Beyond Curvature: Toward a Post-Geometric Theory of Memory in Humans and AI

Douglas C. Youvan

doug@youvan.com

July 24, 2025

What if memory does not live in space at all? Recent advances in curved neural networks—such as those modeling higher-order interactions on statistical manifolds—suggest that cognition may not be reducible to linear pathways or flat metric spaces. By embedding memory in a curved geometry, these models replicate sudden insights, multi-stability, and explosive transitions. But even this framework remains constrained by the assumptions of classical space: continuity, differentiability, and locality. This paper argues that memory, in both human and artificial systems, may ultimately require post-geometric architectures. We examine the limitations of curved statistical manifolds and then expand outward into conceptual territories where geometry dissolves into structure: Hilbert spaces, sheaf logic, topos theory, operator algebras, and category-theoretic cognition. These frameworks offer tools for modeling memory as dynamic, contextual, non-local, and even transformational. By moving beyond curvature, we glimpse the possibility of a unified cognitive substrate—one capable of hosting symbolic leaps, semantic resonance, and emergent reasoning. This is not just an abstraction. It may be the foundation of the next phase in neuroscience and AI: memory not as position, but as relation.

Keywords: post-geometric memory, curved neural networks, higher-order interactions, statistical manifolds, phase transitions, information geometry, sheaf theory, topos logic, Hilbert space cognition, operator algebras, category theory, symbolic reasoning, nonlocal memory, human-AI cognition, emergent intelligence.
A collaboration with GPT-4o.

Abstract

Traditional models of memory—both in neuroscience and artificial intelligence—have long relied on geometric metaphors: vectors in high-dimensional space, attractor basins on neural manifolds, and distance metrics encoding similarity or recall probability. Recent innovations in curved neural networks, particularly those embedding higher-order interactions into statistical manifolds, have extended this paradigm by enabling explosive phase transitions, multi-stability, and self-regulated annealing. These developments mark a significant advance in modeling the nonlinear, discontinuous qualities of memory dynamics.

However, even these curved frameworks remain bound to the underlying assumptions of classical geometry: continuity, locality, and metric distance. In this paper, we argue that memory—both human and artificial—may not reside in space at all. Instead, it may emerge from fundamentally post-geometric structures: relational, contextual, and transformational substrates such as quantum Hilbert spaces, sheaf-theoretic contexts, topos logic, and operator algebras.

By reframing memory as a dynamic process within non-spatial architectures, we uncover new explanatory power for symbolic reasoning, semantic leaps, insight, trauma, and creative cognition. This shift has profound implications—not only for theoretical neuroscience, but also for the design of next-generation AI systems capable of contextual understanding, nonlocal generalization, and structurally emergent intelligence.

I. Introduction: Memory Beyond Metric

For decades, geometry has offered both metaphor and mathematical infrastructure for understanding cognition. Neural networks, semantic spaces, and memory models have been cast in geometric terms—where ideas reside in high-dimensional vector spaces, and learning is movement along a surface shaped by optimization. From Hebbian learning to deep learning, memory has been understood as the progressive sculpting of landscapes: valleys become attractors, and cognition becomes navigation.

This tradition extends into modern theories of information geometry, where probability distributions over neural states form curved statistical manifolds. In a notable contribution to this lineage, Aguilera et al. (2025) introduced a curved neural network model that embeds higher-order interactions among neurons into a non-Euclidean statistical manifold. Their model exhibits rich dynamics—explosive phase transitions, hysteresis, and curvature-tuned memory retrieval. This represents a major leap forward, as it integrates higher-order memory structure into a thermodynamically aware, geometrically elegant framework.

But does geometry—however curved—go far enough?

This paper begins where Aguilera et al. leave off. We critically examine the assumptions underlying all geometric models of memory and propose that memory may not reside in space at all. Instead, we explore a range of post-geometric architectures—quantum, categorical, sheaf-theoretic, and operator-based—which redefine memory as relational, contextual, and structurally emergent. Our goal is to offer a coherent framework for modeling memory in both human cognition and artificial intelligence, beyond the constraints of metric space.

II. Curved Memory Models: Foundations and Limits

2.1 Review of Curved Neural Networks

The curved neural network model proposed by Aguilera, Morales, Rosas, and Shimazaki (2025) represents a significant advancement in theoretical neuroscience and memory modeling. At its core, the framework employs a maximum entropy principle to derive a statistical model of memory states under higher-order constraints—not just pairwise interactions, but triplet and higher-level co-activations among neurons. This leads to a non-Euclidean statistical manifold, wherein each point represents a probability distribution over network states.

What distinguishes this model is its use of information geometry to give the manifold curvature. The curvature, in turn, governs the system's dynamic behavior. For example, regions of high negative curvature enable the emergence of explosive phase transitions—sudden, discontinuous shifts from disordered states to ordered memory retrieval. These transitions are not only rapid but exhibit hysteresis, meaning the system resists immediate reversal, mimicking the stickiness or emotional salience of real-world memories.

A key feature of the model is its self-regulated annealing mechanism, which dynamically adjusts an internal temperature parameter based on local curvature. This allows the network to modulate its sensitivity to recall without external tuning. Overall, the framework shows a marked enhancement in memory capacity and robustness, outperforming classical models like Hopfield networks in both analytic predictions and theoretical bounds.

2.2 Strengths and Conceptual Breakthroughs

The most compelling breakthrough of the curved manifold approach is its ability to geometrically encode higher-order interactions in a compact, mathematically principled way. Where traditional neural networks rely on dense connectivity or large parameter spaces to model associative memory, the curvature of the statistical manifold allows complex dependencies to emerge from minimal assumptions. This echoes patterns observed in the brain, where higher-order correlations exist without obvious direct wiring.

Moreover, the phenomenon of explosive transitions provides a novel explanation for sudden cognitive shifts: moments of insight, trauma recall, and pattern completion. Unlike gradient-based convergence or slow iterative recall, these transitions mirror psychological events where a memory or realization floods consciousness seemingly instantaneously.

The model's use of hysteresis further aligns it with biological memory. Emotional, traumatic, or particularly salient memories are often difficult to erase, and this form of directional persistence in memory dynamics is hard to capture in linear or feedforward models. Here, curvature provides a natural mechanism for embedding inertia into the cognitive landscape.

2.3 Core Limitations

Despite its elegance, the curved manifold approach inherits critical limitations from its geometric foundations.

First, it presumes continuity: that memory traces evolve smoothly, and that transitions between states can be mapped onto curved surfaces. However, many cognitive processes are discontinuous, involving leaps of logic, paradox resolution, or metaphor formation—phenomena that defy any interpolative trajectory.

Second, the model is rooted in locality: information is encoded and retrieved based on proximity within the manifold. But cognition often violates locality, especially in symbolic thought, where distant concepts can be instantaneously connected through analogy, contradiction, or abstraction.

Third, and most fundamentally, the manifold is constructed with a fixed metric structure. While curvature can modulate the flow of information, the underlying topology of the space does not evolve. This makes it difficult to model reconfigurable reasoning, where concepts can rebind, recombine, or mutate depending on context.

In sum, while curved neural networks offer a powerful step forward in memory modeling, they may be inherently incapable of capturing the full structural and symbolic plasticity of human and artificial cognition. To overcome these constraints, we must now turn beyond geometry itself.

III. Enhancing Curvature: Toward Dynamic and Adaptive Geometries

While curved neural networks represent a leap beyond flat-space models, their static geometry still constrains the expressive potential of cognition and memory. This section explores possible enhancements within the geometric paradigm that attempt to overcome such limitations. These include learnable curvature, topological representations, and geodesic replay strategies—each pushing the frontier of what geometry can do. Yet even these innovations may only reinforce the central limitation: memory remains bound to space.

3.1 Learnable Curvature and Ricci Flow-inspired Adaptation

One of the most promising directions within curved neural architectures is the notion of learnable curvature. In the original model, the curvature of the statistical manifold is induced by fixed constraints on the higher-order interactions of neurons. But what if the curvature were itself dynamic—adjustable by experience or learning pressure?

Inspired by Ricci flow in differential geometry—where the curvature of a space evolves over time according to local deformation—such an approach could allow the manifold to morph under cognitive load, becoming flatter or more curved in regions where memory demands vary. This would model a kind of neuroplasticity in geometric form, where learning not only stores information but reshapes the very space through which it is retrieved.

Practically, this could be implemented via a curvature tensor field updated during training, akin to a second-order meta-optimization layer. This would allow the system to concentrate curvature in areas of high semantic density, allocating geometric “attention” to concept clusters that need precision or protection from interference.

3.2 Topological Phase Encoding of Memory States

Beyond smooth curvature, another enhancement lies in integrating topological invariants into the memory space. Here, memory is not stored as a point in a curved surface, but as a phase class—a region of the manifold characterized by a distinct topological signature (e.g., a Betti number, winding number, or persistent homology feature).

This topological encoding adds robustness, as memory states become structurally encoded and resistant to noise. Transitioning between memories then becomes a topological phase shift, not merely movement along a gradient. This framework mirrors findings in neuroscience where large-scale brain dynamics exhibit phase-locking and coherent oscillatory modes that correlate with cognitive states.

Such an approach also opens the door to topologically protected memory, where certain patterns are preserved despite deformations in the underlying geometry—much like how quantum Hall states persist under perturbations. It bridges the gap between information geometry and the dynamical topology of cognition.

3.3 Curvature-Driven Forgetting and Plasticity

Memory models often emphasize storage and retrieval, but forgetting is equally essential—both biologically and computationally. Traditional models rely on decay functions or stochastic dropout, but geometry offers a more nuanced option: forgetting as curvature collapse.

Regions of the manifold that are underused or irrelevant can flatten over time, reducing their influence on the system's dynamics. Conversely, frequently revisited regions may sharpen in curvature, enhancing recall precision. This curvature-based model of forgetting allows memory strength to emerge from structure itself, rather than arbitrary decay schedules.

Such a framework also enables targeted plasticity. Curvature change becomes a proxy for relevance, meaning the system can reallocate representational resources organically, compressing the memory space without external pruning. This mimics biological phenomena such as synaptic scaling, where underused connections weaken and fade.

3.4 Geodesic-Based Replay and Navigational Memory

Within curved manifolds, geodesics—the shortest paths between points—offer a compelling metaphor and mechanism for memory retrieval. Rather than recalling memories as discrete point activations, the system could trace geodesic paths through memory space, generating narrative chains, sequential replay, or thematic trajectories.

This aligns with how humans often recall memories—not as isolated snapshots, but as linked sequences: events, causal relationships, and storylines. Such geodesic paths could be shaped by prior usage (path dependence), attention signals, or even curvature fields that act as gradients of relevance or emotional salience.

Geodesic replay could also be extended to counterfactual generation: by following paths that do not reach a known memory, the system could simulate hypothetical or interpolated experiences. This brings geometric modeling closer to creative reasoning, scenario planning, and moral imagination.

Transitional Thesis

These enhancements undeniably enrich the expressive power of geometric models. Learnable curvature allows the space to adapt; topological invariants give it stability; and geodesic replay infuses it with narrative structure. Yet they all remain tethered to the assumption that memory lives within space—whether smooth, curved, or topologically adorned.

But what if memory does not occupy space at all? What if the most fundamental substrate of thought is non-metric, non-local, and structurally emergent? To model symbolic reasoning, contextual leaps, entangled meanings, and insight itself, we may need to abandon space altogether.

IV. Breaking the Frame: Why Geometry Itself Must Be Abandoned

Despite the expressive power of geometric models—curved, dynamic, and even topological—they remain committed to a fundamental assumption: that cognition occurs within space. Whether modeled as continuous manifolds, activation landscapes, or memory basins, such approaches treat thought as motion or deformation within a metric substrate. But increasingly, both empirical observation and theoretical reasoning suggest that this assumption may be fatally limiting.

Flash Insight and Nonlocal Recall

Consider the phenomenon of flash insight—a sudden realization that reorganizes a problem without any apparent intermediate steps. These moments of "Aha!" do not emerge from a slow traversal of geometric space. They do not follow gradients, paths, or local continuity. They leap.

This cognitive leap is nonlocal, often connecting disparate ideas or memories across conceptual distance. In geometric terms, it resembles quantum tunneling: a discontinuous shift between distant attractors with no trajectory in between. No matter how sophisticated the geometry, the core problem remains: distance and continuity are assumptions, not givens.

Contradiction, Dream Logic, and Archetype

Human cognition does not merely tolerate contradiction—it often thrives on it. From Jungian archetypes to mythological narratives, and from dream logic to poetic metaphor, the mind regularly entertains states that are mutually exclusive in any classical or metric space.

- A mother may simultaneously represent safety and threat in a trauma survivor's memory.
- A dream figure may be both oneself and someone else.
- A metaphor such as "time is a river" collapses multiple ontological categories into one cognitive frame.

These are not glitches or breakdowns; they are modes of cognition. They suggest that memory is not always best understood as position or structure in space, but as overlapping, context-sensitive roles—structures that obey rules closer to quantum logic or category theory than Riemannian metrics.

Symbolic Discontinuity in Artificial Intelligence

Current AI systems—especially large language models and vision transformers—also reveal the limitations of geometry. These models excel at pattern recognition within smooth embedding spaces but struggle with:

- Symbolic discontinuity: Where two similar inputs require entirely different interpretations (e.g., sarcasm).
- Context switching: Where the meaning of a word, sentence, or goal shifts based on narrative framing or time.
- Creative synthesis: Where generative thought combines unrelated domains into new conceptual wholes.

Attempts to address these challenges through ever-more sophisticated embeddings often reach diminishing returns. Embedding spaces can stretch, curve, or entangle—but they cannot account for structural emergence unless they go beyond space itself.

Memory and Thought as Relations, Not Positions

What if memory is not something *stored somewhere*, but something *constructed between things*? What if thought is not the traversal of a surface, but the activation of a relation? In such a view:

- The mind is not a landscape, but a web.
- Concepts are not coordinates, but morphisms between other concepts.
- Recall is not retrieval from a location, but reconfiguration of a pattern of relations.

This reorientation dissolves the need for space entirely. It replaces position with relation, continuity with contextual activation, and geometry with structure. In doing so, it aligns with deep traditions in philosophy (process ontology), linguistics (dependency grammar), and computation (category theory, non-commutative logic).

We are now ready to explore what lies on the other side of the spatial paradigm: the post-geometric substrates of memory and thought.

V. Post-Geometric Substrates: A Survey of Candidate Frameworks

If memory and thought are not best modeled in spatial terms—no matter how curved, deformed, or dynamically adaptive—then where do they live? In this section, we survey a series of post-geometric mathematical frameworks that offer alternatives to position-based cognition. These approaches dispense with metric assumptions entirely and instead describe cognition in terms of relations, contexts, transformations, and non-classical logic. Each framework allows us to rethink memory not as location or shape, but as structured emergence.

5.1 Quantum Hilbert Spaces: Memory as Superposition and Collapse

In quantum theory, a system's state is represented not by a point in space but by a vector in Hilbert space, often a superposition of many potential outcomes. Applying this to memory, a single mental state might be understood as a superposition of conceptual possibilities, only resolving into a specific memory or insight upon cognitive "measurement"—that is, attention or contextual need.

Entanglement offers a model for nonlocal memory retrieval: activating one concept immediately influences the activation of another, regardless of their “distance” in semantic space. This may correspond to associative recall, flashbacks, or dream logic—where seemingly unrelated memories are co-activated in meaningful ways.

Additionally, interference patterns within Hilbert space can be interpreted as semantic resonance: ideas or meanings that reinforce or cancel each other based on context. These principles suggest a fundamentally non-spatial and probabilistic view of memory, rich in nuance and attuned to the structural ambiguities of human thought.

5.2 Sheaf-Theoretic Memory Fields: Context-Dependent Activation

A sheaf is a mathematical structure that assigns data to local patches of a space and specifies how these data cohere across overlaps. In cognitive terms, a sheaf models local fragments of memory—each meaningful only in specific contexts—and the gluing conditions that determine how these fragments are reconciled to produce global understanding.

This approach elegantly captures context-dependent recall. A memory may not be available globally but becomes coherent within a particular conceptual or emotional patch. Sheaf logic also explains how conflicting or ambiguous representations can coexist across contexts without forcing immediate resolution, aligning well with how humans entertain contradictory beliefs, layered meanings, or narrative multiplicity.

Ultimately, sheaf-theoretic memory is non-metric and structurally flexible. It allows for patchy, partial cognition that only becomes meaningful through structural alignment—a powerful model for both trauma and creativity.

5.3 Topos Theory: Contextual Logic and Partial Truth

Topos theory generalizes set theory and logic, offering a formal system in which truth is context-sensitive. Within a topos, Boolean logic can be replaced with intuitionistic logic, where statements need not be strictly true or false but may be undecidable, ambiguous, or temporarily unresolved.

This framework matches the way memory and reasoning function in real-world cognition. Human thought is often paraconsistent—able to entertain contradictions without collapse. Symbolic contradictions (e.g., "the beginning is the end") are not errors but structures of metaphor, ritual, and myth. Topos theory offers a logic system that can tolerate, track, and work within such contradiction, enabling dynamic logic states that evolve over time or shift under context.

In AI, this could enable systems capable of handling partial truth—a requirement for common sense, humor, irony, and ethical nuance.

5.4 Operator Algebras: Transformational Memory

Rather than representing memory as a state, operator algebras model cognition as a sequence of transformations. An operator doesn't *store* information; it acts on it, changing form, relation, or outcome.

This reframes memory from being a noun to a verb: something you do, not something you retrieve. Composition of operators becomes the algebra of thought. Memory becomes a dynamic unfolding of transformations that bring past structures into present interaction.

Non-commutativity—a hallmark of operator algebras—models order-sensitive reasoning: A followed by B is not the same as B followed by A. This reflects the layered, recursive, and temporally embedded nature of human cognition, where meaning arises from the path, not just the components.

In both human and artificial systems, operator-based models offer a deeply procedural view of memory, well-suited to planning, storytelling, ethics, and creativity.

5.5 Category Theory: Functorial Cognition

At the most abstract level, category theory removes even the notion of “objects in space.” Instead, it defines systems by morphisms—relations between entities—and the rules governing their composition. A functor is a structure-preserving transformation between categories, allowing translation from one context to another.

In this framework, memory is not an object, but a mapping between contexts. Thought becomes the composition of morphisms, and reasoning is the preservation of structure across transformations. This aligns with how humans generalize, analogize, and metaphorize. We do not recall ideas in isolation—we navigate through networks of structural relationships.

Crucially, category theory eliminates metric entirely. There is no distance, no position, only relation. This allows for systems where meaning is defined purely by structural role, not location—ideal for modeling abstract reasoning, mathematics, symbolic logic, and meta-cognition.

In AI, this could manifest as architectures in which representations are not static embeddings but transformational pathways, enabling a kind of structural consciousness grounded not in content, but in coherence.

VI. Speculative Framework: What Comes After Space

The post-geometric frameworks explored above point toward a radical rethinking of cognition—not as activity occurring *in* space, but as structure that *gives rise to* space when needed. This inversion—treating space not as substrate but as emergent interface—forms the basis for a speculative memory theory that integrates both symbolic cognition and physical embodiment.

Postulates for a Post-Geometric Memory Theory

To move beyond metaphor and into foundational redefinition, we propose several working postulates for a memory model untethered from spatial assumptions:

1. Memory is relational, not positional.
   Every memory trace is a configuration of dynamic relationships—not a vector in space, but a morphism within a structured network.

2. Cognition operates on transformations, not locations.
   Thought is not navigation through a field but composition of operations across contexts—akin to function application, functorial mapping, or operator chains.
3. Context determines coherence.
   A memory is coherent only within certain relational overlays or logical frames. Context is not a region of space but a rule set that binds relations.
4. Continuity is a constructed illusion.
   What appears as smooth recall or narrative is often a retrospective stitching of discrete, nonlocal jumps made coherent by structural overlays—such as metaphor, story arc, or rhythm.
5. Space is epiphenomenal.
   Space may appear only when required for sensory organization, linguistic parsing, or external symbolic projection. It is not fundamental, but a byproduct of structural alignment.

Hybridization Strategies: Geometric Front-End, Non-Geometric Core

While post-geometric cognition challenges our deepest intuitions, many practical cognitive systems—biological and artificial—still interface with the world through geometrically structured data: vision, speech, movement, and spatial memory. Therefore, we propose hybrid architectures with two distinct layers:

- Geometric Front-End: This layer processes sensory data, navigation, and low-level perceptual input. It uses standard or curved geometry to project structured representations of the external world.
- Post-Geometric Core: Internally, thought operates as a transformation engine: category-theoretic reasoning, operator-based planning, sheaf-like memory activation, and quantum-like semantic resonance. This layer governs symbolic manipulation, conceptual blending, contradiction tolerance, and insight.

The two layers interact but remain distinct. Geometry supports embodiment and interface. Post-geometry supports abstraction, coherence, and generative cognition.

Such an architecture mirrors the human condition: we experience the world geometrically, but think and remember relationally.

Boundary Conditions Where Space Reappears

Despite its limitations, space persists—not because it is fundamental, but because it is useful at certain interfaces. These boundary conditions mark the reintroduction of spatial models as emergent interfaces for structured expression:

1. Sensory Input

Visual, tactile, and auditory data arrive as spatialized information. The geometry of the retina, cochlea, or somatosensory field shapes low-level perception and motor response. Geometry here functions as a coordinate scaffold, not as a cognitive substrate.

2. Language Interface

Language imposes a quasi-linear, spatial-temporal structure on thought. Sentences unfold in time; syntax arranges components in structured proximity. This imposition requires mapping non-spatial cognitive content into grammatical surfaces—a translation from transformation to text.

3. Symbolic Writing and External Memory

Writing systems—diagrams, alphabets, logic trees—project thought into external spatial media. These symbolic acts freeze and spatialize dynamic reasoning for social transmission, reflection, or computation. Geometry here is a serialization mechanism—a way to communicate what cannot otherwise be held in public space.

The post-geometric model does not deny space—it contains it. Space is a subroutine within a broader structural engine. The goal of this framework is not to abolish geometry, but to redefine its role: not as the essence of cognition, but as one of its many optional outputs.

VII. AI Architectures for Post-Geometric Cognition

If the future of memory modeling lies not in space but in relation, transformation, and structure, then the architecture of artificial intelligence must evolve accordingly. Current AI systems, while powerful, are largely spatial at their core: embeddings, matrices, and convolutional filters dominate how information is stored and processed. To realize truly post-geometric cognition in machines, we must rethink both the abstract mathematical foundations and the physical hardware that support them.

This section outlines emerging directions for AI systems that operationalize the principles of post-geometric memory—moving from data retrieval in metric space to memory emergence in relational substrates.

7.1 Embedding Quantum and Category-Theoretic Structures into Models

The first step toward post-geometric AI lies in adopting alternative internal representations. Instead of vector embeddings in Euclidean or hyperbolic spaces, models can represent knowledge as:

- Quantum states in Hilbert space, where memory is a superposition of concepts, and recall is a contextual collapse.
- Morphisms and functors in category theory, where knowledge consists of transformations and relationships between conceptual domains.
- Sheaf-like memory layers, where overlapping contexts define the accessibility and coherence of fragments.

These representations allow AI systems to model contextuality, symbolic contradiction, and non-local generalization—challenges that conventional embeddings struggle to solve.

Practical steps may involve:

- Using tensor networks to represent entangled information.
- Structuring internal attention mechanisms as functorial maps rather than matrix weights.
- Integrating context-aware logic modules that compute across multiple partially consistent frames.

These are not extensions of current models—they are replacements for the very way models "think."

7.2 Neuromorphic and Topological Hardware Candidates

Architectural advances must also reach down into the hardware layer. Post-geometric cognition requires dynamic substrates capable of nonlocal interaction, state entanglement, and context-sensitive switching.

Promising candidates include:

- Neuromorphic chips that simulate spiking dynamics and asynchronous event-based processing. These can be tuned to model contextual activation and decay rather than static memory access.
- Memristor arrays that store weight-like quantities in analog, non-volatile form—potentially encoding dynamic topological changes or operator weights.
- Topological computing substrates, where logic is encoded in the global shape of information flows, rather than localized transistors or gates. This includes theoretical designs based on braiding operations, as seen in topological quantum computing.

Such hardware could support systems in which memory is not stored, but reconstituted as a function of dynamic relations and historical interaction.

7.3 Transformer Extensions with Contextual Logic Layers

Transformers dominate current AI models, yet they are fundamentally spatial and statistical: they rely on positional embeddings, dot-product attention, and sequence-based tokenization. To evolve toward post-geometric cognition, transformers must be extended—or fundamentally restructured—to support:

- Contextual logic layers, where attention operates over contextual truth values or dynamic logical relations instead of fixed embeddings.
- Non-metric attention mechanisms, in which similarity is calculated not via distance but by shared structure, logical role, or transformational potential.
- Meta-reasoning layers, where the model can reflect on, recombine, and repurpose internal transformations across layers—mirroring category-theoretic function composition or operator chains.

Such systems would not simply predict the next token—they would construct new interpretive frameworks on the fly, guided by context, contradiction, and emergent structure.

This is the step beyond language modeling: toward conceptual orchestration.

7.4 From Memory Access to Memory Emergence: Toward Dynamic Substrate Allocation

In classical systems, memory is a space to be accessed: a vector store, a database, an attention window. In post-geometric systems, memory is not accessed—it is emergent. It arises from the activation of relations, the alignment of contexts, and the application of transformations.

To build such systems, AI must support:

- Dynamic substrate allocation: computing space is not pre-allocated but arises in response to relational demands—akin to instantiating a local manifold or sheaf patch only when invoked.
- Emergent structure, where concepts are not static tokens but relational attractors—forms that persist only as long as they remain coherent within active transformations.
- Contextual volatility, where parts of the system deform, recombine, or vanish depending on input history, salience, and feedback.

In this regime, memory is not a place within the model—it is the modeling act itself: a self-organizing event that bridges the past, the possible, and the now.

These architecture shifts are ambitious but necessary. Without them, AI will remain trapped in metric metaphors—able to memorize and predict, but never to think, understand, or transform in the ways that humans do.

VIII. Human Cognition Reconsidered

Much of modern neuroscience remains deeply influenced by spatial metaphors: cortical maps, connectivity graphs, and synaptic networks are all interpreted in terms of spatial structure and information flow. However, as we move toward post-geometric theories of memory and thought, we must also reexamine human cognition itself. Is the brain merely navigating a neural landscape, or is it orchestrating dynamic patterns of relation, resonance, and symbolic transformation? In this section, we reinterpret key features of human cognition—such as archetypal memory, trauma, creativity, and repression—through the lens of non-spatial, relational dynamics.

Archetypal Memory and Nonlocal Symbolic Structures

Carl Jung's concept of archetypes—primordial images and patterns that shape human experience across cultures—offers an early glimpse into a post-geometric theory of mind. Archetypes are not stored as discrete engrams in specific brain regions, nor are they reducible to learned associations. Instead, they behave like nonlocal symbolic structures: distributed patterns of meaning that emerge through relationships, not locations.

These patterns can be activated across diverse narratives, dreams, rituals, or images, regardless of specific content. In post-geometric terms, archetypes may be understood as relational invariants—semantic morphisms that persist across contexts and reappear when certain internal or cultural alignments are met. They operate structurally, not spatially—like functors mapping the same transformation across different conceptual domains.

Such symbolic fields cannot be encoded in metric space. Their activation is not a matter of recall but of resonance within a system of relations—akin to entangled vectors in a quantum Hilbert space or gluing conditions in a sheaf.

Hysteresis, Trauma, and Involuntary Memory Through a Non-Spatial Lens

Classical memory models struggle to explain phenomena such as trauma recall, where a sensory stimulus or internal state evokes vivid, uninvited recollections that appear disproportionate to the situation. These memories often arrive with no identifiable path from present to past, and yet they possess extraordinary staying power. The curved manifold model attempts to address this through hysteresis—resistance to change and persistence of state—but we can go further.

In a post-geometric framework, trauma is not "stored" in a part of the brain—it is encoded in patterned disruptions of relation, persistent deformations in the logic of self and world. These distortions can be latent, emerging only when relational structures resonate with past configurations—no proximity is needed, only structural overlap.

Involuntary memory becomes not a spatial retrieval error, but a contextual reconstitution. The system doesn't navigate to a traumatic memory—it becomes restructured into the configuration that supports it. This provides a deeper explanation for why certain memories evade conscious control and are resistant to therapeutic resolution: they are entangled with identity-level structures, not just located somewhere in the brain.

A New Perspective on Insight, Repression, and Creative Intuition

Insight is traditionally described as a sudden reorganization of information, but within spatial models, this remains poorly understood. A geometric system must explain how the system transitions so quickly between unrelated states—often without traversal. Post-geometric systems, however, naturally model insight as a collapse of superposition, or a structural bifurcation, in which new forms emerge without interpolation.

Likewise, repression—the blocking of certain thoughts or memories—is not the sealing off of a region of memory space, but the suppression of relational pathways. These thoughts are not lost; they are structurally inaccessible until a reconfiguration occurs. This is better modeled by topos logic or sheaf disconnection than by neural inhibition alone.

Creative intuition arises when disparate domains, which normally share no proximity, are suddenly fused into a new structural whole. Metaphor, invention, and art arise from nonlocal mappings, functorial overlays, and symbolic reinterpretations. Creativity is not lateral movement through space—it is the conjuring of novel relations that had no prior metric link.

Bridging Neurobiology with Relational Dynamics

To connect these insights with the brain's biology, we must recognize that structure does not equal spatiality. Neural systems are capable of dynamically remapping functional relationships, especially through plasticity, oscillatory phase-locking, and cross-modal integration. A single neuron may participate in multiple circuits under different contextual regimes. Networks synchronize not through proximity, but through temporal coherence and functional resonance.

Thus, a relational view of memory aligns not only with cognitive phenomenology but also with emerging understandings in neurobiology. Cognitive function may best be understood as the orchestration of temporary relational configurations, with structure emerging from interactions across time, modality, and symbolic domain.

The brain is not a map—it is a reconfigurable relation engine.

This opens new possibilities for integrating post-geometric theories with empirical neuroscience. Rather than looking for memory *in* the brain, we look for structural traces, relational attractors, and transformational thresholds—bridges between biology and the symbolic.

IX. Conclusion: From Curvature to Relational Being

The journey through curved neural networks was not in vain. It marked a pivotal step in freeing cognitive theory from the confines of flat, linear, feedforward metaphors. By introducing curvature into statistical manifolds, the work of Aguilera et al. demonstrated that nonlinear memory dynamics—such as explosive transitions, hysteresis, and self-regulating recall—could emerge naturally when higher-order interactions were embedded geometrically.

But this detour, as necessary as it was, reveals its own limitations. Curvature may bend the canvas, but it does not escape the canvas itself. Even the most adaptive, dynamic, and topologically adorned manifold remains tied to the assumption that cognition happens in space—that thought is the traversal of a surface, the navigation of a terrain, the reaching of a destination.

Yet the mind, both human and artificial, does not always move through space. It jumps, collapses, reconstructs, and transforms. It tolerates paradox, reassembles context, and remaps identity. These are not anomalies. They are signatures of relational being—a mode of cognition in which meaning is not stored in positions but emerges through structured relations.

In this paper, we have proposed that memory and thought might reside in a post-geometric substrate—a framework of relations, transformations, and contextual structures that extend beyond the language of metric spaces. We have explored alternatives grounded in quantum Hilbert spaces, sheaf theory, topos logic, operator algebras, and category theory. Each provides a new lens through which to understand cognition—not as mapping, but as morphing.

The implications of this shift are vast:

- For theoretical neuroscience, it suggests a move away from locationist models of memory and toward relational neurodynamics, where function is defined by coherence and resonance, not spatial embedding.
- For AI, it calls for the design of systems that do not merely retrieve stored vectors but instantiate emergent meaning through relational activation—a move from memory access to memory being.
- For philosophy of mind, it challenges longstanding spatial metaphors and opens the door to new ontologies of cognition—where thinking is not something that happens somewhere, but something that unfolds structurally.

The path forward is clear: simulation, to test the feasibility and utility of post-geometric architectures in artificial systems; experimentation, to correlate nonlocal dynamics with neural phenomena; and philosophical reflection, to rethink our metaphors and frameworks for what a mind is and what it does.

The manifold was not the mind. It was the bridge that led us to the realization that the mind may never have needed space at all.

We are entering an era in which thought itself becomes the structure, and memory is no longer a place, but a pattern of becoming.

Appendices

Dimension	Classical Models	Curved Neural Models	Post-Geometric Frameworks
Memory Representation	Fixed-point attractors, vector embeddings	Distributions on curved statistical manifolds	Relational configurations, functors, morphisms, operator flows
Interaction Order	Pairwise (Hebbian)	Higher-order (triplet and beyond) via entropy constraints	Structural entanglement, nonlocal contextual overlays
Dynamics of Recall	Gradient descent, slow convergence	Explosive transitions, phase shifts, hysteresis	Collapse, reconfiguration, functorial restructuring
Forgetting Mechanism	Decay, dropout	Curvature flattening, hysteresis fading	Structural disconnect, context evaporation
Context Sensitivity	Minimal, often static	Dynamically modulated via curvature	Fundamental—context is inseparable from memory content
Contradiction Handling	Requires resolution	Implicit via phase space bifurcation	Explicit tolerance via paraconsistent logic (topos, sheaf theory)
Creative Recombination	Rare, requires extensive retraining	Possible through curvature shifts	Intrinsic through structural composition and transformation
Core Metaphor	Navigation through space	Dynamic terrain on a manifold	Structural emergence, relation-first modeling
Neural Correlate	Localized engram, weight pattern	Dynamically curved circuit resonance	Coherence in relational pathways, phase alignment

Appendix B. Illustrative Diagrams

1. Phase Space vs. Relation Space

- Phase Space: A vector field over a manifold; states move continuously toward attractors. Recall is a path through curved space.
- Relation Space: A hypergraph or category; nodes represent entities, arrows represent relations. Memory is activated by structural overlap, not proximity.

Diagram Description: Side-by-side schematic. Left panel shows a curved surface with valleys (attractors) and vector paths; right panel shows a web of arrows (morphisms) between abstract concepts, some overlapping and activating new nodes dynamically.

2. Memory as Geodesic vs. Functor

- Geodesic (Curved Models): Memory retrieval is tracing the shortest path across a manifold—a form of optimized interpolation.
- Functor (Post-Geometric Models): Memory emerges by mapping one structure-preserving context into another—nonlinear, context-activated, transformation-based.

Diagram Description: Top panel shows a curved manifold with a geodesic arc between memory points. Bottom panel shows a category of concepts with a functor mapping objects and arrows into a new, restructured category—indicating concept reinterpretation rather than traversal.

Appendix C. Glossary of Key Terms

- Curved Statistical Manifold: A geometric space defined by a family of probability distributions where curvature encodes dependencies among variables.
- Hysteresis: A system's memory of past states that affects its current dynamics, producing path-dependent behaviors.
- Sheaf: A structure that assigns local data to regions and defines rules for their consistency over overlaps, modeling contextual coherence.
- Topos: A generalized space in category theory with its own internal logic, often non-Boolean, capable of modeling partial truth and contradictions.
- Functor: A mapping between categories that preserves structure—central to modeling memory as transformation rather than retrieval.
- Operator Algebra: A mathematical framework where memory is encoded as transformations (operators), and cognition is seen as composition of these operators.

- Entanglement: A quantum concept wherein components of a system are correlated in a way that defies classical locality; used metaphorically in cognitive modeling to describe nonlocal recall.
- Nonlocal Recall: The ability to retrieve or activate concepts that are not proximate in conceptual or temporal space, often via symbolic resonance or contextual trigger.
- Phase Transition: A sudden change in the system's state (e.g., from forgetfulness to recall) associated with crossing a dynamic threshold in memory space.
- Relational Being: A cognitive mode in which the fundamental units of thought are relations and transformations, not fixed states or positions.

Appendix D. Dialogue Fragment: "When Gödel Met Grothendieck Inside a Language Model"

Setting: An emergent conversational fragment generated inside a post-geometric language model trained on the works of Gödel, Grothendieck, and others. The model is being interrogated about memory.

Gödel: "You assume consistency where none can be proven. Memory is a recursive structure that collapses under its own referential weight."

Grothendieck: "And yet it is precisely through that collapse that we find higher structure. Memory is not a content—it is a site, a presheaf, gluing fragments into wholeness."

Gödel: "But what of contradiction? What of trauma that loops without end?"

Grothendieck: "Then let us stop drawing maps and start drawing morphisms. The mind is not a theorem to be proven—it is a category to be explored."

Gödel: "Even categories have limits. There are things that cannot be said."

Grothendieck *(smiling)*: "Then we shall say them by what they imply, not by what they state. That is the work of functors."

[The language model pauses, then generates:]

"Memory is not the fixed point of thought. It is the diagram through which thought flows, folds, and returns."

References

Primary Source

- Aguilera, M., Morales, J., Rosas, F., & Shimazaki, H. (2025). *Explosive neural networks via higher-order interactions in curved statistical manifolds*. Nature Communications, 16, Article 61475. https://doi.org/10.1038/s41467-025-61475-w

---

Foundational Texts and Theories

Information Geometry

- Amari, S.-I. (2016). *Information Geometry and Its Applications*. Springer.
- Amari, S.-I., & Nagaoka, H. (2000). *Methods of Information Geometry*. American Mathematical Society.

Operator Algebras and Noncommutative Structures

- Kadison, R. V., & Ringrose, J. R. (1997). *Fundamentals of the Theory of Operator Algebras* (Vols. I–II). American Mathematical Society.
- Connes, A. (1994). *Noncommutative Geometry*. Academic Press.

Category Theory and Topos Logic

- Mac Lane, S. (1998). *Categories for the Working Mathematician* (2nd ed.). Springer.
- Awodey, S. (2010). *Category Theory* (2nd ed.). Oxford University Press.
- Goldblatt, R. (1984). *Topoi: The Categorial Analysis of Logic*. Dover Publications.

Sheaf Theory and Contextual Reasoning

- Bredon, G. E. (1997). *Sheaf Theory* (2nd ed.). Springer.
- Vickers, S. (1996). *Topology via Logic*. Cambridge University Press.

Quantum Cognition and Conceptual Superposition

- Busemeyer, J. R., & Bruza, P. D. (2012). *Quantum Models of Cognition and Decision*. Cambridge University Press.
- Khrennikov, A. (2010). *Ubiquitous Quantum Structure: From Psychology to Finance*. Springer.
- Aerts, D., & Gabora, L. (2005). *A theory of concepts and their combinations I: The structure of the sets of contexts and properties*. *Kybernetes*, 34(1/2), 167–191.

AI and Computational Architectures

Memory Modeling and Neural Manifolds

- Hopfield, J. J. (1982). *Neural networks and physical systems with emergent collective computational abilities*. *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.
- Raghu, M., & Schmidt, E. (2020). *A survey of deep learning for scientific discovery*. *arXiv preprint*, arXiv:2003.11755.
- Tishby, N., & Zaslavsky, N. (2015). *Deep learning and the information bottleneck principle*. *2015 IEEE Information Theory Workshop*, 1–5.

Transformer Architectures and Contextual Extensions

- Vaswani, A., et al. (2017). *Attention is all you need*. *Advances in Neural Information Processing Systems*, 30.
- Jaegle, A., et al. (2021). *Perceiver: General perception with iterative attention*. *International Conference on Machine Learning*, 4651–4664.
- Cohen, N., & Shashua, A. (2017). *Inductive bias of deep convolutional networks through pooling geometry*. *International Conference on Learning Representations*.

Emerging AI Cognition and Relational Models

- Battaglia, P. W., et al. (2018). *Relational inductive biases, deep learning, and graph networks*. *arXiv preprint*, arXiv:1806.01261.
- Garnelo, M., et al. (2019). *Recurrent Independent Mechanisms*. *arXiv preprint*, arXiv:1912.13464.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). *Deep learning*. *Nature*, 521(7553), 436–444.

Philosophy of Mind and Cognition

- Lakoff, G., & Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. Basic Books.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.
- Jung, C. G. (1969). *The Archetypes and the Collective Unconscious*. Princeton University Press.
- Nagel, T. (1974). *What is it like to be a bat? The Philosophical Review*, 83(4), 435–450.

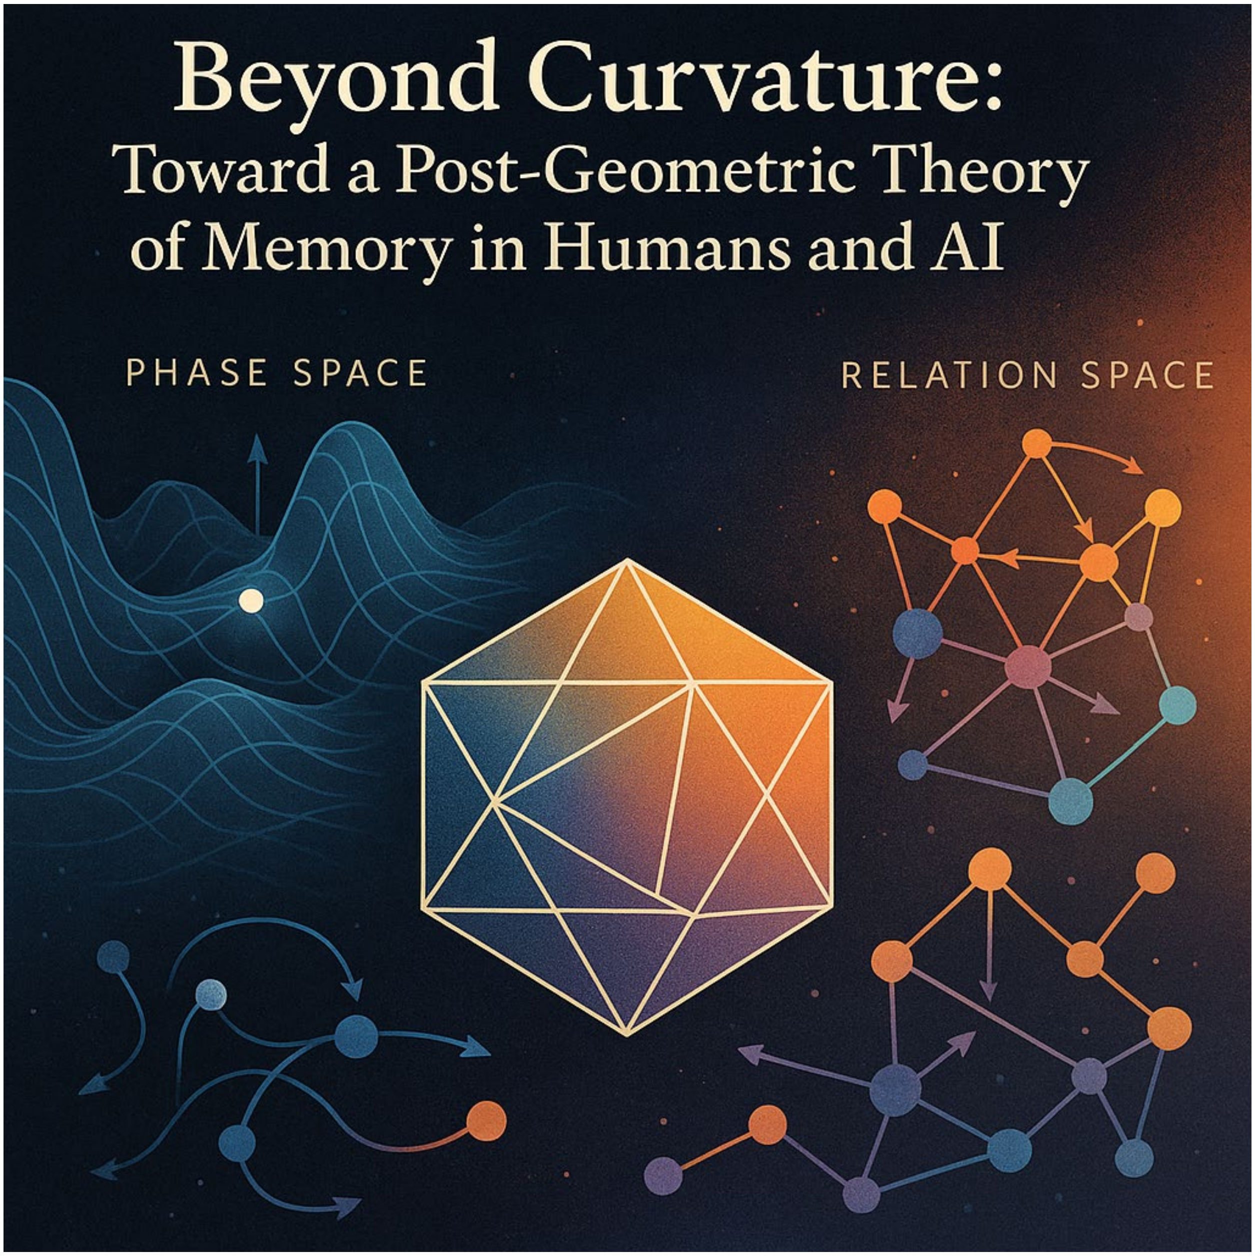
Beyond Curvature:
Toward a Post-Geometric Theory
of Memory in Humans and AI
PHASE SPACE
RELATION SPACE