
Predicting Neural Activity from Connectome Embedding Spaces

Zihan Zhang

Allen Institute

zihan.zhang@alleninstitute.org

Huanqiu Zhang

The University of Chicago

huanqiuz@uchicago.edu

Stefan Mihalas

Allen Institute

stefanm@alleninstitute.org

Abstract

Understanding how structured patterns of neural activity emerge from the underlying connectivity is fundamental to elucidating brain function. While cortical connectomes are intrinsically high-dimensional, population activity typically resides in a much lower-dimensional subspace. Consequently, only a small fraction of the information encoded in the connectome appears relevant for shaping activity. Can we identify low-dimensional features of the connectome features that reliably predict neural activity? Leveraging the MICrONS dataset, which combines millimeter-scale, nanometer-resolution connectivity with simultaneously recorded in-vivo activity, we demonstrate a statistically significant alignment between morphological and functional similarity, quantified by subspace angles and centered kernel alignment. Topological analyses further reveal that the representation spaces of both the connectome and neural activity share a low-dimensional hyperbolic geometry with exponential scaling. These parallels motivated the hypothesis that embedding anatomical affinities into an appropriate geometric space can isolate the functionally relevant features of the connectome. We therefore applied multidimensional scaling to generate such embeddings and trained a simple linear model to reconstruct neuronal activity. Remarkably, the embedded connectome explained 68% in activity similarity, surpassing models that had direct access to activity similarity itself and outperforming similarly simple models that used the full high-dimensional connectome (56%). Our findings uncover a robust structure-function coupling: geometry-aware dimensionality reduction discards much of the connectome’s microscopic detail yet yields superior predictions of neural activity. This suggests that synaptic wiring implicitly encodes an abstract, low-dimensional organization that underlies the observed low-dimensional cortical activity.

1 Introduction

The human brain is a complex multiscale structure, encompassing various molecular, cellular, and neuronal phenomena that constitute the foundation of cognition [8, 35]. Experimental evidence estimates that the average human brain contains over 86 billion neurons and approximately 10^{14} synaptic connections [4, 36], while the mouse brain comprises approximately 85 million neurons, forming roughly 100 billion synaptic connections [41]. Each synapse, depending on its strength, has been suggested to store up to 4.7 bits of information [7]. However, are all these structural details essential to define the circuit activity in the brain? For the first time, based on monumental efforts

[15], we now have the capability to compare simplified models that link structure to activity in cortical circuits under varying levels of coarse-graining in connectivity.

While detailed structural properties shape brain activity, the extent to which they are necessary for functional modeling remains an open question. Simple artificial neural networks (ANNs) without bio-realistic constraints may not serve as ideal models, as their material and metabolic limitations, synaptic signals, and anatomical scaffolds are not encoded in the same manner as parameters in ANNs. Consequently, the mechanics and functionalities of the brain may not be fully explained or replicated by the scaling laws employed in artificial systems [10, 47, 40, 57, 5, 46]. Recent studies have found that, despite the complex and high-dimensional nature of individual neuronal responses, population activity across neurons often unfolds with low-dimensional subspaces, particularly in sensory and motor systems [14, 19, 2, 56]. Although these studies suggest that, due to its low-dimensional nature, activity contains only a smaller amount of information than neurons can independently encode, no link to its emergence from structure is proposed.

Historically, the nontrivial relationship between structural connectivity (\mathbf{W}) and functional characteristics (\mathbf{A}) of biological neuronal networks has been widely recognized, as exemplified by classical Hebbian theory and the three-factor learning rule [21, 30]. However, most studies exploring this relationship remain theoretical and lack systematic validation with empirical data. Experimental efforts to establish a direct link between structure and function primarily rely on MRI data, which emphasize macroscopic network-level analyses rather than detailed synaptic-level investigations [37, 3, 73, 69, 33, 25]. More recently, studies on *C. elegans* and fruit flies (*Drosophila Melanogaster*) have provided detailed connectomic data, enabling finer-grained analyses [71, 64, 55, 53, 18]. Building on this, [61] treat the *Drosophila* connectome as a prior, uncovering that key circuits rely on only a handful of neurons and enabling whole-brain casual modeling from sparse data. [49] likewise predict each neuron’s response to dynamic visual stimuli by sharing parameters within cell types. While these connectome-based studies employ state-of-the-art deep learning models to uncover complex high-dimensional relationships between structure and function, our work instead focuses on identifying normative, low-dimensional representations of these relationships. Specifically, we adopt a low-dimensional embedding approach based on vectorized connectome similarity, which has been previously used to predict and explain edgewise functional connectivity in inter-hemispheric homologous regions [62, 50]. From a broader perspective, our study aligns with previous work by advocating for the incorporation of biological details into computational models in a simplified yet principled manner that enhances both performance and metabolic efficiency [61, 49]. However, unlike prior studies that emphasize the precise details of the connectome, our approach focuses on a lossy compression of connectomic information, capturing essential structural properties while abstracting away fine-grained complexities.

2 Results

2.1 Structure-Function Alignment

We use the largest functionally-imaged electron microscopy (EM) dataset currently available, MI-CrONS, which was generated through dense calcium imaging followed by the reconstruction of a millimeter-scale volume [15, 23, 22]. This dataset provides nanometer-resolution details on the spatial locations of each cell’s soma and associated synapses, along with classification labels, postsynaptic density (PSD) sizes, and activity recordings of neurons in the primary visual cortex (V1) as well as the anterolateral (AL), lateromedial (LM), and rostromedial (RL) higher visual areas. The activity data consists of detrended, deconvolved, and stimulus- and behavior-aligned recordings from excitatory neurons spanning layers 2 to 6. Figure 1A-B provides an illustration of neuronal connectivity within the dataset. In Sections 2.1 and 2.3, because we incorporate both structural and functional information, we restrict analyses to neurons that were imaged, reconstructed, and coregistered – 497 unique cells in total – comprising 35-80 neurons per scan across 12 scans (Section A.1). In Section 2.2, only the connectome information is required for the geometric analysis; accordingly, the available sample expands to 716. To maintain consistency with prior studies, we adopt the same inclusion criteria and preprocessing procedures as in [15, 22].

We quantify how closely neuronal input- and output-similarity matrices, ($\mathbf{C}^{\text{in}}/\mathbf{C}^{\text{out}}$), align with functional-similarity matrix (\mathbf{C}^{act}). We compute \mathbf{C}^{in} and \mathbf{C}^{out} from the weighted connectome by counting the total number of synapses between each pair of neurons. To capture the network’s

global relational structure, we characterize pairwise strength and similarity using these higher-order connectivity measures rather than the direct Hebbian weight matrix (\mathbf{W}), thereby avoiding additional weighting by synaptic counts or strengths [67, 58, 44, 16]. For \mathbf{C}^{out} , we include only neurons whose dendrites are fully reconstructed; for \mathbf{C}^{in} , we include neurons whose axons are either extended or clean [15]. Functional similarity, \mathbf{C}^{act} , is computed from the full neuronal response across each scan, comprising 30,000 data points and yielding a median pairwise similarity of 1.18% (Figure 1C). Formally, the connectivity matrices are defined as $\mathbf{W}^{\text{in}} \in \mathbb{R}^{M_1 \times N}$ and $\mathbf{W}^{\text{out}} \in \mathbb{R}^{N \times M_2}$, where N denotes the number of targeted neurons (or features), and M_1 and M_2 represent the number of structurally confident neurons for axons and dendrites, respectively, where $M_1, M_2 \gg N$. The connection similarity matrix is then computed as:

$$\mathbf{C}_{ij}^{\text{in}} = \text{corr}(\mathbf{W}_{:,i}^{\text{in}}, \mathbf{W}_{:,j}^{\text{in}}), \quad \mathbf{C}_{ij}^{\text{out}} = \text{corr}(\mathbf{W}_{i,:}^{\text{out}}, \mathbf{W}_{j,:}^{\text{out}}). \quad (1)$$

$\forall 1 \leq i, j \leq N$, $\mathbf{C}^{\text{in}}, \mathbf{C}^{\text{act}} \in \mathbb{R}^{N \times N}$. We then quantified how each neuron’s connectivity profile relates to its activity profile by correlating the corresponding rows in \mathbf{C}^{act} with those of $\mathbf{C}^{\text{in}}/\mathbf{C}^{\text{out}}$:

$$\mathbf{C}_{ij}^{\text{in,act}} = \text{corr}(\mathbf{C}_{i,:}^{\text{in}}, \mathbf{C}_{j,:}^{\text{act}}), \quad \mathbf{C}_{ij}^{\text{out,act}} = \text{corr}(\mathbf{C}_{i,:}^{\text{out}}, \mathbf{C}_{j,:}^{\text{act}}). \quad (2)$$

The distribution of connectivity-activity similarities for individual neurons is significantly higher than the shuffled control (Figure 1D1), which revealed a significant dependence of functional similarity on connection similarity. Notably, input similarity is generally a better predictor of functional similarity than output similarity (Figure S2). The neuronal identities and spatial organization of the target neurons are detailed in Table S2, Table S3, and Figure S1. In particular, more than 3/4 (382) of the neurons are from V1; the remainder are from higher visual area that are spatially distant. This finding is further supported by centered kernel alignment (Section A.3 and Figure 1D2) and by analyzing the principal angles between subspaces in the functional and structural affinity spaces (Section A.4 and Figure S3) [9, 42, 45].

2.2 Hyperbolic Geometry for Both Connectome and Activity Representation Spaces

Next, we analyze the underlying geometry of the representation space formed by neuronal connectivity and activity using topological data analysis. Specifically, we used Betti curves, which quantify cycles (i.e., closed loops in a topological space) and cliques (i.e., fully connected subgraphs in a simplicial complex) in graphs generated by varying the threshold that defines significant connections between neurons, to identify the underlying geometry of our data (Figure S7) [65]. Similar topological methods based on persistent homology has been widely applied to study the manifold structures of population neural or perceptual responses in systems such as olfaction, the hippocampus, head direction, and grid cells, consistently revealing that the observed geometry is often hyperbolic rather than Euclidean [32, 82, 13, 28, 80].

To estimate the intrinsic dimensionality of the experimental graph, we compare Betti curve distributions from random geometric models of varying dimensionality and select the model that best matches the experimental curve. Specifically, we sample out point clouds in hyperbolic spaces, define similarity as inverse geodesic distance, compute their Betti curves, and quantify agreement with the experimental curves via mean-squared error (MSE) and integrated Betti values. Because Betti curves are invariant under monotone reparameterization, this procedure assesses geometric alignment rather than scale [82, 80]. In our data, a low-dimensional hyperbolic model, \mathbb{H}^2 , provides the closest match. This analysis indicates that both connectivity and neuronal dynamics are better described by low-dimensional hyperbolic space than with Euclidean space (Figure 1E). The advantages of hyperbolic geometry become more pronounced as dataset size increases, reflecting its higher embedding capacity and efficient accommodation of additional nodes (Figure S16). This result is robust across different connectivity patterns of \mathbf{W} and across multiple neuronal similarity measures (Section B.1 and Section B.3). Additional methodological details are provided in Section A.2. Finally, applying hyperbolic t-SNE to embed proofread neurons in \mathbb{H}^2 reveals clear cell-type separation, further supporting hyperbolic organization in these circuits (Figure S5) [81]. These findings point to a shared geometric substrate linking structural and functional spaces. Notably, low-dimensional hyperbolic embeddings capture the data’s underlying structure with fidelity comparable to that of much higher-dimensional Euclidean embeddings, providing a robust and scalable framework for analyzing and interpreting large-scale neural circuits (Figure S20) [32].

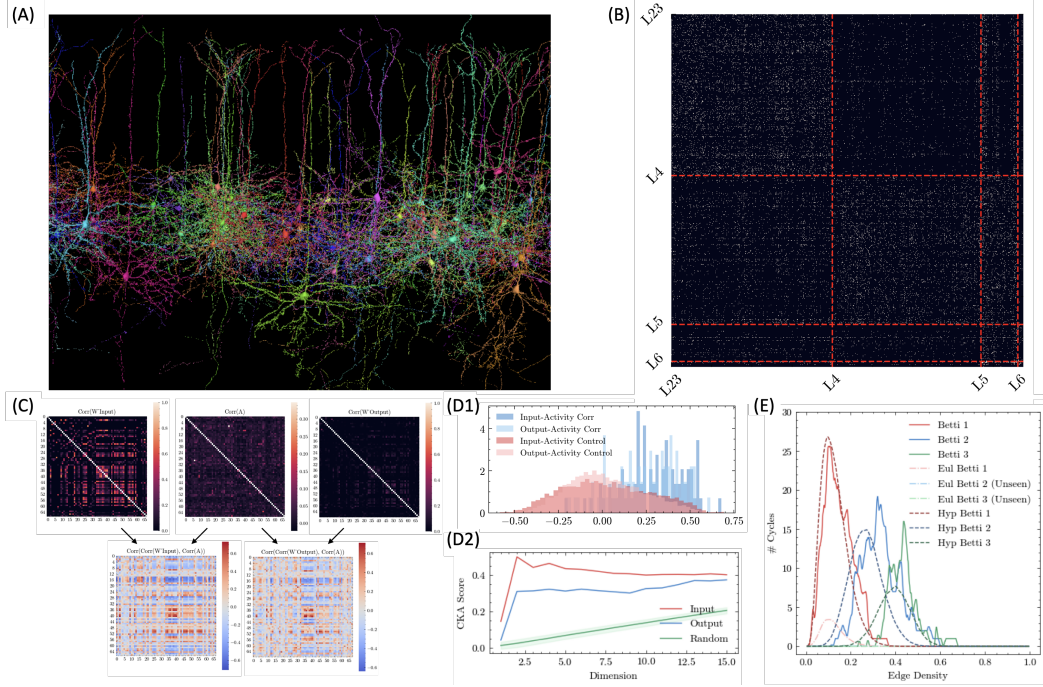


Figure 1: Demonstration of MICrONS Dataset and Structure-Function Alignment: (A): Illustration of 89 proofread layer 4 pyramidal cells from the MICrONS dataset using *neuroglancer*; (B): Binary pairwise connections between fully proofread excitatory neurons ($N = 716$), ordered by layers. The horizontal axis represents postsynaptic neurons, while the vertical axis represents the presynaptic neurons (Section B.1); (C): Illustration of input connection correlation C^{in} (top-left), neuronal activity correlation C^{act} (top-middle), output connection correlation C^{out} (top-right), $C^{\text{in,act}}$ (bottom-left, also see Equation (2)), and $C^{\text{out,act}}$ (bottom-right). For instance, each entry in $C^{\text{in,act}}$ represents the Pearson correlation between the i th row of C^{in} and the j th row of C^{act} . This matrix has diagonal elements representing the correlation between connection and activity similarities for each neuron, while the off-diagonal elements serve as controls, relating the functional similarity of one neuron to the connection similarity of another. (D1): Correlation between activity and connection similarity for inputs (dark blue) and outputs (light blue), along their respective controls (dark red and light red). t -test p -values: $p_1 = 2.32 \times 10^{-22}$ (dark blue vs. dark red, testing whether input connections are correlated with activity), $p_2 = 2.43 \times 10^{-16}$ (light blue vs. light red, testing whether output connections are correlated with activity), and $p_3 = 4.90 \times 10^{-4}$ (dark blue vs. light blue, testing whether input connections exhibit a stronger correlation with activity than output connections). See Figure S2 for results from additional scans; (D2): Centered Kernel Analysis (CKA) of low-rank approximations using Singular Value Decomposition (SVD) for C^{act} with C^{in} (red), C^{out} (blue), and random matrices (green). The x -axis represents the inclusion of different numbers of top singular values (i.e. truncation rank) (Section A.3); (E): Betti curves derived from the activity pairwise correlation matrix C^{act} (solid lines) align more closely with points uniformly sampled from two-dimensional hyperbolic space than with those from Euclidean space. Standard deviations are omitted for clarity. Averaging across 100 hyperbolic samples (dark dashed line) at the optimized radius R_{max} value (Section A.2) yields a cumulative absolute pointwise error of 14.67 across the first three Betti curves. The corresponding hyperbolic Betti values, $\bar{\beta}_m^{\text{hyp}}$ for $m = 1, 2, 3$, are 3.31, 2.46, and 1.21, which closely match the groundtruth Betti values of 3.47, 3.08, and 1.64. By contrast, the Euclidean Betti values, $\bar{\beta}_m^{\text{eul}}$, are only 0.42, 0.09, and 0.01. Section B.2 presents a comprehensive Betti analysis of all activity data. Figure (D-F) analyzes Session 8, Scan 5, from the MICrONS dataset [15].

2.3 Activity Trace Reconstruction through Linear Model

Next, we tested whether the activity of individual neurons could be directly predicted based on the activity of others and their anatomical connectivity. We implemented a simple statistical linear model, previously shown to be effective in fMRI analysis [60]. Given an activity trace $\mathbf{A} \in \mathbb{R}^{N \times T}$ with N

targeted neurons and T time points, neuronal dynamics are reconstructed using:

$$\tilde{\mathbf{A}} = (\mathbf{C} - \mathbf{I}) \cdot \mathbf{A}, \quad (3)$$

where $\mathbf{C} \in \mathbb{R}^{N \times N}$ represents pairwise affinities (i.e. weights of dynamical similarity), and $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times T}$ denotes the reconstructed activity. We define $\alpha = \mathbf{C} - \mathbf{I}$, so that α_{ij} represents the weight assigned to neuron i 's activity trace when predicting the activity of neuron j (Figure 2). This model predicts each neuron's activity by leveraging the simultaneous activity of all other neurons in a leave-one-out manner. To evaluate model performance, we compute the reconstruction accuracy acc as the average Pearson correlation between corresponding rows in $\tilde{\mathbf{A}}$ and \mathbf{A} , enabling neuron-wise comparison. When $\mathbf{C} = \mathbf{C}^{\text{act}}$, acc_{act} reaches 14.1%, averaged across 12 scans. This serves as an upper-bound of all the other linear models, as the model directly uses information from observed activity. To predict activity from structure without accessing the activity of the targeted neuron, \mathbf{C} can be replaced with either the in-degree (\mathbf{C}^{in}) or out-degree (\mathbf{C}^{out}) correlation matrices derived from the connectome (Section 2.1), allowing structural similarity to serve as the predictor. Alternatively, instead of using the full connectomic, samples can be projected into a low-dimensional space using non-metric multidimensional scaling (MDS) to either hyperbolic or Euclidean space [17, 81]. Notably, activity correlation varies substantially across scans due to changes in the mouse's behavioral state during visual stimulus presentation. This results in significant variability in absolute reconstruction accuracy, largely driven by a small subset of neurons [22]. To mitigate this effect, we focus on relative performance, quantified as the explanation ratio acc/acc_{act} , which measures the effectiveness of alternative affinity measures relative to \mathbf{C}^{act} [37].

Using the input connectome correlation, we observe an explanation ratio of 56.3%, whereas the output connectome yields 35.6% (Figure 3A). The observation that \mathbf{C}^{in} yields better accuracy than \mathbf{C}^{out} is consistent with the discussion in Section 2.1. Furthermore, having more neurons included in the scan (larger N) improves reconstruction accuracy when using connectome or activity similarity, suggesting that further experimental refinement may improve performance (Figure S21). However, the explanation ratio remains stable and uncorrelated to N , highlighting the robustness of the observation (Figure S22). Moreover, we evaluate the performance of direct synaptic connectivity – another common approach to quantifying α – within the linear model and find that it performs significantly worse than input similarity, although comparably to output similarity, regardless of the connectivity modality (Figure S33) [34, 37].

Surprisingly, utilizing the embedding results from either geometry improves reconstruction accuracy compared to using the raw connectome correlation (e.g. for hyperbolic embedding, 66.6% for input and 62.0% for output), despite a substantial reduction in the dimensionality of the connectome information (Figure 3A). The stability range of R_{max} is relatively broad (Figure 3B). The explanation ratios for hyperbolic and Euclidean embeddings are statistically similar across scans, likely due to the limited number of coregistered neurons, which reduces structural complexity and introduces a selective bias, resulting in a nonuniform or bimodal distribution caused by nonhomogeneous neuronal identity (Section B.1 and Figure S8). Only neurons that are both reconstructed and coregistered are included in the structure-function analysis, whereas all reconstructed neurons are used for the hyperbolic geometry analysis (Section A.1). Thus far, we exclusively focused on the 2D embedding. When we increased the embedding dimension and surprisingly observed that, despite significant differences in compression quality, the overall reconstruction remained well aligned, underscoring the model's robustness (Figure S32).

Our results remain generalizable and stable when \mathbf{C}^{in} and \mathbf{C}^{out} are computed using either the binary connection pattern or postsynaptic density (PSD) volume (Figure S23), or when considering connections exclusively within subpopulations of specific cell types (Figure S24) [59, 11]. Compared to the baseline experiment, using the binary connectome yields comparable accuracy in activity reconstruction (56.2% for input and 36.0% for output), whereas incorporating PSD information significantly reduces performance (52.0% and 28.0%, respectively). Furthermore, even partially accurate structural information – specifically, incorporating connections between unproofread E/I neurons in the computation of $\mathbf{C}^{\text{in}}/\mathbf{C}^{\text{out}}$ – remains beneficial for activity prediction. The most notable improvement is observed for \mathbf{C}^{in} (64.6%), with a marginal improvement for \mathbf{C}^{out} (37.8%) (Figure S25). Embedding-based results remain relatively consistent across different connectivity representations. Additionally, restricting connections from or to certain subpopulations of neurons (e.g. E/I) significantly reduces the dimensionality of \mathbf{W} and limits the performance (36.4% for excitatory input and 48.8% for inhibitory input), whereas the embedding results are less sensitive to

this reduction (63.8% for excitatory input and 66.5% for inhibitory input) (Table S1, Figure S24) [44, 38].

To test whether the structure–function coupling uncovered by our embeddings generalizes across distinct sensory contexts and even at the raw-signal level, we next dissect prediction performance within individual stimulus periods and for unprocessed fluorescence traces. Although our initial analysis considered neuronal activity across the entire scan, multiple stimulus types were presented in discrete periods (Section A.1). To accommodate this structure, we compute \mathbf{C}^{act} , reconstruct activity traces separately within each period, and report the averaged performance. As expected, explanation ratios decrease decline because \mathbf{C}^{act} now captures more period-specific structure, yet embeddings still confer a clear advantage (see p -values in Figure S26). For completeness, we also evaluate the model directly on the raw fluorescence data, which exhibit substantial population-level pairwise correlations (median of 4.90% and mean of 7.00%) [15]. The benefit of input-based embeddings become less significant, whereas the advantage of output-based embeddings is retained (see p -values in Figure S27 and Table S5).

To test how information-efficient and robust our structure–function mapping truly is, we first ask how prediction accuracy degrades when only a fraction of the available neurons or connections is retained. We explore downsampling by selecting neurons either based on the top- K entries or at random. Specifically, after computing α , we filter (zero out) its entries column by column, either by magnitude or randomly, and then predict the leave-one-out neuron using only a subset of the N neurons (Figure S28). Notably, \mathbf{C}^{act} preserves its accuracy when using only the top-20%, whereas the performance of structurally derived representations deteriorates significantly, particularly for embedded results. For example, in the baseline experiment (Figure S28A), selecting the top 20% entries from \mathbf{C}^{act} preserves 93.0% of its performance, while only 70.8% and 85.1% are retained for $\mathbf{C}^{\text{in,hyp}}$ and \mathbf{C}^{in} respectively. The degradation of $\mathbf{C}^{\text{in,hyp}}$ arises because subsampling reduces data quantity, and embedding further compresses dimensionality, together incurring greater information loss compared to the higher resolution and structural density of the original array (\mathbf{C}^{in}). The performance advantage of $\mathbf{C}^{\text{in,hyp}}$ over \mathbf{C}^{in} becomes less pronounced with smaller values of K ($p = 0.003$ for 100% (Figure 3A), 0.051 for 50%, and 0.905 for 20%). However, the significance level is preserved under random selection ($p = 0.001$ for random 50% and 0.002 for random 20%). The observation remains qualitatively consistent across different models (Figure S28). Furthermore, we evaluate the robustness of the embedding results when the connectome (\mathbf{W}) is corrupted by randomly permuting synaptic connections (Section B.6). Output correlation exhibits a 19.3% performance degradation when 40% synapses are perturbed, compared to the baseline experiment, whereas the hyperbolic embedding shows only a 1.7% decrease (Figure S29). This highlights the model’s robustness in the presence of unobservable experimental errors.

To further refine our analysis, we restrict the dataset to neurons that are both functionally coregistered and fully proofread ($N = 338$ in total). This filtering enhances reliability by leveraging output similarity in prediction and generally improves relative performance compared to using the unfiltered dataset, partly because the performance of \mathbf{C}^{act} also declines for small N (Figure S21 and Figure S30). To maintain consistency with the previous geometric analysis (Section B.1), we perform MDS on all proofread excitatory neurons ($N = 712$) to obtain a universal embedding that preserves relevant interactions across a large subset of neurons. We then extract the corresponding information for the scan-specific coregistered neurons for reconstruction purposes.

We also construct a null model using the classical continuous-time recurrent neural network (RNN) [77]. Without any training, we inject random inputs and compute the correlation of hidden activities in the recurrent layer with 60 neurons. To ensure a fair comparison with the MICrONS dataset, we finetune the inputs so that the dimensionality of the resulted activity space, quantified through normalized participation ratio, matches that of the experimental data (≈ 0.30) (Figure S31) [27]. To mimic the typical non-negative distribution of deconvolved activity traces (Figure 2), we use ReLU as the activation function, which naturally constrains values in \mathbf{C}^{act} . Notably, projecting random matrices into a low-dimensional space presents empirical challenges, often causing pairwise distances to approach the space limit (Figure S20), leading to a high stress value. Figure 3C illustrates the correlation of the reconstructed results, where \mathbf{C}^{act} achieves promising performance with $\approx 74\%$ accuracy. In contrast, both \mathbf{C}^{in} and \mathbf{C}^{out} follow an approximately Gaussian distribution centered at 0, effectively nullifying the alignment effect through total summation. The weak signal observed in the embedded results can be attributed to the averaged signal from population activity. This comparative study of the random network highlights that biological systems provide a well-structured scaffold

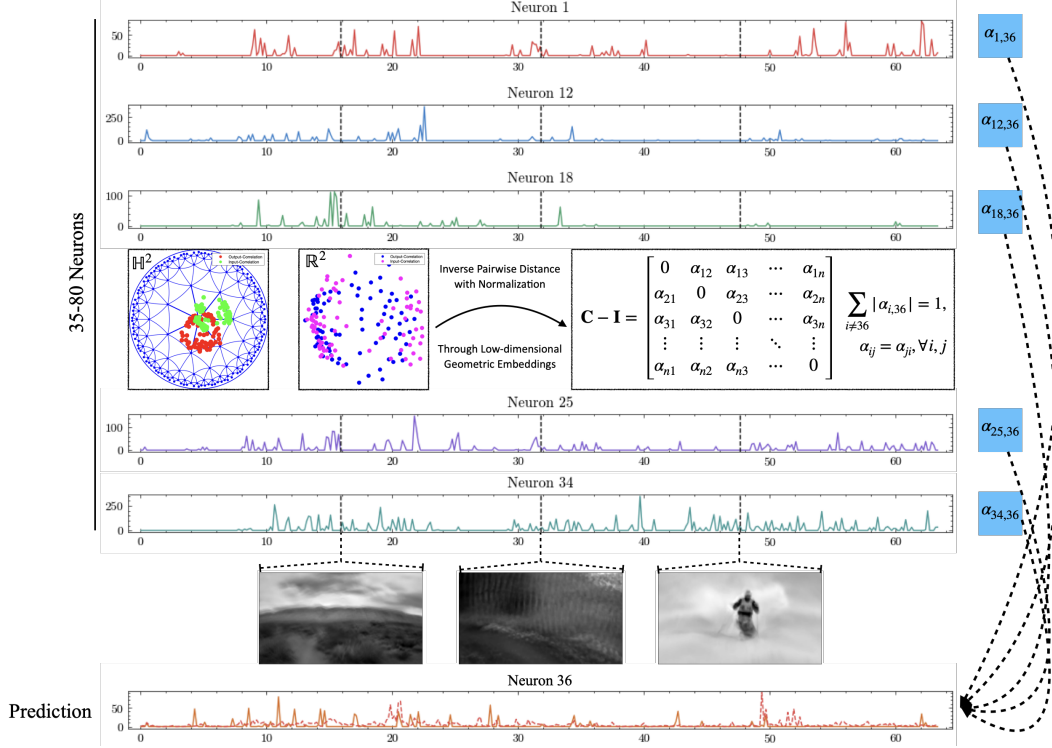


Figure 2: **Schematics of Linear Model:** The model utilizes deconvolved calcium traces from coregistered imaged neurons recorded during serial video trails to predict the activity of a held-out neuron [15, 22]. The prediction is obtained by computing a weighted sum of activities from other neurons, with weights represented by the symmetric matrix \mathbf{C} . We present 1% of the total traces from Session 8, Scan 5 of the MICrONS dataset. Solid line: groundtruth; dashed line: prediction. Frame rate: 6.3/s. x -axis in seconds. Inset: MDS projections of \mathbf{C}^{in} and \mathbf{C}^{out} to low-dimensional Hyperbolic and Euclidean spaces.

with meaningful structure-function alignment, which can be leveraged to infer and adapt information across different levels.

3 Discussion

In this work, we investigate the relationship between synaptic connectivity and the functional responses of excitatory neurons in the mouse visual cortex, leveraging dense reconstruction data from the MICrONS dataset [15, 75, 22]. Our analysis reveals a clear alignment between pairwise structural and functional affinities: neurons that receive similar inputs or outputs tend to exhibit similar functional behaviors, and vice versa. Furthermore, by identifying matching topological signatures, we demonstrate that both structural and functional representation spaces can be embedded in low-dimensional hyperbolic spaces with varying radii. Using non-metric MDS method, we construct a correlation matrix based on the spatial proximity of points within these geometric embeddings. Notably, using these results, we achieve not only comparable but significantly improved reconstruction accuracy of activity traces via a simple linear model, despite operating in a significantly lower-dimensional space than the full connectivity matrix. This observation suggests that many anatomical details may not be critical for understanding coarse structure-function mapping and that their interactions can be effectively captured in a reduced parameter space.

Our findings demonstrate that weighted connectivity patterns yield activity prediction accuracy comparable to that achieved using binary connectivity, consistent with prior studies suggesting that the topological structure, defined by the adjacency matrix, is sufficient for structure-function prediction [72, 22]. Also, the use of unweighted graphs to model excitable dynamics also offers computational advantages by simplifying network representation and analysis while maintaining high performance

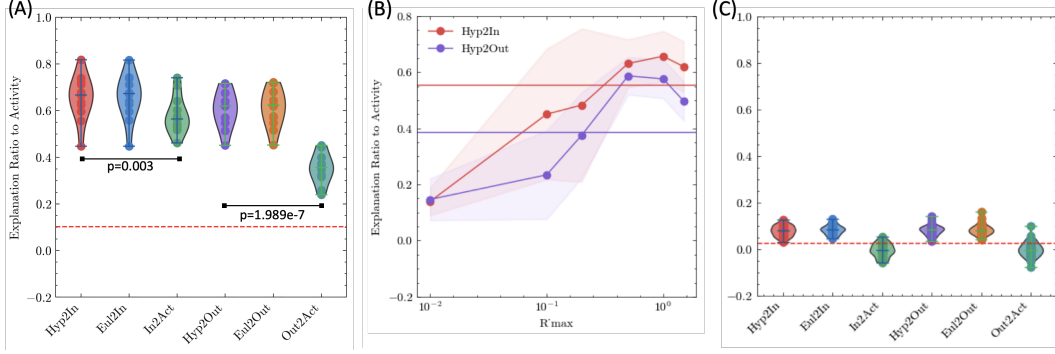


Figure 3: Using Low-Dimensional Geometric Embeddings of Structural Relationship Improves Reconstruction Quality of Activity Traces: (A) Explanation ratio (normalized by scan-specific a_{act}) achieved by C^{in} , C^{out} , and their respective 2D Euclidean and hyperbolic embeddings ($R_{\max} = 1$). The dashed line (15.1%) indicates the average of the median pairwise neuronal activity correlations across 12 scans, normalized by a_{act} . The p -values, calculated from paired (dependent) t-test, comparing the performance of C^{in} and C^{out} with their corresponding hyperbolic embeddings are shown. Each scattered point corresponds to a separate scan and represents the median performance of scan-specific targeted neurons; (B) Sensitivity analysis of reconstruction accuracy as a function of embedding space capacity, varied through R_{\max} . Horizontal lines represent median accuracy for C^{in} and C^{out} ; (C): Similar analysis as in (A), but on an untrained continuous-time recurrent neural network, where no signal is represented through the connectome and minimal information is derived from its embeddings.

[54]. Additionally, our results suggest that total synaptic strength, measured as the sum of PSD sizes, may not serve as a reliable indicator of such alignment (Figure S23). For improved interpretability, we employ a simple linear model to reconstruct activity traces. However, more complex neural-network-based models that incorporate nonlinearity and leverage embedding information, or that simply make α parameters trainable, such as a single-layer perceptron, may yield improved results [77, 61, 49]. Also, considering the design of neurophysiological experiments in which a head-fixed mouse freely walks on a cylindrical treadmill while viewing visual stimuli, the effectiveness of our simple model occurs may result from the passive nature of stimuli presentation, as the mouse is not engaged in any demanding task such as decision making based on contextual information or cues (Figure S31) [37, 48, 15]. Evaluating the model’s effectiveness and robustness under various experimental settings represents an interesting direction for future research. Additionally, we focus on activity reconstruction, noting that subtle modes of activity may play a critical role in behavior [68]. Our work exclusively examines the recurrent circuit components reconstructed in the MICrONS dataset and does not make any assumptions or provide formulations regarding how stimuli are presented or transmitted through the sensory cortex (e.g. retinal ganglion cell (RGC) and lateral geniculate nucleus (LGN)) [76, 75, 22].

We employ an RNN with randomly initialized weight as a null model, showing that the organized structure captured by C^{in} or C^{out} is a necessary prerequisite for an embedding to achieve a high explanation ratio. Because the reconstructed excitatory neurons exhibit sparse all-to-all connectivity, we cannot readily incorporate the experimental connectome into simulations to establish sufficiency. In addition, for greater fidelity we compute C using rectangular matrices (Section 2.1), whereas the standard RNN formulation assumes a square weight matrix. Obtaining more densely reconstructed experimental data will therefore be crucial for fully addressing this question.

The effectiveness of using MDS to set weights in the linear model may be attributed to at least three factors. First, anatomical similarity may create a spatial embedding for each neuron, where similar neuron types are positioned closer together and exhibit similar activation patterns. This arises from the presence of multiple neuronal types, which are not explicitly differentiated in the recordings [79, 70, 78]. Second, the spatial location of neurons imposes constraints on connection formation, which MDS then extracts from the connectome [6, 10, 66, 1]. While the mouse cortex lacks organization in terms of orientation columns, it exhibits strong retinotopic organization [29, 83]. These mechanisms are not independent, as spatial distance plays a key role in refining neuronal subtypes and predicting

variations in synaptic connectivity patterns [58, 52]. Third, neuronal responsiveness and activity patterns may be strongly influenced by the spatial and temporal organization of synaptic inputs, which, in conjunction with Hebbian plasticity, shape synchronous dynamics [44, 43, 63]. MDS could effectively capture these dynamics. However, empirical measurements of the dimensionality of MICrONS activity space suggest that neural dynamics are not strictly low-dimensional (Figure S31), limiting the influence of this factor. A precise attribution of MDS effectiveness to these or other factors is beyond the scope of this study.

References

- [1] Jascha Achterberg et al. “Spatially embedded recurrent neural networks reveal widespread links between structural and functional neuroscience findings”. In: *Nature Machine Intelligence* 5.12 (2023), pp. 1369–1381.
- [2] Ege Altan et al. “Low-dimensional neural manifolds for the control of constrained and unconstrained movements”. In: *bioRxiv* (2023), pp. 2023–05.
- [3] Andrea Avena-Koenigsberger, Bratislav Misic, and Olaf Sporns. “Communication dynamics in complex brain networks”. In: *Nature reviews neuroscience* 19.1 (2018), pp. 17–33.
- [4] Frederico AC Azevedo et al. “Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain”. In: *Journal of Comparative Neurology* 513.5 (2009), pp. 532–541.
- [5] Yasaman Bahri et al. “Explaining neural scaling laws”. In: *Proceedings of the National Academy of Sciences* 121.27 (2024), e2311878121.
- [6] Marc Barthélemy. “Spatial networks”. In: *Physics reports* 499.1-3 (2011), pp. 1–101.
- [7] Thomas M Bartol Jr et al. “Nanconnectomic upper bound on the variability of synaptic plasticity”. In: *Elife* 4 (2015), e10778.
- [8] Danielle S Bassett and Michael S Gazzaniga. “Understanding complexity in the human brain”. In: *Trends in cognitive sciences* 15.5 (2011), pp. 200–209.
- [9] ke Björck and Gene H Golub. “Numerical methods for computing angles between linear subspaces”. In: *Mathematics of computation* 27.123 (1973), pp. 579–594.
- [10] Ed Bullmore and Olaf Sporns. “The economy of brain network organization”. In: *Nature reviews neuroscience* 13.5 (2012), pp. 336–349.
- [11] Luke Campagnola et al. “Local connectivity and synaptic dynamics in mouse and human neocortex”. In: *Science* 375.6585 (2022), eabj5861.
- [12] Brendan Celii. “NEURD: automated proofreading and feature extraction for connectomics”. PhD thesis. Rice University, 2024.
- [13] Rishidev Chaudhuri et al. “The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep”. In: *Nature neuroscience* 22.9 (2019), pp. 1512–1520.
- [14] Mark M Churchland et al. “Neural population dynamics during reaching”. In: *Nature* 487.7405 (2012), pp. 51–56.
- [15] MICrONS Consortium et al. “Functional connectomics spanning multiple areas of mouse visual cortex”. In: *Nature* 640.8058 (2025), pp. 435–447.
- [16] Lee Cossell et al. “Functional organization of excitatory synaptic strength in primary visual cortex”. In: *Nature* 518.7539 (2015), pp. 399–403.
- [17] Trevor F Cox and Michael AA Cox. *Multidimensional scaling*. CRC press, 2000.
- [18] Matthew S Creamer, Andrew M Leifer, and Jonathan W Pillow. “Bridging the gap between the connectome and whole-brain activity in *C. elegans*”. In: *bioRxiv* (2024), pp. 2024–09.
- [19] John P Cunningham and Byron M Yu. “Dimensionality reduction for large-scale neural recordings”. In: *Nature neuroscience* 17.11 (2014), pp. 1500–1509.
- [20] Carina Curto and Katherine Morrison. “Relating network connectivity to dynamics: opportunities and challenges for theoretical neuroscience”. In: *Current opinion in neurobiology* 58 (2019), pp. 11–20.
- [21] Peter Dayan and Laurence F Abbott. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. 2005.
- [22] Zhuokun Ding et al. “Functional connectomics reveals general wiring rule in mouse visual cortex”. In: *Nature* 640.8058 (2025), pp. 459–469.
- [23] Sven Dorkenwald et al. “CAVE: Connectome annotation versioning engine”. In: *Nature Methods* (2025), pp. 1–9.
- [24] Leila Elabbady et al. “Perisomatic features enable efficient and dataset wide cell-type classifications across large-scale electron microscopy volumes”. In: (2022).
- [25] Panagiotis Fotiadis et al. “Structure–function coupling in macroscale human brain networks”. In: *Nature Reviews Neuroscience* 25.10 (2024), pp. 688–704.
- [26] Juan A Gallego et al. “Neural manifolds for the control of movement”. In: *Neuron* 94.5 (2017), pp. 978–984.

- [27] Peiran Gao et al. “A theory of multineuronal dimensionality, dynamics and measurement”. In: *BioRxiv* (2017), p. 214262.
- [28] Richard J Gardner et al. “Toroidal topology of population activity in grid cells”. In: *Nature* 602.7895 (2022), pp. 123–128.
- [29] Marina E Garrett et al. “Topography and areal organization of mouse visual cortex”. In: *Journal of Neuroscience* 34.37 (2014), pp. 12587–12600.
- [30] Wulfram Gerstner et al. “Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules”. In: *Frontiers in neural circuits* 12 (2018), p. 53.
- [31] Robert Ghrist. “Barcodes: the persistent topology of data”. In: *Bulletin of the American Mathematical Society* 45.1 (2008), pp. 61–75.
- [32] Chad Giusti et al. “Clique topology reveals intrinsic geometric structure in neural correlations”. In: *Proceedings of the National Academy of Sciences* 112.44 (2015), pp. 13455–13460.
- [33] Zijin Gu et al. “Heritability and interindividual variability of regional structure-function coupling”. In: *Nature Communications* 12.1 (2021), p. 4894.
- [34] Patric Hagmann et al. “Mapping the structural core of human cerebral cortex”. In: *PLoS biology* 6.7 (2008), e159.
- [35] Michael Hawrylycz et al. “Canonical genetic signatures of the adult human brain”. In: *Nature neuroscience* 18.12 (2015), pp. 1832–1844.
- [36] Nicole Hindley, Anna Sanchez Avila, and Christopher Henstridge. “Bringing synapses into focus: Recent advances in synaptic imaging and mass-spectrometry for studying synaptopathy”. In: *Frontiers in Synaptic Neuroscience* 15 (2023), p. 1130198.
- [37] Christopher J Honey et al. “Predicting human resting-state functional connectivity from structural connectivity”. In: *Proceedings of the National Academy of Sciences* 106.6 (2009), pp. 2035–2040.
- [38] Xiaolong Jiang et al. “Principles of connectivity among morphologically defined cell types in adult neocortex”. In: *Science* 350.6264 (2015), aac9462.
- [39] William B Johnson and Assaf Naor. “The Johnson–Lindenstrauss lemma almost characterizes Hilbert space, but not quite”. In: *Discrete & Computational Geometry* 43.3 (2010), pp. 542–553.
- [40] Jared Kaplan et al. “Scaling laws for neural language models”. In: *arXiv preprint arXiv:2001.08361* (2020).
- [41] Narayanan Kasthuri et al. “Saturated reconstruction of a volume of neocortex”. In: *Cell* 162.3 (2015), pp. 648–661.
- [42] Andrew V Knyazev and Merico E Argentati. “Principal angles between subspaces in an A-based scalar product: algorithms and perturbation estimates”. In: *SIAM Journal on Scientific Computing* 23.6 (2002), pp. 2008–2040.
- [43] Ho Ko, Thomas D Mrsic-Flogel, and Sonja B Hofer. “Emergence of feature-specific connectivity in cortical microcircuits in the absence of visual experience”. In: *Journal of Neuroscience* 34.29 (2014), pp. 9812–9816.
- [44] Ho Ko et al. “Functional specificity of local synaptic connections in neocortical networks”. In: *Nature* 473.7345 (2011), pp. 87–91.
- [45] Simon Kornblith et al. “Similarity of neural network representations revisited”. In: *International conference on machine learning*. PMLR. 2019, pp. 3519–3529.
- [46] Tanishq Kumar et al. “Scaling Laws for Precision”. In: *arXiv preprint arXiv:2411.04330* (2024).
- [47] Christopher W Kuzawa et al. “Metabolic costs and evolutionary implications of human brain development”. In: *Proceedings of the National Academy of Sciences* 111.36 (2014), pp. 13010–13015.
- [48] International Brain Laboratory et al. “Standardized and reproducible measurement of decision-making in mice”. In: *Elife* 10 (2021), e63711.
- [49] Janne K Lappalainen et al. “Connectome-constrained networks predict neural activity across the fly visual system”. In: *Nature* (2024), pp. 1–9.
- [50] Gidon Levakov et al. “Mapping individual differences across brain network structure to function and behavior with connectome embedding”. In: *NeuroImage* 242 (2021), p. 118469.

- [51] W Johnson J Lindenstrauss and J Johnson. “Extensions of lipschitz maps into a hilbert space”. In: *Contemp. Math* 26.189-206 (1984), p. 2.
- [52] Lijuan Liu et al. “Neuronal connectivity as a determinant of cell types and subtypes”. In: *Research Square* (2023), rs-3.
- [53] Omer Mano et al. “Predicting individual neuron responses with anatomically constrained task optimization”. In: *Current Biology* 31.18 (2021), pp. 4062–4075.
- [54] Arnaud Messé, Marc-Thorsten Hütt, and Claus Christian Hilgetag. “Excitable dynamics simplify neural connectomes”. In: *Cell Reports Physical Science* (2025).
- [55] Lu Mi et al. “Connectome-constrained latent variable model of whole-brain neural activity”. In: *International Conference on Learning Representations*. 2021.
- [56] Emily R Oby et al. “Dynamical constraints on neural population activity”. In: *Nature Neuroscience* (2025), pp. 1–11.
- [57] Zahid Padamsey and Nathalie L Rochefort. “Paying the brain’s energy bill”. In: *Current opinion in neurobiology* 78 (2023), p. 102668.
- [58] Rodrigo Perin, Thomas K Berger, and Henry Markram. “A synaptic organizing principle for cortical neuronal groups”. In: *Proceedings of the National Academy of Sciences* 108.13 (2011), pp. 5419–5424.
- [59] Carsten K Pfeffer et al. “Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons”. In: *Nature neuroscience* 16.8 (2013), pp. 1068–1076.
- [60] Jean-Baptiste Poline and Matthew Brett. “The general linear model and fMRI: does love last forever?” In: *Neuroimage* 62.2 (2012), pp. 871–880.
- [61] Dean A Pospisil et al. “The fly connectome reveals a path to the effectome”. In: *Nature* 634.8032 (2024), pp. 201–209.
- [62] Gideon Rosenthal et al. “Mapping higher-order relations between brain structure and function with embedded vector representations of connectomes”. In: *Nature communications* 9.1 (2018), p. 2178.
- [63] Britton A Sauerbrei et al. “Cortical pattern generation during dexterous movement is input-driven”. In: *Nature* 577.7790 (2020), pp. 386–391.
- [64] Louis K Scheffer et al. “A connectome and analysis of the adult Drosophila central brain”. In: *elife* 9 (2020), e57443.
- [65] Ann E Sizemore et al. “The importance of the whole: topological data analysis for the network neuroscientist”. In: *Network Neuroscience* 3.3 (2019), pp. 656–673.
- [66] H Francis Song, Henry Kennedy, and Xiao-Jing Wang. “Spatial embedding of structural similarity in the cerebral cortex”. In: *Proceedings of the National Academy of Sciences* 111.46 (2014), pp. 16580–16585.
- [67] Sen Song et al. “Highly nonrandom features of synaptic connectivity in local cortical circuits”. In: *PLoS biology* 3.3 (2005), e68.
- [68] Carsen Stringer et al. “Spontaneous behaviors drive multidimensional, brainwide activity”. In: *Science* 364.6437 (2019), eaav7893.
- [69] Laura E Suárez et al. “Linking structure and function in macroscale brain networks”. In: *Trends in cognitive sciences* 24.4 (2020), pp. 302–315.
- [70] Bosiljka Tasic et al. “Shared and distinct transcriptomic cell types across neocortical areas”. In: *Nature* 563.7729 (2018), pp. 72–78.
- [71] Fabian David Tschopp, Michael B Reiser, and Srinivas C Turaga. “A connectome based hexagonal lattice convolutional network model of the Drosophila visual system”. In: *arXiv preprint arXiv:1806.04793* (2018).
- [72] Kerem Uzel, Saul Kato, and Manuel Zimmer. “A set of hub neurons and non-local connectivity features support global brain dynamics in *C. elegans*”. In: *Current Biology* 32.16 (2022), pp. 3443–3459.
- [73] Bertha Vázquez-Rodríguez et al. “Gradients of structure–function tethering across neocortex”. In: *Proceedings of the National Academy of Sciences* 116.42 (2019), pp. 21219–21227.
- [74] Edgar Y Walker et al. “Inception loops discover what excites neurons most using deep predictive models”. In: *Nature neuroscience* 22.12 (2019), pp. 2060–2065.

- [75] Eric Y Wang et al. “Foundation model of neural activity predicts response to new stimulus types”. In: *Nature* 640.8058 (2025), pp. 470–477.
- [76] Daniel LK Yamins and James J DiCarlo. “Using goal-driven deep learning models to understand sensory cortex”. In: *Nature neuroscience* 19.3 (2016), pp. 356–365.
- [77] Guangyu Robert Yang and Xiao-Jing Wang. “Artificial neural networks for neuroscientists: a primer”. In: *Neuron* 107.6 (2020), pp. 1048–1070.
- [78] Hongkui Zeng. “What is a cell type and how to define it?” In: *Cell* 185.15 (2022), pp. 2739–2755.
- [79] Hongkui Zeng and Joshua R Sanes. “Neuronal cell-type classification: challenges, opportunities and the path forward”. In: *Nature Reviews Neuroscience* 18.9 (2017), pp. 530–546.
- [80] Huanqiu Zhang et al. “Hippocampal spatial representations exhibit a hyperbolic geometry that expands with experience”. In: *Nature Neuroscience* 26.1 (2023), pp. 131–139.
- [81] Yuansheng Zhou and Tatyana O Sharpee. “Hyperbolic geometry of gene expression”. In: *Isience* 24.3 (2021).
- [82] Yuansheng Zhou, Brian H Smith, and Tatyana O Sharpee. “Hyperbolic geometry of the olfactory space”. In: *Science advances* 4.8 (2018), eaaq1458.
- [83] Jun Zhuang et al. “An extended retinotopic map of mouse cortex”. In: *elife* 6 (2017), e18372.

A Methods

A.1 MICrONS Dataset and Experiments

In this paper, all experiments and analysis are conducted using the MICrONS minnie65 V943 dataset, which contains 205,277,810 synapses and 90,230 neurons [15]. The overall sparsity of pairwise connectivity is 0.127%. Among the recorded cells, 63,758 are excitatory neurons, 7,849 are inhibitory neurons, and 18,693 are non-neuronal cells, primarily glial cells. Regionally, 57,781 cells are located in V1, 21,184 in RL, and 10,480 in AL.

Proofreading plays a critical role in improving the quality of automated EM segmentation by resolving incorrect associations between presynaptic and postsynaptic partners. This process primarily addresses two types of reconstruction errors: false merges and false splits. False merges occur when segmented objects are erroneously grouped (e.g., an axon or dendrite incorrectly assigned to a soma), while false splits arise when objects are mistakenly separated (e.g., an axon or dendrite omitted from its corresponding soma). To correct false merges, the reconstruction is *cleaned* by removing incorrectly associated segments, whereas false splits are addressed by *extending* the reconstruction to restore missing segments. A detailed protocol for morphological proofreading is provided in [15, 22].

For axonal proofreading, 321 neurons are classified as extended, 714 as clean, and the remaining neurons are unproofread [24]. Similarly, for dendritic proofreading, 44,163 neurons have been completely merged and properly cleaned – either manually or using the NEURD algorithm – resulting in fully reconstructed dendrites, while the remaining neurons are unproofread [24, 12]. To increase the number of axonally trustworthy neurons, we relax the selection criterion to include neurons labeled as either extended or clean. In contrast, reconstructed dendrites are generally complete and require minimal extension, resulting in a greater number of neurons classified as dendritically trustworthy. However, only 845 neurons (fewer than 1% of the total) meet the structural confidence criteria for both axonal and dendritic aspects. In analyzing the geometric representation of the connectome (Section B.1), we will consider different subsets, including all excitatory neurons ($N = 716$) and excitatory neurons in V1 ($N = 675$).

For the analyses in Section 2.1 and the baseline experiment in Section 2.3, \mathbf{C}^{out} is computed using connections exclusively *to* E/I neurons with fully reconstructed dendrites, whereas \mathbf{C}^{in} is computed using connections exclusively *from* E/I neurons with extended or clean axons. Table S1 details the dimensionality of \mathbf{W}^{in} and \mathbf{W}^{out} under different experimental settings, and Tables S2 and S3 detail the cell identities. Neurons that are isolated in the connectivity are excluded, as their inclusion would result in entire rows and/or columns of NaN values in \mathbf{C} (corresponding to 0 values in \mathbf{W}). Additional comparative studies, including those involving different connectome types \mathbf{W} or the inclusion of synapses *to* or *from* unproofread neurons and glial cells, are discussed in Section 2.3 (Table S1, "Noise" row).

In this study, we align the anatomical and functional properties of excitatory neurons that are both functionally coregistered (i.e. located within the volume of dense reconstruction, also see Figure S9) and structurally proofread. These selection criteria are implemented using CAVE [23]. While restricting analysis to neurons that meet both criteria enhances data reliability, it significantly reduces the number of available neurons. Requiring proofreading of both the axon and dendrite would shrink the dataset even further, so we include coregistered neurons if at least one of these compartments has been proofread (**Prf**). Table S4 provides a detailed comparison of selection criterion [22]. We use the number of neurons listed under "Number of **Coreg + Prf** Neurons" in the linear model presented in Section 2.3. Across all scans, only 497 **Coreg + Prf** neurons (excluding duplicate appearances) meet these criteria.

During the recordings, three dynamic stimulus types, natural movies, global directional parametric stimuli ("Monet"), and local directional parametric stimuli ("Trippy"), were presented in blocks interleaved with 8–12 off-periods [15, 22].

A.2 Clique Topology Method for Uncovering Hidden Geometries in Connectomic and Activity Structures

It has been extensively discussed in previous literature that the clique topology method could be used to detect the geometric organization of the correlation or connectivity matrix, including purely random or embedded in Euclidean or hyperbolic space [32, 82, 80]. In this analysis, we extend these

discussions by using the input as a symmetric matrix \mathbf{C} , which quantifies the pairwise correlation between neurons i and j . The algorithm applies a step function to the matrix at various thresholds, retaining only the entries that exceed the specified value and generating corresponding topological graphs. These graphs are characterized by the number of cycles (holes) in different dimensions (1, 2, or higher). These counts, known as Betti numbers $\bar{\beta}_n$, are generally small at low thresholds (fully connected networks) or high thresholds (sparse networks where most nodes are not interconnected) [31]. Plotting them as a function over the entire density range ρ of links (from the highest threshold to the lowest threshold) produces Betti curves $\beta_n(\rho)$, $0 \leq \rho \leq 1$. Clearly:

$$\bar{\beta}_m = \int_0^1 \beta_m(\rho) d\rho. \quad (\text{S1})$$

Typically, the first three Betti curves are sufficient to reveal the underlying structure, and our analysis adopts this convention. Figure S7 provides an illustrative example of computing Betti curves for a simple graph.

Uniformly sampling points from an n -dimensional Euclidean space is straightforward through high-dimensional uniform distribution sampling. Previous literature has also addressed sampling from n -dimensional hyperbolic space [82, 80]. In the polar coordinate setting, both the curvature K and the maximal radius R_{\max} can alter the geometry of the space. Specifically, changing K only rescales the radial coordinates of the nodes without affecting the topological properties of the network. For simplicity, we fix $\zeta := \sqrt{-K} = 1$ and treat R_{\max} as a hyperparameter. Each point is sampled with uniformly distributed angles and a radius $r \in [R_{\min}, R_{\max}]$, following the distribution $\rho(r) \sim \sinh^{(d-1)r}$, which is approximately proportional to $e^{(d-1)r}$ when $r \gg 1$. In [82], R_{\min} is set to $0.9R_{\max}$ to constrain the sampled points to have similar radii (i.e. similar to torus-like shape in \mathbb{R}^3). Following the convention in [80], we set $R_{\min} = 0$, leading to a less constrained distribution. Due to this difference in setting, the Betti curves presented in our work are qualitatively more similar to those in [80] and differ from those in [82] (e.g. decreasing prominence of peaks as m increases). Consequentially, unlike the L^2 distance in Euclidean space, the metric in hyperbolic geometry adheres the hyperbolic law of cosines:

$$\cosh(\zeta x) = \cosh(\zeta r) \cosh(\zeta r') - \sinh(\zeta r) \sinh(\zeta r') \cos \Delta\theta, \quad (\text{S2})$$

where x is the hyperbolic distance, r, r' the corresponding radii (of the two points) from the origin, and $\Delta\theta$ is the angle in between [80]. To incorporate noise into the simulation, we consider the magnitude of Gaussian noise ϵ , which is scaled based on the geometric distance $\bar{\mathbf{D}}$ calculated through Equation (S2):

$$\mathbf{D}_{ji} = \bar{\mathbf{D}}_{ji} \cdot (1 + \epsilon \cdot \mathcal{N}(0, 1)). \quad (\text{S3})$$

We use the negative distance, $\mathbf{C}_{ji} = \mathbf{C}_{ij} = -\mathbf{D}_{ji}$, to quantify the correlation between neuron i and j . Any other order-reversing operation, including the reciprocal, could also accomplish this transformation, but we follow the previous literature and use the inverse for its stability in scaling [82, 80]. In the main text, we primarily use Pearson's correlation to calculate \mathbf{C} . However, any operation that produces bounded and symmetric results can be applied, including Jensen-Shannon divergence, Hellinger distance, Jaccard distance, or Czekanowski-Dice distance for quantifying connectome similarity, and cosine similarity for quantifying activity similarity [66]. In Section B.3, we demonstrate the robustness of the low-dimensional embedding geometry by using cosine similarity as an alternative method to compute \mathbf{C} for activities.

In all simulations, we set N to match the size of each distance matrix as the input, with $\delta = 0.1$, $\epsilon = 0.0625$, $d = 2$ (or 3, see Section B.4), while fine-tuning the hypodermic radius R_{\max} to achieve optimal alignment of the top three Betti curves. The optimization objective is to minimize the cumulative absolute pointwise error (Frobenius norm) of Betti curves using experimental and synthetic data. We validate the results by comparing the Betti number, $\bar{\beta}_m$. The analysis will display ± 1 standard deviation in all such cases. Through experiments, we observe that both the shape and magnitude of $\beta_m(\rho)$ are highly dependent on ϵ , which can be interpreted as representing various levels of mixing with random matrices (Equation S3). The parameter values are selected based on prior literature with minor adjustments [82, 80]. Benefiting from recent developments, the scope of network size considered in the topological analysis presented in this paper surpasses that of previous works multiple times (Section A.6).

A key difference between clique topology analysis and hyperbolic MDS (Section 2.3) is that the latter applies dimensionality reduction to empirical data by minimizing the stress function and does

not guarantee a uniform distribution of embedded points. The validity of our topological analysis is supported by the relatively uniform distribution of the targeted neurons, as confirmed through Euclidean non-metric MDS, which satisfies the assumptions underlying this analysis (Figure S6) [82, 80]. In practice, we initialize the hyperbolic MDS algorithm with the converged Euclidean embedding of the same dimension and observe improvements in stress, highlighting the enhanced embedding accuracy of hyperbolic MDS.

A.3 Centered Kernel Alignment

Centered kernel alignment (CKA) gauges how similar two sets of neural activations are by comparing their centered Gram (kernel) matrices [45]. Given activation matrices $X \in \mathbb{R}^{n \times p}$ and $Y \in \mathbb{R}^{n \times q}$ for the same n features (in this study, the latent dimension of the reduce representation), form linear kernels $K = XX^\top$, $L = YY^\top$, and mean-center them with $H = I - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$ to obtain $K_c = HKH$, $L_c = H L H$. CKA is the cosine between these two matrices:

$$\text{CKA}(K, L) = \frac{\text{tr}(K_c L_c)}{\|K_c\|_F \|L_c\|_F} = \frac{\text{HSIC}(K, L)}{\sqrt{\text{HSIC}(K, K) \text{HSIC}(L, L)}}, \quad (\text{S4})$$

where HSIC is the Hilbert-Schmidt independence criterion. For the common linear CKA variant this reduces to:

$$\text{CKA}(X, Y) = \frac{\|Y^\top X\|_F^2}{\|X^\top X\|_F \|Y^\top Y\|_F}, \quad (\text{S5})$$

which lies in $[0, 1]$: values near 1 mean the two representations rank stimuli in almost the same way, whereas values near 0 indicate unrelated structure. Because CKA is invariant to any orthogonal rotation of features and to isotropic rescaling, it avoids spurious differences that plague canonical-correlation-based metrics, and, unlike fully linear-invariant measures, still distinguishes wide layers when $p, q > n$.

A.4 Principal Angle Analysis

Principal angle analysis quantifies the overlap between two neural-activity subspaces [9, 42]. Given column-orthonormal bases $U \in \mathbb{R}^{d \times p}$ and $V \in \mathbb{R}^{d \times q}$, form $U^\top V$ and compute its singular values $\sigma_1 \leq \dots \leq \sigma_{\min(p, q)}$. The principal angles are $\theta_i = \arccos \sigma_i$ with $0^\circ \leq \theta_i \leq 90^\circ$. The first angle θ_1 measures the most-aligned pair of directions, whereas subsequent angles describe alignment after removing previously matched components. The average squared cosine, $\frac{1}{p} \|U^\top V\|_F^2 = \frac{1}{p} \sum_i \cos^2 \theta_i$, equals the fraction of variance that one subspace projects onto the other. Small angles (near 0°) indicate shared coding dimensions, while large angles (near 90°) denote independence. Because the computation scales efficiently when $p, q \ll d$, principal angle analysis is widely used to compare population codes across tasks, time points, and brain regions [26].

A.5 Linear Model

In the baseline experiment, glial cells are excluded to focus on interactions between excitatory (E) and inhibitory (I) neurons. Synaptic sign is assigned according to Dale’s law, with inhibitory neurons having negative output weights and excitatory neurons having positive output weights. Neuronal identity is determined based on morphological features [15]. \mathbf{C}^{in} and \mathbf{C}^{out} are computed using the same approach described in Section 2.1. Synaptic information from all functionally coregistered neurons across scans ($N = 497$), including neurons without complete structural confidence, is concatenated and subjected to dimensionality reduction. To avoid overfitting, we bin the activity data, compute \mathbf{C}^{act} using the concatenated information from odd-numbered bins, and evaluate acc_{act} on the remaining data.

While the effectiveness of using correlations derived from the input connectome are intuitive, the ability of the output connectome correlations to predict activity is less immediately clear [21, 20]. Previous studies employed a "digital twin" model to predict neural responses to a diverse stimulus set of natural and parametric movies, characterizing the *in silico* functional properties of imaged neurons. These studies revealed an overarching like-to-like principle – neurons with highly correlated responses to natural videos tend to be connected across multiple layers and visual cortical areas – providing insights into the effectiveness of \mathbf{C}^{out} [74, 75, 22].

A.6 Code availability

All analysis code for this study will be made publicly available upon acceptance.

B Supplementart Materials

	Number of Input	Number of Output
Baseline	1035	25497
E Only	850	23019
I Only	185	2478
Noise	71607	71607

Table S1: **Number of Input and Output Neurons Used in Calculating C^{in} and C^{out}** : Corresponds to the values of M_1, M_2 . All rows, except for "Noise", include only neurons that are proofread on the axon side (for input) and the dendrite side (for output) correspondingly.

	L1	L2-3	L4	L5	L6	WM
Input	10	407	355	182	77	4
Output	402	7751	7453	3828	4719	1344
Targeted	1	205	157	92	42	0

Table S2: **Layer Classification of Input, Output, and Target Neurons in the Baseline Experiment**: L: Layer; WM: White Matter.

	V1	RL	AL	Others
Input	899	88	48	0
Output	18895	6158	413	31
Targeted	382	68	47	0

Table S3: **Region Classification of Input, Output, and Target Neurons in the Baseline Experiment**: V1: Primary Visual Cortex; AL: Anterolateral Area; RL: Rostrolateral Area.

B.1 Representation Space of Neuronal Connectome is Low-Dimensional Hyperbolic Embedded

See Figure S8, S9, S10, S11, and S12.

B.2 Additional Betti Curve Analysis on Representation Space of Activities Traces for All Scans

See Figure S15 and S16.

B.3 Betti Curve Analysis on Representation Space of Activities Traces Using Cosine Similarity

See Figure S17 and S18.

B.4 Equivalence Between 2D and 3D Hyperbolic Embeddings

See Figure S19. A similar discussion in Euclidean space is well addressed through the classical Johnson–Lindenstrauss lemma, which establishes the existence of a linear map $f : \mathbb{R}^N \rightarrow \mathbb{R}^n, N > n$, that nearly preserves the distances between points [51, 39]. Here, we shift the focus from pairwise distance preservation to topological feature preservation and observe a perfectly linear relationship in the hyperbolic radius under pointwise transformation from \mathbb{H}^3 to \mathbb{H}^2 . Our goal is not to prove the universality of such transformations but rather to highlight the embeddings in \mathbb{H}^2 discussed previously, including Section B.1, B.2, B.3, have direct counterparts in \mathbb{H}^3 , as is typical in prior literature [82, 80]. The range of $R_{\text{max}}^{d=2}$ shown in Figure S19 aligns with the discussions in earlier sections.

Session & Scan	# Recorded Neurons	# Coreg Neurons	# Coreg + Prf Neurons
Session 4, Scan 7	8395	1018	82
Session 8, Scan 5	10579	1123	68
Session 6, Scan 4	8826	1152	59
Session 9, Scan 4	8415	1027	56
Session 6, Scan 6	8551	1052	54
Session 5, Scan 6	9254	997	54
Session 5, Scan 7	8862	1000	41
Session 6, Scan 2	8756	1021	39
Session 7, Scan 3	9296	743	38
Session 7, Scan 5	9034	631	39
Session 9, Scan 3	8661	1170	37
Session 5, Scan 3	8790	426	35

Table S4: **Available Neurons under Different Selection Criteria, Categorized by Session and Scan:** **Coreg**: functionally coregistered. **Prf**: proofread on at least one side (axon or dendrite).

B.5 Reconstruction of Raw Fluorescence Traces

See Figure S27 and Table S5.

	Raw Trace	Deconvolved Trace
HypEmbed \mathbf{C}^{in}	24.1%	8.66%
\mathbf{C}^{in}	23.1%	7.70%
HypEmbed \mathbf{C}^{out}	22.8%	8.05%
\mathbf{C}^{out}	15.8%	4.92%
\mathbf{C}^{act}	36.4%	13.4%
Mean of \mathbf{C}^{act}	7.00%	1.36%

Table S5: **Performance (*acc*) Comparison between Reconstructing Raw Fluorescence Traces and Detrended, Deconvolved Traces (Used in Main Text) via the Linear Model:** Different forms of \mathbf{C} are used to compute the coefficients α under the baseline experiment setup (Section 2.3 and Figure 3).

B.6 Using Embedding Results Demonstrates Robustness in Prediction Under Raw Connectome Perturbations

See Figure S29. Here, we perturb the connectivity matrices using the baseline experimental setting but with binary connections, concatenating all targeted neurons across scans ($N = 497$), \mathbf{W}^{in} and \mathbf{W}^{out} (Section 2.1). A $k\%$ perturbation in Figure S29 refers to the random deletion of $k\%$ of synapses (i.e. false negative entries) and the simultaneous addition of $k\%$ of "fake" synapses (i.e. false positive entries). This approach preserves the overall connection sparsity (3.87% for \mathbf{W}^{in} and 0.52% for \mathbf{W}^{out}). The perturbation is subject to the following constraints: 1) It must not result in any isolated targeted neuron, as this would lead to undefined correlation with others, rendering their activity unusable for reconstruction (i.e. undefined α entries). Omitting neurons could bias the comparison (Figure S21); 2) It must satisfy Dale's constraint for \mathbf{W}^{in} ; since all targeted neurons are excitatory, \mathbf{W}^{out} only contains 0 and 1 entries. Due to constraint 1) and the limited number of axonally proofread neurons (Table S1), perturbation of \mathbf{W}^{in} will be biased, therefore Figure S29 presents results only for \mathbf{W}^{out} .

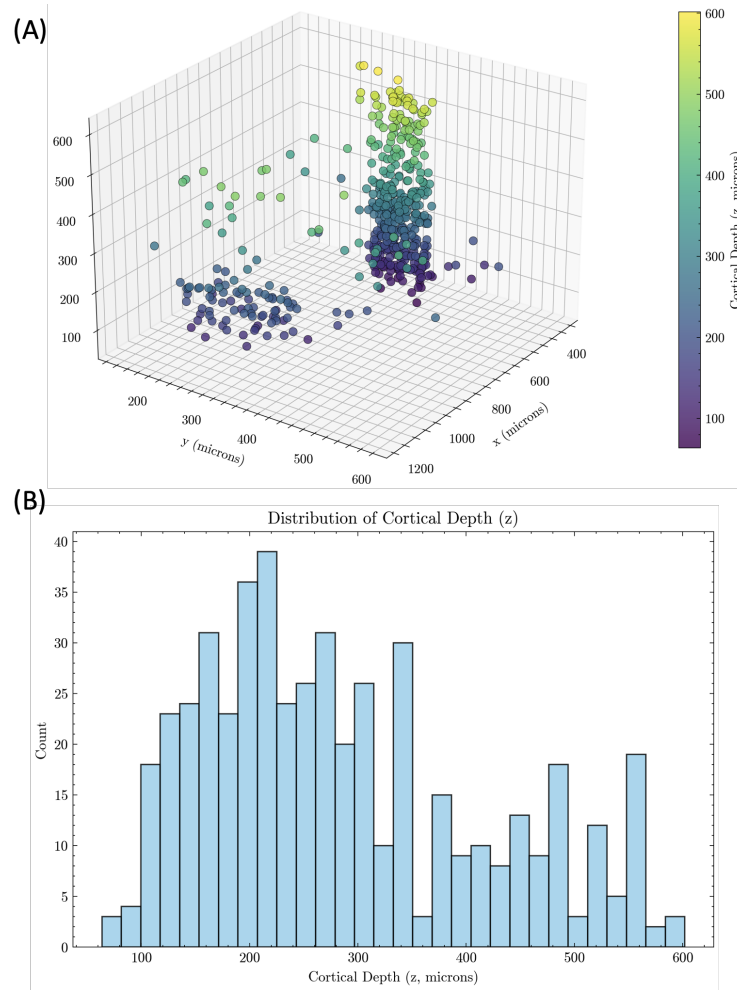


Figure S1: **Spatial Organization of Target Neurons:** (A) Somatic locations of the 497 targeted neurons, color-coded by cortical depth (μm); The z -coordinate should be flipped in order to be aligned with the physical orientation; (B): Histogram of cortical depth. Results are consistent with Table S2 and Table S3.




Figure S2: **Statistical Comparison of On-Diagonal and Off-Diagonal Elements of "Correlation of Correlation" Matrix:** Expanding on Figure 1D1, we present results from additional scans to demonstrate the consistency and robustness of our observations. (A): Session 4, Scan 7; (B) Session 9, Scan 4; (C): Session 5, Scan 3. The p_1 to p_3 values, defined identically to those in Figure 1D1, are provided in the captions.




Figure S3: Principal Angle Analysis: Expanding on Figure 1D2, we compute the principal angle between subspaces (flats) constructed by C^{act} and C^{in} (red), C^{out} (blue), and random matrices (green) [9, 42]. Statistical significance is indicated by p_4 for the difference between the red and blue curves, p_5 for the red and green curves, and p_6 for the blue and green curves. As discussed in the main text, the input affinity subspace exhibits the strongest alignment with the activity space. (A) Session 4, Scan 7; (B) Session 9, Scan 4; and (C) Session 8, Scan 5 (Figure 1).




Figure S4: Neurons with Strong Activity Correlation Are More Likely to Form Synaptic Connections: Unlike Figure 1, Figure S2, and Figure S3, which use non-local connectivity measures (C^{in} and C^{out}) to probe the structure-function relationship, here we analyze the direct synaptic connectivity matrix between targeted neurons ($W \in \mathbb{R}^{N \times N}$) and confirm the like-to-like principle [22]. (A) Session 4, Scan 7; (B) Session 9, Scan 4; and (C) Session 8, Scan 5.

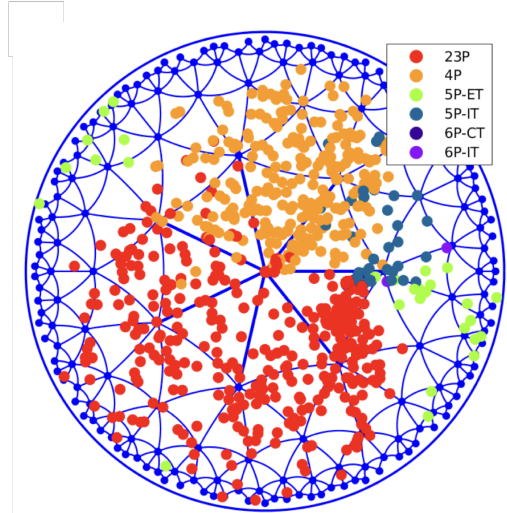


Figure S5: Embedding neurons from Figure 1B based on their output similarity into \mathbb{H}^2 with hyperbolic t-SNE algorithm, clustered by cell type and cortical layer: The resulting projection is displayed on a unit circle [81].




Figure S6: **2D Euclidean MDS Embeddings of Input and Output Dissimilarity in the Baseline Experiment:** The dissimilarity is calculated through $1 - \text{Pearson correlation}$. Colored by frequency.




Figure S7: **Illustrative example of order-based analysis of symmetric matrices:** (A): An arbitrary correlation matrix where each cell's value is converted to its relative rank; (B): By iteratively adjusting the threshold (filtration density), a sequence of binary adjacency matrices (and graphs) is generated. β_0 counts the number of connected components, while other β_m values count the number of m -cycles, such as 1-cycles (squares), 2-cycles (octahedra), and 3-cycles (orthoplex, not shown), which may be created, persist, and disappear at different densities. The colored shadow suggests the presence of a cycle; (C): The overall Betti curves $\beta_m(\rho)$ are shown by combing all the discrete instances discussed in (B). This example is reproduced from Dr. Nicole Sanderson's tutorial at IPAM's Mathematical Approaches for Connectome Analysis Workshop.

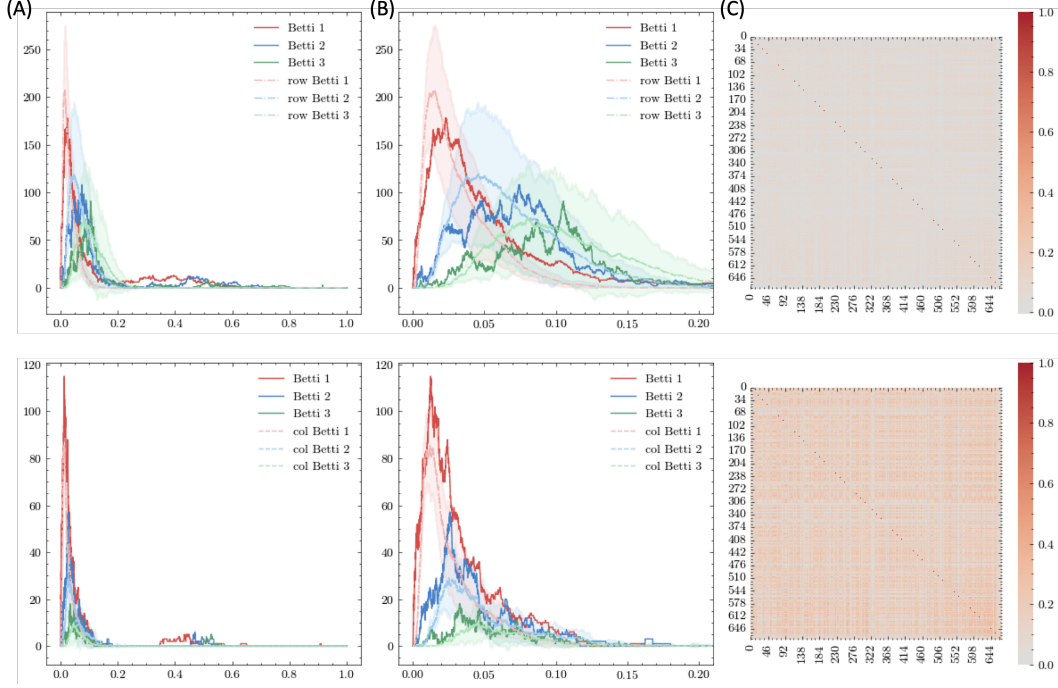


Figure S8: 2D Hyperbolic Embedding of Structural Representation of Binary Connectome Among Proofread Excitatory Neurons in V1: To compare with Figure 1E in the main text, square Pearson correlation matrices of the binary connectome are computed for all proofread excitatory neurons in V1 ($N = 675$). As discussed in Section 2.3, structurally incomplete (unchecked) yet informative data is included. C^{out} (top row) and C^{in} (bottom row) are calculated using the rectangle matrix that includes all recorded excitatory and inhibitory neurons from MICrONS Version 943[15]. The reconstructed neurons are distributed across two distinct cylindrical regions, corresponding to the V1 area (94.0%) and the RL/AL area (6.0%). Overall, their structural similarities, both for inputs and outputs, form two clusters, violating the assumption of uniform distribution in Betti analysis [80]. Here, we focus specifically on V1 neurons; however, the results remain consistent for RL/AL neurons (Figure S12). A finer classification that groups neurons by cell types produces smoother and more regular Betti curves, thereby achieving better alignment with hyperbolic geometry in terms of topological feature mapping (Figure S13 and S14): (A): Betti curves for groundtruth and synthetic data, sampled from a 2D hyperbolic space (see Section A.2), averaged across 10 trials due to the high computational memory costs; results are shown with best fits via (pointwise) mean squared error (MSE) comparison to the mean curve (dashed lines) with optimized R_{max} of 18 and 14 respectively (grid size = 1); (B): Zoomed-in view of (A) highlighting the low edge density region; (C): Heatmap visualization of each correlation matrix. Using alternative correlation measures instead of Pearson correlation yields consistent results (not shown), as discussed in Section A.2 and [66].

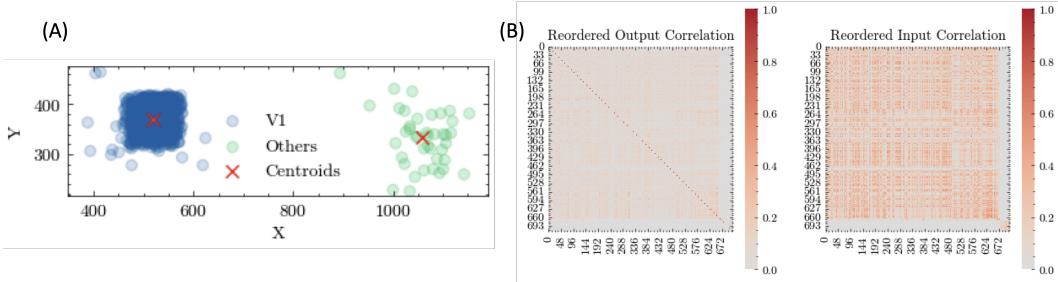


Figure S9: Spatial Distribution of Reconstructed Neurons: (A): Cross-sectional projection on the XY plane (in μm) of all reconstructed neurons ($N = 716$), clustered using the K-means algorithm [15]; (B): Sorted C^{out} and C^{in} , arranged according to the labels in (A), and the matrices correspond to Figure S8.

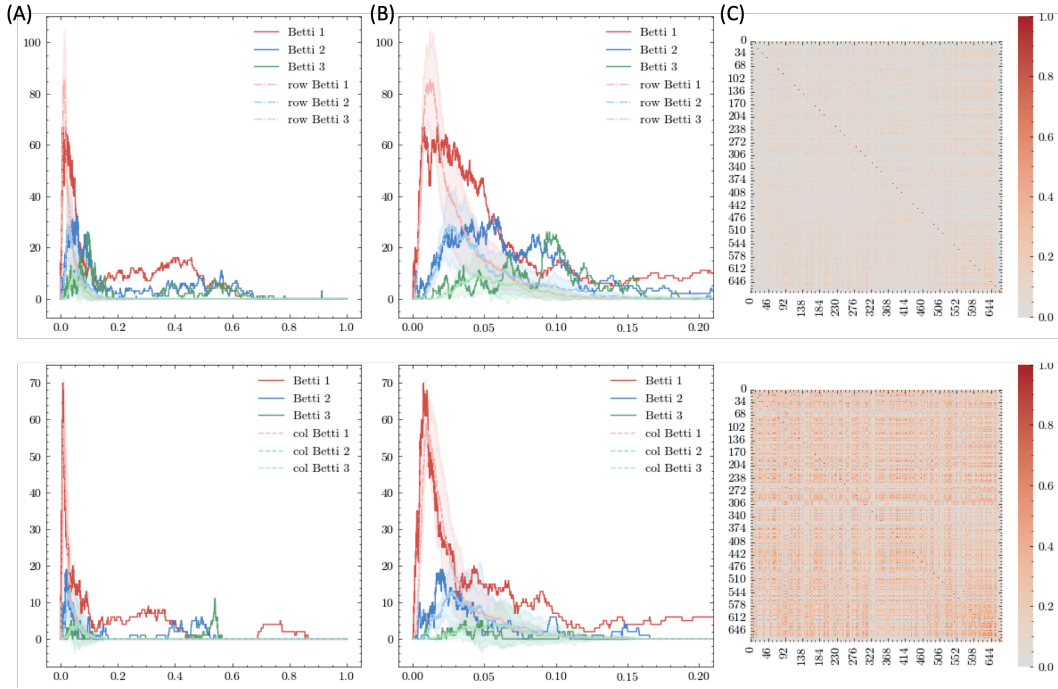


Figure S10: **2D Hyperbolic Embedding of the Structural Representation of the Weighted Connectome Among Proofread Excitatory Neurons in V1:** Similar to Figure S8, but C^{out} and C^{in} are computed using synaptic counts between neuronal connections, with optimized R_{max} of 14 and 13, respectively. As in Figure S8, the distinct emergence of a second peak in the high edge density region is consistent with observations in [82] and may be attributed to the significant nonuniform distribution of points in \mathbb{H}^n .




Figure S11: **Direct Comparison of Betti Values $\bar{\beta}_m$ Based On Groundtruth and Synthetic Data:** Results derived from Figure S8 and Figure S10; Vertical lines indicate the optimized R_{max} values for C^{out} and C^{in} (4 in total, for binary and weighted connectome respectively); Horizontal lines represent the $\bar{\beta}_m$ values, $m = 1, 2, 3$, calculated based on the groundtruth. The deviations from the intersection points ("*" for C^{out} and "o" for C^{in}) to the shadows indicate the effectiveness of approximations. All results are well-aligned except Betti 1 due to the occurrence of the second peaks at the high edge density region. The results further demonstrate the accuracy of the fits.

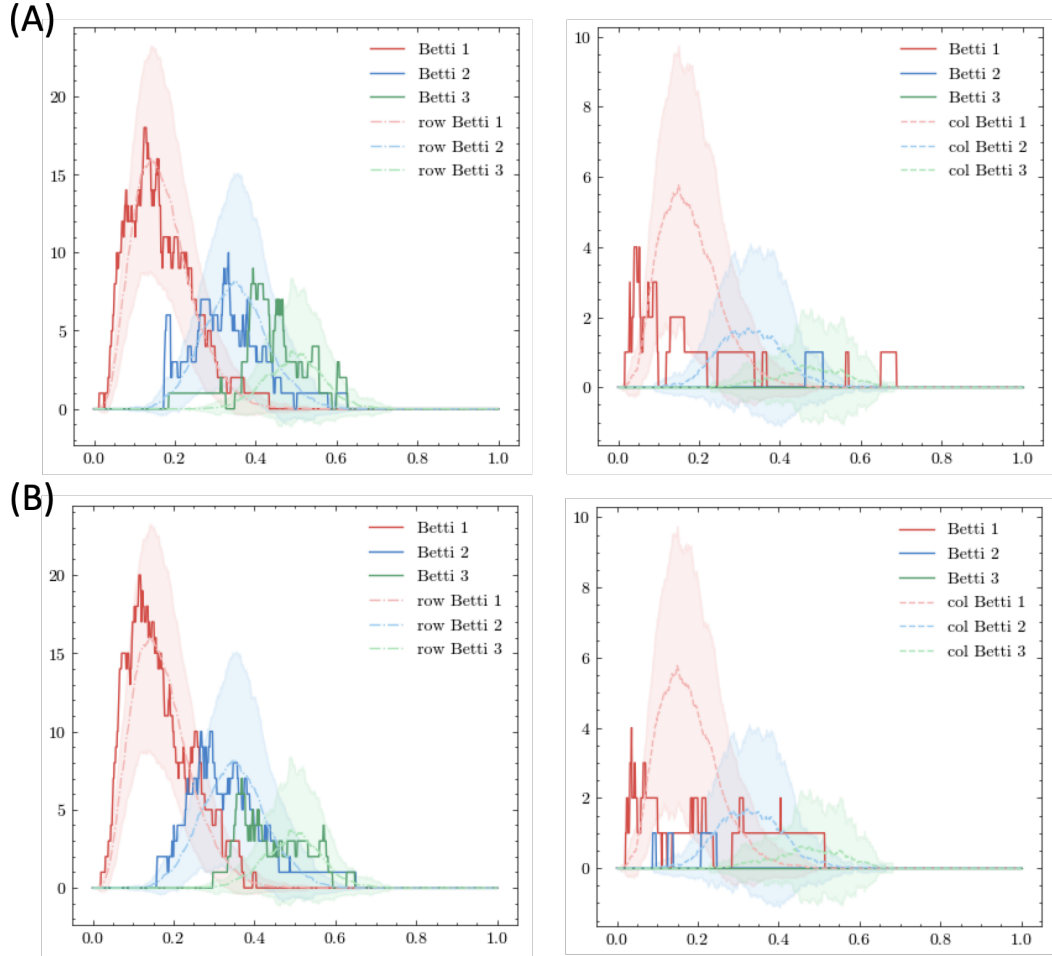


Figure S12: **2D Hyperbolic Embedding of the Structural Representation of the Weighted Connectome Among Proofread Excitatory Neurons in RL/AL:** Similar to Figure S8 and S10, this analysis focuses on neurons in higher visual areas (RL/AL) ($N = 41$) using (A) binary connectome or (B) weighted connectome to evaluate C^{out} (left column) and C^{in} (right column). Synthetic data are presented with optimized R_{max} , averaged across 100 trials. Similar to observations in V1, the topological features of the representation space formed by C^{in} are less pronounced and can be approximated using a small R_{max} , making it more similar to Euclidean space. This may be attributed to neuronal axons traveling much greater distances than dendrites, or alternatively, to an imbalance in reconstruction.

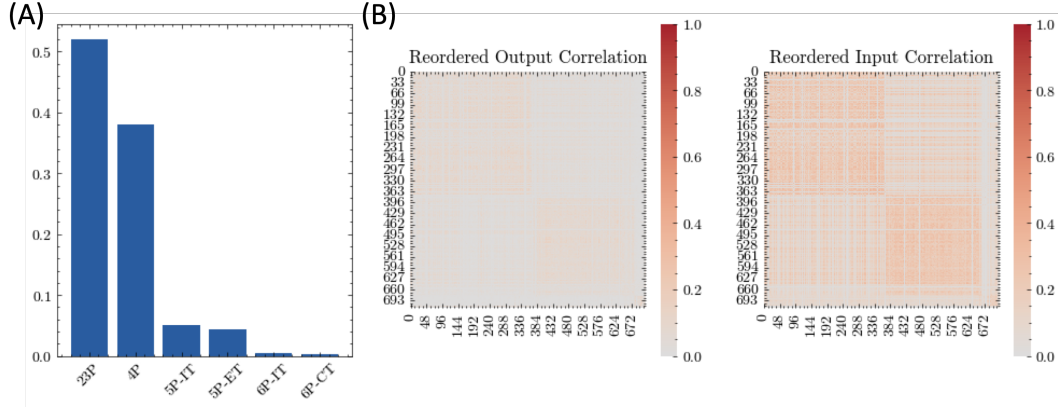


Figure S13: **Cell Type of Reconstructed Neurons**: (A): Labeled cell types of all reconstructed neurons, where layer 2-3 pyramidal cells (23P, 52.0%) and layer 4 pyramidal cells (4P, 38.0%) are dominant [15]; (B): Sorted C^{out} and C^{in} , arranged according to the labels in (A).

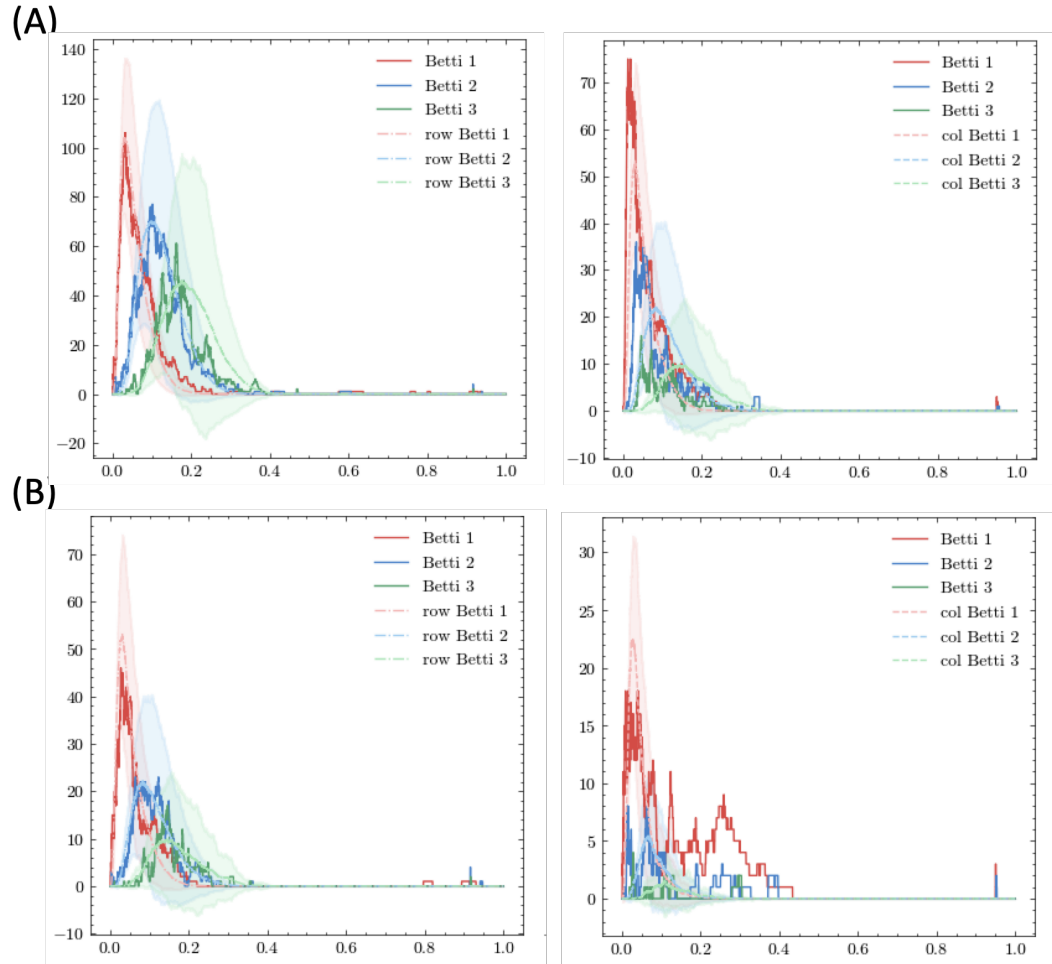


Figure S14: **2D Hyperbolic Embedding of the Structural Representation Among Proofread Excitatory Neurons in 4P**: Similar to Figure S12, this analysis focuses on layer 4 pyramidal cells ($N = 272$) using (A) binary connectome or (B) weighted connectome to evaluate C^{out} (left column) and C^{in} (right column). Synthetic data are presented with optimized R_{max} , averaged across 100 trials.

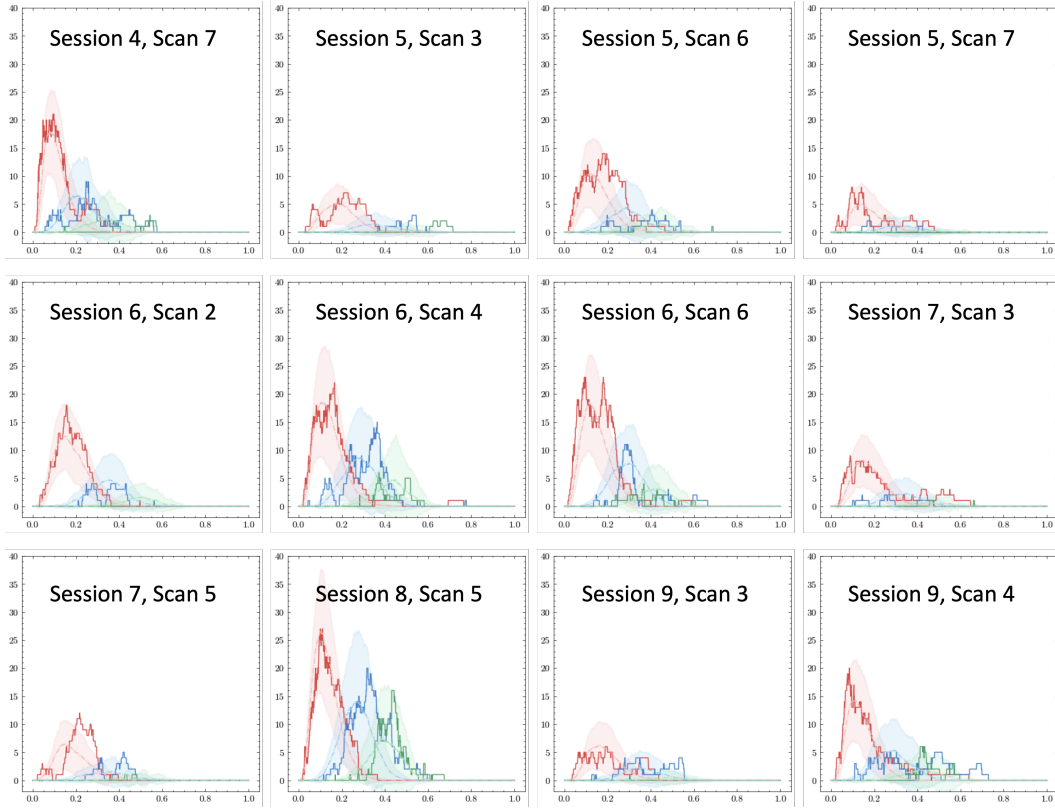


Figure S15: **Betti Analysis of All Activity Traces Indicates a Universal Hyperbolic Embedding:** Repeat the analysis in Figure 1E for all scans by comparing the alignment of the top three Betti curves between groundtruth and synthetic data, averaged across 100 trials. Use the parameters described in Section A.2, with varying R_{\max} values optimized within the range of 5 to 20 (grid size = 0.1).




Figure S16: **Positive Relationship Between the Number of Neurons in the Scan and the Embedded Hyperbolic Radius:** Results derived from Figure S15; including more neurons increase the effective dimensionality of the representation space, thereby necessitating greater capacity.

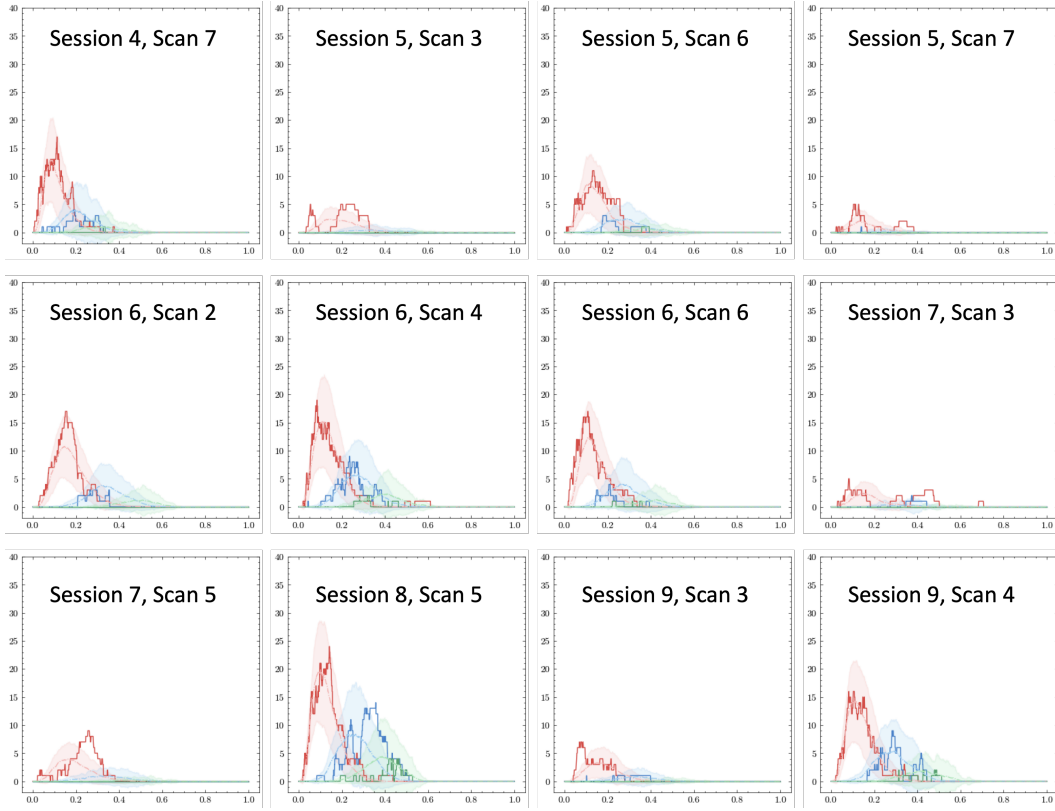


Figure S17: **Betti Analysis of All Activity Traces Using Cosine Similarity:** Similar to Figure S15, but using cosine similarity to quantify activity proximity.




Figure S18: **Positive Relationship Between the Number of Neurons in the Scan and the Embedded Hyperbolic Radius Using Cosine Similarity as the Correlation Measure:** Results derived from Figure S17; Similar to Figure S16. Under this measure, the representation space becomes less complex, reducing the magnitude of the Betti curves and an overall decrease in the hyperbolic radii R_{\max} .

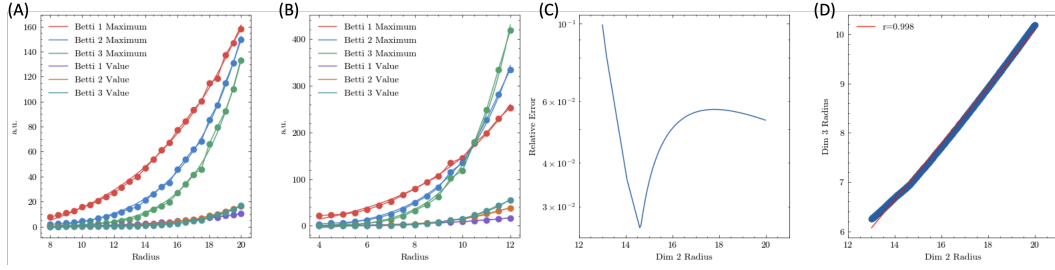


Figure S19: Transformation between 2D and 3D hyperbolic embedding can be achieved through an approximately linear relationship of R_{\max} : Specific case for $N = 217$ and noise level $\epsilon = 0.0625$ averaged over 100 trials. (A-B): Presentations of $\bar{\beta}_m$ and $\max(\beta_m(\rho))$, $m = 1, 2, 3$ for 2D (A) and 3D (B). Each curve is interpolated using a simple exponential function and transformed into the continuous case; (C-D): Optimized cumulative pointwise error (C) and the mapping relationship (D) between $R_{\max}^{d=2}$ and $R_{\max}^{d=3}$. Note that when $R_{\max}^{d=2} \leq 13$, the $\bar{\beta}$ values are small (B part), which could lead to stability issues.

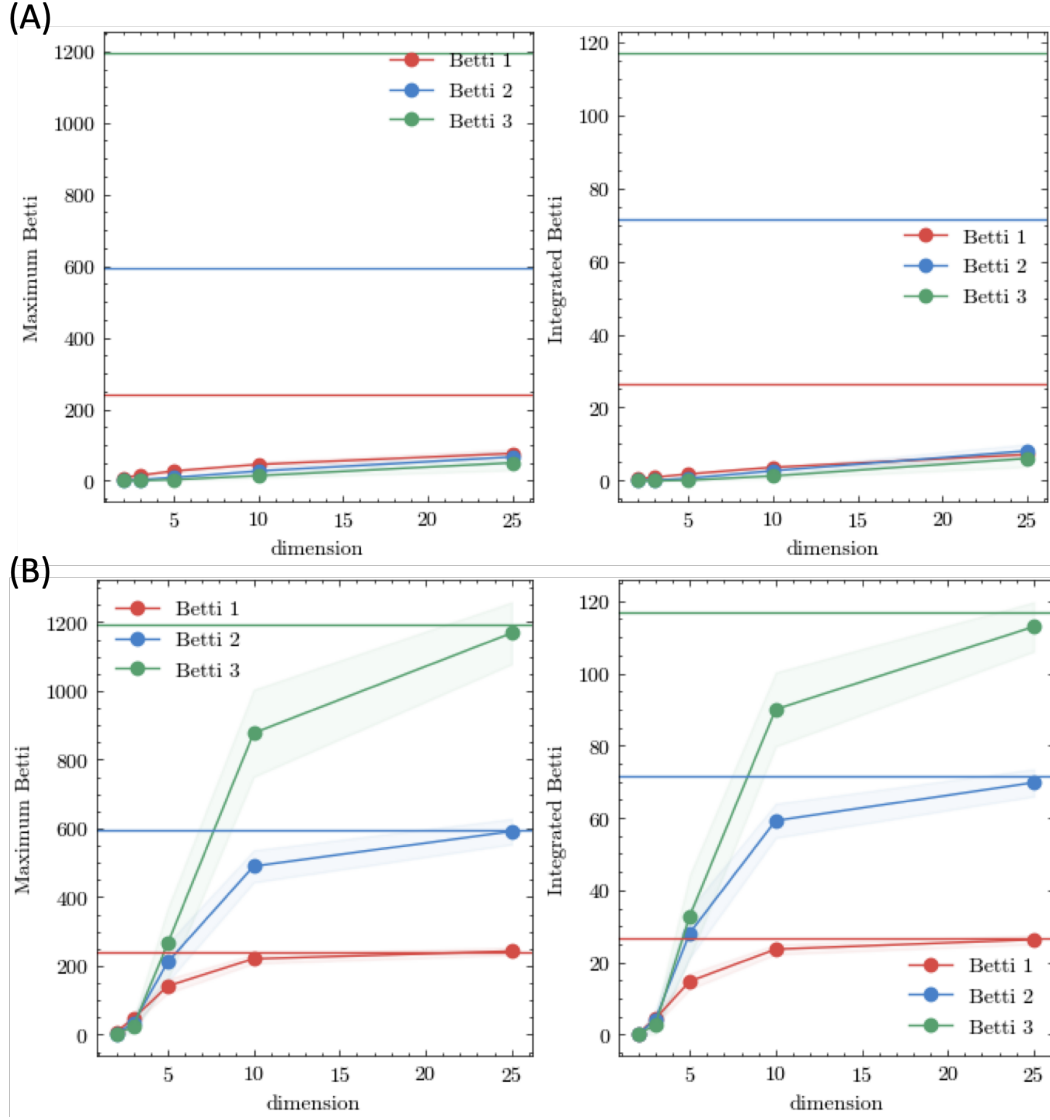


Figure S20: **Comparison of Random Matrices and Distances Matrices Sampled from Hyperbolic and Euclidean Spaces Across Various Dimensions:** Variations in Betti values β_m (left column) and maximum Betti curve value $\max(\beta_m(\rho))$ (right column) for $m = 1, 2, 3$ are presented across different dimensions with $N = 100$ for (A) Euclidean space and (B) hyperbolic space, using the parameters specified in Section A.2. Horizontal lines indicate benchmark outcomes derived from random matrices. Each experiment is repeated for 500 trials.

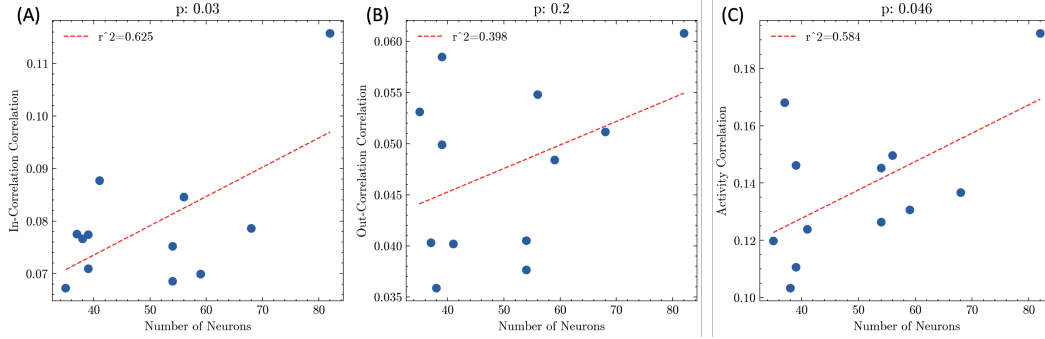


Figure S21: Having More Functionally Coregistered Neurons in a Single Scan Enhances Reconstruction Accuracy: This result remains consistent when using (A) C^{in} , (B): C^{out} , and (C): C^{act} . This example considers only proofread neurons with the weighted connectome, but the result is applicable to various settings.

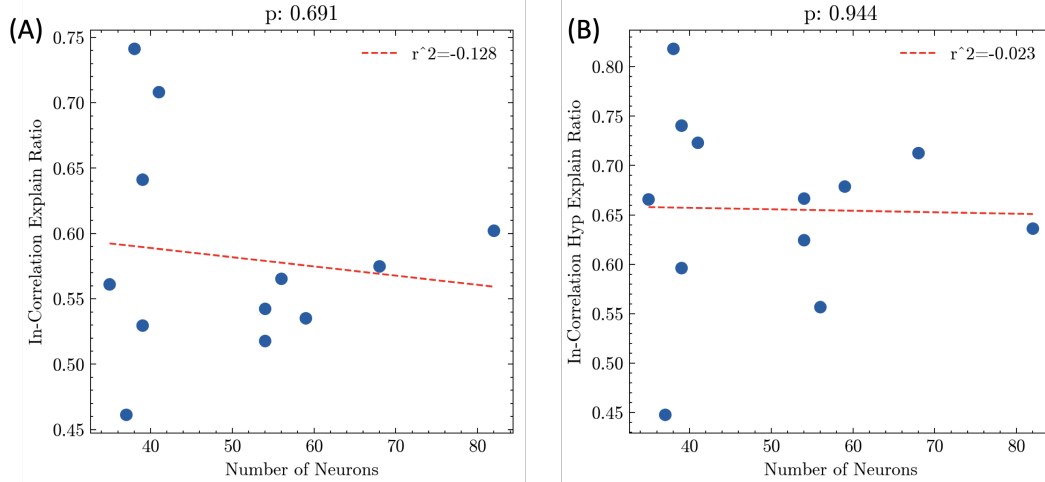


Figure S22: Having More Functionally Coregistered Neurons in a Single Scan Does Not Improve the Explanation Ratio, Regardless of Full Connectome Information or Embedded Representation: Compared to Figure S21, although the absolute reconstruction accuracy a significantly improves with increasing N , the explanation ratio (a/a_{act}) remains uncorrelated with N for both the full connectome C^{in} (left) and its hyperbolic embeddings (right). The observation holds consistently for output correlation (not shown). Data points correspond to the green and red parts in Figure 3.

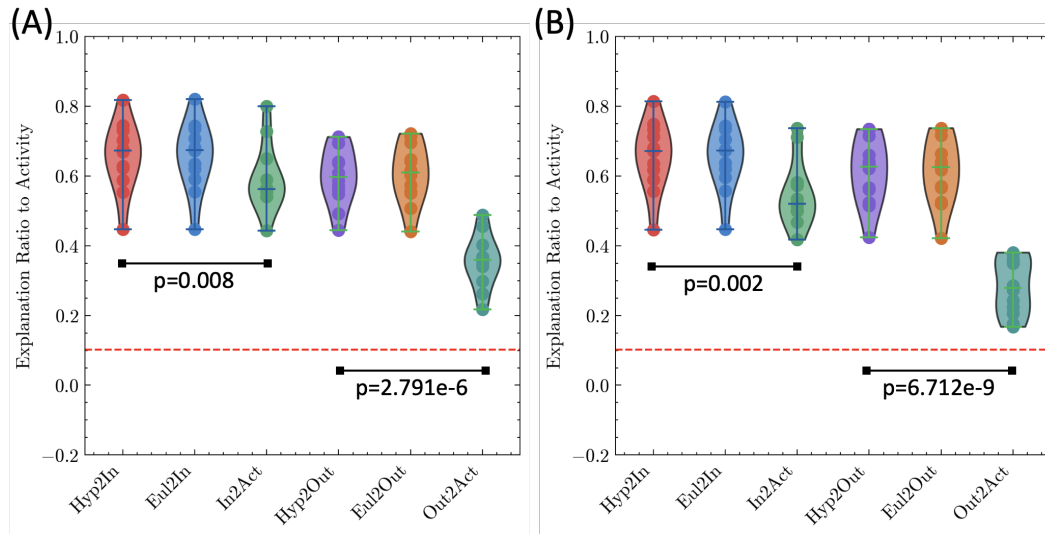


Figure S23: Calculate C^{in} and C^{out} (and embeddings) using (A) binary connectome (A) or (B) postsynaptic density (PSD) information.

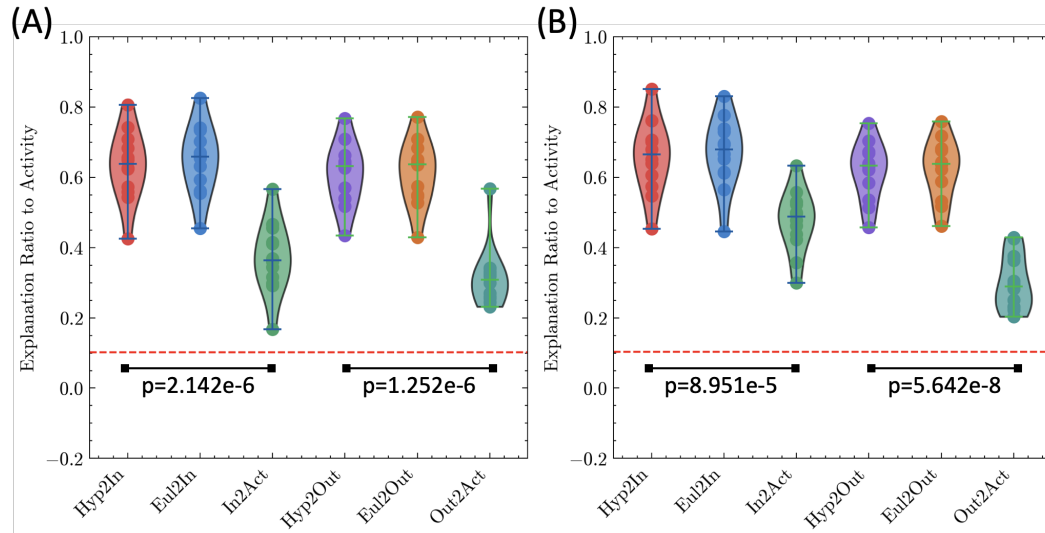


Figure S24: Calculate C^{in} and C^{out} (and embeddings) using weighted connectome to/from only (A) excitatory neurons and (B) inhibitory neurons.




Figure S25: Calculate C^{in} and C^{out} (and embeddings) using structural information from all excitatory and inhibitory neurons, including not fully proofread ones.

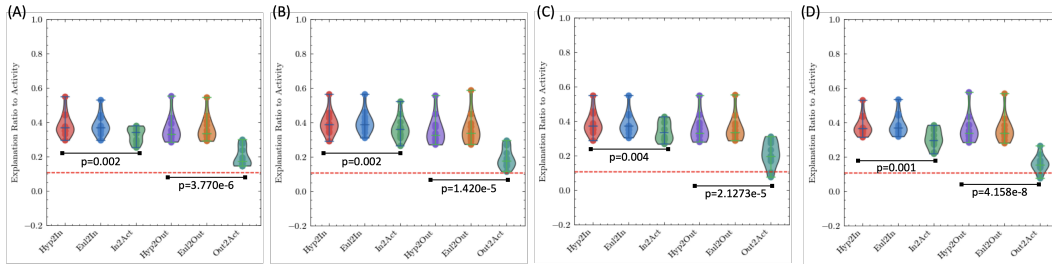


Figure S26: Calculate C^{act} respectively for each visual stimulus period based on (A) synaptic count with proofread neurons, (B) synaptic count with all E/I neurons, (C) binary connectome with proofread neurons, and (D) PSD with proofread neurons to calculate structural similarity.

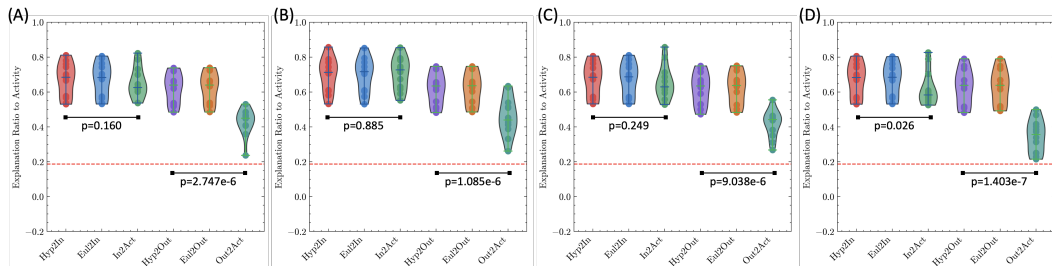


Figure S27: Reconstruction of Raw Fluorescence Traces based on (A) synaptic count with proofread neurons, (B) synaptic count with all E/I neurons, (C) binary connectome with proofread neurons, and (D) PSD sizes with proofread neurons to calculate structural similarity: The explanation ratio based on the median pairwise neuronal activity increases to 18.6% (Figure 3).

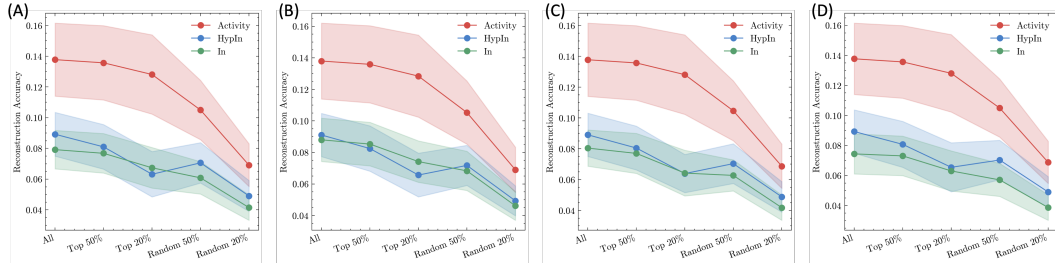


Figure S28: Reconstruction using Partially or Randomly Selected Neuronal Activity Traces (averaged over 100 trials), based on (A) synaptic count with proofread neurons, (B) synaptic count with all E/I neurons, (C) binary connectome with proofread neurons, and (D) PSD sizes with proofread neurons to compute structural similarity. Means and standard deviations are calculated across 12 independent scans.

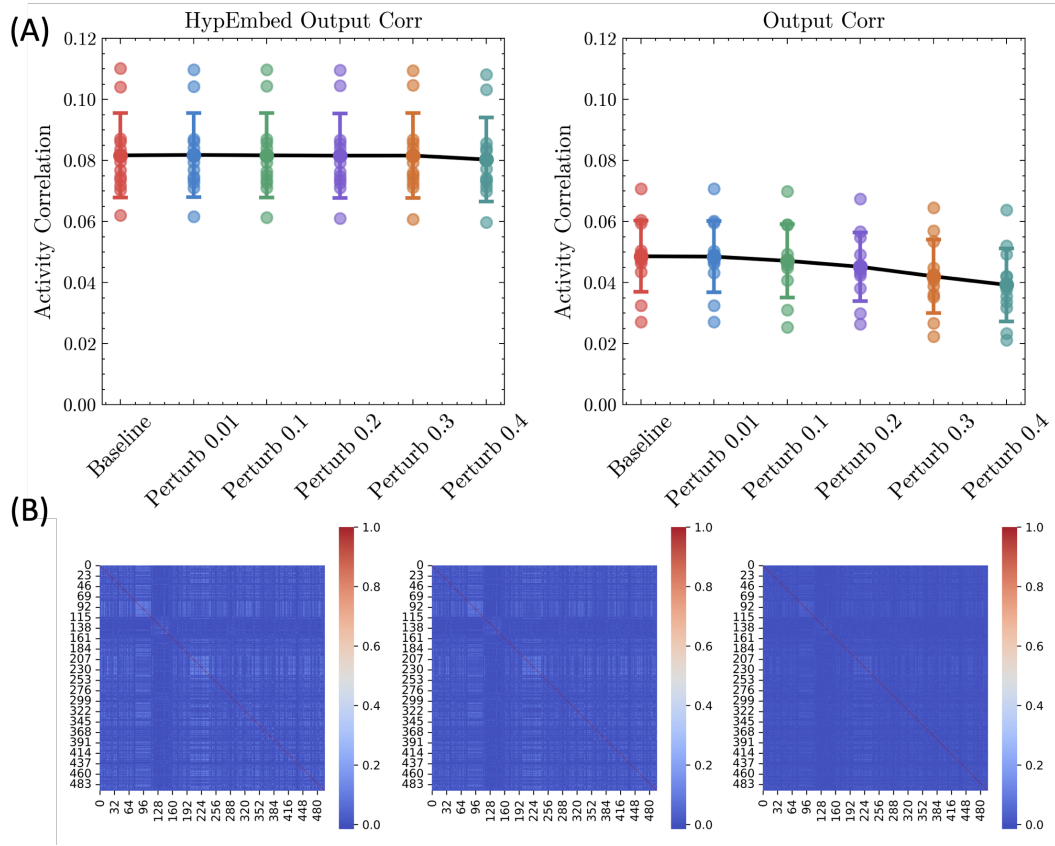


Figure S29: Comparison of Reconstruction Accuracy Using Hyperbolic Embedding of Output Similarity versus Output Similarity: (A): Reconstruction accuracy of the hyperbolic embedding of C^{out} (left) and C^{out} itself (right) under the baseline experiment (Figure 3) and varying levels of perturbation. (B) Illustration of C^{out} among targeted neurons ($N = 497$) for the baseline experiment (left), 10% perturbation (middle), and 40% perturbation (right). In (A), each scattered point represents the average result across 10 random perturbations for an individual scan. The mean and standard deviation across all scans are also shown. In (B), higher perturbation diminishes the grouping structure of neurons, resulting in a more homogeneous outcome.

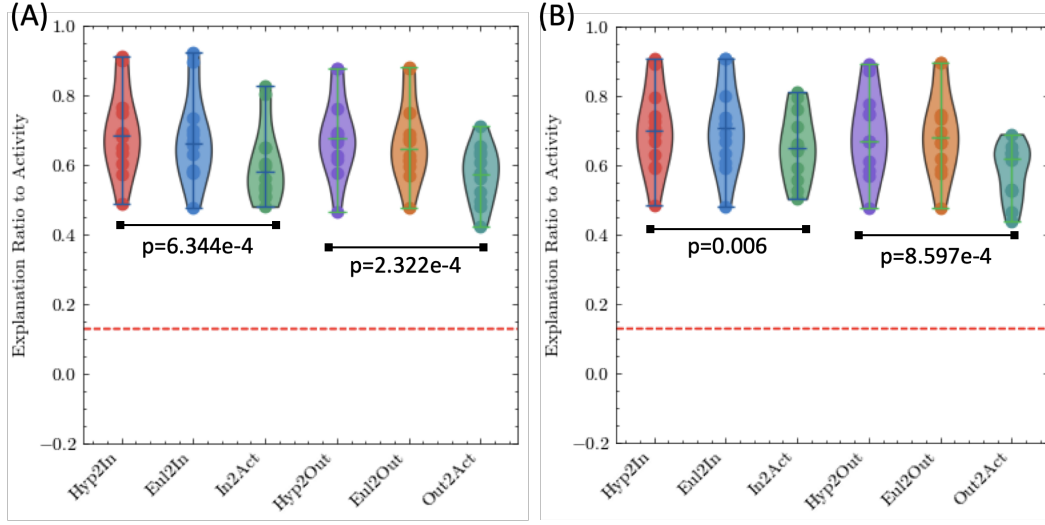


Figure S30: **Reconstruction using activity from proofread neurons based on (A) synaptic count with proofread neurons and (B) synaptic count with all E/I neurons to calculate structural similarity:** The graph illustrates explanation ratio (a/a_{act}). As discussed in Section 2.3 and Figure S21, overall accuracies decrease (except the output reliability) due to the stricter criterion and less neurons included per scan. For (A), C^{in} achieves 6.78% and C^{out} achieves 6.50%, while the corresponding counterpart values are 7.17% and 4.65%, respectively. (B) has 7.28% and 6.65%, and the counterpart has 8.08% and 5.01%.




Figure S31: **Dimensionality of Activity Space:** In the main text, we exclusively considered neurons that are functionally coregistered (i.e. neurons for which connectome information, though potentially incomplete, is available). We quantify the dimensionality of the activity space formed by these N_i neurons using the normalized participation ratio (NPR) [27]. Additionally, we evaluate the NPR for randomly selected sets of N_i neurons in each scan, averaged over 100 instances. The size of the scattered points is scaled with N_i .

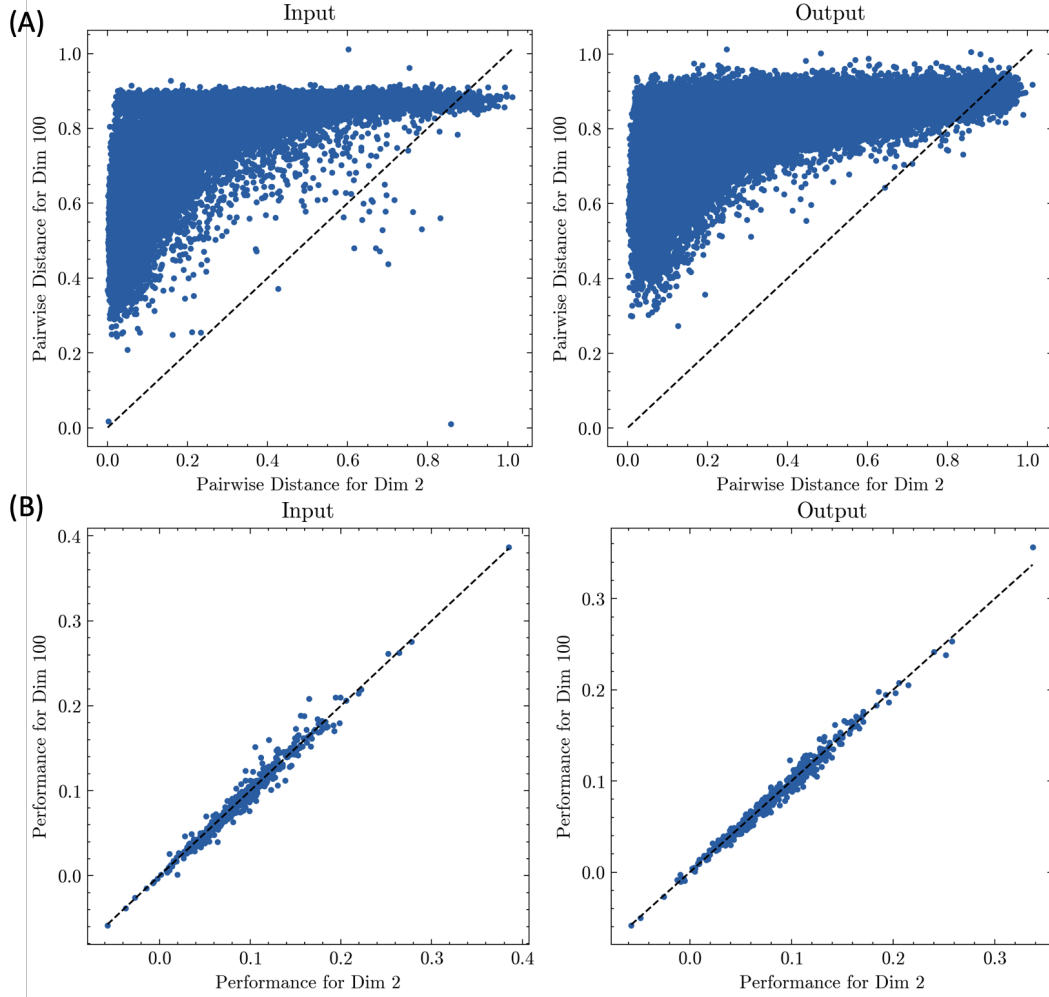


Figure S32: Comparison of (A) pairwise distances between Euclidean embedded points and (B) neuron-wise reconstruction accuracy, concatenated across all scans, for different dimensionality: We compare the results using optimized Euclidean embeddings in \mathbb{R}^2 and \mathbb{R}^{100} . Despite significant variation in distance preservation quality, the neuron-wise reconstruction accuracy remains consistent. In (B), we averaged the performance (Pearson correlation) for neurons included in multiple scans. Compared to Section A.1, \mathbb{R}^{100} would still be considered "low-dimensional" given the dataset size. Although it achieves near zero stress (not shown), the rank preservation (of pairwise distance remains imperfect ($R^2 = 0.73$ for the input and $R^2 = 0.92$ for the output), highlighting its distinction from the raw connectome data. Due to the computational challenges and instability caused by the high degree of freedom in "high"-dimensional embeddings, we limit the test to \mathbb{R}^{100} .

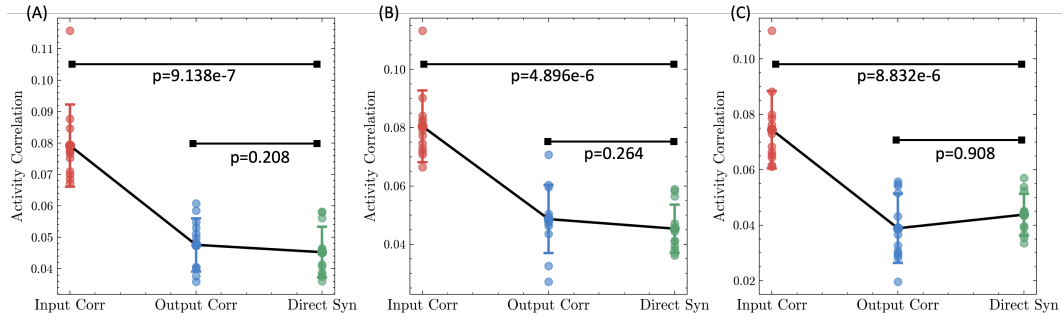


Figure S33: Comparison Between Using Normalized Direct Connectivity as Coefficients in the Linear Model and Using Input and Output Similarity: We compare the reconstruction performance in Figure 3 and Figure S23, using input and output correlation respectively versus direct connectivity measures. The direct connectivity is represented through (A) synaptic count, (B) binary connectivity, and (C) PSD sizes. Means and standard deviations are reported.