

Siara (00:00)

10% of US adults said that they use an AI chatbot for emotional support or advice, 4% said that they use them for companionship. And I know that that number is

AI creep is the gradual normalization of relying on AI for parts of your life you once never would have considered. The boundary moves so incrementally that you barely notice the shift.

I think that you should decide.

What are your boundaries with AI are before you're asked to cross them?

today I want to make that a lot easier for you. I want to walk through six boundaries that I've set with AI that I think you should consider.

My first rule is that

Hey y'all, happy Thursday. I have been thinking about a term that I heard the other day. You've probably heard of the term lifestyle creep.

It's the phenomenon where as your income rises, things that once felt like luxuries gradually start to feel like necessities, and you don't really feel that extra income that you've earned.

There's a version of this in technology that they are calling AI creep. AI creep is what happens when AI slowly makes its way into more and more parts of your life, eventually leading to an over reliance on AI and potentially negative impacts on your life. First, you might use it to help you write something. Then you may use it to organize your schedule. You decide to give it access to your calendar, your email, your files. Seemingly not a big deal. But then you find yourself considering sharing your health.

data with the tool. and all of a sudden, you're telling it about the intimate conversation you had with your best friend. Over time, things you once would have considered invasive or unnecessary start to feel completely normal.

Four years ago, you might have thought it would be insane for me to connect my health data to an LLM. But by 2026, you've been using the app for years, you've already given it access to other parts of your life, and when the new feature gets presented to you as a convenient and potentially useful tool, suddenly it just doesn't seem so crazy anymore. The consequences, of course, haven't disappeared. If anything, there's more of them now than there was before. But for some reason, it just doesn't feel like a big deal.

AI creep is the gradual normalization of relying on AI for parts of your life you once never would have considered. The boundary moves so incrementally that you barely notice the shift. I believe that the easiest moment to decide whether you're comfortable with a new technology is before the technology is presented to you.

A while ago, I started keeping a running list of things I never wanted to do with AI or technology in general. Some items came from mistakes I'd already made and wanted to roll back. Other things were things that I could see coming and wanted to make a decision about ahead of time before the tool could make the decision for me.

The other day, for example, ChatGPT prompted me to connect my financial accounts. This was an easy no for me, but I know that many other users had to make the same decision and it might not have been so easy for them.

It was easy for me to make because I had already written down a long time ago that I was not comfortable with sharing my banking statements with an LLM or another AI tool. At least not yet. Especially with a tool like ChatGPT getting into advertising. I know that giving it my purchasing data would expose me to highly targeted ads based on my purchasing decisions. Imagine a company learns that I stopped paying for a subscription and then they never leave me alone with ads. Or imagine I start ordering puppy food.

And so the app starts pushing puppy toy ads my way or something. I don't know. I just didn't want any of that. I also know that LLMs don't have the perfect cybersecurity record, and so I just didn't want to create a new vulnerability for myself. Clean, simple, and easy.

So here is my bid to you. I think that you should have some sort of AI constitution. You should decide.

What are your boundaries with AI are before you're asked to cross them? Because once a technology is deeply embedded in your life, rolling it back can be really difficult. So today I want to make that a lot easier for you. I want to walk through six boundaries that I've set with AI that I think you should consider. Some of these boundaries are things that I've learned from experience, some are lines that I've always known that I don't want to cross. And you know, over time it may get longer, it may get stricter, it may get looser. All I know is that I

I want to write my AI constitution down now while I still have the luxury of deciding what role I want this technology to play in my life. So here are the boundaries I'm setting with AI before convenience decides them for me.

My first rule is that I will not give AI blanket access to my digital life. I believe that every single AI permission needs to earn its place one at a time and for a very specific reason. I plan to be extremely intentional with what I grant AI access to.

I think AI is going to ask you to connect something new every week until eventually your digital life and non digital life is fully immersed. And some people honestly love the idea of this. I get it, it's kind of cool, like the idea of it, so go off if that's your thing. But I have heard too many horror stories and learned too many things about these companies to ever let that be me. Just

because connecting another service makes AI more useful does not mean the trade off for your privacy, autonomy, or

Safety is automatically worth it. There is an opportunity cost to every connection that you make, and I truly believe that. For me, things like my health data and my financial services are not worth risking in an LLM. But I think that this list of non-negotiables would differ from person to person. I think it will highly depend on your life and your values and what is most sensitive for you. For instance, if I were LeBron James,

I would not be putting like my training schedule or my diet into an LLM because that feels like proprietary information that he wouldn't necessarily want to leak. Or if I was

As a politician, I know connecting work email it's it's honestly probably not even an option, it's probably illegal because even they know that's not a secure place for your communication. So I would adjust this boundary based on what's important to you. I would also periodically revoke access to systems that the AI is not currently using. Like let's say you granted something in the past because it was relevant for a temporary project, but now it's not being used, or maybe you tried to make a data.

source useful to you, but it just didn't really work, I would just definitely disconnect things that aren't in use or not bringing you value right now.

My second rule is that I will not outsource my judgment or my responsibility to

I want to be very careful about how I use AI at work or when I'm making life decisions. I don't want AI to replace my judgment at work. And this is more than just about my integrity and my quality of work, which is very important to me, but it's also about preserving my mental agency and my ability to think.

There's growing research that shows when we offload too much thinking to AI, our cognitive engagement actually drops. And so over time that can weaken skills like critical thinking, problem solving, memory recall. And so I don't want that to happen to me, and I really don't want that to happen to you.

While I do use AI at work, and a lot of us do, it's always important to remember that AI does have limits. So I always like to treat AI kind of like an intern. Now, it could be a really great intern or it could be a really terrible intern. But at the end of the day, like an intern, it is going to be limited in its understanding and in its experience. An intern, for instance, may have access to all of the company knowledge and all of the public data in the world. The intern may be great at repeatable, monotonous tasks.

But you better make sure your instructions to that intern are perfectly comprehensive and understand that they are not quite going to be able to connect the dots the way that an expert can. AI is exactly the same way. It has limits.

that's not to say that AI won't enhance the work, but it's not gonna have the same experience, value, or empathy of a human. Soft skills and nuance are everything. So to me, human work and work completed by AI are just not equal.

I have a rule that nothing unseen, unreviewed, or untouched by a human should be seen by anyone but the person who prompted it. because if you think about it, it's actually quite rude to share untouched AI output with a colleague or a team member, and definitely the public, because you're sort of outsourcing the work of sifting through slop to someone else, which is not very considerate.

another good reminder is that AI doesn't just have a chance of hallucinating, it will hallucinate. Hallucination is an inherent failure mode of current generative models, so you should build any AI workflows you have around the assumption that mistakes will happen rather than treating them like freak accidents. Even if it's not at work and you're simply making a judgment call about your life or someone else's life, if the answer matters enough that being wrong would meaningfully affect you or that person, you need to verify it elsewhere.

I can share an example of this. The other day I was helping to solve a problem for a friend at an airport and I asked Chat GPT. It was a very helpful answer. It was actually an accurate answer. But ultimately, I did call TSA to verify the claim because a hallucination in this case would have resulted in that person potentially missing their flight. Imagine I tell someone the wrong information, now they're looking at me like I just missed my flight because AI hallucinated at you. That's not gonna be me.

I think a good rule of thumb here is that AI should be more of an advisor than a decision maker if your name is on the work.

Okay, my third rule, which I think a lot of people will disagree with me on this one, but that's okay because these are your boundaries to set.

I will not let AI record my life with always on recording tools. So that means AI or any tool will never automatically have access to all of my private conversations. No matter how convenient it becomes to record or transcribe or summarize every interaction that I have, I simply will not be doing that. I'm noticing a huge trend right now of these always-on recording tools being pushed. And I just think it's important to remember that there is an enormous

Business incentive to this market. A random conversation between people gets remembered in their head and then it kind of disappears, right? But a recorded conversation becomes another group of data points about your relationships, your preferences, your work, your business, your habits, your plans, your interests, your feelings, your thoughts. And if you stack thousands of those files together, you will have an extraordinarily detailed model of a person, and I just think that's really dangerous.

One that I keep hearing about is called Pocket. They actually have great marketing, so I'll give them credit for that. But they brand themselves as an always-on AI personal assistant. So it's hardware will record phone calls and in-person conversations, it'll create transcripts and summaries. It'll let you search across your previous conversations. But I think that it's

Kind of weird. Like I think I would be weirded out by someone who I was speaking with and who randomly had their like pocket out. There's also Plaud's notepin. This is another wearable recorder that you can wear, but it could be like a pin or a necklace or a wristband or a clip. And so I think an interesting wrinkle there is that these devices are becoming increasingly difficult to detect, which I find weird. there's limitless pendant, their thing is like remember what you say throughout the day.

So that could be meetings, impromptu conversations, personal thoughts. Meta actually acquired Limitless in December of 2025. This is kind of like not surprising because they've gotten into the whole AI-enabled wearable thing. They're not selling to new customers right now, but I feel like that acquisition tells you a little bit about where the market is heading. And then of course there's MetaGlasses, which we're already seeing an issue with people recording interactions secretly.

So I think that these tools are a huge privacy risk to me, but they're also a massive privacy risk for every person around me who did not consent to having their conversation ingested by AI. And then also, can you imagine using any of these products and then them getting leaked? I mean, that is a major invasion of privacy. actually.

I

actually think we're gonna see a lot more stories about public figures conversations getting leaked and it's gonna get messy. So I'm watching out for that.

but yeah I think inherently this is also a bit of a risk to humanity because for every conversation you have to become searchable data. It just sounds like a recipe for disaster. I don't want every bad joke or hot take or vulnerable conversation I have to become a file somewhere. It's just like a dystopian nightmare to me. So I'm opting out of that trend for sure.

My next rule is that I will never consent to or leverage an AI clone of myself.

I have watched enough Black Mirror episodes to know that this concept makes me deeply uncomfortable. I don't want a chatbot with my likeness, I don't want an avatar speaking as me. I wouldn't even want an AI agent sending messages on my behalf without me reviewing them first.

If someone believes that they're communicating with me, I think the words they receive should have actually passed through me. I don't believe that there is actually a right amount of data that

can exist that can accurately duplicate me or clone me as a human. I have free will, not to get all philosophical, but my opinion or my response.

or my mood may change with the way that the wind blows. And so I do not trust that any collection of code can accurately portray me. And so this is just a boundary for me.

and if you think this idea is far-fetched, it's not. There's a company called Hey Gen that literally sells digital twins that can look like you, sound like you, and use real-time avatar technology to generate videos of you without you actually recording them. A company called Delphi goes even further. It lets you create what is called a "Digital You" that learns from your writing, your content, and other data. It speaks in your voice, it can actually take calls, answer questions, and interact with people on your behalf. Reed Hoffman.

The co-founder of LinkedIn already has an AI twin called Reed AI. It has talked to reporters, appeared in interviews, it judged a Stanford business competition, it helps deliver speeches in multiple languages. And so, you know, there's a degree of that that's cool, there's a degree of that that's kind of uncool. If that's your thing, go off. Tell me that I'm communicating with your digital twin. If so, I'd like to know if I'm not talking to a human. But for me, this is just a

Boundary that I'm not willing to cross for myself. AI can help me communicate, but I will not allow it to impersonate me.

My next rule is that I will not replace human relationships with AI. I refuse to build emotional connections or intimacy with an AI. And I will, as much as I can control it, not replace human connection with AI. Yes, that covers any form of companionship, of course. But also coaches, doctors, therapists, and anything of that nature. Again, as much as I can control. And if you think I'm being dramatic, this is already.

Much more mainstream than I think people realize. In Pew's February 2026 survey, 10% of US adults said that they use an AI chatbot for emotional support or advice, 4% said that they use them for companionship. And among 18 to 29-year-olds, that jumps to about 20% for emotional support and 7% for companionship. And I know that that number is growing. I never want to condition myself to reach for a machine instead of another human when I need connection. I like humans, I'm gonna stick with them. I've seen a lot of

Companies attempt to create AI companions as replacement for human companions. There's this company called Friend that you might have heard of. They sell an AI wearable that you can wear and talk to and interact with 24-7.

Another example that I was kind of surprised by that was announced recently was an AI life coach. It's actually meant to represent the likeness of Mark Manson. You might have heard of his book, *The Subtle Art of Not Giving a F*. So Mark Manson's new company, Purpose, is an AI coach designed to learn users' patterns, values, needs, relationships, and goals, and provide

personalized guidance on demand. His own description says it is intended to understand users quickly and challenge or guide them over time.

You know, this may be useful to someone, but my personal preference is that if I wanted Mark Manson's perspective, I would rather read what he said and interpret it myself or speak with him directly than develop a relationship with software that's pretending to be him. I really believe that in something as delicate and sensitive as an emotional connection, a human is always going to outperform an AI. So that's where I'm at there.

Okay, my last rule, rule number six, that I feel very strongly about, is that I will never willingly support synthetic artists.

AI slop is not art. A bot cannot be an artist.

So for me, if I know an artist is synthetic, I am not interested. This applies to music, visual art, movies, books, creators, influencers, actors, whatever other synthetic entertainment gets invented.

I am out. Now, of course, we can only opt out of the content that we know of. So, for instance, if Spotify were to sneak AI music into a playlist and you didn't know, there's nothing you can really do about that. If there's an AI actor in a film that you don't notice, there's nothing that you can really do about that. I would hope that they would communicate that, but I think there's probably already been many instances where they don't. I just believe that once you know that the artist is an AI persona.

You can make a point to direct your attention or your money towards a human instead.

For me, this isn't really about

Whether AI is technically capable of making something that sounds or looks good, I would just genuinely rather spend my time and resources supporting human creativity. I believe that this is an attempt to replace humans in a function that didn't need to be replaced, and I would argue can't really be replaced. I really think that art is precious and we shouldn't take it for granted or outsource it because we might begin to forget what art really is. Not to get too deep, but I just really do feel that way.

so that's my list. in most of these rules I'm really just trying to stay as connected to my humanity and to others as much as I can. I know AI will increasingly offer to take over tasks and functions, but I believe everyone should ask themselves if that specific use case has much value beyond efficiency. Because sometimes human effort is kind of the point.

I'll probably revisit this list every year. And I do reserve the right to adjust them as I go because technology's gonna change and my understanding of the risks will change. like I mentioned,

they might get more strict, they might get less strict, we'll see, but I do think that we should always be checking ourselves when it comes to AI use.

Okay, so moving on, I do want to share a few quick news bits. The first one pertains to rule number six. So if you do happen to want to avoid synthetic art, you will want to note this one. Spotify is working on a feature that would let paying users create licensed AI-generated covers and remixes of existing songs.

Artists do need to opt in to letting their songs be remixed, which I think is a nice safeguard, and they will get compensated, which is an even better one.

Spotify also has AI credits in beta that can disclose whether AI was used for the vocals, the lyrics, the instrumentation, or the production. So if you want to check whether a song used AI, you can open the track credits in Spotify and look for those AI disclosures. I think the important caveat here is that the AI credit system is optional and still limited. So a missing AI label does not automatically mean that no AI was used.

While I think it's interesting that Spotify will let users manipulate certain tracks with AI, I do think that the AI credits is one step towards transparency. There have been some sneaky AI artists that have come up in recent years, and people have not taken too kindly to that. So overall, I think this is good news.

Another story in AI transparency. Last week I explained the new EU AI transparency rules, which requires AI generated or manipulated text, images, audio, and video to be marked in a machine readable way so that it can be detected as AI generated.

Anthropic responded to this regulation by announcing that Claude models will now embed an invisible watermark directly into AI-generated text and generated files will carry signed metadata showing that Claude was involved. The watermark is designed to travel with the text, even when you copy and paste it,

says heavy rewriting or mixing it with other texts can make it harder to detect. But from what I can tell, this will make a major difference and hopefully make it easier for people to detect AI slop. They're applying this globally, not just to people using CLaude in Europe. So I think that this is also a net positive.

One more quick note for iOS users. There is a feature on iCloud that you might have heard of called Private Relay. It's Apple's privacy feature for Safari, and it's designed to hide your real IP address from websites and prevent any one party from seeing both who you are and what sites you visit. Researchers, unfortunately, recently found that under certain conditions, websites could still potentially expose a user's real IP address. So if you're using private relay, first make sure it's actually turned on under settings, but also note that



I would never really treat private relay as a full VPN. It mainly protects Safari browsing and certain network requests rather than all of the traffic coming from your device.

So, if you want broader protection across your iPhone, I would definitely invest in a reputable VPN that encrypts more of your internet traffic. I use NordVPN, for example, but there are many other well-established VPN providers that do the same thing.

That is all for this week, y'all. I'm really hoping that you'll consider adopting some of these AI boundaries for yourself. And if you have any that I didn't think of, please let me know. I'm so curious. but I really think that future you will be proud of yourself for restraining, and I genuinely believe you'll be better off and have the upper hand in life because of it.

most importantly, I hope you get a chance to log out this weekend. Bye.