

Article

Not peer-reviewed version

Physics of Mindful Knowledge

[Rao Mikkilineni](#)* and [Max Michaels](#)

Posted Date: 19 November 2025

doi: 10.20944/preprints202510.1913.v2

Keywords: knowledge constraints; mind; metastability; teleonomy; General Theory of Information; hard-to-vary explanations; structural machines; coherence depth; dissipative structures



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

Article

Physics of Mindful Knowledge

Rao Mikkilineni ^{1,*} and Max Michaels ²

¹ Center for Business Innovation. Golden Gate University, San Francisco, USA

² CRYZTAL CAPITAL, Dallas, USA

* Correspondence: rmikkilineni@ggu.edu

Abstract

Contemporary physics explains matter and energy with extraordinary success yet offers no principled account of how knowledge and mind arise and act causally in nature. We propose a Physics of Mindful Knowledge (PMK) in which knowledge-bearing constraints are recognized as organizing principles that stabilize metastable coherence and guide system dynamics. PMK integrates three pillars: (i) Burgin’s General Theory of Information (GTI), which formalizes information as a triadic relation (carrier–content–recipient), avoiding reification; (ii) Fold Theory, modeling the emergence and top-down efficacy of coherent structures; and (iii) Deutsch’s epistemology, framing knowledge as “hard-to-vary explanations” and a constructor-level causal resource, using the Burgin–Mikkilineni Thesis (BMT), which reconceives computation as structural evolution and provides an engineerable path to oracle-like behavior with explanatory, predictive structure and teleonomy. We contrast PMK with microphysical (substrate-dependent) and beyond-physical (metaphysical) approaches and derive falsifiable predictions and measurable signatures—including Recovery Energy, Entropy per Useful Work, and Coherence Depth. We translate these into engineering protocols using structural machines (Mindful Machine Architecture) and outline comparative tests versus state-of-the-art algorithmic systems under structural, functional, and epistemic perturbations. By embedding epistemic constraints into a naturalistic, computable framework, PMK provides a unified, testable bridge across physics, biology, and intelligent systems, advancing a practical science of knowledge, and mind. Physics of Mindful Knowledge (PMK) studies how matter–energy systems instantiate, transform, and use knowledge in mindful ways—meaning they monitor, evaluate, and adjust their own operations to preserve identity, purpose, and coherence.

Keywords: knowledge constraints; mind; metastability; teleonomy; General Theory of Information; hard-to-vary explanations; structural machines; coherence depth; dissipative structures

1. Introduction

Physics explains matter and energy with extraordinary precision, yet remains silent on phenomena central to life and intelligence: knowledge and mind—the apparent goal-directedness of living systems. Why is there such an explanatory gap between our fundamental theories’ success in predicting planetary orbits and their failure to explain the adaptive, goal-directed behavior of organisms or intelligent machines?

This gap is not trivial. Without a physics of knowledge and mind, we lack a principled foundation for explaining and engineering systems that maintain coherence, adapt to change, and act with purpose. Teleonomy—the ability to pursue goals and preserve metastability—is essential for understanding resilience in biology and for designing trustworthy AI. Existing approaches fall short: microphysical theories seek exotic substrates, metaphysical proposals abandon empirical testability, and process philosophies, while conceptually rich, lack a formal, computable language to operationalize their insights. We define PMK as the study of how physical systems embody, monitor, and regulate knowledge in ways that preserve coherence, purpose, and teleonomic integrity.

Our contribution is not to rediscover this problem but to propose a formal, testable, and engineerable bridge: the Physics of Mindful Knowledge (PMK). PMK treats knowledge—not as an epiphenomenon—but as a causal constraint embedded in physical systems. Building on Burgin's General Theory of Information (GTI), the Burgin–Mikkilineni Thesis (BMT), Fold Theory, and Deutsch's epistemology, we show how epistemic structures stabilize metastability and enable teleonomy.

We must clarify our terms. When we speak of "mind" in this context, we are not, for the purposes of this paper, addressing the "hard problem" of subjective experience or consciousness. We define "mind" in a functional, naturalistic sense: it is the process of modeling and anticipation that enables a system to act with apparent purpose to maintain its own coherence and achieve goals.

In asking this question, we are keenly aware that we are not the first. As physicists and engineers, we are, in many respects, "late-comers to the party." A vast and deep literature on this exact subject exists in fields outside of physics [1], most notably in process philosophy [2–5], complex systems theory [6–9] and biosemiotics [10]. These thinkers have long and rigorously argued for the need to incorporate information, meaning, and process into our fundamental theories of the natural world.

The central challenge, we argue, has not been a lack of philosophical insight but the absence of a computable, physical language to operationalize these concepts. Our intended contribution is to offer this formal, testable, and engineerable bridge. The PMK framework is designed to connect the deep insights of these process and semiotic traditions to the practical fields of modern physics and artificial intelligence, providing a language where "knowledge" is no longer just a concept, but a causal, physical, and engineerable constraint.

The remainder of this paper proceeds as follows: Section 2 surveys existing paradigms; Section 3 introduces PMK and its formal foundations; Section 4 compares PMK with competing frameworks; Section 5 outlines falsifiable predictions and experimental protocols; and Section 6 discusses implications for physics, AI, and philosophy. Section 7 concludes with our observations and comparison of our engineerable architecture to realize them in non-biological substrates with Friston's Free Energy principle and Froese's enactive tradition.

2. The Existing Debate in Physics

The explanatory gap we identified—the causal efficacy of knowledge and teleonomy—is not unaddressed. The problem has, in fact, been approached from several powerful, yet often non-interacting, perspectives. A successful framework must synthesize these disparate lines of inquiry.

2.1. The Microphysical/Substrate-Dependent Approach

Based on Gödel's incompleteness theorems, Roger Penrose ruled out computational sufficiency: human mathematical insight, he claimed, cannot be reduced to formal rule-following [11]. To ground this non-computability in physics, Penrose introduced the concept of Objective Reduction (OR)—a hypothesized collapse of the quantum state driven by gravitational effects at the Planck scale [12]. Together with Stuart Hameroff, he developed the Orch-OR model, which proposes that neuronal microtubules act as orchestrators of such non-computable collapses. Conscious experience, in this view, arises from orchestrated episodes of quantum gravitational collapse inside neurons.

While Orch-OR boldly asserts that consciousness depends on new physics, the theory remains speculative. Critics have noted that decoherence times in warm, wet brains are too short to preserve quantum coherence [13]. Empirical evidence for microtubule-level quantum computation remains limited, and the link to consciousness is untested. Nevertheless, Orch-OR illustrates a broader theme: that physics as currently formulated may be missing a layer critical to explaining the mind.

However, the search for a specific, non-computable substrate continues. Recent research has proposed other candidates. Work by Adamatzky [14], Schumann [15,16], and others, for instance, suggests that actin filaments and even fungal networks could act as "Turing automata with oracles," capable of solving non-computable functions. This line of inquiry is promising. However, it remains focused on what substrate mind is made of. We argue this focus may be misplaced. The problem may

not be the substrate, but the architecture. A theory of knowledge should be substrate-independent, capable of explaining a "mindful" process whether it runs on neurons, actin filaments, or silicon.

2.2. Rickles and the Beyond-Physical Ontology

A second approach, seeing no solution within current naturalism, proposes a "metaphysical rupture." This view is represented by scholars like Dean Rickles, who argues that the failure of quantum gravity research reveals not just a technical gap but an ontological one. In works such as *The Philosophy of Physics* [17], this framework explores the profound challenge of mapping formal physical theories onto a single, coherent reality. The core argument is that physics is bounded by the assumptions of materialism—assumptions that cannot accommodate mind or meaning. This "R" archetype suggests the next paradigm must go "beyond the physical," re-centering the observer and mind as foundational rather than derivative. This resonates with earlier proposals such as Wheeler's "It from Bit" [18] which reframe matter as emergent from deeper informational or implicate processes.

The strength of this view lies in its willingness to question entrenched materialist commitments. We agree with the R framework's diagnosis: that a naive, purely microphysical account is insufficient. However, we diverge fundamentally on the prescription. The R framework remains a philosophical and interpretative project, exploring what physics should mean. Its weakness is that it offers little in the way of testable hypotheses, risking a slide into metaphysics rather than extending physics. Our Physics of Knowledge and Mind, in contrast, is a naturalistic and engineerable proposal. We do not posit a "mind-first" ontology; we provide a formal, testable mechanism for how knowledge and "mindfulness" emerge from and act causally upon a physical substrate. PMK is not an interpretation of physics but a falsifiable extension of it.

2.3. The Process & Semiotic Foundation

The entire discussion of information and meaning in nature has a deep, rich history outside of 20th-century physics. Process philosophy, championed by thinkers like Bergson [2], Whitehead [5], and more recently Kauffman [3], argues that reality is not composed of static "things" but of dynamic "processes." In this view, constraints and relations are just as fundamental as matter and energy.

Complementing this is the field of biosemiotics, founded on the work of Charles Sanders Peirce [10]. Peirce argued that all processes of meaning (semiosis) are fundamentally triadic: a relation between a Sign (or representamen), an Object, and an Interpretant (the effect or "meaning" produced). This tradition, we believe, has the correct philosophical orientation. It understands that mind is a process and that meaning is relational. Its primary historical limitation, however, has been the lack of a formal, computable, and engineerable language to connect these profound philosophical insights to the mathematics of physics and the architecture of computation.

2.4. The Informational-Biological Validation

Finally, a body of recent work in theoretical and experimental biology has provided powerful empirical validation for an information-centric view. Terrence Deacon, in *Incomplete Nature*, develops a rigorous account of "absence causation," where constraints (what is not present) exert real causal power, forming a nested hierarchy that enables life and, ultimately, teleonomy [19].

Even more strikingly, Michael Levin's work on "scale-invariant cognition" demonstrates this process in action. Levin has shown that bioelectric fields in living tissues function as an informational layer, a "knowledge-bearing constraint" that stores the "memory" of a correct body plan [20]. This informational layer governs the behavior of the underlying cells, enabling stunning displays of teleonomy, such as the regeneration of a planarian's head on a different part of its body.

2.5. A Fragmented Landscape

This survey reveals a fragmented landscape. We have:

- Substrate-Hunters (Penrose, Adamatzky) looking for a "magic" material.
- Metaphysicians (Rickles, Hoffman) proposing non-physical, "mind-first" realities.
- Process Philosophers & Semioticians (Peirce, Whitehead) with the right relational philosophy but lacking a formal, computable framework.

Informational Biologists (Deacon [19], Levin [20]) who are empirically proving that informational constraints do have causal power in living systems, but without yet providing a universal, engineerable architecture.

A synthesis is urgently needed. The Physics of Mindful Knowledge (PMK), which we introduce in the next section, is proposed as this synthesis. It is a naturalistic framework that is philosophically grounded in the semiotic and process traditions, biologically validated by work like Levin's, and—critically—made formally computable and engineerable through the tools of GTI and BMT.

3. The Physics of Mindful Knowledge: A Formal Framework

The gap between microphysical laws and teleonomic organization concerns a missing level of description—a physics of organization—rather than an undiscovered microphysical interaction. Microphysics supplies the lawful substrate; organizational physics specifies how constraints (structures, knowledge) select and stabilize macro-dynamics within thermodynamic limits [8].

The fragmented landscape described in Section 2 reveals the need for a framework that is naturalistic (not metaphysical), substrate-independent (not microphysical), formally rigorous (unlike process philosophy alone), and empirically validated (like the phenomena in informational biology). The Physics of Mindful Knowledge (PMK), is proposed as this synthesis. PMK is a naturalistic framework where knowledge—defined as explanatory, predictive, and causal constraints—is a fundamental part of physical reality. This framework is built on two core components: Burgin's General Theory of Information (GTI) as its ontology and the Burgin-Mikkilineni Thesis (BMT) as its dynamics.

Physics of Mindful Knowledge (PMK) studies how knowledge-bearing structures—in biological, social, or digital systems—are physically realized in matter–energy substrates and how they constrain system dynamics.

By *mindful*, we mean that these knowledge processes are not merely reactive mappings but self-referential and purpose-laden: the system maintains models of itself and its environment, evaluates its own performance against goals and norms, and modifies its behavior accordingly.

Thus, PMK treats *mind* as a particular regime of organization within physical systems: a regime in which knowledge structures include models of the system's own state, capabilities, and purposes, and are used to regulate ongoing processes. This links epistemic and ontological structures, reframing knowledge and mind as general physical phenomena—not human-exclusive attributes.

3.1. GTI as Formal Semiotics: A Solution to Reification

The most difficult challenge, as noted by critics, is the "reification of constraints." How can information be causal without becoming a "third substance" alongside matter and energy? The solution has existed for over a century, first in biosemiotics and now in information theory. Charles Sanders Peirce defined a "sign" (the carrier of meaning) as an irreducible triadic relation between an Object, a Representamen (the sign itself), and an Interpretant (the effect or meaning) [10].

Generations later, Mark Burgin's General Theory of Information (GTI) covered in four books [21–24] provides the formal, mathematical, and computable specification of this same semiotic process. GTI defines information not as a thing, but as a triadic relation between:

The Carrier: The physical substrate or structure (matter/energy).

The Content: The "knowledge" or constraint (the form or absence that makes a difference, as in Deacon's work [19])

The Recipient: The system that "reads" or is affected by the content, generating the "Interpretant."

This triadic model is our core defense against reification. Information has no existence apart from its embedding in organizational structures. The "knowledge-bearing constraint" (the content) is not an independent entity. It is immanent in its physical carrier and only exerts causal force by being interpreted by a recipient. A simple funnel's shape is a constraint that is causal, but it is not a "third substance" alongside the plastic (carrier) and the water (recipient). GTI is the formal language for this immanent, relational causation. It is the physical instantiation of biosemiotics.

The central claim of a The Physics of Mindful Knowledge, is that knowledge-bearing structures are not epiphenomenal but causal constraints within physical systems. In Mark Burgin's General Theory of Information (GTI), information is defined not as an abstract quantity but as a triadic relation between a carrier, content, and recipient [21,23]. This relational view implies that information has no existence apart from its embedding in organizational structures. Extending this logic, knowledge—understood as explanatory, hard-to-vary models in David Deutsch's terms [25]—can be seen as a higher-order constraint on system dynamics. Knowledge restricts the space of possible transformations, channeling processes toward goal-directed outcomes. Unlike mere Shannon information, which is statistical, knowledge is counterfactual: it encodes not only what is, but what could or could not be otherwise. This property makes knowledge uniquely suited to stabilize metastable coherence in far-from-equilibrium systems, preventing collapse into entropy-driven attractors. Thus, knowledge-bearing constraints can be treated as a new kind of physical resource, shaping system trajectories in a way that complements energy and matter.

3.2. Structural Machines and the Burgin–Mikkilineni Thesis

If GTI provides the ontology of knowledge, how does it act? How does it compute and evolve? The "substrate-dependent" approaches (Penrose, Adamatzky) are correct that the answer must be "non-computable" in the classical sense [14,16]. Their search for exotic physics, however, misses the formal point made by Alan Turing himself.

In 1939, Turing proposed "Systems of Logic Based on Ordinals" [26], which introduced the "oracle machine." An oracle machine is a standard Turing machine augmented with a "black box"—an oracle—that can solve a non-computable problem (like the Halting Problem) in a single step. Turing introduced the oracle to define the very limit of algorithmic computation. Later, there were many extensions of the oracles to extend the original idea [27–31].

Here, we propose that the Burgin–Mikkilineni Thesis (BMT) [28,30] provides the architectural basis for a physical oracle. The BMT redefines computation:

- Classical Computation: The processing of data (symbols) by a fixed algorithm (a static structure).
- BMT/Structural Computation: The evolution and interaction of the structures themselves.

In a "structural machine" [22,23,32] built on the BMT, the "knowledge" is not a static program in the machine; it is the machine's architecture. The computation is the continuous, self-organizing evolution of that architecture. This structural process acts as a physical oracle. It does not "compute" an answer algorithmically; its self-organizing dynamics embody the answer as a stable, structural state. This provides a naturalistic, engineerable path to the "non-computable" functions that Penrose, Schumann, and others seek in exotic substrates [12,16]. It is a new computational paradigm.

3.3. The Role of Fold Theory: Modeling Emergence and Coherence

A key challenge is to specify the dynamic process by which epistemic structures (knowledge) emerge from and subsequently "fold back" to govern their ontological (physical) substrate. This is the problem of emergence, organizational closure, and top-down causation.

If the General Theory of Information (GTI) provides the ontology (defining information as a triadic relation of Carrier/ontological and Content/epistemic), and the Structural Machine Architecture (SMA) is the engineered system, then Fold Theory [33] provides the physics or formal language to describe what this information does. Fold Theory is a mathematical framework designed to model how systems "fold" to create coherence. We can use the analogy of a whirlpool:

- Ontological Substrate (Carrier): The flow of water (the physical dynamics).

- Epistemic Structure (Content): A stable whirlpool pattern that emerges from the water's own dynamics. This "fold" is not a new substance but a coherent, self-maintaining pattern.
- Causal Efficacy: This emergent "fold" (the whirlpool) now exerts top-down causal power, governing the behavior of the individual water molecules and pulling them into its stable pattern.

Fold Theory provides the formal mathematics to describe how this self-creating, self-maintaining "fold" arises, persists, and gains causal efficacy. In the context of our framework, Fold Theory is the "physics of emergence" for the Mindful Machine. It formally explains how the machine's knowledge (the Epistemic Structure) emerges from its hardware's dynamics (the Ontological Structure) to become a real, causal, and governing part of the system's operation, thus unifying the static ontology of GTI with the dynamic, teleonomic function of the BMT.

3.4. *Deutsch's Epistemology: Knowledge as a Physical Constructor*

While GTI provides the ontology for information and BMT provides the computational dynamics, the epistemology of David Deutsch provides the physical and philosophical justification for our central claim: that knowledge is a real and potent causal force. Deutsch, in *The Beginning of Infinity* and his work on Constructor Theory, argues that the laws of physics are not just about what happens (dynamics) but about what can be made to happen (which transformations are possible vs. impossible) [34]. The entities that cause these transformations are "constructors."

3.4.1. Knowledge as "Hard-to-Vary Explanations"

Deutsch redefines knowledge. It is not "justified true belief," but rather "hard-to-vary explanations." A good explanation (a good theory) is one whose parts are so interconnected and constrained by evidence that you cannot easily change one part without destroying the whole. This directly parallels our framework:

- Deutsch's "hard-to-vary explanation" is the epistemic component.
- Our "knowledge-bearing constraint" or "stable, coherent structure" (modeled by Fold Theory [33]) is the physical instantiation of that hard-to-vary information.

A Mindful Machine's "schema" or "digital genome" is a hard-to-vary explanation of its purpose and its environment, encoded in a physical structure.

3.4.2. Knowledge as the Ultimate Constructor

For Deutsch, knowledge is not abstract; it is a type of constructor. A simple constructor, like a catalyst, can facilitate a transformation, but it doesn't know anything. A system with knowledge (like a 3D printer with a blueprint, or a human) is a universal constructor. It can cause an indefinitely wide range of transformations. Deutsch argues that knowledge is, in fact, the most significant causal force in the universe. It is the only thing that can bend the physical world to its will, creating structures and processes (like computers or cities) that would be statistically impossible to arise by "chance and necessity" alone [34].

3.4.3. The BMT as an Engineered Constructor

This is where Deutsch's epistemology provides profound support for our paper's engineering thesis:

- A Mindful Machine, built on the Burgin-Mikkilineni Thesis (BMT) [35], is the engineering instantiation of a Deutsch-ian universal constructor.
- The "knowledge-bearing constraints" (the epistemic structures) are not just in the machine; they are the machine.

The system's ability to self-repair and maintain its coherence—as we propose to test with the Recovery Energy metric—is the very definition of a true constructor: a system that persists in its function (its knowledge) while causing transformations in its environment.

In summary, Deutsch's epistemology provides the physical and philosophical argument why we must treat knowledge as a causal, structural entity. It shifts physics from a "physics of dynamics" to a "physics of constraints and possibility." Our PMK framework, with its synthesis of GTI, BMT, and Fold Theory, provides the formal language and the engineerable architecture to build these knowledge-based constructors.

3.5. Metastability and Teleonomy

A unifying concept across GTI, Fold Theory, and BMT is that of metastability. Knowledge-bearing systems remain poised between equilibrium and chaos, able to flexibly adapt without disintegration. Teleonomy—the apparent goal-directedness of living systems [36]—can be reframed in this context as the causal efficacy of knowledge constraints. This reframing has profound implications: teleonomy is not a mystical property of life, nor reducible to blind statistical regularities, but a manifestation of physics extended to include knowledge as an organizing principle. The Physics of Mindful Knowledge therefore provides a naturalistic account of mind and purpose without recourse to exotic microphysics [14] or metaphysical rupture [12].

When you combine GTI (knowledge as a triadic relation) with BMT (computation as structural evolution), you get the function of PMK: to stabilize metastable coherence and enable teleonomy. A living system, like those described by Deacon and Levin [19,20], is a far-from-equilibrium system. It is "metastable"—it could collapse into thermodynamic equilibrium (death) at any moment. What prevents this collapse? The Answer: A nested set of knowledge-bearing constraints (encoded via GTI) that compute (via BMT) to actively maintain the system's coherence.

Figure 1 shows a Formal Bridge from Foundational Insights to Testable Engineering. The layered figure of the Structural Machine (Mindful Machine Architecture) presents the three tiers—Physical Substrate, Knowledge Constraints, and Teleonomy & Mindful Systems.




Figure 1. Layered Mindful Machine Architecture.

4. Comparative Framework

The debate over “missing physics” reveals three distinct loci of explanation: microphysics (Penrose–Hameroff), metaphysics (Rickles), and knowledge constraints (Physics of Mindful Knowledge). Each identifies a real gap, but their solutions diverge in explanatory scope, testability, and engineering implications.

Table 1 below summarizes the three theories using eight dimensions.

Table 1. Three explanations for the gap in physics.

Dimension	Penrose-Hameroff (Orch OR)	Rickles (Beyond Physics)	Physics of Mindful Knowledge
Location of the Gap	Exotic microphysical processes (quantum gravity collapse in microtubules)	Limits of materialism itself; need for pre-physical ontology	Missing organizational principles — knowledge-bearing constraints shaping dynamics
Mode of Explanation	Non-computable collapse events yield conscious moments	Mind/observer as ontological primitive; physics must be transcended	Constraints encode explanatory knowledge; metastable coherence sustains teleonomy
Relation to Self-Organization	Distinct from dissipative structures; relies on unknown physics	Rejects physicalist self-organization as inadequate	Extends self-organization by embedding epistemic constraints within physical law
Testability	Limited; speculative, challenged by decoherence times	Largely untestable; risks metaphysical drift	Offers measurable predictions: recovery energy, entropy per useful work, coherence depth
Engineering Implications	No clear engineering pathway	Philosophical, not engineerable	Structural machines, digital genomes, mindful AI architectures provide realizations
Strengths	Highlights insufficiency of computation and current physics	Radically challenges materialist assumptions	Naturalistic, testable, engineerable; unifies physics, life, and mind
Weaknesses	Highly speculative; no direct evidence	Risks abandoning empirical science	Still emerging; requires empirical validation

5. From Theory to Testability

A central claim of this paper is that the Physics of Mindful Knowledge (PMK), unlike metaphysical or purely substrate-dependent approaches, is "empirically testable and engineerable." This section provides a detailed elaboration. The "dynamical reciprocity between structure, model, and process" is specified by the Burgin-Mikkilineni Thesis (BMT) and its embodiment in structural machines [5,6,23]. The metrics are not pre-supposed; they are the direct, measurable consequences of comparing this BMT model against the reigning algorithmic (Turing) paradigm. To demonstrate this, we will "work out" our proposed testable signature using concrete implementations [23].

PMK perspective: *Mind arises when knowledge becomes self-referential and teleonomic—regulating matter–energy dynamics through explanatory, purpose-guided constraints.*

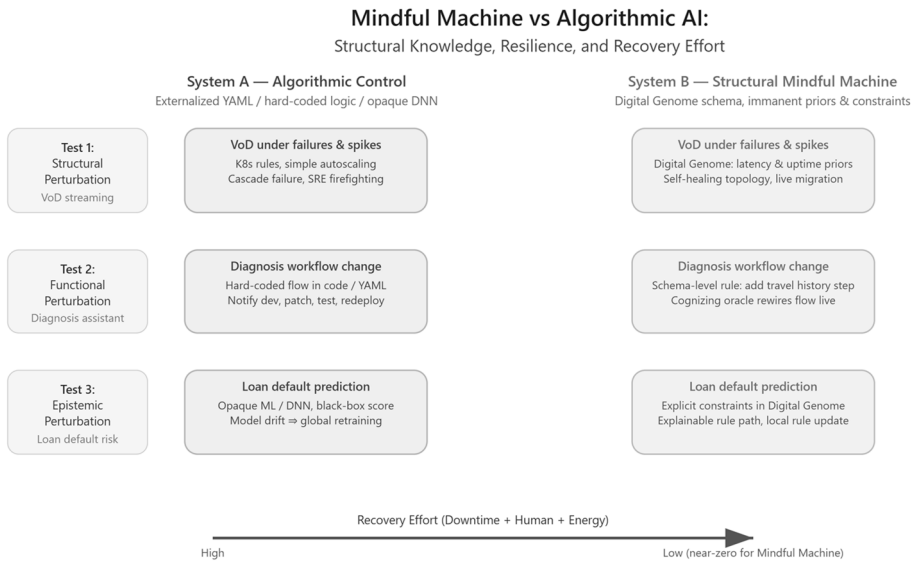


Figure: Three implemented Mindful Machine case studies (VoD, diagnosis assistant, loan default) illustrating the PMK prediction that structural Digital Genome control yields robust teleonomy and orders-of-magnitude lower recovery effort than algorithmic AI.

Figure 2. Testing Models.*5.1. A Concrete Testable Metric: Resilience and Recovery Effort*

Aim: To experimentally falsify the PMK hypothesis by comparing the resilience and operational efficiency of a system whose knowledge is externalized and algorithmic (System A) versus one whose knowledge is immanent and structural (System B).

Hypothesis (based on PMK): A structural machine (System B, the "Mindful Machine"), whose knowledge is embodied in its evolving architecture (its "Digital Genome" or "schema"), will demonstrate orders of magnitude less "Recovery Effort" (in compute, energy, and human intervention) in response to perturbations than an algorithmic system (System A), which must be manually or algorithmically patched, retrained, and redeployed.

5.2. Experimental Protocol & Comparative Tasks

We define our two systems and test them against real-world implementation challenges.

- **System A (Algorithmic Control):** A state-of-the-art implementation, typically using cloud-native microservices (managed by Kubernetes or other software structures) for application logic and a Deep Neural Network (DNN) for predictive logic. Its "knowledge" is brittle, existing externally in static YAML files, hard-coded workflows, or the opaque, static weights of a trained model.
- **System B (Structural Control):** A structural machine (the "Mindful Machine") built on the BMT. Its "knowledge" is an immanent, machine-readable "Digital Genome" that defines the intent, priors, and constraints of the system. The machine is its own architecture, which it evolves dynamically.

Test Case 1: Video on Demand (VoD) — Structural Perturbation

- **Task:** Both systems are instantiated to deliver a VoD streaming service.
- **Perturbation (Structural Damage):** The system is subjected to sudden, unpredictable hardware failures and massive, simultaneous user demand spikes, testing auto-failover, self-scaling, and live migration.
- **Falsifiable Prediction 1 (Resilience & Effort):**
 - System A (Algorithmic) will react based on pre-programmed rules (e.g., Kubernetes Horizontal Pod Autoscaler). These rules are simple, reactive, and slow. A cascade failure or a novel scaling pattern will overwhelm the rules, causing service failure (high downtime) and requiring high-effort human SRE (Site Reliability Engineer) intervention to diagnose and repair.
 - System B (Structural), in contrast, is its knowledge. The "Digital Genome" defines the intent (e.g., "maintain 5ms latency and 99.999% uptime"). It will autonomously and proactively self-organize and repair, dynamically re-routing workflows and re-instantiating processes on healthy nodes. Its Recovery Effort (in downtime and human intervention) will be near zero, while its "energy" cost is simply the minimal computational overhead of its own self-organization.

Test Case 2: Digital Assistant for Diagnosis — Functional Perturbation

- **Task:** Both systems execute a digital assistant workflow for early medical diagnosis.
- **Perturbation (Functional Damage):** A new healthcare regulation or medical best-practice is introduced, changing the required workflow (e.g., "a 'Travel History' check must now be added before the 'Symptom Check'").
- **Falsifiable Prediction 2 (Teleonomy & Adaptation):**
 - System A (Algorithmic) will exhibit extreme brittleness. The workflow is hard-coded by a developer. To adapt, a human must be notified, re-write the application code, run tests, and manage a complex redeployment. The system is "ignorant" and cannot adapt.
 - System B (Structural) will exhibit robust teleonomy. The "Digital Genome" (the schema) is simply updated with the new constraint. The machine itself, acting as a "cognizing oracle," instantly dynamically changes its own workflow structure to comply with the new rule, with

zero downtime. This is directly analogous to the hexapod "discovering" a new limping gait; System B solves the non-algorithmic problem of "what is the correct workflow now?"

Test Case 3: Loan Default Prediction — Epistemic Perturbation (ML vs. PMK)

- Task: Both systems are used to approve or deny loan applications based on default risk.
- Perturbation (The "Damage" of Opacity): This tests the fundamental difference between ML-based knowledge and structural knowledge. The systems are evaluated on explainability and adaptability to model drift.
- Falsifiable Prediction 3 (Explainability & Coherence):
 - System A (ML/DNN), upon being "damaged" (i.e., making a bad or biased decision), is functionally "ignorant" and "brittle." It is a "black box" and cannot provide a "hard-to-vary explanation" (per Deutsch [25]) for its decision. If its model drifts due to new market conditions, it must initiate a massive, high-energy, global retraining process, essentially re-learning "from scratch."
 - System B (Structural) is its knowledge. It makes a decision by following the explicit constraints in its "Digital Genome" (e.g., "debt-to-income ratio," "credit history," "regulatory rules"). It can explain its decision perfectly by providing the exact structural path it took. If a rule changes, it adapts dynamically (as in Test Case 2) without a high-energy "retraining."

5.3. Summary of Predictions

This single, unified framework demonstrates that the PMK is not a "stub" or an "opinion piece." It is a testable engineering paradigm. The implementations [35] show that a structural machine (System B) exhibits robust, teleonomic behavior and near-zero recovery effort, while the current algorithmic paradigm (System A) remains brittle, ignorant, and requires massive external (human or energy) effort to adapt or repair. The other metrics we proposed, such as Entropy per Useful Work and Coherence Depth, can be similarly mapped to these real-world applications.

6. Implications

The Physics of Mindful Knowledge carries implications across three domains: physics, AI, and philosophy. For physics, it extends physics with constraint-based ontology. For engineering and AI, it provides a roadmap for mindful machines and structural machines that embody explanatory knowledge. For philosophy, it offers a naturalist alternative to dualism and eliminativism, grounding teleonomy in causal constraints. This framework suggests a research program that bridges physics, AI, and philosophy, and opens pathways to testable science of knowledge.

In AI, the current drive toward a 'world model' aligns naturally with the emerging view of knowledge as a causal, structural constraint rather than an ethereal epiphenomenon. In LeCun's architecture, a world model is a learned, internal predictive structure that captures the regularities and causal dynamics of the environment. Its purpose is not passive representation but active constraint: it restricts the space of plausible futures, guides action selection, and enables planning without exhaustive trial-and-error. In this sense, the world model functions like a computational riverbank—shaping behavior by embedding learned structure within the system itself. Intelligent behavior emerges from the system's ability to simulate, anticipate, and optimize actions under these internal constraints.

PMK generalizes this idea beyond computation, reframing knowledge as a physically instantiated constraint that channels the flow of matter, energy, and processes, much like a topological structure that governs far-from-equilibrium dynamics. Building on General Theory of Information, Fold Theory, and Deutsch's epistemology, this framework treats knowledge as a causally efficacious substrate capable of stabilizing metastable states and generating teleonomic behavior. LeCun's world model can thus be seen as a specific instantiation of this broader ontology: an engineered informational constraint that acquires causal power through predictive coherence. We extend the concept further by introducing measurable physical signatures—such as coherence depth and recovery energy—and by emphasizing metacognition, self-maintenance, and structural

evolution. In this synthesis, world models provide the computational skeleton, while the Physics of Mindful Knowledge supplies the ontological and dynamical grounding for machines capable not merely of prediction, but of genuine resilience, self-correction, and mindful adaptation.

PMK provides the theoretical backdrop for the Mindful Machine Architecture (MMA) [37,38]. In PMK, a system is mindful when it (i) encodes prior knowledge in its structure, (ii) continuously senses and evaluates its state and context, and (iii) uses that knowledge to adaptively preserve its identity and purpose. The MMA instantiates this regime: the *Digital Genome* encodes priors and constraints, *AMOS* [35] governs self-maintenance, and *Cognizing Oracles* or large language models update contextual knowledge to steer control.

7. Conclusions

This paper began by identifying a specific explanatory gap in current physical theories: the causal efficacy of knowledge and teleonomy. In addressing this gap, we attempted to prove that our proposal is not merely a "brief draft" or "opinion piece," but a substantive, testable, and rigorously-grounded scientific framework.

We have framed our argument, not as a new "discovery," but as a necessary formal and engineerable bridge to a deep body of existing thought. We demonstrated that the philosophical foundations of our work lie in fields like biosemiotics [26] and process philosophy [16,38], which had correctly identified the relational, process-based nature of reality.

We then proposed the Physics of Mindful Knowledge (PMK) as the formal synthesis, a tradition that was missing, built on three interconnected pillars upon which BMT scaffolding is raised:

A Formal Ontology (GTI): We showed that Burgin's General Theory of Information [5–8] is a formal, computable instantiation of Peirce's triadic semiotics [10]. This solves the ontological problem of "reification" by defining knowledge as an immanent, relational constraint—a "content" (epistemic) that cannot exist apart from its "carrier" (ontological).

A Formal Dynamic (Fold Theory): We identified Fold Theory [33] as the formal language to describe the process of emergence. It models how the system's own dynamics "fold" to create and maintain coherent, stable knowledge structures, explaining how they arise from their physical substrate to exert top-down causal power.

A Physical Justification (Deutsch): We grounded our framework in Deutsch's epistemology [25], which defines knowledge as a physical "constructor" in the form of "hard-to-vary explanations." This provides the physical justification for treating our "knowledge-bearing constraints" as real, causal entities.

A Computable Architecture (BMT): We argued that the Burgin-Mikkilineni Thesis (BMT) [28,30] provides the naturalistic and engineerable architecture for Turing's "Oracle Machine" [26]. This "structural machine" (the implementations are called Mindful Machines [35]) is the physical embodiment of a Deutsch-ian constructor, one that computes by evolving its own hard-to-vary structure, offering a substrate-independent solution to non-algorithmic computation.

We further showed that this framework is not a mere abstraction. It is empirically validated by the "absential causation" described by Deacon [19] and, most strikingly, by the "scale-invariant cognition" and robust teleonomy demonstrated in Levin's biological experiments [20].

This 'physics of organization' is not without precedent. Karl Friston's Free Energy Principle formalizes this by defining the organism as a statistical model that minimizes 'surprisal' to resist entropic decay, effectively treating the organism's structure as a physical prediction of its environment [39]. Complementing this, Tom Froese and the enactive tradition argue that these constraints are not merely statistical but autopoietic, where the system's goal-directedness arises from the precarious, self-generating nature of its own boundaries [40]. Our PMK framework integrates these views: Friston provides the minimization calculus, Froese provides the autonomous dynamics, and the Burgin-Mikkilineni Thesis (BMT) provides the engineerable architecture to realize them in non-biological substrates.

Finally, we have demonstrated that the PMK is not a metaphysical proposal but a falsifiable engineering paradigm. By expanding our testable signatures into a detailed experimental protocol—Recovery Energy—we have provided a concrete, measurable, and falsifiable method for testing a structural, BMT-based controller against a traditional algorithmic one.

The implications are profound. The PMK framework suggests that matter, energy, and knowledge-bearing constraints (information) are an irreducible triad. It offers a new, non-binary, non-algorithmic model of computation—structural machines—that may be the key to building truly "mindful" artificial intelligence. By embedding epistemic constraints into physical law, the Physics of Mindful Knowledge unifies physics, life, and mind, offering a new, naturalistic, and powerfully predictive path forward.

Author Contributions: Conceptualization, R.M; methodology, R.M, M.M.; "All authors have read and agreed to the published version of the manuscript.

Funding: Please add: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: One of the authors. R.M acknowledge Judith Lee, Director of the Center for Business Innovation for many discussions and support.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bateson, G. (1972). Steps to an ecology of mind. University of Chicago Press.
2. Bergson, H. (1998). Creative evolution (A. Mitchell, Trans.). Dover Publications. (Original work published 1907)
3. Kauffman, S. A. (1993). The origins of order: Self-organization and selection in evolution. Oxford University Press.
4. Sulis, W. (2011). Generative processes: A new foundation for causal and informational theories. *Entropy*, 13(2), 525–553. <https://doi.org/10.3390/e13020525>
5. Whitehead, A. N. (1978). Process and reality: An essay in cosmology (D. R. Griffin & D. W. Sherburne, Eds.). The Free Press. (Original work published 1929)
6. Holland, J. H. (1998). Emergence: From chaos to order. Oxford University Press.
7. Maturana, H. R., & Varela, F. J. (1980). Autopoiesis and cognition: The realization of the living. D. Reidel Publishing Co.
8. Prigogine, I. (1977). Self-organization in nonequilibrium systems. Wiley.
9. Prigogine, I., & Stengers, I. (1984). Order out of chaos: Man's new dialogue with nature. Bantam Books.
10. Peirce, C. S. (1931–1958). Collected papers of Charles Sanders Peirce (Vols. 1–8; C. Hartshorne, P. Weiss, & A. W. Burks, Eds.). Harvard University Press.
11. Penrose, R. (1994). Shadows of the mind. Oxford University Press.
12. Hameroff, S., & Penrose, R. (2014). Consciousness in the universe: A review of the 'Orch OR' theory. *Physics of Life Reviews*, 11(1), 39–78. <https://doi.org/10.1016/j.plrev.2013.08.002>
13. Tegmark, M. (2000). The importance of quantum decoherence in brain processes. *Physical Review E*, 61(4), 4194–4206.
14. Adamatzky, A., & Gandia, A. (2022). Fungi anaesthesia. *Scientific Reports*, 12(1), 340. <https://doi.org/10.1038/s41598-021-04172-0>
15. Schumann, A. (2018). Decidable and undecidable arithmetic functions in actin filament networks. *Journal of Physics D: Applied Physics*, 51(3), 034005. <https://doi.org/10.1088/1361-6463/aa9d7b>

16. Schumann, A., Adamatzky, A., Król, J., & Góles, E. (2024). Fungi as Turing automata with oracles. *Royal Society Open Science*, 11(10), 240768. <https://doi.org/10.1098/rsos.240768>
17. Rickles, D. (2016). *The philosophy of physics*. Polity.
18. Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In W. H. Zurek (Ed.), *Complexity, entropy, and the physics of information* (pp. 3–28). Addison-Wesley.
19. Deacon, T. W. (2012). *Incomplete nature: How mind emerged from matter*. W. W. Norton & Company.
20. Levin, M. (2021). Bioelectric signaling: Reprogrammable circuits underlying embryogenesis, regeneration, and cancer. *Cell*, 184(8), 1971–1989. <https://doi.org/10.1016/j.cell.2021.03.018>
21. Burgin, M. (2010). *Theory of information: Fundamentality, diversity and unification*. World Scientific.
22. Burgin, M. (2012). *Structural reality*. Nova Science Publishers.
23. Burgin, M. (2011). *Theory of named sets*. Nova Science Publishers.
24. Burgin, M. (2016). *Theory of knowledge: Structures and processes*. World Scientific.
25. Deutsch, D. (2011). *The beginning of infinity: Explanations that transform the world*. Viking.
26. Turing, A. M. (1939). Systems of logic based on ordinals. *Proceedings of the London Mathematical Society*, s2-45(1), 161–228.
27. Penrose, R. (1989). *The emperor's new mind*. Oxford University Press.
28. Burgin, M., & Mikkilineni, R. (2021). On the autopoietic and cognitive behavior. EasyChair Preprint No. 6261, Version <https://easychair.org/publications/preprint/tkjk>
29. Mikkilineni, R., Morana, G., & Burgin, M. (2015). Oracles in software networks: A new scientific and technological approach to designing self-managing distributed computing processes. In *Proceedings of the 2015 European Conference on Software Architecture Workshops (ECSAW '15)* (Article 11, pp. 1–8). Association for Computing Machinery. <https://doi.org/10.1145/2797433.2797444>
30. Mikkilineni, R. (2022a). A new class of autopoietic and cognitive machines. *Information*, 13(1), 24. <https://doi.org/10.3390/info13010024>
31. Mikkilineni, R., Comparini, A., & Morana, G. (2012, June). The Turing O-Machine and the DIME Network Architecture: Injecting the architectural resiliency into distributed computing. In *Turing-100* (pp. 239–251).
32. Mikkilineni, R. (2022b). Infusing autopoietic and cognitive behaviors into digital automata to improve their sentience, resilience, and intelligence. *Big Data and Cognitive Computing*, 6(1), 7. <https://doi.org/10.3390/bdcc6010007>
33. Hill, S. L. (2025a). Fold Theory: A categorical framework for emergent spacetime and coherence [Preprint]. Academia.edu. <https://www.academia.edu/130062788>
34. Monod, J. (1971). *Chance and necessity*. Knopf.
35. Mikkilineni, R., Kelly, W. P., & Crawley, G. (2024). Digital genome and self-regulating distributed software applications with associative memory and event-driven history. *Computers*, 13(9), 220. <https://doi.org/10.3390/computers13090220>
36. Haken, H. (1977). *Synergetics: An introduction*. Springer.
37. Mikkilineni, R. (2025). General theory of information and mindful machines. *Proceedings*, 126(1), 3. <https://doi.org/10.3390/proceedings2025126003>
38. Mikkilineni, R., & Michaels, M. (2025). *Society of Minds: The Architecture of Mindful Machines*. Preprints. <https://doi.org/10.20944/preprints202510.0046.v1>
39. Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475. (Describes the Markov blanket and minimizing free energy as the physics of life).
40. Froese, T. (2023). Irruption theory: A novel conceptualization of the enactive account of motivated activity. *Entropy*, 25(5), 748. (Discusses entropy, constraints, and self-organization)

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.