

Unraveling network dynamics via Riemannian filtering of network interactions

Yonatan Kleerekoper¹, Kesem Shapira¹, Yonatan Keselman¹, Sandra Nestler², Mohammad Kurtam¹, Shai Abramson¹, Dori Derdikman¹, Yitzhak Schiller¹, Simon Musall³, and Hadas Benisty^{1,*}

¹*Technion, Israel Institute of Technology, Haifa, Israel*

²*Institute for Advanced Simulation (IAS-6), Jülich Research Centre, Jülich, Germany*

³*Institute of Biological Information Processing, Department for Bioelectronics (IBI-3), Jülich Research Centre, Jülich, Germany*



(Received 31 December 2025; accepted 10 August 2026; published 4 September 2026)

Functional connectivity (FC) is fundamentally nonstationary, undergoing continuous reconfigurations that track shifting behavioral and cognitive states. Despite the importance of these transitions, existing analytical frameworks struggle to reconcile the high-dimensional nature of these reconfigurations with the need for structured trajectories and mechanistic interpretability. Specifically, identifying the precise network components responsible for driving these dynamics remains a significant challenge. To bridge this gap, we introduce RONI (Riemannian filtering of network interactions), a geometry-aware, unsupervised framework that treats time-varying FC as a signal evolving on the manifold of symmetric positive semidefinite (SPSD) matrices. By defining intrinsic filtering operators along the SPSD geodesics, RONI produces a multiresolution decomposition of FC dynamics directly on the manifold. This representation allows for the isolation of “dynamic drivers”—specific network elements that dominate coherent connectivity modes across distinct temporal scales. We demonstrate the versatility of RONI by applying it to a diverse array of large-scale neural recordings, including hippocampal electrophysiology data, cortex-wide and dendritic calcium imaging, and human electroencephalography. Across these varied modalities and spatiotemporal scales, RONI identifies biologically meaningful subnetworks that shape the geometry of the connectivity trajectory, providing a unified, interpretable, and quantitative framework for studying the evolution of distributed neural interactions.

DOI: [10.1103/xzdr-sh34](https://doi.org/10.1103/xzdr-sh34)

I. INTRODUCTION

Understanding how brain connectivity reorganizes over time is a longstanding goal in systems and computational neuroscience [1–5]. Functional connectivity (FC), typically estimated through correlations between neural activity, is widely used to relate network interactions to behavior, state, and cognition [6–9].

FC is often analyzed edge by edge, treating correlations as a collection of pairwise values rather than as a structured matrix encoding population-level organization [10,11]. Moreover, many traditional analyses assume FC is effectively stationary within a behavioral condition, averaging activity over long epochs and thereby neglecting its intrinsically dynamic nature [8,10,12,13]. Such averaging collapses rich temporal structure into a single summary and can obscure transient, task-related, or learning-induced reconfigurations of network organization.

A common alternative, particularly in fMRI, is to model FC as a time-varying graph whose edges fluctuate over time [14–24]. While effective for capturing certain aspects

of network dynamics, these approaches often emphasize individual edges or summary graph statistics and may therefore miss coordinated, low-dimensional reconfigurations expressed in the joint structure of the full FC matrix [25].

A growing body of evidence suggests that FC dynamics form structured trajectories linked to changes in brain state, learning, and pathology [26–28]. Importantly, correlation matrices are symmetric positive semidefinite (SPSD) and thus do not form a Euclidean vector space: Operations such as averaging, interpolation, and filtering must preserve both positivity and symmetry. Riemannian geometry provides principled tools for data-driven analysis of such data and for extracting low-dimensional trajectories. This technique has been used to study FC dynamics in calcium imaging [29,30] and fMRI [31], but it does not provide a straightforward mechanistic interpretation.

Relatedly, supervised learning on Riemannian manifolds has been applied to electroencephalography (EEG) (e.g., brain-computer interface pipelines) [32–35] and to fMRI [36–38]. While these methods can improve predictive performance and offer some interpretability, their reliance on labels limits their ability to uncover intrinsic, data-driven mechanisms underlying FC reorganization. Interpretable and unsupervised Riemannian approaches to functional connectivity have also been proposed. Still, they are typically designed to compare connectivity across discrete states or conditions rather than to analyze continuous, multiscale sequences of FC matrices [39–41].

*Contact author: hadasbe@technion.ac.il

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

A notable step toward an interpretable and unsupervised framework for FC dynamics that respects SPSP geometry was introduced by Shnitzer *et al.* [42]. Their approach is based on two operators defined on SPSP matrices: a similarity operator that preserves shared structure and a difference operator that captures deviations. However, while the similarity operator preserves the SPSP property, the difference operator does not. Because its output does not remain within the SPSP cone, the framework supports iterative smoothing. Still, it does not yield a complete multiresolution decomposition, thereby limiting its ability to resolve dynamics across temporal scales.

In this work, we introduce RONI (Riemannian filtering of network interactions), a geometric, interpretable, and unsupervised framework for analyzing FC as a high-dimensional temporal signal evolving on the cone of SPSP matrices. Our central contributions are (1) a novel high-pass operator that remains within the SPSP cone, (2) a recursive filtering framework for extracting a multiresolution representation for FC dynamics, and (3) an algorithm that attributes frequency-resolved connectivity changes to specific components of the underlying neural signals, which we term *dynamic drivers* (DDs). This construction yields a principled multiresolution decomposition that is both scalable and geometrically consistent. Crucially, by operating directly on SPSP matrices, RONI preserves the structure of FC throughout the analysis and links dominant frequency-specific modes back to the original neural components, providing interpretability that is often absent in unsupervised geometric embeddings.

We present a rigorous mathematical formulation of the framework together with an algorithm for detecting dynamic drivers. We apply RONI to large-scale neural recordings spanning multiple spatiotemporal scales, from cortical networks to dendritic dynamics, and multiple modalities, including calcium imaging, electrophysiology, and EEG. Across all modalities, RONI consistently recovers interpretable and biologically meaningful drivers of functional connectivity dynamics.

II. METHODS

A. Problem formulation

Let $\mathbf{x}[t] \in \mathbb{R}^d$ denote a multivariate time series of neural activity. We assume that the underlying FC of the network evolves slowly in time, so that pairwise interactions between the components of $\mathbf{x}[t]$ are approximately constant within a short window of T_w samples. We slide this window along the recording and compute the sample correlation in each window, yielding a sequence of SPSP matrices, $\{\mathbf{C}_n\}_{n=1}^N$, where n indexes window centers. Each $\mathbf{C}_n$ is SPSP, and the sequence $\{\mathbf{C}_n\}$ therefore traces a trajectory on the SPSP manifold. This trajectory describes the FC temporal dynamics of the observed network.

This setup gives rise to two central questions:

(1) How can we characterize the temporal dynamics of the FC sequence, $\{\mathbf{C}_n\}_{n=1}^N$ while respecting the manifold geometry?

(2) Which components of $\mathbf{x}[t]$ drive these FC dynamics?

A commonly used approach for the first question is to apply nonlinear manifold learning to obtain a low-dimensional

coordinate system that preserves the intrinsic geometry of correlation matrices [35,36,43]. We use diffusion maps [44] with affinities derived from Riemannian distances on the SPSP manifold. Specifically, we compute pairwise distances $d(\mathbf{C}_n, \mathbf{C}_m)$ using the geodesic-distance extension for SPSP matrices [45] (see Appendix A 1 for more details) and define the diffusion kernel

$$K_{nm} = \exp\left(-\frac{d^2(\mathbf{C}_n, \mathbf{C}_m)}{\varepsilon}\right), \quad (1)$$

where $\varepsilon > 0$ is the kernel bandwidth. Let $\mathbf{D}$ be the diagonal matrix with $D_{nn} = \sum_{m=1}^N K_{nm}$ and define the row-stochastic diffusion operator $\mathbf{P} = \mathbf{D}^{-1}\mathbf{K}$. Denoting by $\{(\lambda_\ell, \psi_\ell)\}_{\ell=0}^{N-1}$ the eigenpairs of $\mathbf{P}$ (ordered as $1 = \lambda_0 > \lambda_1 \geq \dots \geq 0$), the diffusion-map embedding is

$$\mathbf{C}_n \mapsto \Psi(n) = (\lambda_1 \psi_1(n), \lambda_2 \psi_2(n), \dots, \lambda_q \psi_q(n))^T \in \mathbb{R}^q, \\ \times q \ll N. \quad (2)$$

This approach has been used to obtain low-dimensional representations of time-varying FC and relate them to behavior or learning [12,29,30].

The second question is more challenging: Geometry-aware methods (including diffusion maps) typically provide no explicit inverse mapping from the low-dimensional coordinates $\Psi(n)$ back to specific entries of $\mathbf{C}_n$, and thus do not directly identify which components of $\mathbf{x}[t]$ drive the observed FC changes. Our approach, RONI, addresses both questions simultaneously, providing a multiresolution decomposition of $\{\mathbf{C}_n\}_{n=1}^N$ and detecting components of $\mathbf{x}[t]$ that consistently dominate in informative frequency bands, which we term *dynamic drivers*.

B. Similarity and difference operators on the SPSP manifold

The core building blocks of our method, RONI, are two geometric operators that act on pairs of SPSP matrices: a *similarity* operator, which emphasizes shared structure, and a *difference* operator, which accentuates changes. Both are defined via the affine-invariant Riemannian geodesic connecting two matrices.

Let $\mathbf{C}_1$ and $\mathbf{C}_2$ be SPD. The geodesic $\gamma(p)$ satisfying $\gamma(0) = \mathbf{C}_1$ and $\gamma(1) = \mathbf{C}_2$ is

$$\gamma(p) = \mathbf{C}_1^{1/2} (\mathbf{C}_1^{-1/2} \mathbf{C}_2 \mathbf{C}_1^{-1/2})^p \mathbf{C}_1^{1/2}, \quad 0 \leq p \leq 1. \quad (3)$$

Bonnabel and Sepulchre [45] extended this construction to SPSP matrices; we denote the resulting SPSP geodesic by $\tilde{\gamma}(p)$ (see Appendix A 1). In particular, they prove the following:

Proposition 1. For any pair of SPSP matrices $\mathbf{C}_1$ and $\mathbf{C}_2$, the geodesic $\tilde{\gamma}(p)$ is well defined for all $p \in \mathbb{R}$. Moreover, the SPSP manifold is geodesically complete [45].

Following Shnitzer *et al.*, the similarity operator is defined as the geodesic midpoint:

$$\mathbf{C}_1 \# \mathbf{C}_2 \triangleq \tilde{\gamma}(0.5). \quad (4)$$

This midpoint operation emphasizes the structure shared between $\mathbf{C}_1$ and $\mathbf{C}_2$.




FIG. 1. (a) A schematic illustration of the geodesic connecting two matrices where the similarity operator is positioned at the midpoint (light blue), while the difference operator (magenta) extends the geodesic such that C_2 becomes the midpoint between $C_1\%C_2$ and C_1 . (b) Frequency response of the low-pass filter. (c) Frequency response of the high-pass filter. (d) An illustration of the proposed filtering framework: recursive application of low-pass and high-pass filters, followed by downsampling by a factor of 2.

We next introduce a *difference operator* that extrapolates along the same geodesic beyond C_2 :

$$C_1\%C_2 \triangleq \tilde{\gamma}(2). \quad (5)$$

Proposition 1 guarantees that for any given SPSP inputs, the outcome is also SPSP; therefore, the operator preserves the manifold structure. Figure 1(a) presents a geometric illustration of the difference operator as it extends the geodesic so that C_2 becomes the midpoint between C_1 and $C_1\%C_2$. The difference operator is not symmetric: In general, $C_1\%C_2 \neq C_2\%C_1$. The two operators $\#$ and $\%$ are mutually inverse (Proposition A1; Appendix A4) and form the basis for our geometric filtering scheme. In the next section, we show that the difference operator selectively amplifies eigenmodes that differ between C_1 and C_2 , thereby emphasizing changes in FC structure.

C. Spectral analysis

This section characterizes how the similarity ($\#$) and difference ($\%$) operators transform the *spectral structure* (eigenvalues and eigenvectors) of their inputs. When C_1 and C_2 are FC estimates from adjacent time windows, we expect many large-scale network modes to remain approximately stable, while changes in FC may appear as changes in the *strength* of particular modes. We formalize this intuition by analyzing the action of $\#$ and $\%$ on shared (or approximately shared) eigenmodes.

We begin with an idealized case in which C_1 and C_2 share an eigenvector ψ but assign it different eigenvalues.

Theorem 1. Let ψ be an eigenvector of both C_1 and C_2 with eigenvalues λ_1 and λ_2 , respectively: $C_1\psi = \lambda_1\psi$, $C_2\psi = \lambda_2\psi$. Then, ψ is also an eigenvector of the similarity operator $C_1\#C_2$, with eigenvalue $\lambda_{\#} = \sqrt{\lambda_1\lambda_2}$.

Theorem 1 (follows from prior work by Shnitzer *et al.* [42]) shows that the similarity operator preserves shared eigenmodes while combining their strengths via a geometric mean, consistent with its interpretation as a geometry-aware “averaging” operation.

We next state the corresponding result for our difference operator, which is central to isolating changes in FC.

Theorem 2. Under the same assumptions of Theorem 1 and with $\lambda_1 > 0$, ψ is also an eigenvector of $C_1\%C_2$, with eigenvalue $\lambda_{\%} = \lambda_1^{-1}\lambda_2^2$.

Proof. See Appendix A2.

Theorem 2 reveals how $\%$ emphasizes changes in mode strength: the mapping $\lambda_2 \mapsto \lambda_2^2/\lambda_1$ amplifies increases in λ_2 relative to λ_1 and suppresses decreases. In other words, if a coordinated network pattern ψ becomes more prominent from one window to the next ($\lambda_2 > \lambda_1$), then applying $C_1\%C_2$ accentuates that mode. This property underlies our use of $\%$ to isolate temporally evolving connectivity features and, ultimately, to attribute those changes to specific neural components—dynamic drivers.

Exact eigenvector sharing is a strong assumption. In practice, $\#$ and $\%$ are applied to matrices from adjacent time windows where eigenspaces might be only approximately aligned. Appendix A2 provides a perturbation analysis showing that the above spectral behavior is stable under small deviations, supporting the robustness of these operators for slowly varying FC.

D. Similarity and difference operators as low- and high-pass filters

We now leverage the spectral properties in Sec. II C to define intrinsic, Riemannian analogs of low-pass and high-pass filtering for an SPSP matrix sequence. For consecutive FC matrices, we define the low-pass and high-pass outputs as

$$C_n^{\text{LP}} = C_n\#C_{n+1}, \quad n = 1, \dots, N-1, \quad (6)$$

$$C_n^{\text{HP}} = C_n\%C_{n+1}, \quad n = 1, \dots, N-1. \quad (7)$$

The similarity operator $\#$ computes the geodesic midpoint, so C_n^{LP} provides a geometrical smoothing of FC across adjacent windows. In contrast, the difference operator $\%$ accentuates deviations between C_n and C_{n+1} , acting as a high-pass operation on connectivity dynamics.

Classical filters for one-dimensional signals are characterized by their frequency response, defined as the relative output amplitude obtained when the input is a sinusoid at a single frequency. We adopt the same operational definition here. We probe the SPSP filters with a controlled input sequence in which only a small subnetwork exhibits sinusoidal connectivity dynamics at frequency f , while all remaining entries are unstructured noise. Specifically, we generate a sequence $\{C_n\}_{n=1}^N$ of 10×10 SPSP matrices such that correlations

within components 1–3 oscillate sinusoidally,

$$C_n(k, l) = 0.5 \sin(2\pi f n) + \varepsilon_{k,l},$$

$$\varepsilon_{k,l} \sim U(-0.01, 0.01) \text{ i.i.d.}, \quad k, l \in \{1, 2, 3\}, \quad (8)$$

and all other off-diagonal entries are noise, with unit diagonal:

$$C_n(k, l) = \varepsilon_{k,l}, \quad \text{otherwise,} \quad C_n(k, k) = 1, \quad (9)$$

where $U(a, b)$ denotes the uniform distribution on $[a, b]$.

To quantify an “amplitude” for the filtered matrix sequence, we exploit that the structured oscillation is low dimensional and is expressed primarily through the dominant spectral mode of $\mathbf{C}_n$. We therefore compute, at each time index n , the fraction of total variance explained by the leading eigenvalue,

$$r_n \triangleq \frac{\lambda_{\max}(\mathbf{C}_n)}{\text{Tr}(\mathbf{C}_n)}. \quad (10)$$

Large r_n indicates that a single coherent connectivity mode dominates the matrix, consistent with a structured subnetwork contribution, whereas smaller r_n indicates a more spectrally diffuse (noiselike) matrix.

We summarize the sequence by the average complement of this dominance ratio,

$$D(\{\mathbf{C}_n\}_{n=1}^N) \triangleq \frac{1}{N} \sum_{n=1}^N (1 - r_n) \quad (11)$$

and refer to $D(\cdot)$ as a *spectral dispersion index*. Finally, given an input sequence $\{\mathbf{C}_n^{\text{in}}\}$ and its filtered output $\{\mathbf{C}_n^{\text{out}}\}$, we define the frequency response at frequency f as the ratio of dispersion indices (in decibels),

$$H(f) \triangleq 20 \log_{10} \left(\frac{D(\{\mathbf{C}_n^{\text{out}}\})}{D(\{\mathbf{C}_n^{\text{in}}\})} \right). \quad (12)$$

With this construction, $H(f)$ reports how strongly the filter attenuates (or preserves) structured sinusoidal connectivity dynamics at frequency f . Figures 1(b) and 1(c) show the frequency responses of the resulting low-pass and high-pass filters, computed numerically from $N = 256$ matrices generated according to Eq. (8). As expected, the low-pass filter suppresses higher frequencies, whereas the high-pass filter suppresses lower frequencies.

E. RONI: Riemannian filtering of network interactions

In this section, we introduce our proposed framework, RONI, that captures time-varying FC across multiple temporal scales and identifies network components that consistently drive structured changes in connectivity. Starting from an SPSD FC sequence $\{\mathbf{C}_n\}_{n=1}^N$, RONI performs a geometry-preserving multiresolution decomposition on the SPSD cone and then extracts *dynamic drivers* by aggregating evidence across frequency-resolved connectivity modes.

Riemannian multiresolution filtering on the SPSD manifold. Building on the Riemannian low-pass and high-pass operators in Eqs. (6) and (7), we construct a multiresolution decomposition analogous to a wavelet-packet transform for scalar signals, but computed intrinsically on the SPSD manifold. Given an input sequence of N SPSD correlation matrices, at

the first scale, we apply the low-pass and high-pass filters to adjacent pairs and then downsample by a factor of 2. This yields two SPSD sequences of length $N/2$: The low-pass branch captures slowly varying FC structure shared across adjacent windows, whereas the high-pass branch captures faster FC changes emphasized by the difference operator. We then apply the same filtering-and-downsampling procedure again to each branch. This yields four SPSD sequences, each of length $N/4$, corresponding to a finer partition of four frequency bands and a reduction in temporal resolution by an additional factor of 2 [see Fig. 1(d)]. For N a power of 2, we recursively iterate $\log_2 N$ steps of filtering and downsampling until each final node contains a single SPSD matrix. We denote these terminal coefficients by $\hat{\mathbf{C}}_f$, where the index f labels the frequency bin associated with that terminal node. At this maximal depth, the representation attains the finest frequency resolution, while temporal evolution is no longer represented within a bin.

Dynamic drivers detection from frequency-resolved connectivity modes. Each coefficient $\hat{\mathbf{C}}_f$ summarizes FC structure concentrated within frequency bin f . For each bin, we compute the dominant eigenpair

$$\hat{\mathbf{C}}_f \psi_f = \lambda_f \psi_f \quad (13)$$

and interpret ψ_f as the leading frequency-specific connectivity mode. In particular, $|\psi_f[i]|$ reflects the contribution of component i to that dominant mode.

Not all frequency bins are informative: Some coefficients are dominated by diffuse, noiselike structure. To quantify whether ψ_f is concentrated on a small subset of components, consistent with a structured subnetwork, we compute the normalized entropy of its magnitude profile:

$$h_f \triangleq -\frac{1}{\ln d} \sum_{i=1}^d p_f[i] \ln p_f[i], \quad p_f[i] = \frac{|\psi_f[i]|}{\sum_{j=1}^d |\psi_f[j]|}. \quad (14)$$

The normalization by $\ln d$ ensures $h_f \in [0, 1]$. Low entropy indicates that the dominant mode is concentrated on a few components, whereas high entropy indicates a more distributed mode.

We retain a set of informative frequency bins using two criteria: (1) concentration, $h_f \leq h_{\text{cutoff}}$, and (2) a frequency-range constraint, $f < f_{\text{cutoff}}$, which biases the analysis toward structured, slower fluctuations and away from the highest-frequency bins that are more likely to reflect noise. Denoting the retained set by

$$\mathcal{F}^* \triangleq \{f : f < f_{\text{cutoff}}, h_f \leq h_{\text{cutoff}}\}. \quad (15)$$

We next convert each retained eigenvector into a binary membership pattern over components. Specifically, for each $f \in \mathcal{F}^*$, we apply k -means clustering ($k = 2$) to the values $\{|\psi_f[i]|\}_{i=1}^d$, partitioning components into high-amplitude and low-amplitude groups. This yields an indicator $\mathbf{1}_f[i] \in \{0, 1\}$, where $\mathbf{1}_f[i] = 1$ if component i belongs to the high-amplitude cluster for bin f .

ALGORITHM 1. Dynamic drivers detection using RONI: Riemannian filtering of network interactions.

Require: SPSD matrix sequence $\mathbf{C}_n$ of length N , cutoff frequency f_{cutoff} , cutoff entropy h_{cutoff} , and threshold M

- 1: Construct a multiresolution decomposition:
- 2: **for** $k = 1$ to $\log_2(N)$ **do**
- 3: **for** each node in level $k - 1$ **do**
- 4: Apply low-pass and high-pass filtering using (6), (7)
- 5: Downsample the resulting signals by 2
- 6: **end for**
- 7: **end for**
- 8: **for** each frequency bin f **do**
- 9: Extract eigenvector corresponding to the highest eigenvalue (13)
- 10: Compute normalized entropy h_f using (14)
- 11: **end for**
- 12: Retain the informative subset of frequency bins using eq. (15)
- 13: Calculate the driver score using eq. (16) and identify dynamic drivers using eq. (17).

Finally, we aggregate evidence across frequencies by defining a voting score for each component,

$$D_{\text{score}}(i) \triangleq \frac{1}{|\mathcal{F}^*|} \sum_{f \in \mathcal{F}^*} \mathbf{1}_f[i], \quad (16)$$

and identify *dynamic drivers* as those components whose score exceeds a threshold M :

$$DD \triangleq \{i : D_{\text{score}}(i) > M\}. \quad (17)$$

The three user-defined parameters of RONI control complementary aspects of the dynamic-driver identification procedure. The entropy threshold, h_{cutoff} , determines the degree of concentration required for a frequency-specific connectivity mode to be considered informative: Lower values preferentially retain highly localized subnetworks, whereas larger values allow more distributed modes to contribute. The frequency threshold, f_{cutoff} , restricts the analysis to lower-frequency bins to reduce the contribution of rapidly fluctuating components that are more likely to reflect noise. Finally, the voting threshold M determines how consistently a component must participate across informative frequency bins to be classified as a dynamic driver. Dynamic drivers are thus components that repeatedly dominate leading connectivity

modes across multiple temporal scales. The full procedure is summarized in Algorithm 1.

We illustrate RONI on a controlled toy example that contains two temporally distinct subnetworks and background noise. Specifically, we generate a sequence of $N = 256$, 20×20 SPSD matrices $\{\mathbf{C}_n\}_{n=1}^N$ in which correlations among components 1–3 oscillate at a low frequency $f_1 = 4/256$ cycles/sample, correlations among components 7–9 oscillate at a higher frequency $f_2 = 66/256$ cycles/sample, and all other off-diagonal entries are unstructured noise [see Fig. 2(a)]:

$$C_n(k, l) = 0.5 \sin(2\pi f_1 n) + \varepsilon_{k,l}, \quad k, l \in \{1, 2, 3\},$$

$$C_n(k, l) = 0.5 \sin(2\pi f_2 n) + \varepsilon_{k,l}, \quad k, l \in \{7, 8, 9\},$$

$$C_n(k, l) = \varepsilon_{k,l} \text{ otherwise, } \quad C_n(k, k) = 1,$$

$$\varepsilon_{k,l} \stackrel{\text{i.i.d.}}{\sim} U(-0.25, 0.25),$$

To ensure an SPSD sequence, each matrix was projected to the nearest SPSD matrix. Figure 2(b) shows the diffusion embedding of this sequence as a noisy cloud where the expected structure of two distinct domains corresponding to the two halves of the sequence is barely recognizable. Applying Algorithm 1 yields low-entropy leading eigenvectors in the frequency bins 4 and ~ 70 , respectively [Fig. 2(c), top]. The




FIG. 2. Toy example: a simulated sequence of SPSP matrices. (a) Schematic representation of the toy example entries, where entries 1–3 (light blue) exhibit slow dynamics and entries 7–9 (dark blue) exhibit fast dynamics; gray areas indicate unstructured noise. (b) Diffusion embedding of the sequence, computed using all components. (c) (Top) Entropy computed across frequency bins. (Bottom) Entry values of the leading eigenvector across frequency bins. (Right) Leading eigenvectors associated with frequency bins 4 and 70, demonstrating that dynamic components are preferentially highlighted. (d) Voting scores across components, light and dark blue—true dynamic drivers, gray—nondrivers. Diffusion embedding of the sequence of matrices considering: (e) dynamic drivers and (f) nondrivers only. (g) Dynamic-driver identification performance as a function of signal-to-noise ratio (SNR). Performance is quantified using the F1 score relative to the ground-truth dynamic drivers. RONI (blue) is compared with PCA-based driver identification (black), the framework of Shnitzer *et al.* using the zerofrequency bin (yellow) or the extended version of using the outputs of all high-pass filters and the zerofrequency bin (red). RONI consistently achieves superior recovery of the dynamic drivers at moderate and high SNR values. (h) Robustness analysis of hyperparameters. Heatmaps show dynamic-driver identification performance (F1 score) as a function of the entropy threshold (h_{cutoff}), frequency cutoff (f_{cutoff}), and voting threshold (M) across different SNR conditions. Dynamic-driver detection remains stable across a broad parameter regime, indicating that the method does not require parameter fine-tuning.

eigenvectors corresponding to these frequency bins point to components 1–3 and 7–9, respectively. Aggregating across informative bins via the voting score in Eq. (16) assigns the highest scores to components 1–3 and 7–9 [Fig. 2(d)], correctly identifying them as dynamic drivers.

To test whether the identified drivers preserve the intrinsic geometry of the FC dynamics, we compare the diffusion-map embeddings computed from the whole matrix sequence [Fig. 2(b)] to (1) a driver-restricted sequence obtained by retaining only rows and columns corresponding to detected drivers, and (2) the complementary sequence containing only nondrivers (non-DDs). The driver-restricted sequence recovers a clean, low-dimensional manifold consistent with the generative structure [Fig. 2(e)], whereas the nondriver sequence yields an unstructured embedding [Fig. 2(f)].

To quantitatively evaluate the proposed framework under different signal-to-noise (SNR) conditions, we compared

RONI against several alternative approaches for identifying dynamic drivers. First, we implemented a PCA-based approach in which principal components were extracted from the vectorized correlation matrices (lower triangular entries), and dynamic drivers were identified based on the corresponding component loadings. Second, we compared against the framework of Shnitzer *et al.* [42]. In its original formulation, this method identifies drivers from a single low-frequency mode obtained by recursively applying the similarity operator and downsampling by a factor of 2 until a single SPSP matrix remains. In addition, we implemented a multiscale extension of this approach by constructing the Riemannian analog of a wavelet decomposition, retaining all high-pass outputs together with the final low-pass component.

Performance was quantified using the F1 score relative to the known ground-truth dynamic drivers across a range of SNR values. As shown in Fig. 2(g), RONI consistently

achieved the highest recovery accuracy for moderate and high SNR values (>2 dB), whereas all methods performed poorly at lower SNR levels.

We further examined the sensitivity of the framework to the entropy threshold (h_{cutoff}), frequency cutoff (f_{cutoff}), and voting threshold (M). Dynamic-driver recovery remained stable over a broad range of parameter values surrounding the working point of $h_{\text{cutoff}} = 0.1$, $f_{\text{cutoff}} = 256$, and $M = 0.2$ [Fig. 2(h)], indicating that RONI is robust to moderate parameter variations. In practice, we recommend selecting the largest values of h_{cutoff} and f_{cutoff} that do not degrade performance, thereby retaining as many informative connectivity modes as possible while excluding noise-dominated components. The voting threshold M should be chosen to identify components that participate consistently across multiple informative frequency bins rather than those appearing only sporadically.

Together, these results demonstrate that RONI isolates components carrying structured multiscale connectivity dynamics while suppressing unstructured noise.

III. RESULTS

RONI provides an interpretable view of how FC reorganizes over time and which components predominantly shape these dynamics. We demonstrate the framework across multiple recording modalities and spatial scales, from large-scale multiarea electrophysiology and cortex-wide imaging to dendritic calcium signals and human EEG, highlighting both generality and interpretability.

A. Hippocampal FC encodes spatial position through dynamic drivers

We first asked whether short-timescale FC dynamics within hippocampal-entorhinal circuits encode spatial position during navigation. We analyzed chronic Neuropixels recordings from freely moving mice ($n = 4$) exploring multiple arena geometries (38 arenas; 20 min per arena), targeting subiculum, medial entorhinal cortex (MEC), and temporal association cortex (TEA). We estimated firing rates in 0.25 s bins and computed short-term correlations using a sliding window (4 s length; 2 s hop), yielding a sequence of SPSD correlation matrices $\{C_n\}_{n=1}^N$ over each exploration session. Examples of firing rate traces along with the xy position as a function of time, and as trajectories in three selected arenas are presented in Figs. 3(a) and 3(b), respectively. For a detailed description of the recording setup and preprocessing, see Appendix B 1.

As a baseline, we trained a linear decoder to predict the animal's (x, y) position directly from the neuronal ensemble activity. The decoder achieved high performance [Fig. 3(c)], consistent with the well-established role of hippocampal and entorhinal neuronal populations in encoding spatial position. We next asked whether comparable spatial information could also be extracted from the dynamics of functional connectivity, rather than from neuronal activity itself. We embedded the sequence of FC matrices using diffusion maps and trained a linear decoder on the leading ten diffusion coordinates. Position could again be decoded with high accuracy.

Having established that position can be robustly decoded from FC, we next asked which neurons contribute most strongly to this FC-based representation. For each recording session, neurons were partitioned into DD and non-DD using RONI. FC matrices were then recomputed separately for the DD and non-DD populations, embedded with diffusion maps, and used to train new position decoders based on the corresponding low-dimensional coordinates.

As controls, we repeated the same procedure using randomly selected neuronal populations matched in size to the DD population. We additionally compared RONI against two standard (Euclidean) approaches applied directly to vectorized correlation matrices (i.e., the lower-triangular entries): PCA- and LASSO-based sparse feature selection.

Dynamic-driver subnetworks retained most of the position-related information present in the full network and significantly outperformed the nondriving population ($p = 3.8 \times 10^{-6}$, paired t -test). As expected, randomized partitions achieved intermediate results between drivers and nondrivers. Importantly, the decoders based on embedded correlations (using all components, DD, non-DD, or randomized partitions) all outperformed the non-Riemannian standard approaches (PCA- and LASSO-based representations) [Fig. 3(c)]. Figure 3(d) illustrates the dramatic effect of using DD compared to non-DD, where the corresponding animal trajectory was colored according to the value of the first diffusion coordinate related to DD (top) or non-DD (bottom). The DD-derived embedding exhibited a smooth spatial gradient across the arenas, indicating that the dominant mode of FC variation closely tracked the animal's position. In contrast, the embedding obtained from the non-DD population produced a substantially less structured spatial organization. Together, these results indicate that dynamic drivers preferentially capture the FC dynamics most relevant to spatial coding.

To better understand the biological nature of the identified driver populations, we compared dynamic drivers derived from correlation dynamics with drivers obtained directly from activity dynamics. Activity drivers were identified by training a linear decoder with $L1$ regularization to promote sparsity and selecting the components with the top 50% of the absolute decoder weights. Spatially tuned neurons, including allocentric cells, egocentric cells, spatial-information cells, and head-direction cells, were preferentially identified as activity drivers rather than dynamic drivers [Fig. 3(e), $p = 10^{-4}$, paired t -test with false discovery rate (FDR) correction for multiple comparisons]. In contrast, speed cells were represented similarly in both populations.

These findings suggest that the information captured by correlation dynamics is not merely a reflection of classical single-neuron tuning properties or even the collective dynamics of neural activity. Instead, dynamic drivers represent neurons whose contribution emerges through coordinated network interactions.

We next examined the anatomical distribution of activity drivers and dynamic drivers. Neurons located in the medial entorhinal cortex were more frequently classified as activity drivers, whereas neurons from the subiculum and temporal association cortex were more evenly distributed between the two populations [Fig. 3(f); $p = 0.013$ for MEC and $p = 0.03$




FIG. 3. Spatially informative hippocampal dynamics. (a) (Top) Example firing-rate traces of neurons during a 60 s segment extracted from a 20 min recording session. (Bottom) Corresponding mouse position trace during the same 60 s interval. (b) Three example arena outlines (black) with the full mouse trajectories overlaid in blue. (c) Decoding accuracy (R^2 , tenfold cross validation) for predicting the position of mice from ensemble activity (orange), correlation dynamics using: all cells (gray), driving cells (blue), randomly selected subpopulations, nondriving cells (gray), LASSO (linear model with L1 regularization) (gray), and PCA (gray) [mean $\pm$ SEM (scanning electron microscopy), $n = 38$ arenas]. (d) Mouse (x, y) position colored by the first diffusion coordinate for drivers (top) and nondrivers (bottom) in three example arenas, illustrating that driver-restricted embeddings form smooth spatial gradients, whereas nondrivers lack structure. (e) Percentage of drivers in spatially recognized cells. (f) Percentage of drivers across brain regions.

for the unclassified cells, paired t -test with FDR correction for multiple comparisons].

Finally, we quantified the relationship between activity drivers and dynamic drivers. On average, $25 \pm 1\%$ of neurons were classified as both activity drivers and dynamic drivers, whereas $23 \pm 1\%$ were identified exclusively as dynamic drivers. In contrast, $28 \pm 1\%$ were classified only as activity drivers, while the remaining $24 \pm 1\%$ were assigned to neither group. Thus, although the two classifications exhibit a partial overlap, a substantial fraction of neurons is uniquely identified by only one of the two approaches. These findings suggest that activity dynamics and functional-connectivity dynamics capture complementary aspects of spatial representation, with dynamic drivers reflecting information that cannot be fully explained by activity dynamics alone.

B. Learning-related cortical reorganization is dominated by a reproducible set of driver regions

We used RONI to study learning-related reorganization of cortex-wide FC. We analyzed widefield calcium imaging data from mice performing an auditory decision-making task [46]

(see Appendix B 2 for a detailed description of the experimental setup). We included four animals, each contributing ~ 1400 successful trials over ~ 30 sessions, and restricted the analysis to successful detection trials to reduce behavioral confounds. Activity was parcellated into 19 cortical areas (one hemisphere; Allen CCF), and FC sequence was computed using a sliding window of 20 trials.

To visualize how cortical interactions evolve throughout learning, we compared FC matrices obtained during early (naive) and late (expert) stages of training. Substantial changes in correlation structure were observed across learning [Fig. 4(a)], suggesting that learning is accompanied by large-scale reorganization of cortex-wide interactions.

We next asked whether this reorganization follows a structured trajectory. To this end, we embedded the FC dynamics using diffusion maps and compared the resulting representation with a standard PCA-based embedding of vectorized correlation matrices. While both approaches revealed changes across training, the Riemannian diffusion embedding produced a substantially smoother and more coherent trajectory aligned with learning progression [Fig. 4(b)].




FIG. 4. Cortical regions driving learning-related FC reorganization. (a) Schematic illustration of the correlations across cortical areas in naive and expert animals and their differences. (b) Cortex-wide FC trajectory during auditory-decision training visualized via PCA of correlations (left) and diffusion-map embedding based on Riemannian distances (right); color indicates session index (training progression). (c) Decoding accuracy (R^2 , fivefold cross validation) for predicting training stage (session index) from ensemble activity, correlation dynamics, dynamic drivers, randomized subnetworks, nondynamic drivers, and standard approaches applied to vectorized correlations (bottom triangle), LASSO (a linear model with L1 regularization), and PCA. (d) Dynamic driver scores highlight cortical regions that shape correlational dynamics. (e) Activity-based scores were evaluated using the weights of a linear model to decode the session index (mean $\pm$ SEM, $n = 4$).

To quantify this relationship, we trained linear decoders to predict session index from either cortical activity or the first ten diffusion-map coordinates derived from FC dynamics [Fig. 4(c)]. Session index could be reliably decoded from FC dynamics and exceeded activity-based decoding, suggesting that learning is more strongly reflected in the reorganization of network interactions than in changes of regional activity alone.

We next investigated whether the learning-related signal present in the FC trajectory is preferentially concentrated within the dynamic drivers identified by RONI. For each animal, we partitioned cortical regions into dynamic drivers and nondrivers using RONI and repeated the FC embedding and decoding analysis using only the selected subsets. As controls, we performed the same analysis using randomly selected subsets matched in size to the dynamic-driver population. We additionally compared RONI against two standard approaches applied directly to vectorized correlation matrices (i.e., the lower-triangular entries): PCA- and LASSO-based sparse feature selection.

Dynamic-driver subnetworks retained most of the learning-related information present in the full network and substantially outperformed nondrivers ($p = 0.02$, paired t -test). As expected, the randomized set achieved intermediate results between drivers and nondrivers. All decoders based on embedded FC (using the full network or subnetworks) outperformed the decoders trained using the PCA- and LASSO-based representations [Fig. 4(c)]. Overall, these findings indicate that learning-related FC reorganization is preferentially concentrated within the cortical regions identified by RONI.

Finally, we compared the driver score achieved by RONI (in an unsupervised fashion) to the activity weights obtained by training the decoder for session index. Correlation-derived dynamic drivers consistently highlighted retrosplenial cortex, medial frontal areas, and specific somatosensory regions [Fig. 4(d)], suggesting that learning-related reconfiguration is dominated by a reproducible subset of cortical areas. In contrast, activity-derived drivers identified broadly similar regions but exhibited substantially greater variability across animals, making it difficult to distinguish informative




FIG. 5. Dendritic subnetworks support task-relevant motor encoding. (a) Example dendritic activity from a single trial in an expert animal. A tone is given at 4 s, instructing the mouse to pull the lever. (b) Corresponding lever trajectory for the same trial as in panel (a), in an expert animal. (c) Decoding accuracy (R^2 , fivefold cross validation) for predicting lever position from the activity of all ROIs, evaluated per training session. (d) Decoding power of subpopulations as percentage of the full-model accuracy using dynamic drivers identified by RONI, non-DD, and randomized subpopulation. (e) Percentage of neurons classified as homogeneous, defined as having at least 80% of their dendrites either identified as dynamic drivers or as nondrivers. For panels (c)–(e), mean $\pm$ SEM, $n = 8$ animals).

regions from the surrounding cortex [Fig. 4(e)]. Together, these findings suggest that the progression of learning is more consistently reflected in the evolving interaction structure among cortical regions than in regional activity alone, revealing a stable network substrate of cortical plasticity.

C. Dendritic dynamic drivers preferentially support emergence of task-relevant motor encoding

To probe FC dynamics at finer spatial scales, we analyzed dendritic calcium imaging from primary motor cortex (M1) during acquisition of a lever-pull task [8 mice; 15 training days; Figs. 5(a) and 5(b)]. The same field of view (FOV) was tracked across days, and regions of interest (ROIs) were annotated along layer-5 apical tuft dendritic branches (about 12 cells per animal; see Appendix B 3 for more details). Although dendritic ROIs could be assigned to parent neurons based on anatomical reconstruction, we intentionally did not use this information as prior knowledge.

First, we investigated how the neural representation of the task evolved throughout learning. We trained a linear decoder (per animal; fivefold cross validation) to predict lever position from the activity of the full dendritic population. Decoding performance increased progressively across training, with the largest improvements occurring during the initial stages of learning [Fig. 5(c)], indicating that task-relevant information becomes increasingly prominent as animals acquire the motor skill.

We next asked whether the dynamic drivers identified from FC dynamics correspond to dendritic populations that are particularly informative about behavior. To test this, we applied RONI to sequences of trial-based correlation matrices to partition the ROIs into dynamic drivers and nondrivers.

Unlike the hippocampal and cortex-wide analyses, where decoding was performed on low-dimensional representations of FC trajectories, here we evaluated the activity of the identified populations directly. Specifically, we trained separate decoders using activity from dynamic drivers, nondrivers, and randomly selected populations matched in size to the dynamic-driver population.

Dynamic-driver populations consistently achieved the highest decoding performance and exhibited the largest improvement across learning, whereas nondriver populations showed substantially weaker performance [Fig. 5(d); $p = 0.013$ for the 10th training session, paired t test with FDR correction over all training sessions]. Randomly selected populations yielded intermediate results. These findings indicate that dendrites identified as dynamic drivers of FC reorganization are preferentially involved in encoding task-relevant motor variables.

Finally, we asked whether dynamic-driver labels were randomly distributed across dendritic branches or clustered within individual neurons. Grouping ROIs according to their parent neuron revealed a strongly bimodal organization: most neurons were homogeneous, with at least 80% of their dendritic branches classified either as dynamic drivers or as nondrivers [Fig. 5(e)]. This organization suggests that the FC reconfiguration identified by RONI primarily




FIG. 6. Contralateral sensorimotor organization in human EEG. Scalp maps of aggregated driver scores for left-hand (red) and right-hand (blue) imagery across subjects. Driver topographies exhibit contralateral dominance, matching known sensorimotor organization.

reflects neuron-level specialization rather than isolated branch-specific effects.

D. RONI recovers contralateral sensorimotor structure in human EEG during motor imagery

Finally, we asked whether RONI can recover an interpretable driver structure in a noninvasive human recording modality. We analyzed a public EEG motor-imagery dataset in which participants imagined left-hand or right-hand movements (BCI Competition IV, dataset IIa [47]). This dataset provides a useful validation case, as motor imagery is known to engage contralateral sensorimotor networks, yet EEG recordings are substantially noisier than the datasets considered above.

For each subject and movement class, we computed trial-wise FC matrices. Trials were band-pass filtered (8–30 Hz), and correlation matrices were computed over a 3 s window starting 0.5 s before the go cue, yielding one SPSP matrix sequence per movement class. We applied RONI separately to left-hand and right-hand imagery trials to detect which electrodes are identified as DD. Importantly, dynamic drivers were identified in an unsupervised manner from FC dynamics alone, without using electrode labels or hemispheric information.

The resulting driver topographies recapitulated the expected contralateral organization of motor control: Electrodes over the left hemisphere contributed more strongly during right-hand imagery, whereas electrodes over the right hemisphere contributed more strongly during left-hand imagery (Fig. 6). Thus, even in human scalp EEG, RONI identified a physiologically interpretable driver structure consistent with known sensorimotor organization.

IV. CONCLUSION

We introduce RONI, a multiresolution framework for analyzing temporal sequences of FC matrices that explicitly respects their underlying Riemannian geometry. By defining recursive low-pass and high-pass filtering operations directly

on the manifold, RONI enables a multiresolution decomposition of structured correlation dynamics. Spectral analysis of the decomposed components enables the identification of dynamic drivers as subnetworks that shape temporal changes in the correlation structure. The successful application of RONI across multiple experimental modalities and spatial scales demonstrates its robustness and general utility for studying neural dynamics.

Our results further suggest that functional connectivity dynamics constitute a distinct level of neural representation. While activity-based analyses characterize how individual neural signals encode behavioral variables, they do not capture how interactions among those signals are continuously reorganized over time. By identifying the components that drive these evolving network configurations, RONI provides a complementary perspective on how distributed neural systems encode behavior.

Several promising directions remain for future work. One is to extend RONI beyond the dominant eigenmode currently used to characterize each frequency bin, allowing multiple eigenmodes to capture complementary patterns of network reorganization. Another is to relax the requirement for consistent components over time, enabling applications to settings where tracking identical units across sessions is infeasible. More broadly, integrating RONI with deep learning approaches may enable analysis of dynamic correlation structures in the presence of partial, noisy, or mismatched observations.

As RONI operates on temporal sequences of SPSP matrices, rather than neural recordings specifically, the framework is readily applicable to other domains involving evolving covariance or correlation structures, including finance, climate science, and spatiotemporal graph learning.

ACKNOWLEDGMENT

This work was supported by the Israel Science Foundation (ISF; Grant No. 2418/24) to H.B.

DATA AVAILABILITY

The implementation of RONI is publicly available at [48]. The data that support the findings of this article are available from the corresponding authors upon request.

APPENDIX A: MATHEMATICAL FOUNDATIONS AND THEORETICAL ANALYSIS

1. Approximating the geodesic between two SPSP matrices

We summarize here the extension of the affine-invariant geometry from the SPD manifold to the SPSP cone following Ref. [45].

Let $\mathbf{C}_1 = \mathbf{V}_1 \mathbf{\Lambda}_1 \mathbf{V}_1^\top$ and $\mathbf{C}_2 = \mathbf{V}_2 \mathbf{\Lambda}_2 \mathbf{V}_2^\top$ be two SPSP matrices of dimension $d \times d$, where $\mathbf{V}_k \in \mathbb{R}^{d \times r}$ contains orthonormal eigenvectors and $\mathbf{\Lambda}_k$ are diagonal matrices with positive eigenvalues.

First, we align the subspaces by computing the SVD $\mathbf{V}_1^\top \mathbf{V}_2 = \mathbf{O}_1^\top \mathbf{\Sigma} \mathbf{O}_2$, where $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_r)$, $r = \min(\text{rank}(\mathbf{C}_1), \text{rank}(\mathbf{C}_2))$, and $\theta_i = \arccos(\sigma_i)$, and define $\mathbf{\Theta} = \text{diag}(\theta_1, \dots, \theta_r)$, $\mathbf{U}_1 = \mathbf{V}_1 \mathbf{O}_1$, and $\mathbf{U}_2 = \mathbf{V}_2 \mathbf{O}_2$.

We then interpolate the subspaces along the Grassmannian by defining

$$\mathbf{U}(t) = \mathbf{U}_1 \cos(t\boldsymbol{\Theta}) + \mathbf{X} \sin(t\boldsymbol{\Theta}), \quad (\text{A1})$$

where $\mathbf{X} = (\mathbf{I} - \mathbf{U}_1 \mathbf{U}_1^\top) \mathbf{U}_2 (\sin(\boldsymbol{\Theta}))^\dagger$ and $(\cdot)^\dagger$ is the Moore-Penrose pseudoinverse. Next, we interpolate the SPD component by rotating the eigenvalue matrices $\mathbf{R}_1 = \mathbf{O}_1^\top \boldsymbol{\Lambda}_1 \mathbf{O}_1$ and $\mathbf{R}_2 = \mathbf{O}_2^\top \boldsymbol{\Lambda}_2 \mathbf{O}_2$, and forming the affine-invariant geodesic

$$\mathbf{R}(t) = \mathbf{R}_1^{1/2} (\mathbf{R}_1^{-1/2} \mathbf{R}_2 \mathbf{R}_1^{-1/2})^t \mathbf{R}_1^{1/2}. \quad (\text{A2})$$

Finally, the approximate SPSPD geodesic is reconstructed as

$$\tilde{\gamma}_{\mathbf{C}_1 \rightarrow \mathbf{C}_2}(t) = \mathbf{U}(t) \mathbf{R}(t) \mathbf{U}(t)^\top. \quad (\text{A3})$$

From this, we can approximate the squared geodesic distance as

$$d^2(\mathbf{C}_1, \mathbf{C}_2) \approx \|\boldsymbol{\Theta}\|_F^2 + \|\ln(\mathbf{R}_1^{-1/2} \mathbf{R}_2 \mathbf{R}_1^{-1/2})\|_F^2, \quad (\text{A4})$$

where the first term is the squared Grassmannian distance and the second is the squared affine-invariant distance on the SPD manifold.

2. Proofs for spectral analysis

The proof is given for SPD matrices. The result extends to SPSPD matrices through the geodesic continuation, defined in Appendix 1.

Theorem A1. Let ψ be an eigenvector of both $\mathbf{C}_1$ and $\mathbf{C}_2$ with eigenvalues λ_1 and λ_2 , respectively: $\mathbf{C}_1 \psi = \lambda_1 \psi$, $\mathbf{C}_2 \psi = \lambda_2 \psi$. Then, ψ is also an eigenvector of the difference operator $\mathbf{C}_1 \% \mathbf{C}_2$, with eigenvalue $\lambda_1^{-1} \lambda_2^2$.

Proof.

$$\begin{aligned} (\mathbf{C}_1 \% \mathbf{C}_2) \psi &= \mathbf{C}_1^{0.5} (\mathbf{C}_1^{-0.5} \mathbf{C}_2 \mathbf{C}_1^{-0.5})^2 \mathbf{C}_1^{0.5} \psi \\ &= \mathbf{C}_1^{0.5} (\mathbf{C}_1^{-0.5} \mathbf{C}_2 \mathbf{C}_1^{-0.5})^2 \sqrt{\lambda_1} \psi \\ &= \mathbf{C}_1^{0.5} (\lambda_1^{-1} \lambda_2)^2 \sqrt{\lambda_1} \psi \\ &= \lambda_1^{-1} \lambda_2^2 \psi. \end{aligned}$$

3. Stability under weak perturbations

In practice, shared structure across matrices may only be approximate. The following analysis demonstrates that our operators are stable under small perturbations in shared eigenvectors, indicating that they behave well in the presence of noise. The ε -pseudospectrum of $\mathbf{C}$ is denoted by $\Lambda_\varepsilon(\mathbf{C})$.

Definition A1. Let $\mathbf{C} \in \mathbb{R}^{d \times d}$, the following definitions of the $\Lambda_\varepsilon(\mathbf{C})$ are equivalent for a small $\varepsilon > 0$ [49]:

- (1) The set of $\lambda \in \mathbb{R}$ such that $\|(\lambda \mathbf{I} - \mathbf{C})^{-1}\| \geq \varepsilon^{-1}$.
- (2) The set of $\lambda \in \mathbb{R}$ such that $\lambda \in \Lambda(\mathbf{C} + \mathbf{E})$ for some $\mathbf{E} \in \mathbb{R}^{d \times d}$, with $\|\mathbf{E}\| \leq \varepsilon$.
- (3) The set of $\lambda \in \mathbb{R}$ such that $\|(\lambda \mathbf{I} - \mathbf{C})\mathbf{v}\| \leq \varepsilon$ for some $\mathbf{v} \in \mathbb{R}^d$, with $\|\mathbf{v}\| = 1$.

The following two theorems extend our spectral results to the case of approximately common components, showing that pseudoeigenstructure is preserved under our operators. We denote by $S = \mathbf{C}_1 \# \mathbf{C}_2$ and $Q = \mathbf{C}_1 \% \mathbf{C}_2$.

Theorem A2 (Ref. [42]). Suppose there exists an eigenpair (λ_k, ψ_k) of $\mathbf{C}_k$ for $k = 1, 2$ so that $\psi_1 = \psi_2 + \psi_\varepsilon$, where $\|\psi_\varepsilon\|_2 \leq \frac{\sqrt{\lambda_2}}{\lambda_{2\max} \sqrt{\lambda_1}} \varepsilon$ for a small $\varepsilon > 0$, where

$\tilde{\lambda}_{2\max} = \|\mathbf{C}_2 - \lambda_2 \mathbf{I}\|$. Then, we have $\sqrt{\lambda_1 \lambda_2} \in \Lambda_\varepsilon(\mathbf{S})$. Specifically, we have $\|(\mathbf{S} - \sqrt{\lambda_1 \lambda_2} \mathbf{I})\psi_1\|_2 \leq \varepsilon$, and ψ_1 is a corresponding ε -pseudoeigenvector of $\mathbf{S}$.

Theorem A3. Suppose there exists an eigenpair λ_k and ψ_k of $\mathbf{C}_k$ for $k = 1, 2$ so that $\psi_1 = \psi_2 + \psi_\varepsilon$, where $\|\psi_\varepsilon\|_2 \leq \frac{1}{(\sigma_{\max}(\mathbf{C}_2 \mathbf{C}_1^{-1}) + \lambda_2 \lambda_1^{-1}) \tilde{\lambda}_{2\max}} \varepsilon$, where $\sigma_{\max}$ denotes the maximum eigenvalue, for a small $\varepsilon > 0$, and where $\tilde{\lambda}_{2\max} = \|\mathbf{C}_2 - \lambda_2 \mathbf{I}\|$. Then, we have $\lambda_2^2 \lambda_1^{-1} \in \Lambda_\varepsilon(\mathbf{Q})$. Specifically, we have $\|(\mathbf{Q} - \lambda_2^2 \lambda_1^{-1} \mathbf{I})\psi_1\|_2 \leq \varepsilon$, and ψ_1 is a corresponding ε -pseudoeigenvector of $\mathbf{Q}$.

Proof. Using Eq. (A6) from Appendix A5, we have

$$\mathbf{Q} = \mathbf{C}_1^{0.5} (\mathbf{C}_1^{-0.5} \mathbf{C}_2 \mathbf{C}_1^{-0.5})^2 \mathbf{C}_1^{0.5} = (\mathbf{C}_2 \mathbf{C}_1^{-1})^2 \mathbf{C}_1.$$

Since ψ_1 is an eigenvector of $\mathbf{C}_1$ with eigenvalue λ_1 , we obtain

$$\mathbf{Q} \psi_1 = (\mathbf{C}_2 \mathbf{C}_1^{-1})^2 \mathbf{C}_1 \psi_1 = \lambda_1 (\mathbf{C}_2 \mathbf{C}_1^{-1})^2 \psi_1.$$

Thus, it suffices to show that ψ_1 is an ε -pseudoeigenvector of $(\mathbf{C}_2 \mathbf{C}_1^{-1})^2$. By direct expansion, we have

$$\mathbf{C}_2 \mathbf{C}_1^{-1} \psi_1 = \frac{1}{\lambda_1} \mathbf{C}_2 \psi_1 = \frac{1}{\lambda_1} \mathbf{C}_2 (\psi_2 + \psi_\varepsilon)$$

$$= \frac{\lambda_2}{\lambda_1} \psi_1 + \frac{1}{\lambda_1} (\mathbf{C}_2 - \lambda_2 \mathbf{I}) \psi_\varepsilon.$$

This implies

$$\begin{aligned} ((\mathbf{C}_2 \mathbf{C}_1^{-1})^2 - \lambda_2^2 \lambda_1^{-2} \mathbf{I}) \psi_1 &= (\mathbf{C}_2 \mathbf{C}_1^{-1} + \lambda_2 \lambda_1^{-1} \mathbf{I}) \\ &\quad \times (\mathbf{C}_2 \mathbf{C}_1^{-1} - \lambda_2 \lambda_1^{-1} \mathbf{I}) \psi_1. \end{aligned}$$

Multiplying both sides by λ_1 , we get

$$(\mathbf{Q} - \lambda_2^2 \lambda_1^{-1} \mathbf{I}) \psi_1 = (\mathbf{C}_2 \mathbf{C}_1^{-1} + \lambda_2 \lambda_1^{-1} \mathbf{I}) (\mathbf{C}_2 - \lambda_2 \mathbf{I}) \psi_\varepsilon.$$

Taking the ℓ_2 norm, we obtain

$$\begin{aligned} \|(\mathbf{Q} - \lambda_2^2 \lambda_1^{-1} \mathbf{I}) \psi_1\|_2 &\leq \|\mathbf{C}_2 \mathbf{C}_1^{-1} + \lambda_2 \lambda_1^{-1} \mathbf{I}\|_2 \\ &\quad \times \|\mathbf{C}_2 - \lambda_2 \mathbf{I}\|_2 \|\psi_\varepsilon\|_2. \end{aligned}$$

Applying the given bound on $\|\psi_\varepsilon\|_2$, we conclude

$$\|(\mathbf{Q} - \lambda_2^2 \lambda_1^{-1} \mathbf{I}) \psi_1\|_2 \leq \varepsilon.$$

Thus, $\lambda_2^2 \lambda_1^{-1}$ is an ε -pseudoeigenvalue of $\mathbf{Q}$, where ψ_1 is a corresponding ε -pseudoeigenvector. These results demonstrate that our operators are robust to perturbations in shared structure—a desirable property when working with noisy real-world data, such as neural recordings.

4. Proof mutually inverse

Proposition A1. Let $\mathbf{C}_1$ and $\mathbf{C}_2$ be SPD matrices, the operators defined in Eqs. (4) and (5) are mutually inverse, meaning

$$\begin{aligned} \mathbf{C}_2 \% (\mathbf{C}_1 \# \mathbf{C}_2) &= \mathbf{C}_1, \\ \mathbf{C}_2 \# (\mathbf{C}_2 \% \mathbf{C}_1) &= \mathbf{C}_1. \end{aligned} \quad (\text{A5})$$

The same relations hold for SPSPD matrices by continuity of the extended geodesic (Appendix A1).

Proof. We will first prove that $\mathbf{C}_2 \% (\mathbf{C}_1 \# \mathbf{C}_2) = \mathbf{C}_1$:

$$\mathbf{C}_1 \# \mathbf{C}_2 = \mathbf{C}_1^{0.5} (\mathbf{C}_1^{-0.5} \mathbf{C}_2 \mathbf{C}_1^{-0.5})^{0.5} \mathbf{C}_1^{0.5} = \mathbf{C}_3,$$




FIG. 7. (a) Schematic of the mouse and Neuropixels probe used for electrophysiological recordings. (b) Coronal brain section showing the probe insertion site and recording trajectory. (c) Brain regions recorded during the experiment. (d) Outlines of all arena geometries used in the study.

$$\begin{aligned}
 C_2 \% C_3 &= C_2^{0.5} (C_2^{-0.5} C_3 C_2^{-0.5})^2 C_2^{0.5} \\
 &= C_2^{0.5} (C_2^{-0.5} C_1^{0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^{0.5} C_1^{0.5} C_2^{-0.5})^2 C_2^{0.5} \\
 &= C_1^{0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^{0.5} C_1^{0.5} C_2^{-1} C_1^{0.5} \\
 &\quad \times (C_1^{-0.5} C_2 C_1^{-0.5})^{0.5} C_1^{0.5}.
 \end{aligned}$$

We now verify that this expression simplifies to C_1 :

$$\begin{aligned}
 C_1 &= C_1^{0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^{0.5} C_1^{0.5} C_2^{-1} C_1^{0.5} \\
 &\quad \times (C_1^{-0.5} C_2 C_1^{-0.5})^{0.5} C_1^{0.5},
 \end{aligned}$$

$$\begin{aligned}
 C_1^{-0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^{-0.5} C_1^{-0.5} C_1 C_1^{-0.5} \\
 \times (C_1^{-0.5} C_2 C_1^{-0.5})^{-0.5} C_1^{-0.5} = C_2^{-1},
 \end{aligned}$$

$$C_1^{-0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^{-0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^{-0.5} C_1^{-0.5} = C_2^{-1},$$

$$C_1^{-0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^{-1} C_1^{-0.5} = C_2^{-1},$$

$$C_1^{-0.5} C_1^{0.5} C_2^{-1} C_1^{0.5} C_1^{-0.5} = C_2^{-1},$$

$$C_2^{-1} = C_2^{-1},$$

$$\Rightarrow C_1 = C_2 \% (C_1 \# C_2).$$

Now, we prove that $C_2 \# (C_2 \% C_1) = C_1$:

$$C_2 \% C_1 = C_2^{0.5} (C_2^{-0.5} C_1 C_2^{-0.5})^2 C_2^{0.5} = C_4,$$




FIG. 8. (a) Schematic of the widefield imaging setup for cortex-wide functional imaging. Imaging was performed through the intact, cleared skull. (b) Imaging data were aligned to the Allen CCF using four anatomical landmarks. Cortical activity, measured as relative changes in measured fluorescence, could then be related to the activity of cortical areas (white lines). (c) Learning curves of mice trained in the auditory discrimination task. Animals achieved expert task performance within 10–20 sessions of behavioral training.




FIG. 9. (a) Experimental setup: A head-fixed mouse performs a lever-pull task during two-photon calcium imaging. Each trial lasts 12 s, where a go cue is given after 4 s. Lever position crossing the 5 mm threshold is rewarded with a sweetened drop of water delivered by a water spout. (b) Example field of view showing manually labeled dendritic ROIs (yellow). (c) Morphological reconstruction and z -plane imaging. (d) Lever-pull success rate (mean $\pm$ SEM) of eight mice across training days, showing an increase in performance with training.

$$\begin{aligned}
 C_2 \# C_4 &= C_2^{0.5} (C_2^{-0.5} C_4 C_2^{-0.5})^{0.5} C_2^{0.5} \\
 &= C_2^{0.5} (C_2^{-0.5} C_2^{0.5} (C_2^{-0.5} C_1 C_2^{-0.5})^2 C_2^{0.5} C_2^{-0.5})^{0.5} C_2^{0.5} \\
 &= C_2^{0.5} ((C_2^{-0.5} C_1 C_2^{-0.5})^2)^{0.5} C_2^{0.5} \\
 &= C_2^{0.5} C_2^{-0.5} C_1 C_2^{-0.5} C_2^{0.5} \\
 &= C_1.
 \end{aligned}$$

5. Equivalent form

Proposition A2. Equivalent form of the filters. We have

$$C_1^{0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^p C_1^{0.5} = (C_2 C_1^{-1})^p C_1.$$

Proof. For some $p \geq 0$, consider the expression

$$C_1^{0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^p C_1^{0.5}.$$

Define the matrices

$$\mathbf{M} = C_1^{-0.5} C_2 C_1^{-0.5}, \quad \tilde{\mathbf{M}} = C_2 C_1^{-1}.$$

Since C_1 and C_2 are positive definite, so is $\mathbf{M}$, which implies that $\mathbf{M}$ and $\tilde{\mathbf{M}}$ are similar via

$$\mathbf{M} = C_1^{-0.5} \tilde{\mathbf{M}} C_1^{0.5}.$$

Let Λ^M and $\mathbf{V}^M$ be the eigenvalue and eigenvector matrices of $\mathbf{M}$, respectively. Then, the eigenvalue matrix of $\tilde{\mathbf{M}}$ satisfies

$$\Lambda^{\tilde{\mathbf{M}}} = \Lambda^M,$$

with right and left eigenvectors given by

$$\mathbf{V}_R^{\tilde{\mathbf{M}}} = C_1^{0.5} \mathbf{V}^M, \quad \mathbf{V}_L^{\tilde{\mathbf{M}}} = C_1^{-0.5} \mathbf{V}^M.$$

Thus, we obtain the matrix powers:

$$\mathbf{M}^p = \mathbf{V}^M (\Lambda^M)^p (\mathbf{V}^M)^T,$$

$$\tilde{\mathbf{M}}^p = C_1^{0.5} \mathbf{V}^M (\Lambda^M)^p (\mathbf{V}^M)^T C_1^{-0.5}.$$

From this, we conclude

$$\tilde{\mathbf{M}}^p = C_1^{0.5} \mathbf{M}^p C_1^{-0.5}.$$

Therefore, we derive

$$C_1^{0.5} \mathbf{M}^p C_1^{0.5} = C_1^{0.5} \mathbf{M}^p C_1^{-0.5} C_1 = \tilde{\mathbf{M}}^p C_1.$$

Substituting back for $\mathbf{M}$ and $\tilde{\mathbf{M}}$, we obtain

$$C_1^{0.5} (C_1^{-0.5} C_2 C_1^{-0.5})^p C_1^{0.5} = (C_2 C_1^{-1})^p C_1. \quad (\text{A6})$$

APPENDIX B: EXPERIMENTAL DATASETS AND METHODOLOGICAL DETAILS

1. The hippocampal network—navigation dataset

Neuropixels probes recorded neural activity from the subiculum, MEC, and TEA in freely moving mice [Fig. 7]. A motorized XY-stage tethering system enabled unrestricted, guided exploration for 20 min in a 60×60 cm² arena. We considered 38 distinct geometries, 4 mice, and approximately 150 neurons per mouse. For each experiment, neurons were filtered by minimum mean rate ≥ 0.1 spikes per second, inclusion of Kilosort-classified “good” units [50], and exclusion of Bombcell-classified “noise” units [51].

2. Cortical network during learning

Comprehensive details of the widefield imaging are provided in Refs. [46,52]. Widefield imaging was done with an inverted tandem-lens microscope and a sCMOS camera (Edge 5.5, PCO) running at 30 frames per second and a field of view of 12.5×10.5 mm² [Fig. 8(a)]. Imaging resolution was 640×540 pixels after $4\times$ spatial binning, resulting in a spatial resolution of about 20 μ m per pixel. Mice were Ai93D;Emx-Cre;LSL-tTA animals, expressing the calcium indicator GCaMP6f in all excitatory cortical neurons. All imaging data were then rigidly aligned to the Allen Mouse Brain CCF, using anatomical landmarks, to allow for inferring the activity of different cortical areas over the course of task learning [Fig. 8(b)]. Accurate atlas alignment was further confirmed using retinotopic mapping, showing high agreement of functionally inferred locations of visual areas.

During widefield imaging, head-fixed mice were trained in an auditory discrimination task, where two sequences of Poisson-distributed click sounds were presented to either the left or right side of the animal for 1–1.5 s. After the stimulus and a subsequent 0.5-s-long delay period, mice had to perform a licking response on the side where more clicks were presented to obtain a water reward. The auditory discrimination task, therefore, required the accumulation of the auditory information over time and working memory to obtain high task

performance [Fig. 8(c)]. Animals were trained between 25 and 40 sessions, and we analyzed changes in cortical correlation structure over the course of task learning. Results here include cortex-wide population activity recorded from all training sessions of four animals.

3. Dendritic networks in the primary motor cortex (M1) during motor learning

In this longitudinal experiment, eight mice were trained to perform a motor task during 15 sessions [Fig. 9]. Each trial lasted 12 s, during which an auditory go cue was given 4 s into the trial, prompting the animals to pull the lever. Pulls that surpassed a preset displacement threshold were rewarded by a drop of sweetened water delivered to the animal's

mouth through a water spout. Lever position was sampled at 150 Hz, yielding millisecond-resolution kinematic traces. Throughout every session, simultaneous two-photon calcium imaging (30 Hz) was conducted to monitor the apical tuft dendrites of 12 (on average) layer-5 Pyramidal Tract (L5 PT) neurons in M1 forelimb region per animal. For stable tracking of dendrites, a sparse labeling technique of L5 PT neurons was applied. In addition, the pons was injected at three depths with a viral mixture comprising (1) a diluted retrograde AAV carrying Cre recombinase, (2) a concentrated retro AAV-CAG-FLEX-jGCaMP7s, and (3) a concentrated retro AAV-FLEX-tdTomato. A chronic cranial window implanted over forelimb M1 enabled repeated imaging of the same FOV across days. After each session, ROIs were manually annotated, preserving the identity of individual dendritic segments across all recordings.

-
- [1] M. Fu, X. Yu, J. Lu, and Y. Zuo, Repetitive motor learning induces coordinated formation of clustered dendritic spines *in vivo*, *Nature (London)* **483**, 92 (2012).
 - [2] J. Cichon and W.-B. Gan, Branch-specific dendritic Ca^{2+} spikes cause persistent synaptic plasticity, *Nature (London)* **520**, 180 (2015).
 - [3] A. J. Peters, J. Lee, N. G. Hedrick, K. O'Neil, and T. Komiyama, Reorganization of corticospinal output during motor learning, *Nat. Neurosci.* **20**, 1133 (2017).
 - [4] M. Kumar, G. Handy, S. Kouvaros, Y. Zhao, L. L. Brinson, E. Wei, B. Bizup, B. Doiron, and T. Tzounopoulos, Cell-type-specific plasticity of inhibitory interneurons in the rehabilitation of auditory cortex after peripheral damage, *Nat. Commun.* **14**, 4170 (2023).
 - [5] H. Safaai, A. Y. Wang, S. Kira, S. B. Malerba, S. Panzeri, and C. D. Harvey, Specialized structure of neural population codes in parietal cortex outputs, *Nat. Neurosci.* **28**, 2550 (2025).
 - [6] Y.-L. Shi, R. Zeraati, A. Levina, and T. A. Engel, Spatial and temporal correlations in neural networks with structured connectivity, *Phys. Rev. Res.* **5**, 013005 (2023).
 - [7] Y.-L. Shi, N. A. Steinmetz, T. Moore, K. Boahen, and T. A. Engel, Cortical state dynamics and selective attention define the spatial pattern of correlated variability in neocortex, *Nat. Commun.* **13**, 44 (2022).
 - [8] R. Jiang, N. Zuo, J. M. Ford, S. Qi, D. Zhi, C. Zhuo, Y. Xu, Z. Fu, J. Bustillo, J. A. Turner, *et al.*, Task-induced brain connectivity promotes the detection of individual differences in brain-behavior relationships, *NeuroImage* **207**, 116370 (2020).
 - [9] L. Varga, V. V. Moca, B. Molnár, L. Perez-Cervera, M. K. Selim, A. Díaz-Parra, D. Moratal, B. Péntek, W. H. Sommer, R. C. Mureşan, S. Canals, and M. Ercsey-Ravasz, Brain dynamics supported by a hierarchy of complex correlation patterns defining a robust functional architecture, *Cell Syst.* **15**, 770 (2024).
 - [10] O. Hazon, V. H. Mincez, D. P. Tomàs, S. Ganguli, M. J. Schnitzer, and P. E. Jercog, Noise correlations in neural ensemble activity limit the accuracy of hippocampal spatial representations, *Nat. Commun.* **13**, 4276 (2022).
 - [11] L. Kunz, B. P. Staresina, P. C. Reinacher, A. Brandt, T. A. Guth, A. Schulze-Bonhage, and J. Jacobs, Ripple-locked coactivity of stimulus-specific neurons and human associative memory, *Nat. Neurosci.* **27**, 587 (2024).
 - [12] S. Lohani, A. H. Moberly, H. Benisty, B. Landa, M. Jing, Y. Li, M. J. Higley, and J. A. Cardin, Spatiotemporally heterogeneous coordination of cholinergic and neocortical activity, *Nat. Neurosci.* **25**, 1706 (2022).
 - [13] A. Gilad and F. Helmchen, Spatiotemporal refinement of signal flow through association cortex during learning, *Nat. Commun.* **11**, 1744 (2020).
 - [14] K. Jiang, Y. Yi, L. Li, H. Li, H. Shen, F. Zhao, Y. Xu, and A. Zheng, Functional network connectivity changes in children with attention-deficit hyperactivity disorder: A resting-state fMRI study, *Int. J. Dev. Neurosci.* **78**, 1 (2019).
 - [15] R. Wang, M. Liu, X. Cheng, Y. Wu, A. Hildebrandt, and C. Zhou, Segregation, integration, and balance of large-scale resting brain networks configure different cognitive abilities, *Proc. Natl. Acad. Sci. USA* **118**, e2022288118 (2021).
 - [16] W. Chang, J. Liu, L. Nie, X. Pang, Z. Lv, and J. Zheng, The degree centrality and functional connectivity in patients with temporal lobe epilepsy presenting as ictal panic: A resting state fMRI study, *Front. Neurol.* **13**, 822253 (2022).
 - [17] L. Song, M. Kolar, and E. Xing, Time-varying dynamic Bayesian networks, in *Advances in Neural Information Processing Systems*, edited by Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta (Curran Associates, Inc., Red Hook, NY, 2009), Vol. 22.
 - [18] R. F. Betzel, K. C. Wood, C. Angeloni, M. Neimark Geffen, and D. S. Bassett, Stability of spontaneous, correlated activity in mouse auditory cortex, *PLoS Comput. Biol.* **15**, e1007360 (2019).
 - [19] T. F. Varley and O. Sporns, Network analysis of time series: Novel approaches to network neuroscience, *Front. Neurosci.* **15**, 787068 (2022).
 - [20] J. Faskowitz, F. Z. Esfahlani, Y. Jo, O. Sporns, and R. F. Betzel, Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture, *Nat. Neurosci.* **23**, 1644 (2020).
 - [21] O. Sporns, J. Faskowitz, A. S. Teixeira, S. A. Cutts, and R. F. Betzel, Dynamic expression of brain functional systems

- disclosed by fine-scale analysis of edge time series, *Network Neurosci.* **5**, 405 (2021).
- [22] F. Z. Esfahlani, L. Byrge, J. Tanner, O. Sporns, D. P. Kennedy, and R. F. Betzel, Edge-centric analysis of time-varying functional brain networks with applications in autism spectrum disorder, *NeuroImage* **263**, 119591 (2022).
 - [23] T. Azevedo, A. Campbell, R. Romero-Garcia, L. Passamonti, R. A. Bethlehem, P. Liò, and N. Toschi, A deep graph neural network architecture for modelling spatio-temporal dynamics in resting-state functional MRI data, *Med. Image Anal.* **79**, 102471 (2022).
 - [24] I. Ricchi, A. Tarun, H. P. Maretic, P. Frossard, and D. Van De Ville, Dynamics of functional network organization through graph mixture learning, *NeuroImage* **252**, 119037 (2022).
 - [25] L. Novelli and A. Razi, A mathematical perspective on edge-centric brain functional connectivity, *Nat. Commun.* **13**, 2693 (2022).
 - [26] M. G. Preti, T. A. Bolton, and D. Van De Ville, The dynamic functional connectome: State-of-the-art and perspectives, *NeuroImage* **160**, 41 (2017).
 - [27] F. Z. Esfahlani, Y. Jo, J. Faskowitz, L. Byrge, D. P. Kennedy, O. Sporns, and R. F. Betzel, High-amplitude cofluctuations in cortical activity drive functional connectivity, *Proc. Natl. Acad. Sci. USA* **117**, 28393 (2020).
 - [28] Y. E. Pollak, R. Sachdev, M. Larkum, and A. Gilad, Cortex-wide laminar dynamics diverge during learning, *bioRxiv* (2025), <https://doi.org/10.1101/2025.07.15.664840>.
 - [29] H. Benisty, D. Barson, A. H. Moberly, S. Lohani, L. Tang, R. R. Coifman, M. C. Crair, G. Mishne, J. A. Cardin, and M. J. Higley, Rapid fluctuations in functional connectivity of cortical networks encode spontaneous behavior, *Nat. Neurosci.* **27**, 148 (2024).
 - [30] A. Ghanayim, H. Benisty, A. Cohen Rimón, S. Schwartz, S. Dabdoob, S. Lifshitz, R. Talmon, and J. Schiller, VTA projections to M1 are essential for reorganization of layer 2–3 network dynamics underlying motor learning, *Nat. Commun.* **16**, 200 (2025).
 - [31] Z. Dai, Z. Zhang, and A. Srivastava, Analyzing dynamical brain functional connectivity as trajectories on space of covariance matrices, *IEEE Trans. Med. Imaging* **38**, 2196 (2019).
 - [32] Y.-T. Pan, J.-L. Chou, and C.-S. Wei, MATT: A manifold attention network for EEG decoding, *Adv. Neural Inf. Process. Syst.* **35**, 31116 (2022).
 - [33] R. Kobler, J.-i. Hirayama, Q. Zhao, and M. Kawanabe, SPD domain-specific batch normalization to crack interpretable unsupervised domain adaptation in EEG, *Adv. Neural Inf. Process. Syst.* **35**, 6219 (2022).
 - [34] X. S. Nguyen, S. Yang, and A. Histace, Neural networks on symmetric spaces of noncompact type, in *The 13th International Conference on Learning Representations* (OpenReview.net, 2025).
 - [35] M. Wang, Y. Wang, and Y. Yang, Dynamic and low-dimensional modeling of brain functional connectivity on Riemannian manifolds, *NeuroImage* **314**, 121243 (2025).
 - [36] T. Dan, Z. Huang, H. Cai, P. J. Laurienti, and G. Wu, Learning brain dynamics of evolving manifold functional MRI data using geometric-attention neural network, *IEEE Trans. Med. Imaging* **41**, 2752 (2022).
 - [37] T. Dan, J. Ding, and G. Wu, Geodynamics: A geometric state-space neural network for understanding brain dynamics on Riemannian manifolds, in *NeurIPS 2025 Conference* (Curran Associates, Inc., Red Hook, NY, 2025).
 - [38] E. Wong, J. S. Anderson, B. A. Zielinski, and P. T. Fletcher, Riemannian regression and classification models of brain networks applied to autism, in *Connectomics in Neuroimaging: 2nd International Workshop (CNI) @ MICCAI* (Springer International Publishing, Cham, Switzerland, 2018), pp. 78–87.
 - [39] K. You and H.-J. Park, Re-visiting Riemannian geometry of symmetric positive definite matrices for the analysis of functional connectivity, *NeuroImage* **225**, 117464 (2021).
 - [40] C. L. Naranjo, F. A. Razzaq, M. Li, Y. Wang, J. F. Bosch-Bayard, M. A. Lindquist, A. G. Mitjans, R. Garcia, A. G. Rabinowitz, S. G. Anderson, G. A. Chiarenza, A. Calzadareyes, T. Virues-Alba, J. R. Galler, L. Minati, M. L. B. Vega, and P. A. Valdes-Sosa, EEG functional connectivity as a Riemannian mediator: An application to malnutrition and cognition, *Hum. Brain Mapp.* **45**, e26698 (2024).
 - [41] G. Simeon, G. Piella, O. Camara, and D. Pareto, Riemannian geometry of functional connectivity matrices for multi-site attention-deficit/hyperactivity disorder data harmonization, *Front. Neuroinf.* **16**, 769274 (2022).
 - [42] T. Shnitzer, H.-T. Wu, and R. Talmon, Spatiotemporal analysis using Riemannian composition of diffusion operators, *Appl. Comput. Harmon. Anal.* **68**, 101583 (2024).
 - [43] M. Venkatesh, J. Jaja, and L. Pessoa, Comparing functional connectivity matrices: A geometry-aware approach applied to participant identification, *NeuroImage* **207**, 116398 (2020).
 - [44] R. R. Coifman and S. Lafon, Diffusion maps, *Appl. Comput. Harmon. Anal.* **21**, 5 (2006).
 - [45] S. Bonnabel and R. Sepulchre, Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank, *SIAM J. Matrix Anal. Appl.* **31**, 1055 (2010).
 - [46] S. Musall, X. R. Sun, H. Mohan, X. An, S. Gluf, S.-J. Li, R. Drewes, E. Cravo, I. Lenzi, C. Yin, *et al.*, Pyramidal cell types drive functionally distinct cortical activity patterns during decision-making, *Nat. Neurosci.* **26**, 495 (2023).
 - [47] C. Brunner, R. Leeb, G. Müller-Putz, A. Schlögl, and G. Pfurtscheller, BCI competition 2008—Graz data set A, *J. Neural Eng.* **5**, 1 (2008).
 - [48] <https://github.com/BenistyLab/RONI>.
 - [49] L. N. Trefethen and M. Embree, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators* (Princeton University Press, 2020).
 - [50] M. Pachitariu, S. Sridhar, J. Pennington, and C. Stringer, Spike sorting with Kilosort4, *Nat. Methods* **21**, 914 (2024).
 - [51] J. M. J. Fabre, E. H. van Beest, A. J. Peters, M. Carandini, and K. D. Harris, Bombcell: Automated curation and cell classification of spike-sorted electrophysiology, data (Version 1.0.0) [Computer software], Zenodo, 2023, <https://doi.org/10.5281/zenodo.8172822>.
 - [52] J. Couto, S. Musall, X. R. Sun, A. Khanal, S. Gluf, S. Saxena, I. Kinsella, T. Abe, J. P. Cunningham, L. Paninski, *et al.*, Chronic, cortex-wide imaging of specific cell populations during behavior, *Nat. Protoc.* **16**, 3241 (2021).