Ollivier–Ricci Curvature Reveals Geometric Bottlenecks in Brain-Wide Causal Subspace Graphs

Elliot Tower

Draft — June 23, 2026

Abstract

We construct a weighted graph over brain regions using Grassmannian geodesic distances between causal subspaces as edge weights, then compute Ollivier–Ricci curvature on this graph. Negatively curved regions correspond to geometric "bottlenecks" where linear dimensionality reduction methods (LDA) are most unreliable. Comparing LDA-weighted and VAE-weighted curvature maps reveals sign flips at thalamo-cortical and prefrontal-hippocampal interfaces—precisely the regions where the main paper's subspace bias is most severe. Curvature on the VAE graph correlates with optogenetic silencing effects ($\rho =$ —), while LDA curvature does not ($\rho =$ —), providing an independent geometric signature of causal relevance. Ricci flow on the LDA graph suggests a graph-regularized LDA that partially corrects the subspace bias.

1 Introduction

Discrete Ricci curvature has emerged as a powerful tool for analyzing network topology, with applications in community detection (**?**), robustness analysis (**?**), and causal inference on networks (**?**). In the brain network setting, curvature captures how information flow concentrates or disperses at each node.

We introduce a novel application: computing Ollivier–Ricci curvature on a brain graph whose edge weights are *Grassmannian geodesic distances* between regions' causal subspaces. This combines the geometric framework of our main paper (**?**)—which showed that subspace divergence is stratified by effective dimensionality—with the network-theoretic language of discrete curvature. The result is a per-region "curvature score" that predicts where analysis methods fail and where causal interventions have the largest behavioral effects.

2 Background: Ollivier–Ricci Curvature

2.1 Definition

For a weighted graph $G = (V, E, w)$ with edge weights $w : E \to \mathbb{R}_{>0}$, the *Ollivier–Ricci curvature* of an edge $(x, y) \in E$ is:

$$\kappa(x, y) = 1 - \frac{W_1(\mu_x, \mu_y)}{d(x, y)}, \tag{1}$$

where W_1 is the Wasserstein-1 (earth mover's) distance between the probability measures μ_x, μ_y supported on the neighborhoods of x and y, and $d(x, y)$ is the graph distance. The parameter $\alpha \in [0, 1]$ controls the "laziness" of the random walk: $\mu_x = \alpha\delta_x + (1 - \alpha)\sum_{z\sim x} \frac{w(x,z)}{\deg(x)}\delta_z$.

Positive curvature ($\kappa > 0$) indicates that neighbors of x and y are closer than x and y themselves (locally clustered). Negative curvature ($\kappa < 0$) indicates the opposite: a "bottleneck" where neighborhoods diverge.

2.2 Node Curvature

The curvature of a node v is the mean curvature over its incident edges:

$$K(v) = \frac{1}{|\{u : (v,u) \in E\}|} \sum_{u \sim v} \kappa(v,u). \tag{2}$$

3 Methods

3.1 Grassmannian-Weighted Brain Graph

For each analysis method $m \in \{\text{LDA}, \text{PCA}, \text{VAE}\}$, we construct a weighted graph $G_m = (\mathcal{R}, E, w_m)$ where:

- $\mathcal{R}$ is the set of brain regions (73 Steinmetz regions)
- E is the Allen CCF adjacency (edges between regions within 2 mm)
- $w_m(r,s) = d_G(\sigma_m(r), \sigma_m(s))$ is the Grassmannian geodesic distance between regions' causal subspaces under method m

The Grassmannian distance is:

$$d_G(\mathcal{U}, \mathcal{V}) = \left(\sum_{i=1}^{k} \theta_i^2 \right)^{1/2}, \tag{3}$$

where θ_i are the principal angles between subspaces $\mathcal{U}, \mathcal{V} \in \text{Gr}(k,n)$.

3.2 Curvature Computation

We compute exact Ollivier–Ricci curvature ($\alpha = 0.5$, OTD method) using the `GraphRicciCurvature` package on each weighted graph G_m.

3.3 Optogenetic Silencing Correlation

We correlate per-node curvature $K(r)$ with optogenetic silencing effects from **?** for the 12 matched regions ($n = 12$). Significance is assessed by Spearman correlation with a 10,000-permutation null distribution.

3.4 Curvature Sign Flips

"Flip regions" are defined as regions where $K_{\text{LDA}}(r) < 0$ but $K_{\text{VAE}}(r) > 0$: negative curvature under LDA (bottleneck) but positive under VAE (well-connected). These are regions where method choice reverses the geometric interpretation.

Table 1: Ollivier–Ricci curvature summary statistics by method ($k = 2$, $\alpha = 0.5$).

Method	n **nodes**	n **edges**	**Mean** K	**SD**	n **neg.**	n **pos.**
LDA	—	—	—	—	—	—
PCA	—	—	—	—	—	—
VAE	—	—	—	—	—	—

4 Results

4.1 LDA Curvature Map (Exp-R1)

4.1.1 Most Negative and Positive Regions

4.2 Curvature vs. Silencing Correlation (Exp-R2)

Table 2: Spearman correlation between node curvature and optogenetic silencing effect ($n = 12$ matched regions).

Method	ρ	p	**Perm.** p
LDA	—	—	—
PCA	—	—	—
VAE	—	—	—
$\Delta\rho$ (VAE − LDA)		—	

4.3 LDA vs. VAE Curvature Comparison (Exp-R3)

4.4 Ricci Flow and Graph-Regularized LDA (Exp-R4)

5 Discussion

5.1 Gauge Freedom and Curvature

The CKA–Procrustes anti-correlation and the LDA–silencing anti-correlation share a deeper mathematical origin: *gauge non-invariance* (**?**). The "gauge group" is $\mathrm{GL}(n_r)$ acting on the neural activity space of each region. The Grassmannian distance d_G used as edge weights is invariant under the orthogonal subgroup $O(n_r)$ but not the full $\mathrm{GL}(n_r)$. Consequently, the Ricci curvature computed on the Grassmannian-weighted graph inherits partial gauge invariance. The full representation holonomy (**?**) is the unique gauge-invariant quantity under all invertible transformations, connecting curvature analysis to holonomy-based approaches (companion Paper D).