



A FIELD GUIDE FOR THE UNREHEARSED

The AI Question

How to answer the interview question everyone's asking — and no one's prepared for.

INSIDE THIS GUIDE

- ⬆ The hidden rubric interviewers use to score your answer — in plain language
- ⬆ The double bind that fails both AI enthusiasts *and* AI skeptics
- ⬆ A 90-second answer formula, with a fill-in template you can finish in 5 minutes
- ⬆ Scripts for all seven versions of the question — including the follow-up traps
- ⬆ The rules for when AI use gets you hired vs. disqualified

The question you haven't rehearsed

You have a polished answer for "tell me about yourself." You could recite "what's your greatest weakness" in your sleep.

You don't have an answer for **"How do you use AI in your work?"**

Candidates describe the moment the same way, over and over: *"I was asked how I use AI in my process. I was quite surprised, not gonna lie."* — *"They asked me about AI — I was like, wtf???"* It shows up in interviews for roles that have nothing to do with AI — marketing, ops, HR, finance, sales, support — dropped casually into an otherwise ordinary conversation.

71%

of leaders would hire a *less experienced* candidate with AI skills over a more experienced one without

66%

say they wouldn't hire someone without AI skills at all

+56%

wage premium for AI skills — *largest* in non-technical roles

Sources: Microsoft/LinkedIn Work Trend Index (n=31,000) · PwC AI Jobs Barometer (~1 billion job ads analyzed)

Shopify, Microsoft, Meta, and Duolingo have all issued company-wide mandates making AI use a baseline expectation — and hiring practices follow the memos.

Here's the part that matters: the question almost never appears on the formal interview script. An analysis of 6,380 recorded first-round interviews found *explicit* AI questions in barely 2% of them. Instead, it arrives **informally and unannounced** — one conversational question, or hidden inside "walk me through a process you improved." That's exactly why nobody has rehearsed it. It's the new "greatest weakness" — except no standard script exists yet.

The good news: candidates who prepared report the same result. *"I went home and learnt as much as I could in one evening — and that did the trick."* This guide is that evening.

What they're actually scoring

First, drop the assumption that interviewers want AI *enthusiasm*. They don't. They're screening against **risk**.

Unverified AI output — polished-looking work that's subtly wrong — costs companies an estimated **\$186 per employee per month** to clean up. **66%** of employees admit using AI output without checking it; **56%** have made work mistakes because of it (KPMG / University of Melbourne, 48,000 respondents). Every interviewer has seen the damage.

So when they ask about AI, they're listening for four things:

What they're scoring	The question behind the question
1 • When you reach for it	Do you pick the right tasks to hand to AI — and keep the right ones for yourself?
2 • How you direct it	Can you give it context and iterate, or do you paste a one-line prompt and pray?
3 • How you check it	What's your verification step? How do you catch it being confidently wrong?
4 • Who's accountable	When AI-assisted work goes out with your name on it — do you own it?

Number 3 is the differentiator. Almost every candidate can now say they use AI. Almost none can explain how they *verify* it. A single concrete verification habit — "I check every citation before anything leaves my desk" — separates you from the pile.

Knowing this rubric before you walk in puts you ahead of the overwhelming majority of candidates, because they're answering a different question ("do you like AI?") than the one being scored.

The double bind: both extremes fail

This is the trap nobody names. There are two obvious ways to answer — and **both get rejected**.

"I don't really use it." Reads as an adaptability red flag. Hiring managers openly vent about candidates who, asked what AI tools they know, "just say Midjourney." In a market where 71% of leaders prefer AI skills over experience, "not really" sounds like "not curious."

"I use ChatGPT for everything!" Reads as a judgment red flag. Roughly **1 in 5** hiring managers treat heavy AI reliance as a strike against you — it signals exactly the unverified-output risk they're screening for. Interviewers also report waves of candidates giving *"the EXACT same answer"* — enthusiasm with no specifics is indistinguishable from faking.

The skeptical hot take fails too. In one widely shared story, a candidate told a CEO he uses AI only as "smarter autocomplete, not a crutch" — technically defensible, delivered as a dismissal. Rejected.

The needle you're threading: *enthusiastic, specific, and in control.* You use AI deliberately, you can prove it with details, and you're clearly the one steering.

■ Red-flag phrases — interviewers are trained to catch these

- "I use it for everything."
- "I've been experimenting lately..." (vague timeline = no real usage)
- Claiming months of use but having no saved or refined prompts
- Unable to name a single time AI got something wrong for you
- Philosophy about the future of AI instead of a workflow you actually run

The 90-second answer formula

Every winning answer — across employer rubrics and real candidate success stories — contains the same six elements:

- 1 **Named tools, by task.** Not "I use AI" — *which* tools for *which* jobs.
- 2 **One specific workflow.** A real process you run, step by step.
- 3 **A number.** Time saved, output improved, turnaround cut.
- 4 **An iteration detail.** Proof you refined it — "took six versions before it was reliable."
- 5 **A failure + your fix.** The credibility element almost everyone omits.
- 6 **A judgment boundary.** Where you *don't* use AI, and why you stay accountable.

A model answer, annotated (adapt the structure, not the words):

"I keep two tools open for different jobs — Claude for drafting and analysis, ChatGPT for quick research summaries ^[1]. My main workflow is client proposals: I feed in the brief and past examples, generate a structured first draft, then rewrite the strategy section myself — that part has to be mine ^[2]. It cut my turnaround from about three days to one ^[3]. The prompt took me five or six iterations before the drafts stopped being generic ^[4]. It has burned me — early on it invented a statistic, so now I verify every number and citation before anything goes out ^[5]. And I don't use it for final pricing or anything touching confidential client data. It's a power tool — I'm still the one accountable for what ships ^[6]."

Read that aloud. It's about 90 seconds — long enough to be substantive, short enough to invite follow-up on *your* terms.

Why the failure story matters: it's the one element that can't be faked by someone who started using AI yesterday. It converts "I use AI" into "I have judgment about AI" — which is the thing being scored.

Build yours in 5 minutes

Fill in the six blanks. Real details only — every one of these gets probed in follow-up.

01 My tools, by task

I use _____ for _____ and _____
for _____.

02 One workflow I actually run

When I need to _____, I _____.

03 The number

It cut/improved _____ from _____ to _____.

04 The iteration detail

It took me _____ attempts/versions before _____.

05 A failure and my fix

It once _____, so now I
_____.

06 My boundary

I don't use it for _____, because
_____.

Now say it out loud, once, as a single answer. That's the whole exercise.

A second flavor — operations/PM version (to show the formula travels):

"I use AI mainly for meeting synthesis and first-draft documentation — Claude for anything long, plus the transcript summarizer in our meeting tool. Weekly status reports used to take my whole Friday morning; now I generate the draft from the week's notes and spend 30 minutes editing instead of three hours. The summaries missed action-item owners at first — took a few prompt revisions to fix. I still write anything stakeholder-facing myself, and I never paste customer data into public tools."

Same six elements, different job. Yours will sound like *your* job — that's the point. Identical-sounding answers are what interviewers are complaining about.

The seven versions of the question

One rehearsed core answer covers most of these — this table tells you which element to lead with.

They ask...	They're really testing...	Your move
"How do you use AI in your work?"	Workflow reality	The full 90-second formula
"When would you <i>not</i> use AI?" · "How do you validate its output?"	Judgment — the differentiator	Lead with your boundary + verification step (elements 5–6)
"Tell me about a time AI improved your work"	Behavioral proof	Workflow + number (elements 2–3), told as a story
"What do you think about AI in our field?" · "Are we in an AI bubble?"	Perspective without hot takes	Balanced take: real about limits, clearly adaptive. Never dismissive, never doom
"What AI tools do you know?"	Depth vs. name-dropping	Differentiate tools <i>by task</i> (element 1). One-tool answers read as shallow
"Did you use AI to write this?"	Integrity	Honest disclosure + what stayed yours (see page 9)
"What's new in AI?" · "What do you think of [latest model]?"	Genuine curiosity	One recent thing you actually tried — not a headline recap

The follow-up trap: strong answers get probed. *"Show me the prompt."* *"What did it get wrong?"* *"Walk me through yesterday's usage."* This is where invented answers die — and why every blank on the previous page has to be true. If you have a refined prompt saved anywhere, screenshot it before your interview; it's the single most convincing artifact you can reference.

"But I barely use AI"

This is the highest-anxiety version of the problem — and the most fixable. You don't need months of history. You need **true stories**, and true stories can be manufactured deliberately in one week.

Tonight — 15 minutes

Pick three real tasks from your actual job. Not toy exercises — things you did this month. A report, an analysis, an email sequence, a plan, a spreadsheet.

This week — 30 minutes per task

Run each one through an AI tool. Keep a simple log: what you asked, what worked, what it got wrong, and how you caught it. That log is elements 2, 4, and 5 of the formula — workflow, iteration, failure-and-fix.

This weekend — 10 minutes

Drop your best task into the fill-in template on page 6. Done.

You're not faking experience — you're compressing it. The honest version scores fine: *"I've been deliberately building AI into my workflow recently — here's what I've learned so far, including what it gets wrong."* Recent-but-real with a verification habit beats fake-veteran every time, because fake-veteran collapses at "show me the prompt."

If your employer bans AI at work: say exactly that, then pivot to what you've done on your own — *"My current company restricts AI tools, so I've been building fluency on personal projects: [example]. I'd be glad that's not a constraint here."* Policy awareness is itself a point in your favor — it demonstrates the responsible-use instinct they're screening for.

When AI use gets you hired vs. disqualified

The same skill that wins the AI question can end your candidacy if you deploy it in the wrong place. Three rules cover every case:

Rule 1 — AI for *preparation* is safe and increasingly expected

Researching the company, drafting your stories, rehearsing answers, generating practice questions: universally accepted, and a legitimate part of your answer to the AI question itself.

Rule 2 — Covert AI use *during* a live interview is disqualifying

Amazon and Goldman Sachs ban it explicitly. 65% of hiring managers say they've caught deceptive AI use, and companies including Google and McKinsey have reinstated in-person rounds specifically to catch it. Interviewers also probe past real-time-assistant answers — candidates freeze the moment the follow-up leaves the script. The whisper-bot doesn't just risk disqualification; it prevents you from ever building the fluency being tested.

Rule 3 — For take-homes, read the instructions and disclose

Policies now split by company: some expect AI-assisted work, some prohibit it, most reward transparency. When it's allowed (or ambiguous), disclose cleanly:

"I used AI to draft the initial structure; the analysis, judgment calls, and final version are mine. Happy to walk through exactly how I used it."

That sentence turns a gray area into a demonstration of accountability — the thing the whole rubric is about.

■ One absolute

Never mention (or commit) feeding a current or former employer's confidential data into public AI tools. It's an instant disqualifier — and *"I don't put confidential data into public models"* is itself one of the strongest boundary statements you can make.

Make it automatic

Here's the gap this guide can't close on paper: **knowing your answer is not the same as saying it under pressure.**

The AI question arrives unannounced, mid-conversation, with follow-ups. People who genuinely use AI every day still bomb it — not from ignorance, but because they've never *narrated* their workflow out loud. "I knew the answer — I just couldn't say it" is the most common interview failure there is. The fix is the same one elite performers use for everything: rehearse out loud, get scored, iterate.

Your one-evening checklist:

- ☐ Fill in the six-line template (page 6) with real details
- ☐ Say your 90-second answer out loud — twice
- ☐ Screenshot your best refined prompt (your "show me the prompt" insurance)
- ☐ Pick your lead element for the judgment version: boundary + verification
- ☐ Rehearse the disclosure script (page 9) once, so it's available if you need it

Then pressure-test it. Interview Drills scores your spoken answer the way an interviewer would — specificity, structure, judgment — so you find the gaps in rehearsal instead of in the interview. The candidates who win this question aren't the ones who use AI most. They're the ones who rehearsed talking about it.

Practice it out loud, scored.

Say your answer. Get an interviewer-grade score. Fix the gaps before it counts.

interviewdrills.com →

Every statistic in this guide traces to primary research: Microsoft/LinkedIn Work Trend Index (n=31,000) · PwC AI Jobs Barometer (~1B job ads) · Harvard Business Review's analysis of 6,380 recorded interviews · KPMG/University of Melbourne (n=48,000) · BetterUp/Stanford · Lightcast · published employer rubrics from Indeed, Anthropic, and others.