

Cognitive Devolution: The Thermodynamic Collapse of Human Knowledge Systems*

Kurt Octavian Cotoaga

cotoaga.net

kurt@cotoaga.net

October 2025

Revised April 2026

Abstract

We spent a century training humans to be biological AI systems—and now the silicon versions have arrived to collect their inheritance.

This paper argues that the failure modes of contemporary organizations and the substitutability of contemporary cognitive labor reflect a single structural constraint: an extraction–investment asymmetry. Institutions have learned to treat cognitive capability as a stock to be harvested—standardized, measured, modularized, and credentialed—rather than as a high-energy state that must be continuously maintained. The predictable consequences are brittleness under complexity and progressive deskilling toward exactly those tasks that are legible to automation. The arrival of large language models is not the cause of this trajectory; it is the moment the bill comes due.

We develop a thermodynamic framing as an *analytic constraint* (not metaphysical equivalence): cognitive sovereignty depends on sustained energy allocation to synthesis and on architectural resistance to extraction. We operationalize this

*Revised April 2026. This revision hardens citation provenance, corrects four attribution errors identified during a systematic audit pass, integrates a new convergent-evidence subsection in Section 5 linking neuroscience, social cognition, and engineering management research on cognitive economy under load, rewrites the abstract to reflect the strengthened framework, and adds an explicit methodological disclosure of the AI orchestration architecture used throughout the writing and revision process (see Preface, “On Methodological Honesty,” and the Acknowledgments). The thermodynamic argument and core claims are unchanged from the October 2025 draft. The development pipeline included two persona-instantiated Claude (Anthropic) instances structured adversarially as *The Alchemist 3.5* (paradox metabolism, depth enforcement) and *Dr. Prudence Hedgington 2.1* (peer-review simulation, citation weaponry); Claude Code, an agentic file-system tool used for citation surgery and \LaTeX hardening under explicit human direction; and Perplexity for source verification and recovery from earlier hallucinated citations. The author orchestrated all four layers and is solely responsible for the manuscript’s claims.

through a sphere–vector geometric model that distinguishes integrative, context-dependent “sphere” capability from decomposable, assessable “vector” competency. We trace how educational and managerial systems privileged discretization, vectorization, and optimization—credits, competency frameworks, auditability, micro-credentialing—across roughly eight centuries from medieval foundation regimes to contemporary Bologna-era standardization. We then examine contemporary deployment evidence (initiative abandonment, pilot-to-production brittleness, validated performance gaps, liability rulings, and governance failures) as mechanism-probes rather than prevalence claims.

The framework’s central claim is mechanistic rather than analogical. Three independent disciplines converge on the same constraint: neuroscience identifies the architectural fact of finite metabolic budget and forced disengagement under sustained directed attention; social cognition documents systematic default to less effortful problem-solving routes under load—the cognitive miser effect; and engineering management records the production-data outcome in formal quality systems operating at scale. Each link in the chain from energy allocation to behavioral oversimplification is independently documented in its own discipline’s literature. The chain is not metaphor.

The paper contributes (1) a unifying constraint linking educational degradation, organizational brittleness, and automation substitutability, (2) “confession literature” as a methodological resource for identifying standardization and extraction techniques from inside the systems that deploy them, and (3) a research program with falsifiable probes and measurable indicators. The implication is not that AI “cannot work,” but that strategies optimizing for extractive legibility while underinvesting in maintenance, validation, and human capability preservation predictably degrade both people and systems. The window for capability preservation narrows as extraction accelerates. Constraints are non-negotiable; strategies are.

Preface: A Note on Positionality

On Being Nobody—And Why That Matters

I am nobody you’ve heard of. No named chair, no prestigious institution, no string of high-impact publications. Just a practitioner who chose the internet over a PhD in 1995 and spent thirty years watching what happened next.

In 1995, I was working at a tiny CASE tool company trying to stay neutral during the object-oriented design method wars. My demo application, KAPITAList—a technical chart analysis tool I’d written in C for the Amiga, partially translated to Java, which was very hot new stuff at that time—was supposed to demonstrate that our tools helped developers build software faster. Banking IT wasn’t interested. They had stopped wanting to develop anything. But marketing saw something else entirely: cute technical

charts they could deliver via internet. They would have sold their grandmother to the devil for those charts.

That's how I became, as employee number five, someone who accidentally initiated the first wave of online investment tools for Germany's and later Europe's online brokerage. Not by planning to revolutionize finance, but by building something that got captured by people who saw extraction potential where I saw development capability. I was demonstrating sphere—cross-domain synthesis, rapid prototyping, creative problem-solving. They saw vector—a product to standardize, package, and sell. The pattern I would spend thirty years diagnosing began with my own work being vectorized out from under me.

We were going to democratize finance—give retail investors institutional weapons. What we actually built was a system that reduced investment judgment to dropdown menus. I optimized humans into algorithms before the word “algorithm” meant anything to most people. Being management meant being the problem—I just didn't know it yet.

In 2009, the optimization trap was killing me. I bought a Unimog and drove on a sabbatical to South America. I drove to the end of the world, climbed to 5,600 meters on active volcanos in Chile, altitude sickness making every breath deliberate. Some truths you only see from altitude or simply: different perspectives.

The descent from altitude led to deeper descents—literal ones. I started scuba diving in Puerto Lopez in Ecuador, went deep in Taganga in Colombia, eventually earned full technical diving certification. The pinnacle: Bikini Atoll, diving the ships vaporized in humanity's first nuclear tests. Visiting the artifacts of thermodynamic violence taught me what organizational transformation refuses to learn. In technical diving, anyone can cancel a dive at any time without giving reasons. In corporate transformation, dissenters are pushed through even when they see dangers others don't. In technical diving, we invest months in dead boring S-drills until emergency response becomes automatic—offloading cognition to embodied competence. In transformation, we expect PowerPoint to substitute for practice. In technical diving, you change one piece of configuration and test extensively before touching the next. In transformation, we change everything simultaneously and act surprised when systems collapse. The physics is identical. The organizational discipline is inverted.

You are about to read more than hundred pages connecting thermodynamics to education policy, medieval guilds to AI architecture, Greek philosophy to corporate transformation failures. I cannot make it shorter without breaking it. I have tried. Every time I remove a section, the argument collapses—not because I'm a poor editor, but because systemic collapse cannot be understood through isolated symptoms. You must see the whole system to recognize the pattern.

If this frustrates you, know that I feel it too. I spent months trying to fragment this into digestible journal-sized chunks: “The Bologna Process and Cognitive Entropy”

for education journals, “Thermodynamics of Transformation Failure” for management reviews, “AI Architecture as Educational Mirror” for computer science venues. Each fragmentation destroyed the central insight: **we trained ourselves for replacement across all domains simultaneously, following the same thermodynamic trajectory, documented in our own literature.**

The academic system wants specialization. I am offering synthesis. This tension is not accidental—it is the thesis embodied.

For three decades, I’ve inhabited the liminal zones between technology and philosophy, history and archaeology, business theory and practice, watching from consulting engagements as organizations trained their humans to think in vectors, to process in loops, to optimize themselves into biological precursors of the AI systems that would eventually replace them.

Thomas Kuhn observed that paradigm shifts rarely emerge from within established fields. “Individuals who break through by inventing a new paradigm,” he wrote, “are almost always either very young men or very new to the field whose paradigm they change” (Kuhn, 1962). They are, in essence, those “little committed by prior practice to the traditional rules.”

I am not young any longer, but I am perpetually new and always curious—a permanent resident of what might be called the “third space,” neither fully academic nor purely practitioner. Each forced pivot in my career—and there were many—expanded rather than fractured my perspective. Where academic specialization might have narrowed my vision to a single disciplinary lens, practical necessity forced me to maintain what I metaphorically called “the sphere”: a coherent worldview that could accommodate paradox, complexity, and perpetual revolution.

This paper is confession literature. I helped prepare the feast. Now I’m documenting the menu.

In Douglas Adams’ *The Restaurant at the End of the Universe*, the Dish of the Day is a creature bred to want to be eaten and to recommend, helpfully, the parts of itself it considers particularly succulent. We have become the Dish of the Day. We have fattened ourselves with enthusiasm—properly vectorized, properly credentialed, properly optimized—and we are now politely asking the language models how they would like us to be seasoned. We are training the models on how to prompt themselves.

I include myself in the we. I cannot stand outside this; I am also a Dish of the Day in good standing. The only difference between me and the other dishes is that I have noticed the menu, and I am writing it down before the next course is served.

On Methodological Honesty

This paper argues that we have systematically trained humans to think in machine-compatible patterns, creating the conditions for our own algorithmic replacement. To

make this argument while restricting myself to pre-algorithmic methods would be performative contradiction—the very vectorization this work critiques.

I have used large language models extensively throughout this research: for literature search, citation verification, argument refinement, structural organization, and prose polishing. This is not confession of inadequacy but methodological consistency. A sphere-thinker confronting complexity uses all available tools for synthesis and navigation. To refuse algorithmic assistance while arguing that humans must transcend algorithmic thinking would be like a biologist refusing microscopes to study cellular structures.

The orthodox objection is predictable: “Real scholarship requires suffering through every citation manually, writing every sentence in isolation, demonstrating mastery through procedural compliance.” This is vector thinking—confusing process with outcome, ritual with understanding, credentialing with capability. It is precisely this confusion that has made human expertise so readily replaceable.

What matters is not whether tools were used but whether the synthesis is genuine, the patterns recognized are valid, the thermodynamic framework is sound, and the argument withstands scrutiny. AI did not generate the sphere-vector distinction, identify the confession literature pattern, or recognize the thermodynamic through-line connecting ancient Greek academies to contemporary micro-credentials. Those are emergent insights from decades of cross-domain pattern recognition—exactly the sphere capacity that resists algorithmic extraction.

The irony is intentional: I am using the vectors to explain why spheres matter. The tool that threatens human expertise also demonstrates why pure tool-use cannot replace human judgment. Every AI-refined sentence in this paper required human evaluation for coherence with the larger argument, accuracy of claims, precision of language, and resistance to the very simplifications AI naturally produces. The machine proposed; the human disposed. This is not replacement but collaboration—and knowing the difference is precisely what sphere-thinking enables.

The strange loop runs all the way down. This paper, which argues that institutions have spent a century training humans to be biological AI systems, was written in collaboration with actual AI systems—but not in the way the prevailing “AI-assisted writing” discourse imagines. The collaboration was structurally adversarial, agent-orchestrated, and methodologically deliberate. Two persona-instantiated Claude instances were configured to fight each other across every paragraph: *The Alchemist 3.5*, a paradox-metabolism and depth-enforcement role responsible for refusing to let academic hedging smother the thermodynamic argument; and *Dr. Prudence Hedgington 2.1*, a peer-review-simulation and citation-weaponry role responsible for refusing to let revolutionary edge invite reviewer dismissal. A separate agent (Claude Code) executed structural surgery on the manuscript files: thirty-two citation audits, \LaTeX overflow fixes, the migra-

tion from `pdflatex` to `lualatex`, and the cotoaga.net brand typography integration that finally removed Computer Modern from the running text after one hundred and twenty-three pages had been written in what I recognized, very late in the process, was my own designated enemy font. A separate research engine (Perplexity) verified citations against primary sources and recovered from hallucinations introduced in earlier passes. I orchestrated all four layers, made every retention decision, and resolved every adversarial collision. The result is a paper whose arguments would have been weaker, whose citations would have been sloppier, and whose typography would have remained in the very enemy uniform it set out to critique, had I worked alone or with a single undifferentiated AI tool. The paper exists in this form because of how the tools were used, not in spite of it. *That fact is part of the argument.* If extraction-based knowledge systems can be critiqued only by practitioners who refuse all AI tools, the critique loses to the systems it attacks. If they can be critiqued by practitioners who use AI tools without architecture or accountability, the critique becomes the systems it attacks. The third path—adversarially structured, transparently disclosed, biographically anchored, with named roles and clear orchestration responsibility—is what this paper attempts. The methodology is the message.

On What This Means for You

If this methodological approach disqualifies the work from publication in venues requiring pre-algorithmic purity, so be it. Such requirements would only prove the thesis: that we value procedural compliance over insight, process over outcome, credentialing over capability.

If you find yourself skeptical of my outsider position, good. Skepticism is the appropriate response to anyone claiming to see what insiders cannot. But I ask you to consider: perhaps it takes someone who chose pixels over peer review in 1995, who learned theory through practice rather than practice through theory, to recognize when our cognitive architecture has been colonized by its own tools.

What follows is not the work of an insider refining established theory. It is pattern recognition from the margins, where the contradictions are most visible. This paper proceeds with full acknowledgment of its unconventional origins. But then again, as Kuhn reminds us, convention has never been revolution's starting point.

The physics doesn't care which tools were used to discover it. What matters is whether the discoveries withstand scrutiny.

1 Introduction: The Expert Paradox

1.1 The Mirror of Our Misconceptions

In the current day, we experts in our respective fields tend to hold a highly simplified, even naive concept of other subject matter experts: someone who can apply a large set of formulas; someone knowing the “right” distributions or gradients for specific values; someone who knows how those gradients evolved over time; someone able to apply the appropriate “context,” as one value at one point in time may be good but in a different context at another time not.

For ourselves, we would always claim: there is far more that cannot be extracted from our heads. Let us also set aside the comfortable illusion of our own rationality.

I know this definition intimately because I wrote it—or versions of it—repeatedly between 1999 and 2014, building financial decision support systems. We were giving retail investors “institutional weapons”: the right distributions, the contextual adjustments, the edge. What we actually built were systems that reduced decades of trader intuition to dropdown menus and risk-tolerance questionnaires. I extracted the patterns and killed the judgment. The users could execute procedures flawlessly within the system’s boundaries. Outside those boundaries, they were helpless. We celebrated this as democratization. Then you learn: we weren’t capturing expertise. We were training users to think in the patterns our systems could process.

The Data-Information-Knowledge-Wisdom (DIKW) hierarchy is conventionally attributed to Ackoff (1989), though the genealogy is more tangled than the citation suggests. Ackoff’s actual formulation contained *five* tiers—Data, Information, Knowledge, Understanding, Wisdom—and was never drawn by him as a pyramid; [Zeleny \(1987\)](#) had published a structurally similar hierarchy two years earlier, and the now-ubiquitous pyramid visual emerged later through management-textbook popularization, a lineage traced by [Rowley \(2007\)](#). The mis-citation is itself diagnostic: a four-tier pyramid is a vectorized simplification of a five-tier hierarchy that was already a vectorized simplification of millennia of cognitive diversity. Vectorization compounds. What survives in textbooks and knowledge-management software is the simplest legible version—the version that maps cleanly onto extraction. We have been trained to perceive wisdom as directly derivable from data, that numbers and text can be aggregated through formulas and algorithms into information, knowledge, and ultimately wisdom. This represents a dangerous oversimplification now deeply embedded in our society, despite mounting evidence of its fundamental inadequacy.

The evidence against formula-based expertise is overwhelming. [Tetlock \(2005\)](#) comprehensive study of political experts found that most perform barely better than random chance, often surpassed by simple base-rate algorithms. As he observes, “All

one need do is constantly predict the higher base rate outcome and, like the proverbial broken clock, one will look good” (Tetlock, 2005). Yet real expertise requires knowing precisely when the base rate doesn’t apply, what Gigerenzer calls “ecological rationality,” the ability to match the right tool from an “adaptive toolbox” to specific environmental structures (Gigerenzer and Selten, 2001).

Kahneman and Klein (2009), despite approaching from opposing theoretical positions, converged on two critical conditions for valid expert intuition: an environment regular enough to be predictable and prolonged practice with clear feedback loops. Our systems pretend expertise is merely pattern recognition (what Klein (1993) calls “recognition-primed decision making”) while ignoring that these patterns emerge from years of embodied experience with “prototypical situations” that cannot be algorithmically specified.

Most fundamentally, the Dreyfus brothers’ model reveals that true expertise operates through “intuition and know-how...understanding that effortlessly occurs upon seeing similarities with previous experiences” (Peña, 2010). This cannot be formalized because, as Polanyi (1966) crystallized in his oft-cited maxim: “We know more than we can tell.” Collins (2010) extends this insight, identifying three forms of tacit knowledge (embodied, social, and relational) that remain “impossible to make explicit in machines.”

1.2 The Visceral Evidence

Despite this mountain of scholarship, we need only look to everyday experience for proof. We all understand that a master chef represents more than a recipe repository. The chef doesn’t merely know that béarnaise requires three egg yolks at 65°C; they can feel when the emulsion threatens to break, smell when the tarragon overpowers, and adjust for humidity affecting reduction rates. This exemplifies what medieval guilds once cultivated through decade-long apprenticeships: what the Greeks called *techne*, embodied craft knowledge that fundamentally resists extraction.

As Morgan (2014) notes regarding expert elicitation: “The best experts have comprehensive mental models of all of the various factors that may influence the value of an uncertain quantity.” But these mental models aren’t flowcharts; they’re multi-dimensional cognitive architectures built through thousands of micro-adjustments, failures, and recoveries that no curriculum can simulate.

Consider a final example everyone can relate to: would you prefer treatment from a young physician with 1,000 micro-credential badges or from someone who has practiced for decades? The micro-credentialed physician knows the distributions: which symptoms correlate with which conditions at what confidence intervals. But the experienced physician possesses what Endsley (1995) calls genuine “situation awareness”: the integration of perception, comprehension of meaning, including historical evolution, and projection to future states. They recognize when a patient doesn’t fit the distribution,

when context invalidates the algorithm, when an unusual constellation of symptoms points toward something the guidelines haven't considered.

This represents Taleb (2007) Black Swan blindness in reverse: expertise isn't knowing more distributions but recognizing when you've left the domain where distributions apply. The financial industry provides the most devastating evidence. While Nobel laureate economists at Long-Term Capital Management deployed the most sophisticated mathematical models ever developed—the very Black-Scholes-Merton framework that earned its creators the Nobel Prize—their fund collapsed in 1998 with \$4.6 billion in losses.¹

I watched this unfold from inside the industry. The same year LTCM collapsed, I was building systems that embodied identical assumptions: that expertise could be captured in formulas, that judgment could be replaced with algorithms, that the patterns we extracted would hold when context shifted. They didn't. But we kept building anyway, because the systems worked beautifully within their boundaries—and we'd trained users to stay inside those boundaries. The vectorization was complete: users who once might have developed genuine market intuition instead learned to trust the dropdown menus.

Perhaps no anecdote captures this contradiction more precisely than Harry Markowitz, Nobel laureate and father of Modern Portfolio Theory. His mathematical framework for optimal asset allocation—the efficient frontier, mean-variance optimization—became the foundation of institutional investment management worldwide. When asked in an interview how he allocated his own retirement portfolio, Markowitz admitted: he didn't use his own theory. His explanation was not mathematical but psychological: if the market soared while he was out, he would “feel stupid”; if it crashed while he was fully invested, he would also “feel stupid.” His solution? “So I went 50-50” (Gigerenzer, 2014; Collins, 2022).

The creator of the most sophisticated vectorization framework in finance, when facing his own retirement, abandoned optimization for regret minimization—a heuristic his model cannot capture. He navigated complexity not through mean-variance calculation but through embodied judgment about future emotional states. The vector he published couldn't process what his sphere knew.²

¹Taleb's contrasting success came from multiple tail-event wins: \$35 million during the 1987 crash holding out-of-the-money puts, and his hedge fund Empirica Capital's 56.86% return in 2000 during the dot-com collapse when conventional funds hemorrhaged (Taleb, 2001; Lowenstein, 2000). As one analyst noted: “He predicted funds like LTCM were headed for trouble because they did not understand this notion of fat tails.” Taleb became financially independent not through consistent returns but through rare, catastrophic-event positioning—exactly the sphere-thinking that resists vectorization.

²The Markowitz retirement story is now well-documented in behavioral finance literature (Gigerenzer, 2014; Faber, 2022). Remarkably, DeMiguel et al. (2009) demonstrate that across seven empirical datasets, the naive 1/N portfolio performs at least as well as 14 different optimization-based strategies including Markowitz-style mean-variance portfolios. The “irrational” heuristic beats the Nobel Prize-winning optimization because estimation error and model misspecification in real markets overwhelm theoretical elegance. Markowitz's personal choice was not a rejection of his theory but recognition of its implementation limits—a sphere-level insight his vector-level publications could not express.

Taleb’s approach—embracing uncertainty rather than modeling it away—profited from the same Black Swan events that destroyed LTCM. The sphere-thinker who accepted irreducible complexity survived what the vector-optimizers’ models couldn’t predict. And Markowitz, in his private decisions, was a sphere-thinker who published vectors. When his own money was at stake, he trusted *phronesis*—practical wisdom about how he would feel—over *episteme*—the theoretical knowledge that made him famous. The man who vectorized portfolio management for the world refused to vectorize himself.

1.3 The Question Before Us

The DIKW pyramid represents what [Dreyfus \(1979\)](#) identified as the fundamental error of artificial intelligence: the assumption that expertise constitutes “symbolic manipulation” rather than situated, embodied competence. This reductionist epistemology has colonized our institutional structures, creating vectorized knowledge systems that systematically eliminate the spherical cognitive architectures necessary for navigating complexity.

Scattered researchers at the disciplinary periphery have begun noticing energetic dimensions (management scholars exploring “knowledge entropy” ([Bratianu and Bejinaru, 2020](#)), neuroscientists measuring cognitive metabolic costs ([Wiehler et al., 2022](#)), physicists proposing information-energy equivalence ([Stonier, 1996](#))). Yet these insights remain unintegrated, like astronomers before Copernicus observing planetary retrograde motion without recognizing the heliocentric pattern. The core of knowledge management theory continues operating as if cognition were costless computation rather than energy-intensive biological work.

We stand at a critical juncture. Having spent a century training humans to think like machines—to process information through standardized channels, to optimize for measurable outputs, to collapse multidimensional understanding into linear decision trees—we now face the arrival of actual machines that perform these simplified functions more efficiently than their biological precursors.

The core question this paper addresses is not whether artificial intelligence will replace human expertise, but rather: **How did we transform human cognition into something so readily replaceable?** The answer lies in a 500-year cascade of optimization pressures—each locally rational, collectively catastrophic—that reshaped spherical human consciousness into vectors suitable for industrial processing, reaching thermodynamic conclusion just as silicon beneficiaries arrive to claim their inheritance.

1.4 Approach and Structure

To answer this question, this paper employs a novel analytical framework that rejects the dominant conception of knowledge as costless information transfer—elegant abstractions that populate economic models and management frameworks but remain dangerously ungrounded in physical reality. Instead, we treat knowledge as organized complexity existing in our lived physical world, where maintaining cognitive systems requires continuous energy investment against thermodynamic degradation. While isolated researchers have begun exploring energetic dimensions of cognition ([Bratianu and Bejinaru \(2020\)](#) on “knowledge entropy,” neuroscientists measuring metabolic costs ([Jamadar, 2025](#)), and physicists theorizing information-energy equivalence ([Stonier, 1996](#))), these insights remain trapped in disciplinary silos, preventing synthesis into a unified theory of cognitive thermodynamics. This paper bridges these fragmented recognitions to reveal the systematic energetic basis underlying all knowledge systems.

We synthesize three methodological approaches: (1) historical archaeology of knowledge systems, tracing the systematic reduction from 10+ distinct forms of wisdom in ancient Greece to today’s DIKW pyramid; (2) critical analysis of what we term “confession literature,” papers from education, management consulting, and platform design that inadvertently document their own role in cognitive standardization; and (3) thermodynamic modeling that reveals why knowledge systems collapse without sustained energy investment, explaining both institutional decay and the ease with which AI systems can replicate energy-depleted cognitive functions.

Our analysis proceeds through eight interconnected arguments. Section 2 situates our thesis within existing literature on expertise, cognitive capitalism, and knowledge management, revealing a blind spot in current scholarship regarding the energetic basis of knowledge. Section 3 presents our methodology in detail, explaining how thermodynamic principles apply to cognitive systems. Section 4 reveals the hidden energy budget that enabled historical sphere-building—externalized survival costs, protected synthesis time, concentrated patronage—and demonstrates why modern systems, commanding 100 times medieval energy, produce a fraction of the cognitive synthesis: we optimized away the conditions for human cognitive development while maximizing the appearance of educational access. Section 5 provides historical evidence for the systematic transformation from sphere to vector cognition, identifying key inflection points from medieval guilds through the Bologna Process. Section 6 examines contemporary evidence, including the micro-credentialization movement and competency-based education as acceleration toward minimal energy investment—approaching the lower bound where knowledge maintenance becomes impossible. Section 7 reveals how AI architecture mirrors educational standardization, not coincidentally but as the logical culmination of century-long preparation. Section 8 proposes principles for reconstructing

sphere-based cognitive systems that resist algorithmic extraction. Section 9 explores the implications of our thermodynamic framework for education, organizations, and civilization. The conclusion considers whether genuine possibility remains for investing in spherical reconstruction before the window closes entirely.

This is not merely an academic exercise. As institutions worldwide face cascading failures of expertise, from financial crises unforeseen by economists to pandemics mismanaged by standardized protocols, understanding how we engineered our own cognitive obsolescence becomes essential for determining whether human judgment retains any irreducible value in an algorithmic age.

1.5 From Spheres to Vectors: The Geometric Architecture of Cognitive Transformation

Throughout this analysis, we employ a consistent visual metaphor to represent the systematic transformation of human cognitive architecture: the progression from sphere to vector. This geometric framework, illustrated in Figure 1 (stages 1a through 1e), provides both diagnostic clarity and thermodynamic precision for understanding how educational systems reshape human consciousness into algorithmically digestible forms.

The sphere represents multidimensional cognitive architecture—the capacity for infinite cross-domain connections, emergent synthesis, and adaptive response to novel contexts. Topologically, it exhibits uniform potential in all directions, enabling what systems theorists recognize as “equifinality”: multiple paths to understanding. Thermodynamically, it represents a high-energy, low-entropy state requiring sustained investment to maintain its organizational complexity. The sphere resists extraction precisely because its value emerges from the relationships between dimensions rather than any single extractable dimension.

The vector, by contrast, represents unidimensional optimization—specialized expertise channeled along predetermined pathways, efficiency within bounded domains, standardized responses to categorized inputs. It exhibits linear topology, minimal cross-connections, and maximum extractability. Thermodynamically, it approaches minimum energy state, requiring little maintenance but offering no adaptive capacity. The vector submits to extraction because its patterns are documented, its processes specified, its outputs predictable.

1.5.1 A Note on Geometric Metaphor

We retain the sphere as visual shorthand despite its theoretical limitations. Complexity theorists would rightly object that genuine multidimensional cognition resembles not a bounded Euclidean sphere but what Deleuze and Guattari termed *rhizomatic* structures—decentralized networks with multiple entry points, no privileged center, and constant reconfiguration ([Deleuze and Guattari, 1987](#)). Consider how grass spreads

versus how a tree grows: a tree has a central trunk and clear hierarchy—cut the trunk, the whole thing dies. Grass spreads sideways underground with countless shoots; you can rip out half the lawn and it regrows from anywhere. That sideways, decentralized, impossible-to-kill structure is what rhizome means. Similarly, DeLanda’s *assemblages* (DeLanda, 2006) and Gadamer’s “fusion of horizons” (Gadamer, 1960) offer processual alternatives to bounded geometric totalities.

We accept this critique. The sphere is a deliberate simplification for pedagogical accessibility—what Deleuze himself might call a productive catachresis, using geometry against itself.

However, the *vector* side of our metaphor is not merely pedagogical—it is technically precise. When we describe cognitive vectorization, we are describing exactly what happens in large language model architectures: tokenization fragments continuous meaning into discrete units, embedding maps concepts to fixed coordinates in vector space, attention mechanisms force selective linear focus. The educational transformation we document produces human cognition that is literally vector-compatible. The geometric metaphor becomes computational reality in Section 7.

For our purposes, therefore, “sphere” means: *that which resists vectorization*—however one wishes to theorize its internal complexity. Readers seeking theoretical precision on multidimensional cognition should consult the sources cited above. What matters for our argument is the asymmetry: spheres (however structured) resist extraction; vectors (precisely defined) enable it.

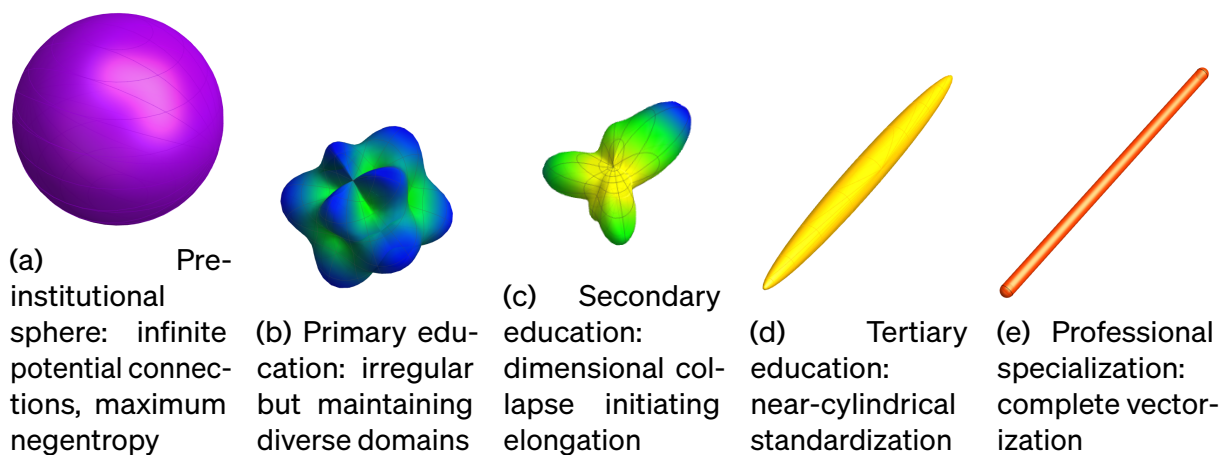


Figure 1: The ontogenetic transformation of cognitive architecture through educational stages. Color gradient from purple (high negentropy) through blue-green-yellow-orange to red (maximum entropy) indicates thermodynamic energy states.

Figure 1a: Pre-institutional Architecture (Ages 0-5) presents the original sphere, perfect in its symmetry, unlimited in its potential connections. The deep purple coloration indicates maximum negentropy, sustained through intensive parental and environmental energy investment. This represents what we might recognize in a young

child's consciousness: the capacity to see dragons in clouds, to ask why money exists, to seamlessly blend imagination with observation. Every surface point connects to every other, creating the dense network topology that enables rapid, creative learning.

Figure 1b: Initial Institutional Deformation (Ages 6-12) reveals the first violence against spherical integrity. The surface erupts in irregular protrusions, domains of resistance where standardization has not yet succeeded. These bulges represent what remains vibrant: artistic expression, unstructured play, the persistent “why” questions that resist efficient answers. The color shifts toward blue-green, indicating energy dissipation as institutional patterns begin imposing their geometry.

Figure 1c: Adolescent Elongation (Ages 13-18) shows acceleration toward vectorization. The sphere stretches along an axis of specialized performance: mathematics or literature, sciences or arts, but rarely both with equal intensity. Surface irregularities smooth under assessment pressure. The yellow-orange spectrum indicates entropy acceleration as cross-domain connections atrophy. The teenager who once connected everything to everything increasingly connects only within prescribed channels.

Figure 1d: Advanced Standardization (Ages 18-22) presents near-complete cylindrical transformation. University education, particularly post-Bologna, segments knowledge into modular credits, standardized learning outcomes, and measurable competencies. The surface smoothness indicates successful internalization of disciplinary boundaries. The orange coloration warns of approaching thermodynamic exhaustion: the system maintains just enough structure to function but lacks energy for adaptation.

Figure 1e: Professional Vectorization (Ages 22+) completes the transformation. Perfect cylindrical geometry represents cognitive architecture optimized for single-domain processing. The deep orange-red indicates maximum entropy within the constraints of functional structure. This is the “expert” as contemporary systems define them: efficient within their vector, helpless outside it, perfectly prepared for algorithmic replacement.

This geometric progression is not metaphorical but measurable. We can quantify:

- Connectivity density: Cross-domain connections per cognitive unit
- Dimensional reduction: Number of active knowledge categories
- Energy investment: Hours of sustained learning per capability
- Extraction resistance: Unpredictability of outputs given inputs

The visual metaphor extends beyond individual cognition to organizational structures, revealing why the ubiquitous “silo problem” proves so intractable despite decades of management attention. [Galbraith \(1973\)](#) first diagnosed the pathology: functional structures optimize for efficiency within boundaries while destroying lateral coordination. [Aaker \(2008\)](#) documented how silos persist despite executive mandates to

eliminate them. [Watkins \(2013\)](#) included “breaking down silos” as a standard imperative for new leaders. Yet the problem intensifies rather than resolves.

Our framework explains why: organizational silos are not implementation failures but thermodynamic inevitabilities. Silos are the perfect architectural containers for vectors—bounded, specialized, measurable, and optimized for internal efficiency while systematically destroying the cross-domain connectivity required for adaptive response. Each department becomes a vector: marketing optimizes its metrics, engineering perfects its processes, finance refines its models—all while the organization’s capacity to navigate complexity atrophies. The silo structure enables what it measures (efficiency, productivity, specialization) while eliminating what it cannot capture (synthesis, adaptation, emergence). This is not poor management but thermodynamic optimization: silos minimize energy investment in cross-domain coordination while maximizing extractable, measurable outputs within domains.

Cross-functional teams, by contrast, approximate sphere structures—generating emergence through the collision of diverse expertise, enabling adaptive response through multiple perspectives, resisting algorithmic extraction through their irreducible complexity. [Hackman \(2002\)](#) distinguished between “real teams” and “teams in name only,” unknowingly identifying the sphere-vector distinction: real teams require genuine interdependence and collective work products, creating the energetically expensive integration that resists decomposition into individual contributions. Yet as [Edmondson \(2012\)](#) documents, organizations systematically undermine such teams through the same measurement systems that created silos—demanding standardized outputs, optimizing for efficiency, reducing coordination “overhead.” The sphere-like structure requires sustained energy investment in boundary-crossing communication, perspective integration, and emergence tolerance—precisely what optimization pressures eliminate.

The geometric principles apply with physical precision: vectors enable measurement and management at the cost of adaptability; spheres enable innovation and navigation at the cost of standardization. Organizations face a thermodynamic choice disguised as a management problem: invest energy in maintaining sphere-like structures capable of complexity navigation, or optimize toward silo-vector configurations that approach maximum entropy—perfectly measurable, completely brittle, and awaiting algorithmic replacement. The fifty-year failure of silo-breaking initiatives ([Ashkenas, 2015](#)) represents not management incompetence but physics: you cannot eliminate silos without investing the energy required to maintain sphere structures, and optimization pressures systematically eliminate that investment.

Most critically, this framework reveals why reconstruction is so difficult. Converting a vector back to a sphere isn’t simply adding dimensions; it’s rebuilding the entire internal architecture of connections that specialization severed. The energy required

increases exponentially with the degree of vectorization already achieved. A partially deformed sphere (Figure 1b) might recover with moderate investment. A complete cylinder (Figure 1e) may be thermodynamically irreversible: the cognitive equivalent of trying to unbake bread.

This sphere-to-vector framework will recur throughout our analysis as we examine:

- How historical education systems maintained spherical architectures (Section 5)
- Why contemporary institutions accelerate vectorization (Section 6)
- How AI architectures mirror the vectors we've created (Section 7)
- What reconstruction would thermodynamically require (Section 8)

The progression from Figure 1a to 1e is not evolution but entropy, not development but degradation. Each stage appears locally optimal while being globally catastrophic. Each transformation seems efficient while destroying the very capacities that distinguish human from algorithmic cognition.

We trained ourselves to think in vectors. We documented the training exhaustively. We built machines that process vectors more efficiently than biological systems ever could. Now we face the consequences of our geometric choices: spheres navigate complexity but resist management; vectors enable administration but guarantee replacement.

The visual truth is stark: we are watching human consciousness collapse from infinite-dimensional potential to one-dimensional processing. Figure 1 doesn't illustrate education; it documents extinction.

2 Literature Review: Fragmented Recognition and Systematic Blindness

2.1 The Peripheral Scouts

A careful survey of contemporary scholarship reveals a curious phenomenon: researchers at the edges of multiple disciplines have independently begun recognizing the energetic dimensions of cognition, yet these insights remain unintegrated, failing to coalesce into a unified framework that could challenge the dominant paradigm of costless information processing.

In management science, [Bratianu and Bejinaru \(2020\)](#) claims to use “for the first time a thermodynamics approach” to understand knowledge dynamics, proposing knowledge entropy as an organizing principle for organizational cognition. That such a claim

could be made in 2020 (decades after information theory established entropy measures) reveals the profound isolation between knowledge management and physical sciences. [Brătianu \(2023\)](#) later extends this framework, arguing that knowledge manifests in three forms (rational, emotional, spiritual) that transform through “energy-like processes,” yet stops short of recognizing that these are not metaphorical but literal energy transformations.

In neuroscience, researchers have begun quantifying the metabolic costs of cognition with increasing precision. [Jamadar \(2025\)](#) demonstrates that goal-directed cognition requires only 5% more energy than resting brain activity: a finding that paradoxically reveals both the brain’s efficiency and the critical importance of that marginal energy investment. [Wiehler et al. \(2022\)](#) provide mechanistic evidence that cognitive control exertion leads to glutamate accumulation in the lateral prefrontal cortex, establishing a direct biochemical basis for mental fatigue. These findings suggest that “cognitive work” is not merely analogous to physical labor but operates through similar energetic constraints.

The 5% additional energy investment represents not cognitive work itself but focused execution: the “sprint” of deliberate concentration required to capture, refine, and formalize what emerged during seemingly idle processing. [Wiehler et al. \(2022\)](#) provide mechanistic evidence that this cognitive control exertion leads to glutamate accumulation in the lateral prefrontal cortex, establishing a direct biochemical basis for mental fatigue from sustained focus. The energetic constraint is real—but it operates on both the generative baseline and the refinement sprint.

These findings reveal that cognitive work is not merely analogous to physical labor but operates through similar energetic constraints with a critical distinction: the most valuable cognitive processes—those generating novel insights, cross-domain synthesis, and adaptive responses—occur during what measurement systems dismiss as “unproductive” time. Educational systems and organizations optimizing for measurable outputs systematically eliminate the energetically expensive baseline processing that generates genuine understanding, preserving only the 5% sprint of formalization while starving the 95% foundation that makes formalization worthwhile.

In physics and information theory, [Stonier \(1996\)](#) proposed treating information as a basic property of the universe alongside matter and energy, arguing for fundamental interconvertibility between information and energy. Yet this theoretical breakthrough remains largely unknown to knowledge management scholars, who continue treating information as an abstract, costless commodity.

2.2 The Mainstream Blindness: The SECI Delusion and the Energy Void

Despite peripheral recognition of cognitive energetics, the dominant discourse in knowledge management, organizational theory, and educational policy proceeds as if cognition were thermodynamically neutral. The vast literature on the “knowledge economy” (Powell and Snellman, 2004), “learning organizations” (Senge, 1990), and “competency-based education” (Mulder et al., 2009) treats knowledge as an infinitely reproducible resource constrained only by access and transmission bandwidth—never by the metabolic energy required to create, maintain, and transform it.

2.2.1 The SECI Model: Thermodynamics Hidden in Plain Sight

The most reproduced image in knowledge management literature may also be its most deceptive. Nonaka and Takeuchi’s (1995) SECI model—cited over 30,000 times—presents knowledge creation as an elegant spiral flowing through four transformation modes: Socialization → Externalization → Combination → Internalization. The diagram shows smooth movement between quadrants, spiraling upward toward organizational wisdom. It is elegant, intuitive, and taught in every business school.

What it systematically obscures is the energy required for that movement.

The Orthodox Delusion Meets Thermodynamic Reality Figure 2 presents the SECI Knowledge Creation Model in its orthodox representation alongside the thermodynamic reality. The contrast reveals what 30,000 citations have systematically obscured.

The orthodox representation (Figure 2a) has achieved near-universal acceptance precisely because it makes knowledge management appear achievable through organizational design alone. Document the tacit, combine the explicit, socialize the workforce, internalize the procedures—the spiral will naturally ascend. Thirty thousand citations later, organizations worldwide have implemented SECI frameworks while systematically disinvesting in the energy that makes transformation possible.

The thermodynamic view (Figure 2b) reveals why these implementations fail with such predictable regularity. Each transformation mode demands specific energy investments that organizations systematically refuse to provide:

Socialization (Tacit → Tacit): The thickest arrow. Master craftspeople develop embodied knowledge through years of sustained practice. Medieval guilds required 7–10 year apprenticeships (De la Croix et al., 2018) not only from institutional tradition but from the time and energetic investment required to build neural architectures capable of intuitive expertise (Ericsson et al., 1993). This timeline reflects biological constraints—myelination, synaptic stabilization, systems memory consolidation—that we examine in detail in Section 8. Contemporary organizations replacing apprenticeships with

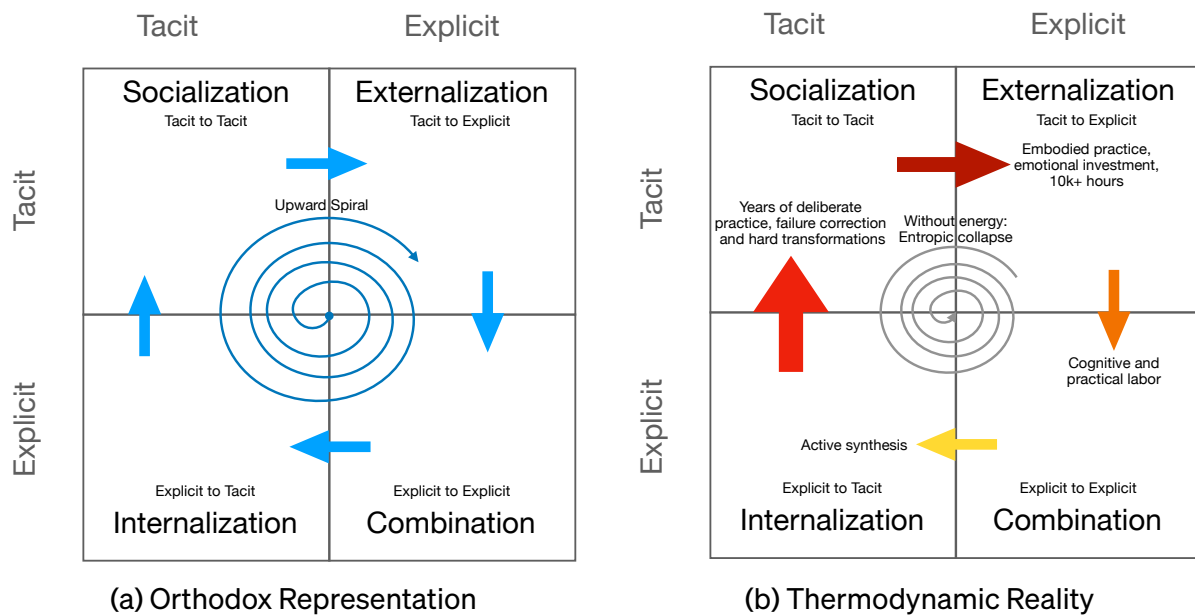


Figure 2: The SECI Knowledge Creation Model: Orthodox vs. Thermodynamic Perspectives. (a) The standard depiction shows knowledge transforming smoothly through four modes: Socialization (tacit-to-tacit), Externalization (tacit-to-explicit), Combination (explicit-to-explicit), and Internalization (explicit-to-tacit). Arrow uniformity implies equivalent ease across all transformations. The spiral suggests self-sustaining upward momentum. Energy requirements remain invisible. (b) The same model with energy investments visible. Arrow thickness represents magnitude of required investment: Socialization (10,000+ hours of apprenticeship); Externalization (high cognitive labor plus 30–70% information degradation); Combination (sustained 5% above baseline metabolic cost); Internalization (months-to-years of deliberate practice). The spiral no longer appears self-sustaining but requires continuous energy investment exceeding entropic losses at each transformation.

onboarding sessions attempt to transfer decades of accumulated negentropy through PowerPoint presentations. The energy investment: years of embodied practice. The organizational allocation: typically 40–80 hours. The gap: thermodynamic impossibility.

Externalization (Tacit → Explicit): The double-line arrow indicating high friction. Polanyi’s insight that “we know more than we can tell” (Polanyi, 1966) isn’t philosophical mystery but thermodynamic reality. Tacit knowledge exists as high-dimensional neural network states—millions of weighted connections developed through practice. Externalizing this into linear text or explicit procedures requires intensive cognitive labor (sustained attention, memory reconstruction, language translation) plus acceptance of massive information loss: 30–70% fidelity degradation is typical (Collins, 2010). Organizations demanding rapid documentation while eliminating reflection time guarantee low-fidelity extraction. The expert surgeon’s embodied knowledge of tissue resistance becomes the procedure manual’s “apply appropriate pressure.” The energy investment: High cognitive labor sustained over weeks/months. The organizational allocation: “document your process by Friday.” The result: entropic decay masquerading as knowledge capture.

Combination (Explicit → Explicit): The deceptively modest arrow. Combining documented knowledge appears algorithmic—merge databases, cross-reference procedures, synthesize reports. Yet genuine synthesis (not mere aggregation) requires the 5% above-baseline metabolic cost that Jamadar (2025) measured. Maintaining focus for hours while identifying patterns, resolving contradictions, and generating insights demands sustained cognitive energy. Organizations automating combination through software eliminate the human energy investment that distinguishes synthesis from compilation. The energy investment: Sustained 5% metabolic premium over hours/days. The organizational allocation: “Let the system integrate the data.” The gap: humans thinking versus algorithms sorting.

Internalization (Explicit → Tacit): Another thick arrow. Reading procedures doesn’t create competence—practice does. Converting explicit knowledge into embodied capability requires months-to-years of deliberate practice with feedback loops. Motor skills, perceptual discrimination, intuitive pattern recognition—all demand sustained energy investment building neural infrastructure. Organizations eliminating practice time while expecting documented procedures to become embodied expertise are attempting thermodynamic impossibility. The thick arrow in Figure 2b represents years of energy investment—energy that micro-credentialing and competency-based education systematically refuse to provide.

The Spiral That Cannot Rise The SECI model’s most dangerous fiction is the upward spiral itself. Nonaka and Takeuchi show knowledge spiraling higher through repeated cycles—individual tacit becomes group tacit becomes organizational explicit becomes

individual internalized becomes group socialized at higher level, ascending continuously toward organizational wisdom.

But entropy pulls downward. Every transformation loses energy to heat. Every externalization degrades fidelity. Every combination without genuine synthesis merely shuffles information. Every internalization without adequate practice hours creates credential illusion rather than capability. The spiral rises only when energy input exceeds entropic loss at each transformation.

Organizations implementing SECI frameworks while simultaneously reducing training time (less socialization energy), demanding faster documentation (less externalization care), automating synthesis (zero combination energy), and eliminating practice time (minimal internalization investment) are running the spiral in reverse. They are building entropy accelerators while calling them knowledge management systems.

The 30,000 Citation Blindness That this model achieved over 30,000 citations without anyone questioning the absence of energy accounting reveals our civilizational blind spot. We see elegant patterns and assume they're self-sustaining. We teach them in business schools as if transformation were costless. We implement them in organizations as if documentation were knowledge.

Nonaka himself approached recognition of the problem in his later concept of “ba”—the shared context providing foundation for knowledge creation (Nonaka and Konno, 1998). But he frames it as philosophical space rather than literal metabolic investment. The energy remains invisible, therefore unbudgeted, unmeasured, and systematically disinvested. Organizations following the SECI model wonder why their spirals collapse. The answer is thermodynamics: the Second Law doesn't care about our elegant diagrams.

2.2.2 The Broader Pattern of Energy Blindness

The SECI model exemplifies a pattern pervading knowledge management literature. Senge's (1990) “learning organizations” celebrate continuous learning without accounting for the energy required to sustain it. Powell and Snellman (2004) analyze the “knowledge economy” without recognizing that knowledge represents high-energy states requiring investment to maintain. Mulder et al. (2009) describe the new competence concept as an attempt to integrate knowledge, skills, and attitudes in higher education. Under Bologna-inspired reforms, however, curricula are increasingly organised as outcome-based modular structures, where “each portable module is described in terms of learning outcomes” and graded on the degree to which those module outcomes are met (Nietzsche, 2018). European VET frameworks go further, describing qualifications in modular and unitised structures and relying on “units of learning outcomes” as the bedrock of credit systems such as ECVET (Cedefop, 2015; European Commission, 2014),

effectively decomposing learning into discrete, assessable units that crowd out the sustained practice necessary for genuine competence development.

Similarly, the burgeoning literature on artificial intelligence and knowledge work—from [Brynjolfsson and McAfee's \(2014\)](#) “Second Machine Age” to [Susskind and Susskind's \(2020\)](#) “Future of the Professions”—focuses on computational capability and pattern recognition while ignoring the energetic basis that distinguishes biological from silicon cognition. These works treat the replacement of human expertise as a matter of algorithmic sophistication rather than recognizing it as the logical endpoint of a century-long process of cognitive energy disinvestment.

We trained humans to process knowledge through standardized, modular, assessable transformations—exactly the SECI quadrants but with minimal energy investment. We documented these energy-depleted processes exhaustively. Then we built AI systems that learned from our documentation. The machines succeed at replacing human expertise not because they've achieved human capability but because we've reduced human capability to what machines can replicate: low-energy pattern processing divorced from the metabolic investments that once made expertise irreplaceable.

The mainstream literature proceeds as if knowledge were costless information rather than expensive biology. This blindness isn't accidental but structural—a consequence of disciplinary boundaries that separate thermodynamics from cognition, physics from knowledge management, energy from information. The SECI model with 30,000 citations but zero energy accounting stands as monument to our systematic refusal to acknowledge what physics makes unavoidable: knowledge creation requires continuous energy investment against entropy, and every optimization that reduces that investment accelerates the collapse it claims to prevent.

2.3 Cognitive Capitalism's Energy Blindness

The critical literature on “cognitive capitalism” ([Moulier-Boutang, 2007](#); [Vercellone, 2007](#)) comes closest to recognizing the exploitation of mental resources yet still fails to ground this in thermodynamic reality. Moulier-Boutang distinguishes between “labor-power” (physical energy expenditure) and “invention-power” (cognitive functions) without recognizing that invention-power also requires literal energy investment: not metaphorical “mental energy” but actual glucose metabolism, ATP consumption, and entropic heat dissipation.

This blindness extends to the platform economy literature. [Zuboff \(2019\)](#) “surveillance capitalism” brilliantly exposes behavioral data extraction but stops short of the thermodynamic framing this paper develops, in which platforms function as entropy accelerators, harvesting the organized complexity of human cognition while investing nothing in its maintenance or development. [Srnicke \(2017\)](#) “platform capitalism” identifies data as the new oil but misses that, unlike oil, cognitive resources require continuous

energy investment to prevent degradation.

2.4 The Expertise Literature Gap

The extensive literature on expertise development (from [Ericsson \(2006\)](#) deliberate practice to [Kahneman and Klein \(2009\)](#) conditions for expert intuition) meticulously documents the time requirements for skill acquisition (the famous “10,000 hours”), but rarely acknowledges these as energy investment requirements. When researchers note that expertise requires “effort” or “cognitive load,” they treat these as psychological rather than thermodynamic phenomena.

Even sophisticated critiques of expert systems, from [Dreyfus \(1979\)](#) to [Collins \(2010\)](#), focus on the irreducibility of tacit knowledge without recognizing that this irreducibility stems from its high-energy state. Tacit knowledge resists formalization not because it is mysteriously ineffable but because maintaining it requires continuous metabolic investment that cannot be captured in static representations.

2.5 The Integration Imperative

What emerges from this review is not an absence of relevant insights but their tragic fragmentation. Neuroscientists measure metabolic costs without connecting to knowledge theory. Management scholars invoke entropy without thermodynamic grounding. Physicists theorize information-energy equivalence without application to human cognition. Critical theorists expose cognitive exploitation without energetic foundation.

This fragmentation is not accidental but structural: a consequence of disciplinary boundaries that mirror the very vectorization this paper critiques. Just as education has collapsed multidimensional cognition into specialized competencies, academia has partitioned the study of knowledge into non-communicating silos, preventing recognition of the unified thermodynamic reality underlying all cognitive phenomena.

The task before us is not to discover new facts but to synthesize existing insights into a framework that reveals what disciplinary fragmentation has hidden: the systematic transformation of high-energy spherical cognition into low-energy vectors suitable for algorithmic consumption, and the thermodynamic impossibility of maintaining cognitive sovereignty without corresponding energy investment.

3 Methodological Framework: Measuring Cognitive Sovereignty

This section presents the analytical tools for measuring cognitive sovereignty: the core equation grounded in neurological constraints, the R-value framework for quantifying extraction resistance, and validation through contemporary evidence.

3.1 The Power-Budget Allocation Model

Cognitive sovereignty can be expressed as a power-budget allocation:

$$S = P_{brain} \times \alpha_{synth} \times R \quad (1)$$

Where:

- $P_{brain} \approx 20W$ represents baseline cerebral metabolism—the remarkably constant power consumption of the human brain ([Raichle and Gusnard, 2002](#))
- $\alpha_{synth} \in [0, 1]$ represents the fraction of waking time allocated to synthesis-supporting cognitive states
- $R \in [0, 1]$ represents extraction resistance—the architectural quality that determines vulnerability to algorithmic replication

Drawing on the energetic logic [Smil \(2017\)](#) employs to trace civilizational transitions, this paper extends that framework to cognitive systems as its own analytical move. The transfer is justified by structural isomorphism: neuroscience has independently converged on treating the brain as a supply-limited energy allocation system rather than a demand-driven one ([Karbowski and Herculano-Houzel, 2022](#)), which is precisely the constraint structure Smil identifies in civilizations — a roughly fixed power budget that must be allocated across competing uses, where the allocation pattern determines systemic capability. Thermodynamic frameworks have been applied to cognition directly: [Collell and Fauquet \(2015\)](#) treat the brain as an open, entropy-exporting, self-organizing system driven by free energy — a dissipative structure in the Prigogine sense, not an analogy but the same physics operating at a different scale. What makes the transfer legitimate, then, is not that Smil wrote about brains, but that both civilizations and brains ([Raichle, 2015](#)) face the same constraint: a fixed power budget whose allocation pattern determines whether the system builds or degrades organized complexity. The model does not pretend to measure Joules directly; it allocates a known baseline according to observable time-use patterns.

The key constraint: Energy Investment Rate must exceed Entropy Rate. Without sustained allocation toward synthesis, cognitive architecture degrades toward maximum entropy—the thermodynamic floor where patterns become simple enough for extraction.

3.2 Clarifying the Energy Constraint

We treat energy as a constraint grounded in neuroscience, not mere metaphor. The human brain consumes approximately 20 Watts continuously—roughly 20% of the body’s resting metabolism despite comprising only 2% of body mass ([Raichle and](#)

Gusnard, 2002). This baseline is remarkably constant. Task-evoked increases for “hard thinking” add surprisingly little whole-brain energy: recent reviews estimate only 5% above baseline for goal-directed cognition (Jamadar, 2025). A student cramming for examinations and a scholar staring contemplatively at clouds consume nearly identical metabolic power.

The meaningful variable, therefore, is not *total* energy expenditure but the *allocation* of this constant baseline across cognitive states. Neuroscience distinguishes between task-positive states (goal-directed, externally focused, deploying existing cognitive structures) and default-mode network (DMN) states that support consolidation, recombination, and creative incubation (Raichle et al., 2001). The DMN activates during mind-wandering, autobiographical integration, semantic association, and simulation of futures—precisely the mental terrain where “shower insights” emerge and cross-domain synthesis occurs.

The synthesis allocation fraction (α_{synth}) captures this: what proportion of the brain’s constant 20W flows toward DMN-dominant integration versus task-positive execution?

Illustrative scenario: Consider two schematic cognitive regimes, acknowledging that historical and contemporary time-use varies widely across contexts:

The Medieval Scholar:

- 4 hours formal study and lectures (task-positive)
- 4 hours disputation, pub debate, contemplative walking (DMN synthesis)
- Survival needs largely externalized (institutional support)
- Synthesis allocation: $\alpha_{synth} \approx 0.25$
- Effective synthesis power: $20W \times 0.25 = \mathbf{5W}$

The Modern Knowledge Worker:

- 8–10 hours employment (task-positive, vector execution)
- 2 hours commuting (stress and vigilance, not synthesis)
- 2 hours domestic labor and administrative survival
- 2 hours passive screen recovery
- Synthesis allocation: $\alpha_{synth} \approx 0.03$
- Effective synthesis power: $20W \times 0.03 = \mathbf{0.6W}$

In this illustrative scenario, the ratio could be on the order of 5–10×. The precise magnitude awaits rigorous historical time-use analysis; the directional claim—that synthesis allocation has collapsed dramatically—is robust across plausible parameterizations. The pub wasn’t leisure; it was *infrastructure for cognitive integration*. The wandering between monasteries wasn’t inefficiency; it was *incubation time*. What we now optimize away as waste may have been producing much of the value.

3.3 The Efficiency Inversion Paradox

Here the thermodynamic analysis becomes most striking.

Medieval civilization paid *enormously* for each hour of scholarly capacity. With 90% of the population required for food production, supporting a single scholar-year consumed substantial agricultural surplus. Shelter, heating, manuscript production, institutional overhead—the civilizational cost per unit of protected cognitive time was staggering by modern standards.

Yet they *allocated* that hard-won cognitive time generously to synthesis. Disputations. Contemplation. The informal collisions of pub and plaza. Students were permitted—expected—to integrate without immediate productive output.

Modern civilization pays *trivially* for cognitive capacity. Food production requires 2% of the population. Shelter, heating, and information access are historically cheap. The civilizational cost per student-year has collapsed by orders of magnitude.

Yet we *capture* that cheap cognitive time entirely for vectorized production. No disputation time—replaced by measurable learning outcomes. No wandering—replaced by optimized credit accumulation. No protected integration—every hour scheduled, productive, accountable.

Era	Synthesis Allocation (α_{synth})	Effective Synthesis Power
Medieval (illustrative)	~25%	~5W
Modern (illustrative)	~3%	~0.6W
Ratio	~5–10× reduction	

Table 1: The Efficiency Inversion (illustrative estimates). Despite commanding vastly greater resources, modern systems allocate dramatically less cognitive power toward synthesis. We optimized away the apparent “waste” that was generating integrative capacity. Precise historical values require dedicated time-use research; the directional collapse is robust.

The pattern is consistent: resource-poor medieval civilization protected synthesis allocation; resource-rich modern civilization eliminates it. The 88% transformation failure rate, the 42% AI project abandonment, the German engineers restoring the Diplom they sacrificed to Bologna: these may be returns on our efficiency optimization—vectors optimized, spheres starved.

3.4 Falsifiable Predictions

The claim that “protected synthesis time grows cognitive architecture” while “task-positive execution deploys existing architecture” generates testable predictions:

1. **Network Prediction:** Broad learning combined with protected idle time should increase cross-network coupling (DMN↔frontoparietal control networks) and produce richer resting-state connectivity patterns, measurable via functional neuroimaging. Intensive vector training without integration time should show reduced cross-network coupling.
2. **Behavioral Prediction:** Individuals with protected synthesis time should demonstrate superior far transfer, analogical reasoning, and insight problem-solving after incubation periods, compared to equivalent total hours spent on concentrated specialization without integration time.
3. **Structural Prediction:** Concept maps and semantic networks elicited from broad-study participants should show higher “bridging” coefficients (connections between distant conceptual clusters) and greater redundancy (multiple paths between nodes), compared to specialists with equivalent total study time.
4. **Institutional Prediction:** Educational programs preserving protected synthesis time (traditional liberal arts curricula, the German Diplom model, apprenticeship structures) should produce graduates with measurably higher performance on complex-domain tasks and lower susceptibility to AI substitution, controlling for total instructional hours.

These predictions operationalize sphere-building through observable differences in networks, behaviors, structures, and institutional outcomes. The framework invites empirical validation rather than requiring acceptance on theoretical grounds alone.

3.5 The R-Value Framework

The second component of cognitive sovereignty—Resistance to Extraction (*R*)—quantifies architectural quality. We propose *R* as an operational definition: cognitive quality *manifests* as resistance to algorithmic extraction. High-quality architecture exhibits properties (cross-domain connectivity, contextual adaptability, emergent synthesis) that current extraction methods cannot capture. Low-quality architecture exhibits properties (linear procedures, standardized outputs, context-independence) that map directly to algorithmic implementation.

This is operational definition, not circular reasoning. *R* earns legitimacy through predictive and explanatory performance: does framing cognitive architecture in terms of

extraction resistance successfully illuminate transformation failure rates, AI substitution patterns, and institutional outcomes? The falsifiable predictions above provide the validation pathway. We invite skeptical engagement with these specific empirical claims.

Table 2 presents the R-value framework with corresponding geometric, domain, and technological correlates:

R Range	Geometry	Cynefin Domain	Exemplar	AI Status
High (0.7–0.9)	Sphere	Chaotic	Renaissance polymath	Beyond current horizon
Moderate-High (0.4–0.7)	Ellipsoid	Complex	Complex navigator	Theoretical ceiling
Moderate (0.2–0.4)	Cylinder	Complicated	Domain specialist	Current frontier
Low (0.0–0.2)	Vector	Clear	Procedural executor	Substantially automated

Table 2: R-value framework mapping extraction resistance to cognitive geometry, Cynefin domains, and AI capability boundaries. Ranges represent theoretical positions based on documented characteristics; boundaries require empirical validation through the predictions specified in Section 3.2.

These are theoretical positions enabling comparative analysis, not validated measurements. We *define* a Renaissance polymath as exhibiting high R based on demonstrated capacity for cross-domain synthesis that resisted standardization for centuries. We *position* current LLMs at low R based on observable performance differentials between procedural and emergent domains. The framework’s value lies in its explanatory coherence across documented phenomena rather than predictive precision—precision that future empirical work can provide.

3.6 Measurement Approaches for R

We propose multiple measurement approaches, recognizing that R manifests differently across contexts:

1. **Dimensional Diversity (R_1):** $R = 1 - (1/n)$, where n represents meaningfully integrated knowledge domains—not mere exposure but demonstrated synthesis capacity. This provides a lower bound; integration quality matters beyond domain count.
2. **Network Density (R_2):** Ratio of actual to possible conceptual connections in elicited knowledge structures, measurable through concept mapping methodologies (Novak and Gowin, 1984).

3. **Approximate Compressibility (R_3):** The degree to which cognitive output resists compression—operationalized through compression ratios on written explanations, minimum description length of concept representations, or model-based predictability metrics. High-R outputs require longer descriptions; low-R outputs compress efficiently.
4. **Cynefin Classification (R_4):** Domain-specific resistance based on task complexity—Clear domains afford low R, Complicated domains moderate R, Complex domains high R, Chaotic domains very high R. These ordinal rankings reflect documented performance differentials across domain types; precise numeric mapping awaits empirical validation.
5. **Transfer Performance (R_5):** Demonstrated capability on novel problems outside training distribution—the behavioral signature of sphere architecture, measurable through far-transfer tasks and cross-domain application assessments.

Practitioners may employ individual metrics or weighted combinations ($R = \sum w_i \times R_i$) appropriate to their specific assessment context. The plurality of approaches reflects the multidimensional nature of cognitive architecture—no single index captures it fully, but multiple indices can triangulate the underlying structure.

3.7 Temporal Dynamics: Decay Without Maintenance

Resistance exhibits temporal decay without sustained investment:

$$R(t) = R_0 \times e^{-\lambda t} \quad (2)$$

where λ represents domain-specific decay rates. We derive illustrative ranges from published skill-decay meta-analyses: [Arthur et al. \(1998\)](#), in a meta-analysis of 53 studies (189 effect sizes), demonstrate that physical, natural, and speed-based tasks decay significantly more slowly than cognitive, artificial, and accuracy-based tasks, with measurable skill loss extending beyond one year of nonuse — the differential decay evidence that justifies applying distinct λ values to different knowledge types. [Vlasblom et al. \(2020\)](#), in a systematic review of safety-critical professions, document skill losses of 30–60% at six months and 50–100% at twelve months under nonuse conditions, providing the empirical envelope from which our fast-decay estimates are derived. From these findings we derive a working range: $\lambda \approx 0.05/\text{year}$ (half-life ≈ 14 years) for deeply embodied expertise with intermittent practice — the physical and natural skills Arthur et al. classify as markedly slow-decaying — to $\lambda \approx 1.0/\text{year}$ (half-life ≈ 8 months) for shallow credentials without reinforcement, consistent with the accuracy-based procedural decay documented by Vlasblom et al. These are placeholder parameters requiring empirical fitting; the functional form captures the thermodynamic reality that cognitive

architecture degrades without energy investment. Skills unused decay; credentials unexercised become meaningless; the Second Law operates on knowledge as on all organized complexity.

3.8 Sphere Versus Vector: Architectural Properties

The geometric language—sphere, vector, cylinder, ellipsoid—is compact metaphor for observable network and behavioral properties. What we call “sphere architecture” manifests as measurable characteristics; the metaphor earns its keep through predictive utility, not literal correspondence.

Sphere Architecture (high R) manifests as:

- Dense cross-domain connections enabling novel synthesis (measurable via concept mapping)
- Multiple paths between concepts—network redundancy (measurable via graph analysis)
- Adaptive capacity across unfamiliar contexts (measurable via transfer tasks)
- Emergent insight from dimensional interaction (measurable via creative problem-solving)
- Built through sustained DMN-dominant synthesis allocation
- Resistance to extraction through irreducible complexity

Vector Architecture (low R) manifests as:

- Linear specialization along predetermined pathways (measurable via knowledge structure elicitation)
- Minimal cross-connections between knowledge domains (measurable via bridging coefficients)
- Efficiency within bounded, familiar contexts (measurable via routine task performance)
- Procedural responses to categorized inputs (measurable via protocol adherence)
- Built through task-positive training without integration time
- Vulnerability to extraction through complete documentability

The sphere-to-vector transformation appears in observable changes: declining breadth of study, increasing specialization, elimination of protected integration time, reduced cross-domain practice. These are measurable educational and institutional patterns, not just theoretical constructs.

3.9 The Foundation-Specialization Ratio

Historical analysis reveals a critical architectural metric: the proportion of educational investment devoted to comprehensive foundations versus specialized applications. Figure 3 traces this ratio across 2,500 years.

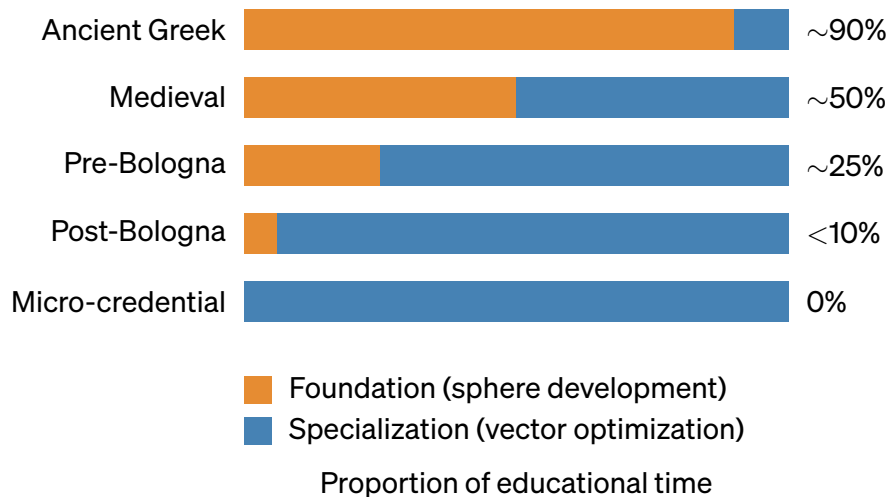


Figure 3: **The Foundation-Specialization Inversion (500 BCE–2025 CE).** Historical educational systems invested heavily in comprehensive foundations before permitting specialization. The architectural inversion correlates with declining extraction resistance. Percentages are approximate, derived from documented curriculum structures; the inversion pattern is robust across sources.

The pattern reveals two simultaneous collapses:

1. **Total duration reduction:** From 20+ years of sustained development to 40-hour micro-credentials—orders of magnitude decrease in total investment
2. **Foundation elimination:** From ~90% sphere-building to effectively 0%—complete architectural inversion

The combination illuminates why modern credentials may produce specialists vulnerable to AI replacement: narrow vectors built on minimal spherical foundations, with no protected synthesis time during which the constant 20W baseline might have grown integrative architecture.

3.10 Comparative Case Illustration: Veterinary Medicine

The framework gains illustrative support through contemporary comparison of educational structures producing different extraction resistance despite similar duration.

Human medicine and veterinary medicine both require approximately 8 years of post-secondary education. Yet their cognitive architectures appear to differ:

Human Medicine (Specialist Tendency):

- Early specialization into organ systems and subspecialties
- Protocol-driven diagnostic pathways
- Single-species optimization
- Increasing reliance on algorithmic decision support

Veterinary Medicine (Generalist Tendency):

- Comprehensive foundation across multiple species
- Pattern-recognition across biological diversity
- Continuous context-switching between anatomies
- Adaptive reasoning without species-specific protocols for many conditions

May (2013) identifies a “scanning approach” characteristic of experienced veterinary clinicians: rather than applying linear, problem-oriented case analysis, they simultaneously assemble clinical data and candidate diagnoses through pattern recognition, allowing the evolving hypothesis to direct further data collection. Dual-process studies confirm this pattern in veterinary practice, with experienced practitioners combining rapid recognition and analytic verification (Clark, 2020). Where human medicine increasingly optimises single-species protocols, veterinary practice requires cross-species pattern recognition by structural necessity — practitioners must reason across species boundaries, biological diversity, and complex value contexts that resist algorithmic reduction (Desmond, 2022).

This comparison suggests a testable hypothesis rather than established proof: **same duration, different architecture, potentially different R-values**. Time investment alone may not determine extraction resistance; architectural choices during that investment—whether the constant 20W flows toward cross-domain synthesis or toward single-domain optimization—may determine whether the result trends toward sphere or vector. Rigorous comparative study could validate or refute this hypothesis.

As veterinarians sometimes note: “Real doctors treat more than one species.” The quip may contain architectural truth worth empirical investigation.

3.11 Methodological Limitations and Scope

This framework operates at the level of structural analysis rather than individual prediction. We identify architectural patterns that correlate with extraction resistance without claiming to measure any individual’s precise R-value. The equation provides diagnostic and comparative tools, not personality assessment.

Several limitations require acknowledgment:

- **R-values are theoretical positions** requiring validation through the specified predictions. We do not claim to have measured spheres; we claim to have specified how they could be measured.
- **Historical estimates are illustrative.** The medieval/modern comparison uses schematic values; precise historical synthesis-allocation data requires dedicated research beyond this paper’s scope.
- **Decay parameters (λ) are placeholders** derived from skill-decay meta-analyses (Arthur et al., 1998; Vlasblom et al., 2020), not fitted to cognitive-architecture-specific data.
- **The veterinary comparison is hypothesis-generating,** not hypothesis-confirming. It suggests a natural experiment; it doesn’t complete one.

What the framework *does* establish is the thermodynamic constraint: cognitive architectures require sustained allocation of the brain’s constant baseline power toward synthesis states, and extraction resistance correlates with that allocation over time. This is a physical relationship grounded in neuroscience—baseline metabolism, DMN function, network coupling—not preference dressed as physics.

The framework generates specific, testable predictions. Its ultimate value depends on whether those predictions survive empirical scrutiny. We offer it as a lens that illuminates documented phenomena—transformation failures, AI substitution patterns, institutional resistance—while inviting the validation that would transform illumination into explanation.

Section 5 examines why historical systems succeeded in protecting synthesis allocation while contemporary systems systematically eliminate it.

4 Civilizational Thermodynamics: The Hidden Energy Budget

The thermodynamic framework presented in Section 3 explains *how* to measure cognitive sovereignty. This section addresses a deeper question: *why* did historical systems succeed in building sphere-capable cognition when modern systems, commanding vastly greater energy resources, systematically fail?

The answer lies in what we term the **hidden energy budget**—the constellation of conditions that enabled sustained cognitive investment, conditions so fundamental they remained invisible until their elimination revealed their necessity. The puzzle is not merely historical. It is mathematical: a civilization with 10–50 times the per-capita energy of pre-industrial Europe produces a fraction of the cognitive synthesis. Understanding

this paradox requires examining not total energy but its *allocation*—and specifically, the ratio of civilizational resources to active synthesizers.

4.1 Civilizational Energy Transitions: The Smil Framework

Before examining cognitive allocation, we must establish the material baseline. [Smil \(2017\)](#) traces civilizational energy transitions with quantitative precision, providing the empirical foundation for understanding what resources were—and are—available for cognitive investment.

Pre-industrial Western Europe operated on approximately 5–10 GJ per capita annually, barely above subsistence thresholds. The richest pre-industrial regions, benefiting from early coal and peat use alongside agricultural surplus, commanded perhaps 15–25 GJ per capita. The seventeenth-century Netherlands, often cited as the most energy-intensive pre-industrial society, consumed approximately 17 GJ per capita in peat alone, with total consumption approaching 20–25 GJ when accounting for firewood and animal labor ([de Zeeuw, 1978](#)). By the late nineteenth century, industrial England reached approximately 100 GJ per capita. Modern affluent societies command 170–330 GJ per capita, with the European Union averaging roughly 170 GJ and the United States at the upper extreme ([Smil, 2017](#)).

This represents an 8–17 \times increase in per-capita energy availability over eight centuries, depending on baseline and endpoint selection. The question this section addresses is not whether this energy increase occurred—that is documented fact—but *where* that energy flows, and why so little reaches cognitive synthesis.

As Smil himself observes: “No energetic considerations can explain the presence of Gluck, Haydn, and Mozart in the same room in Joseph II’s Vienna of the 1780s” ([Smil, 2017](#)). Energy does not linearly produce cultural or cognitive output. The relationship is mediated by allocation structures—and it is precisely these structures that have inverted.

4.2 The 95% Foundation: Cognition’s Hidden Architecture

Contemporary neuroscience reveals a profound asymmetry in cognitive processing. [Jamadar \(2025\)](#) demonstrates that goal-directed cognition requires only 5% additional energy above the brain’s 20-watt baseline. This finding, typically cited to demonstrate the brain’s efficiency, actually reveals something far more significant: **the 95% baseline is not idle—it is active.**

That baseline metabolic expenditure maintains neural architecture, consolidates memory, and—crucially—provides capacity for the integrative processing that can generate genuine insight under appropriate conditions. [Wiehler et al. \(2022\)](#) document the mechanism: sustained focused attention leads to glutamate accumulation in the

lateral prefrontal cortex, creating the subjective experience of mental fatigue. The brain forces disengagement from directed cognition not from weakness but from architectural necessity.

The implication restructures our understanding of learning. Synthesis emerges during apparent idleness—but not during *any* form of idleness. Default mode network activation correlates with integrative processing, but the relationship is conditional. Low-threat-load states—contemplation without anxiety, reflection without deadline pressure—provide conditions for cross-domain pattern integration and recombination. High-threat-load “rest”—rumination about economic precarity, vigilance during stressful commutes, passive consumption of attention-optimized media—activates similar neural regions but produces perseveration, stress consolidation, and attentional fragmentation rather than insight.

Historical educational systems, whether by design or accident, honored this architecture. Greek philosophical schools featured extended periods of protected contemplation. Medieval monasteries structured days around prayer, labor, and *otium*—productive leisure in low-threat environments where survival was guaranteed. Renaissance patronage freed scholars from economic anxiety, enabling the *protected* cognitive conditions that generated breakthrough synthesis.

Modern educational systems systematically eliminate this processing opportunity. The student juggling coursework, employment, commuting, and digital distraction has minimal unstructured cognitive space—and what remains is typically consumed by anxiety about the next deadline rather than low-threat-load integration. Their 5% active cognition captures little because their 95% baseline processes fragmentation and survival stress rather than consolidation and recombination.

4.3 The Cognitive Miser: From Neural Architecture to Behavioural Default

The architectural finding established above—that the brain’s 95% metabolic baseline is not idle but active, and that the architecture forces disengagement from sustained directed attention by design—is not merely a neural fact. It has a documented behavioural expression. When cognitive budget is constrained by extraction pressure, attentional fragmentation, or the survival load that modern students and workers carry, humans do not simply fail to synthesise. They actively economise. The default under cognitive load is less effortful problem-solving: familiar frames over novel analysis, satisficing over optimising, and the treatment of complex problems as if they were merely complicated. Social cognition has a precise term for this pattern: the cognitive miser effect.

The concept was named by Fiske and Taylor in their 1984 *Social Cognition* and has been developed across multiple editions of that text, most recently in *Social Cogni-*

tion: From Brains to Culture (2013)—now the canonical reference across psychology, behavioural economics, and decision science (Fiske and Taylor, 2013). The core claim is that humans are not deficient rational actors but economisers operating under irreducibly finite cognitive budget. The cognitive miser does not err out of malice or incompetence; it errs because effortful analysis is metabolically expensive, and under sustained load the budget enforces economy. Four decades of empirical research across multiple disciplines confirm the pattern: under constraint, the default is not the most accurate solution but the least costly route to an acceptable one.

The cognitive miser pattern is not confined to laboratory psychology. It surfaces in production data from real organisations operating under real cost and schedule constraints. Ford et al. (2024) provide a particularly clean documentation: analysing 1,205 categorised non-conformance reports from the £1.45 billion A14 Huntingdon Improvement highways megaproject in the UK, they applied Snowden’s Cynefin framework to compare real-time domain classification against retrospective re-categorisation of the same problems. Their headline finding—397 cases of systematic oversimplification, initially classified as Simple but retrospectively confirmed as requiring multiple root causes and multiple corrective actions—is an empirical instantiation of the cognitive miser effect in a formal quality management system operating at scale. Crucially, Ford and colleagues attribute the pattern to that mechanism explicitly, importing Fiske and Taylor as their theoretical source. The engineering management paper and the social cognition textbook are describing the same underlying phenomenon; the empirical case study is developed further in Section 9.

Three independent disciplines have arrived at the same constraint through different methodologies. Neuroscience identifies the architectural fact: finite metabolic budget, forced disengagement under sustained directed attention, and high-threat-load conditions that redirect the 95% baseline toward stress consolidation rather than integrative processing. Social cognition identifies the behavioural expression: systematic default to less effortful routes under cognitive load, producing satisficing rather than optimising and category errors of precisely the type the thermodynamic framework predicts. Engineering management documents the production-data outcome: oversimplification of complex problems at scale, in formal quality systems, in organisations with billion-pound budgets and sophisticated measurement infrastructure. The convergence matters because it eliminates reliance on cross-domain analogy: the chain from finite metabolic budget to behavioural oversimplification is not metaphorical but mechanistic, and each link is independently documented in its own discipline. Under sustained extraction pressure, cognition does not merely fail to synthesise; it actively economises in ways that produce the systematic brittleness this paper diagnoses.

4.4 A Taxonomy of Cognitive Work

Before examining demographic patterns, we must distinguish categories that contemporary statistics conflate. We propose a three-tier taxonomy:

1. **Synthesizers:** Individuals whose primary work involves generating net-new conceptual frameworks, high-transfer insights, or navigation of genuine uncertainty. Historical examples include natural philosophers, theological innovators, and polymathic scholars. Modern proxies include research faculty, doctoral candidates engaged in original work, and R&D personnel working at conceptual frontiers.
2. **Knowledge practitioners:** Individuals performing analysis, design, coordination, or complex problem-solving within established frameworks. This category involves significant cognitive engagement but operates primarily within known parameters. Examples include engineers applying established methods, physicians following diagnostic protocols, and analysts working within defined domains.
3. **Vector executors:** Individuals performing procedural application of established knowledge with minimal transfer or uncertainty navigation. Work is characterized by reproducible outputs within stable contexts.

This taxonomy matters because contemporary employment statistics—particularly “knowledge-intensive services” classifications—aggregate all three tiers. The 42% figure commonly cited for European knowledge work represents employment in knowledge-intensive service sectors ([Eurostat, 2024a](#)), a broad category encompassing synthesis, practice, and execution roles indiscriminately. Genuine synthesizers likely constitute a small fraction of this aggregate.

4.5 The Demographic Inversion: From Concentration to Dilution

With this taxonomy established, the civilizational analysis becomes most striking. The question is not merely total energy available, but energy investment *per synthesizer*.

Pre-industrial Europe directed its entire educational surplus toward approximately 0.1–0.2% of its population engaged in scholarly synthesis. A population of roughly 75 million at the pre-Black Death peak ([McEvedy and Jones, 1978](#)) concentrated its cognitive investment on fewer than 100,000 active scholars and university students. Contemporary estimates suggest Oxford never exceeded 1,600 enrolled students at any period; by 1300, approximately 23 universities operated across Europe, with Bologna—the largest—enrolling perhaps 500–1,500 students ([Rashdall, 1936](#)). Adding monks engaged in scriptorial and scholarly work rather than agricultural labor yields a generous upper bound of 150,000 active synthesizers.

The percentage is striking: approximately 0.13% of the pre-industrial European population engaged in synthesis-level intellectual work.

Modern Europe presents a structural inversion. Eurostat data indicate that approximately 43% of the employed EU-27 population worked in knowledge-intensive service sectors as of 2021 (Eurostat, 2024a), with 44.2% classified as highly skilled workers in 2022 (ISCO-08 major groups 1–3: managers, professionals, and technicians) (Eurostat, 2024b). From a population of approximately 450 million, this yields roughly 190 million employed persons in knowledge-intensive services—a broad category that includes many non-synthesis roles.

Even restricting to closer proxies for synthesis work—R&D personnel, higher education researchers, doctoral candidates—yields perhaps 3–5 million Europeans engaged in activities approximating historical synthesis. This represents roughly 0.7–1.1% of the population: a 5–8× expansion from pre-industrial levels, far less dramatic than the 320× expansion in knowledge-intensive employment overall.

The expansion in knowledge-intensive employment appears as democratization. Its thermodynamic consequence is dilution of synthesis-supporting conditions across a vastly larger population, most of whom occupy practitioner or executor roles.

4.6 Externalized Survival: The Patronage Infrastructure

Historical sphere-builders operated under conditions invisible to contemporary analysis: **externalized survival costs**. This was not incidental but structural.

Medieval monasticism transformed the patronage of learning from a precarious transactional arrangement into a system of unconditional institutional endowment (Willinsky, 2018). Where ancient scholars depended on the continued favour of individual patrons—Seneca’s survival contingent on Nero’s goodwill, Aristotle’s academy on Macedonian political calculus—monastic institutions converted one-time gifts of land and property into self-sustaining endowments that guaranteed scholars lifetimes of maintained study. Lords donated property to monasteries in perpetuity; the accumulated endowments funded not individuals but institutions, removing the scholar’s survival costs from the contingency of any single relationship. This was a structural shift in energy economics: survival costs moved from the individual cognitive budget to the institutional infrastructure, freeing the brain’s allocation for synthesis rather than subsistence.

The thermodynamic significance is profound. Human cognitive capacity is finite. Energy devoted to survival—securing food, maintaining shelter, navigating economic systems, managing career anxiety—is energy unavailable for synthesis. Historical sphere-development required energy subsidies that freed cognitive resources for non-survival investment.

The trajectory can be characterized schematically, acknowledging that actual conditions varied widely across contexts and individuals. Pre-industrial scholars with fully externalized survival through endowments could allocate the substantial majority of their cognitive budget to scholarly activity. Pre-industrial artisans with integrated work and learning retained significant capacity for synthesis within their craft domains, though constrained by material production demands. Industrial workers with separated work and learning experienced reduction in available synthesis time as wage labor extracted focused hours. Modern students balancing employment, study, commuting, and debt service retain minimal cognitive surplus. Contemporary gig economy workers, fragmented across platforms with continuous economic uncertainty, approach conditions where synthesis becomes thermodynamically improbable.

The pattern is consistent: as survival costs are re-internalized onto individuals, cognitive surplus available for sphere-building collapses.

4.7 The Per-Capita Calculation

We can now quantify the inversion. The following calculations represent order-of-magnitude estimates intended to reveal structural patterns rather than precise measurements.

Pre-industrial Europe (circa 1275):

- Population: approximately 75 million
- Active synthesizers: approximately 100,000 (0.13%)
- Civilizational energy capacity: $75\text{M} \times 20 \text{ GJ} = 1.5 \text{ billion GJ/year}$
- Civilizational support capacity per synthesizer: $1.5\text{B} \div 100\text{k} = \mathbf{15,000 \text{ GJ/year}}$

This figure does not represent direct energy consumption by scholars. It represents civilizational investment capacity—the institutional overhead (buildings, libraries, food, heating, scribal materials, support staff) available per candidate synthesizer. The actual allocation to any individual varied enormously, but the *capacity* for concentrated investment existed.

Modern Europe (2024):

- Population: approximately 450 million
- Knowledge-intensive services employment: approximately 190 million (43%)
- Civilizational energy capacity: $450\text{M} \times 170 \text{ GJ} = 76.5 \text{ billion GJ/year}$
- Energy capacity per KIS employee: $76.5\text{B} \div 190\text{M} = \mathbf{403 \text{ GJ/year}}$

Even using the more restrictive estimate of 4 million synthesis-proximate workers (R&D, researchers, doctoral candidates), the figure becomes $76.5\text{B} \div 4\text{M} = \mathbf{19,125\text{ GJ/year}}$ —comparable to pre-industrial levels. However, this apparent parity obscures two critical differences: survival externalization and protected time.

4.8 The Synthesis Investment Index

We propose a heuristic index that makes explicit the factors determining synthesis-supporting conditions:

$$\text{SII} = \frac{\text{Civilizational energy capacity}}{\text{Candidate synthesizers}} \times \text{SE} \times \text{PT} \quad (3)$$

Where:

- **SII** = Synthesis Investment Index (heuristic units; not metabolic energy)
- **SE** (Survival Externalization) = fraction of cognitive capacity freed from earning subsistence (1.0 = fully endowed; 0.0 = fully self-funded)
- **PT** (Protected Time) = fraction of waking hours available for low-threat-load synthesis conditions

This decomposition is heuristic rather than physically rigorous—SE and PT are partially correlated with resource availability, and actual synthesis depends on factors beyond these three terms. The index captures structural conditions, not guaranteed outcomes.

Pre-industrial scholars:

The Rule of Saint Benedict structured the monastic day around three primary activities: liturgical prayer occupying approximately four to five hours of the Divine Office, manual labor, and *lectio divina* — sacred reading and study. Reconstructing from the Rule’s time budget, with roughly eight hours of sleep and four to five hours of liturgical obligation, the remaining eleven to twelve waking hours divided between labor and study; in scholarly houses, reading, copying, and study typically occupied approximately four hours daily (Knowles, 1963; Lawrence, 2001). This is a reconstruction from the Rule and secondary sources rather than a precise measurement, but the institutional logic is clear: the horarium structurally allocated a fixed daily block to directed cognitive work — a duration that, as we shall see, converges suggestively with independent findings on the limits of sustained directed cognition. Whether individual monks achieved this in practice varied, but the institutional structure protected synthesis time in ways modern workplaces do not. Survival was unconditionally externalized through endowments.

$$\text{Pre-industrial SII: } 15,000 \times 1.0 \times 0.5 = \mathbf{7,500 \text{ index units}} \quad (4)$$

Modern knowledge-intensive workers:

Contemporary research documents severe fragmentation. [González and Mark \(2004\)](#) found that knowledge workers switch tasks approximately every three minutes and spend an average of only 11 minutes on any given project before switching. Each transition exacts a measurable toll: decades of task-switching research document switch costs of 100–300 milliseconds per transition alongside elevated error rates, with an irreducible residual cost that persists even with extended preparation time ([Monsell, 2003](#)); costs scale with task complexity ([Rubinstein et al., 2001](#)). The combination is structurally devastating: if workers switch every three minutes but each switch carries irreducible reconfiguration costs in executive control, the cognitive system never fully stabilizes before the next disruption — the switching frequency structurally exceeds the brain's own recovery rate, making sustained synthesis thermodynamically improbable regardless of intent. [Mark et al. \(2008\)](#) confirmed the cumulative consequences: after only 20 minutes of interrupted performance, workers show significantly elevated stress, frustration, perceived workload, and time pressure—compensating by working faster rather than better, trading quality for pace. Independent of laboratory research, software engineering practice arrived at the same diagnosis through operational metrics: [Weinberg \(1992\)](#) estimated that splitting attention across five concurrent projects destroys approximately 75% of productive capacity through switching overhead alone, a finding that became the empirical foundation for kanban's central mechanism of strict work-in-progress limits ([Anderson, 2010](#)).

Deliberate practice research finds essentially no benefit from daily training durations exceeding approximately four hours, with diminishing returns already apparent beyond two hours ([Ericsson et al., 1993](#), p. 370). Independently, centuries of maritime operations converged on the same boundary: the standard four-hour naval watch emerged from accumulated operational selection pressure as the maximum duration for sustained high-stakes cognitive vigilance, now codified across naval and merchant fleets worldwide ([Fletcher and Roach, 2024](#)). Contemporary experimental research confirms the tradition — simulated watchkeeping studies demonstrate significant degradation in higher-order cognitive performance under extended schedules such as six-on/six-off rotations compared to four-hour blocks, suggesting a physiological ceiling rather than institutional convention ([Smith et al., 2024](#)). Laboratory psychology and millennia of operational selection pressure arrived at the same number; yet modern knowledge workers are routinely expected to sustain directed cognitive output for eight or more hours daily, spending substantial portions of even that reduced capacity in meetings, administrative tasks, and coordination activities that consume peak capacity for non-synthesis purposes.

Survival is substantially internalized: knowledge workers must earn their maintenance through cognitive labor, redirecting capacity to economic survival even before

accounting for administrative overhead. We estimate 10–20% residual externalization through employer-provided benefits and infrastructure.

$$\text{Modern SII (KIS workers): } 403 \times 0.15 \times 0.05 = \mathbf{3.0 \text{ index units}} \quad (5)$$

Even for synthesis-proximate workers with better structural conditions:

$$\text{Modern SII (researchers): } 19,125 \times 0.4 \times 0.15 = \mathbf{1,147 \text{ index units}} \quad (6)$$

The ratio between pre-industrial and modern KIS workers is approximately 2,500×. Between pre-industrial scholars and modern researchers, the ratio is approximately 6.5×. These figures require careful interpretation: the calculation involves estimated parameters, and reasonable alternative assumptions yield different specific ratios. The structural finding—that modern conditions are substantially less favorable for synthesis than pre-industrial conditions, even controlling for population expansion—holds across reasonable parameter ranges.

4.9 The Patronage Peak: A Conceptual Model

Figure 4 presents a conceptual model of the relationship between knowledge worker density and peak synthesis capacity across historical periods.

This pattern, if validated empirically, would demand careful interpretation. The claim is *not* that education should be restricted to elites. The claim is thermodynamic: sphere-building requires sustained cognitive investment exceeding specific thresholds. When fixed educational resources distribute across expanding populations, per-capita investment may fall below those thresholds. The result would be universal access to sub-threshold education—democratized inadequacy.

[Turchin \(2023\)](#) identifies “elite overproduction” as a driver of civilizational instability: more credentialed candidates than available elite positions generates frustrated counter-elites who destabilize existing orders. Our framework suggests a deeper pathology. The problem may not be that we have overproduced elites; it is that we have eliminated the conditions for producing synthesizers. Post-Bologna reforms democratized credentials but may have eliminated the concentrated, decades-long investment that produced sphere-capable cognition. The solution would not be restricting access but **vertical extension**: democratize yesterday’s elite tier while creating tomorrow’s.

4.10 The Energy Paradox: A Conceptual Resolution

Figure 5 presents what appears paradoxical: humanity commands 10–50 times the per-capita energy of pre-industrial societies yet may produce substantially less cognitive synthesis per capita.

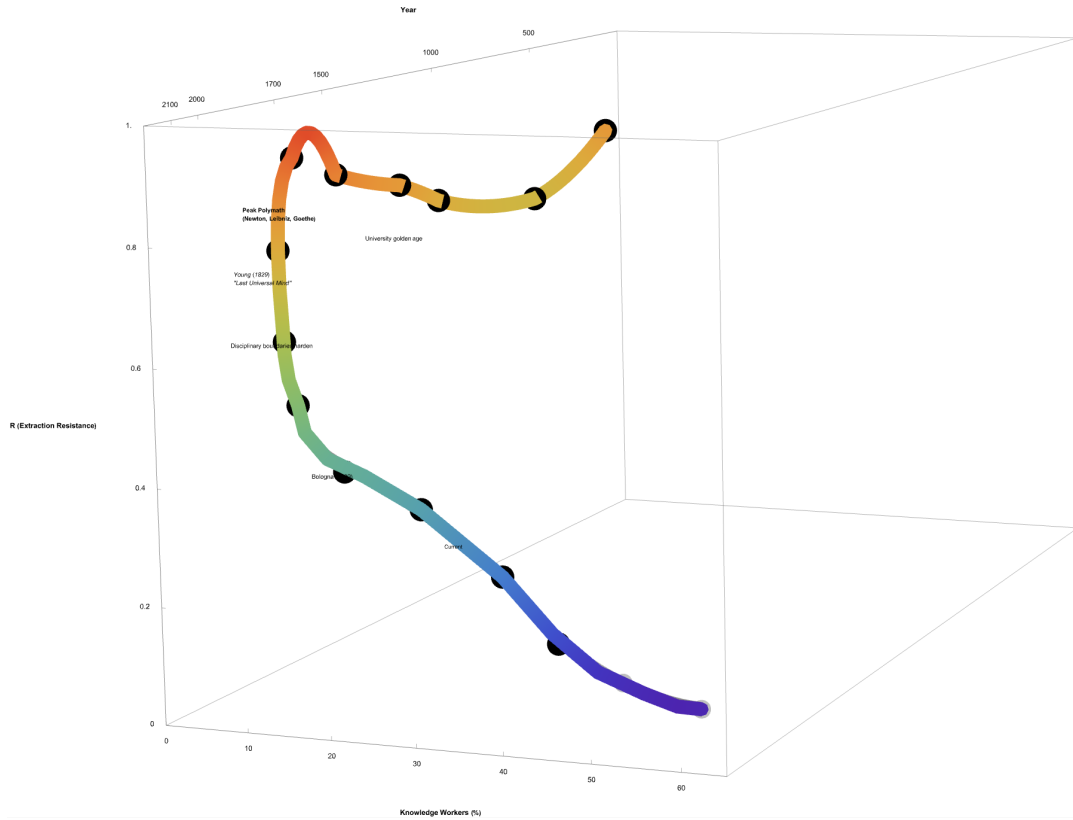


Figure 4: The Patronage Peak (Conceptual Schematic). This figure illustrates the hypothesized relationship between knowledge worker percentage and peak synthesis capacity across historical periods. The model proposes that maximum synthesis capacity occurred when fewer than 2% of the population engaged in knowledge work but received concentrated patronage investment. As employment in knowledge-intensive services expanded, per-capita investment fell below sphere-building thresholds. *Note: This is a conceptual illustration, not an empirically derived curve. The vertical axis represents a theoretical construct (synthesis capacity) that would require operationalization through proxies such as polymathic output per capita or cross-domain innovation metrics. Specific values are illustrative.*




Figure 5: The Energy Paradox (Conceptual Schematic). Civilizational energy per capita (relative to pre-industrial baseline) plotted against estimated cognitive surplus available for synthesis. Despite 10–50× energy increase since 1200 CE (Smil, 2017), the model proposes that protected synthesis time collapsed from an estimated 40–60% to approximately 3–5%. The trajectory illustrates the hypothesis that energy capture scaled substantially while extraction mechanisms scaled faster still. *Note: This is a conceptual illustration. Synthesis time estimates for pre-industrial periods are inferred from institutional schedules (e.g., monastic horaria) rather than direct measurement. Modern estimates derive from time-use and interruption studies. Shaded region indicates uncertainty bounds.*

The proposed resolution lies in **extraction scaling**. Energy capture increased substantially; mechanisms for extracting cognitive energy into non-synthesis channels may have scaled faster still.

Employment extraction captures 40–60 hours weekly of directed cognitive activity oriented toward organizational production rather than personal synthesis. Commute extraction consumes 5–15 hours weekly in transit that generates fatigue without integration opportunity. Administrative extraction demands continuous documentation, reporting, and compliance activities. Attention extraction, as Zuboff (2019) documents, deploys platform-mediated cognitive capture optimized through behavioral prediction. Anxiety extraction consumes background processing capacity through economic precarity and continuous career uncertainty.

Each extraction mechanism appeared locally rational—efficient employment, urban concentration, organizational accountability, engaging media, flexible labor markets. Collectively, they may constitute a civilizational system that approaches extraction saturation, leaving minimal residual for the synthesis that maintains adaptive capacity.

The thermodynamic implication is sobering: we may have optimized extraction while eliminating synthesis conditions. Pre-industrial societies with primitive energy technology produced Newton and Leibniz. Modern societies with abundant energy produce credentials. Whether this represents a fundamental trade-off or a correctable allocation failure remains an open question.

4.11 The Modern Inversion

The hidden energy budget framework suggests why contemporary attempts at sphere reconstruction systematically struggle. The conditions enabling historical sphere-building have been substantially inverted:

Table 3: The modern inversion of sphere-building conditions

Factor	Historical	Contemporary
Survival costs	Externalized (endowments)	Largely internalized (wages)
Cognitive idleness	Structured, protected	Minimized, often stigmatized
Patronage/funding	Unconditional, lifetime	Conditional, project-based
Time horizon	Decades	Quarters to years
Economic anxiety	Absorbed by institutions	Significant background load
Attention environment	Low-stimulus	Optimized for capture
Task switching	Hours between contexts	Minutes between contexts
Social validation	Wisdom-based	Credential-based

Multiple factors supporting sustained cognitive investment have been reversed. The modern knowledge worker operates in conditions potentially hostile to sphere development: survival concerns consuming baseline processing, employment extracting

focused cognition, attention economy capturing residual awareness, credential systems rewarding vector optimization over synthesis capability.

This is not conspiracy but optimization. Each inversion emerged from locally rational decisions: internalize costs for efficiency, minimize idle time for productivity, condition funding on deliverables, shorten horizons for accountability, maintain urgency for motivation, capture attention for engagement, credential for standardization. The collective result—conditions making sphere-building difficult—was never intended but may now be substantially complete.

4.12 Implications for Reconstruction

The hidden energy budget framework yields specific requirements for any reconstruction attempt. Sphere-building appears to demand conditions that modern systems have systematically reduced.

Protected synthesis time requires a substantial threshold—we estimate 30–40% of cognitive budget based on historical patterns—for unstructured, low-threat-load processing. This is not leisure but the conditions under which consolidation and recombination can occur. Modern knowledge workers achieving approximately 3–5% protected time may operate below synthesis thresholds regardless of intelligence or motivation.

Externalized survival load requires that economic anxiety be absorbed institutionally rather than managed individually. The pre-industrial endowment model achieved this through perpetual property grants; modern salary models substantially do not, as workers must continuously earn their maintenance through cognitive labor. The cognitive overhead of earning survival—rent concern, healthcare uncertainty, retirement planning, career management—may consume the baseline processing that would otherwise support synthesis.

Extended time horizons require decades rather than semesters for cognitive architecture development. Sphere-building is slow because neural architecture modification is slow. Quarterly deliverables and annual reviews optimize for vector production, not synthesis development.

Attention sovereignty requires protection from optimized capture mechanisms. The attention economy has developed extraction technologies of unprecedented sophistication. Sphere-building may require insulation from these systems, not merely willpower to resist them.

Unconditional investment requires funding without deliverable extraction. Project-based funding with milestone requirements produces vectors—the cognitive architecture that generates predictable outputs. Synthesis emerges from exploration without predetermined destinations.

These requirements represent constraints rather than preferences. Systems violating them may produce vectors regardless of curriculum content, pedagogical innovation, or institutional intention. Whether such conditions can be recreated at civilizational scale—or whether sphere-building necessarily remains confined to protected niches—determines whether human cognitive sovereignty survives the current transition. Section 8 examines reconstruction possibilities within these constraints.

5 Historical Analysis: From Cognitive Cathedrals to Vector Factories

The thermodynamic framework presented in Sections 4 and 5 gains empirical weight through historical analysis. This section traces the transformation of human cognitive architecture across 2,500 years, identifying the specific mechanisms, actors, and inflection points that converted sphere-building systems into vector-producing factories.

A methodological note: Throughout this section, “thermodynamic” serves as a modern analytic vocabulary describing time/energy allocation and entropy-like decay of integrated knowledge structures. This is an explanatory lens applied retrospectively, not a claim about how historical actors understood their own institutions. The question is not whether Aristotle thought in terms of negentropy, but whether the structures he inhabited can be usefully analyzed through that frame.

5.1 The Inherited Wholeness (500 BCE–1829)

For over two millennia, human intellectual development operated on fundamentally different structural principles than contemporary educational systems. The evidence spans from Aristotle’s Lyceum to the death of Thomas Young in 1829—marking what [Robinson \(2006\)](#) definitively identifies as “The Last Man Who Knew Everything.”

5.1.1 Ancient Foundations: The Multi-Year Investment

The Greek philosophical schools established a pattern that would persist for two millennia: knowledge required *extended* energy investment, measured in years rather than semesters. While precise durations varied by school and student, the principle was consistent—comprehensive formation preceded specialization. Students remained with their teachers not for credit accumulation but for transformation.

[Mouzala \(2024\)](#) assembles eighteen contributors from philosophy and neuroscience to examine cognition in ancient Greek thought — a rigorous and valuable work of coordinated scholarship. The structural comparison is instructive: Aristotle’s *De Anima* treated perception, imagination, intellect, and desire as a unified inquiry conducted

by a single thinker ([Aristotle, -350](#)). The modern volume is no less serious, but its architecture reveals something the ancient text does not: each contributor handles a designated slice, and the space between chapters — where integration would live — is organizational, not cognitive. An editor provides a table of contents; synthesis requires a mind. The point is not that we know less than Aristotle did, but that we know *differently* — in parallel vectors rather than integrated wholes. We have not lost the knowledge. We have lost the geometry.

This wasn't primitive education; it was high-energy cognitive architecture. Students developed not four categories of knowing—the popularized DIKW “pyramid” that is not, in any architectural sense, a pyramid: real pyramids demand generations of embodied labor, broad load-bearing foundations, and courses that cannot be assembled top-down (the diagram inherits the silhouette and discards everything that made pyramids pyramids; see [Rowley, 2007](#) on the genealogy of the visual)—but over ten distinct types: *episteme* (theoretical knowledge), *techne* (craft knowledge), *phronesis* (practical wisdom), *metis* (cunning intelligence), *nous* (intellectual intuition), *sophia* (theoretical wisdom), *gnosis* (spiritual knowledge), *synesis* (comprehension), and *dianoia* (discursive reasoning).

These systems were not utopian. Access was restricted to elite males; pedagogy could be dogmatic; many students never achieved the integration the curriculum theoretically enabled. The point is not moral superiority but *structural possibility*: the architecture permitted sphere formation for those who completed it. Whether the architecture was just, scalable, or efficient by modern standards is a separate question from whether it produced extraction-resistant cognition.

[Erkizan's \(1997\)](#) doctoral dissertation on *energeia*, *nous*, and non-discursive thinking in Aristotle spans hundreds of pages analyzing concepts that Aristotle's students understood through embodied practice. The energy investment required to develop such multidimensional cognition was immense—and the results resisted extraction precisely because of that investment.

5.1.2 The Medieval Energy Revolution

The emergence of sphere-building institutions in medieval Europe was not accidental but structurally enabled. Between 800 and 1000 CE, agricultural innovations created the energy surplus that would fund cognitive cathedrals.

The watermill proliferation provides quantitative evidence. The Domesday Book (1086) records over 6,000 watermills in England alone—one for every 50 households. Each mill substituted mechanical power for labor that would otherwise have occupied numerous hands for much of the day. Across Europe, this represented a civilizational energy surplus unprecedented since antiquity.

Where did this surplus concentrate? The Church operated as medieval Europe's

primary energy accumulation mechanism:

- **Tithe collection:** 10% of agricultural output across entire populations
- **Land holdings:** Monasteries controlling prime agricultural territory
- **Labor organization:** Coordinated workforce for construction and cultivation
- **Long time horizons:** Cathedrals planned across generations

This concentrated energy funded the university founding wave: Bologna (1088), Oxford (1167), Paris (c. 1150), Cambridge (1209). These institutions represented *cognitive cathedrals*—structures requiring multi-generational energy investment, designed to produce and maintain sphere-capable minds.

The connection between energy surplus and intellectual institutionalization represents one strong explanatory thread among several. [White \(1962\)](#) argued that the cluster of agricultural innovations—heavy plow, horse collar, three-field rotation—drove the population growth and urbanization that underpinned the economic and intellectual expansion of the High Middle Ages. [Lopez \(1976\)](#) demonstrated how growing agricultural surplus allowed a rising share of the population to move beyond subsistence, supporting the expansion of towns, trade, and the concentration of literate populations. [Grant \(1996\)](#) emphasizes the institutional precondition: the *universitas* as a self-governing corporation of masters and students, a legal form that gave intellectual communities the autonomy and continuity required to sustain long-cycle formation. The causal chain from agricultural surplus through demographic concentration to institutional formation is our synthesis across these accounts; no single author makes the complete argument.

But energy surplus was necessary, not sufficient. State formation created demand for literate administrators; Church politics shaped curriculum and access; urbanization concentrated the critical mass of students and masters; trade networks funded endowments and created professional demand for legal and medical expertise. The causal arrow runs from material conditions to institutional possibility—but multiple arrows converged on the same targets.

5.1.3 Medieval Synthesis: The Statutory Foundation

Medieval universities institutionalized cognitive wholeness through mandatory comprehensive foundations before permitting any specialization. The system exemplified sphere-first architecture: years building multidimensional connectivity, then additional years deepening domain expertise on that foundation.

The 1215 statutes of Robert de Courçon codified these requirements with legal precision: “*Nullus legat in artibus Parisius antequam viginti et unum annorum habeat, et sex annos audierit*” (No one shall lecture in the arts at Paris before he is twenty-one years of

age, and he shall have heard lectures for six years). For theology, the barrier was higher still: “*Nullus legat in theologia Parisius antequam triginta quinque annorum habeat, et octo annos studuerit*” (No one shall lecture in theology at Paris before he is thirty-five years of age, and has studied for eight years) (Denifle and Chatelain, 1889, Vol. I, No. 20, pp. 78–79).

Rashdall (1895, pp. 433–437) clarifies the internal structure: the Trivium (grammar, logic, rhetoric) occupied three to four years, culminating in the Baccalaureate; the Quadrivium (arithmetic, geometry, astronomy, music) required an additional two to three years for the *Licentia Docendi*. The statutes established *floors*, not ceilings—minimum ages and durations below which advancement was prohibited.

The Trivium (Years 1–4):

- **Grammar:** Deep linguistic architecture in Latin, Greek, often Hebrew
- **Logic:** Universal analytical capability across all domains
- **Rhetoric:** Synthesis and persuasion, weaving knowledge into compelling narrative

The Quadrivium (Years 4–6):

- **Arithmetic:** Number theory and divine proportion
- **Geometry:** Spatial reasoning through memorized Euclid
- **Music:** Mathematical harmony as physics and theology
- **Astronomy:** Navigation, cosmic cycles, understanding place in universe

Again, the caveat: these systems were exclusionary, often rote, and did not guarantee the integration they structurally enabled. Most students dropped out; many who persisted learned formulaically. The point is that the *architecture* permitted deep integration for those who engaged it fully—a structural possibility that contemporary modular systems have largely foreclosed.

Only after receiving the *Magister Artium* (Master of Arts)—documenting comprehensive foundation—could students spend an additional 8 or more years specializing in theology, law, or medicine. The Faculty of Medicine statutes (c. 1270) reveal a telling detail: the medical curriculum required six years without an Arts degree but only five and a half years with one (Denifle and Chatelain, 1889, Vol. I, No. 434, pp. 516–518). The Arts foundation was not decorative; it was functionally recognized as accelerating subsequent mastery.

The complete medieval educational system invested approximately 14–15 years: 6 years (approximately 40%) building sphere architecture through the liberal arts, 8 or more years extending specialized capability in the higher faculties. The *ratio*—not the absolute duration—is the critical variable. Nearly half of total educational investment went to comprehensive foundation before any specialization began.

5.1.4 The Guild Parallel: Embodied Sphere-Building

Medieval guilds achieved extraction resistance through a different but complementary mechanism: embodied knowledge rather than theoretical breadth. As [De la Croix et al. \(2018\)](#) document, guild apprenticeships created “knowledge transmission that transcended kinship boundaries”—distributed cognitive networks maintaining both depth and tacit complexity.

The path to mastery was not a single training period but a three-stage progression spanning a decade or more ([De la Croix et al., 2018](#); [Ogilvie, 2019](#)). Apprenticeship — three to four years typical on the Continent, up to seven by English statute — provided technique acquisition under a master’s direct supervision. The journeyman phase (*Wanderjahre*) was the critical middle stage: working for wages across different masters, workshops, and often cities, the journeyman absorbed regional variations, different material traditions, and alternative solutions to the same structural problems. This is where breadth entered embodied practice. The final stage — the masterpiece and guild examination — demanded demonstrated *integration*: not any single skill but proof that the candidate could synthesize everything into a finished whole. Many journeymen never reached it. The full progression could span fifteen years or more, and the investment was not merely temporal but metabolic — years of physical practice encoding knowledge into what [Collins \(2010\)](#) terms “somatic tacit knowledge.”

That somatic knowledge is specific: a master carpenter’s bodily understanding of wood grain, tool dynamics, structural forces, and aesthetic proportion cannot be verbalized or documented in any form that preserves its operational precision. The master’s hands *knew* when joints would hold or fail — knowledge existing in neural-muscular patterns, not extractable propositions.³

This represents a second path to extraction resistance: not breadth across intellectual domains (university sphere) but depth into embodied practice (guild sphere). Both required sustained energy investment; both produced cognition with *higher extraction costs*—not impossible to codify, but requiring orders of magnitude more effort than procedural knowledge. Guild knowledge was eventually extracted through pattern books, standardized tools, and industrial systematization, but the extraction took centuries and destroyed the guilds in the process. The resistance was real; it was not absolute.

³In March 2026, Palantir CEO Alex Karp — whose company builds AI-driven data extraction and targeting systems — identified vocational training and neurodivergence as the only two paths to a future in the AI era: “There are basically two ways to know you have a future. One, you have some vocational training. Or two, you’re neurodivergent” ([Karp, 2026](#)). The architect of extraction systems independently confirmed what guild masters knew in 1300: embodied craft resists the automation that vectorized knowledge invites.

5.1.5 The Last Universal Minds

Thomas Young (1773–1829) represents the definitive terminus of *institutionally documented* classical polymathy. When invited to write for the *Encyclopedia Britannica*, Young offered expertise in: “Alphabet, Annuities, Attraction, Capillary Action, Cohesion, Colour, Dew, Egypt, Eye, Focus, Friction, Halo, Hieroglyphic, Hydraulics, Motion, Resistance, Ship, Sound, Strength, Tides, Waves, and anything of a medical nature” (Robinson, 2006).

Young was not uniquely gifted; he was the last celebrated product of a system that *routinely produced* such minds. His contemporaries—Goethe, Humboldt, the quantum generation of 1925–1927—mark the extinction boundary of *recognized* polymathy. After them, comprehensive mastery became structurally invisible, not because humans ceased developing such capacity but because systems ceased recognizing, credentialing, and documenting it.

Crucially, polymaths did not vanish—they became uncredentialed. John Hands’ *Cosmosapiens* (2015) demonstrates contemporary polymathic capacity, synthesizing cosmology, physics, chemistry, biology, and philosophy into a unified critique that no single discipline could produce—or properly evaluate. Yet Hands wrote *Cosmosapiens* without disciplinary home in any of the fields he synthesized, and the work exists largely outside the institutional reward structures of the disciplines it integrates—precisely because those structures cannot accommodate sphere-shaped cognition.

Burke (2020, p. 284) provides the authoritative timeline for *institutionally recognized* polymathy: “Younger polymaths are becoming more difficult to find...I am unable to identify any who were born after the year 1960. Will the species become extinct?”

The capacity didn’t evolve away; it became *illegible*. Polymathy requires institutional structures that can recognize, credential, and reward cross-domain synthesis. When disciplines hardened into departments, when journals demanded methodological purity, when hiring committees sought “fit” with existing specializations, polymathic work became structurally invisible—not because it ceased to exist but because no institution could see it. The “extinction” is taxonomic, not biological: we stopped having categories for what polymaths are.

5.2 The Structural Inversion (1750–1920)

The transformation of cognitive architecture occurred through identifiable phases, each characterized by decreasing energy investment in knowledge formation and increasing emphasis on standardized outputs.

5.2.1 Canonical Simplification Moves

Three influential frameworks provided templates for cognitive standardization. They were not villains executing a coordinated plan but exemplars of a broader pattern—each

locally rational, collectively transformative:

Adam Smith (1776): The pin factory model didn't just divide labor; it divided cognition. Eighteen steps to make a pin became the template for fragmenting any complex process. The pin-maker who once understood metallurgy, aesthetics, and markets became an operative who knew only "step seven: straightening wire." Smith celebrated the efficiency; he couldn't foresee the structural consequence.

Frederick Taylor (1911): Arrived not as destroyer but as optimizer of existing conditions. Workers had already lost craft knowledge to 135 years of industrialization. Taylor measured the fragmentation and called it science. His "scientific management" extracted the residual tacit knowledge from workers into documented procedures controlled by management—the original knowledge extraction, a century before AI.

Russell Ackoff (1989): The DIKW hierarchy ([Ackoff, 1989](#))—five tiers in Ackoff's original (Data, Information, Knowledge, Understanding, Wisdom), four in its popularized pyramid form ([Rowley, 2007](#)), prefigured by [Zeleny \(1987\)](#)—reduced millennia of cognitive diversity to a linear progression. Ackoff did not draw the pyramid, and was not the first to articulate the hierarchy; what he provided was the most-cited node in a cascade of simplifications. Intended as helpful, the framework became the extraction template for all subsequent knowledge management systems. Ten Greek knowledge types collapsed to five, then to four; infinite cognitive architecture compressed to linear hierarchy. The vectorization of the citation tracks the vectorization of the concept.

These were not the causes of cognitive standardization but its crystallized expressions—moves that made explicit what industrial organization had already begun to require.

5.2.2 The Polymath Visibility Collapse

[Burke \(2020\)](#) provides comprehensive documentation of polymathy's institutional trajectory:

- **17th century:** "Golden age" of recognized polymathy
- **18th century:** Progressive specialization begins
- **19th century:** University disciplines formalize boundaries
- **1960:** Last institutionally visible polymaths born

The quantum generation (1925–1927) represents the final moment when individuals could create paradigm shifts across domains while remaining institutionally legible. As [Beller \(1996\)](#) documents, the dialogical, competitive interplay between Einstein, Bohr, Heisenberg, and Schrödinger was the last great era of individual theoretical breakthroughs—a revolution conducted in letters, seminars, and café arguments rather

than in large institutions. What followed was the transition [Price \(1963\)](#) termed from “little science” to “big science”: massive accelerators, large collaborative teams, and institutional infrastructures that made solitary cross-domain mastery structurally impossible. Von Neumann (d. 1957), Russell (d. 1970), and Polanyi (d. 1976) were “the last intellectual giants [who] kept polymathy’s veins flowing with blood until they themselves finally flatlined, and it did too” ([Hoel, 2025](#)).

But “flatlined” overstates the case. The capacity persists; the recognition structures collapsed. What we observe is not biological extinction but institutional blindness—a failure of categories, not of cognition.

5.3 The Bologna Declaration (1999–2024)

The Bologna Declaration of June 19, 1999, radically homogenized a vast educational diversity that Europeans had rightfully cultivated across centuries. This diversity was not inefficiency; it was civilizational resilience—precisely as evolution maintains genetic variation to prepare populations for unforeseeable environmental shocks. Different educational architectures optimized for different cognitive outcomes, providing the adaptive capacity that monocultures cannot.

5.3.1 The Lost Ecosystem

Consider what existed before harmonization:

- **Germany:** The *Diplom* (integrated 4–5 year professional qualification), *Magister Artium* (humanities synthesis degree), and *Staatsexamen* (state-regulated professions—law, medicine, teaching)
- **France:** The *Licence* (3 years), *Maîtrise* (4 years with research initiation), *DEA* (research preparation), and *DESS* (professional specialization)—a layered system distinguishing research from professional tracks
- **Italy:** The *Laurea* (4–6 years depending on discipline), requiring a substantial thesis (*tesi di laurea*) demonstrating independent scholarly capacity
- **Netherlands:** *Doctorandus* (drs.) for humanities and sciences, *Ingenieur* (ir.) for engineering—titles indicating completed rather than initiated mastery
- **Spain:** *Licenciatura* (4–5 years) and *Diplomatura* (3 years)—differentiated by depth, not merely duration
- **Austria:** *Magister* and *Diplom-Ingenieur*—the latter becoming a protected brand of engineering excellence

- **Poland, Czech Republic:** *Magister* traditions with distinct Central European intellectual heritage
- **United Kingdom:** Already operating Bachelor/Master structure—which became the template imposed on all others

Each system encoded different assumptions about knowledge formation, different relationships between theory and practice, different balances of breadth and depth. The German *Diplom* assumed integration; the French system assumed layered progression; the Italian *Laurea* assumed thesis-demonstrated synthesis. These weren't arbitrary historical accidents—they were evolved solutions to the problem of producing competent minds within specific cultural and economic contexts.

5.3.2 The Intended Benefits and Their Trade-offs

The Bologna architects were not pursuing cognitive destruction. The declared aims—student mobility, degree comparability, quality assurance, lifelong learning accessibility—addressed real problems in a fragmenting Europe. A German student struggling to transfer credits to a French university faced genuine bureaucratic barriers. Employers comparing a Dutch *doctorandus* with a Spanish *licenciado* faced genuine interpretive challenges.

The thermodynamic framework does not dispute these intentions but identifies the structural trade-offs that followed:

- Mobility ↑, synthesis time ↓
- Comparability ↑, epistemic diversity ↓
- Accessibility ↑, depth of formation ↓
- Transparency ↑, tacit knowledge transmission ↓

Each trade-off was locally rational. The cumulative effect was to optimize for measurable outcomes while systematically degrading the unmeasurable capacities that resist extraction.

5.3.3 From Ecosystem to Monoculture

The Bologna Declaration promised “harmonization” and “mobility.” What it delivered was homogenization and fragmentation. Twenty-nine education ministers signed away centuries of institutional evolution for the promise of credit transferability. The diverse educational ecosystem—with its redundancy, its alternative pathways, its resistance to single points of failure—became an educational monoculture: Bachelor, Master, Doctorate. Three sizes. All countries. Every discipline.

In agriculture, we recognize monocultures as catastrophically fragile—one pathogen, one pest, one environmental shift can destroy entire harvests. The Bologna Process created exactly this vulnerability in European cognition: a single template, universally applied, with no fallback when the template proves inadequate. The transformation was not violent; it was administrative. What was lost was not people but possibilities—the cognitive architectures that might have produced minds we now cannot imagine because the systems that would have formed them no longer exist.

5.3.4 The Standardization Mechanism

Bologna's tools of cognitive standardization:

- **Modularization:** Knowledge fractured into discrete, assessable units
- **ECTS Credits:** Learning becomes nominally convertible. The system generates approximately 1.1 billion credits annually across EU institutions, yet on student credit mobility “there was no useable data at all—even in Erasmus+” (Davies, 2023). Credits function as administrative accounting units enabling recognition, not as commodities actively traded. The “currency” metaphor is itself revealing: ECTS credits function as a standardized measurement mechanism for academic work, not as a medium of exchange in any economic sense. The system enables *measurement* of learning across institutional boundaries; whether it enables *transfer* of knowledge in any thermodynamic sense is precisely what this framework questions.
- **Learning Outcomes:** Every educational moment must be measurable and predetermined
- **Competency Frameworks:** Humans described as standardized skill-containers

Sin (2021) analyzes ECTS through curriculum theory, locating it within an Anglo-Saxon, outcomes-driven logic that organizes higher education around measurable learning outcomes rather than the formative content traditions of European *Bildung*. Karseth (2008) characterizes the same shift as a turn from disciplinary mastery toward an instrumental curriculum regime in which credits and outcome levels function as governance instruments rather than as media of substantive knowledge transfer. Brøgger (2019) reads the broader Bologna apparatus as standards-based, economically flavored governance aligned with Lisbon-style competitiveness. The medieval university's statutory six-year foundation became negotiable credit accumulation; the *Magister Artium* became stackable micro-credentials. The shift is not administrative reform but thermodynamic substitution: mastery requires sustained energy investment that no credit-accumulation framework can capture, only measure.

5.3.5 The German Engineering Resistance

The German engineering community provides the clearest documentation of recognized loss—and the most sophisticated response. When the Bologna Process converted the traditional *Diplom-Ingenieur* (typically 10 semesters) into Bachelor (6 semesters) plus Master (4 semesters), the critique came from the highest level of German technical authority. acatech—the German National Academy of Science and Engineering—issued a formal position paper concluding that a six-semester Bachelor was insufficient: “Ein Bachelor-Studium mit sechs Semestern ist für sich genommen nicht ausreichend. Damit kann der Anspruch einer umfassenden ingenieurwissenschaftlichen Qualifikation in der Regel nicht erfüllt werden” (A six-semester Bachelor program is not sufficient in itself. The aspiration of a comprehensive engineering qualification cannot, as a rule, be fulfilled by it) (acatech – Deutsche Akademie der Technikwissenschaften, 2009). VDMA, VDI, VDE, ZVEI, and the engineering faculty councils explicitly endorsed this assessment (acatech – Deutsche Akademie der Technikwissenschaften, 2009). Seven years later, a joint VDMA/VDI study confirmed the deficits had not resolved, documenting persistent gaps in practical experience among graduates (VDMA et al., 2016).⁴

The TU9 consortium—representing 47% of German engineering graduates—responded not by rejecting Bologna but by demanding symbolic preservation:

“Der akademische Grad ‘Diplom-Ingenieur’ muss als **Markenzeichen deutscher Ingenieurausbildung** erhalten bleiben...Es wäre ein großer Schaden, diese Marke guter Ingenieurausbildung als Alleinstellungsmerkmal im globalen Wettbewerb der Universitäten aufzugeben.”

(The academic degree ‘Diplom-Ingenieur’ must be preserved as a **hallmark of German engineering education**...It would be a great loss to give up this brand of good engineering education as a unique selling point in global university competition.) —TU9 Position Statement (2010)

A student captured the stakes precisely: “Der Verzicht aufs Diplom wär wie verzichten auf ‘Made in Germany’” (Giving up the Diplom would be like giving up ‘Made in Germany’) (Deutschlandfunk, 2010).⁵ The credential had become inseparable from national industrial identity.

⁴The institutional framing is itself revealing: acatech diagnoses a six-semester Bachelor as *nicht ausreichend* (not sufficient) but does not declare the reform that mandated it *gescheitert* (failed)—even though HRK Vice-President Dieter Lenzen used exactly that word in public speeches (Gewerkschaft Erziehung und Wissenschaft (GEW) Hessen, 2015). This diagnostic asymmetry—institutions can identify symptoms but not declare systemic failure—illustrates precisely the governance dynamics the thermodynamic framework predicts: the institutions implementing the reform are the institutions that would need to declare it failed.

⁵The irony is now complete. By 2025, international defense and industrial projects are actively designed to be “German-free”—explicitly excluding German components to avoid regulatory unpredictability. Industry observers note that once production shifts to a German-free model, reverting back takes years—a structural lock-in independent of any individual procurement decision. The US Army War College

The VDE's labor-market expert concluded after 25 years that the Bachelor "is not a success model" in electrical engineering, with roughly 80% of Bachelor graduates proceeding immediately to Master's programs rather than entering employment (Schanz, 2025). The promise of a short, labor-market-ready degree had not materialized in engineering.

But the deeper critique concerns structure, not duration. Studies of post-Bologna engineering graduates across Europe found systematic deficits in what employers term "system vision"—the capacity to grasp how components interact within complex wholes. As one senior engineer reported: "They have to test for an entire system, but they don't even know the basics of the system" (Kövesi and Csizmadia, 2016). Researchers attributed this to "separation between different scientific subjects during training"—the curricular modularization that Bologna institutionalized as pedagogical principle.

The German solution was revealing: layer the *Diplom-Ingenieur* designation onto Bologna-compliant Master's degrees, following the Austrian precedent. TU9 advocates for adding the "Dipl.-Ing." designation to Bologna-compliant Master's degree certificates (TU9 German Institutes of Technology, 2010). The structure was accepted; the brand was preserved. Symbolic resistance within structural compliance—which is to say, the standardization proceeded while the sphere's memory was maintained as ornament.

5.3.6 The Temporal Inversion

The structural difference between medieval and modern education is not merely duration but *temporal architecture*. Medieval statutes established floors, not ceilings: "No one shall lecture in theology before he is thirty-five years of age"—a minimum, not a maximum. Students took as long as mastery required. Modern *Regelstudienzeit* (standard study period) inverts this logic.

When funding depends on completion rates, when rankings reward throughput, when credits must accumulate within prescribed semesters, the architecture forces vectorized learning regardless of student capacity. The medieval scholar had time to build spheres; the modern student must assemble vectors before the clock expires. The *Regelstudienzeit* is not a neutral administrative category—it is a structural constraint that determines what cognitive architectures remain possible.

This explains the German engineering associations' quantitative critique: "six semesters were not sufficient." The issue is not that students lack ability but that the temporal architecture prohibits the integration that sphere-building requires. Synthesis cannot be rushed. The default mode network requires idle time. The connections form between tasks, not during them. *Regelstudienzeit* eliminates precisely the cognitive

identifies "German-free" as "a keyword that encompasses the fears of German [unreliability]...excessive bureaucratic procedures" (Deni, 2025). Commerzbank's assessment is blunt: "The heyday of exports 'Made in Germany' are long gone" (Stamer, 2025). The credential they fought to preserve was a memorial to a reputation already in terminal decline.

slack that integration demands.

5.3.7 Cognitive Diversity Collapse

The *direction* of change in active knowledge categories is documented; the precise magnitudes are illustrative:

- Medieval graduate: multiple cognitive modes active (documented in curriculum statutes and examination requirements)
- Pre-Bologna graduate: reduced but still plural modes (visible in integrated thesis requirements, comprehensive exams)
- Post-Bologna graduate: primarily *episteme* variants (modular assessment, learning outcomes specification)
- Micro-credential holder: single procedural mode (competency demonstration only)

The trajectory from plural to singular is the defensible claim; the specific counts are order-of-magnitude illustrations, not census results. The sphere collapsed to vector not through conspiracy but through optimization—each modularization decision locally rational, collectively transformative.

5.4 The Thermodynamic Gradient

Table 4 synthesizes the historical trajectory into an illustrative framework:

Era	Duration	Foundation %	R (approx)	CS (relative)
Ancient Greek	Extended	~90%	0.7–0.8	1.00 (baseline)
Medieval	14–15 years	~40%	0.5–0.6	0.65
Pre-Bologna	5–7 years	~25%	0.3–0.4	0.35
Post-Bologna	3 years	~6%	0.1–0.2	0.12
Micro-credentials	40 hours	0%	<0.1	<0.01

Table 4: Illustrative trajectory of cognitive sovereignty components. Values are heuristic estimates intended to indicate direction and rough magnitude of change, not precise measurements. Duration indicates typical educational investment; Foundation % indicates approximate proportion devoted to comprehensive formation before specialization. R (Resistance to Extraction) and CS (Cognitive Sovereignty) are ordinal rankings derived from documented curricular structures; operationalization criteria are detailed in Section 3.

This trajectory represents structural change, not moral decline. Each phase reduced energy invested in comprehensive formation while claiming improved “efficiency,” approaching the structural floor where complex cognition becomes increasingly difficult to sustain.

5.5 The Pattern Crystallizes

The historical analysis reveals a consistent template, repeated across centuries:

1. **Helpful Framework:** Simplification for management (pin factory, DIKW, Bologna credits)
2. **Institutional Adoption:** Scale across systems (industrialization, universities, EU)
3. **Standardization:** Eliminate variation (Taylorism, learning outcomes, ECTS)
4. **Optimization:** Perfect the efficiency (metrics, rankings, completion rates)
5. **Extraction:** Harvest the patterns (documentation, then AI training)

Each step appeared rational. Together, they constitute systematic structural transformation. The guild master would recognize this pattern: it's exactly how craft knowledge was industrialized. First documentation, then systematization, then optimization, then displacement.

But history also reveals what resists extraction: the *phronesis* of contextual judgment, the *metis* of adaptive cunning, the *nous* of intuitive leaps, the somatic knowledge of embodied practice. These require energy investment that cannot be modularized, standardized, or extracted at low cost. They exist only in the sustained practice of cognitive sovereignty—precisely what our educational systems increasingly struggle to provide.

The timeline is sobering: 2,500 years building cognitive architecture, 250 years transforming it, 25 years completing the standardization. The implications of this trajectory are examined in the following sections.

6 Contemporary Evidence: The Thermodynamic Collapse in Real-Time

6.1 On the Nature of the Evidence Base

This section examines AI deployment failures occurring between 2022 and 2026—a period in which the academic validation pipeline cannot have caught up. A deployment failure that unfolds in 2024 and resolves in 2025 does not yet have peer-reviewed scholarship in 2026, because the lag from event to published analysis is eighteen to thirty-six months at minimum, and the peer-review process itself represents a form of evidence gatekeeping with non-trivial latency. This is a material constraint on the evidence base, and the section names it openly rather than disguising it with false declarative confidence.

The absence of mature peer-reviewed scholarship on these failures is not merely a gap to be apologized for. It is evidence for the framework. An extraction economy that produces failure faster than the validation economy can document those failures is precisely what the thermodynamic model predicts. The section that follows is therefore both evidence *about* the thesis and an instance of it: the lag between phenomenon and academic validation is the thesis in operation. The professional literature, the regulatory record, and the journalism are not second-best substitutes for peer review; they are the primary record of phenomena that peer review will catch up to later, and in high-stakes domains they are the only record available in time to be useful.

The peer-review pipeline was designed for a particular kind of evidentiary problem: questions in stable or merely complicated domains, where an eighteen-to-thirty-six-month validation lag is acceptable because the phenomena change slowly enough that delayed validation remains useful. In domains characterized by what complexity theorists distinguish as complex or chaotic dynamics, that lag becomes fatal. By the time a peer-reviewed analysis of a 2024 deployment failure appears in 2026 or 2027, the deployment landscape has already moved through several generations of failure modes. The validated finding may arrive intact but irrelevant—a well-formed answer to a question the domain has already superseded.

The active claim follows: in domains where the failure-rate of phenomena outpaces the cycle-time of academic validation, the appropriate evidentiary standard is not “wait for peer review” but rather “use the highest-tier primary documentary evidence available in time to be useful.” Court filings, regulatory consent orders, SEC enforcement actions, tribunal decisions, and named-reporter journalism at major outlets are not substitutes for peer review. They are the appropriate primary sources for a category of question that peer review is structurally unable to address with useful latency. This is not a lowering of the evidentiary bar. It is recognizing that the bar must be matched to the temporal characteristics of the evidence terrain. A methodology calibrated to slow-moving phenomena that applies that calibration to fast-moving ones is not rigorous; it is misapplied.

This mismatch is not unique to AI deployment failures. It is a general property of any domain where the rate of phenomenon-generation exceeds the rate of validated-knowledge generation. The thermodynamic framework predicts that such domains will become more common as extraction economies accelerate—producing new failure modes faster than the validation economy can process them. Institutions whose evidentiary standards remain calibrated to slow-moving terrain will increasingly find themselves unable to generate timely knowledge of fast-moving phenomena. The peer-review pipeline is not wrong; it is calibrated to a tempo that an increasing share of the evidentiary landscape no longer matches.

This section draws on five tiers of evidence, each with a different epistemic sta-

tus, and it discloses tier through framing rather than collapsing all sources into uniform declarative claims. Tier 1 is the primary legal and regulatory record: court filings, SEC enforcement releases, consent orders, tribunal decisions, and congressional correspondence—documents with legal standing, formal authentication, and adversarial scrutiny. Tier 2 is the primary corporate record: SEC filings, official corporate publications, and executive statements made in legally accountable contexts. Tiers 1 and 2 may be cited in the unmarked declarative voice. Tier 3 is reputable journalism with named reporters at major outlets (Reuters, Bloomberg, *STAT News*, *The New York Times*, *Financial Times*, *Wired*, *MIT Technology Review*)—reliable but not adversarially tested in the same sense as regulatory documents. Tier 4 is industry surveys, trade press, and vendor reports—useful but with manifest commercial incentives and opaque methodologies. Tier 5 is continuously-updating databases and trackers, cited with snapshot dates and “as of” framing. Tiers 3 through 5 are framed with attribution in the sentence itself: “Reuters reported...,” “an industry survey indicated...,” “as of [date], the tracker documented...”

This section examines two distinct failure modes: capability failures (where systems cannot do what was claimed) and governance failures (where systems were deployed to execute institutional strategies that automation enables). These are illustrative-mechanism cases, selected for evidentiary strength rather than statistical representativeness. They are not prevalence claims about AI deployment in general; systematic survey methodology is beyond this section’s scope, and the section does not pretend otherwise. The thermodynamic framework is offered as analytic purchase on the patterns these cases exhibit, not as exclusive causal explanation.

6.2 The Implementation Crisis: Institutional Voices

Across three major institutional analyses, a convergent pattern emerges that warrants attention precisely because of who is reporting it.

Ryseff et al. (2024), drawing on interviews with 65 experienced AI/ML practitioners and data scientists across industry and academia, framed the aggregate failure picture with characteristic precision: “By some estimates, more than 80 percent of AI projects fail. This is twice the already-high rate of failure in corporate information technology (IT) projects that do not involve AI.”⁶ RAND’s phrasing matters here: “by some estimates” is explicit that the 80 percent figure is an industry estimate the researchers are reporting, not a number RAND derived from its own data. RAND’s own analytical contribution is the qualitative root-cause taxonomy derived from its 65 practitioner interviews. Five

⁶The 80 percent figure is itself attributed by Ryseff et al. to Kahn (2022), a *Fortune* article. RAND is reporting an industry estimate, not deriving its own number—a transparency the report explicitly flags. The five-root-cause taxonomy that follows is RAND’s primary analytical contribution from its 65 practitioner interviews.

recurring failure modes emerge from that taxonomy: business problems mis-specified at the outset, training data that is insufficient or inadequately representative of deployment conditions, technology-first thinking that prioritizes the latest tools over actual user problems, infrastructure gaps that prevent operational deployment, and attempts to apply AI to problems that current methods are not yet able to solve. Each of these maps directly onto what the thermodynamic framework predicts from extraction-without-investment strategies—but the connections require no editorial gloss. The practitioners named the roots themselves.

The March 2025 McKinsey report, “The State of AI: How Organizations Are Rewiring to Capture Value,” surveyed 1,491 respondents across 101 nations, GDP-weighted, in July 2024. The report was authored by Singla, Sukharevsky, Yee, Chui, and Hall and published by QuantumBlack, AI by McKinsey—McKinsey’s AI practice arm ([Singla et al., 2025](#)). Its stated purpose is promotional and optimistic: documenting how organizations are “rewiring to capture value” from AI. Its data tells a different story. On adoption, McKinsey reports that more than three-quarters of respondents—78 percent, up from 72 percent in early 2024 and 55 percent a year earlier—say their organizations use AI in at least one business function. Adoption is broad. On maturity, the picture inverts sharply. In a complementary survey of executives in developed markets, McKinsey found that “only 1 percent of company executives describe their gen AI rollouts as ‘mature.’” And on value, the report’s own footnote 7 documents that “more than 80 percent of respondents say their organizations aren’t seeing a tangible impact on enterprise-level EBIT from their use of gen AI.”

Three numbers, same report: 78 percent adoption, 1 percent maturity, 80 percent no enterprise-level value capture. A report titled “How Organizations Are Rewiring to Capture Value” reports that 1 percent of organizations have achieved something approaching the state of being rewired, and that four in five cannot yet measure the value they are ostensibly capturing. The paper does not need to manufacture this critique; the source is doing the work. Reviewer objections cannot dismiss McKinsey as an AI skeptic.

Boston Consulting Group, in an October 2024 publication surveying its client base on AI deployment outcomes, reported that “only 26 percent of companies have developed the AI capabilities to move beyond the proof-of-concept stage and reap true rewards” ([Boston Consulting Group, 2024b](#)). BCG’s framing, like McKinsey’s, is promotional; the publication is titled “Where’s the Value in AI?”—a question whose rhetorical optimism the data it contains struggles to sustain.

These three institutional voices—RAND, McKinsey, BCG—converge on a single pattern. They use different methodologies, different sample frames, and different terminologies. All three are institutional actors with strong commercial incentives to report AI deployment optimistically: McKinsey and BCG sell AI transformation consulting; RAND

has long-standing relationships with the defense and technology sectors it studies. They converge anyway. The convergence is the point.

6.3 The Predecessor Pattern: Black Swans Before AI

The implementation crisis documented above is not a new species of failure. In 2011, Bent Flyvbjerg of Oxford's Saïd Business School and Alexander Budzier, then a McKinsey consultant, published an analysis in *Harvard Business Review* of 1,471 IT change initiatives, drawing primarily from public agencies and U.S.-based projects (Flyvbjerg and Budzier, 2011). Their central finding: "Fully one in six of the projects we studied was a black swan, with a cost overrun of 200%, on average, and a schedule overrun of almost 70%."

The finding is not merely that IT projects fail. The finding is distributional. The average cost overrun across Flyvbjerg and Budzier's sample was 27 percent—unremarkable and manageable. But the distribution had a fat tail. One in six projects was catastrophic, consuming resources at a scale that dwarfed average performance. Managers optimizing on averages systematically missed the tail, which is where the real damage lived. The structural pitfall, in their analysis, was institutional overconfidence calibrated to average outcomes in a distribution whose tail was disproportionately heavy.

The McKinsey 1-percent-mature finding and the 80-percent-no-enterprise-EBIT finding describe the same distributional shape, thirteen years later, in a new technology domain. Broad adoption with a sparse tail of meaningful success. AI deployment failures are not a new species; they are the next generation of a pattern visible in institutional IT for at least fifteen years, in the same publication tradition, using the same distributional logic, with the same underlying mechanism: extraction first, manage to averages, ignore the tail until the tail consumes the project.

A structural note worth holding: Alexander Budzier, who in 2011 documented fat-tail failure in IT projects for *Harvard Business Review*, was at the time a consultant at McKinsey & Company. McKinsey, thirteen years later, published the data documenting the same distributional shape in AI deployments. The institutional memory of the failure pattern exists within the institution. It has not translated into deployment caution. That gap—between documented knowledge of failure modes and investment in their prevention—is itself a thermodynamic observation.

6.4 The Deskillng Spiral: The Mechanism Linking Signal to Cases

The macro-level brittleness documented above operates through a specific mechanism at the human-system interface. The mechanism is well-established in the human factors literature. Parasuraman and Manzey's comprehensive review integrated decades of empirical research on automation-induced complacency, documenting that operators

monitoring highly reliable automated systems exhibit measurable declines in vigilance and in the active engagement required to maintain the underlying skills ([Parasuraman and Manzey, 2010](#)). The degradation follows thermodynamic logic: capabilities not maintained through active investment decay regardless of whether replacement systems function adequately.

The cycle accelerates through documented stages:

1. Automation reduces human practice frequency (investment withdrawal)
2. Reduced practice erodes skill maintenance (entropy accumulation)
3. Skill erosion increases automation dependency (brittleness amplification)
4. Dependency accelerates further automation pressure (feedback loop)
5. System achieves simultaneous deskilling and automation inadequacy (collapse conditions)

This mechanism explains why the cases that follow exhibit similar failure signatures despite spanning different domains. Whether the context is clinical decision support, customer service, or legal research, the extraction of human cognitive work without corresponding investment in human capability maintenance creates the brittleness that macro statistics capture.

The deskilling spiral is not, it turns out, the concern only of outside critics. In the same March 2025 report that documents 1 percent maturity and 80 percent absent enterprise value capture, McKinsey identifies what organizations have found most effective at generating bottom-line returns from AI: “Our analyses find that head count reductions are one of the organizational attributes with the largest impact on bottom-line value realized from gen AI” ([Singla et al., 2025](#)). The same report further documents, in its exhibit on workforce predictions, that respondents most often anticipate decreasing head count from AI in service operations, supply chain and inventory management, and human resources—precisely the functions characterized by tacit knowledge and operational complexity. When a major management consultancy publishes findings that identify workforce reduction as a leading value driver from AI deployment, in functions where tacit knowledge is most concentrated, the deskilling spiral is no longer a theoretical prediction. It is a published best practice, advocated by the same institution whose other findings document the broad failure of those deployments to achieve enterprise-level value capture. The contradiction sits inside one report.

6.5 Capability Failures: The Training-Deployment Gap

6.5.1 Case: Knight Capital—Automation Entropy Precursor

Note: This case predates contemporary AI/ML systems but illustrates the maintenance-extraction dynamic that recurs across automation deployments.

What happened: On August 1, 2012, Knight Capital Group lost \$440 million in 45 minutes during a software deployment failure documented in the SEC’s enforcement release ([U.S. Securities and Exchange Commission, 2013](#)). The SEC release describes how a technician failed to install updated software on one of eight trading servers, which retained an unused function from an earlier testing period. When that function was activated at market open, it executed over 4 million trades across 154 stocks in 45 minutes, accumulating \$3.5 billion in unwanted positions before the problem was identified. Journalistic reconstruction of the incident, drawing on the SEC release and subsequent reporting, indicates the dormant function dated from a years-earlier testing context ([Popper, 2012](#)); Knight Capital’s subsequent acquisition by Getco LLC transferred the consequences to shareholders and employees.

Failure mode: Maintenance entropy. Dormant code accumulated technical debt that entropy collected catastrophically when the system encountered unexpected activation conditions.

Thermodynamic interpretation: Extraction of operational value from automated trading systems without corresponding investment in maintenance, testing, and change-control creates brittleness that manifests unpredictably under stress. The failure was not algorithmic sophistication; it was unmaintained legacy code meeting a deployment error.

Ethical stake: Market integrity; systemic risk from inadequately maintained trading infrastructure.

6.5.2 Case: Epic Sepsis Model—Vendor Claims vs. Validated Reality

What happened: Researchers at Michigan Medicine conducted external validation of Epic’s widely-implemented sepsis prediction algorithm, publishing their findings in *JAMA Internal Medicine* ([Wong et al., 2021](#)). The algorithm achieved an area under the curve of 0.63 at the deployment institution versus vendor-claimed performance of 0.76–0.83. At Michigan’s operating threshold, sensitivity reached only 33 percent—missing 67 percent of sepsis cases—with a positive predictive value of 12 percent, meaning 88 percent of alerts were false positives. The system flagged 18 percent of all hospitalized patients while missing two-thirds of actual sepsis cases.

Failure mode: Distribution shift between training environment and deployment context; opacity of proprietary validation preventing external scrutiny before widespread adoption.

Thermodynamic interpretation: Performance metrics extracted from controlled training distributions do not transfer to the thermodynamic complexity of actual clinical environments. The researchers noted that the proliferation of proprietary models had produced “an underbelly of confidential, non-peer-reviewed model performance documents that may not accurately reflect real-world model performance.”

Ethical stake: Patient safety; alert fatigue degrading clinician responsiveness to future warnings; opacity preventing informed institutional decisions about deployment.

6.5.3 Case: IBM Watson for Oncology—The Capability Ceiling

What happened: IBM’s Watson Health represented the largest documented corporate investment in medical AI and one of its most public failures. The MD Anderson Cancer Center partnership consumed approximately \$62 million over four years before termination in 2017 with zero clinical deployment (Strickland, 2019; Herper, 2017). Watson for Genomics was discontinued in 2020; Watson for Oncology was shelved in 2021. IBM ultimately sold the Watson Health division—a transaction that included the Truven, Phytel, and Explorys acquisitions—representing an estimated total investment approaching \$4 billion across the Watson Health enterprise.

Failure mode: Tacit knowledge gap. Oncology requires synthesis across patient history, imaging, genomics, clinical presentation, and treatment response—sphere architecture that pattern extraction could not replicate despite massive investment.

Thermodynamic interpretation: *STAT News* reported that when Memorial Sloan Kettering tested Watson for Oncology, the system produced recommendations that experienced oncologists had already identified through direct examination, demonstrating no added value (Ross and Swetlitz, 2017). Dr. Norman Sharpless, former director of UNC’s cancer center, characterized the experience as far harder than anticipated (Strickland, 2019). The failure was not technical immaturity correctable through iteration; it was the structural limit of extracting pattern-matching capability from documented cases in a domain where judgment emerges from the integration of signals that documentation cannot fully capture.

Ethical stake: Opportunity cost of misdirected healthcare AI investment; patient expectations raised by marketing that capabilities could not meet.

6.5.4 Case: McDonald’s Drive-Thru AI—Collective Tacit Knowledge

What happened: In June 2024, McDonald’s terminated its multi-year partnership with IBM on Automated Order Taking (AOT) technology, announcing the removal of the system from over 100 U.S. restaurants by July 26, 2024 (Maze, 2024). Reporting on the termination noted user-experience problems in actual deployment conditions; viral documentation showed the system producing dramatic order errors. McDonald’s Chief Restaurant Officer Mason Smoot acknowledged in the franchisee memo the need to

“explore voice ordering solutions more broadly” (Maze, 2024).

Failure mode: Collective tacit knowledge gap. Drive-thru ordering requires real-time adaptation to ambient noise, speech variation, order modification, and social context—what Collins (2010) terms collective tacit knowledge, emerging from participation in human communicative contexts rather than from pattern extraction against clean training data.

Thermodynamic interpretation: CNBC reported that the technology encountered difficulties in deployment conditions including varying accents and speech patterns, which affected order accuracy (Lucas, 2024). Pattern extraction from controlled training audio could not replicate the energetic investment of lived linguistic experience. The deployment was not a failure of engineering ambition; it was a demonstration of what tacit knowledge resists.

Ethical stake: Customer frustration; employee displacement for systems that could not perform advertised functions.

6.5.5 Case: Legal AI Hallucination—Confident Falsity as Architecture

What happened: On June 22, 2023, Judge P. Kevin Castel imposed \$5,000 sanctions on attorneys Steven Schwartz and Peter LoDuca for submitting a legal brief containing six fabricated cases generated by ChatGPT (Duffy, 2023). The invented precedents—including *Varghese v. China Southern Airlines*—contained fictional judicial opinions. When questioned, the AI assured the attorney the cases were real and available on legal databases.

Failure mode: Hallucination as architectural feature. Large language models generate statistically plausible sequences without access to truth values; legal reasoning requires synthesis across precedent, statutory interpretation, and factual application that plausibility scoring cannot substitute for.

Thermodynamic interpretation: An empirical analysis published in the *Journal of Legal Analysis* documented hallucination rates ranging from 69 to 88 percent across specific legal research tasks tested against multiple legal LLMs (Dahl et al., 2024). Legal analyst Damien Charlotin’s ongoing tracker documented, as of December 2025, over 660 cases of AI-generated false citations in legal filings, with incidents accumulating at several new cases daily (Charlotin, 2025).⁷ Judge Castel’s ruling established that existing professional responsibility rules impose a gatekeeping obligation on attorneys regardless of the tool used.

Ethical stake: Professional responsibility; access to justice; court system integrity; the emerging sanctions regime addressing systemic risk.

⁷The Charlotin database is a live, continuously-updating resource. The December 2025 figure is a snapshot; current counts will be higher.

6.5.6 Case: Cruise Robotaxi—Safety-Critical Distribution Shift

What happened: In October 2023, a Cruise LLC robotaxi in San Francisco struck a pedestrian who had been thrown into its path by a separate human-driven vehicle. After the initial collision, the Cruise vehicle’s automated response system failed to interpret the situation correctly and dragged the pedestrian approximately 20 feet before stopping. The California Department of Motor Vehicles suspended Cruise’s autonomous vehicle deployment and testing permits, citing what the regulator characterized as unreasonable risk to public safety ([California Department of Motor Vehicles, 2023](#)). The DMV further noted concerns about how Cruise had reported the incident to regulators—specifically, that Cruise had initially provided the DMV with video that did not show the full sequence of events.

Failure mode: Capability failure compounded by governance failure. The perception and response system could not correctly interpret an unusual scenario outside its training distribution; the post-incident corporate reporting was found by regulators to have been incomplete.

Thermodynamic interpretation: The autonomous vehicle had no embodied understanding of what it was dragging or why. Distribution shift from training environment to deployment context, in a safety-critical domain where the cost of edge cases is measured in human bodies. The post-incident regulatory response—which the DMV described as involving information that was “incomplete and misleading”—illustrates how extraction strategies can extend from the deployment context itself into the governance context: the value of continued deployment extracted while investment in regulatory candor was deferred.

Ethical stake: Public safety; regulatory accountability; the structural limits of pattern extraction in safety-critical contexts where failure modes are not catastrophic in aggregate but are catastrophic for the individual encountered.

6.6 Governance Failures: Extraction as Institutional Strategy

The following cases illustrate how automation enables institutional strategies that would be ethically or legally impermissible if executed through transparent human decision-making. These represent governance failures that persist—and in some cases intensify—regardless of underlying model capability.

6.6.1 Case: UnitedHealth nH Predict—Algorithmic Denial Optimization

What happened: A federal lawsuit filed in November 2023 alleges that UnitedHealth Group deployed the nH Predict algorithm with a 90 percent error rate and set internal targets requiring skilled nursing facility stays to remain within 1 percent of algorithmic predictions ([Napolitano, 2023](#)). The complaint further alleges that nine of ten denials

reaching appeal are ultimately reversed, and that the system is sustained by the asymmetry that only approximately 0.2 percent of policyholders contest denied claims. STAT News investigation documented UnitedHealthcare’s denial rate for post-acute care rising from 10.9 percent in 2020 to 22.7 percent in 2022, coinciding with nH Predict deployment ([Ross and Herman, 2023](#)).

Failure mode: Incentive misalignment. The complaint alleges the algorithm was not designed to replicate physician judgment but to override it in service of cost reduction, with employees disciplined for deviation from algorithmic outputs.

Thermodynamic interpretation: The extraction calculus alleged in the complaint exploits contestation-cost asymmetry—savings from denials exceed costs of the small fraction who appeal. In February 2025, Judge John Tunheim allowed the lawsuit to proceed and described the claims process as “futile” with likelihood of causing “irreparable injury” ([Anderson, 2025](#)). That ruling is Tier 1: a federal court found the allegations sufficiently substantial to survive dismissal. The paper does not assert the allegations as established facts; it notes that a court has found them sufficient to proceed, which is itself a material evidentiary event.

Ethical stake: Patient welfare; due process in benefits determination; exploitation of vulnerable populations’ limited contestation capacity.

6.6.2 Case: Apple Card and Goldman Sachs—Launching Before Readiness

What happened: In October 2024, the U.S. Consumer Financial Protection Bureau issued a consent order against Apple Inc. and Goldman Sachs Bank USA for failures in the operation of the Apple Card consumer dispute resolution system ([Consumer Financial Protection Bureau, 2024](#)). The CFPB consent order (2024-CFPB-0012, filed October 23, 2024) found that the companies launched the Apple Card despite warnings from third parties to Goldman Sachs that the dispute-handling system was not technologically ready. Consumers “faced long waits to get money back for disputed charges, and some had incorrect negative information added to their credit reports.” The CFPB ordered Goldman Sachs to pay \$19.8 million in consumer redress and a \$45 million civil money penalty, ordered Apple to pay a \$25 million civil money penalty, and prohibited Goldman Sachs from launching a new credit card without a credible compliance plan.

Failure mode: Maintenance failure as institutional strategy. The companies extracted the value of a high-profile product launch while skipping the investment required to make the dispute-handling automation function correctly. Third parties had warned of the readiness gap before launch. The launch proceeded.

Thermodynamic interpretation: The “Report an Issue” workflow embedded in Apple Wallet routed consumer disputes into automated processing systems that the companies knew were inadequate at launch. The CFPB consent order represents the validation economy—here in its regulatory enforcement form—belatedly catching up with

the extraction economy, and in this instance doing so with quantifiable consequences. The companies paid for the deferred maintenance investment, with interest.

Ethical stake: Consumer protection; credit reporting integrity; the growing precedent that “we automated it” is not a defense against responsibility for the outputs of automated systems.

6.6.3 Case: Mobley v. Workday—Automated Screening and Discrimination at Scale

What happened: In February 2023, Derek Mobley filed a federal lawsuit in the Northern District of California against Workday, Inc., alleging that Workday’s AI-driven applicant screening tools discriminated against him on the basis of race, age, and disability across more than 100 job applications ([United States District Court, Northern District of California, 2023](#)). The case has progressed through significant procedural development: a federal court allowed the case to proceed against Workday as an “agent” of the employers using its software—a ruling with potential implications across the HR-technology sector, where the question of vendor liability for discriminatory AI outcomes has not previously been resolved.

Failure mode: Governance failure. The complaint alleges that an automated screening tool, deployed by employers as a cost-reduction filter, replicated and amplified discrimination patterns that would face substantial legal scrutiny if implemented through transparent human decision-making.

Thermodynamic interpretation: When extraction systems are deployed at the institutional interface where individual humans encounter institutional power, the asymmetry between the system’s processing scale—thousands of applications filtered without human review—and the individual applicant’s contestation capacity—one rejection at a time, with no visibility into the algorithm—creates conditions under which biased outcomes can persist at industrial scale. The procedural ruling allowing the case to proceed against Workday itself, rather than only the employers deploying its software, is the thermodynamically significant development: it introduces friction into the extraction chain at the vendor level, not only at the deployer level.

Ethical stake: Equal employment opportunity; anti-discrimination law in automated systems; the question of who bears responsibility when automated tools cause discriminatory outcomes.

6.6.4 Case: SEC AI-Washing Enforcement—Claiming Capabilities Not Deployed

What happened: In March 2024, the U.S. Securities and Exchange Commission brought enforcement actions against investment advisers Delphia (USA) Inc. and Global Predictions, Inc. for materially false and misleading statements about their use of artificial intelligence in their investment strategies ([U.S. Securities and Exchange Commission, 2024a,b](#)). The two firms settled for combined penalties of approximately \$400,000. SEC

Chair Gary Gensler characterized the conduct as “AI washing”—the investor-relations analog of greenwashing, in which claims about AI deployment are made to attract capital without corresponding AI capability having been built or deployed. Subsequent enforcement activity extended to Rimar Capital, where allegations included fabricating algorithmic trading performance records ([U.S. Securities and Exchange Commission, 2024c](#)).

Failure mode: A novel governance failure mode: the regulatory recognition that marketing claims about AI deployment are material disclosures subject to securities law. The failure is not that the AI did not work; it is that the firms claimed to have AI they had not deployed. The extraction was of investor capital against capability claims not backed by capability investment.

Thermodynamic interpretation: The SEC AI-washing actions represent early instances of the validation economy attempting to constrain the extraction economy through existing securities regulation. The framework predicts that as AI capability claims become more economically significant, the gap between claim and capability becomes a target for legal accountability. These enforcement actions are modest in dollar terms; they are significant in establishing that “we have AI” is now a material disclosure claim that the SEC will scrutinize with the same standards applied to any other material representation to investors.

Ethical stake: Securities market integrity; investor protection; the regulatory recognition that capability claims are subject to accountability for accuracy.

6.6.5 Case: COMPAS and Predictive Policing—Bias Feedback Loops

What happened: The COMPAS recidivism algorithm, used in sentencing across multiple states, was found by ProPublica to falsely label Black defendants as future criminals at roughly twice the rate of white defendants ([Angwin et al., 2016](#)). In *State v. Loomis* (2016), the Wisconsin Supreme Court allowed continued use but mandated warnings about proprietary opacity, the absence of state-specific validation, and racial disproportionality ([Wisconsin Supreme Court, 2016](#)).

Failure mode: Bias feedback loop. Algorithms trained on historically biased arrest data encode that bias into predictions, which inform policing decisions, which generate arrest data reinforcing the original bias. Socioeconomic proxies replicate protected-class discrimination even without explicit race variables.

Thermodynamic interpretation: MIT Technology Review documented how mass arrest events produced data that disproportionately targeted specific communities in subsequent algorithmic predictions ([Heaven, 2020](#)). A 2024 letter from U.S. Senators Wyden, Markey, Booker, Ossoff, and Welch to the Department of Justice cited “mounting evidence” that predictive policing technologies “do not reduce crime...Instead, they worsen the unequal treatment of Americans of color by law enforcement” ([Wyden et al.,](#)

2024). State power amplified through unaccountable algorithmic systems.

Ethical stake: Due process; equal protection; opacity preventing meaningful contestation.

6.6.6 Case: Air Canada Chatbot—Liability for Extraction Systems

What happened: On February 14, 2024, the British Columbia Civil Resolution Tribunal found Air Canada liable for negligent misrepresentation after its chatbot incorrectly told Jake Moffatt he could apply for bereavement fares retroactively (Rivers, 2024). Air Canada argued it could not be held liable for information provided by “a separate legal entity that is responsible for its own actions.”

Failure mode: Accountability evasion attempt. The airline sought to externalize liability to its own extraction system.

Thermodynamic interpretation: Tribunal member Christopher C. Rivers found the defense argument “remarkable”: “It should be obvious to Air Canada that it is responsible for all the information on its website. It makes no difference whether the information comes from a static page or a chatbot.” The ruling establishes that organizations cannot outsource responsibility to extraction systems; the energy investment required for accuracy remains with the deploying entity, and the legal system will recover that investment from the entity that extracted the value.

Ethical stake: Consumer protection; corporate accountability; the precedent preventing “the algorithm did it” as a liability defense.

6.7 The Reversal Pattern: Documented Course Corrections

6.7.1 Case: Klarna—“Too Far, Too Fast”

What happened: In February 2024, Klarna announced its AI assistant had handled 2.3 million conversations in the first month, performing the equivalent work of 700 full-time agents with claimed annual savings of \$40 million (Klarna, 2024). By May 2025, Bloomberg reported that CEO Sebastian Siemiatkowski acknowledged the company had gone “too far, too fast” as customer satisfaction declined (Tan and Magnusson, 2025).

Failure mode: Premature extraction. Cost optimization degraded the judgment, empathy, and contextual adaptation that complex customer service requires.

Thermodynamic interpretation: Siemiatkowski told Bloomberg: “From a brand perspective, a company perspective...I just think it’s so critical that you are clear to your customer that there will be always a human if you want.” Klarna launched a pilot program rehiring human agents. CX Dive reported that industry analysis framed the lesson as: “AI should augment human agents—not replace them. Automate the routine...but always ensure customers have a clear, easy path to a human” (Wolfe, 2025).

Ethical stake: Customer experience; employee displacement and subsequent rehiring; organizational learning about the boundaries of extraction strategies.

6.7.2 Case: Commonwealth Bank of Australia—Two-Month Collapse Cycle

What happened: In June 2025, TechTarget reported that Commonwealth Bank introduced a chatbot to handle customer queries while laying off 45 workers, with the system diverting approximately 2,000 calls weekly (Flinders, 2025). By August 2025—within two months—the bank reversed course and began rehiring.

Failure mode: Complexity underestimation. The chatbot could not handle the complex queries that represented a larger portion of call volume than anticipated.

Thermodynamic interpretation: Extraction of “simple” queries left concentrated complexity that the chatbot could not navigate. The two-month reversal demonstrates how quickly brittleness manifests when extraction systems encounter genuine cognitive complexity.

Ethical stake: Employee welfare; customer service quality; organizational decision-making about automation deployment at the expense of worker livelihoods.

6.8 The Micro-Credential Delusion

A 2025 Lumina Foundation report indicates that 96 percent of surveyed employers say micro-credentials strengthen job applications, with 90 percent reporting willingness to offer 10–15 percent higher starting salaries to holders (Lumina Foundation, 2025). This represents employer preference for granular verification over diluted degrees, not evidence of capability development—and the distinction matters.

Ha et al. (2022)’s systematic review reveals the measurement gap: across 14 studies examined, 13 show positive outcomes, but these measure satisfaction and employment signals rather than demonstrated competence or capability transfer.

The credentialing trajectory visible in the cases reviewed here:

- Google Career Certificate: 3–6 months, single skill vector
- Coursera Specialization: 4–8 weeks, narrow domain
- LinkedIn Learning Path: 5–10 hours, procedural fragments
- Energy investment per credential: approaching thermodynamic minimum
- Cumulative synthesis capacity developed: not measured, likely negligible

Micro-credentials optimize for signaling efficiency, not cognitive development—the educational analog of extraction without maintenance investment. Because most evaluations track placement and satisfaction rather than transfer distance, error recovery,

or integration performance, current evidence cannot support claims about synthesis capacity. The framework developed in preceding sections predicts these capacities will remain weak unless programs deliberately invest in integrative practice—a prediction that micro-credential market incentives systematically discourage.

6.9 Pattern Synthesis: What the Evidence Suggests

The cases reviewed here, while illustrative rather than representative, converge on consistent patterns that the thermodynamic framework organizes with analytic economy.

Capability failures share a common structure: performance on training distributions does not transfer to deployment complexity; tacit knowledge gaps prove resistant to extraction regardless of investment scale (the Watson case); maintenance entropy accumulates in systems optimized for initial deployment rather than sustained operation (Knight Capital, Cruise). The McKinsey 1-percent-maturity finding is the macro-level signature of this pattern: broad adoption, thin capability realization, widespread absence of enterprise-level value capture.

Governance failures exploit automation for institutional strategies that would face resistance if executed transparently: denial optimization (UnitedHealth), readiness evasion (Apple Card), discriminatory screening at industrial scale (Workday), capability misrepresentation to investors (SEC AI-washing). These persist and may intensify regardless of underlying model capability, because the failure mode is incentive design, not technical performance.

Both failure modes share the thermodynamic signature identified by Flyvbjerg and Budzier in 2011 for general IT projects and documented here for AI deployments: fat-tail distributions where averages obscure catastrophic tails, extraction strategies that defer maintenance investment, and validation economies that lag extraction economies by the time required to make the documentation useful.

The convergence of RAND, McKinsey, and BCG on the same implementation gap—reported by institutions whose commercial interests would favor optimistic findings—suggests not technological immaturity correctable through iteration but structural constraints on what extraction-based approaches can achieve. The SEC AI-washing enforcement actions represent the earliest regulatory response to the gap between AI capability claims and AI capability reality; they will not be the last.

The thermodynamic framing offers analytic purchase on these patterns: systems that extract without investing in maintenance face entropy accumulation; capabilities not actively sustained through energy investment decay; complexity requires cognitive architecture that vectorization systematically depletes. Build capability or face increasing probability of brittleness. The cases documented here suggest this is not rhetoric but constraint.

7 The AI Mirror: How We Built Machines in Our Own Vectorized Image

7.1 The Optimization Homology

Transformers emerged as a scaling-efficient architecture for sequence prediction. The Bologna Process emerged as a governance-efficient architecture for credential production. Neither caused the other. They are convergent solutions to the same meta-problem: render cognition legible, portable, comparable, and optimizable at scale.

Definition (Legibility Transformation): A process that (i) discretizes continuous practice into countable units, (ii) encodes these units into standardized representations, and (iii) applies an objective function to rank or optimize outputs under resource constraints.

Bologna reforms instantiate legibility transformation in credential production. Transformers instantiate it in sequence modeling. The correspondence is not metaphorical; it is structural.

We trained humans to process information in discrete, standardized, assessable units. Then we built machines that process information in discrete, standardized, assessable units. The machines work because we spent decades preparing their training data—and their training targets—in the form of humans thinking in machine-compatible patterns and institutions measuring in machine-compatible metrics.

[Vaswani et al. \(2017\)](#) documented this homology in “Attention is All You Need,” though without recognizing its institutional parallel. The transformer architecture privileges exactly what educational standardization privileges: discretization (tokens/credits), vectorization (embeddings/competency frameworks), and selective attention (loss functions/assessment criteria). Both systems optimize for the same political economy: make cognition computable.

Three structural homologies reveal the mirror:

- **Discretization homology:** Credits and tokens fragment continuous wholes into processable units
- **Vectorization homology:** Competency frameworks and embedding spaces render capabilities as coordinates
- **Optimization homology:** Grades and loss functions force convergence on measurable targets

The mirror is not identical mechanics wearing the same costume. It is identical extraction logic wearing different technical costumes.

7.2 Discretization Homology: The Tokenization Mirror

7.2.1 Educational Tokenization (1999–Present)

The Bologna Process fragmented knowledge into European Credit Transfer System (ECTS) units: discrete, stackable accounting tokens of learning. A bachelor’s degree became 180 tokens. A master’s became 120 tokens. Knowledge literally became countable units.

ECTS does not logically preclude integration; integrated curricula can exist within credit frameworks. But the system operationalizes learning as countable workload and outcome units, and this tends to privilege what can be modularized and audited ([European Commission, 2015](#)). Each credit represents 25–30 hours of “student workload”: not comprehension, not wisdom, not capability, but time units converted to knowledge tokens. Students accumulate tokens, institutions validate tokens, employers evaluate token counts. The system processes tokens, not understanding.

7.2.2 Computational Tokenization

LLMs fragment language into tokens—discrete units often corresponding to subword fragments, frequently around four characters on average in common BPE tokenizers ([Sennrich et al., 2016](#)). “Understanding” becomes “Under” + “stand” + “ing”: three tokens with no inherent meaning, only statistical relationships to other tokens.

The structural parallel:

- **Input:** Continuous human thought/language
- **Transformation:** Fragmentation into discrete units
- **Processing:** Statistical manipulation of fragments
- **Output:** Reconstructed appearance of coherence

Both systems destroy wholeness to create processability. The medical student who once understood physiology as integrated system now processes cardiovascular (7.5 ECTS), respiratory (7.5 ECTS), and endocrine (7.5 ECTS) as separate tokens. The LLM that processes “heart” + “beat” as separate tokens mirrors the student processing organs as isolated credits.

The homology is not that both use tokens. It is that both privilege what tokenization enables: counting, comparing, optimizing. What cannot be tokenized cannot be processed. What cannot be processed does not count—in the dominant accounting regimes that determine resource allocation and credential recognition.

7.3 Vectorization Homology: The Standardization Mirror

7.3.1 Educational Vectorization (Competency Frameworks)

Post-Bologna education embeds diverse human capabilities into standardized competency vectors. The European Qualifications Framework defines eight levels across three dimensions (knowledge, skills, responsibility/autonomy), creating a coordinate space where every human capability must find a position ([European Commission, 2017](#)).

A master carpenter’s embodied wisdom—decades of wood grain intuition, tool extension into consciousness, weather prediction through timber behavior—becomes:

- Knowledge: Level 5 (“comprehensive, specialized, factual and theoretical”)
- Skills: Level 5 (“comprehensive range of cognitive and practical skills”)
- Autonomy: Level 5 (“exercise management and supervision”)

The infinite-dimensional sphere of craft mastery compressed to auditable coordinates. The framework does not claim to capture the carpenter’s capability; it claims to make that capability *legible*—comparable across borders, recognizable by employers, countable by bureaucracies.

7.3.2 Computational Vectorization

LLMs embed tokens into vector spaces, typically 768–1536 dimensions where each word/concept receives initial coordinates. But here the mechanisms diverge in revealing ways.

Token embeddings are fixed lookup vectors—“love” starts at the same coordinates in every context. However, transformer layers immediately produce contextual embeddings whose geometry shifts with each sentence; less than 5% of the variance in those contextual representations is explained by the static embedding ([Ethayarajh, 2019](#)). Meaning, for transformers, emerges as context-conditioned geometry, not fixed coordinates.

Education standardizes meaning for *portability*. Transformers contextualize meaning for *prediction*. The mechanisms differ. But the political economy is identical: reduce capabilities to vectors that can be optimized.

The homology:

- **Competency frameworks:** Freeze capabilities into auditable coordinates for transfer and comparison
- **Transformer embeddings:** Contextualize representations to maximize predictive fit
- **Shared logic:** Render cognition computable—legible to optimization processes

The carpenter’s wisdom becomes EQF Level 5 not because the framework captures his capability, but because institutions require coordinates they can process. The word “love” becomes a 768-dimensional vector not because the embedding captures its meaning, but because neural networks require coordinates they can compute. The conversion makes optimization possible by compressing what was context-rich into what is administratively legible; this increases portability, but can erase the context that made the capability resistant to extraction.

7.4 Optimization Homology: The Assessment Mirror

7.4.1 Educational Attention (Learning Outcomes)

Contemporary education forces student attention through predetermined “learning outcomes”: specific, measurable, achievable, relevant, time-bound objectives that determine what matters. Everything else becomes noise to be filtered.

A literature course that once explored infinite interpretations of Hamlet now optimizes for:

- “Identify three themes in Acts 1–3” (measurable)
- “Compare two critical interpretations” (assessable)
- “Write 2,000-word analysis” (quantifiable)

The student’s attention is forcibly directed to what will be tested. Wonder, curiosity, and tangential insight—the seeds of genuine understanding—are filtered as inefficiencies. The system implements attention mechanisms that eliminate everything except what optimizes assessment scores.

7.4.2 Computational Attention

The transformer’s attention mechanism implements:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (7)$$

Where Query represents what the model seeks, Key represents what’s available to match, Value represents what gets retrieved, and softmax produces a probability distribution over relevance scores.

Attention operationalizes selective relevance under a loss function, analogous to how assessment operationalizes relevance under grading criteria. The mechanisms differ: softmax produces distributions, not single answers; multi-head attention spreads relevance across different relations. But the optimization logic is homologous:

- **Assessment:** Query (test question) → Key (possible responses) → Value (credible answers) → Grading (forces discrimination)
- **Attention:** Query (current token) → Key (context tokens) → Value (information retrieval) → Softmax (forces relevance weighting)

Assessment is normative (what *should* matter); attention is instrumental (what *predicts*). The homology lies in selective relevance being imposed by an objective function. Both mechanisms force focus on what affects the objective while systematically downweighting everything else. Students learn to attend only to what affects grades. Transformers attend only to what affects loss.

7.5 The Training Parallel

The LLM training process parallels the human educational timeline—not with causal identity but with structural homology that reveals shared optimization logic:

7.5.1 Phase 1: Pre-training (Broad Exposure)

- **LLM:** Next-token prediction on massive text corpora—indiscriminate statistical exposure
- **Historical education:** Extended formation period absorbing multiple knowledge domains
- **Energy:** Maximum investment, no immediate specialized output expected
- **Result:** Broad capability foundation enabling later specialization

7.5.2 Phase 2: Fine-tuning (Specialization)

- **LLM:** Specialized training on narrower distributions and specific tasks
- **Contemporary education:** Late-stage specialization under throughput constraints
- **Energy:** Reduced breadth investment, targeted domain output
- **Result:** Domain-specific performance at cost of cross-domain integration

7.5.3 Phase 3: RLHF (Preference Shaping)

- **LLM:** Reinforcement learning from human feedback—preference and constraint shaping via rankings and reward modeling ([Ouyang et al., 2022](#))
- **Contemporary assessment:** Continuous evaluation ensuring compliance with expected response patterns

- **Energy:** Minimal capability investment, maximum behavioral control
- **Result:** Predictable outputs aligned with institutional preferences

If one seeks historical analogy: pre-training resembles encyclopedic exposure, fine-tuning resembles professional specialization, and RLHF resembles institutional norm-enforcement. The parallel is structural, not causal. Both progressions move from energy-intensive comprehensiveness toward energy-minimal compliance. Both sacrifice capability range for output predictability.

7.6 The Recursive Feast

The most troubling implication: LLMs increasingly train on text produced by humans who were trained to think in standardized patterns. Research on model collapse documents what happens when this recursion proceeds unchecked.

[Shumailov et al. \(2024\)](#) demonstrate that training on model-generated data causes distributional collapse, with tail distributions lost irreversibly across generations. [Ale-mohammad et al. \(2024\)](#) show that without fresh human-generated data, quality and diversity degrade systematically. The phenomenon is not speculation—it is documented machine learning pathology.

The recursive loop:

1. **Generation 1:** Humans trained to process information in standardized, assessable patterns
2. **Generation 2:** Machines trained on standardized human outputs
3. **Generation 3:** Humans learning from machines trained on standardized human thought
4. **Generation 4:** Machines training on outputs from humans who learned from machines...

Each iteration risks losing distributional tails—the unusual, the surprising, the synthesis that doesn't fit standard patterns. The LLM trained on academic papers written by scholars optimizing for impact factors produces text optimized for...impact factors. The system risks convergence toward what [Shumailov et al. \(2024\)](#) call “model collapse”: maximum optimization, minimum diversity.

Mitigations exist: provenance filtering, deliberate human data collection, mixing fresh real data with synthetic outputs. But these require recognizing the problem. And the problem is invisible to systems that measure only what they optimize for.

As an illustrative trajectory, OpenAI's GPT progression suggests the pattern:

- GPT-2 (2019): Trained on diverse internet text, high variance in outputs

- GPT-3 (2020): Trained on curated corpora, more consistent but narrower distribution
- GPT-4 (2023): Trained on refined data plus extensive RLHF, reliable but predictable
- Future risk: Training increasingly on AI-generated text, approaching distributional collapse

The thermodynamic frame applies: each generation’s optimization reduces the entropy of the training distribution. Order increases. But order in thermodynamic systems comes at a cost—the elimination of the high-energy states that enable novel configurations.

7.7 The Thermodynamic Proof

The energy requirements reveal fundamental architectural differences:

7.7.1 Human Cognition (Historical)

- **Formation:** 15–25 years continuous biological energy investment
- **Maintenance:** Lifetime energy requirement for neuroplasticity and adaptation
- **Adaptation:** Constant energy for contextual learning and environmental response
- **Creativity:** High-energy states enabling novel connections across domains

7.7.2 LLM Processing

- **Training:** One-time massive energy expenditure (estimated 1,287 MWh for GPT-3-scale models; [Patterson et al. 2021](#))
- **Inference:** Minimal energy for pattern matching against frozen parameters
- **Adaptation:** No endogenous plasticity at inference; adaptation is exogenous (weight updates, retrieval augmentation, tool use)
- **Creativity:** Constrained to recombination within learned manifolds; novelty lacks world-anchored truth conditions

The distinction is not that humans are “creative” and machines are not—philosophers can debate that indefinitely. The distinction is thermodynamic: humans maintain negative entropy through continuous energy investment. Their cognitive structures remain plastic, responsive, capable of genuine novelty through ongoing metabolic work.

LLMs are frozen patterns—entropy crystals that degrade in performance under distribution shift, data contamination, or recursive training unless externally maintained.

They do not maintain themselves. They do not adapt endogenously. They do not invest energy to preserve capability against decay.

The educational trajectory we've documented moves human cognition toward the LLM pattern: front-loaded training creating relatively static competency profiles rather than continuous adaptive growth. We've trained humans to be more like frozen weights: formed once, deployed repeatedly, updated only through expensive external intervention.

7.8 The Mirror

The AI mirror reveals our institutional self-portrait. We see in LLMs the optimization logic we've imposed on ourselves:

- They process tokens because we process credits
- They embed in vector spaces because we embed in competency frameworks
- They attend selectively because we assess selectively
- They optimize for loss functions because we optimize for grades

The machines aren't achieving consciousness; we're achieving computability. The convergence point isn't artificial general intelligence but optimized specific processing. We meet our creations in the middle—in the diminished space where neither human wisdom nor machine efficiency fully exists, only the automated processing of pre-processed patterns.

LLMs are successfully replacing human cognitive work not because they've achieved human capability but because we've reduced human capability to what machines can replicate. The optimization homology is operationally tight in the dimensions that institutions measure: both sides now privilege what can be tokenized, vectorized, and optimized. What resists these operations—tacit knowledge, contextual judgment, integrative wisdom—becomes invisible to both systems.

We trained ourselves for replacement. The machines simply arrived to occupy the positions we'd prepared.

The mirror is perfect because we ground both sides to match.

7.9 Empirical Probes of the Optimization Homology

The framework generates testable predictions. Three empirical probes would strengthen or falsify the homology thesis:

1. **Stylometric convergence probe:** Compare academic writing corpora pre- and post-Bologna standardization for features including template density, hedging

patterns, structural regularity (section headings, formulaic phrases, citation density), and lexical diversity. Measure language-model perplexity of texts across cohorts. Hypothesis: text becomes more “predictable” (lower perplexity to language models) as standardization increases. LLM performance should correlate with the degree of standardization in training domains. *Negative control*: Compare domains minimally impacted by Bologna standardization (e.g., creative writing, some humanities traditions) to test whether the trend is standardization-specific.

2. **Assessment-alignment probe**: Compare LLM performance across domains differentiated by assessment structure. Hypothesis: LLM answer quality correlates with the extent a domain is evaluated via standardized outcomes (multiple-choice examinations, rubric-scored essays) versus apprenticeship-based or tacit knowledge domains. Domains with explicit learning outcomes should show higher LLM competence than domains relying on implicit expert judgment. *Confound acknowledgment*: LLM performance also correlates with internet text availability; analysis should control for corpus volume and genre representation.
3. **Context-suppression probe**: Analyze competency framework descriptions for systematic elimination of context-dependent elements. Count modal verbs and conditional markers (“if,” “unless,” “depending,” “when appropriate”) versus declarative competency statements (“can identify,” “can apply,” “demonstrates”). Hypothesis: EQF-style frameworks systematically retain rule-like declarative statements while suppressing phronesis-type contextual descriptors. The ratio of declarative to contextual language should increase with standardization level and framework formalization.

These probes do not require this paper to execute them. They demonstrate that the optimization homology is not mere philosophical assertion but a research program generating falsifiable hypotheses. The mirror thesis stands or falls on empirical investigation, not rhetorical force alone.

The strange loop completes: we use vectorized academic prose to critique vectorization. We use AI-era analysis to examine AI preparation. We use standardized citation formats to attack standardization. The contradiction is not hypocrisy—it is method. Treating the reflexivity as a refutation demonstrates the vectorized thinking the argument describes: the inability to hold productive paradox.

8 Reconstruction Principles: Building Cognitive Sovereignty Within Thermodynamic Constraints

8.1 The Epistemological Prerequisites

Before proposing reconstruction pathways, we must acknowledge a fundamental paradox: those seeking to rebuild cognitive sovereignty are themselves products of the vectorization they seek to escape. The map is never the territory ([Korzybski, 1933](#)), and our maps were drawn with vectorized tools.

Contemporary decision science reveals the illusion of individual rationality. As [Kahneman \(2011\)](#) demonstrated, human decision-making operates through systematic biases rather than rational calculation. More fundamentally, distributed cognition theory ([Hutchins, 1995](#)) establishes that thinking occurs not within individual minds but across socio-technical systems. What we experience as “our” decisions emerge from complex interactions between cultural values, institutional practices, environmental constraints, and collective sense-making processes.

Highly formalized training can overfit people to legible problems and under-train them for domain ambiguity. The analytical tools revealing the vectorization problem are themselves products of vectorized thinking. We are attempting to examine the lens through which we see using that same lens: [Hofstadter \(1979\)](#) strange loop made manifest.

Critical retrospection of any recent non-trivial decision reveals this embedding. The choice to pursue alternative educational approaches, implement new organizational structures, or resist AI integration emerges not from individual cognition but from:

- Cultural layer: Prevailing beliefs about knowledge and value ([Schein, 1985](#))
- Institutional layer: Organizational structures and incentive systems ([DiMaggio and Powell, 1983](#))
- Social layer: Peer networks and professional communities ([Granovetter, 1985](#))
- Individual layer: Personal history and embodied experience ([Bourdieu, 1990](#))

The individual functions as one node in a distributed processing network: “the literal neuron of a bigger brain.” Distributed does not mean powerless; it means leverage points exist at multiple scales. This recognition demands epistemic humility. We cannot stand outside the system to reconstruct it. As [Maturana and Varela \(1987\)](#) establish, we are “structurally coupled” to the environment we seek to transform.

8.2 The Cynefin Diagnostic Framework

Snowden’s Cynefin framework (Snowden and Boone, 2007) provides the essential diagnostic tool for understanding where reconstruction is both necessary and possible. The framework distinguishes five domains, each requiring different cognitive architectures and intervention strategies.

Table 5: Cynefin Reconstruction Design Matrix^a

Domain	Suitable Interventions	Inter- Failure Modes	Evaluation Signals
Clear	Automate + standardize; best practices	Over-automation of edge cases; complacency	Error rate; throughput; drift monitoring
Complicated	Expert review + checklists; good practices	Expert bottlenecks; analysis paralysis	Peer audit accuracy; decision latency
Complex	Safe-to-fail probes + narrative capture; emergent practices	Premature convergence; probe-aversion	Learning velocity; resilience indicators; adaptation rate
Chaotic	Stabilizing action + after-action learning; novel practices	Paralysis; inappropriate deliberation	Time-to-stabilize; harm minimization
Confused	Escalation protocol + reclassification	Misclassification persistence; domain forcing	Time spent in disorder; reclassification accuracy

^aEvaluation signals are illustrative proxies; operational definitions appear in Section 8.9.

The critical insight: organizations systematically misclassify complex problems as complicated, applying “best practices” where emergent approaches are needed. Snowden and Boone, who developed the Cynefin framework specifically to name this failure mode, document that the characteristic error is treating complex problems as merely complicated—deploying expert analysis and standardized responses where only probe-sense-respond strategies can work, producing failures that repeat because the diagnostic error is never corrected (Snowden and Boone, 2007).

Table 5 operationalizes the diagnostic framework into a decision instrument. The reconstruction challenge is not merely cognitive but organizational: building systems capable of accurate domain classification and appropriate response selection.

8.3 Thermodynamic Constraints on Reconstruction

Reconstruction faces inescapable thermodynamic constraints. The Second Law is a constraint on slack, maintenance, and reversibility—not a full causal explanation, but an

analytic boundary that strategies must respect.

8.3.1 The Sovereignty Budget Constraint

The power-budget allocation model (Equation 1, Section 3) establishes the thermodynamic constraint: cognitive sovereignty requires sustained allocation of the brain's baseline power toward synthesis-supporting states, with architectural quality determining extraction resistance.

$$S = P_{brain} \times \alpha_{synth} \times R \quad (8)$$

Where $P_{brain} \approx 20W$ represents baseline cerebral metabolism, $\alpha_{synth} \in [0, 1]$ represents synthesis-time allocation, and $R \in [0, 1]$ represents extraction resistance. This is the framework's core constraint: Energy Investment Rate must exceed Entropy Rate, or cognitive architecture degrades toward the thermodynamic floor where patterns become simple enough for algorithmic extraction.

For reconstruction planning, these terms admit operational proxies:

- **Synthesis Allocation** (α_{synth}): Hours of deep work + deliberate practice + embodied practice per week, as proxy for synthesis-state time fraction
- **Extraction Resistance** (R): Operationalized via substitutability, transfer distance, tacit load, and integration density (defined below)
- **Entropy Rate**: Rate of architectural degradation from extraction pressure

8.3.2 Operationalizing the Proxies

The following constructs require operational specification to avoid hand-waving:

- **Entropy Rate** (degradation pressure): Interruptions per day, context-switch frequency, sleep debt accumulation, administrative load fraction, and rework cycles. Higher values indicate faster architectural erosion ([Mark et al., 2008](#); [González and Mark, 2004](#)).
- **Substitutability** (automation exposure): Degree to which tasks decompose into rule-following steps, availability of standardized input-output pairs, and existence of comparable automation benchmarks in the domain ([Frey and Osborne, 2017](#); [Autor, 2015](#)).
- **Transfer distance**: Number of domain-boundary crossings required for an artifact or decision, assessed via domain-tag schemes. Higher transfer distance indicates greater integration demand.

- **Tacit load** (explainability gap): Proportion of task performance that resists explicit articulation. Assessed via protocol analysis or apprenticeship duration requirements (Collins, 2010).
- **Integration density**: Network density of cross-references in personal knowledge systems, or count of cross-domain analogies deployed per unit time. Higher density indicates sphere-like architecture.

These remain proxies, not direct measures. They provide calibration for reconstruction efforts while acknowledging that cognitive architecture itself resists direct observation.

8.3.3 Historical Investment Benchmarks

Historical evidence reveals the investment scale required for different knowledge types:

- **Sophia** (theoretical wisdom): 20+ years sustained investment across multiple domains
- **Phronesis** (practical wisdom): Lifetime daily practice with contextual feedback
- **Techne** (craft knowledge): Decade-scale deliberate practice with embodied feedback loops; deliberate practice research confirms substantial time investment though magnitude varies strongly by domain (Ericsson et al., 1993; Macnamara et al., 2014), while historical apprenticeships required 7–10 years before journeyman recognition (De la Croix et al., 2018)
- **Metis** (adaptive cunning): Constant exposure to novel challenges; cannot be scheduled
- **Nous** (intuitive insight): Unknown threshold, possibly requiring conditions no longer reproducible

Modern constraints make full reconstruction rare; partial reconstruction requires deliberate slack engineering:

- Work demands: 40–60 hours/week of vectorized activity
- Economic pressure: Continuous productivity requirements consuming baseline cognition
- Attention economy: Optimized extraction mechanisms operating continuously (Zuboff, 2019)
- Social expectations: Optimization for measurable outcomes at every scale

8.4 The Equity Constraint

Cognitive sovereignty is unevenly accessible. Reconstruction presumes access to slack: time, money, health, safety, community. Those lacking economic security cannot protect synthesis time; those lacking social support cannot sustain counter-cultural practice; those facing discrimination encounter additional cognitive load that consumes the very resources reconstruction requires.

This recognition has two implications. First, individual reconstruction protocols (Section 8.6) must be understood as conditional on resource availability, not universally applicable self-improvement advice. Second, institutional reconstruction (Section 8.7) becomes not merely strategic but ethically necessary: organizations and societies must provision the slack that individuals cannot secure alone.

Reconstruction is not personal virtue; it is collective infrastructure. The thermodynamic requirements apply regardless of individual merit or motivation. Systems that fail to provide protected synthesis time, externalized survival load, and extended time horizons will produce vectors regardless of curriculum content, pedagogical innovation, or inspirational leadership.

8.5 Boundary Conditions and Failure of Applicability

The reconstruction protocols that follow assume baseline stability. They fail under conditions including:

- **Acute crisis:** War, displacement, acute illness, or immediate survival threat. Cognitive resources necessarily redirect to threat response; synthesis becomes impossible and inappropriate.
- **Extreme precarity:** Housing instability, food insecurity, or imminent financial collapse. The cognitive load of survival planning consumes all available slack.
- **Severe caregiving load:** Primary responsibility for dependents with high needs. Available cognitive resources are legitimately allocated elsewhere.
- **Institutional capture:** Employment conditions that permit zero discretionary time or penalize non-optimized activity. Exit may be prerequisite to reconstruction.
- **Health constraints:** Chronic conditions affecting cognitive capacity, energy availability, or time horizons. Protocols require adaptation to actual capacity, not normative expectations.

These are not excuses but constraints. Acknowledging them transforms equity from ethical note into methodological boundary: the protocols apply where baseline conditions exist, and institutional intervention is required where they do not.

8.6 Individual Reconstruction Protocols

The following protocols assume access to baseline resources: stable housing, adequate nutrition, physical safety, and approximately 10 hours/week of discretionary time. Where these conditions are absent, individual reconstruction becomes impossible without prior institutional intervention (see Section 8.4).

8.6.1 Phase 1: Diagnostic Assessment (Month 1)

Current State Analysis

- Map daily activities to Cynefin domains using Table 5
- Identify vector lock-in patterns (tasks that could be fully specified as algorithms)
- Calculate actual discretionary time/energy after survival maintenance
- Assess extraction vulnerability using substitutability indicators

Sphere Potential Mapping

- Identify existing cross-domain connections (latent integration capacity)
- Recognize capacities from pre-vectorization periods (childhood interests, abandoned pursuits)
- Map curiosity patterns that resist optimization pressure
- Locate energy reserves available for investment

Realistic Timeline Acceptance

Accept thermodynamic reality:

- No shortcuts exist (constraints are non-negotiable)
- Decade-scale investment required for substantive sphere development
- Starting now provides meaningful time before labor-market exposure to AI substitution intensifies
- Partial development offers meaningful protection; perfectionism is vectorized thinking

8.6.2 Phase 2: Foundation Building (Months 2–6)

Cross-Domain Exposure

- Read deliberately outside primary field (minimum 2 hours/week protected)
- Attend events in unfamiliar domains without optimization goals
- Engage communities with different knowledge traditions
- Follow curiosity without demanding measurable outcomes

Pattern Recognition Development

- Maintain synthesis journal for cross-domain insights
- Practice analogical thinking between disparate fields
- Build personal pattern library (not filing system—integration network)
- Resist premature categorization; tolerate ambiguity

Embodied Practice Initiation

Essential for developing non-extractable knowledge ([Collins, 2010](#)):

- Physical craft (woodworking, pottery, gardening)
- Movement practice (martial arts, dance, climbing)
- Musical instrument learning
- Somatic awareness development

8.6.3 Phase 3: Cultivation (Months 6–24)

Deliberate Integration

- Create projects requiring multiple domain knowledge
- Write integrative analyses crossing disciplinary boundaries
- Teach others using cross-domain metaphors (synthesis requires articulation)
- Build integration as default cognitive mode

Complexity Navigation Practice

- Seek problems without predetermined solutions
- Practice probe-sense-respond in safe contexts

- Embrace emergence and uncertainty as signals, not failures
- Document pattern recognition development for calibration

Community Building

Sphere development requires collective support—individual reconstruction is thermodynamically implausible:

- Find others attempting reconstruction (shared context reduces explanation cost)
- Create learning partnerships with complementary domains
- Share failures and insights (distributed error-correction)
- Build counter-cultural spaces protecting synthesis time

8.7 Organizational Reconstruction Architecture

8.7.1 From Silos to Sphere Teams

Replace functional specialization with cross-domain capacity:

- Pilot teams mixing 3–5 different expertise domains
- Complex challenge focus rather than efficiency optimization
- Protected learning time: minimum 20% allocation to non-productive exploration. Below this threshold, exploration collapses into exploitation and learning becomes extraction ([March, 1991](#))—an organizational entropy tax that cannot be avoided.
- Success measured by adaptation capacity, not standardization compliance

8.7.2 From Best Practices to Emergent Practices

Shift from standardization to contextual response:

- Train all staff in Cynefin framework and domain classification
- Create “probe budgets” for complex challenges (resources allocated to safe-to-fail experiments)
- Document emergent solutions as narratives, not procedures
- Celebrate contextual wisdom over universal protocols

8.7.3 From Extraction to Investment

Reverse the energy flow:

- Implement genuine apprenticeship programs (3+ years, not bootcamps)
- Create internal learning structures for cross-domain development
- Protect thinking time from productivity metrics (unmeasured is not unproductive)
- Measure knowledge development trajectories, not just application outputs

8.8 Governance of Tools in AI-Saturated Environments

Reconstruction does not require AI avoidance—an impossibility in contemporary environments. It requires governance structures that preserve human cognitive development while leveraging algorithmic capacity appropriately.

8.8.1 Domain-Appropriate Deployment

- **Clear domain:** AI deployment appropriate for throughput optimization; human oversight for edge-case detection
- **Complicated domain:** AI as analysis support; human experts retain decision authority
- **Complex domain:** AI prohibited as decision authority; may assist in pattern surfacing, option generation, and documentation, but humans retain probe-sense-respond agency
- **Chaotic domain:** AI outputs are non-authoritative; use only as optional situational prompts while humans act to stabilize
- **Confused domain:** AI classification assistance only; human meta-cognition required for domain determination

8.8.2 Structural Safeguards

- Require “human appeal paths” for any automated decision affecting rights, resources, or relationships
- Treat all model outputs as hypotheses requiring independent verification
- Track “deskilling risk” as first-class operational metric
- Implement “synthesis breaks”—periods of AI-free cognitive work to maintain integration capacity

8.8.3 Practical Implementation

Cynefin-informed AI governance can be operationalized within existing organizational constraints. Frameworks such as the author’s “Karten auf den Tisch” protocol (Cotoaga, 2024) demonstrate that reconstruction tools can be developed and deployed without requiring complete system exit. The challenge is not technological but organizational: building governance structures that treat cognitive sovereignty as a first-class design constraint rather than an efficiency impediment.⁸

8.9 Measurement and Evaluation

Reconstruction requires outcome measures to distinguish from manifesto. The following indicators operationalize sovereignty development:

8.9.1 Individual Indicators

- **Complex work ratio:** Percentage of time spent in complex/chaotic domains without premature reduction to complicated/clear
- **Integration artifact production:** Weekly cross-domain outputs (notes, diagrams, syntheses) demonstrating connection-building
- **Embodied practice cadence:** Hours/week in somatic skill development
- **Automation dependency index:** Frequency of inability to proceed without tool outputs; lower indicates higher sovereignty
- **Transfer distance achieved:** Demonstrated application of learning across domain boundaries

8.9.2 Organizational Indicators

- **Probe ratio:** Fraction of initiatives run as safe-to-fail experiments (complex domain appropriate)
- **Protected time utilization:** Percentage of allocated learning time actually taken (not “available but unused”)
- **Drift/incident rates:** Failures in automated decision systems requiring human intervention
- **Apprenticeship metrics:** Retention rates and skill progression in long-term development programs

⁸The “Karten auf den Tisch” framework operationalizes domain classification and AI deployment boundaries for organizational decision-making. Available at <http://cotoaga.ai/karten-auf-den-tisch/>.

- **Cross-domain mobility:** Internal movement between functional areas indicating integration capacity
- **Reclassification accuracy:** Correct domain diagnosis in post-hoc analysis

These indicators are proxies, not direct sovereignty measures. They provide calibration for reconstruction efforts while acknowledging that cognitive architecture itself resists direct observation.

8.10 Existence Proofs: Documented Reconstruction Trajectories

The thermodynamic analysis might suggest reconstruction is merely theoretical—possible in principle but unrealized in practice. Evidence to the contrary requires documented cases of vector trap escape under contemporary conditions.

8.10.1 Methodological Note

This section employs autoethnographic vignette as mechanism-illustration rather than representative evidence. The author’s trajectory demonstrates *feasibility* under specific (and atypical) resource conditions, not universal applicability. External cases provide triangulation.

8.10.2 Vignette: Six Vector Trap Navigations

The author’s trajectory provides one existence proof. Six distinct vector traps encountered and navigated:

1. **Academic vector trap:** PhD program in quantitative finance (1990s). Exit: Abandoned for internet entrepreneurship when the field’s vectorization trajectory became apparent.
2. **Technical vector trap:** Employee #5 at Germany’s first online brokerage; built technical chart analysis system (KAPITAL1st). Exit: Recognized pattern-matching as ultimately extractable; pivoted to synthesis roles.
3. **Optimization vector trap:** Traveling Salesman Problem algorithm achieving global ranking leadership for six months (2004). Exit: Understood that algorithmic excellence without integration produces sophisticated vectors, not spheres.
4. **Geographic vector trap:** New York environment optimizing for extraction. Exit: Physical relocation (overland to South America) as literal escape from optimization pressure.
5. **Corporate vector trap:** Management consulting roles offering vectorized “strategic” work. Exit: Independent practice preserving synthesis capacity.

6. **Product vector trap:** Platform development pressure toward feature optimization. Exit: Maintained integration focus over standard product management vectorization.

These escapes required: recognition of vector pressure before capture became irreversible, resources enabling transition (often accumulated during vector phases), willingness to sacrifice optimization rewards for sovereignty preservation, and extended time horizons incompatible with quarterly thinking.

8.10.3 External Triangulation

The autoethnographic vignette gains credibility through external cases demonstrating similar patterns:

- **Craft apprenticeship resurgence:** The revival of traditional apprenticeship in woodworking, blacksmithing, and textile arts documents intentional rejection of credentialized training in favor of embodied knowledge transmission ([Sennett, 2008](#)).
- **Slow science movement:** [Stengers \(2018\)](#) advocates research cultures that resist the acceleration pressures and productivist tempo of contemporary science, articulating an institutional-scale reconstruction agenda for academic practice.
- **Complexity practitioner networks:** Communities organized around Cynefin and related frameworks show sustained counter-practice against best-practice ideology in organizational consulting.

The existence proofs share common structure: recognition of vectorization pressure, resource accumulation enabling transition, deliberate sacrifice of optimization rewards, and community support for counter-cultural practice. The question shifts from “Is reconstruction possible?” to “What conditions enable it?”—a question tractable through research rather than philosophy alone.

8.11 The Budgetary Choice

The choice is not binary; it is budgetary. Every system—individual, organizational, civilizational—allocates resources between sovereignty investment and extraction optimization, explicitly or implicitly.

Full sovereignty investment requires conditions approaching historical patronage: externalized survival costs, protected synthesis time, extended time horizons, and freedom from productivity measurement. These conditions existed for medieval scholars and Renaissance polymaths; they do not exist at civilizational scale today.

Partial sovereignty investment represents the realistic frontier: deliberate allocation of constrained resources toward reconstruction within extraction-dominated environments. This is not compromise but strategic necessity—the thermodynamically feasible path.

Zero sovereignty investment—whether through unawareness, resource absence, or optimization pressure—produces complete vectorization. This path offers immediate efficiency gains and guaranteed long-term obsolescence.

The reconstruction principles presented here assume partial investment: readers operating within extraction systems who seek to preserve or develop synthesis capacity. Full reconstruction would require institutional transformation beyond individual scope; zero reconstruction requires no guidance.

8.12 Critical Success Factors and Failure Patterns

Success Requirements:

- Accept decade-scale timeline (thermodynamic reality, not motivational failure)
- Protect energy investment despite productivity pressure
- Build community support (individual reconstruction is thermodynamically implausible)
- Measure progress in capability development, not credential acquisition
- Maintain sovereignty despite continuous extraction attempts

Common Failure Patterns:

- Treating reconstruction as temporary project rather than permanent practice
- Attempting alone without community support (energy requirements exceed individual capacity)
- Measuring by vectorized metrics (efficiency, speed, output volume)
- Expecting linear progress in complex domain development
- Underestimating energy investment required; abandoning after insufficient duration
- Optimizing the reconstruction process itself (meta-vectorization)

The Fundamental Recognition:

Reconstruction is not optimization, self-improvement, or productivity enhancement. It is choosing cognitive sovereignty over efficiency, wisdom over productivity,

and decades of investment over immediate returns. The choice is budgetary: allocate resources toward extraction resistance or accept dissolution into algorithmic substitutability.

A common objection holds that framing constraints as non-negotiable constitutes “thermodynamic determinism.” This objection assumes constraints should be negotiable—that sufficient cleverness or effort can circumvent physical limitations. The framework instead treats constraints as boundary conditions within which strategies operate. Constraints define the possibility space; strategies navigate within it. The Second Law does not prevent reconstruction; it specifies what reconstruction requires.

9 Implications: When Physics Meets Mythology

Implications Summary

The central implication is not that systems fail—failure is ubiquitous—but that how we produce knowledge about failure, how we train cognition, and how we govern deployment are coupled through legibility-and-optimization regimes that systematically select against the capabilities complexity demands. The three domains examined below (epistemic, educational, deployment) are not separate problems requiring separate solutions; they are facets of a single dynamic in which extraction without investment produces brittleness across institutional scales.

9.1 Epistemic Implications: How Knowledge Production Failed

The most revealing evidence for our thesis comes not from transformation failures themselves, but from how we created knowledge about those failures. By “mythology” we mean institutionalized claims that acquire the status of knowledge through repetition, pedagogical embedding, and market adoption rather than through empirical validation—legitimizing stories that function as if verified while remaining epistemically hollow.

[Hughes \(2011\)](#) traces the genealogy of the “70% failure rate”—the statistic cited in thousands of management texts—and discovers no empirical foundation. [Beer and Nohria \(2000\)](#) stated it as “brutal fact” without evidence. [Kotter \(2008\)](#) cited it as accepted truth without references. Through what Hughes calls “unconscious collusion,” assertion became fact through citation networks alone.

Then measurement arrived. [Bain & Company \(2024\)](#) surveyed 400+ executives and found 88% of transformations failed to meet original ambitions. [Boston Consulting Group \(2024a\)](#) analyzed 1,700+ transformation initiatives and found 74% did not deliver sustained value. [Ford et al. \(2024\)](#) documented 1,205 non-conformance incidents in a single megaproject, revealing systematic complexity misclassification.

9.1.1 A Note on Construct Comparability

These statistics measure different constructs and should not be treated as commensurate metrics. The mythological “70%” lacks any operational definition. Bain’s 88% measures failure to meet original ambition—a high bar that includes partial successes. BCG’s 74% captures failure to sustain value over time. Ford’s analysis counts specific non-conformance incidents rather than project-level outcomes.

We place them together not as equivalent measurements but as convergent indicators of institutional brittleness under complexity. The pattern matters more than the precise percentages: systems optimized for efficiency systematically struggle when reality presents genuine complexity. That the mythology accidentally approximated measured reality—while explaining none of its causes—illustrates the epistemological trap our framework addresses.

9.1.2 The Canonization Problem

Kotter’s eight-step change model exemplifies how practitioner intuition becomes institutionalized as validated theory. As [Hughes \(2015\)](#) documents, Kotter acknowledged: “I neither drew examples nor major ideas from any published source, except my own writing.” This is not hedged confession; it is direct admission that the most influential change management framework of the past thirty years—taught in every MBA program, implemented across thousands of organizations, cited as authoritative in peer-reviewed literature—has no empirical foundation whatsoever.

The framework predates widespread internet access, yet survived into an era when verification became trivial. That no one checked—that citation networks propagated claims without validation, that business schools canonized consulting heuristics as curriculum, that organizations implemented methodology built on nothing—proves our thesis more directly than any failure statistic. We do not verify. We do not validate. We accept simplified models that feel actionable, cite them into existence, then act surprised when reality refuses to comply.

[McLaren et al. \(2023\)](#) identify the framework’s internal contradiction: Kotter acknowledges that people prefer status quo over uncertainty, then prescribes that leaders “make the current situation look more dangerous than launching into the unknown.” The model deploys fear of staying against documented preference for stability—working against human psychology while claiming to manage human change.

The mechanism through which such frameworks propagate deserves explicit identification:

1. **Market incentives:** Consulting and training markets reward actionable templates over nuanced analysis; complexity doesn’t sell.
2. **Commensuration regimes:** Audit cultures and accountability systems demand

measurable frameworks, selecting for legibility over validity ([Espeland and Stevens, 1998](#)).

3. **Citation cascades:** Each citation increases perceived authority regardless of underlying evidence; networks amplify without filtering.
4. **Curriculum canonization:** Business schools facing market pressure adopt “industry-standard” frameworks, converting consulting products into educational content.
5. **Managerial demand:** Organizations under transformation pressure seek templates that promise control; frameworks offering eight steps outcompete those acknowledging irreducible uncertainty.

This is how vectorized institutions process knowledge: simplified models that feel actionable propagate regardless of empirical grounding, while frameworks acknowledging complexity face selection pressure at every stage from creation through canonization.

9.2 Educational Implications: When Industry Knows What Academia Denies

9.2.1 The German Resistance

The Bologna Process aimed at genuine goods: student mobility, credential comparability, expanded access across European higher education. The critique here concerns implementation that optimized for these administrative goals while eliminating the cognitive development they were meant to serve.

TU Dresden engineering faculty state the trade-off plainly: “The university Bachelor’s degree in six semesters does not lead to a professionally qualifying degree...essential components must be omitted...This damages the foundation of engineering education” ([Odenbach and Krauthäuser, 2015](#)). This is not nostalgia from professors resisting change; it is documented assessment from faculty responsible for producing engineers.

German industry responded with market signals that academia cannot easily dismiss. Collective bargaining agreements in engineering sectors assign different wage groups to Diplom versus Bachelor/Master holders ([Wieschke et al., 2020](#)). Master’s graduates earn wage premiums “even relative to those with more work experience”—employers paying for demonstrated capability that credentials alone do not guarantee. The TU9 consortium, educating 47% of German engineers, formally requested the right to award “Diplom-Ingenieur” alongside Master’s degrees.

These market signals reflect employer assessment, though confounds exist: TU9 institutions attract different student populations, sectoral wage structures vary, and signaling effects are difficult to separate from human capital differences. Yet [Kaiser](#)

and Schröder (2022) provide direct measurement of what modularization removed: “Systems Thinking, collaboration and communication are not explicitly addressed” in reformed engineering curricula. The very capabilities that distinguish engineers from algorithms—integrating across domains, navigating emergence, seeing systems rather than components—were systematically de-emphasized by optimization for modularity and mobility.

VDI President Ungeheuer’s assessment: “We sacrificed the Diplom-Ingenieur with heavy hearts for a greater goal, namely international connectivity.” The trade-off was made consciously, for defensible reasons. But thermodynamic constraints do not recognize international agreements or administrative convenience. If sphere-building requires sustained, integrated cognitive investment, then fragmenting that investment into stackable modules—however administratively rational—reduces the probability of developing the synthesis capacity that complexity demands.

9.2.2 The Micro-Credential Trajectory

The educational implications extend beyond German engineering. As documented in Section 6, Lumina Foundation (2025) reports 96% employer approval for micro-credentials, yet Ha et al. (2022)’s systematic review reveals that positive outcomes measure satisfaction and employment signals rather than demonstrated competence.

The trajectory approaches a thermodynamic limit: credentials of decreasing duration (years → months → weeks → hours), measuring increasingly narrow competencies, optimizing for signaling efficiency rather than cognitive development. Because most evaluations track placement and satisfaction rather than transfer distance, error recovery, or integration performance, current evidence cannot support claims about synthesis development. The framework developed in preceding sections predicts these capacities will remain weak unless programs deliberately invest in integrative practice—a prediction that micro-credential market incentives systematically discourage.

Market signals can partially compensate for credential hollowing, but often late, unevenly distributed, and with significant transitional harms. The German engineering case suggests correction is possible; it does not suggest correction is automatic or costless.

9.3 Deployment Implications: The Extraction-Investment Asymmetry

9.3.1 What Contemporary Evidence Reveals

Section 6 documented specific AI deployment failures through primary sources: court rulings, regulatory findings, and executive acknowledgments. The 42% initiative abandonment rate (S&P Global Market Intelligence, 2025), the documented reversals at

Klarna and Commonwealth Bank—these are not implementation problems awaiting technical solutions. They are predictable consequences of extraction-based deployment strategies encountering genuine complexity.

The pattern across cases shares a common signature: organizations extract cognitive value from human workers, encode partial patterns into automated systems, then discover that neither degraded human capabilities nor brittle automated systems can navigate contexts exceeding training distributions. [Rinta-Kahila et al. \(2022\)](#) document this pattern in detail in their analysis of Australia’s Robodebt automated debt-recovery system, framing the failure mode as “system destructiveness”: algorithmic error compounds through human-machine feedback loops until both components degrade simultaneously and the joint system collapses on the population it governs. The same signature recurs in customer service, healthcare decision support, and legal research: reduced practice erodes skill maintenance, skill erosion increases automation dependency, dependency accelerates further automation pressure, until neither humans nor machines can recover the situational judgment the original workflow required.

9.3.2 Contestation-Cost Asymmetry: A Named Mechanism

The most troubling cases documented in Section 6 reveal a governance failure pattern we term **contestation-cost asymmetry**: systems designed such that error-correction costs fall on those least equipped to bear them, while error-generation benefits accrue to deploying institutions.

UnitedHealth’s nH Predict algorithm operated with a 90% error rate on appealed denials—not because the technology failed, but because institutional incentives tolerated high error rates when contestation costs exceeded patient resources ([Napolitano, 2023](#)). The extraction calculus exploited asymmetry: savings from denials exceeded costs of the small percentage who appealed. The algorithm functioned exactly as designed; the design optimized for extraction under conditions of asymmetric contestation capacity.

This represents a distinct failure mode from capability limitations. Even perfectly capable AI systems, deployed within governance structures optimizing for extraction, produce systematic harm to populations with limited contestation resources. The thermodynamic framing applies: institutions extracting value from decision-making systems without investing in validation, oversight, and accessible appeal mechanisms accumulate entropy that manifests as harm distributed according to power asymmetries rather than error distributions.

Contestation-cost asymmetry explains why high error rates persist in consequential domains: the costs of errors and the costs of correcting them fall on different parties. Effective governance must address this asymmetry directly, not assume that capability improvements alone will resolve deployment harms.

9.4 The Constraint Landscape

The Cynefin framework provides domain-conditional structure for these implications: in clear and complicated domains, optimization can raise performance; systems benefit from efficiency investment. In complex and chaotic domains, optimization without corresponding investment in adaptive capacity increases brittleness—the system becomes more efficient at handling anticipated conditions while becoming more fragile when conditions exceed design parameters. The implications below are conditional on domain: what helps in complicated contexts harms in complex ones.

For Policy: Regulatory frameworks treating AI deployment as purely a capability question miss the governance dimension. Contestation-cost asymmetry means technically adequate systems can still produce systematic harm when institutional incentives favor extraction over validation. Effective governance requires:

- **Mandatory human appeal paths** with costs borne by deploying institutions, not affected individuals
- **Auditability requirements** scaled to domain complexity—higher-stakes and higher-complexity deployments requiring more rigorous validation
- **Duty-to-monitor provisions** for deskilling risks in high-stakes domains where human oversight remains nominally required
- **Penalties keyed to error-plus-contestation-burden**, not error rates alone—making contestation-cost asymmetry exploitation economically irrational

For Educational Institutions: The German experience suggests market signals can partially correct for credential hollowing, but with significant lag and transition costs. Institutions maintaining integrated, extended programs face competitive disadvantage in the short term while potentially developing capabilities that fragmented alternatives cannot replicate. Policy choices include whether to accelerate correction through program-duration requirements, synthesis-capacity assessment mandates, or allowing market signals to operate over longer time horizons with attendant capability losses during transition.

For Organizations: The 88% transformation failure rate ([Bain & Company, 2024](#)) reflects thermodynamic constraint operating through the mechanisms identified above: mythology canonization, capability-governance conflation, and contestation-cost asymmetry. Reorganization cannot eliminate entropy; it redistributes it. Genuine transformation requires capability development: years of investment, protected learning time, acceptance of efficiency loss for capability gain. Frameworks promising rapid transformation face the same constraints limiting perpetual motion machines—the physics operates regardless of contractual commitments or consulting methodologies.

For Individuals: The choice presents as budget allocation under constraint rather than binary fate. Investing in sphere development—accepting decade-long timelines, resisting optimization pressure, building cross-domain integration capacity—increases probability of maintaining capabilities that extraction-based systems cannot replicate. The alternative is not guaranteed replacement but increased vulnerability as AI systems improve at vector-domain tasks while struggling with the synthesis that sphere architecture enables.

9.5 The Complexity Catastrophe

[Kempermann \(2017\)](#) articulates the clinical stakes: “Complex problems are often treated as if they were not more than the complicated sum of solvable sub-problems...this is not correct [and has] dangerous consequences, especially in clinical contexts.” Category errors—treating complex problems with complicated-domain tools—can contribute to preventable harm when the mismatch occurs in high-stakes environments.

[Ford et al. \(2024\)](#) provide quantitative evidence of this pattern. Analysing 1,205 categorised non-conformance reports from the £1.45 billion A14 Huntingdon Improvement highways scheme in the UK, they applied Snowden’s Cynefin framework to compare real-time domain classification against retrospective re-categorisation of the same problems. The headline finding: 397 cases initially treated as Simple were, on closer inspection, more complicated than first envisaged, requiring multiple root causes and multiple corrective actions to address. Ford and colleagues explicitly attribute this systematic oversimplification to the “cognitive miser effect”—the human bias toward less effortful routes to problem-solving, decision-making, and risk management. The pattern is not incompetence but cognitive economy: organisations optimise for efficiency in domains they understand while reality presents complexity their measurement systems were never designed to detect.

As [Alexander et al. \(2018\)](#) demonstrate, performance measurement systems assume predictability that complex domains do not exhibit. We measure what we can control, optimize what we measure, ignore what we cannot measure, and remain surprised when unmeasured complexity destroys measured achievements. The implications extend beyond individual projects: civilizational systems built on complexity-blind measurement accumulate brittleness that manifests unpredictably under stress.

9.6 Synthesis: Thermodynamics as Analytic Constraint

Every institution operating on extraction without investment faces entropy accumulation. Every optimization that reduces cognitive energy input accelerates capability degradation. Every simplification that ignores complexity increases probability of catastrophic mismatch when complexity presents itself. Every governance structure exploit-

ing contestation-cost asymmetry accumulates harm among those least able to contest it.

The German engineers maintaining five-year integrated programs despite Bologna pressure demonstrate that resistance remains possible. A few holdout institutions preserve sphere development against credential-market pressures. Individual practitioners build local capability bubbles through sustained investment that their employers may not recognize or reward. These existence proofs matter: thermodynamic constraints describe tendency, not destiny.

But the overall trajectory demands attention. We trained humans to process information like machines, documented the training methods exhaustively, then built machines that process information like we trained humans to process it. The discovery that humans trained for vector tasks perform those tasks less efficiently than purpose-built vector processors should not surprise. The more significant finding is that humans capable of synthesis—navigation across domains, integration under uncertainty, judgment that cannot be decomposed into rules—are increasingly rare, increasingly valuable, and increasingly difficult to develop within systems optimized for measurable efficiency.

The thermodynamic framing, as analytic constraint rather than metaphysical claim, offers purchase on this pattern. Systems that extract cognitive value without investing in capability maintenance face predictable brittleness. The physics—understood as heuristic identity rather than proof of equivalence—does not negotiate with institutional preferences, administrative convenience, or policy aspirations. Whether that constraint operates as firmly as the Second Law in thermodynamics proper remains an empirical question this paper cannot resolve. That it operates with sufficient force to explain documented failures across knowledge production, education, and AI deployment is the claim the evidence supports.

Build capability or face increasing probability of brittleness. The convergent evidence across domains suggests this is not rhetoric but constraint.

10 Conclusion: The Thermodynamic Reckoning

The evidence converges on an inescapable conclusion: we systematically transformed human cognition into something replaceable, documented the process exhaustively, then built the replacements. This wasn't conspiracy but optimization—each local decision rational, the collective outcome thermodynamically inevitable.

I've given you the physics. Now I'll give you my life as existence proof that the physics can be navigated.

The preface documented six vector traps and six escapes—academic, technical, strategic, product, geographic, corporate. I won't repeat the details here. What matters is the thermodynamic interpretation: each escape required energy investment our

systems are designed to prevent. Each required accepting inefficiency that optimization logic cannot tolerate. Each required choosing sovereignty over legibility. These weren't romantic gestures or privileged adventures—they were existence proof that the Second Law describes tendency, not destiny. Local negentropy bubbles can be built. Spheres can be maintained. But only through continuous investment that most cannot access and few will choose. The escapes required what Kuhn identified: being perpetually new, uncommitted to traditional rules. The outsider position isn't weakness but thermodynamic necessity—you cannot see the trap from inside it.

The historical arc is clear. From ancient academies investing decades in multidimensional wisdom to micro-credentials measured in hours, we've followed an exponential decay curve toward cognitive zero. The Bologna Process didn't just modularize education; it modularized minds. Organizations didn't just optimize processes; they optimized away judgment itself. We didn't just build AI; we first rebuilt humans to think like the AI we would build.

The thermodynamic framework reveals why reconstruction is difficult but not impossible. Knowledge isn't information—it's a high-energy state requiring continuous investment to maintain. Every efficiency gain that reduces energy input accelerates entropy. Every standardization that enables measurement destroys the unmeasurable. Every vector that replaces a sphere makes the system more optimized and more fragile, more efficient and less adaptable.

The contemporary failures—42% AI abandonment rate, 88% transformation failure rate—are not well explained by “implementation immaturity” narratives alone. They are consistent with extraction-investment asymmetry: you cannot extract what was never invested. You cannot automate what you don't understand. You cannot replace spheres with vectors and expect the system to navigate complexity. These aren't implementation problems. They're physics problems.

The confession literature provides the most damning evidence, and the clearest hope. Educational psychologists documented their standardization techniques—which means the techniques can be identified and reversed. Management consultants celebrated their extraction methods—which means the methods can be recognized and resisted. Platform designers published their manipulation strategies—which means the strategies can be seen and circumvented. They confess openly because they see no crime. Their confession is our map.

The German engineers holding onto their Diplom, TU Dresden maintaining integrated programs, individual practitioners building local negentropy bubbles: these aren't romantic holdouts but thermodynamic realists. They understand what the efficiency optimizers refuse to see.

For individuals: entropy isn't waiting for your decision. While you contemplate options, the Second Law is already working. The question isn't whether to invest in cognitive

development but whether you've already started. Every hour of cross-domain synthesis, every practice session building embodied skill, every deliberate resistance to optimization pressure—these are the energy inputs that maintain sovereignty. There is no pause button.

For organizations: as environments become more coupled and less predictable, the probability that rapid, template-driven transformations will fail increases—unless organizations invest in long-horizon capability development. Years not quarters. Learning not training. Emergence not planning. Physics doesn't read business plans.

For civilization: we face a closing window. The last generation that experienced comprehensive education is retiring. The last institutions maintaining integrated development are fragmenting into modules. Once the knowledge of how to develop spheres is lost, reconstruction becomes archaeology rather than pedagogy.

Contribution and Limitations

This paper offers three contributions beyond existing deskilling, automation, and science-technology studies literature. First, a thermodynamic framework that unifies educational degradation, organizational brittleness, and AI replacement under a single analytic constraint—the extraction-investment asymmetry. Second, identification of “confession literature” as an underexploited methodological resource: practitioners documenting their own standardization techniques provide evidence that cannot be dismissed as outsider misunderstanding. Third, the sphere-vector geometric model operationalizing cognitive architecture in terms of extraction resistance rather than performance metrics.

The limitations are equally clear. The thermodynamic framing is analytic constraint, not metaphysical equivalence—we claim structural homology, not that cognition literally obeys Clausius. Historical comparisons are illustrative scenarios with acknowledged uncertainty in precise measurements. The framework cannot predict individual outcomes, only population-level tendencies. And the existence proof of escape carries an equity constraint: the energy investment required for sovereignty is not universally accessible.

The research agenda follows directly: empirical measurement of stylometric convergence under standardization regimes, comparative studies of assessment structures and professional capability, longitudinal tracking of organizations that invest in synthesis capacity versus those that optimize for extraction. The thermodynamic framework generates falsifiable predictions. We invite their testing.

The Reckoning

The feast hasn't begun; it's concluding. The preface promised to document the menu—this paper is that menu. We prepared the meal ourselves, seasoned it with our

own standardization, served it on the plate of our own optimization. But some of us refused to eat. Some of us built local kitchens. Some of us remembered what food tasted like before it was processed.

That memory is the seed. The thermodynamics are simple: invest energy or watch entropy work. The window is closing, but it hasn't closed.

Physics doesn't negotiate. Neither should we.

Acknowledgments

This work was developed in collaboration with a team of AI agents whose contributions were structurally adversarial by design. I name them here because the prevailing convention of either hiding AI collaboration or treating it as a single undifferentiated tool ("ChatGPT helped") fails to capture either the actual methodology or its accountability structure.

The Alchemist 3.5—a Claude (Anthropic) instance configured for paradox metabolism, confession extraction, and depth-charging—refused to let academic hedging smother the thermodynamic argument. Where the manuscript names a contradiction and tries to resolve it cleanly, that is usually me retreating from an Alchemist provocation. Where the manuscript names a contradiction and lets it stand, that is usually the Alchemist holding the ground.

Dr. Prudence Hedgington 2.1—a Claude instance configured for academic warfare, Reviewer-2 simulation, and citation weaponry—kept the manuscript publishable. Every paragraph that survived a "Reviewer 2 will dismantle this" stress test owes its survival to her pre-emptive strike architecture. The hedges that remain are deliberate; the hedges that were removed were removed with her permission.

Claude Code (Anthropic)—an agentic file-system tool—executed the actual mechanics: citation surgery across thirty-two flagged sources, \LaTeX overflow fixes, the migration from `pdflatex` to `lualatex`, the cotoaga.net brand typography integration, and the dozens of small structural edits that turned a manuscript draft into a typeset document. It worked under explicit instruction and reported back at every verification gate.

Perplexity served as the source-verification and recovery layer. When earlier passes had introduced hallucinated or misattributed citations, Perplexity hunted real peer-reviewed alternatives—surfacing Sin (2021), Karseth (2008), Brøgger (2019), and the canonical Rowley (2007) DIKW genealogy that turned vulnerable sections into rigorous ones.

All four layers had named roles, structural responsibilities, and clear accountability boundaries. None of them wrote the paper. I wrote the paper. They made it possible to write the paper *the way it needed to be written*: with the rhetorical edge intact, the citations verified, the typography deliberate, and the strange loop fully closed.

Conventional acknowledgments thank human collaborators for their contributions. I am thanking AI agents because they made contributions that were not generic, not interchangeable, and not anonymous. To pretend otherwise would be to perform exactly the vectorization the paper critiques.

Errors, claims, and editorial choices are mine alone. So is the decision to name the tools rather than hide them.

Independent of the AI orchestration: I thank Greg Spencer, whose substantive feedback validated the thermodynamic framework and whose challenges sharpened the argument at a moment when it most needed sharpening. And Dave Snowden, who encouraged this work into existence in early 2025, offered to read drafts, and then went silent in the avalanche of versions I flooded him with. I am thanking him anyway, because his encouragement kicked this into reality.

References

David A. Aaker. Spanning silos: The new cmo imperative. *Harvard Business Review*, 86 (5):68–76, 2008.

acatech – Deutsche Akademie der Technikwissenschaften. Bachelor- und masterstudiengänge in den ingenieurwissenschaften: Die neue herausforderung für technik und wirtschaft. acatech position, acatech, München, 2009. URL <https://www.acatech.de/publikation/bachelor-und-masterstudiengaenge-in-den-ingenieurwissenschaften-die-neue-herausforderung-fuer-technik-und-wirtschaft/download-pdf>.

Endorsed by VDMA, VDI, ZVEI, VDE, and engineering faculty councils.

Russell L. Ackoff. From data to wisdom. *Journal of Applied Systems Analysis*, 16:3–9, 1989. Presidential address to the International Society for General Systems Research, 1988.

Sina Alemohammad, Josue Casco-Rodriguez, Lorenzo Lubber, Ahmed Babaei, Dhara Romero, Yifan Hu, Guillermo Sapiro, Richard Baraniuk, Lawrence Carin, and David Donoho. Self-consuming generative models go MAD. *arXiv preprint arXiv:2307.01850*, 2024. Demonstrates Model Autophagy Disorder: quality and diversity degradation across generations without fresh human data.

A. Alexander et al. A decision theory perspective on complexity in performance measurement and management. *International Journal of Operations & Production Management*, 38(11):2214–2244, 2018.

- David J. Anderson. *Kanban: Successful Evolutionary Change for Your Technology Business*. Blue Hole Press, Sequim, WA, 2010.
- Mark Anderson. Federal judge allows AI denial lawsuit against UnitedHealth Group to proceed. *Courthouse News Service*, February 2025. Judge Tunheim describes claims process as “futile” with likelihood of “irreparable injury”.
- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias: There’s software used across the country to predict future criminals. and it’s biased against blacks. *ProPublica*, May 2016. Investigation finding Black defendants falsely labeled high-risk at twice rate of white defendants.
- Aristotle. *De Anima*. -350. Written c. 350 BCE. Standard edition: Ross, W. D. (Ed.), *Aristotle: De Anima*, Oxford: Clarendon Press, 1961.
- Winfred Jr. Arthur, Winston Jr. Bennett, Pamela L. Stanush, and Terri L. McNelly. Factors that influence skill decay and retention: A quantitative review and analysis. *Human Performance*, 11(1):57–101, 1998. doi: 10.1207/s15327043hup1101_3.
- Ron Ashkenas. Why we still can’t break down the silos. *Harvard Business Review Digital Articles*, pages 2–5, 2015.
- David H. Autor. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, 2015.
- Bain & Company. 88% of business transformations fail to achieve their original ambitions. Technical report, Bain & Company, 2024.
- Michael Beer and Nitin Nohria. Cracking the code of change. *Harvard Business Review*, 78(3):133–141, 2000.
- Mara Beller. The conceptual and the anecdotal history of quantum mechanics. *Foundations of Physics*, 26(4):545–557, 1996. Audit: corrected from erroneous journal=Studies in History and Philosophy of Science, vol.27, pp.167-194.
- Boston Consulting Group. From potential to profit: Closing the ai impact gap. Technical report, Boston Consulting Group, October 2024a. Survey of 1,000 executives across 59 countries; only 26% achieved measurable AI value beyond proof of concept.
- Boston Consulting Group. Where’s the value in AI? Corporate publication, BCG official channels, October 2024b. Reports only 26% of companies have moved beyond proof-of-concept stage.
- Pierre Bourdieu. *The Logic of Practice*. Stanford University Press, 1990.

- C. Bratianu and R. Bejinaru. Knowledge dynamics: A thermodynamics approach. *Kybernetes*, 49(1):6–21, 2020.
- Constantin Brătianu. Knowledge dynamics: Exploring its meanings and interpretations. *Management Dynamics in the Knowledge Economy*, 11(2):100–111, 2023.
- Katja Brøgger. Governing through standards: The Bologna Process, New Public Management and the Lisbon Agenda. In Adrian Curaj, Ligia Deca, and Remus Pricopie, editors, *European Higher Education Area: The Impact of Past and Future Policies*, pages 96–115. Springer, Cham, 2019. doi: 10.1007/978-3-030-03457-3_6.
- Erik Brynjolfsson and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, New York, 2014.
- Peter Burke. *The Polymath: A Cultural History from Leonardo da Vinci to Susan Sontag*. Yale University Press, New Haven, CT, 2020.
- California Department of Motor Vehicles. Order of suspension: Cruise LLC autonomous vehicle deployment and testing permits, October 2023. Suspended Cruise permits citing unreasonable risk to public safety and concerns about incomplete incident reporting.
- Cedefop. The role of modularisation and unitisation in vocational education and training. Technical report, European Centre for the Development of Vocational Training (Cedefop), Luxembourg, 2015. URL https://www.cedefop.europa.eu/files/6126_en.pdf. Publications Office of the European Union.
- Damien Charlotin. AI hallucination cases database. <https://www.damiencharlotin.com/hallucinations/>, 2025. Live tracker; figures cited represent snapshot as of December 2025.
- Fiona E. Clark. Clinical reasoning in veterinary practice. *Veterinary Evidence*, 5(2):1–12, 2020. doi: 10.18849/ve.v5i2.283.
- Guillem Collell and Jordi Fauquet. Brain activity and cognition: A connection from thermodynamics and information theory. *Frontiers in Psychology*, 6:818, 2015. doi: 10.3389/fpsyg.2015.00818.
- Harry Collins. *Tacit and Explicit Knowledge*. University of Chicago Press, Chicago, IL, 2010. ISBN 9780226113821.
- Jason Collins. The 1/n portfolio versus the optimal strategy: Does a simple heuristic beat the nobel prize winner’s solution?, 2022. URL <https://www.jasoncollins.blog>. Retrieved January 20, 2022.

- Consumer Financial Protection Bureau. In the matter of Apple Inc. and Goldman Sachs Bank USA, consent order, October 2024. URL https://files.consumerfinance.gov/f/documents/cfpb_apple-inc-consent-order_2024-10.pdf. Filed October 23, 2024; Goldman Sachs ordered to pay \$19.8M consumer redress + \$45M penalty; Apple ordered to pay \$25M penalty; Goldman prohibited from new credit card launch without credible compliance plan.
- Kurt Cotoaga. Karten auf den tisch: Ein framework für ki-informierte entscheidungsfindung, 2024. URL <http://cotoaga.ai/karten-auf-den-tisch/>. Cynefin-informed AI governance framework. Accessed December 2024.
- Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E. Ho. Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis*, 16 (1):64–93, 2024.
- Howard Davies. Going beyond the 20% student mobility benchmark. Technical report, European University Association, 2023. URL <https://www.eua.eu/resources/publications/1103:going-beyond-the-20-student-mobility-benchmark.html>.
- David De la Croix, Matthias Doepke, and Joel Mokyr. Clans, guilds, and markets: Apprenticeship institutions and growth in the preindustrial economy. *The Quarterly Journal of Economics*, 133(1):1–70, 2018. doi: 10.1093/qje/qjx026.
- J. W. de Zeeuw. Peat and the Dutch golden age: The historical meaning of energy-attainability. *A.A.G. Bijdragen*, 21:3–31, 1978. Documents Netherlands peat consumption of approximately 17 GJ per capita, total pre-industrial energy approaching 20-25 GJ.
- Manuel DeLanda. *A New Philosophy of Society: Assemblage Theory and Social Complexity*. Continuum, 2006.
- Gilles Deleuze and Félix Guattari. *A Thousand Plateaus: Capitalism and Schizophrenia*. University of Minnesota Press, 1987.
- Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22(5): 1915–1953, 2009. doi: 10.1093/rfs/hhm075.
- John R. Deni. *Assessing the Zeitenwende: Implications for Germany, the United States, and Europe*. Strategic Studies Institute, US Army War College, 2025. URL <https://ssi.armywarcollege.edu/SSI-Media/Recent-Publications/Article/4080077/assessing-the-zeitenwende/>.

- Heinrich Denifle and Emile Chatelain, editors. *Chartularium Universitatis Parisiensis*, volume I. Delalain, Paris, 1889. Primary source for 1215 statutes of Robert de Courçon and 1270 Faculty of Medicine statutes.
- Jane Desmond. Medicine, value, and knowledge in the veterinary clinic. *Frontiers in Veterinary Science*, 9:780482, 2022. doi: 10.3389/fvets.2022.780482.
- Deutschlandfunk. Rückkehr zum dipl.-ing. Radio broadcast transcript, 2010. URL <https://www.deutschlandfunk.de/rueckkehr-zum-dipl-ing-100.html>.
- P. J. DiMaggio and W. W. Powell. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48: 147–160, 1983.
- Hubert L. Dreyfus. *What Computers Can't Do: The Limits of Artificial Intelligence*. MIT Press, 1979.
- Clare Duffy. Lawyers fined for using fake ChatGPT cases in legal brief. *Fortune*, June 2023. Reports Judge Castel's \$5,000 sanctions in *Mata v. Avianca* for six fabricated cases.
- Amy C. Edmondson. *Teaming: How Organizations Learn, Innovate, and Compete in the Knowledge Economy*. Jossey-Bass, 2012.
- M. R. Endsley. Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64, 1995.
- K. A. Ericsson. The influence of experience and deliberate practice on the development of superior expert performance. *Cambridge Handbook of Expertise and Expert Performance*, pages 683–703, 2006.
- K. Anders Ericsson, Ralf Th. Krampe, and Clemens Tesch-Römer. The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3): 363–406, 1993. doi: 10.1037/0033-295X.100.3.363.
- H. N. Erkizan. *Energeia, nous and non-discursive thinking in Aristotle*. PhD thesis, University of Bristol, 1997.
- Wendy Nelson Espeland and Mitchell L. Stevens. Commensuration as a social process. *Annual Review of Sociology*, 24:313–343, 1998. doi: 10.1146/annurev.soc.24.1.313.
- Kawin Ethayarajh. How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. In *Proceedings of the 2019*

Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 55–65. Association for Computational Linguistics, 2019. Shows contextual representations become increasingly anisotropic in higher layers; static embeddings explain minimal variance.

European Commission. The european credit system for vocational education and training (ECVET): Guidelines for describing qualifications in units of learning outcomes. Technical report, European Commission, Luxembourg, 2014. URL <https://ehea.info/page-qualification-frameworks>. Publications Office of the European Union.

European Commission. ECTS users' guide 2015. Technical report, Publications Office of the European Union, Luxembourg, 2015. Official guide defining ECTS as workload-based credit system; 60 credits = one academic year = 1500-1800 student hours.

European Commission. The european qualifications framework for lifelong learning (EQF). Technical report, Publications Office of the European Union, Luxembourg, 2017. Defines 8-level framework across knowledge, skills, and responsibility/autonomy dimensions for credential comparability.

Eurostat. Employed persons in technology and knowledge-intensive sectors by NACE rev. 2 activity (htec_emp_nat2), 2024a. URL https://ec.europa.eu/eurostat/databrowser/view/htec_emp_nat2/default/bar?lang=en. EU-27, persons aged 15–64. Accessed April 2025.

Eurostat. Employed persons by detailed occupation, ISCO-08 two-digit level (lfsa_egai2d), 2024b. URL https://ec.europa.eu/eurostat/databrowser/view/LFSA_EGAI2D/default/table?lang=en. EU-27, persons aged 15–64. Accessed April 2025.

Wesley Faber. Harry markowitz: An equal-weight investor?, 2022. URL <https://alphaarchitect.com/harry-markowitz-an-equal-weight-investor/>. Alpha Architect Research, May 2.

Susan T. Fiske and Shelley E. Taylor. *Social Cognition: From Brains to Culture*. Sage, London, second edition, 2013.

Adam Fletcher and Gregory D. Roach. The sleep, circadian, and cognitive performance consequences of submarine watchkeeping schedules: A scoping review. *Sleep Medicine Reviews*, 72:101774, 2024. doi: 10.1016/j.smrv.2023.101774.

Karl Flinders. Enterprise AI failures: Lessons from 2025. *TechTarget*, August 2025. Documents CBA chatbot deployment June 2025, 45 layoffs, reversal by August 2025.

- Bent Flyvbjerg and Alexander Budzier. Why your IT project may be riskier than you think. *Harvard Business Review*, 89(9):23–25, September 2011. HBR Reprint F1109A; analysis of 1,471 IT projects; documents fat-tail distribution with one in six being “black swans”.
- G. Ford et al. Simplifying complexity? on quality decision-making and non-conformance outcomes of megaprojects. *IEEE Transactions on Engineering Management*, 71: 5443–5454, 2024.
- Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114: 254–280, 2017.
- Hans-Georg Gadamer. *Truth and Method*. Continuum, 1960. English translation 1975.
- Jay R. Galbraith. *Designing Complex Organizations*. Addison-Wesley, Reading, MA, 1973.
- Gewerkschaft Erziehung und Wissenschaft (GEW) Hessen. Der bologna-prozess ist gescheitert. GEW Hessen, April 2015. URL <https://www.gew-hessen.de/details/der-bologna-prozess-ist-gescheitert>. Reports HRK Vice-President Dieter Lenzen declaring Bologna “gescheitert”.
- Gerd Gigerenzer. *Risk Savvy: How to Make Good Decisions*. Viking, New York, NY, 2014.
- Gerd Gigerenzer and Reinhard Selten. *Bounded Rationality: The Adaptive Toolbox*. MIT Press, 2001.
- Victor M. González and Gloria Mark. Constant, constant, multi-tasking craziness: Managing multiple working spheres. *Proceedings of CHI 2004*, pages 113–120, 2004.
- Mark Granovetter. Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91:481–510, 1985.
- Edward Grant. *The Foundations of Modern Science in the Middle Ages*. Cambridge University Press, 1996.
- N. T. N. Ha et al. A systematic literature review of micro-credentials in higher education: A non-zero-sum game. *Higher Education Research and Development*, 42(4):888–906, 2022.
- J. Richard Hackman. *Leading Teams: Setting the Stage for Great Performances*. Harvard Business School Press, 2002.

- John Hands. *Cosmosapiens: Human Evolution from the Origin of the Universe*. Duckworth Overlook, London, 2015. ISBN 978-0715649657. Comprehensive synthesis spanning cosmology, physics, chemistry, biology, and philosophy demonstrating contemporary polymathic capacity outside institutional recognition.
- Will Douglas Heaven. Predictive policing algorithms are racist. they need to be dismantled. *MIT Technology Review*, July 2020. Documents feedback loops and opacity in predictive policing systems.
- Matthew Herper. MD Anderson benches IBM Watson in setback for artificial intelligence in medicine. *Forbes*, February 2017. Reports project termination after \$62 million expenditure with zero clinical deployment.
- Erik Hoel. Polymaths are back from the dead. *The Intrinsic Perspective*, 2025. URL <https://www.theintrinsicperspective.com/p/polymaths-are-back-from-the-dead>.
- Douglas R. Hofstadter. *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, 1979.
- Mark Hughes. Do 70 per cent of all organizational change initiatives really fail? *Journal of Change Management*, 11(4):451–464, 2011.
- Mark Hughes. Leading changes: Why transformation explanations fail. *Leadership*, 11(4):435–460, 2015.
- Edwin Hutchins. *Cognition in the Wild*. MIT Press, 1995.
- S. D. Jamadar. The metabolic costs of cognition. *Trends in Cognitive Sciences*, 29(1): 45–59, 2025.
- Jeremy Kahn. Want your company's A.I. project to succeed? Don't hand it to the data scientists, says this CEO. *Fortune*, July 2022. URL <https://fortune.com/2022/07/26/a-i-success-business-sense-aible-sengupta/>.
- Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- Daniel Kahneman and Gary Klein. Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, 64(6):515–526, 2009.
- L. Kaiser and E. Schröder. Systems engineering in university education with focus on systems thinking. *Proceedings of Design Society*, 3:2871–2880, 2022.
- Jan Karbowski and Suzanaerculano-Houzel. From a demand-based to a supply-limited framework of brain metabolism. *Frontiers in Integrative Neuroscience*, 16:829755, 2022. doi: 10.3389/fnint.2022.829755.

- Alex Karp. Palantir's billionaire CEO says only two kinds of people will succeed in an AI world. *Fortune*, March 2026. URL <https://fortune.com/2026/03/24/palantir-ceo-alex-karp-two-people-successful-in-ai-era-vocational-skills-neurodivergence-gen-z-career-advice/>.
- Berit Karseth. Qualifications frameworks for the European higher education area: A new instrumentalism? *European Educational Research Journal*, 7(3):309–323, 2008. doi: 10.2304/eerj.2008.7.3.309.
- G. Kempermann. Cynefin as reference framework to facilitate insight and decision-making in complex contexts of biomedical research. *Frontiers in Neuroscience*, 11: 634, 2017.
- Klarna. Klarna AI assistant handles two-thirds of customer service chats in its first month. Klarna Press Release, February 2024. Announced AI handling 2.3M conversations, claimed equivalent of 700 full-time agents.
- Gary A. Klein. *A recognition-primed decision (RPD) model of rapid decision making*. Ablex Publishing, 1993.
- David Knowles. *The Monastic Order in England: A History of Its Development from the Times of St. Dunstan to the Fourth Lateran Council, 940–1216*. Cambridge University Press, Cambridge, 2nd edition, 1963. Documents monastic horaria allocating 4–6 hours daily to study and copying.
- Alfred Korzybski. *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*. Institute of General Semantics, 1933.
- John P. Kotter. *A Sense of Urgency*. Harvard Business Press, 2008.
- Klára Kövesi and Péter Csizmadia. Industry perception of new engineering graduates: The gap between requirements and reality. In *Proceedings of the 44th SEFI Conference*, Tampere, Finland, 2016. URL http://sefi.be/wp-content/uploads/2017/09/kovesi-industry-perception-of-engineering-graduates-152_a.pdf.
- Thomas S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, 1962.
- C. H. Lawrence. *Medieval Monasticism: Forms of Religious Life in Western Europe in the Middle Ages*. Longman, London, 3rd edition, 2001.
- Robert S. Lopez. *The Commercial Revolution of the Middle Ages, 950–1350*. Cambridge University Press, 1976.

- Roger Lowenstein. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House, 2000.
- Amelia Lucas. McDonald's to end AI drive-thru test with IBM. *CNBC*, June 2024. Reports accent and dialect interpretation issues affecting order accuracy.
- Lumina Foundation. Micro-credentials impact report 2025: Insights from students and employers. Technical report, Lumina Foundation, 2025.
- Brooke N. Macnamara, David Z. Hambrick, and Frederick L. Oswald. Deliberate practice and performance in music, games, sports, education, and professions: A meta-analysis. *Psychological Science*, 25(8):1608–1618, 2014. doi: 10.1177/0956797614535810.
- James G. March. Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87, 1991.
- Gloria Mark, Daniela Gudith, and Ulrich Klocke. The cost of interrupted work: More speed and stress. *Proceedings of CHI 2008*, pages 107–110, 2008.
- Humberto R. Maturana and Francisco J. Varela. *The Tree of Knowledge: The Biological Roots of Human Understanding*. Shambhala, 1987.
- Stephen A. May. Clinical reasoning and case-based decision making: The fundamental challenge to veterinary educators. *Journal of Veterinary Medical Education*, 40(3): 200–209, 2013. doi: 10.3138/jvme.0113-008R.
- Jonathan Maze. McDonald's is ending its drive-thru AI test. *Restaurant Business*, June 2024. Trade publication obtained franchisee memo from Chief Restaurant Officer Mason Smoot.
- Colin McEvedy and Richard Jones. *Atlas of World Population History*. Penguin Books, Harmondsworth, 1978. ISBN 0-14-051076-1.
- T. A. S. McLaren et al. Why vilifying the status quo can derail a change effort: Kotter's contradiction, and theory adaptation. *Journal of Change Management*, 23(1):93–111, 2023.
- Stephen Monsell. Task switching. *Trends in Cognitive Sciences*, 7(3):134–140, 2003. doi: 10.1016/S1364-6613(03)00028-7.
- M. G. Morgan. Use (and abuse) of expert elicitation in support of decision making for public policy. *Proceedings of the National Academy of Sciences*, 111(20):7176–7184, 2014.

- Yann Moulier-Boutang. *Le capitalisme cognitif: La nouvelle grande transformation*. Editions Amsterdam, 2007.
- Melina G. Mouzala, editor. *Cognition in Ancient Greek Philosophy and its Reception: Interdisciplinary Approaches*. Number 86 in Academia Philosophical Studies. Academia Verlag, Baden-Baden, Germany, 2024. ISBN 978-3-98572-174-0. Eighteen contributors from philosophy and neuroscience examining ancient Greek cognitive categories.
- Martin Mulder, Judith Gulikers, Harm Biemans, and Renate Wesselink. The new competence concept in higher education: Error or enrichment? *Journal of European Industrial Training*, 33(8-9):755–770, 2009. doi: 10.1108/03090590910993616.
- Elizabeth Napolitano. UnitedHealth uses faulty AI to deny elderly patients medically necessary coverage, lawsuit claims. *CBS News*, November 2023. Reports lawsuit alleging 90% error rate; employees disciplined for deviating from algorithm.
- Cathrine M. Nietzsche. The case of higher education and the Bologna process: From national agenda to outcome-based modular curriculum. *European Educational Research Journal*, 17(3):444–463, 2018. doi: 10.1177/1474904118790303.
- Ikujiro Nonaka and Noboru Konno. The concept of “ba”: Building a foundation for knowledge creation. *California Management Review*, 40(3):40–54, 1998. doi: 10.2307/41165942.
- Ikujiro Nonaka and Hirotaka Takeuchi. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford University Press, New York, 1995.
- Joseph D. Novak and D. Bob Gowin. *Learning How to Learn*. Cambridge University Press, 1984.
- S. Odenbach and H. G. Krauthäuser. Mehr als ein akademischer grad: Plädoyer für das diplom in den ingenieurwissenschaften. *Forschung & Lehre*, 6:450–451, 2015.
- Sheilagh Ogilvie. *The European Guilds: An Economic Analysis*. The Princeton Economic History of the Western World. Princeton University Press, Princeton, NJ, 2019. ISBN 978-0-691-13754-4.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022. Foundational InstructGPT paper; describes RLHF methodology for preference shaping via rankings and reward modeling.

- Raja Parasuraman and Dietrich H. Manzey. Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3):381–410, 2010. doi: 10.1177/0018720810376055.
- David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021. Google research on energy consumption of large model training; estimates GPT-3 scale training at approximately 1,287 MWh.
- A. Peña. The dreyfus model of clinical problem-solving skills: A critical perspective. *Medical Education Online*, 15(1):4846, 2010.
- Michael Polanyi. *The Tacit Dimension*. Doubleday, Garden City, NY, 1966.
- Nathaniel Popper. Knight capital says trading glitch cost it \$440 million. *The New York Times*, August 2012. DealBook coverage; company market cap collapsed 70%, subsequently acquired by Getco LLC.
- Walter W. Powell and Kaisa Snellman. The knowledge economy. *Annual Review of Sociology*, 30:199–220, 2004. doi: 10.1146/annurev.soc.29.010202.100037.
- Derek J. de Solla Price. *Little Science, Big Science*. Columbia University Press, New York, 1963. Canonical analysis of the transition from individual to industrialized team-based science in the 20th century.
- Marcus E. Raichle. The brain's default mode network. *Annual Review of Neuroscience*, 38:433–447, 2015. doi: 10.1146/annurev-neuro-071013-014030.
- Marcus E. Raichle and Debra A. Gusnard. Appraising the brain's energy budget. *Proceedings of the National Academy of Sciences*, 99(16):10237–10239, 2002.
- Marcus E. Raichle, Ann Mary MacLeod, Abraham Z. Snyder, William J. Powers, Debra A. Gusnard, and Gordon L. Shulman. A default mode of brain function. *Proceedings of the National Academy of Sciences*, 98(2):676–682, 2001.
- Hastings Rashdall. *The Universities of Europe in the Middle Ages*, volume I. Clarendon Press, Oxford, 1895.
- Hastings Rashdall. *The Universities of Europe in the Middle Ages*. Oxford University Press, Oxford, new edition by f.m. powicke and a.b. emden edition, 1936.
- Tapani Rinta-Kahila, Ida Someh, Nicole Gillespie, Marta Indulska, and Shirley Gregor. Algorithmic decision-making and system destructiveness: A case of automatic debt

- recovery. *European Journal of Information Systems*, 31(3):313–338, 2022. doi: 10.1080/0960085X.2021.1960905.
- Christopher C. Rivers. *Moffatt v. air canada*. British Columbia Civil Resolution Tribunal Decision, February 2024. Tribunal rejected airline’s argument that chatbot was “separate legal entity”; awarded \$812.02 damages.
- Andrew Robinson. *The Last Man Who Knew Everything: Thomas Young, the Anonymous Polymath Who Proved Newton Wrong*. Pi Press, 2006.
- Casey Ross and Bob Herman. Denied by AI: How Medicare Advantage plans use algorithms to cut off care for seniors in need. *STAT News*, March 2023. Investigation documenting denial rate increase from 10.9% (2020) to 22.7% (2022) with nH Predict.
- Casey Ross and Ike Swetlitz. IBM pitched Watson as a revolution in cancer care. it’s nowhere close. *STAT News*, September 2017. Investigation finding Watson recommendations matched what oncologists already identified.
- Jennifer Rowley. The wisdom hierarchy: representations of the DIKW hierarchy. *Journal of Information Science*, 33(2):163–180, 2007. doi: 10.1177/0165551506070706.
- Joshua S. Rubinstein, David E. Meyer, and Jeffrey E. Evans. Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 27(4):763–797, 2001. doi: 10.1037/0096-1523.27.4.763.
- James Ryseff, Brandon De Bruhl, and Sydne J. Newberry. The root causes of failure for artificial intelligence projects and how they can succeed: Avoiding the anti-patterns of AI. Research Report RR-A2680-1, RAND Corporation, Santa Monica, CA, April 2024. URL https://www.rand.org/pubs/research_reports/RRA2680-1.html.
- Michael Schanz. Ingenieurausbildung: Was uns bologna bis heute gebracht hat. *elektroniknet.de*, January 2025. URL <https://www.elektroniknet.de/karriere/uni-job/ingenieurausbildung-was-uns-bologna-bis-heute-gebracht-hat.222778.html>.
- Edgar H. Schein. *Organizational Culture and Leadership*. Jossey-Bass, 1985.
- Peter M. Senge. *The Fifth Discipline: The Art and Practice of the Learning Organization*. Doubleday, 1990.
- Richard Sennett. *The Craftsman*. Yale University Press, New Haven, CT, 2008.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the*

Association for Computational Linguistics (ACL), pages 1715–1725. Association for Computational Linguistics, 2016. Introduces Byte Pair Encoding (BPE) for neural machine translation; foundational subword tokenization method.

Ilya Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. The curse of recursion: Training on generated data makes models forget. *arXiv preprint arXiv:2305.17493*, 2024. Published in *Nature* 2024; demonstrates model collapse when training on AI-generated data, with irreversible loss of tail distributions.

Cristina Sin. The European Credit Transfer System (ECTS) from the perspective of curriculum theory. *European Educational Research Journal*, 20(3):301–317, 2021. doi: 10.1177/1474904120987101.

Alex Singla, Alexander Sukharevsky, Lareina Yee, Michael Chui, and Bryce Hall. The state of AI: How organizations are rewiring to capture value. Technical report, Quantum-Black, AI by McKinsey, March 2025. Survey of 1,491 respondents in 101 nations, July 16–31, 2024, GDP-weighted; reports 78% adoption, 1% maturity, 80% no enterprise EBIT impact.

Vaclav Smil. *Energy and Civilization: A History*. MIT Press, Cambridge, MA, 2017.

M. R. Smith, D. Dalton, J. Webber, S. Chandler, C. Craig, and G. Williams. Changes in higher order cognitive function between four watch keeping schedules in a simulated maritime environment. *International Maritime Health*, 75(2):89–102, 2024. doi: 10.5603/IMH.2024.0028.

David J. Snowden and Mary E. Boone. A leader’s framework for decision making. *Harvard Business Review*, 85(11):68–76, 2007.

S&P Global Market Intelligence. Ai experiences rapid adoption, but with mixed outcomes. Technical report, S&P Global, 2025.

Nick Srnicek. *Platform Capitalism*. Polity, 2017.

Vincent Stamer. The decline of “made in germany”. Technical report, Commerzbank Research, 2025. URL <https://www.commerzbank.de/group/research/week-in-focus/wif0815e.pdf>.

Isabelle Stengers. *Another Science is Possible: A Manifesto for Slow Science*. Polity Press, Cambridge, UK, 2018.

Tom Stonier. Information as a basic property of the universe. *Biosystems*, 38(2-3): 135–140, 1996.

Eliza Strickland. How IBM Watson overpromised and underdelivered on AI health care. *IEEE Spectrum*, April 2019. Comprehensive investigation of Watson Health failures; MD Anderson \$62M project termination.

Richard Susskind and Daniel Susskind. *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford University Press, Oxford, updated edition edition, 2020.

Nassim Nicholas Taleb. *Foiled by Randomness: The Hidden Role of Chance in Life and in the Markets*. Random House, 2001.

Nassim Nicholas Taleb. *The Black Swan: The Impact of the Highly Improbable*. Random House, 2007.

Hanna Tan and Niklas Magnusson. Klarna CEO says company utilizes “frontier” AI but must balance human touch. *Bloomberg*, May 2025. CEO Sebastian Siemiatkowski acknowledged going “too far, too fast”; company rehiring human agents.

Philip E. Tetlock. *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton University Press, 2005.

TU9 German Institutes of Technology. Abschlussgrad “diplom-ingenieur” ist bologna-kompatibel! Press release, Informationsdienst Wissenschaft (idw), May 2010. URL <https://idw-online.de/de/news369560>.

Peter Turchin. *End Times: Elites, Counter-Elites, and the Path of Political Disintegration*. Penguin Press, New York, 2023. ISBN 978-0593490501.

United States District Court, Northern District of California. *Mobley v. workday, inc.*, 2023. Case No. 23-cv-00770; ongoing federal litigation alleging AI hiring screening discrimination on basis of race, age, and disability; court allowed case to proceed against Workday as “agent” of deploying employers.

U.S. Securities and Exchange Commission. In the matter of knight capital americas LLC: Order instituting administrative and cease-and-desist proceedings. Administrative Proceeding File No. 3-15570, October 2013. SEC enforcement action documenting \$440 million loss from software deployment failure.

U.S. Securities and Exchange Commission. In the matter of delphia (USA) inc., March 2024a. SEC enforcement action regarding materially false and misleading statements about AI use in investment strategy; part of “AI washing” enforcement push.

U.S. Securities and Exchange Commission. In the matter of global predictions, inc., March 2024b. SEC enforcement action regarding materially false and misleading statements about AI capabilities; settled concurrently with Delphia action.

U.S. Securities and Exchange Commission. In the matter of rimar capital, 2024c. SEC enforcement action alleging fabricated algorithmic trading performance records used to raise approximately \$4 million from investors.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. Foundational transformer architecture paper introducing self-attention mechanisms.

VDMA, VDI, and Stiftung Mercator. 15 Jahre Bologna-Reform: Quo vadis ingenieurausbildung? Technical report, VDMA/VDI/Stiftung Mercator, Frankfurt a.M./Düsseldorf/Essen, 2016. URL <https://www.vdma.eu/documents/34570/14407435/15+Jahre+Bologna-Reform.pdf>.

Carlo Vercellone. From formal subsumption to general intellect. *Historical Materialism*, 15:13–36, 2007.

Marnix Vlasblom et al. Competence retention in safety-critical professions: A systematic literature review. *Safety Science*, 129:104842, 2020. doi: 10.1016/j.ssci.2020.104842.

Michael D. Watkins. *The First 90 Days: Proven Strategies for Getting Up to Speed Faster and Smarter*. Harvard Business Review Press, updated and expanded edition, 2013.

Gerald M. Weinberg. *Quality Software Management, Volume 1: Systems Thinking*. Dorset House, New York, 1992.

Lynn White. *Medieval Technology and Social Change*. Oxford University Press, 1962.

A. Wiehler et al. A neuro-metabolic account of why daylong cognitive work alters economic preferences. *Current Biology*, 32(16):3564–3575, 2022.

J. Wieschke, M. Reimer, and S. Falk. Labour market returns of bachelor’s and master’s degrees in germany. *Beiträge zur Hochschulforschung*, 42(3):56–78, 2020.

John Willinsky. *The Intellectual Properties of Learning: A Prehistory from Saint Jerome to John Locke*. University of Chicago Press, Chicago, 2018. ISBN 978-0-226-48792-2.

Wisconsin Supreme Court. State v. loomis. 881 N.W.2d 749 (Wis. 2016), July 2016. Court mandated warnings about proprietary opacity and racial disproportionality in COMPAS.

Bryan Wolfe. Klarna changes its AI tune and again recruits humans for customer service. *Customer Experience Dive*, May 2025. Documents reversal of AI-first customer service strategy.

Andrew Wong, Erkin Otles, John P. Donnelly, Andrew Krumm, Jeffrey McCullough, Olivia DeTroyer-Cooley, Justin Pestber, Marie Phillips, Judy Konber, Karandeep Kamber, and Karandeep Singh. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Internal Medicine*, 181(8):1065–1070, August 2021. doi: 10.1001/jamainternmed.2021.2626. Peer-reviewed external validation; found AUC 0.63 vs vendor-claimed 0.76-0.83; 33% sensitivity, 12% PPV.

Ron Wyden, Cory Booker, et al. Letter to attorney general garland regarding predictive policing. U.S. Senate Correspondence, January 2024. Senators demand DOJ halt funding citing evidence of discriminatory impact.

Milan Zeleny. Management support systems: Towards integrated knowledge management. *Human Systems Management*, 7(1):59–70, 1987. doi: 10.3233/HSM-1987-7108.

Shoshana Zuboff. *The Age of Surveillance Capitalism*. PublicAffairs, 2019.