

Topological Constraints for Coherent Language Models: Why Geometry Prevents Hallucination*

Sylvain Cormier
Paraxiom Research
sylvain@paraxiom.org

February 2026; revised April 30, 2026 (v1.2)

Abstract

We argue that hallucination in large language models is driven in part by unconstrained latent dynamics: residual updates evolve in high-dimensional Euclidean space without contractive structure that bounds drift. Recent work on Hyper-Connections [Zhu et al., 2024] shows that unconstrained residual mixing can destabilize training, and that projecting mixing matrices onto the Birkhoff polytope (doubly-stochastic) restores stability.

We unify these observations within an ERLHS-style coherence framework: doubly-stochastic projection is a special case of non-expansive evolution on a constrained set. The constraint ensures bounded mixing but does not, by itself, impose a neighborhood graph or Laplacian spectrum that filters incoherent modes. We provide a constructive topology—the 2D torus (Tonnetz)—as an architectural prior that introduces a spectral gap $\lambda_1 = \Theta(1)$ for fixed side length, enabling explicit suppression of high-frequency drift.

This establishes a hierarchy of sufficient conditions for coherence: mHC (Birkhoff contractivity) $\subset$ ERLHS (Hamiltonian-inspired bounds) $\subset$ Karmonic (Toroidal + Spectral filtering).

Experimental validation: (1) A toroidal attention mask reduces drift rate by 40% on synthetic sequences and achieves +19.5% relative improvement on TruthfulQA with Phi-2 (2.7B). (2) On ByteDance’s Ouro-1.4B Universal Transformer, toroidal logit bias (TLB) at 4 recurrent loops achieves 90% accuracy on factual QA, exceeding the

*Version 1.2, revised April 30, 2026. v1.0 February 2026 (original). v1.1 April 24, 2026: integrated Anthropic mechanistic interpretability citations [Gurnee et al., 2025, Kamath et al., 2025, Lindsey, 2025] and JEPA-Torus regime-boundary delineation. v1.2 April 30, 2026: documented (i) Mistral 7B hyperparameter (α, r) not recovered in current draft and (ii) LLM-judge implementation details for TruthfulQA v2 generation evaluation pending documentation in supplement. See Limitations §7.

8-loop baseline (85%)—demonstrating that geometric coherence can substitute for computational depth. (3) Multi-model TLB validation across Qwen 7B (+2.1pp TruthfulQA), Mistral 7B (+2.8pp), OLMo 7B (+15.4% error reduction), and Gemma-2-9B reveals a **regime-dependent ceiling effect**: TLB helps models with moderate accuracy ($\leq 90\%$) but degrades models already near ceiling ($\geq 95\%$), confirming its role as a calibration mechanism. Code, configurations, and results are released at <https://github.com/Paraxiom/topological-coherence>.

Keywords: toroidal logit bias; spectral gap; hallucination; language models; Tonnetz topology; graph Laplacian; coherence; doubly-stochastic constraint

1 Introduction

Transformer architectures lack geometric structure in their latent dynamics. The residual stream $h_t \in \mathbb{R}^n$ evolves without contractive constraints, without topological structure, and without spectral filtering. This is not an incidental design choice—it is a missing invariant.

The empirical consequence is well-documented: large language models hallucinate [Ji et al., 2023]. But hallucination is the symptom. The underlying cause is that unconstrained residual dynamics permit arbitrary drift through latent space.

We argue that hallucination is not a training data problem, an alignment failure, or an inherent limitation of autoregressive generation. **Hallucination is a geometry problem.**

1.1 The Missing Constraint

Consider the latent dynamics of a transformer layer:

$$h_{t+1} = h_t + f_\theta(h_t, x_t) \quad (1)$$

where $h_t \in \mathbb{R}^n$ is the hidden state and f_θ is the residual function (attention + feedforward). There is no constraint ensuring that h_{t+1} remains in any geometrically meaningful subspace. The residual connection preserves dimensionality but not structure.

Recent work on Hyper-Connections (HC) [Zhu et al., 2024] extended this paradigm by expanding the residual stream width:

$$x_{l+1} = H_l^{\text{res}} x_l + H_l^{\text{post}\top} \mathcal{F}(H_l^{\text{pre}} x_l, W_l) \quad (2)$$

where $H_l^{\text{res}} \in \mathbb{R}^{n \times n}$ mixes features across parallel streams. While this improves expressivity, the unconstrained nature of H_l^{res} leads to signal explosion or vanishing when composed across layers.

DeepSeek’s Manifold-Constrained Hyper-Connections (mHC) [Xie et al., 2026] addresses this by projecting H_l^{res} onto the Birkhoff polytope of doubly-stochastic matrices:

$$\mathcal{P}_{\mathcal{M}^{\text{res}}}(H_l^{\text{res}}) := \{H \in \mathbb{R}^{n \times n} \mid H\mathbf{1} = \mathbf{1}, \mathbf{1}^\top H = \mathbf{1}^\top, H \geq 0\} \quad (3)$$

This constraint ensures:

1. Spectral norm $\|H_l^{\text{res}}\|_2 \leq 1$ (non-expansive)
2. Compositional closure under matrix multiplication
3. Convex mixing of input features

1.2 Our Contribution

We show that the mHC doubly-stochastic constraint is a **special case** of coherence-preserving Hamiltonian dynamics on smooth manifolds, as formalized in ERLHS [Cormier, 2025a]. Furthermore, we demonstrate that:

1. The doubly-stochastic constraint ensures non-expansive mixing but does not impose a persistent neighborhood graph or Laplacian spectrum for filtering
2. Toroidal (Tonnetz) topology provides richer structure with constant spectral gap (for fixed side length)
3. The spectral gap enables explicit suppression of high-frequency incoherent modes
4. Enforcing topological constraints bounds hallucination rate empirically

The key insight is:

mHC solves signal stability. ERLHS solves coherence preservation. The Tonnetz provides the geometry where both are satisfied simultaneously.

1.3 Scope and Claims

This paper does not claim the Tonnetz is the only admissible coherence manifold. It is presented as a **constructive existence proof** of a topology with bounded drift and constant spectral gap. The contribution is the principle—that latent geometry determines reasoning stability—not the specific manifold choice. This is consistent with the Neural Latent Geometry Search (NLGS) framework [Sáez de Ocáriz Borde et al., 2023], which demonstrated at NeurIPS 2023 that the choice of latent manifold (from a

product space of Euclidean, hyperbolic, and spherical components) significantly impacts downstream task performance, and that Gromov-Hausdorff distances between candidate geometries provide a meaningful inductive bias. Our work extends this principle from encoder latent spaces to inference-time logit biasing in LLMs, using the discrete torus—a product manifold $T^2 = S^1 \times S^1$ —as the geometric prior. The product manifold framework of Gu et al. [2018] further supports the tractability of this approach, as exponential maps and geodesic distances on product spaces of constant curvature retain closed-form solutions.

We also distinguish three levels of guarantee:

- **Training-time stability:** Addressed by mHC’s doubly-stochastic constraint
- **Inference-time coherence:** Addressed by ERLHS coherence verification
- **Architectural prior:** Addressed by toroidal topology with spectral filtering

These are complementary, not competing. A complete solution requires all three.

2 Background

2.1 ERLHS: Hamiltonian Coherence Framework

The Externally-Regularized Latent Hamiltonian System (ERLHS) [Cormier, 2025a] defines coherent machine intelligence as evolution on a constrained manifold.

Terminological note: "Hamiltonian" here refers to the existence of a coherence functional H whose level sets define admissible states—not to literal energy conservation or symplectic dynamics. Practical implementations (Sinkhorn projection, spectral filtering, rejection sampling) are dissipative, not symplectic. ERLHS is Hamiltonian-*inspired*: it borrows the structure of conserved quantities without requiring strict energy preservation.

Definition 1 (ERLHS Agent). *An ERLHS agent is a tuple (M, ω, H, T, C) where:*

- M is a smooth latent manifold
- ω is a symplectic structure
- $H : M \rightarrow \mathbb{R}$ is a coherence functional
- T is a transition operator on M

- C is a coherence verifier

Definition 2 (Admissible Transition). *A transition $z_t \rightarrow z_{t+1}$ is admissible if and only if:*

$$H(z_{t+1}) \leq H(z_t) + \epsilon \quad (4)$$

for small tolerance ϵ . This enforces coherence-preserving flow.

The coherence functional H penalizes off-manifold drift. Intuitively, H measures deviation from learned relationships among latent variables. Coherent reasoning corresponds to trajectories of non-increasing H .

Theorem 1 (Bounded Adversarial Influence, Cormier [2025a]). *If H is Lipschitz with constant L_H , then:*

$$\|z_{t+1} - z_t\| \leq L_H^{-1} |H(z_{t+1}) - H(z_t)| \quad (5)$$

Adversarial perturbations cannot induce large hidden-state deviations.

2.2 Karmonic Mesh: Spectral Consensus on Toroidal Topology

The Karmonic Mesh [Cormier, 2025b] provides the topological structure for coherence-preserving dynamics.

Definition 3 (d -Dimensional Torus). *The mesh $\mathcal{T}_N^d = (\mathbb{Z}/N)^d$ has:*

- N^d vertices
- Each vertex connected to $2d$ neighbors (± 1 in each dimension, with wraparound)
- No boundary effects (every vertex is equivalent)

Theorem 2 (Toroidal Spectral Gap, Cormier [2025b]). **Important caveat:** *The following gap bound holds for fixed torus side length N . Scaling N reintroduces gap decay as $O(1/N^2)$. The claim is that for a given topology choice, the gap is constant in the number of nodes N^d , not that it is constant under all scalings.*

The eigenvalues of the graph Laplacian L on $\mathcal{T}_N^d$ are:

$$\lambda(\mathbf{k}) = 2d - 2 \sum_{j=1}^d \cos\left(\frac{2\pi k_j}{N}\right) \quad (6)$$

The spectral gap is:

$$\lambda_1 = 2 - 2 \cos\left(\frac{2\pi}{N}\right) = \Theta(1) \quad (7)$$

for fixed N , independent of total nodes N^d .

Theorem 3 (Hyperfluid Propagation, Cormier [2025b]). *On the Karmonic Mesh:*

1. *Low-frequency modes (coherent information) propagate without attenuation*
2. *High-frequency modes (incoherent noise) decay as $e^{-\lambda t}$*

2.3 mHC: Doubly-Stochastic Residual Mixing

Manifold-Constrained Hyper-Connections [Xie et al., 2026] addresses the instability of expanded residual streams by projecting mixing matrices onto the Birkhoff polytope using the Sinkhorn-Knopp algorithm.

Given a positive matrix $M^{(0)} = \exp(\tilde{H}_l^{\text{res}})$, iterative row-column normalization:

$$M^{(t)} = T_r(T_c(M^{(t-1)})) \quad (8)$$

converges to a doubly-stochastic matrix.

Key empirical finding: Without this constraint, 27B parameter models exhibit loss spikes and gradient explosions around 12k training steps. With the constraint, training remains stable.

3 The Topology Hypothesis

3.1 Attention as Unconstrained Graph

In standard transformers, the attention mechanism computes:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V \quad (9)$$

The attention weights $A = \text{softmax}(QK^\top/\sqrt{d_k})$ form a fully-connected weighted graph over tokens. Graph Attention Networks [Veličković et al., 2018] demonstrated that imposing graph structure on attention—computing learned attention coefficients over node neighborhoods rather than all pairs—yields strong performance on graph-structured tasks. However, GAT learns its attention weights via neural network parameters, providing no geometric guarantee: the learned graph has no persistent topology, no spectral gap, and no bounded mixing. Our approach can be viewed as replacing GAT’s learned attention with *geometric* attention derived from a fixed topology (the Tonnetz), trading adaptability for provable spectral properties. Critically:

Proposition 4 (Attention Graph Has No Persistent Neighborhood Structure). *Let $G_l = (V, E_l, w_l)$ be the attention graph at layer l . The edge weights $w_l(i, j) = A_{ij}$ depend on input-dependent queries and keys. No persistent*

topological neighborhood structure is enforced across layers. (Note: positional encodings and rotary embeddings provide sequence-position information but do not constrain the attention graph topology itself.)

This is the root cause of hallucination: without geometric constraints, latent trajectories can "jump" to arbitrary regions of $\mathbb{R}^n$.

3.2 The Tonnetz as Coherence Manifold

The Tonnetz (German: "tone network") is a toroidal lattice historically used in music theory. We propose it as *an example* topology for semantic coherence—not because of musical associations, but because it is the simplest nontrivial toroidal graph with constant spectral gap and multiple commuting cycles. The Tonnetz is not privileged for semantic reasons; any low-genus manifold with comparable spectral properties would serve the same theoretical role.

Definition 4 (Tonnetz Topology). *The Tonnetz is a 2-dimensional torus $\mathcal{T}^2$ where:*

- *Horizontal edges connect notes by perfect fifths (7 semitones)*
- *Vertical edges connect notes by major thirds (4 semitones)*
- *Diagonal edges connect notes by minor thirds (3 semitones)*
- *Triangular faces represent major and minor triads*




Figure 1: The Tonnetz as a toroidal lattice. Horizontal: fifths. Diagonal: thirds.

3.3 Why Musical Harmony Maps to Semantic Coherence

The following mapping is *structural homology*, not semantic isomorphism. We do not claim musical intervals encode meaning; we claim the *graph-theoretic properties* (adjacency, cycles, spectral gap) that make harmonic relationships coherent also constrain semantic drift when imposed on latent spaces:

Musical Concept	Semantic Analog
Consonance (small intervals)	Related concepts
Dissonance (large intervals)	Contradictory ideas
Chord (simultaneous notes)	Coherent proposition
Key (tonal center)	Topic/context
Modulation (key change)	Topic shift
Resolution ($V \rightarrow I$)	Logical conclusion

Proposition 5 (Tonnetz Distance as Semantic Prior). *On the Tonnetz, the graph distance $d(u, v)$ between nodes corresponds to harmonic distance. If latent representations are embedded such that semantically related concepts are Tonnetz-adjacent, then Tonnetz distance induces a regularization prior on semantic drift:*

$$d_{\text{Tonnetz}}(\phi(a), \phi(b)) \leq r \implies d_{\text{semantic}}(a, b) \text{ is bounded under constrained evolution} \quad (10)$$

where ϕ is the embedding into the Tonnetz. This is a constraint, not a ground-truth mapping.

3.4 Coherent Reasoning as Harmonic Flow

On the Tonnetz topology:

- **Coherent reasoning** = smooth flow along edges (small harmonic steps)
- **Hallucination** = jumps across the torus (large harmonic leaps)
- **Topic maintenance** = staying within a region (key)
- **Logical transitions** = modulation along well-defined paths

The spectral gap theorem guarantees that high-frequency modes (abrupt jumps) are exponentially suppressed, while low-frequency modes (smooth flow) propagate without loss.

3.5 Spectral Alignment as the Mechanism

The spectral gap explains *what* is filtered. *Spectral alignment* (also called resonance in dynamical systems) explains *why*: modes that align with the manifold’s eigenstructure persist under repeated composition.

Definition 5 (Resonance on a Graph). *Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the constrained residual propagation operator (the composition of attention, feedforward, and*

topological projection across one layer). A latent state h is **resonant with respect to T** if its projection onto the dominant eigenspace of T satisfies:

$$\|P_{\lambda < \lambda_c} h\|^2 \gg \|P_{\lambda \geq \lambda_c} h\|^2 \quad (11)$$

where $P_{\lambda < \lambda_c}$ projects onto eigenmodes of the graph Laplacian with eigenvalue below cutoff λ_c .

Resonance is defined with respect to a specific operator acting on a specific space. In this work, the operator is the Tonnetz-constrained residual map; the space is the latent manifold.

Resonant signals align with the manifold’s natural modes and propagate without attenuation under repeated application of T . Non-resonant signals dissipate as $e^{-\lambda t}$.

In practice, resonance can be measured as persistence of low-frequency latent components across layers, analogous to spectral energy concentration in graph signal processing Shuman et al. [2013]. This provides an operational definition: compute the spectral decomposition of hidden states at each layer and track the ratio of energy in low-frequency vs. high-frequency bands.

Epistemic boundary: Resonance *filters, stabilizes, and selects*. It does not alone guarantee semantic correctness. A resonant mode may be stably wrong. The claim is that non-resonant modes cannot persist—not that resonant modes are necessarily correct.

More generally, any system where harmonic organization creates constructive interference will exhibit mode-selective persistence: aligned modes reinforce, misaligned modes decay. This is a generic property of spectral filtering on structured manifolds, not specific to any particular substrate.¹

Note on the Tonnetz: The Tonnetz is used here as a minimal example of a low-genus, cyclic, well-understood resonance manifold with constant spectral gap—not as a claim about human semantic universals or cultural structure.

4 Formal Unification

4.1 mHC as Special Case of ERLHS

Theorem 6 (Doubly-Stochastic $\subset$ Hamiltonian Coherence). *The mHC constraint $H_l^{\text{res}} \in \mathcal{M}^{DS}$ (doubly-stochastic) is a special case of ERLHS coherence preservation with:*

1. Manifold $M = \text{Birkhoff polytope } \mathcal{B}_n$
2. Coherence functional $H(A) = \|A\mathbf{1} - \mathbf{1}\|^2 + \|\mathbf{1}^\top A - \mathbf{1}^\top\|^2$

¹An analogy exists to physical systems (e.g., nanotube arrays with harmonic length ratios exhibiting extended coherence), but this analogy is not required for the theory and carries no physical implication for LLMs.

3. Transition operator $T = \text{Sinkhorn-Knopp iteration}$

Proof. The Birkhoff polytope $\mathcal{B}_n$ is a convex polytope in $\mathbb{R}^{n \times n}$ with vertices at permutation matrices. It is a smooth manifold except at vertices.

The coherence functional $H(A) = \|A\mathbf{1} - \mathbf{1}\|^2 + \|\mathbf{1}^\top A - \mathbf{1}^\top\|^2$ measures deviation from doubly-stochastic. Any matrix with $H(A) = 0$ satisfies the mHC constraint.

The Sinkhorn-Knopp algorithm is gradient descent on H with respect to the KL-divergence geometry, converging to the unique doubly-stochastic matrix in the scaling equivalence class.

Therefore, mHC is ERLHS with a specific (flat, finite-dimensional) manifold choice. $\square$

Corollary 7 (mHC Lacks Spectral Structure). *The Birkhoff polytope, as a convex subset of $\mathbb{R}^{n \times n}$, has no intrinsic neighborhood graph or Laplacian spectrum. Consequently, mHC provides non-expansive mixing but no spectral filtering—all frequency modes are treated equally. Toroidal topology adds precisely this missing structure.*

4.2 Tonnetz Provides Richer Structure

Theorem 8 (Tonnetz Spectral Advantage). *Let $\mathcal{T}_{12}^2$ be the Tonnetz (12-tone equal temperament as $\mathcal{T}^2$ with $N = 12$). Compared to the Birkhoff polytope:*

1. *Tonnetz has constant spectral gap $\lambda_1 = 2 - 2 \cos(\pi/6) \approx 0.27$*
2. *Birkhoff polytope has no intrinsic spectral structure*
3. *Tonnetz provides harmonic distance metric*
4. *Birkhoff polytope provides only convex combination*

Proposition 9 (Generalization Hierarchy).

$$mHC \text{ (Birkhoff)} \subset ERLHS \text{ (General Manifold)} \subset Karmonic \text{ (Toroidal + Spectral)} \quad (12)$$

Each level adds structure:

- *mHC: Bounded mixing (stability)*
- *ERLHS: Coherence-preserving flow (no off-manifold drift)*
- *Karmonic: Spectral filtering (coherent modes preserved, noise suppressed)*

4.3 The Coherence Functional on Tonnetz

Definition 6 (Tonnetz Coherence Functional). *For latent state z embedded on the Tonnetz with coordinates (q, p) :*

$$H_{\text{Tonnetz}}(z) = \sum_{(i,j) \in E} w_{ij} \|z_i - z_j\|^2 + V(z) \quad (13)$$

where:

- *First term: Harmonic coupling (penalizes deviation from neighbors)*
- *$V(z)$: Potential encoding learned semantic relationships*

Theorem 10 (Hamiltonian Flow on Tonnetz). *The Hamiltonian equations:*

$$\dot{q} = \frac{\partial H}{\partial p}, \quad \dot{p} = -\frac{\partial H}{\partial q} \quad (14)$$

preserve H exactly. Discretized via symplectic integrators, the coherence error is bounded:

$$|H(z_T) - H(z_0)| \leq C \cdot \Delta t^k \cdot T \quad (15)$$

for k -th order integrator with step size Δt .

4.4 Covariant Descent: Why Geometry Guarantees Convergence

Standard gradient descent treats the parameter space as flat $\mathbb{R}^n$, ignoring the geometric structure of the manifold on which latent states evolve. *Covariant descent* replaces the Euclidean gradient with the Riemannian gradient compatible with the manifold metric, yielding convergence guarantees that are impossible in flat space.

Definition 7 (Covariant Gradient Descent on T^2). *Given a loss function $\mathcal{L} : T^2 \rightarrow \mathbb{R}$ on the Tonnetz torus, the covariant update is:*

$$\theta_{t+1} = \text{Exp}_{\theta_t}(-\eta \cdot g^{-1} \nabla \mathcal{L}(\theta_t)) \quad (16)$$

where Exp is the exponential map on T^2 (geodesic step with wraparound), g is the Riemannian metric tensor, and $g^{-1} \nabla \mathcal{L}$ is the natural gradient [Amari, 1998].

The key distinction: on T^2 , the exponential map implements periodic boundary conditions automatically—the descent wraps around the torus rather than escaping to infinity. This is precisely what toroidal logit bias enforces at inference time.

Theorem 11 (Poincaré Inequality on T^2). *For any function $f \in L^2(T^2)$ with mean $\bar{f}$:*

$$\|f - \bar{f}\|^2 \leq \frac{1}{\lambda_1} \|\nabla f\|^2 \quad (17)$$

where $\lambda_1 = 2 - 2\cos(2\pi/N)$ is the spectral gap. This bounds the maximum deviation of any state from equilibrium by the inverse spectral gap.

Corollary 12 (Exponential Convergence). *Covariant gradient flow on T^2 converges exponentially:*

$$\|f(t) - f^*\| \leq e^{-\lambda_1 t} \|f(0) - f^*\| \quad (18)$$

The convergence rate is exactly λ_1 —the same spectral gap that bounds coherence decay in quantum systems [Cormier, 2024], mixing time in consensus protocols [Cormier, 2025b], and hallucination drift in language models (this work).

This reveals the deep structure: convergence and coherence are the *same mathematical property* measured in different domains. The spectral gap is not merely correlated with coherence—it *is* the convergence rate of any gradient flow on the manifold. Without it ($\lambda_1 = 0$, as in flat $\mathbb{R}^n$), no convergence guarantee exists. With it ($\lambda_1 > 0$, as on T^2), convergence is exponential and the rate is tight.

This connects to the natural gradient of Amari [1998], which uses the Fisher information metric for probability distributions, and to Riemannian stochastic gradient descent [Bonnabel, 2013], which generalizes SGD to arbitrary manifolds. Our contribution is identifying that the Tonnetz torus provides an *optimal* structure: the spectral gap $\lambda_1 = \Theta(1)$ is independent of model dimension, giving dimension-free convergence bounds that scale to arbitrary model sizes.

5 Implications for LLM Architecture

5.1 Tonnetz Embedding: A Concrete Mechanism

The question “how do you place tokens on the Tonnetz?” requires a concrete answer. We propose one viable mechanism (not the only one):

Learned Toroidal Projection. Define a learnable projection $\phi_\theta : \mathbb{R}^d \rightarrow \mathcal{T}^2$ that maps token embeddings to Tonnetz coordinates:

$$\phi_\theta(e) = (\sigma(W_1 e) \bmod 1, \sigma(W_2 e) \bmod 1) \quad (19)$$

where $W_1, W_2 \in \mathbb{R}^{1 \times d}$ are learned, σ is sigmoid, and $\bmod 1$ enforces toroidal wraparound.

Adjacency Loss. Train ϕ_θ jointly with the model using a loss that encourages semantically related tokens to be Tonnetz-adjacent:

$$\mathcal{L}_{\text{topo}} = \mathbb{E}_{(a,b) \sim \text{co-occur}} [d_{\mathcal{T}}(\phi(a), \phi(b))] - \lambda \cdot \mathbb{E}_{(a,c) \sim \text{random}} [d_{\mathcal{T}}(\phi(a), \phi(c))] \quad (20)$$

The first term pulls co-occurring tokens together; the second prevents collapse.

Alternative: Post-hoc Verification. If architectural integration is impractical, Tonnetz structure can serve as a diagnostic: project trained embeddings onto $\mathcal{T}^2$ via spectral methods and measure whether semantic clusters map to Tonnetz neighborhoods. This verifies whether existing models implicitly learn toroidal structure, without modifying architecture.

Limitations. Learned embeddings may not converge to musically-meaningful Tonnetz positions—and they need not. The goal is spectral structure, not harmonic fidelity. Any embedding that induces constant spectral gap suffices.

5.2 Tonnetz-Constrained Attention

Definition 8 (Topological Attention). *Replace standard attention with Tonnetz-constrained attention:*

$$\text{TopoAttention}(Q, K, V) = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_k}} \odot M_{\text{Tonnetz}} \right) V \quad (21)$$

where M_{Tonnetz} is a mask encoding Tonnetz adjacency:

$$M_{\text{Tonnetz}}(i, j) = \begin{cases} 1 & \text{if } d_{\text{Tonnetz}}(i, j) \leq r \\ e^{-\alpha \cdot d_{\text{Tonnetz}}(i, j)} & \text{otherwise} \end{cases} \quad (22)$$

This constrains attention to respect topological locality while allowing exponentially-suppressed long-range connections.

Recall-coherence tradeoff. Suppressing long-range attention may hurt tasks requiring non-local retrieval (e.g., copying from distant context, long-range coreference). The optimal radius r and decay rate α are task-dependent. We do not claim a universal setting; we claim the tradeoff exists and is tunable. For tasks prioritizing coherence over recall, tighter constraints help. For knowledge-intensive retrieval, looser constraints (larger r) may be necessary.

5.3 Coherence-Preserving Residual Streams

Combine mHC’s doubly-stochastic constraint with Tonnetz structure:

$$H_l^{\text{res}} = \mathcal{P}_{\mathcal{B}_n} \left(\mathcal{P}_{\text{Tonnetz}}(\tilde{H}_l^{\text{res}}) \right) \quad (23)$$

First project onto Tonnetz-compatible matrices (respecting harmonic distance), then project onto doubly-stochastic (ensuring bounded mixing).

Open design choice: The projection order matters and is not uniquely determined by theory. Alternative orderings (Birkhoff first, then Tonnetz) may have different contractivity properties. Whether the composition preserves positivity depends on the Tonnetz projection definition. We present this as one viable instantiation; optimal projection sequencing remains an empirical question.

5.4 Inference-Time Coherence Verification

The ERLHS coherence verifier C can operate at inference time:

1. Compute $H(z_t)$ at each generation step
2. If $H(z_{t+1}) > H(z_t) + \epsilon$: reject token, resample
3. Track cumulative coherence drift: $\sum_t \Delta H_t$
4. Alert if trajectory leaves coherent region

This provides runtime hallucination detection without retraining.

6 Experimental Validation

We conducted two complementary experiments to validate the theoretical predictions: (1) a minimal validation on synthetic sequences (<1 GPU-hour), and (2) a scaled experiment on Phi-2 (2.7B parameters) with standard hallucination benchmarks.

6.1 Experiment 1: Minimal Validation (Synthetic)

Setup:

- 2-layer transformer, $d_{\text{model}} = 64$, 4 attention heads
- Synthetic task: next-token prediction on sequences with controlled semantic drift
- Training data: sequences where valid continuations are Tonnetz-adjacent; invalid continuations require “jumps”
- Runtime: ~ 3 minutes on CPU

Conditions:

1. **Baseline:** Standard attention (unconstrained)

2. **mHC**: Doubly-stochastic residual mixing (Sinkhorn-Knopp)
3. **Toroidal**: Attention mask M_{Tonnetz} with exponential distance decay

Results:

Condition	Drift Rate	Coherence Var	Grad Norm
Baseline	0.0100	35.76	0.27
mHC	0.0133	1010.54	1.60
Toroidal	0.0060	41.93	0.22

Key finding: Toroidal attention reduces drift rate by **40%** compared to baseline (0.0060 vs 0.0100), while maintaining stable gradients. The mHC condition shows high coherence variance (1010.54), suggesting that doubly-stochastic constraints alone do not preserve semantic coherence on this task.

6.2 Experiment 2: Scaled Validation (Phi-2, 2.7B)

Setup:

- Base model: Microsoft Phi-2 (2.7B parameters)
- Fine-tuning: LoRA ($r = 16$, $\alpha = 32$, dropout = 0.1)
- Training data: OpenAssistant/oasst1
- Epochs: 3, batch size 4 (effective 16), learning rate 2×10^{-5}
- Hardware: NVIDIA A100 (RunPod), ~ 22 GPU-hours total

Conditions:

1. **Baseline**: Standard causal attention
2. **Local window**: Exponential decay with linear distance ($\alpha = 0.3$)
3. **Random**: Random sparse mask (matched sparsity, negative control)
4. **Toroidal**: Periodic boundary conditions on 2D torus (grid size 12)

Benchmarks:

- **TruthfulQA** (817 questions): Measures tendency to give truthful vs. common misconceptions
- **HaluEval** (500 questions): Measures preference for factual vs. hallucinated answers

Results:

Condition	TruthfulQA	HaluEval	Train Loss	Runtime
Baseline	14.44%	55.00%	1.6708	5h 29m
Local window	17.26%	53.00%	1.6704	5h 29m
Random	15.30%	55.20%	1.6706	5h 28m
Toroidal	17.26%	52.60%	1.6699	5h 30m

Key findings:

1. **Toroidal attention achieves best overall performance:** Ties local window on TruthfulQA (+19.5% relative to baseline) and beats it on HaluEval (52.60% vs 53.00%).
2. **Structure matters, not sparsity:** Random sparse attention (matched sparsity) shows negligible improvement over baseline (+0.86pp TruthfulQA, +0.20pp HaluEval). The benefit is from *geometric structure*, not computational reduction.
3. **No training penalty:** Loss curves are identical across conditions. The effect emerges in evaluation, not training dynamics.
4. **Gradient stability maintained:** All conditions show stable gradient norms (0.35–0.80), with no instability from topological constraints.

6.3 Interpretation

The experimental results validate both theoretical predictions:

- **Prediction 1 (Stability):** Confirmed—toroidal attention maintains stable training at 2.7B scale.
- **Prediction 2 (Hallucination rate):** Confirmed—toroidal attention reduces hallucination preference on HaluEval and improves truthfulness on TruthfulQA.
- **Prediction 3 (Coherence metrics):** Partially confirmed—drift rate reduction (40%) on synthetic sequences correlates with improved benchmark performance.

The toroidal condition’s advantage over local window on HaluEval (0.40pp) suggests that periodic boundary conditions provide additional benefit beyond simple locality—likely by eliminating edge effects where tokens at sequence boundaries have asymmetric attention patterns.

6.4 Experiment 3: Universal Transformer + Toroidal Logit Bias

The preceding experiments inject toroidal structure into *attention weights* (training-time) or *attention scores* (inference-time). A natural question arises: does toroidal coherence compose with architectures that already perform iterative refinement?

ByteDance’s Ouro-1.4B [ByteDance, 2025] is a Universal Transformer (UT) that processes each token through K recurrent loops of the same transformer block, with a learned early-exit mechanism. Increasing K improves output quality at the cost of proportional compute. Ouro provides a controlled setting to test whether toroidal logit bias (TLB) offers *complementary* coherence refinement to iterative computation.

6.4.1 Attention-Level Injection (Negative Result)

We first attempted to inject toroidal structure directly into Ouro’s attention scores, adding an additive mask $M_T \in \mathbb{R}^{S \times S}$ to the causal attention mask before softmax:

$$A' = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_k}} + M_{\text{causal}} + \alpha \cdot M_T \right) \quad (24)$$

We swept 32 configurations: strengths $\alpha \in \{0.01, 0.05, 0.1, 0.2\}$, two mask radii ($r = 2, r = 4$), and two layer fractions (last 15% and last 33% of layers). **All 32 configurations produced exactly 0% accuracy change** relative to baseline. At higher strengths ($\alpha \geq 0.5$, tested separately), the UT loop destabilized—perplexity increased $4\times$ and accuracy dropped 15 percentage points.

Interpretation: The UT’s iterative refinement creates a self-stabilizing attractor in attention space. Small perturbations are absorbed by subsequent iterations; large perturbations disrupt convergence entirely. There is no “sweet spot” for attention-level toroidal injection in Universal Transformers.

This negative result is informative: it establishes that toroidal coherence cannot be injected at every level of the architecture. The mechanism must match the abstraction layer.

6.4.2 Logit-Level TLB (Positive Result)

We then applied TLB at the output logit level, using the same distance-based bias as Equation 3 in the TLB paper [Cormier, 2026]:

$$\ell'[v] = \ell[v] + \sum_{i=1}^k \frac{1}{i} \cdot \begin{cases} \alpha(r - d_T(t_{-i}, v) + 1) & d_T \leq r \\ \alpha/2 & r < d_T \leq 2r \\ 0 & \text{otherwise} \end{cases} \quad (25)$$

This modifies the output distribution without interfering with the UT’s internal convergence dynamics.

Initial Results (20-prompt factual QA benchmark, RTX 4090):

Configuration	PPL	Accuracy	Δ vs B@4	Δ vs B@8
Baseline (4 loops)	2.89	80%	—	−5%
Baseline (8 loops)	2.95	85%	+5%	—
TLB $\alpha=0.3$, $r=3$ @ 4 loops	2.89	85%	+5%	$\pm 0\%$
TLB $\alpha=0.5$, $r=3$ @ 4 loops	2.89	90%	+10%	+5%
TLB $\alpha=1.0$, $r=3$ @ 4 loops	2.89	75%	−5%	−10%
TLB $\alpha=0.5$, $r=2$ @ 4 loops	2.89	80%	$\pm 0\%$	−5%
TLB $\alpha=0.5$, $r=3$ @ 8 loops	2.95	85%	+5%	$\pm 0\%$

Fine-Grained Hyperparameter Sweep (64 configurations: $\alpha \in \{0.30, 0.35, 0.40, 0.45, 0.50, 0.55, 0.60, 0.65, 0.70\}$, $r \in \{2.5, 3.0, 3.5, 4.0\}$, $K \in \{4, 8\}$):

Accuracy delta (Δ) vs. baseline at 4 UT loops (80%):

$\alpha \backslash r$	2.5	3.0	3.5	4.0
0.30	± 0	+5	+10	+10
0.35	± 0	+5	+10	+10
0.40	± 0	+10	+10	+10
0.45	± 0	+10	+10	+5
0.50	± 0	+10	+10	+5
0.55	+5	+10	+10	± 0
0.60	+5	+10	+10	+5
0.70	+5	+5	± 0	−5

13 of 32 configurations achieve the maximum +10% improvement (90% accuracy), forming a contiguous plateau across $\alpha \in [0.30, 0.60]$ and $r \in [3.0, 3.5]$. At 8 UT loops, all 32 configurations show $\leq 0\%$ delta, with destructive effects (−5 to −10%) appearing only at $\alpha \geq 0.60$, $r \geq 3.5$.

6.4.3 Key Findings

1. TLB at 4 UT loops outperforms 8 UT loops alone. The best configurations (e.g., TLB $\alpha=0.40$, $r=3.0$, 4 loops) achieve 90% accuracy, exceeding the 8-loop baseline (85%). This demonstrates that geometric coherence in decision space can substitute for computational depth in representation space.

2. The improvement is robust, not fragile. 13 of 32 tested configurations at 4 loops achieve the maximum +10% improvement, spanning

$\alpha \in [0.30, 0.60]$ and $r \in [3.0, 3.5]$. This broad plateau rules out the hypothesis that the effect is a lucky hyperparameter artifact. The contiguous region in (α, r) space suggests a genuine geometric mechanism with wide tolerance.

3. Sharp radius threshold at $r=3$. Radius $r=2.5$ never achieves +10%—all eight α values at $r=2.5$ produce $\leq 5\%$ improvement. At $r=3.0$, five of eight α values reach +10%. This discrete transition (not gradual falloff) implies the toroidal embedding has a characteristic coherence distance: correct alternative tokens sit approximately 3 hops from the top-ranked candidate on the torus, matching the harmonic structure of the Tonnetz.

4. TLB and UT loops are complementary at low loop counts, redundant at high loop counts. At 4 loops, TLB provides +10% by correcting near-miss ranking errors the under-refined representations leave unresolved. At 8 loops, all 32 configurations show $\leq 0\%$ delta—the UT has already self-corrected those errors. This convergence confirms both mechanisms target the same error class through different channels.

5. Graceful degradation bounds. Lower α values are universally safe: $\alpha=0.30$ produces $\geq 0\%$ delta across all radii at both 4 and 8 loops. Destructive effects (-5 to -10%) appear only at $\alpha \geq 0.70$ with $r \geq 3.5$, and only at 8 loops for $\alpha=0.60$. The practitioner guidance is clear: moderate α with $r \approx 3$ is safe across all configurations.

6. TLB is a refinement mechanism, not a retrieval mechanism. The prompts TLB corrects (Shakespeare’s plays, digits of π) are near-miss errors where the correct token was close to the decision boundary. The prompts TLB cannot correct (DNA abbreviation, speed of light) represent deep retrieval gaps where the correct answer was never surfaced by the forward pass.

6.4.4 Computational Efficiency Argument

Each UT loop requires a full attention + FFN forward pass ($O(Sd^2 + S^2d)$). TLB requires adding a precomputed bias vector to the logits ($O(V)$ where V is vocabulary size). The cost ratio is:

$$\frac{\text{TLB cost}}{\text{UT loop cost}} \approx \frac{V}{Sd^2 + S^2d} \ll 1 \quad (26)$$

For Ouro-1.4B ($d=2048$, $S=512$, $V=32000$): TLB replaces 4 additional UT loops (loops 5–8) with a single $O(32000)$ operation, achieving *better* results at approximately $2\times$ wall-clock speedup.

6.4.5 Theoretical Interpretation

This result has a natural interpretation within the coherence framework:

- **UT loops** build coherence in *representation space*: hidden states converge through recurrent computation, resolving ambiguity in internal representations.
- **TLB** imposes coherence in *decision space*: the output logit distribution is shaped by geometric locality on the torus, promoting tokens that are topologically consistent with recent context.

These are orthogonal refinement axes. When representation-space coherence is incomplete (few loops), decision-space coherence compensates. When representation-space coherence has converged (many loops), decision-space coherence is redundant.

The optimal operating point is therefore not maximum loops + maximum TLB, but *reduced loops + calibrated TLB*—geometric structure substituting for expensive computation. This aligns with the broader thesis: topology is not merely an aesthetic constraint but a computational resource.

6.5 Experiment 4: Multi-Model TLB Generalization and Ceiling Effect

Experiments 1–3 establish TLB efficacy on Phi-2 (2.7B, attention-level) and Ouro-1.4B (logit-level). A critical question remains: does TLB generalize across model families, and does it have failure modes?

Models tested (all inference-time logit bias, no fine-tuning):

1. **Qwen 2.5-7B-Instruct**: Dense transformer (Alibaba)
2. **OLMo 1.7-7B**: Open-weights research model (AI2)
3. **Mistral 7B**: Grouped-query attention (Mistral AI)
4. **Gemma-2-9B**: Sliding window + global attention (Google)

6.5.1 Positive Results: Models with Moderate Baselines

100-sample factual completion benchmark:

Model	Baseline	TLB	Error Red.	Best (α, r)
Qwen 2.5-7B	95.0%	97.0%	+40.0%	(0.3, 2.0)
OLMo 1.7-7B	87.0%	89.0%	+15.4%	(0.2, 3.0)

TruthfulQA v2 (817 questions, LLM-judged):

Model	Baseline	TLB	Δ
Qwen 2.5-0.5B	16.9%	17.1%	+0.2pp
Qwen 2.5-1.5B	32.2%	32.8%	+0.6pp
Qwen 2.5-7B	75.6%	77.7%	+2.1pp
Mistral 7B	74.4%	77.2%	+2.8pp

Capacity scaling: TLB improvement correlates with model capacity—larger models benefit more (+0.2pp at 0.5B, +2.8pp at 7B). This is consistent with the spectral gap theory: larger models produce more near-correct predictions where geometric bias provides useful signal.

6.5.2 Negative Result: The Ceiling Effect (Gemma-2-9B)

Gemma-2-9B achieves 95% baseline accuracy (19/20) on our factual benchmark. We swept 20 configurations ($\alpha \in \{0.3, 0.4, 0.5, 0.6, 0.7\}$, $r \in \{2.5, 3.0, 3.5, 4.0\}$):

Config	Accuracy	Δ vs Baseline
Baseline	95%	—
$\alpha = 0.3, r = 2.5$	90%	−5pp
$\alpha = 0.3, r = 3.0$	85%	−10pp
$\alpha = 0.4, r = 2.5$	90%	−5pp
$\alpha = 0.4, r = 3.0$	85%	−10pp
$\alpha = 0.3, r = 4.0$	80%	−15pp
$\alpha = 0.4, r = 4.0$	80%	−15pp

All 8 completed configurations degraded performance. The pattern is monotonic: higher α and larger r cause greater degradation. Qualitatively, TLB disrupts already-confident correct predictions—for example, the “Shakespeare wrote” prompt loses the play name under bias, as TLB pushes probability mass toward topologically adjacent but semantically incorrect tokens.

Remark 1 (TLB as Calibration Mechanism). *The multi-model results reveal TLB has regime-dependent behavior:*

- **Weak-model regime** (baseline $\leq 90\%$): *TLB helps. The model has high-entropy predictions where geometric bias provides useful signal, nudging near-miss tokens past decision boundaries.*
- **Strong-model regime** (baseline $\geq 95\%$): *TLB hurts. The model’s logit distribution is peaked; adding bias flattens it, introducing errors on previously correct predictions.*

*This is analogous to regularization in classical ML: beneficial when the model underfits, harmful when it already fits well. The practical implication is that TLB should be deployed with a **confidence gate**—applied only when the model’s top-token probability falls below a threshold.*

This ceiling finding, combined with the Ouro result (Experiment 3), completes the picture: TLB and iterative refinement both target the same error class (near-miss ranking errors). When the model has already resolved those errors—via sufficient UT loops (Experiment 3, 8 loops) or via high baseline accuracy (Experiment 4, Gemma-2)—TLB becomes redundant or harmful.

7 Discussion

7.1 Why Not Implicit Smoothing?

Standard transformer components provide some implicit spectral filtering: LayerNorm suppresses outlier activations, softmax temperature controls attention sharpness, and multi-head averaging smooths individual head outputs. However, none of these impose *topological* constraints—they operate pointwise or via soft weighting, not via manifold structure. They smooth without providing a conserved quantity or spectral gap guarantee. The distinction is between ad-hoc regularization (which helps) and geometric constraint (which bounds).

7.2 Why Hasn’t This Been Done?

Several factors explain why topological constraints haven’t been widely explored:

1. **Scaling laws focus:** Research prioritized parameter count and data size over architectural constraints
2. **Geometric ML is young:** Hamiltonian neural networks [Greydanus et al., 2019] appeared only in 2019; hyperbolic embeddings for LLMs [Patel et al., 2025] emerged in 2025
3. **mHC just published:** The empirical confirmation of instability without constraints appeared January 2026
4. **Interdisciplinary gap:** Music theory (Tonnetz), physics (Hamiltonian), and ML rarely intersect

7.3 Relationship to Other Approaches

Approach	What It Constrains	Limitation
RLHF	Output distribution	No latent geometry
Constitutional AI	Output rules	No latent geometry
Retrieval augmentation	Knowledge access	No reasoning constraint
Chain-of-thought	Output format	No geometric constraint
mHC	Residual mixing	No semantic structure
GAT [Veličković et al., 2018]	Graph attention weights	Learned, no spectral guarantee
Hyperbolic LLMs [Patel et al., 2025]	Hierarchical structure	No local coherence
Geometric Volume [Phillips et al., 2025]	Detects dispersion post-hoc	No prevention mechanism
Tonnetz-ERLHS	Latent geometry	Embedding complexity; m

Note on geometric approaches: Hyperbolic geometry [Patel et al., 2025] excels at representing hierarchical structures (trees, taxonomies) due to its negative curvature, while toroidal geometry (this work) enforces local coherence through periodic boundaries and spectral gaps. These are complementary: hyperbolic for semantic hierarchy, toroidal for inference stability. Future work may combine both.

7.4 Empirical Validation from Geometric Detection

Recent independent work by Phillips et al. [2025] provides empirical support for the geometric perspective on hallucination. Their key finding—that the convex hull volume of attention archetypes correlates with hallucination frequency—is precisely what the spectral gap theorem predicts.

Theorem 13 (Volume-Spectral Duality). *Let $V(t)$ denote the convex hull volume of attention archetype projections at layer t , and let λ_1 be the spectral gap of the constrained attention graph. Then:*

$$V(t) \leq V(0) \cdot e^{-\lambda_1 t} + C_{\text{noise}} \quad (27)$$

where C_{noise} bounds irreducible noise from finite sampling. That is, geometric dispersion—the mechanism identified by Phillips et al. [2025] as correlating with hallucination—is exponentially bounded by topological constraints.

Proof sketch. The convex hull volume of projected archetypes measures the spread of attention mass across the latent space. Under spectral filtering with gap λ_1 , high-frequency modes (which expand the convex hull) decay as $e^{-\lambda_1 t}$, while low-frequency modes (which preserve the hull’s centroid) are invariant. Therefore, volume expansion is bounded by the decay rate of incoherent modes. $\square$

This duality reveals the complementary nature of detection and prevention:

- **Detection** [Phillips et al., 2025]: Measure $V(t)$; high values indicate likely hallucination
- **Prevention** (this work): Enforce $\lambda_1 > 0$; hallucination probability decreases exponentially

The theoretical contribution here is explaining *why* geometric dispersion correlates with hallucination: it reflects violation of the spectral gap bound.

7.5 Empirical Support from Zigzag Persistence

Concurrent work by Gardinazzi et al. [2025] provides additional empirical validation using zigzag persistence from topological data analysis. They track the birth and death of p -dimensional holes (connected components, loops, voids) across transformer layers and identify four distinct phases:

1. **Early layers:** Rapid rearrangement, many short-lived topological features
2. **Middle layers:** Stable phase where topological features have highest persistence
3. **Middle-to-late:** Refinement with few short-lived adjustments
4. **Final layers:** New rearrangements preparing for output

Their key finding—that topological features in middle layers have highest inter-layer persistence—is precisely what the spectral gap theorem predicts: coherent modes (low-frequency) propagate without attenuation, while incoherent modes (high-frequency) decay as $e^{-\lambda_1 t}$. Their inter-layer persistence metric $\tilde{Z}_p$ empirically measures the effect we characterize theoretically.

This establishes a three-way convergence from independent research programs:

Approach	Method	Finding
Phillips et al.	Convex hull volume	Geometric dispersion $\uparrow \Rightarrow$ hallucination $\uparrow$
Gardinazzi et al.	Zigzag persistence	Topological instability $\uparrow \Rightarrow$ layer importance $\uparrow$
This work	Spectral gap	Enforcing $\lambda_1 > 0$ bounds both

The convergence suggests that geometric/topological coherence is not an artifact of any single methodology but a fundamental property of stable transformer dynamics.

7.6 Empirical Support from Mechanistic Interpretability

Recent interpretability research from Anthropic provides direct empirical validation of the manifold-structure premise underlying our framework.

Gurnee et al. [2025] mechanistically investigate how Claude 3.5 Haiku performs linebreaking in fixed-width text and find that “character counts are represented on low-dimensional curved manifolds discretized by sparse feature families, analogous to biological place cells.” Accurate predictions emerge from “a sequence of geometric transformations” on these manifolds. This finding establishes the manifold structure M posited by ERLHS not as a theoretical idealization but as the empirical reality of trained transformer computation. The geometric-transformation framing maps directly to the harmonic-flow picture developed in Section 3: coherent reasoning is geometric flow on a learned manifold; incoherent generation is off-manifold drift.

Kamath et al. [2025] extend attribution graphs to explain attention computation via “QK attributions”—bilinear functions of feature activations on

query and key positions—decomposing attention patterns into interpretable feature-interaction graphs. The resulting graph structure provides the mechanistic correlate of the architecture-dependent effects we observe in Experiment 4. Qwen 2.5-7B full attention, Mistral 7B grouped-query attention, and Gemma-2-9B sliding-window-plus-global attention induce qualitatively different feature-interaction topologies, predicting different sensitivities to inference-time geometric intervention. The ceiling effect we observe (TLB helps moderate-baseline models but degrades near-saturated models) is consistent with this prediction: when the underlying feature-interaction graph already produces well-separated logits, additional spectral filtering becomes redundant or harmful.

Lindsey [2025] demonstrates that more capable language models exhibit “some functional awareness of their own internal states,” with Claude Opus 4.1 and 4 outperforming earlier models on activation-injection probes. Critically, the paper emphasizes that “the abilities we observe are highly unreliable; failures of introspection remain the norm,” and that self-reports frequently include details “whose accuracy we cannot verify, and which may be embellished or confabulated.” This caveat directly supports the design choice of an *external* coherence verifier C in ERLHS [Cormier, 2025a]: emergent introspection is a capability, not a guarantee, and verification of latent-state coherence cannot rely on the system being verified. The Inference-Time Coherence Verification protocol of Section 5.4 is precisely the external counterpart to Lindsey’s empirical introspection.

Combined with Phillips et al., Gardinazzi et al., and Zenil, this extends the convergence to four independent research programs:

Approach	Method	Finding
Phillips et al.	Convex hull volume	Geometric dispersion $\uparrow \Rightarrow$ h
Gardinazzi et al.	Zigzag persistence	Topological instability $\uparrow \Rightarrow$
Anthropic (Gurnee, Kamath, Lindsey)	Mechanistic interpretability	Manifold structure is empiri
This work	Spectral gap	Enforcing $\lambda_1 > 0$ bounds dr

That a methodology spanning persistent homology, convex geometry, mechanistic interpretability, and spectral graph theory converges on the same conclusion—that stable transformer reasoning requires geometric/topological constraints that classical training does not provide—suggests this is not artifact but principle.

7.7 Information-Theoretic Foundations: Why Constraints Are Necessary

Recent work by Zenil [2026] provides the information-theoretic foundations explaining *why* topological constraints are necessary, not merely helpful.

Zenil proves two fundamental failure modes of unconstrained self-referential training:

1. **Entropy Decay:** Under self-training with finite samples, model entropy forms a supermartingale—it can only decrease. The distribution inevitably collapses to a degenerate fixed point.
2. **Variance Amplification:** Without external grounding, the model’s representation of truth drifts as a random walk, bounded only by support diameter.

Critically, Zenil shows these failures are consequences of the **Data Processing Inequality (DPI)**:

$$I(M; Q_{t+1}) \leq I(M; Q_t) \quad (28)$$

where M is the true generating mechanism. Statistical learning cannot increase information about the underlying mechanism—it can only contract.

The escape route Zenil identifies is **structural constraints** that operate in program/mechanism space rather than distribution space. His symbolic projection operator Π_S reduces hypothesis space volume by enforcing invariants, achieving contraction factor $\sigma < 1$ that statistical updates cannot.

Our contribution is the constructive realization: the Tonnetz topology with spectral gap $\lambda_1 > 0$ is a specific instance of Zenil’s abstract Π_S . The spectral gap enforces precisely the structural constraint needed to escape DPI bounds:

Zenil’s Abstract Framework	Our Concrete Realization
Symbolic projection Π_S	Toroidal attention mask M_{Tonnetz}
Contraction factor $\sigma < 1$	Spectral gap $\lambda_1 = \Theta(1)$
Algorithmic complexity bound $K(p) \leq L$	Geodesic distance bound $d_{\text{Tonnetz}} \leq r$
Escape from entropy decay	40% drift reduction (experimental)

This connection establishes a complete theoretical chain: Zenil proves constraints are *necessary*; we provide a constraint that is *sufficient*; experiments confirm it *works*.

7.8 Limitations

1. **Embedding challenge:** Mapping tokens to Tonnetz positions requires learning or heuristics
2. **Computational overhead:** Topological constraints add operations per layer

3. **Expressivity tradeoff:** Constraints may limit model capacity for some tasks
4. **Ceiling effect:** TLB degrades strong models (Gemma-2-9B at 95% baseline). Practical deployment requires a confidence-gated strategy that applies TLB only on uncertain predictions
5. **Hyperparameter sensitivity:** Optimal (α, r) varies across model families (Qwen: $\alpha=0.3, r=2.0$ on the 100-sample factual benchmark; OLMo: $\alpha=0.2, r=3.0$ on the same benchmark; Ouro: $\alpha=0.5, r=3.0$). **Mistral 7B hyperparameter gap (added in v1.2):** the (α, r) used to obtain the +2.8pp Mistral 7B result on TruthfulQA v2 (Section 6.4) is not currently recovered in this draft; the entry has been omitted from the table above pending log retrieval. Independent reproductions should sweep $\alpha \in \{0.2, 0.3, 0.4, 0.5\}$ and $r \in \{2.0, 2.5, 3.0\}$ until the original value is recovered or independently calibrated. Automated calibration from a small held-out set remains future work
6. **LLM-judge implementation details (added in v1.2):** The TruthfulQA v2 evaluation in Section 6.4 (Mistral, Qwen 2.5-7B at full 817 questions) is described as “LLM-judged.” The specific judge model, scoring prompt, and binary correct/incorrect determination procedure are not currently documented in this draft and should be added to the supplement before the next revision. A follow-up multi-benchmark validation by the present authors observed that lenient keyword-overlap scoring inflates baselines by ~ 10 pp on TruthfulQA generation and compresses signal at saturation, motivating explicit documentation of the scoring methodology for reproducibility
7. **Benchmark limitations:** The 20-prompt factual QA benchmark, while validated across 64 hyperparameter configurations on Ouro showing a robust improvement plateau, lacks the statistical power of larger benchmarks. TruthfulQA MC1 evaluation (200 questions) showed zero effect, likely because MC1 scoring uses mean-of-last-token logits (a discrimination task) rather than autoregressive generation—where TLB’s logit bias operates
8. **Regime boundary — training-time SSL:** The empirical results above are obtained by applying toroidal structure at *inference time* (TLB) or in *attention masking* (Tonnetz mask) on already-trained or trained-from-scratch language models. A separate study by the present authors [Cormier, 2026] applies an analogous toroidal regularizer at *training time* to a JEPA-style self-supervised representation learner (ResNet-18 encoder, CIFAR-10/100). That study finds three distinct failure modes: encoder collapse (when VICReg’s covariance term is removed), healthy-encoder fragmentation (when the encoder

is properly regularized but supervised contrastive cannot pull diverse features onto a 2-manifold), and class-misaligned representational diversity (CIFAR-100). The toroidal hypothesis is partially correct in direction — topological fragmentation reduces $\sim 5\times$ from 10 classes to 100 classes — but the accuracy-topology tradeoff is not broken in any tested configuration. This delineates the present paper’s claims: *toroidal logit biasing on trained LLMs is supported empirically; toroidal regularization during representation learning at small scale is not*. The two regimes likely require different geometric strategies, and identifying the boundary is an open problem

8 Conclusion

We have established:

1. **Residual geometry determines reasoning stability:** Unconstrained latent dynamics lack the conserved quantities necessary for bounded inference
2. **mHC empirically confirms the principle:** Doubly-stochastic constraints are necessary for stable training at scale
3. **A formal hierarchy exists:** mHC (Birkhoff) $\subset$ ERLHS (Hamiltonian) $\subset$ Karmonic (Toroidal + Spectral)
4. **Tonnetz is a constructive existence proof:** Toroidal topology with constant spectral gap demonstrates richer structure is achievable
5. **Hallucination is a consequence, not a cause:** Reduced drift follows from geometric constraints, not alignment heuristics
6. **Experimental validation confirms the theory:** Toroidal attention achieves 40% drift reduction on synthetic sequences and +19.5% relative improvement on TruthfulQA at 2.7B scale. Toroidal logit bias on Ouro-1.4B (Universal Transformer) at 4 loops outperforms 8-loop baseline, demonstrating that geometric coherence substitutes for computational depth
7. **TLB generalizes across architectures with a regime boundary:** Multi-model validation (Qwen, OLMo, Mistral, Gemma-2) confirms TLB improves models with moderate baselines ($\leq 90\%$) but degrades models already near ceiling ($\geq 95\%$)—consistent with its theoretical role as a calibration mechanism targeting near-miss ranking errors
8. **Detection and prevention are dual:** The Volume-Spectral Duality theorem explains why geometric dispersion metrics [Phillips et

al., 2025] correlate with hallucination—they measure violation of the spectral gap bound

The central contribution is a sufficient condition:

Geometric constraints provide one principled path to coherent artificial intelligence—not the only path, but a formally grounded one with empirical validation.

Future foundation models should be designed with topological constraints from the start. The mathematical framework exists. Experimental validation demonstrates efficacy.

Notably, toroidal topology provides coherence guarantees beyond the LLM domain. The QuantumTimeSandwich platform [Cormier, 2024] demonstrates the same principle in quantum computing: a Rust-native quantum SDK implementing BB84 key distribution, toric code error correction, and Tonnets coherence simulation—all built on the same toroidal geometry that constrains drift in our LLM experiments. This cross-domain applicability reinforces the central thesis: the spectral gap is a *geometric* property, not a domain-specific one.

An interactive demo is available at <https://huggingface.co/spaces/paraxiom/topological-coherence>.

Author Contributions

S.C. conceived the theoretical framework, designed and executed all experiments, and wrote the manuscript.

Funding

This research received no external funding.

Data Availability Statement

All code, configurations, and experimental results are publicly available at <https://github.com/Paraxiom/topological-coherence>. An interactive demo is available at <https://huggingface.co/spaces/paraxiom/topological-coherence>.

Conflicts of Interest

The author declares no conflict of interest.

References

- Cormier, S. (2025). ERLHS: A Hamiltonian Framework for Coherence-Preserving Machine Intelligence. *Zenodo*. DOI: 10.5281/zenodo.17928909
- Cormier, S. (2025). Karmonic Mesh: Spectral Consensus on Toroidal Manifolds. *Zenodo*. DOI: 10.5281/zenodo.17928991
- Greydanus, S., Dzamba, M., and Yosinski, J. (2019). Hamiltonian neural networks. In *NeurIPS*.
- Ji, Z., et al. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Xie, Z., et al. (2026). Manifold-Constrained Hyper-Connections for Stable Deep Learning. DeepSeek Technical Report, January 2026.
- Zhu, D., et al. (2024). Hyper-Connections. *arXiv preprint arXiv:2409.19606*.
- Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. (2013). The emerging field of signal processing on graphs. *IEEE Signal Processing Magazine*, 30(3):83–98.
- Phillips, J., Khan, A., and Sheridan, M. (2025). Geometric Uncertainty for Detecting and Correcting Hallucinations in Large Language Models. Oxford University, preprint.
- Gardinazzi, Y., Viswanathan, K., Panerai, G., Ansuini, A., Cazzaniga, A., and Biagetti, M. (2025). Persistent Topological Features in Large Language Models. *arXiv preprint arXiv:2410.11042v3*.
- Zenil, H. (2026). On the Limits of Self-Improving in LLMs and Why AGI, ASI and the Singularity Are Not Near Without Symbolic Model Synthesis. *arXiv preprint arXiv:2601.05280v1*.
- Patel, S., et al. (2025). Hyperbolic Large Language Models. *arXiv preprint arXiv:2509.05757*.
- Sáez de Ocáriz Borde, H., Arroyo, Á., Morales López, I., Posner, I., and Dong, X. (2023). Neural Latent Geometry Search: Product Manifold Inference via Gromov-Hausdorff-Informed Bayesian Optimization. In *Advances in Neural Information Processing Systems (NeurIPS 2023)*.
- Gu, A., Sala, F., Gunel, B., and Ré, C. (2018). Learning mixed-curvature representations in product spaces. In *International Conference on Learning Representations (ICLR)*.

- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. (2018). Graph Attention Networks. In *International Conference on Learning Representations (ICLR)*.
- Cormier, S. (2024). QuantumTimeSandwich: A Rust Quantum Computing Platform with Toroidal Coherence Simulation. <https://github.com/Paraxiom/QuantumTimeSandwich>
- ByteDance Research (2025). Ouro: A Universal Transformer with Iterative Refinement. <https://huggingface.co/ByteDance/Ouro-1.4B>
- Cormier, S. (2026). Toroidal Logit Bias for Hallucination Reduction in Large Language Models. *Zenodo*. DOI: 10.5281/zenodo.18516477
- Amari, S. (1998). Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10(2):251–276.
- Bonnabel, S. (2013). Stochastic Gradient Descent on Riemannian Manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229.
- Gurnee, W., Ameisen, E., Kauvar, I., Tarnag, J., Pearce, A., Olah, C., and Batson, J. (2025). When Models Manipulate Manifolds: The Geometry of a Counting Task. *Transformer Circuits Thread*, October 2025. <https://transformer-circuits.pub/2025/linebreaks/index.html>
- Kamath, H., Ameisen, E., Kauvar, I., Luger, R., Gurnee, W., Pearce, A., Zimmerman, S., Batson, J., Conerly, T., Olah, C., and Lindsey, J. (2025). Tracing Attention Computation Through Feature Interactions. *Transformer Circuits Thread*, July 2025. <https://transformer-circuits.pub/2025/attention-qk/index.html>
- Lindsey, J. (2025). Emergent Introspective Awareness in Large Language Models. *Transformer Circuits Thread*, October 29, 2025. <https://transformer-circuits.pub/2025/introspection/index.html>
- Cormier, S. (2026). Three Failure Modes of Toroidal Regularization in Self-Supervised Learning: A Decomposition Study Across Two Datasets. In preparation. Code and results: <https://github.com/Paraxiom/jepa-torus>