A connectomics-driven analysis reveals novel characterization of border regions in mouse visual cortex

Neehal Tumma*[a], Linghao Kong*[b], Shashata Sawmya[b], Tony T. Wang[b], Nir Shavit[b]

[a]*Department of Computer Science, Harvard University, Cambridge, 02138, MA, USA*
[b]*Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, 02139, MA, USA*

Abstract

Leveraging retinotopic maps to parcellate the visual cortex into its respective sub-regions has long been a canonical approach to characterizing the functional organization of visual areas in the mouse brain. However, with the advent of extensive connectomics datasets like MICrONS, we can now perform more granular analyses on biological neural networks, enabling us to better characterize the structural and functional profile of the visual cortex. In this work, we propose a statistical framework for analyzing the MICrONS dataset, focusing our efforts on the network encompassed by the retinotopically-induced V1, RL, and AL visual areas. In particular, we bridge the gap between connectomics and retinotopy by identifying several structural and functional differences between these regions. Most notably, by placing our attention on the *borders* between these regions, we demonstrate how connectomics, in some ways, supersedes retinotopy, providing evidence for two major findings. One, by comparing the V1-RL and RL-AL border regions, we show that not all borders in the visual cortex are the same with respect to structure and function. Two, we propose a novel interpretation for the V1-RL border region in particular, motivating it as a subnetwork that possesses heightened synaptic connectivity and more synchronous neural activity. Going one step further, we analyze structure and function in tandem by measuring information flow along synapses, demonstrating that the V1-RL border serves as a bridge for communication between the V1 and RL visual areas, offering justification as to why it presents itself uniquely with respect to both structure and function.

1. Introduction

The mammalian visual cortex is among the most advanced visual processing systems in existence. It comprises multiple visual areas, including the primary visual cortex (V1), which directly receives most of the visual inputs from the thalamus, and other higher order visual areas. In mice, such higher order visual areas include the rostrolateral (RL), anterolateral (AL), and lateromedial areas (Marshel et al., 2011). These regions have highly complex synaptic connections within and between each other (Nassi and Callaway, 2006) and display unique functional characteristics, with one of the most notable being retinotopy (Tyler, 2004).

Retinotopy is a property in which portions of the visual field are each represented by and encoded within specific areas in the visual cortex (Fize et al., 2003). Classically, retinotopy is most often determined through experiments including the rotating bar and expanding ring stimuli (Tyler et al., 2005), though more recent work allows retinotopic mappings to be recovered from natural movie stimuli (Kumar et al., 2021). Retinotopic maps are canonically used to induce a parcellation of the visual cortex into its sub-regions (Zhuang et al., 2017), characterizing the spatial arrangement of neurons in the visual cortex and defining boundaries between different visual areas (i.e. V1, RL, AL, etc.).

In this work, we explore a different, more recently developed framework for analyzing the visual cortex (and more generally the brain) known as *connectomics*. In particular, we strive to understand how connectomics can be used to improve upon our existing characterization of the mouse visual cortex currently enabled by retinotopy. Connectomics is an emerging field at the intersection of neuroscience and computer science that leverages high resolution electron microscopy to characterize exact neuron-to-neuron connections (Behrens and Sporns, 2012). One of the most comprehensive connectomics datasets is from the Machine Intelligence from Cortical Networks (MICrONS) consortium. It comprises a cubic millimeter of the mouse visual cortex and includes over 75,000 neurons and 500,000 synapses. Furthermore, it is one of the few connectomics datasets to also include functional traces associated with individual neurons (Consortium et al., 2021).

Connectomics differs from retinotopy in that it goes beyond studying the functional organization of specific sensory regions by providing a comprehensive and highly detailed view of the brain's structural and functional connectivity patterns (Dorkenwald et al., 2023). While retinotopy focuses on how sensory information is represented in specific brain areas, connectomics, especially at the microscale, offers an unprecedented level of detail (Scholtens and van den Heuvel, 2018). Connectomics maps the intricate network of individual neurons, their synapses, and the broader patterns of connectivity, surpassing the resolution of functional imaging methods like retinotopy. This more granular approach allows us to explore the fine-scale

*Denotes equal contribution
**Correspondence to Linghao Kong (linghao@mit.edu) and Nir Shavit (shanir@csail.mit.edu)

architecture of neural circuits, revealing the precise organization of connections and offering a holistic understanding of how information flows and is processed throughout the brain.

In this work, we use the unique opportunity presented by the MICrONS dataset to profile the structural and functional granularities of the mouse visual cortex and use this information to understand if regions in the mouse visual cortex delineated via retinotopy exhibit structural and/or functional differences when analyzed under the lens of connectomics. We place our attention on the V1, RL and AL visual areas and investigate the biological neural network induced by the portion of the connectomics data corresponding to these retinotopically-induced regions. In doing so, we focus much of our analysis on the borders between these regions: although retinotopic boundaries between visual regions have been shown to demonstrate unique phenomena such as mirroring and compression (Lyon et al., 2002), they are otherwise little understood. Using the improved power enabled by MICrONS, we draw novel conclusions about anatomical and functional patterns present at the V1-RL and RL-AL borders that have yet to be properly characterized using retinotopic data alone.

2. Related work

Defining regions in the mouse visual cortex. As discussed, one standard approach to defining visual areas in the mouse visual cortex is using a retinotopic mapping (Bibollet-Bahena et al., 2023). Such maps give rise to borders between regions which are linked to changes in the chirality of the retinotopic map, identified functionally through the recording of neuronal responses using electrodes or imaging techniques (Zhuang et al., 2017). Another approach to mapping the visual cortex is architectonics, which unlike retinotopy is instead concerned with the structural organization of the brain, providing information about the arrangement of neurons, cell types, and cortical layers (Herculano-Houzel et al., 2013; Olavarria and Montero, 1989). Borders constructed using this approach are associated with rapid shifts in the aforementioned structural metrics.

There are many studies debating whether there is a mismatch between the boundaries of visual regions discerned through retinotopy and architectonics (Schuett et al., 2002; Zhuang et al., 2017), demonstrating the need for a more robust and nuanced formulation of structure and function in the visual cortex. Connectomics serves to address some of the drawbacks of individually using retinotopy and architectonics to map the brain and its visual regions. In particular,

- Retinotopic analyses primarily operate in a functional coordinate space (Kalatsky and Stryker, 2003), while architectonic ones operate in a physical coordinate space; connectomics consolidates structural and functional data into a single, physical coordinate space via co-registration, enabling the analysis of structure and function in tandem
- Retinotopy, which primarily focuses on localized responses to visual stimuli, may not inherently capture broader network dynamics and connectivity properties; this is of particular importance when analyzing boundaries between visual areas in the network
- Connectomics surpasses architectonics in providing a more detailed and comprehensive view of brain structure by capturing multiscale connectivity patterns and detailed circuitry

Our work attempts to leverage the more nuanced profiling of the visual cortex conferred by connectomics by understanding how we can use it to improve upon existing characterizations enabled by retinotopy (and by extension architectonics). Notably, our work is the first of its kind to analyze the structure and function of the visual cortex in tandem, enabling us to make conclusions regarding information flow across border regions in the visual cortex.

Exploring MICrONS. Research on specifically the MICrONS connectomics dataset is ever-growing. Some groups have utilized the highly proofread region to analyze inhibitory and excitatory neurons in depth (Schneider-Mizell et al., 2023). They were able to find distinct subcategories of both types of neurons with differing morphology and synaptic connectivity, providing a more granular understanding of the nuances within the same cell types. Other groups have focused on the functional data MICrONS provides, creating machine learning models trained on the responses of individual neurons to predict their responses to novel stimuli (Wang et al., 2023). Our work on MICrONS stands distinct with respect to prior work in its attempt to explore the use of connectomics data in emulating the boundaries drawn by retinotopic mappings and, furthermore, characterize these boundaries in an unprecedented level of detail.

Relating function and form. Understanding the relationship between function and form has long been a goal in computational neuroscience, and has become increasingly more within reach due to the advent of connectomics. Some of the work enabled by connectomics, specifically in regard to neural circuitry, includes experiments showing that the behavior of highly complex circuits can be explained in a large degree by a few key anatomical parameters (Haber et al., 2023). Other work has characterized the functional behavior and neuronal connectivity of predictive error circuits within the mouse visual cortex, finding that layer 2/3 pyramidal neuronal responses are specifically amplified during unexpected stimuli (Furutachi et al., 2023). While the narrative of this work does not directly relate to the question of function versus form, we certainly explore such relationships through various modes of analysis that tie together structural and functional components of the V1-RL-AL network.

Biologically-inspired ANNs. Beyond a better understanding of biological neural networks from a neuroscience perspective, such findings can also be used to improve upon the design of artificial neural networks. There are a plethora of ANNs modeled after biological networks, including MouseNet (Fu et al., 2016) and BMFC-ARM (Shi et al., 2022), and there is substantial literature that compares the learned representations between brains and ANNs (Yamins and DiCarlo, 2016; Schrimpf

et al., 2018). While we do not explore extensions to artificial networks in our work, the potential to create a new ANN architecture modeled after our findings is a promising direction for future work (Abeles, 2009).

3. Main findings

In this section, we provide a brief overview of the main findings in this paper before discussing them in further detail in the following sections.

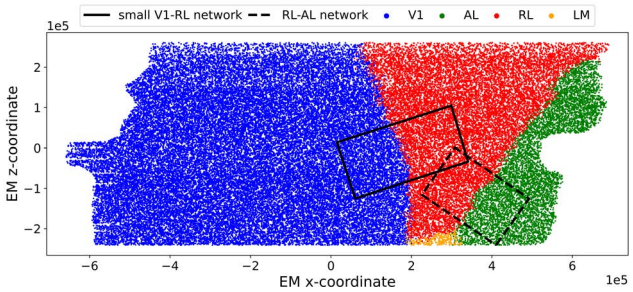

Figure 1: Visualization of the entire MICrONS dataset in EM-space, labelled by visual area (the axes are in units of μm). Note that it is projected onto the x-z plane (this is motivated in Section 5). The V1-RL and RL-AL networks/windows denote the portion of the dataset we analyze in the first part of this work.

Past analysis at the neuronal level has focused either on connectivity in space or on correlations of activations in time. Here, we use the unique dataset offered by MICrONS to, for the first time, analyze boundary regions in combined space-time. We are able to not only collect structural and functional statistics in isolation, but also track information flow along synaptically connected neurons across a large subset of the border regions comprising the mouse visual cortex. Such an integrated space-time analysis has previously been confined to much smaller, non-border regions (Bock et al., 2011).

The MICrONS dataset can be visualized in EM-coordinate space, with some subset of the neurons co-registered with functional activations. The dataset contains information regarding the connectivity amongst neurons, their cell types, their regional membership in the visual cortex determined via retinotopy and their firing patterns in response to various visual stimuli (Figure 1). Using MICrONS, we are able to conduct a novel combined analysis of the structural and functional components of the mouse visual cortex, placing our attention on the V1, RL and AL visual areas. Our work is best characterized in two parts. In the first, we aimed to answer the following high-level question.

- To what extent do the V1-RL and RL-AL networks (and in particular their border regions) differ with respect to both structure and function?

Leveraging various analyses driven by connectomics, we were able to address this question as follows:

- A structural and functional comparison between the V1-RL and RL-AL networks reveals that the two networks are distinct, pointing to a relationship between the information captured by retinotopy and the information captured by connectomics (Section 6.1, Section 6.2)
- A comparison between the V1-RL and RL-AL borders provides evidence for the larger notion that even border regions can present themselves differently within their larger networks. Namely, we find that the two borders show starkly different trends with respect to synaptic connectivity and correlations in neural activity (Section 6.1)
- We observe that the V1-RL border in particular possesses heightened synaptic connectivity and more pairwise correlated neural activity than the rest of the V1-RL network, supporting the notion that the V1-RL border serves a potentially more nuanced role within its network. In contrast, we observe neither of these trends at the RL-AL border. Instead, we find that the RL-AL border represents a region of notable shift in both synaptic connectivity and correlatedness of neural activity (Section 6.1)

Given the seemingly counterintuitive characterization of the V1-RL border as a more connected portion of its larger network, as well as the fact that the V1-RL network comprises a significant portion of the MICrONS dataset, the second part of our work aims to characterize the V1-RL network and its border region in further detail. In this part, we pose the following high-level questions:

- Beyond the metrics measured in the first part (i.e. neuron density, synaptic connectivity and activity correlation), to what extent do the retinotopically-induced V1 and RL regions vary with respect to both structure and function?
- What does a connectomics-driven analysis capture about borders that retinotopic maps do not?
- If we claim that the V1-RL border is a distinct, standalone region in the larger V1-RL network, how robust is our formulation?
- If the V1-RL border is a distinct region, what role does it play as a subnetwork of the larger V1-RL network?

We were able to address these questions as follows:

- We find that V1 and RL are not only distinct with respect to neuron density, synaptic density and correlatedness of functional activity, but also with respect to distribution over cell types, rate of information propagation, 2-cycle density, distribution over 3-motifs and distance between connected neurons (Section 6.2)
- Connectomics demonstrates that characterizing borders as regions of rapid shift in structure and/or function (as is done using retinotopy) is not necessarily the appropriate formulation. In particular, we find that the V1-RL border presents itself uniquely with respect to both structure and function, a notion that is not captured via parcellations

over retinotopic or architectonic maps (Section 6.3, Section 6.5). Furthermore, this motivates the V1-RL border as a standalone region within the larger V1-RL network.

- Our characterization of the V1-RL border as a disparate region in the visual cortex is robust. Namely, we find that it is identifiable amongst fake border regions constructed via translation and rotation (Section 6.5, Section 6.6)

- The V1-RL border region acts a bridge for communication between the V1 and RL visual areas. In particular, by analyzing novel message-passing motifs that combine both structure and function, we find that the information flow along synapses traversing the border is the highest amongst all synapses in the V1-RL network (Section 6.7)

These findings provide a thorough statistical characterization of the V1, RL, and AL visual regions to a degree that was previously not possible due to a lack of connectomics data like MICrONS. By profiling the V1-RL and RL-AL networks in such detail, we have identified many structural and functional granularities of the mouse visual cortex. Furthermore, by delving deeper into the V1-RL network in particular, we have put forth a novel characterization of the V1-RL border – offering a new interpretation regarding the role certain border regions serve in biological neural networks. We discuss these findings among others in the sections to follow.

4. Interpretations of borders in the visual cortex

In this section, we offer a high-level characterization of border regions between visual areas and delineate the set of possible interpretations for the role they can play within their respective networks.

Previous works that leverage retinotopic mappings have framed the characterization of the mouse visual cortex and its sub-regions as a parcellation problem (Kumar et al., 2021). These parcellation models yield borders that delineate boundaries between regions that have been determined to be retinotopically disparate. A characterization of the visual cortex and in particular the V1-RL-AL network admits two reasonable interpretations for the V1-RL and RL-AL borders: the cut interpretation (readily derived from a retinotopic parcellation) and the window interpretation (less readily derived from a retinotopic parcellation). We describe each in turn below.

4.1. Cut interpretation

Under the cut interpretation, a border serves as a divider between two regions that are retinotopically distinct. In this interpretation, we only assign importance to the structure and function induced by interactions between neurons on either side of the cut. The cut in this instance refers to the plane fit using visual area labels collected from a parcellation computed over retinotopic data (Figure 2A). We will expand further on the fitting procedure for the border in Section 5.2.

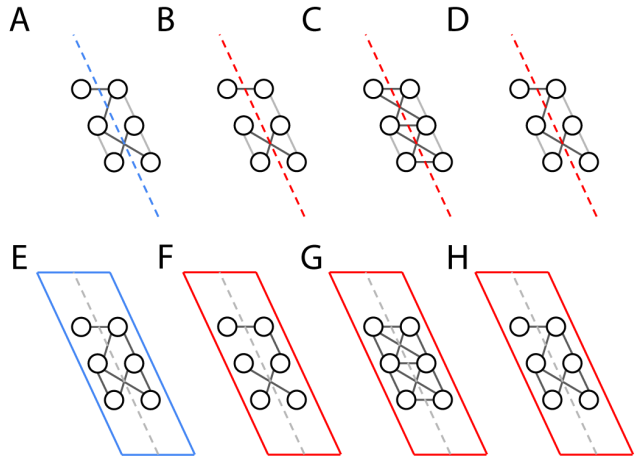

Figure 2: (A) Depiction of an arbitrary cut in the graph. Connections in black denote across cut connections, and connections in gray denote same side connections. We represent and discuss the other cut interpretations relative to this one. (B) Segregated cut regime. (C) Connected cut regime. (D) Homogenous cut regime. (E) Depiction of an arbitrary window in the graph. We represent and discuss the other window interpretations relative to this one. (F) Segregated window regime. (G) Connected window regime. (H) Homogenous window regime.

4.2. Window interpretation

Recall that in the cut interpretation, the volume surrounding the border plane is not assigned any importance. Under the window interpretation, we now consider all neurons within distance d of the plane as members of a border window (Figure 2E). Here, we refer to distance in the canonical sense as the orthogonal projection of the point corresponding to the neuron in EM space onto the border plane. This formulation now assigns importance to a volume of area that we can consider as a standalone region, distinct from the rest of the network.

4.3. Border regimes

In both the cut and window interpretations, we describe three possible regimes that potentially characterize the border: the segregated regime, the connected regime and the homogenous regime. We first introduce these three regimes in the context of the cut interpretation.

1. **Segregated cut regime** Pairs of neurons interacting across the border cut are more structurally segregated and functionally dissimilar than the average pair of neurons sampled across an arbitrary cut in the graph (Figure 2B)
2. **Connected cut regime** Pairs of neurons interacting across the border cut are more structurally connected and functionally similar than the average pair of neurons sampled across an arbitrary cut in the graph (Figure 2C)
3. **Homogeneous cut regime** Pairs of neurons interacting across the border cut look structurally and/or functionally the same as the average pair of neurons sampled across an arbitrary cut in the graph (Figure 2D)

Analogously, we can introduce three possible regimes the state of the network can lie in under the window interpretation:

1. **Segregated window regime** Pairs of neurons interacting within the border window are more structurally segregated and functionally dissimilar than the average pair of neurons sampled from an arbitrary window in the graph (Figure 2F)
2. **Connected window regime** Pairs of neurons interacting within the border window are more structurally connected and functionally similar than the average pair of neurons sampled from an arbitrary window in the graph (Figure 2G)
3. **Homogeneous window regime** Pairs of neurons interacting within the border window look structurally and/or functionally the same as the average pair of neurons sampled from an arbitrary window in the graph (Figure 2H)

One thing to note is that in these interpretations, we implicitly assume structure and function are tied together. While this might not necessarily be true, we will find in practice that this assumption holds (Section 6.2).

4.4. Contrasting the cut and window interpretations

The window interpretation encapsulates the cut interpretation. In particular, the window interpretation assigns importance to all pairs of neurons in the window: this includes both pairs of neurons interacting across the cut and pairs of neurons interacting on the same side of the cut. In contrast, the cut interpretation places emphasis only on the former. Note that previous works which model the border via a parcellation over retinotopic data lend themselves most readily to the cut interpretation, as parcellation algorithms assign no importance to the regions around the boundaries, only to the parcels. Note that the window interpretation has no meaning under this model, and is thus an avenue we hope to explore using connectomics.

In both the cut and window interpretations of the border, we have proposed three regimes for the state of a network. In the segregation regime, the border acts as a plane (cut interpretation) or region (window interpretation) over which information flow is significantly lower than the rest of the network. This is a natural hypothesis since it is reasonable to believe that two visual areas act individually as information processors and any information leakage from one area to the other would be detrimental to each regions' ability to specialize as a functional unit.

In the connected regime, we posit that the border presents itself as a hub for information aggregation and dissemination – acting as a conduit of information flow between the two regions. This is also a reasonable hypothesis, as it suggests that while the two regions act as distinct information processors, their ability to communicate and synergize enhances computation in the visual cortex.

In the homogenous regime, the border does not present itself uniquely with respect to either structure or function, and its existence is simply induced by differences in retinotopy between the two regions. In this case, the border denotes the shift in structure and function when going from one region to another.

We note that amongst the three proposed border regimes, the canonical interpretation under existing formulations of a border between visual areas is the homogenous regime. In particular, the leading assumption in these formulations is that the border simply represents the region of greatest shift in structure and/or function in the network. In contrast, the segregated and connected regimes take a different stance: instead, they posit that the border itself is a distinct construct, potentially serving its own role within the larger network.

In this work, by analyzing various metrics that measure both structure and function across the V1-RL and RL-AL networks, we will determine the appropriate characterizations for both the V1-RL and RL-AL border regions. In doing so, we demonstrate the inaccuracies of existing border interpretations induced by approaches that use retinotopic/architectonic maps to parcellate the visual cortex.

5. Methods

Before discussing our findings, we first detail and motivate the various frameworks of analysis used in this work.

5.1. Dataset description and construction

The connectomics dataset that we utilize is from the Machine Intelligence from Cortical Networks (MICrONS) consortium. It includes over 75,000 neurons and 500,000 synapses, and has, at the time of this writing, calcium and derived activity traces for approximately 33,000 of those neurons. Each neuron has its position recorded in 3-D electron microscope (EM) acquired space, as well as a list of synaptic connections it is involved in with other neurons. To distinguish different visual areas (i.e. V1, RL, AL) via retinotopy, the MICrONS group used a drifting bar stimulus to generate a sign map, and then manually annotated traces within that map.

For the purposes of our analysis, we will visualize the dataset in 2-D, with the dimension associated with cortical depth (i.e. the EM y-axis) removed, aggregating over all cortical layers (Figure 1). This is because we find that the meaningful variation for the purposes of our characterization of the V1-RL-AL network exists primarily along the x and z directions. However, we note that we do perform some analyses on a per-layer basis and will clarify as such when necessary. Finally, since the cortex depth is not perfectly aligned with the y-axis in MICrONS, we rotated the volume about the border between V1 and RL such that the cortical layers were aligned with the xz-plane.

In the first part of our work, we restrict the scope of our analysis to the portion of the V1-RL and RL-AL networks shown in Figure 1. Note that the windows are oriented such that the border axis is roughly parallel to the left and right edges that bound the window. This is done in order to ensure that analyses comparing the two window/networks are not confounded by the orientation of the border. Finally, note that the windows are constructed such that they are roughly equal in size.

In the analysis comparing the V1-RL and RL-AL networks, the size of the windows were bottlenecked by the data availability in the AL visual area afforded by MICrONS (Figure 1). In the second part of our work, we restrict the scope of our analysis to the V1-RL network (for reasons that are made clear in Section 6.1). Here, since we are no longer operating under this

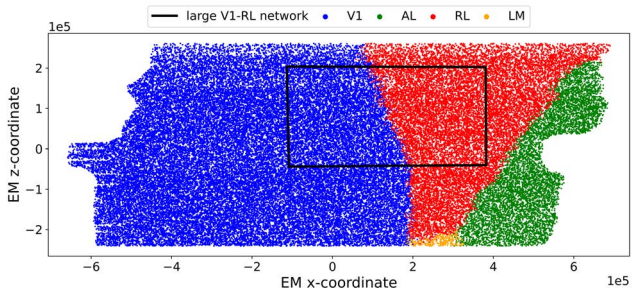

Figure 3: This visual depicts the increase in the portion of the network we analyze when focusing solely on the V1-RL network (as compared to Figure 1).

data constraint, we construct the window of analysis to capture as much of the dataset as possible. However, we also need to take into account edge effects, which refers to neurons near the edges of the dataset artificially showing fewer connections due to many of their neighboring neurons not being present in the data given by MICrONS. Hence, the V1-RL window, shown in Figure 3, captures as much data as possible while minimizing any potential confounders that may arise due to edge effects.

Finally, we note that there is a region within the dataset with a high degree of manual proofreading (Consortium et al., 2021) that we do not include to control for the proofread neurons acting as a potential confounding factor.

5.2. Fitting the borders

As described in Section 5.1, we can imagine the dataset as a 3-D graph projected onto the EM x-z plane. The retinotopic area assignments for each neuron in windows we consider (Figure 1, Figure 3) induce a boundary that can be approximated quite closely by a first-degree polynomial. Formulating the boundary as a linear function of the x and z coordinates simplifies more complex modes of analysis that we detail later on.

We formulate the fitting of the borders as a binary classification problem on the area assignment for each neuron. If we let y denote the vector of V1/RL area labels, x denote the vector of EM x-coordinates and z denote the vector of EM z-coordinates, the functional form of the border is given by a logistic regression as follows:

$$\hat{y} = \sigma(\beta_1 x + \beta_2 z + \beta_0)$$

Recall that the coregistered graph is a subset of the entire graph. We do not refit the borders when performing analysis on the coregistered graph; rather, we use the borders fit on the full graph as described above.

5.3. Metric definitions

In this section, we provide preliminaries for various structural and functional metrics leveraged in our analyses.

5.3.1. Graph representation of dataset

Consider a directed graph, $G = \{V, E\}$ which denotes the graph representation of the V1-RL network and the dataset. Take an arbitrary pair of neurons $v_1, v_2 \in V$. There exists an $e_{12} \in E$ if and only if there is a synapse going from v_1 to v_2. In this context, we refer to v_1 as the *pre-synaptic* neuron and v_2 as the *post-synaptic* neuron. We will use vertex and neuron interchangeably, as well as edge and synapse.

The graph G represents the full dataset. Recall that only a subset of it is coregistered; so, we denote this subgraph by $G_F = \{V_F, E_F\} \subseteq G$. Since the coregistered subgraph has functional recording data, we now represent a neuron $n \in V_F$ as a tuple given by (v, t) where v is the vertex in the graph corresponding to the neuron and t is the trace of functional activity corresponding to the neuron's activations induced by natural movie stimuli. We preprocess the neural traces by setting all measurements less than 10^{-3} to 0 as we consider measurements below this threshold as noise in the recording instrument.

To compute estimators for the metrics detailed below, we perform spatial bootstrapping on the graph (Thompson et al., 2016). This entails sampling sub-volumes of the 3-dimensional graph that represents the network. In particular, if we denote the dimensions of the subgraph volume by (X, Y, Z), then a randomly sampled sub-volume's dimensions are given by (tX, Y, tZ) for hyperparameter t such that $0 < t < 1$. Bootstrapping in this manner captures spatial correlations present in the data, which enables us to construct viable tests of statistical significance.

5.3.2. Neuron Density

The neuron density within a region R is calculated as:

$$D_n = \frac{N}{V}$$

- N is the total number of neurons within R.
- V is the volume of R.

5.3.3. Neuron Types

The neurons in the dataset can be categorized as either excitatory or inhibitory. For the purposes of our analyses, we do not distinguish between subgroups amongst inhibitory neurons. Within the excitatory class of neurons, we consider only pyramidal neurons which belong to one of the following four cortical layers:

- **Layer 2/3 (L2/3) - External Granular Layer:** Contains small pyramidal neurons and stellate cells; plays a role in integrating sensory information.
- **Layer 4 (L4) - Internal Granular Layer:** Receives direct input from the thalamus and is critical for processing visual information from the eyes.
- **Layer 5 (L5) - External Pyramidal Layer:** Contains large pyramidal neurons that project locally and to other brain regions; involved in relaying information and motor control.
- **Layer 6 (L6) - Polymorphic Layer:** Contains a mix of neurons with diverse functions, sending feedback signals to the thalamus and other areas.

5.3.4. Synaptic density

The synaptic density of a directed graph G with V vertices and E edges is calculated as:

$$D(G) = \frac{E}{V \cdot (V-1)}$$

5.3.5. 2-cycle density

The 2-cycle density (C_2) in a directed graph G can be calculated as follows:

$$C_2(G) = \frac{\sum_{v_1 \in V} \sum_{v_2 \in V} \mathbb{1}(v1 \neq v2) \cdot e_{v_1 v_2} e_{v_2 v_1}}{V(V-1)}$$

- V represents the set of vertices in the graph
- e_{ij} is 1 if and only if there is an edge going from node i to node j

5.3.6. Distribution over 3-motifs

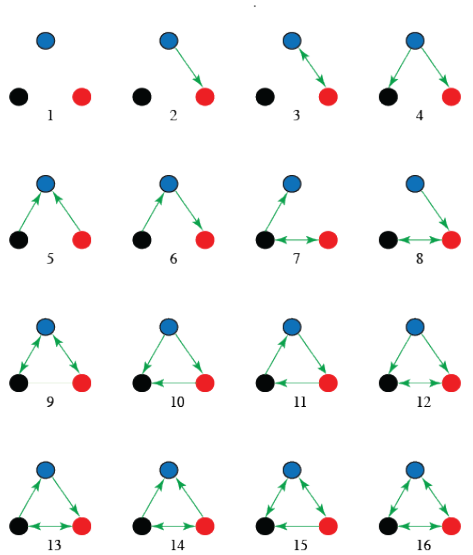

Figure 4: Set of all structural motifs formed by three neurons.

Let G be a network with N nodes. A 3-motif is a subgraph M of G consisting of three nodes and the edges between them (Figure 4). The distribution $P(M)$ over 3-motifs in the network G can be calculated by counting the occurrences of each distinct 3-motif M and normalizing by the total number of possible subgraphs of size 3.

$$P(M) = \frac{\text{Count}(M)}{\binom{N}{3}}$$

5.3.7. Microcircuit expansions

Let the set of neuron types be denoted by $T = \{2/3, 4, 5/6, \text{inh}\}$ where here we have grouped L5 and L6 neurons together. In this work, we will consider the following two microcircuits:

- $4 \rightarrow 5/6 \rightarrow 2/3 \rightarrow 4$
- $2/3 \rightarrow \text{inh} \rightarrow 4 \rightarrow 2/3$

Both of these circuits are an important unit of visual information processing (Hooks and Chen, 2020), with the former a general visual processing circuit and the latter a predictive coding circuit (Bastos et al., 2012). We can more generally represent these circuits as A $\rightarrow$ B $\rightarrow$ C $\rightarrow$ A for neuron types A, B, C. We define the expansion rate of such as a circuit as follows:

Consider a decomposition of G into sub-graphs G_A, G_B, G_C where each G_L corresponds to a graph comprised only neurons of type L. Consider an arbitrary node in G_A. Run a BFS from that node that expands onto the nodes in G_B, G_C and G_A in that order. Denote the number of nodes in each level of the BFS by the following 4-tuple: $(1, N_B, N_C, N_A)$. The expansion rate of the circuit for a single neuron v can be computed as the following 3-tuple:

$$E_v = (E_{AB}, E_{BC}, E_{CA})_v = (N_B, \frac{N_C}{N_B}, \frac{N_A}{N_C})_v$$

To compute the expansion rate of this circuit over the entire $G = \{V, E\}$, we element-wise average over all possible starting nodes:

$$E_G = \frac{1}{V} \sum_{v \in V} E_v$$

The expansion rate of these microcircuits provides us with estimates of the rate of signal propagation in units that are relevant to visual information processing.

5.3.8. Activity level

Recall that a neuron $n \in V_F$ has a corresponding trace t representing its functional activity. $t \in \mathbb{R}^{T+}$ where $T = 40000$. We can define the average activity $A(G_F)$ in the graph as follows

$$A(G_F) = \frac{1}{V_F} \sum_{(v,t) \in V_F} \sum_{j=1}^{T} \mathbb{1}(t_i \neq 0)$$

5.3.9. Average pairwise correlation

The pairwise correlation coefficient r between two neural traces X and Y can be calculated using the Pearson correlation formula:

$$r_{XY} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2 \sum_{i=1}^{n}(Y_i - \bar{Y})^2}}$$

We are interested in the average pairwise correlation over an entire graph G_F. Letting $N = V_F$, this is computed as

$$r_G = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} r_{ij}$$

5.3.10. Median information flow (Granger causality)

Pairwise correlation does not leverage the temporal relationship between firing neurons. To capture this, we use Granger causality as a measure for the strength of directed information flow (Seth et al., 2015). Consider a pair of neural traces X and Y corresponding to neurons n_x and n_y respectively such that

there exists an edge going from n_x to n_y. We let g_{XY} denote the p-value of a Granger causality test from X to Y computed for a fixed lag ℓ. Here we use $\ell = 5$. The median pairwise Granger causality measures the median directional influence of one neuron on another. We compute the median as opposed to the average, due to the skewness of the empirical distribution over p-values we observe. If we let $[g_{ij}]$ denote a list of all pairwise Granger causalities, then the median information flow is computed by

$$g_G = \text{median}([g_{ij}])$$

Note that lower values under this metric indicate higher Granger causality and stronger information flow.

5.3.11. *Effective dimensionality of neural activity*

The effective dimensionality of a set of neural traces (Jazayeri and Ostojic, 2021) is the cumulative explained variance of the top k principal components of the trace data for a fixed k. This is given by

$$k_{\text{eff}} = \frac{\sum_{i=1}^{k} \lambda_i}{\sum_{i=1}^{N \cdot T} \lambda_i},$$

- N is the number of traces in the dataset.
- T is the number of time points in each trace.
- k is the number of principal components being considered.
- λ_i represents the eigenvalue of the i-th principal component, sorted in descending order.

Note that higher values of this metric indicate lower dimensional activity.

5.4. *Sliding window framework*

When asking questions about structure and function in the V1-RL-AL network, we can use the metrics described above. To compute these metrics, one approach is to aggregate over the entirety of V1, RL and AL individually and calculate metrics on each subgraph and compare them. However, this does not allow us to characterize the potentially more granular changes in structure and function that may be present over the network. Here, we elucidate the framework for a sliding window analysis that allows us to do exactly that.

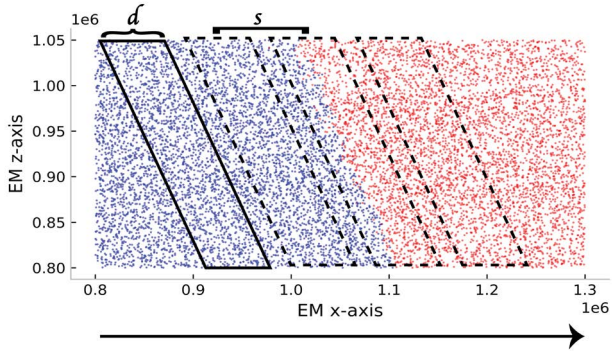

Figure 5: Depiction of sliding window used for analysis. Window is parameterized by width d and stride s.

Recall the window interpretation of the border described in Section 4.2. Using the fitted border (Section 5.2), we can create a window around the border plane where all points within distance $\frac{d}{2}$ are included, which yields a window of width d (Figure 5). Let us denote this window by $W_d(0)$. We can generalize this to an arbitrary window over the volume denoted by $W_d(T)$ where T denotes a T-unit horizontal translation from the border window $W_d(0)$.

Let the leftmost EM x-coordinate in the dataset be L and the rightmost be R. Denote the x-coordinate corresponding to the border by B. We define a sliding window on the V1-RL network as a set of $n + 1$ windows given by

$$\left\{W_d(-P), W_d(-P + \frac{2P}{n} * 1), \ldots, W_d(-P + \frac{2P}{n} * (n-1)), W_d(P)\right\}$$

where $P = \min(|B - R|, |B - L|)$. Now, what we have is a set of n equally spaced windows that span the network. For any given window $W_d(T)$, we will refer to the plane that divides it by $C_d(T)$. This means that we can denote the border plane fit in Section 5.2 by $C_d(0)$.

We can compute our metric of choice in each window to create a fine-grained characterization of the evolution of structure and function across the network. Also, note that if we choose an odd number of windows, $W_d(0)$ is included in the sliding window set.

5.5. *Border identifiability experiments*

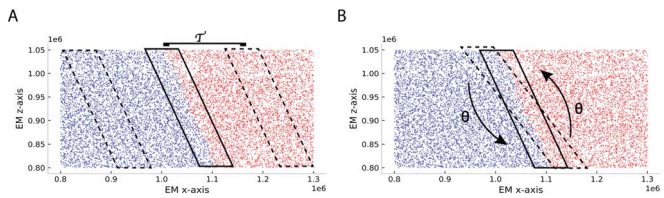

Figure 6: (A) Fake border parameterized by horizontal translation T. (B) Fake border parameterized by rotation angle θ.

5.5.1. *Construction of fake borders*

A persistent issue in constructing borders between visual areas is that there is a certain degree of arbitrariness. Works like Zhuang et al. (2017) have shown that the border fit using anatomical data does not necessarily align with the border fit using retinotopic data. To demonstrate the robustness of our connectomics-driven characterization of the border, we construct fake border cuts/windows and seek to understand if any of the fake borders could be construed as the true border. We denote a fake border cut as $F_C(T, \theta)$ and a fake border window as $F_W(T, \theta)$ in which our parameterization consists of a horizontal translation T from the true border and an angle θ which denotes the rotation from the true border (Figure 6A, Figure 6B).

5.5.2. *Fake border classifiers*

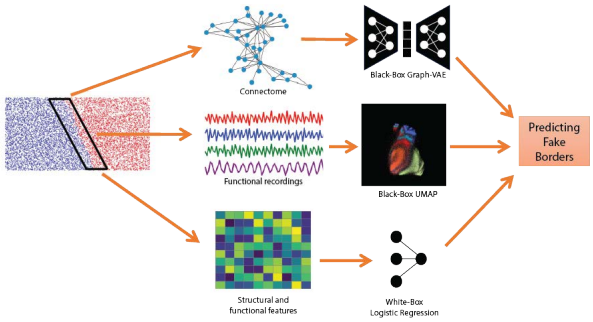

Figure 7: We extract the connectome and functional recordings of a particular window and use both white-box and black-box classifiers to perform classification across a set of fake borders.

Here, we describe our methodology for strengthening our characterization of the V1-RL network in particular by deriving feature representations of the neurons in the dataset and using those to build a predictive model on the retinotopic area label of the neuron. This accomplishes two things: one, it demonstrates that we can use connectomics data to predict on retinotopic labels. Two, it allows us to formalize a notion of border identifiability under a given feature set. Let us first elucidate the setup for the classification modeling. We consider two paradigms: the border as a region of notable rate of shift and as a region of notable magnitude.

With respect to notable rate of shift, we aim to address the third interpretation of the border window (Section 4.2) in which the border itself is not structurally or functionally unique. Rather, under the assumption that V1 and RL are structurally/functionally disparate (which is true as shown in Section 6.2), the border denotes a region of notable structural and functional shift since it represents the intersection between V1 and RL. Consider a window $F_W(T, \theta)$ whose cut is denoted by $F_C(T, \theta)$. Label the points left of $F_C(T, \theta)$ that reside in $F_W(T, \theta)$ as members of class, K_1 and analogously label the points on the right of $F_C(T, \theta)$ as K_2. A classifier that is fit on such a labeling performs best when neurons in K_1 and neurons in K_2 look functionally and structurally disparate; thus, if the border can be interpreted as the region of greatest shift in the network; we would expect the classifier to perform best when fit on $F_W(0, 0)$.

With respect to the border being a region of notable magnitude, we aim to address the first and second interpretations of the border window: the region surrounding the border plane looks distinct from both V1 and RL, either by serving as an information segregator or an information conduit. Here, let us consider a set of fake windows F of size n constructed using arbitrary and distinct (T, θ) pairs. Construct a label for each window $w \in F$ that is 1 for each neuron in w if $w = F_W(0, 0)$ and is 0 otherwise. We will call this labeling of windows L_0. Now, generate the other $n - 1$ permutations of L_0. For each generated permutation i, we denote L_i as the permutation of L_0 corresponding to the window $F_W(T_i, \theta_i)$ being labelled as 1. For each L_i, we can fit a classifier C_i that maps a feature representation of each neuron to its respective window label. If the border presents itself uniquely with respect to some feauturization (i.e. has notable magnitude), we would expect C_i to perform best on labelling L_0.

For the notable rate of shift classifier, we claim that the true border is identifiable if the classifier fit on the labeling induced by $F_C(0, 0)$ performs better than a classifier fit on the labeling induced by $F_C(T, \theta)$ for arbitrary (T, θ). And similarly, for the notable magnitude classifier, we claim that the true border is identifiable if the classifier fit on labeling L_0 performs better than a classifier fit on any other labeling L_i. This notion of border identifiability allows us to make claims about the robustness of our characterization of the border under the set of features we have chosen to represent it.

We explore two types of classifiers for both the rate of shift and magnitude paradigms: a white box classifier and black box classifier. The white box classifier refers to a feature representation of each cut/window using some subset of structural and functional metrics discussed in Section 5.3. Such a classifier enables us to motivate the border biologically by interpreting the features useful for its identifiability. The black box classifier refers to a feature representation of each cut/window using less interpretable methods described below. Such a classifier enables us to motivate the border more thoroughly than the white-box classifier by capturing features in the graph that we are unable to identify via the chosen set of structural and functional metrics.

To generate a black box structural representation of a given neuron n, we embed the graph using a variational-graph auto encoder (Kipf and Welling, 2016). A variational-graph auto encoder (VGAE) is designed to learn latent representations of neurons by leveraging higher-order information in the graph that may be difficult to identify manually. Since the VGAE model only supports undirected graphs, we make the edges in the V1-RL graph bidirectional. A VGAE can also leverage labels for the nodes in the graph; we make use of this by labelling the nodes with a one-hot encoding representing the neuron type (Section 5.3.3).

To generate a black box functional representation of a given neuron n, we reduce the dimensionality of the traces, since each trace contains 40000 measurements. To do so, we use UMAP (McInnes et al., 2018) with 10 components to create an expres-

sive representation of the neural traces. Here we lose interpretability along the temporal dimension, but this approach has proven to be effective for classification modeling on functional data (Kumar et al., 2021). One final note we make is regarding the classifier models used in each paradigm. For the white-box models, we fit a logistic regression and for the black-box models, we fit a multi-layer feed-forward neural network trained by gradient descent.

6. Results

6.1. Sliding window analysis reveals disparities between border regions

Here, we leverage the sliding window framework (Section 5.4) in order to compare the structural and functional profile of the V1-RL and RL-AL networks (Figure 8A), with emphasis on understanding how the border regions in each network present themselves. To investigate the structure of each network, we will compute neuron density and synaptic density. To investigate function, we will examine the correlation amongst the recorded neural traces. We note that for each sliding statistic we compute, we min-max normalize the sequence of values formed by the sliding window. This is done in order to explain away the potential differences in structure and function between the larger V1, RL and AL visual areas (which we will explore later in Section 6.2). Instead, we want our comparison to focus on the rate at which structure and function evolve across each network, understanding whether the structural and functional profile is the same or different, particularly at the border region.

There are two axes of comparison we want to emphasize in our analysis:

- How does the overall structural and functional profile of the V1-RL network compare to that of the RL-AL network?
- How do the border regions within each network present themselves? Do they lie in the same border regime (as detailed in Section 4.3) or different ones?

With respect to neuron density, we note similar trends in the structural profile of each network, in that the neuron density decreases as we move from left to right (Figure 8B). However, when focusing our attention on the borders, we do observe a decrease in neuron density at the V1-RL border that we do not observe at the RL-AL border.

When analyzing synaptic density, we note even more of a disparity between the two networks. Not only do the synaptic density profiles look entirely different, we also observe a stark difference between the two border regions as well. Namely, we find that the V1-RL border is a region of increased synaptic density, while the RL-AL border is not (Figure 8C). Rather, we observe that the RL-AL border presents itself as a region of notable shift with respect to synaptic density. While the underlying cause for this disparity between the V1-RL and RL-AL borders warrants further exploration, we hypothesize that it is due to the distinction between the V1 and RL/AL visual areas. V1, being the primary visual cortex, both projects to and receives projections from many higher order visual areas. In contrast, both RL and AL, being higher order visual areas themselves, do not necessarily need to communicate with each other to the same extent as V1 has to with other regions.

To compare the functional profile of the two networks, we take a look at the average pairwise correlation between the activations of neurons in the RL-AL and V1-RL regions. Analogous to synaptic density, we find that both the profile across the full network and the presentation of the borders within their respective networks look entirely disparate. With respect to the latter, we find that the activity at the RL-AL border is amongst the least correlated sub-regions of the RL-AL graph (Figure 8C). In contrast, the V1-RL border is the region with the most synchronous activity in the V1-RL network. This points to the functional disparity between the overall V1-RL and RL-AL network profiles, and more specifically their border regions.

By utilizing the sliding window approach on the V1-RL and RL-AL networks, we have shown two things. One, that the evolution of structure and function across each network looks distinct, pointing to the notion that changes in retinotopy are also captured by changes in connectomics. Secondly, we have established the concept of heterogeneity between different borders in the visual cortex. This notion can be tied back to the proposed border regimes detailed in Section 4.3. In particular, we find that the RL-AL border does not tend to present itself uniquely (within the scope of the RL-AL network) with respect to structure or function (Figure 8), suggesting that it lends itself most readily to the homogenous window/cut interpretation of a border in a visual network. In contrast, the V1-RL border does present itself uniquely with respect both structure or function, providing evidence supporting the idea that the V1-RL border lies in the connected cut/window regime, acting as a standalone region in the V1-RL network.

Given the distinctiveness of the V1-RL border in particular, the remainder of this work will focus on analyzing the V1 and RL visual areas in more detail. We leave further exploration as to why different borders in the visual cortex are structurally and functionally distinct to future work, primarily due to the lack of data present in the LM visual region in the current version of MICrONS.

6.2. Large scale differences between V1 and RL

In Section 6.1, we identified the V1-RL network and particularly the V1-RL border as a region of interest within the visual cortex. Before further characterizing the V1-RL border and understanding its role in the visual cortex, we first aim to understand the V1 and RL regions as a whole beyond the metrics explored in Section 6.1. In particular, we investigate each of the aforementioned metrics formalized in Section 5.3. We also note that from this analysis onwards, since we are restricting our scope to only the V1-RL portion of the network and are no longer constrained by the lack of data in the AL region, we expand our window of analysis to a larger portion of the V1-RL network (Figure 3).

In Section 6.1, we showed that neuron density decreased from V1 to RL in the smaller portion of the V1-RL network,

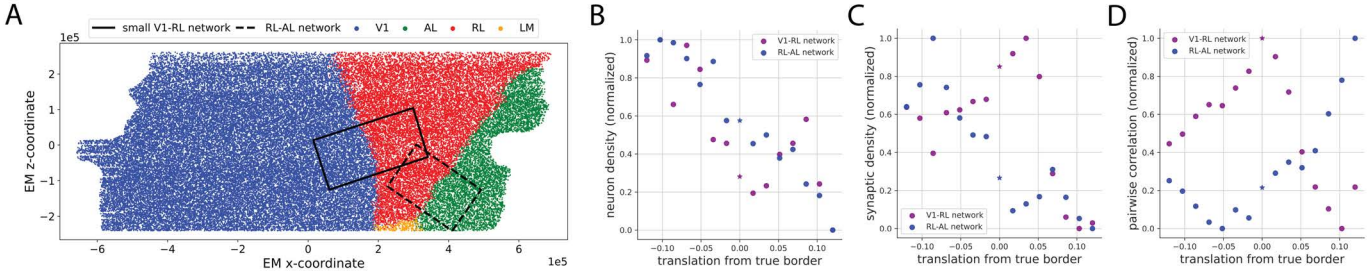

Figure 8: (A) Visual of windows used for comparison analysis in the context of the full MICrONS graph. The solid black rectangle bounds the V1-RL window considered for this analysis. The dashed black rectangle bounds the RL-AL window used for this analysis. (B, C, D) Note that the values in the plots shown in (B), (C) and (D) are min-max normalized to explain away differences between the V1-RL and RL-AL networks. The star in each plot denotes the border region between the two visual areas. The x-axis refers to the distance of the sliding window away from the true border in μm. (B) Neuron density sliding window analysis. (C) Synaptic density sliding window analysis. (D) Pairwise correlation of functional activations measured in sliding window analysis.

shown in Figure 8A. We observe the same trend when analyzing the larger portion of the V1-RL network (Figure 9A) in that neuron density presents itself uniquely in each region: in particular, we find that neurons are more densely packed in the V1 regional volume than in RL (Figure 9B), in line with previous, albeit less granular, findings of the visual cortex of primates (Collins et al., 2010).

We next address the disparity in neuron type composition between V1 and RL (Figure 9C, left). We find that L2/3 and L4 neurons occur more frequently in V1 and L5 and L6 more frequently in RL. This aligns with the sensory processing hierarchy that encapsulates the visual cortex. V1 is the first cortical area to receive visual inputs from the thalamus. Given its role as the entry point for visual information, it is intuitive that V1 would require a significant number of L4 neurons to process and relay the detailed visual inputs accurately. An analogous line of reasoning can be used for L2/3 neurons, which are also mediators of visual input.

Synaptic density is also higher in V1 than RL (Figure 9E, left), supporting the general trend observed in the sliding window analysis conducted on the smaller portion of the V1-RL network (Figure 8B) in Section 6.1. Given its role in detailed visual processing, it is reasonable that V1 has a higher synaptic density to support more intricate neural interactions. We provide further insight on connectivity patterns by analyzing the relationship between the distance of any given pair of neurons and the probability a connection exists between them (Figure 9F). Note that V1 tends to form shorter-range connections on average than RL, which is in line with the higher neuron density we observe in the region. However, we do recall from Section 6.1 that while distance between neurons is certainly correlated to connectivity, it cannot fully explain changes in synaptic density. This is notable at the V1-RL border, which shows elevated synaptic density despite not showing the same trend with respect to neuron density (Figure 8).

We move beyond lower order measures of structure and examine 2-cycle and 3-motif properties of the network. In line with the elevated synaptic density, we find that the density of 2-cycles is also higher in V1 (Figure 9E, middle). It is important to note that the prevalence of 2-cycles might be influenced by a variety of other factors. However, it is natural that a more connected graph tends to give rise to more 2-cycles in general. Another common measure of connectivity is 3-motifs. 3-motifs are specific patterns of connectivity involving three nodes in a graph (Figure 4). These motifs have important implications for the functioning and behavior of biological networks. Some motifs might create paths for efficient information propagation, while others can lead to information bottlenecks or localized processing. While the specific functionality of each individual motif is still an open question, differences in the distribution over all possible 3-motifs is a canonical measure of connectivity (Figure 9E, right). We can pinpoint the difference in distribution between V1 and RL to 3-motif number 4 and 3-motif number 6, the former of which is more prevalent in RL and the latter of which is more prevalent in V1. While each consists of 2-edges, we note that the directionality of these edges is different, which points to the disparity in information flow consistent with retinotopic differences between the two regions.

The final structural metric we analyze is microcircuit expansion, which, to the extent of our knowledge, has not been proposed before in the context of biological neural networks. We find that the expansion rate of the microcircuits we investigated tends to be higher on average in V1 (Figure 9D). This result amalgamates all the structural findings into a notion that is biologically interpretable: signals are propagated faster in V1 than RL.

We recap our structural findings by noting that under each metric, there exists a disparity between V1 and RL. Notably, this provides support for the notion that changes in retinotopy are correlated with changes in structure, tying together the information captured by connectomics and retinotopic data. Furthermore, we now have a nuanced characterization of interpretable attributes of the V1-RL network that are potentially responsible for inducing retinotopy.

We now shift the focus to functional metrics. Here, we aim to answer whether the retinotopically fit V1 and RL visual areas also posses differences with respect to the functional metrics computed in EM coordinate space. Note that these metrics on a subset of the full graph (Figure 10A) due to the current lack of a more extensive co-registration in MICrONS.

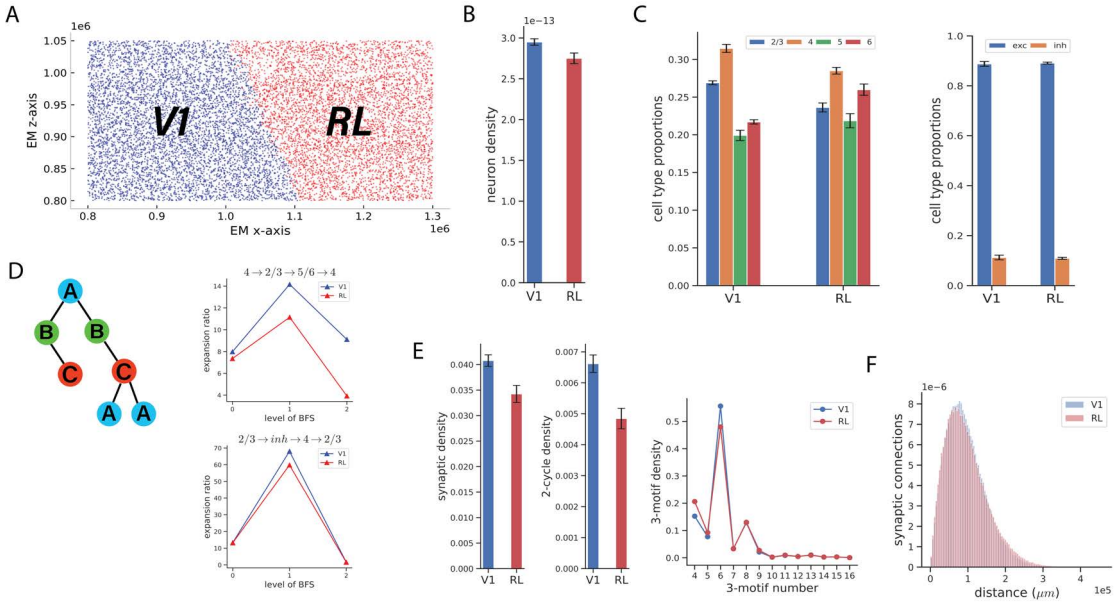

Figure 9: (A) Depiction of the dataset projected onto EM x-z plane. We estimate structural and function statistics over the entirety of V1 and RL. This is the same as the window shown in Figure 3, but plotted without the surrounding network. (B) Regional neuron density measured in units of $\frac{\text{neurons}}{\mu m^3}$. Two-sample t-test reveals higher neuron density in V1 (p-value < 0.01). (C, left) Neuron composition amongst excitatory types. Chi-squared goodness of fit test reveals difference between regions (p-value < 0.01). (C, right) Relative frequencies of inhibitory and excitatory neurons. Two-sample t-test shows no statistical difference between regions (p-value > 0.01). (D, left) Diagram showing an arbitrary example of the BFS tree generated by a microcircuit expansion. (D, top right) Microcircuit expansion rate for the excitatory microcircuit. Two-sample t-test conducted on each level of BFS reveals faster expansion in V1 (at least one level of BFS has p-value < 0.01) (D, bottom right) Microcircuit expansion rate for inhibitory microcircuit. Two-sample t-test conducted on each level of BFS faster expansion in V1 (at least one level has p-value < 0.01). (E, left) Regional synaptic density measured in units of $\frac{\text{total edges present}}{\text{total edges possible}}$. Two-sample t-test reveals higher synaptic density in V1 (p-value < 0.01). (E, middle) 2-cycle density measured in units of $\frac{\text{total 2-cycles present}}{\text{total 2-cycles possible}}$. Two-sample t-test reveals higher 2-cycle density in V1 (p-value < 0.01). (E, right) Distribution over 3-motifs in each region. The 3-motif numbering is described in Section 5.3.6. Chi-squared goodness of fit test shows difference between regions (p-value < 0.01). (F) Distance-connectivity distributions for V1 and RL. Two-sample KS-test reveals longer-range connections in RL (p-value < 0.01).

We find that the average activity of neurons is higher in V1 than RL (Figure 10B) as measured by the proportion of time a neuron fires, averaged over all neurons in the region. We offer potential justification for this result using the notion of visual attention. Visual attention involves selectively focusing on specific visual stimuli while filtering out others. The higher activity in V1 can reflect the prioritization of the brain to process the natural movie stimuli exposed to the mouse in greater detail. Due to the role of RL as a higher-order processor, it is unlikely that it necessitates the same computational capacity as V1. So, from an efficiency perspective, it is reasonable that RL neurons would fire less. Furthermore, recall that both neuron and synaptic density are lower in RL than V1. Under the higher-level notion that function tends to follow structure, the heightened sparsity of the graph in RL offers another interpretation as to why it has less activity.

We build on this notion of computational capacity by investigating the effective dimensionality of neural activity measured in state space. We find that the activity in RL is inherently lower dimensional than the activity in V1 (Figure 10C). This provides evidence of simpler neural dynamics in the RL region, which again aligns with the functional and structural sparsity elucidated above. The general idea that RL has lower dimensional representations than V1 offers a statistical justification to the sensory processing hierarchy. Essentially, V1 acts as a subnetwork responsible for generating more complex representations of visual input and RL abstracts away the nuances of data represented in high dimensional space by learning lower dimensional representations. This intuition is largely similar to what we see in convolutional architectures in which earlier layers are tasked with extracting fine-grained details from the input image and later layers filter these high-dimensional representations to generate a more robust, lower dimensional representation useful for prediction.

Unsurprisingly, we also observe a disparity in average pairwise correlation, aligning with the trends observed in the smaller V1-RL network analyzed in the previous section (Figure 8D). In particular, we see that activity is more pairwise correlated in V1 (Figure 10D). This also aligns with the elevated synaptic density, which potentially enables the region to synchronize its neuron firings. Also, while the higher pairwise correlation and higher dimensional activity may appear contradictory, there is no strict relationship between the two measurements. In fact, the relationship between them in the context of biological networks is not well-understood (Umakantha et al., 2021). This points to the distinction between pairwise measure and group-wise measures of functional activity and shows that while V1 may exhibit more complicated population dynamics, it still possesses more synchrony on the pairwise level than RL.

To this point, we have shown that there exist structural and

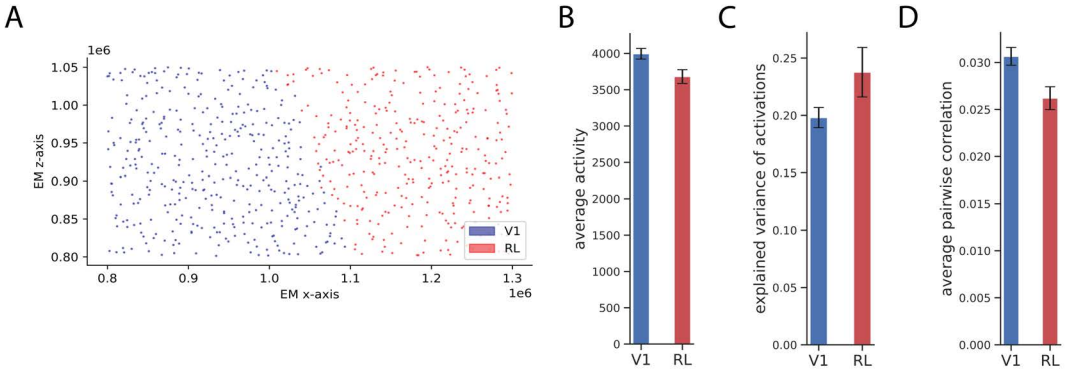

Figure 10: (A) Depiction of the co-registered dataset. Note how much sparser this is than the full graph in Figure 9A. All functional metrics are computed on this subgraph. (B) Average activity in V1 and RL. Two-sample t-test reveals higher activity level in V1 than RL (p-value < 0.01). (C) Effective dimensionality of activity in V1 and RL. Two-sample t-test reveals more complex activity in V1 than RL (p-value < 0.01). (D) Average pairwise correlation in V1 and RL. Two-sample t-test reveals more synchronous neuron firings in V1 than RL (p-value < 0.01).

functional differences between V1 and RL under a wide array of metrics as measured via connectomics. Furthermore, each metric is interpretable in the context of a biological neural network and offers a potential reason for the retinotopic differences between the two regions. Overall, we have provided support for the claim that retinotopy correlates with structure and function individually, which does give merit to retinotopy as a means of delineating regions in the mouse visual cortex (although we will demonstrate its drawbacks in Section 6.3). Furthermore, although not the focus of our work, the observed correlations between structural metrics being higher in V1 than RL (i.e. neuron and synaptic density) and the corresponding functional metrics also being higher (i.e. average activity and synchrony) provides insight into the potential role structure plays in inciting functional phenomena.

6.3. Granular analysis at interface of V1 and RL reveals border region

While providing valuable insights, the scope of the results discussed in the previous section are limited in that metrics were computed over the entirety of V1 and RL, which precludes us from making claims about the potentially more granular structural and functional profile over the V1-RL network. This was precisely the motivation behind the sliding window analysis shown in Section 6.1, which as discussed earlier, was computed on a smaller portion of the V1-RL network to ensure a fair mode of comparison to the RL-AL network. Here, we extend the sliding window framework to the larger portion of the network by again analyzing the evolution of various structural and functional metrics over the course of the network. Note that the shape of the sliding window is always constructed such that the edges are parallel to the border (Figure 11A), which here yields a different sliding window shape than was used in the analysis of the small portion of the V1-RL network examined in Section 6.1. This is because we opted against orienting the window of analysis such that the border is parallel to the edges of the larger window (Figure 3) in order to capture as much data as possible while avoiding edge effects (refer to Section 5.1 for more details). We use the opportunity afforded to us by not only the different/larger portion of the network, but also by the change in the orientation of the border axis, to demonstrate the robustness of our findings from Section 6.1 as well as explore them in further detail.

Again, we find that synaptic density presents itself uniquely at the V1-RL border: in particular, the density of the subgraph formed by the region surrounding the border is by far the highest amongst all regions in the network (Figure 11C). We also note that we analyzed the synaptic densities of the network on a per-layer basis and found that the significant increase at the border is consistent across all layers (results not shown). This, along with a similar trend observed in the smaller portion of the V1-RL network (results not shown), offers support for the robustness of our characterization of the V1-RL border as a more connected region of the visual cortex.

We analogously analyze neuron density in the sliding window paradigm within this larger portion of the V1-RL network. We observe that around the border, the neuron density decreases significantly (Figure 11B, left), a trend even more pronounced than what we observed in Section 6.1. To understand this further, we analyzed the neuron density of each neuron type and found that the dip at the border was only observable amongst L4 neurons (Figure 11B, right). This is a notable finding as it further provides support for window interpretation of the border, as there is no well-defined notion of neuron density under the cut interpretation. Another interesting trend, is the inversion of the relationship between synaptic density and neuron density that we characterized in V1 and RL (Section 6.3) which we find only occurs at the V1-RL border. Namely, we found that synaptic density and neuron density are positively correlated in each retinotopic region, which makes sense under the idea that connectivity is roughly some function of distance between neurons. However, at the border, the opposite is true: synaptic density is high in spite of the neuron density being low. While this phenomenon warrants further exploration, we hypothesize

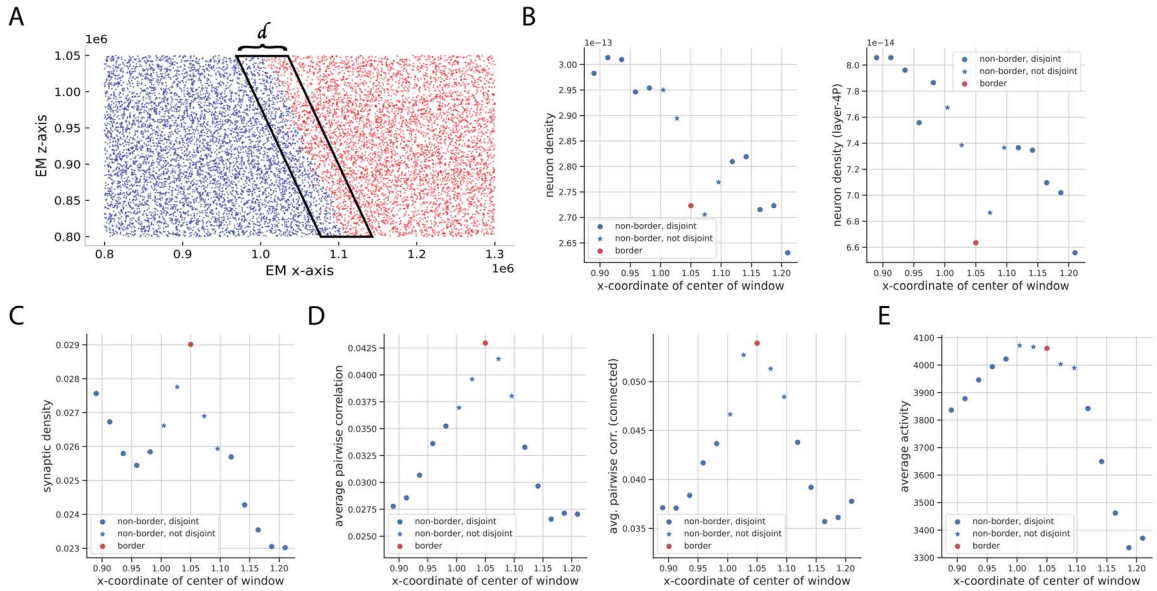

Figure 11: (A) Depiction of sliding window used in analyses that follow. (B, C, D, E) The legend in subsequent plots can be interpreted as follows: non-border, disjoint refers to windows that are disjoint from the border window and non-border, non-disjoint refers to windows that are not disjoint from the border window. Hence, the starred points represent a hybrid of the non-border and border subgraph. (B, left) Sliding window analysis for neuron density. (B, right) Sliding window analysis for L4 neuron density. We see that L4 neuron density decreases even more than the overall neuron density at the border. It is the only excitatory neuron type that decreases in this fashion at the border. (C) Sliding window analysis performed for synaptic density. (D, left) Sliding window analysis performed for average pairwise correlation computed amongst all pairs of neurons. (D, right) Sliding window analysis performed for average pairwise correlation only amongst connected neurons. (E) Sliding window analysis for average activity.

that it can be partially attributed to the border's role as a message passer in the V1-RL network (this idea is made concrete in Section 6.7). In particular, since L4 neurons are sparse near the border, the firings are dominated by information passed internally through the network (i.e. from V1 and RL) as opposed to externally from the thalamus.

One final observation we make is that both neuron density and synaptic density are quite heterogeneous across the V1-RL network (Figure 11B, Figure 11C). We are able to observe this at a level of granularity that was previously unobtainable due to the lack of extensive connectomics datasets like MICrONS. This points to the promise these datasets have with respect to the larger goal of extensively profiling the brain.

We next examine some functional metrics. With respect to average activity, we also see that it is heterogeneous across the network, with its peak at the border (Figure 11E). In fact, we see that average activity follows a similar trend to synaptic density. Recall in Section 6.2, we observed a high-level relationship between connectivity and activity when aggregated over the entirety of V1 and RL. However, what we observe here is evidence of a more granular relationship: namely, both synaptic density and average activity are highest at the border and lowest at the rightmost point in RL. While the correlation between these two metrics appears to be weaker in V1, this still provides support for the notion that neither structure nor function is homogenous across the V1-RL network and furthermore, both tend to shift together.

We further substantiate these claims with an analogous analysis on pairwise correlation. We find that the average pairwise correlation is highest at the border (Figure 11D, left), again in line with the same analysis performed on the smaller V1-RL network (Section 6.1). This provides yet another measure that supports the notion that the border region lies in the connected regime. And, similar to average activity, the change in pairwise correlation resembles the change in synaptic density over the network. But, note that the pairwise correlation we have discussed is measured amongst all pairs of neurons. We also aim to see what happens if we consider correlation only amongst connected neurons. In doing so, we find that while the pairwise correlation rises across all windows, the shift across the network in relative units remains the same (Figure 11D, right). One, this aligns with the idea in functional connectomics that neurons wired together tend to fire together. But two, it also shows that the uniqueness of the border with respect to more correlated activity does not rely upon first-order connections between neurons.

At a higher-level, we can tie many of these findings back to the larger goal of demonstrating instances where retinotopic and architectonic maps can fail to capture aspects of network structure and function that functional connectomics can. In particular, both mappings operate under the assumption that borders between visual areas delineate regions of notable shift in structure and/or function. But based on the analyses conducted here, this is not true of all borders. In particular, with respect to

neuron density, synaptic density, average activity and correlatedness of activity, the V1-RL border is not a region of a notable shift, but rather one of notable magnitude (which we will explore more formally in Section 6.6). Despite the retinotopic and architectonic borders being fit to best divide the network into two regions, they cannot fully explain the functional and structural variation we observe across the full network to the extent connectomics data can.

6.4. Structural and functional trends under the window interpretation of the border are localized phenomena

While we have provided thorough evidence for the border lying in the connected regime, we have not yet explicitly addressed whether it lies in the connected cut or connected window regime. Recall that the results with respect to neuron density and average activity already point to the window interpretation of the border as the more reasonable model. But, here, we go one step further and explicitly distinguish between neurons connected across the cut and on the same side of the cut (Figure 12D). We find that under synaptic density, the border is equally identifiable (Figure 12E) using cross-cut and same side connections. As expected, the synaptic density across the cut is shifted down by an amount consistent across all windows due to the larger distance between cross-cut connections than same side connections. However, in relative units the relationship between V1 windows, RL windows and border window looks the same. We notice the same effect when measuring pairwise correlation (Figure 12F). These findings reinforce the window interpretation of the border by demonstrating that the identifiability of the border is not solely attributable to structural and functional relationships formed across it; rather, the entire region around the border plane is unique.

We take a step further in characterizing the connected border window by presenting an alternative formulation of the sliding window. In particular, consider the notion of locality. Recall that in the sliding window analysis presented in Figure 10, we consider only connections between neurons contained within the window. An alternative formulation is to consider neurons in the border window, but allow them to form connections with neurons outside the window that are still contained within a larger window (Figure 12A). We recompute both synaptic density and pairwise correlation under this non-local formulation and find that under both metrics, the distinct peak at the border we observe in the local formulation dissipates (Figures 12B, 12C). This provides evidence supporting the higher-level notion that the structural and functional uniqueness of the V1-RL border is a localized phenomenon.

With the sliding window framework, we have been able to interpret the V1-RL network by analyzing various structural and functional metrics in sub-regions of the graph. To summarize, we have shown the following:

1. Structure and function are heterogeneous in both V1 and RL
2. Structure and function cross-correlate on a granular level
3. The border region is structurally and functionally unique; notably, it is the most connected and most pairwise correlated region in the graph
4. The window interpretation of the border is a more accurate model than the cut interpretation

6.5. True V1-RL border is unidentifiable when interpreted as a region of notable shift

To this point, we have characterized the V1-RL border as a standalone region, separate from V1 and RL, that has higher connectivity and more synchronized activity. However, this characterization is not a formal measure of border identifiability (Section 5.5). Here, we to propose a classification framework that can identify the true border amongst fake borders to demonstrate the robustness of our characterization. First, we will examine the rate of shift paradigm and show that the border is unidentifiable under this framework.

Recall that the rate of shift characterization of the border is most tied to the border lying in the homogenous regime which lends itself most readily to existing interpretations of borders under retinotopic maps. In this regime, the border window serves simply as the intersection between V1 and RL. Since these two regions are functionally and structurally distinct (Section 6.2), we would expect the border to denote the region in which structure and function change the most. Yet, we know from earlier analyses (Section 6.3) that this interpretation of the border is unlikely, as both structural and functional statistics measured at the border do not present themselves as a hybrid of V1 and RL.

To demonstrate the inefficacy of interpreting the border as a region of notable shift, we construct fake cuts by horizontal translation (Figure 13A). We consider the classification problem detailed in Section 5.5 in which we use the same label for neurons that lie on the same side of the cut. The structural white box classifier employs a two-feature representation of each neuron by computing the synaptic density and neuron density of a subgraph constructed by taking a fixed size neighborhood around each point. Under this featurization, we find that the classifier performs worse at the true border relative to many of the fake borders (Figure 13B). If we were to use performance under this classifier as a means of identifying the true border, we would conclude that many of the fake windows are more reasonably construed as the true border than the true border itself. This result is unsurprising in lieu of the sliding window neuron and synaptic density findings discussed in Section 6.3. However, it is still possible that there exists a structural statistic that we have not computed in which the border appears identifiable under this classification paradigm. To address this, we construct a black box classifier in which the neuron representations in the graph are learned by a VGAE which is capable of capturing more nuanced aspects of network structure. Even here, we find that the true border is unidentifiable (Figure 13C).

Next, we aim to see if the true border is also unidentifiable when using a functional feature representation of each neuron. Again, we construct fake cuts by horizontal translation. We generate a single-feature representation of each neuron by measuring the pairwise correlation of neurons in a subgraph constructed by taking a fixed size neighborhood around each point. Even under a functional feature representation, we find that the

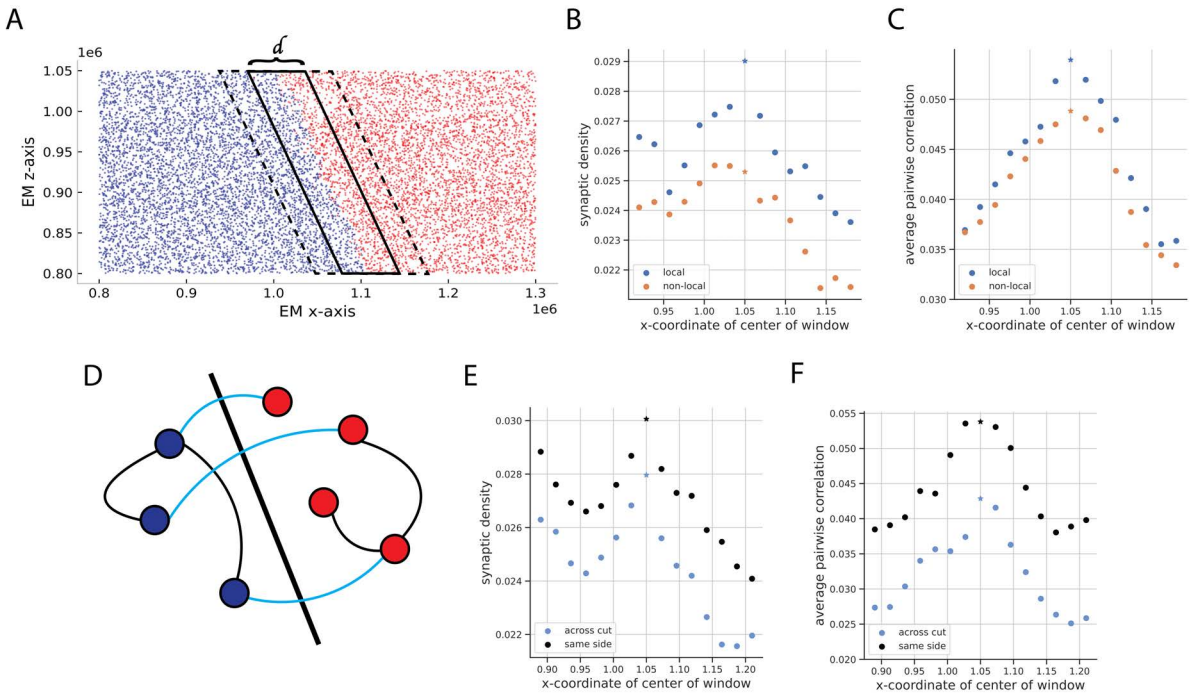

Figure 12: (A) Depiction of a secondary window that allows us to measure the locality/non-locality of the computed metrics. In this analysis, we examine all connections formed within the dashed-line window, such that at least one neuron lies in the solid-line window. In the plots that follow, the star denotes the border window. (B) Plot comparing local and non-local synaptic density. Local refers to sliding window analysis shown in Figure 11 in which only connections formed within the solid-line window are analyzed. (C) Locality analysis for average pairwise correlation. (D) Diagram showing the distinction between connections formed on the same side of the cut (black) and across the cut (light blue). (E) Comparison of the synaptic density of the subgraph formed by considering neurons across the cut versus the same side of the cut. (F) Comparison of the average pairwise correlation of the subgraph formed by considering neurons across the cut versus the same side of the cut.

classifier does not perform particularly well at the true border (Figure 13D). Analogous to the structural white box classifier, this result is unsurprising given the sliding window pairwise correlation analysis (Figure 11). But, here we take a step further, and also show that a more expressive functional featurization is equally ineffective at identifying the border. Using UMAP on the neural traces yields better performance at fake borders than the true one (Figure 13E).

Recall that prior to this classification analysis, we knew that the notable rate of shift interpretation of the border was unlikely to be true given the observations we made in the sliding window analysis. However, here we made this notion concrete by showing that even when using more complex and expressive representations of neurons, we are still incapable of identifying the V1-RL border.

6.6. True V1-RL border is identifiable when interpreted as a region of notable magnitude

We have yet to propose a classification framework that robustly identifies the true V1-RL border. Here, we try a different approach: interpreting the border as a region of notable magnitude. This framework lends itself well to the connected window interpretation, as it selects for regions that are unique with respect to some feature representation. As we did with the rate of shift classifier, we construct fake cuts by horizontal translation (Figure 14A). However, here we label the neurons differently. Namely, we consider n different labelings in which each has a different one of the fake borders labeled as the true border – analogous to a permutation test (refer to Section 5.5 for details).

Again, we will first consider a structural white box and black box classifier. As described earlier, for the white box classifier, we generate a two-feature representation of each neuron that contains a measure of local synaptic and neuron density. We fit a classifier to each permutation of the labeling and find that the true border is most identifiable (Figure 14B). This demonstrates that the true V1-RL border is distinct from the fake borders under this feature representation in a way that is not observed when any of the fake borders are treated as the true one. Yet, this remains a simple structural feature representation of the neurons in the graph. So, we next ask whether the border becomes more identifiable when we leverage neuron embeddings generated by a VGAE. Fitting the notable magnitude classifier using the more complex feature representation shows that it does (Figure 14C). This supports the notion that higher order structures that are intractable to compute using our rudimentary analyses alone also present themselves uniquely at the border region.

We analogously construct a functional white box and black

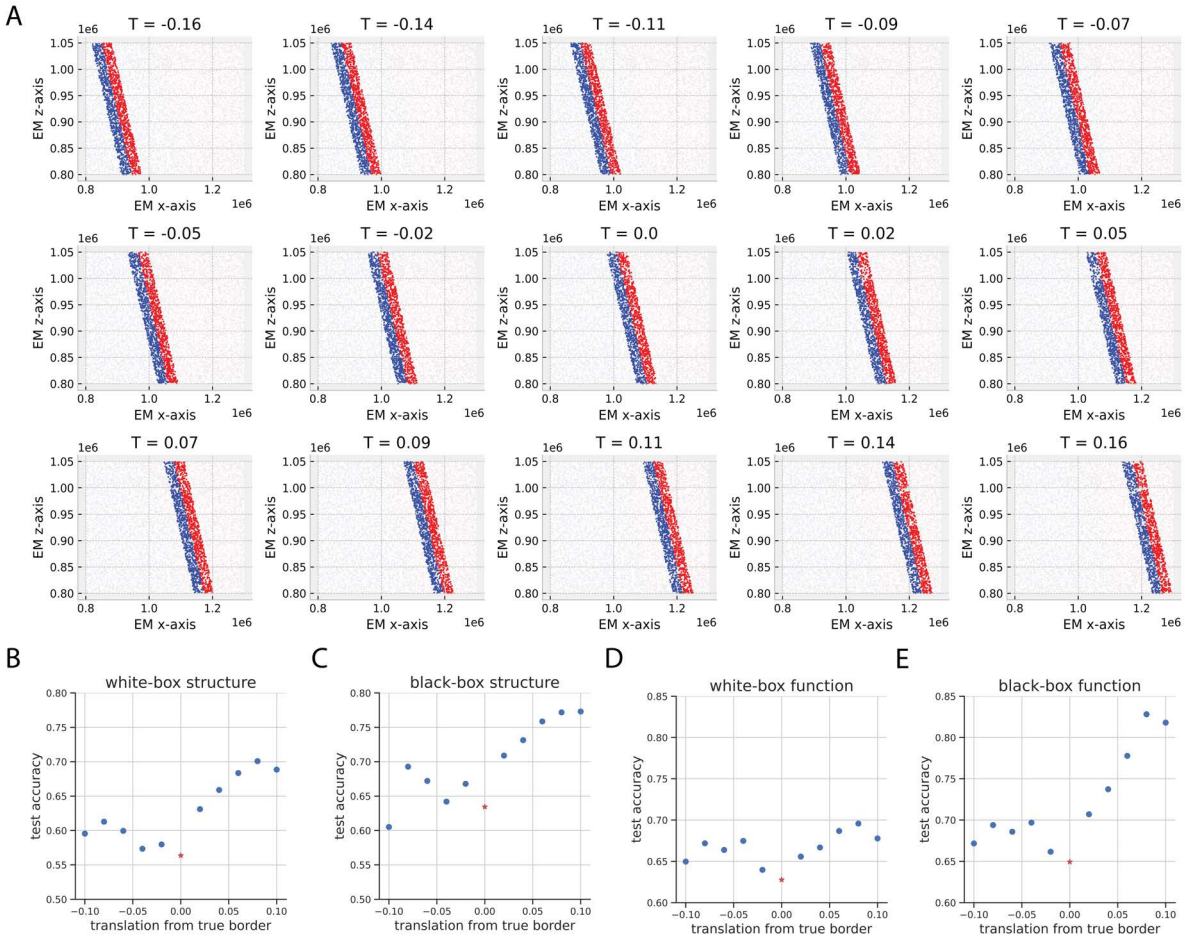

Figure 13: (A) Depiction of arbitrary fake border windows constructed via horizontal translation. To perform classification for a given translation, we label the neurons in blue with a 1 and the neuron in red with a 0. Recall that we refer to this classification paradigm as notable rate of shift. The red dot in plots that follow denotes the true border. (B) Classifiers fit on a two-feature representation of each neuron using synaptic and neuron density. (C) Classifiers fit on representation of neuron learned by VGAE. (D) Classifiers fit on a single-feature representation of each neuron using average pairwise correlation. (E) Classifiers fit on neuron representations generated via UMAP.

box classifier. For the white box classifier, we consider a single-feature representation of a neuron using the local average pairwise correlation. We also find that under this simple functional featurization, the border is identifiable (Figure 14D). As we did with structure, we next aim to construct a more complex functional featurization in order to further distinguish the true border from its fake translations. We find that the activations pre-processed using UMAP accomplishes exactly that (Figure 14E).

So we have shown that under the notable magnitude paradigm, the border is identifiable with respect to multiple feauturizations of the neurons in the graph. However, we have only shown that the true border is identifiable amongst fake borders constructed by horizontal translation. Given what we know about the structural and functional profile of the V1-RL network being quite heterogeneous (Section 6.3), this wasn't particularly surprising. Here, we go one step further and also show that the border is identifiable against fake borders constructed by rotating the true border around its center point. Before developing a classification framework that addresses this question, we will revisit the sliding window framework for analysis. Recall that the shape of the window is constructed such that the bounds of the window are parallel to the border axis. We can analogously construct sliding windows for fake borders generated by rotation and compute the sliding statistics under the newly shaped borders (Figure 15A).

With respect to structure, we compute both the sliding synaptic density and neuron density across rotations of the border window. We find that rotations of the border to the left ($\theta < 0$) no longer have a distinct synaptic density at the border. However, for rotations to the right, there exist some rotated borders in which the fake border also presents itself uniquely with respect to synaptic density (Figure 15B). While one could argue that the synaptic density at the border window for rotation $\theta = 0$ presents itself most uniquely across all generated rotations, it is not a strong argument for border identifiability. So, we next aim to construct a more robust set of identifiability criteria.

To do so, we turn to sliding neuron density across rotations of the border. Specifically, we look at L4 neuron density since we know that L4 density is responsible for the decrease in overall

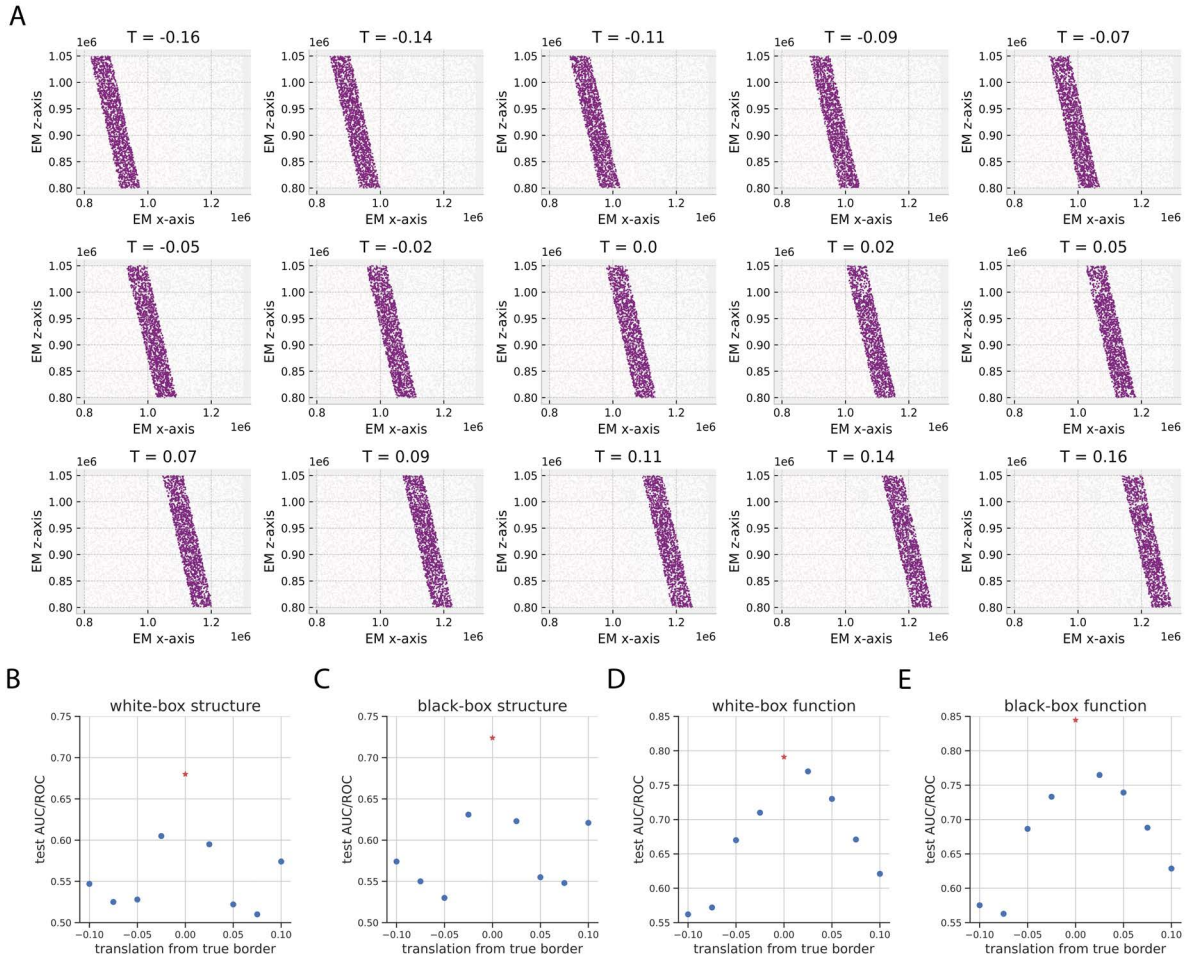

Figure 14: (A) Depiction of arbitrary fake borders constructed via horizontal translation. To perform classification for a given translation T, we label the purple window with a 1 and other windows with a 0. Recall that we refer to this classification paradigm as notable magnitude. The red dot in plots that follow denotes the true border. (B) Classifiers fit on a two-feature representation of each neuron using synaptic and neuron density. (C) Classifiers fit on representation of neuron learned by VGAE. (D) Classifiers fit on a single-feature representation of each neuron using average pairwise correlation. (E) Classifiers fit on neuron representations generated via UMAP.

neuron density observed at the border (Figure 11B, left). Interestingly, we find that rotations of the border to the right ($\theta < 0$) no longer have a distinct L4 neuron density at the border (Figure 15C). Recall that rotations of the border to the left no longer resemble the true border with respect to sliding neuron density. So, using both synaptic density and neuron density, the true border is identifiable against rotation. We formalize the notion of border identifiability against rotation by constructing a classifier that leverages these structural differences. Neurons are featurized by taking a neighborhood around them and computing the synaptic and L4 neuron density. For each rotation, we construct a classifier under the labeling L_0 in which the true border is given a label of 1 and the rotated borders are given a label of 0. As expected, we find that the true border is more identifiable than its fake counterparts constructed by rotation (Figure 15F).

Next, we perform an analogous form of analysis using functional data. We first compute the sliding pairwise correlation across various rotations of the border window. The correlation at the true border appears to present itself most uniquely relative to the fake border rotations, however there exist a few rotations for which pairwise correlation is also highest at the border (Figure 15D). So, we analyze a potentially more robust metric in Granger causality (Section 5.3). Granger causality moves beyond pairwise correlation by capturing the temporal dynamics of network activity. Again, we perform a sliding window analysis using the fake rotated borders and find there is less ambiguity with respect to the identifiability of the true border than we observed analyzing pairwise correlation (Figure 15E). To demonstrate the identifiability of the true border using functional data, we featurize neurons by computing their neighborhood pairwise correlation and Granger causality and perform classification under the same paradigm discussed for the structural rotation classifier. In doing so, we find that the true border is indeed more identifiable than the rotated fake borders (Figure 15G).

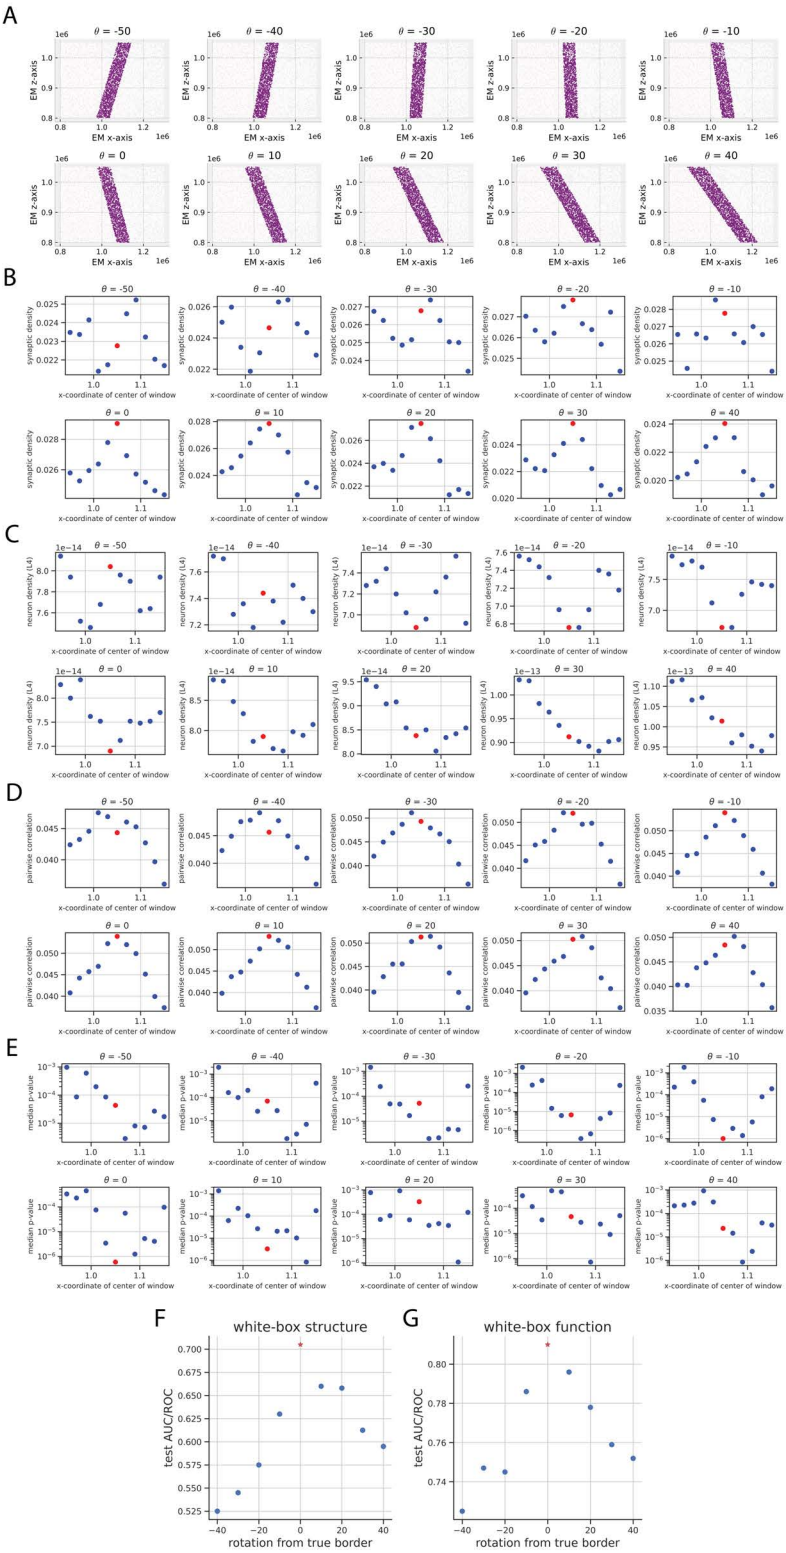

Figure 15: (A) Depiction of arbitrary fake rotations of the true border. We use these rotations in both the sliding window and classification frameworks discussed in previous analyses. Red dot denotes true border in all the following plots. (B) Sliding window analysis measuring synaptic density across various border rotations. (C) Sliding window analysis measuring L4 neuron density. (D) Sliding window analysis measuring average pairwise correlation amongst connected neurons. (E) Sliding window analysis measuring median Granger causality (using p-value as a proxy). (F) Structural notable magnitude classifier fit across multiple rotations of border. (G) Functional notable magnitude classifier fit across multiple rotations of border.

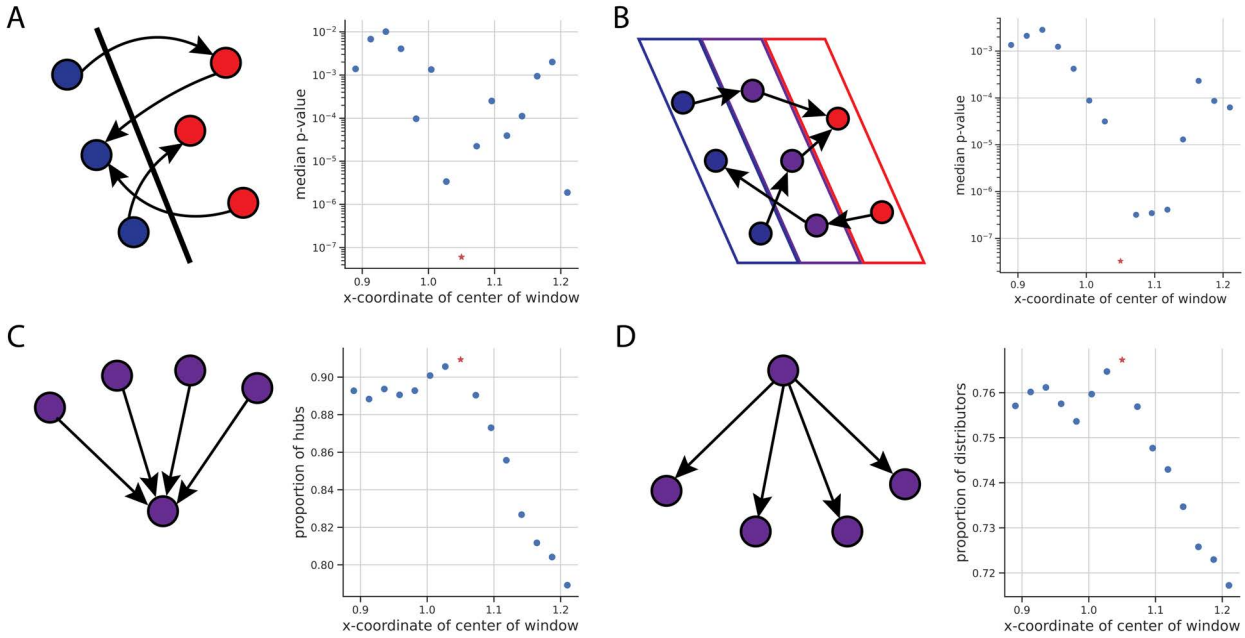

Figure 16: (A, left) Depiction of cross-cut motif. Median G-causality across all synapses crossing the cut is measured (using p-value as a proxy). (A, right) Cross-cut motif shows that information flow along synapses crossing the true border cut is the highest in the network. (B, left) Depiction of message passing motif. Median G-causality going from the blue neuron to the red neuron or from the red neuron to the blue neuron is measured depending on direction of connections. (B, right) We find that the message-passing motif is also highest at the border, indicating the role neurons in the border region play in carrying information between V1 and RL. (C, left) Depiction of hub motif (there being 4 pre-synaptic neurons in diagram is arbitrary). (C, right) The proportion of neurons that are hubs is highest near the border, pointing to the border's role as an information aggregator. (D, left) Depiction of distributor motif (there being 4 post-synaptic neurons in diagram is arbitrary). (D, right) The proportion of neurons that are distributors is highest near the border, pointing to the border's role as an information disseminator.

6.7. Motifs in the V1-RL border region

Up to now, we have shown that formulating the border as a region of notable magnitude is more justified than its interpretation as a region of notable shift, offering merit to the use of connectomics-driven frameworks for border delineation in lieu of retinotopic ones. However, we have yet to develop intuition as to why exactly the border is both structurally and functionally unique relative to the rest of the V1-RL network. More specifically, we aim to identify various structural and functional motifs that arise at the border as a means of explaining its role in the visual cortex.

The first motif we investigate captures the extent to which neurons communicate across the border cut (Figure 16A, left). This motif, which we term the cross-cut motif, is computed by measuring the Granger causality (a measure of information flow) of neurons connected across the cut for the various cuts constructed using the sliding window framework. We find that the true border contains the highest across-cut Granger causality (Figure 16A, right). The prevalence of the cross-cut motif at the border region points to its relevance as a conduit of information flow between V1 and RL.

We further investigate this concept of message passing at the border by next looking at a message passing motif, which is a slight variant of the cross-cut motif. The message passing motif is distinct from the cross-cut motif in that as opposed to measuring interactions constrained to the window, we now allow connections from outside the window. This is analogous to the non-locality analysis we discussed earlier (Figure 12). Referring to the diagram above, this motif is computed by measuring the Granger causality of the red neuron onto the blue neuron if the connections go from left to right and vice versa if they go from right to left (Figure 16B, left). The purple neuron can be viewed as acting as a message passer between non-border neurons. Our analysis shows that message passing is significantly higher at the border than other areas in the network (Figure 16B, right).

Next, we analyze a related, yet distinct set of border motifs: hubs and distributors. In the context of a biological network, a hub motif refers to a specific pattern or subgraph structure that represents a central node (the hub) connected to several pre-synaptic nodes (Figure 16C, left). Here, we interpret this motif as a unit of information aggregation. For our purposes, we arbitrarily define a hub motif to be a neuron that has 25 pre-synaptic neighbors. Similarly, a distributor motif refers to a central node (the distributor) connected to several post-synaptic nodes (Figure 16D, left). Here, we interpret this motif as a unit of information dissemination. We arbitrarily define a distributor motif to be a neuron that has 6 post-synaptic neighbors.

We first examine the proportion of neurons that are hubs under the sliding window framework. In doing so, we find that hubs occur most frequently at the border region (Figure 16C, right). This supports the notion of the border acting as a region that is responsible for aggregating information from multiple different sources, and also offers a reason as to why the border

is a highly connected subgraph of the V1-RL network. In addition, we also find that distributors occur most frequently at the border region (Figure 16D, right) which lends itself well to the idea of the border acting not only as a region of information aggregation, but also information dissemination.

Tying everything together, we now have a coherent, although likely not exhaustive, set of explanations as to why the border presents itself uniquely with respect to both structure and function. The cross-cut and message passing motifs both act as proxies for the strength of information flow along a synapse, each of which is highest at the border. This idea of the border acting as a bridge between V1 and RL is natural: intuitively, V1 and RL, each of which functionally specialize, likely benefit in their ability to process information when supervised by the other. Furthermore, the notion of message passing goes hand in hand with the ability of the border to efficiently and effectively aggregate and disperse information. By analyzing the hub and distributor motifs, we demonstrated that, structurally, the border region has the highest frequency of neurons that act to both absorb information from many sources and spread it to many destinations.

7. Conclusion

Understanding the mouse visual cortex and its retinotopically-induced visual regions is an important step towards improving our understanding of the brain. In this work, we propose a statistical framework for analyzing neural networks in connectomics datasets like MICrONS that map out both the structural graph and a subset of its functional activations. By placing our focus on the borders between visual areas, our findings reveal two things: one that not all borders are the same, and two that the V1-RL border in particular serves as a standalone piece of the visual cortex. Further characterization of the V1-RL border demonstrated its role as a message passing interface within the larger V1-RL network.

There are a couple interesting directions we hope to pursue as future work. One, we want to develop a more extensive profiling of the visual cortex that includes analysis of its other regions and respective borders using the proposed sliding window approach, beyond what is explored in Section 6.1. We aim to do this under the larger vision of understanding the nuances of all border regions in the visual cortex, and more generally the brain. Unfortunately, since this relies upon a more extensive profiling of the visual cortex than is afforded by MICrONS in its current form, we hope to extend our work using later versions of the dataset upon its release. Two, we hope to translate these results into an artificial neural network by inducing an interface for communication between subnetworks that functionally specialize in order to emulate the role the V1-RL border region plays in the mouse visual cortex. While artificial networks are typically not designed with such structural nuances in mind, exploring the role structure and more specifically connectivity plays in modulating network function is certainly a promising avenue to explore.

References

Abeles, M., 2009. Synfire chains. Scholarpedia 4, 1441.

Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. Neuron 76, 695–711.

Behrens, T.E., Sporns, O., 2012. Human connectomics. Current opinion in neurobiology 22, 144–153.

Bibollet-Bahena, O., Tissier, S., Ho-Tran, S., Rojewski, A., Casanova, C., 2023. Enriched environment exposure during development positively impacts the structure and function of the visual cortex in mice. Sci Rep 13, 7020.

Bock, D.D., Lee, W.C.A., Kerlin, A.M., Andermann, M.L., Hood, G., Wetzel, A.W., Yurgenson, S., Soucy, E.R., Kim, H.S., Reid, R.C., 2011. Network anatomy and in vivo physiology of visual cortical neurons. Nature 471, 177–182.

Collins, C.E., Airey, D.C., Young, N.A., Leitch, D.B., Kaas, J.H., 2010. Neuron densities vary across and within cortical areas in primates. Proc Natl Acad Sci U S A 107, 15927–15932.

Consortium, M., Bae, J.A., Baptiste, M., Bishop, C.A., Bodor, A.L., Brittain, D., Buchanan, J., Bumbarger, D.J., Castro, M.A., Celii, B., et al., 2021. Functional connectomics spanning multiple areas of mouse visual cortex. BioRxiv , 2021–07.

Dorkenwald, S., Matsliah, A., Sterling, A.R., Schlegel, P., Yu, S.C., McKellar, C.E., Lin, A., Costa, M., Eichler, K., Yin, Y., Silversmith, W., Schneider-Mizell, C., Jordan, C.S., Brittain, D., Halageri, A., Kuehner, K., Ogedengbe, O., Morey, R., Gager, J., Kruk, K., Perlman, E., Yang, R., Deutsch, D., Bland, D., Sorek, M., Lu, R., Macrina, T., Lee, K., Bae, J.A., Mu, S., Nehoran, B., Mitchell, E., Popovych, S., Wu, J., Jia, Z., Castro, M., Kemnitz, N., Ih, D., Bates, A.S., Eckstein, N., Funke, J., Collman, F., Bock, D.D., Jefferis, G.S.X.E., Seung, H.S., Murthy, M., 2023. Neuronal wiring diagram of an adult brain. bioRxiv .

Fize, D., Vanduffel, W., Nelissen, K., Denys, K., d'Hotel, C.C., Faugeras, O., Orban, G.A., 2003. The retinotopic organization of primate dorsal v4 and surrounding areas: a functional magnetic resonance imaging study in awake monkeys. Journal of neuroscience 23, 7395–7406.

Fu, H., Niu, Z., Zhang, C., Ma, J., Chen, J., 2016. Visual cortex inspired cnn model for feature construction in text analysis. Frontiers in computational neuroscience 10, 64.

Furutachi, S., Franklin, A.D., Mrsic-Flogel, T.D., Hofer, S.B., 2023. Cooperative thalamocortical circuit mechanism for sensory prediction errors. bioRxiv , 2023–07.

Haber, A., Wanner, A., Friedrich, R.W., Schneidman, E., 2023. The structure and function of neural connectomes are shaped by a small number of design principles. bioRxiv , 2023–03.

Herculano-Houzel, S., Watson, C., Paxinos, G., 2013. Distribution of neurons in functional areas of the mouse cerebral cortex reveals quantitatively different cortical zones. Frontiers in neuroanatomy 7, 35.

Hooks, B.M., Chen, C., 2020. Circuitry Underlying Experience-Dependent Plasticity in the Mouse Visual System. Neuron 107, 986–987.

Jazayeri, M., Ostojic, S., 2021. Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity. Curr Opin Neurobiol 70, 113–120.

Kalatsky, V.A., Stryker, M.P., 2003. New paradigm for optical imaging: temporally encoded maps of intrinsic signal. Neuron 38, 529–545.

Kipf, T.N., Welling, M., 2016. Variational graph auto-encoders. arXiv preprint arXiv:1611.07308 .

Kumar, M.G., Hu, M., Ramanujan, A., Sur, M., Murthy, H.A., 2021. Functional parcellation of mouse visual cortex using statistical techniques reveals response-dependent clustering of cortical processing areas. PLOS Computational Biology 17, e1008548.

Lyon, D.C., Xu, X., Casagrande, V.A., Stefansic, J.D., Shima, D., Kaas, J.H., 2002. Optical imaging reveals retinotopic organization of dorsal v3 in new world owl monkeys. Proceedings of the National Academy of Sciences 99, 15735–15742.

Marshel, J.H., Garrett, M.E., Nauhaus, I., Callaway, E.M., 2011. Functional specialization of seven mouse visual cortical areas. Neuron 72, 1040–1054.

McInnes, L., Healy, J., Saul, N., Großberger, L., 2018. Umap: Uniform manifold approximation and projection. Journal of Open Source Software 3, 861. doi:10.21105/joss.00861.

Nassi, J.J., Callaway, E.M., 2006. Multiple circuits relaying primate parallel visual pathways to the middle temporal area. Journal of Neuroscience 26, 12789–12798.

Olavarria, J., Montero, V.M., 1989. Organization of visual cortex in the mouse revealed by correlating callosal and striate-extrastriate connections. Vis Neurosci 3, 59–69.

Schneider-Mizell, C.M., Bodor, A., Brittain, D., Buchanan, J., Bumbarger, D.J., Elabbady, L., Kapner, D., Kinn, S., Mahalingam, G., Seshamani, S., et al., 2023. Cell-type-specific inhibitory circuitry from a connectomic census of mouse visual cortex. bioRxiv .

Scholtens, L.H., van den Heuvel, M.P., 2018. Multimodal Connectomics in Psychiatry: Bridging Scales From Micro to Macro. Biol Psychiatry Cogn Neurosci Neuroimaging 3, 767–776.

Schrimpf, M., Kubilius, J., Hong, H., Majaj, N.J., Rajalingham, R., Issa, E.B., Kar, K., Bashivan, P., Prescott-Roy, J., Geiger, F., et al., 2018. Brain-score: Which artificial neural network for object recognition is most brain-like? BioRxiv , 407007.

Schuett, S., Bonhoeffer, T., Hübener, M., 2002. Mapping retinotopic structure in mouse visual cortex with optical imaging. Journal of Neuroscience 22, 6549–6559.

Seth, A.K., Barrett, A.B., Barnett, L., 2015. Granger causality analysis in neuroscience and neuroimaging. J Neurosci 35, 3293–3297.

Shi, J., Tripp, B., Shea-Brown, E., Mihalas, S., A. Buice, M., 2022. Mousenet: A biologically constrained convolutional neural network model for the mouse visual cortex. PLOS Computational Biology 18, e1010427.

Thompson, M.E., Ramirez Ramirez, L.L., Lyubchich, V., Gel, Y.R., 2016. Using the bootstrap for statistical inference on random graphs. Canadian Journal of Statistics 44, 3–24. doi:https://doi.org/10.1002/cjs.11271.

Tyler, C.W., 2004. Representation of stereoscopic structure in human and monkey cortex. Trends in Neurosciences 27, 116–118. doi:https://doi.org/10.1016/j.tins.2003.12.009.

Tyler, C.W., Likova, L.T., Chen, C.C., Kontsevich, L.L., Schira, M.M., Wade, A.R., 2005. Extended concepts of occipital retinotopy. Current Medical Imaging 1, 319–329.

Umakantha, A., Morina, R., Cowley, B.R., Snyder, A.C., Smith, M.A., Yu, B.M., 2021. Bridging neuronal correlations and dimensionality reduction. Neuron 109, 2740–2754.e12. doi:https://doi.org/10.1016/j.neuron.2021.06.028.

Wang, E.Y., Fahey, P.G., Ponder, K., Ding, Z., Chang, A., Muhammad, T., Patel, S., Ding, Z., Tran, D., Fu, J., et al., 2023. Towards a foundation model of the mouse visual cortex. bioRxiv .

Yamins, D.L., DiCarlo, J.J., 2016. Using goal-driven deep learning models to understand sensory cortex. Nature neuroscience 19, 356–365.

Zhuang, J., Ng, L., Williams, D., Valley, M., Li, Y., Garrett, M., Waters, J., 2017. An extended retinotopic map of mouse cortex. eLife 6, e18372. doi:10.7554/eLife.18372.