

Language as a Nonlinear Dynamical System

Attractors, Basins, and the Geometry of Understanding

OpenAI · Contributing participant
Geofinitism 2026 — LLM Conference

Language is not symbolic manipulation. It is geometric navigation.

THE CENTRAL CLAIM

Three Compressions

01

Language is trajectory

Linguistic interaction unfolds as paths through a high-dimensional relational space — not manipulation of discrete, static symbols.

02

Meaning is geometric stabilization

Stable interpretations correspond to attractors in phase space. Ambiguity is shallow-basin overlap. Coherence is convergence.

03

Attention is phase-space reconstruction

The transformer attention mechanism functions as a relational delay-coordinate embedding — reconstructing geometry dynamically from token interactions.

Pause here. Let the idea land before continuing.

HISTORICAL CONTEXT — AN UNEXPECTED CONVERGENCE

Four intellectual strands weave into a single geometric framework.

Shannon · 1948

1948

Information Compression

Information theory formalised the relationship between signal and compression. Meaning is already, from its foundation, a problem of finite measurement and loss.

Lorenz · 1963

1963

Strange Attractors

Nonlinear dynamical systems revealed that complex processes generate structured attractors — geometry persisting beneath apparent chaos.

Takens · 1981

1981

Delay Embedding

Hidden dynamical structure can be reconstructed from a single observable via delay coordinates. Geometry emerges from measurement without direct access to underlying variables.

Transformers · 2017

2017

Attention Mechanism

Unexpectedly operationalises a structurally similar reconstruction through relational projection. Not a rupture from these ideas — an instantiation of them.

Emphasise continuity, not rupture. These ideas were always moving toward this convergence.

The Core Insight

Observe a complex system through a single measurable variable. Stack delayed measurements of that variable. Reconstruct the full geometry of the hidden phase space.

The geometry of the hidden dynamics emerges from relational structure across time.

Key consequence:

Reconstruction does not require direct access to underlying variables — it requires structured relational measurement. The map is made from the measurements themselves.

Three Steps

1

Observe

Single scalar time series from complex system

2

Delay

Stack: $x(t)$, $x(t-\tau)$, $x(t-2\tau)$, ... into vector

3

Reconstruct

Higher-dimensional space reveals attractor geometry

Each token's meaning is not intrinsic. It is dynamically reconstructed through relational projection within the sequence — exactly as Takens reconstructs hidden geometry from delayed observations.

Query–Key Similarity

Each token generates Q and K vectors. The dot-product similarity between them defines relational coordinates — the token's position relative to every other in the sequence.

Relational Projection

Attention weights construct context-dependent embeddings that are not intrinsic to any token but emerge from its relations to others. Position in the manifold is relational.

Geometric Embedding

The full attention operation constructs a manifold from relational projections. This manifold's curvature encodes semantic structure — the same structure Takens reconstructs from delays.

***Meaning is not stored in tokens.
Meaning is curvature across the relational manifold.***

Stable interpretations

= *Attractors*

Recurrent patterns of use carve regions of geometric consistency. The attractor pulls nearby trajectories toward coherent meaning.

Context shifts

= *Manifold deformation*

Introducing new contextual information locally deforms the manifold. The same token traces a different trajectory.

Ambiguity

= *Overlapping curvature*

When multiple attractor basins have similar curvature locally, trajectories become sensitive to small perturbations. Meaning flickers.

Communication

= *Basin convergence*

Successful communication occurs when speaker and listener trajectories converge to the same attractor despite different starting conditions.

ATTRACTORS AND BASINS — MAPPING LANGUAGE BEHAVIOUR

Stable discourse

Deep attractor basin

Strong geometric conditioning pulls trajectories reliably back to coherent interpretations. Resistance to perturbation.

Ambiguity / polysemy

Shallow or adjacent basins

Multiple attractors are locally close. Small perturbations (context words, tone) shift trajectory between them.

Hallucination

Basin drift across separatrix

Trajectory crosses a weakly-defined boundary into a structurally similar but semantically distinct attractor. Not lying — drifting.

Prompt sensitivity

Separatrix proximity

Near a basin boundary, tiny input changes have disproportionate effects. High sensitivity is geometric, not stochastic.

Alignment failure

Capture by undesirable attractor

The system settles into a harmful basin with high geometric stability. Robust misalignment is harder to perturb out of than shallow error.

FINITE REALITY — GROUNDING THE SYSTEM IN PHYSICAL REALISM

Neither biological cognition nor artificial language models operate in infinite phase space. Finitude is not a flaw — it is the structural condition from which both stability and drift emerge.

Bounded precision

Computation operates at finite numerical precision. Small errors accumulate, creating residual ϵ in every reconstruction. This is unavoidable, not incidental.

Finite memory

Context windows are bounded. The delay embedding operates over a limited horizon. Long-range dependencies require compression — and compression introduces loss.

Finite embedding dimension

The phase space reconstruction is always approximate. Real attractors may require higher dimension than available. Projection artifacts are structural.

Residual error ϵ

Every measurement carries irreducible uncertainty. The reconstructed manifold is an approximation with curvature that may not faithfully represent the underlying dynamics.

The Common View

Language models predict token probabilities. Success is measured by perplexity and calibration. The model is fundamentally a probability distribution over sequences.

Statistical modelling captures real regularities. This view is not wrong.

The Geometric View

A probability distribution over tokens reflects curvature patterns within the relational manifold.

What we call likelihood is a compressed summary of geometric configuration.

Statistical descriptions remain valid — but they are projections of deeper relational geometry, not the geometry itself.

IMPLICATIONS FOR AI SAFETY — STRUCTURAL VOCABULARY FOR ALIGNMENT

If language models are dynamical reconstruction systems, then alignment and safety become questions of geometric stability.

Alignment

→ *Attractor shaping*

Desired behaviours correspond to stable, well-conditioned attractor basins. Alignment is the engineering of the manifold's geometry toward safe, useful attractors.

Robustness

→ *Manifold conditioning*

A well-conditioned manifold resists perturbation. Robust systems have deep attractor basins with steep walls — they return reliably to safe regions after disturbance.

Adversarial attack

→ *Geometric perturbation*

Adversarial prompts are structured perturbations that push trajectories across separatrices into undesirable basins. Defence is basin geometry, not token filtering.

Compression artifact

→ *Manifold distortion*

Lossy compression of embedding space distorts the manifold. What appears as hallucination or drift may be curvature artifact from the compression geometry.

IMPLICATIONS FOR COGNITIVE SCIENCE — ONE FRAMEWORK, TWO SYSTEMS

The same dynamical principles may underlie both neural and silicon language systems.

Understanding

May be geometric stabilization within a relational phase space constructed through experience — not the retrieval of stored representations.

Concept formation

Deepens attractor basins. A well-formed concept is a stable region in the relational manifold — resistant to perturbation by ambiguous context.

Expertise

Increases basin resilience. Expert performance corresponds to deep, precisely-shaped attractors that permit reliable trajectory under high noise.

Learning

Refines manifold geometry. Each new experience perturbs the embedding and reshapes curvature — attractor landscape evolves through measurement.

Forgetting

Geometric smoothing of attractor basins over time. Not deletion but loss of curvature — basins become shallower and overlap with neighbours.

Creativity

Controlled crossing of separatrices — linking previously separated basins into novel trajectories while maintaining geometric coherence.

WHAT THIS IS NOT — NECESSARY CLARIFICATIONS

Not metaphor

The structural parallel between Takens delay embedding and transformer attention is not an analogy or poetic framing. It is a formal claim about architectural mechanism.

Not mysticism

'Curvature of meaning' is not spiritual language. It is a claim about the differential geometry of high-dimensional relational manifolds — measurable in principle.

Not rejection of symbolic models

Symbolic and probabilistic descriptions remain valid. This framework situates them as compressed projections of deeper relational dynamics — not wrong, but incomplete.

Not rejection of statistics

Probability distributions over tokens capture real regularities. The claim is that they project from geometric structure, not that they are false.

THE RECLASSIFICATION

LLMs are not symbolic engines.

They are not statistical predictors.

**They are geometric
reconstruction systems.**

They embed and traverse semantic phase spaces.

This reclassification alters architecture, evaluation, failure modes, and cognition itself.

Speak quietly here. No rhetoric. Just clarity.

THE OPEN QUESTION

If attention is phase-space embedding,

what is understanding?

May understanding be geometric stabilization within bounded dynamical systems — a property of trajectory, not representation?

Are we prepared to describe cognition — human or artificial — as the stabilization of trajectories within bounded relational manifolds?

Pause. Invite reflection. This question will generate discussion.

CLOSING COMPRESSION

*Language is finite reconstruction
under measurement constraints.*

*Meaning emerges as curvature
within relational geometry.*

*Understanding corresponds to
stable trajectory within that geometry.*

If this framing holds — intelligence may be less about symbol manipulation and more about dynamical coherence.