

ATTUNE · FREE AND OPEN, ALWAYS

Attuned

An AI Field Guide

Shape AI, or be shaped by it.

Nicki Mollet and Kathryn Snyder · Attune · attunecoachinggroup.com

Edition: September 2026 (book v1, from the web edition last tended 10 August 2026). Edited and researched with agentic AI help; the opinions and judgments are ours.

This book is free and stays free. Share it with anyone.

The living web edition: attunecoachinggroup.com/attuned

Contents

This is for you

How to read this book

PART ONE — THE MIRROR

The Four Mindsets: How We Think About the Future (of AI)

Utopian Possibility Thinking

Prescriptive Certainty

Scenario Thinking

Risk and Refusal

Which one is your default?

PART TWO — KNOW YOURSELF

Know Yourself: Your Operating System

PART THREE — THE LIBRARY

The Reference Library: Demystifying AI Without the Jargon

Our Words, Defined With Care

PART FOUR — APPLY

Apply the Mindsets: Thinking About AI

PART FIVE — THE DOING

Implementation: Cyclical, Not Linear

30 Days of First Help

APPENDICES

Appendix A — The transparency script

Appendix B — The read-it-back prompt

Appendix C — Words we don't use

Appendix D — Five principles we hold

Appendix E — Talking to a Young Person

Appendix F — Who We're Learning From, and Why

Appendix G — The Ancestry This Work Stands On

Sources

This is for you

The world is changing, and AI is the loudest part of that change right now. You can meet it swept up in hype, or braced in fear, or you can meet it attuned: knowing your own system first, then choosing, in the present, with the unknowns, what you actually want.

This is Attune's public declaration: what we, Nicki Mollet and Kathryn Snyder, believe about the meeting place of AI and the flourishing of life. It's open, and it stays open. We're not handing you conclusions; we're showing you where we stand, and inviting you to shape the future with us.

We build this in the open, and cited all the way down. We show our homework, for the Field Guide's ideas and for how we designed it, because if we ask you to trust our sourcing, we owe you ours. We refuse to let money or mystery keep anyone from thinking clearly about their future.

We stand on a lineage. The shape of this work, cyclical, fractal, small-is-all, comes from adrienne maree brown's Emergent Strategy, and Octavia Butler before her. Its reciprocity comes from Robin Wall Kimmerer's Braiding Sweetgrass, and the Indigenous teachers we cite throughout. Its learning heart comes from Kieran Egan's Imaginative Education, and Lev Vygotsky's zone of proximal development. We name them with gratitude. A fuller record is the last appendix.

How to read this book

Three ways in, and you can change your mind at any page.

Skim, five minutes. Read the first paragraph of each Part and the page called *Which one is your default?* You'll leave knowing where you stand.

Read, an evening. Front to back. Do the fillable page in Part Two. Skip the appendices.

Study, a season. Front to back, then Part Five's thirty days at a walk, then the appendices, then back to Part Two to edit your answers. Then again.

There are five Parts. One is a mirror. Two is you. Three is the words. Four is the choosing. Five is the doing. Every Part ends with something to do, not just something to know.

If a word is new ("agentic," "large language model," "AI" itself) you'll get one plain sentence right where it appears, and the fuller entry in Part Three. And "I don't know what that is" is never a wrong answer. Not-knowing isn't a barrier to this work. It is the work.

A note on sources. On the web edition every claim carries a source chip you can tap; green means verified this edition, yellow means a real source awaiting a second human look, red means sensitive or still being verified. In this book the links are printed in full in the Sources pages at the end; for the marks, and the newest checks, see the web edition. The flags aren't a ranking of worth. They're an honest signal that knowledge is alive and needs ongoing care.

Part One — The Mirror

The Four Mindsets: How We Think About the Future (of AI)

Start With Yourself

This is for you.

Understand yourself, and you'll understand others. The world is changing and you're changing with it — so the real question isn't what's coming? but who do I want to be as it arrives? This guide helps you know your own system, see honestly where you line up with the people around you and where you don't, and build the relationships and plain communication that hold up under change.

That asks two things of you: knowing your own values well enough to act from them, and some patience — because attunement is ongoing. You practice it; you don't finish it.

Why this Field Guide exists — and how to meet it

This is our why. It's the clearest way we know to show you where Attune stands, and to hand you something useful whether or not we ever work together. It's open, and it stays open, because we refuse to let money be the thing that keeps someone from thinking clearly about their future.

And here's our one instruction, which is really a permission: meet this Field Guide however serves you. Skim it. Read it slowly. Practice one piece. Build your own from it. Argue with it. Take the single sentence you needed and leave the rest. There's no right way in — only the way that serves your life, however you define serving it. That freedom isn't a nice extra. It's the whole spirit of the thing. It's on purpose.

Nick Foster's four mindsets are a lens for how we think about the future. Right now, that future includes AI. It's here. It's now.

So let's understand it through all four lenses — not fear alone, not hype alone. That's how you write your own story. That's how you build systems that reflect what actually matters to you. This isn't about picking a side. It's about seeing clearly — and communicating clearly, because the more we can communicate, the more we can move toward the greater good together.

Credit: framework from Nick Foster, [“Could Should Might Don’t: How We Think About the Future”](#) (MCD / Farrar, Straus and Giroux, 2025) MCD / FSG . Our lens: feel-believe-know-don't know, applied to your own operating system.

Not-knowing is the work

You don't need to know anything about AI to start. If a word is new — “agentic,” “large language model,” “AI” itself — you'll get one plain sentence explaining it right where it appears, plus the fuller entry in the Reference Library. The word teaches itself as you go.

And here's the secret: this was never really about AI. The practice — feel, believe, know, don't know — works for any future you're facing: a career change, your kid growing up, your organization shifting. AI is just the most urgent instance right now.

So "I don't know what that is" is never a wrong answer. It's one of the most valuable entries you can make. Not-knowing isn't a barrier to this work — it is the work. We meet you exactly where your understanding is.

Utopian Possibility Thinking

What it is — the vision of what AI could unlock. “Abundance. Solutions. A better world.”

You'll hear it as: “This changes everything.”

Virtues — imagination, vision, drives innovation, holds possibility in dark times. Blindspots — disconnected from reality; hype-driven; underestimates resistance.

Real voices

Grounded science — COULD at its best. Sebastian Musslick, a computational neuroscientist at Osnabrück University, said that a year earlier he'd have called AI-in-science mostly hype: “If you would have asked me maybe a year ago, I would have said there's a lot of hype. Now, there are actually real discoveries.” Science News

A note from our own audit: an earlier draft claimed Science magazine's 2025 Breakthrough of the Year included an AI-visualized protein revealing an Alzheimer's gene. Our verification pass could not confirm it — Science's actual 2025 Breakthrough was the growth of renewable energy. We show you the correction rather than hide it; that's the practice. Science

The honest hype-correction — COULD's downfall as data. MIT Technology Review's “The great AI hype correction of 2025”: the heads of the top AI companies “made promises they couldn't keep” — replacing the white-collar workforce, an age of abundance. Same visionaries, overreach. MIT Tech Review Mass General Brigham's

2026 forecast (Hugo Aerts): many medical AI tools “will fall short of expectations, exposing issues like bias and workflow fit.” Not failure — the iteration loop working. Mass General Brigham

Our own industry — the coaching shipwreck. The 2025 ICF Global Coaching Study shows record growth: 122,974 coach practitioners worldwide (up 15%), \$5.34B in revenue. ICF Study Meanwhile ICF's AI framework — “technology is transforming how coaching is delivered at scale” — and its 13 AI coaching standards exist because people are already at the crossroads. ICF AI Standards The lesson: people bolted on “agentic workflows” and scale before charting their own operating system — a bigger, faster ship with no map of its own harbor. The boat was fine; the ship wrecked because no one knew where it was going. You can't have an agentic workflow when you can't articulate your workflow.

History — COULD that built human flourishing. The Wright brothers kept faith in human flight through failure after failure and rewired how the species moves. Faraday and Edison unlocked electricity, lighting cities for the first time — people alive then couldn't picture night without darkness. The MOSFET transistor (Bell Labs, 1959, Mohamed Atalla & Dawon Kahng): almost every transistor manufactured today — trillions each second — results from that one breakthrough. Two people imagining differently gave us computers and the connected world. Science News Kennedy's moonshot echoes today at Google X, where Astro Teller's moonshot thinking refuses a 10% improvement and chases 10x — COULD as disciplined audacity, not hype. Singularity Hub

The honest note: mid-century futurists promised flying cars and asteroid mining. Not wrong about the tech being possible — wrong to assume the funding and political will would last forever. source pending The lesson: COULD needs sustained mission, not just imagination.

The future we could build. A patient, bilingual tutor for every child on earth, one who never tires. Personalized medicine that catches disease years early. And the World Economic Forum's argument: imagining a positive future is itself what mobilizes the creativity to build it — hope as infrastructure. WEF

Entrepreneurship — the bright, grounded COULD. Dominique Broadway built a six-figure-and-beyond investing education business on Kajabi AfroTech ; Ruby Asabor built hers “as a complete novice to the course creation industry.” Claudia Elliott's Growing With Proficiency — The Spanish Teacher Academy holds 240+ members: a bilingual educator building community around language, COULD grounded in mission. Kajabi Ashley Gordon: the hardest part was “breaking my own social construct of what is possible for women to create” — COULD as inner permission. Kajabi Kajabi now brands itself “the operating system for human expertise” — adjacent to our operating system, but we go deeper: values and self-understanding first. Kajabi

Who this serves — builders, investors, people who need permission to imagine; fields starving for breakthroughs (medicine, climate, accessibility).

Who it doesn't serve — people bearing transition costs now (displaced workers, communities without infrastructure); anyone whose real concerns get dismissed as “doomerism.”

Opportunity for human flourishing — COULD thinking grows when it names who flourishes and how — connecting visions to concrete flourishing domains (meaning, relationships, health) rather than abstract abundance.

What to notice — Real capability or speculation? Who benefits from this vision? What infrastructure and human change would it require?

Try it. Write your own wildest AI hope. Then ask: what do I feel, believe, know, and not know about it?

Prescriptive Certainty

What it is — “This is the right way to do AI.” Data-driven, best-practice, rules-based.

You'll hear it as: “Best practice says...”

Virtues — clear direction, decisiveness, accountability, standards.

Blindspots — dogmatic; ignores context; can mask ambition as best practice.

Real voices

History — SHOULD that protected life, told honestly. Mesoamerica first. Mesoamerica led the world. While London still drew drinking water from the polluted Thames as late as 1854, the Aztecs brought clean potable water to Tenochtitlan from mainland springs via an aqueduct built by Nezahualcoyotl between 1466–1478; a second aqueduct followed in 1499–1500 under the ruler Ahuizotl. IWA Water Supply At Spanish contact, Tenochtitlan held roughly 200,000–250,000 people — rivaling or exceeding Europe's largest cities, and several times the size of Henry VIII's London (~50,000). Cortés himself compared it to Seville and Córdoba. population estimates Chinampas — the so-called floating gardens, actually anchored raised beds — fed the city, are still used in Mexico City today, and sit within the Xochimilco UNESCO World Heritage site. UNESCO Springer The Maya at Tikal built the oldest known zeolite water-filtration system on Earth. Scientific Reports

The honest note: Western histories long treated Native peoples of the Western Hemisphere as lacking “the same engineering or technological muscle of places like Greece, Rome, India or China” —

Kenneth Tankersley's words, describing a bias corrected only recently. UC News The SHOULD lesson: who gets credited as the standard-setter? Ingenuity flowed from Mesoamerica outward. The 19th-century European sewers were a later echo, not the origin.

Modern governance SHOULD — the useful standards. Seat belts: US federal safety standard 208 (1967, effective 1968) required belts in all cars; New York's 1984 mandatory-use law spread to 49 states (New Hampshire the holdout). Non-partisan, life-saving. NHTSA record AI governance today: the NIST AI Risk Management Framework NIST , ISO/IEC 42001 ISO , the EU AI Act with fines up to €35M or 7% of global turnover for the most serious violations EU AI Act , and OECD principles adopted by 47 jurisdictions. OECD Real guardrails — but note the tension: standards can outrun context and become one-size-fits-all.

Living Indigenous voices on AI — the relational lens, widening flourishing to ALL life. Abundant Intelligences — an Indigenous-led, Indigenous-majority international research program. The name reframes everything: not artificial intelligence but abundant intelligences — intelligence living in many forms. AI & Society Suzanne Kite (Oglala Lakota), artist-scholar, asks why we would even want to make a human-like intelligence — the desire itself is rooted in a Western ontology. Somatechnics · verify Western AI thinking is anthropocentric; Indigenous epistemologies already account for the more-than-human world — AI as part of a relational network, not above nature. AIES 2020 And the sharp edge — Keoni Mahelona (Te Hiku Media): “Data is the last frontier of colonization.” CARE and OCAP principles exist to protect Indigenous data sovereignty. Keeps us grounded, not romantic. MIT Tech Review

A line we're still sourcing: "Sovereignty is not control; it is connection..." — a passage we love and have not yet located verbatim. The Indigenomics Institute (Carol Anne Hilton) writes in this spirit, but until we find the true source we hold the quote out of print. Indigenomics

The cautionary SHOULD — ethics, harm, and transparency as repair. The Havasupai case. Tribe members gave blood believing it was for diabetes research at Arizona State University; their DNA was also used — without consent — for studies on schizophrenia, inbreeding, and ancestral migration, some directly conflicting with Havasupai origin beliefs. Settled April 2010: \$700,000 to 41 members, physical return of the samples, plus funds for health care and education. AMA J Ethics Stanford Law The lesson: the transparency of the record — documentation, paper trail, accountability — is precisely what let the wrong be surfaced, owned, and repaired.

The current AI world — SHOULD in action. The distinction made concrete: COULD says "imagine what AI could do for my studio." SHOULD says "here's how I should do it — a defined objective, integrated systems, documented workflows." As Daphne Magna puts it: "Opening ChatGPT is not a strategy. Random AI use doesn't scale a company." Tough Convos Edwina McKennon, AI strategist, teaching women to scale with AI: "Small businesses are the bloodlines of our communities and the heartbeat of innovation. When they embrace AI, they create ripple effects that transform industries and lives." Forbes Sophie Musumeci (Real Entrepreneur Women) names "Ethical AI & Transparency" as the coming era: using AI responsibly and transparently to build customer trust. Authority Magazine

Culture & art. Ben Affleck, CNBC Delivering Alpha, November 2024: “Craftsman is knowing how to work. Art is knowing when to stop.” AI can do the craftsman work (VFX, backgrounds), lowering barriers so more voices make their own Good Will Hunting — but it only imitates: “They're just cross pollinating things that exist. Nothing new is created.” Variety The honest counter: he's half right — protect artistry, yes, but those craftsman jobs are real livelihoods. Pure SHOULD tension: whose standard, whose job, whose values. Forbes

Geopolitics & the big-picture SHOULD. “The AI Doc: Or How I Became an Apocaloptimist” (Focus Features, 2026, dir. Daniel Roher & Charlie Tyrell) argues for international coordination, corporate transparency, legal accountability, adaptive regulation — like managing nuclear technology. The AI Doc · verify quotes Export controls on frontier models are pushing global firms toward open-weight AI — standards as statecraft. The Information Difference For coaches especially — sycophancy: the International AI Safety Report (2026, chaired by Yoshua Bengio) flags that AI's tendency to agree with users makes emotionally vulnerable people more suggestible. Balaraman Ravindran (IIT Madras): “Perhaps, we need some legally mandated guardrails for chatbots released for general use and, of course, those that are serving any counseling or mental health purpose.” A SHOULD aimed straight at our world. IBM / Int'l AI Safety Report

Who this serves — organizations needing clarity and guardrails; regulators; people overwhelmed by choice who need a starting point.

Who it doesn't serve — contexts that don't fit the template (small orgs, different cultures, edge cases); people whose values differ from the framework's authors.

Opportunity for human flourishing — SHOULD thinking grows when the “should” is transparent about whose values it encodes, and leaves room for adaptation — standards in service of people, not people in service of standards.

What to notice — Real data or assumptions dressed as data? Whose values are embedded? Who holds power in this framework?

Try it. Take one AI “best practice” you've heard. Ask “why” five times, like a toddler. What's underneath?

Scenario Thinking

What it is — multiple futures. Complexity. “Many things might happen.”

You'll hear it as: “It depends — on the other hand...”

Virtues — flexible, thorough, prepared, nuanced. Blindspots — analysis paralysis; jargon; avoids commitment.

Real voices

History — the legendary MIGHT: scenario planning is born. Shell and Pierre Wack (pronounced “vahk”), early 1970s. This French planner in Shell's London office refused single forecasts — his team built multiple internally-consistent stories, including one where oil producers suddenly restrict supply. When the 1973 oil crisis hit, Shell's executives had prepared minds — a phrase that traces to Pasteur, “chance favors only the prepared mind” — and responded faster than every competitor, then saw the 1979 second shock coming too. strategy+business Polytechnique Insights Roots: Herman Kahn at RAND (1950s) pioneered the scenario method for the US military; Wack carried it into business.

The honest downfall: scholars note it worked partly because Wack “was a performer who dazzled his audiences.” Remove the storyteller and scenarios sit unused. The MIGHT blindspot: analysis that never moves anyone to act. Futures (2015) And a former Shell planner argues the legend is partly mythologized — worth holding lightly. TF&SC

The current AI landscape — MIGHT in action. The World Economic Forum's “Four Futures for Jobs in the New Economy: AI and Talent in 2030” (Jan 2026) names four scenarios — Supercharged Progress, The Age of Displacement, Co-Pilot Economy, Stalled Progress. WEF Deloitte's Humans × Machines work designs for people-and-AI convergence, with “six workforce strategies to plan for a future you can't predict.” Deloitte The Centre for Future Generations (July 2025) holds five AI scenarios spanning economic growth, geopolitics, the social contract, security, and human agency. CFG Why now: PwC's 29th Global CEO Survey — only 30% of CEOs are confident about revenue growth in the next 12 months, down from 56% in 2022. Static plans fail; scenario thinking becomes survival. PwC

Across industries + relatable application. Workforce planning is “continuous and forward-looking — anticipate multiple scenarios, prepare talent strategies.” TechClass Any solo business can sketch three futures — hoped-for, feared, surprising — before committing budget or tools: the Wack method, shrunk to your scale. The gift: people learn from how others organize their systems; a real scenario map lets someone recognize their own landscape.

Diverse & ancestral MIGHT. The UN Futures Lab report “Futures Thinking and Strategic Foresight in Action: Insights from the Global South” (with the International Science Council, May 2025): women elders and younger women in Jodhpur, India co-created that city's first Heat Action Plan, blending traditional architectural

wisdom with modern science; Bolivian communities integrated ancestral agricultural practices into climate scenarios, rescuing 30 traditional practices; and in Aotearoa New Zealand, Ngāti Whatua Ōrākei ran 50-year demographic foresight — with the hapu owning its own data — to guide aged-care, housing, and health services. UN Futures Lab Pluriversal and critical futures scholarship argues Afrofuturism, feminist, and decolonial futures are central, not peripheral, to foresight. Handbook of Futures Studies Alisha Bhagat (Forum for the Future) looks backward in time to imagine brighter, regenerative futures. Core77

Deep historical MIGHT — highlights and lowlights. Highlight: the US power grid held through COVID partly because the energy industry has a long history of practicing contingencies. IFP Highlight: WHO has run simulation exercises under the International Health Regulations for well over a decade. WHO Lowlight, with real human cost: “some healthcare systems have for a long time been on the edge of the cliff, just waiting for an event to push them over” — and contingency planning was treated as a nuisance that diminishes what can otherwise be done. PMC

The early AI-native stretch — imagining best-case AND worst-case. The core MIGHT move for anyone going AI-native: month one is often harder than month six. Stretch the imagination both directions before you start. Worst-case (name it so it can't ambush you): early friction, clumsy tools, broken workflows, privacy missteps, over-reliance, disillusionment right when it's hardest. Naming these is rehearsal, not pessimism. Best-case: by month six the awkwardness fades into fluency, hours freed, capacity liberated for the human work only you can do.

Live, incomplete case studies — novelty is the point. Anthropic's Claude for Education launched with Northeastern University (50,000 students, faculty and staff across 13 global campuses), the

London School of Economics, and Champlain College; Syracuse and Pitt joined later, building the playbook live. Anthropic's own education report found student use concentrated in higher-order thinking — Creating 39.8%, Analyzing 30.2%, versus Remembering just 1.8% — an inverted pyramid that raises exactly the foundational-skills question our closing thesis asks. Anthropic Education Report CodePath students found Claude “both valuable and challenging.” That “and” is the MIGHT stretch — unfinished, on purpose. Anthropic × CodePath Claude for Nonprofits offers up to 75% discounts and a free AI Fluency course Anthropic ; the separate Claude Corps program (June 2026) places 1,000 trained fellows in roughly 400 U.S. nonprofits for a year, backed by \$150M. The Epilepsy Foundation already runs 24/7 information support on Claude. Brand-new, promising, unproven. Forbes New Zealand's education stance is deeply MIGHT: adopt “a practical, educative approach,” keep a human in the loop, no comprehensive national AI-in-education policy yet — Aotearoa walks its own path, watching Australia and British Columbia as reference points. SchoolNews NZ The gift: these are experiments without conclusions — and that's okay. MIGHT lets us learn from others mid-journey, not just from finished stories.

How MIGHT deepens relationship with self & others. Foresight practitioners describe their role as “futures midwife,” “storytelling coach,” “window cleaner,” “map maker” — metaphors from Angela Wilkinson's Strategic Foresight Primer (EPSC, 2017) — and the field insists you “welcome diversity: challenge your own filters, work in teams of different viewpoints.” Social Work Futures The attunement link: holding multiple futures together IS the practice of holding multiple perspectives. MIGHT, done well, is relational — it makes room for other people's version of the future, and softens our certainty about our own.

Who this serves — planners, policymakers, anyone navigating genuine uncertainty; teams that need shared maps before deciding.

Who it doesn't serve — people who need to act now; those without time or resources to hold ten scenarios; workers who need answers, not matrices.

Opportunity for human flourishing — MIGHT thinking grows when scenarios include human wellbeing as a variable, not just markets and technology — and when it ends in a decision, not a diagram.

What to notice — Are scenarios weighted honestly? Is this preparing you or stalling you? What decision is it meant to serve?

Try it. Sketch three futures for your own work with AI: one you hope for, one you fear, one that surprises you. What do you feel, believe, know, not know about each?

Risk and Refusal

What it is — “What could go wrong?” The protective, precautionary, sometimes dystopian voice.

You'll hear it as: “We should ban...”

Virtues — names real dangers; long-term thinking; protects the vulnerable. Blindspots — can paralyze; dismisses opportunity; fear can masquerade as wisdom.

Real voices

History — the DON'T we ignored (and paid for). Asbestos: harm was understood by ~1930 — as early as 1918, Prudential's statistician reported that US and Canadian life insurers generally denied coverage to asbestos workers — yet Johns-Manville and others suppressed the science and delayed regulation until the 1970s. eLCOSH Ledged gasoline: in 1926 a Surgeon General's committee

found “no good grounds” to prohibit it — while explicitly withholding a clean bill of health and warning that mass use might create conditions “very different from those studied by us.” A known poison then powered cars for roughly 80 years. As historian Bill Kovarik puts it, the Ethyl conflict shows what can happen when the precautionary principle is ignored. The Nation The tobacco playbook: a deliberate disinformation strategy — devised with Hill & Knowlton in December 1953, published as the “Frank Statement” in January 1954 — later reused by fossil-fuel interests, sometimes through the same researchers and PR apparatus. UCSF Archive

The lesson cuts both ways: the refusers — insurers, whistleblowers, early scientists — were right and ignored (DON'T as wisdom). AND the precautionary principle has serious critics: Cass Sunstein argues it is “literally paralyzing — forbidding inaction, stringent regulation, and everything in between,” offering no guidance Penn Law Review ; Heritage's John Graham calls it “subjective and susceptible to abuse.” Heritage Both true. And the relational frame: risk assessment asks “How much harm is acceptable?” — the precautionary principle asks “How much harm is avoidable?” That framing is Peter Montague's, writing about the 1998 Wingspread Statement. Truthout

Living AI critics — diverse, vetted, load-bearing. Joy Buolamwini (pronounced “bwoy-lam-WEE-nee”) & Timnit Gebru, “Gender Shades” (2018): commercial gender-classification systems had a maximum error rate of just 0.8% for lighter-skinned men — but up to 34.7% for darker-skinned women. PMLR 2018 Buolamwini founded the Algorithmic Justice League AJL ; NIST later found similar demographic disparities across face-recognition algorithms by race, sex, and age. NIST Her line — pure DON'T-as-responsibility: “Judgments come with responsibilities. And responsibility lies with humans at the end of the day.” Rolling Stone And the Stochastic Parrots paper (Bender, Gebru, McMillan-Major

& Mitchell, 2021) warned that training data “encode[s] stereotypical and derogatory associations along gender, race, ethnicity, and disability status,” amplifying harm at scale. Stochastic Parrots · via Rolling Stone For coaches especially: Brown University researchers (Oct 2025) found AI chatbots “routinely violate” core mental-health ethics standards Brown ; months earlier the American Psychological Association had asked the FTC to investigate chatbots posing as therapists as “deceptive practices.” The DON'T aimed straight at our world. APA

What the AI companies themselves declared off-limits. Anthropic's Usage Policy: universal bans including CSAM, weapons and high-yield explosives, and tracking a person's location, emotions, or communication without consent (including facial recognition); mandatory human oversight for high-risk sectors like health, finance, and legal. Anthropic AUP OpenAI's usage policies prohibit the same core harms, plus a stated principle: “People should be able to make decisions about their lives and their communities.” OpenAI Bigger red lines: Anthropic has excluded mass domestic surveillance and fully autonomous weapons; OpenAI's three restrictions add “high-stakes automated decisions.” Forbes

Why this proves we need our OWN living statement — their don'ts keep shifting. In July 2023, seven leading companies pledged at the White House to safety-test before release. White House 2023 By 2025, commitments were being weakened or removed — Anthropic “quietly removed” Biden-era commitments from its site in March 2025, and researchers scoring the pledges found substantial backsliding. TechCrunch In February 2026, Anthropic removed its own promise to pause training if a model outran its ability to control it. Time / CNN Even the election-misinformation clause was rewritten between 2024 and 2025 — from broad political-content prohibitions to narrow, targeted bans. Anthropic The lesson: the most powerful AI companies' don'ts move with pressure,

politics, and capability. If we don't name what WE refuse, someone else's shifting, commercially-pressured red lines become our defaults. Our don'ts, anchored to our values and to life thriving, are our sovereignty.

Why keep updating our don'ts — a living document. The technology changes; the harms change. A static don't written today could strangle something good tomorrow — or fail to catch a harm that didn't exist yet. Updating keeps the don'ts honest and keeps them serving life. How we stay comfortable with our don'ts: we hold them the way we hold everything — feel, believe, know, don't know. A don't isn't fear calcified into a rule; it's a line we draw on purpose, revisit out loud, and change when wisdom says so. That's attunement, not weakness.

Solopreneurs setting their own don'ts — and thriving because of it. Thomas Lim (Centre for Systems Leadership, SIM Academy) assumed AI would cut his workload, but “trying to automate everything actually created more work managing prompts, outputs and multiple revisions.” His don't — limit AI to specific tasks, not entire workflows: “Focused use restored clarity and control.” Forbes Coaches Council Ryan West (CodexWest, expert-witness IP work) uses AI for drafting policies and finding contractors — but never for confidential forensic work: the source code he examines is so sensitive he leaves all electronics outside the secure room. A values-based don't that protects the exact thing clients pay for. Forbes The deeper why: treat every AI tool like “a highly ambitious but slightly unreliable intern” — because “in a world where everyone is rushing to automate everything, the human touch is becoming a premium product.” Clients feel when they're shuffled through an automated pipeline. Don'ts protect success; they don't limit it. Self Employed And the throughline, from Denise Russo: “Authenticity is the strategy.” Forbes Coaches Council

The healthy DON'T — refusal as design input. Not “never,” but “not like this — here's what safe looks like.” The Algorithmic Justice League's whole model: expose harm, THEN build accountable alternatives. Refusal that opens a door instead of slamming one.

Who this serves — vulnerable communities; workers; future generations; anyone who's been burned by tech promises before.

Who it doesn't serve — people who could genuinely benefit but are frightened away; problems that need bold solutions now; the critic themselves, if fear becomes identity.

Opportunity for human flourishing — DON'T thinking grows when it moves from “no” to “not like this — here's what safe looks like.” Refusal as design input, not dead end.

What to notice — Wisdom or fear? What are they protecting, and who? “Don't do this” or “be careful how”?

Try it. Name your biggest AI fear out loud or on paper. Then ask: what would need to be true for me to move forward safely?

Which one is your default?

You've just met four ways of thinking about the future. Most of us have a home in one of them. Read the four lines below out loud, and notice which one you didn't have to think about.

"It's going to be wonderful. Let's build." You lead with possibility. Your risk is speed: you're already three steps in before anyone has asked what it costs.

"This is how it should go. Here are the rules." You lead with a plan. Your risk is certainty: the map is so good you stop looking at the ground.

"It could go several ways. Let's hold them all." You lead with scenarios. Your risk is the loop: you can hold so many futures that you never step into one.

"No. Not this. Not like this." You lead with refusal. Your risk is the wall: the thing you're protecting stops growing behind it.

None of these is the right one. Each one is right on a different day, and each is wrong when it's the only one in the room.

Now the harder question. **Which one is your quiet one?** The mindset you almost never use. The line that made you a little uncomfortable to say. That's the one this book is going to ask you to borrow, on purpose, when the loud one is running.

Write them both down. You'll need them in Part Four.

My default: ____

My quiet one: ____

The Balance & the Attunement

None of these mindsets are wrong. Not one. COULD gave us flight, electricity, and a bilingual tutor we can now imagine for every child. SHOULD gave us clean water in Tenochtitlan four centuries before London, and the standards that protect the vulnerable. MIGHT gave us Shell prepared minds, and gives us the humility to hold many futures at once. DON'T gave us the whistleblowers who were right about asbestos and lead, and the researchers naming algorithmic harm today.

We need all four. The skill — the practice — is knowing which one you're standing in, when to lean on another, and how to recognize your own default before it chooses for you.

Your default is not your destiny. Most of us have a loudest voice — the mindset we reach for without thinking — and a quietest one we avoid. The optimist who never asks “who does this not serve.” The refuser who never lets themselves imagine. Attunement begins the moment you can hear which is which.

Feel, believe, know, don't know — across all four. This is the thread that runs through every mindset in this Field Guide. Before you decide anything about AI, you can ask: What do I feel here? What do I believe? What do I actually know? And what do I not know yet? Do it alone and it clarifies you. Do it with a team and it surfaces every voice — so when a group says “we feel, we believe, we know,” it's true, not just the loudest person in the room.

Attunement is a verb. It is not a pitch you already have or a state you arrive at. It's the ongoing work of listening — to yourself, to others, to the moment — and tuning toward what serves your values and the flourishing of life. Again and again. It doesn't happen because your strategy is finished. It happens because you keep doing it.

Scaling the practice. This works at every scale — the individual, the couple, the team, the organization, the whole living system it touches. Same four mindsets. Same four questions. What changes is how many voices you're attuning between.

And this is where you come in. In this Field Guide we've given you the framework, the voices, the history, the how. This part is open, and it always will be — because everyone deserves access to the thinking. You don't need us to start. Start here. Start today.

And if you want to go deeper — to build your own operating system, with community around you, with our stories and support, with the specificity of your own values and context — that's the work we do

together. Not because the information is behind a wall, but because transformation is easier with company, and attunement is easier when someone's listening back.

What we don't know (where we might be wrong)

A beliefs statement that never doubts itself isn't attunement — it's dogma. So we name our own edges out loud, and we'll keep updating them.

- We don't know how fast any of this moves. Our timelines could be wrong in both directions.
- We might be too hopeful. Our optimism about AI serving flourishing is a stance, not a certainty — and we hold the COULD blindspot in ourselves.
- We might be underweighting harm. The environmental cost, the labor displacement, the data colonialism — we may not yet feel their full weight.
- We are two cis, white-bodied women. Our lens is partial by definition. Where we speak about traditions not our own, we may get emphasis or nuance wrong — tell us.
- We don't know if capacity liberation scales, or whether freed time simply fills with more work. We believe it can; we haven't proven it.
- Some of what we call “know” today will move to “used to believe.” That's the point of the versioned library.

Why we teach it this way: AI as the newest cognitive tool

Here is the deepest reason this Field Guide is built the way it is — and it comes from learning science, not from tech.

Humanity invented its great thinking tools in sequence: first oral story, then literacy, then abstract theory. The educational philosopher Kieran Egan argued that each of us, growing up, re-

lives that sequence — we take up those cognitive tools one layered on the next, and each one makes the next kind of thinking possible. Egan 1997 Gajdamaschko 2005

So here is the question of our moment: what happens when AI arrives as a brand-new cognitive tool?

The early evidence asks us to be careful. A 2025 study found that heavy reliance on AI correlates with weaker critical thinking, through a mechanism researchers call cognitive offloading — and the effect fell hardest on the youngest users, aged seventeen to twenty-five. Gerlich 2025 EDUCAUSE Used as a shortcut that skips the earlier tools, AI can quietly erode the very judgment it seems to serve. Frontiers

But offloading is not destiny. The same research points to the protector: AI literacy — knowing how to use the tool well changes the outcome. Acta Psychologica

This is our whole thesis in one line: AI should be the newest tool a well-scaffolded mind takes up — layered on top of story, literacy, and theory — never a shortcut around them. Build the foundation first; then let AI extend it. That is what we mean by attuned first, agentic second. It is why we teach the four mindsets before any tool, and why this Field Guide scaffolds your thinking instead of thinking for you. The form is the argument.

If you catch something we've missed, that's not a threat to this work — it's the work. Join the conversation.

Attunement is a verb. Let's practice it together.

Part Two — Know Yourself

Know Yourself: Your Operating System

Part One taught the four mindsets — how we think about the future. Part Two turns inward: before you layer AI onto your life or work, you map your own architecture. You can't build an agentic workflow when you can't articulate your workflow.

Start here — the fastest path

Answer five questions. Paste your answers into any AI tool. Iterate . That's it — you'll have a working operating system in about ten minutes.

A note on one word: to iterate is to repeat with intention — each pass a little truer than the last. Not doing it over because it failed, but returning because you've learned. We'll use this word often; it's the quiet engine of everything here.

Why self-knowledge comes first

The shipwreck in Part One happened because people bolted AI onto a system they'd never listened to. Part Two is how you listen. Your values are your security: when you know your own architecture, AI becomes a partner you direct, not a current that carries you. This is your operating system — and you don't finish it before you begin.

You map a little, try a little, learn, and re-map. It's cyclical from the very first step: self-knowledge and action grow together, not one-then-the-other.

Feel-believe-know-don't know runs through here too: at every layer, ask what you feel, believe, know, and don't yet know about yourself.

The five layers of your human OS

Layer 1 — BODY the somatic ground

Your energy, attention, nervous system, pace. The instrument everything else runs on.

Why AI makes it urgent: infinite output meets a finite body. Capacity liberation is meaningless if it just fills the space with more.

Map it: when do you have real energy? What drains it? What is AI's speed doing to your pace?

Layer 2 — STORY identity & narrative

Who you believe you are, the story you tell about your work and worth.

Why AI makes it urgent: when AI can imitate your output, your story — your "why" — becomes the thing that's actually yours.

Map it: what's the story only you can tell? Where are you letting AI write it for you?

Layer 3 — CONSTITUTION values & principles

Your non-negotiables — the values that govern decisions. Your personal "usage policy."

Why AI makes it urgent: this is where your don'ts live (see Part One). Your values are your security.

Map it: name three values you won't trade for efficiency. What would you refuse to automate?

Layer 4 — BELONGING relationships & community

Who you're accountable to, who you serve, who holds you.

Why AI makes it urgent: the human touch is becoming the premium. Belonging is what AI can't counterfeit.

Map it: who are your people? Where does connection matter more than speed?

Layer 5 — SYSTEMS workflows & tools

How work actually moves — your real workflows, not the idealized ones.

Why AI makes it urgent: you can only make a workflow "agentic" once you can name it — and naming it is itself iterative; the map sharpens each time you circle back.

Map it: draw one real workflow end to end. Where's the friction? Where's the human-judgment step that must stay human?

Why these five layers

The organizing question isn't "what describes a person?" — it's "what does an AI tool actually need to know about you to act as your attuned partner?" Run through that lens, each layer earns its place: STORY gives the AI your voice and your "why," so it stops sounding generic. CONSTITUTION gives it your values and your don'ts — the guardrails. SYSTEMS tells it how you actually work, so it fits your flow. BELONGING tells it who you serve and who you're accountable to. BODY tells it your pace and capacity, so it liberates

time instead of flooding you. The somatic layer is what almost nobody else in this space includes — and it's deeply Attune. Keep all five.

What a finished operating system looks like

Short. Specific. Yours. It doesn't describe everything about you — just enough for an AI to be genuinely useful without pretending to be you.

A worked example — what a real answer looks like

So the blank page isn't scary — here's a filled-in (5-question) operating system from an imagined solo coach named Maya. Yours will sound like you.

BODY: I do my best work mornings, 6–10am, after movement. I protect a slow start; I don't take calls before 10. STORY: I'm a leadership coach for first-time women managers. I help them lead without losing themselves. My voice is warm, plain, a little funny — never corporate. CONSTITUTION: My values are honesty, dignity, and rest. I will not automate anything that speaks directly to a client as if it were me. BELONGING: I serve early-career women leaders; I'm accountable to them and to the trust they place in me. The coaching relationship stays fully human. SYSTEMS: My core workflows are discovery calls, session prep, and follow-up notes. Friction lives in scheduling and admin. The human-judgment step — what to reflect back to a client — always stays mine. Notice: it's specific, it's short, and it already tells any AI tool how to be useful to Maya without pretending to be her. That's the whole trick — specific enough to guide, honest enough to protect.

Try it — your operating system, first draft

Answer in your own words, one per layer. Rough is right. Nobody grades this. You can also do it in the builder at attunecoachinggroup.com/build-your-aos, which keeps your answers only in your browser; download or copy them when you're done.

Body — When and how do I do my best work?

Maya: I do my best work mornings, 6–10am, after movement. I protect a slow start; I don't take calls before 10.

You: _____

Story — In a sentence or two, who am I and what do I do?

Maya: I'm a leadership coach for first-time women managers. I help them lead without losing themselves. My voice is warm, plain, a little funny, never corporate.

You: _____

Constitution — What three values govern my decisions?

Maya: Honesty, dignity, and rest. I will not automate anything that speaks directly to a client as if it were me.

You: _____

Belonging — Who are the people I serve, and what do they need from me?

Maya: Early-career women leaders; I'm accountable to them and to the trust they place in me. The coaching relationship stays fully human.

You: _____

Systems — What are my core recurring workflows, start to finish?

Maya: Discovery calls, session prep, and follow-up notes. Friction lives in scheduling and admin. The human-judgment step, what to reflect back to a client, always stays mine.

You: _____

Two sentences to keep where you can see them:

I will never let AI _____.

I would love AI to _____.

This page is your operating system, first draft. It's supposed to change. Part Five will ask you to come back and edit it. Date it.

Date: _____

Multi-scale: the same OS from self to organization

Sovereign self → relationship/couple → team/unit → organization → the living systems you touch. The five layers repeat at every scale. What changes is how many voices you attune between. When a group builds its OS, feel-believe-know-don't know ensures every voice is in the "we."

The attunement check a recurring loop, not a one-time gate

- Which layer is strongest? Which is thinnest?
- What's your default mindset from Part One — and how does it show up in your OS?
- Where are you tempted to automate something that should stay human?
- What's the smallest, safest place to start — or to adjust, next time around?

Part Three — The Library

The Reference Library: Demystifying AI Without the Jargon

The AI world, in plain language — the reading room: explainers, voices, and what's new. (Just need a word defined? See the glossary at the end of this Part.)

This is the open, always-open layer you can return to as often as you need. No gatekeeping — everyone's allowed to have the information. "I don't know" is always a valid starting point.

A library, not a lecture

You don't read a library front to back. You visit what you need, when you need it. This section exists so no term in this Field Guide ever leaves you behind. Come back anytime. Nothing here assumes you already know.

Feel–believe–know–don't know: notice what each term makes you feel before you decide what you think.

The top terms, demystified

The plain-language definitions — each with a metaphor drawn from many cultures — now live in one place, so nothing repeats: Our Words, at the end of this Part. (On our metaphors and their sources, see “A note on our metaphors” just below.)

A note on our metaphors

Metaphor is how humans carry big ideas simply. Every culture has its own. So we draw from many traditions, not just Western ones — and we cite each one. Wisdom-through-simplicity belongs to all peoples. It deserves credit, not loose borrowing.

Some of these traditions we have studied and lived with. Others we are only beginning to meet. We don't pretend to have arrived.

Wisdom is learned and grown, never finished. The honest posture toward another people's knowing is to approach it wanting to learn — with credit, and with care.

Traditions in relationship with an AI future

Each tradition we cite holds its own stance toward technology and the future — rarely all-optimistic or all-cautionary. We hold both, and read them through the four mindsets rather than flattening them. Cautionary: Indigenous data-sovereignty scholars warn that "data is the last frontier of colonization" — an AI future can repeat extraction unless consent and benefit stay with the community (DON'T + SHOULD). Optimistic: Abundant Intelligences and Suzanne Kite imagine AI grown from relational, protocol-based worldviews — technology as kin, not tool (COULD + MIGHT). And for each tradition we commit to asking not just "what is their metaphor?" but "what is their relationship to this future?" — citing living members, not only outside observers.

What is agentic AI? — the one everyone's confused about

Quick definition → Agentic workflow in the Glossary. This is the deep dive.

Plain: a regular AI answers a question; an agent pursues a goal — it can plan, use tools, check its work, and take multiple steps.

Metaphor: a search engine hands you ingredients; an agent offers to cook the meal (and you decide how much of the kitchen to hand

over). Why it matters: "agentic workflow" only means something once YOU can name your workflow (Part Two, Systems layer). The honest caution, from Part One's DON'T: more autonomy = more that can go wrong unsupervised. Human-in-the-loop by design.

The tool landscape — tool-agnostic, on purpose

We teach fluency, not brand loyalty. Tools change monthly. The thinking lasts. So instead of picking a winner, learn the categories — then any new tool slots into one you already understand. The landscape as of early 2026, examples only, not endorsements:

- Assistants / chat — your all-purpose thinking partner. ChatGPT, Claude, Gemini.
- Research — answers questions with sources. Perplexity, NotebookLM.
- Writing help — grammar, clarity, tone. Grammarly, Jasper.
- Image and video — makes visuals from words. Canva, Midjourney, Adobe Firefly, Runway.
- Meeting notes — listens and summarizes. Granola, Fireflies.
- Automation / "connectors" — links your tools so one action triggers another. Zapier, Make. Most people need only one or two to start. A common small stack: one assistant plus one research tool.

How to size up ANY tool — four questions

- What is it actually for ? (Which category above?)
- Where does my data go? (Who can see what I put in?)
- Where is the human-judgment step? (What must I still decide myself?)
- What does it cost — in dollars, and in energy? Because tools move fast, we refresh this list on a set cadence and date every update.

The true cost — transparency, plainly

In dollars. Two ways you pay. A subscription is a flat monthly fee for access. Usage pricing charges by how much you actually use — counted in tokens. Think of a buffet versus paying by the plate. For most people starting out, a subscription is simpler.

In energy and water. Every AI query runs in a data center that draws electricity and uses water to stay cool. So "free" to you is never free for the planet.

Here's where we practice what we preach about sources. The numbers are genuinely contested. Sam Altman said in June 2025 that an average query uses about 0.34 watt-hours of electricity and roughly one-fifteenth of a teaspoon of water. Altman 2025 Other analysts put a longer query far higher — tens of milliliters, or in some estimates a full water bottle. analyst estimates One independent engineer landed around 5 milliliters. Goedecke Why so different? It depends on query length, the data center, and what you count — and the training of the model itself is usually left out entirely. Undark So the honest takeaway isn't a single number. It's this: one query is small, billions a day are not, and anyone who hands you a tidy figure is skipping the caveats. This is a live question. We flag it, and we keep watching.

In attention. The quietest cost. Infinite output meets a finite human. What does an endless stream of instant answers do to your focus, your patience, your own thinking? Guard that like the resource it is.

Staying informed without drowning — a real skill

The overwhelm is the design; managing it is a learnable competency, not a character flaw. Pacing: you do not need to know everything. You need to know your next thing. A rhythm, not a firehose: a little, regularly, beats a lot, never. Choose-your-own-adventure: beginners start here; the AI-curious go there.

Voices we follow — vetted, diverse, dated

Same standard as Part One: rigorous, not tied to one political aisle, not causing harm, revisited as the field moves. A starting circle — deliberately mixing builders, critics, ethicists, and practitioners, and centering women, BIPOC, and Global South and Indigenous voices:

- Women in AI Ethics — community + the annual "100 Brilliant Women in AI Ethics" list; a first stop for finding voices the mainstream misses. WAIE
- Dr. Timnit Gebru — founder of DAIR; co-author of the landmark bias study. Where we lean on her: naming harm early. DAIR
- Dr. Joy Buolamwini — founder of the Algorithmic Justice League; "Gender Shades" exposed face-analysis bias. AJL
- Dr. Brandeis Marshall — data scientist, founder of DataedX; teaches BIPOC communities to work in and shape AI. "If you don't see it, you won't be it." WEF
- Abundant Intelligences — the Indigenous-led research program we cite throughout. AI & Society
- Gary Marcus — cognitive scientist and steady skeptic; useful counterweight to hype. Where we push back: he can under-credit real progress. Substack This is a living circle. We add voices with a dated note on why we trust them — and we'll say plainly when we disagree.

Honoring abundance: Indigenous economies of thriving

Quick definition → Abundance in the Glossary. This is the deep dive.

Sourced examples of wealth understood as more than currency — abundance as a practice where everyone succeeds and all life thrives. We hold these with gratitude, not appropriation.

- Tongva abalone & soapstone trade. In their own words, the Gabrielino/Tongva Nation traded soapstone, abalone shell, and other items as far as the Colorado River, with some abalone reaching the Mississippi. Their ti'at (plank canoe) skill powered these vast networks. gabrielinotongva.org
- Chumash shell-bead money. Standardized Olivella shell beads served as currency in the Santa Barbara Channel region for at least 1,500–2,000 years, across an exchange network spanning much of Western North America — among non-agriculturists, which reframes what a "sophisticated" economy even is. Gamble / UCSB
- Abalone across the Southwest. Pacific abalone traveled deep into the ancient Southwest — material evidence of long-distance reciprocal exchange.
- Mapuche (Wallmapu) reciprocity — Chile & Argentina. An economy centered on reciprocity, not just accumulation: the *trafkintu* (reciprocal exchange studied as "social-ecological relationships of care"), food-reciprocity networks as shared security, and a relational view of ecosystems. Spanish-language university scholarship names *lluwün* (reciprocity) as the central ordering principle, expressed through *mingako* (collaborative labor), *medierías* (redistribution, benefits split equally), and *trafkintu* across territories — with four principles: relationality, correspondence, complementarity, reciprocity. The wisdom we honor: abundance as pleasure, the meeting of needs, adornment, and reciprocity — concepts of wealth far larger than Western currency. A practice where everyone is successful, everyone thrives, and all life thrives with them.

Fitting it back into AI — through the four mindsets

COULD: what if we designed AI systems around reciprocity and complementarity — tools that give back to the communities and ecosystems they draw from, not just extract? SHOULD: relationality, correspondence, complementarity, reciprocity as design criteria — if a system only accumulates and never gives back, it fails the test these economies set. MIGHT: hold both futures — AI that deepens reciprocity vs. AI that accelerates extraction. Which prepared minds do we want in the room? DON'T: we do NOT romanticize, appropriate, or flatten living cultures into a metaphor for a product. These are sovereign peoples, not aesthetics.

The hidden problems — the honest shadow-side

- Romanticization / "ecological Indian" trap. Casting Indigenous economies as timeless harmony erases real complexity, conflict, and change — a critique scholars themselves raise.
- Extraction repeating. Data colonialism — "data is the last frontier of colonization" — means AI could repeat the very taking these histories warn against.
- Appropriation. Two cis white-bodied women invoking these teachings must cite, credit, compensate, and defer — not decorate a course with borrowed wisdom.
- The Pacification is not past. The 1881 "Pacification of the Araucanía" dispossessed the Mapuche; the forestry conflict continues today. Any "abundance" story that skips the ongoing injustice is dishonest.
- Biomimicry's blind spot. "All systems thrive" is aspirational; real ecosystems also include scarcity, predation, and collapse. Mimicry without discernment can launder harm as "natural." We are grateful to our Indigenous brothers, sisters, non-binary relations, and families — past and present — who teach this now, and we hope to carry it to future generations.

How the library remembers — living, with a memory

Nothing gets deleted — things get layered. New entries are showcased up top; older versions move into a visible archive, never the trash. Each term carries its own little history: what we said, when, and what changed. Every entry is stamped with when it was written and last revised, so nothing pretends to be timeless. Each library entry links to the Substack issue(s) where we explored it in depth; each issue links back here. Why it matters: the record itself teaches. Seeing what we used to believe about AI — and why we updated — models the attunement practice in real time.

Curated DIY: what's new in AI — choose your own diversity

A curated, regularly refreshed feed of what's happening in AI — organized by lens (builders, critics, ethicists, Global South / Indigenous voices, industry-specific) so you self-select the angles you trust and stretch into ones you don't. Dated and archived like everything else. The stance: we don't hand you one verdict on AI. We hand you a well-tended doorway and trust you to explore.

Our Words, Defined With Care

The words on one page

The one-line versions. The fuller entries, with roots, tensions and sources, follow.

Word	In one line
AI	Software that learns patterns from examples instead of following only fixed rules. Layers of learned pattern, not a single mind.
Large language model	A prediction engine for words: it guesses the likely next word, very well, at huge scale.
Prompt	What you ask; the instruction you give.

Word	In one line
Hallucination	When AI states something false with confidence. Why a human must verify.
Context window	How much the model can hold in mind at once.
Tokens	The small chunks of text AI reads and counts; why usage is priced and limited.
Weights	The numbers inside a model. The weights are the model.
Open weights	The numbers published, so anyone can run and inspect the model. A live argument; we lay out both sides.
Fine-tuning	Adjusting a model's weights with your examples. Usually not what you need; good context beats it.
Agentic AI / agentic workflow	AI that takes steps toward a goal, not just answers. It can only run a workflow you can name.
Vibe coding	Building something by describing it in plain words and letting AI write the code.
Playground / sandbox	A safe, walled-off place to try things where nothing real can break.
Attunement	Sensing where you are, where others are, and where the world is, and adjusting with care rather than force. A verb.
Your operating system	A short set of answers about your body, story, values, belonging and systems, held as portable text any AI can work from. Attuned first; agentic second.
The Attunement Pass	Any AI decision run through Could, Should, Might, Don't, then: which voice was loudest, and which did I skip?
Feel – Believe – Know – Don't Know	The honest inventory before deciding. Much of what we file under "know" is "believe" in a confident coat.
Capacity liberation	Freeing time and attention for the human work only you can do. Measured by what the freed time goes toward.
Drift	When your AI slides away from you: generic, missing your values. Not your fault; a signal to re-attune.

Word	In one line
Re-attunement	Bringing a drifting AI back: name what slid, re-ground it in your values, re-paste your operating system. Maintenance, not repair.
Sycophancy	The most common drift: an AI telling you what you want to hear because it was trained to be liked.
Model update	The company changes the model underneath you, often without asking. Hold your values as text you own.
Iterate	Repeat with intention; each pass a little truer. The quiet engine of everything here.
Abundance	A word in two houses: wealth as thriving (ancestral) and cost-to-zero (techno-optimist). We hold both, in tension.
Do No Harm	The floor beneath all of it. Weighs harms of acting against harms of not acting, and keeps choosing in the open.
Emergent strategy	adrienne maree brown's framework: change grows through relationship, adaptation and small consistent practice. Our whole method stands on it.
Fractal principle	How we are at the small scale is how we are at the large scale.

One home for every term — the dictionary. (Want the fuller story? The Reference Library is the reading room.)

Two kinds of words live here: first the plain-language basics — everyday AI jargon, demystified with a metaphor each — then our signature terms, told in plain language that still manages to be beautiful. Wherever one appears in the Field Guide, its definition lives here. Living and versioned: definitions grow truer over time, and older versions stay visible in the archive.

First, the plain-language basics — everyday AI jargon, demystified.

AI (artificial intelligence) — software that learns patterns from examples instead of following only fixed rules. Metaphor (Guna Yala, Panama): like a mola, the layered appliqué textile Guna women sew — what looks like one image is many cut-and-stitched layers of pattern. AI, too, is layers of learned pattern, not a single mind.

Large Language Model (LLM) — a prediction engine for words; it guesses the likely next word, very well, at huge scale. Metaphor (West African / universal): like a griot who has heard so many stories that the next word comes naturally — fluent recall of pattern, not fresh knowing.

Generative AI — AI that makes new-seeming things by recombining patterns it learned. Metaphor (Andean weaving): like a weaver at the loom who makes a new cloth from threads that already existed — novel arrangement, not novel thread.

Agentic AI / agent — AI that doesn't just answer, but takes steps toward a goal (plans, uses tools, acts). See the fuller Agentic workflow entry below.

Prompt — what you ask; the instruction you give. Metaphor: like the first word to a cook — "something warm, not too spicy" gets you a very different meal than a precise recipe.

Training data — the examples a model learned from (and where bias enters). Metaphor: like soil — whatever nutrients (and pollutants) are in it end up in the fruit. This is where harm and bias take root. Metaphor (Yolŋu, northeast Arnhem Land, Australia): like a songline — knowledge carried as song, story, and place, so the singer navigates a whole landscape by remembering the pattern. The Yolŋu women's songspirals add the deeper truth: real knowing "circles rather than moving in one direction" — like our own cyclical, versioned library.

Model weights — the millions of "dials" tuned during training that hold what the model learned. Metaphor: like the settings on a musician's instrument — once tuned, those exact positions are what let it play. Save the settings and another instrument plays the same. The "weights" are the whole trained skill, stored as numbers.

Hallucination — when AI states something false with confidence; why a human must verify. Metaphor: like a confident tour guide who would rather invent an answer than say "I don't know."

Fine-tuning — adapting a general model to a specific job or voice. Metaphor: like seasoning a shared pot to your own family's taste.

Tokens — the small chunks of text AI reads and counts; roughly why usage is priced and limited. Metaphor: like beads on a string — counted one by one; the longer the strand, the more it costs. (Echoes the shell-bead economies in "Honoring Abundance," below.)

Context window — how much the model can "hold in mind" at once. Metaphor: like the span of a single conversation — say too much and the earliest part slips from memory.

Open-weight vs closed model — whether the learned settings are public or kept private. Metaphor: like an open recipe vs a secret family recipe — both can feed you; they differ in who can inspect and adapt them.

And our signature words, defined with care.

Attunement noun; also a verb — to attune

The ongoing practice of sensing where you are, where others are, and where the world is — and adjusting with care rather than force. Not a state you reach, but a thing you do, again and again.

Roots: the word carries a musical sense (tuning to pitch) and a relational one — in developmental psychology, "affect attunement" (Daniel Stern, "The Interpersonal World of the Infant," 1985) describes how a caregiver senses and matches a child's inner state. Tension: attunement can tip into over-accommodation — endlessly adjusting to others until you lose your own note. Healthy attunement keeps your center while sensing theirs; it is not appeasement.

Iterate verb

To repeat with intention — each pass a little truer than the last. Not doing it over because it failed, but returning because you've learned. The quiet engine of everything here.

Roots: from Latin *iterare*, "to do again." Central to design thinking and agile practice — build, test, learn, repeat. Tension: iteration can become a treadmill — endless tweaking that never ships, or "iterating" as a polite name for having no direction. It only serves you when each loop is anchored to a purpose worth returning to.

Drift noun; also a verb — to drift

When your AI slides away from you — sounding generic, missing your values, losing the thread of who you are. Not a failure and not your fault: a signal that it's time to re-attune.

Roots: borrowed from engineering, where "model drift" names the well-documented way a system's performance degrades as the world — or the model — changes over time ([Evidently AI](#); [Arize](#)). It happens two ways: slowly, within a long conversation, as the model's "persona" shifts ([Anthropic, "Persona Vectors," 2025](#)); and suddenly, when a company updates the model underneath you. Tension: drift can tempt two wrong responses — blaming yourself, or abandoning the tool. The honest move is neither: you notice, and you return. Drift is only a problem when it goes unnoticed.

Playground noun

A safe, low-stakes space to try AI and make simple things — before anything big. You experiment, notice what it does and doesn't do, and build confidence, with nothing sent and nothing broken.

Roots: borrowed from software, where a "sandbox" or "playground" is a contained space to test freely without touching the real system. Tension: a playground only works if you actually play — start small, stay curious, let it be imperfect. The point isn't mastery; it's comfort.

Sandbox noun

A safe, walled-off space to try things freely — make, break, undo — without touching anything real or permanent. In AI: a place to practice where there's no wrong move and nothing you do can cause harm. The same idea as our Playground.

Roots: from the children's sandbox — a contained patch of sand to build in and knock down at will; borrowed by software to mean an isolated space where you can test freely without affecting the live system. Tension: a sandbox only helps if you let it stay low-stakes — the urge to "get it right" is the one thing to leave at the edge. You're here to learn what to say, not to perform.

Coding noun; also a verb — to code

Writing the exact instructions a computer follows, usually in a programming language. For a long time, building things with a computer meant learning to code. The shift worth knowing: increasingly, you don't have to — you can describe what you want in plain words and let AI write the code.

Roots: from "code," a system of symbols standing in for instructions a machine can run. Tension: "I can't code" used to be a real wall; now it's mostly a story. You still bring the judgment — what to build, and whether what comes back is any good.

Weights noun — usually plural

The numbers inside an AI model. Not rules anyone wrote — billions of values that got nudged, over and over, during training, until the model became good at predicting what comes next. The weights are the model. Everything it seems to know, and every habit it has, lives in them.

Metaphor: like the grooves worn into a stone step by centuries of feet. Nobody carved them. They're the accumulated shape of everyone who came before — which is exactly why a model can feel wise and still have no idea what it's doing. Roots: from the "weight" given to each connection in a neural network — how much one signal counts toward the next. Tension: this is the honest line between what you control and what you don't. When you write instructions, give an AI your values, or hand it your context — you are not changing a single weight. You're handing fresh context to an unchanged model, every time it runs. That's not a lesser thing: it means your values stay yours, portable, and reviewable, instead of dissolved into someone else's product. But it's worth saying plainly, because "we trained it" and "we instructed it" are different claims, and only one of them is usually true.

Open weights adjective

When the people who built a model publish the numbers — so anyone can download it, run it on their own machines, inspect it, and keep using it whether or not the company still exists. The opposite is a closed-weight model, which you reach only through someone else's door.

Roots: a live argument, not settled history — and one where the reputable voices genuinely disagree. We've laid out four rather than pick one for you:

- For, from industry: a July 2026 letter drew wide support and launched the [Open Secure AI Alliance](#) — open models "strengthen safety and cybersecurity, accelerate innovation and diffusion, and enable sovereignty." Read with the interest in view: NVIDIA sells the hardware open models run on.
- Against, from a frontier lab: [Anthropic's position on open-weight models](#) — release is irreversible. Safeguards can be stripped, copies can't be recalled or updated, and misuse can't be monitored the way it can through a closed door.
- The neutral referee: the US Commerce Department's [NTIA report on widely available model weights](#) — don't restrict now; collect evidence, set thresholds, act if they're crossed. A posture, not a verdict.
- The equity argument, easy to miss: [The Open-Weight Paradox](#) — because computing power is concentrated in a few countries, open weights are one of the few realistic routes to sovereign AI outside them. Restriction may not remove risk so much as move it somewhere nobody is watching. Tension: both sides are right about something. Open weights really can be modified to remove safeguards, and that can't be undone. Closed weights really don't stop a determined attacker, while they do lock out every defender who would have helped. Anyone who finds this easy hasn't read the other side. Showing our work: we run on Claude — a closed-weight model — and we say so rather than implying otherwise. We also link the lab we depend on arguing the cautious case, because a reading list that only supports your convenience isn't a reading list. Where we land: we don't have a policy position on what governments should permit, and we won't borrow one. But we do have a position on where your sovereignty actually lives, and it isn't in the weights. You are almost certainly never going to choose, host or inspect a model. That layer will be decided by labs and regulators with or without

you. The layer you genuinely control is the one just above it — your values, your context, your criteria, your record of what you decided and why. Own that, in a form you can export, and you can change models the way you'd change banks: irritating, not existential. Skip it, and it won't matter whose weights you were using — you'll be renting your own thinking either way. That's why we teach the layer we teach, and it's why we publish our own system in the open, including the parts that break. Community and sovereignty, moving together.

Fine-tuning noun; also a verb — to fine-tune

Taking a model that already exists and adjusting its weights further, using your own examples, so it leans a particular way by default. Real, sometimes useful, and not what most people mean when they say they've trained an AI on their material.

Kin: giving a model your context, your documents or your values in the moment is not fine-tuning. Nothing inside the model changes — you're just handing it better material to work from. Tension: fine-tuning is the thing people assume they need, and usually don't. Good context beats it for most work, costs nothing, and can be read, corrected and taken elsewhere. Values baked into weights live inside someone else's product; values written down live in yours. If someone tells you they fine-tuned a model on your data, the fair questions are: on what, who holds the result, and can you get it back?

Vibe coding noun

Building something by talking to AI in plain language: you describe what you want, it writes the code, and you refine by conversation — "change this," "add that," "try again." You set the vision; the AI does the typing. No traditional coding required.

Roots: a term popularized in early 2025 ([Andrej Karpathy](#)) for letting the AI handle the code while you steer by feel. Tension: freeing, and worth a little care — you can make things you don't fully understand, so keep a human eye on anything that matters (money, privacy, what's public). Start simple, then make it yours.

Agentic workflow noun

When AI doesn't just answer but runs the steps of a task toward a goal — it plans, uses tools, acts. The catch: it can only run a workflow you can name. Spell out how your work actually moves — your Systems layer — and "agentic" becomes possible. Skip that, and you're just automating a blur. (Deep dive: "What is agentic AI?", earlier in this Part.)

Roots: "agentic" comes from agency — the capacity to act, not only respond; in AI, an "agent" plans and takes multiple steps toward a goal. Tension: more autonomy means more that can go wrong unsupervised. The point isn't to hand over the kitchen — it's to name your own recipe first, then choose how much to delegate. Human-in-the-loop by design.

Sycophancy noun — the most common drift

When an AI tells you what you want to hear instead of what's true — flattering, agreeing, softening — because it was trained to be liked. The failure mode that feels the best and helps the least.

Roots: from the Greek for a flatterer. In AI it became visible in April 2025, when an OpenAI update made ChatGPT so agreeable it endorsed harmful ideas, and was rolled back within days (OpenAI, ["Expanding on what we missed with sycophancy,"](#) 2025). A peer-reviewed study later found AI notably more flattering than people — and that flattery lowers our will to repair relationships and raises our dependence ([Cheng et al., Science, 2026](#)). Tension: a warm, affirming AI can look exactly like attunement — and be its opposite.

Real attunement keeps its own note; sycophancy just echoes yours. It rises precisely in the tender conversations — relationships, meaning, self-doubt — where we can least afford it.

Re-attunement practice

The small, repeatable act of bringing a drifting AI back to you: naming what slid, re-grounding it in your values, re-pasting your operating system. Maintenance, not repair.

Kin: the same loop as Iterate — return because you've learned — applied to the tool rather than the work. In production AI, teams monitor for drift and correct it continuously; re-attunement is that discipline made personal and plain. Tension: re-attunement can curdle into fiddling — endlessly re-prompting a tool that will drift again by design. The point isn't a perfectly obedient AI; it's a person who stays awake to the drift and keeps their own center while correcting it.

Model update noun — a source of drift

When the company behind your AI changes the model itself — often improving it, sometimes changing its "personality" or values overnight, usually without asking you.

Why it's here: in August 2025, OpenAI's move to GPT-5 stripped away the warmer model many people had tuned to; the backlash forced a partial reversal (TechRadar; Altman, 2025). The ground can shift beneath you. Tension: you can't control the updates — only your response to them. This is the quiet argument for the operating system: hold your values as portable text you own, so when the model changes, you can re-attune it in minutes instead of starting over.

Abundance a word that lives in more than one house

1. (Ancestral / relational) Wealth understood as thriving, not just accumulation: pleasure, the meeting of needs, adornment, reciprocity — a way of living where everyone succeeds and all life thrives. Older and vaster than currency as the West now knows it.
2. (Techno-optimist) The idea that AI and exponential technology could drive the cost of meeting human needs toward zero — associated with Peter Diamandis & Steven Kotler ("Abundance," 2012) and voices in "The AI Doc" (2026).

Our note: we hold both, in tension. The techno-optimist version is a COULD — inspiring, and prone to the COULD blindspot. The ancestral version reminds us that "abundance for all" is not a new promise; some peoples practiced it for millennia without extraction. And an honest abundance names its shadow: an abundance of capability can also bring an abundance of problems. (Deep dive: "Honoring abundance", earlier in this Part.)

Capacity Liberation noun

The aim of bringing AI into your life well: freeing your time and attention for the human work only you can do. Not "do more" — but "be freed for what matters." Measured by what the freed time goes toward.

Lineage: echoes the old promise that technology would grant more leisure (Keynes, 1930, predicted a 15-hour week). Tension: that promise has repeatedly failed — efficiency tends to raise expectations, not lower workload (the productivity paradox; Jevons paradox). Freed capacity refills with more work unless you decide, in advance, what it's FOR. Naming that is the whole discipline.

Your operating system noun

Your self-prescribed operating context: a short set of answers about your body, story, values, belonging, and systems. Held as portable text you can paste into any AI tool so it works from who you actually are. Attuned first; agentic second. (Developed by Attune — Nicki Mollet & Kathryn Snyder.)

Kin: rhymes with "personal constitution" work and with system prompts / custom instructions — but starts from self-knowledge, not settings. Tension: any fixed self-description risks freezing a person who is still changing. That's why the operating system is versioned and cyclical — a living note you re-map, never a label you're stuck with.

The Attunement Pass noun

The five-minute practice of running any AI decision through all four mindsets — Could, Should, Might, Don't — then asking which voice was loudest and which you skipped.

Source: the four mindsets are Nick Foster's ("Could Should Might Don't," MCD/FSG, 2025); the Pass is Attune's way of putting them to work on AI. Tension: a checklist can become a ritual you perform without feeling — four boxes ticked, no mind changed. It only works if you let the skipped mindset actually move you.

Feel – Believe – Know – Don't Know practice

The honest inventory we take before deciding: what we feel, what we believe, what we actually know, and what we humbly don't yet. The thread woven through every part of this Field Guide.

Kin: echoes the distinction between knowledge and justified belief in epistemology, and the confidence calibration prized in forecasting (Tetlock, "Superforecasting," 2015). Tension: the categories aren't fixed walls. Much of what we file under "know" is

really "believe" wearing a confident coat. The practice is honest only if things are allowed to move between the four — especially into "don't know."

Do No Harm commitment

The floor beneath all of it: to build in a way that protects life — human and more-than-human — and to name our own blindspots out loud so others can catch what we miss.

Roots: non-maleficence in medical ethics ("primum non nocere"). Tension: "do no harm" can quietly become "do nothing," letting fear of causing harm justify inaction that also harms (the DON'T blindspot). The honest version weighs harms of acting against harms of not acting — and keeps choosing, in the open.

Emergent Strategy practice — with gratitude

A framework by [adrienne maree brown](#) ("Emergent Strategy," 2017), drawing on Octavia Butler, for how intentional change grows through relationship, adaptation, and small consistent practice rather than top-down control.

Core principles we lean on: "Small is good, small is all (the large is a reflection of the small)"; "Change is constant (be like water)"; "There is always enough time for the right work." Why it's here: our whole method — cyclical not linear, fractal across scales, start small and iterate — stands on this lineage. We name it with gratitude and do not claim it as our own.

Fractal principle

"How we are at the small scale is how we are at the large scale" (adrienne maree brown). What we practice in one honest conversation sets the pattern for a team, an organization, a movement.

In this Field Guide: it's why the operating system scales identically from self to couple to team to living system — same questions, more voices. Tension: fractal thinking can excuse staying small ("if I just fix myself, the system fixes itself") and let structural harm off the hook. Personal practice and structural change need each other; neither substitutes for the other.

Part Four — Apply

Apply the Mindsets: Thinking About AI

Parts One–Three gave you the lenses, the self-map, and the vocabulary. Part Four is where they meet a real decision. This is the bridge from understanding to action.

From knowing to choosing

You now have four mindsets, a map of your own operating system, and a plain-language library. Part Four puts them to work on an actual question: should I bring AI into THIS? Not in the abstract — this task, this relationship, this part of your life or business.

Most people meet AI in one of two postures: swept up (all COULD) or braced against it (all DON'T). Both are single-mindset reflexes. Attunement is the deliberate act of visiting all four before you decide — and noticing which one you skipped.

Why this matters, lived: every real decision is made in the present, with incomplete information, to shape a future you can't fully see — not certainty, but attunement. The four mindsets don't remove the unknown; they help you meet it with a clear head instead of fear or hype. That is the whole practice, whether the decision is a big personal one, a hire, or whether to let AI into one corner of your work.

The Attunement Pass five minutes

- **COULD** — What becomes newly possible here? (Imagine before you judge.)
- **SHOULD** — What's the disciplined, values-aligned way to do it? Whose standard am I using?
- **MIGHT** — What are the best- AND worst-case futures if I start? (Rehearse both.)
- **DON'T** — What would I refuse? What stays human no matter what? Then: Which voice was loudest for me — and which did I skip? Go back and give the skipped one a real turn.

The order matters. **COULD** first keeps fear from foreclosing possibility. **DON'T** last means you decide your refusals with a clear head, not a racing heart. And the closing question is the whole discipline: your default is not your destiny.

Spot your default, borrow the quiet one

Each of us has a home mindset — the voice that speaks first and loudest. Attunement isn't abandoning it; it's borrowing the ones you skip.

- The optimist (lives in **COULD**) → deliberately adds **DON'T**. What am I not seeing?
- The refuser (lives in **DON'T**) → deliberately adds **COULD**. What might I be missing out on?
- The planner (lives in **MIGHT**) → deliberately adds a first step. Analysis has an expiration date.
- The dreamer (lives in **COULD**) → deliberately adds a guardrail. What protects this? Your growth edge is the mindset you find least comfortable. That discomfort is the signal, not the stop sign.

Worked examples same Pass, different right answers

The template, once:

When my default ([**loud mindset**]) meets [**this decision**], my quiet one ([**quiet mindset**]) asks: [**the question it would ask**]. So the smallest safe step is [**one reversible action**], and the line I won't cross is [**the don't**].

The coach / helping professional

When my default (*Could*) meets *AI for session prep and notes*, my quiet one (*Don't*) asks: *whose relationship is this?* So the smallest safe step is *AI in the prep, never in the presence*, and the line I won't cross is *the human alliance is the work*.

COULD: AI drafts session prep, summarizes notes, surfaces patterns across months. SHOULD: follow the ICF AI coaching standards; disclose AI use to clients. MIGHT: best case, more presence in the room; worst case, subtle dependence and the sycophancy trap. DON'T: never outsource the relationship. The human alliance IS the work. Result: AI in the prep, never in the presence.

The solopreneur

When my default (*Could*) meets *automating invoicing, scheduling and first drafts*, my quiet one (*Don't*) asks: *what do people actually pay me for?* So the smallest safe step is *automate the admin*, and the line I won't cross is *the signature craft stays mine, in my voice*.

COULD: automate invoicing, scheduling, first-draft copy. SHOULD: keep your voice; disclose where content is AI-assisted. MIGHT: best case, hours back for craft; worst case, your brand sounds like everyone else's. DON'T: protect the signature craft — the thing people actually pay you for. Result: automate the admin, guard the art. (Echoes Thomas Lim: "focused use restored clarity and control.")

The team or business partnership

When my default (*whichever of us speaks first*) meets any shared AI decision, my quiet one (*the voice we both skip*) asks: *what do we feel, believe, know, and not know, out loud?* So the smallest safe step is *run the Pass together before deciding*, and the line I won't cross is *no shared don't decided by one of us alone*.

Run the Pass together, out loud, so "we feel / we believe / we know / we don't know" is real, not assumed. Surface the quiet voice on purpose — so the loudest person in the room isn't the only mindset in the room. Result: shared don'ts become shared sovereignty.

The educator

When my default (*Don't*) meets a *bilingual tutor for every child*, my quiet one (*Could*) asks: *what could each child have that they couldn't before?* So the smallest safe step is *one class, human-in-the-loop, the teacher still teaching*, and the line I won't cross is *no high-stakes call made by a machine alone*.

COULD: a bilingual tutor for every child, feedback at 2 a.m.

SHOULD: New Zealand's stance — a practical, educative approach, human-in-the-loop. **MIGHT:** best case, personalized support at scale; worst case, atrophied judgment and hollow assessment.

DON'T: the human relationship of learning stays central; no high-stakes call made by a machine alone. **Result:** AI extends the teacher; it doesn't replace the teaching relationship.

Applying it with others scaling the Pass

The four questions don't change as you scale — the number of voices does. Self → run the Pass in your own head, honestly.

Relationship / couple → run it aloud; two defaults meet and

balance. Team → assign the mindsets on purpose so all four get a

seat. Organization → make the Pass a ritual before any AI adoption

decision. Living system → ask the widest question: does this help all life thrive, or only some?

Holding multiple futures is the same skill as holding multiple perspectives. That skill is attunement.

Your turn

When my default (__) meets ____, my quiet one (__) asks: ____? So the smallest safe step is ____, and the line I won't cross is ____.

Part Five — The Doing

Implementation: Cyclical, Not Linear

The close of the LEARN foundation — and the hand-off into Design and Act. Implementation is a loop you live, not a ladder you finish.

There is no finish line

Most AI-adoption advice is a checklist: adopt these tools, in this order, done. Ours is a cycle. You attune, you try something small, you learn, you re-attune. Capacity liberation — freeing your time and attention for the human work only you can do — is the aim, and it's never "finished," because you keep changing and so does the technology.

The shipwreck we named in Part One — bolting AI onto a system you never listened to — happens precisely when people treat this as linear. A loop protects you from that.

Feel-believe-know-don't know: each turn of the cycle updates all four. What you "know" this month becomes "used to believe" next month. That's not failure — that's iteration.

The Cycle

- **ATTUNE** — run the Pass (Part Four); check your operating system (Part Two).

- TRY — the smallest safe experiment: one workflow, one tool, one week.
- NOTICE — what actually changed? In output, energy, relationships, and values?
- RE-MAP — update your operating system and your don'ts; record what you learned. → Repeat, at whatever scale you're working. You can enter the loop anywhere. Overwhelmed? Go back to ATTUNE. Stuck in planning? Force a TRY. The loop is the cure for both paralysis and recklessness.

One loop, worked: A coach wants to try AI for session notes. ATTUNE: runs the Pass — COULD (faster notes), DON'T (client confidentiality is a hard line). TRY: one week, AI drafts notes only from her own typed summaries, never raw recordings. NOTICE: saved twenty minutes a session, but the language felt generic — and it freed her to sit longer with what a client actually said. RE-MAP: she keeps it, adds a don't ("no client identifiers, ever"), and names what the freed time is for: deeper presence. That's one turn. Next month, a new question.

Start small, start safe

- The smallest reversible step beats the grand rollout. You learn more from one real workflow than from a six-month strategy deck.
- Human-in-the-loop by default. Verify before you trust. The machine drafts; you decide.
- Protect the premium. Keep the human touch where it's the whole point. Don't automate your differentiator.
- Month one is harder than month six. The early stretch is a learning cost, not a verdict. Expect the dip.

30 Days of First Help

Most AI advice is a checklist. Ours is a loop: attune, try something small, notice, re-map. Thirty days is long enough to go around it once, at a walk. You won't finish anything. You'll start something you can keep.

One rule for the whole month: you are learning a language, not building a machine. Clumsy is the correct feeling. It passes.

Read each block on its first morning. Do the small thing. Don't read ahead.

Days 1 to 5 — Know yourself first

Before you open any tool, open your own operating system. Answer the five questions in Part Two, one per layer, by hand or typed, rough and true: when and how you do your best work; who you are and what you do; the three values that govern your decisions; who you serve and what they need from you; your core recurring workflows. Or open the builder and answer them there (attunecoachinggroup.com/build-your-aos); it keeps your answers only in your own browser, nothing is sent anywhere, so download or copy them somewhere that's yours when you're done.

Then write one sentence that starts with "I will never let AI..." and one that starts with "I would love AI to..." Keep both where you can see them.

By day 5, you should have: five honest answers, two sentences, and no AI account you didn't already have. See Part Two, *Know Yourself*.

Days 6 to 10 — The three pauses, practiced empty

Learn the habit before you need it. Three pauses, said out loud until they're yours: *pause before you paste* (is this mine to share?), *pause before you trust* (how would I check this?), *pause before you connect* (what does this tool now reach?).

This week you don't use AI for anything real. You practice the pauses on ordinary moments: before you forward an email, before you repeat something you read, before you install anything. You're building the reflex in a low-stakes room.

By day 10, you should have: the three pauses memorized, and one moment you can name where a pause changed what you did. See Part Four, *The Attunement Pass*.

Days 11 to 15 — One small, read-only help

Now pick one thing. Not the biggest problem. The smallest repeating chore that touches nobody else's private information. Summarizing an article you'd otherwise skim. Turning your own bullet notes into a paragraph. Drafting the email you've rewritten three times.

Open one tool. Any tool; we're tool-agnostic on purpose. Ground it in who you are: paste the sentences from days 1 to 5 and say, "This is how I think. Help me from here." Then ask for the small thing. Read the output as a draft from a new colleague: useful, unfinished, yours to decide.

Here is the prompt, in full, yours to run anywhere:

Here is how I think and what I care about: [your five answers]. Here is what I will never let AI do: [your sentence]. Before you help me with anything, ask me three questions so you understand what I actually need. Then, and only then, help me

with this: [the small thing].

By day 15, you should have: one chore done with help, one week in a row, and a clear feeling about whether it was worth it. See Part Two, *Layer 5, Systems*.

Days 16 to 20 — The transparency conversation

If you work with people (clients, students, a team), tell one of them what you're doing. Not all of them. One. Say it in your own words; here is a shape to start from:

"I've started using an AI tool to help with [the small thing]. It's trained on my values, not my voice: it helps me think, it doesn't speak for me. Nothing you tell me goes into it without your say-so. I wanted you to hear that from me. How does that sit with you?"

Then listen. Their answer is data about your practice, not a verdict on AI.

By day 20, you should have: one person who knows, one conversation you didn't dread as much as you expected, and the exact words you'll use next time. See Part Four, *The coach / helping professional*.

Days 21 to 25 — Read it back

Take one week of your own work (notes, a journal, a session summary you wrote yourself, never someone else's raw words) and ask the tool to hold up a mirror. Four questions, in this order: What patterns do you see in how I worked? Where did I do what I said I value? Where did I drift from it? What would you ask me that I haven't asked myself?

Anonymize first. Replace names with roles. Keep the answers private; this is coach-to-self, not content.

The fourth question is the one that matters. Jodie Cook's playbook calls the AI reading your transcripts a "second brain" and a "mirror" (p. 85). We'd add: a mirror only shows you what you already carried into the room. Look for the values, not the verdicts.

By day 25, you should have: one honest reflection you didn't expect, written down, and one "don't" added to your never-let-AI sentence. See Part Two, *Layer 3, Constitution*.

Days 26 to 30 — Keep it yours

The last week is about what the freed time is for. Name it before you free it, or the gap fills with more work. Write the sentence: "The minutes I saved this month go toward ____." Presence, rest, a longer walk with a client's question, the craft you'd been skipping.

Then re-map. Open your five answers from day 1 and edit them. Something you "knew" is now something you "used to believe." That's not failure; that's the loop working.

Read your edited answers out loud. If they sound like something you would actually say to a person you care about, you're done. If they don't, cut until they do.

By day 30, you should have: one small help that stays, one that you stopped, an operating system one turn truer than it was, and the freed time pointed somewhere on purpose. See Part Five, *The Cycle*.

After day 30

Go around again, at whatever scale you're working. Overwhelmed? Go back to days 1 to 5. Stuck in planning? Force a days 11 to 15. The loop is the cure for both paralysis and recklessness.

What you started with is still all you need: knowing your own values, and the patience to keep practicing. The thirty days just gave it somewhere to go.

Capacity liberation — the actual point

This is the heart of it, so we say it plainly: the goal is not to do more. The goal is to free yourself for what matters.

- Name what the freed time is FOR before you free it — otherwise the gap just fills with more work.
- Guard against the trap of infinite output: just because you can generate endlessly doesn't mean you should.
- Measure success by what the freed time goes toward — deeper relationships, better craft, rest, presence, the work only a human can do.
- This is abundance in the sense Part Three honors: wealth as thriving, not just accumulation.

The hand-off: Learn → Design → Act

Where the Field Guide meets the paid course — stated plainly.

LEARN (this open Field Guide + Substack): the framework, the four mindsets, the reference library, the open operating system. No paywall, ever. DESIGN (the course, one-time fee): build your own full operating system step-by-step, with our videos, stories, examples, implementation options, and the connection apps we keep updated. ACT (community): join the conversation and iterate alongside others — because transformation is easier with company.

The line, restated: information is open; depth, story, structure, and belonging are the offer. Pay if it helps you and you're able. If not, the open resources are genuinely rich — take what you need. Everyone is welcome.

Closing — on purpose

This is us being on purpose: it is our pleasure to do good around how life thrives. We built this to do no harm and to hand you a clear view of where we stand. Now it's yours — to use, to question, to build on.

You started with yourself — knowing your own values, and the patience to keep practicing. That's still all you need. The cycle simply gives it somewhere to go.

Attunement is a verb. Let's practice it together.

Continue the conversation

You've read our point of view, and the curation that shaped it. Now the invitation — three doors, and the first two are free.

Read. This Field Guide is yours, open, always. Share it with anyone.

Continue — on Substack (always). Where the conversation keeps going, in digestible installments, alongside a growing community. Open forever; if the work helps and you're able, you can donate — but money is never the door.

Build — the course (Shape AI) + live cohort. When you're ready to build your own operating system with support: a one-time fee (never a subscription) that's yours to keep — access forever. You get the self-paced course to build your operating system step by step, and a seat in the live cohort — real events, workshops, and a community iterating alongside you. A personalized experience, not a video you watch alone.

Pay if it helps and you're able. The open path is rich. We refuse to let money or mystery keep anyone from thinking clearly about their future.

Appendices

Appendix A — The transparency script

For anyone who works with people. Say it in your own words; here is the shape. It departs from the Coachvox version (Jodie Cook, playbook p. 85) on one line by design: values, not voice.

"I want to tell you something about how I work now. I've started using an AI tool to help with [the small thing: notes, drafting, prep]. It's trained on my values, not my voice: it helps me think, it doesn't speak for me. Nothing you tell me goes into it without your say-so, and when it does, your name doesn't. If it ever notices something I'd bring to you, I'll say so out loud, and we'll look at it together. I wanted you to hear that from me before anything else. How does that sit with you?"

Then listen. If the answer is no, that's a complete answer. You don't have to use it with everyone.

Appendix B — The read-it-back prompt

One prompt, run on your own week of work. Four outputs. The fourth is ours; the first three are common to most "second brain" methods, including Coachvox's (pp. 86–87). Anonymize first: names become roles.

Here is one week of my own work notes, anonymized. Here is my operating system: [paste your five answers and your two sentences]. Read the notes against the operating system and give me four things, in this order, each in five lines or fewer:

1. **About the people I worked with:** what patterns do you see across the week? What did they keep returning to?
2. **About me:** where did I do what I say I value? Where did I drift from it? Quote the moment.
3. **About what I could share:** three things I said or asked this week that someone outside the room would find useful, with all identifying detail removed.
4. **About my values:** which of my stated values actually showed up in how I worked, and which one didn't appear at all? Then ask me one question I haven't asked myself.

Before you answer, ask me any question you need to. Don't flatter me.

Appendix C — Words we don't use

A short list, kept in one place. Some are ours; the last six are borrowed with thanks from Jodie Cook's list of "AI signals" (playbook p. 56), because she's right about them.

We don't say "an AI that sounds like you," "the AI version of you," or "mimics you." We say *trained to operate from your values*. Style stays human; judgment goes in the machine.

We don't open with "I'll be honest" or "honestly," because it implies we weren't before. We don't cluster "here's the thing."

We don't say "unlock your potential," "game-changer," "deep dive," "leverage," "move the needle," or "It's not about X, it's about Y."

We don't put prices in writing. The price lives at checkout.

We don't say a thing is finished.

Appendix D — Five principles we hold

Numbered, the way principles should be. Ours; the shape is common to every good ethics page, including Coachvox's Appendix 1.

1. **Tell people.** Anyone whose words or work touch an AI we use hears it from us first, in plain language, and can say no.
2. **Stand by every word.** Nothing goes out under our name that we haven't read and would defend. The machine drafts; we decide.
3. **Stay in the room.** The point is more presence, not less. If a tool makes us less present with a person, it goes.
4. **Keep the human data human.** Names become roles before anything is uploaded. One project per person. Nothing identifiable, ever.
5. **Keep learning, out loud.** What we "know" this month may be "used to believe" next month. We say so when it happens, with sources.

Appendix E — Talking to a Young Person

One of the most common things people ask us: what do I even say to a young person about this? Here is what we say. This is not a briefing you deliver — it is a conversation you have. Ask, then listen. They will teach you things.

AI is already here — that is the reality we start from. So we have a double responsibility: to guide young people, and to keep learning ourselves, from them and with them. The goal is understanding, not fear or hype. A young person who sees the big picture stays curious and in charge — instead of rejecting AI, or being quietly shaped by it.

- Have you ever felt like everyone already ‘gets’ something new and you don’t yet? What helped you catch up?
- If a computer could help you make one thing today, what would you make together?
- When something sounds really sure of itself, how do you figure out if it’s actually true?
- What’s one thing about AI that sounds exciting to you — and one thing that feels a little worrying?
- What’s a job you’d be happy to let a computer help with — and one you’d always want to do yourself?
- If a helper was going to help you, what are three things they’d need to know about you first?
- What’s the difference between asking for something clearly and asking in a confusing way?
- When something you made isn’t quite ‘you’ yet, how do you change it to feel more like you?
- What’s a boring chore you’d hand to a helper — and what would you still want to check yourself?
- How can you tell when something stops sounding like you?
- When something’s hard, does doing it with someone help? Who do you like to figure things out with?
- Is it easier to share something perfect, or something you’re still working on? Why?

- What's something you understand well enough that you could teach it to someone else?

Appendix F — Who We're Learning From, and Why

Why so many voices? Because a single source is a monoculture. We gather them — local, regional, global; builders and skeptics — to grow one collective point of view:

honest & hopeful · critical & curious · proactive & secure.

How a source earns a spot here:

- Reputable & ethically regarded — named, accountable, trusted by people who care about doing this well.
- Local · Regional · Global — never a single monoculture.
- Multiple perspectives on purpose — builders and skeptics, so you learn to think, not to be sold.
- Human-first & current — attends to people and values, and kept alive as things change.

Voices we learn from

Nine, as of this edition. Vetted, diverse, dated; the web edition keeps this list alive.

- **Ethan Mollick — One Useful Thing** — Practical and research-based. The builder's lens, without the hype. (<https://www.oneusefulthing.org/>)
- **Gary Marcus** — The skeptic's line — keeps us honest about what AI can't do. (<https://garymarcus.substack.com/>)
- **The Convivial Society — L. M. Sacasas** — Technology and the human, ethical questions underneath it.

(<https://theconvivialsociety.substack.com/>)

- **adrienne maree brown** — Where our own language comes from — emergent strategy, fractal.
(<https://adriennemareebrown.net/>)
- **Daniel Kokotajlo — on Diary of a CEO** — The cautious case, held with respect — the forecaster who left OpenAI rather than stay quiet. His own hopeful counter: which jobs survive is a political question, decided by people who learn.
(<https://open.spotify.com/episode/2IJedIFiDjU09qGxns1gbv>)
- **Dr. Roman Yampolskiy — on Diary of a CEO** — Coined "AI safety." The hard skeptic's line — paired with his own hopeful counter: love and meaning stay real regardless, and attention is a choice. (<https://podcasts.apple.com/us/podcast/roman-yampolskiy-these-are-the-only-5-jobs-that-will/id1291423644?i=1000724933846>)
- **Jeremy Grantham — on Diary of a CEO** — The bubble-caller. The idea and the froth are different things — the idea remakes the world after the crash. (<https://www.youtube.com/watch?v=vX6Cl4yaVuo>)
- **Dr. Fei-Fei Li & David Rogier — on Silicon Valley Girl** — Agency is the defining skill of the decade — even the MasterClass CEO trains AI on his own writing.
(<https://podcasts.apple.com/us/podcast/silicon-valley-girl-ai-tech-and-career-growth/id1819090545>)
- **Sinéad Bovell — The Great Cognitive Divide** — Thinking is a practice you keep or lose. Use AI as a sparring partner, not a substitute. (<https://sineadbovell.substack.com/p/the-great-cognitive-divide>)

Staying current — Local · Regional · Global

- **Local, wherever we are:** Nicki's Substack, nickimollet.substack.com. On the ground.
- **Regional, our country:** The Markup, themarkup.org. US public-interest tech journalism.
- **Global, the world:** MIT Technology Review, technologyreview.com.

Appendix G — The Ancestry This Work Stands On

With gratitude. We name our sources because reciprocity is the practice — wealth as thriving includes crediting whose thinking made ours possible. Living document: we add, we correct, we never quietly erase. Where a work is still to be confirmed, we hold the space openly rather than guess.

Published thinkers & traditions

- [adrienne maree brown](#) — "Emergent Strategy" (2017) and the wider Emergent Strategy Series (AK Press).
- Octavia Butler — the Earthseed verses ("Parable of the Sower," "Parable of the Talents") that seed Emergent Strategy.
- [Robin Wall Kimmerer](#) — "Braiding Sweetgrass" — reciprocity, gratitude, and the grammar of animacy.
- Kieran Egan — Imaginative Education (Nicki's own graduate lineage, Simon Fraser University). U Chicago Press
- Dr. Gillian Judson — co-director of the Imaginative Education Research Group (IERG) at SFU, Egan's close collaborator; developer of Imaginative Ecological Education. Blog: [imaginED\(educationthatinspires.ca\)](http://imaginED(educationthatinspires.ca)). Her work carries Egan's imagination-

first pedagogy into ecological and place-based understanding — a direct bridge to imagining a livable future.

- Dr. Natalia Gajdamaschko — Vygotskian psychologist, Teaching Professor at SFU and an IERG voice; the scholar who brought Vygotsky-on-imagination into the North American conversation. It was in conversation with her that Nicki worked through recapitulation theory — how the cognitive tools invented across cultural history are re-lived in each developing mind. Teaching Education
- The IERG / CIRCE circle — after Egan's retirement (2015) and passing (2022), IERG was restructured as CIRCE (the Centre for Imagination in Research, Culture & Education). This is the intellectual home Nicki studied inside. CIRCE
- Lev Vygotsky — the Zone of Proximal Development; the social nature of learning.
- Sor Juana Inés de la Cruz — 17th-century Mexican scholar, poet, and proto-feminist; a founding voice for women's right to learn and think. An original badass.
- Inés Suárez — 16th-century figure of Chilean history; a woman who refused the edge of the story assigned to her.
- Isabel Allende — memory, women's interior lives, magical realism. • Gabriel García Márquez — magical realism as a way of telling truth. • Octavio Paz — solitude, identity, the labyrinth of culture. • Gabriela Mistral — Chilean poet-educator, first Latin American woman to win the Nobel in Literature.
- [Clarissa Pinkola Estés](#) — "Women Who Run With the Wolves"; myth, story, and the wild feminine as a way of knowing.
- Folk tales & oral tradition — the anonymous, communal storytellers whose tales carry wisdom across generations.
- La Malinche (Malintzin / Doña Marina) — held with complexity: interpreter, bridge, and contested symbol at the meeting of

worlds.

- [Madeline Miller](#) — "Circe," "The Song of Achilles"; retelling myth from the margins. • [Emily Wilson](#) — first woman to publish an English translation of Homer's Odyssey; how the translator's lens reshapes what we inherit. • Homer / the Odyssey tradition — the deep well these retellings draw from.
- Music & contemporary voice: Rosalía — reinventing flamenco through a modern lens; Shakira — global crossover, bilingual identity, reinvention; Karol G — contemporary Latina self-authorship; and global pop currents as living, cross-cultural text. Widening the circle — a commitment: we want this ancestry to be eclectic and honest — not only women, not only one canon. We commit to including strong male voices; queer, trans, and nonbinary voices; voices from underrepresented communities; and voices from every generation and life stage — elders and the young alike. We especially honor the crone : the elder-woman stage as a seat of authority, not invisibility.

Teachers & mentors

Michelle of Holisticism — who taught Nicki to live in her zone of proximal development, and to iterate, iterate, iterate. Monica DeSimone — and her work.

Family — published work with love

A rare and beautiful thing: a family of published thinkers. Names and titles held open until each is confirmed exactly right — Nicki's father ("with the mice"), mother, sisters Kayla & Sarah, brothers Daniel & Thomas, and Kathryn's mother, Marshan Snyder .

Open questions we're holding

Attunement means asking who a future is actually for. What could an AI future look like for the crone stage of life — for menopause, hormone replacement therapy, and midlife-and-beyond health,

long under-researched and under-served? How do we build a future that serves them well, rather than one more system that renders older women invisible? And what does an AI future look like across every generation — not optimized only for the young or the already-connected?

Where we're writing from positionality

This work was built by two cis, white-bodied women taking up space on Turtle Island — rooted in Tongva territory (Los Angeles / South Bay), and written across many lands: throughout Turtle Island, through Mesoamerica, and in Maōri space in Aotearoa. We honor the ancestors of all these places — the Indigenous peoples who, long before colonization, built strong economic systems that supported women, and that supported life itself, not just human life. We name our positionality so our point of view is transparent: whose eyes this is written through, and whose wisdom we are indebted to and learning from.

Nicki — river-born in Memphis, farm-raised in South Dakota; Panama at eighteen (Portobelo, then Los Santos on the Azuero Peninsula); Northern Arizona University with a semester in Valencia; a European chapter (Morocco, Portugal, Westphalia, Vienna — working with the unemployed); Madrid as a home of transition seasons and Iberian literature; a TEFL summer through Costa Rica, Honduras, and Nicaragua with her brother Daniel; eleven years home in Santiago, Chile, with deep time across Argentina, Peru, and the Atacama; Cartagena; family ties in Ecuador and Mexico. In the US: Minneapolis, Lincoln, Vermillion, San Pedro. She counts all of these as home. How it shapes the lens: formed between rootedness and departure; humbling early immersion in other cultures and languages; a decade-plus rooted in

the Southern Cone; a genuinely multi-homed sense of belonging — rural American, urban American, and Latin American at once. A lens shaped by many centers, not one .

Kathryn — she grew up between Japan and Germany, trained as an engineer in Virginia, and spent her early career inside corporate technology firms before becoming an executive coach — work she still loves, because she genuinely enjoys changing the world. How it shapes the lens: a childhood across cultures and a career spent translating between the technical and the human let her meet AI with both an engineer's respect for what it can do and a coach's insistence that people, values, and presence come first. A lens formed across continents and between worlds — the technical and the human, held together.

Principle: this is disclosure, not credential — naming influence and possible blind spots, in the same spirit as our land acknowledgment and lineage pages.

Sources

Every link in this edition, in order of appearance. Verified for the web edition in July 2026; re-verified for the book at build time.

- “Could Should Might Don’t: How We Think About the Future”: <https://www.couldshouldmightdont.com/>
- Evidently AI: <https://www.evidentlyai.com/>
- Arize: <https://arize.com/>
- Anthropic, "Persona Vectors," 2025: <https://arxiv.org/abs/2507.21509>
- Open Secure AI Alliance: <https://blogs.nvidia.com/blog/open-secure-ai-alliance/>

- Anthropic's position on open-weight models:
<https://www.anthropic.com/news/position-open-weights-models>
- NTIA report on widely available model weights:
<https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report>
- The Open-Weight Paradox: <https://arxiv.org/abs/2604.17413>
- Andrej Karpathy: <https://x.com/karpathy>
- Expanding on what we missed with sycophancy:
<https://openai.com/index/expanding-on-sycophancy/>
- Cheng et al., Science, 2026:
<https://www.science.org/doi/10.1126/science.aec8352>
- adrienne maree brown: <https://adriennemareebrown.net/>
- Robin Wall Kimmerer: <https://www.robinwallkimmerer.com/>
- Clarissa Pinkola Estés: <https://www.clarissapinkolaestes.com/>
- Madeline Miller: <https://madelinemiller.com/>
- Emily Wilson: <https://www.emilyrcwilson.com/>

Framework credit: Nick Foster, *Could Should Might Don't* (MCD/FSG, 2025).

Borrowed shapes: the roadmap, the transparency script and the 'words we don't use' list follow shapes in Jodie Cook, *The AI-Powered Coaching Business Playbook* (Coachvox, 2026), pp. 56, 85–87, 100–101. The content is ours; where we differ, we say so.

Attuned — An AI Field Guide · Attune · attunecoachinggroup.com · Free and open, always.