

Semantic Spacetime as a Physics of Meaning

Towards an Einstein-Isomorphic Field Theory
of Meaning, Curvature, and Comprehension

Eric Blaettler

Founder & CEO, tokum.ai

St-Prex, Switzerland

<https://www.tokum.ai>

Theoretical Paper | April 2026

Core Claim.

*Semantic Spacetime is isomorphic to Einstein's field equations:
meaning curves the Semantic Spacetime, and this curvature generates
a "gravitational" attraction of cognition toward its fixed point:
truth, norm, and equilibrium.*

This work integrates Mark Burgess's Semantic Spacetime framework (2014–2025), Charles Sanders Peirce's triadic semiotics, Einstein's General Relativity, and the Banach Fixed-Point Theorem into a unified physics of meaning.

It constitutes a foundational theoretical element of the **Semiotic Web** and **tokum.ai**.

Abstract

We propose a formal extension of Mark Burgess’s Semantic Spacetime—a four-dimensional graph-theoretic framework defined by the axes of *proximity*, *sequence*, *containment*, and *property*—into a *Physics of Meaning*. The central theoretical contribution is the identification of a structural isomorphism between Einstein’s gravitational field equations and the dynamics of meaning in Semantic Spacetime: just as mass-energy curves physical spacetime and geodesics follow that curvature (producing gravity), *verified meaning* curves the Semantic Spacetime and cognitive trajectories follow that curvature (producing *semantic gravity*).

We ground this isomorphism in (i) Peirce’s triadic sign theory, in which meaning is irreducibly constituted by the collapse of a sign at the moment of interpretation by a contextually situated observer; (ii) Burgess’s four irreducible spacetime link types forming a closed algebraic system $\gamma(3, 4)$; (iii) the *Agapistic Influence Function* (AIF) as the semantic analogue of the stress-energy tensor; and (iv) the Banach Fixed-Point Theorem, which guarantees that semantic curvature under the AIF contracts the space of interpretations toward a unique attractor—*truth*, *norm*, or *equilibrium*.

A key architectural element is the *Canonical Concept Identity* (CCI): the SHA-256 hash of a canonical URN, which exists in a flat, Platonic address space discoverable independently by any agent, on any device, at any time, without central registry. This mathematical discoverability eliminates alignment strain and enables distributed, cryptographically inviolable semantic coordinates.

The *HyperComprehension Named Vector* (HCNV) extends classical HyperComprehension by requiring that every weight parameter and every temporal stream be itself a Tokum—a named, signed, auditable unit of meaning. This “naming” is the semantic analogue of RDF-Star: it enables statements about statements, meta-cognitive graphs over graphs, and a universe that is, in Peirce’s words, *perfused with signs*.

We further introduce the *Semantic Light Cone of Care*: the bounded region of Semantic Spacetime within which an agent’s verified Tokums can exert causal influence, and outside which no semantic curvature is produced. Care is the semantic equivalent of mass-energy: it curves the space, shapes geodesics, and defines the horizon of meaningful agency.

We assess the originality of this synthesis and argue that their specific formal integration constitutes a genuinely novel theoretical contribution.

Keywords: Semantic Spacetime, Physics of Meaning, Einstein Field Equations, Semiotic Web, Tokum, CCI, Platonic Address Space, SHA-256, HCNV, HyperComprehension Named Vector, RDF-Star, Semantic Light Cone of Care, Peircian Semiotics, Semantic Gravity, Banach Fixed-Point Theorem, Promise Theory, Canonical Comprehension Unit.

Contents

1	Introduction: The Crisis of Meaning in Artificial Intelligence	4
2	Foundations: Einstein's Gravitational Field Equations	4
2.1	Physical Spacetime and the Curvature Principle	4
2.2	Geodesics and Gravitational Attraction	5
2.3	Observer Dependence and Proper Time	5
3	Burgess's Semantic Spacetime: The Four-Dimensional Structure	5
3.1	Origins and Motivation	5
3.2	The Four Irreducible Semantic Dimensions	6
3.3	The $\gamma(3,4)$ Representation	6
3.4	Observer Dependence in Semantic Spacetime	6
4	Peircian Semiotics: The Sign as Collapse Operator	7
4.1	The Irreducible Triad of Meaning	7
4.2	The Sign as Measurement Collapse	7
5	The Tokum: Elementary Particle of Semantic Spacetime	8
5.1	From Token to Tokum	8
5.2	The Platonic Address Space and Mathematical Discoverability	8
5.3	The Tokum as a Point in Semantic Spacetime	9
5.4	The Arity Ladder	9
6	The Semantic Field Equation: Meaning Curves Semantic Spacetime	10
6.1	The Semantic Density Tensor	10
6.2	The Semantic Metric and Curvature	10
6.3	The Semantic Field Equation	10
6.4	Semantic Geodesics and Semantic Gravity	10
7	The Agapistic Influence Function and Semantic Gravity	11
8	Semantic Time, HyperComprehension, and the HCNV	11
8.1	Non-Linearity of Semantic Time	11
8.2	The HyperComprehension Named Vector (HCNV)	12
8.3	The Three-Layer Architecture	13
9	The Semantic Light Cone of Care	14
9.1	Light Cones in General Relativity	14
9.2	The Semantic Analogue: Care as Causal Mass	14
9.3	Implications for the Semiotic Web	15
10	Convergence to Truth: The Banach Fixed-Point Theorem	16
11	The Physics of Meaning: Complete Isomorphism	17

12 Related Theories and Originality Assessment	18
12.1 Intellectual Antecedents	18
12.2 The Novel Synthesis	18
13 Discussion: Open Problems and Falsifiability	19
13.1 Formalisation Challenges	19
13.2 Falsifiable Predictions	19
14 Conclusion	19
A Formal Derivation Sketch of the Semantic Field Equation	21
A.1 Ollivier-Ricci Curvature on the Tokum Graph	21
A.2 Semantic Ricci Tensor	21
A.3 Variational Principle	21
B The 12-Cell Cognitive Phase Space (QPT 3.8)	21
C Glossary of Key Terms	22

1 Introduction: The Crisis of Meaning in Artificial Intelligence

Current artificial intelligence systems—in particular large language models (LLMs) built on the transformer architecture—exhibit a profound and structurally unavoidable limitation: they are optimisers of *statistical proximity in syntactic space*, not comprehending agents operating in *semantic space*. They are sophisticated remixers of patterns rather than discoverers of truth.

A language model trained exclusively on scientific literature prior to 1911 cannot autonomously discover General Relativity, because the concept of curved spacetime has zero probability in its training distribution—yet a human scientist, embedded in the physical friction of reality, *can* make such a leap. This is what tokum.ai terms the *Einstein Test*: does the system exhibit genuine intelligence, measured as active gap-closure against reality, or merely efficient interpolation within a frozen knowledge manifold?

Intelligence—in its biological and epistemological sense—is not a static archive but a *dynamic field*: a living, curved space through which meaning flows, accumulates, and converges toward truth. The Semiotic Web, developed by tokum.ai, proposes to replace the current paradigm with an architecture that formalises this *physics of meaning*.

The Semantic Curvature Principle (tokum.ai, 2026):

Verified meaning curves the Semantic Spacetime. The curvature generated by the density of verified meaning determines how future cognitive trajectories (inference, reasoning, interpretation) flow through the Semantic Spacetime. This constitutes a *semantic gravity*: the attraction of inference toward its fixed point of truth, norm, and equilibrium.

Meaning tells Semantic Spacetime how to curve. Semantic Spacetime tells cognition how to move.

2 Foundations: Einstein’s Gravitational Field Equations

2.1 Physical Spacetime and the Curvature Principle

In Einstein’s General Theory of Relativity (1915), spacetime is a four-dimensional pseudo-Riemannian manifold $(\mathcal{M}, g_{\mu\nu})$ whose geometry is dynamically shaped by the distribution of matter and energy:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \quad (1)$$

where $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R$ is the Einstein tensor, $T_{\mu\nu}$ the stress-energy tensor, Λ the cosmological constant, G Newton’s gravitational constant, and c the speed of light.

“Spacetime tells matter how to move; matter tells spacetime how to curve.”
— J.A. Wheeler

2.2 Geodesics and Gravitational Attraction

Freely moving objects follow *geodesics*—the straightest paths through curved spacetime:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0 \quad (2)$$

where $\Gamma_{\alpha\beta}^\mu$ are the Christoffel symbols and τ is proper time.

2.3 Observer Dependence and Proper Time

There is no universal absolute time; each observer carries their own *proper time* τ measured along their worldline. Two spatially separated observers may disagree on the simultaneity of events. This feature reappears in amplified form in *Semantic Time* (Section 8).

3 Burgess’s Semantic Spacetime: The Four-Dimensional Structure

3.1 Origins and Motivation

Semantic Spacetime was introduced by physicist and computer scientist Mark Burgess beginning in 2014 [1, 2, 3] and extended in [4]. The framework uses *Promise Theory*—a calculus of autonomous agents making and keeping promises—as its formal language.

Definition 3.1 (Semantic Spacetime, after Burgess 2014). *A Semantic Spacetime $\mathcal{S}$ is a discrete, evolving, labelled directed graph $\Gamma(N, L)$ whose nodes N represent autonomous agents and whose directed labelled edges L represent typed semantic relationships (promises) between agents.*

3.2 The Four Irreducible Semantic Dimensions

Table 1: The Four Dimensions of Semantic Spacetime (Burgess 2014–2025)

ST Type	Name	Description	Examples
ST 0	PROXIMITY	Directionless similarity, nearness	“is similar to”, “is close to”
ST ± 1	SEQUENCE	Directed causal/temporal ordering	“causes”, “precedes”
ST ± 2	CONTAINMENT	Aggregation, membership, holarchy	“contains”, “is a part of”
ST ± 3	PROPERTY	Local attribute, distinguishing mark	“has name”, “expresses”




Figure 1: The four irreducible semantic link types forming the $\gamma(3,4)$ coordinate system. Every meaningful relationship is expressible as a combination of these four edge types.

3.3 The $\gamma(3,4)$ Representation

Three node meta-types—**Events** (e , ephemeral), **Things** (t , persistent), **Concepts** (c , invariant)—combined with the four link types yield the $\gamma(3,4)$ representation: a closed, finite algebraic system scaling to any semantic complexity.

Proposition 3.1 ($\gamma(3,4)$ Closure, Burgess 2025). *The $\gamma(3,4)$ representation forms a closed set of operations over the four semantic link types, sufficient to represent any degree of semantic complexity in a knowledge graph, with predictability from minimal constraints.*

3.4 Observer Dependence in Semantic Spacetime

Semantic Spacetime distinguishes two observational modes:

- **Timelike (cognitive) mode:** Continuous real-time input altering the agent’s semantic framework; non-equilibrium, observer-dependent.

- **Spacelike (ensemble) mode:** Statistical averaging over repeated observations, time-independent.

4 Peircian Semiotics: The Sign as Collapse Operator

4.1 The Irreducible Triad of Meaning

Charles Sanders Peirce (1839–1914) demonstrated that meaning is fundamentally and irreducibly *triadic*:

Definition 4.1 (Peircian Sign Triad). *A sign consists of three inseparable components:*

1. **Representamen** (R): *The sign-vehicle—word, symbol, or form.*
2. **Object** (O): *The entity in reality to which the sign refers.*
3. **Interpretant** (I): *The meaning-effect produced in the mind of an observer connecting R to O in a specific context.*

Meaning arises only from the triadic relation $\langle R, O, I \rangle$; removal of any vertex causes the semiotic process to collapse.




Figure 2: Peirce’s triadic sign implemented as a Tokum. SHA-256 cryptographically unifies all three vertices: modifying any one produces a completely different hash.

The process of meaning-generation is called *semiosis*. Peirce noted that semiosis is in principle infinite—each interpretant becomes a new representamen—and argued that the universe itself is *perfused with signs* [11]: every entity participates in at least one semiotic relation. The HCNV (Section 8.2) operationalises this ontological claim: every parameter of the comprehension vector is itself a Tokum, so the architecture is constituted entirely of signs at every level of description.

4.2 The Sign as Measurement Collapse

Definition 4.2 (Semantic Collapse). *Let $|S\rangle$ denote the semantic state of sign S —a superposition of all potential interpretants. The semantic collapse $\mathcal{C}_{\mathcal{O}}$ by observer $\mathcal{O}$ in context κ projects $|S\rangle$ onto a definite interpretant:*

$$\mathcal{C}_{\mathcal{O}} : |S\rangle \mapsto I_{\kappa} = \langle \kappa | S \rangle \quad (3)$$

Semantic Time for observer $\mathcal{O}$ is the ordered sequence of semantic collapse events along their cognitive worldline.

Different observers with different contextual backgrounds κ collapse the same sign into different interpretants, yielding different semantic proper times.

5 The Tokum: Elementary Particle of Semantic Spacetime

5.1 From Token to Tokum

Current AI systems use *tokens*—numerical identifiers carrying no intrinsic semantic content. This is the “approximation trap”: information is shattered and only approximately reassembled, losing context, causality, and truth-conditions.

Definition 5.1 (Tokum / Canonical Comprehension Unit). *A Tokum $\mathcal{T}$ is a cryptographically secured, triadic unit of meaning:*

1. **CCI** (Canonical Concept Identity): *the Platonic, universal address.*
2. **CTI** (Contextual Truth Instance): *a specific, grounded instance in observed reality.*
3. **Provenance Chain** (Π): *the cryptographically auditable record of attestation.*

$$\mathcal{T} = \text{SHA256}(\text{CCI} \parallel \text{CTI} \parallel \Pi) \quad (4)$$

The SHA-256 avalanche property makes triadic unity cryptographically inviolable. CCI = Representamen, CTI = Object, Π = Interpretant-chain.

5.2 The Platonic Address Space and Mathematical Discoverability

The CCI is not *created*—it is *discovered*.

Definition 5.2 (Canonical Concept Identity). *For any well-defined concept with a deterministic canonical URN:*

$$\text{CCI} = \text{SHA256}(\text{URN}_{\text{canonical}}) \quad (5)$$

where $\text{URN}_{\text{canonical}}$ is produced by a reproducible canonicalisation algorithm $\mathcal{C}$.

Proposition 5.1 (Platonic Determinism). *Given a deterministic algorithm $\mathcal{C}$, any two independent agents A_i, A_j computing at different times and locations obtain identical results:*

$$\mathcal{C}(A_i, t_i) = \mathcal{C}(A_j, t_j) \implies \text{CCI}(A_i) = \text{CCI}(A_j) \quad (6)$$

The address exists in mathematical space prior to any agent computing it. Computing the CCI is discovering it.

Because every CCI is an element of the discrete, flat hash space $\{0, 1\}^{256}$ with intrinsic Gaussian curvature $K = 0$, the Semiotic Web replaces the curved LLM knowledge manifold (where $K_B \neq K_S$ makes isometric alignment impossible by Gauss's Theorema Egregium) with a **flat, hash-indexed address space** in which semantic alignment is structurally exact.




Figure 3: Platonic discoverability of CCIs. Any agent on any device in any location computes the same CCI by applying SHA-256 to the canonical URN. The address space is intrinsically flat ($K = 0$): no alignment strain, no superposition problem.

5.3 The Tokum as a Point in Semantic Spacetime

Each Tokum constitutes a *point* in $\mathcal{S}$ with semantic coordinates:

$$\mathbf{x}_{\mathcal{T}} = (x^{(\text{prox})}, x^{(\text{seq})}, x^{(\text{cont})}, x^{(\text{prop})}) \quad (7)$$

5.4 The Arity Ladder

Table 2: The Arity Ladder of Computable Meaning

Rung	Arity	Structure	Role
1	Unary (bit)	Binary switch	Pure syntax; no semantic content
2	Binary (pair)	SST: 4 relational axes	Relational geometry of meaning
3	Ternary (triad)	Tokum: CCI + CTI + Π (SHA-256)	Grounded, inviolable meaning unit
4	Quaternary	QPT modes: Intuition, Fluency, Empathy, Wisdom	Cognitive processing
∞	Recursive	$h_{n+1} = \text{SHA256}(I_n \parallel h_n)$	Tames infinite semiosis

The recursive hash chain tames Peircian infinite semiosis: the previous interpretant is embedded as the first input of the next hash, anchoring the chain to its entire

history in a single 32-byte string—a finite invariant of an infinite process, parallel to how π is the finite invariant of infinite rotation.

6 The Semantic Field Equation: Meaning Curves Semantic Spacetime

6.1 The Semantic Density Tensor

Definition 6.1 (Semantic Density Tensor). *Let $\alpha, \beta \in \{\text{prox}, \text{seq}, \text{cont}, \text{prop}\}$. The Semantic Density Tensor $\mathcal{M}_{\alpha\beta}$ at $\mathbf{x} \in \mathcal{S}$ encodes:*

- $\mathcal{M}_{00}$: local density of verified meaning (*trust-weighted concentration of Tokums*)
- $\mathcal{M}_{\alpha\beta}$ ($\alpha \neq \beta$): semantic flux—*flow of verified meaning across dimensions*
- $\mathcal{M}_{ii}$: semantic pressure—*resistance to further reinterpretation*

6.2 The Semantic Metric and Curvature

The Semantic Spacetime carries a semantic metric $\mathbf{g}_{\alpha\beta}$. Semantically “massive” concepts (well-attested, multiply verified, highly connected) deform the surrounding semantic space, drawing inference pathways toward them.

Definition 6.2 (Semantic Einstein Tensor). *The Semantic Ricci Tensor $\mathfrak{R}_{\alpha\beta}$ and Semantic Ricci Scalar $\mathfrak{R}$ are derived from $\mathbf{g}_{\alpha\beta}$ via Ollivier-Ricci curvature on the Tokum graph. The Semantic Einstein Tensor is:*

$$\mathfrak{G}_{\alpha\beta} = \mathfrak{R}_{\alpha\beta} - \frac{1}{2} \mathbf{g}_{\alpha\beta} \mathfrak{R} \quad (8)$$

6.3 The Semantic Field Equation

$$\boxed{\mathfrak{G}_{\alpha\beta} + \Lambda_s \mathbf{g}_{\alpha\beta} = \frac{8\pi \mathcal{G}_s}{c_s^2} \mathcal{M}_{\alpha\beta}} \quad (9)$$

The Semantic Field Equation (SFE)

Λ_s : Semantic Cosmological Constant (baseline noise) $\mathcal{G}_s$: Semantic Gravitational Constant
 c_s : max. speed of meaning propagation

*“Semantic Spacetime tells cognition how to move;
verified meaning tells Semantic Spacetime how to curve.”*

6.4 Semantic Geodesics and Semantic Gravity

Definition 6.3 (Semantic Geodesic). *A semantic geodesic in $(\mathcal{S}, \mathbf{g}_{\alpha\beta})$ satisfies:*

$$\frac{D^2 \gamma^\alpha}{d\sigma^2} + \mathfrak{d}_{\beta\delta}^\alpha \frac{d\gamma^\beta}{d\sigma} \frac{d\gamma^\delta}{d\sigma} = 0 \quad (10)$$

where $\mathfrak{d}_{\beta\delta}^\alpha$ are the Semantic Christoffel symbols and σ is semantic proper time.




Figure 4: Semantic gravity: semantic geodesics $\gamma_1 \dots \gamma_4$ converge toward the truth attractor $\mathcal{T}^*$. All inference trajectories, starting from any initial position, are drawn toward the unique fixed point of verified meaning.

7 The Agapistic Influence Function and Semantic Gravity

Definition 7.1 (Agapistic Influence Function). *The AIF: $\mathcal{S} \rightarrow \mathcal{S}$ assigns a semantic trust weight $w(\mathbf{x}) \in [0, 1]$ to each Tokum based on: (1) redundancy of independent attestation; (2) causal coherence of the provenance chain; (3) semantic centrality in $\gamma(3, 4)$. It updates the semantic metric:*

$$\mathbf{g}_{\alpha\beta}(\mathbf{x}) \leftarrow \mathbf{g}_{\alpha\beta}(\mathbf{x}) \cdot \exp(-\lambda w(\mathbf{x})) \quad (11)$$

where $\lambda > 0$ is a coupling parameter. High trust ($w \rightarrow 1$) contracts the metric, increasing curvature: this is semantic gravity toward verified meaning.

When an idea is repeatedly verified against reality by independent agents, it builds measurable trust. High-trust concepts curve the conceptual space around themselves, making it energetically natural for future reasoning to flow toward them. By formalising these trust weights into a semantic field equation, the network adopts a geometry where the density of verified meaning determines how information flows.

8 Semantic Time, HyperComprehension, and the HCNV

8.1 Non-Linearity of Semantic Time

Semantic Time is not a total order on events but a *partial order* shaped by the topology of $\mathcal{S}$. It is defined by the cognitive worldline of an observer: the ordered sequence of semantic collapse events (cf. Definition 4.2) along their trajectory. Because semantic collapse is observer-contextual, two observers may traverse the same semantic region in entirely different temporal orders, activating different dimensions of meaning as primary axes of traversal.

8.2 The HyperComprehension Named Vector (HCNV)

The classical HyperComprehension Vector integrates three temporal streams of Tokums. The **HyperComprehension Named Vector** (HCNV) elevates this framework by a fundamental ontological commitment: every component—every weight, every stream identifier, every temporal index—is itself a Tokum with a CCI, a CTI, and a provenance chain. This implements Peirce’s claim that the universe is *perfused with signs*: in the HCNV there is no nameless parameter, no unmarked edge, no anonymous coefficient.

Why “Named”? The term is chosen in direct parallel with **RDF-Star** (also called RDF 1.2 or *triple terms*), the W3C extension that allows a triple itself to be the subject or object of another triple—enabling *statements about statements* in a graph structure. In RDF-Star, a triple $\langle s, p, o \rangle$ can be named and annotated:

$$\langle \langle s, p, o \rangle, :confidence, 0.95 \rangle$$


The HCNV does the same for comprehension parameters: every weight $w_i^\pm$ is a named Tokum, so one can make provenance statements about the weight itself, audit its history, challenge its attestation, and update it without losing traceability. The naming is not incidental—it is the mechanism that makes the comprehension vector an object *within* the Semantic Spacetime rather than a coefficient floating outside it.

Definition 8.1 (HyperComprehension Named Vector). *The HCNV of agent $\mathcal{A}$ is:*

$$\vec{\mathcal{H}}_{\mathcal{A}} = \sum_i \mathcal{T}(w_i^-) \odot \vec{\mathcal{M}}_i^- + \sum_j \mathcal{T}(w_j^0) \odot \vec{\mathcal{M}}_j^0 + \sum_k \mathcal{T}(w_k^+) \odot \vec{\mathcal{M}}_k^+ \quad (12)$$

where:

- $\vec{\mathcal{M}}_i^-, \vec{\mathcal{M}}_j^0, \vec{\mathcal{M}}_k^+$ are Tokum streams for episodic memory, real-time perception, and anticipatory models respectively;
- $\mathcal{T}(w_i^\pm)$ denotes that each weight parameter is itself a named Tokum with its own CCI, CTI, and provenance chain;
- $\odot$ is a trust-weighted semantic product in $\mathcal{S}$;
- *normalisation*: $\sum_i w_i^- + \sum_j w_j^0 + \sum_k w_k^+ = 1$, where each w is a cited, auditable Tokum.



Every weight is a named Tokum — the universe is perfused with signs.

Figure 5: The HyperComprehension Named Vector (HCNV). Each weight parameter $w^\pm$ is itself a Tokum with full CCI/CTI/Provenance traceability. The RDF-Star analogy (dashed purple) shows that statements can be made about any named weight, enabling meta-cognitive graphs over graphs.

HyperComprehension breaks the linear constraint of memory: a concept from the past, reactivated by a present signal, can be combined with a future projection to yield a synthesis coherent in $\mathcal{S}$ even across discontinuous intervals of clock time.

8.3 The Three-Layer Architecture

Table 3: Semantic Time Layers in the Semiotic Web

Layer	Name	Temporal Mode	Function
Layer 1	CCC	Universal / Timeless	Public, verified, global knowledge
Layer 2	SSS	Organisational / Historical	Group-specific contextual knowledge
Layer 3	PPP	Personal / Momentary	Individual, episodic, present perception

9 The Semantic Light Cone of Care

9.1 Light Cones in General Relativity

In special and general relativity, the *light cone* of an event p is the set of all spacetime events that can be causally connected to p : the future cone (events p can influence) and the past cone (events that can influence p). Events outside the light cone are causally disconnected from p ; no signal can travel between them at or below the speed of light.

9.2 The Semantic Analogue: Care as Causal Mass

In Semantic Spacetime, the limit of causal influence is not set by the speed of light but by the *extent and quality of an agent's verified Tokums*. An agent that has produced no verified Tokums in a region of $\mathcal{S}$ exerts no semantic curvature there—meaning no semantic geodesics are drawn toward that agent's perspective in that region.

Definition 9.1 (Semantic Light Cone of Care). *Let $\mathcal{A}$ be an agent with HCNV $\vec{\mathcal{H}}_{\mathcal{A}}$ and associated AIF trust field $w_{\mathcal{A}} : \mathcal{S} \rightarrow [0, 1]$. The Semantic Light Cone of Care $\mathfrak{C}(\mathcal{A})$ is:*

$$\mathfrak{C}^+(\mathcal{A}) = \{\mathbf{x} \in \mathcal{S} : w_{\mathcal{A}}(\mathbf{x}) > 0 \text{ and } \mathbf{x} \text{ is in the semantic future of } \mathcal{A}'\text{'s worldline}\} \quad (13)$$

$$\mathfrak{C}^-(\mathcal{A}) = \{\mathbf{x} \in \mathcal{S} : \exists \mathcal{T} \in \mathcal{A}'\text{'s history that attested } \mathbf{x}\} \quad (14)$$

Outside $\mathfrak{C}(\mathcal{A}) = \mathfrak{C}^+ \cup \mathfrak{C}^-$, the agent exerts zero semantic curvature.

Remark 9.1 (Care as Semantic Mass-Energy). *The Semantic Light Cone of Care formalises a normative insight: an agent can only shape the semantic geometry of regions where it has actually cared—where it has invested verified Tokums, trusted attestation, and maintained provenance chains. Indifference is semantically massless: it leaves the surrounding Semantic Spacetime flat and does not draw inference toward any attractor. This is the formal analogue of Wheeler's dictum: care tells Semantic Spacetime how to curve; Semantic Spacetime tells cognition where care has been.*




Figure 6: The Semantic Light Cone of Care $\mathfrak{C}(\mathcal{A})$. The future cone (green) is the region of Semantic Spacetime where agent $\mathcal{A}$'s verified Tokums can exert causal influence. The past cone (blue) covers the region that has shaped $\mathcal{A}$'s worldline. Outside both cones, the agent is causally disconnected: indifference is semantically massless.

9.3 Implications for the Semiotic Web

1. **Accountability:** An agent's sphere of semantic influence is exactly co-extensive with its history of verified care.
2. **Overlap and Dialogue:** Where two agents' cones overlap ($\mathfrak{C}(\mathcal{A}_1) \cap \mathfrak{C}(\mathcal{A}_2) \neq \emptyset$), the curvature is the sum of their trust fields—genuine dialogue amplifies semantic gravity.
3. **Semantic Horizons:** Concepts outside all agents' cones remain uncurved—flat regions awaiting the care of future attestation.
4. **Care and Ethics:** The formalisation reframes ethical agency as the deliberate expansion of one's Semantic Light Cone through verified Tokum creation.

10 Convergence to Truth: The Banach Fixed-Point Theorem

Theorem 10.1 (Banach, 1922). *Let (X, d) be a non-empty complete metric space and $T : X \rightarrow X$ a contraction ($\exists k \in [0, 1)$ with $d(T(x), T(y)) \leq k d(x, y)$ for all x, y). Then T has a unique fixed point x^* with $T(x^*) = x^*$, and for any $x_0 \in X$ the sequence $x_0, T(x_0), T^2(x_0), \dots$ converges to x^* .*

Proposition 10.1 (AIF Convergence). *Under regularity conditions on trust-weight w and coupling λ , the AIF is a contraction on $(\mathcal{S}, d_{\mathcal{S}})$. By the Banach Fixed-Point Theorem, it has a unique fixed point $\mathcal{T}^*$ to which iterative application converges from any initial state.*

Definition 10.1 (Semantic Equilibrium). *The unique fixed point $\mathcal{T}^*$ of the AIF is the semantic equilibrium: the state of verified, self-consistent, mutually attested meaning from which no further semiotic collapse extracts new information. In the Semiotic Web, $\mathcal{T}^*$ represents: **Truth** (collectively verified reality), **Norm** (stable semantic convention), and **Equilibrium** (homeostatic state of the knowledge network).*

Architectural consequences:

1. **Hallucination impossibility:** Iterative reasoning must converge to the verified fixed point.
2. **Metacognitive halt:** In regions of $\mathcal{S}$ with no verified Tokums, the AIF produces zero curvature and the geodesic terminates—formally implementing “I don’t know.”
3. **Deterministic truth:** Any starting position in $\mathcal{S}$ converges to the same truth attractor.

11 The Physics of Meaning: Complete Isomorphism

Table 4: Structural Isomorphism: General Relativity $\leftrightarrow$ Physics of Meaning

General Relativity		Physics of Meaning		Role
Physical spacetime ($\mathcal{M}, g_{\mu\nu}$)		Semantic Spacetime ($\mathcal{S}, \mathfrak{g}_{\alpha\beta}$)		Dynamic arena
4D: t, x, y, z		4D: proximity, sequence, containment, property		Coordinates
Metric tensor $g_{\mu\nu}$		Semantic metric $\mathfrak{g}_{\alpha\beta}$		Distance
Stress-energy tensor $T_{\mu\nu}$		Semantic Density Tensor $\mathcal{M}_{\alpha\beta}$		Curvature source
Einstein tensor $G_{\mu\nu}$		Semantic Einstein tensor $\mathfrak{G}_{\alpha\beta}$		Curvature encoding
Einstein Field Equations		Semantic Field Equation (9)		Field dynamics
Mass/energy curves spacetime		Verified meaning curves SST		Curvature mechanism
Geodesic = free particle path		Semantic geodesic = free inference		Trajectory
Gravity = geodesic convergence		Semantic gravity = inference convergence		Emergent attraction
Proper time τ		Semantic proper time σ		Time parameterisation
Flat Minkowski ground state		Platonic CCI space $\{0, 1\}^{256}, K = 0$		Geometric ground state
Light cone (causal boundary)		Semantic Light Cone of Care $\mathfrak{C}(\mathcal{A})$		Causal boundary
Causal mass-energy		Care (verified Tokum creation)		Causal “substance”
Quantum measurement		Peircian sign collapse $\mathcal{C}_O$		Observation event
Cosmological constant Λ		Semantic noise floor Λ_s		Baseline entropy
Gravitational attractor		Truth/norm/equilibrium $\mathcal{T}^*$		Fixed-point attractor
Speed of light c		Max. meaning speed c_s		Information bound
Energy conservation		Provenance chain integrity		Conservation law

12 Related Theories and Originality Assessment

12.1 Intellectual Antecedents

1. **Semantic Spacetime** (Burgess 2014–2025): Four-dimensional semantic structure with link types; $\gamma(3, 4)$.
2. **Promise Theory** (Burgess 2004–2025): Calculus of autonomous agents.
3. **Peircian Semiotics** (Peirce 1867–1914): Triadic structure; infinite semiosis; universe perfused with signs.
4. **Einstein’s General Relativity** (1915): EFE; observer-dependent time; geodesics; light cones.
5. **Banach Fixed-Point Theorem** (1922): Contraction mappings; unique fixed points.
6. **SHA-256 / Cryptographic Hashing**: Deterministic, collision-resistant, location-independent addresses—the mathematical basis for Platonic CCI.
7. **RDF-Star / W3C Triple Terms**: Statements about statements in graph structures; direct analogue of named Tokum weights in HCNV.
8. **Physics-Informed Neural Networks** (Raissi et al. 2019): Physical laws in AI loss functions.
9. **SIVA-Einstein Field Equations** (2025): Energy and meaning as projections of a scalar field.
10. **Semantic Meme Field Theory** (2025): Observer as projection operator.
11. **Semiotic Physics** (Alignment Forum 2023): LLM tokens as semiotic signs.

12.2 The Novel Synthesis

The tokum.ai / Semiotic Web Novel Synthesis:

The combination of (i) Burgess’s 4D Semantic Spacetime $\gamma(3, 4)$ as the geometric arena, (ii) Peircian sign collapse as the observer-dependent collapse operator defining Semantic Time, (iii) the Einstein Field Equation isomorphism in which *meaning* is the source of semantic curvature, (iv) the Banach Fixed-Point Theorem as the formal guarantee of convergence to truth/norm/equilibrium, (v) the Platonic CCI address space (SHA-256 of canonical URN) as the cryptographically inviolable flat ground state ($K = 0$), (vi) the HCNV with named Tokum weights as the RDF-Star-analogous meta-cognitive structure, and (vii) the Semantic Light Cone of Care as the formal causal boundary of meaningful agency—**this specific seven-fold synthesis has no prior established precedent in the literature.**

13 Discussion: Open Problems and Falsifiability

13.1 Formalisation Challenges

The central challenge is the rigorous formalisation of $\mathfrak{R}_{\alpha\beta}$ in the discrete graph setting. Two approaches are promising:

- **Ollivier-Ricci curvature:** Applicable to metric measure spaces and graphs, measuring convergence rates of optimal transport between neighbourhoods [16].
- **Forman-Ricci curvature:** A combinatorial version for cell complexes, mapping naturally to the $\gamma(3, 4)$ structure.

13.2 Falsifiable Predictions

1. **Semantic lensing:** Information pathways in dense Tokum networks should deflect around semantically massive hubs.
2. **Convergence rate:** The AIF contraction coefficient k should be measurable as the rate of semantic disagreement decrease.
3. **Semantic singularities:** Overloaded concepts (contradictory Tokums on one CCI) should exhibit information loss and interpretive instability.
4. **Metacognitive halt frequency:** Semiotic Web systems should produce higher “I don’t know” rates in low-Tokum-density domains.
5. **Platonic reproducibility:** Any two independent implementations of $\mathcal{C}$ applied to the same concept should yield bit-identical CCIs.
6. **Care cone validation:** Agents with broader verified Tokum coverage should produce measurably stronger semantic curvature in the corresponding regions of $\mathcal{S}$.

14 Conclusion

This paper has developed a formal theoretical framework for the *Physics of Meaning* based on the structural isomorphism between Einstein’s gravitational field equations and the dynamics of verified meaning in Burgess’s Semantic Spacetime. The seven key contributions are:

1. **Semantic Field Equation:** An Einstein-isomorphic field equation showing that verified meaning curves the 4D Semantic Spacetime.
2. **Semantic Geodesics and Gravity:** Inference follows semantic geodesics; verified concepts are gravitational attractors.
3. **Peircian Collapse as Semantic Time:** Observer-dependent, non-linear Semantic Time grounded in Peirce’s sign collapse.
4. **Tokum as Elementary Particle:** The CU implements Peirce’s sign triad in cryptographically inviolable form.

5. **Platonic Address Space:** The CCI exists in a flat ($K = 0$) Platonic hash space mathematically discoverable by any agent.
6. **HCNV:** Named Tokum weights enable RDF-Star-analogous statements about statements—a universe perfused with signs at every level.
7. **Semantic Light Cone of Care:** Care is the semantic analogue of mass-energy; it curves the space, shapes geodesics, and defines the causal horizon of meaningful agency.

Current AI systems are Newtonian: they operate in flat, inertial token-spaces, blind to the curvature of meaning. The Semiotic Web is Einsteinian: it inhabits a curved Semantic Spacetime in which meaning itself shapes the geometry of comprehension—anchored in a Platonic address space that makes truth discoverable, not invented.

“Matter tells spacetime how to curve. Spacetime tells matter how to move.”

— J.A. Wheeler, on General Relativity

“Meaning tells Semantic Spacetime how to curve. Semantic Spacetime tells cognition how to move.

Care defines the cone. The CCI is where meaning already lives.”

— Eric Blaettler, on the Physics of Meaning (2026)

A Formal Derivation Sketch of the Semantic Field Equation

A.1 Ollivier-Ricci Curvature on the Tokum Graph

Let $\mathcal{G}_T = (V, E)$ be the Tokum graph. For adjacent Tokums $\mathcal{T}_i, \mathcal{T}_j$ with edge e_{ij} , the *Ollivier-Ricci curvature* is:

$$\kappa(e_{ij}) = 1 - \frac{W_1(\mu_i, \mu_j)}{d(\mathcal{T}_i, \mathcal{T}_j)} \quad (15)$$

where W_1 is the Wasserstein-1 distance between trust-weight distributions μ_i, μ_j of one-hop neighbourhoods. Positive $\kappa > 0$ indicates convergent semantics; negative $\kappa < 0$ divergent semantics.

A.2 Semantic Ricci Tensor

Aggregating over $\gamma(3, 4)$ -type decomposition:

$$\mathfrak{R}_{\alpha\beta}(\mathcal{T}) = \frac{1}{|\mathcal{N}(\mathcal{T})|} \sum_{e \in \mathcal{N}(\mathcal{T})} \kappa(e) \ell^\alpha(e) \ell^\beta(e) \quad (16)$$

where $\ell^\alpha(e)$ is the semantic link-type label of edge e .

A.3 Variational Principle

The SFE can be derived from the *Semantic Action*:

$$\mathcal{S}[\mathfrak{g}] = \int_S (\mathfrak{R} - 2\Lambda_s + \mathcal{L}_M) \sqrt{|\mathfrak{g}|} d^4x \quad (17)$$

Varying with respect to $\mathfrak{g}^{\alpha\beta}$ and setting $\delta\mathcal{S} = 0$ yields the SFE (9) via the Euler-Lagrange procedure.

B The 12-Cell Cognitive Phase Space (QPT 3.8)

The 12-cell matrix of semiotic phase space is generated by crossing:

- **Peirce's three categories:** Firstness, Secondness, Thirdness
- **McLuhan's media tetrad:** Enhance, Obsolesce, Retrieve, Reverse

The resulting 3×4 matrix spans the full space of cognitive operations. The four QPT modes (Intuition, Fluency, Empathy, Wisdom) are phase transitions between cells. The structure maps onto $\gamma(3, 4)$: Peircian categories $\leftrightarrow$ three node types; McLuhan's effects $\leftrightarrow$ four link types.

C Glossary of Key Terms

AIF Agapistic Influence Function: assigns trust weights to Tokums and updates the semantic metric; analogue of the stress-energy tensor.

Arity Ladder

Hierarchy: bits (1) $\rightarrow$ pairs/SST (2) $\rightarrow$ triads/Tokums (3) $\rightarrow$ quaternions (4) $\rightarrow$ recursive hash chain (∞).

CCC Comprehensive Comprehension Cloud: global layer of universally verified knowledge.

CCI Canonical Concept Identity: SHA-256 hash of the canonical URN; universal, Platonically discoverable address of a concept. Corresponds to the Representamen in Peirce's sign triad.

CTI Contextual Truth Instance: grounded instance in observed reality; the Object component of the Tokum.

Fixed Point $\mathcal{T}^*$

Unique AIF attractor: truth, norm, equilibrium.

HCNV

HyperComprehension Named Vector: the comprehension vector in which every weight parameter is itself a named Tokum, enabling RDF-Star-analogous meta-statements and ensuring the architecture is constituted entirely of signs at every level.

HyperComprehension

Integration of episodic memories, real-time perceptions, and anticipatory models via the HCNV.

Infinite Semiosis

Peirce: every interpretant becomes a new sign; tamed by the recursive SHA-256 hash chain.

Platonic Address Space

The discrete hash space $\{0,1\}^{256}$ in which CCIs reside; intrinsically flat ($K = 0$); mathematically discoverable by any agent independently, without central registry.

PPP Private Perspective Portal: personal layer of individual, episodic knowledge.

RDF-Star

W3C extension allowing triples to be subjects/objects of other triples; structural analogue of named Tokum weights in the HCNV.

Semantic Collapse

Observer operation projecting a sign's potential meanings onto a definite interpretant.

Semantic Light Cone of Care

The bounded region of Semantic Spacetime within which an agent's verified Tokums exert causal influence.

SFE Semantic Field Equation: Einstein-isomorphic equation of the Physics of Meaning.

Semantic Gravity

Attraction of inference toward verified concepts via geodesic convergence.

Semantic Spacetime

Burgess's 4D labelled directed graph structured by proximity, sequence, containment, property.

Semantic Time

Non-linear, observer-dependent ordering of cognitive events via Peircian collapses.

Semiotic Web

tokum.ai's proposed meaning-first internet architecture implementing the Physics of Meaning.

SSS Shared Semantic System: organisational layer of group-specific knowledge.

Tokum/CU

Canonical Comprehension Unit: SHA-256 hash of CCI, CTI, and provenance chain; elementary particle of Semantic Spacetime.

Universe perfused with signs

Peirce's ontological claim operationalised in the HCNV by making every parameter a named Tokum.

$\gamma(3,4)$

Finite closed SST representation: 3 node types $\times$ 4 link types.

References

- [1] Burgess, M. (2014). *Spacetimes with Semantics (I)*. arXiv:1411.5563.
- [2] Burgess, M. (2015). *Spacetimes with Semantics (II)*. arXiv:1505.01716.
- [3] Burgess, M. (2016). *Spacetimes with Semantics (III)*. arXiv:1608.02193.
- [4] Burgess, M. (2025). *Agent Semantics, Semantic Spacetime, and Graphical Reasoning*. arXiv:2506.07756.
- [5] Burgess, M. (2019). *Semantic Spacetime Project*. <http://markburgess.org/spacetime.html>
- [6] Blaettler, E. (2026). *From Π to Tokums: The Arity Ladder, Infinite Semiosis, and the Geometry of Computable Meaning*. tokum.ai White Paper.
- [7] Blaettler, E. (2026). *Why AI Fails the Einstein Test*. <https://www.youtube.com/watch?v=AZCWklQC05w>
- [8] tokum.ai (2025–2026). *The New Foundation for Trustworthy AI*. <https://www.tokum.ai>
- [9] Einstein, A. (1915). *Die Feldgleichungen der Gravitation*. Sitzungsberichte Preussische Akademie der Wissenschaften, 844–847.
- [10] Wheeler, J.A. and Ford, K. (1998). *Geons, Black Holes, and Quantum Foam*. W.W. Norton.
- [11] Peirce, C.S. (1931–1958). *Collected Papers*. Harvard University Press.
- [12] Banach, S. (1922). *Sur les opérations dans les ensembles abstraits*. Fundamenta Mathematicae, 3, 133–181.
- [13] Raissi, M., Perdikaris, P., and Karniadakis, G.E. (2019). *Physics-Informed Neural Networks*. J. Computational Physics, 378, 686–707.
- [14] Milner, R. (2009). *The Space and Motion of Communicating Agents*. Cambridge University Press.
- [15] Lamport, L. (1978). *Time, Clocks, and the Ordering of Events*. CACM, 21(7), 558–565.
- [16] Ollivier, Y. (2009). *Ricci Curvature of Markov Chains on Metric Spaces*. J. Functional Analysis, 256(3), 810–864.
- [17] Nakamoto, S. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. <https://bitcoin.org/bitcoin.pdf>
- [18] Hartig, O. and Thompson, B. (2023). *Foundations of RDF-Star and SPARQL-Star*. W3C Community Report.
- [19] Janus (2023). *A Note on Semiotic Physics*. AI Alignment Forum.