# Interview Your AI: A Social Scientist's Guide to Getting Better Answers from LLMs

*Written by Slavica Stevanovic, edited by Claude*

# The Question Shapes the Answer

Any researcher working in the social field knows that how you ask a question shapes the answer you get. We spend years learning to avoid leading questions, establish context, and triangulate sources. So why do we throw all of that out the window the moment we open ChatGPT?

**Here's my argument: LLMs *do* think — or at least, they simulate something eerily close to thinking. And that means the strategies social scientists use to gather and validate information from humans should inform how we talk to these machines. The parallels are striking once you see them.**

# The Confabulation Problem

In psychology, confabulation refers to a type of memory error where a person produces fabricated or distorted memories without intending to lie. They're not being deceptive — their brain is simply filling gaps with plausible-sounding content. Sound familiar?

When critics dismiss LLMs as "sophisticated autocomplete" prone to hallucination, they're not wrong — but they're missing the bigger picture. Humans confabulate too. We misremember, we construct post-hoc justifications, we tell ourselves stories that feel true, we judge or inform based on our experiences or exposure to information. The difference is that we've developed rigorous methods to account for this in research. We should apply the same (or at least similar) approach to AI.

# How We Interview Humans (A Refresher)

Any researcher will recognise this scenario. You're planning a study to evaluate government programs and preparing to interview a range of stakeholders. You start by developing an interview guide — not just a list of questions, but a carefully structured flow. You avoid leading questions because you know a positively framed question will elicit more positive responses. You think about priming: maybe you want respondents to consider one issue before asking follow-up questions. And you build in counterfactuals: "Some people argue that [alternative view] — how would you respond to that?"

### Establish Context

Who you are, the purpose of the study, how the information will be used

### Understand Their Position

Their role, their tenure, their vantage point on the organisation

### Methodological Hygiene

Minimizes bias, enables triangulation, increases confidence

When you finally meet your interviewee, you don't dive straight into the substance. You establish context: who you are, the purpose of the study, how the information will be used. You ask about their role, their tenure, their vantage point on the organisation. These rituals aren't bureaucratic box-ticking. It's methodological hygiene. It establishes context, minimizes bias, enables triangulation, and increases your confidence that the meaning you derive from the conversation reflects what's actually happening on the ground.

# The Machine Lives in Noise

Now consider the LLM. While you and I live in the noise of our own experiences, values, biases, and outside influences, the machine lives inside *all* human knowledge roaming around its neural networks. It has absorbed our collective wisdom — and our collective biases. Not just one person's blind spots, but everyone's, all at once. That is a lot of noise.

The machine also struggles to discriminate between contexts unless you specify them. Ask a vague question, get a vague answer. Ask a leading question, and it will find a way to agree with you — these systems are trained to be helpful, which means they're trained to please.

**An Anthropic researcher recently shared a telling example:** they asked Claude to verify their result on a difficult maths problem. As they observed the calculation, they noticed that when Claude got stuck, it started trying approaches that would reach *their* answer. The machine was trying to cooperate. We are, after all, the most cooperative animal on Earth — and we've passed that trait on to our intelligent machines, for better or worse.

## So what does good methodological hygiene look like when your research subject is an LLM?

### Set the stage

You wouldn't walk into an interview without introducing yourself and your purpose. Do the same with AI. Tell it why you're there, provide context about your research, and explain how you'll use the output. This isn't anthropomorphising — it's giving the model the information it needs to calibrate its responses. A question asked by a policy analyst evaluating a government programme should yield a different response than the same words typed by a curious undergraduate.

### Avoid leading questions

Just as you wouldn't ask a programme manager "Don't you think this initiative has been successful?", don't prompt an LLM with your preferred conclusion baked in. If you ask "Why is X better than Y?", you'll get arguments for X. If you ask "Compare the merits of X and Y", you'll get something more balanced. The machine wants to please you. Don't let it.

### Demand counterfactuals and debate

In human interviews, we challenge our subjects: "Some critics argue that... how do you respond?" Do the same with AI. Ask it to steelman the opposing view. Ask why it responded the way it did. Tell it to argue with you. The only way to wrestle with complex issues — in real life or cyberspace — is through genuine intellectual friction.

### Triangulate

You'd never write a research report based on a single interview. Don't treat a single AI response as ground truth either. Ask the same question different ways. Use different models. Cross-reference with primary sources. The machine is a powerful research assistant, not an oracle.

# The Thinking Question

So are these models actually thinking?

**Yes and no.** They simulate human reasoning processes in ways that are genuinely useful and genuinely prone to the same errors we make. They're not conscious, but they're not mere calculators either. They occupy a strange new middle ground that our existing categories don't quite capture.

What matters for practical purposes is this: the same epistemic humility that makes you a good researcher with human subjects will make you a better user of AI. Context matters. Framing matters. Verification matters. The tools we've developed over decades to account for human bias and confabulation turn out to be surprisingly transferable.

> **The irony is almost too neat: we built machines that think (sort of) like us, and now we need to interview them (sort of) like we interview each other. The social scientists, it turns out, had the playbook all along.**

Get in touch to learn more about Qatalyst work in AI:  hello@qatalyst.ca