

The AI Bootcamp

23

How to Fix a Bad Answer

In this chapter you'll learn

The named failure shapes, their real causes, and the one question that says whether to rewrite or start again

The AI Bootcamp

THE PROBLEM

Most AI education is a pile of disconnected tips. A LinkedIn carousel here, a YouTube tutorial there, a \$2,000 cohort course that goes out of date before the refund window closes. None of it is built to be read in order, so nothing compounds. You finish twelve videos and can recite twelve facts, not perform one new skill.

THE FIX

The AI Bootcamp is built the other way. 90 skills, packed 2 to 3 per chapter, 36 chapters, 3 tracks. Every chapter states what it needs from you and what it hands to the chapter after it, so skill 47 only works because skill 12 already landed. Read it in order and by the end you are not collecting facts about AI, you are running work through it.

LETTER FROM THE EDITOR

We built this because we kept meeting smart people who felt behind on AI for no reason other than nobody had organized the material for them.

This book does not assume you are technical, and it does not water anything down either. Every chapter earns its place: a real skill, a real exercise, a one-page companion you can flip back to in month three when you have forgotten the details but still need the answer.

If a chapter loses you, or a chapter changes how you work, I want to hear about it either way.

Kris

Editor, AI Central

Write to me at editor@thecentral.ai

THE 3 TRACKS

AI 101

Get running

Pick a tool, set it up, write a clear instruction, use it safely at work.

AI 202

Prompt properly

The framework behind every good instruction, and catching AI when it is wrong.

AI 303

Build without watching it

Connect AI to your tools, and automations that run while you are not looking.

IN THIS LESSON YOU'LL LEARN

By the end of this chapter you can look at a bad answer and name what went wrong in a phrase, say whether the fault is in your ask or in the conversation, and run the one repair that fits instead of retyping the same request with more force.

You can now tell a bad answer from a good one, which means you meet more of them. What most people do next is retype the request louder, and that fails often enough to be worth explaining: different failures have different causes, and a repair aimed at the wrong cause costs you an afternoon. Naming the failure is most of fixing it.

CONTENTS	23.1	A bad answer is evidence, and it points at its own repair
	23.2	5 shapes that mean your request was the problem
	23.3	4 shapes that mean the conversation is the problem
	23.4	Some failures are ceilings, and no rewording moves them

TOPICS REQUIRED

Chapter 22. You need your quality bar, and one recent answer that failed it, still sitting open in your history. Bring the whole conversation it came from, not just the last reply.

AFTER THIS CHAPTER

Chapter 24 assembles the whole of AI 202 into a system you run without thinking about it. It assumes you can already name a failure, because a playbook with no repair section is a playbook for the days when everything works.

23.1 A bad answer is evidence, and it points at its own repair

Nothing you do to a request matters until you know what went wrong with the last one. Almost every bad answer belongs to 1 of 3 families, and each family has a different repair.




Fig 23.1 Three families, three repairs. The bar heights are visual rhythm, not a frequency count: in practice the first two are close, and the third is rarer than people assume once they stop misfiling refusals under it.

23.1

Ask failures. Repaired where you are standing, in the same conversation, by saying the thing you did not say.

Thread failures. Not repairable in place. The conversation is carrying a wrong turn and will keep carrying it.

Ceiling failures. Refusals and invented specifics. Neither rewriting nor restarting touches them.

Here is the problem this chapter exists to fix. You now have a quality bar, so you can tell when an answer is not good enough. What you probably do next is retype the request with a bit more force, and the reason that so often fails is that you skipped a step: you noticed the failure without naming it.

A bad answer is evidence. It has a recognisable shape, that shape has a cause, and the cause decides the repair. Retyping is what you do in place of diagnosis, and it is expensive. Three rounds of it costs 10 minutes and usually produces a slightly longer version of the same failure.

Nearly everything that comes back wrong belongs to 1 of 3 families. Some failures are yours: you left out a piece of context, or your instruction ran against what the assistant would have done anyway. Some belong to the conversation: it went wrong several messages ago and everything since has been built on top of that. And some are ceilings: it will not do the thing, or cannot do it accurately, and no rewording moves either.

One question separates the first 2, and it is worth more than the rest of this chapter. Was the very first answer already wrong in this way, or did it start out fine and go wrong later? Already wrong points at your ask, and you fix it where you are. Went wrong later points at the thread, and the fix is to leave.

Getting that backwards is the standard way to lose an afternoon. Correcting a conversation that has already committed to a wrong assumption feels productive, because every reply acknowledges the correction politely. Then the next answer quietly reverts, so you correct again, and now you are 6 messages deep in a thread that was never going to recover.

So before you read on, open the answer that failed your bar and read the whole conversation, not just the last reply. You are looking for the message where it first went sideways. If that message is the first one, section 23.2 is yours. If it is not, go straight to 23.3.

23.2

5 shapes that mean your request was the problem

These are the failures you caused, which is the good news: they are repaired where you are standing, in the same conversation, usually by adding 1 sentence.




Fig 23.2 Read down until a line matches what you are looking at. Every shape here is repaired in the conversation you are already in. None of them is a reason to start over, and starting over will not help.

23.2

Padding is a signal. Across 14 models tested on 5 datasets, verbose answers were measurably likelier to be wrong than direct ones. Read it as uncertainty, not as style.

The rule it drops. The instruction most often ignored is the one that opposes what it would have done unprompted. Forbid the habit by name.

It cannot count. A word count is estimated as it writes, never measured. Give a budget per part instead: 4 bullets, 1 sentence each.

Five shapes mean the fault is in your request rather than anywhere else. All 5 are repaired in the conversation you are already in, and most of them cost you a single sentence.

Padding is the most common and the most misread. It restates your question back to you, adds headings nobody asked for, lists considerations, and finishes without landing anywhere. Most people read that as the assistant being wordy. It is closer to a tell. A study across 14 models found the pattern showed up in about half of one leading model's answers, and that it tracked uncertainty: the padded answers were measurably likelier to be wrong than the direct ones. Padding usually means it does not have what it needs and is covering the gap. Supply the missing input, or ask for the answer in 1 sentence first and the working afterwards.

The second shape wastes the most senior time. You asked for something specific and what came back is the standard version of that job, competently done, for somebody else entirely. It recognised the category and produced the category's usual answer. There is now a good measurement of this. An August 2026 evaluation scored 12 leading models on rules that deliberately opposed what each model would have done unprompted, and every single one was worse on those rules than on rules it agreed with anyway, by between 3.6 and 7.4 points. The instruction most likely to be dropped is the one that fights its habits. So name the habit and forbid it: not just shorter, but no summary paragraph at the top and no closing paragraph at all.

Third: a rule you stated plainly, broken. 200 words comes back as 340. The word you banned appears twice. This is not defiance and it is not carelessness. It does not count words as it writes them, so a length is an estimate rather than a measurement, and a strict format competes for the same effort as the thinking. Two repairs work. Give a budget per part instead of a total, because 4 bullets of 1 sentence each is something it can actually track, where 200 words is not. Then check the constraint yourself, in the document, rather than asking whether it complied.

Fourth: you asked which one, and got 3 options, a paragraph on how it depends, and no recommendation. Two different things cause that. Either the answer genuinely turns on a fact it does not have, in which case give it the fact. Or you never authorised it to take a position, in which case do so out loud: assume the budget is fixed, pick one, give me the single strongest reason. It will pick. It hedges by default because hedging is rarely wrong, and rarely wrong was not what you needed.

Fifth, and quickest: the content is right and the form is not. You asked for an email and got headed sections. You asked for flowing prose and got bullets. This is the second shape again in its mildest version, because headed and bulleted is what it produces when nothing stops it. State the form as a rule rather than a preference, put it at the end of the ask where it sits closest to the writing, and say what it must not do. One thing worth knowing: a strict output format costs it some of its reasoning, so when the thinking is hard and the shape is rigid, get the thinking first and the shape second. That is Chapter 16's territory.

23.3 4 shapes that mean the conversation is the problem

These do not respond to better wording, because the thread is already carrying a wrong turn. The repair is to leave, and there is a right way to leave.

Rewrite the ask

The conversation is sound. One piece of your instruction was missing, or was quietly fighting its habits.

TRAITS

- The very first answer was already wrong in this way
- It has not committed to an assumption you disagree with
- You can name the sentence you failed to write

PROS

- + Costs 1 sentence and keeps everything the thread knows
- + You find out immediately whether you were right

CONS

- Does nothing at all once a wrong turn is baked in
- Tempting to try it 4 times before admitting that

Start a new thread

It took a wrong turn earlier, is still building on it, and repeating yourself in place is not clearing it.

TRAITS

- It obeyed, then reverted a few messages later
- It has forgotten something you settled earlier
- Your second correction landed worse than your first

PROS

- + Recovers close to a clean first attempt
- + Forces you to consolidate what you actually want

CONS

- You lose what it learned, unless you carry it across
- Feels like giving up, which is why people will not do it

Fig 23.3 The most expensive mistake in this chapter is treating the right-hand column as the left. Once a conversation has taken a wrong turn, more corrections inside it do not clear it, however politely each one is received.

23.3

Length alone does damage.

A single stated constraint loses more than 11% of its accuracy as a conversation grows longer. Several constraints at once lose more than 40%.

It does not reconsider. Across 200,000+ simulated conversations, an early wrong assumption was rarely recovered from. Corrections land on top of it, not instead of it.

Recap is not restart.

Summarising everything again inside the same thread recovered about 2 thirds. Carrying that summary into a fresh one recovered about 95%.

Four shapes mean the conversation itself has gone wrong, and not one of them is repaired by trying again inside it.

The first is drift. It obeys, and then a few messages later it quietly stops obeying, usually the constraint you cared about most. This is measured and it is not subtle. A May 2026 benchmark built from real conversations found accuracy on a single stated constraint falling by more than 11% simply as the conversation got longer, and following several constraints at once taking accuracy down by more than 40%. Adding or changing a constraint partway through cost another 9%. Nothing about your wording caused that. Length caused it.

The second is the expensive one. It made an assumption early, produced a full answer on top of it, and is now defending that answer rather than reconsidering it. A study of more than 200,000 simulated conversations found that when information arrives turn by turn instead of all at once, performance drops by an average of 39% across 6 kinds of task, and that almost all of the damage is this: it commits early and cannot get back. The authors' own summary is worth memorising. When it takes a wrong turn in a conversation, it gets lost and does not recover.

The third is amnesia. Something you settled 20 messages ago is simply gone, and it is confidently producing the version you already rejected. The fourth is the agreement spiral, where every push-back is accepted, every draft is now excellent, and it has stopped working as any kind of check on you. That last one has a cure and it is Chapter 21's, not this chapter's. Your job here is only to catch it early: if it has agreed with your last 3 corrections without once pushing back, believe the pattern rather than the compliments.

Starting again is not giving up, and it is not the same as retyping. The same research tested both moves. Recapping everything at the end of a conversation that had gone wrong recovered roughly 2 thirds to 3 quarters of what a clean attempt would have produced. Consolidating the same information into 1 fresh instruction recovered about 95%. A dead thread stays expensive even after you have said everything again inside it.

So do it properly. Before you close the conversation, ask it to summarise everything you have agreed so far in one paragraph: the requirements, the decisions, the things ruled out. Read that paragraph, correct it, then paste it into a new conversation along with your original request. It costs you 30 seconds and you keep everything the thread was worth. Chapter 11 is where this stops being a rescue and becomes permanent, because the requirements you keep retyping are exactly the ones that belong in saved instructions.

23.4

Some failures are ceilings, and no rewording moves them

A refusal or an invented number is not evidence that you asked badly. It is the edge of what the tool will do, and the skill is spotting one before you have rephrased it 4 times.




Fig 23.4 Four ceilings. Rewording moves none of them, so the value of recognising one is the 20 minutes you do not spend rephrasing. Bar heights are visual rhythm, not a ranking of how often each occurs.

23.4

Refusals are not improving.

One August 2026 evaluation found a model scoring 99.3 on safety alignment while refusing 1 harmless query in 3. Stronger alignment buys more of this, not less.

The label, not the task. The identical action gets refused more often when it is described with an alarming word. Rename it in neutral language and ask again.

There is no tell. Confidently wrong and confidently right look the same across 15 things a reader can actually see. That is why Chapter 8 gave you a habit rather than a checklist.

Two shapes are not about your wording at all, and knowing them on sight saves you the 20 minutes you would otherwise spend rephrasing.

The first is a refusal on legitimate professional work. You ask for a competitive teardown, a redundancy letter, a security review of your own product, a plain summary of a drug interaction for a briefing, and you get a polite decline or a watered-down answer that avoids the actual question. This has not improved with newer models, and it is worth being blunt about that. An evaluation published in August 2026 found a model scoring 99.3 on safety alignment while refusing 1 in 3 entirely harmless queries, and named the trade directly: stronger alignment systematically produces more of this.

What is useful is that the refusal usually attaches to your words rather than to your request. Work published the same month found that guardrails decline an identical action noticeably more often when it carries an alarming label instead of a neutral one. So the moves are practical. Rename the thing: attack becomes failure mode, kill becomes discontinue, exploit becomes weakness. Say who you are and what the output is for, in one clause. And ask for the part it can do, which is often 90% of what you wanted. If 3 attempts get you nowhere, it is a ceiling, and the answer is to write that part yourself.

The second is the confident specific: a number, a date, a name, a citation, delivered in exactly the same voice as the parts that are true. The unnerving property is not that it happens, it is that nothing on the page separates it from the rest. A 2026 audit tried to tell confidently-correct answers from confidently-wrong ones using 15 features a reader can actually see, and got essentially nothing. That is why Chapter 8 gave you a verification habit instead of a list of warning signs. There are no warning signs.

Be straight with yourself about which of these are getting better, because a great deal of advice in circulation is describing 2023. Fluency, single-step arithmetic, obeying 1 clear instruction in a short exchange: all much better than they were. Long conversations, several constraints at once, refusals on professional work, and specifics invented with total confidence: all still here, all measured this year, and refusals arguably worse than before. Plan around the second list.

Two more ceilings belong on your list without belonging in this chapter. It will do arithmetic in prose and get it wrong in ways that read as right, and that is Chapter 19. It cannot see the file, the page or the week you assumed it had, and that is Chapter 20. Both feel like failures of wording. Neither is. Go and add the 2 ceilings you have personally hit to the companion page, because the list that matters is yours, not this one.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 24.

Do these before Chapter 24

3 exercises. About 15 minutes

Open the answer that failed your bar. Do not touch the request yet. Read the whole conversation and fill in these 3 boxes first.

THE SHAPE, IN 2 OR 3 WORDS

ASK, THREAD, OR CEILING

THE ONE MOVE THAT FOLLOWS

If you wrote ask: add the sentence you never wrote and send it in the same conversation. Note what changed, and whether the shape you named actually went away or just got longer.

If you wrote thread: ask it to summarise everything you have agreed in one paragraph, correct that paragraph, and paste it into a new conversation with your original request. Note which of the two answers you would actually use.

Chapter 23 on one page

THE TRIAGE QUESTION

Was the first answer already wrong this way, or did it go wrong later? Already wrong is an ask failure and you fix it here. Went wrong later is a thread failure and you leave

REWRITE THE ASK WHEN

It padded and restated · it did the standard version of the job · it broke a rule you stated · it hedged instead of choosing · it came back in the wrong form

START AGAIN WHEN

It obeyed then reverted · it forgot something you settled · it is defending an early assumption · it agrees with everything now (Chapter 21). Leave properly: get a one-paragraph summary of what you agreed, correct it, paste it into a fresh conversation. Recapping inside the dead thread recovers far less

NEITHER WILL FIX

A refusal (rename the thing in neutral words, say who it is for, ask for the part it can do) · an invented specific (Chapter 8) · arithmetic (Chapter 19) · what it cannot see (Chapter 20)

NEEDED BEFORE CHAPTER 24

One failure you named, and the repair that cleared it

BRING TO CHAPTER 24

The failure you named and the repair you ran. Chapter 24 turns AI 202 into a playbook you use without thinking, and the repair section is the half you reach for on a bad day.

BOTTOM LINE

You can name what went wrong in a phrase, say whether the fault is in your ask or in the thread, and run the repair that fits.

NEXT CHAPTER

Your Prompt Playbook

Chapter 24 closes the track by assembling it: the framework, the examples, the templates, the checks and the repairs, as one system you run without thinking about it. You leave with a page you actually use rather than a chapter you half remember.

BRING WITH YOU


- ☐ The failure shape you named here, and the repair that cleared it
- ☐ Your quality bar from Chapter 22, and the prompts you already reuse

GO FURTHER

The Ultimate AI Library

Every guide, prompt pack and playbook AI Central has published, in one place. The reference shelf this course keeps pointing you to.

explore the library




WHO MAKES THIS



AI Central Media is a London-based editorial team. Our flagship publication covers practical AI for senior professionals, and 100+ AI companies, SaaS platforms and growth teams have advertised with us since 2023.

100+ AI companies, SaaS platforms, education brands and growth teams have advertised with us since 2023	613K Accounts reached every month, across the 7 channels we operate	151 Countries our readers are in, plus all 50 US states
--	--	--

PROUD PARTNER OF



We turn attention into pipeline.

visit thecentral.ai

➤