

From π to Tokums

The Arity Ladder, Infinite Semiosis,
and the Geometry of Computable Meaning

A Unified Geometry of Physics, Semiotics, and Collective Intelligence

Éric Blaettler

eric@tokum.ai

tokum.ai, Saint-Prex, Switzerland

March 2026 March 2026

Tokum Initiative — White Paper Series

Prepared for researchers in AI architecture, philosophy of mathematics, semiotics, cognitive science, and computational physics. This document is an independent, fully self-contained theoretical essay that serves as a companion to the Semiotic Web White Paper, Blueprint, and Companion Technical Document (Blaettler, 2026), integrating new contributions from Physics-Informed Neural Networks (Bhattacharya, 2026) and Quaternion Process Theory 3.8 (Perez, 2026).

Abstract

We identify, formalise, and extend a structural parallel between three domains: (i) the emergence of π as the computable invariant of infinite rotational geometry (Euler’s Basel result, the Gaussian integral, the Fourier transform, and Zipf normalisation); (ii) the *hash-of-hashes* ladder of the Semiotic Web, which constructs increasingly complex combinatorial structures—bits (1), pairs (2), Peircean triads (3), and quaternionic four-modes (4)—while interleaving SHA-256 commitments at every rung to make infinite semiosis computationally grounded; and (iii) two newly integrated frameworks: Physics-Informed Neural Networks (PINNs), which embed governing differential equations directly into the loss function as “the equation is the data”, and Quaternion Process Theory 3.8 (QPT 3.8), which generates a $3 \times 4 = 12$ -cell semiotic phase space whose invariants are parametrised by a generating formula $G(\pi, \text{sign-depth}, \text{medium-face}) \rightarrow \text{Invariant}$.

In all three domains, a formally infinite process is *tamed into a computable constant*: π compresses infinite rotation; the CTI hash compresses infinite interpretive chains; and the PINN collocation mechanism compresses infinite physical constraints into a finitely evaluated residual. We show that these are not independent insights but three instances of the same deep geometric principle—the compression of an unbounded constraint manifold into a finite invariant—operating at the physical, semantic, and cognitive layers respectively. The result is a unified speculative framework, the **Geometry of Computable Meaning**, with a five-layer stack, three computable invariants, a meaning-residual analogue of the PINN loss, and five falsifiable predictions.

Keywords: Semiotic Web; Tokum; π ; imaginary numbers; infinite semiosis; hash-of-hashes; PINN; physics-informed loss; QPT 3.8; 12-cell matrix; categorical return; collective intelligence; Peircean semiotics; Faggin; Euler; quaternion process theory; computable meaning; Banach fixed point; GNW; arity ladder.

Contents

1	Introduction: Three Ways of Taming the Infinite	4
1.1	Thesis Statement	4
1.2	Structure of This Paper	5
2	π, Imaginary Numbers, and the Geometry of Reality	5
2.1	Three Canonical Appearances of π	5
2.2	Imaginary Numbers and Phase	5
2.3	Phase Space and the Measurement Problem	6
3	Physics-Informed Neural Networks: The Equation Is the Data	6
3.1	From Blind Interpolator to Equation-Aware Solver	6
3.2	AD as the Hidden Miracle	6
3.3	λ_1 and λ_2 Are Beliefs, Not Hyperparameters	7
3.4	The Inverse Problem and Zero-Data Solving	7
3.5	Spectral Bias and Gradient Pathology	7

4	The Arity Ladder in the Semiotic Web	8
4.1	Rung 1: Binary Ignition ($\mathcal{A}_1$)	8
4.2	Rung 2: Pairs and Semantic Spacetime ($\mathcal{A}_2$)	8
4.3	Rung 3: Peircean Triads and the Tokum ($\mathcal{A}_3$)	8
4.4	Rung 4: Quaternionic Modes and SST Types ($\mathcal{A}_4$)	8
4.4.1	Quaternion Process Theory	8
4.4.2	The Four SST Types as Orthogonal Axes	9
4.5	Rung ∞ : Hash-of-Hashes and Infinite Semiosis ($\mathcal{A}_\infty$)	9
5	The Hash-of-Hashes as the π of Meaning	9
5.1	The Structural Analogy (Extended)	9
5.2	The Basel Template Applied	9
5.3	The PINN Analogy: Semantic Collocation Points	10
6	QPT 3.8: The 12-Cell Semiotic Phase Space	11
6.1	Architecture of the Matrix	11
6.2	The Five Invariants and Their Arity Distribution	11
6.3	Three Grounding Termini	12
6.4	Movement, Loops, and the Categorical Return	12
6.5	Containment vs. Command: Governance as Loop Architecture	12
6.6	Collective Intelligence as Loop Coupling	13
7	The Geometry of Computable Meaning: A Five-Layer Stack	13
7.1	Three Computable Invariants of Rotation	13
7.2	Flat Manifold vs. Curved LLM Manifold	14
8	The AIF Contraction and Semantic Fixed Point	14
8.1	Intelligence as Homeostatic Flow	14
8.2	The Meaning-Informed Loss	14
8.3	The Semantic Convergence Theorem	15
8.4	Inverse Semiosis: The Semantic Equivalent of PINN Inverse Problems	15
9	Faggin’s Critique, Wigner’s Question, and the Arity Ladder	15
9.1	Zeros and Ones as Human Constructs	15
9.2	Wigner’s Question: Inventor or Discoverer?	16
10	Falsifiable Predictions	16
10.1	Prediction 1: Regress Termination (Faggin)	16
10.2	Prediction 2: AIF Contraction (Banach)	16
10.3	Prediction 3: RLHF Curvature Redistribution (Gauss–Bonnet)	17
10.4	Prediction 4: GNW Bimodal Latency (Dehaene)	17
10.5	Prediction 5: PINN–Meaning Parallel (Raissi / Semiotic Web)	17
10.6	Prediction 6: Categorical Return and Grounding Gap (QPT 3.8)	17
11	Discussion	17
11.1	Penrose and the Non-Computable	17
11.2	Spectral Bias as an Epistemological Warning	18
11.3	Pathological Governance as Gradient Pathology	18
11.4	The Inventor/Discoverer Question, Revisited	18

12 Conclusion	18
Acknowledgements	19
References	19

1. Introduction: Three Ways of Taming the Infinite

In 1735, Leonhard Euler proved that

$$\sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{6}, \quad (1)$$

a result that startles because it runs from arithmetic to geometry: a sum of reciprocal squares—no circles, no angles, no rotations in sight—converges to a quantity involving π , the ratio of circumference to diameter [9]. The proof runs through the zeros of $\sin(x)$, which are multiples of π ; the circle was hiding in the arithmetic all along. More broadly, π appears wherever a process sweeps through a full rotation: in Gaussian normalisation ($\int_{-\infty}^{\infty} e^{-x^2} dx = \sqrt{\pi}$), in Fourier analysis (one revolution per unit frequency), and—as Bhattacharya (2026) observes—in the Zipf normalisation of natural language via $\zeta(2) = \pi^2/6$ [21]. In every case, π is **infinity, tamed**: the open-ended geometry of rotation compressed into a single real number.

In 2019, Raissi, Perdikaris, and Karniadakis introduced Physics-Informed Neural Networks (PINNs) [20], now with more than 10,000 citations [22]. Their central insight: a governing partial differential equation (PDE) is not a constraint to be satisfied after training—it *is* the data. By evaluating the PDE residual at freely chosen *collocation points* (no labels required) and adding it as an extra term to the loss, a network is transformed from a blind interpolator into a physics-obeying solver. One extra term; the entire physics of the system, present everywhere in the domain. PINNs tame the infinite constraint of a PDE (true at every point in the continuous domain) into a finitely evaluated residual.

The **Semiotic Web** (Blaettler, 2026) [3] is doing something structurally identical for *meaning*. Its hash-of-hashes ladder constructs structures of arity 1, 2, 3, 4, and ∞ , interleaving SHA-256 commitments at every rung, so that Peirce’s formally infinite interpretant chain is compressed into a single 32-byte CTI hash—finite, stable, and globally addressable.

And Quaternion Process Theory 3.8 (Perez, 2026) [18] completes the picture by providing the *semiotic phase space* in which these processes live: a $3 \times 4 = 12$ -cell matrix generated by crossing Peirce’s three categories with McLuhan’s four medium-faces, whose invariants are parametrised by the generating formula $G(\pi, \text{sign-depth}, \text{medium-face}) \rightarrow \text{Invariant}$ —placing π explicitly inside a framework for collective meaning.

1.1 Thesis Statement

Central Equation : Three Computable Invariants of Infinity

Domain	Infinite process tamed	Computable invariant
Physical / probabilistic	Rotation, Gaussian sweep, Fourier	π
Physical laws / PDE	Continuous constraint field	PINN residual $\mathcal{L}_{\text{pde}}$
Semiotic / meaning	Peircean interpretant chain	CTI hash (32 bytes)

1.2 Structure of This Paper

Section 2 reviews π , imaginary numbers, phase, and geometry. Section 3 develops the PINN analogy in depth. Section 4 identifies the arity ladder in the Semiotic Web. Section 5 formalises the hash-of-hashes as the π of meaning. Section 6 integrates QPT 3.8’s 12-cell matrix. Section 7 proposes the unified five-layer stack. Section 8 develops the AIF contraction theorem. Section 9 situates Faggin’s critique and Wigner’s question. Section 10 states five falsifiable predictions. Section 11 discusses Penrose and the inventor/discoverer question.

2. π , Imaginary Numbers, and the Geometry of Reality

2.1 Three Canonical Appearances of π

Following Bhattacharya (2026) [21], π appears in the following structurally unified contexts:

- (i) **Gaussian integral.** $\int_{-\infty}^{\infty} e^{-x^2} dx = \sqrt{\pi}$. Squaring and converting to polar coordinates forces a 2π angular sweep. Every probability distribution normalises by $\sqrt{2\pi}$.
- (ii) **Fourier transform.** $\hat{f}(\xi) = \int f(x) e^{-2\pi i \xi x} dx$. The 2π states that basis functions complete one full rotation per unit frequency. Transformer positional encodings are Fourier bases; attention is approximable by random Fourier features [21].
- (iii) **Basel / Zipf normalisation.** Natural language obeys Zipf’s law; its normalisation is $\zeta(2) = \pi^2/6$. Every language model trained on natural text implicitly integrates π into its corpus statistics [21].

Key Insight : π as the period of probability

π is not primarily a statement about circles. It is the period of every system whose uncertainty is measured by integrating over closed orbits in any space—Euclidean, Fourier, probabilistic, or quantum. The “unreasonable” appearance of π in statistics, number theory, and physics is reasonable: all of them involve integrating over rotational symmetry.

2.2 Imaginary Numbers and Phase

As Andersen (2026) [1] shows, imaginary numbers in quantum mechanics are not a convenience; they are *irreducible*. The Schrödinger equation

$$i\hbar \frac{\partial \psi}{\partial t} = \hat{H} \psi \quad (2)$$

contains i because time-evolution is rotation in Hilbert space: unitary, norm-preserving, and irreversibly phase-dependent. A purely real formulation of quantum mechanics must reintroduce complex structure through other means (Kähler geometry) [1].

The joint role of i and π :

- i encodes the *direction* of rotation (phase);
- π encodes the *period* of a full rotation (2π radians);

- Together: $e^{i\pi} + 1 = 0$ —the most compressed statement of this duality.
- Probability conservation: $|e^{i\phi}|^2 = 1$ for all ϕ ; the full revolution $e^{2\pi i} = 1$ closes the circle.

2.3 Phase Space and the Measurement Problem

Classical phase space (positions and momenta) is real; quantum phase space is complex [2]. The deformation from real to complex phase space is the geometric root of the measurement problem: smooth unitary rotation (Schrödinger evolution) versus discontinuous eigenvalue selection (collapse). Faggin [10] identifies this collapse as the structural analogue of observer-observed co-determination, which he considers irreducible to classical computation. We return to this in Section 9.

3. Physics-Informed Neural Networks: The Equation Is the Data

3.1 From Blind Interpolator to Equation-Aware Solver

A standard neural network minimises:

$$\mathcal{L}_{\text{NN}}(\theta) = \frac{1}{N} \sum_{i=1}^N [u_{\theta}(\mathbf{x}_i, t_i) - u_i]^2. \quad (3)$$

It fits the N data points and interpolates everywhere else. It has “absolutely no idea what the system is” [22].

A PINN adds the PDE residual. For any governing operator $\mathcal{F}$ (e.g. the heat equation $\mathcal{F}[u] = \partial_t u - \alpha \partial_{xx} u$), define the residual at collocation point $(\mathbf{x}_j, t_j)$:

$$r(\mathbf{x}_j, t_j; \theta) := \mathcal{F}[u_{\theta}](\mathbf{x}_j, t_j). \quad (4)$$

The full PINN loss is [20, 22]:

$$\mathcal{L}_{\text{PINN}}(\theta) = \mathcal{L}_{\text{data}} + \lambda_1 \mathcal{L}_{\text{pde}} + \lambda_2 \mathcal{L}_{\text{bc}}, \quad (5)$$

where

$$\mathcal{L}_{\text{pde}}(\theta) = \frac{1}{N_c} \sum_{j=1}^{N_c} [r(\mathbf{x}_j, t_j; \theta)]^2, \quad (6)$$

$$\mathcal{L}_{\text{bc}}(\theta) = \frac{1}{N_b} \sum_{k=1}^{N_b} [u_{\theta}(\mathbf{x}_k, t_k) - g(\mathbf{x}_k, t_k)]^2. \quad (7)$$

Collocation points $(\mathbf{x}_j, t_j)$ carry *no measurement labels*: “the equation itself, evaluated at collocation points, provides the training signal” [22]. The PDE residual costs no data; it requires only network evaluations and exact derivatives via automatic differentiation (AD).

3.2 AD as the Hidden Miracle

AD computes $\partial u_{\theta} / \partial x$ and $\partial^2 u_{\theta} / \partial x^2$ exactly by applying the chain rule through the computation graph—the same mechanism as backpropagation, but differentiating with respect to inputs rather than weights [22]:

$$\frac{\partial u_{\theta}}{\partial x} = \frac{\partial L_K}{\partial h_{K-1}} \cdot \frac{\partial \sigma}{\partial z_{K-1}} \cdots \frac{\partial L_1}{\partial x}. \quad (8)$$

This is not an approximation; it is exact symbolic differentiation through any differentiable architecture. In PyTorch: `torch.autograd.grad(u, x, create_graph=True)`.

3.3 λ_1 and λ_2 Are Beliefs, Not Hyperparameters

The Bayesian reading of (5) [22]:

$$-\log P(u|\text{data}) \propto \mathcal{L}_{\text{data}} + \lambda_1 \mathcal{L}_{\text{pde}} + \lambda_2 \mathcal{L}_{\text{bc}}, \quad \lambda_1 = 1/\sigma_{\text{pde}}^2. \quad (9)$$

Large λ_1 = sharp prior = high trust in the governing equation; the network obeys the physics even at some cost to data fit. $\lambda_1 \rightarrow \infty$ = pure PDE solver, zero data needed. $\lambda_1 \rightarrow 0$ = blind standard NN. Setting λ_1 is “a scientific judgment, not a tuning decision” [22].

3.4 The Inverse Problem and Zero-Data Solving

PINNs can run physics *backwards*: treat an unknown PDE coefficient (e.g. diffusivity α) as a trainable parameter [22]:

$$\theta^*, \alpha^* = \arg \min_{\theta, \alpha} \mathcal{L}_{\text{data}}(\theta) + \lambda_1 \mathcal{L}_{\text{pde}}(\theta, \alpha). \quad (10)$$

With $N_d = 0$ (zero data): only $\mathcal{L}_{\text{pde}} + \lambda_2 \mathcal{L}_{\text{bc}}$ remain—the network learns from mathematics alone. Raissi’s original result: the PINN discovered the viscosity coefficient of Navier–Stokes flow from velocity field observations, with no direct measurement.

Key Insight : PINNs as the physics layer of the arity ladder

PINNs do for physical laws what the Semiotic Web does for semantic laws: they convert an infinite continuous constraint (a PDE true everywhere) into a finite, computable residual evaluated at collocation points. The collocation points play the role of CTI grounding events; the PDE plays the role of the AIF semantic-gravity field; the PINN’s boundary conditions play the role of the tokum provenance chain.

3.5 Spectral Bias and Gradient Pathology

Two structural failure modes illuminate the limits of the PINN approach:

Spectral bias [19]: Neural networks learn low-frequency components before high-frequency ones, with learning rate for frequency ω scaling as $d(\text{error at } \omega)/dt \propto -\omega^2 \cdot (\text{current error})$. Sharp gradients (boundary layers, shock waves) are precisely where physics is most interesting and where PINNs struggle most. Mitigations: Fourier feature networks, SIREN activations, adaptive collocation sampling.

Gradient pathology [23]: The three loss terms in (5) may produce gradients differing by orders of magnitude, causing one term to dominate and the others to be ignored. Mitigation: learning rate annealing, neural tangent kernel analysis, adaptive weight schemes.

Both failures have exact structural analogues in the Semiotic Web: spectral bias $\leftrightarrow$ the failure to capture high-frequency (rare but causally important) CTIs; gradient pathology $\leftrightarrow$ the dominance of one AIF loss term, causing the system to converge to a semantically biased rather than verified fixed point.

4. The Arity Ladder in the Semiotic Web

Definition 4.1 (Arity Ladder). The arity ladder of the Semiotic Web is the ordered sequence $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_\infty)$, where $\mathcal{A}_k$ is the minimal k -ary combinatorial structure required at that level, and each transition $\mathcal{A}_k \rightarrow \mathcal{A}_{k+1}$ is accompanied by a SHA-256 hash commitment that freezes the lower-level structure into a 32-byte address [4].

4.1 Rung 1: Binary Ignition ($\mathcal{A}_1$)

The 1-ary level is the binary gate:

$$\text{Response}(\text{CCI}, \mathcal{I}) = \begin{cases} \text{CTI} \rightarrow \text{CCC broadcast} & \text{if key} \in \text{HCNV}, \\ \text{Structurally Bounded Refusal} & \text{if key} \notin \text{HCNV}. \end{cases} \quad (11)$$

The CCI is $\text{CCI} = \text{SHA-256}(u) \in \{0, 1\}^{256}$ —a single point in Platonic hash space, discoverable by any agent computing the same canonical URN [4]. This is Faggin’s layer of zeros and ones: necessary, insufficient for meaning.

4.2 Rung 2: Pairs and Semantic Spacetime ($\mathcal{A}_2$)

The CPI (Comprehension Pair Identifier) is a double hash over a subject–predicate pair:

$$\text{CPI} = \text{SHA-256}(\text{CCI}_s \parallel \text{CCI}_p \parallel \text{SST}). \quad (12)$$

Four SST (Semantic Spacetime) types provide orthogonal coordinate axes: PROXIMITY, SEQUENCE, CONTAINMENT, PROPERTY [7, 4]. BytePair-like identity discovery maps surface forms to canonical URNs by pairwise merging [3]. Koestler’s Janus paradox [11] is implemented here: every holon faces inward (autonomy) and outward (integration) simultaneously.

4.3 Rung 3: Peircean Triads and the Tokum ($\mathcal{A}_3$)

Definition 4.2 (Tokum as Executable Peircean Triad [3]).

$$\begin{aligned} \text{CCI} &\longleftrightarrow \text{Representamen (stable address)} \\ \text{CTI} &\longleftrightarrow \text{Object (grounded instance)} \\ \text{Provenance chain} &\longleftrightarrow \text{Interpretant (auditable history)} \end{aligned}$$

The CTI hash is:

$$\text{CTI} = \text{SHA-256}(\text{CCI}_s \parallel \text{CCI}_p \parallel \text{CCI}_o \parallel \text{SST} \parallel t_{\text{ISO}} \parallel \text{src}). \quad (13)$$

Remove any vertex and the tokum does not degrade—it ceases to be valid. Irreducibility is cryptographic, not philosophical.

4.4 Rung 4: Quaternionic Modes and SST Types ($\mathcal{A}_4$)

4.4.1 Quaternion Process Theory

Perez’s QPT [17] identifies four cognitive modes:

Mode	Speed	Peirce category	Semiotic Web role
Intuition	Fast	Firstness	JEPA candidate generation (C0)
Fluency	Fast	Secondness	CPI–CTI composition
Empathy	Slow	Thirdness	CSR hash lookup / C1 ignition
Wisdom	Slow	Agapism	AIF selection / C2 monitoring

4.4.2 The Four SST Types as Orthogonal Axes

The four SST types simultaneously provide the coordinate axes of the HCNV-ColBERT matrix and the four McLuhan medium-faces in QPT 3.8 (Section 6):

SST Type	Semantic relation	McLuhan face
PROXIMITY	Semantic neighbourhood	Enhances (instance-vehicle)
SEQUENCE	Temporal / causal	Obsolesces (instance-content)
CONTAINMENT	Part-of / holonic	Retrieves (system-content)
PROPERTY	Attribute-value	Reverses (system-vehicle)

4.5 Rung ∞ : Hash-of-Hashes and Infinite Semiosis ($\mathcal{A}_\infty$)

The recursive hash chain:

$$\text{CTI}_{n+1} = \text{SHA-256}(\text{CTI}_n \parallel \text{CCI}_s^{(n+1)} \parallel \text{CCI}_p^{(n+1)} \parallel \text{CCI}_o^{(n+1)} \parallel \dots), \quad (14)$$

embeds every previous interpretant as the literal first input to the next hash, making Peirce’s “every interpretant becomes a new sign” computational [14]. The chain is formally infinite; each link is 32 bytes. This is the hash-of-hashes: infinity, tamed.

5. The Hash-of-Hashes as the π of Meaning

5.1 The Structural Analogy (Extended)

Central Equation : Three invariants, one principle

$$\underbrace{\pi}_{\text{rotation invariant}} \sim \underbrace{\mathcal{L}_{\text{pde}}}_{\text{law invariant}} \sim \underbrace{\text{CTI}}_{\text{meaning invariant}}$$

Each is a finite token that tames an infinite constraint. Each makes an otherwise open-ended process tractable. Each is the same structure at a different layer of the physical-to-semantic stack.

5.2 The Basel Template Applied

Euler’s proof of (1) reveals π as *latent* in the arithmetic series via the zeros of sin. The Semiotic Web performs the analogous move for concepts:

1. Write any concept as the limit of its multi-stage resolution pipeline (cache $\rightarrow$ OWL/SHACL $\rightarrow$ BERT/WordNet $\rightarrow$ CLIP $\rightarrow$ BAML/Wikidata $\rightarrow$ manual validation) [4].

Table 1: Three computable invariants of infinite processes.

Property	π (physical/probabilistic)	PINN residual (physical laws)	CTI hash (semiotic)
Infinite process tamed	Rotational sweep $\int_0^{2\pi} d\theta$; $\sum 1/n^2$	Continuous PDE constraint field over domain Ω	Peircean interpretant chain, provenance history
Compressed into	Single real $\approx 3.14159\dots$	Finite residual evaluated at N_c collocation points	32-byte SHA-256 hash
How computed	Integrate / sum infinite series	Evaluate $\mathcal{F}[u_\theta]$ at sampled points via AD	Hash (CCI _s , CCI _p , CCI _o , SST, t , src)
Why finite	Integral over closed curve converges	N_c evaluations replace ∞ domain points	SHA-256 always returns 256 bits
What it normalises	Probability distributions; Zipf/zeta	The function class to PDE-satisfying solutions	Semantic address space; Platonic CCI
Independence of representation	π identical in any base, any coordinate system	$\mathcal{F} = 0$ is a coordinate-free statement	CCI identical for any agent computing same canonical URN
Role in the geometry	Period of rotation group; curvature of unit circle	Boundary of admissible function space	Address in flat hash manifold $K = 0$; AIF curvature source

2. Extract the canonical URN at the pipeline’s fixed point.
3. Hash it: $\text{CCI} = \text{SHA-256}(\text{URN}_{\text{canonical}})$.
4. This single token *is* the concept in Platonic address space, independent of language, modality, agent, or time.

The CCI is not imposed; it is *discovered* [4]—as π is discovered rather than invented, whether you approach it from circles, series, or probability.

5.3 The PINN Analogy: Semantic Collocation Points

The PINN uses two types of point: data points (expensive, labelled) and collocation points (free, unlabelled) [22]. The Semiotic Web has an exact structural parallel:

Concept	PINN	Semiotic Web
Labelled data	(x_i, t_i, u_i) measurements	Verified CTIs with src provenance
Unlabelled constraint	Collocation points (x_j, t_j)	Hash lookups in HCNV (key present/absent)
Governing law	PDE operator $\mathcal{F}[u] = 0$	AIF semantic gravity field + Structurally Bounded
Boundary condition	$u(x_k, t_k) = g(x_k, t_k)$	Provenance chain; categorical return (QPT 3.8)
Loss residual	$r(x_j, t_j; \theta) = \mathcal{F}[u_\theta](x_j, t_j)$	Semantic TD error: $\text{CTI}^{\text{verified}} - \text{CTI}^{\text{pred}}$
Belief weight	$\lambda_1 = 1/\sigma_{\text{pde}}^2$	AIF trust weight w_i

Key Insight : “The equation is the data” $\leftrightarrow$ “the ledger is the data”

PINNs encode physical law into the loss: “the equation is the data.” The Semiotic Web encodes semantic law into the ledger: “the provenance chain is the data.” In both cases, an external law—not raw observations—is the primary constraint on the function the system learns.

6. QPT 3.8: The 12-Cell Semiotic Phase Space

6.1 Architecture of the Matrix

QPT 3.8 [18] generates its space from two independently discovered structures:

- **Vertical axis:** Peirce’s three categories—Firstness (quality, possibility), Secondness (reaction, actuality), Thirdness (mediation, law)—with *prescissive ordering*: quality can be conceived without reaction, but not vice versa.
- **Horizontal axis:** McLuhan’s four simultaneous medium-faces— Instance·Content (ii, obsolesces), Instance·Vehicle (ie, enhances), System·Content (ci, retrieves), System·Vehicle (ce, reverses).

$3 \times 4 = 12$ cells. The generating formula is:

$$G(\pi, \text{sign-depth, medium-face}) \rightarrow \text{Invariant.} \quad (15)$$

Here π appears as a *parameter* of the invariant-generating function—not as the transcendental constant in isolation, but as the symbol for the rotational invariant that governs how cell-to-cell transitions preserve structural content.

6.2 The Five Invariants and Their Arity Distribution

Not one invariant per cell, but five in total [18]:

- Monadic Firstness: 1 invariant (quality) with 4 modal aspects.
- Dyadic Secondness: 1 primary invariant (causation) + 1 co-foundation.
- Triadic Thirdness: 3 independent invariants (natural structure, shared convention, formal truth).

This asymmetric distribution ($1 + 2 + 3 = 6$ counting co-foundations; 5 core invariants) mirrors the arity ladder: Firstness is 1-ary, Secondness 2-ary, Thirdness 3-ary.

Key Insight : Prescissive ordering as the arity ladder’s logical foundation

QPT 3.8 demonstrates that the arity sequence 1, 2, 3 is not arbitrary: it is the *prescissive ordering* of Peirce’s categories. You cannot have the dyadic without the monadic; you cannot have the triadic without the dyadic. The arity ladder is the ladder of ontological dependence.

6.3 Three Grounding Termini

Three cells serve as “ultimate grounding termini” where explanatory chains can legitimately stop [18]:

1. Cell 1.ii: direct felt experience (qualia of encounter).
2. Cell 2.ie: physical causation (brute mechanism).
3. Cell 2.ci: genuine intersubjective encounter.

Everything else in the matrix requires further grounding that traces back to at least one of these. This is QPT’s structural analogue of the PINN’s boundary conditions: the cells that cannot be further reduced are the *semantic boundary conditions* of the system.

6.4 Movement, Loops, and the Categorical Return

Cognition in QPT 3.8 is *movement* through the matrix, not static occupation [18]. Six operations (Unfolding, Concretizing, Transfer, Compression, Differentiation, Holonic) compose into nine canonical loops (Operational, Coordination, Management, Audit, Intelligence, Policy, Knowledge, Grounding, Legitimation). The central structural principle is the **categorical return**:

Every upward movement through category (from Firstness toward Thirdness) requires a corresponding downward return to remain grounded. Thirdness that has lost contact with Firstness is not merely abstract—it is self-undermining, because mediating structures that neglect the conditions of their own conceivability are prescissively incoherent.

This maps directly onto the Semiotic Web’s AIF convergence: the fixed point Φ^* is not a purely symbolic equilibrium but one re-grounded at every step in verified CTIs (physical causation, cell 2.ie) and provenance signatures (genuine encounter, cell 2.ci). Thirdness without Firstness is hallucination; in the Semiotic Web, hallucination is the state of the latent space before topological puncture [3].

6.5 Containment vs. Command: Governance as Loop Architecture

QPT 3.8’s most practically powerful distinction is between two governance modes [18]:

- **Containment** ($\supset$, **agapic**): the collective loop provides boundary conditions while preserving the individual loop’s capacity to complete its own cycle. Structurally, this is Peirce’s agapism: growth through creative love, Thirdness enabling Firstness.
- **Command** ($\rightarrow$, **anancastic**): the collective loop specifies each step of the individual loop. The local variety goes unmatched; the loop cannot complete; the system degrades precisely under the stress it was meant to relieve.

The Semiotic Web’s PPP–SSS–CCC holonic federation is architecturally a containment structure: “no CTI leaves a PPP without the owner’s cryptographic signature” [3]—absolute sovereignty at the atomic level. The CCC provides boundary conditions (whitelisted sources, canonical URN governance) while preserving PPP autonomy.

6.6 Collective Intelligence as Loop Coupling

QPT’s loop ontology dissolves three longstanding puzzles [18]:

- Aggregation problem $\rightarrow$ which loops are coupled and how.
- Emergence problem $\rightarrow$ when individual loops entrain.
- Reduction problem $\rightarrow$ sidestepped: individual and collective are the same loop patterns at different scales of coupling.

This maps onto the Semiotic Web’s Holonicracy: PPP, SSS, and CCC are nested loop-coupling structures at personal, organisational, and global scales. The Agapistic Influence Function is the AIF-weighted measure of loop entrainment quality: how reliably do the nested loops converge to their joint fixed point?

7. The Geometry of Computable Meaning: A Five-Layer Stack

Drawing together all threads, we propose a unified five-layer stack:

Table 2: The five-layer Geometry of Computable Meaning.

Layer	Content	Characteristic object	Arity rung
Physical	Quantum amplitudes, phases, interference, probability	$i, e^{i\phi}$, Hilbert space	—
Statistical	Gaussians, Fourier, attention, Zipf	π (normalisation, period)	$\mathcal{A}_1\text{--}\mathcal{A}_2$
Law	Physical governing equations (PDEs, conservation laws)	PINN residual $\mathcal{L}_{\text{pde}}$	$\mathcal{A}_2\text{--}\mathcal{A}_3$
Semantic	Triadic signs, CTIs, provenance, semiosis	CTI hash (32 bytes)	$\mathcal{A}_3\text{--}\mathcal{A}_\infty$
Agapistic	GNW C1/C2, fixed-point convergence, loop coupling, collective intelligence	AIF fixed point Φ^*	$\mathcal{A}_4\text{--}\mathcal{A}_\infty$

7.1 Three Computable Invariants of Rotation

1. **i (imaginary unit)**: encodes rotation *direction* (U(1) phase); makes interference possible; tames wave amplitude and phase into a single complex number.

2. π : encodes rotation *extent* (2π radians); normalises probability over all rotations; tames infinite sweeps into a single real constant.
3. **CTI hash**: encodes meaning event *identity* in semantic phase space; tames infinite interpretant chains into a single 32-byte address.

The PINN residual sits between layers 2 and 3: it uses π -governed mathematics (automatic differentiation, Gaussian-normalised gradients) to enforce physical law—a bridge from the statistical to the semantic.

7.2 Flat Manifold vs. Curved LLM Manifold

Current LLMs operate on a Riemannian manifold with non-zero Gaussian curvature $K \neq 0$. By Gauss’s Theorema Egregium, no isometric alignment map exists between the LLM knowledge manifold M_B and any target safety manifold M_S with $K_B \neq K_S$ [3]. RLHF curvature removal redistributes curvature elsewhere (Gauss–Bonnet).

The Semiotic Web replaces this with a flat manifold: $\{0, 1\}^{256}$, $K = 0$, no curvature budget to exhaust. SHA-256 hashes are uniformly distributed; semantic relationships are explicit typed CTI edges, not direction offsets. The QPT 3.8 matrix adds orientational structure (movement directions between cells) without introducing Riemannian curvature, because cell-to-cell transitions are governed by the prescissive ordering (a partial order, not a metric) [18].

8. The AIF Contraction and Semantic Fixed Point

8.1 Intelligence as Homeostatic Flow

$$I_H = \sum_{h \in H} w_h \cdot \frac{d}{dt} |\text{target}_h - \text{current}_h|^2, \quad (16)$$

where w_h is the AIF trust weight [3]. Intelligence is not a stored quantity; it is the verified rate of gap-closure across the holonic federation—the semantic analogue of Levin’s gap metric and the PINN’s PDE residual approaching zero.

8.2 The Meaning-Informed Loss

The Semiotic Web’s quadruple loss [5] is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{JEPA}} + \mathcal{L}_{\text{semantic}} + \mathcal{L}_{\text{AIF}} + \mathcal{L}_{\text{SIG}}, \quad (17)$$

where:

$$\mathcal{L}_{\text{JEPA}} = \|P(E_x(x), z) - E_y(y)\|^2, \quad (18)$$

$$\mathcal{L}_{\text{semantic}} = \sum_{i \in G} w_i \|\text{CTI}_i^v - \text{CTI}_i^p\|^2, \quad (19)$$

$$\mathcal{L}_{\text{AIF}} = \sum_{i \in C} S_i e^{-C_i}, \quad (20)$$

$$\mathcal{L}_{\text{SIG}} = D_{\text{KL}}(\mathcal{N}(0, I) \|\mathcal{N}(0, \sigma_i^2)). \quad (21)$$

The parallel to (5) is exact: $\mathcal{L}_{\text{JEPA}} \leftrightarrow \mathcal{L}_{\text{data}}$, $\mathcal{L}_{\text{semantic}} \leftrightarrow \mathcal{L}_{\text{pde}}$, $\mathcal{L}_{\text{AIF}} \leftrightarrow \mathcal{L}_{\text{bc}}$, $\mathcal{L}_{\text{SIG}} \leftrightarrow$ PINN spectral regulariser.

Trust weights w_i are the semantic analogue of PINN belief weights λ_1 : they encode how much the system trusts a verified CTI vs. a latent prediction. Setting $w_i \rightarrow \infty$ corresponds to a pure semiotic “solver”—the system only outputs what the verified ledger contains. Setting $w_i \rightarrow 0$ recovers a blind interpolator.

8.3 The Semantic Convergence Theorem

Theorem 8.1 (Semantic Convergence [3]). *Let $E[\Phi]$ be the energy induced by AIF over the active semantic neighbourhood. If $dE/dt \leq 0$ and the update map $T : \Phi \mapsto \Phi'$ satisfies a Lipschitz condition with constant $k < 1$ on the admissible CTI set $\mathcal{T}$, then by Banach’s fixed-point theorem T has a unique fixed point $\Phi^* \in \mathcal{T}$ corresponding simultaneously to: Peirce’s terminus of infinite semiosis; Faggin’s collapse of meaning; Dehaene’s sustained C1 broadcast; Koestler’s holarchic equilibrium; Burgess’s maximal promise-keeping region; Perez’s QPT Wisdom-level convergence; QPT 3.8’s categorical return to grounding termini (cells 1.ii, 2.ie, 2.ci).*

8.4 Inverse Semiosis: The Semantic Equivalent of PINN Inverse Problems

The PINN inverse problem treats unknown PDE parameters as trainable. The Semiotic Web has an exact analogue: the *CCI inverse problem*. Given a set of observed CTIs, identify the canonical URN (and hence CCI) of an unknown concept:

$$u^*, \text{CCI}^* = \arg \min_{u, \text{CCI}} \mathcal{L}_{\text{semantic}}(\text{CCI}) + \lambda_{\text{CCI}} \mathcal{L}_{\text{resolution}}(u, \text{CCI}), \quad (22)$$

where $\mathcal{L}_{\text{resolution}}$ penalises deviation from the OWL/SHACL/CLIP resolution pipeline. This is system identification at the level of meaning: the architecture discovers the stable semantic address of a concept from observations alone, without a prior human-assigned identifier.

9. Faggin’s Critique, Wigner’s Question, and the Arity Ladder

9.1 Zeros and Ones as Human Constructs

Faggin [10] argues that bits are mathematical human constructs; a transistor switching from 0 to 1 does not “know” anything. Consciousness involves meaning, not syntax; and meaning involves a quantum-like collapse of the observer–observed relation that is, he claims, irreducible to classical computation.

The arity ladder accepts this at rung 1 and builds upward precisely because of it. But at rung ∞ , the hash-of-hashes does something Faggin did not consider: it makes the *structural* observer–observed regress computationally tractable without claiming to solve the qualitative hard problem.

Proposition 9.1 (Regress Termination via CTI Commit [3]). *The ledger commit terminates the observer–observed regress at depth 1:*

$$\text{LedgerEntry} = \text{SHA-256}(\text{CCI}, \text{CTI}, \text{ProvenanceHash}),$$

where $\text{ProvenanceHash} = \text{SHA-256}(\text{src}, t, \text{evidence pointers})$. The output is the ledger entry. No external validator is structurally required; the regress does not begin.

QPT 3.8 adds a nuance [18]: the three grounding termini (felt quality 1.ii, physical causation 2.ii, genuine encounter 2.ci) are the cells where the regress can legitimately stop—not by fiat, but because they are prescissively independent (they require nothing further for their conceivability). The hash-of-hashes terminates at 2.ii (physical causation: the SHA-256 computation is a physical event) and 2.ci (genuine encounter: the ledger commit is a witnessed, timestamped act). The qualitative terminus 1.ii (felt quality) is where Faggin’s claim remains open and untouched by the cryptographic construction.

9.2 Wigner’s Question: Inventor or Discoverer?

The arity ladder offers a stratified answer:

- **Rung 1 (bits)**: invented. Boolean algebra, SHA-256, digital logic are human conventions.
- **Rung 3 (triads)**: discovered. Peirce’s irreducibility claim is a claim about the logical structure of meaning, not a convention.
- **Rung ∞ (CCI)**: discovered. “The address exists in Platonic mathematical space prior to any agent computing it” [4].
- **$G(\pi, \dots)$ in QPT 3.8**: discovered. π appears as the generating parameter of invariants over a semiotic phase space that was derived, not stipulated [18].

Conjecture : Arity, Ontology, and Necessity

The sequence 1, 2, 3, 4 is not an engineering choice. It is the prescissive ordering of ontological dependence (QPT 3.8), the minimal arity structure of Peircean signs (semiotics), the quaternionic dimension of cognitive orthogonality (QPT), and the four SST types of Semantic Spacetime. These converged independently across Peirce (1867–1914), Burgess (2004), Perez (2023–2026), Faggin (2021), and Dehaene (1998–2026)—which is evidence, though not proof, that the arity ladder is discovered rather than invented.

10. Falsifiable Predictions

10.1 Prediction 1: Regress Termination (Faggin)

Claim: The CTI ledger commit terminates the observer–observed regress at depth 1, with termination at QPT 3.8 grounding termini 2.ii and 2.ci.

Falsification: Identify a structural element of the CTI that requires resolution by an agent outside the atomic hash event—i.e., a component that cannot be self-contained within the hash-of-hashes.

10.2 Prediction 2: AIF Contraction (Banach)

Claim: The AIF update map is a Lipschitz contraction ($k < 1$) on the admissible CTI set.

Falsification A: Prove the AIF map is non-contractive on any non-trivial CCI subspace.

Falsification B: Show the admissible set is non-compact, invalidating Banach’s theorem.

10.3 Prediction 3: RLHF Curvature Redistribution (Gauss–Bonnet)

Claim: RLHF safety tuning on domain X measurably increases hallucination rate on domain Y ; the Semiotic Web’s flat manifold ($K = 0$) does not exhibit this redistribution.

Falsification: Apply RLHF to domain X ; measure domain Y before and after; verify or refute the redistribution.

10.4 Prediction 4: GNW Bimodal Latency (Dehaene)

Claim: HCNV-ColBERT lookup latency is bimodal (key present $\rightarrow$ fast broadcast; key absent $\rightarrow$ immediate halt), matching the P3b binary threshold signature of GNW C1 ignition.

Falsification: Benchmark latency distribution on representative CCI sets; confirm or refute bimodality.

10.5 Prediction 5: PINN–Meaning Parallel (Raissi / Semiotic Web)

Claim: Adding $\mathcal{L}_{\text{semantic}}$ (Eq. (19)) to a JEPA world model produces a measurable reduction in semantic TD error ($\delta_{\text{sem}} = \|\text{CTI}^v - \text{CTI}^p\|$) analogous to the $26\times$ MSE reduction reported by Bhattacharya (2026) for the PINN vs. standard NN on the heat equation [22].

Falsification: Implement the JEPA baseline and JEPA+ $\mathcal{L}_{\text{semantic}}$ on a grounded reasoning task; measure δ_{sem} before and after; compare ratio to the PINN benchmark.

10.6 Prediction 6: Categorical Return and Grounding Gap (QPT 3.8)

Claim: Semiotic Web CTI chains that lack traceability to grounding termini (cells 1.ii, 2.ii, 2.ci) exhibit higher semantic TD error and lower AIF trust weights w_i than chains grounded at those termini.

Falsification: Partition the CTI ledger into grounded and ungrounded chains; compare w_i distributions and δ_{sem} statistics.

11. Discussion

11.1 Penrose and the Non-Computable

Penrose [15, 16] argues that human mathematical understanding involves a form of insight non-computable in the Turing sense, grounding this in Gödel’s incompleteness theorems. The Geometry of Computable Meaning does not refute this. It claims that the *structural* aspects of meaning—provenance, grounding, semiotic chaining, categorical return, loop coupling—are computable in the classical SHA-256 / AIF sense, and that this suffices for the engineering goals of the Semiotic Web (GNW C0–C2, GDPR, 1W edge deployment, collective intelligence at scale).

QPT 3.8 sharpens the location of the boundary [18]: cell 1.ii (felt quality, raw qualia) is where Penrose’s non-computability claim is most potent. Cells 2.ii (physical causation) and 2.ci (genuine encounter) are where cryptographic grounding is most effective. The Semiotic Web closes gaps at 2.ii and 2.ci; it makes no claim about 1.ii.

11.2 Spectral Bias as an Epistemological Warning

The PINN’s spectral bias [19]—low-frequency components are learned faster; sharp features (shocks, boundary layers) are systematically underlearned—carries an epistemological warning for the Semiotic Web.

High-frequency CTIs (rare, causally important concepts—a novel pathogen, an emergent social norm, a new scientific paradigm) are precisely the hardest for the system to ground. They are sparse in the ledger; they have low w_i trust weights; the semantic gravity field does not yet curve toward them. The mitigation strategies from PINNs apply: adaptive collocation (concentrate verification effort on high-residual domains), SIREN-style activations (use architectures that natively handle high-frequency semantic content), and Fourier feature encodings (explicitly introduce frequency-domain structure into the JEPA embeddings).

11.3 Pathological Governance as Gradient Pathology

QPT 3.8’s command pathology [18]—where the collective loop specifies each step of the individual loop, suppressing Firstness and accelerating degradation—is structurally isomorphic to the PINN gradient pathology [23]: when $\|\partial\mathcal{L}_{\text{pde}}/\partial\theta\| \gg \|\partial\mathcal{L}_{\text{data}}/\partial\theta\|$, the physics term dominates and the network ignores measurements.

In governance: when the collective loop dominates too strongly (command), local variety (individual Firstness) is ignored. The cure is the agapic containment structure—exactly as the PINN cure is adaptive weight balancing between loss terms.

11.4 The Inventor/Discoverer Question, Revisited

The convergence of five independent research traditions on the arity sequence 1, 2, 3, 4 (Peirce, Burgess, Perez, Faggin, Dehaene), and the independent appearance of π in QPT 3.8’s generating formula $G(\pi, \dots)$, in Zipf normalisation, in Gaussian machine learning, and in the Basel series, strongly suggests that these are discoveries of a shared underlying structure rather than coincidental inventions.

The most parsimonious explanation: the arity ladder reflects the prescissive ordering of ontological dependence (QPT 3.8), which is itself a structural feature of the universe—not of any particular human formalisation of it. And π is the invariant that surfaces whenever a system integrates over the rotational symmetry of that structure, at whatever layer.

12. Conclusion

We have unified four independent lines of inquiry into a single framework:

1. **Euler’s Basel result and π :** infinity tamed by rotation invariance.
2. **PINNs:** infinity tamed by governing law residuals, with “the equation is the data” as the central epistemological move.
3. **The Semiotic Web’s hash-of-hashes:** infinity tamed by cryptographic commitment at every rung of the arity ladder.
4. **QPT 3.8’s $G(\pi, \dots)$:** infinity tamed by the prescissive ordering of semiotic phase space, with π as the rotational invariant of invariant generation itself.

The result is a five-layer stack (physical / statistical / law / semantic / agapistic), three computable invariants (i , π , CTI hash), and a meaning-informed loss function that is the semantic analogue of the PINN loss.

The arity ladder—1, 2, 3, 4, ∞ —is not an engineering choice. It is the prescissive ordering of ontological dependence, surfacing independently in every framework that has seriously asked: what is the minimal structure required for meaning?

The Universe is perfused with signs, if it is not composed exclusively of signs.
—C. S. Peirce, 1906

A PINN is the same neural network you already know. With one extra term in the loss function. That term is the entire difference between a network that is blind to physics and a network that has been taught to see it.
—Swarnendu Bhattacharya, 2026

Intelligence is not what you accumulate. Intelligence is what flows through you toward the reduction of stress in the system you serve.
—Éric Blaettler, *The Semiotic Web*, 2026

Acknowledgements

The author thanks Swarnendu Bhattacharya (DrSwarnenduAI) for his definitive guide to Physics-Informed Neural Networks, which provided the PINN analogy developed in this paper; Carlos E. Perez for QPT 3.8 and its generating formula $G(\pi, \dots)$; Tim Andersen for his treatment of imaginary numbers in quantum mechanics; and the communities around Peirce, Burgess, Dehaene, Faggin, LeCun, Levin, and Juarrero whose independent contributions converge in the Semiotic Web.

References

- [1] T. Andersen, “Imaginary numbers in quantum mechanics may point to a deeper reality,” *The Infinite Universe*, Medium, February 2026.
- [2] T. Andersen, “A visual introduction to classical mechanics in phase space,” Medium, July 2020.
- [3] É. Blaettler, “The Semiotic Web: Eighteen convergent traditions through the lens of the Global Neuronal Workspace (Codename: Falcon Architecture),” *tokum.ai* preprint, February–March 2026. <https://www.tokum.ai>
- [4] É. Blaettler, “Tokumspace architecture specification: Semiotic Web Blueprint,” *tokum.ai*, Version 2.0, March 2026.
- [5] É. Blaettler, “The Semiotic Web Falcon Architecture: Companion Technical Document,” *tokum.ai*, March 2026.
- [6] M. Burgess, *Promise Theory: Principles and Applications*. Xtics, 2015.
- [7] M. Burgess, *Semantic Spacetime*. Xtics, 2017.

- [8] S. Dehaene, H. Lau, and S. Kouider, “What is consciousness, and could machines have it?” *Science*, vol. 358, no. 6362, pp. 486–492, 2017.
- [9] L. Euler, “De summis serierum reciprocarum,” *Opera Omnia*, Series I, Vol. 14, 1740 (presented 1735).
- [10] F. Faggin, *Silicon: From the Invention of the Microprocessor to the New Science of Consciousness*. Waterside Productions, 2021.
- [11] A. Koestler, *The Ghost in the Machine*. Hutchinson, 1967.
- [12] Y. LeCun, “A path towards autonomous machine intelligence,” OpenReview, 2022.
- [13] M. Levin, “Technological approach to mind everywhere,” *Frontiers in Systems Neuroscience*, vol. 13, 768201, 2022.
- [14] C. S. Peirce, *Collected Papers*, Vols. 1–8. Harvard University Press, 1931–1958.
- [15] R. Penrose, *The Emperor’s New Mind*. Oxford University Press, 1989.
- [16] R. Penrose, *Shadows of the Mind*. Oxford University Press, 1994.
- [17] C. E. Perez, “Quaternion process theory: A framework for artificial cognition,” Preprint, 2024.
- [18] C. E. Perez, “The architecture of thinking together: QPT 3.8 as a novel framework for modeling collective intelligence,” *Intuition Machine*, Medium, March 2026.
- [19] N. Rahaman *et al.*, “On the spectral bias of neural networks,” *ICML*, 2019.
- [20] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.
- [21] S. Bhattacharya (DrSwarnenduAI), “ π is not a circle. It is the shape of uncertainty. And it is running inside every model you have ever used,” *Towards Data Science*, Medium, March 2026.
- [22] S. Bhattacharya (DrSwarnenduAI), “Physics-informed neural networks: The complete guide to making neural networks obey the laws of physics,” *Data Science Collective*, Medium, March 2026.
- [23] S. Wang, Y. Teng, and P. Perdikaris, “Understanding and mitigating gradient flow pathologies in physics-informed neural networks,” *SIAM Journal on Scientific Computing*, vol. 43, no. 5, 2021.
- [24] E. P. Wigner, “The unreasonable effectiveness of mathematics in the natural sciences,” *Communications on Pure and Applied Mathematics*, vol. 13, pp. 1–14, 1960.