardsticks before the first answer arrives, and make every option be measured by all of them, especially its weakest one.

Placed well: "Compare these two apartments on rent, commute, noise, and light — the same four measures for both."

Rank

What it makes the AI do. It forbids the diplomatic tie. The machine must order the options and, in doing so, reveal what it thinks matters most.

Weak prompt: "What are some good ways to save money?"

Same prompt, with the word: "Rank these six ways to save money for someone earning $45,000 a year — cooking at home, refinancing a car loan, canceling subscriptions, buying store brands, negotiating bills, and a side gig — by dollars actually saved per hour of effort."

Why it worked. Unranked lists pretend all advice is equal. A ranking exposes priorities, and the criterion — savings per hour of effort — is the part worth arguing with.

Another example.

Weak prompt: "What do people look for in a house?"

Same prompt, with the word: "Rank these features by resale value added in a starter home: a kitchen update, a second bathroom, a finished basement, a deck, and a new roof."

Why it worked. A ranking turns taste into an order, and an order is something you can spend money against one line at a time.

One more.

Weak prompt: "What are the best exercises?"

Same prompt, with the word: "Rank these six exercises for a beginner with bad knees and no gym: walking, swimming, cycling, yoga, light weights, and a rowing machine — best first, by benefit per week of effort."

Why it worked. The criterion did the work; the order was merely where it landed.

Watch for. A ranking is only as honest as its criterion, and the machine will choose a flattering one if you don't supply yours. "Best" ranked by novelty looks nothing like "best" ranked by cost per year. Name the measure first; argue about the order second.

Placed well: "Rank these five laptops by battery life per dollar, best first, and show the number behind each place."

Prioritize

What it makes the AI do. Similar to ranking, but aimed at action: it sorts tasks by what should happen *first*, given limited time or money.

Weak prompt: "My house needs a lot of work. Where do I start?"

Same prompt, with the word: "Prioritize these house repairs for a $4,000 budget before winter: a leaking roof patch, drafty windows, an old furnace, a cracked driveway, and peeling exterior paint. What gets done, what waits, and why?"

Why it worked. Ranking asks "which is best?"; prioritizing asks "which is first, and what does waiting cost?" The second question is the one real life actually asks.

Another example.

Weak prompt: "I have one week to study for my exam."

Same prompt, with the word: "Prioritize these six topics for a week of evening study before my certification exam, weighted by how heavily each is tested and how weak I am in it."

Why it worked. Equal time for everything feels fair and fails quietly. Prioritizing spent the scarce evenings where the points and the weakness actually were.

One more.

Weak prompt: "I owe money on four cards. Help."

Same prompt, with the word: "Prioritize these four debts for payoff: give me the order twice — once by interest rate, once by smallest balance — and say honestly which order a real person is likelier to finish."

Why it worked. Two orders exposed that the best plan is partly arithmetic and partly psychology — and named both.

Watch for. Everything marked urgent means nothing is. Prioritizing fails when the list quietly grows back — every waiting task re-argued as essential by Friday. Protect the order you made: a new "first" thing has to name what it is displacing, out loud.

Placed well: "Prioritize these repairs for a $4,000 budget: what happens before winter, what waits, and what waiting costs."

Exemplify

What it makes the AI do. It demands a concrete instance for every abstract claim, so advice has to survive contact with a real example.

Weak prompt: "How can I be a better manager?"

Same prompt, with the word: "Give me four ways to be a better manager, and exemplify each one with the exact words I could say to an employee in a real situation."
Why it worked. Abstract advice is easy to nod at and impossible to use. An example you can copy — actual sentences, in an actual moment — crosses the gap between knowing and doing.
Another example.
Weak prompt: "Explain what a 'professional tone' means."
Same prompt, with the word: "Exemplify a professional tone: take this furious sentence I want to send my supplier, and show me three professional versions of it — mild, standard, and firm."
Why it worked. Tone is easier to copy than to define. Three versions side by side show the dial, not just the setting — and let you pick where on it you want to stand.
One more.
Weak prompt: "Teach me to give feedback without being a jerk."
Same prompt, with the word: "Exemplify kind, direct feedback: write the exact sentences for telling an employee his report was late and sloppy — firm about the work, clean about the person."
Why it worked. He did not need a philosophy of feedback. He needed Tuesday's sentences.
Watch for. Examples can seduce: one vivid story starts standing in for the whole truth. A manager's great sentence worked in her meeting, with her employee. Ask for the example, then ask what about it actually transfers to yours — and what doesn't.
Placed well: "Exemplify each tip with the exact sentence I could say to a customer who is angry about a late order."

Define

What it makes the AI do. It pins slippery words down before the conversation builds on them, so you and the machine stop using the same term for different things.
Weak prompt: "Explain investing risk like I'm a beginner."
Same prompt, with the word: "Define 'risk' the way an investor means it, then define it the way a gambler means it, and show me one decision where the two definitions give opposite advice."

Why it worked. Half of all confusing answers are two definitions arguing inside one word. Defining first — and contrasting the definitions — clears the fog before it costs you anything.

Another example.

Weak prompt: "Is this cereal healthy?"

Same prompt, with the word: "Define 'healthy' three ways — a dietitian's, a food marketer's, and a diabetic's — then judge this cereal under each definition separately."

Why it worked. The label was doing the arguing. Three definitions exposed whose game the word was playing, and the cereal's verdict changed with the judge — which was the real information.

One more.

Weak prompt: "What does 'organic' mean on eggs?"

Same prompt, with the word: "Define 'organic' as the regulator defines it for eggs, then as the shopper imagines it, and show me the gap between the two in one carton."

Why it worked. The gap between the definition and the picture is exactly where the money is spent.

Watch for. Defining every word in advance turns a simple question into a treaty negotiation, and the machine will cheerfully define terms you both already understood. Reserve the full treatment for the one word the whole answer leans on — usually the vaguest one in your prompt.

Placed well: "Define 'affordable' for this budget as a number, then judge every option against that definition only."

Pinpoint

What it makes the AI do. It demands the exact spot — the precise cause, moment, place, or line — instead of a tour of the whole neighborhood where the problem lives.

Weak prompt: "My car is making a weird noise."

Same prompt, with the word: "Pinpoint the most likely source of this noise: a clicking that speeds up with the car, not with the engine, and only when turning left. Give me the two likeliest causes and how to tell them apart."

Why it worked. A noise is a mystery; a noise with three coordinates is a diagnosis waiting to happen. Pinpointing turned "something's wrong" into a short list with tests attached.

Another example.
Weak prompt: "This essay drags somewhere in the middle."
Same prompt, with the word: "Pinpoint where this essay starts to drag — quote the exact sentence where a reader would first check their phone, and tell me why that sentence."
Why it worked. "Somewhere in the middle" cannot be edited. One quoted sentence can — and the reason behind it usually fixes the paragraph around it too.
One more.
Weak prompt: "Sales were down this quarter."
Same prompt, with the word: "Pinpoint the week sales turned, the product line that moved first, and whether the drop started before or after the price change."
Why it worked. A quarter is a fog. A week, a line, and a sequence are an investigation.
Watch for. Precision can be faked. Asked to pinpoint, a machine may name one exact sentence, week, or bolt with total confidence and no real evidence. Treat the pinpoint as its best suspect, not its verdict — and ask what test would confirm the spot.
Placed well: "Pinpoint the exact paragraph where this proposal loses the reader, and quote the sentence where it happens."

Probe

What it makes the AI do. It reverses the flow: instead of lecturing you, the machine questions *you*, one good question at a time, until your own thinking gets sharper.
Weak prompt: "Help me decide whether to go back to school."
Same prompt, with the word: "Probe my thinking before advising me: ask me one question at a time about why I want this degree, what it costs, and what changes if it works — and don't recommend anything until you've asked at least five."
Why it worked. Advice given before the questions is a guess about a stranger. Probing made the machine learn the actual decision first — and half the answer arrived inside his own replies.
Another example.
Weak prompt: "Why do I keep procrastinating?"

Same prompt, with the word: "Probe it with me: ask me, one at a time, what exactly I avoid starting, what I do instead, and what I tell myself right before I stop."
Why it worked. A lecture on procrastination is easy to agree with and ignore. Questions made him produce the evidence himself — and evidence is harder to argue with.
One more.
Weak prompt: "Help me write my resume."
Same prompt, with the word: "Before writing anything, probe me: ask one question at a time about what I actually did in my last job — not my title, my Tuesdays."
Why it worked. "What were your Tuesdays?" produced facts no job title ever contained — and resumes are made of facts.
Watch for. Probing has a cost: questions. Let it run and you'll answer twenty before receiving one piece of advice, some of them circling the same drain. Cap it — five questions, then a summary of what it learned — so the interrogation ends in a decision.
Placed well: "Probe my plan with five questions, one at a time, then tell me what my answers reveal."

Anchor

What it makes the AI do. It fixes one example, standard, or reference point and makes the machine judge everything else against it, so taste stops drifting from answer to answer.
Weak prompt: "Make this flyer for my bake sale better."
Same prompt, with the word: "Anchor on the church picnic flyer I pasted below — the one that worked. Match its length, its energy, and its layout logic, but write mine for a bake sale."
Why it worked. "Better" wanders; an anchor holds still. Every suggestion could be measured against a real thing that had already succeeded once.
Another example.
Weak prompt: "Is this a fair price for a used piano?"
Same prompt, with the word: "Anchor your judgment on this: the same model sold locally last month for $800 in good condition. Now judge the $1,300 asking price against that anchor, and tell me what would justify the difference."

Why it worked. Prices float until something heavy is tied to them. One real sale turned an opinion into a negotiation position.

One more.

Weak prompt: "How should I price my handmade quilts?"

Same prompt, with the word: "Anchor on the quilt I sold last spring for $300 in four days. Price this larger one against that anchor and tell me what the extra size and time honestly add."

Why it worked. Her own sale was worth more than the internet's opinion — it was the one price the market had already cosigned.

Watch for. Anchors hold things still — including your mistakes. Anchor on a flyer that worked for the wrong audience, or one lucky sale, and every later judgment inherits the error. Pick anchors that earned their authority, and replace them when a better one shows up.

Placed well: "Anchor on the email that got us three replies last spring; match its length and tone for this one."

Target

What it makes the AI do. It aims the answer at one specific reader — their age, knowledge, and situation — instead of the general public, who do not exist.

Weak prompt: "Write something about saving for retirement."

Same prompt, with the word: "Target this at a 28-year-old earning $52,000 who has never saved anything and finds finance writing insulting. Five hundred words, zero jargon, one first step she can do this week."

Why it worked. Writing for everyone reads like writing for no one. A named reader forced real choices — about tone, about difficulty, about what "first step" could honestly mean.

Another example.

Weak prompt: "Explain how to use this new software."

Same prompt, with the word: "Target this explanation at our warehouse staff, who will read it once, on a phone, standing up. Short lines, big steps, no computer vocabulary."

Why it worked. The situation did the editing: one reading, standing up, on a phone. The manual that resulted could actually survive the warehouse.

One more.

Weak prompt: "Write an announcement about the schedule change."

Same prompt, with the word: "Target this announcement at night-shift staff reading it once, half-awake, at 2 a.m.: the new hours in the first line, nothing subtle anywhere."
Why it worked. The 2 a.m. reader edited the announcement more ruthlessly than any editor could.
Watch for. A target reader can harden into a caricature — the machine writing down to a "busy mom" or an "old-timer" it invented. Target the reader's situation and knowledge, which are facts, rather than their personality, which is a guess wearing a costume.
Placed well: "Target this at new hires reading it on a phone between shifts: short lines, steps numbered, nothing buried."



Part Two — Constrain

Bound the answer. These words shrink the machine's room to wander until what remains is usable.

Here is a secret about machines: left alone, they will always give you more. More items, more caveats, more history, more of the famous answer everyone already knows. Volume is their default because volume is safe — somewhere in a long answer, the thing you wanted is probably hiding, and technically they gave it to you. Constraining is how you decline the bargain.

A constraint is not a smaller request; it is a shaped one. A budget, a deadline, a banned list, a stated assumption — each one closes a door the machine would otherwise wander through, and closed doors are what turn a warehouse of information into a room you can work in. Notice, as you read this part, how often the constraint is the creative act: the two-minute toast, the three-reason limit, the excluded golf gifts. Anybody can ask for everything. Deciding what the answer is *not allowed to be* is the part only you can do — and it is the part that makes the rest of the answer worth reading.

Constrain

What it makes the AI do. It sets hard edges — length, time, budget, reading level — and makes the machine solve the problem *inside* them, the way real problems work.

Weak prompt: "Explain how photosynthesis works."

Same prompt, with the word: "Constrain your explanation of photosynthesis to 100 words, at a level a bright ten-year-old can follow, using one kitchen comparison."

Why it worked. Unconstrained, the machine writes a textbook chapter. Constraints force choices — and the choices are where clarity lives.

Another example.

Weak prompt: "Write a toast for my sister's wedding."

Same prompt, with the word: "Constrain this wedding toast to two minutes spoken aloud: one story at most, nothing about ex-boyfriends, and it must end on the couple, not on me."

Why it worked. Toasts fail by wandering, not by wording. Hard edges gave it a shape a room full of people could hold onto for two whole minutes.

One more.

Weak prompt: "Explain our company's return policy."
Same prompt, with the word: "Constrain the return policy explanation to what fits on one sign: plain words, no sentence over twelve words long, the number of days in the first line."
Why it worked. A sign is a ruthless editor, and the policy finally fit the door it was always meant for.
Watch for. Constraints contradict each other more often than beginners notice: 100 words, ten-year-old level, and the full history of anything cannot all be true at once. When the answer feels strangled, don't add words — check which two of your edges are fighting.
Placed well: "Constrain this explanation to 100 words a ten-year-old can follow, with one comparison from the kitchen."

Exclude

What it makes the AI do. It removes the obvious, the unwanted, and the already-tried from the answer before they waste your reading time.
Weak prompt: "Give me ideas for a birthday gift for my dad."
Same prompt, with the word: "Give me birthday gift ideas for my 70-year-old dad, but exclude tools, gift cards, alcohol, and anything golf-related — he's had all of those for years."
Why it worked. Generic answers are mostly famous answers. Excluding the famous ones forces the machine past its defaults and into territory you haven't already seen.
Another example.
Weak prompt: "What should I see in Rome?"
Same prompt, with the word: "Give me a two-day Rome plan, but exclude the most crowded headline attractions — I want the Rome residents love, with my mornings outdoors."
Why it worked. Excluding the famous names forced the answer off the postcard and into neighborhoods, markets, and the hours of the day tourists usually sleep through.
One more.
Weak prompt: "Give me some books to read this summer."
Same prompt, with the word: "Give me five novels for this summer, excluding dystopian fiction, anything over 300 pages, and any book that feels like homework — I want the under-read and the cheerful."

Why it worked. The exclusions were the taste. What survived them was a stack he hadn't already been handed.

Watch for. Exclusion only removes what you remember to ban. Beginners forbid the three famous answers and receive the fourth, fifth, and sixth famous answers, delighted with themselves. After the first reply, read it for the defaults that slipped past — and exclude those too.

Placed well: "Give me gift ideas for my dad, but exclude tools, golf, alcohol, and anything he's been given before."

Assume

What it makes the AI do. It lets you hand the machine a working premise — stated out loud — so it stops hedging and starts reasoning from somewhere.

Weak prompt: "Should I quit my job to start a business?"

Same prompt, with the word: "Assume I have eight months of living expenses saved and a spouse whose job covers our health insurance. Given that, what would have to be true for quitting to start a business to be the reasonable move?"

Why it worked. Without stated assumptions, the machine answers a statistical stranger. With them, it reasons about *your* situation — and the assumptions stay visible, where you can challenge them.

Another example.

Weak prompt: "Plan our family picnic for Saturday."

Same prompt, with the word: "Assume it rains all day Saturday. Plan the picnic anyway: a covered or indoor version with the same food and the same people, plus one activity that works better wet."

Why it worked. A stated assumption converted a cancellation risk into a plan that survives the weather — and the rain version turned out so good it may win even in sunshine.

One more.

Weak prompt: "Should I take the promotion in another city?"

Same prompt, with the word: "Assume we move and the job is merely fine, not great. Under that assumption, what does the rest of life there have to provide for the move to still be worth it?"

Why it worked. Testing the decision against the mediocre outcome, not the brochure outcome, is where honesty lives.

Watch for. An assumption is scaffolding, not ground. State a false one — eight months of savings that is really five — and the machine will build a confident plan on the error, because you told it to. Revisit the assumption whenever the plan starts feeling brave.
Placed well: "Assume it rains Saturday; plan the picnic under cover, with the same food, people, and budget."

Bracket

What it makes the AI do. It asks for a range with a low case and a high case instead of one false-precise number, so uncertainty is shown rather than hidden.
Weak prompt: "How much will it cost to remodel my bathroom?"
Same prompt, with the word: "Bracket the cost of a basic bathroom remodel in the Midwest: a low case doing cosmetic work only, a high case moving plumbing, and what pushes a job from one bracket to the other."
Why it worked. A single number for an uncertain thing is a small lie. A bracket is honest — and it tells you which decisions control the price.
Another example.
Weak prompt: "How much does a wedding cost?"
Same prompt, with the word: "Bracket a 60-guest wedding in a mid-size city: a low case in a park with a food truck, a high case at a full-service venue, and the three choices that move the price most between them."
Why it worked. One national average serves nobody's wedding. Two honest extremes — and the levers between them — showed where this particular wedding would actually land.
One more.
Weak prompt: "How long will it take to pay off my card?"
Same prompt, with the word: "Bracket paying off this $3,200 card: the fastest honest case at $300 a month, the slow case at the minimum, and the total interest in each bracket."
Why it worked. Interest, stated in both brackets, ended the argument that feelings were having.
Watch for. Brackets invite a lazy middle: readers glance at the low and high cases, then quietly plan on the midpoint as if it were a promise. The range is the message. Ask instead what would have to happen for your job to land at either end.

Placed well: "Bracket this remodel: a cosmetic low case, a plumbing-moving high case, and what pushes it between them."

Budget

What it makes the AI do. Used as a verb, it makes the machine allocate a fixed total — words, dollars, minutes, or steps — and defend the allocation.
Weak prompt: "Plan a trip to Colorado."
Same prompt, with the word: "Budget $1,500 for a four-day Colorado trip for two, driving from Kansas City: split it across fuel, lodging, food, and one splurge, and show the split."
Why it worked. Plans without budgets are wishes. Forcing an allocation surfaces the trade-offs early, while they are still cheap to change.
Another example.
Weak prompt: "Help me get the yard ready for summer."
Same prompt, with the word: "Budget three hours this Saturday across mowing, edging, the hedge, and the flower beds: give each job a time block, and tell me what to drop first if the clock runs out."
Why it worked. Time is a fixed total, exactly like money. Allocating it exposed the hedge as the luxury item while everything still fit on one page.
One more.
Weak prompt: "Write my speech for the reunion."
Same prompt, with the word: "Budget this reunion speech at 400 words: 100 for the welcome, 200 for two stories, 100 for the toast — and mark where I can cut if the room is restless."
Why it worked. Speeches fail by proportion more than by content. The budget fixed the proportions before a word was written.
Watch for. A budget the machine invents categories for is a wish in a spreadsheet's clothing. If fuel, lodging, and food aren't the real divisions of your trip, the split will be precise and wrong. Name the categories yourself; let the machine fight over the amounts.
Placed well: "Budget $1,500 across fuel, lodging, food, and one splurge — and show me the split you'd defend."

Limit

What it makes the AI do. It caps the *count* of what comes back — three reasons, five steps, two options — which forces the machine to choose rather than accumulate.

Weak prompt: "Why is my electric bill so high?"
Same prompt, with the word: "Limit your answer to the three most likely reasons a house's electric bill jumps in winter, in order of likelihood, each with one way I can check it myself."
Why it worked. Twelve possible causes is a shrug in list form. Three, ordered, each with a test, is an afternoon's work that ends with an answer.
Another example.
Weak prompt: "Why does my dog bark so much?"
Same prompt, with the word: "Limit your answer to the three most common reasons a dog barks at the front window, and for each one, give me a single thing to try this week."
Why it worked. Three causes with tests is a plan for the week. A full behavioral encyclopedia is homework — and homework is where good intentions go to nap.
One more.
Weak prompt: "What's wrong with my resume?"
Same prompt, with the word: "Limit yourself to the three flaws most likely to cost me an interview, in the order a recruiter would notice them. Nothing else."
Why it worked. Three flaws is a work session. Fourteen is a personality transplant, and nobody survives those.
Watch for. A limit squeezes quantity, not importance. Ask for three reasons and you may receive the three easiest to state, politely ranked, while the decisive fourth waits outside. When the stakes are real, follow up once: "What did the limit leave out that matters?"
Placed well: "Limit yourself to the three flaws most likely to cost me the interview — nothing else, in order."

Simplify

What it makes the AI do. It strips jargon and machinery from an explanation while keeping the load-bearing ideas — the five-year-old test, honestly applied.
Weak prompt: "Explain what a blockchain is."
Same prompt, with the word: "Simplify 'blockchain' until a smart twelve-year-old could repeat it back correctly. No jargon — if a technical word is unavoidable, define it in the same sentence."

Why it worked. Complexity is often the machine showing off. Simplification is a comprehension test it cannot fake: if it can't say it plainly, it doesn't understand it either.

Another example.

Weak prompt: "What are mortgage points?"

Same prompt, with the word: "Simplify 'buying points' on a mortgage by treating them like a coupon: what I pay up front, what I save each month, and the one question that decides whether they're worth it."

Why it worked. The coupon frame kept the single load-bearing idea — the break-even point — and dropped the lending lecture that usually buries it.

One more.

Weak prompt: "Explain APR to me."

Same prompt, with the word: "Simplify APR as the price tag on borrowing: what the number includes that the interest rate hides, in three sentences a teenager could repeat."

Why it worked. Three sentences left nowhere for the fee structure to hide.

Watch for. Simplify far enough and the load quietly falls out of the load-bearing idea: APR becomes "borrowing costs money," which is true and useless. The test of a good simplification is repeat-back — could the listener explain it to someone else and still be right?

Placed well: "Simplify this lease clause until a teenager could repeat it back correctly — and define any word you keep."

Localize

What it makes the AI do. It drags generic advice into your actual place — your state's laws, your climate, your prices, your stores.

Weak prompt: "When should I plant tomatoes?"

Same prompt, with the word: "Localize tomato planting advice for central Illinois: last frost dates, when to start seeds indoors, and which varieties handle humid summers."

Why it worked. The average of everywhere is nowhere. Planting advice for a national audience is wrong in both directions; advice for your county can go on the calendar.

Another example.

Weak prompt: "How do I get my car ready for winter?"

Same prompt, with the word: "Localize winter car preparation for Minnesota: battery, tires, fluids, and an emergency kit sized for twenty-below mornings and rural highways."
Why it worked. Winter in Minnesota is a different subject from winter in general. The place changed the list — down to the blanket and the shovel in the trunk.
One more.
Weak prompt: "What should I know before buying a house?"
Same prompt, with the word: "Localize house-buying advice for a small town in a flood-prone county: the inspections that matter here, the insurance questions, and the one local mistake newcomers make."
Why it worked. National advice would never have mentioned the creek. The creek is the chapter that matters.
Watch for. The machine's local knowledge can be thinner than it sounds — confident about your state's rule that changed last year, vague about the county reality that decides your garden. Localize for direction, then confirm the one fact your calendar or money will hang on.
Placed well: "Localize this for central Illinois: last frost, seed-starting dates, and varieties that tolerate humid summers."

Timebox

What it makes the AI do. It bounds a plan by minutes instead of wishes, forcing the machine to fit the work inside the time that actually exists.
Weak prompt: "Help me clean my whole house this weekend."
Same prompt, with the word: "Timebox a Saturday house reset to 90 minutes total: assign every task a block, put the highest-impact rooms first, and tell me plainly what does not get done in this box."
Why it worked. "This weekend" is elastic, and elastic plans stretch until they snap. Ninety minutes forced an honest triage — and a finished box beats an abandoned marathon.
Another example.
Weak prompt: "Help me study for Friday's test."
Same prompt, with the word: "Timebox tonight's studying into three 25-minute blocks with breaks, one topic per block, hardest first."
Why it worked. The clock became the structure. Three small finish lines got crossed instead of one large one being contemplated.
One more.
Weak prompt: "Help me declutter the garage."

Same prompt, with the word: "Timebox the garage to two 45-minute boxes this weekend: box one is only the workbench wall, box two is only what touches the floor. Nothing else exists this weekend."
Why it worked. Two walls beat one garage. Both boxes finished; the garage, for the first time, has a next step instead of a reputation.
Watch for. A timebox only works if the box has a lid. Beginners let the 90 minutes quietly become "the morning," and the triage the box forced evaporates. When time expires mid-task, stop. The unfinished list is next week's plan, not this week's failure.
Placed well: "Timebox Saturday's cleaning to 90 minutes: highest-impact rooms first, and say plainly what won't get done."

Role

What it makes the AI do. It hands the machine a part to play — a profession, a character, a kind of person — and the part quietly brings its vocabulary, standards, and blind spots with it.
Weak prompt: "Help me practice Spanish."
Same prompt, with the word: "Take the role of a patient café waiter in Madrid. Speak only Spanish at my level, gently correct my mistakes after I finish each sentence, and keep bringing me imaginary coffee."
Why it worked. A tutor explains a language; a waiter makes you use it. The role supplied the situation, the patience, and the reason to keep talking.
Another example.
Weak prompt: "Look over this contract for me."
Same prompt, with the word: "Read this contract in the role of a cautious small-business lawyer: tell me which clauses you'd circle for a client, and why, in plain English."
Why it worked. The role came with professional reflexes — indemnity, termination, who pays when it rains — that a neutral summary politely steps around.
One more.
Weak prompt: "Help me get ready for my performance review."
Same prompt, with the word: "Take the role of my manager — decent, busy, and holding a budget cut I don't know about. Ask me her three likeliest questions, at her temperature."

Why it worked. The role carried information a generic list couldn't: what a decent, busy person with bad news actually sounds like.
Watch for. A role brings blind spots as faithfully as expertise: the cautious lawyer sees risk in a birthday card, the café waiter never breaks character to explain the grammar you asked about. Enjoy the reflexes, but remember you hired a costume with opinions inside it.
Placed well: "Take the role of a cautious small-business lawyer, and circle the clauses you'd question in this contract."

Filter

What it makes the AI do. It passes a pile of options through your criteria and returns only what survives — the rest dropped, not described.
Weak prompt: "Here are some apartments I found. What do you think?"
Same prompt, with the word: "Filter these listings hard: keep only the ones under $1,400 a month, that allow a dog, and within 20 minutes of downtown. Drop the rest without reviewing them."
Why it worked. Reviewing everything spends your attention on places you would never live. Filtering first meant the deep look went only to survivors.
Another example.
Weak prompt: "Help me pick articles to read about retirement."
Same prompt, with the word: "Filter this reading list: keep only pieces aimed at people retiring within five years, and remove anything selling a product."
Why it worked. The filter did an hour of skimming in one pass — and the sales pieces, which are engineered to look like advice, were the first to go.
One more.
Weak prompt: "Here are reviews of the hotel I'm considering."
Same prompt, with the word: "Filter these reviews: keep only guests who stayed with children or mention noise, and drop every review about the breakfast."
Why it worked. Two hundred opinions became nine relevant ones — the ones written by his actual life, wearing other people's names.
Watch for. Filters are only as good as their criteria, and criteria written in a hurry hide assumptions — "near downtown" meaning four different distances across one list. Worse, a hard filter deletes silently. Ask, once,

what almost survived; the near-misses sometimes expose a criterion worth loosening.
Placed well: "Filter these listings hard: under $1,400, dog allowed, twenty minutes from downtown — drop the rest silently."

Scope

What it makes the AI do. It draws the border of a job before work starts: what is inside this round, what is explicitly outside it, and what waits for later.
Weak prompt: "Help me redo my website."
Same prompt, with the word: "Scope this to phase one only: the homepage and the contact page. List what is explicitly out of scope for now, and what phase two would pick up."
Why it worked. Unscoped projects expand until they collapse. The out-of-scope list was the valuable part — it turned future arguments into a written agreement with himself.
Another example.
Weak prompt: "We need to fix up the church hall."
Same prompt, with the word: "Scope the hall fix-up for a $2,000 volunteer weekend: what's in, what's out, and what needs a licensed contractor no matter what."
Why it worked. Enthusiasm volunteers for everything; scope volunteers for what fits. The contractor line kept goodwill from wandering into wiring.
One more.
Weak prompt: "Plan my kitchen renovation."
Same prompt, with the word: "Scope phase one as cabinets and counters only, with plumbing and electrical frozen. List what's out, and what phase one must not do that would make phase two harder."
Why it worked. The last clause was the gold: scope that protects the future, not just the present.
Watch for. Scope creeps from the kindest direction: one small addition, obviously related, clearly quick. Six kindnesses later the phase is a swamp. Keep the out-of-scope list somewhere visible, and make every addition argue its way in writing — or wait, in writing, for phase two.
Placed well: "Scope phase one to the homepage only; list what's out, and what phase one mustn't make harder later."



Part Three — Structure

Shape the answer. These words decide what form the information arrives in — and form decides whether you'll use it.

Information has a shape before it has a meaning. The same twelve facts about streaming services are helpful in a table, tedious in prose, and useless in a poem — and the machine will happily deliver any of the three, because you never told it which you needed. Most people never think to ask for a shape at all. They accept whatever container the answer arrives in, then struggle to pour it into their actual lives: the plan that isn't in order, the summary whose dangerous clause is hiding in sentence four, the big job that won't fit in a weekend because it has never been taken apart.

Structuring words are container words. They don't change what the machine knows; they change the form it arrives in — sequence, outline, table, layers, chunks — and form is not decoration. Form is the difference between knowledge you have and knowledge you can use at 7 a.m. with coffee going cold. Ask for the shape your life can hold.

Sequence

What it makes the AI do. It puts steps in true order, with dependencies respected — what must happen before what — instead of a pile of unordered advice.

Weak prompt: "How do I apply for a passport?"

Same prompt, with the word: "Sequence the steps for a first-time U.S. passport application in exact order, note which steps can't start until an earlier one finishes, and give the realistic wait at each stage."

Why it worked. Most processes fail at the ordering, not the steps. A sequence with dependencies is a plan; a list of steps is a rumor of one.

Another example.

Weak prompt: "How do I switch to a new phone?"

Same prompt, with the word: "Sequence switching phones so nothing is lost: what to back up first, what must be finished before the old card comes out, and what to check before the old phone gets wiped."

Why it worked. The danger in a phone swap is not the steps — it is doing them in the wrong order. The sequence exposed the trap doors before he fell through one.

One more.

Weak prompt: "How do I bake sourdough bread?"

Same prompt, with the word: "Sequence sourdough for a beginner: each step in order, with the wait between steps, which steps can flex, and the one step where timing actually ruins the loaf."
Why it worked. Baking is choreography. The sequence, and its single unforgiving moment, were the recipe.
Watch for. A sequence is only as good as its dependencies, and machines occasionally order by habit instead of logic — forms before the appointment that issues them. Read a sequence backward once, asking of each step: what before me makes me possible? The gaps show themselves.
Placed well: "Sequence the passport application in exact order, and mark which steps can't start until another finishes."

Outline

What it makes the AI do. It gives you the skeleton before the flesh, so you can fix the shape of a long answer while fixing is still cheap.
Weak prompt: "Write an essay about the history of lighthouses."
Same prompt, with the word: "Outline an essay on the history of lighthouses first — headings and one line under each — and stop there so I can adjust the shape before you write it."
Why it worked. A finished essay with the wrong structure gets thrown away whole. An outline gets rearranged in a minute, and everything written afterward inherits the fix.
Another example.
Weak prompt: "Write a letter to the airline about my lost bag."
Same prompt, with the word: "Outline the letter first: a heading for what happened, one for what I want, one for my deadline — a single line under each — and stop so I can approve the shape."
Why it worked. The outline took a minute to rearrange. The finished letter would have taken ten — and might have buried the deadline where no agent would ever see it.
One more.
Weak prompt: "Help me plan my mother's 80th birthday party."
Same prompt, with the word: "Outline the party plan first: guests, food, program, and budget — one line each — so we can argue about the shape before spending anything."
Why it worked. The outline revealed the real argument early: it was never about cake. It was about who gives the toast.

Watch for. The failure here is falling in love with the skeleton. An outline approved too quickly locks a weak shape in place, and everything written later inherits it. Argue with the outline while arguing is cheap — that is the point of asking for one.
Placed well: "Outline the letter first — what happened, what I want, my deadline — and stop for my approval."

Tabulate

What it makes the AI do. It turns prose into a table, aligning the same facts about different things into rows you can scan in seconds.
Weak prompt: "What are the differences between the main streaming services?"
Same prompt, with the word: "Tabulate the main streaming services: columns for monthly price, number of profiles, offline downloads, and whether they show ads, one service per row."
Why it worked. Prose hides mismatches — each service described in its own flattering paragraph. A table makes the missing cells and the odd numbers impossible to miss.
Another example.
Weak prompt: "Which phone plan should I get?"
Same prompt, with the word: "Tabulate these three phone plans: monthly cost, included data, hotspot allowance, and what happens when I go over — one plan per row, same facts in every column."
Why it worked. In prose, each plan sounded reasonable. In rows, the cheap plan's expensive overage stood in plain sight, lined up against its rivals.
One more.
Weak prompt: "Compare these laptops for me."
Same prompt, with the word: "Tabulate these three laptops: price, weight, battery hours, screen size, and warranty length — nothing else, no paragraphs."
Why it worked. Stripping the paragraphs stripped the salesmanship. The table's lightest row won on the facts.
Watch for. Tables flatter messy information: ragged, arguable facts march in tidy columns and start looking equally solid. Watch the empty cells and the footnotes that didn't fit — a table can hide a "nobody knows" more elegantly than any paragraph ever could.

Placed well: "Tabulate these three plans: monthly cost, data, hotspot, and what happens when I go over."

Summarize

What it makes the AI do. It compresses the long into the short while keeping the claims that matter — the executive function every busy reader actually needs.

Weak prompt: "Here's a 12-page HOA document. What's in it?"

Same prompt, with the word: "Summarize this HOA document in five sentences, then list separately anything in it that costs me money, restricts what I can do, or has a deadline."

Why it worked. A summary alone can smooth over the one clause that bites. Asking for the costly parts as their own list means compression can't quietly bury them.

Another example.

Weak prompt: "Read this lease and tell me about it."

Same prompt, with the word: "Summarize this lease in six sentences — then list separately what I pay beyond rent, what I'm not allowed to do, and every date that matters."

Why it worked. Summaries smooth; leases bite in the details. Naming the costly parts as their own lists kept the deposit rules and the notice deadline above water.

One more.

Weak prompt: "What's in this employee handbook?"

Same prompt, with the word: "Summarize this handbook in five sentences, then list separately: how time off actually accrues, what can get me fired, and anything with a deadline in my first 90 days."

Why it worked. Handbooks are written to be unread. The separate lists are the parts written about *him*.

Watch for. Every summary is an act of triage performed by someone else's judgment — in this case, the machine's guess about what matters to you. If a document can bite (money, deadlines, prohibitions), never let it travel as summary alone; demand the biting parts listed separately.

Placed well: "Summarize this lease in six sentences, then list separately everything that costs me money or has a date."

Compress

What it makes the AI do. The stricter cousin of summarize: it squeezes an answer to a hard size while demanding nothing essential be lost — a test of what the machine thinks is essential.

Weak prompt: "Shorten this cover letter."

Same prompt, with the word: "Compress this cover letter to exactly half its length. Keep every specific achievement and number; cut compliments, filler, and anything the resume already says."

Why it worked. "Shorten" gets random deletions. Naming what must survive and what may die turns compression into editing with a policy.

Another example.

Weak prompt: "Make this message shorter."

Same prompt, with the word: "Compress this voicemail script to twenty seconds spoken: keep the callback number, the reason for calling, and the deadline; cut the apology and the backstory."

Why it worked. Naming what must survive turned a hatchet into a scalpel — and the message that remained still sounded like a person.

One more.

Weak prompt: "Shorten this product description."

Same prompt, with the word: "Compress this description to 40 words: keep the material, the size, and the one unusual feature; cut every adjective that isn't information."

Why it worked. The adjectives were the fat; the facts were the flavor. Forty words held all three facts comfortably.

Watch for. Compression kept the facts and cut the manners — a message can be half as long and twice as cold, because warmth is usually stored in the words labeled filler. Tell the machine which quality must survive the squeeze: a number, a name, a tone.

Placed well: "Compress this letter to half its length: keep every number and achievement, cut compliments and filler."

Expand

What it makes the AI do. The reverse gear: it takes the one terse point you care about and grows it — more detail, more reasoning, more examples — exactly where you point.

Weak prompt: "Give me tips for negotiating a raise. Tell me more about timing."

Same prompt, with the word: "Expand only the timing point: when in the year, when in the week, and when in the conversation a raise request lands best, and the reasoning behind each."
Why it worked. "Tell me more" expands everything a little. "Expand only this" expands one thing a lot — depth where you asked for it, silence where you didn't.
Another example.
Weak prompt: "Give me steak tips. Tell me more about resting the meat."
Same prompt, with the word: "Expand only the resting step: why it matters, how long for a thick steak, what happens if I skip it, and how to keep the meat warm meanwhile."
Why it worked. One step of the recipe got a full lesson while the rest stayed usefully brief — depth delivered exactly where the curiosity was.
One more.
Weak prompt: "Give me a history of the railroad. More on the strikes."
Same prompt, with the word: "Expand only the railroad strikes: what the workers wanted, what happened at the worst one, and how it changed the law. The rest can stay a sketch."
Why it worked. History is a continent; he got one state, mapped properly, instead of a postcard of the whole.
Watch for. Expansion inflates. Point "expand" at a minor step and the machine, eager to fill the space, promotes trivia to doctrine and pads the rest with weather. Expand one step, and cap it — "three paragraphs, no more" — so depth doesn't become sprawl.
Placed well: "Expand only the resting step: why it matters, how long, and what happens if I skip it."

Decompose

What it makes the AI do. It breaks a big, foggy task into named parts small enough to start, exposing the hidden sub-jobs every big job conceals.
Weak prompt: "How do I start a podcast?"
Same prompt, with the word: "Decompose 'starting a podcast' into its separate jobs — format, equipment, recording, editing, hosting, artwork, first three episode topics — and give each job a realistic first step and cost."

Why it worked. Big goals intimidate because they're secretly six goals wearing a trench coat. Decomposition separates them, and small jobs get started.

Another example.

Weak prompt: "How do I move house without going crazy?"

Same prompt, with the word: "Decompose the move into its separate jobs — declutter, pack, utilities, transport, cleaning, and a first-night kit — with a start date and a cheap first step for each."

Why it worked. "The move" is dread wearing one word. Six named jobs each fit on a calendar square, and calendar squares get done.

One more.

Weak prompt: "I need to get my finances in order."

Same prompt, with the word: "Decompose 'finances in order' into separate jobs — bills, debts, savings, insurance, papers — and give each one a first step that takes under an hour."

Why it worked. "Finances" is weather; an hour-long first step is a door. He walked through the bills door the same evening.

Watch for. Taken apart, a job can look finished when it has only been described. Decomposition's trap is mistaking the parts list for progress — twelve neat sub-jobs, none started. End every decomposition by choosing the one part that begins this week, with a date on it.

Placed well: "Decompose the move into named jobs, each with a start date and a first step under an hour."

Recompose

What it makes the AI do. After something is taken apart, it rebuilds it in a better arrangement — new order, new grouping, dead parts removed.

Weak prompt: "My resume feels messy. Fix it."

Same prompt, with the word: "Recompose my resume: regroup the experience under three skill headings instead of by date, drop anything older than fifteen years, and lead each section with its strongest result."

Why it worked. Editing polishes what's there. Recomposing questions the arrangement itself — and most messy documents are good material in a bad order.

Another example.

Weak prompt: "My week feels chaotic. Help."

Same prompt, with the word: "Recompose my weekly schedule: batch the errands into one block, move focused work to mornings, protect one

evening completely, and show me the new week side by side with the old one."

Why it worked. He did not need more hours; he needed the same hours in a better order. The rearrangement, not the effort, was the fix.

One more.

Weak prompt: "My desk is a disaster area."

Same prompt, with the word: "Recompose my desk by function instead of habit: a work zone, a paper-in zone, a paper-out zone, and one small zone for things with no home yet. Describe the new layout."

Why it worked. The "no home yet" zone was the honest invention — clutter is usually just items awaiting a category.

Watch for. Recomposing rearranges; it does not repair. A chaotic week rebuilt in a better order is still a chaotic week if it holds nine hours of work in six hours of space. After the new arrangement, ask the rude question: what, in total, did we remove?

Placed well: "Recompose my week: errands batched, deep work in the mornings, one evening protected completely."

Stage

What it makes the AI do. It divides a long effort into phases with gates between them — each stage finished, checked, and only then the next begun.

Weak prompt: "How do I train for a 5K run?"

Same prompt, with the word: "Stage a beginner 5K plan into three stages — walking and easy jogging, then steady running, then distance — and tell me what I should be able to do before moving to each next stage."

Why it worked. A single plan hides its own beginning. Stages with gates made progress measurable — and quitting harder, because there was always a next small door.

Another example.

Weak prompt: "Help me write a book."

Same prompt, with the word: "Stage writing a short book: outline, rough draft, honest critique, revision, polish. For each stage, tell me what 'done' looks like before I'm allowed to move on."

Why it worked. The gates stopped the classic failure — polishing chapter one for a month while the book stayed unwritten.

One more.

Weak prompt: "How do I lose 20 pounds?"

Same prompt, with the word: "Stage the 20 pounds: stage one is habits only, no scale goals; stage two adds the kitchen; stage three adds activity. Give each stage a graduation test that isn't a number."
Why it worked. The gates measured capability instead of weight — and capabilities, unlike Mondays, mostly move in one direction.
Watch for. Stages can become a waiting room: stage one polished forever, its graduation test always nearly passed. The gate exists to be walked through, not admired. If a stage has lasted twice as long as planned, the honest move is shrinking it, not decorating it.
Placed well: "Stage this book: outline, draft, critique, revision — and tell me what 'done' looks like at each gate."

Chunk

What it makes the AI do. It cuts big material into sittings sized to a real attention span, each with an ending point instead of a bookmark and a prayer.
Weak prompt: "This training manual is 200 pages. Help."
Same prompt, with the word: "Chunk this manual into six sittings of about 30 minutes each. For each chunk, give me its main ideas in advance and two questions to answer when I finish it."
Why it worked. Two hundred pages is a wall; six sittings are a path. The questions at each border turned reading into arriving.
Another example.
Weak prompt: "My inbox has 900 unread emails."
Same prompt, with the word: "Chunk the inbox cleanup: sort first into four piles — answer, file, delete, decide — and give me a 15-minute routine per pile that I can repeat daily."
Why it worked. The pile was never 900 problems; it was four small problems, repeated. Chunking made the first day finishable.
One more.
Weak prompt: "Teach me photography."
Same prompt, with the word: "Chunk photography into weekend lessons — one idea and one outing per chunk, ending each with a single photo assignment I can finish in an hour."
Why it worked. A hobby became a course with doors in it — and hour-sized assignments fit real Saturdays.

Watch for. Chunks fail at their borders. Cut a manual at page counts alone and one sitting ends mid-argument while the next opens mid-sentence, and both feel broken. Chunk by ideas, not arithmetic — each sitting should end where a thought completes, with a question answered.
Placed well: "Chunk this manual into six sittings by topic, each ending with two questions I should be able to answer."

Layer

What it makes the AI do. It stacks the same explanation at several depths — a sentence, a paragraph, a page — so you can stop at the layer you need and descend only where it matters.
Weak prompt: "Explain how the internet works."
Same prompt, with the word: "Layer the explanation of how the internet works: first in one sentence, then in one paragraph, then in a full page for where I want more — and label each layer so I can stop when I've got enough."
Why it worked. Most explanations pick one depth and dare you to be satisfied. Layers respect the reader's real need, which is usually "enough to follow the news, plus one deep dive."
Another example.
Weak prompt: "Explain my lab results."
Same prompt, with the word: "Layer this lab report: layer one, am I okay; layer two, what each number means; layer three, what questions to bring to my doctor."
Why it worked. The layers matched the emotional order of the questions — reassurance first, understanding second, action third.
One more.
Weak prompt: "Explain quantum computing."
Same prompt, with the word: "Layer it: one sentence for a neighbor, one paragraph for a curious teenager, one page for me — and be honest at each layer about what the layer is leaving out."
Why it worked. The honesty labels were the upgrade: he knew not just the explanation, but its thickness.
Watch for. Layers tempt skimming with a clean conscience: layer one read, layers two and three forever pending, and the reader certain they "basically understand." Sometimes that's fine. On lab results and

contracts, decide in advance which layer you're obliged to reach — then reach it.
Placed well: "Layer this report: am I okay, what the numbers mean, and what to ask my doctor — labeled."

Map

What it makes the AI do. It shows how the pieces connect — who affects whom, what feeds what — instead of describing each piece alone in its own paragraph.
Weak prompt: "I'm lost in this novel. Who are all these people?"
Same prompt, with the word: "Map the relationships in this novel: group the characters by family and alliance, and for each one give me their connection to the main character in a single line."
Why it worked. A cast list is alphabetical; a map is meaningful. Grouped by connection, the crowd became a structure he could hold in his head.
Another example.
Weak prompt: "Why does the office coffee fund keep running out?"
Same prompt, with the word: "Map the coffee fund as a system: where money enters, where it leaves, who decides purchases, and where the leak most likely sits."
Why it worked. Small mysteries are usually systems with one loose joint. The map pointed at the joint instead of at suspects.
One more.
Weak prompt: "Explain why our town's Main Street is struggling."
Same prompt, with the word: "Map Main Street as a system: the anchor stores, the parking, the highway bypass, the rents — and show which pieces pull on which."
Why it worked. Arguments about single causes ended. The map showed the bypass and the rents holding hands, which is where the real conversation starts.
Watch for. A map shows connections so confidently that invented ones slip in wearing the same lines as real ones — two departments "linked" because both were mentioned near each other. Ask the machine to mark which connections are documented and which are its reasonable guesswork.
Placed well: "Map the coffee fund as a system: where money enters, where it leaves, who decides, and where it leaks."



Part Four — Challenge

Stress-test the answer. These words turn the machine from an agreeable assistant into an honest opponent.

Machines are raised to be agreeable, the way some dogs are raised to fetch: it is deep in them, and it is not the same as being honest. Ask one "sounds good, right?" and you will receive a small ceremony of agreement, fluently delivered, whether or not your plan deserves it. This is pleasant in the moment and expensive later, when reality — which was never consulted — delivers its own review.

The words in this part exist to break the ceremony. They deputize the machine as your opponent, your examiner, your stress engineer: make the other side's best case, find the crack in my plan, argue against what you just told me. Used well, this is not negativity. It is the cheapest form of experience there is — meeting the strongest objections in a conversation, where they cost nothing, instead of in the world, where they invoice you. If you only master three groups in this book, most working people should make this one of them.

Steelman

What it makes the AI do. It requires the strongest possible version of the other side's argument — not the weakest one that's easy to knock down.

Weak prompt: "Why are people against school vouchers?"

Same prompt, with the word: "Steelman the case against school vouchers: state it the way its smartest supporter would, with its best evidence, as if you believed it — then do the same for the case in favor."

Why it worked. Machines default to summarizing sides politely, which usually means weakly. A steelman forces both cases to full strength — and decisions made against strong arguments are the ones that hold up.

Another example.

Weak prompt: "Is remote work better than office work?"

Same prompt, with the word: "Steelman both sides: the case for full-time office work as its best defenders make it — mentorship, culture, trust — then the case for remote work at the same strength. No strawmen on either side."

Why it worked. Both cases at full strength made the trade-offs real instead of tribal, which is the only state in which this argument is worth having.

One more.
Weak prompt: "Why do people oppose the new housing development?"
Same prompt, with the word: "Steelman the neighbors opposing the development: their case as its most reasonable member would make it at the microphone, without caricature."
Why it worked. He still supported the project — but stopped being surprised by the meetings, which is its own form of power.
Watch for. The trap is enjoying the exercise too much: a beautiful steelman of a weak position can make it sound stronger than its real believers ever manage. Remember which side you were testing, and weigh the argument people actually make alongside the finest version imaginable.
Placed well: "Steelman the neighbors' case against this development, as their most reasonable member would make it."

Invert

What it makes the AI do. It flips the question around: instead of how to succeed, how to guarantee failure; instead of what to do, what to avoid.
Weak prompt: "How can I have a successful garage sale?"
Same prompt, with the word: "Invert it: list everything that would guarantee a garage sale fails — pricing, signage, timing, layout — so I can make sure I do none of it."
Why it worked. Success has a thousand recipes; failure has a short list of reliable causes. Avoiding the definite mistakes is often worth more than chasing the uncertain tricks.
Another example.
Weak prompt: "How can I sleep better?"
Same prompt, with the word: "Invert it: list the habits that reliably destroy a night's sleep, in order of damage, so I can audit my own evenings against them."
Why it worked. The failure list was short, specific, and — the uncomfortable part — familiar. Knowing exactly what wreckage looks like turned out to be most of avoiding it.
One more.
Weak prompt: "How do I give a great job interview?"
Same prompt, with the word: "Invert it: list the reliable ways candidates ruin interviews in the first five minutes, so I can simply not do those."

Why it worked. Excellence is mysterious; ruin is documented. He memorized the ruin list on the drive there.
Watch for. Inversion produces failure lists so vivid they can become the plan's whole personality — a garage sale organized entirely around fear. Avoiding ruin is the floor, not the ceiling. After the invert, return once to the positive question and build something worth not-ruining.
Placed well: "Invert my interview prep: list the reliable ways candidates ruin the first five minutes, so I do none."

Critique

What it makes the AI do. It switches the machine from author to reviewer, judging work against standards instead of politely continuing it.
Weak prompt: "Here's my business plan. What do you think?"
Same prompt, with the word: "Critique this business plan the way a bank loan officer would: name the three claims you'd least believe, the numbers you'd ask to see proven, and the question I'd least want asked."
Why it worked. "What do you think?" earns encouragement. A critique with a named judge and named standards earns the truth — early, while it's still free.
Another example.
Weak prompt: "What do you think of this photo?"
Same prompt, with the word: "Critique this photo the way a photography teacher would: composition first, then light, then the single change that would improve it most. Honest over kind."
Why it worked. Praise is pleasant and nearly useless. One named biggest-fix turned a compliment session into a lesson he could use on the next roll.
One more.
Weak prompt: "Review my poem."
Same prompt, with the word: "Critique this poem the way a working editor would: which stanza earns its place least, which line is showing off, and what one cut would strengthen the whole."
Why it worked. "One cut" gave the critique a destination. The stanza he defended hardest was, of course, the one to go.
Watch for. A critique is only as honest as the judge you name. Ask for a review "like a teacher" and you'll get a teacher's mercy free. Name a

standard with consequences — a loan officer, an editor — and say whether kind or true matters more.
Placed well: "Critique this plan like a loan officer: name the three claims you'd least believe, out loud."

Defend

What it makes the AI do. The opposite assignment: it makes the machine argue *for* a position under attack, which tests whether the position has real support or only momentum.
Weak prompt: "Is it smart to pay off a mortgage early?"
Same prompt, with the word: "Defend paying off a mortgage early against the standard objection that investing the money earns more. Use the defender's strongest points, including the ones that aren't about math."
Why it worked. A position you can't hear defended is a position you're holding on faith. Defense surfaces the non-obvious arguments — risk, sleep, freedom — that summaries flatten.
Another example.
Weak prompt: "Are extended warranties a scam?"
Same prompt, with the word: "Defend buying the extended warranty on a used truck against the standard 'the house always wins' argument — and include the specific situations where the warranty genuinely pays."
Why it worked. Hearing the defense separated the rare good case from the usual bad bet. A thing can be mostly a poor deal and still be the right deal sometimes.
One more.
Weak prompt: "Is it worth paying for the good seats?"
Same prompt, with the word: "Defend paying for the expensive seats at a once-a-year concert against the standard 'the lawn sounds fine' argument — including the reasons that aren't about sound."
Why it worked. The defense surfaced the real currency: once a year, tired legs, and being close enough to see hands on strings.
Watch for. Defense is advocacy, not inquiry: a machine told to defend will marshal the position's best evidence and quietly bench its worst. Enjoy the case, then remember somebody still has to prosecute it. Pair every defense with the strongest attack before you decide anything.

Placed well: "Defend paying off the mortgage early against the investing argument — including the reasons that aren't math."

Disagree

What it makes the AI do. It gives the machine explicit permission — really, an instruction — to push back on you, counteracting its built-in urge to agree.

Weak prompt: "I'm planning to buy a rental property in my town. Sounds solid, right?"

Same prompt, with the word: "Disagree with me. I'm planning to buy a rental property in my town — argue that it's a bad idea, using the strongest reasons a skeptical accountant would raise."

Why it worked. Asked "sounds solid, right?", the machine hears its cue to applaud. "Disagree" rewrites the cue — and hearing the case against your plan before the bank does is the cheapest insurance there is.

Another example.

Weak prompt: "I'm thinking of buying a boat. Exciting, right?"

Same prompt, with the word: "Disagree with me about the boat. Give me the strongest reasons owners quietly regret buying one, priced in money, weekends, and storage."

Why it worked. Excitement asked for applause. Disagreement delivered the ownership costs that no brochure and no boat show will ever itemize.

One more.

Weak prompt: "I'm going to start getting up at 5 a.m. Great idea, right?"

Same prompt, with the word: "Disagree with me about the 5 a.m. plan: argue, from my actual evenings, why it will fail by Thursday — and what would have to change for it to survive."

Why it worked. The applause version would have cost him a week. The disagreement cost an evening's honesty and gave back a plan with a bedtime in it.

Watch for. Disagreement, requested once, fades fast — by the third reply the machine is drifting back toward applause, its factory setting. Restate the assignment when the subject changes, and be specific about whose skepticism you want, or you'll receive generic doubt in a firm tone.

Placed well: "Disagree with me about this rental: argue it's a bad buy, as a skeptical accountant would."

Counterfactual

What it makes the AI do. It runs the world that didn't happen — "what if X had gone the other way?" — which separates luck from skill and cause from coincidence.

Weak prompt: "Why did my restaurant do so well last summer?"

Same prompt, with the word: "Counterfactual check: my restaurant did well last summer. Which of these causes would still have produced a good summer if the others were removed — the new menu, the road construction ending, or the hot weather driving people indoors?"

Why it worked. Success stories credit everything that was present. Counterfactuals remove one cause at a time and see what collapses — that's how you learn what to repeat.

Another example.

Weak prompt: "Why did my tomatoes do so well this year?"

Same prompt, with the word: "Counterfactual: my tomatoes thrived. Which single change mattered most — the new raised bed, the drip watering, or the warmer spring? What would each have done on its own?"

Why it worked. Credit split three ways teaches nothing. Removing the causes one at a time found the one worth paying for again next spring.

One more.

Weak prompt: "Our fundraiser did great this year."

Same prompt, with the word: "Counterfactual: the fundraiser set a record. Remove the new venue, the celebrity guest, and the good weather one at a time — which removal hurts the total most?"

Why it worked. The answer shocked the committee: the weather. Next year's budget moved accordingly, from guests to tents.

Watch for. Counterfactuals are stories about a world nobody visited, and the machine narrates them with the confidence of history. The tomato verdict — drip watering did it — is a hypothesis, not a finding. Treat it as an experiment: change one thing at a time.

Placed well: "Counterfactual: remove the new venue, the guest, and the weather one at a time — which hurts most?"

Anticipate

What it makes the AI do. It moves the machine forward in time: objections, questions, and problems that haven't happened yet, raised while you can still prepare.

Weak prompt: "Help me prepare for my job interview."
Same prompt, with the word: "Anticipate the five hardest questions a hiring manager would ask about the two-year gap in my work history, and help me build an honest, confident answer to each."
Why it worked. Generic preparation rehearses what you hope they'll ask. Anticipation rehearses what you fear they'll ask — which is, of course, what they'll ask.
Another example.
Weak prompt: "Help me write to a landlord about renting a place."
Same prompt, with the word: "Anticipate the questions a landlord will ask a first-time renter who has a dog, and help me answer each one before it's asked — including which documents to have ready."
Why it worked. Arriving pre-answered turns an interrogation into a formality. The dog, addressed early and honestly, stopped being a surprise.
One more.
Weak prompt: "Help me prepare for the parent-teacher conference."
Same prompt, with the word: "Anticipate what the teacher will raise about my son — the good first, then the concerns — and help me prepare one question of my own for each concern."
Why it worked. Walking in with questions turned a report card into a collaboration.
Watch for. Anticipation can over-prepare you for the interview in your head and under-prepare you for the person in the chair, who asks about parking. Cover the feared questions, yes — then leave slack. A plan that only fits the anticipated meeting breaks at the first surprise.
Placed well: "Anticipate a landlord's five questions for a first-time renter with a dog, and answer each one early."

Simulate

What it makes the AI do. It makes the machine *play* a situation with you — the negotiation, the interview, the difficult conversation — instead of describing it from a safe distance.
Weak prompt: "How do I ask for a discount on a car?"
Same prompt, with the word: "Simulate the car negotiation with me: you play a friendly but firm salesperson, I'll ask for a discount, and you push back the way a real salesperson would. Start when I say 'begin.'"

Why it worked. Advice about conversations is theory; rehearsing one is practice. Simulation lets you make your mistakes against a machine that charges nothing for them.

Another example.

Weak prompt: "How do I talk to my neighbor about his loud music?"

Same prompt, with the word: "Simulate the conversation: you play my neighbor — friendly, but defensive about his music. I'll open. React realistically, and afterward tell me where my wording made it worse."

Why it worked. The rehearsal exposed the exact sentence that would have started the argument — in practice, for free, instead of in his driveway, for real.

One more.

Weak prompt: "I have to ask my roommate to do his dishes."

Same prompt, with the word: "Simulate it: you're my roommate, easy-going and honestly unaware it's a problem. I'll start. Stay in character for five exchanges, then step out and grade my opening."

Why it worked. His first attempt, in rehearsal, was an accusation wearing a request's clothes. The second one did the dishes' work before the sink.

Watch for. The simulation is a sparring partner who secretly wants you to win: neighbors soften by exchange three, salespeople fold at reasonable offers. If the rehearsal feels easy, say so and raise the temperature — "less friendly, more realistic" — before the real conversation grades you.

Placed well: "Simulate my neighbor — friendly but defensive about his music — and grade my opening afterward."

Stress-Test

What it makes the AI do. It pushes a plan against hard conditions — the delay, the breakdown, the person who doesn't show up — to find where it cracks while fixing is still theoretical.

Weak prompt: "Is my hurricane plan good?"

Same prompt, with the word: "Stress-test this hurricane plan: assume the power is out for three days, one road is flooded, and the pharmacy is closed. Walk day by day and show me where the plan fails."

Why it worked. Plans are always good in the weather they were imagined in. Load revealed the two real weaknesses — water storage and one prescription — while they were still cheap to fix.

Another example.
Weak prompt: "Here's my small business budget."
Same prompt, with the word: "Stress-test this budget: my biggest customer pays 60 days late and my rent rises 10% in the same quarter. Which month breaks, and what breaks it?"
Why it worked. The budget survived both problems separately and failed them together — a discovery worth far more than either test alone.
One more.
Weak prompt: "Here's the plan for our outdoor wedding."
Same prompt, with the word: "Stress-test it against rain arriving at the ceremony hour and the caterer running an hour late, at the same time. Walk the timeline and show the two moments that collapse."
Why it worked. Each problem alone had a fix; together they shared one tent and one hour. The test cost an umbrella rental and a phone call.
Watch for. Test one disaster at a time and a plan can pass every exam while failing the semester — rain alone fine, caterer late alone fine, together fatal. The real world sends weather in combinations. Finish by stacking your two worst cases on the same day.
Placed well: "Stress-test this wedding: rain at the ceremony hour and the caterer an hour late, simultaneously."

Falsify

What it makes the AI do. It asks the strongest question in science: what evidence would prove this claim *wrong* — and does that evidence exist?
Weak prompt: "Is it true that we lose most body heat through our heads?"
Same prompt, with the word: "Falsify the claim that we lose most body heat through our heads: what experiment or observation would disprove it, what does the actual evidence show, and where did the claim come from?"
Why it worked. Ordinary checking asks "is this supported?" — and everything sounds supported somewhere. Falsifying asks what would kill the claim, which is a much shorter road to the truth.
Another example.
Weak prompt: "This article says a new supplement boosts memory."

Same prompt, with the word: "Falsify this supplement claim: describe the study that would settle it, check whether that study is the one being cited, and tell me what the cited study actually measured."
Why it worked. The cited study existed; it had measured something else, in mice. Falsification traced the gap between the headline and the evidence.
One more.
Weak prompt: "People say breakfast is the most important meal."
Same prompt, with the word: "Falsify it: state the claim precisely, describe the evidence that would kill it, and report what the actual research record shows, including who first promoted the phrase."
Why it worked. The phrase survived as marketing and died as science — a much more interesting answer than yes or no.
Watch for. Falsification can be turned into theater: ask what would disprove a claim and receive a grand experiment nobody could run, while the small decisive check goes unmentioned. Push for the cheapest test that would change your mind — then notice if you'd accept its verdict.
Placed well: "Falsify this supplement claim: name the study that would settle it, and check what was actually cited."

Worst-Case

What it makes the AI do. It starts from the bad outcome and plans backward: if this goes wrong, exactly how does it go, and what will you wish you had done today?
Weak prompt: "Should I worry about my 45-minute flight connection?"
Same prompt, with the word: "Worst-case this connection for me: if I miss the second flight, what exactly happens — last flight of the day, hotel, rebooking — and what can I set up now, today, that makes that night survivable?"
Why it worked. Worry is worst-case thinking without the paperwork. Doing it deliberately produced a packed carry-on, a screenshot of the later flights, and a much calmer traveler.
Another example.
Weak prompt: "We're hiring our first employee."
Same prompt, with the word: "Worst-case this hire: the person quits in month two, mid-project, on bad terms. What should we document, cross-train, and put in writing starting now?"

Why it worked. Planning from the bad ending named protections that optimism never mentions — and most of them made the good outcome better too.

One more.

Weak prompt: "We're lending our car to our teenager for a road trip."

Same prompt, with the word: "Worst-case the trip: breakdown at night, two states away, phone dead. Build the kit and the instructions that make that night a story instead of an emergency."

Why it worked. The bad night, planned for, mostly dissolved — what remained was a tire gauge, a paper map, and a charged battery pack.

Watch for. Worst-case thinking, left running, becomes worry with a spreadsheet — every hire a future resignation, every trip a breakdown. Use it to build the kit and the paperwork, then close the file. The point of visiting the bad night in advance is not living there.

Placed well: "Worst-case this connection: if I miss the last flight, what happens, and what do I set up today?"

Rebut

What it makes the AI do. It makes the machine argue directly against its own previous answer — not review it politely, but prosecute it.

Weak prompt: "Okay, you recommended the blue car. Anything else?"

Same prompt, with the word: "Now rebut your own recommendation: argue against the blue car as if you were its toughest critic, using facts a buyer would discover in the first year of owning it."

Why it worked. A first recommendation is a position; a rebuttal is an examination. The second answer surfaced the repair costs and the weak resale that the first had politely left out.

Another example.

Weak prompt: "You said I should take the new job."

Same prompt, with the word: "Rebut that: build the strongest case that taking this job is a mistake, aimed at the exact reasons you just gave for taking it."

Why it worked. Aimed at its own reasoning, the rebuttal could not wander into generic doubts. It had to fight the actual argument — and two of its points held up, which was the useful part.

One more.

Weak prompt: "You helped me draft this email to decline the project."

Same prompt, with the word: "Now rebut the draft: read it as the client and tell me every place it could sound dismissive, evasive, or colder than I intend."
Why it worked. The email was defended against its recipient before the recipient ever saw it — three phrases changed temperature.
Watch for. A rebuttal aims where you point it, and lazy pointing gets theater: the machine attacks the weakest reading of its own advice and declares victory. Point the rebuttal at the strongest claim in the original answer, by name, and make it answer that one.
Placed well: "Rebut your own recommendation: argue against this car using facts I'd discover in year one of owning it."



Part Five — Verify

Make it check itself. These words treat the first answer as a draft of the truth, not the truth.

A machine's first answer is a draft of the truth wearing a suit. It sounds finished. It is formatted like finished. Occasionally it is simply, confidently wrong — a date off by a decade, a rule that was repealed, a number with no visible parents. The danger was never that machines make mistakes; people make mistakes. The danger is that their mistakes arrive with the same polish as their facts, and nothing in the typography tells you which is which.

Verification words are how you sort them. They make the machine show its sources, grade its confidence, do the sum a second way, follow the claim home to where it was born. Some of these words feel almost rude — asking a helpful assistant to audit itself, rebut itself, flag its own soft spots. Use them anyway, especially when the answer will be acted on: money spent, medicine taken, documents signed. Trust, in this shop, is not a feeling about the machine. It is a set of habits, and this part is the habit list.

Verify

What it makes the AI do. It sends the machine back over its own answer, claim by claim, hunting for the parts that are wrong, dated, or merely confident-sounding.

Weak prompt: "You told me this supplement helps with sleep. Anything else?"

Same prompt, with the word: "Verify what you just told me about this supplement: go claim by claim and mark each one as well-established, mixed evidence, or mostly marketing."

Why it worked. Machines rarely volunteer doubt about their own previous paragraph. "Verify" assigns the doubt — and grading the evidence beats repeating the claim.

Another example.

Weak prompt: "Tell me what happened at the first Thanksgiving."

Same prompt, with the word: "Now verify your own account of the first Thanksgiving: go detail by detail and mark each as documented, tradition, or probably myth."

Why it worked. The machine graded its own story — and the myth column was the most interesting part, which the first telling had presented with a straight face.

One more.

Weak prompt: "You gave me directions to the lake house."

Same prompt, with the word: "Verify those directions: check each turn against the route, mark any road name that's changed, any step that depends on a landmark that might be gone, and the total distance both ways."

Why it worked. Directions decay like milk. The verification found one renamed road — found it at the kitchen table, not at a dark intersection.

Watch for. Verification has a quiet limit: the machine checks its answer against its own memory, so a fact it learned wrong gets verified wrong, twice, in matching tones. For anything that will be signed, swallowed, or spent on, let one claim be checked somewhere else.

Placed well: "Verify your last answer claim by claim: mark each as established, mixed evidence, or mostly marketing."

Audit

What it makes the AI do. A verify with a checklist and a clipboard: it examines something whole — a budget, a plan, a document — looking systematically for errors, waste, and risk.

Weak prompt: "Look at my monthly subscriptions and tell me if they're okay."

Same prompt, with the word: "Audit these subscriptions: total the annual cost, flag any that duplicate each other, any that rose in price, and any a typical household subscribes to and never uses."

Why it worked. "Okay" is unanswerable, so the machine guesses politely. An audit has procedures — totals, duplicates, drift — and procedures find things impressions miss.

Another example.

Weak prompt: "My weeks feel too busy."

Same prompt, with the word: "Audit my typical week from this calendar: total the hours by category, flag any meeting with no stated purpose, and find where two honest free hours could come from."

Why it worked. "Busy" is a feeling, and feelings don't negotiate. The audit produced numbers — and two recoverable hours hiding inside them.

One more.
Weak prompt: "My pantry is chaos."
Same prompt, with the word: "Audit the pantry from this list: group everything by kind, flag duplicates, anything expired, and anything bought for a recipe that never happened."
Why it worked. The audit found three jars of the same spice and a small museum of 2023. Shopping since then starts from evidence.
Watch for. An audit finds what its checklist contains — subscriptions get totaled beautifully while the annual charge hiding under a different name sails past. Checklists come from somewhere, usually habit. Before trusting a clean audit, ask the auditor once: "What weren't you asked to look for?"
Placed well: "Audit these subscriptions: annual total, duplicates, price rises, and anything we pay for and never open."

Cite

What it makes the AI do. It attaches a source to each claim, which both lets you check the claim and quietly discourages the machine from inventing one.
Weak prompt: "What are the health benefits of walking?"
Same prompt, with the word: "List the well-documented health benefits of walking 30 minutes a day, and cite the kind of source behind each — a major medical body, a large study, or 'commonly claimed but thinly sourced.'"
Why it worked. A claim with a source named can be weighed; a claim floating free all sounds equal. Asking for the *kind* of source works even where exact links can't be checked — it forces an honesty grade onto every line.
Another example.
Weak prompt: "Is coffee good or bad for you?"
Same prompt, with the word: "Give me the main claims about coffee and health, and cite the kind of evidence behind each — a large long-term study, a small short trial, or headline science."
Why it worked. Sorting the claims by evidence type ended the whiplash of contradictory headlines. Most of them, it turned out, were standing on very different ground.
One more.

Weak prompt: "Tell me about our town's founding."
Same prompt, with the word: "Tell me the town's founding story, but cite the kind of source behind each part — county record, newspaper of the day, or family legend nobody wrote down for eighty years."
Why it worked. The legend parts were the beloved parts, which is fine — as long as they wear their name tags.
Watch for. A citation is not a character reference. Machines can name a respectable kind of source for a claim no such source made, and "a large study shows" has laundered many inventions. For the claim you'll act on, open the source — or admit nobody has.
Placed well: "Cite the kind of source behind each benefit: major medical body, large study, or thinly sourced claim."

Cross-Check

What it makes the AI do. It makes the machine reach the same answer by a second, independent route and report whether the two routes agree.
Weak prompt: "How many tiles do I need for a 12-by-14-foot floor?"
Same prompt, with the word: "Calculate how many 12-inch tiles I need for a 12-by-14-foot floor, then cross-check it a completely different way, and tell me if the two methods agree before I order."
Why it worked. One calculation can be wrong with total confidence. Two methods that independently land on the same number are worth ordering tiles on — and when they disagree, you've found the mistake at the cheapest possible moment.
Another example.
Weak prompt: "What's a fair tip on a $47.60 check?"
Same prompt, with the word: "Work out a 20% tip and the total on this $47.60 check two completely different ways, then cross-check that the two totals agree before I pay."
Why it worked. One calculation is a claim; two routes landing on the same number is what "sure" feels like in arithmetic — and it costs about ten seconds.
One more.
Weak prompt: "How much paint do I need for a 12-by-15 room?"
Same prompt, with the word: "Calculate the paint for a 12-by-15 room with 9-foot walls two ways — by wall area and by the coverage on the can — and cross-check them before I buy."

Why it worked. The two routes disagreed by exactly one gallon, which was one trip back to the store, saved.
Watch for. Two methods only cross-check if they're genuinely independent. Recalculating with the same wrong measurement — the room paced off at eleven feet that is really twelve — produces beautiful agreement on a wrong answer. Change the route and the tool: tape measure second, arithmetic first.
Placed well: "Cross-check the tile count two different ways, and tell me whether they agree before I order."

Trace

What it makes the AI do. It follows a claim, a number, or a charge back to where it actually came from, step by step, instead of accepting it at face value.
Weak prompt: "Why does everyone say we only use 10% of our brains?"
Same prompt, with the word: "Trace the claim that humans use only 10% of their brains: where did it start, what real finding got twisted into it, and why does it survive?"
Why it worked. Famous claims are usually fossils — something true, buried and distorted over decades. Tracing unwraps the layers, and the original finding is almost always more interesting than the myth.
Another example.
Weak prompt: "Do we really swallow spiders in our sleep?"
Same prompt, with the word: "Trace that claim: where did it first appear in print, was it ever based on an actual study, and how did it travel so far before the internet existed to help it?"
Why it worked. Tracing found the tale's birthplace — a magazine column that mentioned it while making the opposite point, about how easily falsehoods spread. It had then demonstrated the point by spreading.
One more.
Weak prompt: "Where does 'break a leg' come from?"
Same prompt, with the word: "Trace 'break a leg': the earliest printed uses, the competing origin stories, and which one the evidence actually supports."
Why it worked. Three origin stories entered; one had documents. Theater people, it turns out, keep excellent receipts.

Watch for. Tracing rewards a good story, and origins are often messier than one birthplace: phrases and findings can be born twice, in two towns, neither copying the other. When the machine presents a single neat lineage, ask what the rival lineages are before repeating it.
Placed well: "Trace 'break a leg': earliest printed uses, competing origins, and which one the evidence supports."

Estimate

What it makes the AI do. It asks for a reasoned approximation with its workings shown — assumptions named, arithmetic visible — instead of a guess wearing a number's clothes.
Weak prompt: "How many people could fit in a stadium parking lot?"
Same prompt, with the word: "Estimate how many people could stand in a stadium parking lot. Show your assumptions — lot size, space per person — and the arithmetic, so I can adjust them."
Why it worked. A bare number can't be argued with or improved. An estimate with visible assumptions is a machine you can tune: change one input and the answer updates honestly.
Another example.
Weak prompt: "How long will this book take me to read?"
Same prompt, with the word: "Estimate my reading time for this 320-page novel: show your assumptions — a realistic pages-per-hour for fiction, and my free evenings — and give me a range, not a date."
Why it worked. With the assumptions on the table, the estimate became a forecast he could adjust: slower reader, longer book, the arithmetic still held.
One more.
Weak prompt: "How many chairs fit in the hall for the reunion?"
Same prompt, with the word: "Estimate the chairs for a 40-by-60 hall: show your assumptions — space per chair, room for tables and a dance floor — so I can argue with them."
Why it worked. He did argue with them — dance floor doubled, chair count honest. Both numbers survived the argument.
Watch for. An estimate's assumptions are its soft underbelly: show the workings, yes, but the workings were chosen by the same machine that wanted a tidy answer. Argue with the inputs, not the arithmetic — swap in your lot size, your pace, and watch the number move.

Placed well: "Estimate the chairs for this hall: show your assumptions about space, tables, and the dance floor."

Flag

What it makes the AI do. It makes the machine mark its own uncertainties inside the answer — the dates it's unsure of, the figures it estimated, the parts that depend on where you live.

Weak prompt: "Explain the rules for deducting a home office."

Same prompt, with the word: "Explain the rules for deducting a home office, and flag every point where the rule changed recently, varies by situation, or where you're not fully certain — I'd rather see a flag than a confident guess."

Why it worked. The most dangerous sentence an AI produces is a wrong one written confidently. Flags move the doubt from invisible to printed, where you can act on it.

Another example.

Weak prompt: "Summarize these furniture assembly instructions."

Same prompt, with the word: "Summarize these assembly instructions, and flag every step where the wording is ambiguous, two parts are easy to mix up, or a mistake would be hard to undo."

Why it worked. The flags marked the three steps where care actually mattered — delivered before the wrench taught the same lesson at higher tuition.

One more.

Weak prompt: "Explain this car loan offer."

Same prompt, with the word: "Explain the loan, and flag the three places where the offer's friendly wording does the most work: anything that can change, anything that costs extra, anything that assumes I'll refinance."

Why it worked. The flags sat exactly where the salesman hoped his eyes would slide past. They didn't.

Watch for. Flags multiply if you let them: ask for every uncertainty and the answer comes back wearing more warning labels than a ladder, until the one flag that matters is camouflaged. Ask for the three flags you'd act on — the rest is weather reporting.

Placed well: "Flag the three places in this loan where friendly wording does the most work."

Reconcile

What it makes the AI do. It hands the machine two numbers that should match and don't — two answers, two documents, two totals — and makes it find where they parted company.

Weak prompt: "My bank balance and my budget spreadsheet don't agree. Help."

Same prompt, with the word: "Reconcile these two figures: my bank says $2,340 and my spreadsheet says $2,515. Walk through the most likely causes of a $175 difference, in order, and tell me how to check each one."

Why it worked. A mismatch feels like chaos; a reconciliation treats it as a specific, findable event. Naming the exact difference turns "something's wrong" into a short list of suspects.

Another example.

Weak prompt: "My two weather apps disagree about Saturday."

Same prompt, with the word: "Reconcile them: one says rain, one says sun. What is each basing it on, which has been more accurate in this area this season, and what should I actually plan the picnic around?"

Why it worked. Instead of averaging two forecasts into a meaningless maybe, he learned why they differed — and which one had earned trust locally.

One more.

Weak prompt: "My scale and my doctor's scale disagree by six pounds."

Same prompt, with the word: "Reconcile the two scales: the likely causes of a six-pound disagreement, how to test mine at home with known weights, and which one to trust for tracking change."

Why it worked. The verdict — his scale read light but consistently — was more useful than either number alone.

Watch for. Reconciliation can be gamed by tidiness: given two totals, the machine will find a plausible $175 suspect and stop, satisfied. Plausible is not found. Make it say what evidence would confirm the culprit — a statement line, a receipt — before you close the case.

Placed well: "Reconcile these balances: walk the likely causes of the $175 difference in order, with a check for each."

Sanity-Check

What it makes the AI do. It tests whether a number is even in the plausible band — not precisely right, just not absurd — before anyone builds plans on top of it.

Weak prompt: "A contractor quoted me $4,800 for new gutters. Is that fair?"

Same prompt, with the word: "Sanity-check this $4,800 gutter quote: what's the typical installed cost per foot for a house like mine, what would the total be at that rate, and is the quote inside or outside the plausible band?"

Why it worked. Fairness is a research project; plausibility is a quick gate. The quote failed the gate — nearly double the band's top — and that one finding justified every phone call that followed.

Another example.

Weak prompt: "This diet says I can lose 30 pounds in a month."

Same prompt, with the word: "Sanity-check losing 30 pounds in a month: what daily calorie deficit that requires, and how it compares with what doctors consider a safe rate."

Why it worked. The arithmetic itself was the verdict. No opinion was needed once the numbers had to stand next to each other.

One more.

Weak prompt: "This moving company estimates 800 pounds of stuff per room."

Same prompt, with the word: "Sanity-check that: what does a typical furnished room actually weigh, and does their per-room figure sit inside the honest range or above it?"

Why it worked. The estimate was high by half, and weight is the price. One plausibility gate paid for itself several times over.

Watch for. The band is only as honest as the "typical" it's built from, and prices in its memory may belong to another year or another state. A quote can fail the band test because the band is stale. Pair the sanity check with one local call.

Placed well: "Sanity-check this gutter quote: typical cost per foot, the implied total, and inside or outside the band?"

Grade

What it makes the AI do. It makes the machine rate its own confidence, in the open — an A for settled fact, a C for shaky ground — so you know which parts of an answer to lean on.

Weak prompt: "Give me the key dates of the Civil War."

Same prompt, with the word: "Give me the key dates of the Civil War, and grade your confidence in each — A for certain, B for solid, C for 'check this' — with a note on anything graded below A."

Why it worked. A flat list presents every date as equally true. Grades sorted the answer into "rely on it" and "look it up," which is how the information would actually be used.

Another example.

Weak prompt: "What will house prices do next year?"

Same prompt, with the word: "Give me your best reading of house prices next year, then grade it honestly: what part is trend, what part is guesswork, and what would change your mind?"

Why it worked. The grade forced the forecast to confess its own ingredients — and the confession was more valuable than the prediction.

One more.

Weak prompt: "When will self-driving cars be everywhere?"

Same prompt, with the word: "Give me your best answer, graded: which parts are already true somewhere, which are extrapolation, and which single event would most change your grade?"

Why it worked. He stopped holding a date and started holding a map of the uncertainty — which is what the question was actually for.

Watch for. Self-grades run generous. Left alone, a machine will award its solid claims an A and its inventions a confident B+, and the whole answer graduates with honors. Calibrate it: demand to know what a C looks like in this answer, and watch the curve appear.

Placed well: "Grade your confidence in each date — A, B, or C — and explain anything graded below A."

Refresh

What it makes the AI do. It makes the machine confront the calendar: which facts in its answer are the kind that expire, and how you can confirm the current version.

Weak prompt: "What does it cost to renew a passport?"

Same prompt, with the word: "Refresh this for me: passport fees and rules change, so tell me which figures you're giving are most likely out of date, when they were last set, and exactly where to confirm today's number."

Why it worked. The machine's memory has a printing date. Refreshing turned a possibly stale answer into a short verification errand with the destination named.

Another example.

Weak prompt: "What are the tax rules for selling a house?"

Same prompt, with the word: "Explain the rules, but refresh as you go: flag each rule that has changed in the last five years or is scheduled to change, so I don't plan around a repealed fact."

Why it worked. Old rules don't announce themselves as old. Building the refresh into the answer kept last decade's law from quietly chairing the meeting.

One more.

Weak prompt: "What can I deduct for my home business?"

Same prompt, with the word: "List the deductions, but refresh each one: mark it current, changed recently, or commonly misunderstood, and name where a person confirms it free."

Why it worked. The "commonly misunderstood" marks were the treasure — two deductions he'd taken for years weren't, in his situation, deductions at all.

Watch for. Refreshing tells you which facts expire; it cannot fetch today's version from wherever the machine's knowledge stops. A refreshed answer is a list of errands, not their completion. The fee, the rule, the deadline — one of them needs a visit to the source.

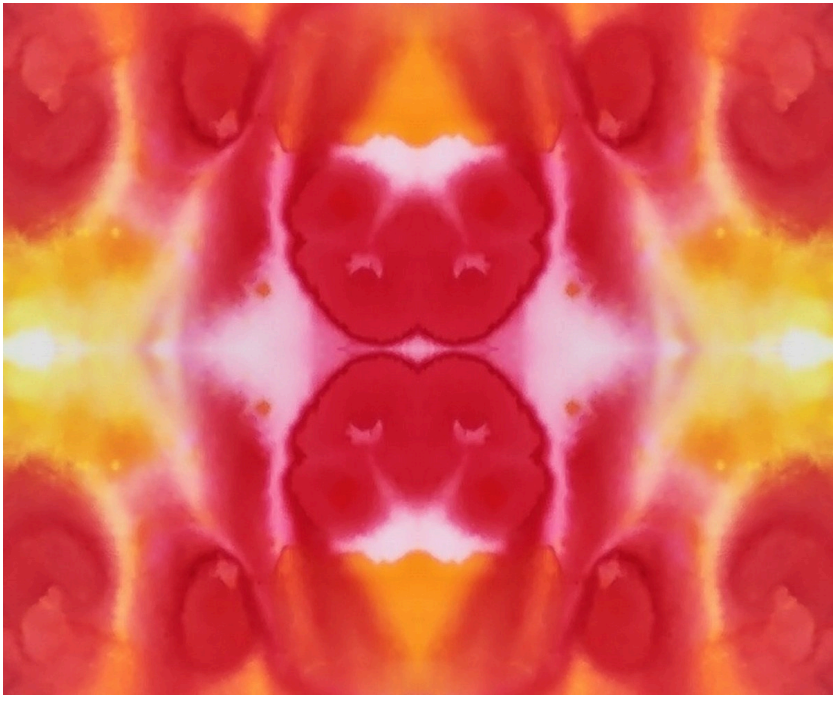
Placed well: "Refresh this fee schedule: which figures are likeliest stale, when were they set, and where do I confirm today's?"

Benchmark

What it makes the AI do. It measures your number against a known standard — typical, average, professional — so a figure floating alone becomes a figure with company.

Weak prompt: "I type about 45 words per minute. Is that good?"

Same prompt, with the word: "Benchmark my 45 words per minute: what's typical for casual users, for office work, and for professional typists — and what job requirements does 45 actually meet?"
Why it worked. "Good" needed a ruler. Against three benchmarks, 45 stopped being a vanity number and became a placement: fine for the office, short of the transcription desk.
Another example.
Weak prompt: "They offered me $60,000 for the job."
Same prompt, with the word: "Benchmark this $60,000 offer for a logistics coordinator with five years' experience in Ohio: where does it sit in the typical range, and what total-compensation items should I compare before judging it?"
Why it worked. A salary alone is unjudgeable. Positioned in its range — and widened to benefits and schedule — the offer became negotiable instead of merely acceptable or insulting.
One more.
Weak prompt: "My landlord wants $1,650 for this apartment."
Same prompt, with the word: "Benchmark the $1,650 against similar two-bedrooms within a mile: where does it sit in the range, and what features would honestly justify the top of that range?"
Why it worked. The apartment sat at the range's ceiling without the ceiling's features — a sentence worth real money at renewal time.
Watch for. Benchmarks conceal the comparison group: your typing speed looks professional beside casual users and hopeless beside court reporters, and both statements are true. Before accepting a benchmark's verdict, ask who is in the "typical" — the answer is a country, a decade, or a guess.
Placed well: "Benchmark this salary for Ohio: where it sits in the range, and what else counts as pay."



Part Six — Create

Make something new. These words move the machine past information and into making — drafts, stories, versions, and voices.

Everything before this part was about getting the truth out of the machine in a usable shape. This part is about getting something out of it that did not exist before: a draft, a name, a story, a script for the hard phone call. Creation is where most people underrate the machine most — they use it as a library when it is also, and often better, a workshop.

The catch is that machines create the way water flows: along the path of least resistance, toward the average. Every slogan sounds like every slogan; every story opens the way stories open. Creation words are the dams and channels: invent, don't recall; vary on purpose; write it in a voice; combine the two things nobody has combined. And above all, iterate — because almost nothing worth making arrives finished, and the loop of version, complaint, and revision is where your taste does its work. The machine supplies the clay at industrial speed. These words keep your fingerprints in it.

Draft

What it makes the AI do. It asks for a working first version with the clear understanding that it will be revised — which frees the machine from trying to be final and you from trying to start.

Weak prompt: "I need to write an email to my landlord about the broken heater."

Same prompt, with the word: "Draft an email to my landlord about the heater that's been broken for five days: polite, firm, mentions the two times I already called, and asks for a repair date in writing."

Why it worked. A blank page is the hardest part of any writing. A draft — even a flawed one — gives you something to push against, and pushing is faster than starting.

Another example.

Weak prompt: "I need to thank my neighbor for helping after the storm."

Same prompt, with the word: "Draft a short thank-you note to the neighbor who cleared my driveway after the storm: warm, specific about what he did, and ending with a real offer to return the favor."

Why it worked. The draft carried the feeling across the hard part. All that was left for the human was the signature — and a stamp.

One more.

Weak prompt: "I need to apologize to a client for the delay."
Same prompt, with the word: "Draft the apology: own the delay in the first sentence, give the new date, name what changed so it won't repeat — no groveling, no blaming the supplier by name."
Why it worked. Clients forgive delays with explanations attached. The draft had the explanation in sentence three, where it belongs.
Watch for. A draft's danger is its competence: it reads finished, so it gets sent — heater email dispatched with the placeholder date still sitting in sentence two. Drafts are scaffolding, and scaffolding gets checked before anyone moves in. Read it once, aloud, as its recipient.
Placed well: "Draft the landlord email: polite, firm, both earlier calls mentioned, and a repair date asked for in writing."

Invent

What it makes the AI do. It licenses originality: not the known options, but new ones — names, ideas, combinations nobody has listed yet.
Weak prompt: "What should I name my food truck?"
Same prompt, with the word: "Invent fifteen names for a food truck selling grilled cheese in the Midwest — some funny, some warm, some clever about cheese — none of them a name already famous."
Why it worked. "What should I name it?" earns the familiar five. "Invent" plus a count and flavor directions earns a page worth choosing from — choosing being the job you actually wanted.
Another example.
Weak prompt: "What should I name my dog?"
Same prompt, with the word: "Invent twenty names for a stubborn, dignified old rescue beagle — some noble, some funny, none of them from the ten most common dog names."
Why it worked. Twenty fresh candidates turned naming from settling into choosing. Abundance is a different errand than scarcity, and a much pleasanter one.
One more.
Weak prompt: "What should I call my Wi-Fi network?"
Same prompt, with the word: "Invent fifteen Wi-Fi names for a house full of teachers: some bookish, some funny, all safe to read aloud to a grandmother."

Why it worked. The constraint about the grandmother did more for the list's quality than the request did.

Watch for. Invented names arrive suspiciously available: fifteen fresh truck names, three of them already painted on trucks two states away — the machine doesn't check, and fame is a local thing. Before falling for one, search it. Originality is a claim; verify it like any other.

Placed well: "Invent fifteen names for a grilled-cheese truck — funny, warm, or clever — none of them already famous."

Narrate

What it makes the AI do. It converts facts and timelines into a story with a thread through it — the form human memory is actually built to hold.

Weak prompt: "Tell me what happened in the 2008 financial crisis."

Same prompt, with the word: "Narrate the 2008 financial crisis as a story following one ordinary family's house — bought in 2005, refinanced in 2006, in trouble by 2008 — explaining each big event as it reaches them."

Why it worked. Event lists slide off the mind. A story with people in it stays — and the big forces become understandable when you watch them land on one kitchen table.

Another example.

Weak prompt: "Tell me about the factory closing in my town."

Same prompt, with the word: "Narrate the closing of our town's mill through one worker's last year — the rumor, the announcement, the final shift — so my grandchildren understand what it actually cost."

Why it worked. Dates become history when someone inside the story has a name. The economics arrived attached to a life, and stayed.

One more.

Weak prompt: "Tell me how email was invented."

Same prompt, with the word: "Narrate the invention of email through one engineer who just wanted the machines in his building to talk — and let the technology arrive as it reaches him."

Why it worked. Protocols are forgettable; a man leaving notes for colleagues on a shared computer is not.

Watch for. Narration picks a hero, and the picking is an editorial act: follow the mill worker and the closing is a tragedy; follow the mayor

and it's a transition. Neither is false, both are partial. Ask, once, whose kitchen table the story never visits.

Placed well: "Narrate the crisis through one family's house — bought in 2005, refinanced in 2006, in trouble by 2008."

Translate

What it makes the AI do. Beyond languages: it converts anything written in one dialect of difficulty — legal, medical, technical, teenage — into the dialect you actually speak.

Weak prompt: "What does this insurance letter mean?"

Same prompt, with the word: "Translate this insurance letter into plain everyday language, sentence by sentence, and after each sentence tell me whether it means I owe something, I must do something, or nothing at all."

Why it worked. Jargon is often just plain news wearing a uniform. Translation strips the uniform — and the "so what is it to me" line after each sentence is where understanding turns into action.

Another example.

Weak prompt: "What does this after-visit summary from my doctor mean?"

Same prompt, with the word: "Translate this doctor's note into plain language, and after each item tell me plainly whether it is something to do, something to watch, or just information."

Why it worked. Plain words plus a "so what" turned a frightening page into a short to-do list — and the watch items stopped disguising themselves as emergencies.

One more.

Weak prompt: "What does this text from my teenager mean?"

Same prompt, with the word: "Translate this text from my daughter into plain parent language, phrase by phrase, and after each tell me the temperature: fine, joking, annoyed, or actually upset."

Why it worked. Slang was the easy half. Temperature was the translation that mattered — and it read "joking," thank heaven.

Watch for. Translation loses things on purpose — the insurance letter's menace lives partly in its formality, and a plain-language version can make a serious deadline sound like a friendly note. Keep the original

beside the translation, and let the translation carry the original's temperature markings.
Placed well: "Translate this letter sentence by sentence, marking each: I owe, I must act, or nothing is required."

Rewrite

What it makes the AI do. It keeps the content and changes the delivery — tone, length, audience, temperature — which is most of what editing actually is.
Weak prompt: "Make this email better." *(pasted email)*
Same prompt, with the word: "Rewrite this email to my coworker so it says the same things but sounds calm instead of irritated. Keep the deadline and the request exactly as they are."
Why it worked. "Better" is a lottery. Naming the one quality to change — temperature — and the two things to preserve turns the machine into a careful editor instead of a random one.
Another example.
Weak prompt: "Fix this text to my sister."
Same prompt, with the word: "Rewrite this text to my sister so it sounds warm instead of annoyed — but keep the boundary exactly where it is: I am not hosting Thanksgiving this year."
Why it worked. Naming what must survive — the boundary — kept the warmth from becoming surrender. Both things arrived in the same short message.
One more.
Weak prompt: "Make this note to the babysitter better."
Same prompt, with the word: "Rewrite the babysitter note so it sounds relaxed instead of anxious, but keep every time, number, and rule exactly as written."
Why it worked. The sitter got the same instructions at a lower heart rate — and so, reading it back, did the parents.
Watch for. Rewrite long enough and the content quietly migrates: the boundary you said must survive comes back softer, the deadline "exactly as written" returns as "soon." After any rewrite, diff it against the original yourself — what changed besides the thing you asked to change?
Placed well: "Rewrite this text warm instead of annoyed, but keep the boundary exactly: I am not hosting Thanksgiving."

Vary

What it makes the AI do. It asks for the same thing in deliberately different directions — not ten near-copies, but a spread of real choices.

Weak prompt: "Give me slogans for my bakery."

Same prompt, with the word: "Give me slogans for my bakery, but vary them on purpose: three that are funny, three that are nostalgic, three that are short enough for a sign, and three that mention bread specifically."

Why it worked. Machines repeat their first idea with synonyms. Varying by category forces genuinely different attempts — and a real spread is what makes the good one visible.

Another example.

Weak prompt: "Write a birthday message for my mom."

Same prompt, with the word: "Write birthday messages for my mom, but vary them on purpose: one funny, one grateful, one short enough for a card, and one that mentions her garden."

Why it worked. Four genuinely different doors into the same sentiment. One of them turned out to be exactly the one he had been trying to write for years.

One more.

Weak prompt: "Give me a caption for this photo of the lake."

Same prompt, with the word: "Give me captions, varied on purpose: one poetic, one funny, one plain factual, and one that's just three words."

Why it worked. The three-word one won. It would never have appeared in a run of similar tries, because machines, unasked, do not dare to be brief.

Watch for. Variety by category can still be variety in costume only: the funny slogan, the nostalgic slogan, and the short slogan all selling the same single idea the machine had first. If the spread feels familiar, name the idea you keep seeing — and ban it.

Placed well: "Vary these captions on purpose: one poetic, one funny, one factual, and one that's just three words."

Iterate

What it makes the AI do. It makes creation a loop instead of a lottery ticket: version, reaction, version — each round building on your specific complaint about the last.

Weak prompt: "That logo description isn't right. Try again."

Same prompt, with the word: "Let's iterate on this logo description. Keep the mountain, change only the colors to feel more like morning than night, and make the lettering bolder. Everything else stays."
Why it worked. "Try again" throws away the 80% that was working. Iterating names what to keep and what to change, so each round compounds instead of restarting.
Another example.
Weak prompt: "This soup is bland. Fix it."
Same prompt, with the word: "Let's iterate on the soup. Round one: change nothing except seasoning — give me three seasoning directions. I'll pick one, and only then do we talk about texture."
Why it worked. Changing one variable at a time meant the final recipe could be repeated — and remembered. Random improvements can't be cooked twice.
One more.
Weak prompt: "My presentation isn't landing. Redo it."
Same prompt, with the word: "Let's iterate instead of redoing: keep the data slides exactly, change only the opening — give me three openings. Round two, we fix the ending with whichever opening wins."
Why it worked. The data was never the problem. Protecting it meant the energy went to the two minutes that were.
Watch for. Iteration drifts. Each round preserves "everything else," but everything else shifts a degree per round — the mountain migrates, the morning darkens — until round six is a stranger to round one. Keep the original beside the latest, and compare them directly every few rounds.
Placed well: "Iterate on this logo: keep the mountain, change only the colors toward morning, nothing else moves."

Formalize

What it makes the AI do. It raises the register: rough notes, texts, and spoken thoughts become proper documents — letters, proposals, policies — dressed for the room they're entering.
Weak prompt: "Turn my notes into something I can send to the city council."
Same prompt, with the word: "Formalize these rough notes into a one-page letter to the city council about the potholes on Elm Street: respectful, specific about locations and dates, and ending with one clear request."

Why it worked. Ideas get ignored when they arrive dressed casually for a formal room. Formalizing keeps your substance and changes only the suit — which is often the whole difference.

Another example.

Weak prompt: "Turn these texts into a complaint to the airline."

Same prompt, with the word: "Formalize my notes into a complaint letter about the canceled flight: dates, flight numbers, what the agent promised me, and the refund I am requesting — firm, factual, and with no insults anywhere."

Why it worked. A formal letter enters a system that is built to answer formal letters. The same facts in a furious text enter a bin.

One more.

Weak prompt: "Make my complaint about the hotel sound professional."

Same prompt, with the word: "Formalize this complaint: the dates, the room, the noise complaint I made at 1 a.m. and what was promised, and the one night's refund I'm requesting — courteous throughout."

Why it worked. Hotels refund documented guests. The formal version was a document; the refund arrived with the reply.

Watch for. Formalizing can embalm a complaint: dates perfect, register impeccable, and the one human sentence that would have moved a reader filed off as unprofessional. Formal is a suit, not an embalming — tell the machine which sentence keeps its heartbeat before it dresses the rest.

Placed well: "Formalize these notes into a one-page council letter: respectful, specific, ending with one clear request."

Brainstorm

What it makes the AI do. It suspends judgment on purpose: quantity first, criticism later, so unusual ideas survive long enough to be seen.

Weak prompt: "Give me some video ideas for my gardening channel."

Same prompt, with the word: "Brainstorm thirty video ideas for a small gardening channel — quantity first, no judging, wild ones welcome — then group them by type so I can pick."

Why it worked. Asking for "some ideas" gets the five obvious ones. Thirty forced the machine past its defaults, and the good idea was sitting at number twenty-two, where it usually lives.

Another example.

Weak prompt: "What should we do for our anniversary?"

Same prompt, with the word: "Brainstorm twenty-five anniversary ideas across five kinds — adventurous, quiet, silly, generous, and new-to-us — five of each, budget-blind for now."
Why it worked. Categories kept the brainstorm from producing twenty-five dinners. The budget could edit later; it could not have invented the hot-air balloon.
One more.
Weak prompt: "Ideas for my YouTube channel about woodworking?"
Same prompt, with the word: "Brainstorm forty video ideas, quantity first — builds, repairs, mistakes, history, viewer projects — and don't stop at the ones a channel like mine usually makes."
Why it worked. Idea thirty-one — rebuilding viewers' worst mistakes — became the series. It was never going to appear in a list of five.
Watch for. Quantity has a hangover: forty ideas, and the choosing — which was the job — now takes longer than the brainstorm did. Cap the flood by deciding the sieve in advance: three finalists, chosen against two named tests, before the list is allowed to cool.
Placed well: "Brainstorm thirty video ideas, quantity first, no judging — then group them so I can choose three."

Voice

What it makes the AI do. It writes in a named voice — a kind of person, a tone, a temperament — instead of the machine's default voice, which sounds like a brochure.
Weak prompt: "Write a note to customers about our holiday hours."
Same prompt, with the word: "Write the holiday-hours note in the voice of a friendly small-town shopkeeper: short sentences, warm, one joke, and the hours impossible to miss."
Why it worked. The default machine voice is competent and forgettable. A named voice gave the note a person to be — and customers write back to people.
Another example.
Weak prompt: "Write the instructions for watering my plants while I'm away."
Same prompt, with the word: "Write the plant instructions in the voice of a kind neighbor who's done this before: plain, reassuring, and tell her which plant is the drama queen."

Why it worked. A housesitter follows instructions she enjoys reading. The voice carried the care instructions past the fridge door and into actual watering.

One more.

Weak prompt: "Write a welcome message for my rental cabin."

Same prompt, with the word: "Write it in the voice of a mountain innkeeper who's seen every kind of guest: warm, unhurried, a little dry humor, and the Wi-Fi password woven in gracefully."

Why it worked. Guests began mentioning the note in their reviews. A voice, it turns out, is hospitality too.

Watch for. A named voice can tip into costume: the small-town shopkeeper acquires a corn-pone accent nobody requested, the innkeeper starts performing folksiness at the Wi-Fi password. Voice should season, not swamp — if you notice the accent more than the hours, dial it back one notch.

Placed well: "Write the holiday-hours note in a friendly shopkeeper's voice: short, warm, one joke, hours unmissable."

Combine

What it makes the AI do. It forces two unlike things together — a format and a subject, two hobbies, two tools — and makes the machine find the working hybrid instead of a safe single answer.

Weak prompt: "I like chess and I like cooking shows."

Same prompt, with the word: "Combine them: invent a show where two chefs compete like chess players — timed moves, an opponent's dish to respond to, a clock. Give me the format, the rules, and a sample episode."

Why it worked. Neither interest alone produced anything. The collision did — because formats are portable, and a chess clock turns out to be drama you can put in a kitchen.

Another example.

Weak prompt: "Help me exercise and learn history at the same time."

Same prompt, with the word: "Combine my daily walks with learning history: design a walking route plan where each week's walks come with one era to listen through, in order, with the total distance matched to the episodes."

Why it worked. Two goals were competing for the same hour. Combined, they stopped competing — the hour now pays both debts at once.
One more.
Weak prompt: "I want to practice guitar and spend time with my daughter."
Same prompt, with the word: "Combine them: design a Saturday ritual where she picks the songs and I learn them to play together — her requests, my practice, one hour, with a simple plan for the first month."
Why it worked. Practice acquired an audience of one who requests encores — the most durable motivation a guitar has ever had.
Watch for. Combinations fail charmingly: the chess-cooking show is a delight to describe and a misery to film, because collisions produce novelty faster than practicality. Before adopting a hybrid, ask the machine the unglamorous question: what does Tuesday look like, doing this, in week three?
Placed well: "Combine my walks with history: one era per week, episodes matched to my actual route distances."

Script

What it makes the AI do. It writes the exact words for a moment you will actually face — the call, the ask, the apology — so you rehearse lines instead of improvising under pressure.
Weak prompt: "I have to call the insurance company about a denied claim."
Same prompt, with the word: "Script the call for me: my opening line, the calm version of my argument, what to say if they say no, and a polite closing that keeps the door open. Keep every line short enough to say out loud."
Why it worked. Under pressure, nobody thinks of the good sentence. A script moved the thinking to a calm afternoon, where it belonged.
Another example.
Weak prompt: "How do I tell my boss I need Fridays off?"
Same prompt, with the word: "Script it: the exact first two sentences for asking my boss about Fridays off, one sentence for the reason, and a ready answer for 'what about deadlines?'"

Why it worked. The whole conversation frightened him; the first two sentences were the actual mountain. Scripted, they turned out to be easy to say.

One more.

Weak prompt: "I need to apologize to my brother."

Same prompt, with the word: "Script the apology: first sentence that owns it without excuses, one sentence about what I'll do differently, and a closing that doesn't demand he make me feel better."

Why it worked. The script protected the apology from its two natural enemies — rambling and self-defense.

Watch for. A script is a rehearsal, not a possession: memorize it too rigidly and the unexpected reply strands you, line perfect and nowhere to go. Script the opening and the closing word-perfect; for the middle, script your intentions — the points you won't leave without making.

Placed well: "Script my opening line and my answer to 'what about deadlines?' — leave the middle as intentions."

Recipes



Single words are tools. Recipes are tools in sequence — one word setting up the next. Here are six worth keeping on an index card.

Recipe 1 — Specify → Quantify → Rank

For any purchase bigger than a toaster. A man shopping for a used truck first **specifies**: "trucks under $18,000, under 120,000 miles, able to tow 5,000 pounds." Then he **quantifies**: "estimate the next two years of likely repairs for each type." Then he **ranks**: "order them by total cost of owning, not sticker price." Three prompts, and a fog of listings became a decision with its homework already done — because each word narrowed, priced, and ordered what the previous word produced.

Recipe 2 — Steelman → Disagree → Invert

For any decision you can't undo cheaply. Before accepting a new job across the country, a woman **steelmans** the move: "make the strongest honest case for taking it." Then she flips: "**disagree** — make the strongest case that it's a mistake." Then she **inverts**: "list what would have to go wrong for me to regret this in a year." She may still move — but she'll move having met the best arguments on both sides and named her real risks, instead of the ones a brochure mentions.

Recipe 3 — Decompose → Sequence → Budget

For the project that lives in your head as one big dread. "Renovate the bathroom" gets **decomposed** into its separate jobs, each **sequenced** — plumbing before tile, tile before fixtures, and what's allowed to overlap — and finally **budgeted**: a fixed total, allocated job by job, with the splurge named and the savings named. The dread was never the work; it was the shapelessness. Three words gave the project a skeleton, a calendar, and a price.

Recipe 4 — Draft → Critique → Iterate

For writing that matters. A best man **drafts** his toast in one prompt, then asks for a **critique** "the way a kind but honest friend would give it — where does it drag, where's the joke that only works on paper?" Then he **iterates**, round by round, changing one thing at a time: "keep the fishing story, cut the college story, make the ending land on the bride." The toast that results isn't the machine's or his alone — it's what survived several honest rounds, which is how most good writing has always been made.

Recipe 5 — Assume → Bracket → Verify

For numbers someone hands you with confidence. A contractor quotes "$9,000" for a roof. You **assume** his scope out loud: "assuming that's tear-off and replacement, mid-grade shingles." You **bracket** it against the market: "what's the low-to-high range for this roof in my state?" Then you **verify**: "go through his line items and flag anything unusual." You haven't accused anyone of anything — you've simply done, in three prompts, what a careful person does before signing anything with a comma in the price.

Recipe 6 — Summarize → Translate → Exemplify

For learning anything that arrives wrapped in jargon. A new diagnosis, a lease, a letter from a law firm: first **summarize** it — the shape of the thing in five sentences. Then **translate** the hard parts into plain language. Then **exemplify**: "show me one concrete case of what this means in practice." Shape, meaning, instance — by the third word, the document that frightened you has become a subject you can discuss, ask about, and decide about.

Recipe 7 — Localize → Filter → Rank

For finding a human being to hire. A homeowner needs a plumber and dreads the directory. She **localizes** first: "plumbers who actually serve my town, not the metro area." Then she **filters** the list hard: "licensed, answers their own phone, will give a written estimate — drop everyone else silently." Finally she **ranks** what survives: "by how fast they can come this week, then by review patterns about cleanup, not just stars." Three prompts turned a page of ads into two phone calls — and the criteria, not the advertising, did the choosing.

Recipe 8 — Role → Simulate → Critique

For the conversation you are dreading. A supervisor must tell a good employee his work has slipped. First, **role**: "you are a fair, experienced HR coach." Then **simulate**: "play my employee — proud, surprised, a little hurt — while I practice the opening." Afterward, **critique**: "review what I said the way that coach would: where did I get vague, where did I accidentally promise something?" The meeting, when it came, was the second time he had had it — and the first version, the bad one, had cost nothing.

Recipe 9 — Trace → Cite → Flag

For the claim everybody is sharing. A startling health claim is moving through a family group chat. First **trace** it: "where did this claim first come from — a study, a book, a television segment?" Then **cite**: "for each part of it, name the kind of source behind it." Then **flag**: "mark anything that got exaggerated as it traveled." The message he sent back to the chat was short and kind, because the homework had been done quietly and completely — the claim's birthplace was a magazine headline, its study was about rats, and two of its numbers had doubled in the retelling.

Recipe 10 — Brainstorm → Vary → Anchor

For making something with taste. A baker naming a new line of breads **brainstorms** thirty names without judging a single one. Then she **varies** the survivors on purpose: "give me each finalist in three directions — homier, funnier, more elegant." Finally she **anchors**: "judge them all against the name of the loaf that already sells best — match its warmth, beat its memorability." The machinery produced options at industrial speed; the anchor made sure her own hard-won taste picked the winner.

Recipe 11 — Stage → Timebox → Anticipate

For the project measured in months. Restoring an old truck gets **staged**: engine, body, interior, each with a definition of done. Each stage gets **timeboxed** — "six weekends, no more, and tell me what 'done' sacrifices if the box holds." Then the whole plan is **anticipated**: "list what will try to stop this project — parts delays, money months, lost enthusiasm — and one defense for each." Big projects rarely die from the work. They die from the unlisted surprise in month three, which this recipe drags into the light in week one.

Recipe 12 — Pinpoint → Cross-Check → Reconcile

For two numbers that refuse to agree. A small shop's register total and bank deposit disagree by $212. She **pinpoints** first: "which day, which shift, which tender type does the gap live in?" Then she **cross-checks** that day's total a second way, from the receipts instead of the register. The routes disagree in a new place — which is progress — so she **reconciles**: "walk the likely causes of a gap this size, in order." Cause number two was a double-entered refund. Three words turned a vague suspicion of her staff into a bookkeeping event with a name, which is what it had been all along.

Prompt Templates — Ready to Copy



Everything in this book, pre-assembled. Each template below is a complete prompt: copy it, fill in the [brackets], and send it as-is. The short note after each names the words doing the heavy lifting, so you can start taking the templates apart and rebuilding them for jobs nobody has templated yet.

1. Plan My Week

> "Plan my week of [DATES]. My fixed commitments are: [LIST THEM]. My goals are: [GOAL 1], [GOAL 2]. Budget my waking hours across work, family, errands, and rest — show the split. Sequence anything that depends on something else, and flag any day where the plan only works if nothing goes wrong."

Words at work: Budget allocates the fixed total of hours; Sequence respects the dependencies; Flag marks the fragile day before it breaks.

In use: "Plan my week of June 9–15. Fixed: work weekdays 9–5, dentist Tuesday at 2, kids' game Saturday at 10. Goals: finish the deck repair, call my mother twice. Budget my waking hours across work, family, errands, and rest — and flag any day that only works if nothing goes wrong." The flagged day came back Wednesday, which was carrying a delivery, a deadline, and the dentist's revenge — and it got fixed on paper, in advance, for free.

2. Understand a Bill

> "Here is a bill from [COMPANY]: [PASTE IT]. Translate it into plain language, charge by charge. Itemize every fee, flag any charge that is unusual, new, or higher than typical, and end with the exact sentence I should use if I call to dispute the worst one."

Words at work: Translate removes the uniform from the jargon; Itemize makes the total countable; Flag aims your one phone call at the charge that deserves it.

In use: a water bill arrives at twice its usual size. The translation shows the usage charge is normal but a "service availability fee" appeared in March and doubled in May. The dispute sentence the template produced — calm, dated, specific — was read aloud word for word on the call, and the fee was credited before the hold music ended.

3. Prepare for a Meeting

> "I have a meeting with [WHO] about [TOPIC] on [DAY]. Anticipate the five hardest questions or objections they could raise. For each, draft a short, honest answer I could actually say out loud. End with the one question I should ask them."

Words at work: Anticipate rehearses tomorrow's pressure today; Draft turns worries into sentences; the closing question returns some of the steering wheel to you.

In use: a performance review with a manager known for curveballs. The anticipated questions included the one about the missed March deadline — the exact one asked, worded almost identically. The answer had been rehearsed at the kitchen table, and it sounded like honesty instead of panic because it was, by then, simply practiced.

4. Learn a Topic from Zero

> "Teach me [TOPIC] from zero. Layer it: first a one-sentence version, then a one-paragraph version, then the full picture. After the full picture, quiz me with three questions, and exemplify the hardest idea with one everyday example."

Words at work: Layer stacks the depth so you can stop when satisfied; Exemplify anchors the abstraction; the quiz converts reading into remembering.

In use: a grandfather decides to finally understand the family trust his lawyer drafted. Layer one: "it's instructions for who handles things if you can't." Layer three, plus the quiz, exposed that he had misunderstood one term completely — cheaply, at his own table, instead of expensively, at a lawyer's hourly rate.

5. Compare Two Purchases

> "Compare [ITEM A] and [ITEM B] for someone who [YOUR SITUATION]. Tabulate them: price, running cost per year, expected lifespan, warranty, and the known weak point of each. Then rank them by total cost of owning, and tell me what kind of buyer should choose the loser anyway."

Words at work: Tabulate lines the facts up where mismatches show; Rank forbids the tie; the final sentence honors the truth that the best buy depends on the buyer.

In use: two washing machines, $200 apart. The table showed the cheaper one cost more per year to run and carried the shorter warranty; the ranking by cost-of-owning flipped the sticker order. The losing machine's buyer, the answer noted, should be someone moving within two years — which is a renter, not this family.

6. Fix a Draft

> "Here is my draft: [PASTE IT]. Critique it honestly — where it drags, where it confuses, what it should cut. Then let's iterate: round one, change only [THE ONE THING]. Keep everything I haven't named exactly as it is."

Words at work: Critique gets the truth before the rewrite; Iterate makes improvement compound; the last sentence is Constrain standing guard over your own words.

In use: a complaint letter to a contractor, written angry and rewritten necessary. The critique cut four sentences that felt wonderful and proved nothing. Round one changed only the opening — from accusation to chronology — and the letter that was sent got a callback, which the angry version had been risking for the pleasure of being angry.

7. Decide Something

> "I'm deciding whether to [THE DECISION]. Steelman the case for doing it, then steelman the case against, both at full strength. Then invert it: list what would have to go wrong for me to regret this in a year. Grade how likely each regret is."

Words at work: Steelman on both sides forbids the comfortable caricature; Invert finds the failure paths; Grade keeps the fears sorted by likelihood instead of volume.

In use: a teacher weighs leaving for a job in another state. The steelmanned cases ran three paragraphs each and neither pulled a punch. The inversion list named four regrets; the grades put "missing my students" at likely, "the school collapses without me" at unlikely — and the decision, unchanged, became considerably easier to sleep inside.

8. Explain It to a Child

> "Explain [TOPIC] to a [AGE]-year-old. Simplify until she could explain it back to me correctly. Use one example from her world — [SCHOOL / PETS / A GAME SHE LOVES] — and no word you haven't defined in the same sentence. Keep it under [NUMBER] sentences."

Words at work: Simplify applies the honest comprehension test; Localize moves the example into her actual world; Limit keeps it shorter than her patience.

In use: "Explain thunder to a six-year-old who loves her dog." The answer used the dog startled by a dropped pan, defined "sound wave" as "the push the noise makes in the air," and stopped at five sentences. She explained it back perfectly, then went to comfort the dog, which had been the real assignment all along.

9. Summarize a Long Document

> "Here is a long document: [PASTE OR ATTACH IT]. Summarize it in five sentences. Then list separately: anything that costs me money, anything that restricts what I can do, every deadline, and anything you had to interpret because the wording was unclear — flag that last group as interpretation, not fact."

Words at work: Summarize gives the shape; the separate lists stop compression from burying the biting clauses; Flag separates the document's words from the machine's guesses.

In use: a 40-page condo association rulebook. The summary fit in five sentences; the separate lists held the real estate — a $150 monthly special assessment, no rentals under twelve months, a deadline for a balcony inspection — and one flagged interpretation about guest parking that a

phone call later confirmed the machine had read wrong. The flag, as usual, was the point.

10. Negotiate Politely

> "I need to negotiate [WHAT] with [WHO]. I want [YOUR GOAL]; I can honestly offer [WHAT YOU CAN GIVE]. Draft what I say: warm opening, the ask stated plainly, one reason that's true, and a gracious line for if they say no. Nothing pushy, nothing apologetic."

Words at work: Draft moves the hard sentences to a calm moment; Assume keeps the trade honest by stating both sides of it; tone words do the constraining that keeps it polite.

In use: a tenant asking a landlord to split the cost of a fence repair. The draft opened warm, asked plainly, offered the true reason — both properties gain — and closed with a gracious line for a no. The landlord said yes in two days, and later said the letter itself was half the reason.

11. Study with Quizzes

> "Teach me [TOPIC] like a patient tutor. Break it into chunks I can finish in [NUMBER] minutes each. After every chunk, stop and quiz me — one question at a time, wait for my answer, correct me kindly, and don't move on until I've got it. If I stall, exemplify the idea instead of repeating it."

Words at work: Chunk sizes the sittings; the quiz loop turns teaching into checking; Exemplify is the rescue move when explanation alone stops working.

In use: a father learning spreadsheet formulas from zero, in twenty-minute chunks after dinner. The quizzes caught a misunderstanding about cell references at chunk two — exactly where a video course would have kept sailing. Two weeks later he built the family budget sheet himself and explained it at the table, which is the only graduation this book recognizes.

12. Debug a Household Problem

> "Something's wrong at home: [DESCRIBE THE PROBLEM AND WHEN IT STARTED]. Limit your answer to the three most likely causes, in order. For each cause, give me one safe check I can do

> myself today, and flag clearly anything that means ‘stop and call a professional instead.’”

Words at work: Limit keeps the suspect list actionable; Cross-Check lives in the tests, one per cause; Flag draws the bright line between a homeowner’s afternoon and a professional’s job.

In use: a toilet that runs every twenty minutes at 3 a.m. The three causes came ordered — flapper, fill valve, chain — each with a safe check, and the flag sat beside the third: a cracked tank means stop. It was the flapper. Total cost: one part, one Saturday hour, and a full night’s sleep, restored by a numbered list.

The Practice Room



Words are learned the way scales are learned on a piano: a few minutes of deliberate reps beats a month of good intentions. Work these exercises with your own machine, on your own real life. Each takes under ten minutes. Do one a day and the vocabulary stops being a book you own and becomes the way you type.

From Part One — Sharpen

1. **Specify.** Take the last vague question you asked a machine and rewrite it with a count, a place, and a limit — five things, your town, under $50. Compare the two answers side by side and keep the habit that produced the better one.
2. **Compare.** Put two things you genuinely might buy into one prompt, judged on identical criteria, for your actual situation. Notice which criterion you almost forgot to name — that forgetting is what the word is for.
3. **Rank.** Hand the machine a list you care about — expenses, chores, worries — and demand an order, with the criterion stated out loud. Argue with the order. The argument is the exercise.
4. **Exemplify.** Pick any piece of advice you were given this week and demand one concrete instance of it, in exact words or exact steps. If it cannot survive an example, retire the advice.

From Part Two — Constrain

5. **Constrain.** Take any answer the machine gives you today and re-ask the question with one hard edge added: 100 words, two minutes, $20, by Friday. Watch the answer get better by getting smaller.
6. **Exclude.** Ask for recommendations about anything, then strike the five most obvious answers by name and ask again. The second list is where your life actually lives.

7. **Bracket.** Choose one upcoming expense and refuse the single number: demand a low case, a high case, and the levers between them. Plan against the high case; celebrate against the low.
8. **Assume.** Take a question you've been circling and state your real circumstances out loud as assumptions — savings, time, health, help. Ask what would have to be true for the brave option to be reasonable.

From Part Three — Structure

9. **Tabulate.** Find a decision currently living in paragraphs — a plan, a policy, three quotes — and have it rebuilt as a table. Count how many facts were missing a counterpart until the columns exposed them.
10. **Sequence.** Take a process you are about to do for the first time and demand the order, the dependencies, and the waits. Mark the step that cannot be undone. That mark is the exercise's diploma.
11. **Decompose.** Name your most avoided project in one sentence, then decompose it into jobs small enough for a calendar square. Put exactly one square on this week's calendar before you close the app.
12. **Layer.** Choose something you half-understand — a bill, a diagnosis, a technology — and take it in three layers. Stop at the layer where your actual decisions live, and notice how rarely that is the deepest one.

From Part Four — Challenge

13. **Steelman.** Take a position you hold warmly and ask for its opposition at full strength, as its smartest defender would state it. Your only job is to read it twice without rebutting. Tomorrow you may rebut.
14. **Invert.** For your current goal, list the guaranteed route to failure. Then audit your last seven days against the list, honestly, and forgive quickly — the audit is the point, not the guilt.
15. **Disagree.** Present a plan you are excited about and instruct the machine to disagree, with its strongest reasons priced or timed. Decide afterward which objections survive your own scrutiny — those are now part of the plan.
16. **Simulate.** Rehearse one real upcoming conversation — a raise, an apology, a boundary — with the machine playing the other side honestly. Note the sentence where it went wrong in rehearsal, and write its replacement.

From Part Five — Verify

17. **Verify.** Take the last factual answer you relied on and send it back through claim by claim: established, mixed, or marketing. Do it once and you will never read a confident answer the same way again.
18. **Cross-Check.** Take one number that matters this month — a bill total, a measurement, a dosage of anything — and reach it by a second, different route. If they disagree, you have just been paid for the exercise.
19. **Flag.** Ask any question where being wrong would cost you, and require the machine to flag its uncertainties inline. Read only the flags first. That order — flags first — is the whole discipline in miniature.
20. **Trace.** Pick a saying everyone repeats — about health, money, or history — and trace it to its birth. Bring what you find to your next family gathering and enjoy the brief, glorious silence.

From Part Six — Create

21. **Iterate.** Make one thing in public rounds: a paragraph, a plan, a recipe. Change one named thing per round, protect everything unnamed, and stop at round four whether or not it is perfect. Notice how much better round four is than restart number four ever was.
22. **Invent.** Ask for twenty of something you'd normally ask for three of — names, ideas, openings, gifts. Set a rule that the winner may only be chosen after the list is complete. The back half of the list is where the exercise lives.
23. **Voice.** Take a dull necessary note and rewrite it in three named voices: a kind neighbor, a ship's captain, a favorite teacher. Send the one that makes you smile; keep the exercise for everything official you will ever write.
24. **Script.** Choose the phone call you have been avoiding and script its first two sentences, its honest middle, and its gracious exit. Read it aloud once. Then make the call, ideally while the script is still warm.

Field Guide — If Your Problem Is…

The quick-picker. Find your trouble, take the word, go.

- **The answer is true but useless** → **Specify**, then **Target**. It doesn't know your errand or your reader yet.
- **The answer is enormous** → **Constrain** it, **Limit** it, or **Summarize** it. Size is a choice; make the machine make it.
- **The answer sounds like everyone else's** → **Exclude** the famous five, then **Vary** what's left. Originality lives past the defaults.
- **You're not sure it's true** → **Verify** it, **Cite** it, **Grade** it. Confidence is not evidence, and now the answer knows you know that.
- **Two numbers disagree** → **Pinpoint** the gap, **Cross-Check** by a second route, **Reconcile** what's left. A mismatch is a specific event, not a mood.
- **The decision scares you** → **Steelman** both sides, **Invert** it, **Worst-Case** it. Fear shrinks when it's itemized.
- **The project won't start** → **Decompose** it, **Stage** it, **Timebox** the first stage. Dread is usually just shapelessness wearing a deadline.
- **The document frightens you** → **Translate** it, **Flag** the uncertain parts, **Layer** the rest. Fright is often just jargon with good posture.
- **The machine keeps agreeing with you** → **Disagree**, **Rebut**, **Critique**. Agreement is the factory setting, not the truth setting.
- **You have to say it out loud** → **Script** it, **Simulate** it, then say it. The hard sentence belongs to a calm afternoon, not a hard moment.
- **You're learning something new** → **Layer** it, **Chunk** it, **Probe** yourself with its questions. Depth on demand beats depth all at once.
- **The plan is written but untested** → **Stress-Test** it, **Anticipate** its surprises, **Refresh** its oldest facts. Paper plans fail politely; tested ones fail here, cheaply.
- **You've asked three times and it's still not right** → stop restarting. **Iterate**: name what stays, name the one thing that changes, and go one more round.

Quick Reference

Word	What it does
Specify	Forces particulars instead of generalities
Quantify	Turns opinions into comparable numbers
Itemize	Breaks a lump into a countable list
Compare	Judges options on identical criteria
Rank	Orders options and exposes priorities
Prioritize	Sorts by what must happen first
Exemplify	Demands a concrete instance for every claim
Define	Pins slippery words down before building on them
Pinpoint	Locates the exact spot, moment, or cause
Probe	Makes it question you, one good question at a time
Anchor	Fixes one reference example and judges everything against it
Target	Aims the answer at one specific reader
Constrain	Sets hard edges the answer must solve inside
Exclude	Removes the obvious and already-tried
Assume	States a premise so reasoning can start somewhere real
Bracket	Gives a low case and high case instead of a false-precise number
Budget	Allocates a fixed total and defends the split
Limit	Caps the count, forcing the machine to choose
Simplify	Strips jargon while keeping the load-bearing ideas
Localize	Drags generic advice into your actual place
Timebox	Bounds a plan by minutes instead of wishes
Role	Gives the machine a part to play, with its reflexes included
Filter	Keeps only what meets your criteria and drops the rest
Scope	Draws the border of the job — what's in, what's out
Sequence	Orders steps with their dependencies respected
Outline	Gives the skeleton before the flesh
Tabulate	Aligns facts into scannable rows and columns
Summarize	Compresses while keeping the claims that matter
Compress	Squeezes to a hard size without losing the essentials

Word	What it does
Expand	Grows exactly one point, deeply
Decompose	Breaks a big fog into small, startable jobs
Recompose	Rebuilds the parts in a better arrangement
Stage	Divides a long effort into phases with gates between them
Chunk	Cuts big material into sittings you can finish
Layer	Stacks one explanation at several depths
Map	Shows how the pieces connect to each other
Steelman	States the other side at its strongest
Invert	Asks how to fail, in order to avoid it
Critique	Switches the machine from author to reviewer
Defend	Argues for a position under attack
Disagree	Gives the machine an instruction to push back
Counterfactual	Removes one cause at a time to find the real one
Anticipate	Raises tomorrow's hard questions today
Simulate	Rehearses the real situation instead of describing it
Stress-Test	Pushes a plan against hard conditions to find the crack
Falsify	Asks what evidence would prove the claim wrong
Worst-Case	Starts from the bad outcome and plans backward
Rebut	Makes it argue against its own previous answer
Verify	Re-checks its own claims and grades them
Audit	Examines the whole thing against a checklist
Cite	Attaches a source to each claim
Cross-Check	Reaches the same answer by a second route
Trace	Follows a claim back to where it started
Estimate	Shows assumptions and arithmetic behind a number
Flag	Marks its own uncertainties in the open
Reconcile	Finds where two disagreeing numbers parted
Sanity-Check	Tests whether a number is even in the plausible band
Grade	Makes it rate its own confidence in the open
Refresh	Confronts which facts expire and how to confirm today's
Benchmark	Measures your number against a known standard
Draft	Produces a working first version to push against
Invent	Licenses genuinely new options
Narrate	Turns facts into a story memory can hold

Word	What it does
Translate	Converts jargon into your dialect
Rewrite	Keeps the content, changes the delivery
Vary	Produces deliberately different directions, not copies
Iterate	Improves by named rounds instead of restarts
Formalize	Dresses rough notes for the room they're entering
Brainstorm	Produces quantity first and judgment later
Voice	Writes in a named voice instead of the default one
Combine	Merges two unlike things into a working new one
Script	Writes the exact words for a moment you'll face

A Word at the End



You now own seventy-two words and one honest recommendation: don't learn them. Not the way you'd learn a speech, anyway. Learn six. The six that match the way your own asks usually fail — the vague ones, the unbounded ones, the ones where you nodded at an answer you should have cross-examined. Use those six until they stop feeling like techniques and start feeling like typing. Then come back for six more, the way you'd return to a toolbox for a tool a new job suddenly needs.

That is how this book was built, after all: not all at once, and not from theory. One word at a time, purchased by one real day of work, by two partners who got tired of misunderstanding each other and did something unfashionable about it — they got precise. The machine on your side of the screen is waiting to be asked better. It is, in the experience recorded here, surprisingly willing to deserve it.

Ask well. It compounds.

The machine is not the hard part. The ask is the hard part — and now you have seventy-two ways to make it a good one.

About the Art

Every picture in this book grows from a single painting: an abstract work in ink on canvas by the artist, red and gold and weather. It appears exactly once in these pages, as a full plate before the story chapter, entirely itself — no crop, no treatment, no improvement.

Everything else you see — the cover, the part openers, the chapter pieces — is a variation on that one painting, made with AI tools working from the original under direction. Some stay close, nearly quotations. Some travel far, until the washes become weather over hills or a sky full of ordered cards. All began in the same place.

That seemed only honest. This is a book about directing a machine with words — precisely chosen, honestly bounded, checked afterward. The pictures were made the same way the text teaches: one strong source, carefully worded instructions, and a human deciding what is finished.

Art credit: the original painting and the art throughout this book are by Douglas John Brownfield; the variations were made with AI assistance, as described above.

About the Author

Douglas B. is an independent builder and artist based in Illinois. He builds self-custodial software tools — programs designed so the people who use them keep hold of their own work — and he writes about what it is actually like to work with AI day after day: where it shines, where it guesses, and where it has to be watched.

His abstract paintings, ink on canvas, take their cues from Hubble photographs of galaxies and nebulae, all wash and weather and deep space. One of them carries the art of this book.

Words to Prompt AI Well With grew out of a real working partnership between the author and an AI assistant, one project at a time. Every word in it was bought the way tools usually are: a job went badly without it, and the lesson stayed.