

Geometric Insights into  
Neural Learning: A  
Riemannian Approach to  
Dynamics of Functional  
Connectivity

Yonatan Kleerekoper







The research thesis was done under the supervision of Dr. Hadas Benisty in the Department of Medicine.

The author of this thesis states that the research, including the collection, processing and presentation of data, addressing and comparing to previous research, etc., was done entirely in an honest way, as expected from scientific research that is conducted according to the ethical standards of the academic world. Also, reporting the research and its results in this thesis was done in an honest and complete manner, according to the same standards.

The generous financial support of the Technion is gratefully acknowledged.



# Contents

<b>Abstract</b>	<b>1</b>
<b>Abbreviations and Notations</b>	<b>3</b>
<b>1 Riemannian Wavelet Analysis</b>	<b>6</b>
1.1 Problem formulation . . . . .	6
1.2 Similarity and Difference operators on the SPSP manifold . .	8
1.3 Spectral analysis . . . . .	9
1.4 Similarity and difference operators as low- and high-pass filters	10
1.5 RONI: Riemannian filtering Of Network Interactions . . . . .	12
1.6 Results . . . . .	17
1.6.1 Hippocampal FC encodes spatial position through dynamic drivers . . . . .	17
1.6.2 Learning-related cortical reorganization is dominated by a reproducible set of driver regions . . . . .	18
1.6.3 Dendritic dynamic drivers preferentially support emergence of task-relevant motor encoding . . . . .	20
1.6.4 RONI recovers contralateral sensorimotor structure in human EEG during motor imagery . . . . .	21
1.7 Reconstruction on the Riemannian manifold . . . . .	23
1.7.1 Riemannian wavelet reconstruction . . . . .	23
<b>2 A Finsler-Randers Framework for Directed Dynamics on the SPD Manifold</b>	<b>26</b>
2.1 Problem formulation . . . . .	27
2.2 A Finsler-Randers metric on the SPD manifold . . . . .	27
2.3 Geodesics and distance under the Finsler-Randers metric . .	29

2.4	Results . . . . .	31
2.4.1	Directional dynamic applied to cortical dynamics . . .	31
<b>3</b>	<b>Conclusion</b>	<b>32</b>
<b>4</b>	<b>Appendix</b>	<b>34</b>
4.1	Approximating the geodesic between two SPSD matrices . . .	34
4.2	Proofs for spectral analysis . . . . .	35
4.3	Stability under weak perturbations . . . . .	35
4.4	Proof mutually inverse . . . . .	37
4.5	Equivalent form . . . . .	38
4.6	Proofs for reconstruction . . . . .	39
4.7	Proofs for Finsler geometry . . . . .	43
4.8	Datasets . . . . .	45
4.8.1	The hippocampal network - navigation dataset . . . .	45
4.8.2	Cortical network during learning . . . . .	45
4.8.3	Dendritic networks in the primary motor cortex (M1) during motor learning . . . . .	47

# List of Figures

1.1	Similarity and difference operations . . . . .	13
1.2	Toy Example . . . . .	16
1.3	Spatially informative hippocampal dynamics . . . . .	19
1.4	Cortical regions driving learning-related FC reorganization . .	20
1.5	Dendritic subnetworks support task-relevant motor encoding	21
1.6	Contralateral sensorimotor organization in human EEG . . .	22
1.7	Classical Reconstruction Pipeline . . . . .	23
1.8	Riemannian Reconstruction Pipeline Illustration . . . . .	25
4.1	Dataset: The hippocampal network - navigation . . . . .	45
4.2	Dataset: Cortical network during learning . . . . .	46
4.3	Dataset: Dendritic networks in the primary motor cortex (M1) during motor learning . . . . .	47



# 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 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, cortex-wide and dendritic calcium imaging, and human EEG. Across these varied modalities and spatiotemporal scales, RONI consistently identifies biologically meaningful subnetworks that shape the geometry of the connectivity trajectory, yielding interpretable low-dimensional representations of learning- and task-related network reorganization.

Beyond symmetric interactions, neural dynamics are inherently directional, reflecting asymmetric influences and causal structure. To capture this aspect, we further introduce a complementary Finsler-Randers geometric framework that incorporates directionality into the analysis of connectivity trajectories. By augmenting the Riemannian geometry of covariance

matrices with a causal bias derived from directed interactions, this approach defines a directed distance that integrates both correlational and causal information. Applied to cortex-wide imaging data, the resulting geometry enhances sensitivity to learning-related dynamics that are not captured by correlation structure alone.

Together, these contributions establish a unified geometric framework for analyzing functional connectivity dynamics, combining multiscale Riemannian analysis with a principled extension to directional interactions.

# Abbreviations and Notations

<b>A</b>	—	Bold capital letters for Matrices
$\mathbb{R}^n$	—	Real vector of dimension $n$
$\mathbb{R}^{n \times n}$	—	Real matrix of dimension $n \times n$
$I$	—	The identity matrix
$\mathbb{1}$	—	Vector of ones
$t$	—	Time point
$S_{++}^n$	—	The symmetric positive definite manifold
$S_+(p, n)$	—	The symmetric positive semi definite manifold of fixed rank $p$
$Gr(p, n)$	—	The Grassmannian manifold with $p$ -dimensional linear subspaces of $\mathbb{R}^n$
$\gamma$	—	Geodesic between two matrices on the manifold
#	—	Low pass operator
%	—	High pass operator
SPD	—	Symmetric Positive Definite
SPSD	—	Symmetric Positive Semi Definite
HP	—	High Pass
LP	—	Low Pass
SVD	—	Singular Value Decomposition
PCA	—	Principal Component Analysis
WPD	—	Wavelet Packet Decomposition
ROI	—	Region of Interest
PT	—	Pyramidal Tract
FOV	—	Field of View
AIRM	—	Affine Invariant Riemannian Metric

*“We do not learn, and that what we call learning is only a process of  
recollection.”*  
— Plato, *Meno*

The question of how the brain learns has captivated thinkers for millennia. In ancient Greece, philosophers such as Plato proposed that learning is not the acquisition of new knowledge, but rather the recollection of truths already latent within the soul, a process he termed *anamnesis*. While this view predates modern neuroscience by centuries, it resonates in an unexpected way with contemporary perspectives: learning is not a passive accumulation of information, but a transformation. Today, learning is understood as a dynamic process in which the brain reshapes its own patterns of interaction, continuously reorganizing neural circuits in response to experience.

Understanding how brain connectivity reorganizes over time is therefore a central goal in systems and computational neuroscience [17, 7, 35, 26, 40]. Because synaptic weights and anatomical connectivity cannot be directly observed in behaving animals, researchers rely on neural recordings, such as calcium imaging, electrophysiology, EEG, or fMRI, and use statistical relationships between signals, most commonly correlations, as proxies for functional connectivity (FC). FC has proven to be a powerful tool for linking network-level interactions to behavior, brain state, and cognition [43, 42, 23, 50].

Traditionally, FC is analyzed either edge by edge, treating correlations as independent pairwise quantities [20, 27], or as a stationary object averaged over long behavioral epochs [23, 20, 28, 19]. Such approaches neglect the intrinsically dynamic nature of connectivity and collapse rich temporal structure into a single summary, potentially obscuring transient, task-related, or learning-induced network reconfigurations.

To address this limitation, FC is increasingly modeled as a time-varying graph whose edges fluctuate over time [22, 54, 6, 47, 3, 51, 16, 48, 13, 1, 39]. While these approaches capture important aspects of network dynamics, they often focus on individual edges or low-level graph statistics, and may therefore miss coordinated, low-dimensional reconfigurations expressed in the joint structure of the full FC matrix [32]. Growing evidence suggests instead that FC dynamics form structured trajectories linked to learning, brain state transitions, and pathology [37, 14, 36].

A fundamental challenge in analyzing such trajectories is that correlation matrices are symmetric positive semi-definite (SPSD) objects. The space of PSD matrices is not a Euclidean vector space, and familiar linear operations, such as subtraction, averaging, or filtering, are either ill-defined or can lead to misleading results. Riemannian geometry provides a principled framework for respecting this non-Euclidean structure, enabling meaningful interpolation, smoothing, and dimensionality reduction of FC data. Riemannian methods have been successfully applied to FC dynamics in calcium imaging [2, 18], fMRI [11], EEG, and brain-computer interface pipelines [34, 25, 31, 53]. However, many existing approaches are either supervised, limiting their ability to uncover intrinsic mechanisms, or designed to compare discrete conditions rather than to analyze continuous, multiscale FC sequences [56, 30, 46].

In this thesis, we introduce **RONI** (Riemannian filtering Of Network Interactions), a geometric, interpretable, and unsupervised framework for analyzing functional connectivity as a high-dimensional temporal signal evolving on the cone of PSD matrices. Our core contribution is a novel filtering construction that remains entirely within the PSD cone, enabling recursive low-pass and high-pass decompositions analogous to wavelet analysis. This yields a principled multiresolution representation of FC dynamics that is both scalable and geometrically consistent. Crucially, the spectral structure of this decomposition allows us to attribute frequency-resolved connectivity changes to specific components of the underlying neural signals, which we term *dynamic drivers*.

Beyond symmetric interactions, learning and cognition are inherently directional processes, shaped by asymmetric influences and causal structure. To capture this aspect, we further introduce a complementary framework based on *Finsler-Randers geometry*, which incorporates directionality into the analysis of connectivity dynamics. This perspective provides a principled way to study asymmetric interactions and directional flow of information, extending the spectral, Riemannian analysis to a richer geometric setting. Together, the Riemannian and Finsler viewpoints offer a unified framework for studying how functional connectivity reorganizes across time, scale, and direction during learning.

# Chapter 1

## Riemannian Wavelet Analysis

In this chapter, we introduce a principled multi-resolution framework for analyzing temporal sequences of symmetric positive semidefinite (SPSD) matrices, such as correlation matrices extracted from time-varying neural data. While Riemannian geometry provides powerful tools for embedding and comparing such matrices, most nonlinear embeddings suffer from limited interpretability, making it difficult to link low-dimensional representations back to the original system components.

To address this limitation, we develop RONI (Riemannian filtering Of Network Interactions) a Riemannian analogue of Wavelet Packet Decomposition (WPD), where low-pass and high-pass filters are redefined through geodesic operators on the SPSP manifold. This construction enables a recursive, multi-scale decomposition that preserves positive semidefiniteness at every level. By coupling this decomposition with a spectral analysis of eigenvectors across frequency bins, we obtain an interpretable method for identifying the components of the system that drive the observed dynamics.

### 1.1 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 con-

stant 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 [52, 53, 12]. We use diffusion maps [9] 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 [4] (see Appendix 4.1 for more details) and define the diffusion kernel

$$K_{nm} = \exp\left(-\frac{d^2(\mathbf{C}_n, \mathbf{C}_m)}{\varepsilon}\right), \quad (1.1)$$

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  $1 = \lambda_0 > \lambda_1 \geq \dots \geq 0$ ), the diffusion-maps embedding is

$$\mathbf{C}_n \mapsto \Psi(n) = (\lambda_1\psi_1(n), \lambda_2\psi_2(n), \dots, \lambda_q\psi_q(n))^\top \in \mathbb{R}^q, \quad q \ll N. \quad (1.2)$$

This approach has been used to obtain low-dimensional representations of time-varying FC and relate them to behavior or learning [28, 2, 18].

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*.

## 1.2 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

$$\gamma(p) = \mathbf{C}_1^{1/2} \left( \mathbf{C}_1^{-1/2} \mathbf{C}_2 \mathbf{C}_1^{-1/2} \right)^p \mathbf{C}_1^{1/2}, \quad 0 \leq p \leq 1. \quad (1.3)$$

Bonnabel et al. [4] extended this construction to SPSD matrices; we denote the resulting SPSD geodesic by  $\tilde{\gamma}(p)$  (see Appendix 4.1). 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 [4].

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

$$\mathbf{C}_1 \# \mathbf{C}_2 \triangleq \tilde{\gamma}(0.5). \quad (1.4)$$

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  $\mathbf{C}_2$ :

$$\mathbf{C}_1 \% \mathbf{C}_2 \triangleq \tilde{\gamma}(2). \quad (1.5)$$

Proposition 1 guarantees that for any given SPSD inputs, the outcome is also SPSD; therefore, the operator preserves the manifold structure. Fig. 1.1A 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 4.4) 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.

### 1.3 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  $\psi$  but assign it different eigenvalues.

**Theorem 1** *Let  $\psi$  be an eigenvector of both  $\mathbf{C}_1$  and  $\mathbf{C}_2$  with eigenvalues  $\lambda_1$  and  $\lambda_2$ , respectively:  $\mathbf{C}_1\psi = \lambda_1\psi$ ,  $\mathbf{C}_2\psi = \lambda_2\psi$ . Then  $\psi$  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 [44]) 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  $\lambda_1 > 0$ ,  $\psi$  is also an eigenvector of  $\mathbf{C}_1\%\mathbf{C}_2$ , with eigenvalue:  $\lambda_{\%} = \lambda_1^{-1}\lambda_2^2$ .*

*Proof.* See Appendix 4.2.

Theorem 2 reveals how  $\%$  emphasizes changes in mode strength: the mapping  $\lambda_2 \mapsto \lambda_2^2/\lambda_1$  amplifies increases in  $\lambda_2$  relative to  $\lambda_1$  and suppresses decreases. In other words, if a coordinated network pattern  $\psi$  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 4.2 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.

## 1.4 Similarity and difference operators as low- and high-pass filters

We now leverage the spectral properties in Sec. 1.3 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

$$\mathbf{C}_n^{\text{LP}} = \mathbf{C}_n\#\mathbf{C}_{n+1}, \quad n = 1, \dots, N - 1, \quad (1.6)$$

$$\mathbf{C}_n^{\text{HP}} = \mathbf{C}_n\%\mathbf{C}_{n+1}, \quad n = 1, \dots, N - 1. \quad (1.7)$$

The similarity operator  $\#$  computes the geodesic midpoint, so  $\mathbf{C}_n^{\text{LP}}$  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  $10 \times 10$  SPSD matrices such that correlations within components 1–3 oscillate sinusoidally,

$$C_n(k, l) = 0.5 \sin(2\pi f n) + \varepsilon_{k,l}, \quad \varepsilon_{k,l} \sim U(-0.01, 0.01) \text{ i.i.d.}, \quad k, l \in \{1, 2, 3\}, \quad (1.8)$$

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

$$C_n(k, l) = \varepsilon_{k,l}, \quad \text{otherwise,} \quad C_n(k, k) = 1, \quad (1.9)$$

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,

$$r_n \triangleq \frac{\lambda_{\max}(\mathbf{C}_n)}{\text{Tr}(\mathbf{C}_n)}. \quad (1.10)$$

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,

$$D(\{\mathbf{C}_n\}_{n=1}^N) \triangleq \frac{1}{N} \sum_{n=1}^N (1 - r_n), \quad (1.11)$$

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 dB),

$$H(f) \triangleq 20 \log_{10} \left( \frac{D(\{\mathbf{C}_n^{\text{out}}\})}{D(\{\mathbf{C}_n^{\text{in}}\})} \right). \quad (1.12)$$

With this construction,  $H(f)$  reports how strongly the filter attenuates (or preserves) structured sinusoidal connectivity dynamics at frequency  $f$ . Fig 1.1B and C show the frequency responses of the resulting low- and

high-pass filters, computed numerically from  $N = 256$  matrices generated according to Eq. (1.8). As expected, the low-pass filter suppresses higher frequencies, whereas the high-pass filter suppresses lower frequencies.

## 1.5 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. (1.6)–(1.7), 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. 1.1D). 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 matrices 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 evolution is no longer represented within a bin.




Figure 1.1: (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  $C_2$  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 down-sampling by a factor of two.

### Dynamic drivers detection from frequency-resolved connectivity modes.

Each matrix  $\hat{C}_f$  summarizes FC structure concentrated within frequency bin  $f$ . For each bin we compute the dominant eigenpair

$$\hat{C}_f \psi_f = \lambda_f \psi_f. \quad (1.13)$$

and interpret  $\psi_f$  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 are dominated by diffuse, noise-like structure. To quantify whether  $\psi_f$  is concentrated on a small subset of components—consistent with a structured subnetwork—we compute the normalized entropy of its magnitude profile:

$$h_f \triangleq -\frac{1}{\log d} \sum_{i=1}^d p_f[i] \log p_f[i], \quad p_f[i] = \frac{|\psi_f[i]|}{\sum_{j=1}^d |\psi_f[j]|}. \quad (1.14)$$

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

$$\mathcal{F}^* \triangleq \{ f : f < f_{\text{cutoff}}, \quad h_f \leq h_{\text{cutoff}} \}, \quad (1.15)$$

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,

$$D_{\text{score}}(i) \triangleq \sum_{f \in \mathcal{F}^*} \mathbf{1}_f[i], \quad (1.16)$$

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

$$DD \triangleq \{ i : D_{\text{score}}(i) > M \}. \quad (1.17)$$

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

We illustrate RONI on a controlled toy example that contains two temporally distinct subnetworks and background noise. Specifically, we generate a sequence of  $N = 256$ ,  $20 \times 20$  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. 1.2A):

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

---

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

---

**Require:** SPSD matrix sequence  $\mathbf{C}_n$  of length  $N$ , cutoff frequency  $f_{\text{cutoff}}$ , cutoff entropy  $h_{\text{cutoff}}$ , 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 (1.6), (1.7)
  - 5:         Downsample the resulting signals by 2
  - 6:     **end for**
  - 7: **end for**
  - 8: **for** each frequency bin  $f$  **do**
  - 9:     Extract the eigenvector corresponding to the highest eigenvalue (1.13)
  - 10:     Compute normalized entropy  $h_f$  using (1.14)
  - 11: **end for**
  - 12: Retain the informative subset of frequency bins using (1.15)
  - 13: Calculate the driver score using (1.16) and identify dynamic drivers using (1.17)
- 

Fig. 1.2B 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  $\sim 70$  respectively (Fig. 1.2C, top). The 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. (1.16) assigns the highest scores to components 1–3 and 7–9 (Fig. 1.2D), 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. 1.2B) 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. 1.2E), whereas the non-driver sequence yields an unstructured embedding (Fig. 1.2F). Together, these re-

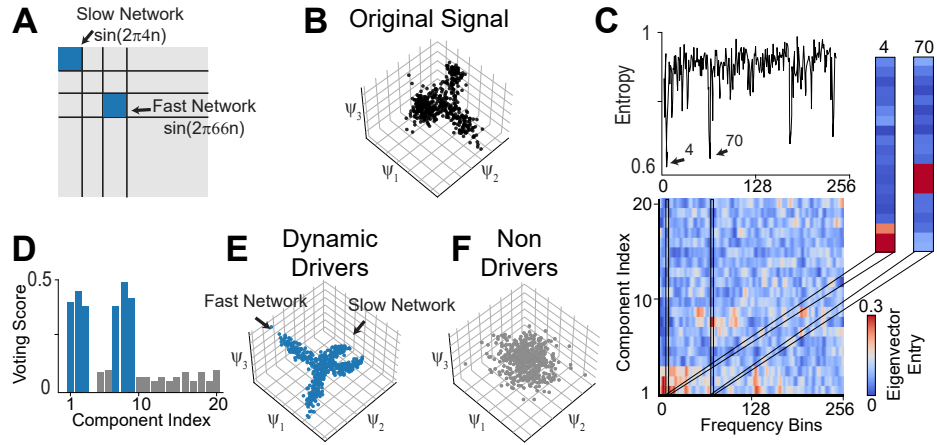


Figure 1.2: Toy example: a simulated sequence of SPSD matrices. (A) Schematic representation of the toy example entries, where entries 1–3 exhibit slow dynamics and entries 7–9 exhibit fast dynamics; gray areas indicate unstructured noise. (B) Diffusion embedding of the sequence, computed using all components. (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.

sults demonstrate that RONI isolates components carrying structured multi-scale connectivity dynamics while suppressing unstructured noise.

## 1.6 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.

### 1.6.1 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 4.8.1). 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. 1.3A and B, respectively.

For a detailed description of the recording setup and pre-processing, see Appendix 4.8.1. 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. (1.1)–(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  $h_{\text{cutoff}}$  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. 1.3C 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. 1.3D). Together, 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.

### 1.6.2 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 [29] (see Appendix 4.8.2 for a detailed description of the experimental setup). We included four animals, each contributing  $\sim 1,400$  successful trials over  $\sim 30$  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. 1.4A). 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. 1.4B).

Finally, we aggregated driver scores across animals to obtain a region-

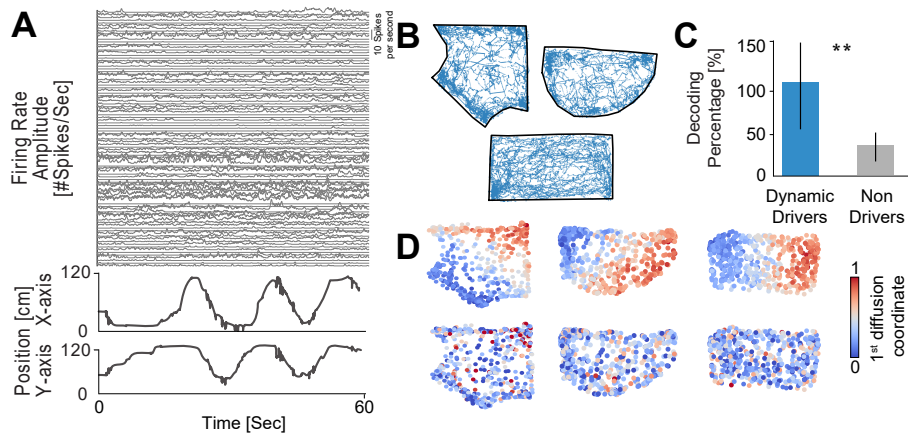


Figure 1.3: **Spatially informative hippocampal dynamics.** (A) Top: Example firing-rate traces of neurons during a 60s segment extracted from a 20 min recording session. Bottom: Corresponding mouse position trace during the same 60s interval. (B) Three example arena outlines (black) with the full mouse 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  $\pm$  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.




Figure 1.4: **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  $\pm$  SEM,  $n = 4$ ). (C) Driver scores highlight cortical regions that shape dynamics.

level summary of contributions to learning-related FC dynamics. Driver scores consistently highlighted retrosplenial cortex, medial frontal areas, and specific somatosensory regions (Fig. 1.4C), 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.

### 1.6.3 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. 1.5A). 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 4.8.3 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. 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




Figure 1.5: **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 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).

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. 1.5C), 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. 1.5D). 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.

#### 1.6.4 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 [5]) in which subjects imagined left- versus right-hand movement. Trials were band-pass filtered (8–30Hz),




Figure 1.6: **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.

and correlation matrices were computed over a 3sec window starting 0.5sec before the go cue, yielding one SPSD matrix sequence per movement class.

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. 1.6). These findings illustrate that RONI can extract an interpretable driver structure in human EEG that aligns with known neurophysiology.

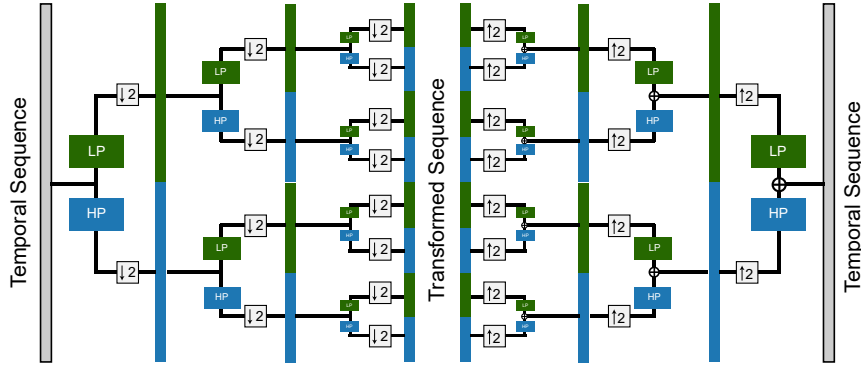


Figure 1.7: Illustration of the filtering scheme for analysis and synthesis of 1-dimensional signals

## 1.7 Reconstruction on the Riemannian manifold

In the previous section, we showed that applying the similarity and difference operators onto temporal sequences of matrices is equivalent to low- and high-pass filtering. Applying them recursively, along with down-sampling by two, yields a multiresolution decomposition of a sequence of SPSD matrices into frequency bins. For 1-dimensional signals, synthesis of the original signals from their decomposition is achieved using a pair of synthesis filters, applied recursively with up-sampling, under conditions for perfect reconstruction that are well defined. In this section, we present the theoretical basis for the synthesis of SPSD matrices in the Riemannian framework. We present the mathematical formulation, along with the necessary steps and conditions for achieving perfect reconstruction.

### 1.7.1 Riemannian wavelet reconstruction

Fig. 1.7 presents a block diagram for analysis and synthesis of 1-dimensional signals. In the previous section, we presented the Riemannian analogue for the analysis part (left). To design a parallel scheme for Riemannian synthesis, several key operations must be defined in the geometric context: upsampling, i.e., insertion of zeros, synthesis filters, and the summation operation (which in the 1D case is a simple pointwise addition). We will present a pipeline that yields perfect reconstruction when the eigenvectors

of subsequent matrices are identical.

**Theorem 3** *Given two SPSD matrices  $\mathbf{W}_1$  and  $\mathbf{W}_2$  with common eigenvectors and possibly different eigenvalues, and their correspond HP and LP filter results.  $\mathbf{C} = \mathbf{W}_1 \# \mathbf{W}_2$ ,  $\mathbf{D} = \mathbf{W}_1 \% \mathbf{W}_2$  in order to reconstruct the original matrices from  $\mathbf{C}$  and  $\mathbf{D}$  follow the following steps:*

$$\mathbf{E} = \mathbf{C} \# \mathbf{I} = \gamma_{\mathbf{C} \rightarrow \mathbf{I}}(0.5)$$

$$\mathbf{F} = \mathbf{D} \% \mathbf{I} = \gamma_{\mathbf{D} \rightarrow \mathbf{I}}(2)$$

$$\mathbf{G} = \mathbf{I} \% \mathbf{D} = \gamma_{\mathbf{I} \rightarrow \mathbf{D}}(2)$$

$$\mathbf{H} = \gamma_{\mathbf{E} \rightarrow \mathbf{F}}(1/9)$$

$$\mathbf{J} = \gamma_{\mathbf{E} \rightarrow \mathbf{G}}(1/9)$$

$$\mathbf{W}_1 = \gamma_{\mathbf{I} \rightarrow \mathbf{H}}(3)$$

$$\mathbf{W}_2 = \gamma_{\mathbf{I} \rightarrow \mathbf{J}}(1.5)$$

*Proof.* See Appendix 4.6. A geometric illustration of Thrm 3 is presented in 1.8.

Thrm. 3 shows that we can reconstruct the original two matrices from their low- and high-pass outcomes. Analogue to the 1-dimensional case, insertion of zeros is replaced here by the insertion of the identity matrix. The synthesis filters are the same low- and high-pass filters defined in Sec. 1.5. The "summation" operation for interleaving the low and high frequencies translates in the Riemannian case to two stages: (i) yielding  $\mathbf{H}$  and  $\mathbf{J}$  and then  $\mathbf{W}_i$ . Applying these operations recursively to the decomposed sequence will yield the reconstructed signal.

Further investigation is required to better understand the limitations of this approach, specifically, what is the mathematical relationship between the reconstruction error and the difference between eigenvectors of subsequent matrices? This scheme has the potential to enhance the interpretability of RONI, the selection of informative frequency bins, and different temporal scales. These important issues remain outside of the scope of this work.




Figure 1.8: Illustration of the Riemannian reconstruction pipeline. **Blue** matrices represent the original inputs:  $\mathbf{W}_1$ ,  $\mathbf{W}_2$ , and the identity matrix  $\mathbf{I}$ . **Green** matrices show the results of applying the low-pass (LP) and high-pass (HP) filters, yielding  $\mathbf{C}$  and  $\mathbf{D}$ . **Orange** matrices correspond to further filtering operations applied to  $\mathbf{C}$ ,  $\mathbf{D}$ , and  $\mathbf{I}$ , resulting in intermediate matrices  $\mathbf{E}$ ,  $\mathbf{F}$ , and  $\mathbf{G}$ . **Red** matrices indicate the results of the "addition" operations: combining  $\mathbf{E}$  with  $\mathbf{F}$ , and  $\mathbf{E}$  with  $\mathbf{G}$ , which reconstruct approximations of the original inputs up to a scaling factor.

## Chapter 2

# A Finsler-Randers Framework for Directed Dynamics on the SPD Manifold

In this chapter, we extend the geometric analysis of correlation matrices by incorporating directional dynamics into the manifold framework. While symmetric positive definite (SPD) matrices naturally capture correlational structure and admit a well-studied Riemannian geometry, directional influences such as those derived from Granger causality [45] introduce asymmetries that are not easily represented within this space.

To address this challenge, we propose a Finsler-Randers [55, 41, 38] inspired approach that combines the Affine-Invariant Riemannian Metric (AIRM) with an additional linear “wind” term reflecting directional interactions. This construction separates the problem into two complementary components: a Riemannian metric governing the intrinsic geometry of the SPD manifold, and a Euclidean perturbation encoding the directional forces. The resulting framework enables the definition of a directed distance that integrates both symmetric and asymmetric components of the system’s dynamics, and provides a principled tool for downstream tasks such as dimensionality reduction and visualization.

## 2.1 Problem formulation

In addition to correlational structure proposed in chapter 1, we assume that the system exhibits a certain directionality in its temporal evolution. Specifically, we postulate the existence of a matrix  $A$  such that the dynamics are approximately governed by

$$x_t \approx \sum_{i=1}^{\tau} \mathbf{A}_i x_{t-i}.$$

Where  $\tau \geq 1$  is some delay coefficient. Our objective is to detect and characterize these dynamics while respecting the underlying geometry of the problem. A key limitation arises from the fact that the matrices associated with directional dynamics are not necessarily symmetric or positive definite, making the formulation presented in 1 unfit for this case.

To address this challenge, we propose a Finsler geometric framework that separates the problem into two complementary components:

1. A Riemannian metric, which defines the manifold on which the dynamics evolve.
2. A Euclidean "force" induced by the directional dynamics, which perturbs and drives the system's movement along the manifold.

This decomposition allows us to capture both the symmetric correlational structure and the asymmetric directional influences within a unified geometric framework.

## 2.2 A Finsler-Randers metric on the SPD manifold

Let  $S_{++}^n$  denote the manifold of  $n \times n$  symmetric positive definite (SPD) matrices, which can be endowed with the natural Affine-Invariant Riemannian Metric (AIRM). For a tangent vector  $\mathbf{V} \in T_S S_{++}^n$  at  $\mathbf{S} \in S_{++}^n$ , the AIRM norm is

$$\|\mathbf{V}\|_S = \sqrt{\text{Tr}(\mathbf{S}^{-1} \mathbf{V} \mathbf{S}^{-1} \mathbf{V})}.$$

To incorporate directional information derived from the system dynamics, we introduce a directed matrix  $\mathbf{G} \in \mathbb{R}^{n \times n}$ , which will act as a functional on

tangent vectors. Combining the AIRM with this directional term gives a Finsler-Randers type metric:

$$F(\mathbf{S}, \mathbf{V}, \mathbf{G}) = \underbrace{\|\mathbf{V}\|_S}_{\text{AIRM term}} + \underbrace{\mu \|\mathbf{G}^\top \mathbf{V}\|_F}_{\text{wind term}}$$

where  $\mu > 0$  is a scalar controlling the strength of the directional influence. Intuitively:

- The first term measures the “intrinsic length” of the tangent vector  $\mathbf{V}$  according to the SPD manifold’s geometry.
- The second term biases motion in the direction of  $\mathbf{G}$ , allowing the metric to incorporate observed directional dynamics.

**Theorem 4** *Let  $\mathbf{S} \in \mathcal{S}_{++}(n)$  and  $\mathbf{G} \in \mathbb{R}^{n \times n}$ . Define*

$$F(\mathbf{S}, \mathbf{V}, \mathbf{G}) := \underbrace{\|\mathbf{V}\|_S}_{\text{AIRM term}} + \underbrace{\mu \|\mathbf{G}^\top \mathbf{V}\|_F}_{\text{wind term}}, \quad \mathbf{V} \in T_{\mathcal{S}_{++}}(n)$$

where  $\|\mathbf{V}\|_S := \sqrt{\text{Tr}(\mathbf{S}^{-1} \mathbf{V} \mathbf{S}^{-1} \mathbf{V})}$  is the AIRM norm. Define the constant

$$c := \sup_{\mathbf{V} \neq 0} \frac{\|\mathbf{G}^\top \mathbf{V}\|_F}{\|\mathbf{V}\|_S} \in [0, \infty]$$

Then, for any  $\mu$  with  $|\mu| < 1/c$ , the function  $F(\mathbf{S}, \cdot)$  is a Finsler metric on  $T_{\mathcal{S}_{++}}(n)$  in the sense of a Minkowski norm: it is positive definite, 1-homogeneous, and satisfies the triangle inequality.

*Proof.* See Appendix 4.7.

This construction generalizes the AIRM to a Finsler-Randers setting, integrating both the manifold geometry and the system’s directional interactions into a unified metric framework.

## 2.3 Geodesics and distance under the Finsler-Randers metric

Given a Finsler-Randers metric  $F(\mathbf{S}, \mathbf{V}, \mathbf{G})$  on the SPD manifold  $S_{++}^n$ , the geodesic is the curve  $\gamma : [0, 1] \rightarrow S_{++}^n$  that minimizes the total length

$$L[\gamma] = \int_0^1 F(\gamma(t), \dot{\gamma}(t)) dt$$

where  $\dot{\gamma}(t) \in T_{\gamma(t)}S_{++}^n$  is the tangent vector along the curve at time  $t$ .

For the SPD manifold endowed with AIRM, the geodesic connecting two points  $\mathbf{C}_1, \mathbf{C}_2 \in S_{++}^n$  is given in closed form by

$$\gamma_{\text{AIRM}}(t) = \mathbf{C}_1^{1/2} (\mathbf{C}_1^{-1/2} \mathbf{C}_2 \mathbf{C}_1^{-1/2})^t \mathbf{C}_1^{1/2}, \quad t \in [0, 1]$$

The tangent vector at  $\mathbf{P}_1$  along this geodesic is

$$\mathbf{V} = \dot{\gamma}_{\text{AIRM}}(0) = \log(\mathbf{C}_1^{-1/2} \mathbf{C}_2 \mathbf{C}_1^{-1/2}) \in T_{\mathbf{P}_1} S_{++}^n$$

where  $\log$  denotes the matrix logarithm.

When incorporating the directional “wind” term from directional matrix  $\mathbf{G}$ , the exact geodesic under the Finsler-Randers metric generally does not admit a closed-form solution. However, an approximation is to evaluate the Finsler norm along the AIRM geodesic, yielding an approximate directed distance:

$$\delta_F(\mathbf{C}_1, \mathbf{C}_2) \approx \underbrace{\|\log(\mathbf{C}_1^{-1/2} \mathbf{C}_2 \mathbf{C}_1^{-1/2})\|_F}_{\text{AIRM distance}} + \mu \underbrace{\|\mathbf{G}^\top \log(\mathbf{C}_1^{-1/2} \mathbf{C}_2 \mathbf{C}_1^{-1/2})\|_F}_{\text{directional bias}} \quad (2.1)$$

- The first term represents the intrinsic distance along the SPD manifold according to the AIRM (and can be extended to the SPSD case).
- The second term biases the distance in the direction of observed causal interactions, effectively shortening or lengthening the path depending on alignment with  $\mathbf{G}$ .

This approach generalizes the classical geodesic distance to account for both the manifold geometry and directional dynamics, resulting in a Finsler-Randers directed distance on the SPSD manifold. See algorithm 2.

---

**Algorithm 2 Finsler-Randers Distance Embedding for Dynamics Detection**

---

**Require:** Signal dataset of length  $T$ , weight parameter  $\mu$

**Ensure:** Low-dimensional embedding  $\mathbf{Y} \in \mathbb{R}^{N \times d}$  preserving geometry and directionality

- 1: Compute correlation matrices  $\{\mathbf{C}_{ij}\}_{i=1}^N \subset S_+^n$  from the signals
  - 2: Compute Directional matrices  $\{\mathbf{G}_i\}_{i=1}^N$  from the same signals
  - 3: **for** each pair  $(i, j)$  with  $1 \leq i, j \leq N$  **do**
  - 4:     Compute the AIRM distance and the directed distance using (2.1)
  - 5: **end for**
  - 6: Normalize  $d_{\text{AIRM}}$  and  $d_{\text{direct}}$  to  $[0, 1]$  using min-max scaling
  - 7: Form the combined Finsler-Randers distance:  
$$D_{\text{full}}(i, j) = d_{\text{AIRM}}(i, j) + \mu d_{\text{direct}}(i, j)$$
  - 8: Construct the diffusion kernel from  $D_{\text{full}}$
  - 9: Apply diffusion maps to obtain the embedding  $\mathbf{Y} \in \mathbb{R}^{N \times d}$
-

## 2.4 Results

### 2.4.1 Directional dynamic applied to cortical dynamics

We evaluated the performance of the proposed method and analyzed wide-field calcium imaging data from mice performing an auditory decision-making task described in Appendix 4.8.2. As in the previous section, we computed undirected correlation matrices across 20 consecutive trials. In addition, we used Granger causality to estimate directional matrices indicating which brain regions are "Granger causal" to which.

We performed dimensionality reduction on both descriptions of the network: (i) diffusion map [8] for the undirected correlation matrices, (ii) magnetic diffusion map [21] for the directional approach combining correlations and Granger causality matrices, as proposed in the previous section. We trained a linear model that maps each embedded representation to the session index, using 20 components of the embedded symmetric correlations and 20 components (10 real and 10 imaginary) for the magnetic diffusion. In all cases, the directional description outperformed the undirected one (see Table 2.1). Therefore, the directional embedding successfully captures dynamic changes in the brain that correlations alone cannot, highlighting the method's potential.

	Undirected	Directed
<b>Mean Pan-Neural</b>	0.74 $\pm$ 0.02	<b>0.79 <math>\pm</math> 0.03</b>
<b>Mean PT Neurons</b>	0.74 $\pm$ 0.10	<b>0.75 <math>\pm</math> 0.11</b>
<b>Mean IT Neurons</b>	0.58 $\pm$ 0.09	<b>0.78 <math>\pm</math> 0.04</b>
<b>Mean</b>	0.69 $\pm$ 0.05	<b>0.77 <math>\pm</math> 0.03</b>

Table 2.1: Accuracy scores ( $R^2$ ) for prediction of session index.

## Chapter 3

# Conclusion

Learning is fundamentally a dynamic process, reflected not only in changes in neural activity but in the evolving patterns of interaction between neural populations. In this thesis, we developed a geometric framework for analyzing such dynamics by treating functional connectivity as a structured temporal signal evolving on the manifold of symmetric positive semi-definite matrices. By respecting the intrinsic geometry of this space, our approach enables principled analysis of connectivity dynamics that is both interpretable and scalable.

Our first contribution, RONI (Riemannian filtering Of Network Interactions), introduces an intrinsic multi-resolution decomposition for SPSD matrices. By defining similarity and difference operators through geodesic constructions, we derived Riemannian analogues of low-pass and high-pass filters that preserve positive semi-definiteness at every stage. This construction enabled a wavelet-packet-like representation of functional connectivity trajectories, providing access to frequency-resolved network structure. Crucially, the spectral properties of this decomposition allowed us to identify dynamic drivers, specific neural components that consistently dominate informative connectivity modes across temporal scales.

Applying RONI across multiple datasets demonstrated its generality and interpretability. In hippocampal recordings, dynamic drivers concentrated spatial information and shaped the geometry of connectivity trajectories during navigation. In cortex-wide imaging, learning-related reorganization was driven by a reproducible subset of cortical regions. At the subcellular scale, dendritic drivers preferentially supported task-relevant motor encod-

ing, revealing a structured organization across neurons. Finally, in human EEG, RONI recovered contralateral sensorimotor organization during motor imagery, aligning with well-established neurophysiology. Across scales and modalities, the framework consistently isolated meaningful subnetworks while suppressing unstructured variability.

Beyond symmetric connectivity, we introduced a complementary Finsler–Randers framework to incorporate directionality and asymmetry into the geometric analysis of neural dynamics. This extension provides a principled way to integrate causal influences into the study of connectivity trajectories, enabling directional distances and embeddings that go beyond traditional Riemannian approaches. Together, the Riemannian and Finsler perspectives form a unified geometric toolkit for studying connectivity across time, scale, and direction.

Several avenues for future work remain. Further theoretical analysis is needed to fully characterize reconstruction guarantees and stability properties of the Riemannian wavelet framework. Methodologically, extending these ideas to other measures of connectivity and to higher-order interactions represents a promising direction. Finally, integrating these geometric tools with predictive models of behavior and learning may offer deeper insight into the mechanisms by which neural circuits adapt.

In summary, this thesis establishes a geometric and interpretable framework for analyzing dynamic functional connectivity. By bridging differential geometry with experimental neuroscience, it provides new tools for understanding how neural networks reorganize during learning and offers a foundation for future advances in the analysis of complex brain dynamics.

# Chapter 4

## Appendix

### 4.1 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 [4].

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

First, we align the subspaces by computing the SVD  $\mathbf{V}_1^\top \mathbf{V}_2 = \mathbf{O}_1^\top \mathbf{O}_2$ , 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

$$\mathbf{U}(t) = \mathbf{U}_1 \cos(t) + \mathbf{X} \sin(t), \quad (4.1)$$

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

$$\mathbf{R}(t) = \mathbf{R}_1^{1/2} \left( \mathbf{R}_1^{-1/2} \mathbf{R}_2 \mathbf{R}_1^{-1/2} \right)^t \mathbf{R}_1^{1/2}. \quad (4.2)$$

Finally, the approximate SPSD geodesic is reconstructed as

$$\tilde{\gamma}_{\mathbf{C}_1 \rightarrow \mathbf{C}_2}(t) = \mathbf{U}(t) \mathbf{R}(t) \mathbf{U}(t)^\top. \quad (4.3)$$

From this we can approximate the squared geodesic distance as:

$$d^2(\mathbf{C}_1, \mathbf{C}_2) \approx \|\cdot\|_F^2 + \left\| \log \left( \mathbf{R}_1^{-1/2} \mathbf{R}_2 \mathbf{R}_1^{-1/2} \right) \right\|_F^2 \quad (4.4)$$

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

## 4.2 Proofs for spectral analysis

The proof is given for SPD matrices. The result extends to SPSD matrices through the geodesic continuation, defined in Appendix 4.1.

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

*Proof.*

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

## 4.3 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  $\varepsilon$ -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  $\varepsilon > 0$  [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 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 5** [44] *Suppose there exists an eigen-pair  $(\lambda_k, \psi_k)$  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  $\psi_1$  is a corresponding  $\epsilon$ -pseudo-eigenvector of  $\mathbf{S}$ .*

**Theorem 6** *Suppose there exists an eigenpair  $\lambda_k$  and  $\psi_k$  of  $\mathbf{C}_k$  for  $k = 1, 2$  so that  $\psi_1 = \psi_2 + \psi_\epsilon$ , where  $\|\psi_\epsilon\|_2 \leq \frac{1}{(\sigma_{\max}(\mathbf{C}_2\mathbf{C}_1^{-1}) + \lambda_2\lambda_1^{-1})\tilde{\lambda}_2^{\max}}\epsilon$ , where  $\sigma_{\max}$  denotes the maximum eigenvalue, for a small  $\epsilon > 0$ , and where  $\tilde{\lambda}_2^{\max} = \|\mathbf{C}_2 - \lambda_2\mathbf{I}\|$ . Then we have  $\lambda_2^2\lambda_1^{-1} \in \Lambda_\epsilon(\mathbf{Q})$ . Specifically, we have  $\|(\mathbf{Q} - \lambda_2^2\lambda_1^{-1}\mathbf{I})\psi_1\|_2 \leq \epsilon$ , and  $\psi_1$  is a corresponding  $\epsilon$ -pseudo-eigenvector of  $\mathbf{Q}$ .*

*Proof.* Using eqn. (4.6) from Appendix 4.5, 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  $\psi_1$  is an eigenvector of  $\mathbf{C}_1$  with eigenvalue  $\lambda_1$ , 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  $\psi_1$  is an  $\epsilon$ -pseudo-eigenvector of  $(\mathbf{C}_2\mathbf{C}_1^{-1})^2$ . By direct expansion, we have:

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

This implies:

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

Multiplying both sides by  $\lambda_1$ , we get:

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

Taking the  $\ell_2$  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  $\varepsilon$  pseudo-eigenvalue of  $\mathbf{Q}$ , where  $\psi_1$  is a corresponding  $\varepsilon$  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.

## 4.4 Proof mutually inverse

**Proposition 2.** Let  $\mathbf{C}_1$  and  $\mathbf{C}_2$  be SPD matrices, the operators defined in (1.4)(1.5) are mutually inverse, meaning:

$$\begin{aligned} \mathbf{C}_2 \% (\mathbf{C}_1 \# \mathbf{C}_2) &= \mathbf{C}_1 \\ \mathbf{C}_2 \# (\mathbf{C}_2 \% \mathbf{C}_1) &= \mathbf{C}_1 \end{aligned} \tag{4.5}$$

The same relations hold for SPSD matrices by continuity of the extended geodesic (Appendix 4.1).

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

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

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

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

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

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

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

## 4.5 Equivalent form

**Proposition 3.** Equivalent form of the filters. We have

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

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

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

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  $\Lambda^M$  and  $\mathbf{V}^M$  be the eigenvalue and eigenvector matrices of  $\mathbf{M}$ , respec-

tively. 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:

$$\begin{aligned} \mathbf{M}^p &= \mathbf{V}^M (\Lambda^M)^p (\mathbf{V}^M)^T, \\ \tilde{\mathbf{M}}^p &= \mathbf{C}_1^{0.5} \mathbf{V}^M (\Lambda^M)^p (\mathbf{V}^M)^T \mathbf{C}_1^{-0.5}. \end{aligned}$$

From this, we conclude:

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

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:

$$\mathbf{C}_1^{0.5} (\mathbf{C}_1^{-0.5} \mathbf{C}_2 \mathbf{C}_1^{-0.5})^p \mathbf{C}_1^{0.5} = (\mathbf{C}_2 \mathbf{C}_1^{-1})^p \mathbf{C}_1. \quad (4.6)$$

## 4.6 Proofs for reconstruction

**Theorem 3.** Given two SPSD matrices  $\mathbf{W}_1$  and  $\mathbf{W}_2$  and their correspond HP and LP filter results.  $\mathbf{C} = \mathbf{W}_1 \# \mathbf{W}_2$ ,  $\mathbf{D} = \mathbf{W}_1 \% \mathbf{W}_2$  in order to recon-

struct the original matrices from  $\mathbf{C}$  and  $\mathbf{D}$  follow the following steps:

$$\begin{aligned}
\mathbf{E} &= \mathbf{C}\#\mathbf{I} = \gamma_{\mathbf{C}\rightarrow\mathbf{I}}(0.5) \\
\mathbf{F} &= \mathbf{D}\%\mathbf{I} = \gamma_{\mathbf{D}\rightarrow\mathbf{I}}(2) \\
\mathbf{G} &= \mathbf{I}\%\mathbf{D} = \gamma_{\mathbf{I}\rightarrow\mathbf{D}}(2) \\
\mathbf{H} &= \gamma_{\mathbf{E}\rightarrow\mathbf{F}}(1/9) \\
\mathbf{J} &= \gamma_{\mathbf{E}\rightarrow\mathbf{G}}(1/9) \\
\mathbf{W}_1 &= \gamma_{\mathbf{I}\rightarrow\mathbf{H}}(3) \\
\mathbf{W}_2 &= \gamma_{\mathbf{I}\rightarrow\mathbf{J}}(1.5)
\end{aligned}$$

*Proof.*

Using the equivalent form 4.5, we get:

$$\begin{aligned}
\gamma_{\mathbf{E}\rightarrow\mathbf{F}}(p) &= (\mathbf{F}\mathbf{E}^{-1})^p \mathbf{E} \\
\mathbf{E} &= ((\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2)^{0.5}, \quad \mathbf{F} = \mathbf{W}_2^{-1} \mathbf{W}_1 \mathbf{W}_2^{-1}
\end{aligned}$$

Using the expressions for  $\mathbf{E}$  and  $\mathbf{F}$ :

$$(\mathbf{W}_2^{-1} \mathbf{W}_1 \mathbf{W}_2^{-1} ((\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2)^{-0.5})^p ((\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2)^{0.5}$$

Let's look at the logarithm of the expression:

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

Since the matrices are positive definite, we can use the Baker-Campbell-Hausdorff (BCH) formula [24] to approximate the matrix logarithm. Because we are dealing with time series, we assume that the matrices have similar eigenvectors (i.e., they almost commute). Empirically, this yields a small reconstruction error.

$$\begin{aligned}
&\approx p \cdot \log \left( \mathbf{W}_2^{-1} \mathbf{W}_1 \mathbf{W}_2^{-1} ((\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2)^{-0.5} \right) + 0.5 \cdot \log \left( (\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2 \right) \\
&\approx p \cdot \log \left( \mathbf{W}_2^{-1} \mathbf{W}_1 \mathbf{W}_2^{-1} \right) + 0.5(1-p) \cdot \log \left( (\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2 \right)
\end{aligned}$$

$$\approx -2p \cdot \log(\mathbf{W}_2) + p \cdot \log(\mathbf{W}_1) + 0.5(1-p) \cdot (0.5 \cdot \log(\mathbf{W}_1) - 0.5 \cdot \log(\mathbf{W}_2) + \log(\mathbf{W}_2))$$

$$= -2p \cdot \log(\mathbf{W}_2) + p \cdot \log(\mathbf{W}_1) + 0.25(1-p) \cdot \log(\mathbf{W}_1) + 0.25(1-p) \cdot \log(\mathbf{W}_2)$$

$$= 0.25(1+3p) \cdot \log(\mathbf{W}_1) + 0.25(1-9p) \cdot \log(\mathbf{W}_2)$$

For  $p = \frac{1}{9}$ , we get:

$$= \frac{1}{3} \cdot \log(\mathbf{W}_1)$$

Taking the exponential:

$$= \mathbf{W}_1^{1/3}$$

So for  $p = \frac{1}{9}$ , we get:

$$\gamma_{\mathbf{E} \rightarrow \mathbf{F}}(1/9) = (\mathbf{F}\mathbf{E}^{-1})^{1/9}\mathbf{E} = \mathbf{W}_1^{1/3}$$

Similarly, define:

$$\gamma_{\mathbf{E} \rightarrow \mathbf{G}}(1/9) = (\mathbf{G}\mathbf{E}^{-1})^p \mathbf{E}$$

where:

$$\mathbf{G} = \mathbf{W}_2 \mathbf{W}_1^{-1} \mathbf{W}_2^2 \mathbf{W}_1^{-1} \mathbf{W}_2$$

Substituting:

$$\left( \mathbf{W}_2 \mathbf{W}_1^{-1} \mathbf{W}_2^2 \mathbf{W}_1^{-1} \mathbf{W}_2 \left( (\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2 \right)^{-0.5} \right)^p \left( (\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2 \right)^{0.5}$$

Taking the logarithm:

$$\log \left( \left( \mathbf{W}_2 \mathbf{W}_1^{-1} \mathbf{W}_2^2 \mathbf{W}_1^{-1} \mathbf{W}_2 \left( (\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2 \right)^{-0.5} \right)^p \left( (\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2 \right)^{0.5} \right)$$

$$\approx p \cdot \log(\mathbf{W}_2 \mathbf{W}_1^{-1} \mathbf{W}_2^2 \mathbf{W}_1^{-1} \mathbf{W}_2) + 0.5(1-p) \cdot \log((\mathbf{W}_1 \mathbf{W}_2^{-1})^{0.5} \mathbf{W}_2)$$

$$\begin{aligned}
&\approx 4p \cdot \log(\mathbf{W}_2) - 2p \cdot \log(\mathbf{W}_1) + 0.25(1-p) \cdot \log(\mathbf{W}_1) + 0.25(1-p) \cdot \log(\mathbf{W}_2) \\
&= 0.25(1-9p) \cdot \log(\mathbf{W}_1) + 0.25(1+15p) \cdot \log(\mathbf{W}_2)
\end{aligned}$$

For  $p = \frac{1}{9}$ , we get:

$$= \frac{2}{3} \cdot \log(\mathbf{W}_2)$$

Taking the exponential:

$$= \mathbf{W}_2^{2/3}$$

Finally after using Proposition 4. we can reconstruct the original matrices.

**Proposition 4.** Let  $\mathbf{W}$  be a SPSD matrix, then:

$$\begin{aligned}
\gamma_{\mathbf{I} \rightarrow \mathbf{W}}(p) &= \mathbf{W}^p \\
\gamma_{\mathbf{W} \rightarrow \mathbf{I}}(p) &= \mathbf{W}^{-p+1}
\end{aligned}$$

*Proof.*

$$\begin{aligned}
\gamma_{\mathbf{I} \rightarrow \mathbf{W}}(p) &= \mathbf{I}^{0.5} (\mathbf{I}^{-0.5} \mathbf{W} \mathbf{I}^{-0.5})^p \mathbf{I}^{0.5} \\
&= \mathbf{W}^p \\
\gamma_{\mathbf{W} \rightarrow \mathbf{I}}(p) &= \mathbf{W}^{0.5} (\mathbf{W}^{-0.5} \mathbf{I} \mathbf{W}^{-0.5})^p \mathbf{W}^{0.5} \\
&= \mathbf{W}^{-p+1}
\end{aligned}$$

**Theorem 7** Consider a vector  $\psi$ , which is an eigenvector of both  $\mathbf{W}_1$  and  $\mathbf{W}_2$ , with possibly different eigenvalues:

$$\mathbf{W}_1 \psi = \lambda_1 \psi, \quad \mathbf{W}_2 \psi = \lambda_2 \psi$$

Then  $\psi$  is also an eigenvector of the points along the geodesic, with the corresponding eigenvalue:

$$\lambda_p = \lambda_1^{1-p} \lambda_2^p$$

*Proof.*

$$\begin{aligned}
\gamma_{\mathbf{w}_1 \rightarrow \mathbf{w}_2}(p)\psi &= \mathbf{W}_1^{0.5} (\mathbf{W}_1^{-0.5} \mathbf{W}_2 \mathbf{W}_1^{-0.5})^p \mathbf{W}_1^{0.5} \psi \\
&= \mathbf{W}_1^{0.5} (\mathbf{W}_1^{-0.5} \mathbf{W}_2 \mathbf{W}_1^{-0.5})^p \sqrt{\lambda_1} \psi \\
&= \mathbf{W}_1^{0.5} \left( \frac{\lambda_2}{\lambda_1} \right)^p \sqrt{\lambda_1} \psi \\
&= \left( \frac{\lambda_2^p}{\lambda_1^{p-1}} \right) \psi \\
&= \left( \lambda_2^p \lambda_1^{1-p} \right) \psi
\end{aligned}$$

This result implies that components which differ between the two matrices (i.e., same eigenvector but different eigenvalues) will be emphasized in the filter, with their relative weight determined by the parameter  $p$ . In contrast, components with similar eigenvalues will remain similarly expressed after filtering.

## 4.7 Proofs for Finsler geometry

**Theorem 4.** Let  $\mathbf{S} \in \mathbb{S}_{++}(n)$  and  $\mathbf{G} \in \mathbb{R}^{n \times n}$ . Define

$$F(\mathbf{S}, \mathbf{V}, \mathbf{G}) := \underbrace{\|\mathbf{V}\|_S}_{\text{AIRM term}} + \mu \underbrace{\|\mathbf{G}^\top \mathbf{V}\|}_F, \quad \mathbf{V} \in T_S \mathbb{S}_{++}(n)$$

where  $\|\mathbf{V}\|_S := \sqrt{\text{Tr}(\mathbf{S}^{-1} \mathbf{V} \mathbf{S}^{-1} \mathbf{V})}$  is the AIRM norm. Define the constant

$$c := \sup_{\mathbf{V} \neq 0} \frac{\|\mathbf{G}^\top \mathbf{V}\|_F}{\|\mathbf{V}\|_S} \in [0, \infty]$$

Then, for any  $\mu$  with  $|\mu| < 1/c$ , the function  $F(\mathbf{S}, \cdot)$  is a Finsler metric on  $T_S \mathbb{S}_{++}(n)$  in the sense of a Minkowski norm: it is positive definite, 1-homogeneous, and satisfies the triangle inequality.

*Proof.*

1. Homogeneity.

For any  $\lambda > 0$ ,

$$F(S, \lambda V, G) = \|\lambda V\|_S + \mu \|\mathbf{G}^\top(\lambda V)\|_F = \lambda \|V\|_S + \mu \lambda \|\mathbf{G}^\top V\|_F = \lambda F(S, V, G)$$

so  $F$  is 1-homogeneous.

2. Positivity.

For  $V \neq 0$ ,  $\|V\|_S > 0$ . By the definition of  $c$ ,

$$\|G^\top V\|_F \leq c\|V\|_S.$$

Hence, for  $0 \leq \mu < 1/c$ ,

$$F(S, V, G) = \|V\|_S + \mu\|G^\top V\|_F \geq (1 - \mu c)\|V\|_S > 0.$$

Also,  $F(S, 0, G) = 0$ . Therefore,  $F$  is positive definite.

3. Triangle inequality.

For any  $U, V \in T_S \mathcal{S}_{++}(n)$ , the AIRM norm satisfies

$$\|U + V\|_S \leq \|U\|_S + \|V\|_S,$$

and the Frobenius norm satisfies

$$\|G^\top (U + V)\|_F \leq \|G^\top U\|_F + \|G^\top V\|_F.$$

Since  $\mu \geq 0$ , multiplying the second inequality by  $\mu$  preserves the inequality direction, and adding the two inequalities gives

$$\begin{aligned} F(S, U + V, G) &= \|U + V\|_S + \mu\|G^\top (U + V)\|_F \\ &\leq (\|U\|_S + \|V\|_S) + \mu(\|G^\top U\|_F + \|G^\top V\|_F) \\ &= F(S, U, G) + F(S, V, G). \end{aligned}$$

Thus the triangle inequality holds.

Since  $F$  is 1-homogeneous, positive definite, and satisfies the triangle inequality, it defines a Minkowski norm on  $T_S \mathcal{S}_{++}(n)$ , and hence a Finsler metric on  $\mathcal{S}_{++}(n)$ .

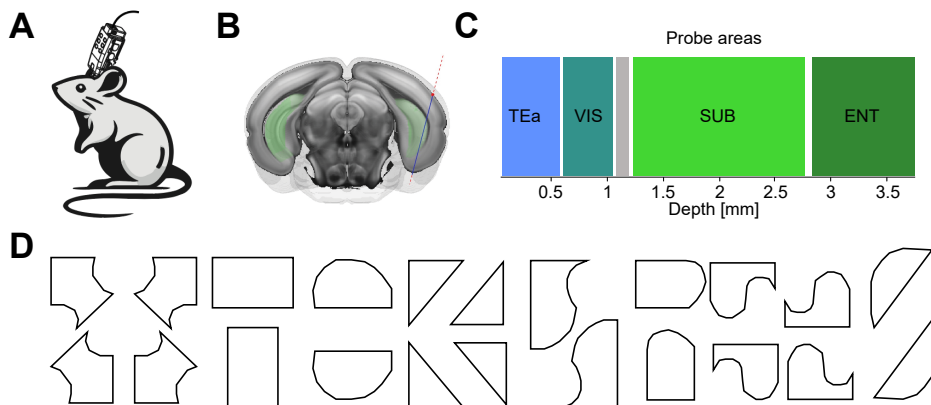


Figure 4.1: (A) Schematic of the mouse and Neuropixels probe used for electrophysiological recordings. (B) Coronal brain section showing the probe insertion site and recording trajectory. (C) Brain regions recorded during the experiment. (D) Outlines of all arena geometries used in the study.

## 4.8 Datasets

### 4.8.1 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  $60 \times 60 \text{ cm}^2$  arena. We considered 38 distinct geometries, four mice, and approximately 150 neurons per mouse. For each experiment, neurons were filtered by minimum mean rate  $\geq 0.1$  spikes per second, inclusion of Kilosort-classified "good" units [33], and exclusion of Bombcell-classified "noise" units [15].

### 4.8.2 Cortical network during learning

Comprehensive details of the widefield imaging are provided in [29] and [10]. 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  $12.5 \times 10.5 \text{ mm}^2$  (Fig. 4.2A). Imaging resolution was  $640 \times 540$  pixels after  $4\times$  spatial binning, resulting in a spatial resolution of about  $20 \mu\text{m}$  per pixel. Mice were Ai93D;Emx-Cre;LSL-tTA animals,




Figure 4.2: (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.

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. 4.2B). 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. 4.2C). 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.




Figure 4.3: (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 $\pm$ SEM) of 8 mice across training days, showing an increase in performance with training.

### 4.8.3 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. 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.

# Bibliography

- [1] Tiago Azevedo, Alexander Campbell, Rafael Romero-Garcia, Luca Passamonti, Richard A.I. Bethlehem, Pietro Liò, and Nicola Toschi. A deep graph neural network architecture for modelling spatio-temporal dynamics in resting-state functional mri data. *Medical Image Analysis*, 79:102471, 2022.
- [2] Hadas Benisty, Daniel Barson, Andrew H Moberly, Sweyta Lohani, Lan Tang, Ronald R Coifman, Michael C Crair, Gal Mishne, Jessica A Cardin, and Michael J Higley. Rapid fluctuations in functional connectivity of cortical networks encode spontaneous behavior. *Nature Neuroscience*, 27(1):148–158, 2024.
- [3] Richard F Betzel, Katherine C Wood, Christopher Angeloni, Maria Neimark Geffen, and Danielle S Bassett. Stability of spontaneous, correlated activity in mouse auditory cortex. *PLOS Computational Biology*, 15(12):e1007360, 2019.
- [4] Silvère Bonnabel and Rodolphe Sepulchre. Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1055–1070, 2010.
- [5] Clemens Brunner, Robert Leeb, Gernot Müller-Putz, Alois Schlögl, and Gert Pfurtscheller. Bci competition 2008–graz data set a. *Journal of Neural Engineering*, 5(1):1–6, 2008.
- [6] Weiwei Chang, Jinping Liu, Liluo Nie, Xiaomin Pang, Zongxia Lv, and Jinou 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.

- [7] Joseph Cichon and Wen-Biao Gan. Branch-specific dendritic ca<sup>2+</sup> spikes cause persistent synaptic plasticity. *Nature*, 520(7546):180–185, 2015.
- [8] Ronald R. Coifman and Stephane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. Special Issue: Diffusion Maps and Wavelets.
- [9] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. Special Issue: Diffusion Maps and Wavelets.
- [10] Joao Couto, Simon Musall, Xiaonan R Sun, Anup Khanal, Steven Gluf, Shreya Saxena, Ian Kinsella, Taiga Abe, John P Cunningham, Liam Paninski, et al. Chronic, cortex-wide imaging of specific cell populations during behavior. *Nature protocols*, 16(7):3241–3263, 2021.
- [11] Zhengwu Dai, Zhenyu Zhang, and Anuj Srivastava. Analyzing dynamical brain functional connectivity as trajectories on space of covariance matrices. *IEEE Transactions on Medical Imaging*, 38(9):2196–2207, 2019.
- [12] Tingting Dan, Zhuobin Huang, Hongmin Cai, Paul J Laurienti, and Guorong Wu. Learning brain dynamics of evolving manifold functional mri data using geometric-attention neural network. *IEEE transactions on medical imaging*, 41(10):2752–2763, 2022.
- [13] Farnaz Zamani Esfahlani, Lisa Byrge, Jacob Tanner, Olaf Sporns, Daniel P Kennedy, and Richard F Betzel. Edge-centric analysis of time-varying functional brain networks with applications in autism spectrum disorder. *NeuroImage*, 280:119591, 2023.
- [14] Farnaz Zamani Esfahlani, Youngheun Jo, Joshua Faskowitz, Lisa Byrge, Daniel P Kennedy, Olaf Sporns, and Richard F Betzel. High-amplitude cofluctuations in cortical activity drive functional connectivity. *Proceedings of the National Academy of Sciences*, 117(45):28393–28401, 2020.
- [15] JMJ Fabre, EH van Beest, AJ Peters, M Carandini, and KD Harris. Bombcell: Automated curation and cell classification of spike-sorted electrophysiology data zenodo, july 2023, 2025.

- [16] Joshua Faskowitz, Farnaz Zamani Esfahlani, Youngheun Jo, Olaf Sporns, and Richard F Betzel. Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nature Neuroscience*, 23:1644–1654, 2020.
- [17] Min Fu, Xinzhu Yu, Ju Lu, and Yi Zuo. Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo. *Nature*, 483(7387):92–95, 2012.
- [18] Amir Ghanayim, Hadas Benisty, Avigail Cohen Rimon, Sivan Schwartz, Sally Dabdoob, Shira Lifshitz, Ronen Talmon, and Jackie Schiller. Vta projections to m1 are essential for reorganization of layer 2-3 network dynamics underlying motor learning. *Nature Communications*, 16(1):200, 2025.
- [19] Ariel Gilad and Fritjof Helmchen. Spatiotemporal refinement of signal flow through association cortex during learning. *Nature communications*, 11(1):1744, 2020.
- [20] Omer Hazon, Victor H Minces, David P Tomàs, Surya Ganguli, Mark J Schnitzer, and Pablo E Jercog. Noise correlations in neural ensemble activity limit the accuracy of hippocampal spatial representations. *Nature communications*, 13(1):4276, 2022.
- [21] Mingzhen He, Fan He, Ruikai Yang, and Xiaolin Huang. Diffusion representation for asymmetric kernels via magnetic transform. *Advances in Neural Information Processing Systems*, 36:53742–53761, 2023.
- [22] Kaihua Jiang, Yang Yi, Lin Li, Hongxin Li, Huijuan Shen, Fangqiao Zhao, Yunpin Xu, and Aibin Zheng. Functional network connectivity changes in children with attention-deficit hyperactivity disorder: A resting-state fmri study. *International Journal of Developmental Neuroscience*, 78:1–6, 2019.
- [23] Rongtao Jiang, Nianming Zuo, Judith M Ford, Shile Qi, Dongmei Zhi, Chuanjun Zhuo, Yong Xu, Zening Fu, Juan Bustillo, Jessica A Turner, et al. Task-induced brain connectivity promotes the detection of individual differences in brain-behavior relationships. *NeuroImage*, 207:116370, 2020.

- [24] S Klarsfeld and JA Oteo. The baker-campbell-hausdorff formula and the convergence of the magnus expansion. *Journal of physics A: mathematical and general*, 22(21):4565, 1989.
- [25] Reinmar Kobler, Jun-ichiro Hirayama, Qibin Zhao, and Motoaki Kawanabe. Spd domain-specific batch normalization to crack interpretable unsupervised domain adaptation in eeg. *Advances in Neural Information Processing Systems*, 35:6219–6235, 2022.
- [26] Manoj Kumar, Gregory Handy, Stylianos Kouvaros, Yanjun Zhao, Lovisa Ljungqvist Brinson, Eric Wei, Brandon Bizup, Brent Doiron, and Thanos Tzounopoulos. Cell-type-specific plasticity of inhibitory interneurons in the rehabilitation of auditory cortex after peripheral damage. *Nature communications*, 14(1):4170, 2023.
- [27] Lukas Kunz, Bernhard P Staresina, Peter C Reinacher, Armin Brandt, Tim A Guth, Andreas Schulze-Bonhage, and Joshua Jacobs. Ripple-locked coactivity of stimulus-specific neurons and human associative memory. *Nature neuroscience*, 27(3):587–599, 2024.
- [28] Sweyta Lohani, Andrew H Moberly, Hadas Benisty, Boris Landa, Miao Jing, Yulong Li, Michael J Higley, and Jessica A Cardin. Spatiotemporally heterogeneous coordination of cholinergic and neocortical activity. *Nature Neuroscience*, 25(12):1706–1713, 2022.
- [29] Simon Musall, Xiaonan R Sun, Hemanth Mohan, Xu An, Steven Gluf, Shu-Jing Li, Rhonda Drewes, Emma Cravo, Irene Lenzi, Chaoqun Yin, et al. Pyramidal cell types drive functionally distinct cortical activity patterns during decision-making. *Nature neuroscience*, 26(3):495–505, 2023.
- [30] Carlos Lopez Naranjo, Fuleah Abdul Razzaq, Min Li, Ying Wang, Jorge F. Bosch-Bayard, Martin A. Lindquist, Anisleidy Gonzalez Mitjans, Ronaldo Garcia, Arielle G. Rabinowitz, Simon G. Anderson, Giuseppe A. Chiarenza, Ana Calzada-Reyes, Trinidad Virues-Alba, Janina R. Galler, Ludovico Minati, Maria L. Bringas Vega, and Pedro A. Valdes-Sosa. Eeg functional connectivity as a riemannian mediator: An application to malnutrition and cognition. *Human Brain Mapping*, 45(7):e26698, 2024.

- [31] Xuan Son Nguyen, Shuo Yang, and Aymeric Histace. Neural networks on symmetric spaces of noncompact type. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [32] Leonardo Novelli and Adeel Razi. A mathematical perspective on edge-centric brain functional connectivity. *Nature Communications*, 13:2693, 2022.
- [33] Marius Pachitariu, Shashwat Sridhar, Jacob Pennington, and Carsen Stringer. Spike sorting with kilosort4. *Nature methods*, 21(5):914–921, 2024.
- [34] Yue-Ting Pan, Jing-Lun Chou, and Chun-Shu Wei. Matt: A manifold attention network for eeg decoding. *Advances in Neural Information Processing Systems*, 35:31116–31129, 2022.
- [35] Andrew J Peters, Jun Lee, Nathan G Hedrick, Keelin O’Neil, and Takaki Komiyama. Reorganization of corticospinal output during motor learning. *Nature neuroscience*, 20(8):1133–1141, 2017.
- [36] Yael E Pollak, Robert Sachdev, Matthew Larkum, and Ariel Gilad. Cortex-wide laminar dynamics diverge during learning. *bioRxiv*, pages 2025–07, 2025.
- [37] Maria Giulia Preti, Thomas AW Bolton, and Dimitri Van De Ville. The dynamic functional connectome: State-of-the-art and perspectives. *NeuroImage*, 160:41–54, 2017. Functional Architecture of the Brain.
- [38] Gunnar Randers. On an asymmetrical metric in the four-space of general relativity. *Physical Review*, 59(2):195, 1941.
- [39] Ilaria Ricchi, Anjali Tarun, Hermina Petric Maretic, Pascal Frossard, and Dimitri Van De Ville. Dynamics of functional network organization through graph mixture learning. *NeuroImage*, 252:119037, 2022.
- [40] Houman Safaai, Alice Y Wang, Shinichiro Kira, Simone Blanco Malerba, Stefano Panzeri, and Christopher D Harvey. Specialized structure of neural population codes in parietal cortex outputs. *Nature Neuroscience*, pages 1–11, 2025.
- [41] Zhongmin Shen. *Lectures on Finsler geometry*. World Scientific, 2001.

- [42] Yan-Liang Shi, Nicholas A Steinmetz, Tirin Moore, Kwabena Boahen, and Tatiana A Engel. Cortical state dynamics and selective attention define the spatial pattern of correlated variability in neocortex. *Nature communications*, 13(1):44, 2022.
- [43] Yan-Liang Shi, Roxana Zeraati, Anna Levina, and Tatiana A Engel. Spatial and temporal correlations in neural networks with structured connectivity. *Physical review research*, 5(1):013005, 2023.
- [44] Tal Shnitzer, Hau-Tieng Wu, and Ronen Talmon. Spatiotemporal analysis using riemannian composition of diffusion operators. *Applied and Computational Harmonic Analysis*, 68:101583, 2024.
- [45] Ali Shojaie and Emily B Fox. Granger causality: A review and recent advances. *Annual Review of Statistics and Its Application*, 9(1):289–319, 2022.
- [46] Guillem Simeon, Gemma Piella, Oscar Camara, and Deborah Pareto. Riemannian geometry of functional connectivity matrices for multi-site attention-deficit/hyperactivity disorder data harmonization. *Frontiers in Neuroinformatics*, 16:769274, 2022.
- [47] Le Song, Mladen Kolar, and Eric Xing. Time-varying dynamic bayesian networks. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009.
- [48] Olaf Sporns, Joshua Faskowitz, Andreia Sofia Teixeira, Sarah A Cutts, and Richard F Betzel. Dynamic expression of brain functional systems disclosed by fine-scale analysis of edge time series. *Network Neuroscience*, 5(2):405–433, 2021.
- [49] Lloyd N Trefethen and Mark Embree. Spectra and pseudospectra: the behavior of nonnormal matrices and operators. *Princeton university press*, 2020.
- [50] Levente Varga, Vasile V. Moca, Botond Molnár, Laura Perez-Cervera, Mohamed K. Selim, Antonio Díaz-Parra, David Moratal, Balázs Péntek, Wolfgang H. Sommer, Raul C. Mureşan, Santiago Canals, and

Maria Ercsey-Ravasz. Brain dynamics supported by a hierarchy of complex correlation patterns defining a robust functional architecture. *Cell Systems*, 15(8):770–786.e5, 2024.

- [51] Thomas F Varley and Olaf Sporns. Network analysis of time series: Novel approaches to network neuroscience. *Frontiers in Neuroscience*, 15:787068, 2022.
- [52] Manasij Venkatesh, Joseph Jaja, and Luiz Pessoa. Comparing functional connectivity matrices: A geometry-aware approach applied to participant identification. *NeuroImage*, 207:116398, 2020.
- [53] Mingyu Wang, Yueming Wang, and Yuxiao Yang. Dynamic and low-dimensional modeling of brain functional connectivity on riemannian manifolds. *NeuroImage*, page 121243, 2025.
- [54] Rong Wang, Mianxin Liu, Xinhong Cheng, Ying Wu, Andrea Hildebrandt, and Changsong Zhou. Segregation, integration, and balance of large-scale resting brain networks configure different cognitive abilities. *Proceedings of the National Academy of Sciences*, 118(23):e2022288118, 2021.
- [55] Oumar Wone. Lectures on finslerian geometry. *arXiv preprint arXiv:2311.06778*, 2023.
- [56] Kisung You and Hae-Jeong Park. Re-visiting riemannian geometry of symmetric positive definite matrices for the analysis of functional connectivity. *NeuroImage*, 225:117464, 2021.

**גאומטריה של תהליכי למידה עצביים:  
גישה רימנית לדינמיקה של למידה  
עצבית**

**יהונתן קרקופר**







המחקר בוצע בהנחייתה של דר. הדס בן-איסטי, בפקולטה לרפואה.  
מחבר חיבור זה מצהיר כי המחקר, כולל איסוף הנתונים, עיבודם והצגתם, התייחסות  
והשוואה למחקרים קודמים וכו', נעשה כולו בצורה ישרה, כמצופה ממחקר מדעי המבוצע  
לפי אמות המידה האתיות של העולם האקדמי. כמו כן, הדיווח על המחקר ותוצאותיו  
בחיבור זה נעשה בצורה ישרה ומלאה, לפי אותן אמות מידה.

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי.



# תקציר

הקישוריות הפונקציונלית במערכת העצבים, בין אם היא נבחנת ברמת אזורי מוח שלמים ובין אם ברמת אוכלוסיות תאים או תאים בודדים, היא תופעה דינמית במהותה. דינמיות זו מהווה מנגנון מרכזי המאפשר למידה, הסתגלות ושינוי התנהגותי בתגובה לגירויים פנימיים וחיצוניים ולשינויים מתמשכים בסביבה. תהליכים קוגניטיביים, תפיסתיים ומוטוריים אינם נשענים על מבנה קישוריות קבוע, אלא על ארגון זמני וגמיש של אינטראקציות בין רכיבי הרשת העצבית, אשר משתנה בהתאם להקשר, למשימה ולמצב הפנימי של המערכת. על אף חשיבותה הרבה של ההשתנות הזמנית בקישוריות הפונקציונלית, מרבית המחקרים בתחום מתייחסים אל הקישוריות כאל ישות סטטית, הנאמדת באמצעות מדדים ממוצעים על פני פרקי זמן ארוכים או ניסויים שלמים, ובכך מאבדים מידע מהותי על הדינמיקה העצבית.

גישות האנליזה המעטות העוסקות בדינמיקה של קישוריות פונקציונלית מתמודדות עם אתגרים מתודולוגיים ומושגיים משמעותיים. ראשית, מדובר בבעיה עתירת-ממדיות, שכן קישוריות בין  $N$  רכיבים מיוצגת באמצעות מטריצות מסדר  $N \times N$ , אשר עצמן משתנות לאורך הזמן. שנית, הקישוריות הפונקציונלית מתוארת לרוב באמצעות מטריצות חצי-חיוביות סימטריות, אשר אינן שוכנות במרחב אוקלידי פשוט, אלא על יריעה בעלת מבנה גיאומטרי לא-ליניארי. אחד האתגרים המרכזיים הנובעים מכך הוא זיהוי דינמיקות משמעותיות בתוך רצף הקישוריות, וכן זיהוי רכיבי הרשת או תת-הרשתות האחראים לשינויים אלו, ובפרט בסקלות זמן שונות.

בעבודה זו אנו מציגים את RONI (Riemannian filtering Of Network Interactions), שיטת אנליזה גיאומטרית, בלתי-מונחית, לניתוח קישוריות פונקציונלית משתנה בזמן. במסגרת גישה זו, הקישוריות הפונקציונלית מיוצגת כאות המתפתח על גבי היריעה של המטריצות הסימטריות החצי-חיוביות (SPSD). ייצוג זה מאפשר לתאר את הדינמיקה של הקישוריות כמסלול גיאומטרי רציף, ולהשתמש בכלים מהגיאומטריה הרימנית לצורך ניתוח מבנה המסלול, קצב השינוי שלו והכיוונים הדומיננטיים בהם הוא מתפתח.

החידוש המרכזי של עבודה זו טמון בהגדרה של אופרטור גיאומטרי המדגיש הבדלים בין מצבי קישוריות עוקבים, תוך כיבוד המבנה הפנימי של היריעה עליה הם מוגדרים. אופרטור זה מאפשר זיהוי שינויים מהותיים בדינמיקה, מבלי להפר את התכונות המתמטיות הבסיסיות של מרחב המטריצות. בשילוב עם אופרטור משלים המדגיש דמיון בין מצבי קישוריות, אנו מציעים מסגרת אנליזה גיאומטרית המבוססת על הפעלה רקורסיבית של אופרטורים אלו. תהליך זה מוביל לפירוק רב-רזולוציות של אות הקישוריות הפונקציונלית, אשר מאפשר להפריד בין דינמיקות איטיות ומהירות, ולזהות את רכיבי הרשת התורמים לשינויים בכל סקלה זמנית.

יתרונות השיטה המוצעת מודגמים באמצעות יישומה על מגוון רחב של רישומים עצביים. ניתוחים אלה כוללים אותות אלקטרופיזיולוגיים מההיפוקמפוס, נתוני הדמיית סידן רחבת-שדה ברמה הקורטיקלית, הדמיית סידן דנדריטית ברזולוציה תאית, וכן נתוני EEG אנושיים. בכל אחד מהמקרים מראה RONI יכולת עקבית לזהות תת-רשתות בעלות משמעות ביולוגית, אשר מעצבות את הגיאומטריה של מסלול הקישוריות. זיהוי זה מוביל לייצוגים נמוכי-ממד של הארגון הרשת, הקשורים באופן מובהק לתהליכי למידה, לביצוע משימה ולשינויים התנהגותיים.

מעבר לאינטראקציות סימטריות, דינמיקות עצביות הן מטבען כיווניות, ומשקפות השפעות אסימטריות ומבנה סיבתי בין רכיבי הרשת. כדי להתמודד עם היבט זה, אנו מצייגים מסגרת גיאומטרית משלימה המבוססת על גיאומטריית Finsler-Randers, המאפשרת שילוב של כיווניות בניתוח מסלולי הקישוריות. במסגרת זו אנו מציעים הרחבה של המטריקה הרימנית בין מטריצות קורלציה, ומגדירים מידת מרחק המשלבת מידע קורלטיבי וכיווני כאחד. יישום שיטה זו על נתוני הדמיה של רשתות קורטיקליות גלובליות מראה כי הגיאומטריה המתקבלת מתארת בצורה מדויקת ועשירה יותר את הדינמיקה הקשורה ללמידה, בהשוואה לניתוח המבוסס על מבנה הקורלציות בלבד.

לסיכום, עבודה זו מציגה מסגרת גיאומטרית מאוחדת לניתוח דינמיקות של קישוריות פונקציונלית, המשלבת ניתוח רימני רב-רזולוציות עם הרחבה עקרונית לאינטראקציות כיווניות. השיטות המוצעות מכבדות את המבנה הגיאומטרי של מרחב הקישוריות, מתמודדות באופן טבעי עם בעיית הממדיות הגבוהה, ומאפשרות הפקת תובנות ביולוגיות חדשות אודות הדינמיקה והארגון של רשתות מוחיות מורכבות.