Unraveling Network Dynamics via
RONI: Riemannian filtering Of Network Interactions

Yonatan Kleerekoper,[1] Mohammad Kurtam,[1] Yonatan Keselman,[1] Shai Abramson,[1]
Dori Derdikman,[1] Yitzhak Schiller,[1] Simon Musall,[2] and Hadas Benisty[1,*]

[1] *Technion, Israel Institute of Technology*

[2] *Juelich, Research Center*

(Dated: December 31, 2025)

Abstract

Functional connectivity (FC) is fundamentally non-stationary, 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 enables a principled 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 EEG. Across these varied modalities and spatio-temporal scales, RONI identifies biologically meaningful sub-networks that shape the geometry of the connectivity trajectory, providing a unified, interpretable, and quantitative framework for studying the evolution of distributed neural interactions.

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

* Contact E-mail: hadasbe@technion.ac.il

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 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].

A notable step toward an interpretable and unsupervised framework for FC dynamics that respects SPSD geometry was introduced by [42]. Their approach is based on two operators defined on SPSD matrices: a similarity operator that preserves shared structure and a difference operator that captures deviations. However, while the similarity operator preserves the SPSD property, the difference operator does not. Because its output does not remain within the SPSD 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 geo-

metric, interpretable, and unsupervised framework for analyzing FC as a high-dimensional temporal signal evolving on the cone of SPSD matrices. Our central contributions are: (i) a novel high-pass operator that remains within the SPSD cone, (ii) a recursive filtering framework for extracting a multiresolution representation for FC dynamics, and (iii) an algorithm that attributes frequency-resolved connectivity changes to specific components of the underlying neural signals, which we term *dynamic drivers*. This construction yields a principled multiresolution decomposition that is both scalable and geometrically consistent. Crucially, by operating directly on SPSD 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 spatio-temporal 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 SPSD matrices, $\{\mathbf{C}_n\}_{n=1}^N$, where n indexes window centers. Each $\mathbf{C}_n$ is SPSD, and the sequence $\{\mathbf{C}_n\}$ therefore traces a trajectory on the SPSD 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 defined by Riemannian distances on the SPSD manifold. Specifically, we compute pairwise distances $d(\mathbf{C}_n, \mathbf{C}_m)$ using the geodesic-distance extension for SPSD matrices [45] (see Appendix V A for more details) and define the diffusion kernel

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

where $\varepsilon > 0$ is the kernel bandwidth. Let $\mathbf{D}$ be the diagonal matrix with $\mathbf{D}_{nn} = \sum_{m=1}^{N} \mathbf{K}_{nm}$ 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 $1 = \lambda_0 > \lambda_1 \geq \cdots \geq 0$), the diffusion-maps embedding is

$$\mathbf{C}_n \longmapsto \Psi(n) = \big(\lambda_1\psi_1(n),\ \lambda_2\psi_2(n),\ \ldots,\ \lambda_q\psi_q(n)\big)^\top \in \mathbb{R}^q, \qquad q \ll N. \tag{2}$$

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 multi-resolution 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 SPSD manifold

The core building blocks of our method, RONI, are two geometric operators that act on pairs of SPSD 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

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

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

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

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

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

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

We next introduce a *difference operator* that extrapolates along the same geodesic beyond $\mathbf{C}_2$:

$$\mathbf{C}_1 \% \mathbf{C}_2 \triangleq \tilde{\gamma}(2). \tag{5}$$

Proposition 1 guarantees that for any given SPSD inputs, the outcome is also SPSD; therefore, the operator preserves the manifold structure. Fig. 1A presents a geometric illustration of the difference operator as it extends the geodesic so that $\mathbf{C}_2$ becomes the midpoint between $\mathbf{C}_1 \# \mathbf{C}_2$ and $\mathbf{C}_1 \% \mathbf{C}_2$. The difference operator is not symmetric: in general, $\mathbf{C}_1 \% \mathbf{C}_2 \neq \mathbf{C}_2 \% \mathbf{C}_1$. The two operators # and % are mutually inverse (Proposition 2, Appendix V D) 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 $\mathbf{C}_1$ and $\mathbf{C}_2$, 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 $\mathbf{C}_1$ and $\mathbf{C}_2$ 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 $\mathbf{C}_1$ and $\mathbf{C}_2$ share an eigenvector ψ but assign it different eigenvalues.

Theorem 1 *Let ψ be an eigenvector of both $\mathbf{C}_1$ and $\mathbf{C}_2$ with eigenvalues λ_1 and λ_2, respectively: $\mathbf{C}_1\psi = \lambda_1\psi, \quad \mathbf{C}_2\psi = \lambda_2\psi$. Then ψ is also an eigenvector of the similarity operator $\mathbf{C}_1\#\mathbf{C}_2$, with eigenvalue: $\lambda_\# = \sqrt{\lambda_1\lambda_2}$.*

Theorem 1 (follows from prior work by [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 $\mathbf{C}_1\%\mathbf{C}_2$, with eigenvalue: $\lambda_\% = \lambda_1^{-1}\lambda_2^2$.*

Proof. See Appendix V B.

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 $\mathbf{C}_1\%\mathbf{C}_2$ 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 V B 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 SPSD matrix sequence. For consecutive FC matrices, we define the low-pass and high-pass outputs as

$$\mathbf{C}_n^{\mathrm{LP}} = \mathbf{C}_n\#\mathbf{C}_{n+1}, \quad n = 1, \ldots, N-1, \tag{6}$$

$$\mathbf{C}_n^{\mathrm{HP}} = \mathbf{C}_n\%\mathbf{C}_{n+1}, \quad n = 1, \ldots, N-1. \tag{7}$$

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

Classical filters for 1-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 SPSD 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 $\{\mathbf{C}_n\}_{n=1}^N$ of 10×10 SPSD matrices such that correlations within components 1–3 oscillate sinusoidally,

$$C_n(k,l) = 0.5\sin(2\pi f n) + \varepsilon_{k,l}, \quad \varepsilon_{k,l} \sim U(-0.01, 0.01) \text{ i.i.d.}, \quad k,l \in \{1,2,3\}, \tag{8}$$

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

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

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,

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

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 (noise-like) matrix.

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

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

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

dispersion indices (in dB),

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

With this construction, $H(f)$ reports how strongly the filter attenuates (or preserves) structured sinusoidal connectivity dynamics at frequency f. Fig 1B and C show the frequency responses of the resulting low- 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 functional connectivity 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 multi-resolution 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)–(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 two. 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 two (see Fig. 1D). For N a power of two, we recursively iterate $log_2 N$ steps of filtering and downsampling, till each final node contains a single SPSD matrix. We denote these terminal coefficients by $\widehat{\mathbf{C}}_f$, 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

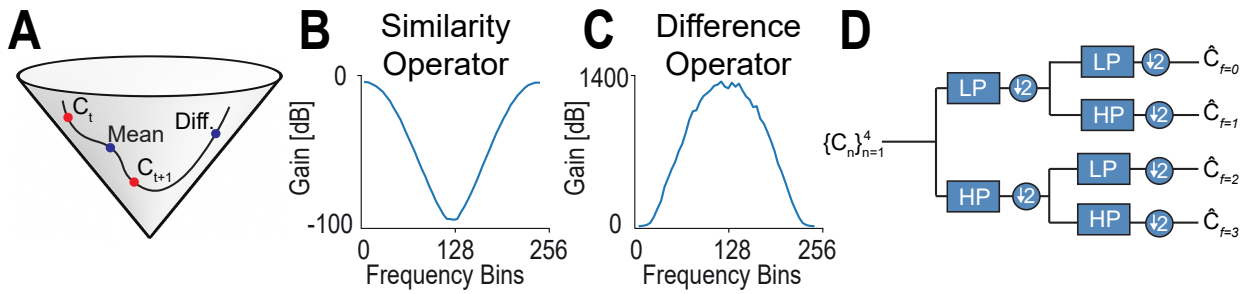

FIG. 1. (A) A schematic illustration of the geodesic connecting two matrices where the similarity operator is positioned at the midpoint, while the difference operator extends the geodesic such that $\mathbf{C}_2$ becomes the midpoint between $\mathbf{C}_1 \% \mathbf{C}_2$ and $\mathbf{C}_1$. (B) Frequency 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 down-sampling by a factor of two.

evolution is no longer represented within a bin.

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

$$\widehat{\mathbf{C}}_f \psi_f = \lambda_f \psi_f. \tag{13}$$

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, noise-like 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:

$$h_f \triangleq -\frac{1}{\log d} \sum_{i=1}^{d} p_f[i] \log p_f[i], \qquad p_f[i] = \frac{|\psi_f[i]|}{\sum_{j=1}^{d} |\psi_f[j]|}. \tag{14}$$

The normalization by $\log d$ 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: (i) **concentration**, $h_f \leq h_{\text{cutoff}}$, and (ii) 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

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

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- 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.

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

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

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

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

Dynamic drivers are thus components that repeatedly dominate leading connectivity modes across multiple temporal scales. The full procedure is summarized in Algorithm 1 (Fig. 2).

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. 3A):

$$\begin{aligned}
&C_n(k,l) = 0.5\sin(2\pi f_1 n) + \varepsilon_{k,l}, \quad k,l \in \{1,2,3\},\\
&C_n(k,l) = 0.5\sin(2\pi f_2 n) + \varepsilon_{k,l}, \quad k,l \in \{7,8,9\},\\
&C_n(k,l) = \varepsilon_{k,l} \text{ otherwise}, \qquad C_n(k,k) = 1, \qquad \varepsilon_{k,l} \overset{i.i.d.}{\sim} U(-0.25, 0.25),
\end{aligned}$$

Fig. 3B 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. 3C, top). The eigenvectors corresponding to these frequency bins point to components 1–3 and 7–9, respectively. Aggregating across informative

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

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

FIG. 2. **Algorithm 1**: Dynamic drivers detection using RONI: Riemannian filtering Of Network Interactions.

bins via the voting score in Eq. (16) assigns the highest scores to components 1–3 and 7–9 (Fig. 3D), 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. 3B) to (i) a driver-restricted sequence obtained by retaining only rows and columns corresponding to detected drivers, and (ii) the complementary sequence containing only non-drivers. The driver-restricted sequence recovers a clean, low-dimensional manifold consistent with the generative structure (Fig. 3E), whereas the non-driver sequence yields an unstructured embedding (Fig. 3F). Together, these results demonstrate that RONI isolates components carrying structured multi-scale connectivity dynamics while suppressing unstructured noise.

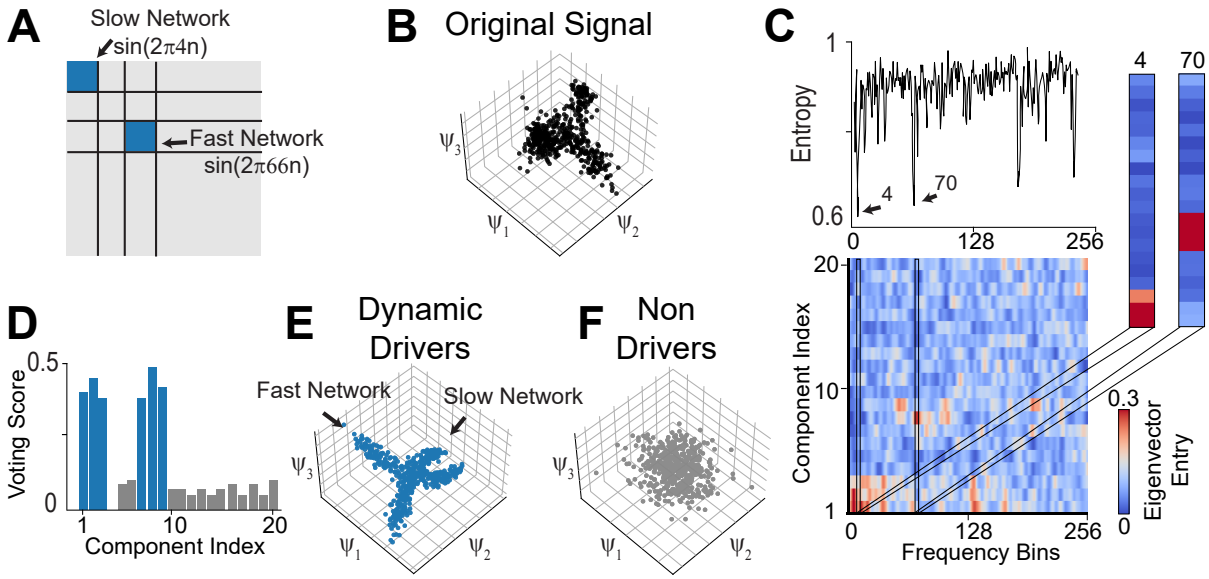

FIG. 3. Toy example: a simulated sequence of SPSD matrices. (A) Schematic representation of the toy example entries, where entries 1–3 exhibit slow dynamics and entries 7–9 exhibit fast dynamics; gray areas indicate unstructured noise. (B) Diffusion embedding of the sequence, computed using all components. (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, blue - true dynamic drivers, gray - non-drivers. (E,F) Diffusion embedding of the sequence of matrices considering: E. dynamics drivers, F. non-drivers only.

III. RESULTS

RONI provides an interpretable view of how functional connectivity (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 multi-area 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) (Fig. 4A,B; Appendix V F). 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 $\{\mathbf{C}_n\}_{n=1}^{N}$ over each exploration session. Examples for firing rate traces along with the x-y position as a function of time, and as trajectories in three selective arenas are presented in Fig. 4A and B, respectively.

For a detailed description of the recording setup and pre-processing, see Appendix V F. We computed short-term correlations using a sliding window (4-sec length; 2-sec hop), yielding a sequence of SPSD correlation matrices over each exploration session.

To quantify whether FC dynamics encode position, we embedded the FC trajectory using diffusion maps (see Methods, eqs. (1)–(2)) and trained a linear decoder to predict the animal's (x, y) location from the 10 leading embedding coordinates. Decoding was robust across arenas ($R^2 = 0.35 \pm 0.04$, mean $\pm$ SEM; 10-fold cross-validation), indicating that the FC trajectory contains spatial information.

We then asked which neurons drive these informative FC dynamics. We applied RONI to each session and identified *dynamic drivers*. We retained frequency bins whose leading eigenvectors had the lowest entropy (with h_{cutoff} set to the median). We labeled neurons as drivers if their scores exceeded the population mean, with a cap of 50% drivers per session. We repeated the same embedding-and-decoding pipeline using only the driver subnetwork or only the complementary non-driver subnetwork. Fig. 4C presents the R^2 obtained by each sub-population divided by the one obtained by the entire population (in percentage). Decoders built from driver-based FC dynamics significantly outperformed those built from non-drivers ($p = 3.8 \times 10^{-6}$), and achieved performance comparable to (and sometimes exceeding) the full population.

This difference was also apparent qualitatively: when coloring position by the first diffusion coordinate, driver-restricted embeddings formed smooth, spatially ordered gradients across arenas, whereas non-driver embeddings lacked organized structure (Fig. 4D). To-

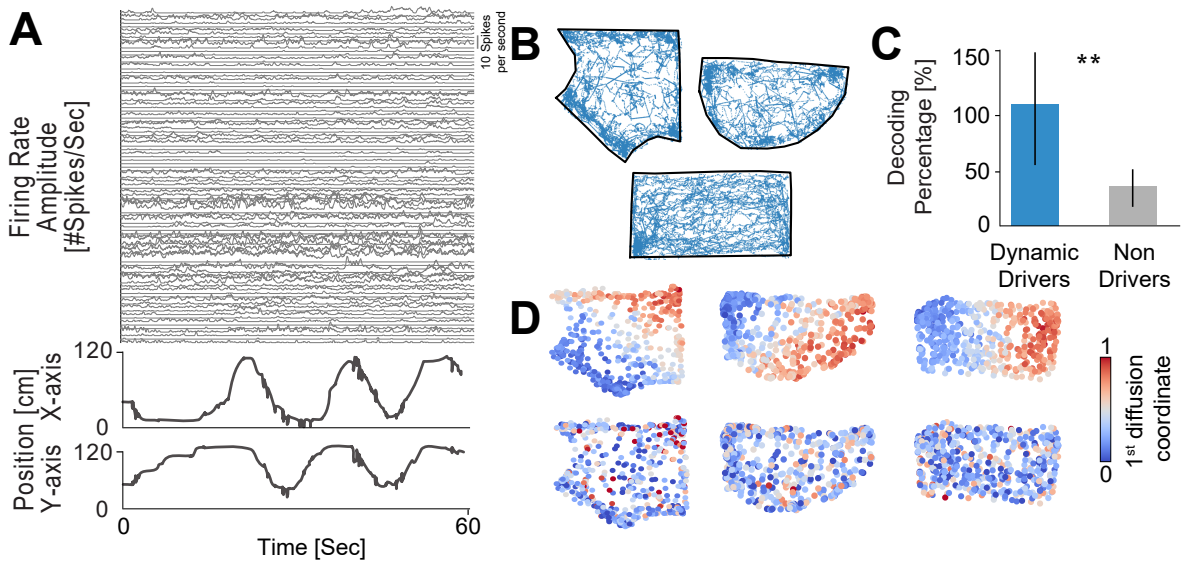

FIG. 4. **Spatially informative hippocampal dynamics.** (A) Top: Example firing-rate traces of neurons during a 60s segment extracted from a 20 min recording session. Bottom: Corresponding mouse position trace during the same 60s interval. (B) Three example arena outlines (black) with the full mouse trajectories overlaid in blue. (C) Decoding power (percentage) calculated as the accuracy (R^2) for predicting location from correlation dynamics using driving cells (blue) and non-driving cells (gray), normalized by the R^2 achieved by the entire network (mean ± SEM, $n = 38$ arenas). (D) Mouse (x, y) position colored by the first diffusion coordinate for drivers (top) and non-drivers (bottom) in three example arenas, illustrating that driver-restricted embeddings form smooth spatial gradients, whereas non-drivers lack structure.

gether, these results show that navigation-related information in short-term FC dynamics is concentrated in a subpopulation of dynamic drivers that shapes the geometry of the FC trajectory.

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

Next, 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 V G for a detailed description of the experimental setup). We in-

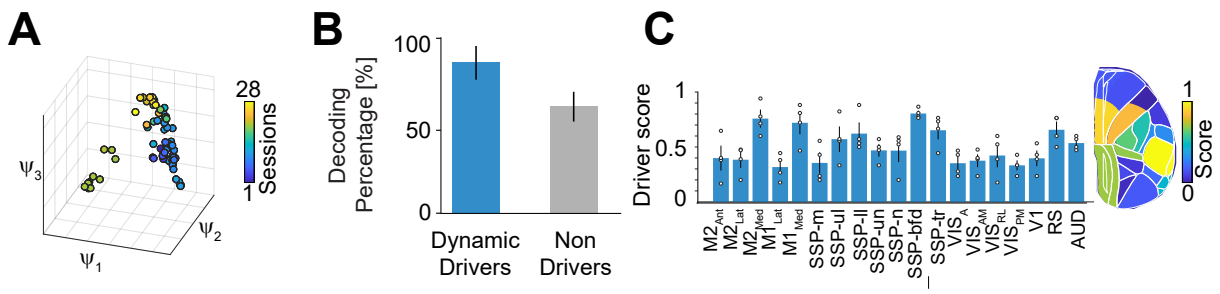

FIG. 5. **Cortical regions driving learning-related FC reorganization.** (A) Cortex-wide FC trajectory during auditory-decision training visualized via diffusion-map embedding; color indicates session index (training progression). (B) Decoding power (percentage) calculated as the accuracy (R^2) for predicting training stage (session index) from diffusion coordinates derived from the driving subnetwork, and non-driving subnetwork (mean ± SEM, $n = 4$). (C) Driver scores highlight cortical regions that shape dynamics.

cluded four animals, each contributing ∼1,400 successful trials over ∼30 sessions, restricting 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.

We embedded each animal's FC trajectory using diffusion maps (Fig. 5A). We trained a linear decoder to predict session index from the embedded coordinates, treating session index as a proxy for training progression. We then applied RONI within each animal and asked whether driver-restricted FC trajectories better captured learning progression than the complementary non-driver set. Across animals, driver-based FC dynamics yielded higher decoding accuracy than the full network, whereas non-drivers consistently performed worse (Fig. 5B).

Finally, we aggregated driver scores across animals to obtain a region-level summary of contributions to learning-related FC dynamics. Driver scores consistently highlighted retrosplenial cortex, medial frontal areas, and specific somatosensory regions (Fig. 5C), suggesting that learning-related reconfiguration is dominated by a reproducible subset of cortical areas. Together, these results indicate that RONI isolates a stable set of regions whose time-varying interactions drive cortex-wide FC reorganization during learning.

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; Fig. 6A). The same field of view was tracked across days, and ROIs were annotated along layer-5 apical tuft dendritic branches (about 12 cells per animal, see Appendix V H 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. For each trial, we computed the full pairwise correlation matrix across ROIs, yielding a trial-by-trial SPSD FC sequence per animal.

Applying RONI to these sequences identified a subset of dendritic ROIs as dynamic drivers in each animal. To test behavioral relevance, we trained linear models to predict lever position from dendritic activity and compared models trained with only driver ROIs with those trained with complementary non-driver ROIs (5-fold cross-validation). Across sessions and animals, driver-based models had greater decoding power than non-driver models (Fig. 6C), indicating that the subnetwork isolated by RONI preferentially supports task-relevant motor encoding during learning.

Moreover, driver labels were not uniformly distributed across dendrites. Grouping ROIs by parent neuron revealed a bimodal organization: most neurons were homogeneous, i.e., composed of driver or non-driver dendrites (Fig. 6). This organization suggests that the major sources of FC reconfiguration occur at the level of whole neurons, rather than being driven by isolated dendritic segments.

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

Finally, we applied RONI to human EEG using a public motor-imagery dataset (BCI Competition IV, dataset IIa [47]) in which subjects imagined left- versus right-hand movement. Trials were band-pass filtered (8–30Hz), and correlation matrices were computed over a 3sec window starting 0.5sec before the go cue, yielding one SPSD matrix sequence per movement class.

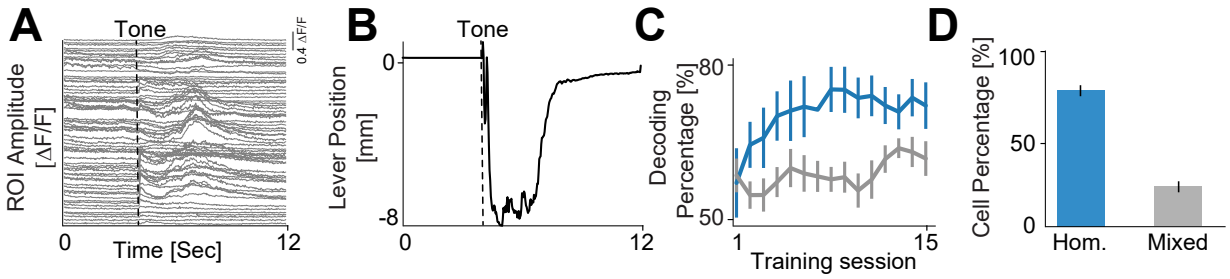

FIG. 6. **Dendritic subnetworks support task-relevant motor encoding.** (A) Example dendritic activity from a single trial in an expert animal. A tone is given at 4 seconds, instructing the mouse to pull the lever. (B) Corresponding lever trajectory for the same trial as A in an expert animal. (C) encoding power [%] for predicting lever position from the activity of sub-populations (driver dendrites vs. non-drivers), normalized by the R^2 achieved using the full neuronal population (mean $\pm$ SEM, $n = 8$ animals). (D) Percentage of neurons classified as homogeneous, defined as having at least 80% of their dendrites either identified as dynamic drivers or as non-drivers (mean $\pm$ SEM, $n = 8$ animals).

We identified dynamic drivers separately for left- and right-hand imagery and aggregated driver scores across subjects. The resulting topography 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. 7). These findings illustrate that RONI can extract an interpretable driver structure in human EEG that aligns with known neurophysiology.

IV. CONCLUSION

We introduce RONI, a novel multi-resolution framework for analyzing temporal sequences of FC matrices that explicitly respects their underlying Riemannian geometry. By defining recursive low- 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 sub-networks that shape temporal changes in the correlation structure. The successful application of RONI to

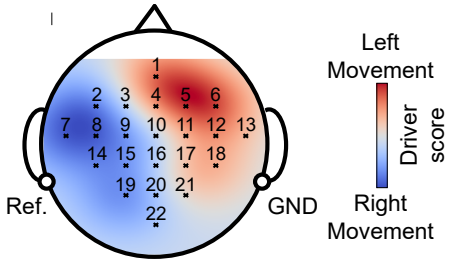

FIG. 7. **Contralateral sensorimotor organization in human EEG.** Scalp maps of aggregated driver scores for left- (red) and right-hand (blue) imagery across subjects. Driver topographies exhibit contralateral dominance, matching known sensorimotor organization.

very diverse experimental modalities and spatial scales demonstrates its robustness and general utility for studying neural population dynamics. Beyond neuroscience, RONI provides a general framework that could be extended to other domains for analyzing evolving structured data, including finance, climate science, and spatiotemporal graph learning. Future directions include addressing the method's reliance on consistent components over time, a challenge in settings such as neuroscience, where tracking identical cells across sessions is often infeasible. Integrating our approach with deep learning may enable analysis of dynamic correlation structures with partial, noisy, or mismatched observations.

---

[1] M. Fu, X. Yu, J. Lu, and Y. Zuo, Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo, Nature **483**, 92 (2012).

[2] J. Cichon and W.-B. Gan, Branch-specific dendritic ca2+ spikes cause persistent synaptic plasticity, Nature **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, Nature neuroscience **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, Nature communications **14**, 4170 (2023).

[5] H. Safaai, A. Y. Wang, S. Kira, S. Blanco Malerba, S. Panzeri, and C. D. Harvey, Specialized structure of neural population codes in parietal cortex outputs, Nature Neuroscience , 1 (2025).

[6] Y.-L. Shi, R. Zeraati, A. Levina, and T. A. Engel, Spatial and temporal correlations in neural networks with structured connectivity, Physical review research **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, Nature communications **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 Systems **15**, 770 (2024).

[10] O. Hazon, V. H. Minces, 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, Nature communications **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, Nature neuroscience **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, Nature Neuroscience **25**, 1706 (2022).

[13] A. Gilad and F. Helmchen, Spatiotemporal refinement of signal flow through association cortex during learning, Nature communications **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, International Journal of Developmental Neuroscience **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, Proceedings of the National Academy of Sciences **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, Frontiers in Neurology **13**, 822253 (2022).

[17] L. Song, M. Kolar, and E. Xing, Time-varying dynamic bayesian networks, in *Advances in Neural Information Processing Systems*, Vol. 22, edited by Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta (Curran Associates, Inc., 2009).

[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 Computational Biology **15**, e1007360 (2019).

[19] T. F. Varley and O. Sporns, Network analysis of time series: Novel approaches to network neuroscience, Frontiers in Neuroscience **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, Nature Neuroscience **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 Neuroscience **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 **280**, 119591 (2023).

[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, Medical Image Analysis **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, Nature Communications **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), functional Architecture of the Brain.

[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, Proceedings of the National Academy of Sciences **117**, 28393 (2020).

[28] Y. E. Pollak, R. Sachdev, M. Larkum, and A. Gilad, Cortex-wide laminar dynamics diverge during learning, bioRxiv , 2025 (2025).

[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, Nature Neuroscience **27**, 148 (2024).

[30] A. Ghanayim, H. Benisty, A. Cohen Rimon, 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, Nature Communications **16**, 200 (2025).

[31] Z. Dai, Z. Zhang, and A. Srivastava, Analyzing dynamical brain functional connectivity as trajectories on space of covariance matrices, IEEE Transactions on Medical Imaging **38**, 2196 (2019).

[32] Y.-T. Pan, J.-L. Chou, and C.-S. Wei, Matt: A manifold attention network for eeg decoding, Advances in Neural Information Processing Systems **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, Advances in Neural Information Processing Systems **35**, 6219 (2022).

[34] X. S. Nguyen, S. Yang, and A. Histace, Neural networks on symmetric spaces of noncompact type, in *The Thirteenth International Conference on Learning Representations* (2025).

[35] M. Wang, Y. Wang, and Y. Yang, Dynamic and low-dimensional modeling of brain functional connectivity on riemannian manifolds, NeuroImage , 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 transactions on medical 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* (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: Second International Workshop (CNI) @ MICCAI* (Springer, 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. Abdul Razzaq, M. Li, Y. Wang, J. F. Bosch-Bayard, M. A. Lindquist, A. Gonzalez Mitjans, R. Garcia, A. G. Rabinowitz, S. G. Anderson, G. A. Chiarenza, A. Calzada-Reyes, T. Virues-Alba, J. R. Galler, L. Minati, M. L. Bringas Vega, and P. A. Valdes-Sosa, Eeg functional connectivity as a riemannian mediator: An application to malnutrition and cognition, Human Brain Mapping **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, Frontiers in Neuroinformatics **16**, 769274 (2022).

[42] T. Shnitzer, H.-T. Wu, and R. Talmon, Spatiotemporal analysis using riemannian composition of diffusion operators, Applied and Computational Harmonic Analysis **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, Applied and Computational Harmonic Analysis **21**, 5 (2006), special Issue: Diffusion Maps and Wavelets.

[45] S. Bonnabel and R. Sepulchre, Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank, SIAM Journal on Matrix Analysis and Applications **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, Nature neuroscience **26**, 495 (2023).

[47] C. Brunner, R. Leeb, G. Müller-Putz, A. Schlögl, and G. Pfurtscheller, Bci competition 2008–graz data set a, Journal of Neural Engineering **5**, 1 (2008).

[48] L. N. Trefethen and M. Embree, Spectra and pseudospectra: the behavior of nonnormal matrices and operators, Princeton university press (2020).

[49] M. Pachitariu, S. Sridhar, J. Pennington, and C. Stringer, Spike sorting with kilosort4, Nature methods **21**, 914 (2024).

[50] J. Fabre, E. v. Beest, A. Peters, M. Carandini, and K. Harris, Bombcell: Automated curation and cell classification of spike-sorted electrophysiology data zenodo, july 2023 (2025).

[51] 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, Nature protocols **16**, 3241 (2021).

V. APPENDIX

A. Approximating the geodesic between two SPSD matrices

We summarize here the extension of the affine-invariant geometry from the SPD manifold to the SPSD cone following [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 SPSD 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 $\mathbf{\Sigma} = \mathrm{diag}(\sigma_1, \ldots, \sigma_r)$, $r = \min(\mathrm{rank}(\mathbf{C}_1), \mathrm{rank}(\mathbf{C}_2))$, and $\theta_i = \arccos(\sigma_i)$, and define $\mathbf{\Theta} = \mathrm{diag}(\theta_1, \ldots, \theta_r)$, $\mathbf{U}_1 = \mathbf{V}_1\mathbf{O}_1$, and $\mathbf{U}_2 = \mathbf{V}_2\mathbf{O}_2$. We then interpolate the subspaces along the Grassmannian by defining

$$\mathbf{U}(t) = \mathbf{U}_1 \cos(t\mathbf{\Theta}) + \mathbf{X}\sin(t\mathbf{\Theta}), \tag{18}$$

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

$$\mathbf{R}(t) = \mathbf{R}_1^{1/2}\left(\mathbf{R}_1^{-1/2}\mathbf{R}_2\mathbf{R}_1^{-1/2}\right)^t\mathbf{R}_1^{1/2}. \tag{19}$$

Finally, the approximate SPSD geodesic is reconstructed as

$$\widetilde{\gamma}_{\mathbf{C}_1\to\mathbf{C}_2}(t) = \mathbf{U}(t)\,\mathbf{R}(t)\,\mathbf{U}(t)^\top. \tag{20}$$

From this we can approximate the squared geodesic distance as:

$$d^2(\mathbf{C}_1, \mathbf{C}_2) \approx \|\mathbf{\Theta}\|_F^2 + \left\|\log\left(\mathbf{R}_1^{-1/2}\mathbf{R}_2\mathbf{R}_1^{-1/2}\right)\right\|_F^2 \tag{21}$$

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

B. Proofs for spectral analysis

The proof is given for SPD matrices. The result extends to SPSD matrices through the geodesic continuation, defined in Appendix V A.

Theorem 2. Let ψ be an eigenvector of both $\mathbf{C}_1$ and $\mathbf{C}_2$ with eigenvalues λ_1 and λ_2, respectively: $\mathbf{C}_1\psi = \lambda_1\psi, \quad \mathbf{C}_2\psi = \lambda_2\psi$. 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.

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

C. 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 $\Lambda_\varepsilon(\mathbf{C})$.

Definition 1 *Let $\mathbf{C} \in \mathbb{R}^{d\times d}$, the following definitions of the $\Lambda_\varepsilon(\mathbf{C})$ are equivalent for a small $\epsilon > 0$ [48]:*

1. *The set of $\lambda \in \mathbb{R}$ such that $\|(\lambda\mathbf{I} - \mathbf{C})^{-1}\| \geq \epsilon^{-1}$*
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 $\|\mathbf{E}\| \leq \epsilon$*
3. *The set of $\lambda \in \mathbb{R}$ such that $\|(\lambda\mathbf{I} - \mathbf{C})\mathbf{v}\| \leq \epsilon$ 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 pseudo-eigenstructure is preserved under our operators. We denote $S = \mathbf{C}_1\#\mathbf{C}_2$ and $Q = \mathbf{C}_1\%\mathbf{C}_2$.

Theorem 3 *[42] Suppose there exists an eigen-pair (λ_k, ψ_k) of $\mathbf{C}_k$ for $k = 1, 2$ so that $\psi_1 = \psi_2 + \psi_\epsilon$, where $\|\psi_\epsilon\|_2 \leq \frac{\sqrt{\lambda_2}}{\tilde{\lambda}_{2\max}\sqrt{\lambda_1}}\epsilon$ for a small $\epsilon > 0$, where $\tilde{\lambda}_{2\max} = \|\mathbf{C}_2 - \lambda_2\mathbf{I}\|$.*

Then we have: $\sqrt{\lambda_1\lambda_2} \in \Lambda_\epsilon(\mathbf{S})$. *Specifically, we have:* $\|(\mathbf{S} - \sqrt{\lambda_1\lambda_2}\mathbf{I})\psi_1\|_2 \leq \epsilon$, *and* ψ_1 *is a corresponding* ϵ*–pseudo-eigenvector of* $\mathbf{S}$.

Theorem 4 *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*-pseudo-eigenvector of* $\mathbf{Q}$.

Proof. Using eqn. (23) from Appendix V E , we have:

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

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-pseudo-eigenvector of $(\mathbf{C}_2{\mathbf{C}_1}^{-1})^2$. By direct expansion, we have:

$$\begin{aligned}\mathbf{C}_2{\mathbf{C}_1}^{-1}\psi_1 &= \frac{1}{\lambda_1}\mathbf{C}_2\psi_1 = \frac{1}{\lambda_1}\mathbf{C}_2(\psi_2 + \psi_\varepsilon).\\ &= \frac{\lambda_2}{\lambda_1}\psi_1 + \frac{1}{\lambda_1}(\mathbf{C}_2 - \lambda_2\mathbf{I})\psi_\varepsilon.\end{aligned}$$

This implies:

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

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:

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

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 pseudo-eigenvalue of $\mathbf{Q}$, where ψ_1 is a corresponding ε pseudo-eigenvector. 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.

D. Proof mutually inverse

Proposition 2. Let $\mathbf{C}_1$ and $\mathbf{C}_2$ be SPD matrices, the operators defined in (4)(5) are mutually inverse, meaning:

$$\begin{aligned}\mathbf{C}_2\%(\mathbf{C}_1\#\mathbf{C}_2) &= \mathbf{C}_1\\ \mathbf{C}_2\#(\mathbf{C}_2\%\mathbf{C}_1) &= \mathbf{C}_1\end{aligned} \tag{22}$$

The same relations hold for SPSD matrices by continuity of the extended geodesic (Appendix V A).

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

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

$$\begin{aligned}\mathbf{C}_2\%\mathbf{C}_3 &= \mathbf{C}_2^{0.5}(\mathbf{C}_2^{-0.5}\mathbf{C}_3\mathbf{C}_2^{-0.5})^2\mathbf{C}_2^{0.5}\\ &= \mathbf{C}_2^{0.5}(\mathbf{C}_2^{-0.5}\mathbf{C}_1^{0.5}(\mathbf{C}_1^{-0.5}\mathbf{C}_2\mathbf{C}_1^{-0.5})^{0.5}\mathbf{C}_1^{0.5}\mathbf{C}_2^{-0.5})^2\mathbf{C}_2^{0.5}\\ &= \mathbf{C}_1^{0.5}(\mathbf{C}_1^{-0.5}\mathbf{C}_2\mathbf{C}_1^{-0.5})^{0.5}\mathbf{C}_1^{0.5}\mathbf{C}_2^{-1}\mathbf{C}_1^{0.5}(\mathbf{C}_1^{-0.5}\mathbf{C}_2\mathbf{C}_1^{-0.5})^{0.5}\mathbf{C}_1^{0.5}\end{aligned}$$

We now verify that this expression simplifies to $\mathbf{C}_1$:

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

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

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

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

$$\mathbf{C}_1^{-0.5}\mathbf{C}_1^{0.5}\mathbf{C}_2^{-1}\mathbf{C}_1^{0.5}\mathbf{C}_1^{-0.5} = \mathbf{C}_2^{-1}$$

$$\mathbf{C}_2^{-1} = \mathbf{C}_2^{-1}$$

$$\Rightarrow \mathbf{C}_1 = \mathbf{C}_2\%(\mathbf{C}_1\#\mathbf{C}_2)$$

Now, we prove that $\mathbf{C}_2\#(\mathbf{C}_2\%\mathbf{C}_1) = \mathbf{C}_1$:

$$\mathbf{C}_2\%\mathbf{C}_1 = \mathbf{C}_2^{0.5}(\mathbf{C}_2^{-0.5}\mathbf{C}_1\mathbf{C}_2^{-0.5})^2\mathbf{C}_2^{0.5} = \mathbf{C}_4$$

$$
\begin{aligned}
\mathbf{C}_2\#\mathbf{C}_4 &= \mathbf{C}_2^{0.5}(\mathbf{C}_2^{-0.5}\mathbf{C}_4\mathbf{C}_2^{-0.5})^{0.5}\mathbf{C}_2^{0.5} \\
&= \mathbf{C}_2^{0.5}(\mathbf{C}_2^{-0.5}\mathbf{C}_2^{0.5}(\mathbf{C}_2^{-0.5}\mathbf{C}_1\mathbf{C}_2^{-0.5})^2\mathbf{C}_2^{0.5}\mathbf{C}_2^{-0.5})^{0.5}\mathbf{C}_2^{0.5} \\
&= \mathbf{C}_2^{0.5}((\mathbf{C}_2^{-0.5}\mathbf{C}_1\mathbf{C}_2^{-0.5})^2)^{0.5}\mathbf{C}_2^{0.5} \\
&= \mathbf{C}_2^{0.5}\mathbf{C}_2^{-0.5}\mathbf{C}_1\mathbf{C}_2^{-0.5}\mathbf{C}_2^{0.5} \\
&= \mathbf{C}_1
\end{aligned}
$$

E. Equivalent form

Proposition 3. Equivalent form of the filters. We have

$$
\mathbf{C}_1^{0.5}\left(\mathbf{C}_1^{-0.5}\mathbf{C}_2\mathbf{C}_1^{-0.5}\right)^p\mathbf{C}_1^{0.5} = \left(\mathbf{C}_2\mathbf{C}_1^{-1}\right)^p\mathbf{C}_1
$$

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

$$
\mathbf{C}_1^{0.5}\left(\mathbf{C}_1^{-0.5}\mathbf{C}_2\mathbf{C}_1^{-0.5}\right)^p\mathbf{C}_1^{0.5}.
$$

Define the matrices:

$$
\mathbf{M} = \mathbf{C}_1^{-0.5}\mathbf{C}_2\mathbf{C}_1^{-0.5}, \quad \tilde{\mathbf{M}} = \mathbf{C}_2\mathbf{C}_1^{-1}.
$$

Since $\mathbf{C}_1$ and $\mathbf{C}_2$ are positive definite, so is $\mathbf{M}$, which implies that $\mathbf{M}$ and $\tilde{\mathbf{M}}$ are similar via:

$$
\mathbf{M} = \mathbf{C}_1^{-0.5}\tilde{\mathbf{M}}\mathbf{C}_1^{0.5}.
$$

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:

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

with right and left eigenvectors given by:

$$
\mathbf{V}_R^{\tilde{M}} = \mathbf{C}_1^{0.5}\mathbf{V}^M, \quad \mathbf{V}_L^{\tilde{M}} = \mathbf{C}_1^{-0.5}\mathbf{V}^M.
$$

Thus, we obtain the matrix powers:

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

$$
\tilde{\mathbf{M}}^p = \mathbf{C}_1^{0.5}\mathbf{V}^M\left(\Lambda^M\right)^p\left(\mathbf{V}^M\right)^T\mathbf{C}_1^{-0.5}.
$$

From this, we conclude:

$$
\tilde{\mathbf{M}}^p = \mathbf{C}_1^{0.5}\mathbf{M}^p\mathbf{C}_1^{-0.5}.
$$

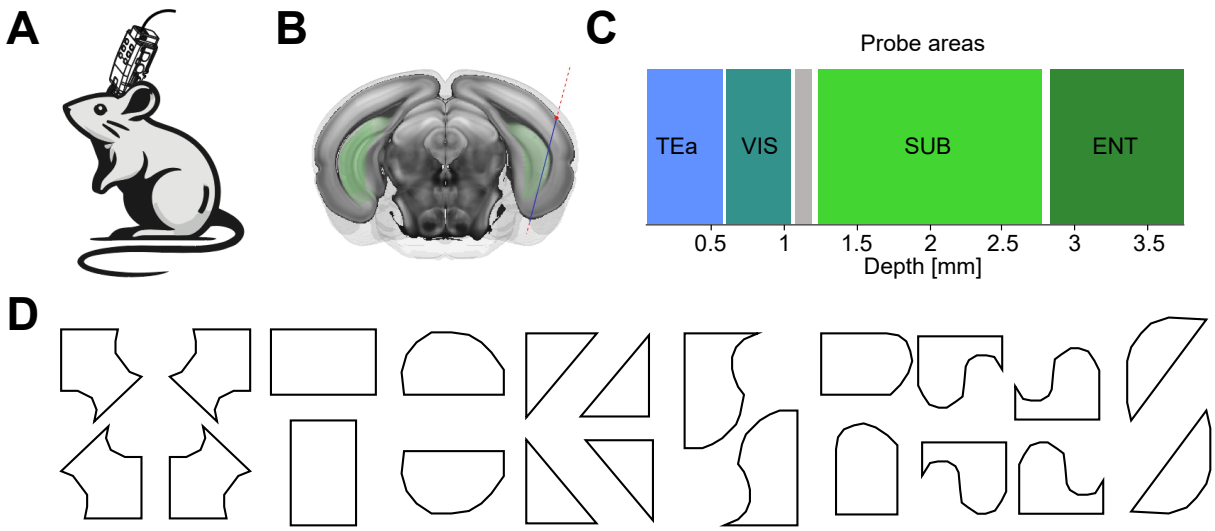

FIG. 8. (A) Schematic of the mouse and Neuropixels probe used for electrophysiological recordings. (B) Coronal brain section showing the probe insertion site and recording trajectory. (C) Brain regions recorded during the experiment. (D) Outlines of all arena geometries used in the study.

Therefore, we derive:

$$\mathbf{C}_1^{0.5}\mathbf{M}^p\mathbf{C}_1^{0.5} = \mathbf{C}_1^{0.5}\mathbf{M}^p\mathbf{C}_1^{-0.5}\mathbf{C}_1 = \tilde{\mathbf{M}}^p\mathbf{C}_1.$$

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

$$\mathbf{C}_1^{0.5}\left(\mathbf{C}_1^{-0.5}\mathbf{C}_2\mathbf{C}_1^{-0.5}\right)^p\mathbf{C}_1^{0.5} = \left(\mathbf{C}_2\mathbf{C}_1^{-1}\right)^p\mathbf{C}_1. \tag{23}$$

F. The hippocampal network - navigation dataset

Neuropixels probes recorded neural activity from the subiculum, medial entorhinal cortex (MEC), and temporal association cortex (TEA) in freely moving mice. A motorized XY-stage tethering system enabled unrestricted, guided exploration for 20 min in a 60×60 cm^2 arena. We considered 38 distinct geometries, four 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 [49], and exclusion of Bombcell-classified "noise" units [50].

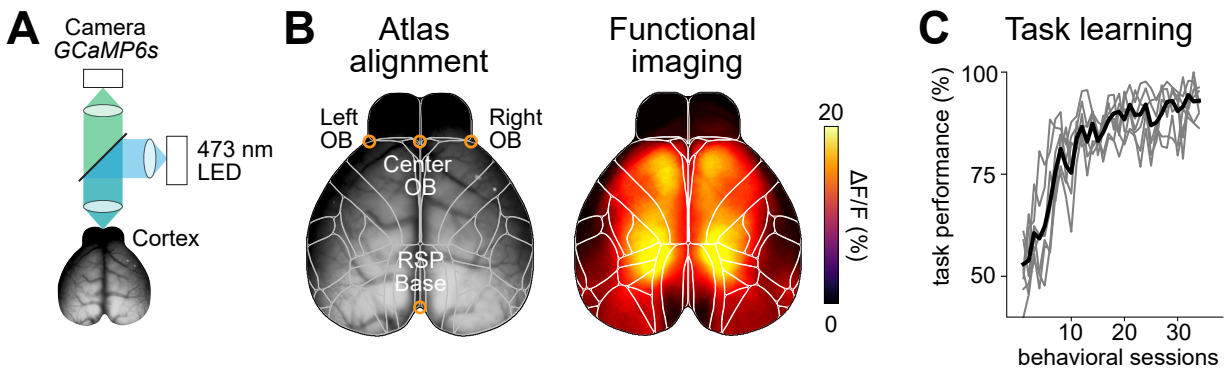

FIG. 9. (A) Schematic of the widefield imaging setup for cortex-wide functional imaging. Imaging was performed through the intact, cleared skull. (B)Imaging data was aligned to the Allen CCF using 4 anatomical landmarks. Cortical activity, measured as relative changes in measured fluorescence, could 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.

G. Cortical network during learning

Comprehensive details of the widefield imaging are provided in [46] and [51]. Widefield imaging was done with an inverted tandem-lens macroscope 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. 9A). Imaging resolution was 640 × 540 pixels after 4× spatial binning, resulting in a spatial resolution of about 20 μm per pixel. Mice were Ai93D;Emx-Cre;LSL-tTA animals, expressing the calcium indicator GCaMP6f in all excitatory cortical neurons. All imaging data was 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. 9B). 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 to 1.5 seconds. After the stimulus and a subsequent 0.5-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. 9C). 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.

H. 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. Each trial lasted 12 seconds, during which an auditory go cue was given four seconds 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 (i) a diluted retrograde AAV carrying Cre recombinase, (ii) a concentrated retro AAV-CAG-FLEX-jGCaMP7s, and (iii) a concentrated retro AAV-FLEX-tdTomato. A chronic cranial window implanted over forelimb M1 enabled repeated imaging of the same field of view (FOV) across days. After each session, regions of interest (ROIs) were manually annotated, preserving the identity of individual dendritic segments across all recordings.

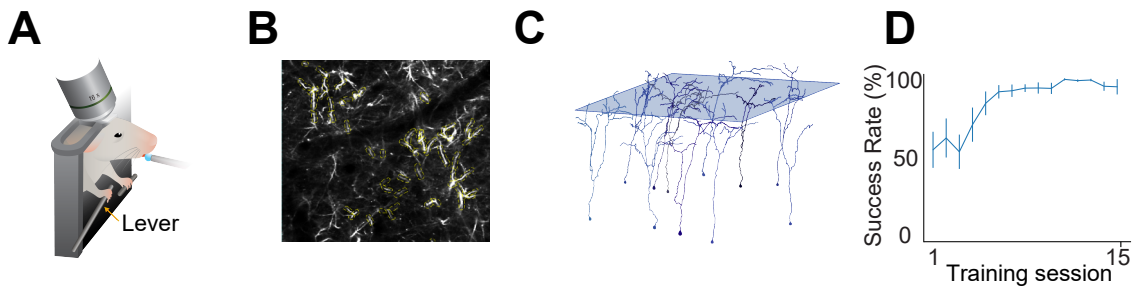

FIG. 10. (A) Experimental setup: a head-fixed mouse performs a lever-pull task during two-photon calcium imaging. Each trial lasts 12 seconds, where a go-cue is given after 4 seconds. Lever position crossing the 5mm threshold is rewarded with a sweetened drop of water delivered by a water spout. (B) Example field of view showing manually labeled dendritic ROIs (yellow). (C) Morphological reconstruction and z-plane imaging. (D) Lever-pull success rate (mean±SEM) of 8 mice across training days, showing an increase in performance with training.