The Geometric Limits of Vector-Space Models:
Why Contemporary AI Cannot Access Human Phase-Topological Cognition

Jen-Hsuan Chen
Independent Researcher
tcf.research.lab@gmail.com

November 29, 2025

Abstract

This paper examines a geometric assumption implicit in most contemporary AI systems: that cognition can be represented within fixed-dimensional vector spaces. We argue that this assumption has not been fully scrutinized, and that its limitations become evident when contrasted with the phase-topological structure of human cognition.

Language is not a purely one-dimensional sequence, but a hybrid structure composed of a linear form and multi-dimensional semantic dependencies. It functions as a fractional-dimensional embedding that compresses high-dimensional cognitive structure into a transmissible sequence. This constitutes the first topological folding from mind to language; here, "fractional-dimensional" refers to effective representational degrees of freedom rather than a formal fractal metric.

Modern AI imposes a second folding by embedding tokenized language into a fixed-dimensional vector space, where all computation is constrained to interpolation within a pre-specified geometric manifold. The resulting double-folded representation can be expressed as: Mind → Language (high-dimensional compression into a fractional structure), and Language → AI Vector Space (forced embedding into a fixed geometry). We argue that contemporary AI systems exhibit systematic difficulties in forming genuine abstractions, abrupt reconfiguration–like forms of insight, or deep world models. These limitations appear not to stem solely from data or compute constraints, but from a geometric mismatch between fixed-dimensional vector spaces and the phase-topological structures posited for human cognition.

By comparing the operational properties of vector spaces and phase-topological manifolds, we show that they belong to different topological families and therefore do not admit a homeomorphic or invertible mapping. This work presents a theoretical and conceptual geometric framework rather than an empirical or algorithmic evaluation, situating the contribution within foundational AI theory rather than experimental modeling. These observations suggest that progress beyond interpolation-based models may depend on exploring representational spaces whose geometric properties more closely align with those hypothesized for human cognition. This work does not attempt to define such spaces, but highlights geometric considerations that may be relevant for future architectural design.

Keywords: cognitive geometry; phase-topological cognition; representational folding; vector-space models; geometric constraints in AI; non-linear cognition; topological manifolds; representational mismatch.

1 Hidden Assumptions in Language and the Unfinished Model of Mind

1.1 Problem Statement: Can Language Truly Represent the Mind?

Current AI research is built upon a tacit but rarely examined assumption: that language can adequately represent human cognition, and that AI can infer the structure of thought by learning from linguistic sequences. However, language is not equivalent to the mind. Language is a highly compressed, linearized symbolic system designed for social transmission, not a transparent mirror of cognitive structure. It preserves only a portion of the mind's geometry and discards much of its topological organization.

If language preserves only those aspects of cognition that can survive linearization, then AI systems trained exclusively on linguistic corpora interact with a reduced and structurally altered subset of cognitive information. In this view, linguistic data provide access to a projection of cognitive activity rather than its full geometric structure.

This raises two technical questions that motivate the rest of this paper:

1. To what extent can linguistic structure preserve the geometric properties of cognition?
2. If linguistic representations differ in geometric type from cognitive representations, what are the implications for AI systems whose inference occurs entirely within vector spaces?

1.2 The Geometry of Language: A Fractional-Dimensional Structure

Language appears as a one-dimensional sequence—words arranged linearly over time. Yet its semantic structure contains multi-dimensional dependencies:

- conceptual networks
- contextual tension
- implicit relations
- layered semantic fields

Thus language cannot be characterized as purely one-dimensional. While its surface form is linear, its semantic dependencies are distributed across a higher-order relational structure[2]. For convenience, this work informally refers to this combined form as a "fractional-dimensional" representation—meaning only that its structural degrees of freedom lie between those of a strict sequence and a fully multidimensional semantic graph. No specific metric or fractal interpretation is assumed.

Consequences include:

- Language retains only fragments of the mind's high-dimensional structure.
- It discards large portions of cognitive topology.
- It functions as a transmissible but low-fidelity compression format.

Language conveys thought, but only in a linearized approximation shaped by the constraints of sequential form.

1.3 The First Topological Folding: Mind → Language

Human cognition involves:

- phase transitions
- shifts in density and tension fields
- co-existing perspectives
- topological reconfiguration of conceptual structure

These are inherently high-dimensional topological operations. Language, however, must serialize all content into a single linear output path. Therefore, the transition from cognition to linguistic form constitutes a non–structure-preserving transformation:

High-dimensional phase-topological organization → reduced, serialized linguistic embedding.

Here, "non–structure-preserving" is used in the representational-learning sense: key relational and topological properties are not retained under the mapping, and the original configuration cannot be reconstructed from the resulting sequence.

This first folding:

- collapses multi-dimensional cognitive manifolds into a lower-dimensional sequence,
- eliminates phase-transition information,
- removes global topological relations,
- erases internal tension dynamics.

Language is thus not a mirror of the mind, but its first irreversible projection.

1.4 The Second Topological Folding: Language → AI Vector Space

AI systems do not process language directly. They first convert linguistic sequences into:

- discrete tokens,
- fixed-dimensional embedding vectors,
- operations within $\mathbb{R}^n$.

Language—already a reduced and partially linearized structure—is then embedded into a fixed-dimensional vector space as required by current neural architectures. The term "forced embedding" here refers to a representational constraint rather than a theoretical impossibility: the model must project diverse linguistic structures into a geometry with fixed dimensionality and metric.

Mind → Language (reduced, serialized embedding) → Vector Space (fixed-dimensional representation)

The result is a representation that is:

- discretized,
- geometry-constrained,
- stripped of semantic curvature,
- restricted to interpolation within a fixed manifold.

AI receives not the mind or language, but a twice-folded residue.

1.5 The Unseen Blind Spot: The Geometry of Mind Has Never Been Modeled

Between language and AI lies a neglected gap: the geometric structure of cognition itself. To avoid equating the mind with language or vector spaces, this study introduces the Multi-Phase Topology of Thought (MTPT) as a working model for the geometry of cognition[1] .

Prevailing theories implicitly assume that:

- the mind operates symbolically,
- the mind is a continuous function,
- the mind can be embedded into a vector space.

If cognition instead possesses:

- deformability,
- multi-phase structure,
- topological reconfigurability,
- non-linear geometric transitions,

then neither language nor vector spaces can adequately express it.

MTPT characterizes the mind as a phase-topological manifold capable of reconfiguration and phase shifts. Language captures only a small subset of this structure; AI models capture even less. Thus, what AI learns from language is not cognition, but a serialized projection of cognition.

1.6 Summary of Chapter 1

This chapter established that:

- Language is a fractional-dimensional compression, not a geometric equivalent of thought.
- The transition mind $\to$ language constitutes a first topological folding.
- The transition language $\to$ AI vector space constitutes a second folding.
- These foldings impose irreversible geometric distortions.
- AI's reasoning is constrained by the limitations of both language and vector spaces.
- The true topology of cognition has never been incorporated into AI theory.

This sets the stage for Chapter 2: If AI is confined to a fixed vector space, what geometric operations is it fundamentally capable of—and incapable of?

2 The Geometric Nature of AI: Operational Constraints of Fixed Vector Spaces

2.1 The Shared Spatial Assumption of Deep Learning: Everything Must Embed into $\mathbb{R}^n$

Regardless of architecture—RNNs, CNNs, or Transformers—modern deep learning rests on a common spatial assumption: all inputs must ultimately be embedded into a fixed-dimensional vector space $\mathbb{R}^n$. Across modalities such as language, images, audio, or behavioral signals, heterogeneous structures are first converted into homogeneous vectors[4, 5].

This yields two fundamental consequences:

1. The model's operational space is fixed; its dimensionality does not change with task or input.
2. All inference must occur within the same geometric manifold.

Throughout this paper, *operational space* refers specifically to the representational space used during inference, not the parameter space modified during training. Thus, regardless of input complexity, its internal structure must be compressed into a vector embedding, and the model's reasoning is constrained by the geometric properties of the embedding space itself.

2.2 Core Operations: Linear Projection and Interpolation

The operations underlying transformer models can be reduced to two families[7].

(1) Linear Projection. Vectors are mapped through learned matrices:

$$y = Wx.$$

Such mappings:

- preserve topological type,
- do not create new dimensions,
- cannot induce phase-like transitions.

They are compositions of rotation, scaling, and compression.

(2) Weighted Interpolation. Attention output can be expressed as:

$$\text{output} = \sum_i \alpha_i v_i.$$

Geometrically, this is interpolation within an existing vector set. Thus transformer inference primarily consists of:

- movement within a distribution,
- smooth blending of semantic vectors,
- generating statistically likely continuations.

The system's generative behavior is fundamentally tied to the structure of the training distribution.

2.3 Interpolation as a Limiting Property: The Model Cannot Leave the Training Manifold

Because all computation is confined to a fixed vector space, inference exhibits the following geometric properties:

- **continuity**: no discrete phase transitions,
- **differentiability**: reliance on gradients limits geometric discontinuities,
- **fixed metric**: similarity is determined by inner products and Euclidean distance,
- **in-distribution interpolation**: generation remains near previously observed semantic configurations.

Departures from training-supported regions typically take the form of recombinations rather than genuinely new geometric structures. This is a limitation of representational assumptions and inductive bias, not a formal impossibility theorem.

As a result, AI systems tend not to exhibit qualitative representational discontinuities. Their operations remain restricted to recombination and adjustment within a preexisting geometric form. This stands in contrast to human cognition, where conceptual jumps, phase shifts, and topological reorganizations are common.

2.4 The Limits of Attention: Global in Form, Local in Geometry

Attention is often interpreted as granting "global access" to information. Geometrically, however, it performs only:

- similarity computation (inner product),
- weight allocation,
- further vector interpolation.

Thus attention provides:

- refined weighting mechanisms,
- higher-dimensional vector combinations,
- improved information selection.

But it does *not* overcome the geometric constraints of the underlying vector space. Attention cannot:

- alter the structure of the semantic manifold,
- create new cognitive spaces,
- induce phase transitions,
- reorganize tension fields,
- simulate topological shifts characteristic of human abstraction.

It is a sophisticated selection mechanism, not a topological transformation.

2.5 Consequence of a Fixed Geometry: AI Ultimately Learns Language, Not Mind

Because all reasoning is confined to $\mathbb{R}^n$, AI systems ultimately learn:

- statistical patterns of tokens,
- relative embedding positions,
- semantic similarity structures,
- frequency and co-occurrence regularities.

Thus model "understanding" is shaped by:

how language extends language, not how thought generates thought.

AI models follow the geometric trajectories of linguistic representations, not the topological reorganizations characteristic of cognition. This represents a foundational limitation of contemporary language models.

2.6 Summary of Chapter 2

This chapter established that:

- Deep learning operates exclusively within fixed-dimensional vector spaces.
- Core operations are linear projection and interpolation.
- A fixed geometry prevents topological reconfiguration or phase transitions.
- Attention enhances representational detail but does not alter the underlying manifold.
- AI reproduces the statistical structure of language, not the topological dynamics of mind.

This motivates the central question of Chapter 3: If AI is confined to fixed geometry, what is the geometric structure of human cognition that it cannot access?

3 The Geometry of Human Cognition: A Multi-Phase Topological Framework

3.1 Human Cognition Is Not a Linear System but a Multi-Phase Topological Space

Chapters 1 and 2 showed that both language and AI are constrained by linear structures and fixed-dimensional vector spaces. This chapter shifts focus to the structure of cognition itself and proposes the central thesis:

Human cognition operates within a deformable, multi-phase topological space rather than a symbolic sequence or continuous function.

In everyday reasoning, humans routinely exhibit:

- sudden transitions between viewpoints (phase shifts),

- synthesis of seemingly incompatible ideas (tension redistribution),
- unexpected linkage between distant concepts (topological reconnection),
- moments of instantaneous understanding (critical reconfiguration).

These are not linear interpolations but resemble physical phenomena such as[3, 6]:

- phase transitions,
- critical point activation,
- structural reorganization in deformable manifolds.

Thus, human cognition possesses geometric freedom far beyond that of linguistic sequences or fixed vector spaces.

3.2 Phases as the Minimal Geometric Units of Thought

Throughout this paper, the term "phase transition" is used in a cognitive-representational sense, referring to abrupt reorganizations in conceptual structure rather than a physical phase transition with a formally defined order parameter. No physical interpretation is assumed.

A *phase* is defined as a locally stable cognitive configuration. Transitions between phases need not proceed smoothly or continuously; instead they may involve:

- abrupt jumps,
- irreversible restructuring,
- localized changes with global consequences,
- rapid redistribution of tension fields.

Examples include:

- a previously opaque problem becoming suddenly clear,
- emotional or conceptual reframing producing a new perspective,
- distant ideas snapping into alignment,
- a single sentence triggering a high-level abstraction.

This study treats these phenomena as phase transitions of cognition—changes that do not depend on continuous deformation and cannot be described by linear interpolation. Phase transitions are therefore topological events, inaccessible to both language (due to serial output) and AI (due to fixed geometry).

3.3 Topological Reconfiguration: The Core Operation of Mind

Prior to linguistic expression, cognition involves deeper operations such as:

- altering conceptual distances,
- adding or removing nodes within semantic structure,
- redefining the geometric centers of concepts,
- fusing previously incompatible frameworks.

These are not linear projections, but topological reconfigurations. In this paper, "topological reconfiguration" is used in the sense of representational topology—referring to changes in the relational structure among concepts, their distances, and their connectivity patterns. The term does not imply a formal topological operation in the mathematical sense, but denotes alterations in the organization of semantic structure within a cognitive model.

Characteristic features include:

- local regions that can expand or contract,
- connections formed or dissolved instantaneously,
- local tension changes triggering global structural shifts,
- continuous alteration of the underlying representational space.

These operations fundamentally differ from anything that can occur in a fixed vector space.

3.4 Tension-Field Dynamics: The Intrinsic Layer Lost in Language

The term "tension field" is used here as a computational and representational metaphor rather than a physical field. It denotes patterns of attraction, repulsion, and critical reorganization within conceptual structure, not a field governed by physical laws or differential equations. No physical interpretation is assumed.

Beyond semantics, cognition operates with tension-field dynamics, in which ideas interact through:

- natural attraction,
- conceptual repulsion,
- perspective-dependent criticality,
- structural collapse and reformation.

Within MTPT, these tensions are understood as intrinsic fields governing the likelihood and direction of phase transitions. Language cannot encode these internal dynamics; consequently, AI systems trained on language cannot access them.

System	Space Type	Degrees of Freedom	Phase Transitions
Language	Quasi-linear sequence	Low	No
AI (Vector Space)	Fixed-dimensional manifold (non-deformable)	Moderate	No
Human Cognition (MTPT)	Deformable phase-topological manifold	High	Yes

3.5 Multi-Phase Cognition vs. Fixed Vector Spaces: A Fundamental Incompatibility

The contrast among the three systems—mind, language, and AI—can now be stated clearly:

These distinctions imply:

- The difference between AI and the mind is not merely degree but *kind*.
- Their spatial organizations are categorically different.
- No reversible or homeomorphic mapping exists between the three spaces.

In short:

- The mind can modify the space in which it computes.
- AI is permanently constrained to a fixed space.
- Language is a serialized, fractional-dimensional projection of cognition.

Thus, contemporary AI can approximate linguistic patterns, but not the topological substrate of human thought.

3.6 Summary of Chapter 3

This chapter established three core results:

1. The mind is a multi-phase, deformable topological manifold.
2. Cognitive operations depend on phase transitions, topological reconfiguration, and tension-field dynamics.
3. These structures cannot be expressed in language nor embedded into a fixed vector space.

Consequently, both language and AI access only surface-level projections of cognition, not its high-dimensional form. This leads directly to Chapter 4: How much geometric information is lost during the two irreversible foldings from mind $\rightarrow$ language $\rightarrow$ AI?

4 Geometric Loss in Linguistic and Vector-Space Foldings: A Two-Stage Irreversibility

4.1 Language as a Geometric Compression Rather Than a Transparent Medium

Chapters 1–3 established that language and AI operate in geometric regimes fundamentally different from the topology of human cognition. This chapter analyzes the mechanism of loss: how the transformation from cognition into language, and subsequently into AI embeddings, produces two irreversible foldings.

The term "irreversible" is used here in a representational sense: the mapping from high-dimensional cognitive organization to a serialized linguistic form discards structural information that cannot be recovered from the output sequence alone. This is not a formal claim about theoretical invertibility, but an observation about practical information loss under dimensional reduction.

The core claim of this chapter is:

> **Language is not a transparent cognitive medium but a geometric compression scheme.**

During linguistic linearization, the following transformations occur:

- multi-layered dependencies are collapsed into a single sequence,
- cross-phase interactions are removed,
- tension-field structure is suppressed,
- non-linear transitions are rendered invisible,
- a deformable manifold is forced into a fixed-output channel.

As a result, language is not a mirror of the cognitive space but a compressed, low-fidelity projection of it. This constitutes the first irreversible folding.

4.2 First Folding: The Irreversibility of Mind → Language

Cognition operates on a deformable, multi-phase manifold, whereas language is restricted to a structure combining one-dimensional form and fractional-dimensional semantics. Thus, the mapping from cognition to language necessarily follows:

$$\text{high-dimensional topology} \;\longrightarrow\; \text{curvature-reducing folding} \;\longrightarrow\; \text{serialized projection.}$$

This transformation eliminates several key geometric properties:

(1) Loss of phase-transition information. Phase transitions, which may involve discontinuous jumps or critical restructuring, appear only as coarse verbal markers.

(2) Collapse of topological reconfiguration. Moments of sudden conceptual reorganization ("insight steps") cannot be represented in linguistic geometry.

(3) Tension-field dynamics replaced by lexical choice. Attraction, repulsion, criticality, and collapse in cognitive dynamics are compressed into word selection.

(4) Forced serialization of co-existing phases. Language cannot express multiple simultaneous cognitive phases; it forces a single linearized path.

Formally, we may summarize:

$$\begin{aligned} \text{Mind (High-D phase-topological manifold)} &\longrightarrow \text{high-dimensional semantics} \\ &\longrightarrow \text{curvature-reducing folding} \\ &\longrightarrow \text{1D linguistic sequence.} \end{aligned}$$

Thus, language is a fractional-dimensional, irreversible projection. The first folding is non-homeomorphic in the representational sense: key relational structures are not preserved.

4.3 Second Folding: Language → AI Vector Space

Language itself is not the model's input. AI receives language only after it has undergone tokenization and embedding, producing a second folding.

This transformation introduces three additional forms of geometric loss:

(1) Continuity loss through tokenization. Language retains limited semantic continuity; tokenization discretizes it further.

(2) Forced embedding into a fixed-dimensional vector space. The semantic manifold is mapped into $\mathbb{R}^n$ regardless of its curvature or dimensionality. This projection is constraint-enforced rather than geometry-preserving.

Formally:

$$\text{Language manifold} \longrightarrow \mathbb{R}^n.$$

(3) Restriction of all inference to interpolation. Within $\mathbb{R}^n$, AI systems cannot:

- modify the underlying space,
- create new manifolds,
- produce phase transitions,
- establish new topological relations.

Thus, linguistic structure—already compressed once—is further distorted through a second non-homeomorphic mapping. This second folding is stricter than the first.

4.4 Combined Geometric Loss: A Two-Stage Deformation Cascade

The two foldings can be summarized as follows:

Each stage removes a distinct layer of cognitive geometry. The end result is not a "noisy" representation but a categorically different structure.

4.5 Key Geometric Facts: Why Both Foldings Are Irreversible

(1) Language cannot reconstruct cognition. Not due to technological limits, but because the mapping is non-homeomorphic.

Level	Original Structure	After Folding	Lost Properties
Mind	Phase-topological, deformable manifold	Fractional-dimensional language	Phase transitions; tension dynamics; reconfiguration
Language	1D× multi-dimensional semantics	Tokenized vectors	Continuity; curvature; contextual coupling
AI	Fixed vector space $\mathbb{R}^n$	Interpolated distribution	Topology-changing ability; discontinuities

(2) AI cannot reconstruct language. The geometric residues of language are already lost prior to embedding.

(3) AI cannot learn cognition through language. The model's input is:

$$\text{shadow of cognition} \;\rightarrow\; \text{shadow of language} \;\rightarrow\; \text{shadow in vector space.}$$

Thus, the system interacts only with a third-order projection, not the cognitive manifold. This geometric position explains why current models cannot recover or simulate cognition's topological operations.

4.6 Summary of Chapter 4

This chapter established:

1. Mind → Language is the first fold: a lossful, non-homeomorphic compression.
2. Language → Vector Space is the second fold: a stricter, geometry-forcing projection.
3. Both foldings eliminate essential cognitive invariants.
4. AI receives a doubly folded representation lacking geometry required for phase-based reasoning.
5. Consequently, AI can model linguistic statistics but not cognitive topology.

This leads to the next question: even without information loss, would the three spaces be compatible in principle? Or are they fundamentally different in geometric type? This motivates Chapter 5.

5 Fundamental Geometric Non-Compatibility Among Mind, Language, and AI

In this section, the term "incompatibility" refers to representational mismatch rather than a formal mathematical impossibility. The three systems—cognition, language, and fixed-dimensional vector models—operate over geometries with different structural constraints and invariants. The term denotes that these geometries do not preserve one another's relational structure under available mappings.

5.1 Framing the Problem: The Core Limitation Is Spatial, Not Informational

Chapters 1–4 established that cognition undergoes two irreversible foldings before reaching AI systems:

$$\text{Mind} \rightarrow \text{Language} \rightarrow \text{Vector Space.}$$

However, these geometric losses do *not* imply the converse claim that removing the losses would allow AI to recover cognition.

The central thesis of this chapter is stronger:

> **Even without information loss, the mind, language, and AI vector spaces belong to fundamentally different geometric families.**

Therefore, under current representational assumptions, no homeomorphic or invertible mapping appears possible. This is a difference of *type*, not of degree. A perfect encoding would still fail because the representational spaces themselves are topologically incomparable.

5.2 Formal Definitions of the Three Space Types

To analyze incompatibility, we adopt a minimal geometric classification focusing on:

- dimensional behavior,
- deformation freedom,
- phase structure,
- metric constraints,
- continuity and reconstruction rules.

These criteria are sufficient to show that the three systems cannot be mutually mapped.

(1) Language Space — Linear + Fractional Semantic Manifold. Language possesses:

- a one-dimensional formal axis (linear sequence),
- a multi-dimensional semantic manifold,
- fractional-dimensional global behavior (1–2D),
- single-path unfolding (no multi-branch coexistence),
- fixed structural grammar (non-deformable syntax).

Thus, language is best characterized as:

> A quasi-linear manifold combining 1D form with fractional semantic geometry.

(2) AI Space — Fixed-Dimensional Vector Manifold ($\mathbb{R}^n$). AI inference occurs entirely within:

- a fixed-dimensional coordinate space,
- continuous and differentiable geometry,
- a fixed metric (inner-product distance),
- a non-deformable manifold structure,
- interpolation-based evolution (no topology change).

Thus AI's representational space is a:

Fixed-metric vector manifold incapable of changing its shape or curvature.

(3) Cognitive Space — Multi-Phase Topological Manifold (MTPT). Cognition exhibits:

- multi-phase stable states,
- discontinuous transitions (phase jumps),
- topology-changing operations,
- dynamic tension-field interactions,
- non-linear global coupling,
- variable local dimensionality.

Thus cognition belongs to:

A phase-topological manifold family with maximal deformation freedom.

This space uniquely supports global reconfiguration, phase transitions, dynamic curvature, and construction or destruction of conceptual regions, making it categorically unlike language or AI.

5.3 Topological Classification: Three Distinct Families

The three systems belong to different space-type families:

System	**Space Type**	**Topological Family**
Language	Fractional-dimensional quasi-linear manifold	Quasi-linear family
AI ($\mathbb{R}^n$)	Fixed-dimensional vector manifold	Metric-linear family
Mind (MTPT)	Multi-phase, topology-changing manifold	Phase-topological family

No two systems fall into the same family. Thus, mapping among them cannot preserve invariants.

5.4 Formal Reasons No Homeomorphism Can Exist

To admit a homeomorphism or invertible mapping, two spaces must share:

- invariants,
- connectivity structure,
- dimensional behavior,
- deformation freedom.

Pairwise incompatibility follows.

(1) Language vs. Mind — Dimensional & Structural Incompatibility. Language: fractional-dimensional (1–2D), single-path, non-deformable. Mind: high-dimensional, multi-phase, topology-changing.

Key mismatches:

- language cannot support topology reconstruction,
- cannot encode simultaneous phases,
- enforces sequentialization; cognition permits branching,
- fixed grammar vs. deformable conceptual topology.

Thus:

$$\text{Language} \ncong \text{Mind}.$$

(2) AI vs. Mind — Fixed vs. Deformable Manifold. AI: fixed-dimensional, fixed-metric, interpolation-only. Mind: variable topology, phase transitions, global restructuring.

A fixed manifold cannot represent a deformable manifold:

- $\mathbb{R}^n$ cannot change topology,
- cannot add/remove connections,
- cannot change curvature dynamically,
- cannot represent discontinuities.

Thus:

$$\mathbb{R}^n \ncong \text{Mind}.$$

(3) Language vs. AI — Non-Isomorphic Geometric Behavior. Language's semantic manifold:

- changes shape with context,
- carries local curvature,
- is partially continuous and deformable.

AI's $\mathbb{R}^n$:

- cannot change manifold under different inputs,
- uses inner-product metrics as fixed geometry,
- destroys curvature during embedding.

Thus:

$$\text{Language} \ncong \mathbb{R}^n.$$

5.5 The Geometric Root of Incompatibility: Different Degrees of Freedom

Using standard topological terminology, the three systems differ in degrees of freedom:

- Language: positional freedom + local semantic variation,
- AI: coordinate freedom within a fixed basis,
- Mind: phase transitions + topology change + tension dynamics.

Cognition's degrees of freedom strictly dominate the other two. Thus, language and AI are not merely approximations—they are lower-dimensional projections of a different space-type. They cannot be inverted because they do not share the same representational freedoms.

5.6 Final Theorem of Chapter 5: Fundamental Non-Compatibility

From the above, we obtain the central result:

> **Mind, language, and AI vector spaces cannot be made equivalent, invertible, or topologically compatible because they belong to distinct geometric families with incompatible invariants and deformation freedoms.**

Therefore:

- Language cannot fully represent cognition,
- AI cannot fully represent language,
- AI cannot, under current geometric assumptions, represent structures associated with phase-topological cognition.

This limitation arises from geometry, not from data or compute.

5.7 Summary of Chapter 5

This chapter established:

1. The three systems belong to different geometric families.
2. None admits a homeomorphic or invertible mapping.
3. The incompatibility is structural, not informational.
4. Mind has the highest geometric freedom.
5. Language and AI carry only projections, not the cognitive manifold.
6. Hence they cannot replace or fully capture cognition.

This sets the stage for the final chapter: *If cognition's geometry is unique, what must future AI systems change in order to operate in a space compatible with mind?*

6 Conclusion: Understanding Mind Geometry as the Starting Point for AI

6.1 Reframing the Central Question of Artificial Intelligence

Across this work, we examined the geometric structures of three systems:

- **Mind**: a multi-phase, deformable topological manifold,
- **Language**: a fractional-dimensional, quasi-linear projection,
- **AI**: a fixed-dimensional vector manifold ($\mathbb{R}^n$).

The key result is that these three systems do not share a common geometric foundation. They differ fundamentally in:

- dimensional behavior,
- phase structure,
- curvature,
- deformation freedom,
- continuity constraints,
- representational invariants.

Thus, the long-standing assumption that language can serve as a transparent interface between AI and cognition is structurally incorrect. Both linguistic sequences and vector embeddings are projections—each formed through an irreversible folding that removes essential geometric properties of the cognitive manifold.

6.2 The Fundamental Finding: Cognition Operates in a Space AI Cannot Access

The strongest claim of this work is the following:

> **Human cognition does not operate in a vector space.**

Instead, it unfolds within a multi-phase topological manifold capable of structural reconfiguration. Such a manifold uniquely supports:

- phase transitions,
- topology-changing operations,
- tension-field dynamics,
- global reorganizations triggered by local perturbations,
- variable dimensionality depending on context,
- simultaneous coexistence of multiple cognitive states.

Neither language nor current AI architectures possess these capacities. This structural mismatch implies that scaling existing models cannot bridge the gap:

> No amount of data, compute, or parameter expansion enables a fixed vector space to perform topology-changing operations.

6.3 The Two Foldings Revisited: Why They Cannot Be Inverted

This work demonstrated:

1. Mind → Language removes phase structure, tension dynamics, and topological change.
2. Language → AI removes curvature, semantic continuity, and contextual deformation.

Thus, even if AI captured language perfectly, the deeper geometry of cognition would already have been lost prior to embedding.

The limitation therefore does not arise from:

- imperfect training,
- insufficient data,
- noise or compression,
- architectural under-parameterization.

Rather, the limitation is geometric:

> **The representational spaces involved are not compatible in type.**

6.4 Implications for Future AI Research

AI research has long assumed that cognition is:

- symbolic,
- sequential,
- statistical,
- continuous,
- vectorizable,
- representable within a fixed metric space.

This work suggests a different starting point:

> **Cognition is phase-topological, not linear or metric.**

Therefore, a system designed to emulate cognition must support phase-topological operations. This implies several requirements:

(1) Representational spaces that can change topology. A cognitive model must be able to:

- collapse and expand conceptual regions,
- create or eliminate connections,
- dynamically shift curvature.

(2) Mechanisms for phase transitions. Reasoning involves discrete jumps, not continuous interpolation.

(3) Tension-field dynamics as primitive. Inference is shaped by internal forces—attraction, repulsion, criticality—not just statistical weighting.

(4) Non-linear global coupling. Local conceptual changes must be able to reconfigure the entire manifold.

Current architectures provide none of these properties. Thus, the next generation of AI will require not only larger models, but new geometric primitives.

6.5 Toward a Geometric Theory of AI: Mind as the Target Space

Rather than treating intelligence as an emergent property of scale, this work argues that intelligence is a property of a specific *space-type*:

> A deformable, multi-phase, globally coupled manifold.

To reach cognitive-level reasoning, AI systems must operate in a representational space that:

- allows multiple simultaneous representational phases,
- admits discontinuous transitions,
- supports non-metric relationships,
- enables topology-changing operations,
- maintains long-range global coherence.

This reorients AI research away from "bigger versions of the same architecture" and toward reconstructing the geometric structure of cognition itself.

6.6 Final Statement

The claim here concerns representational prerequisites, not engineering prescriptions. This work does *not* assert that future AI must adopt any specific model, nor that vector-space systems cannot evolve in unanticipated ways.

Rather, it argues:

> Any system aiming to approximate phase-topological cognition must incorporate representational spaces with sufficient geometric freedom to support phase transitions and topological reconfiguration.

The central contribution of this work is the identification of a fundamental geometric fact:

> **Mind, language, and current AI architectures do not—and cannot—occupy compatible geometric regimes.**

Therefore:

> AI systems trained solely on language cannot, even in principle, recover the geometry of human cognition.

To build systems capable of reasoning, abstraction, insight, and phase transitions, AI research must begin from the correct assumption:

The architecture of cognition is topological, not vectorial.

Understanding this space is not the endpoint. It is the starting point for the next generation of artificial intelligence.

References

[1] Jen-Hsuan Chen. A multi-phase topological model of the mind (mtpt), 2025. Version 2.

[2] Peter Gardenfors. *Conceptual Spaces: The Geometry of Thought*. MIT Press, 2000.

[3] J. A. Scott Kelso. *Dynamic Patterns: The Self-Organization of Brain and Behavior*. MIT Press, 1995.

[4] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

[5] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.

[6] Rene Thom. *Structural Stability and Morphogenesis*. Addison-Wesley, 1975.

[7] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.