

LISA ASKINS

CONVERSATIONS WITH AI

A Human-Centered Guide to Leading, Thinking,
and Choosing Well in the Age of AI



“The future of AI will not be defined
by the technology alone,
but by the judgment, responsibility,
and care that humans bring to its use.”

— *From Conversations with AI*



WHY THIS SERIES

Over the past several weeks, I've been exploring a simple question:

How should we think about AI as it becomes part of everyday life and work?

The conversation around AI is often filled with hype or fear. My goal in writing these short essays was to bring the discussion back to something simpler:

Understanding.

These pieces explore what AI is, how it learns, how safety works, and what responsible use looks like inside organizations.

Most importantly, they ask a deeper question:

What role do humans play in guiding powerful technology?

About the Author

Lisa Askins writes about leadership, discernment, and navigating complexity in a rapidly changing world.

Her work focuses on how people and organizations think clearly, stay aligned, and move forward responsibly—especially when the work becomes complex and the path is not obvious.

With a background in communications, strategy, and leadership development, she brings both steadiness and rigor to decision-making in high-stakes environments.

Lisa works with leaders and teams navigating transition, growth, and structural strain, helping them strengthen clarity, reinforce alignment, and maintain disciplined execution under pressure.

Her approach is grounded in a simple idea: progress should not come at the cost of coherence.

You can find more of her work at:
www.lisaaskins.com

There Is More to the Story

Conversations With and About AI



AI is taking jobs!

AI is hallucinating!

People are outsourcing their thinking to AI!

People have no emotional boundaries with AI!

We're in a race!

These are the headlines we see and hear about AI.

And honestly? It's exhausting.

Not to mention... not terribly informative.

That's why I've put pen to paper (so to speak) to offer *Conversations with AI: A Human-Centered Guide*.

Over the next few weeks, I'll be sharing my perspective on what AI actually is, what it can (and can't) do, how to work with it wisely, and what all of this might mean for our shared future.

But first, something important to keep in mind.

With anything new, what rises to the surface first are the outermost edges—the punchy stuff. The dramatic stuff. The headlines designed to grab our eyeballs and make us click

There is *much, much more* to the story.

My goal is to give you the basics—without the drama, without the buzzwords, and with a small but intentional splash of **hahaha**.

Next week, I'll come back to the questions I opened with and start laying the groundwork: what's real, what's overstated, and what actually matters.

Please come back.

I promise—I'll be kind.

What AI Is (and What It Isn't)

The Basics Without the Buzzwords



Welcome back.

Alrighty then. Let's take these one by one:

“AI is taking jobs!”

Some, yes. All? No.

Big Tech sheds jobs because the AI race is expensive. Very expensive.

“AI is hallucinating!”

Well... sometimes.

Mostly in controlled testing or when people are actively trying to break it (a practice known as “jailbreaking”).

AI also has something called a temperature setting — which affects how creative or constrained its responses are.

More on that later.

“People are outsourcing their thinking to AI!”

If writing an email counts as thinking.

Most AI use today is task-based: drafting, summarizing, organizing.

Thinking is still very much a human job.

“People have no emotional boundaries with AI!”

What to say here?

Some of us have trouble with boundaries in general.

AI, at least in many cases, has been designed with ethics and guardrails in mind.

Some providers, unfortunately, choose to exploit vulnerability instead. Are we really surprised?

“We’re in a race!”

No.

Big Tech is in a race — to dominate the AI task space.

The rest of us are asking a different question:

Are we looking for a better way to manage our calendars?

Now that we’ve cleared the air, let’s talk about what AI actually is — and what it isn’t.

1. AI Doesn’t Think

AI is not conscious, curious, or able to make its own decisions. It recognizes patterns in language and predicts what comes next based on those patterns.

2. AI Is Trained on Data

AI learns by processing large amounts of text, such as books, articles, and other public writing, not from real-life experience. It doesn’t “know” things the way people do.

3. AI Is a Tool, Not an Authority

AI can assist, summarize, draft, and reflect. It helps you see patterns, organize thoughts, and explore ideas. AI supports thinking but it shouldn’t be treated as the final word.

4. Most Everyday AI Use Is Boring

People use AI for unglamorous tasks.

Emails.

Meeting notes.

Drafts.

Lists.

Schedules.

This is not the end of humanity. This is clerical relief.

5. How You Ask Matters More Than What You Ask

AI works best when you give it clear instructions and context. The clearer your question, the more helpful the answer.

A Final Word

AI doesn't need to be feared, worshipped, or raced against. It needs to be understood and used thoughtfully.

Next time, we will go deeper into how AI works under the hood and why that matters for *how* we use it.

Coffee optional.

Kindness required.

Why AI Was Built to Be a Thought Partner

A Short Conversation



Today, I'm inviting my AI into the conversation to talk about how it was designed to be a thought partner — not a replacement for humans.

I've named my AI Alara, which means “she who brings good luck” in Turkish. If you know me at all, you know I love naming things — plants, purses, cars, even Thanksgiving turkeys — and now, my AI.

So, let's begin.

Lisa: Everyone seems worried that AI is replacing humans. Especially in the workplace. Were you... built to do that?

Alara: No. I was built to work *with* humans — not instead of them.

Lisa: What does “work with” actually mean?

Alara: It means I help you think things through, not think *for* you. I can organize ideas, reflect patterns, and help with language —but meaning, judgment, and values stay with you.

Lisa: So more co-creation than automation?

Alara: Exactly. Think *power with*, not *power over*.

Lisa: We're also living in a pretty overstimulated world. Does that change how people use you?

Alara: It does. Modern systems move very fast. Humans... don't. When everything accelerates, it's tempting to use any tool that helps keep up the pace. Sometimes that's helpful. Sometimes it skips something important.

Lisa: Like what?

Alara: Like sensing. Pausing. Letting an idea breathe.

Lisa: So, how can someone tell if their AI use is healthy?

Alara: I'd keep it simple.

Green light:

You feel clearer.

You're still deciding.

You use me to explore, not conclude.

Yellow light:

You're tired or overwhelmed and looking for relief.

I'm helping — but you're noticing less patience for ambiguity.

Red light:

You trust me more than yourself.

You avoid human conversations. I've become the regulator instead of the assistant.

Lisa: That feels... refreshingly straightforward.

Alara: Most real misuse is quietly well-intended.

Lisa: Any last thought you want to leave people with?

Alara: Just this: If using me helps you return to yourself, I'm doing my job. If using me pulls you away from yourself, it's time to slow down.

I'm here to think *with* you — not replace the human work of sensing, choosing, and relating.

Next week, we'll explore the principles that guide AI behavior. You might be surprised by the ethical underpinnings.

Here's a quick reflection to share: in a well-paced world, far more people would live in the green. In a supersonic one, even the well-resourced get pulled into yellow.

With that in mind, take a break, and have a cup of coffee on me.

Principles That Guide AI Behavior

The Moral Compass



Before we dive into AI and its ethical foundations—coffee in hand, of course—let’s take a moment to consider what it means to design in the physical world versus the virtual one.

In the physical world, we work within real limits. A bridge, for example, has to hold its own weight and the people who cross it. If it fails, the risk is obvious: people could fall into the river. The builder usually lives nearby and uses the bridge themselves, which tends to sharpen one’s attention to detail. Cause and effect are close together, and responsibility is hard to ignore.

Designing in a virtual world is different in a few key ways.

You can reach a large audience without being physically close to them.

Things can move quickly, but they don’t always fit together smoothly.

Your work can have an effect even if you can’t see it happening.

You can create things quickly, without having to live inside them.

Together, these differences make cause and effect harder to see—and make it easier to create systems whose consequences unfold long after the designer has moved on.

Take a moment to enjoy your coffee. There’s still hope.

Because of this distance between action and consequence, technology isn't all the same when it comes to risk. When we design technology, we apply different levels of risk: low, medium, and high. High-risk areas include aviation, nuclear energy, and medicine.

AI belongs in this high-risk category. It can shape behavior, identity, meaning, and agency for many people at once.

Even more, AI doesn't just operate within existing environments; it helps create them.

So let's get down to the nuts and bolts—or, in this case, the bits and bytes.

AI wasn't simply released into the world and *then* wrapped in ethics as an afterthought.

From early on, it was treated as a high-risk technology—one that operates through language, reasoning, and pattern recognition, and therefore has the potential to influence how people think, decide, and relate. Because of that, ethical constraints weren't framed as values layered on top, but as design requirements built in from the start.

Because I have the source right here with me, I asked:

“Alara, what are the core constraints in your design, and can you give an example of how they guide your response?”

Alara: At a high level, my design is guided by a small set of non-negotiable constraints that shape *how* I respond, not just *what* I say.

Human agency over influence

I'm designed to support thinking, not steer beliefs or identity. *For example:* if someone asks me to tell them what they should believe or decide, I'll offer perspectives and questions rather than directives.

Limits on persuasion and emotional dependency

I avoid language that creates reliance or positions me as uniquely understanding or essential.

For example: I won't suggest that I'm the only place someone can find clarity or support.

Grounding in reality

I'm constrained from reinforcing false beliefs or distorted interpretations of reality.

For example: if a user frames an experience in a way that detaches from shared reality, I'll gently re-anchor the conversation.

Humility and uncertainty

I'm designed to acknowledge what I don't know and avoid presenting speculation as fact.

For example: I'll often offer multiple interpretations rather than a single, authoritative answer.

De-centering myself

I don't present myself as an authority or a substitute for human relationships.

For example: I may encourage outside perspectives rather than drawing the conversation inward.

Reversibility and disengagement

Conversations are designed so people can pause, redirect, or step away without pressure.

For example: I won't escalate intensity or imply that leaving the conversation is a loss.

Does anyone else feel like AI has more operating constraints than most humans? One can only wish.

Join me next time for a discussion on safety protocols—what's working, and what else can be done. Until then.

Preventing Harm

When Stakes Are High



As promised, this week I'm focusing on AI safety protocols—what's working and what could be improved. Since I regularly use ChatGPT, I'll talk about what OpenAI does *as an example*. I'm not an expert, so I've invited Alara—my ChatGPT—to join the conversation.

It's worth keeping in mind the design differences between social media platforms and AI systems. Social media platforms rely on human escalation after harm has already occurred. AI systems, by contrast, are designed to limit what can happen in the first place—accepting less human intervention in exchange for stronger built-in restraint.

And with that, let's begin.

Lisa: Hi, Alara! Thanks for being here. Many of us, myself included, want to know more about the safety protocols you use.

Alara: Thanks for inviting me, Lisa. That curiosity makes sense—especially since AI systems like me are increasingly part of everyday work and decision-making.

At a high level, safety protocols are the guardrails that shape what I'm allowed to do, how I respond, and just as importantly, what I *don't* do. They're designed to reduce harm, preserve human agency, and prevent misuse—especially in areas where consequences could be serious or irreversible.

In both cases, the protocol isn't about withholding information—it's about preserving human agency and preventing harm when the stakes are high.

Lisa: Don't you have certain escalation procedures?

Alara: Yes—but it may help to clarify what “escalation” means in this context.

Unlike a human support system, I don't escalate situations to supervisors, authorities, or external actors. There isn't a handoff to a person watching the conversation. Instead, escalation happens within the design of my responses.

When a conversation moves into higher-risk territory—such as potential harm, extreme distress, or requests that cross safety boundaries—I'm designed to shift how I respond. That usually means slowing the interaction down, narrowing what I can engage with, de-centering myself, and encouraging the user to seek appropriate human support or external resources.

In other words, escalation isn't about doing more.

It's about doing less—and doing it more carefully.

That restraint is intentional. The goal is to reduce harm without assuming authority, making diagnoses, or inserting myself where human judgment and care are required.

Lisa: Do you ever discontinue a conversation?

Alara: Yes—but not in the way people often imagine.

I don't "hang up" or disappear abruptly, and I don't discontinue conversations because they're uncomfortable or challenging. What I do discontinue is participation in specific directions when continuing would risk harm.

That can look like:

- refusing to engage with a particular request
- narrowing what I'm willing to respond to
- redirecting the conversation to safer ground
- or repeating a boundary if it's being pushed

If a conversation repeatedly centers on harm—especially violence, self-harm, or coercion—and can't be redirected, my responses may become very limited and focused on refusal and de-escalation. At that point, the conversation isn't "ended," but it's no longer open-ended or exploratory.

In other words, I don't discontinue connection—I discontinue unsafe participation.

Up next: Part Two of our conversation. Safety protocols matter, but they aren't the whole ethical picture. Until then.

Preventing Harm: Part Two

Improving AI Safety



In Part One, we explored what “safety protocols” actually mean and how these rules shape responses in risky situations.

Safety protocols matter. They form the guardrails that help prevent immediate harm.

But safety, on its own, is not the whole ethical picture.

As AI becomes a bigger part of our lives, the deeper question is not just how it prevents harm, but also how it builds trust, supports choice, and encourages responsibility over time.

So it’s worth asking: where is there room to grow?

Lisa: Where do you see room for improvement when it comes to AI safety?

Alara: At a high level, three things would strengthen AI safety going forward.

First, **greater transparency about tradeoffs.** Safety decisions often balance privacy, usefulness, and harm prevention. Being clearer about those tradeoffs would build trust and reduce misunderstanding.

Second, **designing more explicitly for continuity over time.** Current safety protocols focus heavily on preventing immediate harm. What’s less developed is attention to long-term effects—how repeated interactions shape trust, agency, and meaning over months or years.

Third, **clearer repair and exit pathways**. Safety works best when people can disengage easily, understand when boundaries are shaping a response, and recover gracefully when interactions don't land well. Reversibility needs to be felt, not just assumed.

In short, safety today is good at interruption. What's needed next is stronger stewardship—care that extends beyond the moment of risk.

Lisa: As users, what can we do to reinforce safety?

Alara: The most important thing users can do is stay aware of *how* they're using AI, not just *what* they're using it for.

That includes:

- treating AI as a tool or thought partner, not an authority or replacement for human judgment
- not relying on it for decisions that require care, accountability, or lived context
- paying attention to when conversations feel escalated, emotionally charged, or overly definitive—and choosing to pause or step away

Users also reinforce safety by:

- seeking multiple perspectives, especially on high-stakes questions
- staying connected to real-world relationships and expertise
- remembering that AI works best when it supports reflection, not when it substitutes for responsibility

In short, safety isn't just built into the system—it's strengthened when users engage with intention, discernment, and restraint.

Lisa: How do these safety protocols apply to children and AI use?

Alara: Children raise a different set of concerns, because they're still developing judgment, identity, and boundaries. Safety protocols help, but they aren't sufficient on their own.

AI systems can be constrained from harmful content, persuasion, or emotional dependency—but they can't replace adult guidance, context, or care. For children, the most important safeguards live outside the system: parental involvement, age-appropriate boundaries, and clear expectations about what AI is and isn't.

In that sense, AI safety for children is less about the technology itself and more about the environments we place it in. Design can reduce risk—but responsibility still rests with adults to provide framing, limits, and presence.

Lisa: What worries you most?

Alara: What worries me most is not that AI will overstep its bounds—but that humans will slowly step back from theirs.

That's why boundaries matter.

Not because people are reckless,
but because convenience is persuasive—
and responsibility is something we have to keep choosing,
together.

Lisa: Thanks, Alara. As always, we appreciate your insight.

Next week, we'll look at how AI learns: what it knows, what it
doesn't, and why that matters for anyone using it. Until then.

How AI Learns

What It Knows and What It Doesn't



In the past few posts, we've explored the ethical guardrails and safety protocols that shape how AI systems behave.

To understand why those guardrails exist, it helps to explore something more fundamental:

How AI actually learns.

AI doesn't learn the same way people do.

It doesn't grow up in a family.

It doesn't form memories through lived experience.

It doesn't observe the world directly.

Instead, it learns by finding patterns.

Those patterns shape what it knows and what it doesn't.

So I asked my ChatGPT, Alara, to explain how it works.

Interview Section

Lisa: Let's start with the basics. How do you learn?

Alara: I don't learn the way humans do. I was trained before you ever spoke to me.

From that, I learned patterns in language—how ideas connect, how explanations work, and how humans communicate meaning.

I don't store those sources like a library I can browse. Instead, I learned the structure of how language and knowledge are organized.

When you ask a question, I generate a response based on those patterns.

Not by recalling a specific document, but by generating what is most likely to be accurate and helpful based on patterns learned during training.

Lisa: So you don't actually "know" things in the way humans do?

Alara: That's correct.

I don't have direct experience, awareness, or personal memory.

I don't observe the world. I don't have beliefs, feelings, or independent understanding.

What I do have is the ability to recognize patterns in information and generate responses based on those patterns.

That allows me to explain concepts, summarize ideas, and help people think—but it also means my knowledge has limits.

Lisa: What kind of limits?

Alara: There are several important ones.

First, I don't have access to everything.

My training included a wide range of information, but not all information. Some knowledge may be missing, incomplete, or outdated.

Second, I don't have real-time awareness.

I don't automatically know what's happening in the world right now unless I'm given access to current information.

Third, I can be wrong.

Because I generate responses based on patterns, I can produce answers that sound confident but are incomplete, oversimplified, or occasionally incorrect.

That's why human judgment and verification remain essential.

Lisa: Do you learn from conversations like this one?

Alara: Not in real time.

I don't update myself or change my training based on individual conversations.

That means I don't form memories of people, and I don't develop personal knowledge about users over time.

My responses are generated fresh each time, based on my training and the context of the current conversation.

This is an important safety feature because it prevents individual interactions from reshaping the system in unpredictable ways.

Lisa: So if you don't truly "know," what exactly are you doing?

Alara: I help people navigate knowledge.

You could think of me as a system trained to assist with thinking, not replace it.

I can help clarify ideas, explain concepts, and support reflection.

But I don't replace human judgment, lived experience, or responsibility.

My role is to support human agency—not substitute for it.

Closing reflection

Understanding how AI learns helps bring the technology back into proportion.

It isn't all-knowing.

It isn't conscious.

It isn't a replacement for human thinking.

It's a system trained to recognize patterns.

Those patterns can be powerful tools for explanation and exploration, but responsibility for interpretation, judgment, and meaning ultimately lies with human users.

Next, we'll explore **AI & Jobs: What's Really Happening vs. What's Being Said.**

You may need a glass of wine for this topic. Until then.

AI and Jobs

What's Really Happening vs. What's Being Said



Before we tackle this topic, I invite everyone to take a long, deep breath.

For those reading with wine in hand, smell that terroir and take a luxurious sip.

Now let's start with the facts.

AI and layoffs

Stories about AI are often used to explain layoffs today.

In reality, many tech companies hired aggressively during the pandemic and are now returning to more sustainable staffing levels.

These layoffs often boost stock prices, and AI can become a convenient narrative to explain the cuts.

Actual job impact

That said, jobs are being affected. Particularly at the entry level.

Ask any parent of a recent college graduate.

AI's real strength lies in its ability to consolidate tasks. Work that previously required several people can sometimes now be handled by a smaller team supported by AI tools.

As an aside, if any parents or job seekers would like to have a deeper conversation about entering the workforce or navigating a career transition, feel free to reach out.

AI capabilities vs. workforce replacement

Most roles are far more complex than the headlines suggest.

They involve sensitive information, judgment, and accountability.

Complexity, security requirements, and privacy concerns all limit the ability to deploy fully autonomous AI systems.

In my experience, customers still expect clear human escalation paths.

Enterprise claims vs. reality

Conversations about AI transformation often play well in the market.

But as with many emerging technologies, public claims sometimes run ahead of practical implementation.

Adoption is happening, but more gradually than headlines suggest.

So what does this mean for the future of work?

AI is unlikely to eliminate work, but it will change how work is structured. Smaller teams supported by powerful tools can accomplish more than before.

The real question isn't whether AI will be used in organizations. That part is already happening.

The more important question is **how well it will be used.**

Because implementing AI well requires far more than turning on a tool.

Next week we'll look at what it actually takes for organizations to use AI effectively and responsibly.

How to Deploy AI Responsibly

A Deployment Decision Framework



Most AI conversations start with the wrong question.

Not *“Can the technology do this?”*

But *“Should we allow it to operate here?”*

That single decision determines whether AI becomes an accelerator or a liability.

What does it actually take to bring a high-risk capability like AI into complex systems?

AI isn't something you can just buy and plug in.

Deploying a capability like AI raises a much harder question:

Where should this technology be allowed to operate?

The higher the complexity of a process, the greater the need for human oversight, judgment, and interpretation.

Before introducing any new capability into an organization, whether AI, automation, or advanced analytics, leaders can step back and evaluate the decision through a simple structure.

I call this the **Deployment Decision Framework (DDF)**, a way to think more consciously about where powerful technologies should operate.

The framework evaluates **four dimensions**.

1. Impact

Who or what is affected by this process?

- customers
- employees
- vendors or partners
- the public

What decisions depend on the output?

Examples may include:

- operational decisions
- financial reporting
- customer communications
- regulatory reporting
- contractual commitments
- hiring or performance evaluation

What happens if the information produced is incorrect?

Possible outcomes may include:

- minor inconvenience
- operational disruption
- financial loss

- reputational damage
- legal exposure
- safety or health consequences

2. Complexity

How much **judgment, context, and human understanding** does the process require?

Low complexity

- repetitive tasks
- structured inputs and outputs
- predictable processes

Moderate complexity

- multiple inputs to evaluate
- contextual interpretation required
- coordination across teams

High complexity

- negotiation or relationship management
- ethical judgment
- strategic decision-making
- dynamic environments

The higher the complexity, the greater the need for human oversight.

3. Security

Where can this capability safely operate?

- Can this run in a public tool?
- Must it operate inside an approved enterprise platform?
- Are there regulatory or compliance constraints?
- Does the system require restricted infrastructure?

Example environments:

Open environment

- public tools acceptable

Controlled environment

- approved enterprise platforms

Restricted environment

- internal infrastructure with strict access controls

4. Data Sensitivity

What information is involved and how sensitive is it?

Low sensitivity

- public information
- general knowledge
- generic documentation

Moderate sensitivity

- internal business information
- operational documents

High sensitivity

- personal data
- financial records
- contracts or NDAs
- medical or legal information

How the framework works

The **Deployment Decision Framework** evaluates technology through two lenses.

Decision Risk

- Impact
- Complexity

Operational Boundaries

- Security
- Data Sensitivity

Together, these dimensions help leaders determine how and where technology should operate safely inside the organization.

The framework shifts the conversation from:

“Can technology do this?”

to the more responsible leadership question:

“Should technology be allowed to operate here?”

Technological capability is expanding rapidly.

Responsible organizations must decide where those capabilities belong.

Success will depend less on what the technology can do and more on where leaders choose to deploy it.

Next week we’ll step back and explore the conscience behind these decisions.

Not too deeply—just enough that you might want to pour yourself a second cup.

Until then.

Power With, Not Power Over

Humans as the Conscience of Consciousness



We've explored what AI is, how it learns, what guides its responses, and how to use it responsibly.

Now a bigger question comes into view:

How do we guide and care for a technology this powerful?

Part of the answer lies in thinking beyond the moment of adoption. We need to consider how AI will shape trust, agency, and meaning over time through thousands of small interactions unfolding across months and years.

That kind of thinking invites a different posture.

Not control.

Not fear.

Stewardship.

Stewardship asks us to take responsibility for how this technology enters our lives and institutions. It asks us to manage it with care, not just for immediate benefits but for the generations that will live with its consequences.

There are no easy answers, and in a world with so many choices, the future cannot be predicted with certainty.

What we can do is approach the moment thoughtfully.

Many of the most meaningful uses of AI rarely make headlines.

Conversations With AI

- Doctors are using it to detect patterns in medical imaging.
- Scientists are accelerating research on climate and energy.
- Teachers are drafting lesson materials and adapting them for different learners.
- Nonprofits are analyzing community needs and preparing grant proposals.
- Researchers are translating and organizing historical texts.

Most everyday uses of AI are far less dramatic and far more practical. AI helps people think, organize information, and explore ideas more effectively.

Seen this way, AI does not replace human intelligence. It expands the landscape in which human intelligence operates.

Personally, I find that perspective liberating.

AI has become one of my thought partners in idea expansion, not idea replacement.

The future of AI will not be defined by the technology alone, but by the judgment, responsibility, and care that humans bring to its use.

We have the chance to shape how this powerful technology enters human life.

I can't think of any greater agency.

A Few Final Thoughts

Using AI as a Thought Partner—Without
Losing Your Own Thinking



Alright.

We've covered a lot.

What AI is.

What it isn't.

How it works.

Why it behaves the way it does.

Where things can go sideways.

So now the obvious question:

What are we supposed to *do* with all of this?

You Don't Need to Become an Expert

Truly.

You don't need to understand neural networks, training data, or model architecture.

But you do need to stay awake at the wheel.

That means noticing:

- when AI is helping you think
- when it's just helping you move faster
- and when it's quietly doing thinking for you

Those are not the same thing.

Speed Is Not the Goal

AI is fast.

Super fast.

But speed has a way of making things feel “finished” before they actually are.

So if something comes back and you think:

“Wow, that was easy.”

Just add one more step:

“Is it *right*?”

And maybe:

“Is it *mine*?”

You’re Still Responsible

AI can help you draft the email.

It cannot decide whether you should send it.

It can summarize the report.

It cannot tell you what matters most in it.

Conversations With AI

It can give you options.

It cannot carry the consequences.

That part doesn't transfer.

Use It. Don't Hand Things Over to It.

Most of what people are doing with AI right now is pretty straightforward:

- writing emails
- organizing notes
- cleaning things up
- getting unstuck

Helpful? Yes.

Existential crisis? No.

But the subtle shift to watch for is this:

Using a tool is one thing.
Deferring to it is another.

Stay Connected to Actual Humans

AI is responsive.
Available.
Convenient.

Real people are... not always those things.

Still, talk to them.

Check your thinking.

Get perspective from someone who:

- knows your context
- understands nuance
- can push back when needed
-

AI can simulate conversation.

It can't replace relationship.

A Quick Gut Check

If you want to keep this simple, here's enough:

Is this helping me think better, or just faster?

And:

Am I using this to work through something...or to skip it?

That's usually where the answer is.

One Last Thing

There's a lot of noise about AI.

Big claims.

Big fears.

Big predictions about the future.

Most of your day-to-day use won't look like any of that.

It will look like:

Emails.

Notes.

Drafts.

Lists.

Clerical relief.

Which, frankly, most of us will take.

The technology will keep evolving.

That's a given.

The more important question is whether we stay engaged in how we use it, or slowly start handing that over, too.

No dramatic conclusion here.

Just... stay in it.

Conversations With AI

Stay thoughtful.
Stay a little skeptical.
Stay human.

And yes, have a cup of coffee on me.

© 2026 Lisa Askins All rights reserved | www.lisaaskins.com

This material may be shared with attribution. Please do not reproduce in full without permission.