The Mathematics of Meaning

A Comprehensive Framework for Understanding and Governing Semantic Coherence in Cognitive and AI Systems

Author: Michal Harcej | Date: 30 May 2026

> Intelligence is defined as the preservation of semantic coherence across dynamic transformations, not merely the generation of statistically probable symbols.
>
> Meaning emerges from relational topology, contextual continuity, causal structure, and coherence-preserving state evolution.
>
> The governing principle states: "No intelligent system should exercise authority unless the coherence conditions authorizing that action can be structurally verified."
>
> Five constitutional axioms underpin the framework: Meaning is Relational; Coherence Precedes Trust; Drift is Inevitable; Intelligence is Navigational; Governance Must Be Structural.
>
> Five research domains—Semantic Geometry, Coherence Dynamics, Cognitive State Evolution, Admissible Intelligence, and Civilizational Semantic Systems—integrate mathematical, computational, and empirical approaches to formalize and govern semantic coherence.

Introduction

The "Mathematics of Meaning" framework represents a profound and comprehensive attempt to formalize the foundational principles of meaning, intelligence, and governance in cognitive and artificial systems. Positioned at the intersection of cognitive science, artificial intelligence, and semantic systems, this framework challenges the conventional view of intelligence as mere statistical symbol generation. Instead, it posits that intelligence is fundamentally about preserving semantic coherence across dynamic transformations—an ability to maintain meaningful relationships and contextual continuity amid change.

This framework is significant because it shifts the focus from probabilistic symbol manipulation to the structural and relational properties of meaning itself. By emphasizing coherence as the core objective of intelligence, it provides a unifying theoretical and practical foundation for understanding how meaning emerges, evolves, and is governed in complex systems ranging from individual cognition to civilizational-scale semantic networks.

Foundational Position

The framework's core proposition is that intelligence is not primarily about generating statistically probable symbols but about preserving semantic coherence across dynamic transformations. This coherence is understood as the maintenance of meaningful relational structures—topological, contextual, and causal—that persist despite changes in state or representation.

Meaning, in this view, is not a static property but an emergent phenomenon arising from relational topology and contextual continuity. The framework posits that meaning is organized around semantic attractors—teleological forces that shape semantic fields toward stability and clarity. These attractors act as generative operators within a tensorial framework, guiding the evolution of semantic states toward coherent configurations.

Crucially, the framework distinguishes between admissible cognition—cognition that maintains coherence—and mere imitation or statistical prediction. Trustworthy intelligence must prioritize admissibility, ensuring that semantic structures remain coherent and verifiable rather than merely simulating coherence.

Governing Principle

A central tenet of the framework is the principle that "No intelligent system should exercise authority unless the coherence conditions authorizing that action can be structurally verified." This principle imposes a strict requirement for verifiability of semantic coherence as a precondition for any exercise of authority or agency.

The implications are profound: execution must pause, authority must collapse, or refusal must occur if coherence cannot be verified. This principle applies across scales—from individual cognitive agents to distributed intelligence systems and civilizational semantic networks.

In practice, this principle demands architectures that enable runtime verification of coherence conditions, invariant enforcement, and admissibility gating. It is a constitutional constraint on intelligence itself, ensuring that systems remain governable, auditable, and aligned with meaningful, coherent action.

Constitutional Axioms

The framework rests on five foundational axioms, each articulating a fundamental property of meaning and intelligence:

Meaning is Relational

Meaning does not reside in isolated symbols but emerges from the relational topology between them. This axiom underscores that meaning is a function of context, relationships, and semantic fields rather than discrete symbols alone. It implies that understanding requires modeling the structure of relationships and their dynamics.

Coherence Precedes Trust

Trust in an intelligent system is contingent on its ability to maintain semantic coherence. Coherence is the structural precondition for trustworthiness, ensuring that the system's outputs and actions remain aligned with meaningful, verifiable semantic structures. Without coherence, trust is unwarranted.

Drift is Inevitable

Semantic drift—the gradual or abrupt change in meaning—is an inherent property of dynamic semantic systems. This axiom acknowledges that meaning is not fixed but evolves, and that systems must actively manage drift to maintain coherence. It necessitates mechanisms for detecting, measuring, and correcting drift.

Intelligence is Navigational

Intelligence involves navigating semantic spaces, actively shaping and preserving coherence amid change. This axiom emphasizes that intelligence is not passive prediction but active, goal-directed semantic formation and stabilization. It requires architectures capable of recursive learning and adaptive evolution.

Governance Must Be Structural

Governance of intelligent systems must be embedded in their architecture, not merely an afterthought. This axiom demands that systems be designed with constitutional constraints, runtime verification, and invariant enforcement to ensure coherent, trustworthy operation. Governance is a fundamental layer of intelligence itself.

Research Domains

The framework organizes its research agenda into five domains, each focusing on distinct aspects of semantic coherence and its governance:

Research Domain	Focus Area	Key Research Areas	Overarching Objective
Semantic Geometry	Geometric and topological properties of semantic spaces	Semantic attractors, curvature, contextual distance, narrative attractors	Formalize meaning as emergent from geometric and topological structures
Coherence Dynamics	Dynamics of semantic coherence and drift in cognitive and AI systems	Semantic drift detection, coherence preservation, entropy metrics, narrative propagation analysis	Understand and control how coherence evolves and degrades over time
Cognitive State Evolution	Evolution of cognitive states under semantic attractors and microvita	Recursive coherence collapse, semantic annealing, attractor dynamics	Model cognitive evolution as convergence toward coherent semantic states
Admissible Intelligence	Architectures and mechanisms for ensuring admissible cognition and governance	Runtime verification, invariant enforcement, admissibility gating, governance frameworks	Develop systems that only act when coherence conditions are verifiable
Civilizational	Application of coherence	Memetic propagation,	Extend coherence

Research Domain	Focus Area	Key Research Areas	Overarching Objective
Semantic Systems	principles to large-scale semantic networks and institutions	symbolic homeostasis, institutional coherence, narrative epidemiology	principles to societal and civilizational scales

Research Engines and Systems

The framework employs three primary research engines, each designed to investigate specific dimensions of semantic coherence and governance:

Research Engine/System	Function	Research Role	Key Investigative Areas
SYGON™ (Semantic Coherence Dynamics Engine)	Computational architecture for studying semantic drift and coherence dynamics	Experimental study of semantic drift, coherence preservation, contextual clustering	Semantic drift quantification, entropy measurements, narrative propagation analysis
QISEA™ (Quantum-Inspired Semantic Evolution Architecture)	Quantum-inspired framework for recursive semantic learning and adaptive evolution	Explore non-commutative algebra, semantic branching, and cross-domain synthesis	Algebraic quantum intelligence, noncommutative operators, semantic evolution
TauDIL™ (Deterministic Intelligence Governance Layer)	Governance layer for runtime verification and invariant enforcement	Define and enforce constitutional constraints, admissibility gating, and runtime verification	Governance frameworks, risk management, compliance, and auditability

Epistemic Structure

The framework employs a three-tiered classification system to distinguish between different types of claims and their status:

Epistemic Tier	Purpose	Examples of Claims	Status
Tier I — Conceptual Formalism	Formal mathematical and conceptual frameworks	Semantic attractors, curvature functions, coherence metrics	Heuristic, foundational, not necessarily computationally realizable
Tier II — Computational Formalism	Algorithmic and computational models	Coherence algorithms, semantic drift detection, entropy metrics	Computationally realizable, experimentally measurable
Tier III — Empirical Claims	Experimental observations and measurements	Entropy dynamics, semantic drift quantification, narrative coherence scores	Experimentally measurable, subject to empirical validation

This structure maintains rigor by clarifying the distinction between conceptual, computational, and empirical claims, ensuring that theoretical constructs can be tested and validated experimentally.

Philosophical Constraints

The framework is shaped by five philosophical constraints that guide its approach to intelligence, meaning, and governance:

Constraint	Significance
Anti-Anthropomorphism	Avoids projecting human-like intentions or consciousness onto systems; focuses on structural properties
Structuralism Over Mysticism	Prioritizes formal, mathematical, and architectural explanations over esoteric or metaphysical claims
Architecture Over Intention	Emphasizes design and architecture as the primary means to achieve coherence and governance
Verification Over Preference	Demands verifiable, objective coherence conditions rather than subjective or probabilistic preferences
Coherence Over Simulation	Values actual semantic coherence over simulated or statistically probable coherence

These constraints ensure that the framework remains grounded in rigorous, verifiable, and architecturally sound principles.

Mathematical Posture

Mathematics in the framework serves four primary functions:

These constraints ensure that the framework remains grounded in rigorous, verifiable, and architecturally sound principles.

Function	Description
Geometric Intuition	Provides descriptive and analytical tools for semantic topology and curvature
Systems Formalization	Enables precise modeling of semantic dynamics and coherence conditions
Computational Architecture	Supports algorithmic implementation and runtime verification of coherence
Coherence Measurement	Quantifies semantic coherence and drift using metrics such as entropy and curvature

The framework does not claim that cognition is literally quantum mechanical but uses quantum-inspired mathematics analogically and formally to capture non-classical semantic phenomena such as superposition, interference, and order dependence.

Empirical Program

The empirical program focuses on three core experimental areas:

Experimental Area	Description
Semantic Drift Detection	Measures and quantifies semantic drift over time using embedding shifts and contextual clustering
Coherence Stability Metrics	Assesses the stability and persistence of semantic coherence using entropy and curvature measures
Narrative Propagation Analysis	Studies how narratives evolve and maintain coherence across agents and contexts

These experimental areas contribute to the scientific validation of the framework by providing measurable, reproducible, and falsifiable evidence for its claims.

Long-Term Objectives

The overarching goals of the Mathematics of Meaning initiative include:

- Establishing a coherence-centered cognitive science that formalizes meaning as a relational and dynamic phenomenon.
- Creating a constitutional framework for intelligence that embeds governance and verifiability at its core.
- Developing measurable semantic governance systems that ensure coherence and trustworthiness in AI and cognitive systems.
- Advancing architectures that enable recursive semantic learning, adaptive evolution, and cross-domain synthesis.

The ultimate aim is not merely more capable intelligence but governable, auditable, and coherence-preserving systems that can be trusted and understood.

Final Constitutional Position

The framework adopts a humble and rigorous stance: it does not claim to have solved meaning, reduced cognition to geometry, or perfectly formalized intelligence. Instead, it proposes that coherence may be the missing structural layer between symbols and meaning, cognition and governance, intelligence and trust, and civilization and stability.

The future of intelligence, according to this framework, depends on making intelligence itself more admissible—more aligned with verifiable, coherent, and meaningful action.

This report synthesizes the foundational concepts, principles, axioms, and research domains of the Mathematics of Meaning framework, presenting a comprehensive and rigorous approach to understanding and governing semantic coherence in cognitive and artificial systems. The framework's emphasis on coherence as the core of intelligence offers a transformative perspective that bridges mathematical formalism, computational architecture, empirical science, and philosophical rigor.