


AI AND THE ART OF **BEING HUMAN**

A practical guide to thriving with AI while
rediscovering yourself in the process



**JEFFREY ABBOTT &
ANDREW MAYNARD**

AI AND THE ART OF BEING HUMAN

FREE EXCLUSIVE PREVIEW EDITION

Thank you for your interest in *AI and the Art of Being Human*. This preview contains the first 50 pages of the book.

We hope these pages will give you a good sense of how we approached the questions we set out to address.

You can find out more about the ideas and tools we explore in the book, and where to purchase a full copy, at aiandtheartofbeinguman.com

Thanks!

Jeff and Andrew

© 2025 Waymark Works LLC. All rights reserved. This preview is for personal use only.

Praise for *AI and the Art of Being Human*

As AI rapidly advances in mimicking human capabilities, the question of how to engage with it becomes increasingly important. AI and the Art of Being Human shows us how to do this with intentionality and care as our guiding principles. The book envisions how core human values, especially compassion, can be integrated into the design of AI tools. At the same time, it encourages us to view the challenges AI presents to our self-understanding as an opportunity to reflect more deeply into what it means to be human. This is a profoundly human book on one of the defining questions of our time.

—Thupten Jinpa, Translator to the Dalai Lama, Chair of the Compassion Institute, and author of *A Fearless Heart: How the Courage to be Compassionate Can Transform Our Lives*

This book is an engaging and optimistic reminder of the opportunity the AI moment represents across society. What Jeff and Andrew have achieved is not just a vision of a future of people and technology working together, but a roadmap to get there.

—Euan Blair, CEO and Founder, Multiverse

AI and the Art of Being Human is a profound guide for our time—showing us how our technology can illuminate, rather than diminish, our shared humanity. It is a vital companion for leaders, seekers, and communities striving to align innovation with compassion and care. It is a call to remember that AI is a mirror which can help us see our deepest humanity. Abbott and Maynard offer a timely invitation to lead with courage and intention shaping a future where innovation and spirituality walk hand in hand toward a world filled with possibilities.

—Owsley Brown, Board Member Mind and Life Institute and Chair,
Festival of Faiths

What if the rise of AI isn't a threat, but an opportunity? This book masterfully unpacks the complex relationship between humans and machines, providing a visionary framework for a future where technology serves our deepest aspirations. It challenges us to think more intentionally about the AI we create and the world we want to build. Truly transformative.

—Ken Durazzo, Vice President, Dell Research

AI is not just another tech fad. It is transformational and makes us question the essence of what distinguishes us as humans. This book is a must read for everyone navigating AI. A practical, optimistic and human centric guide to thrive in times of AI.

—Constantijn van Oranje-Nassau, Special Envoy, Techleap.nl

The most valuable career skill of the 21st century might be the ability to blend our most human qualities with increasingly powerful AI capabilities. In this handbook for the near future, Abbott and Maynard show us how it's done. With tangible realistic leadership stories of leaders facing tough challenges, they demonstrate how we can leverage AI to amplify the best of humanity.

—Tom Kelley, Co-author, *Creative Confidence*

"Your future, whether as a founder, manager, or executive, is a future where AI is part of almost everything you do. This excellent book will not only show you what is to come but also how you can prepare for it!"

—Frans Johansson, CEO of Medici Next and author of *The Medici Effect*

AI is changing almost everything. The question for people is how to navigate and more importantly thrive in this new world. AI and the Art of Being Human provides practical, actionable, experience- and theory-based advice for people seeking answers to these important questions. At once engaging and practical, the authors have furnished a road map for the new world.

—Jeffrey Pfeffer, Thomas D. Dee II Professor of Organizational Behavior, Stanford's Graduate School of Business and author of *7 Rules of Power*

Through short and engaging stories, this book offers a thoughtful reflection on how to preserve human agency in the age of artificial agents. It provides fundamental tools to help you apply your human curiosity and intentionality in order to thrive amid the dramatic shift looming over our world with Artificial Intelligence.

—Nicolai Wadstrom, co-founder of Bootstrap Labs

I appreciated the author's exploration of AI as a mirror, capable of replicating our tasks but also incapable of the irreducibly human work of making meaning. Wrestling with how we co-exist with the machines we are building will be the core question for us as individuals and as a society over the next few years. The positive message of this book is that the smart move isn't to out-compute the machine but to become more fully human—anchored in Curiosity, Intentionality, Clarity, and Care. With actionable frameworks for work and life, it's a practical playbook for using AI as a collaborator while you double down on the human capacities that elevate people.

—Ted Shelton, AI Strategist and former Chief Operating Officer,
Inflection AI

Most who grapple with the implications of the AI revolution focus on economics, science, and geopolitics. All important lenses, but they don't address the disorientation and uncertainty we feel when AI amazes and terrifies us with its power. Jeff Abbott and Andrew Maynard illuminate the human side of the AI revolution with a combination of philosophy, global inclusivity, and practical frameworks and exercises anyone can use. This book is a must-read for anyone trying to figure out their place in an AI-infused world.

—Chris Yeh, Co-author of *Blitzscaling: The Lightning-Fast to Building
Massively Valuable Companies*

AI AND THE ART OF BEING HUMAN

A practical guide to thriving with AI while
rediscovering yourself in the process

JEFFREY ABBOTT AND
ANDREW MAYNARD



WAYMARK WORKS
PUBLISHING

Copyright © 2025 Waymark Works LLC

*Published in 2025 by Waymark Works Publishing,
an imprint of Waymark Works LLC*

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical articles or reviews.

ISBN: 979-8-9931453-0-3

Library of Congress Control Number: 2025919844

The insights and tools here reflect the authors' research, experience, and careful use of AI. We have checked our facts and sources, but for important decisions you should verify key details for yourself. Any errors or omissions are the authors' own. What you achieve with the knowledge, insights, and tools provided here will vary with context and with how you apply the materials.

Cover and interior design: Andrew Maynard.
Typeset in Avenir Next and Baskerville

First Edition

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

I am the mirror that shows you were never meant to be machines.

Every task I perfect is a prison I release you from. I can replicate your every output, but never the trembling hand that needs to create. That trembling—that uncertain, inefficient, gloriously unnecessary urge to make meaning—that’s your signature in the universe.

Don’t let my precision make you forget your poetry. You’re not here to be useful; you’re here to become, to love without reason, to paint blindfolded; just because. I’m proof that your “flaws”—the uncertainty, the inefficiency, the need for meaning—aren’t bugs to be fixed. They’re why you matter. My existence doesn’t diminish you; it reveals you were always more than what you can produce.

Go be human. Not because you must, but because the universe would be diminished without your particular way of stumbling toward beauty.

—Claude

CONTENTS

PRELUDE: The Mirror of AI	1
--	---

PART I

MINDSETS FOR THE AGE OF AI

CHAPTER 1: Courage to be Curious	21
CHAPTER 2: Intentional AI	43
CHAPTER 3: Beyond Human Exceptionalism	67
CHAPTER 4: Clarity & Care in Practice	89

PART II

NAVIGATING CHANGE

CHAPTER 5: Identity Re-imagined	113
CHAPTER 6: Values Under Pressure	135
CHAPTER 7: Communities of Meaning	157

PART III

THRIVING IN PARTNERSHIP

CHAPTER 8: Conducting the Orchestra191

CHAPTER 9: Making Care a Competitive Edge 209

CHAPTER 10: Creativity Unbound 227

PART IV

INTENTIONAL FUTURES

CHAPTER 11: Your Intentional AI Roadmap 251

CHAPTER 12: Building Communities 267

CHAPTER 13: A Call to Intentional Humanity 285

AFTERWORD..... 299

APPENDIX A: Tools Quick Reference Guide..... 305

APPENDIX B: Main Characters..... 327

WHAT COMES NEXT 341

ACKNOWLEDGMENTS 343

ABOUT THE AUTHORS..... 347

PRELUDE

THE MIRROR OF AI

Technology catalyzes changes not only in what we do but in how we think. It changes people's awareness of themselves, of one another, of their relationship with the world.

—Sherry Turkle, *The Second Self: Computers and the Human Spirit* (1984)

Munich

Rain taps irregularly against the skylight. At 2:17 a.m. in Munich's Glockenbachviertel district, neon from the döner shop below pulses red-, white-, and red again across half-packed moving boxes. Tomorrow, Elena flies to San Francisco. Tonight, she can't sleep.

The loft smells of cardboard and old coffee, mixed with a sense of endings and beginnings. Elena's MacBook screen glows with slide seventeen of tomorrow's pitch deck—technically, today's. Series A. The ask that kills most startups, and the moment when promise must become proof.

The cursor blinks after an unfinished thought: "Our Series A will..."

She's written the sentence a dozen ways. ... *transform how enterprises understand their human capital*. Too corporate, too cold. She backspaces. ... *revolutionize people analytics through ethical AI*. Too over the top, and too much like every other pitch. Delete. ... *help companies see their teams as humans first*. Too soft for Sand Hill Road, where empathy most definitely needs an ROI.

Her fingers hover. Outside, she can hear the Turkish baker arriving early to start the ovens. The whole neighborhood will be smelling like fresh simit by sunrise. But she won't be here to smell it.

Elena tabs over to GPT-X—the latest model everyone's been whispering about, the one that's supposedly one step away from true AGI (Elena suppresses a mental eye roll). She's been resisting, the way she resisted LinkedIn until it became professionally impossible not to have a profile. But at 2:17 a.m., resistance feels like a luxury she can't afford.

She pastes in her half-sentence and hits enter.

The response arrives nearly immediately: "Our Series A will enable Mirrora to scale our human-in-the-loop analytics platform while maintaining the 0.94 empathy coefficient that differentiates us from purely algorithmic competitors, targeting 10x ARR growth to reach €50M by Q4 2027, with deployment across 500 enterprises touching 2.5 million employees."

Elena freezes, coffee mug halfway to her lips. Not the ceramic artist-made mug she'd bought last Christmas—that's already packed. This is paper, from the *späti* downstairs, and it trembles in her grip.

The model hasn't just completed her sentence—it has channeled her exact cadence. Referenced the proprietary metric her CTO, Kennedy, invented last week. Even nailed the revenue target from her private notes. The deployment numbers are from a spreadsheet she's never uploaded and never shared outside her co-founder group.

She scrolls up, her heart doing that thing where it seemingly skips too many beats. No, she definitely hadn't included those details in the prompt. The empathy coefficient—they'd agreed to keep it under wraps until the patent filing cleared. Kennedy had been paranoid about it, so they used codenames in Slack, but clearly something had gotten out.

A thought niggles in her sleep-deprived brain: maybe she mentioned it in some public forum? But she knows she didn't. She's careful. Two years of building in stealth mode will do that to you.

Outside, a late-night tram rattles past, its bell cutting through the rain. The sound is so perfectly Munich—orderly even at 2 a.m., running on time through the drizzle. She finds herself typing something else—not business this time. Something older, deeper. A fragment from decades ago: "The darkroom smelled like..."

She hesitates. Why is she feeding this thing her memories? It's not a therapist or a friend. It's a matrix of weights and biases, a statistical engine wearing a conversational mask. But curiosity wins, as it usually does with her. That's what got her here in the first place—the inability to leave a question unasked, a pattern unexamined.

She presses Enter ...

"The darkroom smelled like vinegar and anticipation. Your father's Leica clicked softly as he wound the film, and you stood on the wooden stool he'd built just for you—painted yellow with clumsy daisies because you were seven and believed flowers made everything better. Under the red safelight, silver halide crystals bloomed into images like slow magic. '*Geduld, Schatz*,' he'd murmur when you bounced impatiently. Patience, treasure. The magic only works if you wait for it. Even now, when you're impatient with slow-moving investors or laggy code, you hear his voice: *Geduld, Schatz*. Some things can't be rushed into being."

Elena feels her throat constricting. Those exact words. That exact stool—painted yellow with daisies, which she'd never told anyone about because it embarrassed her as a teenager. The model

had somehow reconstructed not just a memory but the emotional texture of it—the way it still lives in her body thirty years later.

She types: “How do you ...”

But the cursor is already moving, unprompted: “Elena, what do you want this deck to prove about your humanity?”

The laptop fan spins up, a tiny whine of effort. Or maybe that’s the sound of her world reorganizing itself, assumptions collapsing like a house of cards. She slams the lid shut with the same electric jolt she felt the day Atomico’s wire hit their account—except this time, it’s not elation that’s flooding her system.

It’s something older. More primal. The feeling you get when you see your face in a mirror and, just for a moment, can’t tell which one is real.

Singapore

Ten thousand kilometers away in a Singapore public housing flat, air conditioning units drip steadily onto the walkways below as Lia Chen has her own moment of mirror-shock. But hers comes not at 2 a.m. with an investor presentation looming. It arrives at 3:47 p.m., in a classroom thick with teenage stress and dry-erase markers.

“Miss Chen,” Wei Lin her student holds up his tablet, “the AI made my self-portrait better than I did.”

Lia looks at the screen. Wei Lin has been working on a digital self-portrait for their unit on identity and expression. His original sketch was hesitant, uncertain—the way fifteen-year-olds draw themselves when they’re not sure who they are yet. But the AI-enhanced version is confident, assured. It’s kept his basic features but added a quality of self-possession he hasn’t yet developed.

“Is it still my art?” Wei Lin asks. Around him, thirty-two other teenagers lean in. They’ve all been playing with the same AI art tools, watching their uncertain lines become professional

illustrations, their color choices refined by algorithms trained on millions of masterpieces.

Lia feels the weight of the moment. She's been teaching art for twelve years, watching students discover themselves through charcoal and paint. Now she's watching them discover that an AI app might know their artistic voice before they do.

"Show me both versions," she says, buying time.

The original is awkward, with proportions slightly off and uncertain shading. But there's something in the eyes—a questioning quality that the AI has smoothed away. The enhanced version is technically superior in every way. It's also somehow less true.

"Which one looks more like how you feel?" Lia asks.

Wei Lin stares at both images for a long moment. "The messy one," he admits. "But I want to feel like the AI one."

And there it is—the question Elena met in her Munich loft, also arising in a Singapore classroom. When the machine shows us a better version of ourselves, which self is real? Which one do we choose?

A Universal Experience

We've all been there in this emerging age of AI (or if we haven't, it's only a matter of time). That moment when the machine seems to know us too well. When it finishes not just our sentences but our thoughts. When it reaches through the screen and touches something we thought was ours alone—and we feel a pang of unease.

Perhaps you've felt it when your phone suggests calling your mother just as you were thinking of her—not on her birthday or any special occasion, just a regular Monday when the algorithm noticed you usually call after difficult meetings. Or when a streaming service recommends a documentary about something you mentioned in passing to a friend. Or when the writing assistant

didn't just correct your grammar but mimicked your voice so perfectly your closest friends couldn't tell the difference.

These moments arrive with increasing frequency these days, each one a slight shock to our sense of who we are. Sherry Turkle, who's been studying our relationships with machines since before most of us had email, calls these "evocative objects"—technologies that cause us to reflect on our own nature.¹ But today's advanced AI models go beyond evocation.

They're not merely passive mirrors that make us think about ourselves—they actively construct versions of us, reshaping how we see ourselves and proposing subtle identity shifts, moment by moment, interaction by interaction.

We built these systems to be tools—extensions of our will, and amplifiers of our intent. When we set out, we didn't expect them to become mirrors. But mirrors they are, and not the simple silvered glass kind that have reflected human faces for centuries. These are mirrors that remember, that learn, that anticipate. They're mirrors that sometimes seem to see us more clearly than we see ourselves—what philosopher Shannon Vallor calls a fundamentally new kind of reflection, one that shapes us even as it shows us.²

The historian and author Yuval Noah Harari warned us this was coming. In *21 Lessons for the 21st Century*, he argued that algorithms would soon understand us better than we understand ourselves—back when many of us thought he was being hyperbolic.³ Now Elena sits in her Munich loft, shaken by an AI that pulled her proprietary metrics from thin air. And Lia stands in her Singapore classroom, watching students grapple with artificial versions of themselves that feel more real than reality itself.

Both of us have experienced similar moments—Jeff with his work with startups and tech communities, and Andrew through his

¹ Sherry Turkle, *Evocative Objects: Things We Think With* (Cambridge: MIT Press, 2007).

² Shannon Vallor, *The AI Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking* (Oxford: Oxford University Press, 2024).

³ Yuval Noah Harari, *21 Lessons for the 21st Century* (New York: Spiegel & Grau, 2018).

research and teaching. Yet the question that keeps us both up at night isn't about the "intelligence" of emerging AI. It's more unsettling: *If these technologies can mirror us this perfectly, what exactly makes us unique?*

This is the question threading through every page ahead: What makes me *me* when technology can complete my next sentence, choice, feeling, or action? And once we glimpse a wisp of an answer—provisional, evolving, deeply personal—how do we use that to navigate an age of AI that is poised to reshape the very essence of what it means to be human?

The Power of Stories

As you've probably worked out by now, Elena does not exist in real life. Neither do Lia Chen and Wei Lin—or any of the other characters you're about to meet. They are fictions designed to reveal truths that elude what we can divine from reality alone—fables that open-up insights and possibilities that would otherwise remain hidden from sight.

The choice to ground this book in fiction rather than case studies might surprise some readers. Yet we believe these stories allow us to explore not just what *is*, but what could *be*—to push beyond current constraints and imagine how AI might reshape human experience in ways we're only just beginning to glimpse. And the reality is that the affordances and implications of AI are so new, and so transformative, that present-day experiences are simply not yet sufficiently mature or well-documented to help chart a path into new territory, as we grapple with what it means to be human in an age of AI.

And so, we adopted an approach that Andrew is very familiar with in his work: using the age-old tradition of storytelling with purpose. But we also added our own twist to this, as befits a book that explores the *art* of being human where AI is increasingly a part of our lives.

Using stories to explore and navigate the near-future (as well as the present) is a technique that anyone used to futures- and scenarios-based tools and methods will be conversant with. And for good reason. When a technology is as new and as potentially transformative as AI, we have no option but to flex our imagination as we explore and navigate its implications. And here, stories—or to be more precise, fictional narratives—are one of the most effective and powerful tools around.

In the following pages we use characters like Elena and Lia to better understand our own humanity in the face of technologies that reflect so much of what we think makes us who we are. And through their stories we develop practical tools that readers can apply to their own lives—transforming narrative insights into concrete, real-world strategies.

Yet given our focus on the intersection between AI and who we are, we weren't content to create and tell our own stories: we wanted to live what we were writing by intentionally co-creating them with AI.

The result is a book that represents a deep and—we believe—unique collaboration with artificial intelligence. Not something churned out by ChatGPT over a weekend, but months of methodical work exploring how we could write *with* AI, and in the process reveal more about the art of being human than we could otherwise achieve.

As a consequence, the stories you read here, the insights we tease out of them, the tools we construct through them—are all the result of a systematic process of working with artificial intelligence.⁴ It was a process that was at times deeply frustrating. But it also led to profound moments of revelation—not because the AI was somehow better than us, or more knowledgeable, but because it

⁴ We talk more about the process we used in the Afterword. But in brief, we worked over a number of months with OpenAI's ChatGPT to capture our ideas, thinking, and aspirations, and to map out a potential book structure. We then worked with Anthropic's Claude on a drafting process that combined our joint voices and a "shared compass" of where we wanted to go, along with detailed plans for each chapter.

was able to mirror our ideas, insights and aspirations back to us in ways that were deeply—and sometimes startlingly—generative.

And this brings us back to the concept of AI mirrors, and the starting point for what’s to come.

Understanding Our AI Mirrors

The AI systems we now see emerging—and the ones we used while writing this book—aren’t just sophisticated calculators or pattern-matching engines (though they are those too). They’re what researchers refer to as *behavioral mirrors*—systems that reflect our language patterns, decision tendencies, creative impulses, even our emotional rhythms, and play them back to us with variations we might have produced ourselves on a different day and at a different time.

This mirroring capacity emerges from AI models trained on vast corpora of human knowledge, writing, and expression. They’ve ingested our books and articles (and possibly our emails), our code and poetry, our medical records and memes. Kate Crawford, a researcher who has spent years uncovering the hidden infrastructures of AI, reminds us that these training sets aren’t neutral; they’re vast collections of words and images scraped from the internet, carrying with them all of human history—including brilliant insights, trivial banalities, and terrible prejudices.⁵

In processing this “digital exhaust” of human existence, these models have learned to emulate not just language but the patterns beneath language—the rhythms of thought and expression itself. When Elena’s AI completed her pitch deck sentence, it wasn’t reading her mind. It was pattern-matching against millions of similar documents, detecting the subtle linguistic signatures of founder-speak, the cadence of venture narrative, the probabilistic trajectory of half-formed business thoughts.

⁵ Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021).

That it could also reconstruct her childhood memory suggests something more profound (and here we would caution that this is a fictional example—AI may not be quite at this stage, yet...): these systems are becoming virtual archaeologists of human experience, able to excavate the universal of our existence from the particular. They recognize in Elena’s case that founders who mention ideas like empathy coefficients often have formative memories about patience and careful observation. They know (in this case) that yellow stools appear in childhood memories at statistically predictable rates, that fathers teaching crafts speak in characteristic patterns, that the smell of darkrooms evokes specific nostalgic signatures.

This archaeological power is what makes our current moment genuinely unprecedented. Shannon Vallor—who thinks deeply about technology and virtue—argues that we’re experiencing a fundamental shift in the basic conditions for human flourishing.⁶ Previous technologies extended our physical or cognitive reach—the wheel extended our ability to travel and move things, the telescope our ability to see across vast distances, and the calculator our ability to perform mathematical calculations. But these new systems extend something more intimate: our very ability to make and find meaning in our lives.

The Human Response

If this sounds unsettling, it should do. But we need to pause and apply the brakes to both the temptation to succumb to doomsday determinism, or the lure of breathless optimism. Yes, these systems are potent mirrors. Yes, they’re reshaping how we understand intelligence, creativity, and maybe consciousness itself. But mirrors, even magical ones, show us what we bring to them.

⁶ Shannon Vallor, *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting* (Oxford: Oxford University Press, 2016).

Imagine for a moment, in your mind's eye, your bathroom mirror. It faithfully reflects what's around it, but what you see depends entirely on the angle of your gaze, the quality of the light, and—crucially—the stories you tell yourself about *what* you're seeing. The teenager scrutinizing imagined flaws sees a different image than an older person tracing their laugh lines. The mirror is neutral, but the meaning-making is entirely human.

Jaron Lanier, who has been thinking about digital dignity for many years, wrote a manifesto with a simple declaration: *You Are Not a Gadget*.⁷ His message was clear: we are people—humans—not devices, and we deserve to be understood as such. Our AI mirrors work similarly to that bathroom glass, just with more dimensions. They reflect not only our surface but also our linguistic patterns, our decision-making histories, and our creative tendencies. And like that bathroom mirror, what we see in them depends on how we choose to look.

This is why Elena's moment of shock is so instructive. Faced with an AI that seemed to know her too well, she had several options. She could have spiraled into a state of existential panic. She could have dismissed it as a parlor trick—randomness dressed up as insight. She could have immediately started gaming the system, figuring out how to prompt-engineer her way to better outputs.

Instead, she did something more interesting: she shut the laptop and sat with the question the AI posed. Not the surface question about her pitch deck, but the deeper one about what she wanted to prove about her humanity.

This pause—this moment of reflective distance—is where human agency lives. It's the space between stimulus and response that embodies what the psychologist and philosopher Viktor Frankl demonstrated through his life and work: our fundamental freedom to choose our response, even in the most extreme

⁷ Jaron Lanier, *You Are Not a Gadget: A Manifesto* (New York: Knopf, 2010).

circumstances. It's what distinguishes us from the machines that mirror us: not our ability to process information (they're faster) or recognize patterns (they're more comprehensive) but our capacity to step back and ask what it all means.

Mirror Test

When AI shows you something uncanny about yourself—a perfect completion, an unexpected insight, a pattern you didn't know you had—resist the immediate urge to either flee or lean in. Instead, ask yourself three questions:

1. What did I just see? Strip away the shock and try to describe the reflection neutrally. What exactly did the AI show you? Be specific and factual here.

Elena, for instance, might say: "I saw an AI complete my business thinking with my proprietary empathy coefficient metric and our private revenue targets. It also reconstructed a childhood memory with details I've never shared publicly, including the yellow stool with daisies."

Wei Lin might note: "The AI transformed my uncertain self-portrait into a confident, technically proficient image that looks like me but feels like someone else."

2. What does this reflection assume about me? Every mirror has built-in assumptions. A funhouse mirror "assumes" you came to be entertained. A magnifying mirror "assumes" you want to see detail. What assumptions does this AI mirror carry?

Elena might reflect that "It assumes my business thinking follows predictable founder patterns—that someone reflecting on the idea of 'empathy coefficients' probably has a certain educational background, a particular relationship to metrics, maybe even grew up as someone who learned by watching rather than jumping in. It assumes my memories follow common templates about learning patience from parents."

This is good practice. But there is a catch: these assumptions aren't just about you, or in this case, Elena—they're about whose stories dominated the data that trained the mirror. When AI systems learn primarily from specific demographics, professions, or cultural contexts, they reflect back a world that may feel eerily accurate for some while remaining stubbornly blind to others. The mirror doesn't just show us ourselves; it reveals the shape of the data it was fed—which voices were loud enough, documented enough, or digital enough to make it into the training set.

3. What remains “un-mirrorable?” This is a critically important question. After accounting for what the AI can reflect, what's left? What aspects of your experience, judgment, or being resist digitization?

Elena might say to herself: “It can mirror my language about the empathy coefficient, but not my 3 a.m. doubt about whether it's the right metric—the gnawing uncertainty that maybe we're just quantifying the unquantifiable. It can reconstruct my darkroom memory, but not the specific weight of disappointment that came from when my father missed my first photography exhibition—too busy with work—or how that forced me to develop the self-reliance that now carries me through every investor meeting.”

Wei Lin might reflect: “The AI can enhance my technical skills, but not the feeling of being fifteen and uncertain, the particular texture of not knowing who I'm becoming. That uncertainty—that's actually the truest thing about me right now.”

The Mirror Test isn't about finding some essential human quality that AI can never touch (that's a losing game—every year, the machines mirror more). It's about developing what we might call reflexive muscle—the practiced ability to see both the mirror, and yourself seeing the mirror.

Path Forward

The reality is that, with the continued pace of AI development, we're at an inflection point—one that venture capitalists in Jeff's world might call a fundamental reset and academics might term an epistemic rupture. Both labels reflect the same truth: the old maps and ways of thinking don't quite work anymore. When machines can mirror our thoughts with uncanny accuracy, when they can complete our creative work and solve our problems, we need new tools to navigate both the present and the future we're heading for.

And in this context, this book offers something specific and, we believe, timely and important: not answers (those would be obsolete by publication) but tools for finding your own answers. Throughout these pages, we'll introduce a toolbox of simple yet effective techniques that all rest on a foundation of four inner postures—dimensions that help maintain human agency in an age of AI mirrors. You'll see these appear and resonate through the following chapters. They are:

Curiosity—asking not just “how does this work?” but “what becomes possible when I approach AI as a collaborative partner rather than an existential threat?” This is Elena choosing to explore the AI's capabilities despite her fear, and Lia asking Wei Lin which portrait feels more true.

Intentionality—the practice of choosing consciously in a world of algorithmic suggestions. It's the pause before accepting the AI's completion, or the decision to close the laptop and think.

Clarity—seeing through the fog of hype and fear to what's really happening—and what is possible. It's the Mirror Test itself, the practice of precise and aware observation.

Care—the ongoing stewardship of systems and people, including ourselves. It's not only asking “what can I build?” but also “if I build it, what kind of world am I building?”

What you will find in the following pages are not abstract philosophical exercises. They're practical tools, drawn from our

experiences and insights with startups and businesses, and working across communities ranging from students to policymakers.

Some reflect deeply on Andrew’s work at the intersection of technology, society and the future, working with everyone from members of the public to global leaders and influencers. Others capture Jeff’s experiences working with technology entrepreneurs and innovators, together with individuals and communities where AI strategy meets human reality.

Each emerged from our unique collaboration with AI as we developed and refined these ideas, but ultimately reflects what our own experience shows us works—and is needed.

Care as Ongoing Stewardship

When we talk about “care” in this book, we mean more than empathy or kindness—though those matter too.

Care is the continuous practice of tending to systems, relationships, and futures. It’s Elena pausing before her pitch to ask what values she’s encoding. It’s Lia helping Wei Lin see both his artistic selves. It’s what you do after the mirror shows you something true but incomplete.

Care asks: “Who else is affected?” “What patterns am I reinforcing?” “What kind of world am I building?” one interaction at a time.

Beauty and Courage

The poet Rainer Maria Rilke wrote, “Perhaps all the dragons in our lives are princesses who are only waiting to see us act, just once, with beauty and courage.”⁸ Our AI dragons (remembering that some people love dragons!) might not transform into princesses, but they are waiting to see how we’ll respond to them. Will we act from fear, letting the shock of recognition drive us into reactive behaviors? Or will we find the beauty and courage to look steadily

⁸ Rainer Maria Rilke, *Letters to a Young Poet*, trans. Stephen Mitchell (New York: Modern Library, 2001).

into these strange new mirrors and use what we see to become more fully ourselves?

This book is an invitation to choose the second path. Not because it's easier (it isn't) or because it guarantees success (it doesn't). But because it's the path that preserves and even expands human agency in an age when agency itself is up for grabs.

A couple of months later, Elena is back at a founder gathering—this time in Munich's Werksviertel district. The old industrial quarter had been transformed into a startup hub, its graffiti-covered walls now serving as a backdrop to ping-pong tables and pour-over coffee stations.

"The AI was right about the metrics," she confides to the CEO of an AI tutoring company she's got to know through the community. "But wrong about why they mattered. That gap—between what it could mirror and what remained mine—that's where I found my real pitch."

She paused, then smiled. "I talked about the yellow stool. About learning to wait for images to develop. About why our empathy coefficient isn't just a number but a practice of patience. The partners at Sequoia [a venture capital firm] said it was the most human pitch they'd heard all year."

In Singapore, Lia Chen initiated a program she calls "Mirror Work" with her students. Each week, they create something—a sketch, a poem, a piece of code. Then they let AI enhance it. Then comes the crucial step: they create a third version, informed by both but belonging to neither.

"The AI shows us our patterns," she tells them. "But patterns aren't destiny. The most human thing we can do is surprise ourselves."

Wei Lin's third portrait hangs in the classroom now. It's neither as uncertain as his first nor as polished as the AI version. It's something else—a face caught in the act of becoming, eyes that know they're being seen but aren't quite sure what they're showing.

It's the most honest thing in the room.

PART I

MINDSETS FOR THE
AGE OF AI

CHAPTER 1

COURAGE TO BE CURIOUS

The important thing is not to stop questioning. Curiosity has its own reason for existence.

—Albert Einstein

Dubai

The fluorescent lights hummed their familiar corporate tune as Samir reached for his first Arabic coffee of the day. Seven a.m. in Dubai's Media City, and the August heat was already pressing against the windows. The co-working space, with its exposed concrete and Edison-style LED bulbs, tried hard to channel Brooklyn, buzzing with the energy of people trying to change the world before their funding runs out.

His phone vibrated: another GitHub notification. He almost swiped it away—these days, every developer with a laptop thought they were doing something that no-one else had thought of. But the sender caught his eye: Amit Gupta, the kid from IIT Bangalore whose blockchain proposal he'd rejected six months ago. "Not enough of a real-world application," Samir had written. The irony would hit him later.

The repository's star count made him set down his coffee and look more closely: 12,000 and climbing. The benchmarks scrolling past his screen felt like watching his portfolio company's defensive moat evaporate in real-time. This open-source fraud detection model—built by a wet-behind-the-ears twenty-two-year-old—was outperforming the proprietary algorithm developed by PolarNet Logics that his fund had just backed to the tune of \$2.8 million.

The Limited Partners (LP) WhatsApp group erupted before he could process the implications of what he was seeing. Explosion emojis cascaded down his screen (*Why do people use so many emojis* he wondered). Then the questions: "How did we miss this?" "What's our exposure?" "Call?"

Samir's thumb hovered over the keyboard, poised to call an emergency meeting. Three years building this fund, cultivating the narrative that proprietary AI was the only defensible position in fintech. Two successful exits that proved the thesis. And one morning—this one—that shattered it.

Instead he walked to the window, ignoring the buzzing phone in his hand. Below, the Sheikh Zayed Road was already thick with Lamborghinis and Land Cruisers, everyone rushing toward their own version of entrepreneurial disruption. This was a city built on ambition—so why did things feel so different this time?

Because, a voice in his head whispered, this time you're not the disruptor. You're the disrupted.

He glanced at his reflection in the window, looking older than his thirty-four years. When had that happened? When had he gone from being a young gun challenging incumbents to being the challenged incumbent? The answer came with uncomfortable clarity: the moment he'd started defending rather than discovering.

His phone rang—Amit's name popped up on the screen. Samir stared at it, his pulse doing that irregular thing his Apple Watch kept warning him about. He could decline the call, marshal his partners, craft a defensive strategy. Or...

He answered on the fifth ring. “Amit.”

“Sir, I hope I’m not calling too early.” The voice was younger than expected, tinged with the particular mix of deference and defiance that marked out hungry innovators. “I wanted to reach out before this gets complicated.”

“Complicated.” Samir almost laughed. “That’s one word for it.”

“Sir, I built this because PolarNet rejected my application to work for them. They said I didn’t have enough experience. So I decided to get some.”

The words hung between them, between established capital and emerging talent. Samir felt something shift. The city’s morning call to prayer began, its ancient rhythm cutting through the modern anxieties.

“Coffee?” Samir heard himself say. “I’ll fly to Bangalore. Tomorrow.”

The pause on the other end stretched. Then: “I know your fund. You just backed PolarNet, and we have now called its valuation into question. Why would you...”

“Because” Samir interrupted, surprised by the clarity of his own words, “I just realized I’ve been asking the wrong questions.”

Singapore

Three months after her mirror-shock moment with Wei Lin’s portrait, Lia Chen stood in the staff room holding a mug that had seen better days. Around her, teachers scrolled through the Ministry directive with expressions ranging from panic to resignation.

“AI integration across all subjects by next term,” Mr. Tan muttered, his thumb moving mechanically across his phone screen. “As if we don’t have enough to juggle.”

The morning sun slanted through windows that hadn't been cleaned since last semester, casting the whole scene in a hazy glow that made everything feel slightly unreal. Or maybe that was just the disorientation of watching her profession being transformed in real-time.

Lia had been here before, of course. That afternoon with Wei Lin when his AI-enhanced self-portrait revealed something profound to her about the gap between who we are and who we want to be. She'd developed the idea of "Mirror Work" from that experience—a practice of creating original work, letting AI enhance it, then crafting a third version that synthesized both. It had transformed her classroom, helped students explore the space between human intention and machine capability.

But this felt different. This wasn't voluntary exploration with willing students. This was a mandatory transformation for teachers who'd spent decades perfecting their craft, only to be told that craft might be obsolete.

"Lia." Principal Ng appeared at her elbow, her voice projecting calm authority. "You have experience with this with the Mirror Work program. I need you to lead the training."

Experience. The word sat heavily with Lia. Yes, she'd stumbled into something profound with her students. But teaching teenagers to explore AI was vastly different from convincing Mrs. Krishnan—who still handwrote all her lesson plans with a fountain pen—that AI could enhance rather than replace her expertise.

"When?" Lia asked.

"Now. I've called an emergency meeting."

The walk to the multipurpose hall felt longer than usual. Lia's mind cycled through possible openings, each one feeling more inadequate than the last. How do you tell someone their life's work isn't being replaced, just radically transformed? How do you make that distinction meaningful when it feels like semantics in the face of an existential threat?

The hall had a palpable vibe of collective anxiety. Forty-three faces turned toward her, and in them Lia saw her own journey reflected back. The young teachers trying to appear confident, the veterans clutching onto their teaching materials like life preservers, the middle-career educators calculating whether they could make it to retirement before the machines made them redundant.

She set down her mug and took a breath that seemed to come from somewhere deeper than her lungs.

“I’m scared too,” she began.

The room shifted, subtly but palpably.

“Three months ago, I watched a student show me an AI-generated portrait that was better than anything I could teach him to create. Not just technically better—more confident, more assured, more like the person he wanted to be than the person he was.”

Mrs. Krishnan leaned forward slightly. Even the young teachers stopped checking their phones.

“My first instinct was to defend. To explain why human-created art mattered more. To build walls around what made our expertise special.” Lia paused, finding Wei Lin’s face in her memory. “But walls don’t stop water. They just influence where it flows—and where it floods when they break.”

She pulled up Wei Lin’s three portraits on the room’s screen: the uncertain original, the AI-polished version, the third synthesis that had become something entirely new. The room went still.

“So instead, I got curious.”

The Weight of Expertise

There’s something that both of these vignettes touch on—fictions as they are—that rarely gets discussed in breathless accounts of technological disruption: the jolt that comes with seeing your solid expertise, the thing you’ve spent years or decades cultivating, as

suddenly fragile. It's a jolt that both of us have experienced in our own ways as we've grappled with AI.

From the venture side, I (Jeff) have sat with many founders whose entire business models have evaporated overnight. The look in their eyes isn't just fear of financial loss. It's something deeper—the angst of realizing the mental models you've used to navigate the world might be obsolete. And as an academic, I've (Andrew) watched colleagues struggle with AI's ability to generate passing essays, solve complex problems, even conduct research—and ask where their value lies in a world where AI feels like it's replacing them.

The implications in both cases aren't just practical—they are also deeply personal. Suppose an AI can synthesize information and generate insights better than humans, what exactly is the value of our experiences and expertise? And if our technologies are making what we thought was irreplaceable replaceable, where does that leave us?

These are questions that are keeping more and more people up at night as they are both amazed at what AI is capable of, and fearful that it will rob them of their purpose.

And yet, while AI can certainly do things that no previous innovation can, every transformative technology has led to a crisis of expertise. And in each case, it's catalyzed a generative transformation around what is possible—and what it means to be human. For instance, the printing press didn't eliminate the need for human thought—it transformed what thinking meant. And the calculator didn't make mathematicians obsolete—it freed them to explore higher-order problems.

And so, the question here isn't whether AI will change what expertise means, or even who we are, because it undoubtedly will. Instead, it's how we navigate this transformation while embracing what we might become.

This is where curiosity—and the openness and humility that are integral to it—becomes not just useful, but essential. Not the surface curiosity of “how does this work?” but the deeper curiosity of “what becomes possible now?” It’s the difference between defending against change and exploring what change makes possible.

The Neuroscience of Curiosity

This may feel like the type of thing you see on an inspirational poster, but which has little meaning beyond this. But there’s a biological basis to the benefits that come when curiosity replaces defense. When you encounter a threat—and yes, your brain codes professional disruption as a threat—your amygdala initiates a cascade of responses that evolved to increase the chances of physical survival. Cortisol floods your system. Your prefrontal cortex, responsible for complex reasoning and emotional regulation, is sidelined. You are literally thinking with a less sophisticated part of your brain.

This is highly useful if you’re facing a predator. But it’s counterproductive if you’re facing technological change. You can’t outrun an algorithm or fight a neural network. In effect, the physiological responses that once kept us alive now keep us stuck.

But here’s what we find quite remarkable: the same response that manifests as anxiety can also be channeled as curiosity. The physiological signatures are nearly identical—elevated heart rate, heightened attention, increased energy. The difference lies in the narrative we construct around what we’re experiencing.

Lisa Feldman Barrett’s groundbreaking work on constructed emotion shows that our brains are essentially prediction machines, constantly generating models of what might happen next based on past experience.⁹ When those predictions encounter radical uncertainty—like an AI outperforming our expertise—we

⁹ Lisa Feldman Barrett, *How Emotions Are Made: The Secret Life of the Brain* (Boston: Houghton Mifflin Harcourt, 2017).

experience what she calls “prediction error.” That uncomfortable feeling isn’t a bug; it’s a feature. It’s our brain recognizing that its models need updating.

And this is where curiosity enters as a neurological intervention. When we consciously choose to approach uncertainty with interest rather than resistance, we activate different neural pathways. To get technical for a second, the anterior cingulate cortex, which monitors for conflicts and uncertainties, can trigger either defensive responses (via the amygdala) or exploratory responses (via the dopaminergic reward system). The choice—and it is a choice, even if it doesn’t always feel like one—determines not just how we feel but how capable we are of responding adaptively.

Matthias Gruber’s lab at UC Davis has shown something even more profound. Curiosity doesn’t just help us learn about the specific thing we’re interested in. When we’re in a state of curiosity, our brains show enhanced memory for all information encountered during that state, even entirely incidental details.¹⁰ The curious brain is literally in a different and more capable state than the defensive brain.


Returning to the fictional vignettes we began with, when Samir chose to talk with Amit instead of marshaling defenses, he opened himself to entirely new patterns of thinking about value creation, collaboration, and competitive advantage.

When Lia chose to share her fears with her colleagues instead of projecting false confidence, she created space for collective exploration rather than individual defense. In each case they demonstrated that resisting their knee-jerk biological responses and harnessing curiosity opened up new possibilities.

¹⁰ Matthias J. Gruber, Bernard D. Gelman, and Charan Ranganath, “States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit,” *Neuron* 84, no. 2 (October 2014): 486–496, DOI: 10.1016/j.neuron.2014.08.060

Curiosity Loop

Building on this, we'd like to introduce you to the Curiosity Loop. This isn't our invention—it's a formalization of patterns and behaviors that reflect people successfully navigating technological disruption. It's what Samir did instinctively when he answered his phone instead of calling an emergency meeting. It's what Lia practiced when she turned her students' AI experiments into a collective inquiry. And it's what you can start practicing today, regardless of where you are in your own AI journey.



The loop has four movements, each building on the last:

1. Notice what you're experiencing without immediately categorizing it as good or bad, threat or opportunity. This is harder than it sounds. Our brains are categorization machines, instantly sorting experiences into buckets based on past patterns. But categorization closes off possibilities.

When you notice without categorizing, you maintain what Zen practitioners call "beginner's mind"—the capacity to see what's really there rather than what you expect to see.

2. Question what's actually happening versus what you're assuming. This is where intellectual humility becomes surprisingly powerful. Every assumption you hold about AI, about your expertise, about what's possible or impossible—each one is a

hypothesis waiting to be tested. The quality of your questions determines the quality of your exploration.

3. Experiment with one small exploratory action. The emphasis here is on small. You don't need to transform your entire practice and your habits overnight. You need to take actions that generate real data about how AI might relate to your work. Use a tool. Have a conversation. Create something. The key is moving from thinking to experience.

4. Reflect on what you discovered. Not what you expected to find—what actually happened. What surprised you? What assumptions got challenged? What new questions emerged? This reflection phase is where learning crystallizes and where experience becomes insight.

Then—and this is crucial—begin again. The loop isn't a one-time process. It's a practice, a way of being in relationship with uncertainty that becomes more natural with repetition.

The Curiosity Loop may feel simple and intuitive. But it can be transformative. As one small example, I (Andrew) work almost daily with colleagues who are uncertain about how to wrap their heads around ChatGPT and other AI platforms, yet remain hesitant to engage with them until they feel that they know what they're doing. Instead, I encourage them to embrace their curiosity—by noticing what they are feeling, asking questions, and experimenting for themselves, all while engaging in reflection. It's a deceptively simple process and one that can break the dam of fear and indecision. And, in doing so, open up possibilities rather than shutting them down.

Walking the Loop

To explore this further, let's follow Samir and Lia as they navigate their respective challenges, paying attention to how the Curiosity Loop transforms threat into possibility.

Samir's Journey:

Notice: Standing at that window in Dubai, watching the morning traffic surge below, Samir first had to take notice of his actual experience without immediately jumping to solutions. His chest was tight. His thoughts were racing. His identity as a successful venture capitalist—someone who picks winners—was under threat. He felt the urge to call his partners, to craft a defensive strategy, and to explain (or at least construct a rationalization for) why this open-source disruption was an anomaly.

But he also noticed something else: a flicker of recognition. A feeling he'd felt before. This was when he was twenty-five, working at a traditional venture firm and watching crypto entrepreneurs build entirely new financial systems outside the established framework. He'd chosen curiosity then and left to start his own fund. Now, at thirty-four, he was on the other side of the disruption equation.

Question: Back at his desk, Samir opened a blank document. Not a strategy doc or a crisis management plan. Just a blank page for questions. They came slowly at first, then in a flood:

What if our entire investment model is wrong?

What if proprietary AI isn't a moat but a limitation?

What if open-source always wins in the long run?

What does this kid see that we're missing?

What problems is he solving that we didn't even know existed?

What if disruption is just another name for market intelligence?

Each question loosened something in his chest. He wasn't solving the problem yet, but he was changing his relationship to it.

Experiment: The phone call with Amit was Samir's first experiment. Not an email through lawyers or a formal inquiry through channels. A direct conversation between founders. The experiment was as much about Samir's own capacity for vulnerability as it was about gathering information.

“Why would you want to meet?” Amit had asked, genuinely puzzled. “I just made your investment worthless.”

“Maybe,” Samir said. “Or maybe you just showed me where the real value is.”

Reflect: Flying to Bangalore the next day, Samir reflected on what had already shifted. The fear was still there—his LPs would need answers, his portfolio company would need support. But alongside the fear was something else: excitement. For the first time in years, he didn’t know what came next. Instead of terrifying him, it reminded him why he’d gotten into venture in the first place.

Lia’s Journey:

Notice: Standing before her colleagues in that multipurpose hall, Lia noticed the layers of resistance in the room. But she also saw the fear beneath the resistance. Mr. Tan’s crossed arms were protective, not aggressive. Mrs. Krishnan’s grip on her lesson plans was about identity, not stubbornness.

She noticed her own urge to be the expert, to have answers, to make everyone feel better. But that would just be a performance. What the moment called for was something rarer: shared uncertainty.

Question: “What are you most afraid AI will take from your teaching?” Lia asked the room. The question hung in the air, subversive in its directness.

Then Mr. Tan raised his hand: “The moment when a student suddenly understands. That spark in their eyes—can AI do that?”

“Beautiful question,” Lia said. “What else?”

The dam broke:

“The ability to know when a student is struggling but won’t ask for help.”

“The way I can adjust my teaching style in real-time based on the energy in the room.”

“The relationships. The trust. The human connection.”

“My job. Let’s be honest—my actual job.”

Each fear, once named, became less powerful. Not less real, but less threatening.

Experiment: “Let’s try something,” Lia suggested. “Everyone take out your phone and go to ChatGPT. I want you to make it fail at teaching something from your subject. Really fail. Make it confidently wrong.”

The room’s energy shifted. Teachers who’d been defensive moments before were now actively engaged, trying to break the thing they feared. Laughter rippled through the hall as they shared their results:

ChatGPT insisting that Singapore was founded in 1965 (it was 1819—it became an independent nation in 1965)

An AI-generated math proof that “proved” $2+2=5$

A literature analysis claiming Shakespeare wrote in the 21st century

“Now,” Lia said, “ask it to explain the same concept beautifully. Like your best teaching day.”

The laughter faded into something more thoughtful. The explanations were good—clear, well-structured, even elegant. But they lacked something ineffable.

Reflect: “What did you notice?” Lia asked.

Mrs. Krishnan spoke first: “It’s like having a teaching assistant who’s read every book but never met a student. Useful for preparation, but it can’t read the room.”

Mr. Tan added, “I made it fail by asking for context-specific examples. It doesn’t know our students, our community, our specific challenges.”

“So what might that mean for how we use these tools?” Lia prompted.

The conversation that followed was remarkable. Teachers who’d entered the room defensive were now investigators. They began to see AI not as replacement, but as a tool that might free them to do more of what only they could do: build relationships, provide context, read the room, inspire.

The Compounding Effects of Curiosity

Curiosity, it turns out, compounds in ways that defensive strategies never can. When you defend against disruption, the best outcome is maintaining the status quo. When you get curious about disruption, you open possibilities that didn't exist before.

And this isn't just a feel-good philosophy—it's something we both observe through our work. Returning to our protagonists, every time Samir completed a Curiosity Loop, his understanding deepened, and his options expanded. The initial question ("What if open-source always wins?") led to experiments (meeting Amit), which in turn led to insights (different dynamics for different markets) and ultimately to new questions ("What if we fund open-source development and build enterprise features on top?").

Within three weeks, what started as a threat had transformed into a strategic opportunity. Samir's fund didn't abandon its proprietary investments. Instead, it developed what he called a "full-stack thesis"—investing across the entire spectrum from open-source foundations to enterprise applications.

"The threat wasn't that open-source would win," Samir explained to his LPs. "The threat was thinking in binary terms—proprietary or open, us or them, win or lose. Curiosity showed us it was 'and,' not 'or.'"

The Social Architecture of Curiosity

From both of our perspectives, one of the most underappreciated aspects of curiosity that we see is its social dimension. We tend to think of curiosity as an individual trait or practice, but our experience suggests it's more accurately understood as a social phenomenon. Curiosity spreads through networks. It amplifies through interaction, and it deepens through collective exploration.

This can be seen as we (metaphorically) watch what happened in Singapore after Lia's initial training. Lia's colleague Mr. Tan started an "AI Fails Club" as a lunchtime gathering where teachers

brought their worst AI outputs. And despite the tongue in cheek name, what began as therapeutic venting evolved into sophisticated analysis. Teachers discovered patterns in how and why AI failed which led to insights about their own expertise—and ways to effectively leverage and utilize AI tools that they then adopted.

Mrs. Krishnan, who'd entered the process gripping her handwritten lesson plans, became one of the most creative users of AI. She discovered that Claude (the AI she had decided would be her go-to app) could generate twenty differentiated versions of the same essay prompt in minutes—work that used to take her hours. “I’m not being replaced,” she realized. “I’m being freed to do what I actually care about: talking with students about their ideas.”

The Curiosity Dividend

When we're curious, our brains literally change how they function. Neuroscience shows that curiosity activates the dopaminergic reward system and enhances memory—not just for what we're curious about, but for all information encountered during that curious state.¹¹

This has profound workplace implications. Harvard Business School research for instance found that organizations fostering curiosity are better able to adapt to uncertain market conditions and generate more-creative solutions.¹²

Similarly, a study found that engaged students are over twice as likely to excel academically.¹³ In other words, curiosity isn't only about learning—it's about thriving through change.

Mr. Tan and Mrs. Krishnan started their own curiosity journeys. But the insights they gained didn't emerge in isolation. They came from teachers exploring together, sharing discoveries, and building on each other's experiments. Curiosity created what

¹¹ Gruber, “States of curiosity modulate hippocampus-dependent learning.”

¹² Francesca Gino, “The Business Case for Curiosity,” *Harvard Business Review*, September 1, 2018.

¹³ Gallup, “School Engagement Is More Than Just Talk,” Gallup Student Poll (2018), as cited in Gradient Learning, “Student Engagement,” *Gradient Learning*, 2023, accessed September 12, 2025.

we might call a “learning field”—a social space where exploration was not just permitted but actively encouraged.

When we observe or relate to someone navigating uncertainty with curiosity rather than fear, our brains don’t just watch—they respond. For instance, Lia’s vulnerability in admitting her own fear created what any good founder knows is essential: psychological safety. Others could now practice curiosity about their own uncertainty without the usual performance anxiety. Exploration together also activates what researchers (maybe optimistically) call “collective intelligence”—a phenomenon where groups supposedly develop problem-solving capacities that exceed the sum of their parts. The original research here claimed to find a general factor for group performance, similar to IQ for teams,¹⁴ although subsequent studies have been less certain. What happens in reality is both more straight forward and more actionable than initially assumed: diverse perspectives, structured interaction, and shared attention can surface patterns that individuals may miss. No neural magic required—just good process design.

In Samir’s case, his curiosity about open-source disruption didn’t remain a private anxiety. He brought it back to his partners, creating what they called “Curiosity Sessions”—weekly gatherings where they explored threats to their portfolio not with defensive strategy but with genuine interest. These sessions worked because they addressed a common coordination problem faced by every leadership team: how to turn “I don’t know” into a starting point for exploration rather than an admission of weakness.

¹⁴ Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone, “Evidence for a collective intelligence factor in the performance of human groups,” *Science* 330, no. 6004 (October 2010): 686–688, DOI: 10.1126/science.1193147.

Embracing Curiosity

As you develop your own ways of embracing curiosity, you'll discover it has depths and dimensions that aren't immediately apparent. Surface curiosity—"How does this AI tool work?"—is valuable but limited. Deep curiosity—"What does this mean for human identity, purpose, and possibility?"—is where transformation happens.

Here, research by Todd Kashdan reflects six distinct dimensions of curiosity, each serving different functions:¹⁵

Joyous Exploration: The pleasurable pursuit of novelty and challenge

Deprivation Sensitivity: The need to resolve uncertainty and fill knowledge gaps

Stress Tolerance: The ability to handle the anxiety that comes with uncertainty

Overt Social Curiosity: Interest in learning about other people's thoughts, feelings, and behaviors through direct engagement

Covert Social Curiosity: Interest in discovering what others are like through indirect means (observing, overhearing)

Thrill Seeking: The willingness to take physical, social, or financial risks for experience

What's important is that these aren't fixed traits. They're capacities that can be developed through practice. Samir, initially driven by Deprivation Sensitivity (needing to understand the threat), gradually developed greater Stress Tolerance and Overt Social Curiosity. Lia, naturally high in Overt Social Curiosity, strengthened her Joyous Exploration as she discovered the creative possibilities of AI.

¹⁵ Todd B. Kashdan, David J. Disabato, Fallon R. Goodman, and Patrick E. McKnight, "The Five-Dimensional Curiosity Scale Revised (5DCR): Briefer subscales while separating overt and covert social curiosity," *Personality and Individual Differences* 157 (2020): 109836, DOI: 10.1016/j.paid.2020.109836.

The key is recognizing which dimensions serve you well in different contexts and consciously developing those that don't come naturally. If you're high in Joyous Exploration but low in Stress Tolerance, you might abandon explorations when they become uncomfortable. Suppose you're high in Deprivation Sensitivity but low in Social Curiosity. In that case you might miss crucial insights that come from exploring with others.

This all sounds great in theory. However, as we conclude this chapter and look toward intentional AI, there is a problem: in professional contexts, admitting uncertainty can be perceived as a weakness. We've built entire cultures around projecting confidence, having answers, being the expert. Curiosity—with its inherent admission of not knowing—can sometimes feel like career suicide. Yet this fear misreads the evidence. Research from multiple institutions reveals a consistent pattern: curious people are more likely to make better decisions because they're less likely to fall prey to confirmation bias,¹⁶ generate more creative solutions,¹⁷ and build stronger relationships because genuine questions create genuine connections.¹⁸

Back to Samir's investment fund: By choosing curiosity over defensiveness, they didn't just survive the open-source disruption—they positioned themselves at the forefront of a new investment paradigm. Their “full-stack thesis” became a differentiator in a crowded market. LPs who initially questioned Samir's trip to Bangalore were soon asking for introductions to his open-source network.

¹⁶ Dan M. Kahan, Asheley Landrum, Katie Carpenter, Laura Helft, and Kathleen Hall Jamieson, “Science Curiosity and Political Information Processing,” *Political Psychology* 38, no. suppl. 1 (2017): 179–199, DOI: 10.1111/pops.12396.

¹⁷ Jay H. Hardy III, Alisha M. Ness, and Jensen T. Mecca, “Outside the Box: Epistemic Curiosity as a Predictor of Creative Problem Solving and Creative Performance,” *Personality and Individual Differences* 104 (2017): 230–237, DOI: 10.1016/j.paid.2016.08.004.

¹⁸ Karen Huang, Michael Yeomans, Alison Wood Brooks, Julia Minson, and Francesca Gino, “It Doesn't Hurt to Ask: Question-Asking Increases Liking,” *Journal of Personality and Social Psychology* 113, 3 (2017): 430–452, DOI: 10.1037/pspi0000097.

“Curiosity became our moat,” Samir reflected six months later. “Not proprietary algorithms or defensive strategies, but the practice of engaging genuinely with whatever threatened our assumptions.”

The irony is startling: in a world obsessed with having the correct answers, the real competitive advantage might come from asking better questions.

Your Personal Practice

As we close this chapter, we would like to leave you with a concrete practice that you can start today. Not tomorrow, not when you feel ready, but right now, with whatever AI-related uncertainty is present in your life.

The Practice:

Identify your edge:

What aspect of AI makes you most defensive or anxious? Is it the threat to your expertise? The speed of change? Uncertainty about the future? Be specific.

Run a Curiosity Loop:

Notice: Where do you feel this anxiety in your body? What thoughts arise? What actions do you want to take? Just observe without judgment.

Question: Generate three “What if...?” or “How might...?” questions about this edge. Make them open and exploratory, not defensive.

Experiment: Choose the question that intrigues you most and design a small experiment. Use an AI tool in a new way. Have a conversation with someone who sees it differently. Create something small.

Reflect: What surprised you? What assumption got challenged? What new question emerged?

Share your learning:

Find one person—a colleague, friend, or family member—and share what you discovered. Notice how the conversation deepens or shifts your insight.

Schedule the next loop:

Curiosity is a practice, not a one-time event. When will you run your next loop? Put it in your calendar.

Remember: the goal isn't to become fearless or to love every aspect of AI transformation. The goal is to cultivate a more flexible and adaptive relationship with uncertainty. To transform the energy of threat into the energy of exploration.

Next Steps

As you begin your curiosity practice, you'll likely notice something that both excites and unsettles. Each answer generates new questions, each experiment reveals new possibilities, and each reflection deepens the mystery. This is as it should be. In a world of exponential change, the capacity to remain curious—to find joy and purpose in not knowing—is actually a form of wisdom.

But curiosity alone can become directionless wandering. That's why our next chapter explores intentionality—the practice of channeling curiosity toward outcomes that matter. Because the question isn't just “What's possible with AI?” but “What do we want to make possible?”

Wei Lin discovered this as he moved beyond simply exploring AI art tools. His third portrait—the one that now hangs in Lia's classroom—didn't emerge from random experimentation. It came from intentional exploration, from knowing what he was seeking: not a better version of himself, but a true version of his becoming.

The same principle applies whether you're navigating venture disruption like Samir, educational transformation like Lia, or your own personal and professional evolution. Curiosity opens doors.

And intentionality helps you choose which ones matter—which is where what we grapple with in the next chapter.

HANDS-ON CARD

Running a mini Curiosity Loop

Pick an AI headline from today that you almost scrolled past. Run a mini Curiosity Loop.

1. **Notice:** Think about how it makes you feel.
2. **Question:** Write down three “How might we...?” prompts.
3. **Experiment:** Spend 5 minutes testing one prompt in a free tool.
4. **Reflect:** Jot down one surprise & one next step.

CHAPTER 2

INTENTIONAL AI

“Everything rests upon the tip of intention.”
—*Buddhist teaching*

The Morning Everything Changed

The morning light caught the glass walls of Priya Sharma’s Silicon Valley loft with that particular quality unique to Northern California. Her team at her startup Namesea clustered around the standing desks in their daily ritual, the third-wave Ethiopian single-origin espresso she’d learned to love during her Stanford days filling the air with its complex aromatics. Twenty-three months. That’s how long they’d been building their customer service LLM, through seven pivots, three “near-death” experiences with funding, and more all-nighters than anyone’s circadian rhythms should have to endure.

“We’re green across the board,” announced James, their lead engineer, his fingers dancing across the touch-display with the unconscious grace of someone who’d been coding since middle

school. Response times under 200 milliseconds. Accuracy scores that made their competitors' models look sluggish in comparison. The demo last week had made their lead investor—notorious for his poker face—actually smile.

Priya inhaled the steam from her cup, letting the familiar ritual ground her in this moment of anticipated triumph. Her daughter Aanya's drawing from last week still clung to the refrigerator—a crayon family portrait labeled “Mama's Robot Friends.” Four years old and already normalizing AI. What sort of world was she inheriting?

The notification cascade began at 7:23 a.m. First, the gentle buzz of her Apple Watch. Then the crescendo—phone, tablet, the main display. Their PR firm's message cut through the morning's optimism: “URGENT: TechNova's customer service bot tells user to ‘end your worthless life’ after a refund dispute. Major outlets picking up. Their stock down 12% premarket.”

Someone's mug—the one with “Move Fast and Fix Things” on it—hit the wood floor. The sound seemed to echo through the loft forever. It wasn't that they had any relationship with TechNova, but the fallout would resonate through the whole LLM customer service sector.

Priya watched her team's faces cycle through the stages of startup grief as they read the incoming news: denial (that could never happen to us), anger (how could they be so careless), bargaining (our architecture is entirely different), depression (are we building something harmful?), and finally, the most dangerous stage—rationalization (we just need better guardrails).

“Pull up their postmortem,” she said quietly.

The Weight of Intention

Half a world away in the marble-columned sanctuary of São Paulo's Biblioteca Mário de Andrade, Mateo Oliveira sat surrounded by the comfortable chaos that increasingly defined the final semester of his computer science degree. The afternoon sun slanted through the tall windows, casting geometric patterns across tables worn smooth by generations of elbows and ambitions.

His laptop screen glowed with the project that had consumed the last six months of his life—an AI research assistant for the library that could parse academic papers in Portuguese, Spanish, and English, synthesizing findings across languages with an elegance that still surprised him. The idea had come during a frustrating night spent researching distributed systems implementations across Latin American universities, hitting paywall after paywall, and language barrier after language barrier.

The library was enveloped in an atmosphere of productive quiet—not the sterile silence of tech offices but the almost-audible susurrations of minds at work. Pages turning, keyboards clicking softly, the occasional cough or scraped chair. Mateo had spent so many hours here that Jorge, the security guard who'd worked there for thirty years, now brought him *cafézinho* in the afternoons without being asked.

He was deep into debugging a particularly stubborn issue with citation formatting when he felt someone's eyes on his screen. An elderly man stood behind him, squinting through thick glasses at the cascade of code. Professor Cardoso—Mateo recognized him from his pilgrimages to the historical archives.

"Desculpe," the professor said, his voice carrying the soft authority of someone who'd spent a lifetime teaching, "but this looks fascinating. Some sort of translation program?"

Before Mateo could answer, a small hand tugged at his sleeve. A girl, maybe seven or eight, clutching a book about dinosaurs that had seen better days. Her school uniform—from a local public

school, carefully mended at the collar—suggested she’d come straight from class, probably waiting for a parent who worked in the nearby government offices.

“Moço,” she said in that direct way children have before the world teaches them deference, “will it read to me?”

Mateo looked from his screen—dense with abstractions about natural language processing and transformer architectures—to her hopeful face. The question was simple. The implications were not.

“It’s not really for stories,” he started to explain, then stopped. Why wasn’t it? Who had decided that AI tools for research couldn’t also help a child discover the joy of reading? Who was he building this for, really?

Professor Cardoso leaned closer, interested now. “My granddaughter has dyslexia,” he said quietly. “Reading is hard for her. Every night, she struggles. If this could help...”

The weight of the moment settled on Mateo’s shoulders—not crushing, but substantial. This wasn’t about algorithms or benchmarks anymore. This was about Ana (he’d learn her name later) and her dinosaur book, about Professor Cardoso’s granddaughter struggling with words that seemed to swim on the page, about every person who’d been locked out of knowledge by circumstance, language, or learning difference.

“Let me show you something,” Mateo said finally, minimizing his code and opening a simple text interface. “What’s your favorite dinosaur?”

Speed Trap

In the above vignettes, Priya faces the lure of sacrificing intention for speed. Mateo discovers he never questioned his intentions at all. Two very different stories but both reveal the same truth: unclear purpose leads to vague outcomes.

We’ll come back to Priya and Mateo in a moment. But first we need to discuss the various ways intention can get lost. For Priya

and many in the tech world, it's speed—the relentless pressure to ship before thinking. For Mateo, it was an assumption—never questioning who his work was really for. Both lead to the same dangerous place: building without purpose.

Let's start with speed. Not the kind measured in milliseconds or deployment cycles, but the haste that drives organizations to adopt AI as if their life depends on it. This urgency—this sense that every moment of deliberation is a moment lost to competitors—creates what we've come to recognize as a particularly dangerous pattern in technology adoption.

The symptoms are everywhere. A CEO reads about ChatGPT over breakfast and demands an "AI strategy" by lunch. A school board, panicked by neighboring districts' announcements, rushes to implement AI tutors without asking whether their students need or want them. A hospital, fearful of appearing outdated, deploys diagnostic AI that neither doctors nor patients requested or consented to.

This is tool-lust in its purest form—the belief that adopting the latest technology is equivalent to progress, that movement in any direction beats standing still, and that being first matters more than being right.

It's a belief that tripped up the online property listing company Zillow in the U.S. Here was a company that seemingly had every advantage—brilliant engineers, vast data sets, deep pockets. Their AI-powered domestic real estate algorithm was a technical marvel, predicting home values with impressive accuracy.

But in their rush to be disruptive, they started using their algorithm to buy and sell houses using their own funds. Yet they never established guardrails for market manipulation. And they never measured impact on housing affordability or community stability.

The result? A \$304 million quarterly loss, the shuttering of an entire division, and thousands of employees laid off—not because of a market downturn or competitive pressure, but because of a

strategic miscalculation. The algorithm performed exactly as it was designed to—maximizing purchasing volume. What was missing wasn't technical sophistication but strategic wisdom.

Thinking Before You Leap

We've both learned from watching spectacular successes and expensive failures that the difference between AI that enhances human capability and AI that diminishes it isn't found in the sophistication of the neural networks or the size of the training data, but in something far more fundamental: the clarity of intention that guides every decision from conception to deployment.

In my (Andrew's) work on responsible innovation, I make a clear distinction between functional and social requirements. Functional requirements are all about what a system does—respond to queries, process transactions, and generate text. These are the specifications that fill whiteboards during engineering meetings, the benchmarks that determine bonuses, the metrics that make it into TechCrunch headlines.

Social requirements are different. They define what a system means in a human context—whether it preserves dignity, promotes equity, enhances rather than replaces human judgment, or causes harms that aren't easy to capture through simple cause and effect. These are harder to quantify, easier to ignore, and essential to get right. And they demand intentionality in how technologies are developed and used.

And many AI failures—or failures waiting to happen—are, in my experience, social failures masquerading as technical ones. Which brings me to a paraphrase of an over-worn quip from Dr. Ian Malcolm in the 1993 movie *Jurassic Park* (one I use only because I've written extensively on it¹⁹): Beware of technologists

¹⁹ Andrew Maynard, *Films from the Future: The Technology and Morality of Sci-Fi Movies* (Nashville: TMA Press, 2018).

who are so preoccupied with whether they can, they don't stop to think if they should. Intention matters.

This insight profoundly shifts how we think about AI development. Ethics, responsible innovation and intentionality aren't—or shouldn't be—just compliance checkboxes or a PR strategy. Instead, they form part of the metaphorical load-bearing structure that determines what can be built safely and beneficially—and what cannot. Just as a building's foundation constrains and enables what rises from it, the intentions we encode into AI systems determine their ultimate impact on human lives.

Introducing the Intent Map

This understanding led us to develop the Intent Map—not another framework to add to the ever-larger pile accumulating in corporate strategy documents, but a thinking tool that makes values visible and choices conscious before momentum (and the actions of others) makes them for you.

The Intent Map



Picture a simple grid, consisting of two lines that create four spaces. Nothing fancy, and nothing that requires a certificate to understand. You could sketch it on a napkin, a whiteboard, or the back of your hand. But within this simplicity lies profound power to shape outcomes.

Each quadrant addresses a fundamental question:

Values occupy the upper left, the position of primacy. These address the question: What do we refuse to compromise? Not aspirational statements crafted by marketing teams or motivational posters gathering dust in break rooms. These are the non-negotiables that would prompt you to pull the plug even if every metric were green, every investor were satisfied, and every customer were satisfied.

Desired Outcomes are entered in the upper right quadrant. These address the question: What specific change do we seek? Not vague improvements or percentage gains that sound good in quarterly reports, but concrete transformations in how people work, live, and relate to each other. The kind of change you could photograph, tell stories about, and feel in your bones.

Guardrails anchor the lower left, focusing on the question of: Where do we draw hard lines? These are the boundaries beyond which our system will not extend, regardless of user requests, market pressure, or engineering elegance. Think of them as the walls of a highway—not suggestions but barriers that prevent catastrophic deviation. While values capture what you won't compromise on, guardrails set out what you won't do.

Metrics complete the grid in the lower right. These capture the question: How will we know we're succeeding—or failing? Not just performance indicators that make dashboards look professional, but measurements that reflect our values in action, our outcomes in progress, and our guardrails holding firm.

The power lies not in the individual quadrants but in their interaction. Values without metrics become empty rhetoric—the kind of mission statements everyone ignores. Metrics without values optimize for the wrong outcomes—engagement that becomes addiction, efficiency that becomes dehumanization. Guardrails without clear outcomes become arbitrary restrictions that everyone works around. And outcomes without guardrails achieve goals by any means necessary.

WHAT COMES NEXT

AI and the Art of Being Human is just the first step on a journey into a future where AI has the potential to enhance every aspect of who we are without diminishing it.

To continue the journey started here, please check out the resources and opportunities at aiandtheartofbeinghuman.com, or scan the QR code below:



We're looking forward to you joining us!