When and How to Use AI

A Framework for Reflection

Jeff Fossett, PhD

Research Fellow, Technology & Society Initiative

Johns Hopkins Carey Business School

Goals for today

- Generative AI tools are everywhere and can do lots of things — but when and how should we actually use them in our work? (if at all!)
- Today, I will share a framework for reflection: a set of questions that I hope can help you decide when, whether and how to work with generative AI.
- This framework comes from my own context (academic social science research), but the goal is a way of thinking you can be useful for your work.

The point is not the conclusions — it's the framework

I'll share where I land, but your answers may (and probably should!) differ based on your goals, priorities, and tasks.

A task-based view of the work

A research project is really a pile of smaller tasks:

- developing an idea
- searching the literature
- reading related papers
- making a research plan
- collecting data
- cleaning data
- analyzing data
- building charts & tables
- writing up results
- responding to referees ...

For each task, the AI decision can be different.

So: given a task, how do I decide whether — and how — to use AI?

Three questions for any task

Given a task, we can ask...

Q1 Can I?
Can AI help with this task *in principle*? Is using AI even an option here?

Q2 Should I?
Just because I *can*, doesn't mean I *should*.

Q3 How?
If I do use AI — how hands-on or hands-off should I be?

Can I? → Should I? → How?

Q1 — Can AI help in principle?

Q1: Can AI help here in principle?

At a basic level, generative AI can be used to generate text, code, images.

- Some tasks can obviously be reduced to generating code, text, images → AI is an option
- Other tasks can be *facilitated* by generating text or code, even when that is not the task per se.

Example: Running in-person interviews

Generative AI probably can't run my in-person interviews — but it can draft an interview script, give feedback on my approach, suggest references etc.

Q1: Can AI help here in principle?

Two tactics for identifying whether AI can be helpful for a given task:

- **"Go meta" on the task.** AI may not do the task itself — but it might help you *think about how* to approach it. Going up a level can open a door.
- **Watch the context constraint.** Many tasks need specific situational context to do well. Your ability to *give the model that context* is a key limit on how much it can help — at any altitude.

Q1: Can AI help here in principle?

Final point: the answer to "can AI help?" also depends on which model you are using, and its capabilities — and that can change over time.

Takeaway

It's worth experimenting with different tools and models to build personal intuition for when they can be helpful — and to continually re-evaluate this as the landscape changes.

Q2 & Q3 — Should I work with AI? And how?

Q2 & Q3 — Should I work with AI? And how?

Suppose AI *can* help with a task. Should I use it — and if so, how?

Caveats:

- For this presentation, I'll bracket the bigger societal questions around AI — energy & environment, labor markets, IP & attribution, AI safety etc.
- These matter and are worth engaging with — and for some people they settle the issue entirely (much like choosing not to eat meat, or to boycott social media).

Here I'll assume we're *not* categorically opposed to AI, and focus on the narrower question above.

Q2 & Q3 — Should I work with AI? And how?

Four helpful questions to decide whether and how to work with AI for a given task:

(i) How much does accuracy / getting the details right matter for this task?

(ii) How easy is it to verify the task was done correctly?

(iii) How much harder is the task by hand vs. with AI?

(iv) What is lost or gained more broadly by doing this task by hand vs. with AI?

(i) How much does accuracy matter for this task?

This question is important because:

- LLMs are fundamentally fallible, sometimes in non-obvious ways — e.g. *hallucinated* facts or citations, asserted confidently.
- Even setting hallucination aside: delegating a task to AI often involves describing a task/goal and letting the model fill in the details → room for misspecification or misinterpretation.
  - Similar to handing off a task to a human collaborator – risk that they misunderstand, misinterpret, do the wrong thing etc.

AI Usage Implication

The more that accuracy and getting the details right matter, the ↓ less likely I am to delegate the task to AI (or the more hands on I will be).

(ii) How easy is it to verify that the task was done correctly?

This question is an important complement to (i) because:

- If it's easy to audit or verify accuracy, I can hand the task off to AI and then readily confirm the result is correct / as desired — even if accuracy matters a lot.

AI Usage Implication

The easier it is to verify the result, the ↑ more willing I am to delegate the task to AI — since I can confirm afterward that it was done as desired.

Corollary: if accuracy matters for a task, I won't delegate it to AI unless I have the knowledge / expertise / capacity to verify that it was completed accurately.

- *E.g.* I won't hand off a statistical analysis unless I'm confident I can confirm it was done correctly.

(iii) How much harder is it to do this task by hand vs. with AI?

The basic cost–benefit question:

- Will it actually be easier to work with AI for this task — and by how much?
- Seems obvious, but worth keeping in mind.

AI Usage Implication

The *much* easier a task is to do with AI than by hand, the ↑ more willing I am to delegate the task to AI.

(iv) What is lost or gained, more broadly by doing this task with AI?

Three things to think about here:

- **Opportunity cost.** Offloading one task frees my attention for higher-leverage work (e.g. let AI handle the syntactic details so I focus on the design).
- **Long-term skills.** Repeated delegation $\Rightarrow$ skill atrophy. *Is this a skill I want to keep building? Is this something I want to learn?*
- **What the task is for.** Some tasks are valuable *because I do them myself* — a thank-you note, or expressing my own perspective.

A useful thought experiment

How would my life be if I *always* used AI for this? How would the world be if *everyone* did?

Q2 & Q3: Should I work with AI? And how?

There's a lot to say about *how* to work with AI.

For me, the key dimension is the extent to which I abstract myself away from the execution details — i.e. how hands-on vs. hands-off I am.

- This dimension roughly corresponds to different AI tools/paradigms (see next slide)

Q2 & Q3: A spectrum of abstraction

Level of abstraction	Human/AI role	Example
0. Not abstracted (Zero AI)	Human fully completes task without any AI engagement	N/A
1. Not abstracted (AI advisor)	Human fully completes task, but with advice or brainstorming support from AI. No output directly generated by AI	Using ChatGPT to brainstorm an essay idea, then writing yourself
2. Partially abstracted (AI assistant)	Human mainly completes task, but with assistance from AI	Writing / coding in Cursor, and periodically asking AI for help e.g. with formatting, rewording etc.
3. Partially abstracted	AI mainly completes task, but with oversight, direction, feedback from human	Claude Code w/ more detailed oversight & code review; Cursor, implementing and reviewing one step at a time
4. Fully abstracted	AI completes whole task, human reviews output	Claude Code w/ ignore all permissions

The four questions (i–iv) help me place myself on this spectrum:

↑ More willing to delegate ⇒ further down the list (more abstracted).

↓ Less willing to delegate ⇒ higher up the list (more hands-on).

Examples

Example 1: the task — formatting a LaTeX table

The task

Format (or reformat) a table in a research paper using LaTeX.

- LaTeX is a typesetting language: you write code, which is then "compiled" into a polished, professional-looking table.
- The syntax is finicky — but the result is easy to inspect once compiled.

Is this a good task for AI? Let's apply the framework.

Example 1: applying the framework

Question	Answer	AI usage implication
Q1: Can AI help in principle?	**Yes** — translating a table design into LaTeX is fundamentally a code-generation task, which AI tools can readily do.	AI is an option
Q2: Should I work with AI? And How? [hide sub-questions ▾]	**On balance, yes** — the considerations below mostly point toward this being a good task to delegate to AI.	Willing to delegate to AI (level 3–4)
(i) Does accuracy matter?	**Yes, a lot** — the numbers and formatting in a published table need to be exactly right.	↓ Less willing to delegate to AI
(ii) How easy is it to verify?	**Very easy** — I can compile the table and visually confirm the formatting and numbers are correct and as desired.	↑ More willing to delegate to AI
(iii) How much harder is it by hand vs. with AI?	**Much harder by hand** — LaTeX table syntax is finicky, idiosyncratic, and arduous to write manually.	↑ More willing to delegate to AI
(iv) What is lost / gained by hand vs. with AI?	**For me, little is lost by working with AI for this task** — I accumulate less LaTeX formatting expertise, which I am fine with because I don't care about becoming a LaTeX formatting expert; by working with AI, I also free up time to focus on higher-value questions like "what is the best way to format this table?" to clearly communicate the idea.	↑ More willing to delegate to AI

Example 2: the task — cleaning a raw dataset

The task

Clean a raw dataset to get it ready for analysis.

- Handling missing values, recoding and constructing variables, merging files, filtering observations, and so on.
- Involves many small judgment calls — each of which can be consequential for the final results.

Could I hand this task off to AI?

Example 2: applying the framework

Question	Answer	AI usage implication
Q1: Can AI help in principle?	**Yes** — data cleaning is largely a matter of writing code, which AI tools can readily produce.	AI is an option
Q2: Should I work with AI? And How? [hide sub-questions ▾]	**On balance, only minimally** — the considerations below mostly point away from delegating this task.	Minimally willing to delegate to AI (level 0–2)
(i) Does accuracy matter?	**Yes, a lot** — small cleaning choices can be consequential for the final results, and I need to be able to stand behind them.	↓ Less willing to delegate to AI
(ii) How easy is it to verify?	**Hard** — to be confident it is right, I would have to step through essentially every line, since errors can be subtle yet consequential.	↓ Less willing to delegate to AI
(iii) How much harder is it by hand vs. with AI?	**Arduous by hand** — but because I have to verify the output line-by-line anyway, working with AI saves little time; I might as well write the logic myself.	↓ Less willing to delegate to AI
(iv) What is lost / gained by hand vs. with AI?	**A lot can be lost** — doing it myself builds my sense of the data and what is actually happening in the analysis (the difference between reading a book and reading a summary of it).	↓ Less willing to delegate to AI

Rapid-fire

Same framework, quick calls:

- **Feedback / brainstorming / suggestions** — Can be a ↑ good AI task. Why? Quality doesn't matter much – if AI suggests 9 bad ideas but the 10th is good, I can keep the good one and ignore the rest. (But also: risks of skill atrophy?)
- **Finding relevant papers to read** — Can be a ↑ good AI task for similar reasons.
- **Drafting an abstract for a finished paper** — borderline. Part of it is just summarizing (reasonable to delegate), but part is expressing *why the work matters* — a personal, expressive judgment that AI can answer plausibly, but maybe not in *my* voice.

My Typical AI Workflow

What my typical AI workflow looks like

So what does this framework imply about my own workflow as an academic social scientist?

- For most of my research tasks, I land between levels 0 and 3 on the spectrum of abstraction from earlier — somewhere between "no AI" and "partially delegated to AI".
- I *rarely* fully delegate a task to AI, mainly because accuracy matters a lot to me and I want to keep an eye on every step.

My main interface: Cursor

Given this, I mostly work in Cursor, because:

- **I can easily incorporate relevant context.** I typically keep each project in one high-level folder (data, code, results, write-ups, references) and open Cursor at that altitude — shared context for everything.
- **It's flexible across very different AI uses.** E.g. can ask a question vs. revise vs. do fully agentic work in same space → can adapt my AI usage to the task at hand.
- **The UI is useful for hands-on direction & oversight.**
  - Can refer to specific files or lines in prompts → I can point the model precisely at what I mean.
  - Can see clear diffs of what changed → I can keep an eye on the details.

A worked example: how I made these slides

A meta-example: how did I build these slides?

Building this deck was really two tasks — which I approached very differently:

1. **Write a full draft** of my thoughts (as a blog post).
   Did this by hand, with almost no AI. ≈ level 0–1
2. **Turn the draft into slides.**
   Did this in a partially delegated way, with AI. ≈ level 2–3

Task 1: Write the draft by hand

- I wrote the full post myself, with almost no AI (just a little formatting help).
- **Why?** Developing my *own* thoughts is something only I can really do. ("Writing is thinking.")

$\approx$ Level 0 on the spectrum $\rightarrow$ no AI.

Read the full post:

jeffreyfossett.com

Task 2: Turn the draft into slides — partially delegated

After writing the draft by hand, I turned it into slides in a partially delegated way, with AI, working in Cursor.

Three steps:

1. **Gather context** and open the project in Cursor.
2. **Prompt** for a rough first draft of the slides.
3. **Refine** the draft in detail — hands-on, iterating with AI.

$\approx$ Level 2–3 on the spectrum $\rightarrow$ partially delegated.

Task 2 — Step 1: Gather context, open in Cursor

- I put the blog post *and* everything else about the project in one folder:
  - Email thread with Caseysimone about this presentation
  - Her feedback on my post on this topic
  - Guidelines she provided about this course (PDF)
- Open the folder in Cursor → the model can draw on all of it.
- Note: working in plain text as much as possible makes it easier to include as context and reference specific files and lines in prompts.

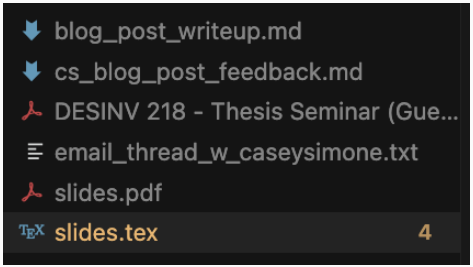

My Cursor sidebar with relevant context files

Task 2 — Step 2: Prompt a rough first draft

- A detailed prompt: referencing specific files, plus my own additional outline of the slides.
- Goal: a *rough draft* of slides based on the blog post for me to refine.

Task 2 — Step 3: Refine, hands-on

- Work through the draft in detail, changing what I don't like — by hand or with AI.
- Often referring to specific parts of the post or other materials.
- **This still takes a long time!** Mostly I use AI so I needn't fuss over LaTeX syntax — and sometimes for feedback or to help turn ideas into concise bullets.
  - AI is helping, but I am still doing the conceptual heavy lifting.

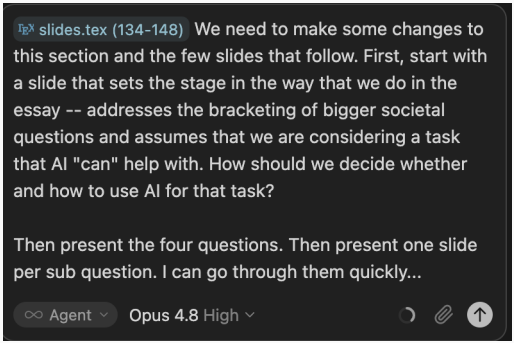

Wrapping up

Recap: a framework for reflection

Working task by task, we asked:

Q1 Can I use AI for this task? — can AI help with this task *in principle*?

Q2 Should I use AI for this task? — weighing accuracy, verifiability, effort saved, and what's gained or lost more broadly.

Q3 How should I use AI for this task? — where to sit on the spectrum of abstraction, from fully "by hand" to fully "delegated" to AI.

We saw this play out in concrete examples, in my own workflow — and even in how these slides were made.

The real takeaway: be intentional

- There's no single right or wrong way to work with AI — and you don't have to agree with where I land on any given task.
- The point of this presentation is to encourage you to be intentional about your own relationship to the technology.

If you take one thing away

Don't reach for AI — or avoid it — on autopilot. For each task, pause and think through the considerations, then make an intentional choice that fits *your* goals and values.

Thank you!