Hebb's Vision: The Structural Underpinnings of Hebbian Assemblies

J Wagner-Carena[1,2,3], S Kate[*1,5], T Riordan[*1], R Abbasi-Asl[1,4], J Aman[1], A Amster[1], AL Bodor[1], D Brittain[1], JA Buchanan[1], MA Buice[1], DJ Bumbarger[1], F Collman[1], NM da Costa[1], DJ Denman[1], SEJ de Vries[1], E Joyce[1], D Kapner[1], CW King[1], JD Larkin[1], J Lecoq[1], G Mahalingam[1], D Millman[1], J Mölter[6], C Morrison[1], RC Reid[1], CM Schneider-Mizell[1], S Daniel[1], S Suckow[1], KT Takasaki[1], M Takeno[1], R Torres[1], D Vumbaco[1], J Waters[1], DG Wyrick[1], W Yin[1], J Zhuang[1], S Mihalas[†1], and S Berteau[†1]

[1] *Allen Institute, Seattle, WA, USA*
[2] *Department of Molecular and Cellular Biology, Harvard University, Cambridge, MA, USA*
[3] *Neuroscience Graduate Program, University of California, San Francisco, San Francisco, CA, USA*
[4] *Department of Neurology, University of California, San Francisco, San Francisco, CA, USA*
[5] *Department of Biomedical Engineering, Georgia Institute of Technology, Atlanta, GA, USA*
[6] *Department of Mathematics, School of Computation, Information and Technology, Technical University of Munich, Germany*

1 Abstract

In 1949, Donald Hebb proposed that groups of neurons that activate stereotypically form the organizational building blocks of perception, cognition, and behavior. He theorized that repeated activations induce the structural changes needed to group neurons in these assemblies. Despite Hebb's enduring influence, testing his predictions at relevant scales has been technically challenging. Here, we test the theory using a novel, large-scale dataset featuring in vivo calcium fluorescence imaging of neural activity with postmortem electron microscopy (EM) for detailed reconstruction of neurons from the same volume of mouse visual neocortex. A coregistration process matches EM-reconstructed neurons to their recorded fluorescence traces. From these traces, we extract neural assemblies from higher-order correlations in neural activity. We find multiple assemblies, many with overlapping neurons, and some neurons that do not participate in any assembly. We then show that these assemblies exhibit properties consistent with Hebb's theory, including more reliable responses to repeated natural movie inputs than size-matched random ensembles and superior decoding of visual stimuli. Using coregistration to probe structural correlates, we find that neurons that participate in assemblies are significantly more integrated into the structural network than those that do not. Contrary to Hebb's original prediction, we do not observe a marked increase in the strength of monosynaptic excitatory connections between cells participating in the same assembly. However, we find significantly stronger indirect feed-forward inhibitory connections targeting cells in other assemblies. Intuitively, the delineation of assemblies can be realized either by internal excitation or external inhibition. Our findings support the latter mechanism. These results show the utility of assemblies in perception and provide a structural underpinning. They lay the foundation for future studies looking at the utility of assemblies in cognition and behavior, as well as the mechanisms for the formation and maintenance of such assemblies.

[*] These individuals contributed equally to this work as second authors
[†] These individuals contributed equally to this work as senior authors

2 Introduction

Since Hebb's 1949 monograph, The Organization of Behavior[1], cell assemblies have retained a persistent place in the imagination of the neuroscience community, both as a prospective unit of functional organization and as a compellingly likely consequence of simple rules of synaptic plasticity, such as the "fire-together wire-together" synapses dubbed 'Hebbian Synapses' by Yves Frégnac in 1986 [2]. Throughout this work, we follow Hebb in using the term assembly to refer to overlapping sets of neurons that activate in a reproducible pattern with high fidelity; in much of the contemporary literature, such groups are also called cell ensembles or simply ensembles. The theory continues to stimulate research on the activity-related aspects of assemblies, and the organization of these processing modules is thought to operate through stable recurrent activity [3, 4]. Incorporating a behavioral approach, experimental evidence implies that the response of discrete cell populations akin to assemblies may also have a causal link with motor functions [5]. In particular, there has been confirmation of sequential activation within neural assemblies of the primary visual cortex (V1) [6] as well as ongoing 'replay' of coactivity in the absence of stimuli [7]. There has also been some limited evidence of cell assemblies' stability over long periods of time [8].

However, research examining the structural aspects of Hebbian cell assemblies has primarily focused on the potentiation of synapses between excitatory neurons. This emphasis is perhaps not surprising, given that this potentiation has a place of prominence in Hebb's original formulation of the theory [1]. In 1963 [9], Hebb himself acknowledged that his excitation-only formulation of assembly theory was a concession to the state of research on inhibitory synapses at that point in time; synaptic inhibition of neural activity was not confirmed until 1952 [10]. Over the following decades, the association between Hebbian assemblies and excitatory potentiation has itself been reinforced by the computational plausibility of excitatory plasticity as a mechanism of assembly formation [11, 12, 13], and the discovery of Hebbian synapses. However, such changes were never a necessary precondition for assembly formation. Mechanistically, there exists a broad range of plausible solutions to assembly formation, ranging from excitatory modulation against a backdrop of relatively stable inhibitory strengths to the opposite, in which formation relies solely on inhibitory modulation contrasting with stable excitatory connections [14].

The approach of this work, using large-scale recordings of individual neurons to identify and analyze assembly function, enters a developing tradition in the literature of population dynamics [15, 16, 17, 18, 19, 20]. However, due to limitations in structural analysis, few studies can relate the structure of assemblies to their function, as reviewed by [14]. Electrophysiological datasets can allow highly reliable inference of connectivity in the case of multi-patch recordings, but with very small numbers of cells in any given study. Extracellular recordings can overcome this limitation, but produce biased connectivity estimates, and therefore cannot be considered a gold standard.

To relate the correlated activity of a large number of neurons to their connections, we used a novel large-scale multi-modal dataset: the Allen Brain Observatory V1 Deep Dive (V1DD) [21] (Fig. 1A). V1DD offers a combination of Ca^{2+} fluorescent recordings and detailed EM reconstruction of neurons and synapses, including post-synaptic density volumes. Taking advantage of advances in optical imaging techniques [22], V1DD provides multiple scans of high-quality simultaneous two-photon calcium imaging (2PCI) recordings of thousands of excitatory neurons within the mouse primary visual cortex (Fig. 1B). In addition to its functional recordings, V1DD also contains electron microscopy (EM) of the same tissue volume, which has uncovered the fine-scale anatomy of the cubic millimeter volume of the brain (Fig. 1C). EM has been employed extensively in large-scale datasets to map drosophila [23, 24], worm [25], mouse [26, 27], and

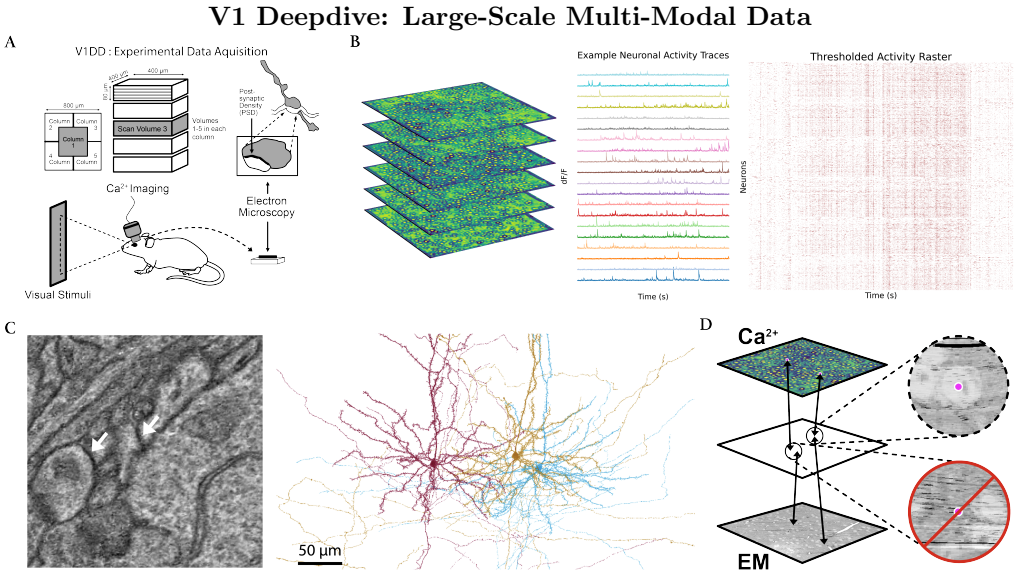

Figure 1: **(A)**. Schematic of experimental data acquisition for the V1 Deepdive dataset (V1DD). V1DD consists of dense thousands of excitatory neurons in an $800 \times 800 \times 800\,\mu m^3$ section of mouse V1, recorded during awake behaving imaging sessions, with no behavioral task. Our work focuses on the center column, particularly the third scan volume (pictured here in grey). Postmortem, the same tissue volume was fixed and imaged via transmission electron microscopy, allowing for reconstruction of synaptic connectivity, including post-synaptic density (PSD) volumes. **(B)** Each scan volume for in vivo imaging consisted of six stacked scan planes. Dense calcium activity allowed for the extraction of individual neuronal traces, with 20 example traces shown in addition to a raster plot of thresholded normalized activity for all 2708 neurons in Scan Volume 3 of Column 1. **(C)** Example of a microscopy view of connected neurons, with white arrows pointing to PSD. Reconstructed pyramidal cells corresponding to the left microscopy view are shown on the right. **(D)** Schematic showing the framework for coregistration of cells between the calcium recordings and the electron microscopy. Identified ROIs were mapped to an interstitial space (see Methods 5.1.4), where the correspondences were manually inspected.

even the human brain [28]. The combination of these two imaging modalities has been applied in very few datasets [29, 30, 27], particularly at this scale.

These advances offer the opportunity to examine the structural correlates of Hebbian cell assemblies at an unprecedented scale. We extract assemblies from a Ca^{2+} fluorescence imaging scan and examine the reliability of their activation and their functional significance in the encoding of visual stimuli. We then analyze the connectivity between neurons based on cells that were coregistered between fluorescence and structural EM scans (Fig. 1D). Deriving hypotheses directly from postulates advanced by Hebb, we test his predictions about synaptic sizes and connectivity within and across assemblies, producing results that suggest a significant role for inhibition in their formation and activation, different from what is traditionally assumed.

3 Results

3.1 Neuronal Organization of Hebbian Assemblies

We analyzed a scan of the optical imaging dataset consisting of 2708 excitatory neurons recorded in parallel. The Similarity Graph Clustering (SGC) algorithm (Fig. 2A) generated 15 assemblies, which we ordered by size with 'A 1' representing the largest assembly ($n = 1016$) and 'A 15' the smallest ($n = 23$). A subset ($n = 748$, 27.6 percent of 2708 total) of neurons was assigned to no assemblies.

Their spatial distributions are shown in (Fig. 2C,D). Although spatial organization is not clearly visible, statistical analysis using the Kolmogorov-Smirnov (KS) [31] test revealed significant differences in the spatial organization of assemblies compared to that of the entire neural space (Fig. 2E). In particular, there was marked organization in nearly every assembly along the x-y plane, indicative of the retinotopic organization observed in studies of the visual cortex [32].

Most, but not all, assemblies shared members with other assemblies (Fig. 2B). For example, neurons assigned to 'A 1' were also assigned to ten other assemblies, highlighting Hebb's proposal that precisely timed phase sequences allow a set of shared member neurons to participate in multiple assemblies.

3.2 Correlation and Sparsity

By construction, assemblies are expected to respond to distinct stimulus features, with a coactivity correlation sufficiently low to prevent their merging into a single assembly. Pearson correlation coefficients between assembly coactivity traces were significantly lower than the coefficients between coactivity traces of random ensembles of the same size distribution (see Methods) (Fig. 3A). As expected from traces derived from the average of population raster activity, both populations were significantly higher in their correlation than individual cell activity raster correlations, regardless of the subset of pyramidal cells being considered (assembly, non-assembly, or all individual cells).

To further characterize the functional properties of assemblies, we computed the Gini coefficient [35] for each assembly's activity trace. The Gini coefficient, a statistical measure of the 'inequality' of signal activity throughout the optical recording, revealed that assemblies exhibited highly sparse activity patterns (Fig. 3B), exhibiting significantly higher Gini coefficients than random ensembles (p-value : $6.23e - 5$). The coefficient for each assembly ranged from 0.55 to 0.83, with particular assemblies with extreme sparsity, such as 'A 13' (0.79) and 'A 15' (0.83), exemplifying a high degree of functional selectivity. Interestingly, this sparsity metric was not solely dependent on assembly size, as intermediate-sized assemblies (e.g., 'A 4' through 'A 12') all exhibited similar Gini coefficients of around 0.70.

3.3 Assemblies Reliability in Stimuli Response

A section of the visual stimuli presented to the mouse consisted of natural movies. Across the hour scan time, twelve unique 15-second movie clips were repeatedly shown eight times. To evaluate the functional reliability of neuronal assemblies, we analyzed their responses to these stimuli. Reliability was quantified using an Oracle score, a leave-one-out correlation metric (see Methods 5.3) that measures the consistency of an activity trace across repeated presentations of the same stimulus. Oracle scores have been used as a measure of the reliability of neuronal response [36]. To provide a baseline with which to compare the functionality of these assemblies to tune to visual stimuli, the oracle scores of assemblies' coactivity traces were compared to the oracle scores of

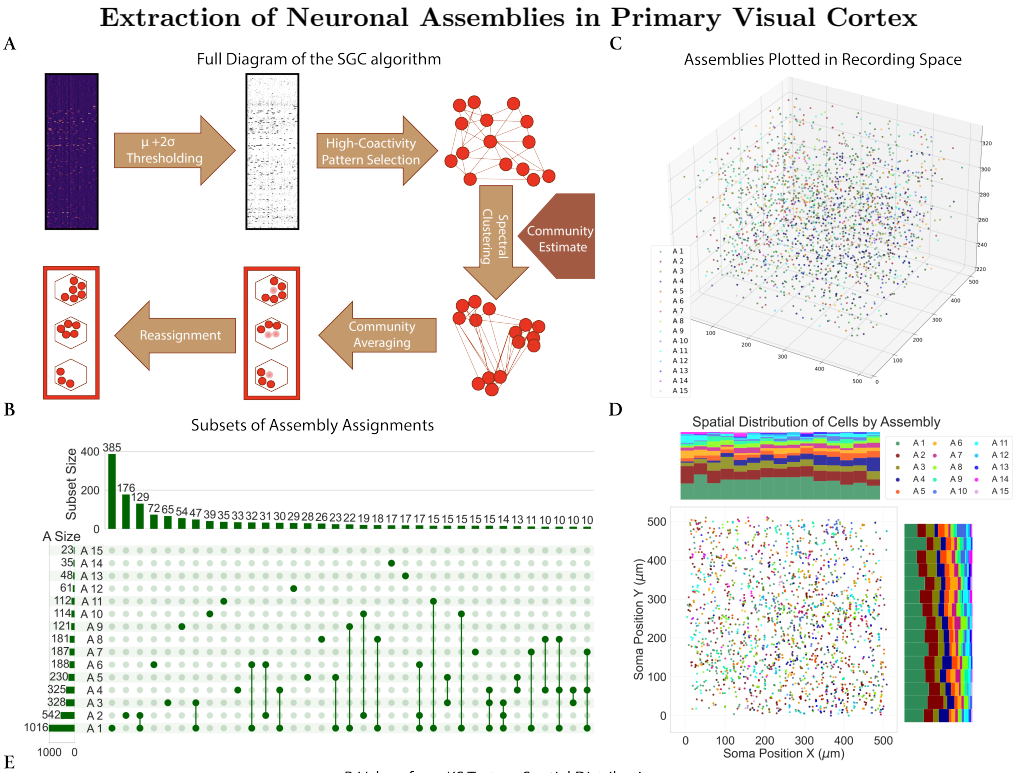

P-Values from KS Test on Spatial Distributions

	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11	A12	A13	A14	A15
Z-Axis	7.9e-2	9.9e-1	2.2e-3	1.6e-4	2.2e-2	2.8e-2	2.3e-2	4.0e-1	8.3e-1	4.0e-1	1.8e-1	9.6e-1	4.7e-2	6.5e-1	1.3e-2
XY-Plane	2.2e-7	2.7e-4	4.4e-2	1e-12	5.8e-4	1.3e-1	1.4e-2	1.6e-5	8.4e-3	7.9e-6	3.3e-2	1.3e-1	1.8e-3	2.4e-6	1.1e-2

Legend: Blue Text = Significant P-Values, Brown Text = Insignificant P-Values

Figure 2: **(A)**. SGC, an extraction algorithm for assemblies uniquely designed for calcium imaging data, groups frames of the calcium fluorescence input to determine when neurons in assemblies are coactivated. Figure adapted from Mölter et al. [33] **(B)**. An UpSet [34] visualization of the subsets formed between assembly assignments. The histogram on the left represents the size of each individual assembly. The top histogram represents the size of the subsets between assemblies. Only subsets of ten neurons or greater were visualized. **(C)**. Spatial positions of fifteen extracted assemblies in the three-dimensional recording field. There are 1960 neurons visualized, including neurons assigned to multiple assemblies (plotted only once) but not including the 748 neurons that were assigned to no assemblies. **(D)**. Spatial distribution of assembly cells projected onto the x-y plane. Histograms for each axis are normalized to provide a per-bin proportional stack of the assembly distributions. **(E)**. A table presenting the KS test results on each assembly's spatial distributions. Values colored in blue represent significant results (p-value < 0.05), while brown signifies insignificant results.

all neuronal traces. In addition, we provide a population-level comparison with random ensembles.

Assemblies exhibited significantly higher Oracle scores compared to the average reliability of individual neurons (p-values < 0.0001), indicating that the assemblies, as populations, respond more consistently to visual stimuli (Fig. 3C). To ensure this result was not merely due to the inclusion of highly reliable neurons within assemblies, we separately calculated Oracle scores for neurons within assemblies and those assigned to no assemblies. We observed no significant difference between the cellular sets (p-values > 0.25). This result suggests that the reliability of these assemblies is derived from their collective activity rather than from the reliability of individual members.

Since population coactivity is expected to be more reliable in the general case than individual neurons, we compared our assemblies to random ensembles (defined as in Methods 5.2). Assembly coactivity traces demonstrated higher Oracle scores than coactivity traces computed in the same way as the assemblies for size-matched random ensembles. Complementary results are seen in the reliability of responses to orientation gratings (see Supp. Fig. 6). Population activity during the presentation of these gratings was used to characterize reliability as a function of orientation [37].

Further analysis revealed consistent patterns of high assembly coactivity during specific visual frames of these natural movies. We defined 'trigger frames' as moments when assembly activity exceeded a baseline threshold (see Methods 5.2) and found that these frames were highly consistent across repeated stimulus presentations. Example mean trigger frame pixel values are presented in (Fig. 3F), along with those of the corresponding size-matched random ensemble, and the squared difference. Visualizations of these frames suggest that assembly activity responds to complex features in the natural movies. These collective results provide evidence of assemblies' ability to serve as functional populations with reliable and specific responses to visual stimuli.

3.4 Decoding Responses from Acute Visual Stimuli

We also assessed the ability of assemblies to decode visual stimuli by implementing a classification framework comparing assemblies to random ensembles (Fig. 3D, E). We employed a Multi-Layer Perceptron Classifier (MLPClassifier) to evaluate how well each grouping could decode the identities of the twelve natural movie clips. These classifiers have been shown to be effective in academic and clinical settings, with high levels of accuracy and flexibility in available hyperparameters [38, 39].

The results of our classifier revealed that assemblies significantly outperformed random ensembles in accuracy. Heatmaps of classification accuracy for natural movie clip identities demonstrated that assemblies provided more reliable decoding across repeated trials. A Mann-Whitney U-test confirmed this finding, with assembly accuracy significantly exceeding random ensembles for both overall performance (u-stat: 7546.0 p-value: $6.06e-5$) and diagonal elements (one-sided, u-stat: 131.5 p-value: 0.00032), indicating enhanced stimulus-specific decoding.

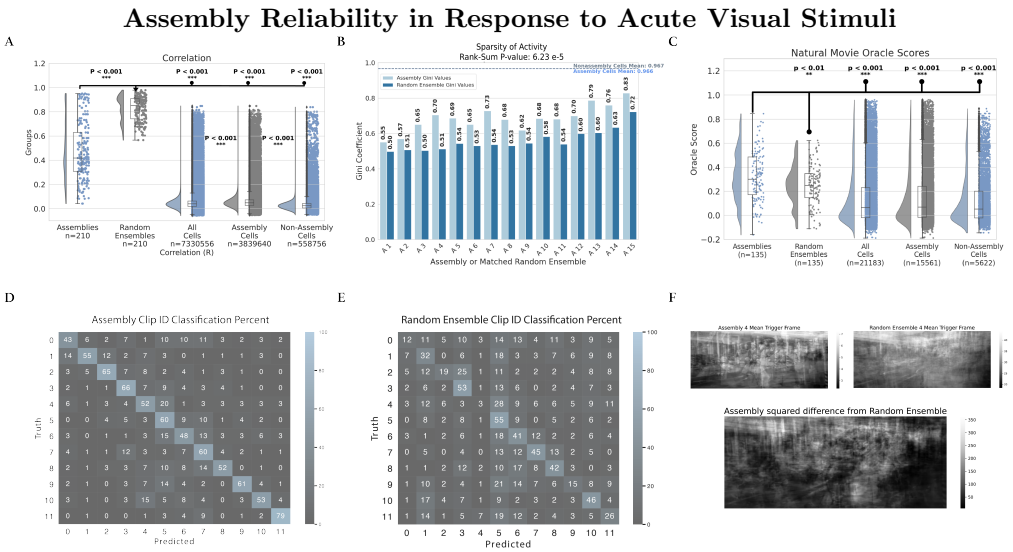

Figure 3: **(A)** Raincloud plot of pairwise correlations of coactivation between assemblies, size-matched random ensembles, and sets of individual cells. Coactivation for an individual cell is equivalent to a binary thresholded activity raster. **(B)**. Grouped bar plot of sparsity (measured by the Gini Coefficient) of coactivity over time in cell assemblies and the null-grouping of size-matched random ensembles. The random ensembles' coefficients are significantly smaller than the set of assembly coefficients (Wilcoxon Rank-Sum p-value: $6.234e-5$). The average sparsity of individual assembly cells and non-assembly cells is also plotted as nearly equal horizontal dashed lines. **(C)**. Raincloud plot illustrating the reliability of activity from assemblies and general neuronal populations in response to natural movies. Oracle scores of each assembly and random ensemble coactivity trace were plotted, as well as the scores of sets of individual cells. These oracle scores are computed for the concatenation of natural movie clips and their responses, rather than individual clips, in order to reduce the likelihood of sparse responses causing an artificially high reliability score. **(D, E)**. Heatmap illustrating decoding accuracy of natural movie clips with Assemblies and Random Ensembles. Heatmap values indicate the accuracy of clip decoding by the percentage of presentation. Clip IDs, indicating a unique natural movie clip, are balanced such that each clip has an equal frequency of presentation. Values in the assembly heatmap are significantly greater than the random ensemble heatmap (Mann-Whitney u-stat: 7546.0 p-value: $6.06e-5$, one-sided on diagonal elements u-stat: 131.5 p-value: 0.00032). **(F)**. Example plots of the mean 'trigger frame' of assemblies and random ensemble during natural movies. Frames were generated by averaging the frames associated with peak coactivity. The natural movie frame was visually better reconstructed by the assembly activity than that of the random ensemble, as signified by the plotted squared difference.

3.5 Structural Organization of Assemblies

Coregistration between recorded activity and EM data provides us with a unique opportunity to explore the structural underpinnings of assemblies in the visual cortex. By mapping neural structure and connectivity at a micrometer resolution, we investigated the anatomical communication and organization of neurons with at least one shared assembly membership (shared assembly cells) compared to those with disjoint membership.

The strength of connections between the two groups was measured by performing a

Wilcoxon Rank Sum test to compare the connection weights (defined as the sum of PSD volumes for all synapses between two cells). A Chi-squared test of Independence was performed to determine whether the pairwise frequency of connections differed between cells with shared assembly membership and those with disjoint assembly membership. To investigate higher-order structural patterns, we performed similar tests on sets of inbound and outbound disynaptic chains, as well as subdividing based on whether the intermediate cell in each chain was inhibitory or excitatory. 5.11.2

Our initial analysis of first-order connectivity surprisingly revealed no significant differences in the probability of direct monosynaptic connections between shared and disjoint sets or the strengths of those connections (Fig. 4D,E). This finding is incongruous with predictions that assemblies are defined by densely interconnected excitatory neurons [1] and suggested that the defining structural characteristics of assemblies might lie beyond simple pairwise connectivity metrics.

To investigate higher-order structural patterns of these assemblies, we conducted a motif analysis. In a neural network, when a significant number of subgraphs containing a small number of interconnected cells repeat a particular pattern (e.g., Cell Type A $\rightarrow$ Cell Type B $\rightarrow$ Cell Type A), they are classified as a motif [40, 41], with each subgraph counted as a motif instance. The frequency with which these motifs persist has been predictive of correlation in similar neural networks [42].

Our analysis concentrated on second-order chain motifs, or structures consisting of three neurons connected by two synaptic links (schematic shown in Fig. 4F,H). Second-order neural motifs are divisible into various types based on the arrangement of their synaptic connections. For this study, we prioritized chain motifs, as they allowed us to utilize aspects of the EM dataset that are not yet coregistered to activity recordings for the central elements in the chain, so long as the first and last cells in the chain are coregistered.

Notably, while our excitatory chain analysis revealed no significant differences between shared and disjoint assembly memberships (Fig. 4F,G), an analysis of inhibitory motifs convey a different story: disjoint assembly memberships exhibited significantly stronger feed-forward inhibitory connections than shared memberships (Fig. 4H,I). Log-scaled plotting of the chain weights (Fig. 4I, inset) revealed a complete lack of overlap between SEM-based confidence intervals, suggesting that this is a result with both reliability and a non-trivial effect size. Notably, this result was insignificant when restricting our analysis to only the inhibitory synapse or only the excitatory synapse within the feed-forward inhibitory chain, suggesting that both connections may play a role.

The significantly stronger feed-forward inhibitory connections suggest a mechanism of mutual inhibition that may regulate the interaction between distinct assemblies, preventing excessive coactivation and ensuring the discriminability of assembly responses to inbound information. Indeed, we find a negative Pearson r correlation coefficient between the disjoint feed-forward inhibitory chain weights and the correlation of the two disjoint assemblies' coactivity traces (r statistic : -0.19, p-value : 0.0048; see Methods 5.11.4).

All together, our structural results reveal that assemblies are not anatomically defined solely by a fundamental increase in local connectivity, but through their higher-order patterns of organization. Stronger feed-forward inhibitory connections in disjoint memberships suggest a possible mechanism for maintaining functional segregation between assemblies. These structural insights provide intuition on the mechanistic basis that drives the functional properties of assemblies, reinforcing their contribution as modular units of sensory processing.

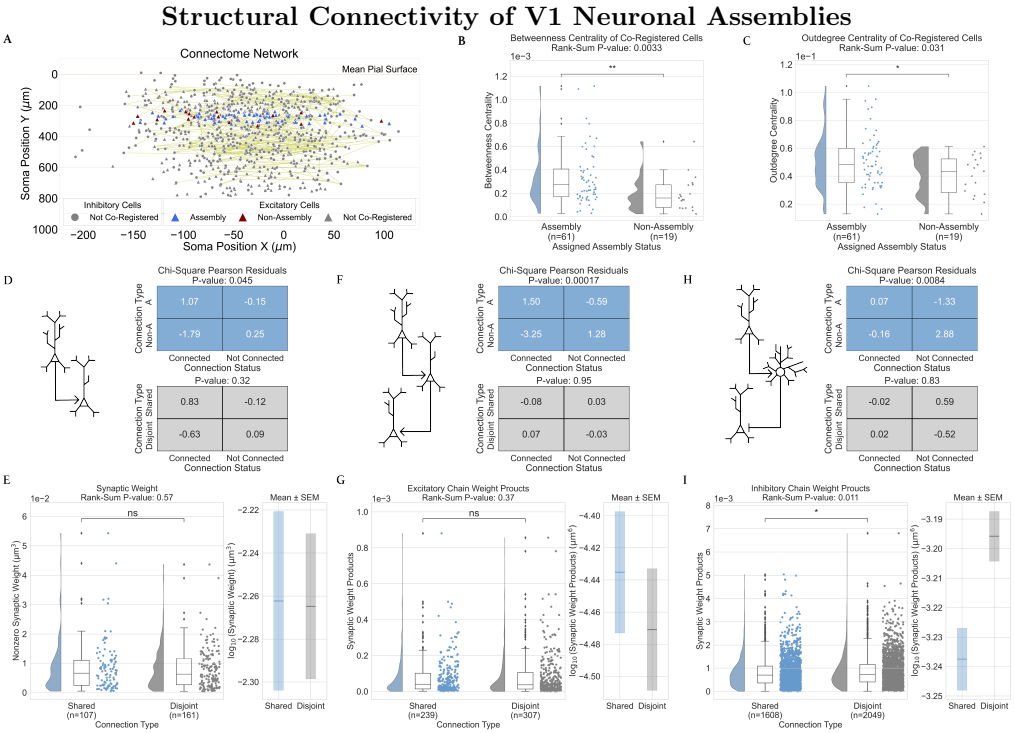

Figure 4: **(A)**. Visualization of the network being analyzed, showing the soma position of cells in the connectome, colored by cell-type and assembly assignment. All coregistered reconstructed neurons were found in layer 2/3 or layer 4 of V1. **(B)** A raincloud plot of betweenness centrality, demonstrating a higher centrality for assembly neurons than those not in assemblies(Wilcoxon Rank-Sum: p-value < 0.01). **(C)** A raincloud plot of outdegree centrality, a mathematical proxy for probability of connection, demonstrating a higher centrality for assembly neurons (Wilcoxon Rank-Sum: p-value < 0.05). **(D,F,H)**. Chi-squared analysis of the likelihood of monosynaptic (D), disynaptic excitatory (F), and disynaptic inhibitory (H) connections (schematic on the left). This comparison was made between assembly neurons and non-assembly neurons (top-right) as well as between neurons that share an assembly membership and neurons that participate in disjoint assemblies (bottom-right). **(E, G, I)** Raincloud plots showing the combined synaptic PSD volume per extant monosynaptic (E), disynaptic excitatory (G), and disynaptic inhibitory (I) connection, each divided between origin and terminus cell pairs which share assemblies and those which participate in disjoint assemblies. Log-scaled plotting of the chain weights with SEM-based confidence intervals is included as an inset plot to the right of each panel.

3.6 Non-Assembly Cells

The same monosynaptic and disynaptic analysis run on pairs of cells with shared assembly memberships and cells with disjoint memberships was also run on pairs of assembly cells (cells that both participate in at least one assembly) and pairs of non-assembly cells (cells that had no assembly memberships for either cell).

We show assembly cells exhibiting a pattern of connectivity distinct from non-assembly cells, a requirement of Hebb's theory. We assessed the higher-order integration of assembly neurons into the broader structural network using centrality metrics. In particular, we found significantly lower betweenness centrality (Fig. 4B), a measure

of a node's importance in mediating communication within a network, and outdegree centrality (Fig. 4C), an analog for probability of outbound connectivity, in neurons outside assemblies compared to neurons within assemblies. Furthermore, we found that the pairwise probability of monosynaptic connections, disynaptic excitatory chains, and disynaptic inhibitory chains was significantly greater for cells sharing assembly membership than for non-assembly cells (p-values : 0.045, 0.0002, and 0.0084, respectively). No significant differences in PSD Volumes between assembly and non-assembly cells were observed. These findings confirm that cells in the connectome participating in at least one assembly are more interconnected than those not participating in any.

Overall, these results indicate reduced participation in the structural framework of the primary visual cortex for non-assembly cells and imply an organizational role of assembly neurons as hubs for information flow. In contrast, non-assembly neurons may play a more peripheral, secondary role in the network, such as noise filtering.

4 Discussion

4.1 Inhibition as the Delineating Basis of Hebbian Assemblies

This work provides the first demonstration of the structural underpinnings of Hebbian assemblies and validates Hebb's assembly theory in a surprising way; while there is strong agreement on functional predictions, our structural findings run counter to Hebb's original proposal. We do not observe a marked increase in the strength or probability of excitatory connections between cells that participate in the same assembly (Fig. 4D-G). Instead, we find significantly stronger feed-forward inhibitory connections targeting cells in other assemblies (Fig. 4I). Hebb and others [43, 9, 3] have hypothesized that targeted inhibition could play a crucial role in the formation and delineation of assemblies. Our findings offer the first clear empirical evidence for this mechanism.

Despite this novel result, our findings remain a validation of Hebb's structural postulates and are consistent with what would be required for assemblies to be delineated by inhibition. Assembly cells were substantially more integrated into the connectome, displaying significantly higher betweenness and outdegree centrality (Fig. 4B,C). This pattern demonstrates that cells participating in our activity-defined assemblies possess markedly different patterns of connectivity from cells that do not, establishing the necessary structure-function relationship postulated by Hebb.

Intuitively, the delineation of assemblies can be realized on a spectrum between internal (shared-assembly) excitation and external (disjoint-assembly) inhibition. Shared-assembly excitation would rely on broad inhibitory connectivity across the entire population of assembly neurons, and disjoint-assembly inhibition would similarly rely on broad excitatory connectivity. We find a significantly greater probability of excitatory connections between cells belonging to any assembly (Fig. 4D,F), however we did not observe significantly stronger or more frequent excitatory connections between neurons with shared assembly membership (Fig. 4D-G). Combined with our finding of significantly greater feed-forward inhibitory weights between neurons with disjoint assembly membership (Fig. 4I), the broad excitation strongly suggests that the extracted assemblies are delineated predominantly via inhibition.

4.2 Functional Consistency of Assemblies

Functionally, ensembles of cells classified as assembly cells behave in ways consistent with Hebbian assemblies. The assemblies demonstrated markedly higher reliability in their responses to naturalistic visual stimuli compared to individual cellular responses of all three sub-populations of cells and size-matched random ensembles, reflected by

significantly elevated Oracle scores (Fig. 3C). This consistent coactivity of the assemblies has been hypothesized to offer resilience to representational drift, allowing them to serve as a substrate for long-lasting representation [8]. Representational drift reflects a functional characterization of fluctuations at the cellular level, such as synaptic turnover [44, 45], even under stable stimulus conditions [46]. This phenomenon has been observed in areas other than V1, such as the piriform cortex [47]. By maintaining consistent and coherent patterns of coactivity, assemblies may offer a general cortical mechanism of stable perceptual representations despite such turnover.

Beyond their reliability in responses, assemblies were also superior in our decoding of natural movie presentations (Fig. 3D), underscoring their efficacy in extracting higher-order visual information. In addition, the sparsity of assembly coactivity was substantially greater than that of random ensembles (Fig. 3B), which, when combined with a lower average correlation 3A), is consistent with assemblies employing a cost-effective encoding strategy [48]. Such a strategy would enhance the capacity for distinct representation of sensory inputs, translating more flexible individual responses into reliable population-level encoding. Finally, cells assigned to assemblies exhibit significantly higher pairwise correlations than non-assembly cells, corroborating the experimental work of Harris and Carandini [49] defining 'choristers' and 'soloists'. However, we found no difference in their average response reliability to natural movies or in their overall signal sparsity. This result implies that a population perspective is required for the encoding of reliable perceptually relevant stimulus features.

4.3 Connective Tissue of Structure and Activity

Some examined ensemble features bridge structure and activity, a crucial aspect of this study, which has not previously been achievable at scale. First, the extensive overlap of neurons across multiple assemblies (Fig. 2D) is consistent with both Hebb's proposal of sub-assemblies [9] and Yuste's review of definitions put forward for ensembles [50], with individual neurons contributing to multiple functional modules. Second, our observation of greater chain weights in feed-forward inhibitory chains offers a mechanistic explanation of how assemblies can retain distinct responses to stimuli without becoming so correlated in their activity that they merge into a single functional ensemble. The plausibility of this explanation is supported by the significant observed negative correlation between inhibitory chain connection strength between disjoint assemblies and the r coefficient of their paired assembly coactivity traces, showing that the greater the inhibitory weight of the connections observed, the less correlated the activation of the two assemblies is 3.5.

4.4 Further Directions

Looking ahead, assemblies exhibit notable variability in size (Fig. 2D). SGC minimizes overestimation of the number of assemblies or neuron assignments, suggesting that these size differences reflect intrinsic properties rather than methodological artifacts. There is therefore an opportunity for future studies to delineate the distinct roles potentially served by larger assemblies, such as 'A 1', compared to their smaller counterparts. In addition, while novel in its scale and multifaceted nature, we see opportunities to reduce some of the current limitations within this dataset. Continued reconstruction, proofreading, and coregistration will allow analyses of more neurons and subsequent new lines of inquiry. For instance, the SGC algorithm flags frequent patterns of activity, potentially analogous to Hebbian phase sequences [6, 7]. We currently do not have sufficiently many coregistered and fully reconstructed cells to allow examination of connections that bridge one pattern to the next. Exploring how these patterns inter-

act could reveal mechanisms of large-scale neural coordination. Other limitations will be more difficult to overcome without the availability of next-generation multi-modal datasets. Most importantly, the temporal resolution of the scans leaves us with no ability to examine activity on some of the most relevant time scales for synaptic plasticity.

These results suggest several potential implications for the field at large. Hebb proposed assemblies as a universal building block, simultaneously addressing the functional and structural sides of perception, cognition, and behavior. A basic unit that bridges structure and function allows one to derive structural predictions from functional characterizations and vice versa. Beyond validating Hebb's theory, the evidence presented here for an underlying inhibitory mechanism provides further support for the analysis of inhibition in cognitive and sensory disorders studied to date primarily through the lens of excitation [51, 52]. In the future, the integration of cell-type-specific genetic information with functional assembly data could also provide deeper insights into the molecular foundations of assembly formation and maintenance in health and disease [53]. Hebb postulated that assemblies form the atoms of cognition, and it has not escaped our notice that the cross-inhibitory mechanism we here demonstrate might be a universal feature of brain-wide assembly organization.

5 Methods

5.1 V1 Deepdive Dataset

5.1.1 Stimuli

Visual stimuli were presented using the same monitor configuration as in de Vries et al. [54]. Imaging sessions were one hour long and offered a wide variety of visual stimuli. Assembly extraction was performed on the fluorescence data from the full session. The remainder of our analysis utilizes only the natural movie clips and the full-field drifting gratings, details of which are provided below. The other stimuli presented during the session included natural images, windowed drifting gratings, and locally sparse noise. The details of these other classes of stimuli can be found in [21]. The natural movies stimulus consisted of 3,600 frames (with 30 Hz frame rate), presented 8 times. The full-field drifting gratings stimulus consisted of a drifting sinusoidal grating at a 1 Hz temporal frequency and 80 percent contrast, presented at 12 different orientations (multiples of 30°) and at 2 spatial frequencies (0.04 and 0.08 cycles per degree). Each condition was presented eight times, in randomized order, with one second of mean luminance grey between presentations.

5.1.2 2/3 Photon Microscopy and Activity Data Processing

The full $800\times800\times800\,\mu m^3$ volume was divided into five columns, each imaged via either 2-photon (2P) or 3-photon (3P) microscopy, depending on depth. The volumes were scanned over the course of several sessions. Each session consisted of the full set of visual stimuli (see below), presented in the same order and with the same timing. Our analysis is concerned only with the central column, centered within the $800\times800\times800\,\mu m^3$ volume, where the largest set of reconstructed, coregistered, and proofread neurons is currently available. Figure 1A shows the arrangement of scan volumes (5 volumes spanning 75 µm to 620 µm depth, each $400\times400\times80\,\mu m^3$) and imaging planes (6 planes within each scan volume, separated by 16 µm at thirty-seven frames per second so that each plane is imaged at 6 Hz.) We selected volume three of the central column for our primary analyses and validated our work with volume four. Both were imaged with 2P microscopy. Details on the other columns and the 3P imaging can be found in [21].

The fluorescence data was preprocessed using the standard LIMS pipeline as used for the Allen Institute's Visual Coding 2P dataset [54], including motion correction, segmentation, demixing, neuropil subtraction, ROI filtering, and df/f calculation. Identified regions of interest were run through a classifier trained on manual labeling data meant to reduce the false classification of artifact ROIs as neuronal somas, and only those with a high confidence score (at least 0.5) were included in our analysis.

5.1.3 Electron Microscopy and Reconstruction

The mouse was transcardially perfused with a fixative mixture of paraformaldehyde and glutaraldehyde. All procedures were carried out in accordance with the Institutional Animal Care and Use Committee at the Allen Institute for Brain Science. The large-volume staining protocol was adapted from [55]. After dissection, the neurophysiological recording site was identified by mapping the brain surface vasculature. A thick (1200 µm) slice was cut with a vibratome and post-fixed in perfusate solution for 12 to 48 h. The tissue was then infused with heavy metals, dehydrated, and embedded in EMS Hard Plus resin. After curing, the samples were epoxy cured to a stub. They were then sliced and placed onto continuous tape by a ATUMtome Automated Tape Collecting Ultramicrotome.

The continuous tape was fed into an automated high-throughput transmission electron microscopy pipeline[56]. Transmission electron microscopy is particularly well suited for automated imaging and preserves very good x-y resolution at the expense of some resolution on the vertical z-axis, and so specialized methodology was deployed during reconstruction[57, 58]. Serial section alignment was performed through a contract with Zetta A.I, followed by stitching [59, 56], segmentation, and automated reconstruction. Proofreading of a subset of cells was performed under contract by Ariadne.ai.

Cell-type predictions were made for single-nucleus objects within 175 microns of the centerline and with a nucleus volume greater than 218, based on dendritic skeleton features adapted from [60]. Segmentation and annotation were stored in a CAVE database for access via CAVEClient [61].

5.1.4 Coregistration

Manual coregistration was performed using the Fiji plugin BigWarp [62, 63]. A structural scan of the vasculature was aligned with the two-photon imaging planes (max intensity projection). Next, a downsampled EM image was aligned to this composite using a thin-plate-spline transform based on manual landmarks. After initial alignment, the transform was used to predict additional correspondences between two-photon ROI centroids and segmented EM cells. Four hundred verified correspondences between a fluorescence imaging ROI with a high classifier score and a corresponding morphologically typed EM-reconstructed cell passed manual inspection by two independent reviewers and were included in this study. 315 of these correspond to cells whose fluorescence was recorded in the scan we analyzed (volume three of the central column). Of these 315, 80 had axon reconstructions that were verified accurate by trained experts ('proofread') to their maximal extension within the scan volume. Coregistration and reconstruction are ongoing at the time of this writing, but the extant data already allow for a relation of the physiology of neural data to its exact anatomy at an unprecedented level.

5.2 Graph Clustering for Assembly Extraction

To extract assemblies, we use the Similarity-Graph-Clustering (SGC) algorithm that has been originally proposed for the detection of neural assemblies during spontaneous activity in the zebrafish optic tectum [64]. The SGC algorithm identifies neuronal assemblies using ideas from graph theory by transforming the problem into one of community detection on some graph in which assemblies correspond to distinct (graph) communities. Unlike traditional methods that rely on pairwise correlations between cells, SGC groups frames of fluorescence indicative of significant coactivity [33]. These moments of higher-order correlation are referred to as potential 'phase sequences', representing when populations of neurons act cohesively as a closed circuit during assembly activation [65].

Importantly, by design, SGC allows cells to participate in multiple assemblies. Assembly overlap has been integral to assembly studies in the past [66]. Hebb originally postulated that the shift between active assembly states could be what is colloquially referred to as a 'train of thought' [1]. Our results revealed a sizable degree of overlap between assemblies, which was visualized through an UpSet plot (Fig. 2D) [34], showing neurons frequently assigned to multiple assemblies (Fig. 2D).

Compared with several other prominent assembly extraction algorithms upon application to both simulated and biological calcium imaging datasets under different conditions, this algorithm was shown to perform best overall [33]. Although it did not yield a perfect reconstruction of the assemblies in the biological dataset, SGC was able to recover the assemblies with higher accuracy than all other algorithms. In part, this performance has been attributed to the computational effort that SGC places in estimating the number of assemblies before defining them.

We used a recent implementation of the SGC algorithm in Python [67]. For completeness, we briefly recall the main steps here: The algorithm commences by thresholding the calcium fluorescence signals (df/f) of the ROIs (neurons) by two standard deviations above the mean to minimize noise (Fig. 2B). Afterwards, activity patterns are selected where the coactivity level of neurons exceeds the significance threshold. This threshold that determines the set of "high-activity patterns" is based on a null model of coactivity obtained from shuffling the thresholded activity signals (significance value: 0.05, rounds: 1000). A k-nearest-neighbor graph is then constructed from the set of high-coactivity patterns based on the similarity between the patterns in the cosine distance. For that, the number of neighbors k is automatically chosen such that the resulting graph is connected. In the next step, the number of communities in this graph is estimated using a statistical inference procedure. For our study, we fixed the hyperparameters with five independent Monte Carlo rounds of $150,000$ steps each. With an estimate for the most likely number of communities, spectral clustering is applied. These clusters of high-coactivity patterns are the first prospective selection of coactivity patterns corresponding to the assemblies. However, the final step is a combination of averaging and reassignment to reject groups that may have been erroneously defined because of a high level of noise in the original signal. This minimizes the likelihood of overestimating the number of assemblies or the neurons that should be assigned to those assemblies. Neurons are assigned to assemblies based on their affinity, the probability that they were active in any of the assemblies' activity patterns (affinity: 0.4). To establish our hyperparameters, we first performed a grid search on an independent functional scan volume (Scan Volume 4, acquired under identical conditions). Candidate values of Monte Carlo rounds, activity thresholds, and affinity cut-offs were evaluated by how well the resulting assembly assignments could fit to a low-dimensional embedding of pairwise activity correlations (see Supp. Fig. 7). The parameters that maximized the AIC value were rounds = $150,000$, Z = 2 sd, and affinity = 0.4. These parameter values were then locked and applied to scan volume 3.

5.3 Random Ensembles

Random ensembles served as a null model and were defined as randomly selected sets of neurons drawn with equal probability and no replacement from the population of all recorded pyramidal cells within scan volume 3. Each set was size-matched to its corresponding assembly, with the same number of neurons.

5.4 Oracle Scores

Oracle scores are a measure of the reliability of the trace response of cell activity to repeated visual stimuli computed through a jackknife mean, or leave-one-out mean [68, 69], of correlations between the activity trace across the repeated visual stimuli. In this work, we treat the entire stereotyped stimulus sequence as a single stimulus, concatenating natural movie clips and their responses, rather than calculating the score for individual clips. This should reduce the risk of sparse responses leading to an artificially high reliability score.

Let $S_i(t)$ denote the activity trace of a neuron or an assembly during the i-th presentation of a stimulus at time t. Let $\bar{S}_{-i}(t)$ represent the mean activity trace of all activity traces excluding the i-th presentation, out of n total presentations of the same stimulus. From this, we can calculate the Oracle score O as

$$O = \frac{1}{n} \sum_{i=1}^{n} \frac{\mathrm{Cov}(S_i(t), \bar{S}_{-i}(t))}{\sqrt{\mathrm{Var}(S_i(t)) \cdot \mathrm{Var}(\bar{S}_{-i}(t))}}.$$

In this formulation, the numerator will calculate the covariance between the activity trace at each presentation to the mean activity of all other repeats, while the denominator scales the magnitude of this covariance by the product of the standard deviations of $S_i(t)$ and $\bar{S}_{-i}(t)$.

5.5 Trigger Frames

To assess the visual stimuli associated with high activity in neuronal assemblies, we computed trigger frames by identifying peak activity times and extracting the corresponding images from a natural movie presentation. From a coactivity trace for assembly k over time $S_k(t)$, we can detect peaks P_k using the `scipy` signal package to define local maxima. We then define a mean trigger frame $\mu_k(x, y)$ for a pixel (x, y) as

$$\frac{1}{|P_k|} \sum_{t \in P_k} I_t(x, y)$$

where $I_t(x, y)$ is the natural movie frame at time t.

This allowed us to visualize the average triggering frame for each assembly. The same process was then repeated for the coactivity traces of the size-matched random ensembles, and the squared difference between the average frame for each assembly and the average frame for its corresponding random ensemble. These computations allowed us to determine which frame of visual stimuli and consistent features in those stimuli were most strongly associated with high assembly coactivity in particular, not merely broad neuronal activation.

5.6 Decoder

To evaluate the ability of assemblies to decode visual stimuli, we implemented a Multi-Layer Perceptron Classifier (MLPClassifier) from the `scikit-learn` library. This classifier was used to differentiate between 15 natural movie clips based on assembly coactivity time traces. Random ensembles of neurons with the same size distribution as the assemblies were used as a null model for comparison.

The MLP is a classical feed-forward neural network composed of an input layer, one or more hidden layers, and an output layer. Each node in a layer is fully connected to every node in the subsequent layer through weighted connections. The final output of the function is determined by a non-linear activation function applied to the weighted sum of its inputs plus a bias term.

The MLPClassifier from `scikit-learn` is an implementation of a Multi-Layer Perceptron (MLP), a type of feed-forward neural network. It operates by mapping input features $\mathbf{x}$ to outputs $\hat{\mathbf{y}}$ through a series of hidden layers. Each layer consists of neurons that perform a weighted sum of their inputs, followed by the application of a non-linear activation function. Defining $\mathbf{W}^{(k)}$ as the weight matrix for the k-th layer, $\mathbf{h}^{(k)}$ as the input, and $\mathbf{b}^{(k)}$ as the bias vector for the k-th layer, we can define with activation function σ

$$\mathbf{h}^{(k)} = \sigma\left(\mathbf{W}^{(k)}\mathbf{h}^{(k-1)} + \mathbf{b}^{(k)}\right).$$

For classification, the final layer uses the softmax activation function to output probabilities for each class. With L denoting the number of layers, we can define these output probabilities as

$$\hat{\mathbf{y}} = \text{softmax}(\mathbf{h}^{(L)}) = \frac{\exp(\mathbf{h}^{(L)})}{\sum_j \exp(\mathbf{h}_j^{(L)})}.$$

The MLPClassifier is trained using backpropagation, optimizing the weights and biases via stochastic gradient descent or adaptive solvers such as Adam. Regularization can be applied through an ℓ_2-penalty term, controlled by a hyperparameter.

The assembly coactivity time traces were paired with corresponding natural movie clip IDs. The data was split into training and test sets using an 80-20 split, and features were scaled to normalize the input. A cross-validated grid search was used to optimize the hyperparameters.

5.7 Gini coefficient

The Gini coefficient [35], a statistical measure that exemplifies the state of inequality within a population. While often applied in economics to evaluate income inequality, this metric has been applied as a valid approximation for signal sparsity [70, 71]. This rendition of the application has been shown to serve as a relatively simple and robust measure [72]. For our study, the coefficient is employed to quantify assembly signal heterogeneity.

The Gini coefficient, G is often calculated with respect to the Lorenz curve, which plots the cumulative distribution of a set (e.g., assembly coactivity trace) against its rank in ascending order. For a given assembly A with coactivity trace $S = [s_1, s_2, ..., s_T]$ where s_t is the proportion of active neurons at time point t, the coefficient for that assembly is then calculated as

$$G_A = \frac{\sum_{i=1}^{T-1} \sum_{j=i+1}^{T} |s_i - s_j|}{T^2 \cdot \bar{s}}.$$

This computation is performed independently for each assembly, providing a metric of signal inequality. A value of 0 for G_A implies all values are identical, while a value of 1 indicates perfect inequality *i.e. in a time series a single time point contains all the activity.* A high value of G_A indicates that coactivity is dominated by a small number of time points, reflecting the temporal sparsity.

5.8 Correlation

We examine coactivity correlations for each pair of assemblies, each pair of random ensembles, and each pair of neurons within each of our three sets of neurons (see Fig. 3A-C). To do so, we computed the Pearson's correlation coefficient r between the two coactivity traces (where the coactivity trace of a single neuron is mathematically equivalent to its thresholded raster activity, per [33]).

5.9 Motif Extraction

Motifs were extracted with the DotMotif Python package which detects subgraphs within a graph based on the principle of subgraph monomorphism. For every graph $H = (N_1, E_1)$ given by the user, the Dotmotif algorithm detects subgraphs G' within the graph $G = (N_2, E_2)$ such that there exists a mapping $f : N_1 \to N_2'$ where $N_2' \subseteq N_2$ and for every edge $(a, b) \in E_1$, the corresponding edge $(f(a), f(b)) \in E_2'$. The matched subgraphs may also contain additional edges within the graph G. This algorithm was used to detect disynaptic chains within the connectome.

To further refine our analysis, chain motifs were classified according to the type of intermediary neuron, distinguishing between excitatory and inhibitory connections, with the latter providing insight into feed-forward inhibition mechanisms.

5.10 Motivating Postulates

Our statistical analyses involve tests of the following postulates. First, excitatory connections between cells that share at least one assembly ('shared' connections) will be stronger than connections between cells which do not participate in any of the same assemblies ('disjoint' connections), due to Hebbian plasticity [1]. This was examined both with regard to the post-synaptic density volume of monosynaptic connections between known coregistered cells within the dataset, and separately in the form of the product of connection PSD volumes in disynaptic excitatory chains which originated and terminated with shared or disjoint cells, allowing us to evaluate indirect excitatory connections where the middle cell had not yet been coregistered.

Second, that excitatory connections will be more frequent within assemblies than between assemblies, due to a combination of Hebbian plasticity and pruning of synapses. A number of computational studies have shown that long-term potentiation [73] and pruning [74] play important roles in effective Hebbian assembly formation. This was examined in the form of per-connection targeting statistics, as well as per-cell inbound and outbound probability of both monosynaptic and disynaptic shared and disjoint connections.

Third, as discussed in the introduction, Hebb himself suggested that the assemblies could have been formulated via modulation of inhibition. Additionally, a number of computational models [75] have relied on inhibition between assemblies to restrict simultaneous activation, enabling competition between assemblies. It was, therefore, taken as an established hypothesis that one would expect inhibition between assemblies to be greater than within a given assembly. As monosynaptic connections between excitatory cells cannot be used to evaluate inhibition, this postulate was examined only in disynaptic chains, where the middle cell was morphologically classified as an inhibitory interneuron. Our examination involved the product of connection PSD volumes in such disynaptic inhibitory chains bridging shared and disjoint assembly neurons, along with a per-cell evaluation of the inbound and outbound probabilities of disynaptic inhibitory chain connections.

Fourth, and finally, we acknowledge that Hebb discusses the reinforcement of sparse connectivity in his accounts of the emergence of cell assemblies [43], and thus this reinforcement might have a significant effect on a sparse subset of connections while producing a minimal difference in the central tendency of the overall set of connections. Dorkenwald et al. [76] demonstrated a bimodality of the log PSD volume of excitatory-excitatory synapses in the similar MICrONS mm^3-dataset [27] (see also Supp. Fig.8 for result in V1DD), suggesting that a subset of such connections is impacted differently by processes determining PSD volume. Combining these two, we decided to test the hypothesis that the larger of the two-component distributions found in [76] would be more likely than chance to involve connections between shared assembly neurons.

5.11 Statistical Methods

In this section, we detail the motivations and specifics of our analysis and methods.

We analyzed differences in connectivity metrics between shared-assembly and disjoint-assembly connection types, focusing on both the probability and strength of connections. Our analysis considered both direct monosynaptic and disynaptic connections between neurons ("by connection"), as well as the sets of inbound and outbound connections grouped by cell ("by cell"). For disynaptic chains, we grouped by whether the intermediate (middle) cell was excitatory or inhibitory.

To evaluate connectivity metrics, we first defined the connection types based on assembly membership. Let A be the set of all assemblies, with A_i denoting the subset of assemblies that include cell i. Formally,

$$A_i = \{a \in A \mid i \in a\}.$$

Using these subsets, we defined the following binary indicators to capture the assembly relationship between pre-cell j and post-cell i:

- $\text{Shared}_{ij} = 1$ if $A_j \cap A_i \neq \emptyset$
- $\text{Disjoint}_{ij} = 1$ if $A_j \cap A_i = \emptyset$
- $\text{Assembly}_{ij} = 1$ if $A_j \neq \emptyset$ and $A_i \neq \emptyset$
- $\text{Non-Assembly}_{ij} = 1$ if $A_j = \emptyset$ and $A_i = \emptyset$

5.11.1 Monosynaptic Metrics

We defined w_{ij} as the pairwise summed post-synaptic density (PSD) between pre-cell j and post-cell i, and b_{ij} as an indicator variable that takes the value 1 if at least one synapse exists between j and i and 0 otherwise. All metrics exclude autapases, as these were not reliably represented in the dataset ($j \neq i$).

The probability of monosynaptic outbound connection for a pre-cell j was calculated as the proportion of realized connections under a given connection type $\mathrm{C} \in \{\text{Shared}, \text{Disjoint}, \text{Assembly}, \text{Non-Assembly}\}$, normalized by the total number of potential post-cell partners for that connection type,

$$b_{\text{out}_{j,\mathrm{C}}} = \frac{\sum_{i|\mathrm{C}_{ij}=1} b_{ij}}{|\{i \mid \mathrm{C}_{ij} = 1, i \neq j\}|}.$$

Similarly, the probability of a monosynaptic inbound connection for a post-cell i was defined as

$$b_{\text{in}_{i,\mathrm{C}}} = \frac{\sum_{j|\mathrm{C}_{ij}=1} b_{ij}}{|\{j \mid \mathrm{C}_{ij} = 1, j \neq i\}|}.$$

For connection strength, we computed the average realized summed monosynaptic outbound PSD for a pre-cell j as the total PSD across all post-cells satisfying the connection type C, normalized by the number of realized ($b_{ij} = 1$) connection under C,

$$w_{\text{out}_{j,\mathrm{C}}} = \frac{\sum_{i|\mathrm{C}_{ij}=1} w_{ij}}{\sum_{i|\mathrm{C}_{ij}=1} b_{ij}}.$$

Similarly, the average realized summed monosynaptic inbound PSD for a post-cell i is given by

$$w_{\text{in}_{i,\mathrm{C}}} = \frac{\sum_{j|\mathrm{C}_{ij}=1} w_{ij}}{\sum_{j|\mathrm{C}_{ij}=1} b_{ij}}.$$

5.11.2 Disynaptic Metrics

In examining the inhibition in sets of cells that share assembly membership, and in sets that do not, we were primarily interested in describing inhibition driven by the excitatory activity of an assembly's member cells. This moved us from the realm of monosynaptic connection analysis into an analysis of chains. Many aspects of our definition remained unaltered. A remained the set of all assemblies, with A_j the subset of assemblies that included pre-cell j, and A_i the subset of assemblies that included post-cell i. The binary indicators indicating the assembly relationship between cells i and j remained unaltered.

But rather than simply using the monosynaptic weight between neurons j and i where $j \neq i$, we defined w_{ikj} as the product of the pairwise summed post-synaptic densities w in a three-cell chain motif with j as the first cell, an interneuron k as the second cell, and i as the third cell. Thus,

$$w_{ikj} = w_{ik} w_{kj}.$$

And similar to our monosynaptic analysis, b_{ikj} was an indicator variable that took the value 1 if at least one disynaptic chain existed between j, k, and i and 0 otherwise.

Building on this definition, we outlined the method of normalization and metrics for disynaptic chain analysis, which accounted for both the intermediate and final (or first) cells in the chain. In this analysis, each pre-cell and post-cell was a coregistered excitatory neuron with an extended axon, consistent with the monosynaptic analysis. The intermediary cell in disynaptic chains, however, did not need to be coregistered or possess an extended axon.

Let n_e represent the number of excitatory cells and n_i represent the number of inhibitory cells from the set of all cells in the all-all connectome. Define $|k|$ be the number of potential middle partners. For inhibitory chains, $|k| = n_e$ as the middle cell is inhibitory. For excitatory chains, the middle cell cannot be the first or final cell, so $|k| = n_e - 2$.

The probability of disynaptic outbound connection for a pre-cell j was calculated as the proportion of realized disynaptic connections under a given connection type $\mathrm{C} \in \{\mathrm{Shared}, \mathrm{Disjoint}\}$, normalized by the total number of potential chains satisfying the connection type,

$$b_{\mathrm{out}_{j,\mathrm{C}}} = \frac{\sum_k \sum_{i|\mathrm{C}_{ij}=1} b_{ikj}}{|k|\,|\{i \mid \mathrm{C}_{ij} = 1, i \neq j\}|}.$$

Similarly, the probability of disynaptic inbound connection for a post-cell i was defined as

$$b_{\mathrm{in}_{i,\mathrm{C}}} = \frac{\sum_k \sum_{j|\mathrm{C}_{ij}=1} b_{ikj}}{|k|\,|\{j \mid \mathrm{C}_{ij} = 1, j \neq i\}|}.$$

For nonzero strength of connection, we computed the average realized summed disynaptic outbound PSD for a pre-cell j as the total PSD across all chains satisfying the connection type C, normalized by the number of realized ($b_{ikj} = 1$) chains under C,

$$w_{\mathrm{out}_{j,\mathrm{C}}} = \frac{\sum_k \sum_{i|\mathrm{C}_{ij}=1} w_{ikj}}{\sum_k \sum_{i|\mathrm{C}_{ij}=1} b_{ikj}}.$$

Similarly, the average realized summed disynaptic inbound PSD for a post-cell i was given by

$$w_{\mathrm{in}_{i,\mathrm{C}}} = \frac{\sum_k \sum_{j|\mathrm{C}_{ij}=1} w_{ikj}}{\sum_k \sum_{j|\mathrm{C}_{ij}=1} b_{ikj}}.$$

To facilitate statistical testing, we defined collections of metrics based on connection type C for both monosynaptic and disynaptic analyses. For disynaptic sets, the indices ij are replaced with ijk appropriately.

5.11.3 Set Definitions

The following sets were defined to evaluate connectivity metrics:

- The set of nonzero pairwise connection strengths under a given connection type C,
$$\{w_{ij} \mid C_{ij} = 1\}.$$
- The set of outbound probabilities of connection for each pre-cell under a given connection type C,
$$\{b_{\mathrm{out}_{j,\mathrm{C}}}\}.$$
Similar collections were defined for inbound probability of connection, and for inbound and outbound nonzero average connection strengths by replacing the metric accordingly.
- For inbound or outbound metrics, paired sets were constructed by including only post- or pre-cells with at least one fulfilled connection under both the Shared and Disjoint connection types. Let
$$\mathcal{I}_{j,C} = \{i \mid b_{ij} = 1 \text{ and } C_{ij} = 1\}$$
be the set of post-cells connected to pre-cell j under connection type C. The set of pre-cells included in the paired comparison was then defined as:
$$\mathcal{J}_{\mathrm{paired}} = \{j \mid \mathcal{I}_{j,\mathrm{Shared}} \neq \emptyset \ \wedge \ \mathcal{I}_{j,\mathrm{Disjoint}} \neq \emptyset\}$$
and the paired set of outbound metrics was:
$$\left\{\left(b_{\mathrm{out}_{j,\mathrm{Shared}}},\ b_{\mathrm{out}_{j,\mathrm{Disjoint}}}\right) \mid j \in \mathcal{J}_{\mathrm{paired}}\right\}$$
Similar paired collections were constructed for inbound probabilities and both inbound and outbound nonzero average connection strengths, as well as for metrics where pre- or post-cells appear in both the Assembly and the Non-Assembly groups.

These sets provided the basis for the statistical tests used to compare metrics across connection types.

5.11.4 Statistical Tests

We performed one-way statistical tests at $\alpha = 0.05$ to compare the Shared and Disjoint groups, as well as to compare the Assembly and Non-Assembly groups. All alternative hypotheses predict Shared > Disjoint or Assembly > Non-Assembly, except for tests involving di-synaptic Inhibitory chain sets, in which the alternative hypotheses predict Shared < Disjoint or Assembly < Non-Assembly. We ran the following tests:

- For unpaired sets, we use a one-sided Wilcoxon Rank-Sum test
- For paired sets, we use a one-sided paired Wilcoxon Signed-Rank test to compare metrics within cells appearing in both groups.

For pairwise binary connectivity, we created a contingency table to compare the frequencies of successful and failed connections across connection types:

Connection Type	Successful Connections ($b_{ij} = 1$)	Failed Connections ($b_{ij} = 0$)
Shared	$\sum_{ij\|\text{Shared}_{ij}=1} b_{ij}$	$\sum_{ij\|\text{Shared}_{ij}=1}(1 - b_{ij})$
Disjoint	$\sum_{ij\|\text{Disjoint}_{ij}=1} b_{ij}$	$\sum_{ij\|\text{Disjoint}_{ij}=1}(1 - b_{ij})$

Then, we performed a Chi-Squared Test of Independence at $\alpha = 0.05$ to determine if pairwise connection frequency differs across connection types.

Finally, to examine the functional correlates of the di-synaptic inhibitory chain findings, we calculated the Pearson's correlation coefficient between the summed feed-forward inhibitory weights of disjoint cells connecting an assembly pair (see Methods 5.11.2) and the correlation scores between the assembly pair's coactivity traces (see Methods 5.8), and examined significance.

5.11.5 Tail Analysis

In addition to the mono-synaptic and di-synaptic analyses, we performed a "tail" analysis to investigate whether the proportion of Shared versus Disjoint connections differs between all pairwise connections and those classified as "tail" connections.

To identify "tail" connections, we modeled the distribution of connection strengths using a Gaussian Mixture Model (GMM) with $k = 2$ components. The model was initialized via k-means clustering to estimate the weights, means, and standard deviations of each component. The decision boundary separating the two Gaussian components was calculated as the intersection of their weighted probability density functions, derived using a quadratic equation based on the GMM parameters. Connections with values greater than or equal to the decision boundary were classified as "tail" connections. We present the model fit and evaluation as well as the tail boundary in the supplemental figures section (see Supp. Fig. 8).

Once the tail connections were identified, we compared the proportions of Shared and Disjoint connections in this subset to their proportions in the full dataset using a Chi-Squared Goodness-of-Fit Test at $\alpha = 0.05$. This test considered only the Shared and Disjoint groups, with expected proportions calculated relative to the total counts of these two groups in the full dataset.

5.11.6 Centrality analysis

Centrality analysis was used to quantify whether in a given network, assembly cells were more likely to be central to the network than non-assembly cells. This analysis gave further insights into the role of assembly cells in higher-order connectivity. To do this, we measured different centrality metrics for assembly and non-assembly cells, namely, in-degree centrality, outdegree centrality, closeness centrality, and betweenness centrality.

In a graph, the centrality of a node refers to its tendency to connect and generally influence other nodes within the network [77]. We developed a directed graph $G = (N, E)$, using the binary connectome such that $|N|$ represents the number of cells in the connectome and E represents the binary, directed connections between all cell pairs. The total number of outbound synaptic connections are given by $\sum_n \deg^+(n)$ and inbound

synaptic connections are given by $\sum_n \deg^-(n)$. Normalizing these connections, in-degree centrality was calculated as

$$I = \frac{\sum_n \deg^-(n)}{|N| - 1}$$

and outdegree centrality was calculated as

$$D = \frac{\sum_n \deg^+(n)}{|N| - 1}$$

where $n \in N$.

A monosynaptic connection, $e = (j, i) \in E$ between a pre-cell j and post-cell i, will have i as the head and j as the tail end of the connection, where j and $i \in N$. The path between a pre-cell j acting as a source neuron and a post-cell i acting as a target neuron, is the alternating sequence of cells and connections starting from j and ending at i, with each cell before a connection in the sequence being a pre-cell and each cell after a connection being a post-cell. The number of monosynaptic connections within the path indicates the length of the path. The shortest path between cells j and i is the minimum length between the two cells. The shortest path, ξ_{ij} , is also referred to as the geodesic path.

Based on this, the closeness centrality for a given cell i was calculated as the reciprocal of the sum of the shortest paths, or distances between the post-cell i and all other pre-cells j in the graph,

$$V(i) = \frac{|N| - 1}{\sum_{j=1}^{N} \xi_{ij}}.$$

For a given cell k, the betweenness centrality [78] was calculated as

$$B(k) = \sum_{(i,j)} \frac{\xi(k)_{ij}}{\xi_{ij}}$$

where $\xi(k)_{ij}$ is the number of shortest paths between pre-cell j and post-cell i that pass through cell k. This value was also normalized to fall between 0 and 1.

6 Code Availability

All code is available at this GitHub link: `https://github.com/AllenInstitute/HebbsVision`

It should be noted that, for all code of our own generation, a fixed random seed is used for reproducibility. Libraries used for analysis and plotting of results include: numpy [79], pandas [80], scikit-learn [81], upsetplot [34], networkx [82], dotmotif [83], matplotlib [84], seaborn [85], and Raincloud [86].

7 Acknowledgements

We wish to thank the Allen Institute for Brain Science founder, Paul G. Allen, for his vision, encouragement, and support. S.B. has been supported by NIH R01EB029813. S.M. has been in part supported by NSF 2223725, NIH R01EB029813, and RF1DA055669 grants.

8 List of Contributions

Name	Contributions
J Wagner-Carena	Conception and Design, Functional and Structural Data Analysis, Data Interpretation, Writing, Revision
S Kate	Data Interpretation, Chain and Graph Theoretic Data Analysis, Writing, Revision
T Riordan	Data Interpretation, Statistical Data Analysis, Writing, Revision
R Abbasi-Asl	Data processing, Revision, Writing
J Aman	Coregistration (Data Analysis), Writing
A Amster	Functional ROI extraction, team PIKA
AL Bodor	Data Acquisition: EM Histology/Sectioning; Revision
D Brittain	Data Acquisition: EM Imaging; Revision
JA Buchanan	Data Acquisition: EM Histology; Revision
MA Buice	Data processing: ROI extraction; Revision
DJ Bumbarger	Data Acquisition: EM Sectioning; Revision
F Collman	Interpretation, Dataset Resources and Infrastructure; Revision
NM da Costa	Data Processing; Interpretation; Writing, Revision
DJ Denman	Data Acquisition: Imaging, stimuli
SEJ deVries	Scientific leadership, optical physiology
E Joyce	Data Processing: Skeleton generation, layer definition, soma&nucleus feature extraction; Revision
D Kapner	Data Processing: EM alignment; Revision
CW King	Software development: allen_v1dd library
JD Larkin	Data Acquisition: Functional Imaging
D Liu	Software development: allen_v1dd library
G Mahalingam	Data Processing: EM Image Processing; Revision
D Millman	Stimulus Creation, Data Analysis
J Mölter	SGC Algorithm, Writing, Revision
C Morrison	Functional ROI extraction, team PIKA
RC Reid	Interpretation; Dataset Resources and Infrastructure, Revision
CM Schneider-Mizell	Interpretation; Dataset Resources and Infrastructure, Revision
D Scott	Functional ROI extraction, team PIKA
S Suckow	Project Management
KT Takasaki	Data Acquisition: Functional Imaging
M Takeno	Data Acquisition: EM Histology/Sectioning; Data Processing: coRegistration; Revision
R Torres	Data Processing: EM Image Processing/Coregistration; Revision; Writing
D Vumbaco	Proofreading Project Management
J Waters	Microscopy team lead
DG Wyrick	Software development: allen_v1dd library
W Yin	Data Acquisition: EM Imaging; Revision
J Zhuang	Data processing
S Mihalas	Conception and Design, Data Interpretation, Writing, Revision
S Berteau	Conception and Design, Data Analysis, Data Interpretation, Writing, Revision

References

[1] Hebb, D. *The Organization of Behavior* (John Wiley & Sons, New York, 1949).

[2] Frégnac, Y. Aplysia: Hebbian or not? *Trends in Neurosciences* **9**, 410 (1986).

[3] Buzsáki, G. Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron* **68**, 362–385 (2010).

[4] Miehl, C., Onasch, S., Festa, D. & Gjorgjieva, J. Formation and computational implications of assemblies in neural circuits. *The Journal of Physiology* **601**, 3071–3090 (2023).

[5] Matheus Gauy, M. *et al.* A Hippocampal Model for Behavioral Time Acquisition and Fast Bidirectional Replay of Spatio-Temporal Memory Sequences. *Frontiers in Neuroscience* **12**, 961 (2018).

[6] Carrillo-Reid, L., Miller, J.-E. K., Hamm, J. P., Jackson, J. & Yuste, R. Endogenous sequential cortical activity evoked by visual stimuli. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* **35**, 8813–8828 (2015).

[7] MacLean, J. N., Watson, B. O., Aaron, G. B. & Yuste, R. Internal Dynamics Determine the Cortical Response to Thalamic Stimulation. *Neuron* **48**, 811–823 (2005).

[8] Pérez-Ortega, J., Alejandre-García, T. & Yuste, R. Long-term stability of cortical ensembles. *eLife* **10**, e64449 (2021). Publisher: eLife Sciences Publications, Ltd.

[9] Hebb, D. O. The semiautonomous process: Its nature and nurture. *American Psychologist* **18**, 16–27 (1963). Place: US Publisher: American Psychological Association.

[10] Brock, L. G., Coombs, J. S. & Eccles, J. C. The recording of potentials from motoneurones with an intracellular electrode. *The Journal of Physiology* **117**, 431–460 (1952).

[11] Marr, D. & Brindley, G. S. Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* **262**, 23–81 (1997). Publisher: Royal Society.

[12] Traub, R. D., Tu, Y. & Whittington, M. A. Cell assembly formation and structure in a piriform cortex model. *Reviews in the Neurosciences* **33**, 111–132 (2022).

[13] Gastaldi, C., Schwalger, T., Falco, E. D., Quiroga, R. Q. & Gerstner, W. When shared concept cells support associations: Theory of overlapping memory engrams. *PLOS Computational Biology* **17**, e1009691 (2021). Publisher: Public Library of Science.

[14] Holtmaat, A. & Caroni, P. Functional and structural underpinnings of neuronal assembly formation in learning. *Nature Neuroscience* **19**, 1553–1562 (2016).

[15] Grinvald, A., Arieli, A., Tsodyks, M. & Kenet, T. Neuronal assemblies: Single cortical neurons are obedient members of a huge orchestra. *Biopolymers* **68**, 422–436 (2003).

[16] Harris, K. D. Neural signatures of cell assembly organization. *Nature Reviews Neuroscience* **6**, 399–407 (2005). Number: 5 Publisher: Nature Publishing Group.

[17] Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G. & Buzsáki, G. Organization of cell assemblies in the hippocampus. *Nature* **424**, 552–556 (2003).

[18] Kamiński, J. *et al.* Persistently active neurons in human medial frontal and medial temporal lobe support working memory. *Nature Neuroscience* **20**, 590–601 (2017).

[19] Wilson, M. A. & McNaughton, B. L. Reactivation of hippocampal ensemble memories during sleep. *Science (New York, N.Y.)* **265**, 676–679 (1994).

[20] Yuste, R., Nelson, D. A., Rubin, W. W. & Katz, L. C. Neuronal domains in developing neocortex: mechanisms of coactivation. *Neuron* **14**, 7–17 (1995).

[21] Sherafati, A. *et al.* Depth-dependant heterogeneity of neuronal responses in mouse primary visual cortex}. *BioRxiv* (2025).

[22] Grienberger, C. & Konnerth, A. Imaging Calcium in Neurons. *Neuron* **73**, 862–885 (2012). Publisher: Elsevier.

[23] Scheffer, L. K. *et al.* A connectome and analysis of the adult Drosophila central brain. *eLife* **9**, e57443 (2020). Publisher: eLife Sciences Publications, Ltd.

[24] Dorkenwald, S. *et al.* Neuronal wiring diagram of an adult brain. *Nature* **634**, 124–138 (2024). Publisher: Nature Publishing Group.

[25] Witvliet, D. *et al.* Connectomes across development reveal principles of brain maturation. *Nature* **596**, 257–261 (2021). Number: 7871 Publisher: Nature Publishing Group.

[26] Gour, A. *et al.* Postnatal connectomic development of inhibition in mouse barrel cortex. *Science (New York, N.Y.)* **371**, eabb4534 (2021).

[27] Bae, J. A. *et al.* Functional connectomics spanning multiple areas of mouse visual cortex. *Nature* **640**, 435–447 (2025). Publisher: Nature Publishing Group.

[28] Shapson-Coe, A. *et al.* A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science* **384**, eadk4858 (2024). Publisher: American Association for the Advancement of Science.

[29] Bock, D. D. *et al.* Network anatomy and in vivo physiology of visual cortical neurons. *Nature* **471**, 177–182 (2011).

[30] Lee, W.-C. A. *et al.* Anatomy and function of an excitatory network in the visual cortex. *Nature* **532**, 370–374 (2016). Number: 7599 Publisher: Nature Publishing Group.

[31] Massey, F. J. The Kolmogorov-Smirnov Test for Goodness of Fit. *Journal of the American Statistical Association* **46**, 68–78 (1951). Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

[32] Tootell, R. B. H. *et al.* Functional analysis of primary visual cortex (V1) in humans. *Proceedings of the National Academy of Sciences* **95**, 811–817 (1998). Publisher: Proceedings of the National Academy of Sciences.

[33] Mölter, J., Avitan, L. & Goodhill, G. J. Detecting neural assemblies in calcium imaging data. *BMC biology* **16**, 143 (2018).

[34] Lex, A., Gehlenborg, N., Strobelt, H., Vuillemot, R. & Pfister, H. UpSet: Visualization of Intersecting Sets. *IEEE Transactions on Visualization and Computer Graphics* **20**, 1983–1992 (2014). Conference Name: IEEE Transactions on Visualization and Computer Graphics.

[35] Dorfman, R. A Formula for the Gini Coefficient. *The Review of Economics and Statistics* **61**, 146–149 (1979). Publisher: The MIT Press.

[36] Walker, E. Y. *et al.* Inception loops discover what excites neurons most using deep predictive models. *Nature Neuroscience* **22**, 2060–2065 (2019). Number: 12 Publisher: Nature Publishing Group.

[37] Green, M. Contrast detection and direction discrimination of drifting gratings. *Vision Research* **23**, 281–289 (1983).

[38] Abdel-aziem, A. H. & Soliman, T. H. M. A Multi-Layer Perceptron (MLP) Neural Networks for Stellar Classification: A Review of Methods and Results. *International Journal of Advances in Applied Computational Intelligence* **Volume 3**, 29–37 (2023). Publisher: American Scientific Publishing Group (ASPG).

[39] Bhatkar, A. P. & Kharat, G. Detection of Diabetic Retinopathy in Retinal Images Using MLP Classifier. In *2015 IEEE International Symposium on Nanoelectronic and Information Systems*, 331–335 (2015).

[40] Alon, U. Network motifs: theory and experimental approaches. *Nature Reviews Genetics* **8**, 450–461 (2007). Number: 6 Publisher: Nature Publishing Group.

[41] Sporns, O. & Kötter, R. Motifs in Brain Networks. *PLoS Biology* **2**, e369 (2004).

[42] Hu, Y., Trousdale, J., Josić, K. & Shea-Brown, E. Motif statistics and spike correlations in neuronal networks. *Journal of Statistical Mechanics: Theory and Experiment* **2013**, P03012 (2013). Publisher: IOP Publishing and SISSA.

[43] Hebb, D. O. Intelligence, Brain Function and The Theory of Mind. *Brain* **82**, 260–275 (1959).

[44] Holtmaat, A. & Svoboda, K. Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience* **10**, 647–658 (2009). Number: 9 Publisher: Nature Publishing Group.

[45] Mongillo, G., Rumpel, S. & Loewenstein, Y. Intrinsic volatility of synaptic connections — a challenge to the synaptic trace theory of memory. *Current Opinion in Neurobiology* **46**, 7–13 (2017).

[46] Aitken, K., Garrett, M., Olsen, S. & Mihalas, S. The geometry of representational drift in natural and artificial neural networks. *PLOS Computational Biology* **18**, e1010716 (2022). Publisher: Public Library of Science.

[47] Schoonover, C. E., Ohashi, S. N., Axel, R. & Fink, A. J. P. Representational drift in primary olfactory cortex. *Nature* **594**, 541–546 (2021). Number: 7864 Publisher: Nature Publishing Group.

[48] Olshausen, B. A. & Field, D. J. Sparse coding of sensory inputs. *Current Opinion in Neurobiology* **14**, 481–487 (2004).

[49] Okun, M. *et al.* Diverse coupling of neurons to populations in sensory cortex. *Nature* **521**, 511–515 (2015). Publisher: Nature Publishing Group.

[50] Yuste, R., Cossart, R. & Yaksi, E. Neuronal ensembles: Building blocks of neural circuits. *Neuron* S0896–6273(23)00967–4 (2024).

[51] Vogels, T. P. & Abbott, L. F. Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature Neuroscience* **12**, 483–491 (2009). Publisher: Nature Publishing Group.

[52] Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C. & Gerstner, W. Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science* **334**, 1569–1573 (2011). Publisher: American Association for the Advancement of Science.

[53] Xu, S. *et al.* Behavioral state coding by molecularly defined paraventricular hypothalamic cell type ensembles. *Science* **370**, eabb2494 (2020). Publisher: American Association for the Advancement of Science.

[54] de Vries, S. E. J. *et al.* A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience* **23**, 138–151 (2020). Publisher: Nature Publishing Group.

[55] Hua, Y., Laserstein, P. & Helmstaedter, M. Large-volume en-bloc staining for electron microscopy-based connectomics. *Nature Communications* **6**, 7923 (2015). Publisher: Nature Publishing Group.

[56] Yin, W. *et al.* A petascale automated imaging pipeline for mapping neuronal circuits with high-throughput transmission electron microscopy. *Nature Communications* **11**, 4949 (2020). Publisher: Nature Publishing Group.

[57] Macrina, T. *et al.* Petascale neural circuit reconstruction: automated methods (2021). Pages: 2021.08.04.455162 Section: New Results.

[58] Phelps, J. S. *et al.* Reconstruction of motor control circuits in adult Drosophila using automated transmission electron microscopy. *Cell* **184**, 759–774.e18 (2021). Publisher: Elsevier.

[59] Mahalingam, G. *et al.* A scalable and modular automated pipeline for stitching of large electron microscopy datasets. *eLife* **11**, e76534 (2022). Publisher: eLife Sciences Publications, Ltd.

[60] Schneider-Mizell, C. M. *et al.* Inhibitory specificity from a connectomic census of mouse visual cortex. *Nature* **640**, 448–458 (2025). Publisher: Nature Publishing Group.

[61] Dorkenwald, S. *et al.* CAVE: Connectome Annotation Versioning Engine (2023). Pages: 2023.07.26.550598 Section: New Results.

[62] Schindelin, J. *et al.* Fiji: an open-source platform for biological-image analysis. *Nature Methods* **9**, 676–682 (2012). Publisher: Nature Publishing Group.

[63] Bogovic, J. A., Hanslovsky, P., Wong, A. & Saalfeld, S. Robust registration of calcium images by learned contrast synthesis. In *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*, 1123–1126 (2016). ISSN: 1945-8452.

[64] Avitan, L. *et al.* Spontaneous Activity in the Zebrafish Tectum Reorganizes over Development and Is Influenced by Visual Experience. *Current Biology* **27**, 2407–2419.e4 (2017). Publisher: Elsevier.

[65] Almeida-Filho, D. G. *et al.* An investigation of Hebbian phase sequences as assembly graphs. *Frontiers in Neural Circuits* **8**, 34 (2014).

[66] Kursin, A., Húsek, D. & Neruda, R. Faster Learning with Overlapping Neural Assemblies. In Kollias, S. D., Stafylopatis, A., Duch, W. & Oja, E. (eds.) *Artificial Neural Networks – ICANN 2006*, Lecture Notes in Computer Science, 226–233 (Springer, Berlin, Heidelberg, 2006).

[67] Mölter, J. & Goodhill, G. J. Detecting Neural Assemblies Through Similarity Graph Clustering. In Carrillo-Reid, L. (ed.) *Identification, Characterization, and Manipulation of Neuronal Ensembles*, 167–176 (Springer US, New York, NY, 2025).

[68] Miller, R. G. A Trustworthy Jackknife. *The Annals of Mathematical Statistics* **35**, 1594–1605 (1964). Publisher: Institute of Mathematical Statistics.

[69] Richter, C. G., Thompson, W. H., Bosman, C. A. & Fries, P. A jackknife approach to quantifying single-trial correlation between covariance-based metrics undefined on a single-trial basis. *NeuroImage* **114**, 57–70 (2015).

[70] Hurley, N. & Rickard, S. Comparing Measures of Sparsity. *IEEE Transactions on Information Theory* **55**, 4723–4741 (2009). Conference Name: IEEE Transactions on Information Theory.

[71] Khosravy, M., Nitta, N., Gupta, N., Patel, N. & Babaguchi, N. Chapter 3 - A descriptive review to sparsity measures. In Khosravy, M., Dey, N. & Duque, C. A. (eds.) *Compressive Sensing in Healthcare*, Advances in ubiquitous sensing applications for healthcare, 43–63 (Academic Press, 2020).

[72] Wright Muelas, M., Mughal, F., O'Hagan, S., Day, P. J. & Kell, D. B. The role and robustness of the Gini coefficient as an unbiased tool for the selection of Gini genes for normalising expression profiling data. *Scientific Reports* **9**, 17960 (2019). Number: 1 Publisher: Nature Publishing Group.

[73] Hetherington, P. A., & Shapiro, M. L. Simulating Hebb cell assemblies: the necessity for partitioned dendritic trees and a post-not-pre LTD rule. *Network: Computation in Neural Systems* **4**, 135–153 (1993). Publisher: Taylor & Francis.

[74] Garagnani, M., Wennekers, T. & Pulvermüller, F. Recruitment and Consolidation of Cell Assemblies for Words by Way of Hebbian Learning and Competition in a Multi-Layer Neural Network. *Cognitive Computation* **1**, 160–176 (2009).

[75] Fransen, E., Lansner, A. & Liljenström, H. A Model of Cortical Associative Memory Based on Hebbian Cell Assemblies. In Eeckman, F. H. & Bower, J. M. (eds.) *Computation and Neural Systems*, 431–435 (Springer US, Boston, MA, 1993).

[76] Dorkenwald, S. *et al.* Binary and analog variation of synapses between cortical pyramidal neurons. *eLife* **11**, e76120 (2022). Publisher: eLife Sciences Publications, Ltd.

[77] Bavelas, A. A Mathematical Model for Group Structures. *Human Organization* **7**, 16–30 (2009).

[78] Freeman, L. C. A Set of Measures of Centrality Based on Betweenness. *Sociometry* **40**, 35–41 (1977). Publisher: [American Sociological Association, Sage Publications, Inc.].

[79] Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362 (2020). Publisher: Nature Publishing Group.

[80] McKinney, W. Data Structures for Statistical Computing in Python. *scipy* (2010).

[81] Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011).

[82] Hagberg, A. A., Schult, D. A. & Swart, P. J. Exploring Network Structure, Dynamics, and Function using NetworkX. *scipy* (2008).

[83] Matelsky, J. K. *et al.* DotMotif: an open-source tool for connectome subgraph isomorphism search and graph queries. *Scientific Reports* **11**, 13045 (2021).

[84] Hunter, J. D. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering* **9**, 90–95 (2007). Conference Name: Computing in Science & Engineering.

[85] Waskom, M. L. seaborn: statistical data visualization. *Journal of Open Source Software* **6**, 3021 (2021).

[86] Micah, A. *et al.* Raincloud plots: a multi-platform tool for robust data visualization. *Wellcome Open Research* (2021).

9 Supplemental Figures

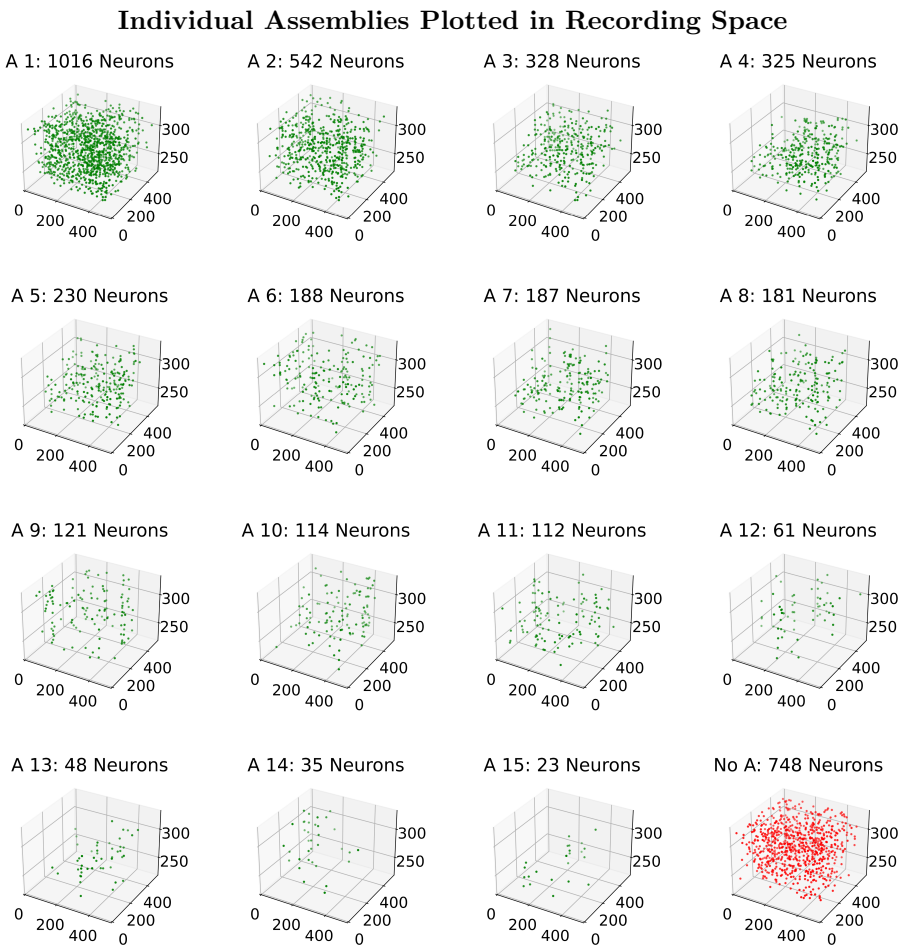

Supplementary Figure 1. Spatial positions of individual extracted assemblies in the three-dimensional recording field. Every subplot provides an isolated view of an assembly, ordered by size, visualized in the optical imaging recording space. Each individual point refers to an identified excitatory neuron. The plot of neurons assigned to no assemblies, 'No A', is also shown in red (bottom-right). Sub-plot axes refer to the three spatial dimensions of the recording field, with units in micrometers. Sub-plot titles also include the size of each assembly. In total, 1960 neurons were assigned to assemblies, and 748 were not. Assemblies are typically spatially spread out throughout the recording field.

Assembly Oracle Scores as a function of Fullscreen Grating Orientations

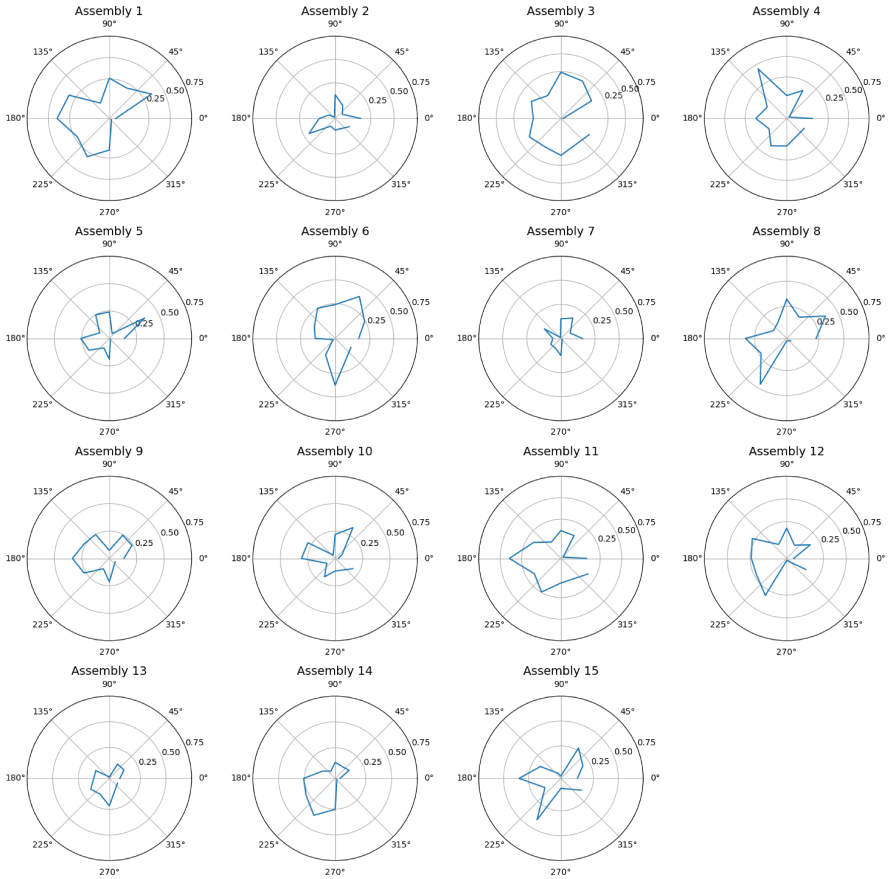

Supplementary Figure 2. Individual Assemblies Oracle Scores to Fullscreen Gratings. Every sub-plot provides an isolated visualization of the reliability in an assembly's response with respect to the orientation of fullscreen gratings. Orientation of gratings is represented by a polar plot. Reliability is measured through the Oracle score metric (see Methods 5.3). The scores of assembly coactivity trace in response to gratings are typically lower than those seen in natural movies (Fig. 3C), but the results are still indicative of tuning properties in these functional populations. Notably, some assembly traces seem to be highly reliable to particular orientations, similar to the orientation receptive fields of simple cells in the primary visual cortex.

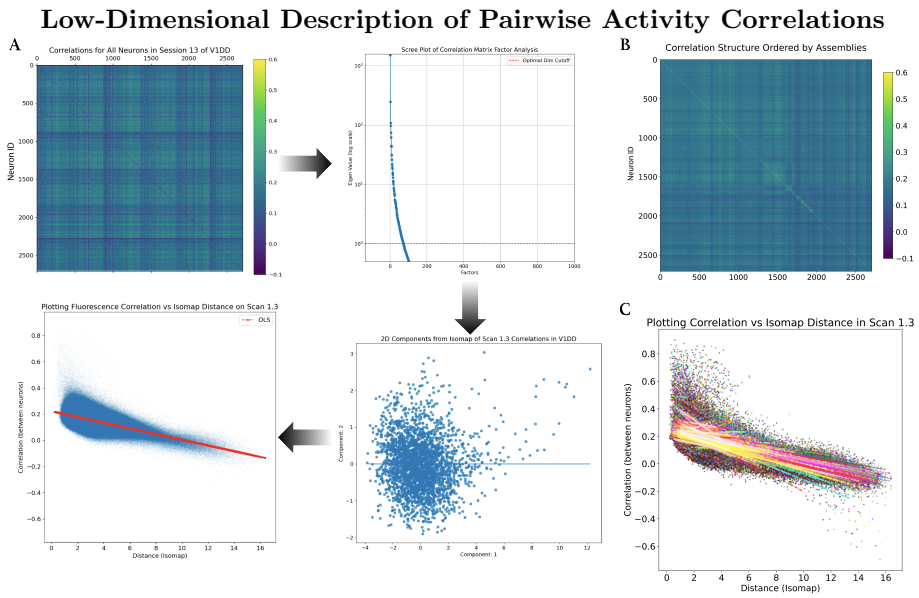

Supplementary Figure 3. Fitting Regressions to establish low-dimensional descriptions of pairwise activity correlations of cells. This framework was used to choose the optimal hyperparameters from Scan Volume 4, with those parameters later being chosen for Scan Volume 3. **(A)**. Pipeline to produce our regression fits to activity correlations. *Top-left*: The procedure begins with the pairwise correlation matrix of the whole scan throughout the hour recording. *Top-right*: Factor analysis reveals which dimensions of the original data matrix explain the most variance. This is done by plotting the eigenvalue of each factor, or dimension. The dotted red line is a horizontal marker of the eigenvalue of one, a threshold for considered dimensions justified by the Kaiser criterion. We ensure that at least ninety percent of the total variance is left explained in the final embedding. Less than fifty factors, or dimensions, are needed to do so. *Bottom-right*: With these dimensions in mind, we apply Isomap to produce a non-linear low-dimensional manifold of the original correlation matrix. This plot illustrates the first and second components from that embedding. *Bottom-right*: Finally, our low-dimensional description of the activity correlation, or our activity correlation space, is produced by plotting the distance of neurons across the Isomap embedding with respect to their pairwise correlations. The OLS fit (solid-red line) is a statistical model exemplifying how well the system continuum is able to describe this activity correlation space. **(B)** Correlation matrix of the whole scan with neurons ordered by assemblies. Cells on each axis were ordered by their largest corresponding assembly, from 'A 1' to 'A 15'. **(C)** OLS fits were developed that corresponded to assembly assignments with pairwise cells. This implies that multiple OLS fits were performed to account for all assignments. Regression fits to different assembly assignments are compared through the Akaike information criterion.

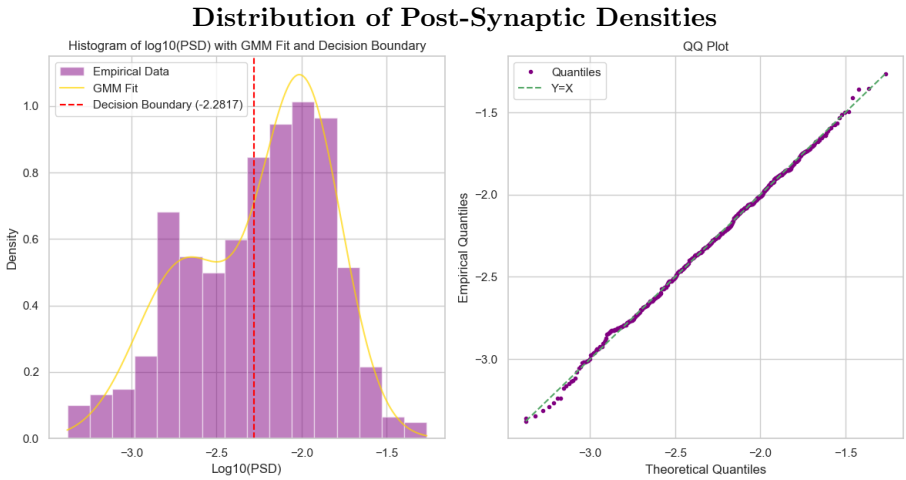

Supplementary figure 4. Histogram of connection strengths (log10-scaled) with GMM fit overlaid, and QQ plot evaluating the fit. The red dashed line indicates the decision boundary separating the two Gaussian components, used to classify "tail" connections.