Specific inhibition and disinhibition in the higher-order structure of a cortical connectome

Michael W. Reimann[1*], Daniela Egas Santander[1], András Ecker[1], and Eilif B. Muller[234*]

[1]**Blue Brain Project, École polytechnique fédérale de Lausanne (EPFL), Campus Biotech, Geneva, Switzerland**
[2]**Department of Neurosciences, Faculty of Medicine, Université de Montréal, Montréal, Canada**
[3]**CHU Sainte-Justine Research Center, Montréal, Canada**
[4]**Mila Quebec AI Institute, Montréal, Canada**
[*]**Co-senior authors**

ABSTRACT

Neuronal network activity is thought to be structured around the activation of assemblies, or low-dimensional manifolds describing states of activity. Both views describe neurons acting not independently, but in concert, likely facilitated by strong recurrent excitation between them. The role of inhibition in these frameworks – if considered at all – is often reduced to blanket inhibition with no specificity with respect to which excitatory neurons are targeted. We analyzed the structure of excitation and inhibition in the MICrONS $1\mathrm{mm}^3$ dataset, an electron microscopic reconstruction of a piece of cortical tissue. We found that excitation was structured around a feed-forward flow in non-random motifs of seven or more neurons. This revealed a structure of information flow from a small number of sources to a larger number of potential targets that became only visible when larger motifs were considered instead of individual pairs. Inhibitory neurons targeted and were targeted by neurons in specific sequential positions of these motifs. Additionally, disynaptic inhibition was strongest between target motifs excited by the same group of source neurons, implying competition between them. The structure of this inhibition was also highly specific and symmetrical, contradicting the idea of non-specific blanket inhibition. None of these trends are detectable in only pairwise connectivity, demonstrating that inhibition is specifically structured by these large motifs. Further, we found that these motifs represent higher order connectivity patterns which are present, but to a lesser extent in a recently released, detailed computational model, and not at all in a distance-dependent control. These findings have important implications for how synaptic plasticity reorganizes neocortical connectivity to implement learning and for the specific role of inhibition in this process.

Keywords: Assemblies, Connectomics, Manifold, Inhibitory control

INTRODUCTION

Assemblies are groups of neurons that tend to fire together, and that have been observed in both hippocampal and cortical activity (Hebb, 1949; Harris et al., 2003; Dragoi and Buzsáki, 2006; Carrillo-Reid et al., 2015). Similarly, in simulation studies neurons are often wired into clusters that produce competing attractor states (Litwin-Kumar and Doiron, 2012; Deco and Hugues, 2012; Lagzi and Rotter, 2015). Such assemblies and clusters increase the reliability of a potential readout by increasing spiking correlations, at the cost of reducing the dimensionality of the activity state space. The models cited above implement the mechanism simply by assigning a stronger connection weight or a higher connection probability to pairs of neurons belonging to the same cluster. However, biological neuronal networks have complex *higher-order* structure that goes beyond connection strengthening on a pairwise level, such as overexpression of triad motifs (Song et al., 2005; Perin et al., 2011). Notably, this structure has been demonstrated to not be captured by models that only assign different strengths for different pathways, such as the ones cited

above (Reimann et al., 2017a; Gal et al., 2017; Reimann et al., 2017b; Gal et al., 2019). Furthermore, Renart et al. (2007) pointed out that the simple approach requires careful fine-tuning of connectivity parameters, and that additional biological realism may be needed for more robustness. This leads to the following question: Are there higher-order structures of connectivity in neuronal networks that support reliable network states, characterized by the activation of sets of assemblies, and transitions between them?

To find possible connectivity mechanisms for this purpose, we analyzed a biological connectome at cellular resolution (The MICrONS Consortium et al., 2021). Specifically, we detected and analyzed *directed simplices* (Reimann et al., 2017b) in the excitatory subnetwork. Directed simplices are large, tightly connected connectivity motifs with directionality giving them an input (source) and an output (target) side. They have been consistently demonstrated to be overexpressed in neural circuits of many organisms and at all scales, ranging from C. elegans (Reimann et al., 2017b; Sizemore et al., 2019; Shi et al., 2021) to rat cortical circuits (Perin et al., 2011; Song et al., 2005; Reimann et al., 2017b), and to human regional functional connectivity (Sizemore et al., 2018). In simulation studies (Reimann et al., 2017b; Nolte et al., 2019) they have been demonstrated to affect network function by increasing spiking correlations and facilitating reliable information transmission from their input to output. Their abundance, strong internal connectivity and correlations make them relevant as a potential structural source of assembly formation, and the directionality of their connectivity a potential mechanism for the temporal transitions between them.

Inhibitory interactions may also play an import role. For example, competition between assemblies may be mediated by inhibitory interneurons (Bathellier et al., 2012). Inhibitory control has been shown to be crucial for a balanced, asynchronous state (Renart et al., 2010), but is often considered to be dense and non-specific (Fino and Yuste, 2011; Packer and Yuste, 2011; Litwin-Kumar and Doiron, 2012; Curto et al., 2019). In the context of assemblies though, more specifically targeted inhibition can be advantageous (Rost et al., 2018) and may even have a computational function. Due to the distributed nature of assemblies (Carrillo-Reid et al., 2015), such specificity would be hard to detect without complete knowledge about which neurons participate in which assemblies. If directed simplices are correlates of assemblies, we may be able to find inhibitory specificity by analyzing connections between inhibitory neurons and these motifs, and specifically which neurons in these motifs are targeted.

At this level of complexity, the analyses have to be carefully controlled by comparing to relevant null models of connectivity, in order to understand the anatomical, morphological and molecular mechanisms giving rise to the results. If results can be explained by the spatial arrangement of neurons and distance-dependent connectivity, the case for their functional relevance is weak. At the same time, the shape of axons and dendrites has been shown to be a contributing factor to the complexity of connectivity (Stepanyants and Chklovskii, 2005) but it remains unclear to what degree these geometrical factors interact with other mechanisms affecting connectivity, such as activity dependent structural plasticity. Controls taking into account some mechanisms that shape connectivity but not others can help us disentangle the mechanisms.

As a data source of ground truth connectivity at the scale required for an exhaustive analysis, we used a connectome extracted from a freely available electron microscopic (EM) reconstruction of neural tissue, the MICrONS 1mm^3 mouse visual cortex dataset (The MICrONS Consortium et al., 2021), from which we extracted locations of 60,048 neurons and their internal synaptic connectivity. We considered multiple centrally placed subvolumes to avoid edge effects and gauge the variability of the results. As controls we used random distance-dependent connectomes on the same neuron locations, and a recently released morphologically detailed model of cortical tissue with highly structured connectivity (Reimann et al. (2022); Isbister et al. (2023); from here on: *nbS1 model*) as a baseline taking into account anatomical and morphological factors, but not molecular or plasticity-related ones. To ascertain the functional relevance of directed simplices, we confirmed the *in silico* finding in Reimann et al. (2017b) that the activity correlation between connected neurons increases if the connection belongs to a simplex motif. To this end, we used calcium imaging data of neuronal activations to visual stimuli that was co-registered with the EM data.

Conducting analyses as outlined above, we found: First, a hidden, lower-dimensional, directed and

divergent network acting as a backbone of the recurrent connectivity and made up of excitatory simplices. Second, specific inhibition that is structured by the divergent network in that it follows the same direction. Third, the inhibition in turn is under targeted control by specialized disinhibitory neurons. Fourth, these trends exist in the higher-order connectivity and are invisible on the pairwise level. Fifth, the trends are partially present in the computational model, indicating that they are morphologically prescribed and strengthened by other mechanisms. A plasticity-inspired rewiring rule makes the excitatory subnetwork of the model more similar to the EM connectome, but not the inhibitory to excitatory connectivity, indicating the presence of additional mechanisms.

1 MATERIALS AND METHODS

Preparation of connectivity data

Electron microscopic dataset

We extracted the internal connectivity between 60,048 neurons in v117 of the "minnie65_public" release of the MICrONS dataset. Specifically we used the table`allen_soma_coarse_cell_class_model_v1__minnie3_v1" to identify neurons and their types in broad classes. Information on synaptic connectivity, including synapse sizes were loaded from the table "synapses_pni_2". Synapses from sources other than one of the 60,048 classified neurons inside the volume were discarded. We have made the extracted data available in a format optimized for the following analyses (see Data and code availability).

Selection of neurons to analyze

We considered 15 centrally located excitatory subnetworks by filtering neurons according to their horizontal positions. They were admitted if they were inside a $500 \times 300 \mu m$ rectangle. For each subnetwork the rectangle was shifted by $50 \mu m$ into another location, resulting in a 5 by 3 grid of subvolumes with significant overlap (Fig. 1A1). Only excitatory neurons were filtered according to these subvolumes; inhibitory neurons were considered for analyses independent of their location inside or outside the rectangle. When results for an individual subnetwork are presented, they are for the most central one. When data for all 15 subnetworks is shown, it provided as a mean and standard deviation taken over the subnetworks.

Distance-dependent controls

To assess significance, we compared results to distance-dependent control connectomes. For each subnetwork exponential control models were fit based on the observed distance-dependent connection probabilities within and between excitatory (E) and inhibitory (I) populations, i.e., one model each for E to E, E to I, I to E, and I to I. Models had the form $P = a \cdot e^{\frac{d}{b}}$, where P was the probability of connection, d was the Euclidean distance between neurons and a and b the fitted parameters. A stochastic instance was then generated, and analyses were conducted as in the baseline case.

Morphologically detailed model control

Additionally, we compared results to a recently released morphologically detailed model with connectivity derived from the shapes of neurons and biological constraints (Reimann et al., 2015; Isbister et al., 2023). This allows us to predict to what degree features of connectivity are explained by neuron morphologies, and to what degree other mechanisms, such as plasticity shape the connectome. For the model, subnetworks were only $300 \times 300 \mu m$. The size was selected such that they contain approximately the same number of excitatory neurons as the MICrONS subvolumes. Inhibitory neurons were inside rectangles 100 μm larger in each of the four directions, yielding approximately the same number of inhibitory neurons.

Topological analyses

Finding directed simplices

Using the `connalysis` (Egas Santander et al., 2023) python package (see Data and code availability), we calculated the list of directed simplices of all dimensions of the excitatory subnetwork. A directed simplex of dimension dim is a motif of $dim+1$ neurons where the connectivity between them fulfills the following criterion: There exists a numbering of neurons, i, j, in the motif from 0 to dim, such that if $i < j$ then there exists a synaptic connection from the ith neuron to the jth neuron (Fig. 1B). We call that numbering the position of the neuron in the simplex. We call neurons numbered near 0 the source side of

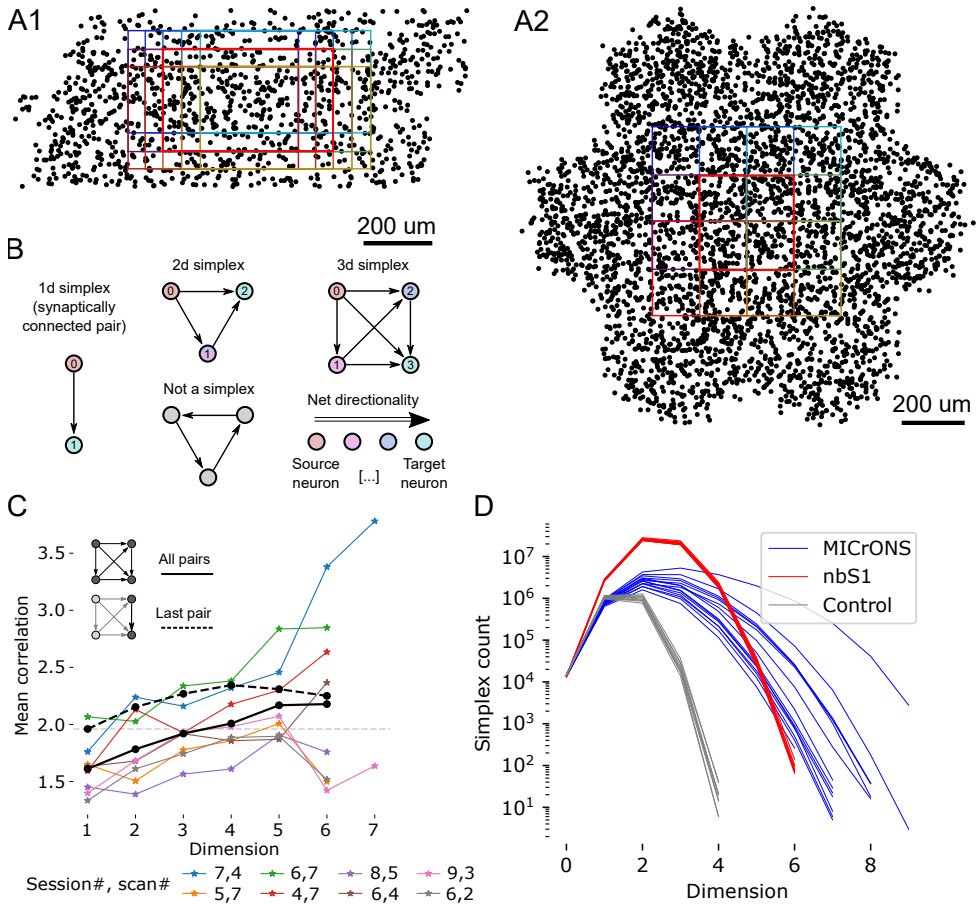

Figure 1. Analyzed connectomes and their simplicial structure. A1: Black dots indicate 2% of the neurons inside the *MICrONS volume*, an electron microscopic reconstruction of around $1mm^3$ of cortical tissue, the contained neurons and their connectivity (The MICrONS Consortium et al., 2021). Colored rectangles show the locations of 15 subvolumes separately analyzed. The thicker red outline indicates the *central subnetwork* that will serve as the example for some of the following analyses. A2: As A1, but for the *nbS1 volume*, a recently released, morphologically detailed model of cortical brain tissue (Isbister et al., 2023). Only 9 subvolumes were analyzed for the nbS1 volume. B: A *directed simplex* is a densely connected, feed-forward motif that generalizes the concept of an edge in a graph. The *dimension* of a simplex is determined by the number of participating neurons. C: Correlations of the activity of pairs of neurons in the MICrONS data against their simplex membership. The x-axis indicates the maximum dimension over simplices the connection participates in. Black lines: mean values over recording sessions. Black dashed line: values when only the last pair in a simplex is considered. Grey dashed line: Overall mean for connected pairs. D: Simplex counts in the MICrONS and nbS1 volumes (blue and red) and control models with distance-dependent connectivity, fit to 9 of the MICrONS subnetworks (grey).

a simplex and neurons near *dim* the target side.

Analysis of inhibition of simplices

We then represented the topological location of all inhibitory neurons relative to the simplices as follows: Let S_{dim} be the number of directed simplices in dimension dim and N_{inh} be the number of inhibitory neurons. We constructed a $S_{dim} \times dim + 1 \times N_{inh}$ tensor $M_{dim,ei}$, where the entry at index i, j, k is the number of synapses from the neuron at position j of excitatory simplex i to inhibitory neuron k. Similarly, in $M_{dim,ie}$ the entry at i, j, k is the number of synapses from the inhibitory neuron k to the neuron at position j of simplex i.

Let $M'_{dim,ei}$ and $M'_{dim,ie}$ be the sums of $M_{dim,ei}$ and $M_{dim,ie}$ over their second dimensions respectively. That

is, $S_{dim} \times N_{inh}$ matrices that count the total number of synaptic connections from / to a simplex to / from an inhibitory neuron. Then we call the matrix product $M'_{dim,ei} \cdot M'^{\top}_{dim,ie} = I_{dim}$ the matrix of disynaptic inhibition between simplices of dimension dim.

Symmetry of simplex-simplex inhibition matrices
We analyzed how symmetrical the matrix of inhibition between simplices was by calculating how much the variance of its entries is reduced when the transpose is subtracted:

$$\mathbb{S}_{dim} = \frac{var(I_{dim} - I^{\top}_{dim})}{var(I_{dim})} \tag{1}$$

Increasing symmetry reduces this measure, until it reaches 0 in the case of perfect symmetry. The expected value if the entry at i, j and at j, i are statistically independent is 2. We calculated controls for S_{dim} by shuffling the columns of $M'_{dim,ei}$ and $M'_{dim,ie}$ before calculating I_{dim} and then S_{dim} as usual.

Analysis of simplex overlap
Simplices can overlap in their constituent neurons by up to $dim+1$ neurons. Let $O_{dim,[j,k]}$ be a $S_{dim} \times S_{dim}$ matrix that counts for each combination of simplices their overlap in neurons in positions between j and k (i.e. the size of the intersection of neurons in positions i with $j \leq i \leq k$).

Clustering of source- and target groups of simplices
Simplices were grouped using the Louvain clustering algorithm (Barber, 2007) from the python `sknetwork` package (Bonald et al., 2020) on the matrices $O_{6,[0,2]}$ and $O_{6,[4,6]}$ of the MICrONS subnetworks, yielding what we denote *source groups* and *target groups* of simplices respectively.

Analysis of disinhibition
For inhibitory neurons, we calculated their *simplex indegree* as the number of connections from excitatory neurons in 6-dimensional simplices in positions 0, 1 or 2. Correspondingly, the *simplex outdegree* was the number of connections to neurons in 6-dimensional simplices in positions 4, 5 or 6. If a neuron participated in more than one simplex in one of the indicated positions its connection would be counted that number of times.
Also, for inhibitory neurons, we calculated their targeting preference for other inhibitory neurons as follows. We calculated for a neuron i its outdegree d_i^{ttl} and outdegree onto inhibitory neurons d_i^{inh}. We then compared against a control that keeps the total outdegrees of individual neurons and the overall fraction of connections onto inhibitory neurons. Under this null hypothesis, the probability to find an inhibitory outdegree as observed or higher is:

$$p_j = 1 - HG(\sum_i d_i^{ttl}, \sum_i d_i^{inh}, d_j^{inh} - 1), \tag{2}$$

where $HG(M,N,n)$ is the cumulative hypergeometric distribution of drawing n times from M objects with N positive objects. We considered a neuron to be inhibitory targeting if $p_j \leq 10^{-6}$. The threshold was selected to yield around 15% of inhibitory neurons, in line with the fraction of inhibitory targeting cells in (Schneider-Mizell et al., 2023), which was 17%.

Neighborhood complexity
We measured the complexity of the graph neighborhoods of individual neurons in the MICrONS data in accordance with Egas Santander et al. (2023). Briefly, we first constructed as a control a configuration model of the central subnetwork, i.e., a random graph that preserves for each neuron its in- and out-degree. We then considered the subgraph given by the neighborhood of a neuron, i.e., the selected neuron and all its afferents and efferents and the connections between them. Specifically, we compared the neighborhood of a neuron with the neighborhood of the same neuron in the configuration model. As the control preserves degrees, both neighborhoods have the same size, but different structure. We compared them using the Wasserstein distance of the distribution of the degrees of their nodes.

Plasticity-inspired rewiring

We rewired the connectivity of the central subvolume of the nbS1 model with a custom algorithm inspired by recent findings, based on simulations and analysis of the MICrONS data (Ecker et al., 2023), showing that plasticity favors connections participating in many high-dimensional simplices. A single rewiring step worked as follows: First, for each excitatory connection $(a \to b)$ the maximum dimension of simplices it participates in was calculated as $D(a \to b)$. Next, for each excitatory neuron a, we counted the number of simplices it participated in, for each dimension and position and denote it by $\mathrm{Par}_{dim,pos}(a)$. A specified number of connections m with a maximum dimension below or equal D were removed independently at random. Finally, the same number m new connections were added between random unconnected pairs of neurons, $a \to b$, according to a probability, $P(a \to b)$. The probability was proportional to the participation of a and b in simplices above dimension D, in positions i and j respectively for $i < j$ belonging to a specified set we denote pos as follows:

$$P(a \to b) \propto \sum_{dim>D} \left(\sum_{i<j \in \mathrm{pos}} \mathrm{Par}_{dim,i}(a) \cdot \mathrm{Par}_{dim,j}(b) \right).$$

We conducted a single rewiring step with $D = 3$ (the middle dimension), $m = 2000$, followed by three steps with $D = 2$, $m = 10000$ each, for all we focused on the source side by using $\mathrm{pos} = \{0, 1\}$.

Analysis of functional MICrONS data

The MICrONS dataset we studied also contains functional data of neurons co-registered with the structural data. Specifically, it contains calcium imaging traces and spike trains derived from a deconvolution of the traces. The Pearson correlation of spike trains of excitatory neurons was computed for 8 sessions. The sessions were selected from the activity recorded in the "minnie65_public" release version 661, such that at least 1000 neurons were scanned in each session and at least 85% of them are co-registered in the structural data we used. While every functionally imaged neuron is part of the EM volume, some were not contained in the table of neuron identifiers we used as the basis of the connectivity data (see above). Neurons that do not have a unique identifier were also filtered out (between 0 and 10.8% of co-registered cells and on average 2.8% across sessions). Next, we average these correlations along groups determined by the underlying structure in two ways. First, we grouped each excitatory connection by the maximum dimension of simplices it participates in. Second, for each dimension n, we extracted the list of the last connections in n-simplices (with repetitions), i.e., the connections between nodes in position $n-1$ to n. Lastly, we averaged the correlations across those connections.

2 RESULTS

2.1 High-dimensional simplices reveal directed and divergent connectivity in the excitatory population

We analyzed the graph structure of synaptic connectivity of the MICrONS dataset (The MICrONS Consortium et al., 2021). To assess its variability, we considered fifteen $500 \times 300 \mu m$ overlapping subnetworks (Figure 1A1). We compared to distance-dependent control models fit to the data and the connectivity of a subvolume of the *nbS1 model* (Reimann et al. (2022); Isbister et al. (2023), Fig. 1A2). Analyzed volumes of the nbS1 model were slightly smaller, at $300 \times 300 \mu m$, to approximately match the neuron counts inside the volumes (14309 ± 708 neurons for MICrONS vs. 13538 ± 289 for nbS1).

Analysis was conducted primarily with respect to *directed simplices* in the graphs representing the respective connectomes. A simplex with *dimension* d is a motif of $d+1$ neurons with directed all-to-all connectivity (Fig. 1B). Each participating neuron has a *simplex position* numbered from 0 (the *source position*) to $d+1$ (the *target position*). An edge must exist from neuron a to b if $b > a$; edges in the opposite direction can exist but are not required. As such, the concept generalizes the concept of a connection between neurons, which is a 1-dimensional simplex.

Simplex motifs have been demonstrated to be overexpressed in virtually all biological connectomes investigated (Perin et al., 2011; Song et al., 2005; Sizemore et al., 2018) and have functional relevance in terms of increasing the spiking correlations of their members (Reimann et al., 2017b; Nolte et al., 2020). We confirmed both findings in the MICrONS data. Previous results on the same dataset found

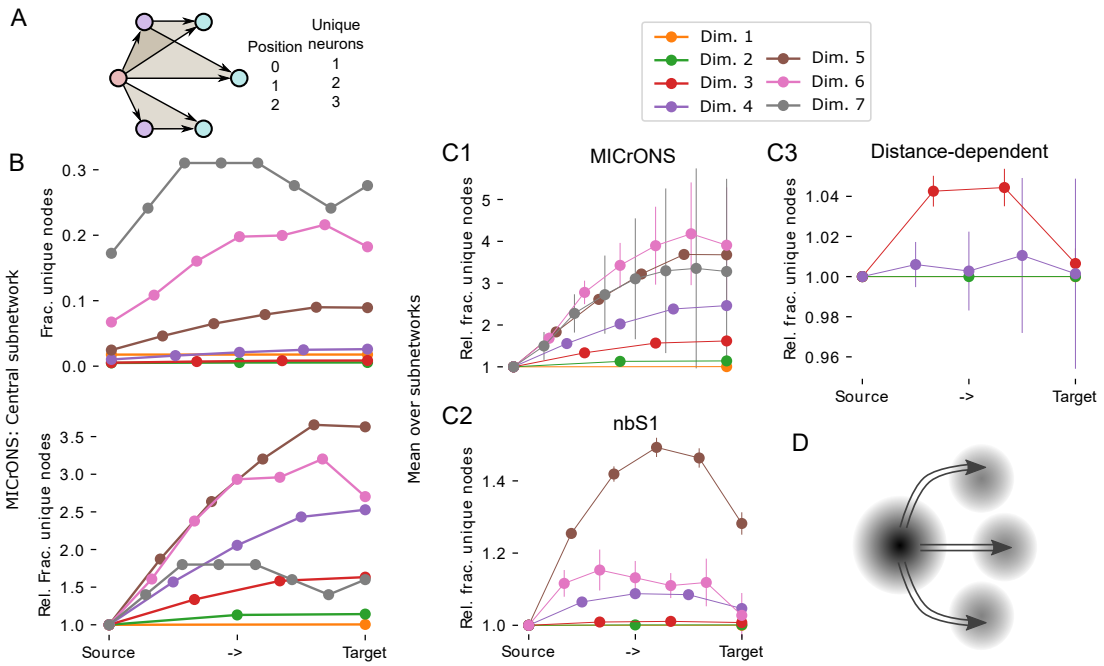

Figure 2. Analysis of simplices reveals directionality of information flow. A: As a neuron can be a member of several simplices, only a fraction of neurons in simplices of a given dimension is unique. B: Top: Fraction of unique neurons in simplices of the central subnetwork of MICrONS, sorted by their position in the simplex from *source* to *target*. Bottom: Same, but normalized to fraction at the source. C1: As B, but means and standard deviation over the 15 subnetworks of the MICrONS data. C2,C3: As C1, but for the 9 subnetworks of the nbS1 model and distance-dependent control models. D: Schematic interpretation of the results: A connectivity structure with activity flow from a compact, relatively small number of sources to a larger number of potential targets.

increased activity correlations for synaptically connected pairs of neurons (Ding et al., 2023). We extend this, showing even higher correlations when the connections are part of a simplex with a dimension above 4 or when it is placed at the target position of the simplex (Fig. 1C). Additionally, simplex counts in dimensions above 4 were even higher than in the nbS1 model, but also more variable (Fig. 1D). Both MICrONS and nbS1 had significantly increased simplex counts compared to the distance-dependent control. We conclude that high-dimensional simplices have both structural and functional relevance in the MICrONS data and continue our analysis of them, with particular interest in their relationship to plasticity and inhibitory network specificity.

Individual neurons can participate in more than one simplex, forming complex, intersecting networks even when only simplices of a single dimension are considered (Fig. 2A). We calculated for the central subvolume of the MICrONS data the fraction of unique neurons separately for different dimensions and positions within the simplex (Fig. 2B, top). Results depended mostly on dimension considered, which is expected from the respective simplex counts being higher than the neuron count in some dimensions. When normalized with respect to the unique count for the source position, we found a strong increase of unique neuron count towards the target position (Fig. 2B, bottom). The strength of the trend mostly increased with dimension considered. When we extended the analysis to all 15 Micron subnetworks (Fig. 2C1), we found the effect to be consistently present and its strength peaking when dimension 6 is considered. At the top dimension (7), results are noisy due to the relatively low motif counts. The trend was also found in the nbS1 model, but significantly weaker, peaking already at dimension 5 with a relative fraction of unique nodes of < 1.5 vs. ≈ 4.0 for microns (Fig. 2C2). It was extremely weak to non-existing in the distance-dependent control (Fig. 2C3).

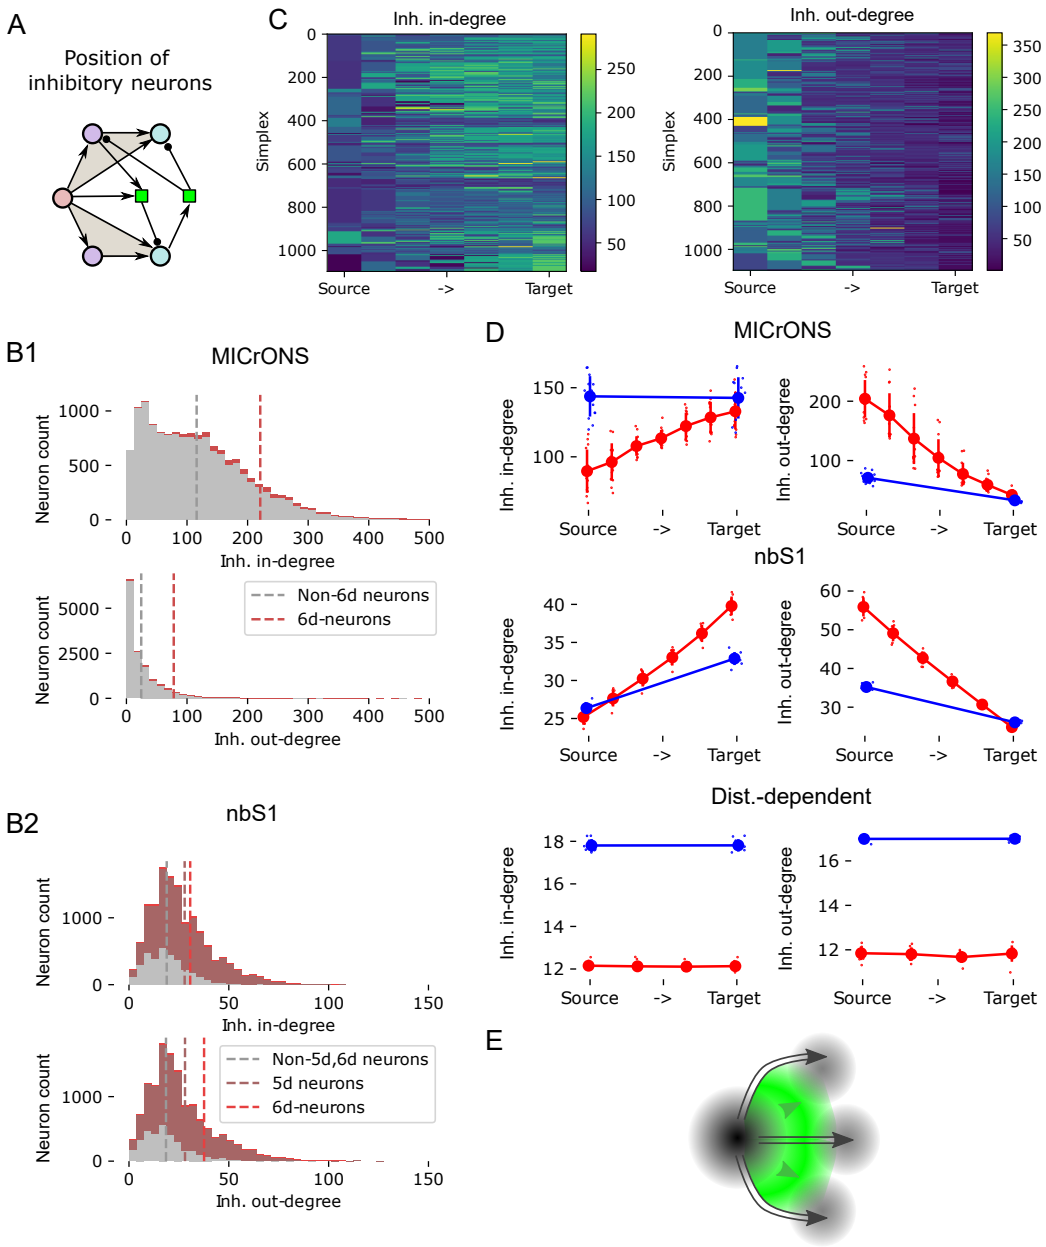

Figure 3. Non-random positioning of inhibition within the simplicial network. A: We consider the graph locations of inhibitory neurons (green) relative to simplices. Specifically, the number of inhibitory neurons innervating the neuron at a given position of a simplex (*inhibitory in-degree*) or innervated by it (*inhibitory out-degree*). B1: Distribution of the inhibitory in- and out-degrees of the central subnetwork of MICrONS. Red: For neurons participating in at least one 6d-simplex. Grey: For other neurons. Dashed lines: respective means. B2: As B1, but for the nbS1 model. Additionally, data for the 5-core is shown (dark red). C: Inhibitory in-degree (left) and out-degree (right) for neurons in all 1093 6d-simplices of the central subnetwork of MICrONS (top to bottom). Participating neurons indicated from left to right, sorted source to target. D: Mean inhibitory in- and out-degrees of neurons in each position of a simplex. Indicated are values for individual subnetworks (small dots) and mean and standard deviation over subnetworks (large dots and error bars). Red: data for the dimension with the highest directionality in Fig. 2C-E, i.e. 6 for MICrONS, 5 for nbS1, 3 for distance-dependent; blue: Data for 1d-simplices, i.e. source and target neurons of individual synaptic connections. E: Schematic interpretation: Inhibition (green) respects the directionality of the high-dimensional excitatory connectivity of Fig. 2.

As connectivity in simplices is directed from the source towards the target neuron, this reveals aspects of the structure of information flow in the network: We found a highly non-random flow of connectivity from a compact source to a larger number of potential targets (Fig. 2D). Notably, this divergence is only visible when high-dimensional simplices are considered as it is entirely absent on the level of individual edges (Fig. 2, orange). We call the dimension where the divergent trend is strongest the *target dimension* of a network (6 for MICrONS, 5 for nbS1, 3 for distance-dependent; see Fig. 2C). Results for the following analyses will focus on and compare the respective target dimensions.

2.2 Inhibition is structured by the simplicial structure of excitatory neurons

Next, we analyzed the graph locations of inhibitory neurons with respect to the simplicial structure (Fig. 3A). First, we found that in the central MICrONS volume excitatory neurons participating in 6d-simplices (i.e., the *6-core* of the central subnetwork) innervate and are innervated by more inhibitory neurons than the rest (Fig. 3B1). We confirmed that this was not merely an edge effect, where neurons in a more central spatial position are simply more likely to be connected to both excitatory and inhibitory neurons, as follows. First, while we detected simplices only for excitatory neuron inside the $500 \times 300 \mu m$ subvolume, we considered for this analysis inhibitory neurons of the entire MICrONS volume, thereby reducing the impact of the border drawn. Second, we compared the distances of neurons from the center of the volume to their inhibitory degrees, finding no dependence at all (Fig. S2). Third, neurons of the 6-core covered almost the entire range of distances from the volume center (Fig. S2, red vs. black). In the central nbS1 volume we conducted a corresponding analysis for its target dimension of 5, but since its 5-core contained more than half of the neurons, we also considered dimension 6 (Fig. 3B2). The increased inhibitory in- and out-degree was also present in the nbS1 model, but weaker: A 60% increase of inhibitory in-degree for neurons in the 6-core vs. 90% for MICrONS; a 100% increase of inhibitory out-degree vs. 200% for MICrONS.

Crucially, this increase was strongly dependent on the position of a neuron in a 6-dimensional simplex (Fig. 3C). It was specifically neurons around the source position that innervated many inhibitory neurons, and neurons around the target position were strongly innervated by inhibitory neurons. Additionally, connections adhering to that principle were also formed by more synapses per connection (Fig. S1) indicating that such connections are selective strengthened through structural plasticity. Once again, the trend was consistently present in all 15 subnetworks of the MICrONS volume (Fig. 3D, top) and also present in the nbS1 data, but absent for the distance-dependent control (Fig. 3D, middle and bottom). Similar to the previous result of a divergent excitatory structure, the trend was almost undetectable when only single edges (1-simplices) are considered (Fig. 3D top, red vs. blue). Except for the nbS1 model, where it was visible but weaker when single edges are considered. We conclude that disynaptic inhibition between excitatory neurons is specific to the connectivity structures formed by high-dimensional simplices, and its directionality follows the direction given by the simplices (Fig. 3E).

2.3 Inhibition implements specific competition between simplicial motifs, moderated by targeted disinhibition

Encouraged by the non-random targeting of disynaptic inhibition we found, we analyzed its structure further. We now considered the strength of disynaptic inhibition between pairs of simplices in the respective target dimensions. Specifically, we counted the number of $E \to I \to E$ paths where the first neuron was member of simplex i and the last neuron of simplex j (Fig. 4A). The resulting matrix of disynaptic inhibition strengths between 6d simplices in the central subnetwork of MICrONS appeared highly structured and symmetrical (Fig. 4B1). To confirm this appearance, we compared it to a random control where we preserved the inhibitory in- and out-degree of each simplex, but shuffled their targets (Fig. 4B2). We found that the actual matrix contained values outside of the distribution given by the control, both significantly weaker and stronger (Fig. 4C, top). Repeating the analysis for all 15 MICrONS subnetworks we found strongly varying results (Fig. S3). Importantly, in each subnetwork we found data points both on the left and on the right side of the control. In some cases, this even manifested in two separate peaks, in others the peak on the right side of the control dominated. This demonstrates that disynaptic inhibition is significantly weaker than expected between some pairs of simplices and significantly stronger between others, i.e., that it is targeted or specific. In terms of matrix symmetry, we computed the normalized variance of the values of the matrix minus its transpose (see Methods, Fig. 4D).

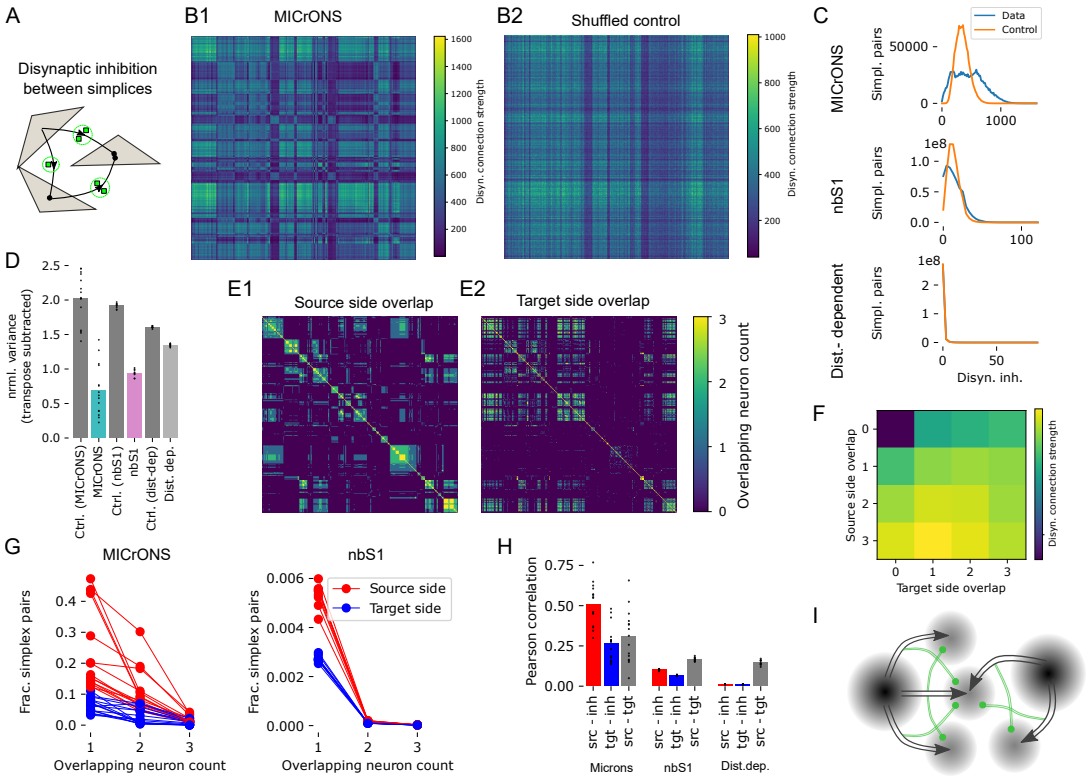

Figure 4. Specific, non-random disynaptic inhibition between simplices. A: We consider the strengths of *disynaptic inhibition* between pairs of high-dimensional simplices. B1: Disynaptic inhibition strength between pairs of 6d-simplices in the central subnetwork of MICrONS. Color indicates the number of paths from a neuron in the simplex along the vertical axis, to an inhibitory neuron, and then to a neuron in the simplex along the horizontal axis. B2: Same, for a random control, where the inhibitory in- and out-degrees of all simplices are preserved, but their sources and targets are shuffled. C, top: Distribution of disynaptic inhibition strength values in B1 (blue) and its control in B2 (orange). Middle: Same, for disynaptic inhibition between 5d-simplices in the nbS1 model. Bottom: For 3d-simplices in the distance-dependent controls. D: As a measure of symmetry of disynaptic inhibition matrices M, we show $\frac{var(M-M^\top)}{var(M)}$. Dots indicate values for individual subnetworks, bars the means for (from left to right): Shuffled controls for MICrONS, MICrONS, shuffled control for nbS1, nbS1, shuffled control for the distance-dependent models, distance-dependent models. E: Number of neurons in common within the first three positions (*source side*, E1) and last three positions (*target side*, E2) of pairs of 6d-simplices in the central subnetwork of MICrONS. F: Mean disynaptic connection strength (as in B1) for pairs of simplices with the source- and target side overlap indicated along the vertical and horizontal axes. Mean over the 15 MICrONS subnetworks. G: Left: Normalized distribution of values in E1 (source side, red) and E2 (target side, blue). Right: Same for 5d simplices in the nbS1 model. Values missing from a sum of 1.0 are for pairs with an overlap of 0. H: Pearson correlations between source side overlap (as in E1), target side overlap (as in E2) and disynaptic inhibition (as in B1). Dots indicate values for individual subnetworks, bars the respective means. For (from left to right): MICrONS, nbS1 and the distance-dependent controls. I: Schematic interpretation: Disynaptic inhibition facilitates mainly competition between different targets activated by the same source.

The expected value of this measure for unstructured data is 2.0. Indeed, this was found in the shuffled control. The value for the actual data was significantly lower, indicating a high degree of symmetry. This finding was consistently present in subnetworks but varied in strength. Both specificity and symmetry were present for the nbS1 model in its target dimension 5, but weaker; and they were not present in the distance-dependent control (Fig. 4C, middle and bottom; 4D).

This led to the question, what are the factors that affect disynaptic inhibition strength? What determined whether a value samples from the high or low peak? We considered *source side overlap* and *target side overlap*: The number of neurons a pair of 6d-simplices has in common in positions 0-2 (but in any order within that range; source side) and in positions 4-6 (target side; for details see Methods). This overlap tended to be stronger on the source side, as expected from the lower fraction of unique neurons in those positions (Fig. 4E, G). We found that the strength of disynaptic inhibition increased with both measures, but more so with source side overlap (Fig. 4F). This was confirmed by calculating the Pearson correlation between the three values (Fig. 4H). Results fluctuated over subnetworks but were consistently largest for correlation between source side overlap and disynaptic inhibition. Once again, the same trend was confirmed but weaker in nbS1 and not present in the distance-dependent control. These findings indicate that inhibition primarily implements a form of symmetric competition between multiple targets that are activated by the same source (Fig. 4I). Note that this differs from established models of competitive inhibition, where largely non-overlapping pools of excitatory neurons compete via inhibitory neurons. Since both source and targets are parts of the same simplices, this leads to a situation where the target neurons are simultaneously excited and disynaptically inhibited by the source population. This may facilitate a balance of excitation and inhibition, where minute fluctuations on top of a high-conductance state determine the outcome, in line with (Renart et al., 2010).

Another function of inhibitory connectivity is *disinhibition*, i.e., the inhibition of other inhibitory neurons. A subclass of Vasointenstinal protein (VIP)-expressing interneurons have been demonstrated to implement this function (Pi et al., 2013) and the presence of a class of inhibitory neurons specifically targeting other inhibitory neurons has been previously demonstrated in the MICrONS dataset (Schneider-Mizell et al., 2023). In the central subnetwork of MICrONS we found indication that the disinhibition preferentially targets the inhibitory neurons mediating competition between simplices (Fig. 5A, top): Not only were inhibitory neurons with many connections to and from 6-dimensional simplices more strongly inhibited themselves (Fig. 5A, middle), but this was specifically mediated by neurons specializing in such disinhibition. The fraction of inhibition coming from neurons with inhibitory targeting preference increased from 20% for neurons weakly interacting with simplices to 45% for the strongest interacting neurons (Fig. 5A, bottom right). We ensured that this is not an edge-effect by comparing the inhibitory indegrees to the distance from the center of the analyzed volume, finding no significant correlation (Fig. S4).

2.4 Spatial and topological structure of disynaptic inhibitory competition

So far, we have found a higher-order network acting as a backbone of the recurrent connectivity with divergent feed-forward excitation from groups of source simplices to groups of target simplices and specific lateral inhibition between them. Here, we describe the structure of that feed-forward network and the neurons participating in it.

As we have demonstrated that overlap in source and target neurons shapes the structure of disynaptic inhibition between simplices, we began by detecting these overlapping groups that inhibit or are inhibited together. We determined *source* and *target groups* of simplices by clustering the matrices of source side and target side overlap respectively using the Louvain algorithm with a resolution parameter of 2.2. Next, we estimate the strength of excitation and inhibition between source and target groups based on all pairs of simplices s_1, s_2 where s_1 is member of the source group and s_2 of the target group. For excitation, we calculate the mean overlap of neurons in any simplex location (Fig. 5B1). For inhibition, we calculate the fraction of pairs with significantly strong inhibition, i.e., with a disynaptic inhibition strength above the 95th percentile of the corresponding shuffled control (Fig. 5B2). The result was a 10×19 connectivity matrix of a feedforward network with relatively sparse excitation and specific inhibition. We note that the numerical values of our measures of excitation and inhibition are not directly comparable, but conclusions may be drawn based on their relative values. Aggregated over all 15 MICrONS subnetworks, the number of source and target groups varies, as does the number of simplices belonging to them. Indeed, the number of simplices in a cluster falls into a wide distribution spanning almost four orders of magnitude, indicating an organization principle at multiple scales (Fig. 5B3).

Taken together, this raises the question: what is the role of the complex higher-order structure within the feedforward structure? If feedforward connectivity is called for, why not implement a simple feedfor-

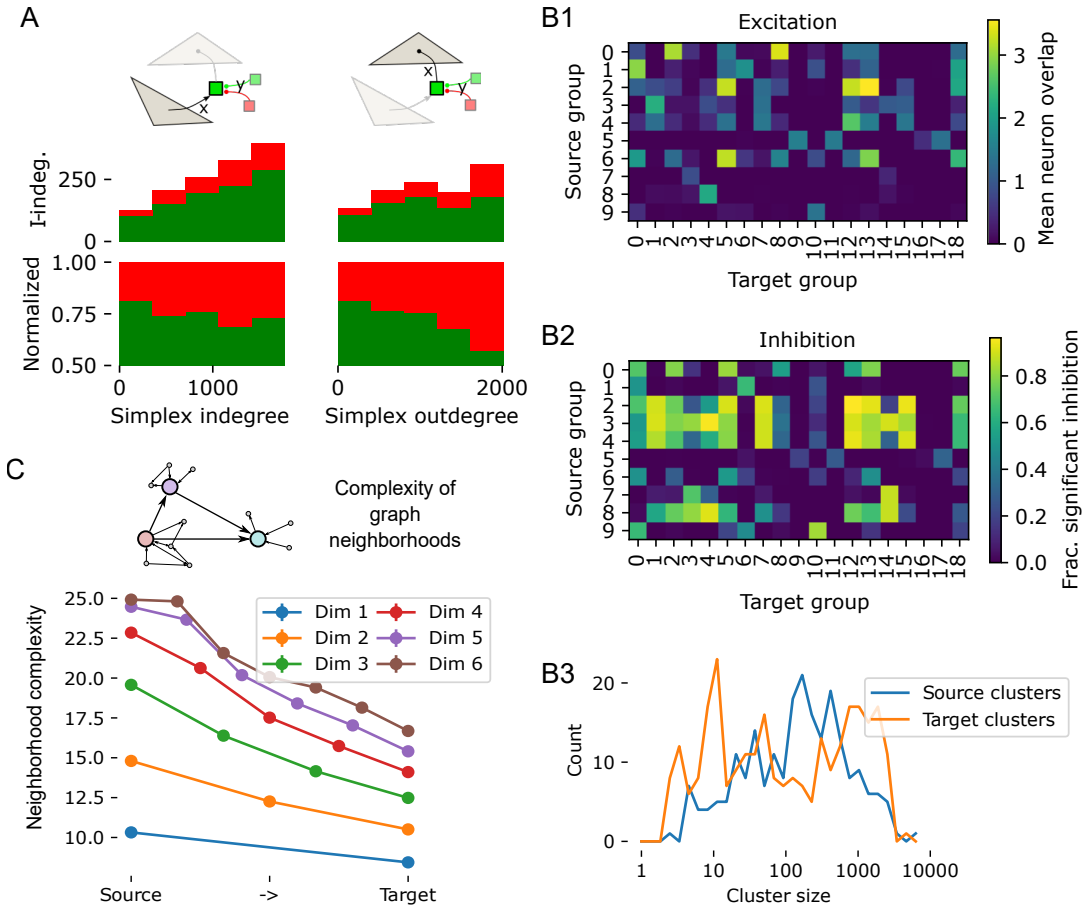

Figure 5. Specific disinhibition and lower-dimensional structure of the competitive simplicial network. A, top: We consider for the x-axis the *simplex degrees* of inhibitory neurons in the microns data. That is, its number of connections from (left, simplex indegree) or to (right, simplex outdegree) 6-dimensional simplices. We consider for the y-axis their indegree from inhibitory neurons with significant targeting preference for other inhibitory neurons ($p \leq 10^{-6}$; 15% of inhibitory neurons; red), or from the remaining inhibitory population (green). Stacked bar plot for absolute (middle) or normalized (bottom) inhibitory indegrees. B: Strengths of excitation (B1) and inhibition (B2) between groups of 6d-simplices in the central subnetwork of MICrONS. Groups were derived by clustering the source side overlap matrix (see Fig. 4 D, left), yielding *source groups* along the vertical axis, and by clustering the target side overlap matrix (Fig. 4D, right), yielding the *target groups* along the horizontal axis. For excitation, we considered the mean number of overlapping neurons in any position over pairs of simplices in the indicated groups. For inhibition we considered the fraction of pairs with disynaptic inhibition strength higher than the 95th percentile of the corresponding shuffled control (Fig. 4A, B). B3: Distributions of source and target group sizes in the 15 MICrONS subnetworks. C: Complexity of subgraphs given by graph neighborhoods of individual nodes in the central subnetwork of MICrONS. Indicated are the mean values for neighborhoods of nodes participating in the indicated position of simplices of the indicated dimension. For details on the complexity measure used see Egas Santander et al. (2023), Methods.

ward network? Simulation results (Egas Santander et al., 2023) predicted a role of higher-order structure in increasing the reliability of the network response. Egas Santander et al. (2023) characterized the graph neighborhoods of neurons in terms of how much their connectivity deviates from random, for example, with overexpression of bidirectional connectivity. They found that the presence non-random neighborhoods increases reliability globally, while more random neighborhoods provide better information readouts. Considering the feedforward structure given by high-dimensional simplices in our results,

we expect the following: As target neurons of simplices are the outputs of the circuit, we would expect them to be members of neighborhoods providing good information readout. Conversely, on the input side given by source neurons, we would expect subnetworks facilitating a reliable response. Indeed, replicating the analysis of Egas Santander et al. (2023) (Fig. 5C), we found the neighborhoods of target neurons to be more similar to random graphs than those of source neurons, with the strength of this trend increasing with simplex dimension.

By considering all neurons participating in at least one simplex of a source group, we can visualize its internal connectivity (Fig. 6A). This illustrates the dense feed-forward connectivity given by the 6-dimensional simplices contained in the group, but also the more complex, highly recurrent nature of this network: 38% of edges between nodes do not participate in the simplices making up this source group. Over all source groups, this is $41\% \pm 10\%$ (mean $\pm$ std). This highlights once more how the trends described in this manuscript exist on the higher-order level of structured connectivity between neuron motifs and not individual pairs. So far, we have soundly refuted the idea that any of the observed trends are caused merely by the relative locations of neurons using distance-dependent control connectomes. However, we still wanted to see to what degree neurons in a source group cluster together spatially. We found neurons in a source group to be spatially compact, but spanning up to $500\mu m$ (Fig. 6B, S5). Furthermore, groups can overlap spatially despite only marginal overlap in participating neurons (e.g. groups 0 vs 2 in Fig. S5). Information flow in local circuits such as these has been traditionally described as a flow between cortical layers (Felleman and Van Essen, 1991). Conversely, we found no layer specificity, with neurons in source positions (Fig. 6B, blue) and target positions (Fig. 6B, red) being present in all layers. This highlights once more that our results reveal a higher-order structure that is largely invisible at the pairwise level. Curiously though, we found that pairs of source neurons tended to be further apart from each other than target neurons (average of $220\mu m$ for pairs of source vs. $130\mu m$ for target neurons; Fig. S3). This may be explained by the source position containing more neurons classified as 5P_IT (Table 1), a class that is known to form long-range corticocortical connections (Harris et al., 2019). Another type overexpressed in the source position is 4P (layer 4 pyramidal cells), which is consistent with the classical view that they are the inputs of a local circuit (Felleman and Van Essen, 1991). Conversely, the 6IT class formed the outputs, participating almost exclusively in the target position.

2.5 Potential rewiring mechanisms shaping connectivity

While present in the nbS1 model, the non-random connectivity features were consistently weaker than in the MICrONS data. One of the mechanisms missed by the model explaining the difference could be structural plasticity. Recent simulation results predicted that plasticity favors connections that participate in many high-dimensional simplices, and this has been confirmed in the MICrONS data (Ecker et al., 2023). Accordingly, we implemented a simple heuristic algorithm for rewiring excitatory connections that removes edges that only participate in low-dimensional simplices, and places edges that are likely to form new high-dimensional simplices (Fig. 6C). To give rise to the observed increase in unique neurons towards the target (Fig. 2C), we focused on placing edges between neurons on the source side (for details, see Methods). We found that rewiring of only 1.2% of connections succeeded in increasing edge participation and simplex counts significantly (Fig. 6D) with a divergent structure at higher dimensions (Fig. 6E). However, this excitatory rewiring alone reduced the specificity of disynaptic inhibition (Fig. 6F, G) between the simplices.

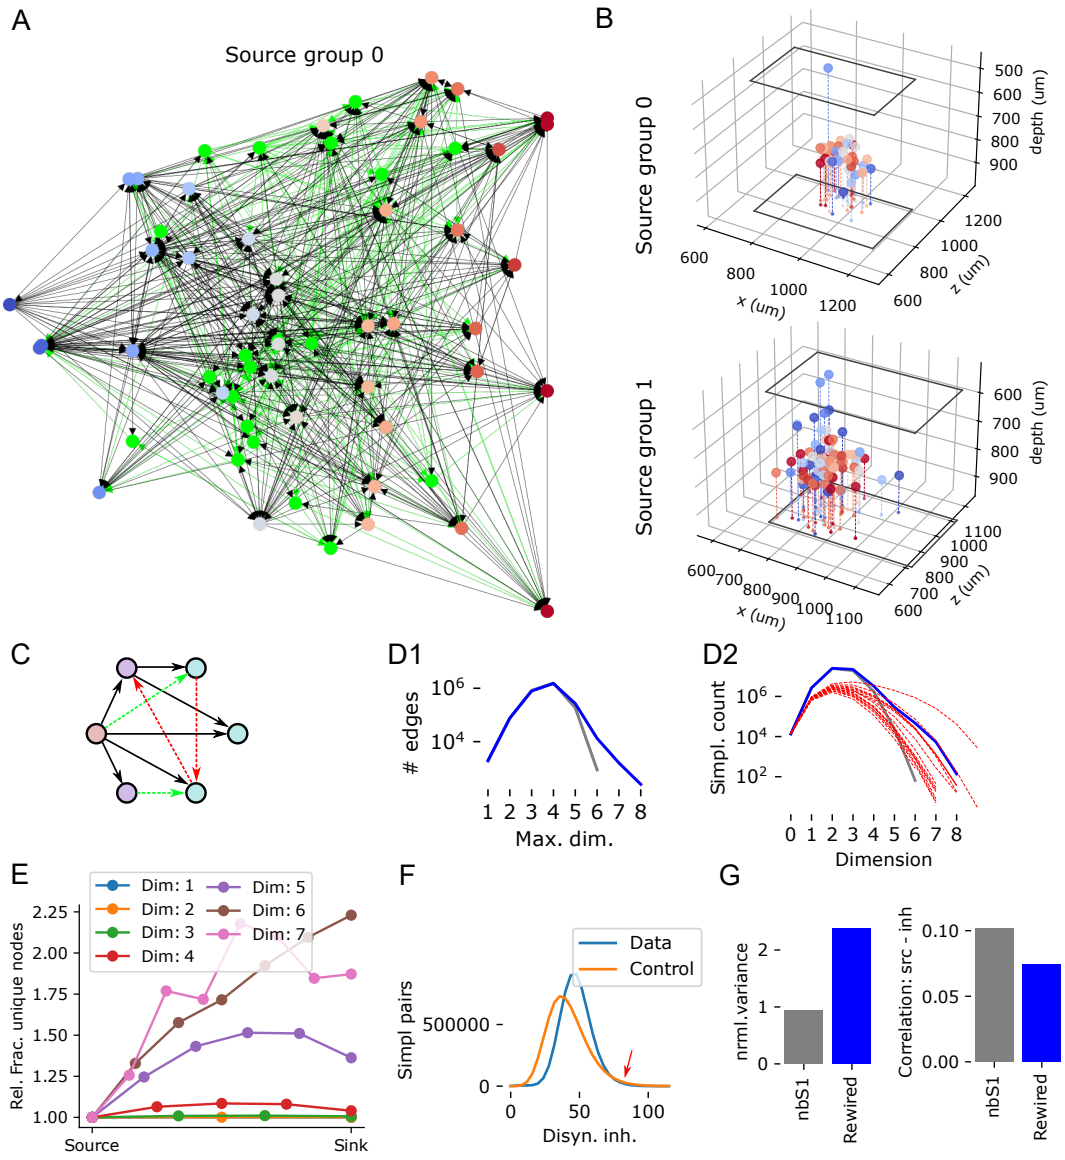

Figure 6. A rewiring rule to better capture the non-random structure of connectivity A: Exemplary connectivity of a source group (group 0 in Fig. 5B) Excitatory neurons that are member of any simplex in the group are placed along the horizontal axis and colored from blue to red, based on their mean position in a simplex. In green: 20 randomly picked inhibitory neurons connected to the simplex group. B: Spatial locations of neurons in source groups 0 and 1. Black outline indicates boundaries of the central subnetwork of MICrONS. See also Fig. S5). C: A rewiring rule removes edges participating only in low-dimensional simplices (red) and adds edges that are likely to form new high-dimensional simplices (green). D1: Maximum dimension of simplices an edge is participating in. Grey: nbS1 model; blue: nbS1 model with 1.2% of edges rewired. D2: Simplex counts as in Fig. 1D. Red dashed lines indicate results for the various MICrONS subnetworks. E: Divergence in terms of the number of unique neurons as in Fig. 2C for the rewired nbS1 model. F: Histogram of disynaptic inhibition strengths as in Fig. 4C for the rewired nbS1 model. Red arrow indicates a region on the right side of the distribution where the value for the control is above the data. G: Symmetry of disynaptic inhibition (left; as in Fig. 4D) and correlation between source side overlap and disynaptic inhibition strength (right; as in Fig. 4H), comparing the nbS1 model (grey) to its rewired version (blue).

3 DISCUSSION

We analyzed the connectivity of an electron microscopic reconstruction of a $1mm^3$ volume of mouse visual cortex. This is considered the gold standard in cortical connectomics due to the size of the EM-reconstructed volume encompassing whole dendritic morphologies and allowing the analysis of local connectivity. Using methods of algebraic topology suited for discovering higher order connectivity motifs, we found an underlying feed-forward network with specific lateral inhibition implementing a mechanism of competition between different outputs. In turn, we found the lateral inhibition is under the control of targeted disinhibition from a specialized class of inhibitory neurons. No aspect of the structure was evident when first-order connectivity between individual neuron pairs was considered; it required specific techniques from the field of algebraic topology (in this context: *neurotopology*) that consider connectivity in motifs of any size, i.e., beyond purely pairwise. This contrasts this work with previous analyses of the same dataset. Schneider-Mizell et al. (2023) found important pairwise principles of inhibitory connectivity at subcellular resolution. Ding et al. (2023) quantified how aspects of neuron function, such as correlations relate to pairwise connectivity. Here, we have demonstrated that increases of correlations are stronger when participation in directed simplices, i.e., densely connected feedforward motifs, is considered. To allow other researchers to conduct these analyses efficiently, we have open-sourced the code in the form of the `conntility` and `connalysis` python packages (Egas Santander et al. (2023); see data and code availability statement).

The non-random features of connectivity we found are possibly related to neural manifolds (Gallego et al., 2017; Stringer et al., 2019), the concept that circuit activity states exists and move around in a comparatively low-dimensional state. In this case, the state would be given by the combination of structurally determined source and target groups that are active, and the flow of information between them. Source and target groups also align with the concept of assemblies (Hebb, 1949; Harris et al., 2003; Dragoi and Buzsáki, 2006; Lopes-dos Santos et al., 2013), i.e., groups of neurons with correlated spiking activity. As neurons participating in a target group are excited as members of the same set of densely connected feed-forward motifs, they are likely to fire together. Within this structure, inhibition implements competition between different potential outputs. Such inhibition has been found to be crucial to achieve a balanced, asynchronous state (Renart et al., 2010), but is often considered to be non-specific blanket inhibition (Fino and Yuste, 2011; Packer and Yuste, 2011). In contrast, Rost et al. (2018) demonstrated that structure in inhibitory connectivity facilitates switching between assemblies without saturating their firing rates. Our results are more in line with the latter result, but additionally describe the internal structure of excitatory assemblies and specific positioning of inhibition within them. The principles of organization described are mostly orthogonal to the concept of information flow between layers, as all layers participate in all simplex positions. Still, our results were in line with established neuron roles: Both 4P and 5P_IT neurons prefer to act in a source position. This is in line with their established roles as initial inputs to local (4P) and distal (5P_IT) cortical circuits (Felleman and Van Essen, 1991; Harris et al., 2019). Our results reinforce the idea of layer 6 as an output layer.

Given all this, we can ask what the purpose of this complicated high-order structure is. And, if we assume that it aligns with the findings relating to the neural manifold, why does activity not explore larger parts of the possible activity space? First, of course the reduced network proposed in Fig. 5 is a dramatic simplification and the actual function of the local circuit is more complicated than anything that can be implemented with such a low number of units. But also, a cortical network with only one neuron per reduced unit would not be able to respond reliably to stimuli. This is largely because of the highly unreliable transmission of cortical synapses, especially *in vivo* (Faisal et al., 2008; Borst, 2010; Renart and Machens, 2014; Nolte et al., 2019), making population codes a necessity. Our findings give this idea more structural specificity: We found that neurons in the source position of high-dimensional simplex motifs are associated with subnetwork topologies that facilitate reliability, while neurons in the target position are associated with topologies that improve read-out of information from their spike trains (Egas Santander et al., 2023).

We compared the results we found in the MICrONS data to a recently released, biologically detailed model of cortical circuitry (Reimann et al., 2022; Isbister et al., 2023). For virtually all analyses, we found the same result: The non-random trends are present in the nbS1 model, but weaker. The nbS1

model captures connectivity trends arising from the dendrite and axon shapes of neurons, along with their placement in the local brain geometry, and additional biological data such as synapses per connection distribution moments and bouton densities (Reimann et al., 2015). However, it lacks additional specificity, such as the strong preference of subclasses of VIP positive interneurons to innervate almost exclusively PV or SST positive neurons (Pi et al., 2013; Reimann et al., 2022). Additionally, connectivity in the model is instantiated randomly within the constraints and unshaped by activity dependent structural plasticity. Our results suggest that important organizational principles of neuronal connectivity are already present in such a naive state (i.e., unaffected by plasticity), and are further enhanced by plasticity. We implemented a heuristic algorithm to approximate the effects of structural plasticity to rewire excitatory connections consistent with results of a recent analysis of the MICrONS data (Ecker et al., 2023), and studied how it affected the connectivity trends. We found that it moved the model closer to the MICrONS data for all analysis focusing on the excitatory subgraph only, but further away for analyses of the structure of inhibitory relative to excitatory connectivity. This could indicate a failing of the rewiring algorithm, or alternatively that it needs to be paired with inhibitory plasticity. This, alongside with our overall results indicates strong specificity of inhibitory connectivity, shaped by structural rewiring of connections to and from inhibitory neurons in cortex. Our findings thus reveal an intricate interplay between structural plasticity of recurrent excitatory motifs, inhibitory mediated competition and disinhibition, and provide important constraints for the development of biologically constrained theories of neocortical plasticity and learning *in vivo*.

REFERENCES

Barber, M. J. (2007). Modularity and community detection in bipartite networks. *Physical Review E*, 76(6):066102.

Bathellier, B., Ushakova, L., and Rumpel, S. (2012). Discrete Neocortical Dynamics Predict Behavioral Categorization of Sounds. *Neuron*, 76(2):435–449.

Bonald, T., Lara, N. d., Lutz, Q., and Charpentier, B. (2020). Scikit-network: Graph Analysis in Python. *Journal of Machine Learning Research*, 21(185):1–6.

Borst, J. G. G. (2010). The low synaptic release probability in vivo. *Trends in Neurosciences*, 33(6):259–266.

Carrillo-Reid, L., Miller, J.-e. K., Hamm, J. P., Jackson, J., and Yuste, R. (2015). Endogenous Sequential Cortical Activity Evoked by Visual Stimuli. *Journal of Neuroscience*, 35(23):8813–8828.

Curto, C., Geneson, J., and Morrison, K. (2019). Fixed Points of Competitive Threshold-Linear Networks. *Neural Computation*, 31(1):94–155.

Deco, G. and Hugues, E. (2012). Neural Network Mechanisms Underlying Stimulus Driven Variability Reduction. *PLoS Computational Biology*, 8(3):e1002395.

Ding, Z., Fahey, P. G., Papadopoulos, S., Wang, E. Y., Celii, B., Papadopoulos, C., Kunin, A. B., Chang, A., Fu, J., Ding, Z., Patel, S., Ponder, K., Muhammad, T., Bae, J. A., Bodor, A. L., Brittain, D., Buchanan, J., Bumbarger, D. J., Castro, M. A., Cobos, E., Dorkenwald, S., Elabbady, L., Halageri, A., Jia, Z., Jordan, C., Kapner, D., Kemnitz, N., Kinn, S., Lee, K., Li, K., Lu, R., Macrina, T., Mahalingam, G., Mitchell, E., Mondal, S. S., Mu, S., Nehoran, B., Popovych, S., Schneider-Mizell, C. M., Silversmith, W., Takeno, M., Torres, R., Turner, N. L., Wong, W., Wu, J., Yin, W., Yu, S.-c., Froudarakis, E., Sinz, F., Seung, H. S., Collman, F., Da Costa, N. M., Reid, R. C., Walker, E. Y., Pitkow, X., Reimer, J., and Tolias, A. S. (2023). Functional connectomics reveals general wiring rule in mouse visual cortex. preprint, Neuroscience.

Dragoi, G. and Buzsáki, G. (2006). Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies. *Neuron*, 50(1):145–157.

Ecker, A., Santander, D. E., Abdellah, M., Alonso, J. B., Bolaños-Puchet, S., Chindemi, G., Isbister, J. B., King, J. G., Kumbhar, P., Magkanaris, I., Muller, E. B., and Reimann, M. W. (2023). Long-term plasticity induces sparse and specific synaptic changes in a biophysically detailed cortical model. preprint, Neuroscience.

Egas Santander, D., Pokorny, C., Lazovskis, J., Santoro, M., Smith, J., Hess, K., Levi, R., and Reimann, M. W. (2023). Efficiency and reliability in biological neural network architectures. *In preparation*.

Faisal, A. A., Selen, L. P. J., and Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303.

Felleman, D. J. and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, 1(1):1–47.

Fino, E. and Yuste, R. (2011). Dense Inhibitory Connectivity in Neocortex. *Neuron*, 69(6):1188–1203.

Gal, E., London, M., Globerson, A., Ramaswamy, S., Reimann, M. W., Muller, E., Markram, H., and Segev, I. (2017). Rich cell-type-specific network topology in neocortical microcircuitry. *Nature Neuroscience*.

Gal, E., Perin, R., Markram, H., London, M., and Segev, I. (2019). Neuron Geometry Underlies Universal Network Features in Cortical Microcircuits. preprint, Neuroscience.

Gallego, J. A., Perich, M. G., Miller, L. E., and Solla, S. A. (2017). Neural Manifolds for the Control of Movement. *Neuron*, 94(5):978–984.

Harris, J. A., Mihalas, S., Hirokawa, K. E., Whitesell, J. D., Choi, H., Bernard, A., Bohn, P., Caldejon, S., Casal, L., Cho, A., Feiner, A., Feng, D., Gaudreault, N., Gerfen, C. R., Graddis, N., Groblewski, P. A., Henry, A. M., Ho, A., Howard, R., Knox, J. E., Kuan, L., Kuang, X., Lecoq, J., Lesnar, P., Li, Y., Luviano, J., McConoughey, S., Mortrud, M. T., Naeemi, M., Ng, L., Oh, S. W., Ouellette, B., Shen, E., Sorensen, S. A., Wakeman, W., Wang, Q., Wang, Y., Williford, A., Phillips, J. W., Jones, A. R., Koch, C., and Zeng, H. (2019). Hierarchical organization of cortical and thalamic connectivity. *Nature*, 575(7781):195–202.

Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, 424(6948):552–556.

Hebb, D. O. (1949). *The Organization of Behavior; A Neuropsychological Theory*. John Wiley & Sons, Inc., New York.

Isbister, J. B., Ecker, A., Pokorny, C., Bolaños-Puchet, S., Santander, D. E., Arnaudon, A., Awile, O., Natali, B.-Z., Alonso, J. B., Boci, E., Chindemi, G., Courcol, J.-D., Damart, T., Delemontex, T., Dietz, A., Ficarelli, G., Gevaert, M., Herttuainen, J., Ivaska, G., Ji, W., Keller, D., King, J., Kumbhar, P., Lapere, S., Litvak, P., Mandge, D., Muller, E. B., Pereira, F., Planas, J., Ranjan, R., Reva, M., Romani, A., Rössert, C., Schürmann, F., Sood, V., Teska, A., Tuncel, A., Van Geit, W., Wolf, M., Markram, H., Ramaswamy, S., and Reimann, M. W. (2023). Modeling and Simulation of Neocortical Micro- and Mesocircuitry. Part II: Physiology and Experimentation. preprint, Neuroscience.

Lagzi, F. and Rotter, S. (2015). Dynamics of Competition between Subnetworks of Spiking Neuronal Networks in the Balanced State. *PLOS ONE*, 10(9):e0138947.

Litwin-Kumar, A. and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*, 15(11):1498–1505.

Lopes-dos Santos, V., Ribeiro, S., and Tort, A. B. (2013). Detecting cell assemblies in large neuronal populations. *Journal of Neuroscience Methods*, 220(2):149–166.

Nolte, M., Gal, E., Markram, H., and Reimann, M. W. (2020). Impact of higher-order network structure on emergent cortical activity. *Network Neuroscience*, pages 1–36.

Nolte, M., Reimann, M. W., King, J. G., Markram, H., and Muller, E. B. (2019). Cortical reliability amid noise and chaos. *Nature Communications*, 10(1):3792. 7 citations (Crossref) [2021-04-01] Springer Nature https://www.biorxiv.org/content/10.1101/304121v2.

Packer, A. M. and Yuste, R. (2011). Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition? *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 31(37):13260–13271.

Perin, R., Berger, T. K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *PNAS - Proceedings of the National Academy of Sciences*, 108(13):5419–5424. 387 citations (Crossref) [2021-09-14] Publisher: National Academy of Sciences.

Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J. I., Huang, Z. J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. *Nature*, 503(7477):521–524.

Reimann, M. W., Bolaños-Puchet, S., Courcol, J.-D., Santander, D. E., Arnaudon, A., Coste, B., Delalondre, F., Delemontex, T., Devresse, A., Dictus, H., Dietz, A., Ecker, A., Favreau, C., Ficarelli, G., Gevaert, M., Herttuainen, J., Isbister, J. B., Kanari, L., Keller, D., King, J., Kumbhar, P., Lapere, S., Lazovskis, J., Lu, H., Ninin, N., Pereira, F., Planas, J., Pokorny, C., Riquelme, J. L., Romani, A., Shi, Y., Smith, J. P., Sood, V., Srivastava, M., Van Geit, W., Vanherpe, L., Wolf, M., Levi, R., Hess, K., Schürmann, F., Muller, E. B., Ramaswamy, S., and Markram, H. (2022). Modeling and Simulation of Neocortical Micro- and Mesocircuitry. Part I: Anatomy. preprint, Neuroscience.

Reimann, M. W., Horlemann, A.-L., Ramaswamy, S., Muller, E. B., and Markram, H. (2017a). Morpho-

logical Diversity Strongly Constrains Synaptic Connectivity and Plasticity. *Cerebral Cortex*, pages 1–16.

Reimann, M. W., King, J. G., Muller, E. B., Ramaswamy, S., and Markram, H. (2015). An algorithm to predict the connectome of neural microcircuits. *Frontiers in Computational Neuroscience*, 9.

Reimann, M. W., Nolte, M., Scolamiero, M., Turner, K., Perin, R., Chindemi, G., Dlotko, P., Levi, R., Hess, K., and Markram, H. (2017b). Cliques of Neurons Bound into Cavities Provide a Missing Link between Structure and Function. *Frontiers in Computational Neuroscience*, 11.

Renart, A., De La Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K. D. (2010). The Asynchronous State in Cortical Circuits. *Science*, 327(5965):587–590.

Renart, A. and Machens, C. K. (2014). Variability in neural activity and behavior. *Current Opinion in Neurobiology*, 25:211–220.

Renart, A., Moreno-Bote, R., Wang, X.-J., and Parga, N. (2007). Mean-Driven and Fluctuation-Driven Persistent Activity in Recurrent Networks. *Neural Computation*, 19(1):1–46.

Rost, T., Deger, M., and Nawrot, M. P. (2018). Winnerless competition in clustered balanced networks: inhibitory assemblies do the trick. *Biological Cybernetics*, 112(1-2):81–98.

Schneider-Mizell, C. M., Bodor, A., Brittain, D., Buchanan, J., Bumbarger, D. J., Elabbady, L., Kapner, D., Kinn, S., Mahalingam, G., Seshamani, S., Suckow, S., Takeno, M., Torres, R., Yin, W., Dorkenwald, S., Bae, J. A., Castro, M. A., Fahey, P. G., Froudakis, E., Halageri, A., Jia, Z., Jordan, C., Kemnitz, N., Lee, K., Li, K., Lu, R., Macrina, T., Mitchell, E., Mondal, S. S., Mu, S., Nehoran, B., Papadopoulos, S., Patel, S., Pitkow, X., Popovych, S., Silversmith, W., Sinz, F. H., Turner, N. L., Wong, W., Wu, J., Yu, S.-c., MICrONS Consortium, Reimer, J., Tolias, A. S., Seung, H. S., Reid, R. C., Collman, F., and Da Costa, N. M. (2023). Cell-type-specific inhibitory circuitry from a connectomic census of mouse visual cortex. preprint, Neuroscience.

Shi, D., Chen, Z., Sun, X., Chen, Q., Ma, C., Lou, Y., and Chen, G. (2021). Computing cliques and cavities in networks. *Communications Physics*, 4(1):249.

Sizemore, A. E., Giusti, C., Kahn, A., Vettel, J. M., Betzel, R. F., and Bassett, D. S. (2018). Cliques and cavities in the human connectome. *Journal of Computational Neuroscience*, 44(1):115–145.

Sizemore, A. E., Phillips-Cremins, J. E., Ghrist, R., and Bassett, D. S. (2019). The importance of the whole: Topological data analysis for the network neuroscientist. *Network Neuroscience*, 3(3):656–673.

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. *PLoS Biology*, 3(3):e68.

Stepanyants, A. and Chklovskii, D. (2005). Neurogeometry and potential synaptic connectivity. *Trends in Neurosciences*, 28(7):387–394.

Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., and Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):eaav7893.

The MICrONS Consortium, Bae, J. A., Baptiste, M., Bishop, C. A., Bodor, A. L., Brittain, D., Buchanan, J., Bumbarger, D. J., Castro, M. A., Celii, B., Cobos, E., Collman, F., Da Costa, N. M., Dorkenwald, S., Elabbady, L., Fahey, P. G., Fliss, T., Froudarakis, E., Gager, J., Gamlin, C., Gray-Roncal, W., Halageri, A., Hebditch, J., Jia, Z., Joyce, E., Joyce, J., Jordan, C., Kapner, D., Kemnitz, N., Kinn, S., Kitchell, L. M., Koolman, S., Kuehner, K., Lee, K., Li, K., Lu, R., Macrina, T., Mahalingam, G., Matelsky, J., McReynolds, S., Miranda, E., Mitchell, E., Mondal, S. S., Moore, M., Mu, S., Muhammad, T., Nehoran, B., Ogedengbe, O., Papadopoulos, C., Papadopoulos, S., Patel, S., Pitkow, X., Popovych, S., Ramos, A., Clay Reid, R., Reimer, J., Rivlin, P. K., Rose, V., Schneider-Mizell, C. M., Seung, H. S., Silverman, B., Silversmith, W., Sterling, A., Sinz, F. H., Smith, C. L., Suckow, S., Takeno, M., Tan, Z. H., Tolias, A. S., Torres, R., Turner, N. L., Walker, E. Y., Wang, T., Wanner, A., Wester, B. A., Williams, G., Williams, S., Willie, K., Willie, R., Wong, W., Wu, J., Xu, C., Yang, R., Yatsenko, D., Ye, F., Yin, W., Young, R., Yu, S.-c., Xenes, D., and Zhang, C. (2021). Functional connectomics spanning multiple areas of mouse visual cortex. preprint, Neuroscience.

ACKNOWLEDGMENTS

We thank Marwan Abdellah, Roberto Araya, Yamine Benabdesselam, Nghi Vuong Nguyen for helpful discussions.

This study was supported by funding to the Blue Brain Project, a research center of the École polytechnique fédérale de Lausanne (EPFL), from the Swiss government's ETH Board of the Swiss Federal Institutes of Technology.

E.B.M. was supported by funding from the Institute for Data Valorization (IVADO), the CHU Sainte-Justine Research Center (CHUSJRC), Fonds de Recherche du Québec–Santé (FRQS), the Canada CIFAR AI Chairs Program, the Quebec Institute for Artificial Intelligence (Mila), and Google. Compute infrastructure was supported in part through a grant of computing time to E.B.M. from the Digital Research Alliance of Canada.

DATA AND CODE AVAILABILITY

- **Electron microscopic connectomics data**. We obtained the EM connectivity data from the pages of IARPA MICrONS (The MICrONS Consortium et al. (2021), `https://www.microns-explorer.org/cortical-mm3`). A version of this data in a format that is optimized for the types of analyses presented here can be found under `https://doi.org/10.5281/zenodo.8364070`

- **Connectivity of a morphologically detailed model**. We obtained the morphologically detailed model of (Reimann et al., 2022; Isbister et al., 2023) from their release under `https://doi.org/10.5281/zenodo.8026353`. A version of this data in a format that is optimized for the types of analyses presented here can be found under `https://doi.org/10.5281/zenodo.10079406`.

- **Loading the morphologically detailed model**. To load data from the morphologically detailed model, we used the `bluepysnap` python package (`https://github.com/BlueBrain/snap`). It is published under `https://doi.org/10.5281/zenodo.10090777`.

- **Representation and basic analysis of connectomes**. Data loading, saving, filtering and basic analyses were conducted using the `conntility` python package (`https://github.com/BlueBrain/ConnectomeUtilities`). It is published under `https://doi.org/10.5281/zenodo.10059227`.

- **Topological analyses**. Detection of simplices and analyses related to them were conducted using the `connalysis` python package (`https://github.com/danielaegassan/connectome_analysis/`). It is based on *flagsercount* (`https://github.com/JasonPSmith/flagser-count`).

- **Custom analysis code**. Plots were generated with custom python code, in jupyter notebooks. The code can be found under `https://github.com/MWolfR/excitation_inhibition_topology`. It requires all the python packages mentioned in the previous points.

SUPPLEMENTARY INFORMATION

		Position						
		0	**1**	**2**	**3**	**4**	**5**	**6**
Class	**23P**	0.91%	9.2%	12.2%	11.9%	8.9%	8.8%	7.2%
	4P	31.5%	30.6%	33.6%	30.2%	29.6%	23.2%	23.0%
	5P_IT	25.7%	26.1%	18.5%	19.6%	8.5%	8.5%	7.1%
	5P_PT	33.8%	30.9%	30.6%	30.3%	35.2%	44.3%	39.1%
	6CT	7.7%	0.2%	0.0%	0.1%	0.1%	0.0%	0.0%
	6IT	0.3%	3.0%	5.0%	7.9%	17.7%	15.2%	23.%

Table 1. 6-dimensional simplex composition in terms of cell types.

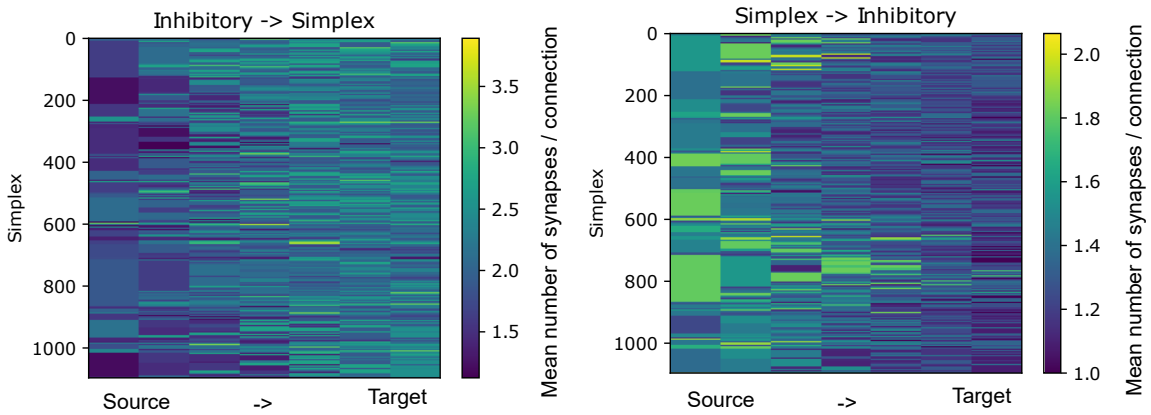

Supplementary Figure 1. Inhibitory connectivity from / to 6-d simplices as in Fig. 3C, but instead of the degree (number of connections) we show the number of synapses per connection.

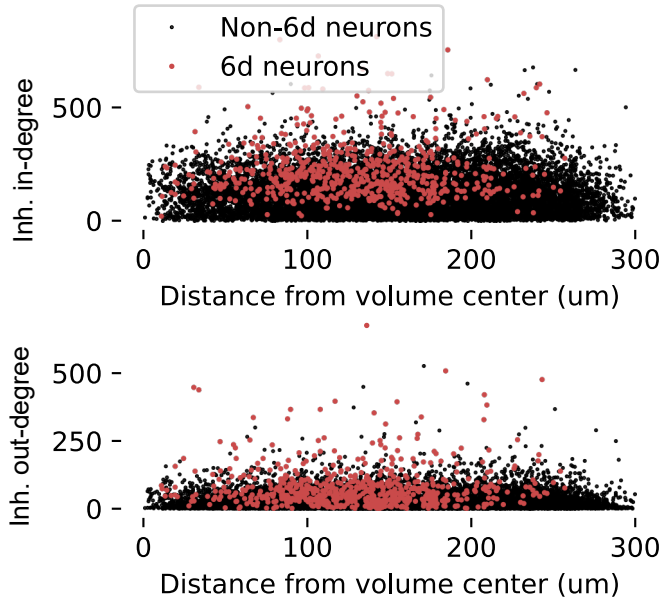

Supplementary Figure 2. Distance from the volume center against inhibitory in- and out-degree for neurons in the central subnetwork of MICrONS. For members of the 6-core (red) and non-members (black).

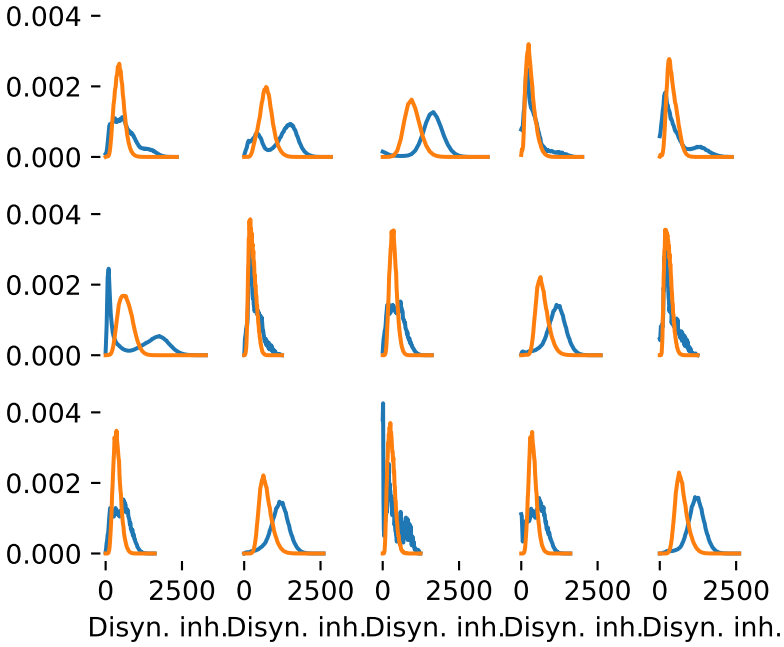

Supplementary Figure 3. Disynaptic inhibition strength comparing the data (blue) against a control (orange) as in Fig. 4C for all subnetworks of MICrONS.

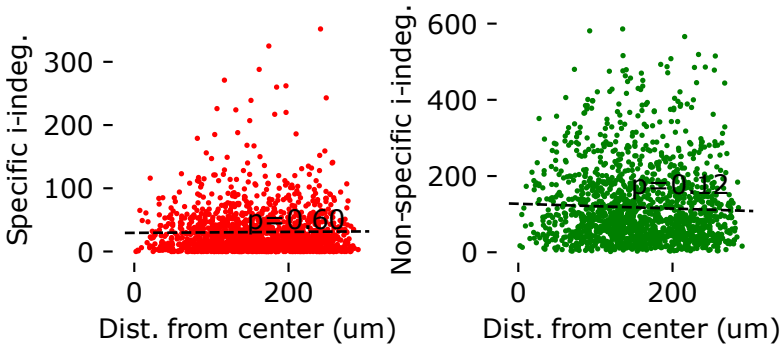

Supplementary Figure 4. Distance from the volume center against inhibitory in-degree of inhibitory neurons in the central subnetwork of MICrONS. Left: from inhibitory neurons with a significant targeting preference for other inhibitory neurons. Correlation is non-significant (pearsonr: 0.013; p=0.60) Right: from neurons without or with weak targeting preference. Correlation also non-significant (pearsonr: -0.04; p=0.12).

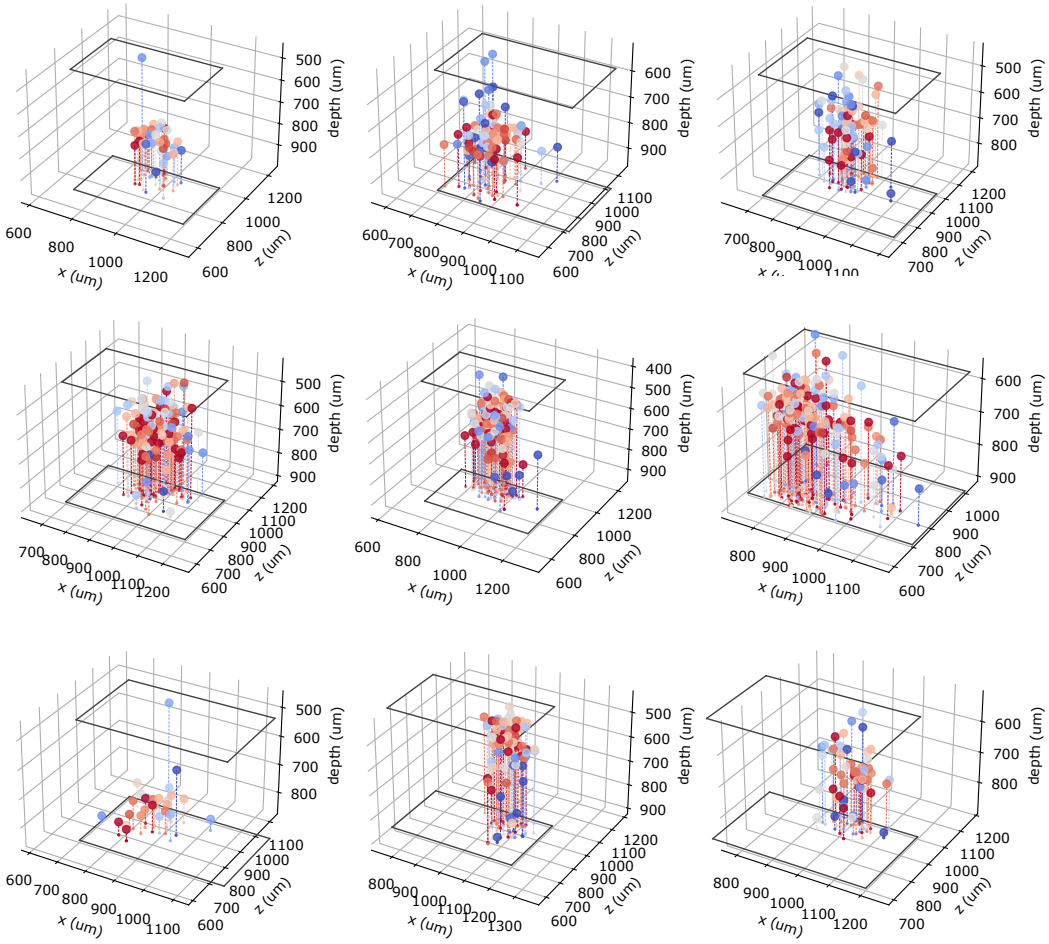

Supplementary Figure 5. Spatial locations of neurons in source groups as in Fig. 6B. Groups 0-8 of Fig. 5B ordered from left to right, and top to bottom.