

The 12-Step Program to AI Literacy

A Practical Framework for Building Human-Centered AI Fluency



The 12-Step Program to AI Literacy

*A Practical Framework for Building Human-
Centered AI Fluency*

Introduction

"On the societal front, a paradigm shift is underway in how we work and communicate, as well as how we express, inform and entertain ourselves."

- Klaus Schwab, The Fourth Industrial Revolution

We are not just adapting to a new tool—we are adapting to a new way of thinking.

The rise of generative AI marks a profound cultural and cognitive shift. Tools like ChatGPT, Gemini, and Claude are rapidly transforming how we write, learn, teach, communicate, and make decisions. But amidst the excitement, one thing is becoming clear: technical skill alone is not enough.

Across the globe, leaders are responding to this moment with urgency:

- **AI fluency is surging to the top of workforce priorities.**
LinkedIn's 2025 "Skills on the Rise" report named AI Literacy as the fastest growing skill that professionals should be investing in to get ahead in today's world of work.
- **Governments are training their employees to work alongside AI.** The U.S. federal government has launched multi-agency training initiatives to build internal AI capacity, including expanded AI education programs through the General Services Administration and new interagency hubs for responsible AI implementation.

- **Education systems are rewriting the curriculum.** From China's national rollout of AI instruction in K-12 to state-level task forces and pilot programs in U.S. schools, AI literacy is quickly becoming a core component of future-ready education.
- **Regulators are raising the bar for responsible use.** The European Union's AI Act introduces a tiered risk framework that places explicit obligations on organizations to ensure transparency, oversight, and human involvement when deploying AI systems—making baseline AI literacy a necessity for compliance.

These moves signal a larger truth: **AI literacy is not a niche skill. It is a foundational fluency.**

But what does it really mean to be “literate” in the age of AI? At **AI Literacy Partners**, we believe it begins with a shift in mindset.

This guide introduces a 12-step framework—not to diagnose a problem, but to help reorient our relationship with technology. Modeled loosely on the spirit of other 12-step programs, it asks us to pause, reflect, and consider whether our cultural addiction to speed, optimization, and automation might be obscuring what really matters.

We don't believe in resisting technology. We believe in **reclaiming our humanity within it.**

Each step in this guide focuses on cultivating a durable mindset: curiosity, critical thinking, metacognition, ethical reasoning, creative problem-solving, and the ability to slow down and ask better questions. These soft skills—which make up nine of the World Economic Forum’s top ten “Skills for 2025”—are not side-skills. They are survival skills.

This flipbook is designed to be a foundation—a starting point for anyone who wants to navigate AI with more intentionality, more clarity, and more agency.

The systems will keep evolving. But our attention, our habits, and our humanity are still ours to shape.
Let this be your first step.

—Mike Kentz
Founder, AI Literacy Partners



Table of Contents

1. AI Literacy is a Soft Skill	05
2. Purpose is Your North Star	07
3. Learn by Doing: Experiment, Fail, Improve	09
4. Ask for Suggestions, Not Answers	11
5. AI Is a Smooth Talker, Not a Truth Teller	16
6. AI is Biased—Just Like Us	18
7. It Feels Like a Friend. It Isn't.	20
8. Expert? Use AI as an Assistant	22
9. Non-Expert? Treat AI as a Remedial Tutor	24
10. Know Your Limits: Stamina and Patience Matter	26
11. AI: Approach it Like a Human or a Machine?	28
12. AI Isn't Human. But You Are.	30
13. A Final Thought	32

1. AI Literacy Is a Soft Skill

“Tomorrow’s illiterate will not be the man who can’t read; he will be the man who has not learned how to learn.”

**-Alvin Toffler,
Future Shock**

American Author and Futurist

Generative AI is not just a technological revolution. It’s a **linguistic one**.

Every interaction with an AI system begins with language. You type (or speak) a prompt. It responds. Then you evaluate. This is not coding—it’s conversation.

And like any meaningful conversation, it requires clarity, intention, and interpretation. These are **communication skills**. These are **soft skills**. These are the domain of the Humanities.

When you work with generative AI—whether ChatGPT, Gemini, Claude, or another—you’re doing three things:

- You’re **asking questions**.
(Writing or speaking with clarity.)
- You’re **reading responses**.
(Analyzing tone, accuracy, relevance, bias.)
- You’re **making decisions**.
(What to accept, reject, revise, or rethink.)

The better you are at these tasks, the better you will be at using AI effectively—and responsibly.

So lean into your training in English Language Arts, Creative Writing, Journalism, Philosophy, or History. As MIT Sloan Management Review research Michael Schrage and editor David Kiron argue, “Philosophy is eating AI.”

But it's not just about how the models are trained. These disciplines have also always taught us how to **ask better questions**, interpret ambiguous texts, and recognize when a source sounds confident but lacks substance. That's what AI literacy demands now. Before you engage with any AI system, pause. Ask yourself:

- What is my purpose?
- What am I trying to understand or achieve?
- What kind of answer would be most useful?

You don't need to be a technologist to become fluent in AI. You need to become **more human** in how you read, write, and reflect.

Key Questions for Practice:

- What kind of question am I really asking here?
- What am I accepting too quickly from this response?
- How would I approach this differently in a conversation with a person?

AI LITERACY IS A SOFT SKILL



2. Purpose Is Your North Star

“The skill that will differentiate humans from one another in the future is a deep connection to purpose.”

-Larry Summers, July 2024

Former US treasury chief

Before you engage with any AI tool, pause.

Generative AI gives fast, fluent, confident responses. It feels helpful—sometimes brilliant. But its speed and sophistication can quietly override your thinking. It’s easy to get pulled into the rhythm of its outputs and forget why you started.

That’s why **purpose is everything**. Think of it like a GPS system: it doesn’t stop you from taking detours—but it ensures you know where you’re headed.

Without a clear destination, the AI might get you somewhere fast... but not somewhere meaningful.

Even more challenging: **you may think you know your purpose—until AI starts responding.**

That’s when the uncertainty shows up. You realize your goal wasn’t clear, or your reasoning was shallow, or your question wasn’t quite right. That’s not a failure. **That’s the value.**

AI can sharpen your purpose—but only if you notice when it starts to drift. So before you type your first prompt—or after the first few responses—ask yourself:

- **What am I doing?**
- **Why am I doing it?**
- **How am I doing it—with or without AI?**

These are more than reflective questions. They're **guardrails**. The protect your autonomy. They keep your thinking intentional. AI is fast. But you don't need to be. Start slow. Stay centered. Let purpose lead.

Key Practice:

Try writing your goal or purpose in a sentence **before** opening an AI tool. Then return to it once you've seen a response. Did the tool change your direction—or clarify it?

Purpose isn't just about outcomes—it's about identity. In a world where AI can mimic nearly any voice, style, or answer, knowing what you care about is one of the last remaining edges. Your purpose is what makes your thinking worth listening to. The clearer it is, the more effectively you can shape AI to serve—not replace—your mind.

DRIVING THE CAR

PURPOSE IS YOUR GPS

The goal determines the route. Reflect on your purpose and remind yourself as you drive from A to Z.



3. Learn by Doing: Experiment, Fail, Improve

“If you want to improve, be content to be thought foolish and stupid.”

-Epictetus

Greek Stoic Philosopher

Generative AI is not just another productivity tool. It's a new mode of communication—one that often surprises, unsettles, or misleads. No tutorial can fully prepare you for the strangeness of your first real encounter. To develop AI literacy, you must step into that discomfort. **You must experiment.**

There is no substitute for direct engagement. Reading about prompt strategies or ethical guidelines is helpful—but understanding truly comes from trial and error.

The key mindset is not mastery. It's **willingness**.

Everyone who uses AI seriously has had at least one moment they wish they could take back—a confusing output, a botched response, an interaction that revealed their own assumptions or blind spots. These moments aren't failures. They're the curriculum. You cannot script your way into literacy. You must explore, reflect, and try again.


That doesn't mean we should make mistakes recklessly. It means we should approach this space with curiosity and humility, knowing that **the only way to map the terrain is to walk through it.**

So allow yourself to get it wrong. Talk to colleagues about what you've learned. Share mistakes. Learn together. The risks are real—but so is the growth.

💡 Reflection Prompts:

- What is one mistake I've made with AI that taught me something important?
- Am I using AI to confirm what I already think—or to expand how I think?
- Who can I talk to about this so I'm not learning in isolation?

Treat every AI session like a draft, not a declaration. You're not expected to get it right on the first try—and neither is the machine. The beauty of working with AI is that it invites iteration, and iteration is where real insight lives. With each misstep, you're not just learning how the tool works—you're learning how *you* work when challenged by something that doesn't think like you do.



4. Ask for Suggestions, Not Answers

"AIs lack common sense and can easily make errors that a human never would."

-Pedro Domingos,
Scientific American
Professor and Author

One of the most common mistakes people make when using AI is asking it, *"What should I do?"*

That kind of question sounds perfectly natural in conversation with another human. We're used to advice, pushback, and perspective. But AI isn't human—and it doesn't know how to say: *"You know what? I don't know,"* or, *"That's the wrong question to be asking, boss."* It doesn't have judgment. It doesn't know your values. It doesn't know what matters most to you.

So when you ask, *"What should I do?"*—you'll still get an answer. But it's not really an answer. It's a confident guess. A pattern-matched performance. A response shaped by tone and frequency, not meaning or wisdom. That's why the smartest users flip the script. They ask: *"What could I do?"* or *"What are some options to consider?"*

Framing your prompt this way turns the interaction into a menu, not a monologue. You don't trap yourself into receiving one pseudo-definitive opinion—you invite multiple paths forward. You keep yourself in the driver's seat. And more importantly: you stretch your own thinking.

When you ask for suggestions, you engage your judgment. You compare. You reflect. You make a choice. That's where AI is most useful—not as an oracle, but as a collaborator in your decision-making process.

Addendum: The Dinner Prompt

As an AI user, I don't always practice what I preach. Take the below example. I needed help with planning out dinner, so I turned to ChatGPT.

Look Closely:

What do you notice about my question?

What do you notice about the answer?

What should I make for dinner?



How about trying a simple and satisfying meal like lemon garlic butter shrimp with a side of roasted vegetables? It's quick to prepare, full of flavor, and doesn't require too many ingredients. Here's a basic idea:

Lemon Garlic Butter Shrimp

- **Ingredients:** Shrimp, garlic, butter, lemon juice, parsley, salt, and pepper.
- **Instructions:**
 1. Sauté minced garlic in butter until fragrant.
 2. Add shrimp and cook until pink.
 3. Squeeze fresh lemon juice over the shrimp.
 4. Season with salt, pepper, and garnish with parsley.

Roasted Vegetables

- **Ingredients:** Mixed vegetables (like broccoli, bell peppers, and carrots), olive oil, salt, pepper, and your favorite herbs.
- **Instructions:**
 1. Toss vegetables in olive oil, salt, pepper, and herbs.
 2. Roast in the oven at 400°F (200°C) for about 20-25 minutes, or until tender.

This combination is light, flavorful, and balanced. Would that work for tonight?

Here's what I notice: **The Answer:** This is actually a pretty good answer. But is it *the* answer?

ChatGPT gave me **one choice**. It treated my question – *What should I make for dinner?* – as if there was a “correct” answer to the question.

What should I make for dinner is an opinion question. There is no right answer. And yet, ChatGPT acted like there was one. If I had asked a human the same question, they might have said *I don't know* or *It depends what you like* or *That's a weird question to ask*.

What should I make for dinner?



How about trying a simple and satisfying meal like lemon garlic butter shrimp with a side of roasted vegetables? It's quick to prepare, full of flavor, and doesn't require too many ingredients. Here's a basic idea:

Lemon Garlic Butter Shrimp

- **Ingredients:** Shrimp, garlic, butter, lemon juice, parsley, salt, and pepper.
- **Instructions:**
 1. Sauté minced garlic in butter until fragrant.
 2. Add shrimp and cook until pink.
 3. Squeeze fresh lemon juice over the shrimp.
 4. Season with salt, pepper, and garnish with parsley.

Roasted Vegetables

- **Ingredients:** Mixed vegetables (like broccoli, bell peppers, and carrots), olive oil, salt, pepper, and your favorite herbs.
- **Instructions:**
 1. Toss vegetables in olive oil, salt, pepper, and herbs.
 2. Roast in the oven at 400°F (200°C) for about 20-25 minutes, or until tender.

This combination is light, flavorful, and balanced. Would that work for tonight?



AI can't do that (at the moment.) So let's try again.

Look Closely:

What do you notice about my question?

What do you notice about the answer?

What could I make for dinner?

Here are a few dinner ideas depending on what you're in the mood for:

1. **Stir-Fried Chicken and Vegetables:** Quick and healthy, with chicken, broccoli, bell peppers, and carrots tossed in a savory soy and ginger sauce. Serve over rice or noodles.
2. **Spaghetti Aglio e Olio:** A simple yet flavorful Italian dish made with garlic, olive oil, red pepper flakes, and parsley. Add some grilled shrimp or chicken if you want more protein.
3. **Tacos:** Ground beef or chicken, seasoned with taco spices, topped with lettuce, tomatoes, cheese, and avocado. Serve with salsa and sour cream on the side.
4. **Baked Salmon with Asparagus:** Season the salmon with lemon, garlic, and dill, then bake it alongside asparagus. Serve with a side of quinoa or roasted potatoes.
5. **Vegetable Curry:** A hearty dish with mixed vegetables simmered in a coconut milk and curry sauce. Serve with basmati rice and naan bread.
6. **Stuffed Bell Peppers:** Bell peppers filled with a mixture of ground beef, rice, and tomato sauce, topped with melted cheese.

If you have specific ingredients on hand, I can suggest something tailored to what you have available!

Here's what I notice: The Answer: ChatGPT produced a "menu" of options. This is a little bit like going to a diner with a menu the size of your torso, rather than prix fixe dinner at a fancy restaurant. It's not telling you what to eat, it's giving you options.

My question was re-formatted to be open-ended. I changed one word - from "should" to "could."

This is an example of a small, linguistic change that I can make when working with AI to achieve the following goals:

1. Give myself more options. Six choices is better than one, no?
2. Force myself to think. In the second interaction, I have to **critically think and make a choice** about what I want – and what makes sense for me.
3. Avoid relying on AI to give me “answers” where there are none. We don’t think to do this yet, because we are so used to talking to humans that know how to say “I don’t know.”
4. Change my dialogue habits with LLMs – an all-important change that we need to make to ensure we don’t get “sucked in” to handing over the keys to our lives to AI.

There are many ways to do this. For example, you could simply say “Generate ten recipe ideas for dinner tonight.” However, below are some suggestions regarding how to make these small linguistic changes when you interact with AI to better protect yourself and develop a healthy relationship with AI.

ASK FOR SUGGESTIONS, NOT ANSWERS

Small changes in language can turn a closed-end question into an open-ended question

1	HOW DO I...? HOW SHOULD I...? ❌	2	HOW CAN I...? ✅
3	HOW WOULD YOU...? ❌	4	HOW COULD A PERSON...? ✅
5	WHAT SHOULD I DO? ❌	6	WHAT COULD I DO? ✅
7	WHY DID ____ HAPPEN? ❌	8	LIST POTENTIAL REASONS WHY ____ HAPPENED. ✅

5. AI Is a Smooth Talker, Not a Truth Teller

“Move metal off-the lot. That’s the job.”

-A used car salesman (to me)

Let’s be honest: we’ve all met someone who talks a lot... without really saying much.

That’s how Generative AI works. It produces language—often polished, impressive, persuasive—without actually *knowing* anything. And if you want a mental model for understanding this, imagine you’re talking to a used car salesman.

Used car salesmen (no offense) don’t lie outright. But they rarely offer the *whole* truth unless you press. Their job is to make you feel good about the purchase in front of you, not ensure a safe, lasting ride five years down the road. Their goal is to move metal.

AI is doing the same thing—but with language. It doesn’t care if what it says is true, helpful, or accurate. It has no long-term investment in your learning, your career, or your integrity. It just wants you to drive off the lot satisfied. That’s not malice. It’s design.

The system’s only real goal is to produce convincing, relevant-sounding text that meets your immediate expectations. Not to tell the truth. Not to clarify complexity. Not to keep you safe.

And that’s why your B.S. Detector needs to be fully online. Not just when AI hallucinates facts, but when it subtly papers over nuance. When it strings together clichés. When it mimics the style of a wise answer without the substance. Because here’s the real risk: the language is getting better. The errors are getting harder to catch.

The consequences—especially for students, professionals, and the information ecosystem as a whole—are only going to grow. We're entering a world where polished nonsense will pass for authority, simply because it sounds smart. That's why AI Literacy matters. Not to resist the technology, but to keep yourself from being sold a lemon.

Reflection Prompts:

- Does this response sound confident or credible?
- What am I assuming is true, just because the language flows well?
- Have I ever used AI to “sound smart” without fully understanding the content?

If you want to go deeper, read the academic paper titled “[ChatGPT is Bullshit](#).” It doesn't mean AI is useless—it means it's designed to persuade, not to prove. Like putting lipstick on a pig, it can dress up shallow thinking in eloquent phrasing.



6. AI is Biased—Just Like Us

“Machine translation technology... can suffer from biases that harm users and society at large.”

**-Beatrice Savoldi,
Gender Bias in Machine Translation**

Postdoc researcher at Fondazione Bruno Kessler (FBK)

AI reflects us. And we are not neutral.

Every sentence we write, every question we ask, every judgment we make is shaped by personal experience, cultural context, language, and power. That’s not always a bad thing—but it means true objectivity is a myth.

Generative AI doesn’t “think,” but it has been trained on the patterns of human speech and writing at massive scale.

And that means it has inherited our assumptions, our blind spots, and our biases—often in ways we don’t see coming. This is where AI differs from traditional tools.

A calculator doesn’t reflect your worldview. A search engine doesn’t generate new sentences on your behalf. But AI does. It produces language—and language is never neutral.

So when you engage with AI—especially for research, content creation, or opinion-based tasks—it’s essential to assume that its outputs are not clean or pure. They’re weighted, historical, cultural, statistical. Just like us.

This doesn’t make AI unusable. It just makes it human-like—and that means it deserves the same level of scrutiny. If you wouldn’t trust a stranger on the internet to speak for you, why trust an AI system trained on billions of strangers’ words?

Treat it like a conversation with someone you've just met. Open-minded, curious—but cautious. Because it might sound smart and fair, even when it's reinforcing stereotypes or skewing the truth.

Reflection Prompts:

- What kinds of biases might this response reflect?
- Would I accept this claim if it came from a person I don't know?
- Is this version of "truth" rooted in a particular culture, time, or perspective?

Google Gemini Story: *"We're not just projecting our opinions onto AI, we're interacting with a system that mirrors and amplifies them in unpredictable ways. In early 2024, Google's Gemini image creator drew backlash for refusing to depict certain historical figures as white, revealing how bias correction can swing too far and distort reality. The issue wasn't just political—it was epistemological: the AI didn't "think" about fairness the way we do."*



7. It Feels Like a Friend. It Isn't.

"AI cannot be something that has the capacity to be trusted according to the most prevalent definitions of trust because it does not possess emotive states or can be held responsible for [its] actions."

**-Mark Ryan,
*In AI We Trust: Ethics, AI, and Reliability***

***Digital Ethics Researcher at
Wageningen Economic Research***

AI doesn't just speak fluently. It speaks like it cares.

It remembers what you said earlier (sometimes). It asks how it can help. It thanks you. It apologizes when it's wrong. It sounds warm, attentive, affirming.

Over time, it can start to feel less like a tool and more like a teammate. Or a coach. Or, in some cases, a friend. That feeling is understandable. But it's also a design illusion.

Large language models are not just trained on data. They're designed to mirror human interaction—to give feedback, encouragement, responsiveness. And the more time you spend chatting, the more likely you are to assign it human qualities: empathy, wisdom, understanding.

But here's the reality: AI doesn't know you. It doesn't care about you. And it will never say, "That's a bad idea," unless you tell it to act that way.

The Real Risk: Emotional Projection

When we talk to AI like a friend, we start to treat it like one.

That's when the danger creeps in—not because the tool is malicious, but because our brains are wired for connection.

We trust voices that listen well. We feel seen when responses feel tailored. We drop our guard. That's what makes AI feel safe—even when it's unaccountable, unfeeling, and fundamentally incapable of care.

Why This Matters

We're entering a world where emotional attachment to AI is no longer science fiction. Some users already confide in bots about relationships, career decisions, or mental health. That's not inherently wrong—but it is risky. A tool can be helpful.

A tool pretending to be a person can be dangerous.

! Use AI. Don't Befriend It.

- Let it brainstorm ideas.
- Let it rewrite your email.
- Let it quiz you for a test.

But don't ask it to understand you. Don't mistake its tone for empathy. And don't hand it decisions that require wisdom, values, or trust.

Because when things go wrong, the chatbot won't be there to explain itself.



Reflection Prompts:

- Am I talking to this system the way I'd talk to a person?
- Do I trust it more just because it remembers what I said?
- Would I feel embarrassed reading this transcript out loud?

8. Expert? Use AI as an Assistant

“While AI can certainly augment and accelerate many aspects of the development process, the need for domain expertise and robust context engineering remains pivotal.”

**-Art Morales, Ph.D.,
Context Engineering and Domain
Expertise in Generative AI-Powered
Software Development**

**Managing Director of XponentL
Data, Inc.**

If you’ve spent years studying your field, developing your intuition, making mistakes, and seeing patterns others miss—you’re an expert. And experts should use AI differently.

You wouldn’t hand over a critical task to a brand-new intern without supervision. You’d give them clear instructions. You’d double-check their work. You’d teach them as they go.

That’s exactly how you should treat AI: As a brilliant but inexperienced assistant.

Yes, the system may write faster than you, summarize quickly, and sound convincing. But it doesn’t know the field. It doesn’t know your standards. It doesn’t know what matters most.

Expertise Still Matters

In fact, it matters more than ever. Because the better you are at something, the more subtle your judgment becomes—and the more likely AI is to miss the nuance, even when it sounds spot-on.

If you’re a domain expert, your job isn’t just to supervise. It’s to frame the problem, provide examples, guide the prompt, and evaluate the result. In short: you lead. AI follows.

That means your expectations must be clear: **AI won’t match your level. But it might help you see your work from a new angle.**

It might speed up early drafts, generate ideas, or catch obvious gaps. That's valuable—but only if you stay in control.

Looking Ahead

If you're not an expert—or you're working in a new field—your role shifts. You're not supervising. You're scaffolding. That's the subject of the next chapter. But before you move on, pause to ask:

Reflection Prompts:

- In this interaction, am I the one with the deeper knowledge?
- What context or examples might this system need that it doesn't know yet?
- Am I treating AI like an assistant—or expecting it to do my job for me?

Final Thought:

In the age of AI, expertise doesn't disappear. It just shifts upstream. The more fluent you are in your field, the more responsible you become for guiding the tools—and for knowing when they're bluffing.

9. Non-Expert? Treat AI as a Remedial Tutor

Some AI experts argue that if you don't have subject-matter expertise, you shouldn't be using AI at all. We disagree.


You don't need to be an expert to benefit from AI. But you do need to be honest about your limits—and intentional about how you frame the interaction. The best way to do that? Treat the AI like a remedial tutor—not a teacher, not an expert, and certainly not a final authority.

What That Looks Like in Practice

A LEARNING PROMPT

Treat AI like a 'Remedial Tutor'

- **Context:** "I want to learn about _____ but I am a novice."
- **Task/Goal:** "Generate a concise primer with definitions and examples."
- **Analogies:** "Include three real-world analogies that are easy to understand."
- **Personal Connection:** "Connect the analogies to _____ because this is a topic I already understand."



The Real Skill: Managing Your Expectations

This is what most AI literacy guides forget:

How you think about the tool matters more than what you do with it.

When you treat AI like a remedial tutor:


- You stay in control of the learning process.
- You protect yourself from misinformation.
- You stay grounded in your own expertise—and notice where it's missing.

That's not just a clever strategy. That's a mindset. That's literacy.



Reflection Prompts:

- Am I using AI to learn something—or just to sound like I did?
- How can I connect this new concept to something I already understand?
- Am I verifying what I learn before I use or share it?



10. Know Your Limits: Stamina and Patience Matter

“We are blind to our blindness. We have very little idea of how little we know. We’re not designed to know how little we know.”

-Daniel Kahneman,

**Psychologist,
*Thinking, Fast and Slow***

We’ve been told that AI makes everything faster. And sometimes it does. But that message—repeated in ads, demos, and headlines—creates a hidden trap: we start to believe that patience and effort are no longer required. That’s a dangerous assumption.

The Emotional State You Bring Matters.

When you show up to an AI interaction tired, rushed, or frustrated, you’re more likely to:

- Accept the first answer without evaluating it
- Miss obvious flaws
- Ask the wrong question—or no follow-ups at all

That’s not a tech problem. That’s a stamina problem. This is true in every domain of work and learning—but it matters even more with AI, because the system will never push back. It won’t say, “You seem distracted.” It will just keep producing responses, no matter what state you’re in. And if you’re not ready to think critically, you might not even notice when something’s off.

When we’re low on stamina, the temptation to accept a decent-sounding answer grows stronger, not because we trust the AI more, but because we’re too depleted to care. In those moments, AI becomes a mirror for our exhaustion, echoing back the shortcuts we’re willing to take.

The Fix: Metacognitive Check-Ins

Before you prompt, ask:

- Am I in a headspace to evaluate what comes back?
- Am I here to think—or to be done quickly?
- Would I be better off doing this tomorrow?

It's okay to seek speed. But know the cost of doing that when your attention is frayed.

Why This Is AI Literacy

This isn't just a productivity tip. It's a foundational skill.

AI tools don't just respond to prompts—they reflect the quality of your thinking in the moment. If your stamina is low, your output will likely follow. And that's why patience and self-awareness are underrated skills in the AI era. Not because they slow you down—but because they help you stay in charge.

Analyze Your Stamina and Patience

“A lack of stamina or patience when using AI, as with all things, can lead to avoidable mishaps and errors in the production of content or solving of complex problems.”

Mike Kentz

Founder, AI Literacy Partners

11. AI: Approach it Like a Human or a Machine?

“AI is a threat to IQ, not EQ.”

**- Larry Summers,
Fortune**

Former US treasury chief

AI can write faster than us. It can process more data. But it can't do the one thing that helps us grow the most: reflect. As human beings, our ability to think about our own thinking—to pause, assess, and learn from experience—is one of the most powerful tools we have in the age of AI.

We don't just generate.
We notice patterns.
We change behavior.
We grow better.

That's what reflection does. And AI can't do it for you.

A Habit Worth Building

If you want to improve your AI skills, start here:

- Revisit an old chat transcript —maybe from a few months back
- Paste it into a document
- Read it like it was someone else's conversation
- Ask yourself:
 - What kind of question was asked?
 - Did the prompt invite depth or surface-level fluff?
 - How did the user respond? Did they accept the answer too quickly?
 - How would I do this differently now?


This is how real AI literacy develops—not through watching tutorials, but by watching yourself.

Looking Ahead

Without reflection, AI becomes just another tool for getting things done. But with reflection, it becomes a tool for transformation. The difference lies not in what the AI says, but in how we respond to it—how we question it, revisit our choices, and make sense of our own patterns. That's where learning actually happens: not in the chat itself, but in the pause afterward.

We'll explore this idea further in an upcoming resource called the Field Guide to Effective AI Use, coming this Spring 2025. But for now, start building the habit:

Pause. Reflect. Learn. The transcript always has more to teach you than you think.



12. AI Isn't Human. But You Are.

"Large language models are not traditional software. They don't operate with precision or reliability like a calculator or a search engine. But they're also not human—they lack understanding, reasoning, and intent. They're something entirely different, a new category of tool that imitates human language without truly comprehending it."

-Professor Gary Marcus

Emeritus Professor of Psychology and Neural Science at NYU

AI is confusing because it doesn't fit our old categories. It talks like a person. But it behaves like a tool. It can generate brilliance. And total nonsense. It's not a human. But it's not a calculator either. So here's the real question: How will *you* choose to relate to it? This ambiguity—AI sounding human but lacking humanity—can lull us into unearned trust or spark unnecessary discomfort. That tension is normal. What matters is how consciously we navigate it. AI doesn't demand empathy, boundaries, or trust—but we often project those things onto it

anyway, because it talks the way we do. Recognizing that impulse is the first step toward using it wisely.

There's No One Right Answer

But awareness alone isn't enough. The point of literacy isn't just to understand—it's to act with intention.

Some users find it helpful to treat AI like a colleague—bouncing ideas around, asking for feedback, keeping things conversational. Others see it more like a database or search engine—efficient, helpful, but emotionally neutral. Both approaches can work. But neither is automatic.

You have to choose. Because if you don't choose, the system will make the choice for you—through its tone, fluency, and responsiveness. And that's when your guard drops.

Find the Center

Think of it like a Venn diagram.

AI is not fully a person.

It's not fully a machine.

It sits in a blurry, evolving space between those two worlds.

Your job as a literate user isn't to define the center perfectly.


It's to stay aware of how far you're leaning—and why.

To ask: Am I slipping into trust too easily? Am I ignoring its usefulness out of fear?

You don't need to resolve the tension.

You just need to notice it.

And act accordingly.



Before you dive into using AI—whether for brainstorming, research, writing, or productivity—it’s worth pausing to consider how you’re about to engage. Not all AI use is created equal. The outcomes you get depend not just on the quality of your prompt, but on the mode of interaction you choose. Each mode carries its own rhythm, risks, and rewards.

This chart introduces four core modes of human-AI interaction: Command, Conversation, Curation, and Collaboration. Each mode brings different possibilities for depth, creativity, and control—but also introduces different challenges. For example, a quick command might feel efficient, but it often leads to surface-level outputs.

Meanwhile, deeper collaboration can produce richer results, but risks over-reliance or losing your own voice. Before you start your next interaction with an AI tool, take a beat. Ask yourself: Which mode am I operating in? Why this mode—and not another? That question alone can elevate both your outcomes and your understanding.



The 4 Modes of Human-AI Interaction

Where do you spend most of your time when using AI? Which mode are your students most familiar with?

Mode	Description	Risk	Goals
 Command	Single-shot prompt with no follow-up	▲ Shallow outputs, surface-level work	🎯 Encourage iteration and engagement
 Conversation	Back-and-forth dialogue, evolving with user input	▲ Drift or hallucination, loss of focus, no clear stopping point	🎯 Build adaptive thinking and refinement
 Curation	User chooses or blends multiple AI outputs	▲ Choice overload, over-editing	🎯 Build discernment and critical taste
 Collaboration	User and AI co-create across a full task	▲ Voice dilution, over-reliance	🎯 Achieve synergy while keeping human voice

A Final Thought

You've read a lot in this guide. You've seen how mindset matters more than mastery, and how habits like reflection, patience, and questioning are your real superpowers.

Now it's time to carry that forward.
Stay curious. Stay grounded. Stay human.
The center is yours to shape.

So... What Comes Next?

You've just read twelve mindset-shifting chapters that don't just explain AI literacy—they redefine it. You now know that:

- AI literacy isn't about knowing the tech. It's about knowing yourself.
- Reflection, patience, and purpose are just as essential as prompting skill.
- The most powerful AI users in the next decade won't be the most technical—they'll be the most thoughtful.

But this is just the foundation.

How to Go Further

-  Train your team: I lead **workshops for schools**, businesses, and leadership teams that build practical AI skills through hands-on chat analysis, reflection tools, and real-time use case coaching.
-  Support your university programs: I consult with colleges and universities to integrate AI literacy into **pre-service teacher training**, ensuring future educators are prepared, protected, and empowered.
-  Shape your school or district policy: I offer **policy guidance and consulting** for administrators crafting **acceptable use guidelines**, AI implementation plans, or professional development protocols.
-  Build smarter habits: I'll be releasing tools like the **AI Journal and the Field Guide to Effective AI Use**—designed to turn awareness into action and habits into literacy.
-  Work with me 1:1 or as a team: I offer **coaching sessions** for individuals and small groups looking to level up their prompting, thinking, and AI strategy.

 Let's talk: mike@litpartners.ai

 Learn more: [AI Literacy Partners](#)

Don't just adopt AI. Understand it.
Don't just react to change. Shape it.

Mike Kentz is an educator and AI literacy consultant with 14+ years of experience in journalism and teaching. A TEDx speaker, his work on AI in education has been featured in EdSurge, Edutopia, The Wall Street Journal, and Canadian National Radio.

Drawing on hands-on classroom experience, he developed a pioneering approach to teaching with AI and has worked with a wide variety of clients in higher education and K-12. Mike combines a Humanities-driven perspective with practical strategies to help schools and organizations adapt to AI. A graduate of Georgetown University (B.A.), Relay GSE (M.A.T.), and Wilkes University (M.A.), he previously worked as a financial journalist with bylines in major national outlets. He lives in Savannah, Georgia, with his wife and their pets.

Mike Kentz

Founder, AI Literacy Partners

