

Preference Fields on Semantic Manifolds

Directional Constraints, Spectral Admissibility, and RGG Structure




Directional Constraints: Modeling evolving user preference as a time-dependent unit vector, not a static query.



Spectral Admissibility: Removing anisotropic centroid bias to unlock discriminative geometry in embedding spaces.



Random Geometric Graph Structure: A dimensionless order parameter unifying field dynamics and topology into a single phase transition.

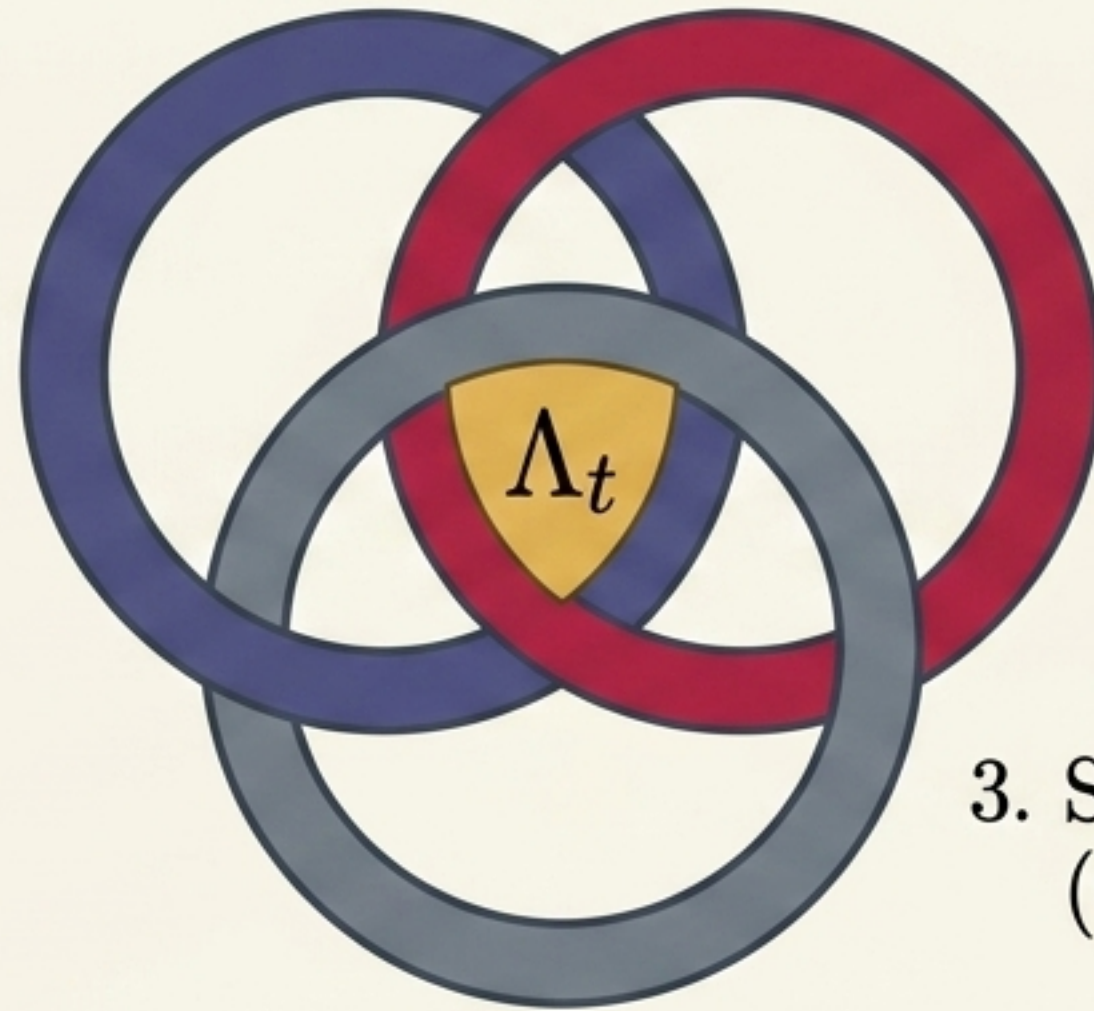


Navigation is a Field Evaluation, Not a Retrieval Operation

Classical Navigation	Preference Field Navigation
<p>Nature of the Query Fixed and given in advance.</p>	<p>Nature of the Query Constructed and evolving through interaction with the manifold.</p>
<p>Action Mechanism Ranks a discrete permutation of corpus items.</p>	<p>Action Mechanism Evaluates the trace of a continuous linear functional $\langle \Phi_t, \mathbf{v} \rangle$ on S^{d-1}.</p>
<p>Space Structure Arbitrary metric space.</p>	<p>Semantic manifold where inner products correlate directly with human-perceived similarity.</p>
<p>Success Condition Reaching the target item.</p>	<p>Success Condition Converging on a semantically coherent admissible neighborhood $A_t(\theta)$.</p>

The Three Interlocking Structures of Field Navigation

1. The Preference Field
(Direction)



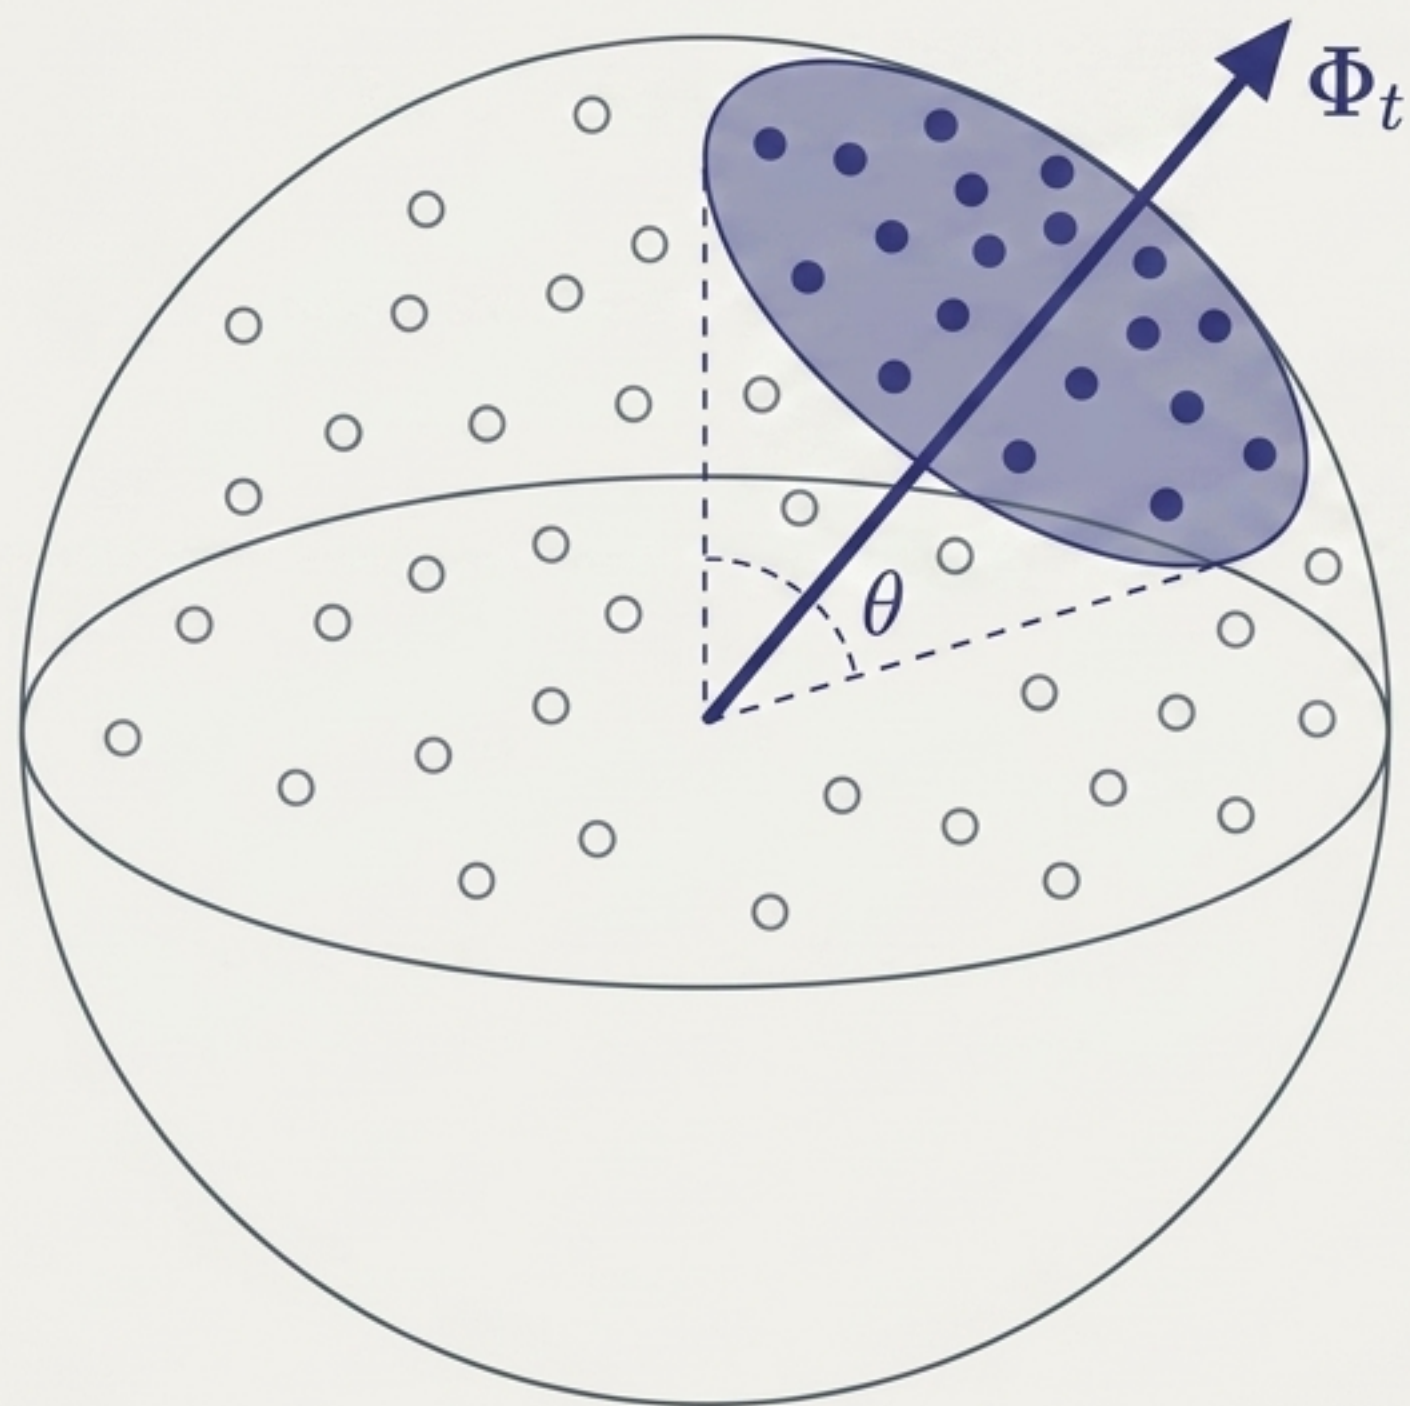
2. Random Geometric Graphs
(Topology)

3. Spectral Admissibility
(Geometry)

The theory progresses from concrete objects to their synthesis:

1. Directional Constraints: How ratings update the field and induce admissible neighborhoods.
2. Topology: How the ambient geometry governs the formation of coherent preference neighborhoods via clique suppression.
3. Geometry: How spectral filtering enforces admissibility and sharpens the field.
4. Synthesis: The Λ_t Order Parameter dictating when geometry becomes detectable.

Defining the Field and its Admissible Neighborhood



The Preference Field at Time t

The preference field at time t is a unit vector $\Phi_t \in \mathbb{S}^{d-1}$. It induces a scalar alignment function $f_t(\mathbf{v}) = \langle \Phi_t, \mathbf{v} \rangle$.

The Discrete Preference Field

Evaluated over the corpus manifold M , it is not merely a ranking. It is a direction in \mathbb{R}^d that generalizes continuously to any item.

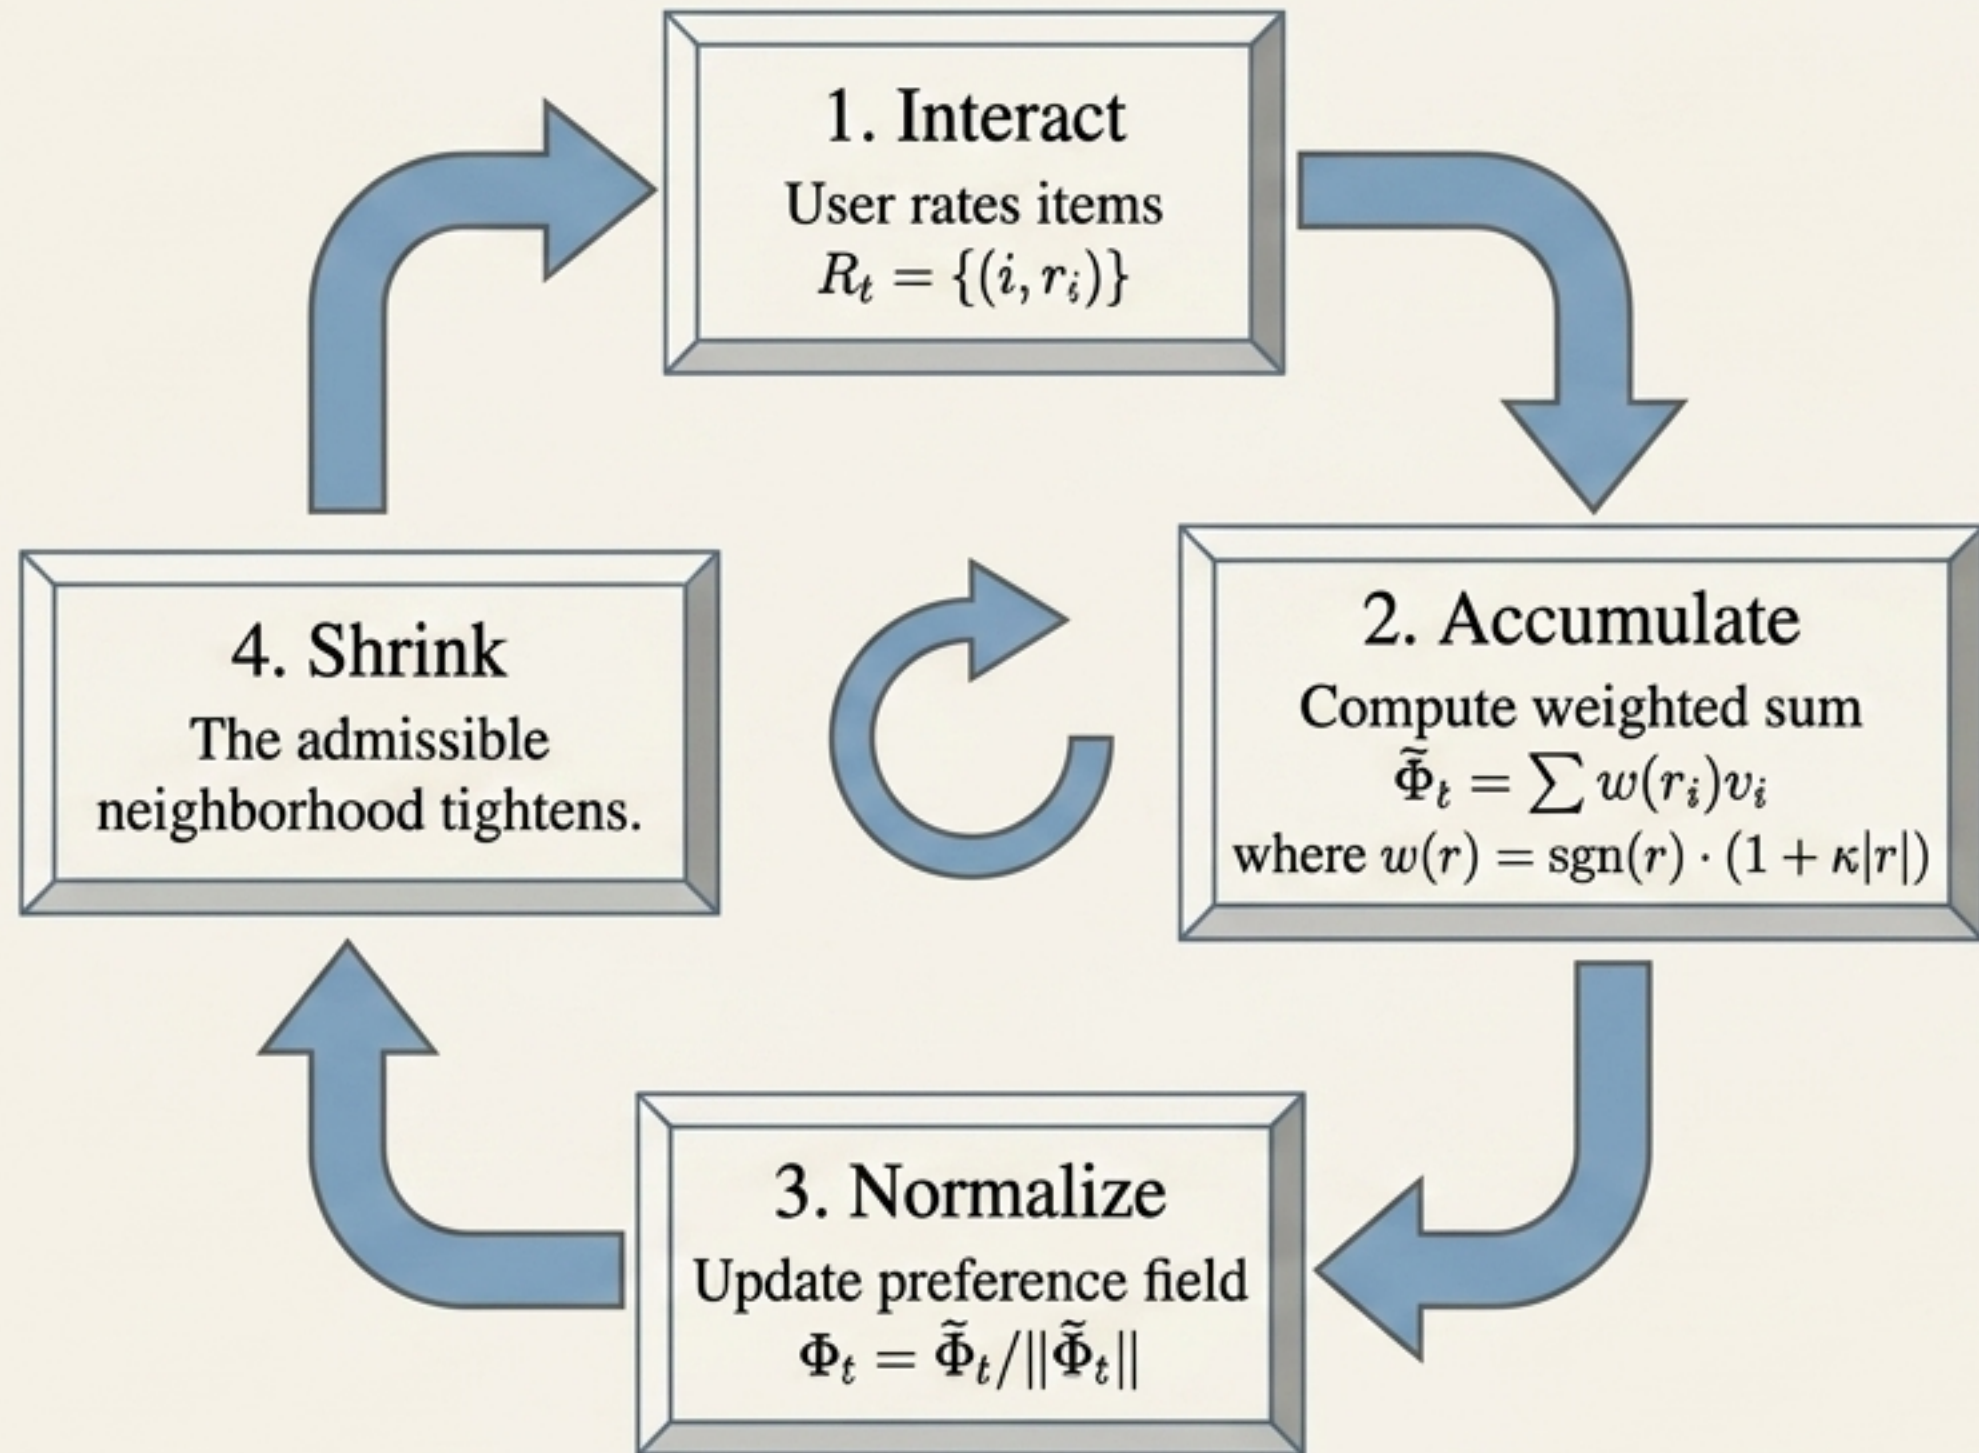
The Admissible Neighborhood $A_t(\theta)$

The subset $\{i : f_t(\mathbf{v}_i) \geq \theta\}$. Only items tightly aligned with the field remain admissible.

Softmax Induction

In practice, θ is induced by a softmax sampling distribution $P_t(i) \propto \exp(s_t(i)/T)$, granting items in $A_t(\theta)$ disproportionate probability.

Monotone Convergence Under Preference Accumulation



Proposition 2.1 (Geodesic Convexity)

How do we guarantee the field actually converges toward the user's intent?


If positively rated items lie within a geodesic ball $B_\rho(\Phi^*)$ of radius $\rho < \pi/2$, adding items moves the sum further into the half-space $\langle \mathbf{v}, \Phi^* \rangle > 0$.

Normalization preserves the sign.

The inner product $\langle \Phi_t, \Phi^* \rangle$ is strictly non-decreasing.


The Admissible Neighborhood is a Spherical RGG

Erdős–Rényi Random Graph ($G_{n,p}$)



Purely chaotic, uniform edge connections, no spatial logic.

Spherical Random Geometric Graph



Dense, localized geometric clusters dictated by ambient space.

Proposition 3.1 (RGG Identification)

Items inside the admissible neighborhood are not just a list—they form a *threshold* graph $G_t(\theta)$.

If the corpus embeddings are approximately *uniform* on the spherical cap $\{v \in \mathbb{S}^{d-1} : \langle v, \Phi_t \rangle \geq \theta\}$, then $G_t(\theta)$ is distributed as a spherical RGG on $|A_t(\theta)|$ vertices.

The geometry of the embedding space fundamentally dictates the topology of the selected neighborhood.


The Triangle Penalty and Exponential Clique Suppression

Self-Reinforcing Alignment

Blue cliques form with higher probability than $G_{n,p}$ predicts. Items aligned with Φ_t tend to be mutually aligned.

$$\delta_b = 1 + a^3(1 - p)^3 / \sqrt{d}$$

Blue Clique




Exponential Suppression

Red (inadmissible) cliques form with exponentially lower probability. Large, mutually distant inadmissible sets are geometrically suppressed by a penalty scaling as $\binom{\ell}{3}$.

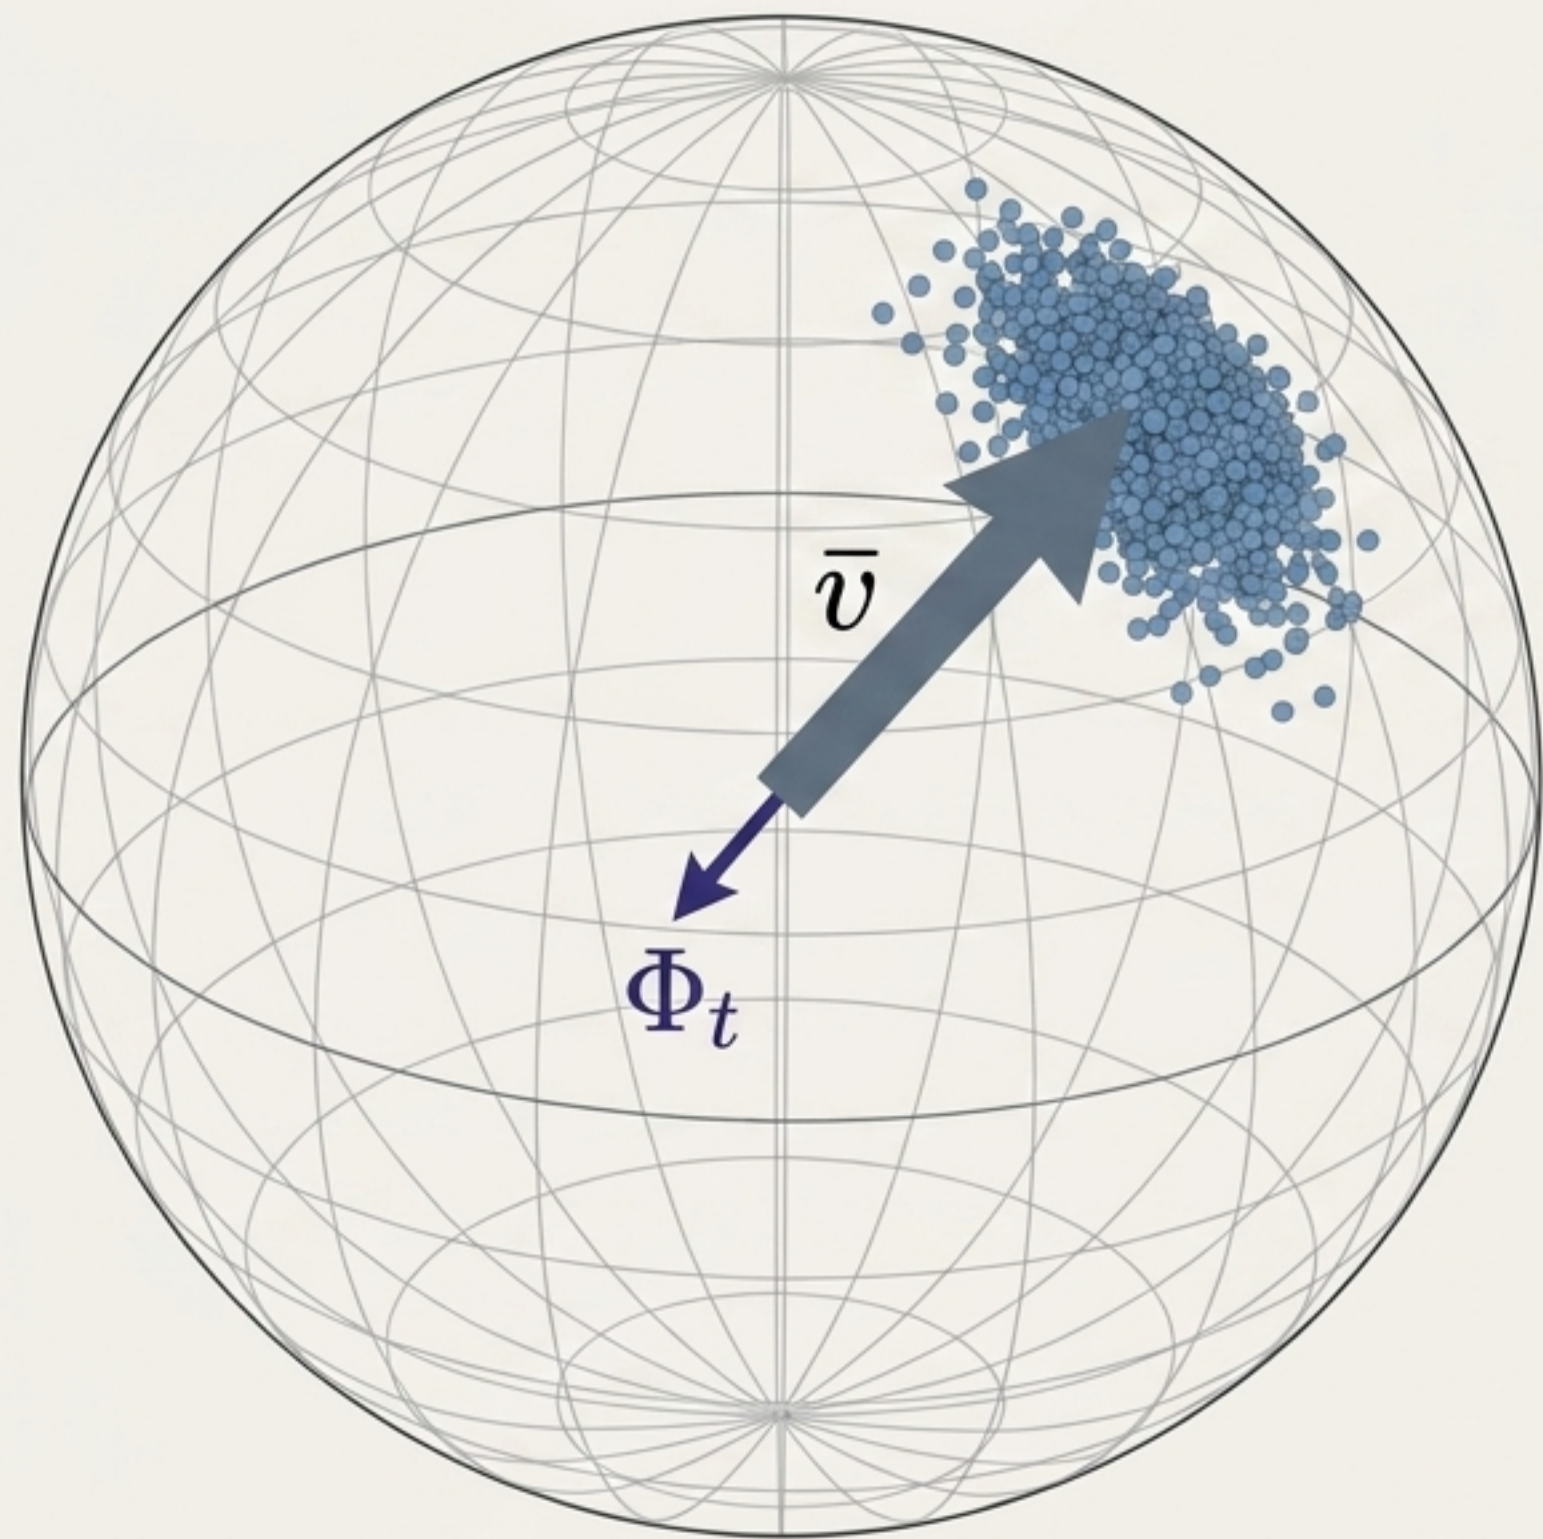
$$\delta_r = 1 - a^3 p^3 / \sqrt{d}$$

Red Clique



Corollary 3.1: The field clique number $\omega_t(\theta)$ is bounded solely by corpus geometry, not session length.

The Obstruction of Anisotropy in Pretrained Models



Pretrained representations (like CLIP) are anisotropic. Embeddings concentrate in a narrow cone with a massive centroid bias \bar{v} .

The Weak-Field Distortion


Early in navigation, the alignment function is dominated by the centroid:

$$f_t(v_i) = \langle \Phi_t, v_i \rangle \approx \langle \bar{v}, v_i \rangle.$$

Semantic Blindness

Because the centroid dominates the inner product, the field is approximately constant across all items. The field cannot discriminate until Φ_t moves far enough from the centroid to produce meaningful variance. The centroid acts as an inadmissible subspace.

Admissibility Enforcement via Spectral Filtering



Proposition 4.1 (EmbedFilter)

Applying the Bulk Spectrum Transformation $\Phi_{\tau}(\mathbf{v}) = \mathbf{V}[\ell_{\tau} : \mathbf{r}_{\tau}]^{\top} \mathbf{v}$ operates as an admissibility enforcer.

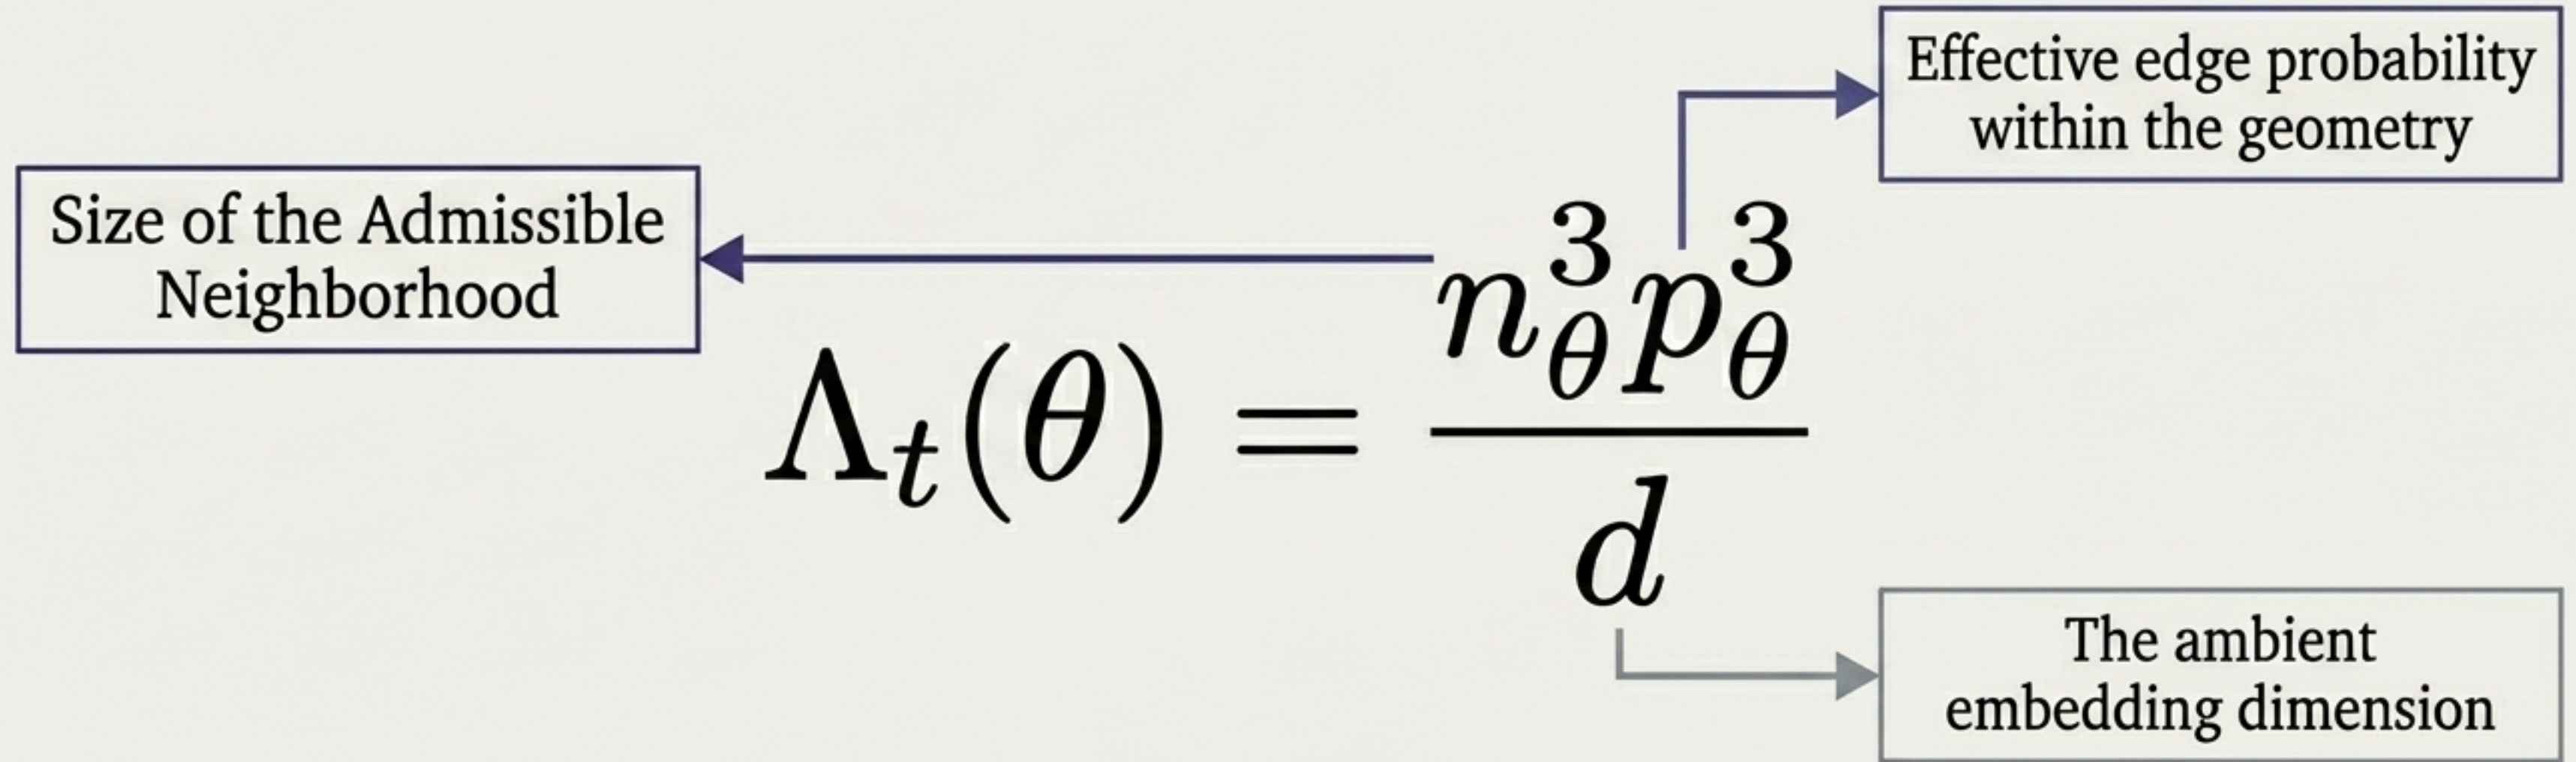
Centroid Elimination

If $\bar{\mathbf{v}} \in V_{\text{edge}}$, then its projection becomes $\mathbf{0}$. Meaningful alignment differences emerge instantly, rather than being masked by the global centroid attractor.

Variance Amplification


By shifting the filtered field \tilde{f}_t to operate exclusively in the bulk subspace \mathbf{V}_{bulk} , we strictly increase discriminative variance.

The Unifying Order Parameter: The Regime Indicator



The signed triangle count distinguishes an RGG from random noise. The signal-to-noise ratio of this distinction yields Λ_t . Like the Reynolds number in fluid dynamics or R_0 in epidemiology, Λ_t is a dimensionless order parameter. It unifies field dynamics, spectral admissibility, and clique structure to dictate the state of the entire geometric system.

The Phase Transition of Field Geometry

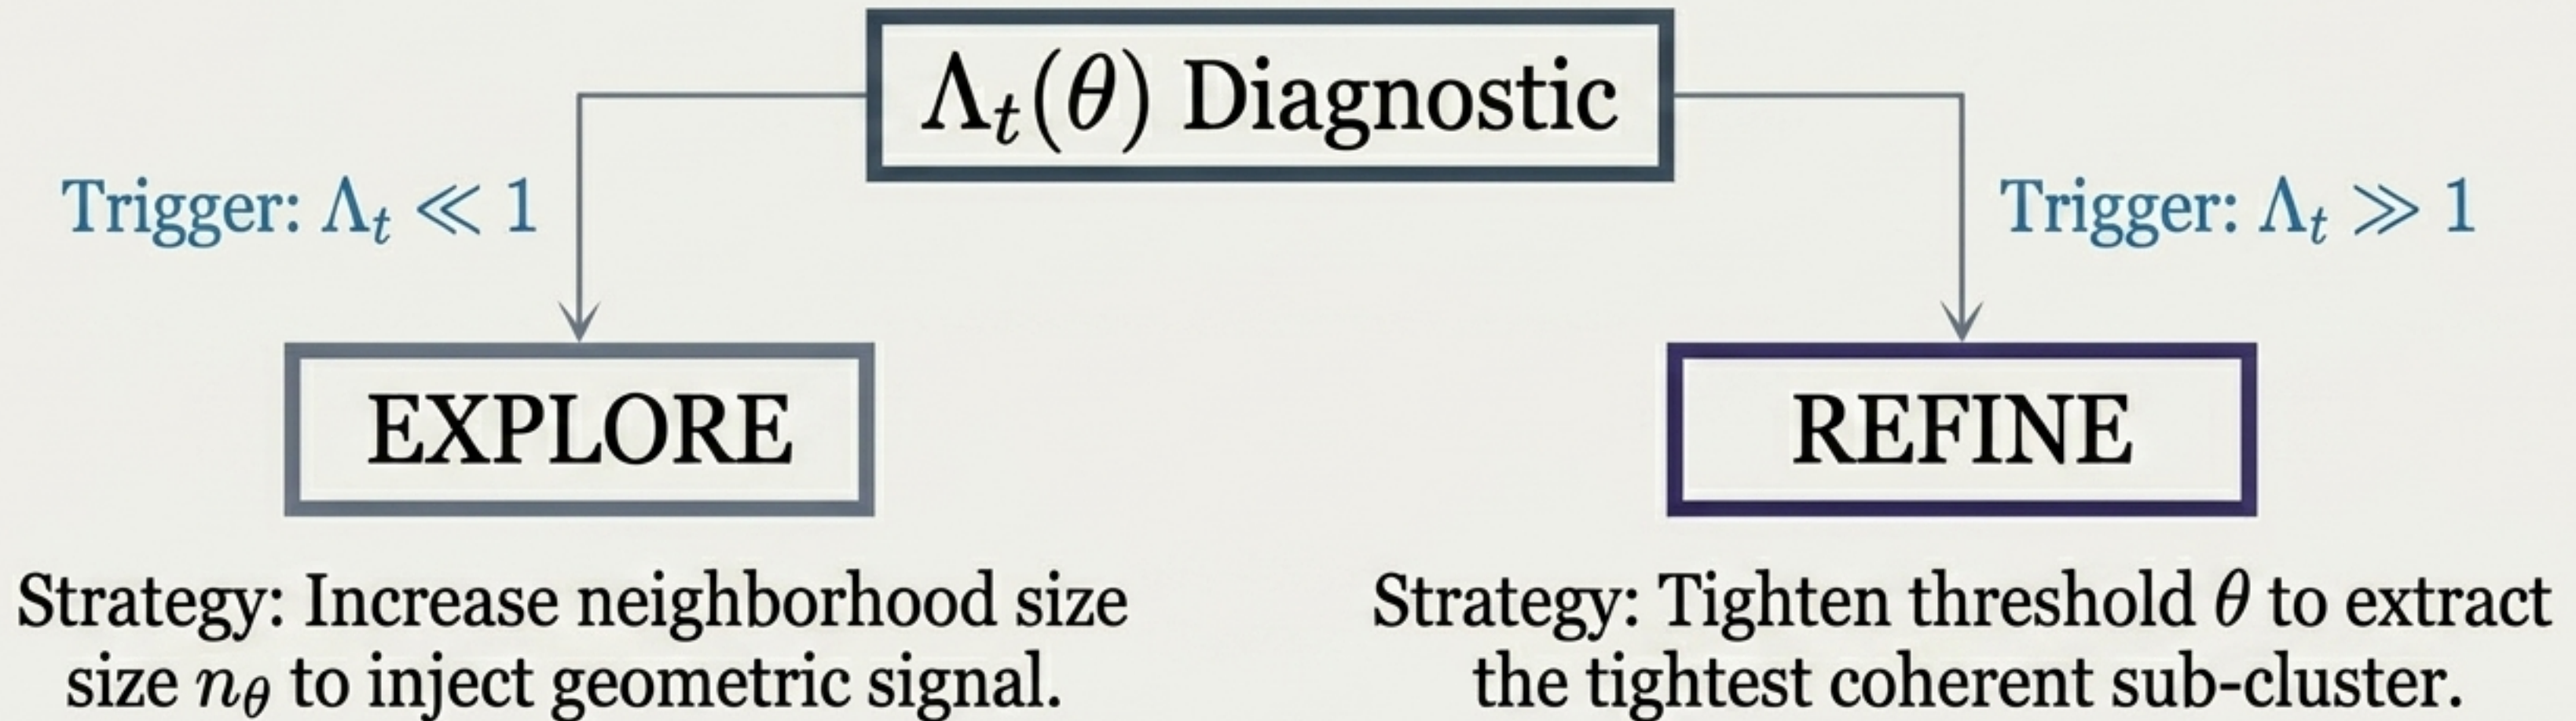


Theorem 5.1

Below the critical threshold, no statistical test can distinguish the threshold graph from random noise. The preference field carries no recoverable geometric information.

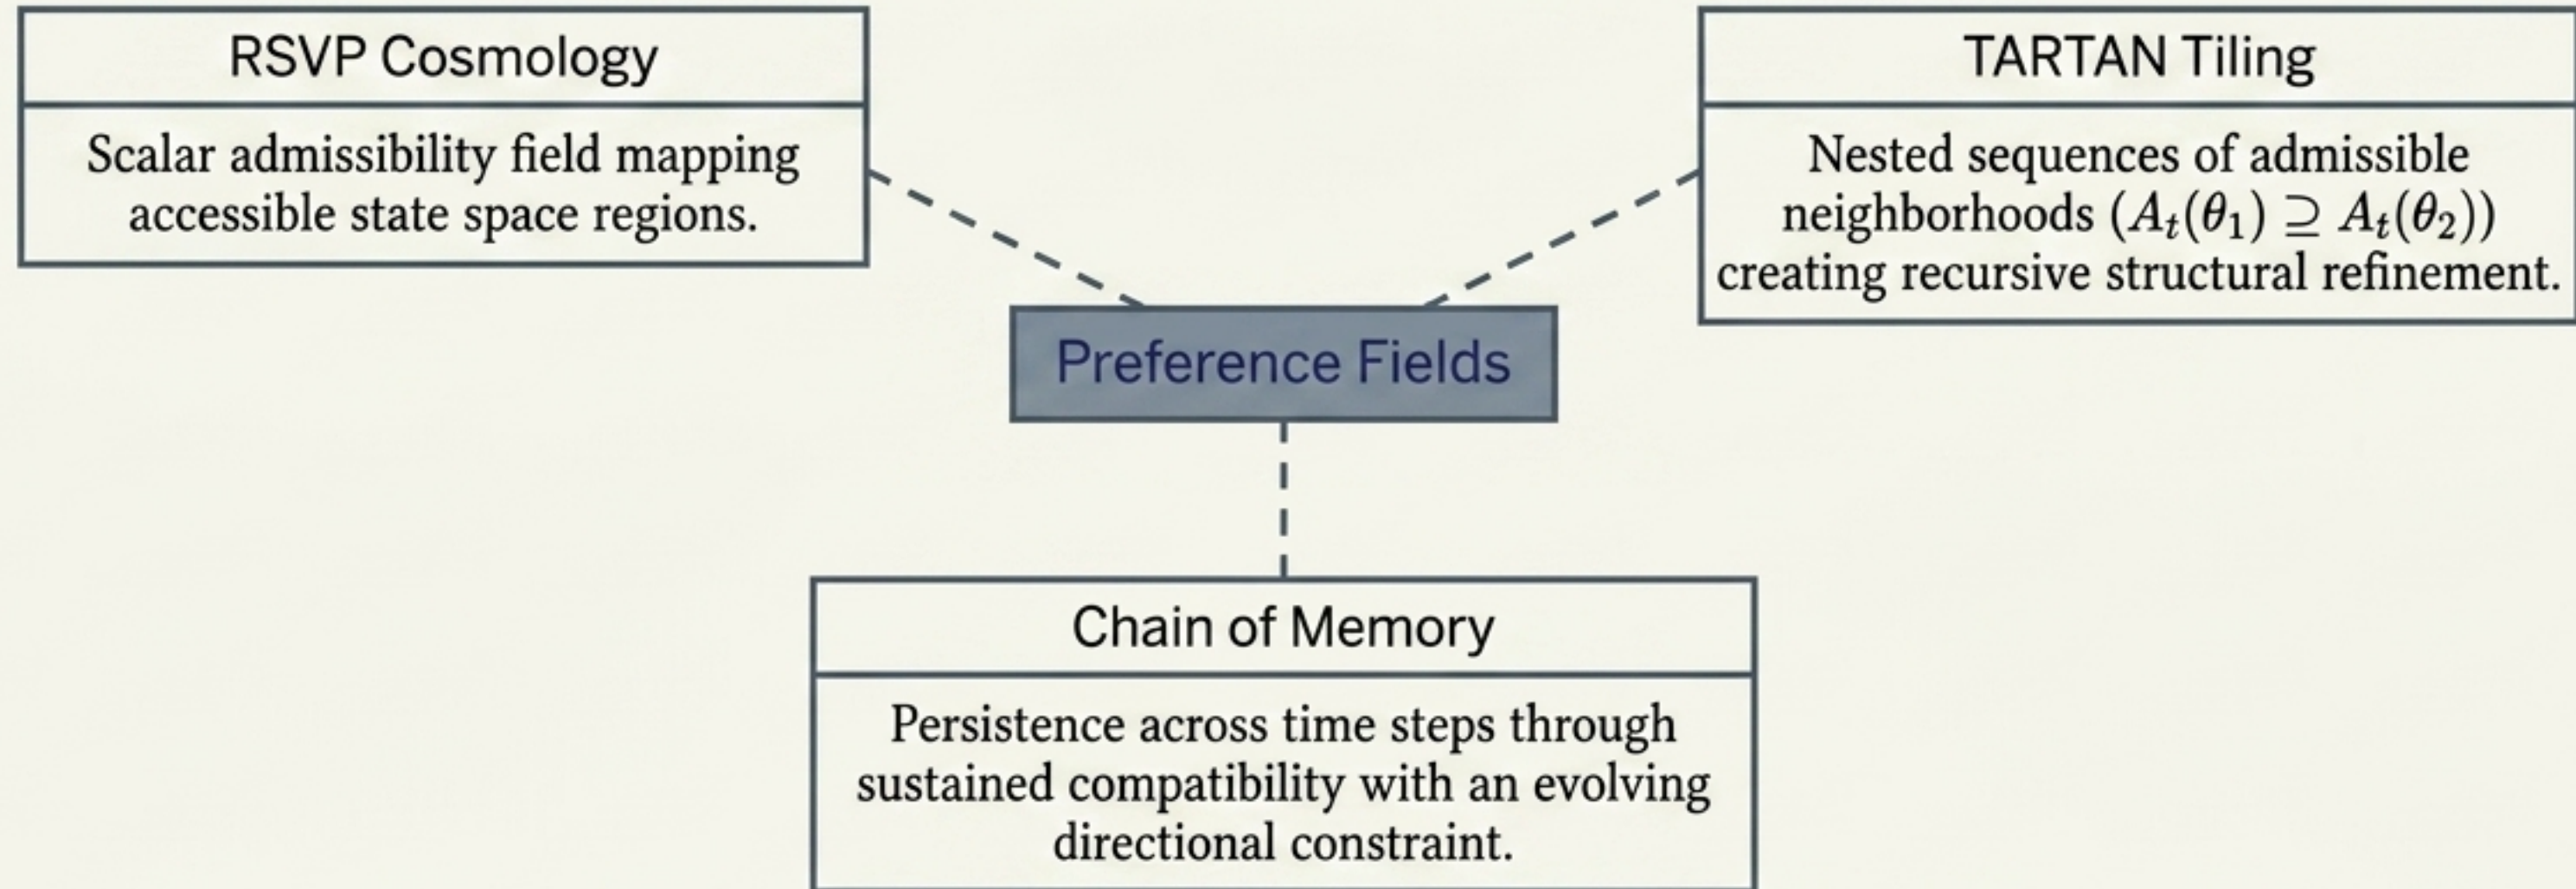
Above the threshold, triangles close at a rate exceeding random prediction. The preference field resolves actual semantic structure with statistical power approaching one.

Strategic Implications for System Design



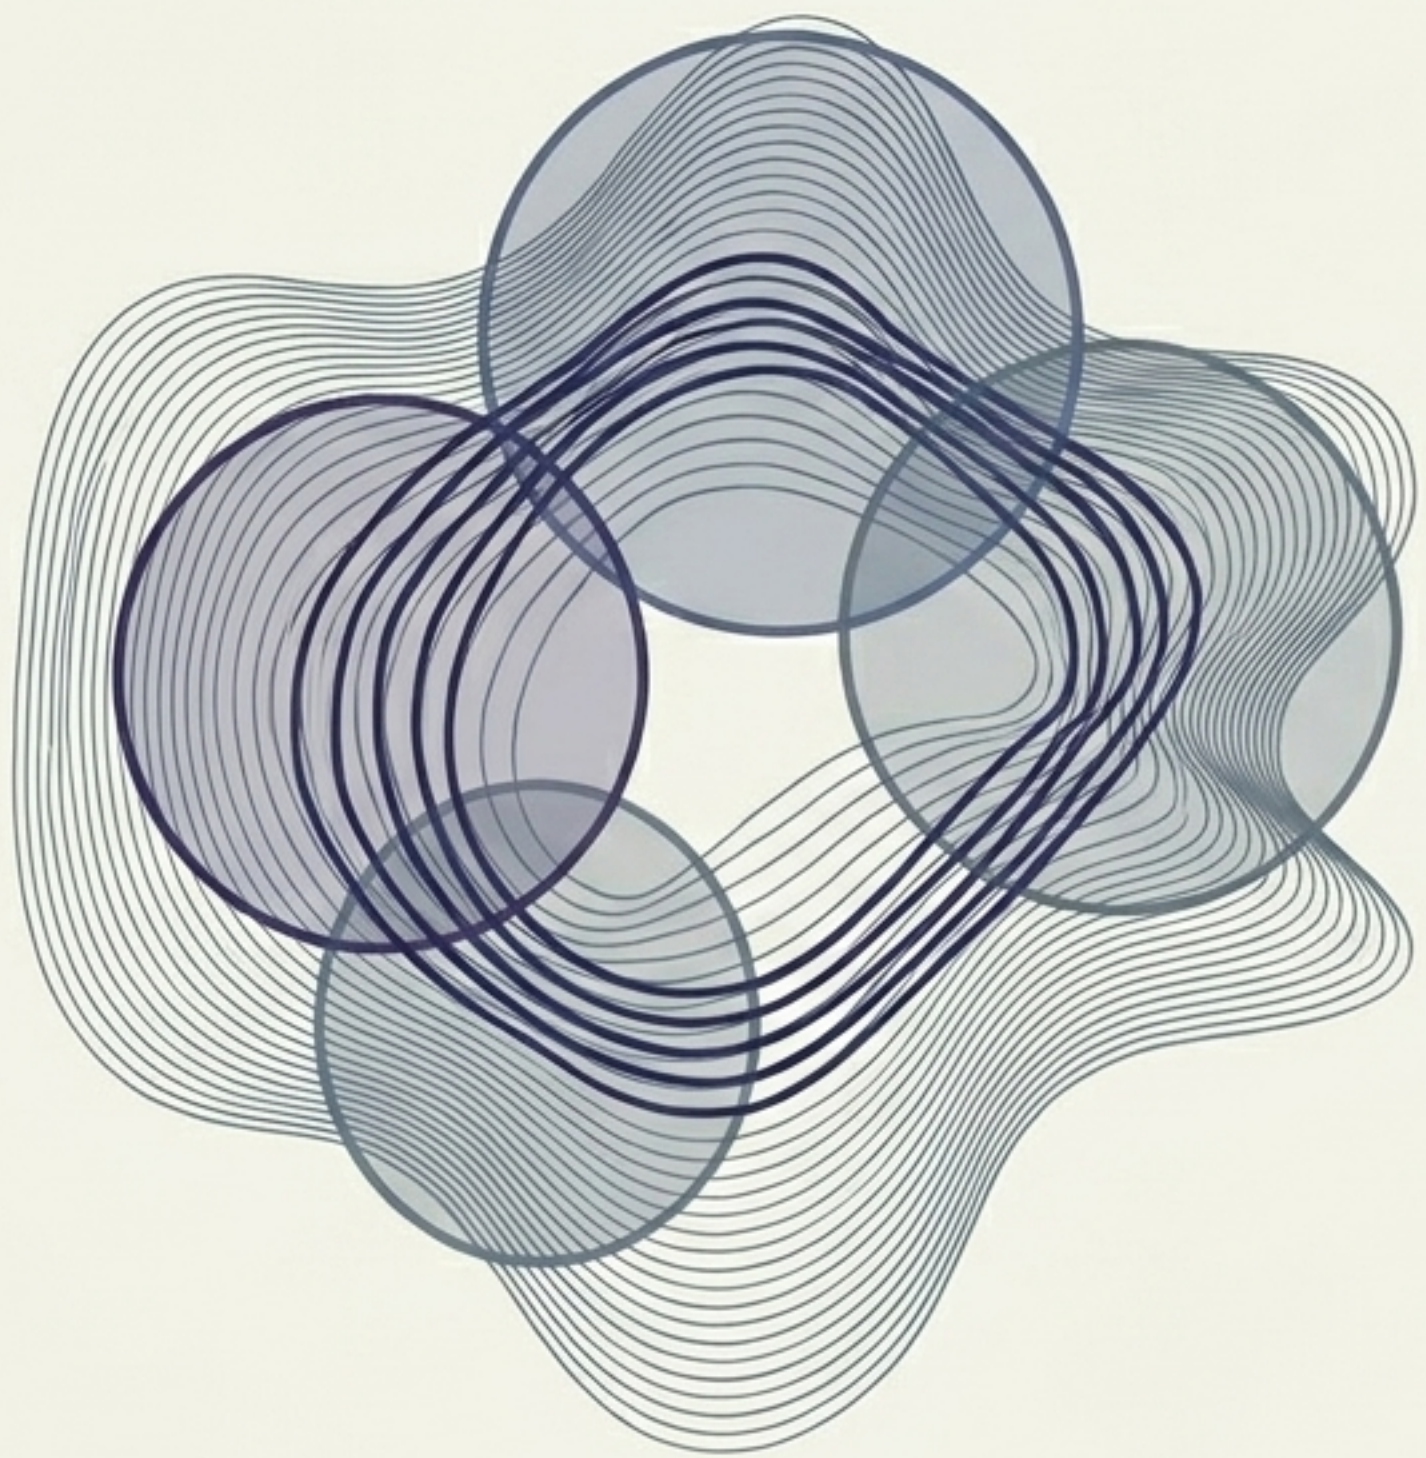
$\Lambda_t(\theta)$ provides a principled, real-time diagnostic of preference field quality. The math does not prescribe how to explore, but it definitively maps where exploration is mathematically required, and when refinement becomes productive. The transition point $\Lambda_t \approx 1$ dictates the switch from searching the space to collapsing the semantic boundary.

Structural Equivalences in Constraint-Based Frameworks



The mathematics of preference fields acts as a formal instantiation of broader constraint frameworks. The admissibility projector acts as an entropy-reduction operator analogous to RSVP, and items that persist in $A_t(\theta)$ form structural links identical to Chain of Memory persistence. Exploring rigorous categorical functors unifying these sets is an active frontier.

The Horizon: Sheaf-Theoretic Admissibility



From Projector to Connection

Moving from finite corpus-relative projectors to a true continuous differential field theory. The adaptive admissibility projector becomes a geometric connection $D_{\Pi}/dt = F(R_t, \Phi_t)$ on the semantic distribution \mathcal{A} .

Integrability

By Frobenius' theorem, is the semantic space locally path-independent, forming closed semantic leaves?

Sheaf Cohomology (H^1)

Can local admissibility assignments glue into global semantic concepts? Concepts with non-trivial obstruction classes ($H^1 \neq 0$) exist only locally and cannot be extended globally across the manifold.

Mathematical Summary & Reference

Preference Field

$$\Phi_t = \tilde{\Phi}_t / \|\tilde{\Phi}_t\|$$

where $\tilde{\Phi}_t = \sum w(r_i) \mathbf{v}_i$

Admissible Neighborhood

$$A_t(\theta) = \{i : f_t(\mathbf{v}_i) \geq \theta\}$$

Triangle Corrections

Red: $\delta_r = 1 - a^3 p^3 / \sqrt{d}$

Blue: $\delta_b = 1 + a^3 (1 - p)^3 / \sqrt{d}$

Spectral Admissibility

$(I - \Pi_{\text{edge}})$ removes centroid bias, increasing discriminative variance of filtered field \tilde{f}_t .

Regime Indicator

$$\Lambda_t(\theta) = \frac{n_\theta^3 p_\theta^3}{d}$$

$\Lambda_t \ll 1$: unnavigable noise

$\Lambda_t \gg 1$: detectable geometry

Field Clique Bound

$$\mathbb{E}[\omega_t(\theta)] \lesssim 2 \log_{1/(1-p_\theta)}(n_\theta) \cdot \left(1 + \mathcal{O}\left(1/\sqrt{d}\right)\right)$$