

Article

The Five-Year Horizon: Quantum Computing and the Philosophy of Technical Becoming

Christophe Jurczak 

Quantonation Ventures, Dallas, TX 75206, USA; christophe@quantonation.com

Abstract

Certain physics-based technologies (quantum computing, fusion energy, advanced materials, brain–computer interfaces) have remained “five years away” for decades. This paper argues that this perpetual horizon is not a forecasting failure but the temporal signature of *Perpetual Five-Year Technologies* (PFYTs): technologies of atoms, not bits, whose technical object and enabling ecosystem must co-develop. Drawing on philosophy of technology and sociology of scientific practice, and grounded in quantitative analysis of cross-platform performance data and high-impact research across hardware architectures, the paper shows that PFYT temporality is endogenous, produced by concretization dynamics and recursive constraint discovery rather than by market failures or insufficient funding. Quantum computing provides the paradigmatic case. Treating the quantum processor as a technical individual in active individuation explains why multiple computing paradigms persist and why each breakthrough resets rather than collapses the horizon. The convergence of machine learning and quantum hardware into pipelines for drug and materials discovery shows how Physical AI reshapes PFYT dynamics, compressing some cycles while introducing new forms of co-individuation between intelligence and matter. The framework generates diagnostics for physics-based frontiers where technical objects and milieus co-produce developmental time.

Keywords: perpetual five-year technologies; quantum computing; Simondon; concretization; technical individuation; physical AI; philosophy of technology

1. The Puzzle

Quantum computing has been “five years away” for more than thirty years. The phrase has become a punchline: shorthand for technological vaporware, for the gap between laboratory demonstration and practical deployment that seems never to close. But this framing misses something important. Each year, experts revise their forecasts; each revision reflects genuine progress. The target is not stationary; it recedes not because nothing happens but because each advance reveals new constraints previously invisible. The horizon moves because the landscape itself is being discovered.

A quantum computing researcher who says “fault tolerance is five years away” is not making a calendar prediction. She is describing a structure of engagement: the problems that must be solved, the dependencies that must align, the unknowns that remain unknown. The “five years” indexes a complexity horizon, not a timeline.

This reframing turns a forecasting question into a philosophical one. To ask when a technology will “arrive” presupposes that we already know what it is; if the identity of a quantum computer is still being settled as the machine develops, that presupposition fails. The five-year horizon is then the symptom of a prior question about the mode of existence



Academic Editors: Fabien Paillusson
and Rossella Lupacchini

Received: 26 April 2026

Revised: 22 June 2026

Accepted: 8 July 2026

Published: 20 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

of a technical object that has not yet fully individuated, about how an object, its associated milieu, and the criteria of its own completion come into being together. It is this problem of technical individuation, not any timeline, that the paper takes as its subject.

The perpetual five-year horizon, this paper contends, is a structural feature of how certain technologies produce their own temporality rather than a forecasting failure. I call these perpetual five-year technologies (PFYTs): physics-based technologies whose components cannot scale independently but must co-individuate with a developing associated milieu, so that each advance in one subsystem reconfigures what counts as progress in the others. The framework grew out of years spent observing how dozens of deep-tech companies progress¹, in leaps rather than lines, through patterns that existing innovation theory does not fully capture. In fact innovation economics offers its own accounts of prolonged development: path dependence and lock-in [1], increasing returns, and coordination failures across supply chains. These frameworks treat temporal delay as exogenous, the result of market failures, misaligned incentives, or insufficient coordination that could, in principle, be overcome by better policy or more capital. The philosophical framework developed here makes a different claim: PFYT temporality is endogenous, emerging from the structure of the technical object itself. The five-year horizon is not a coordination failure to be solved but the temporal signature of concretization, of components becoming mutually conditioning through iterative engagement with material resistance.

The vocabulary draws on Gilbert Simondon's theory of technical individuation [2,3]. Technical objects develop in relation to an associated milieu: the ensemble of subsystems, infrastructure, knowledge, and practices that condition and are conditioned by the object's functioning. Concretization names the process by which initially separate components become mutually conditioning. Individuation proceeds from a metastable pre-individual field charged with incompatible potentials that can resolve in more than one direction; the richer the field, the more individuations it can sustain. This distinction has practical force. A skeptical reader may invoke Manhattan or Apollo as cases where exogenous force collapsed development horizons. But those were engineering achievements within mature or rapidly mactable milieus; the underlying physics was already understood. Contrast superconducting qubits in the 1990s–2000s, where substantial investment failed to deliver because decoherence proved an individuation problem, not an engineering one. Exogenous force compresses engineering timelines, not concretization timelines. Conversely, some phenomena exhibit PFYT-like temporal signatures without the underlying dynamics. Cold fusion showed perpetual deferral without genuine concretization: no associated milieu accumulated, and a well-funded replication effort three decades later found no anomalous effects [4]. The LK-99 superconductivity claim followed the same pattern [5]. The pre-individual field of quantum computing is exceptionally fertile; that of cold fusion was barren, what Langmuir called pathological science [6].

The argument proceeds as follows. Section 2 presents quantum computing as the paradigmatic case and introduces the four PFYT criteria through which the argument operates. Section 3 examines the ontological status of quantum computing as a technical object. Section 4 analyzes the neutral-atom quantum processor as a case of active concretization. Section 5 examines how concretization proceeds through resistance and accommodation. Section 6 turns to physical AI, which reconfigures PFYT dynamics through compression, epistemic substrate, and co-individuation. Section 7 concludes. The paper's claims are limited to physics-based technologies undergoing active concretization; it does not attempt a general theory of innovation temporality or a comprehensive philosophy of AI.

2. Quantum Computing as Paradigmatic Case

A technology qualifies as PFYT when four conditions obtain. First, milieu-dependence: performance is bottlenecked by multiple external subsystems (lasers, vacuum, control electronics, theory, workforce), each with its own innovation curve. Second, concretization dynamics: progress requires integrating components so that improvements become mutually conditioning rather than modular. Third, recursive constraint discovery: each solved bottleneck reveals deeper ones, constraints systematically produced by scale rather than mere “unknown unknowns”. Fourth, high validation latency: iteration cycles are intrinsically slow because verification is expensive, fragile, or physically constrained.

These criteria distinguish PFYTs from related categories. General-purpose technologies diffuse across applications without requiring milieu reconstruction. Platform technologies with modular supply chains can scale components independently. Hype cycles exhibit temporal deferral without underlying concretization dynamics. PFYTs are none of these: they are technologies whose temporality emerges from the structure of their becoming. Quantum computing provides the paradigmatic case.

Quantum computing provides clear evidence for PFYT dynamics because its development is extensively documented across multiple competing platforms. The field offers a natural experiment: different hardware approaches, pursued by different research communities with different institutional supports, all attempting to solve the same fundamental problem. The patterns that emerge illuminate the temporal structure of physics-based innovation.

The empirical claims in Sections 2 and 3 rely on two datasets: quantum-hardware performance data drawn from the Quantum Error Correction Reports 2024 and 2025, and a bibliometric dataset constructed from OpenAlex records of high-impact quantum-computing papers, with hardware modalities classified by architecture and bibliometric assignments verified by manual review (full details in the Data Availability Statement). The figures are used diagnostically rather than predictively: they identify signatures of cross-platform concretization, epistemic momentum, and recursive constraint discovery rather than forecasting platform success.

Quantum computing exploits quantum mechanical phenomena (superposition and entanglement) to process information in ways that classical computers cannot efficiently simulate. But “quantum computing” is not a single technology. Three distinct computational paradigms coexist. Analog quantum simulation uses ensembles of interacting quantum systems to directly emulate physical phenomena (molecular dynamics, magnetic phases, optimization landscapes) without translating them into digital gate sequences. Digital uncorrected (or “noisy intermediate-scale”) quantum computing implements gate-based circuits on qubits with relatively simple encoding. Fault-tolerant quantum computing uses quantum error detection and correction to sustain arbitrarily long computations, but requires error rates below demanding thresholds [7]. These are different ways to control a quantum system, operating at different levels of abstraction from the underlying physics.

2.1. Cross-Platform Individuation in Quantum Computing

Several parameters characterize the performance of a quantum computing platform: qubit count, connectivity, gate speed, coherence time, mid-circuit measurement fidelity, and classical control bandwidth, among others. Of these, qubit number and two-qubit gate error rate are among the most systematically tracked, and for good reason: qubit count sets the ceiling on problem size and encoding redundancy available for error correction, while two-qubit gate fidelity determines whether entanglement, the operational resource that distinguishes quantum from classical computation, can be generated reliably enough to sustain useful circuits in the digital computing framework. Together, these two metrics

define the minimal envelope within which the benefits of using a quantum computer must be demonstrated, making them natural axes for cross-platform comparison.

Consider first the convergence of error rates across platforms (Figure 1a). Six distinct hardware approaches based on different elementary systems (trapped ions; superconducting circuits, both fixed-frequency and tuneable; neutral atoms; silicon spins; photonic; and cat qubits, more generally, microwave bosonic codes) began from radically different starting points. Trapped ions, building on decades of atomic physics, achieved early leads in gate fidelity [8]. Superconducting circuits, leveraging semiconductor fabrication expertise, offered faster gates and manufacturing scalability [9]. Neutral atoms, emerging more recently from laser cooling and optical tweezer technologies, promised natural scalability through identical atoms [10]. Silicon spins offered compatibility with existing semiconductor infrastructure [11], with recent demonstrations achieving above 99% two-qubit gate fidelity in foundry-manufactured devices [12]. Photonic platforms exploited the intrinsic noise resilience of optical photons for modular and networked architectures [13]. Bosonic codes took a fundamentally different approach, encoding logical information in the large Hilbert space of superconducting cavities rather than in two-level systems, enabling hardware-efficient error correction at the single physical unit level [14].




Figure 1. Multi-dimensional development of quantum computing platforms. (a) Two-qubit gate error rates across hardware modalities (trapped ions, superconducting circuits (fixed-frequency and tuneable), neutral atoms, silicon spins, photonic, and bosonic codes), showing converging trajectories toward the practical quantum error correction threshold (10^{-3}). Later platforms descend faster than earlier ones, inheriting accumulated ecosystem knowledge in error correction theory, control electronics, and fabrication methods. (b) Qubit count across the same platforms, showing divergent scaling trajectories. Taken together, the two panels illustrate why crossing the error-correction threshold does not collapse the five-year horizon: platforms that converge on fidelity diverge on scale, and future development requires simultaneous optimization across mutually conditioning parameters. Dashed lines indicate company roadmap projections. Reproduced/adapted from the Riverlane Quantum Error Correction Reports [15,16], with permission from Riverlane, 2026. Data courtesy of Eve Townsend, Joan Campus, Setiawan Wending, Nick Johnson, and Gianluca Aiello (Riverlane).

What the data reveal is striking: despite different starting points and different physical mechanisms, the platforms show converging trajectories toward common error-correction thresholds. More significantly, later platforms descend faster than earlier ones. Neutral atoms and silicon spins, arriving later to the field, have achieved in years what trapped ions and superconducting circuits required decades to accomplish. This acceleration reflects more than improved engineering; it marks the maturation of a shared ecosystem. Later platforms inherit accumulated knowledge in error correction theory, control electronics design, fabrication methods, and measurement techniques. The ecosystem matures across platforms, not within them alone. Recent breakthroughs confirm this pattern: Google's demonstration of performance consistent with below-threshold quantum error correc-

tion [17], Harvard's reported fault-tolerant neutral-atom architecture with 448 qubits [18], IBM's high-threshold code [19], and, most recently, the first experimental validations of quantum simulation of materials against independent laboratory measurements, achieved independently on a neutral-atom platform [20] and a superconducting processor [21], all build on decades of cross-platform knowledge accumulation.

But crossing these thresholds does not collapse the five-year horizon; it transforms it. Error correction below threshold is a necessary condition, not a sufficient one. The next generation of machines must advance not by pushing a single metric further but by concretizing along the full dimensionality of the performance space: qubit count, connectivity, gate speed, mid-circuit measurement fidelity, classical control bandwidth, and integration with error-decoding hardware, simultaneously, because at this level of maturity, these parameters become mutually conditioning. This multi-dimensional character is already visible: neutral atoms and silicon spins offer routes to large qubit counts through identical particles and semiconductor fabrication respectively; photonic architectures and certain superconducting designs achieve faster gate speeds; trapped ions provide all-to-all connectivity that eliminates routing overhead. The convergence in Figure 1a reveals ecosystem maturation (the cross-platform transfer of knowledge in error correction, control, and metrology) while the divergence in Figure 1b confirms that no single platform dominates across all dimensions simultaneously. From the PFYT perspective, this is the point: the technology has not individuated into a single dominant form. Multiple concretization paths remain open, each optimizing the relationship between technical object and associated milieu along different axes. The eventual structure of quantum computing may not be a single winning platform but a differentiated ecology of technical individuals, each concretized for different computational regimes.

2.2. Distribution of High-Impact Research Across Platforms

A second pattern emerges from tracking the distribution of high-impact research across platforms over time (Figure 2). This measures something different from technical performance: it reveals where the field's intellectual resources and attention are flowing, what might be called *epistemic momentum*.

The "trapped ion paradox" is particularly instructive. Ions achieve the best gate fidelities of any platform, yet their share of high-impact research has declined markedly, from approximately 15–19% in 2011–2012 to an average below 6% after 2017, with occasional upticks around specific experimental milestones that do not reverse the structural trend. This divergence between technical maturity and epistemic energy suggests a platform approaching its concretization limits: highly refined, capable of excellent performance, but with diminishing room for the kind of fundamental advances that attract sustained research attention. The platform has individuated successfully but is no longer the site of active becoming.

Neutral atoms present the inverse pattern: rapid technical improvement accompanied by surging research momentum, overtaking superconducting circuits as the dominant research focus by 2022–2023. This indicates a platform in active development, one where both performance gains and conceptual innovations remain available, drawing researchers toward unresolved problems and unexplored possibilities.

Together, these patterns reveal that PFYT development involves more than technical progress: the reorganization of entire fields. Resources, attention, and expertise flow toward platforms in active development; mature platforms, however capable, recede from the frontier. The field as a whole advances through this shifting allocation of intellectual and material resources, a dynamic that cannot be understood by examining any single platform in isolation.

The PFYT framework generates testable predictions. If ecosystem maturation is cross-platform, later entrants should descend faster on error-rate curves by inheriting accumulated knowledge, but this acceleration should *not* occur for platforms requiring entirely novel supply chains. The “epistemic momentum” pattern predicts that capital and labor will migrate toward platforms in active individuation *before* technical superiority is demonstrated, a leading indicator, not a lagging one. The trapped-ion paradox is consistent with this: declining research share despite best-in-class fidelity signals a platform approaching concretization limits. These are not *post hoc* interpretations but diagnostics that can be applied prospectively to emerging platforms.

The recursive structure of this development explains the perpetual horizon. Each breakthrough closes one bottleneck while revealing deeper ones that were previously invisible. Coherence times improve; crosstalk at scale becomes limiting. Error correction advances; qubit leakage reshapes roadmaps [22]. Gate fidelities increase; control system complexity emerges as the constraint. The five-year horizon resets not because nothing changes but because each advance exposes the next layer of constraint. This is not a failure of planning; it is the structure of exploring physical possibility space.




Figure 2. Epistemic momentum across qubit platforms. The distribution of high-impact research papers reveals shifting intellectual investment. Ions achieve the best gate fidelities yet show declining research share. Neutral atoms show the inverse, overtaking superconducting circuits in recent years. The underlying dataset was compiled from OpenAlex, filtering for the 100 most-cited quantum computing papers per year (2010–2025), classified by hardware architecture using a weighted semantic scoring algorithm applied to titles and abstracts; modality classifications were verified by manual review. Trend lines are LOWESS-smoothed (bandwidth 0.3) to reduce sampling noise inherent in annual counts of 24–52 classified papers. Data courtesy of Brooke Abeles (private communication, February 2026).

2.3. The Evidential Signature of PFYT Dynamics

The empirical case for PFYT dynamics rests on a conjunction of patterns rather than on any single figure, and the figures map onto the four criteria. Figure 1 establishes milieu-dependence and concretization: later platforms descend the error-rate curve faster by inheriting a shared ecosystem (Figure 1a), while diverging persistently on scale rather than converging on one architecture (Figure 1b). Figure 2 establishes recursive constraint discovery in the form that bears on this section: research momentum decouples from technical maturity, redistributing toward platforms with more open individuation rather than tracking current performance. The Simondonian vocabulary, the neutral-atom case

study, and the historical contrasts with Manhattan, Apollo, and cold fusion are interpretive and boundary-marking rather than primary evidence. Several rival explanations can account for individual patterns; none accounts for their conjunction, the pairing of accelerating progress with a horizon that does not close. Setting the framework against those alternatives shows what it explains that they do not.

The central claim is that the horizon is endogenous. “Five years away” does not name a fixed set of benchmarks toward which the field converges; it names a benchmark set whose membership the concretization process itself revises, since what counts as a working quantum computer is still under individuation. This is what separates the account from the deflationary reading, on which the horizon is a run of mistaken forecasts about a fixed target, to be cured by better calibration. The separation is empirical, with a discriminating consequence. If the target were fixed but merely hard to estimate, one would expect forecasts to converge and the field to consolidate around a dominant architecture as it matured. Instead, fidelity improves while architectures continue to diverge (Figure 1), so progress accumulates while the horizon holds. That is the signature of a moving target rather than a missed one.² It also dissolves the apparent dilemma between reading “five years” as a single benchmark, which would make the phrase misleading, or as the whole benchmark set, which would make it false: it is neither, because concretization revises which benchmarks belong in the set.

Endogeneity likewise separates PFYT temporality from exogenous accounts, on which development is slow through coordination failure, misaligned incentives, or thin capital, and could in principle be accelerated by policy or funding [1]. Those forces are real, but they bear on a different contrast class. As Section 1 noted, exogenous force compresses engineering timelines within a mature milieu, as in Manhattan and Apollo, while leaving concretization timelines untouched: capital did not purchase coherence for superconducting qubits, and cold fusion deferred perpetually with no milieu accumulating at all [4]. Such accounts explain how fast engineering proceeds once the physics is in hand, not why concretization itself resists compression.

The closest kin to the framework is Lakatos’s methodology of scientific research programmes, which already explains why effort gathers around progressive programmes and drains from degenerating ones [24]; epistemic momentum can be read in those terms. What PFYT adds is a different unit of analysis. Lakatos individuates programmes by theory and grades them by novel empirical content, whereas the platforms in Figure 2 share one quantum mechanics, and what redistributes among them is concretization headroom: the capacity a technical individual retains to keep co-individuating with its milieu. The trapped-ion paradox shows why the difference matters. Trapped ions hold the best gate fidelities and continue to yield strong results, so by a content criterion, they are not degenerating, yet momentum migrates away from them toward platforms with more open individuation. A leading performer losing momentum follows naturally from an account that locates momentum in remaining capacity to individuate rather than in current performance. The same shift in unit clarifies why recursion here is not the generic recursion of all inquiry: the constraints are produced and ordered by scale, each inaccessible until the prior threshold is crossed. Crossing the error-correction threshold is what brings qubit count, connectivity, and decoder integration into view as the next conditioning constraints; the sequence is gated rather than open.

Because the account is endogenous, it makes commitments that could be disconfirmed. Were a single architecture to consolidate and the remaining benchmarks settle into a jointly forecastable set, the horizon would close, and the deflationary reading would be vindicated; were research momentum to track current performance rather than individuation

headroom, the Lakatosian reading would suffice. The framework earns its place if these commitments hold where its rivals expect otherwise.

Taken together, these patterns (converging error rates, divergent scaling trajectories, shifting epistemic momentum, and recursive constraint discovery) give the framework its empirical content. The next section examines what these patterns imply for how quantum computing should be understood as a technical object: why analogies to classical computing history mislead, how “quantum advantage” should be reframed, and what the simultaneous validation of analog and digital quantum simulation reveals about the plurality of computational paradigms.

3. Quantum Computing as a Technical Individual

The empirical patterns documented above raise a prior question: what kind of technical object is a quantum computer? The answer is less obvious than it appears, and getting it wrong distorts both expectations and policy. Actor-network theory [25] illuminates one dimension of this question: a technology stabilizes only when heterogeneous alliances of components, practices, institutions, and expertise have cohered around it. The ontological question asked here is narrower: what is a quantum computer, as a technical individual in Simondon’s sense, prior to that stabilization? In Simondon’s vocabulary, a technical individual is an object that has cohered into a self-sustaining functional unity with its own associated milieu, occupying the level between the isolated element (a single component) and the ensemble (a system of such individuals working together). The question is whether a quantum computer has yet reached that level of unity, or whether it remains a shifting assembly of elements still in search of its individuality.

3.1. *Why Comparison with Classical Computing History Fails*

A common analogy, used in the press but also by researchers and industry practitioners, compares current quantum processors to ENIAC, the primitive ancestor of modern computers [26]. The implication is reassuring: just as ENIAC evolved into smartphones, so quantum computing will mature along a predictable trajectory. But the comparison obscures more than it reveals.

ENIAC inherited a mature technical milieu. Vacuum tubes, by 1945, had been refined through decades of radio and telephony; supply chains existed; manufacturing knowledge was established; trained engineers were available. ENIAC was an application of concretized components to a new problem. Its core abstraction (Boolean logic in switching circuits) was already understood. The path forward was engineering: scale, reliability, miniaturization. Quantum computing faces a categorically different situation. The associated milieu (precision lasers, cryogenic systems, ultra-high vacuum, control electronics, error-correction theory, high-purity materials, and the specialized foundries capable of producing them at the yields required) must be built alongside the processor itself. The components are not mature; they co-individuate with the systems they enable. More fundamentally, the computational abstraction itself remains unsettled. What does “quantum computation” mean? Digital fault-tolerant circuits? Analog simulation? Continuous-variable processing? The answer is not given in advance; it must be discovered through engagement with hardware at scale. The disanalogy runs deeper than an unsettled vocabulary. In a classical machine, the logical bit can be read without disturbing the device, so the digital abstraction floats almost free of the substrate that carries it; this near-independence is what allowed Boolean logic to be specified before the transistor existed. A qubit admits no such separation. Measurement is constitutive: reading the device collapses its state, and the discrete outcome is an observable extracted from an irreducibly continuous substrate [27]. The line von Neumann drew between analog and digital machines does not fall cleanly

on the quantum side; it runs through the middle of every quantum processor, which is analog in its state and digital only at readout. The abstraction therefore cannot be fixed independently of the physics, and its settling is part of the individuation rather than a precondition of it.

The ENIAC comparison assumes technological development is a trajectory through a fixed space, earlier to later, along a known path. PFYT development is radically nonlinear: each advance reconfigures the space of possibilities rather than moving through a stable one, and comparisons across historical eras, however tempting, risk imposing a false continuity on processes that differ in kind rather than in stage. The framework developed here suggests otherwise: quantum computing is undergoing individuation in an open field, where the identity of the technical object, its relationship to its milieu, and even what counts as success emerge through the process itself. The five-year horizon does not mark a position on a known curve; it reflects the structure of a process where each step reshapes what the next step must be.

3.2. Reframing Quantum Advantage

The PFYT framework also reframes debates over “quantum advantage” and “quantum utility.” These terms are typically conceived as threshold events: a moment when quantum outperforms classical on some task. But as LaRose’s survey of the field’s history makes clear, this framing generates perpetual frustration: each claimed demonstration is challenged by improved classical algorithms, and the threshold appears to recede [28,29].

A Simondonian perspective clarifies what is at stake. Eisert and Preskill identify four “gaps” separating current devices from fault-tolerant machines. The PFYT framework explains why these gaps have the structure they do: each is a site of active concretization, a phase during which components must become mutually conditioning, and the technical object must integrate with its milieu. The gaps are recursive; crossing one reveals the next.

What I call *abstraction-level matching* illuminates why certain paradigms demonstrate advantage in specific regimes. Quantum hardware achieves an advantage when its native physical interactions directly implement the target problem, eliminating the translation overhead that plagues general-purpose approaches. This is concretization at the computational level: the technical object becomes optimally integrated with its problem domain when no abstraction layer mediates between implementation and physics. Like early microscopes that revealed cells before anyone proved cells were “impossible to see unaided,” quantum simulators may generate discoveries through systematic exploration of what current hardware accesses naturally, rather than waiting for algorithmic proofs of classical intractability.

Analog quantum simulation operates through a mode of knowledge production that is epistemically distinct from digital computation: it does not *calculate* the behavior of a target system but *enacts* it by constituting an analogous quantum system at a different physical scale. When an array of cold atoms reproduces the magnetic phase diagram of a frustrated quantum material, no algorithm has been executed; instead, one physical system has been made to disclose the behavior of another through structural correspondence. This intersimulation realizes Feynman’s original proposal [30] while collapsing the distinction between instrument and phenomenon in ways he did not theorize.

Recent work provides the first empirical confirmation of this epistemic structure. A 256-atom neutral-atom simulator has reproduced the magnetization, heat capacity, and magnetic phase transitions of the frustrated quantum magnet TmMgGaO_4 in quantitative agreement with independent laboratory measurements: the first one-to-one validation of analog quantum simulation against experimental data on a real material [20]. The simulator further accesses non-equilibrium dynamics on picosecond material timescales, directly

realizing the temporal mediation that analog simulation affords. Independently, a superconducting digital quantum processor has computed the dynamical structure factor of the Luttinger liquid KCuF_3 , benchmarked against inelastic neutron-scattering measurements [21]. That both demonstrations appeared within days of each other, one analog, one digital, on different platforms, validated against different experimental techniques, is itself a PFYT signature: the ecosystem has matured to the point where simulation capabilities emerge cross-platform, but the plurality of computational paradigms persists rather than converging. Each constitutes a distinct technical individual with a distinct mode of engagement with physical reality.

The PFYT perspective suggests that advantage is not a threshold to be crossed but an emergent property of concretization, specific to platform-problem combinations. An analog processor demonstrates an advantage for certain physical dynamics; a fault-tolerant digital machine will demonstrate an advantage for different problems. These are not competing claims but different technical individuals with different associated milieus, each optimally concretized for different problem domains. This is why the reframing bears on the paper's thesis rather than standing apart from it: if advantage is a property of particular object-milieu-problem combinations rather than a single threshold to be crossed once, then no individual comparison closes the horizon, and the persistence of multiple distinct advantages is the same plurality of individuations that keeps the five-year horizon open. "Economic Quantum Advantage", value creation in practice [31], may be the appropriate metric because the technical object individuates with its applications rather than prior to them.

Ihde's postphenomenology clarifies what is philosophically at stake: each platform does more than implement "quantum computing" differently; it constitutes a distinct *technological intentionality*: a different structure of disclosure through which different phenomena become accessible [32,33]. The question "what is a quantum computer?" cannot be answered independently of which platform asks it. Anderson's principle that each level of complexity generates autonomous laws irreducible to the level below [34] suggests why: analog simulation, fault-tolerant digital computation, and hybrid approaches may constitute distinct organizational levels, each with its own emergent regularities and its own associated milieu. The plurality is not a transitional phase awaiting consolidation but the natural structure of a domain whose pre-individual field is rich enough to sustain multiple individuations.

3.3. *Scaling, Thermalization, and the Limits of Individuation*

A still deeper question haunts these debates, one that has shadowed quantum theory since its inception. The boundary between quantum and classical behavior remains one of the deepest open problems in physics, debated by Bohr, Heisenberg, and Schrödinger and confirmed experimentally as a genuine feature of nature rather than a theoretical artifact [35], yet unresolved a century later. Decoherence theory offers the most developed contemporary account: the environment continuously monitors the quantum system, selecting preferred "pointer states" that survive entanglement with their surroundings while destroying coherence between alternatives [36]. Zurek's ambitious extension, quantum Darwinism, proposes that the states we perceive as classical are those whose information has been redundantly imprinted across environmental degrees of freedom, a selective amplification that would produce the appearance of objective reality [37]. The account is compelling but not definitive; the measurement problem persists, and whether decoherence fully explains the emergence of classicality or merely displaces the puzzle remains actively contested. What matters for the PFYT framework is the structural implication: error correction is engineering against the grain of whatever process enforces classicality, maintaining

fragile superpositions that the environment works ceaselessly to destroy. Recent work on quantum many-body scars and thermalization dynamics suggests that at scale, this environmental pressure intensifies: large quantum systems exhibit intrinsic tendencies toward ergodicity that no amount of engineering can simply override [38,39]. Whether the exponential advantages promised by quantum computation survive at the scales required for useful problems remains genuinely open. This is recursive constraint discovery at its most fundamental: the very scaling that engineering pursues may reveal physical boundaries that reshape what quantum computation can mean. The PFYT framework predicts exactly this structure: each layer of progress exposing the next layer of resistance, with the deepest layers invisible until the system is large enough to encounter them.

This raises a question the framework must confront: what happens when resistance cannot be accommodated? If ergodicity limits prove fundamental rather than engineerable, the technical object encounters more than a harder problem; it discovers that the individuation it was pursuing has no viable completion. The PFYT framework accommodates this possibility: not all pre-individual fields resolve into mature technical individuals. Some individuations fail, and the signature of that failure is resistance that redirects rather than yields, forcing not further concretization along the same path but a branching toward alternative individuations, such as the shift from universal fault-tolerant quantum computing toward domain-specific analog simulation.

The same physical process thus operates differently depending on the computational paradigm within which the quantum system is deployed, a paradigm-dependence of resistance that the PFYT framework predicts but that single-metric accounts of quantum advantage cannot capture. Yet thermalization is not a threat to quantum advantage alone; it can also serve as a *validation mechanism*. When a quantum simulator's long-time dynamics converge to thermal equilibrium predictions derived from independent methods, this convergence provides evidence that the simulation is faithful, even in dynamical regimes where no classical method can track the full evolution. The analog simulation of a frustrated quantum magnet discussed in Section 3 confirms this diagnostic pattern: the simulator's convergence to thermal equilibrium validates the quantum simulation against independent macroscopic measurements, while its access to non-equilibrium dynamics probes regimes beyond classical reach [20]. The potential obstacle becomes a diagnostic tool: resistance transformed into accommodation, consistent with the recursive pattern that characterizes PFYT development at every level.

4. Concretization: The Neutral-Atom Quantum Processor

The previous sections established that quantum computing has not individuated into a single dominant form and that its identity as a technical object remains open. To see how this openness is structured, how components move from abstract juxtaposition toward mutual conditioning, we need Simondon's concept of *concretization*: technical objects develop from abstract assemblages of separate components toward concrete unities where each element conditions the others [2].

Consider the historical evolution of the vacuum tube, a case analyzed in detail in the foundational literature on technical individuation. Early tubes were abstract in a precise sense: components were juxtaposed, each serving a single function, with minimal interaction between them. The cathode emitted electrons; the grid controlled their flow; the anode collected them; the envelope contained the vacuum. Each element could be understood, and to some extent designed, independently.

Through decades of development, tubes became increasingly concrete. The grid came to serve multiple functions simultaneously: controlling electron flow, shaping the electric field, and influencing thermal dynamics. Thermal effects that were initially parasitic (un-

wanted heat affecting performance) became integrated into the tube's operation, sometimes exploited for stability. The envelope was no longer mere containment but participated in the tube's thermal and electrical behavior. The mature vacuum tube was a unified technical individual in which no component could be fully understood apart from its relationships to the others.

This concretization occurred not in isolation but in relation to what is called the *associated milieu*: the technical environment within which the tube operated [2]. The tube and its circuit, its power supply, its cooling arrangements, and its mounting hardware co-evolved. Changes in the tube required changes in the circuit; advances in circuit design enabled new tube configurations. The technical individual was never the tube alone but the tube-in-its-milieu.

A methodological point deserves emphasis. Simondon's analyses of concretization (the vacuum tube, the Guimbal turbine, the internal combustion engine) were retrospective: the objects had already individuated, and the philosopher reconstructed the logic of their development. The present analysis attempts something different: applying the framework prospectively to a technology still in active individuation. This is possible because PFYT dynamics make the structure of becoming visible while it unfolds. We can observe neutral-atom processors moving from abstract to concrete configurations in real time, watching the mutual conditioning of components emerge through engineering practice. The philosophical contribution is less an application of Simondon than an extension of his method into the domain of technologies whose identity remains genuinely open. Recent scholarship has begun similar extensions, applying individuation theory to deep learning systems [40], synthetic biology [41], and the co-evolution of technical milieus [42], but the application to physics-based technologies undergoing active concretization remains unexplored. The stakes are ontological rather than methodological: in a PFYT, the identity of the technical object is itself what is at stake. The neutral-atom processor is not a known device being incrementally improved but a technical individual whose nature (analog simulator? digital gate processor? photonic medium?) is being discovered through the process of concretization itself.

The neutral-atom quantum processor exhibits the same developmental logic as Simondon's classical cases, from abstract assemblage toward concrete unity, but with distinctive features that illuminate PFYT dynamics. The analysis focuses on neutral atoms because the platform is in active concretization, abstract enough to exhibit the dynamics of becoming, concrete enough to make them empirically visible. Superconducting circuits, further along in concretization, would illustrate the same structural logic with different specifics; the choice of case reflects analytical clarity, not advocacy. Consider the structure of a contemporary neutral-atom system (Figure 3). Individual atoms, typically rubidium, ytterbium or cesium, are held in arrays of tightly focused laser beams called optical tweezers. Quantum information is encoded in the atoms' internal electronic states. Specific Hamiltonians, the mathematical descriptions of energy and interactions that govern a quantum system's evolution, and eventually quantum gates are implemented by using additional lasers to drive controlled interactions between neighboring atoms. Readout occurs through fluorescence imaging: atoms in one state scatter light; atoms in another do not [10,43,44].

Early neutral-atom systems were abstract in Simondon's precise sense, and deliberately so. Laboratory platforms maximize degrees of freedom: free-space optics that can be realigned, beam paths that can be rerouted, parameters that can be swept independently. This openness is not a deficiency but a necessity: the design space must be explored before it can be committed to. Each subsystem retains its own adjustment logic; integration is provisional, and the architecture remains fluid so that better configurations can be discovered through experiment.

Concretization is the transition from this exploratory openness to a committed unity. In a deployed system, degrees of freedom that were assets in the lab become liabilities: free-space optics drift, unconstrained parameters introduce variance, and flexible beam paths compromise machine availability. The move toward photonic integrated circuits [45] is less miniaturization than the deliberate freezing of design choices into mutual constraints: trapping, control, and readout consolidated onto a single substrate where each element stabilizes the others. The vacuum chamber is no longer mere isolation but constitutive of the computational space; its geometry constrains array configurations, its optical access determines measurement possibilities, and its thermal properties set coherence limits. What was a collection of adjustable knobs becomes a self-conditioning system whose components can no longer be modified independently. This is concretization: not that the system does more, but that it has committed to a mode of existence where each element's identity is defined by its relationship to the others. This trajectory from optical table to chip recapitulates the classical semiconductor story, but with a critical difference: the atoms remain quantum objects whose behavior cannot be fully specified in advance by the chip designer.




Figure 3. Neutral-atom quantum computing: from principle to system. (a) Schematic of a neutral-atom optical tweezer platform. Laser light shaped by a spatial light modulator (SLM) creates an array of tightly focused traps inside an ultra-high vacuum chamber; fluorescence at 780 nm images individual atoms on an EMCCD camera. (b) Fluorescence image of individually trapped atoms arranged in a regular array, a technical image in which decades of accumulated knowledge in atomic physics, laser technology, and quantum control have crystallized into operational capability. (c) Digital-analog quantum circuit on a neutral-atom platform. Digital gate operations (state preparation, characterization) bracket a phase of analog evolution in which the atom array evolves under its native Hamiltonian; the inset reveals the physical atom dynamics underlying both abstractions. This hybrid structure illustrates why quantum computing's identity as a technical object remains unsettled: rather than converging on a single computational paradigm, the platform integrates digital control and analog physics within the same circuit, each conditioning the other. (d) A Pasqal neutral-atom quantum processor deployed at the Jülich Supercomputing Centre as part of the EuroHPC HPCQS project. The sealed enclosure integrates a vacuum system, laser sources, and control electronics. Sources: Pasqal and Antoine Browaeys.

Even measurement, long the simplest operation conceptually, has become a site of active concretization: non-destructive mid-circuit readout, essential for quantum error

correction, now requires spin-resolved techniques that preserve neighboring qubits while extracting syndrome information [18]. The atom array geometry simultaneously determines the connectivity graph for quantum gates, the error profiles arising from inter-atom interactions, the scalability limits imposed by optical aberrations, and the measurement fidelity achievable through fluorescence imaging. No single aspect can be optimized in isolation; each choice ramifies through the entire system.

One feature of neutral-atom systems deserves special emphasis: the atoms themselves are identical. Unlike superconducting qubits, which must be individually fabricated and calibrated, neutral atoms of the same isotope are physically indistinguishable. There is no manufacturing variability at the level of the qubit itself. This might seem to simplify matters, but it shifts the locus of technical individuation. The technical individual is not the atoms but the entire structured relationship: atoms + light + vacuum + control. The processor *is* this relationship; it cannot be separated from its milieu because there is no “processor” apart from the milieu.

A striking consequence of this concretization has recently become visible. When the same quantum simulation, a one-to-one model of a frustrated quantum magnet, is reproduced quantitatively across three geographically distributed neutral-atom processors, two in France and one in Canada, each realizing the same effective Hamiltonian through rescaled array geometry [20], the result confirms that the associated milieu has reached a stage of maturation where the technical individual can be instantiated rather than prototyped. The relationship between atoms, light, vacuum, and control has become sufficiently concrete that it transfers across independent physical implementations. This is milieu portability, a diagnostic of concretization that Simondon’s retrospective case studies could identify only after the fact, but that PFYT analysis can now observe in real time.

Yet the plasticity of the pre-individual field has not been exhausted by this concretization. The same platform can crystallize into a fundamentally different technical individual. Subwavelength arrays of atoms can serve not as qubits but as a collective non-linear medium for photonic quantum information: “the Rydberg atoms are no longer the support of the quantum information, but rather act as a source of non-linearity for photonic qubits” (Section 7 of [10]). The hardware is recognizably similar (optical tweezers, vacuum chambers, laser cooling), yet the computational abstraction is entirely different: quantum information is carried by light, not by atomic states. This is less a repurposing than a distinct individuation from the same pre-individual field, confirming that the field sustains a branching space of possible individuations, the signature of a domain in active becoming.

The associated milieu of a neutral-atom processor extends far beyond the laboratory apparatus. It includes supply chains for precision optical components, laser manufacturers who have adapted their products for quantum applications, vacuum technology companies developing chambers optimized for atom arrays, control electronics specialists creating systems with nanosecond timing precision, software teams developing compilers that translate quantum algorithms into pulse sequences, and the workforce trained at the intersection of atomic physics, photonics, and computer science. Each element has its own developmental trajectory, its own research community, its own manufacturing challenges. The neutral-atom processor matures only as this entire ecosystem matures, and the ecosystem matures only through engagement with the demands of processor development.

This raises a question about boundaries. Simondon distinguished the *associated milieu* from the broader *geographical and social milieu* within which technology exists (Chapter 1 of [2]). For classical technical objects, the distinction was relatively clear: the vacuum tube’s associated milieu was its circuit; the broader milieu was the radio industry. PFYTs blur this boundary. When error-correction theory shapes processor architecture, when workforce training determines which designs are feasible, when funding mechanisms influence which

platforms survive: are these part of the associated milieu or external conditions? The answer may be that PFYT dynamics are characterized by the expansion of what counts as associated milieu: more of the social and institutional environment becomes directly coupled to technical functioning than in conventional engineering, a dynamic that resonates with recent arguments for understanding technical milieus as irreducibly entangled with their broader epistemic and institutional contexts [46]. The five-year horizon indexes more than physical constraints; it indexes the time required for this expanded milieu to cohere.

This co-evolution of technical object and associated milieu explains why concretization cannot be rushed. The stages are not arbitrary delays; they are phases during which the mutual conditioning of components is discovered and engineered. You cannot skip from abstract assemblage to concrete unity because the relationships that constitute concreteness must be found, and finding them requires building, testing, failing, and rebuilding. The five-year horizon reflects the time required for this discovery.

5. Resistance and Accommodation: The Temporal Mechanics of Concretization

Concretization describes the *structure* of technical evolution: what the technical object becomes. But it does not fully explain the *process*: why getting there takes time. For this, we need to understand how engineering actually proceeds: through a dialectic of resistance and accommodation [47]. Simondon's structural account identifies what concretization produces, the increasing mutual conditioning of components, but is less explicit about the temporal mechanics of how each stage is traversed. Pickering's processual vocabulary supplies this: the iterative encounter between human intention and material recalcitrance that explains why each phase of integration takes the time it does.

Technical development is not the straightforward implementation of a design. Engineers approach the material world with goals and plans; the world resists; engineers accommodate by modifying their goals, their instruments, and their conceptual frameworks. This back-and-forth, what has been aptly termed the "mangle" of practice, is how technical knowledge is actually produced.

Consider the development of neutral-atom quantum processors in practice. The goal seems clear: trap atoms, control their quantum states, perform computations. But every step encounters resistance.

The atoms do not behave as expected at scale. At small numbers (tens of atoms) the system is tractable. As arrays grow to hundreds or thousands, new phenomena appear. Atoms interact in ways that single-atom models do not capture. Rydberg blockade, the mechanism that enables entangling gates, produces correlations across the array whose complexity grows beyond what few-atom models capture [43]. Collective effects in fluorescence imaging complicate state readout. Each phenomenon was, in principle, predictable from quantum mechanics, but predicting and observing are different; the resistance of the actual system forces accommodation in measurement protocols, gate sequences, and error mitigation strategies.

The same pattern recurs in every subsystem. Atoms load stochastically into roughly half the available traps, requiring real-time imaging and individual rearrangement by optical tweezers before each computational cycle can begin, an accommodation to probabilistic loading that adds latency and demands its own control infrastructure. Lasers drift; engineers accommodate through active stabilization, but each stabilization system introduces its own complexity and failure modes. Ultra-high vacuum requires managing outgassing dynamics and surface chemistry that destroy atomic coherence; materials that work in other vacuum applications prove unsuitable, and each accommodation must be validated through lengthy testing cycles. Control electronics must achieve nanosecond

timing across hundreds of channels, but faster response demands more power and heat in environments where thermal stability is critical. In each case, the accommodation to one resistance creates new sites of potential resistance.

The pattern repeats at every scale and in every subsystem. Resistance is not an obstacle to be overcome once but a permanent condition of technical development. Accommodation is not a solution but a response that creates new conditions, including new forms of resistance. Concretization proceeds through this mangle: the increasing integration of components occurs via the iterative discovery of how they interact, interfere, and constrain each other.

The resistance-and-accommodation dynamic supplies the mechanism behind the horizon-resetting pattern identified in Section 2: each accommodation creates new conditions from which new resistances emerge, so that the landscape is generated through the act of traversing it. This is not the generic observation that “problems are hard.” It is a structural claim: the specific resistances encountered at each stage are systematically produced by the accommodations that resolved the previous stage. The horizon does not reset because engineers underestimate difficulty; it resets because the difficulty is literally created by progress.

The temporal structure of PFYTs follows from this dynamic. Progress is real: problems are solved, capabilities advance, and knowledge accumulates. But progress is recursive: each advance reconfigures what counts as the problem, what tools are available, what resistances matter. The five-year estimate is always approximately correct because it describes the time required to work through the currently visible layer of resistance. It is always falsified because the work itself generates the next layer. The question is not whether this rhythm can be broken but whether the institutions financing the work can be structured to sustain it.

This recursive structure has direct implications for the institutions that finance PFYT development. Venture capital, originally designed to fund patient technological development through long gestation periods [48], has increasingly compressed its time horizons around software deployment cycles. PFYTs confront investors with a structural tension between two logics. The first pursues short-term value within the recursive cycle itself: each new constraint discovered is also a new problem to be solved, and each solution a potential product. Companies extract value from intermediate capabilities (quantum sensing, calibration tools, simulation services) while the core technology continues to individuate. The second targets the long-term prize directly: quantum computing, commercial fusion, brain–computer interfaces, and novel materials at scale. But this path risks bypassing the intermediate validations that redirect development, forgoing the resistance-and-accommodation cycles through which concretization actually proceeds. The PFYT framework suggests that the first logic is less a funding compromise than a structural necessity: early markets provide the sites of resistance through which the technical object discovers what it must become. Patient capital, in this light, is not simply capital that waits longer; it is capital whose structure aligns with the temporality of concretization.

6. Physical AI and PFYT Dynamics

The preceding sections analyzed quantum computing’s development through the lens of concretization and resistance, dynamics internal to the technical object and its milieu. But quantum computing did not develop in isolation. Its maturation has coincided with the rise of machine learning, and the two technologies are now entering a phase of mutual acceleration that neither trajectory anticipated. The PFYT framework must account for more than the internal dynamics of a single technology: for what happens when independently individuating technologies discover that their capabilities are complementary.

Understanding this convergence, rather than merely exploiting it, is essential because it is reshaping the temporal structure of physics-based innovation in ways that the preceding analysis, focused on a single technical object, does not capture.

I use the term *physical AI* to designate AI systems whose intelligence is grounded in engagement with physical reality, whether through simulation of quantum systems, molecular dynamics, robotic interaction, or learned world models, rather than through pattern extraction from text corpora. The term encompasses both AI trained on physics (foundation models for molecular simulation, quantum control optimizers) and AI that operates in the physical world (autonomous systems, embodied agents). What unifies these is the mode of engagement: physical AI encounters a world that *resists*, where predictions fail, simulations diverge, and material constraints refuse to accommodate computational convenience. Where digital AI operates on knowledge already stabilized into text, with facts black-boxed and uncertainties resolved; physical AI engages with phenomena not yet reduced to propositions, systems whose behavior is still being discovered. LeCun's argument that genuine machine intelligence requires world models grounded in physical reality rather than in language [49] articulates the broader trajectory of which quantum-AI convergence is a specific and technically advanced instance.

To be precise about how physical AI modifies PFYT dynamics: it compresses validation latency for parameter-dominated cycles (control optimization, error characterization, circuit compilation) while leaving material-dominated cycles invariant. No learning algorithm shortens the months required to validate a new vacuum chamber design or qualify a novel photonic integrated circuit. Physical AI expands the associated milieu by coupling AI research communities, their tools, and their data practices to quantum hardware development, creating new sites of co-individuation. And through the recursive loop of simulation, training, and discovery, it introduces a new form of recursive constraint discovery in which the learning system and the physical system mutually reveal each other's limitations. The PFYT conditions are not eliminated but reconfigured: some cycles compress, new dependencies emerge, and the overall temporal structure shifts without collapsing.

Three modes of this transformation deserve analysis.

Compression layer. Foundation models trained on quantum mechanical data (solutions to the Schrödinger equation, not text) can predict molecular properties, reaction pathways, and material behaviors across regions of chemical space that no human researcher could systematically explore [50,51]. For quantum hardware specifically, machine learning compresses calibration, error characterization, and circuit compilation cycles from weeks to hours [52]. Yet compression has limits: it accelerates loops dominated by parameter exploration but cannot compress cycles dominated by materials validation, manufacturing lead times, or the physical constraints of metrology. The five-year horizon reflects these incompressible elements as much as the compressible ones.

Epistemic substrate. The most consequential near-term role of quantum computing may be generating training data for AI rather than performing computation *per se*. Quantum processors, even at current scales, simulate quantum systems natively, producing physically authentic data for strongly correlated systems that classical approximations handle imperfectly [53]. A recent synthesis maps the full architecture of this convergence: foundation models for molecular dynamics achieve quantum-accurate simulations but remain tethered to the limits of classical data generation; hybrid QPU-GPU architectures are positioned as the accelerator that breaks this ceiling [54]. The "hearing aid moment" for quantum computing may not be a specific application domain but a capability: the production of synthetic data that bridges the gap between what classical simulation can access and what AI models need to learn.

Co-individuation. Most significantly, physical AI and quantum hardware are co-individuating, each reshaping the other's developmental trajectory. Co-individuation, in Simondon's terms, is individuation that runs in both directions: neither technology pre-exists the relation as a finished object, and each becomes what it is partly through its coupling to the other. The FeNNix-Bio1 foundation model achieves quantum-accurate molecular dynamics because it was trained on solutions to the Schrödinger equation rather than on empirical force fields, but its accuracy horizon is bounded by the classical methods that generated that data, creating the demand for quantum hardware to extend the training frontier. Conversely, AI models shape which problems get explored and which configurations get sampled, guiding the development of the very hardware that will generate their next generation of training data (Figure 4). This coupling was not designed; it emerged because two independently individuating technologies discovered complementary capabilities that were invisible until both reached sufficient maturity.




Figure 4. The Physics–AI recursive loop. Two classes of data sources feed synthetic training data for foundation models. High-accuracy quantum simulations on classical processors (QMC, sCI, DFT) solve the Schrödinger equation numerically; quantum processors simulate native quantum dynamics at matched scale. The QPU enters this loop not through computational supremacy but through physical fidelity. Foundation models trained on this physics-native corpus accelerate discovery and design of molecules, materials, and reaction pathways. At the center, co-individuation of intelligence and matter. Adapted from [54].

This recursive dynamic resonates with Hui's distinction between mechanical repetition and genuine recursivity: a looping movement that returns to itself while remaining open to contingency at each iteration [55]. The Physics–AI loop is recursive in this sense: each cycle opens regions of possibility space that were previously inaccessible, redirecting the development of the very hardware that generated the training data.

A deeper point emerges. Large language models operate on text corpora, a world already filtered into symbols through human practices of writing. Physical AI engages differently: trained on quantum simulations, it encounters a world not yet reduced to propositions. Digital AI operates on what might be called settled matters, knowledge already stabilized and encoded [56]. Physical AI engages with matters still in contestation, phenomena whose behavior is still being discovered [57]. Whether this constitutes “understanding” in any robust sense remains open. But for PFYT development, the implication is that physical AI may transform the pace and even the character of technical evolution.

7. Conclusions

The perpetual five-year horizon is not a forecasting failure but the temporal signature of a distinctive mode of technical becoming. PFYTs develop through concretization that cannot be rushed, resistance and accommodation at every stage, ecosystem co-individuation, and the recursive discovery of constraints. The framework developed here makes this temporality visible as a structural feature of physics-based innovation rather than a symptom of poor planning or insufficient funding.

Quantum computing provided the paradigmatic case. Cross-platform performance data and the distribution of high-impact research revealed more than technical progress: the reorganization of an entire field. The neutral-atom processor, examined as a technical individual in active concretization, showed how laboratory openness gives way to committed integration, and how milieu portability across geographically distributed machines confirms that concretization has reached a new stage. Physical AI is reconfiguring these dynamics, compressing some development cycles while introducing new forms of co-individuation between intelligence and matter.

The framework generates testable predictions beyond quantum computing, and each domain shows the four PFYT criteria in domain-specific form. Nuclear fusion exhibits all four: milieu-dependence on superconducting magnets, plasma-facing materials, and tritium breeding; concretization as confinement, heating, and fueling become mutually conditioning at reactor scale; recursive constraint discovery as each advance in plasma stability reveals new challenges in neutron damage or heat exhaust; and validation latency measured in years per experimental campaign. Compact particle accelerators and brain-computer interfaces follow the same logic, with domain-specific variations: laser-plasma-beam co-individuation in one case, biological adaptation governing validation latency in the other. Quantum computing is the paradigmatic case partly because it is the most densely benchmarked, so its concretization leaves an unusually legible quantitative trace; the other domains satisfy the same four criteria with sparser instrumentation, which is why the framework is sharpest where development is continuously measured and must rely on more qualitative markers elsewhere. Technologies that have crossed analogous long-horizon thresholds, such as reusable rocketry, offer a retrospective illustration of the same dynamics. Each case deserves the detailed treatment attempted here; the present paper offers the framework and the paradigmatic case, leaving extension as an invitation to scholars at these frontiers.

The framework has clear boundaries. It is a diagnostic of temporal structure, not a predictive model: it identifies PFYT dynamics but does not forecast which technologies will succeed or on what timeline. It does not replace political economy, industrial policy, or geopolitical analysis, all of which shape PFYT trajectories from outside the technical object. Not all physics-based technologies are PFYTs; some inherit mature milieus and scale conventionally. And the framework's extension from Simondon's retrospective mechanical cases to prospective quantum systems, while supported by the evidence presented here, remains an open methodological wager whose validity will be tested by further application.

Philosophy of technology has focused extensively on digital systems. It has paid far less attention to the physics-based frontiers where the relationship between intelligence and matter is being actively contested. Quantum computing, fusion energy, advanced materials, and brain-computer interfaces all raise questions as deep as any posed by digital technology: What is the nature of technical evolution? How do objects and their environments co-constitute each other? What happens when learning systems engage with physical resistance rather than symbolic representation? The PFYT framework is offered as one entry point into these questions. Understanding the temporal structure of physics-based innovation, rather than expecting it to dissolve, is a precondition for judging

these technologies on the timescale they actually occupy: for distinguishing endogenous developmental time from remediable delay, and for calibrating expectation, evaluation, and resource decisions to the structure of concretization rather than to the cadence of software.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The empirical analyses presented here draw on two datasets. Quantum hardware performance data (Figure 1) was provided by the Quantum Error Correction Report 2024 and 2025, available from Riverlane. The epistemic momentum dataset (Figure 2) was constructed from OpenAlex bibliometric data with manual modality classification; the methodology and dataset are available on request from the corresponding author.

Acknowledgments: The author thanks Brooke Abeles for compiling the epistemic momentum dataset (Figure 2) and for her comments, in particular the Manhattan Project and Apollo program comparison. The author also thanks Eve Townsend at Riverlane for kindly providing the error-rate and qubit-count data shown in Figure 1.

Conflicts of Interest: The author is a co-founder and partner of Quantonation, a venture-capital firm that invests in quantum technologies and deep-physics companies, and is a co-founder and shareholder of Pasqal, a neutral-atom quantum-computing company whose systems and data are referenced in this article. These commercial interests are directly relevant to the subject matter of the paper. The analysis, views, and conclusions presented here are the author's own and were not influenced by these interests.

Notes

- ¹ The framework draws on sustained involvement in the deep-tech ecosystem: cross-platform visibility and direct observation of the coupling and resistance dynamics that academic analysis typically encounters only through published accounts. Such a standpoint naturally overweights scale-up bottlenecks relative to political economy, public funding, and geopolitical factors that shape PFYT trajectories from outside the technical object. The framework is offered as complementary to those perspectives, not a replacement. It also predicts failure: ventures that scale marketing faster than their milieu matures will stall regardless of funding.
- ² A structurally similar recession appears in definitions of artificial general intelligence. Turing's criterion of conversational indistinguishability [23], once treated as the threshold, is approximately met by current language models, yet the operative definition of general intelligence has since moved on to properties such as stable identity or autonomous motivation. The goal recedes as the technology approaches it, the same temporal signature, in a conceptual register, that this paper locates in the material development of PFYTs.

References

1. Arthur, W.B. *The Nature of Technology: What It Is and How It Evolves*; Free Press: New York, NY, USA, 2009.
2. Simondon, G. *On the Mode of Existence of Technical Objects*; Malaspina, C., Rogove, J., Translators; Aubier: Paris, France, 1958.
3. Simondon, G. *Individuation in Light of Notions of Form and Information*; Complete Edition of the 1958 Thesis; Adkins, T., Translator; Jérôme Millon: Grenoble, France, 2005.
4. Berlinguette, C.P.; Chiang, Y.M.; Munday, J.N.; Schenkel, T.; Fork, D.K.; Kober, R.; Trevithick, M.D. Revisiting the Cold Case of Cold Fusion. *Nature* **2019**, *570*, 45–51. [[CrossRef](#)] [[PubMed](#)]
5. Garisto, D. LK-99 Isn't a Superconductor—How Science Sleuths Solved the Mystery. *Nature* **2023**, *620*, 705–706. [[CrossRef](#)] [[PubMed](#)]
6. Langmuir, I. Pathological Science. *Phys. Today* **1989**, *42*, 36–48. [[CrossRef](#)]
7. National Academies of Sciences, Engineering, and Medicine. *Quantum Computing: Progress and Prospects*; The National Academies Press: Washington, DC, USA, 2019. [[CrossRef](#)] [[PubMed](#)]
8. Bruzewicz, C.D.; Chiaverini, J.; McConnell, R.; Sage, J.M. Trapped-Ion Quantum Computing: Progress and Challenges. *Appl. Phys. Rev.* **2019**, *6*, 021314. [[CrossRef](#)]

9. Krantz, P.; Kjaergaard, M.; Yan, F.; Orlando, T.P.; Gustavsson, S.; Oliver, W.D. A Quantum Engineer's Guide to Superconducting Qubits. *Appl. Phys. Rev.* **2019**, *6*, 021318. [\[CrossRef\]](#)
10. Henriot, L.; Beguin, L.; Signoles, A.; Lahaye, T.; Browaeys, A.; Raymond, G.-O.; Jurczak, C. Quantum Computing with Neutral Atoms. *Quantum* **2020**, *4*, 327. [\[CrossRef\]](#)
11. Burkard, G.; Ladd, T.D.; Pan, A.; Nichol, J.M.; Petta, J.R. Semiconductor Spin Qubits. *Rev. Mod. Phys.* **2023**, *95*, 025003. [\[CrossRef\]](#)
12. Steinacker, P.; Dumoulin Stuyck, N.; Lim, W.H.; Tantt, T.; Feng, M.; Serrano, S.; Nickl, A.; Candido, M.; Cifuentes, J.D.; Vahapoglu, E.; et al. Industry-Compatible Silicon Spin-Qubit Unit Cells Exceeding 99% Fidelity. *Nature* **2025**, *646*, 81–87. [\[CrossRef\]](#) [\[PubMed\]](#)
13. AbuGhanem, M. Toward Scalable Fault-Tolerant Photonic Quantum Computers. *J. Supercomput.* **2026**, *82*, 51. [\[CrossRef\]](#)
14. Brady, A.J.; Eickbusch, A.; Singh, S.; Wu, J.; Zhuang, Q. Advances in Bosonic Quantum Error Correction with Gottesman–Kitaev–Preskill Codes: Theory, Engineering and Applications. *Prog. Quantum Electron.* **2024**, *93*, 100496. [\[CrossRef\]](#)
15. Riverlane. *The Quantum Error Correction Report 2024*; Riverlane: Cambridge, UK, 2024.
16. Riverlane. *The Quantum Error Correction Report 2025*; Riverlane: Cambridge, UK, 2025.
17. Google Quantum AI. Quantum Error Correction Below the Surface Code Threshold. *Nature* **2025**, *638*, 920–926. [\[CrossRef\]](#) [\[PubMed\]](#)
18. Bluvstein, D.; Geim, A.A.; Li, S.H.; Evered, S.J.; Bonilla Ataides, J.P.; Baranes, G.; Gu, A.; Manovitz, T.; Xu, M.; Kalinowski, M.; et al. A Fault-Tolerant Neutral-Atom Architecture for Universal Quantum Computation. *Nature* **2026**, *649*, 39–46. [\[CrossRef\]](#) [\[PubMed\]](#)
19. Bravyi, S.; Cross, A.W.; Gambetta, J.M.; Maslov, D.; Rall, P.; Yoder, T.J. High-Threshold and Low-Overhead Fault-Tolerant Quantum Memory. *Nature* **2024**, *627*, 778–782. [\[CrossRef\]](#) [\[PubMed\]](#)
20. Leclerc, L.; Julià-Farré, S.; Silva Freitas, G.; Villaret, G.; Albrecht, B.; Béguin, L.; Bourachot, L.; Briosne-Frejaville, C.; Claveau, D.; Cornillot, A.; et al. One-to-one quantum simulation of the low-dimensional frustrated quantum magnet TmMgGaO₄ with 256 qubits. *arXiv* **2026**, arXiv:2603.20372. [\[CrossRef\]](#)
21. Lee, Y.T.; Kumaran, K.; Pokharel, B.; Scheie, A.; Sarkis, C.L.; Tennant, D.A.; Humble, T.; Schleife, A.; Kandala, A.; Banerjee, A. Benchmarking quantum simulation with neutron-scattering experiments. *arXiv* **2026**, arXiv:2603.15608. [\[CrossRef\]](#)
22. McEwen, M.; Kafri, D.; Chen, Z.; Atalaya, J.; Satzinger, K.J.; Quintana, C.; Klimov, P.V.; Sank, D.; Gidney, C.; Fowler, A.G.; et al. Removing Leakage-Induced Correlated Errors in Superconducting Quantum Error Correction. *Nat. Commun.* **2021**, *12*, 1761. [\[CrossRef\]](#) [\[PubMed\]](#)
23. Turing, A.M. Computing Machinery and Intelligence. *Mind* **1950**, *59*, 433–460. [\[CrossRef\]](#)
24. Lakatos, I. *The Methodology of Scientific Research Programmes: Philosophical Papers, Volume 1*; Cambridge University Press: Cambridge, UK, 1978.
25. Latour, B. *Reassembling the Social: An Introduction to Actor-Network-Theory*; Oxford University Press: Oxford, UK, 2005.
26. Riordan, M.; Hoddeson, L. *Crystal Fire: The Birth of the Information Age*; W. W. Norton: New York, NY, USA, 1997.
27. von Neumann, J. *The Computer and the Brain*; Yale University Press: New Haven, CT, USA, 1958.
28. LaRose, R. A Brief History of Quantum vs Classical Computational Advantage. *arXiv* **2024**, arXiv:2412.14703. [\[CrossRef\]](#)
29. Eisert, J.; Preskill, J. Mind the Gaps: The Fraught Road to Quantum Advantage. *arXiv* **2025**, arXiv:2510.19928. [\[CrossRef\]](#)
30. Feynman, R.P. Simulating Physics with Computers. *Int. J. Theor. Phys.* **1982**, *21*, 467–488. [\[CrossRef\]](#)
31. Jurczak, C. Investing in the Quantum Future: State of Play and Way Forward for Quantum Venture Capital. *arXiv* **2023**, arXiv:2311.17187. [\[CrossRef\]](#)
32. Ihde, D. *Technology and the Lifeworld: From Garden to Earth*; Indiana University Press: Bloomington, IN, USA, 1990.
33. Ihde, D. *Instrumental Realism: The Interface Between Philosophy of Science and Philosophy of Technology*; Indiana University Press: Bloomington, IN, USA, 1991.
34. Anderson, P.W. More Is Different: Broken Symmetry and the Nature of the Hierarchical Structure of Science. *Science* **1972**, *177*, 393–396. [\[CrossRef\]](#) [\[PubMed\]](#)
35. Aspect, A.; Dalibard, J.; Roger, G. Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell's Inequalities. *Phys. Rev. Lett.* **1982**, *49*, 1804–1807. [\[CrossRef\]](#)
36. Zurek, W.H. Decoherence, Einselection, and the Quantum Origins of the Classical. *Rev. Mod. Phys.* **2003**, *75*, 715–775. [\[CrossRef\]](#)
37. Zurek, W.H. *Decoherence and Quantum Darwinism: From Quantum Foundations to Classical Reality*; Cambridge University Press: Cambridge, UK, 2025. [\[CrossRef\]](#)
38. Turner, C.J.; Michailidis, A.A.; Abanin, D.A.; Serbyn, M.; Papić, Z. Weak Ergodicity Breaking from Quantum Many-Body Scars. *Nat. Phys.* **2018**, *14*, 745–749. [\[CrossRef\]](#)
39. Serbyn, M.; Abanin, D.A.; Papić, Z. Quantum Many-Body Scars and Weak Breaking of Ergodicity. *Nat. Phys.* **2021**, *17*, 675–685. [\[CrossRef\]](#)
40. Aires, S. On the Individuation of Complex Computational Models: Gilbert Simondon and the Technicity of AI. *AI Soc.* **2025**, *40*, 3109–3122. [\[CrossRef\]](#)

41. Rijssenbeek, J.; Blok, V. New Encounters Between Life and Technology: Simondon and the Case of Synthetic Biology. *Found. Sci.* **2025**. [[CrossRef](#)]
42. Loughnane, C. The Technical Milieu and Its Evolution: Uexküll, Kapp, Cassirer, Simondon. *Humanit. Soc. Sci. Commun.* **2025**, *12*, 1329. [[CrossRef](#)]
43. Saffman, M.; Walker, T.G.; Mølmer, K. Quantum Information with Rydberg Atoms. *Rev. Mod. Phys.* **2010**, *82*, 2313–2363. [[CrossRef](#)]
44. Browaeys, A.; Lahaye, T. Many-Body Physics with Individually Controlled Rydberg Atoms. *Nat. Phys.* **2020**, *16*, 132–142. [[CrossRef](#)]
45. Christen, I.; Propson, T.; Sutula, M.; Sattari, H.; Choong, G.; Panuski, C.; Melville, A.; Mallek, J.; Brabec, C.; Hamilton, S.; et al. An Integrated Photonic Engine for Programmable Atomic Control. *Nat. Commun.* **2025**, *16*, 82. [[CrossRef](#)] [[PubMed](#)]
46. Hui, Y. *The Question Concerning Technology in China: An Essay in Cosmotechnics*; Urbanomic: Falmouth, UK, 2016.
47. Pickering, A. *The Mangle of Practice: Time, Agency, and Science*; University of Chicago Press: Chicago, IL, USA, 1995.
48. Mallaby, S. *The Power Law: Venture Capital and the Making of the New Future*; Penguin Press: New York, NY, USA, 2022.
49. LeCun, Y. A Path Towards Autonomous Machine Intelligence. Position Paper, Version 0.9.2. 27 June 2022. Available online: <https://openreview.net/forum?id=BZ5a1r-kVsf> (accessed on 1 July 2026).
50. Plé, T.; Adjoua, O.; Lagardère, L.; Piquemal, J.P. A Foundation Model for Accurate Atomistic Simulations in Drug Design. *ChemRxiv* **2025**. [[CrossRef](#)]
51. Benali, A.; Plé, T.; Adjoua, O.; Agarawal, V.; Applencourt, T.; Blazhynska, M.; Clay, R., III; Gasperich, K.; Hossain, K.; Kim, J.; et al. Pushing the Accuracy Limit of Foundation Neural Network Models with Quantum Monte Carlo Forces and Path Integrals. *arXiv* **2025**, arXiv:2504.07948. [[CrossRef](#)]
52. Baum, Y.; Amico, M.; Howell, S.; Hush, M.; Liuzzi, M.; Mundada, P.; Merkh, T.; Carvalho, A.R.R.; Biercuk, M.J. Experimental Deep Reinforcement Learning for Error-Robust Gate-Set Design on a Superconducting Quantum Computer. *PRX Quantum* **2021**, *2*, 040324. [[CrossRef](#)]
53. Daley, A.J.; Bloch, I.; Kokail, C.; Flannigan, S.; Pearson, N.; Troyer, M.; Zoller, P. Practical Quantum Advantage in Quantum Simulation. *Nature* **2022**, *607*, 667–676. [[CrossRef](#)] [[PubMed](#)]
54. Ansari, N.; Feniou, C.; Gouraud, N.; Loco, D.; Badreddine, S.; Claudon, B.; Aviat, F.; Blazhynska, M.; Gasperich, K.; Michel, G.; et al. The Convergence Frontier: Integrating Machine Learning and High Performance Quantum Computing for Next-Generation Drug Discovery. *arXiv* **2026**, arXiv:2603.17790. [[CrossRef](#)]
55. Hui, Y. *Recursivity and Contingency*; Rowman & Littlefield: London, UK, 2019.
56. Latour, B. Why Has Critique Run Out of Steam? From Matters of Fact to Matters of Concern. *Crit. Inq.* **2004**, *30*, 225–248. [[CrossRef](#)]
57. Coeckelbergh, M. Democracy, Epistemic Agency, and AI: Political Epistemology in Times of Artificial Intelligence. *AI Ethics* **2023**, *3*, 1341–1350. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.