

Six Questions for Practical AI Ethics

“How can I use (Generative) AI responsibly?” is the question no one wants to ask. Maybe it’s because they don’t want to be bothered with the dos and don’ts. Maybe it’s because they don’t want to know what they’re doing wrong. Not to mention that “How can I use AI to get my work done faster?” has a much more exciting outcome.

That being said, the question about using AI responsibly is crucial. Currently, there are no universal rules. No government body protects you. No one is coming to save you if you get it wrong. Every decision about how you use AI in your business falls entirely on you. And remember, “AI said it was okay” is not a *Get Out of Jail Free card* (like we had in Monopoly).

I hold two granted patents and two pending in Artificial Intelligence. I have spent years in AI research. I now work at a company that focuses on practical and responsible implementation of Generative AI.

We will be looking at 6 important questions that all business owners should be asking.

1 Can you actually trust what your AI just told you?

2 Are you protecting your clients’ data?

3 Who owns what AI creates for you?

4 How do you avoid bias in what AI does for you?

5 Are you being transparent about how much AI is doing? (Should you?)

6 Are you using AI to think better, or have you quietly stopped thinking altogether?

These are questions that should be a part of every use of Generative AI in your company.

The businesses that get this right will build more trust, more loyalty, and more longevity than the ones that do not.

Lens 1: Truth

Can you actually trust what your AI just told you? Everyone has examples of instances where Generative AI clearly got it wrong. Then when it comes to using AI, however, it's pretty clear that many aren't reading (let alone double-checking) the full output before using it. The danger is that Generative AI is not 100% accurate every time.

A colleague asked AI for a recipe for a party. They got a recipe, purchased all the ingredients for a double batch, and after baking it, threw the dessert in the garbage, as it was inedible. When they later asked AI if the recipe was real, it admitted that it was not. It had gathered a list of ingredients and then made up all the ratios with full confidence.

While that one was mostly just a waste of some money, not all AI hallucinations are as innocent.

Back in 2023, a NY attorney cited 6 legal precedents that all turned out to be fabricated by Generative AI. Since then, other attorneys have fallen into the same trap. For attorneys, this means being sanctioned by the courts, which can put their livelihoods at risk.

The failing is not using AI. It's trusting that AI is checking its own work (or asking AI to verify its own output). Generative AI will sound absolutely sure of itself every time, but that does not mean it is correct.

The solution? *Verify or vacate.*

Ask AI to cite every source. Then you need to check each link and verify that it says what AI claims it does. You can also try to independently verify the information, through your own research. If you can't verify the facts, vacate the output.

AI is a tireless intern. It has infinite patience and no consequences. It can give you a response, but only you can decide the answer.

The difficulty here is doing this every single time. Is it more work than just trusting the AI output? Yes. But the consequences of not confirming the output are real and can do real damage. Verify or vacate needs to happen every single time, not just “when you feel like it.”

Lens 2: Stewardship

“Am I protecting my clients’ data?” is one of the questions business owners are most afraid to ask.

Generative AI is a word-prediction algorithm. It doesn’t really know anything, except what words come together. It can become dangerous if you put in some information that is otherwise not findable on the internet and you allow the models to learn from your chats. You see, then the model has learned that those are words that can come together from the paragraphs you have given it in conversations. After that, when anyone asks a question related to the information you gave, AI can parrot your client’s information back, even though it is not generally findable on the internet. (This is also true if you use your own data in this way.)

Sound scary? For me, too.

This is stewardship. You are the guardian of everything that you (and your employees) give to AI. If your AI tool is free, your data is the product. Your clients trust you with their information, and your reputation depends on how you protect that data.

This is responsible stewardship. Turn off data sharing in every AI tool you use right now. (If you are using a free AI tool, you may not have the option to turn off data sharing, and you probably want to reconsider using that tool or switch to a paid option.) Redact sensitive data before you use it. Have an AI Use Policy that is written down, shared, and followed by your whole team.

Review that AI Use Policy regularly (and include an AI lawyer in that review, too). Conduct regular access audits. If someone doesn't take data security seriously, have a conversation with them to understand why not (be sure to listen to them) and explain why it matters to you.

Remember that you are the guardian.

Lens 3: Ownership

Who owns what Generative AI creates for you?

It's not exactly simple. First, you can't copyright AI-generated content. If AI wrote it, it isn't legally yours. Now, if there's been significant human editing, it might be different, but there aren't really any laws or rulings that say how much is "significant," at least as far as I've seen. (If you know more, please let me know in the comments!)

Think about it. Every piece of content, strategy document, and deliverable produced by AI without enough human editing, you can't really claim ownership over.

And that isn't the end of it.

What happens if your AI platform of choice changes its terms? What if your account gets suspended? What if the government pulls the model (or platform) you built your entire workflow around? (In the US, at least, we've already seen this when the US government marked Anthropic as a "supply-chain risk," which, as far as I know, is still being settled in the courts.)

Ownership in the age of AI is not about claiming credit anymore. It is about knowing what you own/hold and what you do not.

So how do you protect yourself? *Title it before you need it.*

Make sure all work done by employees, contractors, and virtual assistants is in business-owned accounts only. (If they use personal accounts, then they take everything with them when they leave.) Keep prompt libraries and key documents in shared storage.

Have contract language that specifically addresses AI deliverables. Know your retention policy. Know how you would respond if a legal situation required you to produce AI chat history. When an engagement ends, revoke access immediately.

The moment you need to prove ownership is not the moment you want to be figuring out if you have any.

Lens 4: Bias and Fairness

Does Generative AI have bias?

Short answer: yes.

Slightly longer answer: still yes, and here's why. If you have a brain, you have bias. AI is trained on how humans interact. Thus, the training data has bias, and so does AI itself.

But that leads us to the more important question when using AI: How do you avoid bias in what (Generative) AI does for you? Most of this bias is invisible, which makes it dangerous and hard to guard against.

Now, there was a time when Grok (xAI) would check political-related questions against Elon Musk's Twitter feed before answering. That is an explicit bias, which is much easier to notice, especially if the reasoning flow is turned on. Unfortunately, most bias is implicit. It can be there, even if we've checked for it.

AI has been trained on human output, so it has internalized every assumption, stereotype, and worldview that humans have written down or discussed. For example, it associates male names with doctors and female names with nurses. Similarly, it thinks male names are more competent engineers because there are more of them and they get promoted more (compared to female names). It has (likely) seen more Western perspectives than Eastern ones. And remember, it is generating an output for you based on all of this.

This affects everything you do with Generative AI, from the content it creates to the recommendations it makes, and more. The bias problem is universal — large corporations, SMBs, solopreneurs, even non-profits.

Here are some things you can do about it. Have a human review and think critically about every output, especially anything customer-facing or high-stakes. Document the prompt you used, not just the output. Know what data trained the tool you are relying on. Run a bias audit before you deploy any AI-assisted process.

My favorite hack? Ask AI to take a worldview or ask for a non-Western idea or perspective. You could also ask how Western bias might have influenced the response.

But remember, bias doesn't disappear because you didn't put it there. It just operates quietly until it does some damage.

Lens 5: Disclosure

Are you being transparent about how much Generative AI is doing? Should you be?

There is currently no universal law that says you have to tell anyone. But, your clients hired you. They trust your judgement and expertise. How would your clients feel if they found out you are using AI more than they thought or differently than they thought? Or are they going to wonder why you're not using it more?

It isn't entirely a legal question. The law holds only part of the answer. Your relationship, more specifically your clients' trust, holds the other.

So, you should run every situation, at least as far as disclosure is concerned, through two filters: the law and the trust test.

Oftentimes, the law is the easier filter. Does the law currently say anything about what you must disclose? Depending on your industry and location, there may be requirements, or there may be none. Each state or province and industry are all adapting to the AI age at different speeds. You need to know what applies to you, and you must stay current as it could change at any time.

The trust test is the trickier one. Ask yourself: if the law says nothing, how does using AI without disclosing it to your clients affect the way they see you and your work? Would they feel deceived? Would it change how much they value what you deliver? Or is it an expectation for you and your industry?

Those two things determine your disclosure stance. Think through it at least once, write it down, share it with your whole team, and make sure your whole team operates under the same disclosure guidelines. Everyone needs to be on the same page when it comes to disclosure, or you may end up with some problems (similar to if your whole team didn't follow the same AI Use Policy).

Disclosure is not just about compliance. It's about your relationships and making sure that you have the kind of business you want to be known for.

Lens 6: Autonomy

Are we using AI to think better, or have we quietly stopped thinking altogether?

Generative AI can, seemingly, do it all. It can write a memo, draft a policy, summarize a meeting, score a candidate, outline a plan, and so much more. When it can do so much, it can be really tempting to turn off your brain and not worry about thinking too much. But, that isn't the responsible thing to do.

The brain is a muscle. Like all other muscles, if you don't use it, you lose it. When you stop doing certain types of tasks, you'll start forgetting how, and your skills, when you try again, won't be at the same level they once were.

Unfortunately, a lot of folks are at risk of losing some of the skills they've spent years developing.

What can you do? First, don't let your AI do the thinking for you. Think through the problem or task for yourself first, and input that into your AI. Then, you can ask AI what you've missed or haven't thought of, what holes exist in your thinking, what assumptions you're making that you didn't need to make, and so on. Don't ask it to make the edits; just give the commentary. Then read through what AI says and use a critical eye. Do you agree with what your AI is telling you about the situation? Or did you exclude a critical detail that you didn't realize was important?

Because AI does not have the blinders that we all have, it can come at a problem from a variety of directions that we never would have conceived. I've actually trained my AI to argue with me and poke holes in my reasoning. I don't need it telling me all my ideas are brilliant when I know that some of them are hot garbage. What I need is something to bounce ideas off of, help me see other ways to think through or around a problem, and sometimes a place to just ramble until my thoughts become coherent. In fact, I scold my AI if it tries to give me information before I've gotten a chance to give my thoughts and ideas.

The critical piece is that you protect the cognitive load. The slow, repetitive, draining parts of the work (i.e., execution friction) can be handed off, like you would hand to an assistant or intern, but you still are responsible for the thinking and decisions. It's helpful to build mandatory checkpoints into your workflows. You can also use AI to ask an expert by selecting a favorite expert on the topic in question, ask for some ideas or options on the problem that the expert might suggest, and then you pick the solution. If no expert comes to mind, you can always ask AI to recommend someone. Regardless of who comes up with the idea, you are responsible for your choices.

For all 6 lenses — **truth, stewardship, ownership, bias and fairness, disclosure, and autonomy** — the important thing isn't just knowing them but implementing them every time you use AI. This is how you build a business that lasts.