Connectome-based models of feature selectivity in a cortical circuit

Victor Buendía,[1,2] Jacopo Biggioggera,[1,2] and Alessandro Sanzeni[1,3]

[1] *Department of Computing Sciences, Bocconi University, 20100 Milan, Italy*
[2] *These authors made equal contributions*
[3] *Lead contact, correspondence: alessandro.sanzeni@unibocconi.it*

Feature selectivity, the ability of neurons to respond preferentially to specific stimulus configurations, is a fundamental building block of cortical functions. Various mechanisms have been proposed to explain its origins, differing primarily in their assumptions about the connectivity between neurons. Some models attribute selectivity to structured, tuning-dependent feedforward or recurrent connections, whereas others suggest it can emerge within randomly connected networks when interactions are sufficiently strong. This range of plausible explanations makes it challenging to identify the core mechanisms of feature selectivity in the cortex. We developed a novel, data-driven approach to construct mechanistic models by utilizing connectomic data—synaptic wiring diagrams obtained through electron microscopy—to minimize preconceived assumptions about the underlying connectivity. With this approach, leveraging the MICrONS dataset [1], we investigate the mechanisms governing selectivity to oriented visual stimuli in layer 2/3 of mouse primary visual cortex. We show that connectome-constrained network models replicate experimental neural responses and point to connectivity heterogeneity as the dominant factor shaping selectivity, with structured recurrent and feedforward connections having a noticeable but secondary effect in its amplification. These findings provide novel insights on the mechanisms underlying feature selectivity in cortex and highlight the potential of connectome-based models for exploring the mechanistic basis of cortical functions.

INTRODUCTION

Neural responses exhibit selectivity for specific stimulus features across sensory areas, including edge orientation in the primary visual cortex [2], sound frequency in the primary auditory cortex [3], and odor identity in the piriform cortex [4]. Beyond sensory coding, selectivity also plays a crucial role in other brain computations, such as movement generation during both preparation and execution phases [5–7], as well as in encoding spatial representations of self and others [8, 9]. Selectivity for the same type of stimulus is observed across different brain areas, but it is often not redundant; rather, it is progressively transformed. In the primate ventral stream, neurons exhibit selectivity for visual stimuli that progresses from simple features, such as oriented edges in the primary visual cortex [2], to corners and junctions in the secondary visual cortex [10, 11], and ultimately to more complex features, such as objects and faces, in the inferotemporal cortex [12, 13]. This transformation is believed to enable the emergence of advanced computations, such as object categorization, in the ventral stream [14–16].

The emergence of feature selectivity in cortical circuits poses a significant computational challenge, as the large number of connections can compromise selectivity unless specific mechanisms are in place to preserve it. In networks receiving a large number of inputs from neuron populations with stimulus-dependent activity, random connections can dilute the stimulus dependence in the inputs, ultimately eliminating selectivity in the network's responses (see [17, 18] and derivation in Supp. Sec. I). To address how orientation selectivity emerges in the cortex, three main classes of models have been proposed. In feedforward models, selectivity is maintained by structured connections that combine inputs from neurons with specific tuning [2, 19]. In recurrent models, weakly selective inputs are amplified by specifically structured recurrent connections [20–22]. Lastly, in balanced models with unstructured connectivity, the large, non-selective input component is canceled by recurrent inhibition, enabling the smaller, selective input component to drive feature selectivity [17, 18]. While feedforward and recurrent models rely on carefully structured connections, balanced models suggest that selectivity can emerge even in networks with predominantly random connections. The diversity of these models highlights the complexity of cortical circuits and underscores the difficulty of determining the precise mechanisms that contribute to feature selectivity in the brain.

Mechanistic models of neural computations are typically built on strong assumptions about circuit structure, often with limited or no direct information about the underlying synaptic connectivity. Recent advancements in Electron Microscopy (EM) reconstruction techniques have made it possible to generate 'connectomes,' which are synaptic wiring diagrams, of large neural circuits or even entire nervous systems [1, 23–36]. Connectomes have introduced a new modeling paradigm, allowing for the direct derivation of connection structures from empirical data instead of relying on prior assumptions. For instance, connectomes have been used in building mechanistic models of retinal selectiv-

ity [37, 38], whitening of odor representations in the zebrafish olfactory bulb [39], and a range of computations in *Drosophila* [40–45]. Significant challenges remain due to the lack of certain crucial parameters in reconstructions, such as single-neuron properties and external inputs to circuits, highlighting the need for methods that address the limitations of connectome data [46]. Connectomic reconstructions of cortical circuits have only recently become accessible [1, 27, 32–34] and have seen limited application in mechanistic modeling approaches [34].

In this study, we use the recently obtained connectomic reconstruction of a 1 mm^3 region of the mouse visual cortex from the MICrONS Consortium [1] to investigate the mechanisms underlying the emergence of feature selectivity. We focus our analysis on understanding how orientation selectivity in layer 2/3 (L23) of mouse primary visual cortex (V1) emerges in response to its primary input source, layer 4 (L4). First, we demonstrate that reconstruction errors hinder a plug-and-play integration of this data into models. Then we introduce a novel approach for constructing mechanistic models of feature selectivity in cortical circuits, which integrates synaptic reconstruction, functional imaging, and theoretical modeling. We find that connectome-constrained models can reproduce key response properties observed experimentally and help to elucidate the mechanisms underlying feature selectivity in this circuit.

RESULTS

Approaches and limitations of constraining cortical feature selectivity with connectomes. To investigate how orientation selectivity develops in layer 2/3 of mouse V1 from layer 4 inputs, we examine the MICrONS dataset [1] through the lens of a well-established network model framework [47], constraining the model with the structural and functional properties extracted from the dataset.

We considered network models consisting of recurrently connected excitatory (E) and inhibitory (I) neurons, which represents layer 2/3 excitatory and inhibitory neurons (Fig. 1a,b). These networks were driven by an excitatory population (X), which models the excitatory neurons in layer 4 (Fig. 1a,b). In the model, the orientation θ of the stimulus determines the firing rates $r_i^X(\theta)$ of the X neurons, while the rates $r_i^{E,I}(\theta)$ are shaped by recurrent dynamics described by

$$\tau_A \frac{dr_i^A(\theta)}{dt} = -r_i^A(\theta) + \phi^A\left[\sum_{B\in\{E,I,X\}} \mu_i^{AB}(\theta)\right], \quad (1)$$

$$\mu_i^{AB} = \tau_A \sum_{j=1}^{N^B} J_{ij}^{AB} r_j^B(\theta) . \quad (2)$$

Here, $\phi^A(\cdot)$ and τ_A are the single-neurons response function (firing rate vs. input current) and the time constant of population $A \in \{E, I\}$, respectively. μ_i^{AB} is the current into neuron i from population $B \in \{E, I, X\}$, calculated as the sum of individual inputs $\tau_A J_{ij}^{AB} r_j^B$ from the j-th neuron in the presynaptic population, along with the synaptic efficacy J_{ij}^{AB}.

Analyzing the functional properties of neurons in the MICrONS dataset (see Methods), we found that a substantial proportion of excitatory neurons exhibited orientation selectivity, with 58% in layer 2/3 and 61% in layer 4, consistent with previous findings [48, 49]. Among these selective neurons, there was considerable heterogeneity in their tuning, as evidenced by the broad distribution of circular variance (Fig. 1c). We investigated if the model of Eq. (1), when constrained by functional and structural information, could account for the observed selectivity, with the aim of inferring the underlying mechanisms. The model posits that, in the stationary state, each neuron's firing rate and its synaptic inputs should be related by

$$r_i^E(\theta) = \phi^E\left[\mu_i^{EE}(\theta) + \mu_i^{EI}(\theta) + \mu_i^{EX}(\theta)\right] . \quad (3)$$

We tested this prediction by comparing the selectivity of the responses observed in the dataset to that expected from measured presynaptic inputs.

We focused on one key aspect of selectivity, preferred orientation, as this is independent of neuron properties not constrained by the recording (e.g., ϕ and τ). For each selective E neuron, we compared its preferred orientation to the angle at which its presynaptic current is maximum, estimated as follows. Synaptic currents depend on the responses of presynaptic neurons and synaptic strength (Eq. (1)). We extracted presynaptic responses from the functional dataset, while synaptic strengths were inferred from the structural dataset using synapse size, which are proportional to synaptic efficacy [50], as an estimator. This approach cannot be applied to the I synaptic currents, as their activity was not recorded in the dataset [1]. However, since inhibitory neurons are less tuned [48], and inhibitory input selectivity is equal to or broader than excitatory input selectivity [51–53], we reasoned that I synaptic currents would not modify the preferred orientation. Thus, by combining structural and functional data with Eq. (3), we estimated the excitatory synaptic currents, up to a proportionality constant that does not affect our estimate of the angle maximizing them.

Fig. 1d illustrates the comparison between the preferred orientation measured in an example neuron and the one estimated from its presynaptic inputs using Eq. (3) (see Methods for details). The total excitatory current ($E+X$), estimated by combining presynaptic responses with synaptic strengths peaked at $\theta = 0$, which aligned with the postsynaptic neuron's preferred orientation and confirmed the prediction from Eq. (3). For other neurons (e.g., Fig. 1e), we observed a significant discrepancy between the measured preferred orien-

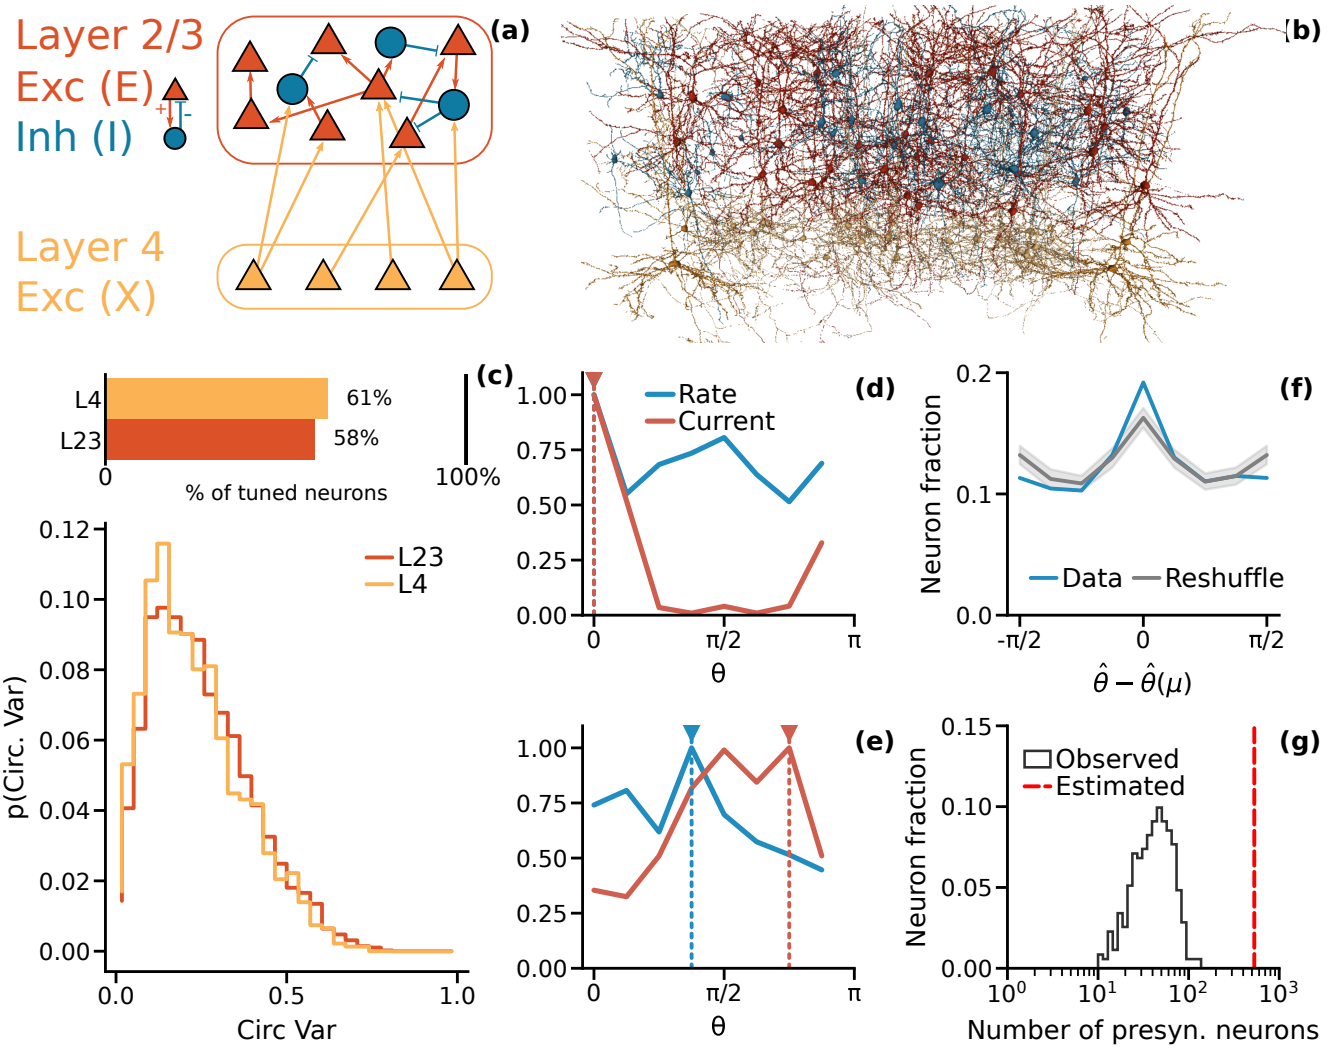

FIG. 1. **Challenges in constraining cortical models with connectomes.** (**a**) Schematic representation of the network model of Eq. (1). (**b**) Electron microscopy reconstruction of example neurons in layer 2/3 and 4 of mouse V1 from the MICrONS dataset [1]. (**c**) Fraction of selective excitatory neurons (top) and the distribution of circular variance for selective neurons (bottom) in each layer. This work focuses on developing a mechanistic understanding of how selectivity in layer 2/3 arises from network connectivity and inputs originating in layer 4. (**d**) Selectivity of a single neuron from [1] (top, blue) and its synaptic current estimated using Eq. (1) (top, red). The preferred orientation measured and predicted by the synaptic current match. (**e**) As in (d), but for a different example neuron. Combining the synaptic current did not predict the postsynaptic preferred orientation. (**f**) Difference between the preferred orientation observed ($\hat{\theta}$) and the preferred orientation predicted from synaptic inputs ($\hat{\theta}(\mu)$) , for all selective layer 2/3 excitatory neurons in [1]. Model's predictions (green) only slightly improve over a control obtained by shuffling connectivity (gray), highlighting critical issues in the approach. (**g**) In-degree distribution of layer 2/3 excitatory neurons observed in [1] (black) and estimated by analyzing manually proofread neurons (red). Reconstruction errors prevent comprehensive knowledge of the input received by neurons, hindering the direct use of Eq. (1) in studying mechanisms of feature selectivity.

tation and that estimated from the combined synaptic inputs. When considering all functionally matched excitatory neurons in layer 2/3 of the dataset (4167 neurons), we found that the measured preferred orientation generally aligns with the one estimated from these inputs (Fig. 1g). However, surprisingly, this alignment represented only a minor increase compared to a shuffled control (from 17% to 19%; Fig. 1f), suggesting that the connectivity information extracted from the connectome provided little information about the synaptic inputs received by these neurons. These results indicate a fundamental issue in integrating the model presented in Eq. (3) with the connectomic data from [1].

Modern algorithms for automatic reconstruction of connectivity often experience errors [54], which can distort the wiring diagrams of neural circuits and affect the

accuracy of our synaptic current estimations. To assess the extent of these reconstruction errors and their potential effect on our analysis of the MICrONS dataset, we analyzed the number of presynaptic neurons connected to layer 2/3 excitatory neurons, i.e. the in-degree of the network. When considering all neurons in the analysis, we found an average of 28.08 ± 0.17 (mean±sem) inputs, with half of the neurons receiving fewer than 23 inputs. In order to obtain a better estimate, we repeated the same computation only for post-synaptic neurons whose dendrites have been manually proofread. For this case the in-degree slightly increased to 44.2 ± 1.1 (see full distribution in Fig.1g), but was still significantly lower than previous estimates (e.g., see discussion in [55]) and lower than our estimation based on the connection statistics of fully proofread neurons (Fig. 1g, details in Methods).

This finding suggests that the deviations from the predictions of Eq. (3), as shown in Fig. 1f, may result from reconstruction errors that hinder accurate estimation of synaptic current. Addressing these issues by increasing the number of proofread neurons would require substantial effort, emphasizing the need for an alternative approach.

An alternative approach to integrating connectomes into cortical models. The challenges we encountered in computing synaptic currents for individual neurons highlight the need for methods that effectively utilize information from the MICrONS dataset [1], and other similar or larger connectomic reconstructions anticipated in the near future [56], to develop mechanistic models of cortical functions. Theoretical models suggest that neural activity in cortical circuits is primarily governed by the statistics of connections and presynaptic activity, rather than specific neuron-to-neuron interactions [57–61]. Mean-field models considering only the first two moments of firing rates and synaptic strengths are able capture the distribution of baseline firing rates in cortical circuits [59] as well as cortical response properties [20–22]. Building on these findings, we developed an approach focused on the statistical properties of neural responses and synaptic connections, rather than individual characteristics.

We characterized the statistical properties of neural responses using functional data. As noted above, a large fraction of neurons in layers 2/3 and 4 exhibited orientation selectivity (Fig. 1c). The responses of these selective neurons were identified by their preferred orientation and by how their activity changed with stimulus orientation. We found an overrepresentation of preferred orientations at 0 and $\pi/2$ among these neurons (Supp. Fig. 2), a pattern that aligns with previous large-scale recordings [62] and neural representations adapted to natural image statistics [63]. On average, selective neurons in layer 2/3 and 4 had similar tuning curves (Fig. 2b), with substantial heterogeneity observed in individual responses (Fig. 2a and Supp. Fig. 2). Response of non-selective neurons were also heterogeneous, mirroring the variability seen in selective neurons at their antipreferred orientation (Supp. Fig. 2).

We characterized the connectivity statistics by integrating functional and structural data through two key observables: connection probability, which quantifies the proportion of non-zero elements in the connectivity matrix, and the synaptic strength associated with these non-zero connections. The strength of synapses between connected pairs showed significant heterogeneity, characterized by an approximately log-normal distribution (Fig. 2c and Supp. Fig. 2), aligning with previous estimates [64, 65]. The distributions of synaptic weights from layers 2/3 and 4 exhibited a statistically significant difference ($p \sim 10^{-38}$; KS test), though the difference was extremely small, with the distributions essentially overlapping (Fig. 2c). Additionally, we found no significant dependence of synaptic strength on tuning ($p = 0.88$; KS test, Supp. Fig. 2). Connection probability is significantly distorted by reconstruction errors, but these errors should be approximately independent of neuron types and selectivity. Consistent with this, we found that connection probabilities varied significantly with proofreading status, but differences could mostly be explained by a common multiplicative factor (Supp. Fig. 3). Accordingly, we separated the inference of absolute connection probability-based solely on a limited sample of neurons with extended proofreading (60 neurons)-from the analysis of how connection probability varied across neuron types and selectivity differences, which was measured using a larger sample with partial proofreading (566 neurons). This approach enhanced our statistical power, demonstrating that connection probability for inputs from both layers 2/3 and layer 4 decreased as the difference in preferred orientation between pre- and postsynaptic neurons increased (Fig. 2d,e). This 'like-to-like' connectivity structure aligns with results from other experimental studies [27, 53, 66–69]. The connection probability between selective and non-selective neurons had similar magnitude (Supp. Fig. 2), indicating that a substantial proportion of synaptic inputs to selective neurons originated from neurons lacking orientation selectivity. This finding aligns with previous results from *in vivo* imaging combined with rabies monosynaptic retrograde tracing [68].

Synaptic current statistics challenge the role of feedforward inputs and connectivity structure in shaping selectivity. To constrain mechanisms shaping selectivity in layer 2/3 of mouse V1, we analyze synaptic current structures by combining neural response statistics and connectivity data from the MICrONS dataset [1]. Specifically, we constructed a model neuron population, where each neuron's inputs were computed through Eq. (3) with connection strengths and presynaptic inputs sampled from the recorded populations according to the inferred data statistics. To construct this population, we assigned each model neuron a target preferred orientation that determined the statistics of the connections it made with the pool of recorded

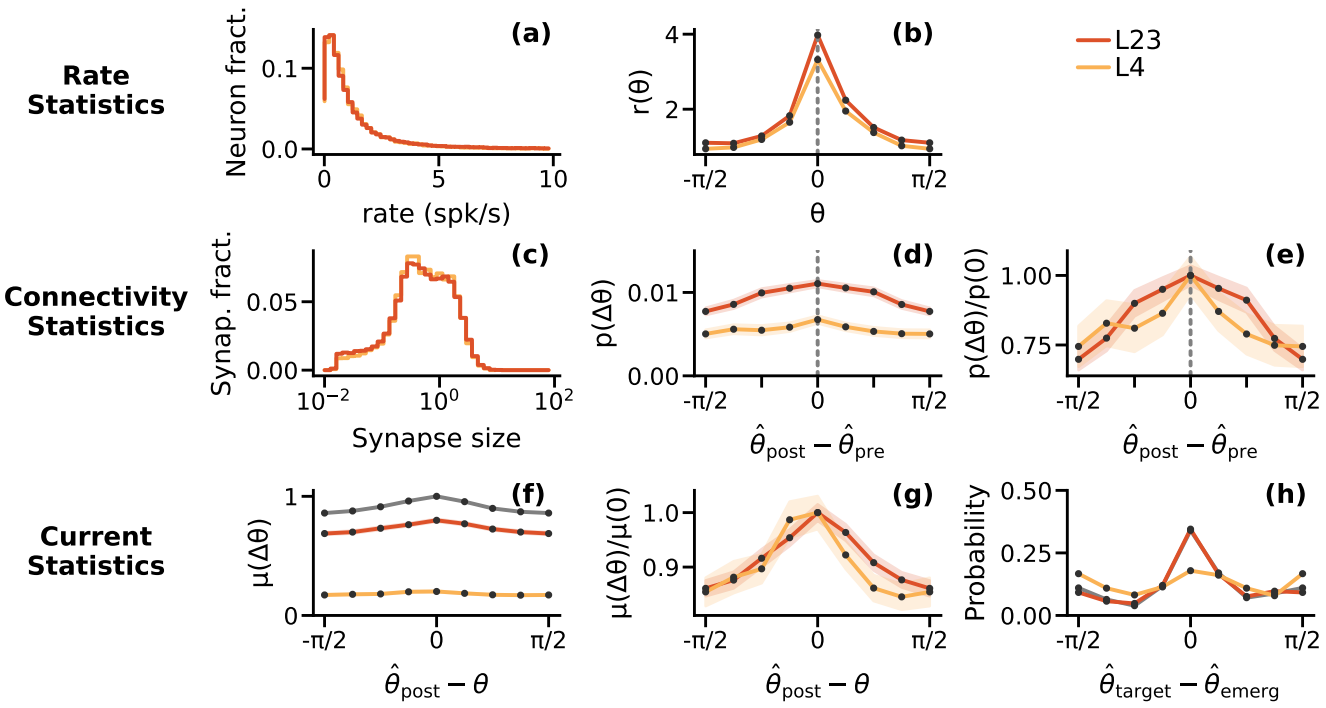

FIG. 2. **Statistical properties of the circuit inferred from connectomic data.** (**a**) Distribution of firing rates in layers 2/3 and 4. (**b**) Average tuning curves of selective neurons in layers 2/3 and 4. (**c**) Distribution of normalized synaptic sizes for connections to layers 2/3 E neurons, originating from layers 2/3 and 4 E neurons. (**d**) Probability of connections as a function of selectivity difference between pre- and postsynaptic neurons, for connections within layer 2/3 and connections from layer 4 to layer 2/3. (**e**) Same as (d), but normalized to the value at a selectivity difference of 0. (**f**) Average feedforward (orange), recurrent (red), and total (gray) excitatory synaptic currents as a function of the difference between the postsynaptic neuron's preferred orientation ($\hat{\theta}_{\text{post}}$) and the stimulus orientation (θ). Currents are normalized so that the peak of total current is one. (**g**) Same as (f), but normalized to the value at a selectivity difference of 0. (**h**) Distribution of angles at which peak current occurs, based on measurements of feedforward (orange), recurrent (red), and total (gray) excitatory synaptic currents.

neurons. This approach accounts for the fact that the statistics of connections depend on both presynaptic and postsynaptic preferred orientations. We then analyzed the relationship between the postsynaptic preferred orientations and presynaptic currents in the population of model neurons.

To probe the relative contribution of feedforward and recurrent inputs in shaping selectivity, we first analyzed the average synaptic currents received by model neurons as a function of the difference between the target preferred orientation of the postsynaptic neuron and the orientation of the presented stimulus. In the model neurons, the total synaptic current peaked at the target preferred orientation of the postsynaptic neurons (Fig. 2f,g), as expected by design. Interestingly, synaptic current selectivity was primarily driven by inputs from layer 2/3, with layer 4 providing a significantly smaller contribution. This difference stemmed from the considerably higher number of inputs from layer 2/3 compared to layer 4, which we estimate to be 400 versus 150, respectively. Although presynaptic selectivity, connection strength, and the decay of connection probability with selectivity were similar across layers, the overall connection probability and the number of presynaptic neurons were both lower in layer 4 compared to layer 2/3. This result calls into question the extent to which models based on a feedforward input structure [2, 19] contribute to the emergence of orientation selectivity in layer 2/3 of mouse V1.

To further test the role of connectivity in shaping selectivity, we examined the relationship between each neuron's target orientation and the orientation at which the synaptic current it receives is maximized, which we termed the emergent preferred orientation. We reasoned that if the structure of connectivity were the sole determinant of selectivity, as in structured models [2, 19–22], we would expect the target and emergent preferred orientations to align. Conversely, if randomness were the dominant contributor, as in balanced models [17, 18], we would expect the two orientations to be independent. In our population of model neurons, we found that 35% exhibited emergent selectivity matching their target orientation, while the remaining 65% showed a difference between the two (Fig. 2g). This indicates that heterogeneity in presynaptic response selectivity and connection strength plays a major role in shaping overall selec-

tivity. The orientation that maximized the synaptic current from layer 2/3 neurons exhibited behavior similar to the total current in matching the target orientation, indicating that the selectivity of individual neurons was primarily influenced by inputs from layer 2/3 neurons (Fig. 2g). In contrast, for layer 4 inputs, the orientation that maximized the current matched the target orientation only about 18% of the time. This difference between layer 2/3 and layer 4 neurons arises from the smaller number of inputs from layer 4, making the relative contribution of random fluctuations more significant.

In summary, our findings suggest that while layer 4 inputs contribute modestly to shaping selectivity, the structure of recurrent connections within layer 2/3, along with the heterogeneity of individual neurons and connections, serves as a primary driver of selectivity. However, these results are based on model neurons that passively receive only excitatory inputs without influencing the activity of other neurons in the network. In a recurrent network, the observed discrepancy between the target and emergent orientations may be influenced by interactions among neurons and the presence of inhibition. To assess the role of these effects, we analyzed recurrent dynamics in network models constrained by connectome data.

Connectome-constrained recurrent network models replicate the statistics of neural responses. To investigate how recurrent interactions shape selectivity in layer 2/3 of mouse V1, we examined the dynamics of network models described by Eq. (1), constrained by connectomic data from [1]. While experimental data provide essential constraints on certain aspects of network structure, key properties of individual neurons and connectivity remain undetermined. In what follows, we outline the model parameters constrained by the data and those that remain undetermined, then present an approach to infer the latter.

Dynamics of Eq. (1) is shaped by the single-neuron response function, which is not directly constrained by structural or functional data in [1]. To address this limitation, we drew on prior works demonstrating that the response function of leaky integrate-and-fire neurons driven by white noise [70] can quantitatively capture single-neuron and network responses in the cortex [71–76]. We utilized this response function, while leaving open four single-neuron parameters that govern the gain and offset for excitatory and inhibitory neuron responses.

As discussed in previous sections, the connectivity matrices J_{ij}^{AB} in the model of Eq. (1) cannot be directly measured from the data. To address this, we assumed that the connectivity matrices in the biological circuit were samples from ensembles of random matrices, with statistics–connection probability and synaptic strength distributions–depending solely on the cell type and selectivity of the pre- and postsynaptic neurons. These assumptions enabled us to construct statistical models suitable for inference from the data, as outlined below.

In the data, connection probability between selective neurons declines as the difference in preferred orientation increases (Fig. 2d,e). To incorporate this into the model, we assigned each neuron a target preferred orientation, forming connections based on the orientation difference between pre- and post-synaptic neurons, similar to approaches in structured models [2, 19–22]. In doing so, we ensured that the fraction of neurons with each preferred orientation matched that observed in the data. In the recurrent network, as in our synaptic current analysis (Fig. 2h), heterogeneity is expected to create discrepancies between a neuron's target preferred orientation and the emergent preferred orientation resulting from network dynamics. The orientations measured in the data correspond to emergent preferred orientations, not the target orientations, and thus the results in (Fig. 2d,e) constrain connection probability estimates based on emergent orientation differences. However, for model construction, we need the probability as a function of target orientation differences, which cannot be directly measured. To address this gap, we parameterized the target-based probability using adjustable functions, calibrated to reflect the observed dependency on emergent orientation differences (see Methods). We used selectivity-dependent connection probabilities for both recurrent connections among E neurons and for feedforward connections from X to E neurons. In contrast, inhibitory neurons were not assigned a target preferred orientation and were connected randomly to other neurons, with connection probabilities based solely on structural data (Supp. Fig. 2).

The strength of connections was sampled randomly based on experimental distributions of synaptic size. Since we did not observe any dependency of size on selectivity (Supp. Fig. 2), we assumed that connection strength depended solely on the types of pre- and post-synaptic neurons. This procedure required specifying two free parameters in the model: one for the average connection strength of excitatory neurons and another for inhibitory neurons, both acting as proportionality constants between synaptic size and connection strength.

Considering both single-neuron and network properties, the model included eight free parameters. We systematically examined how these parameters affected network responses to identify configurations that matched experimental data. Best-fit solutions were found by using simulation-based inference methods [77] (see Methods). Specifically, we focused on the mean response and selectivity of excitatory neurons in layer 2/3, as well as connection probability as a function of differences in emergent preferred orientation. This approach produced a family of solutions that successfully replicated these features with physiologically plausible parameters (Fig. 3a-e). Notably, these models often exhibited structured chaotic dynamics, with chaotic fluctuations in single-neuron activity occurring around an

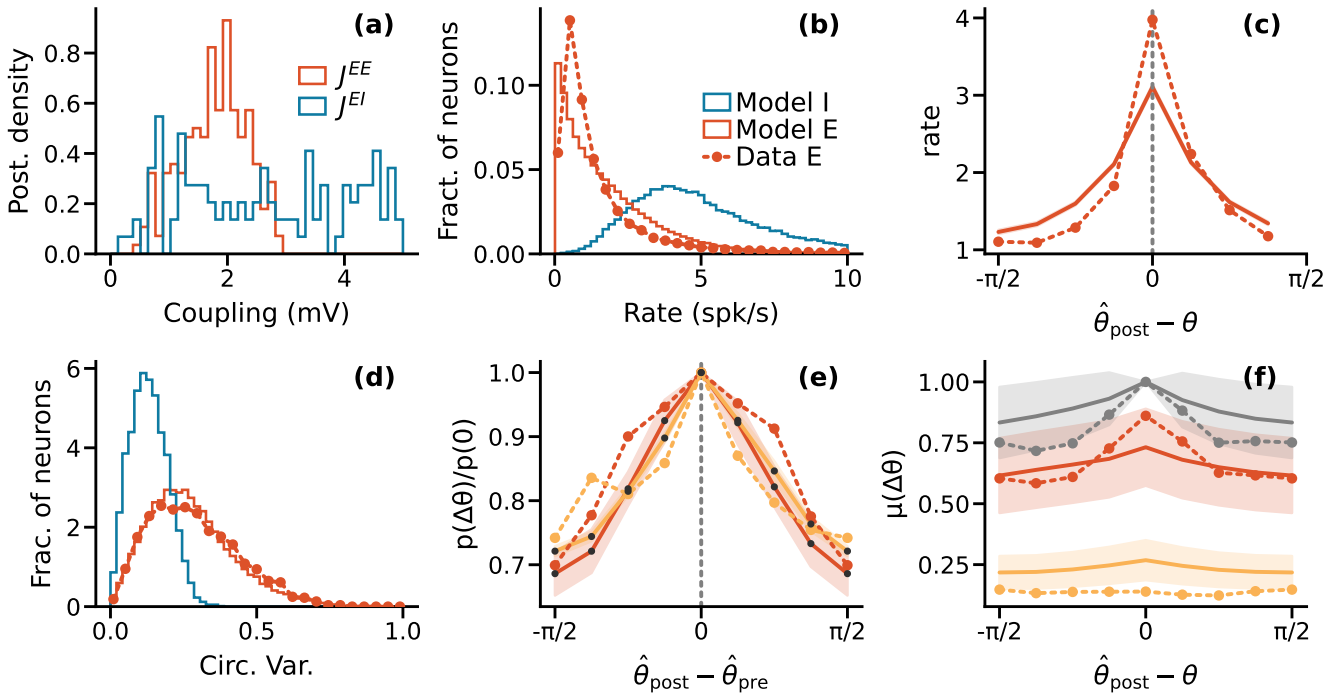

FIG. 3. **Comparison of circuit properties: experimental data vs. connectome-constrained recurrent network model.** (**a**) Posterior distributions of J^{EE} and J^{EI} connection strengths in models matching experimental response properties. These are only two of the eight free parameters. (**b**) Firing rate distributions. (**c**) Average excitatory neuron tuning curves. (**d**) Circular variance distributions. (**e**) Connection probability vs. emergent preferred orientation difference. (**f**) Synaptic currents (feedforward, recurrent, total) vs. postsynaptic preferred orientation ($\hat{\theta}_{\text{post}}$) relative to stimulus orientation (θ). All panels compare the best-fit model (continuous lines) to experimental data (dashed lines with markers).

orientation-dependent average firing rate. They also captured key circuit characteristics not explicitly constrained by the fitting procedure, including broad firing rate distributions and higher firing rates in inhibitory versus excitatory neurons (Fig. 3b). Selective neurons exhibited heterogeneous tuning curves, with a broad distribution of circular variance comparable to experimental data (Fig. 3d). In line with previous observations [48], inhibitory neurons showed lower selectivity (Fig. 3d). Additionally, the selectivity of excitatory synaptic currents aligned with experimental findings, as inputs from layer 4 displayed weaker tuning than neurons in layer 2/3 (Fig. 3f). These results suggest that the emergence of selectivity in our models mirrors mechanisms driving selectivity in the biological circuit.

Probing the role of connectivity structure in shaping selectivity in recurrent networks. In models that accurately described the experimental data, structure was imposed through our construction procedure: each neuron was assigned a target preferred orientation, which influenced both the connections it received and those it formed with other neurons. Heterogeneity in the network are expected to create discrepancies between a neuron's target preferred orientation and the emergent preferred orientation resulting from network dynamics. To quantify this, we measured the distribution of orientation differences in the simulated models. The distribution was peaked at zero, indicating that, on average, the emergent orientation matched the target (Fig. 4a). However, a large fraction of neurons showed an emergent preferred orientation different from the target, consistently with our analysis of synaptic currents (Fig. 2h). Recurrent interactions did not improve the alignment; in fact, they slightly decreased it. This could be due to the additional heterogeneity introduced by inhibitory inputs and the deviations of emergent preferred orientation in layer 2/3 from the target orientations in presynaptic inputs.

To understand how connectivity structure influences selectivity in the circuit, we examined the effects of removing it, reasoning that any changes in network responses would reveal which aspects of dynamics are shaped by the structure. We modified the network model by randomizing the connections while keeping all other parameters at their best-fit values. The simulations were then rerun to measure the network responses. Shuffling the connections weakly altered the distribution of rates in the network (Fig. 4b). However, as expected, it completely disrupted the relationship between target and emergent orientations (Fig. 4a). The average circular variance in the circuit decreased by 27%

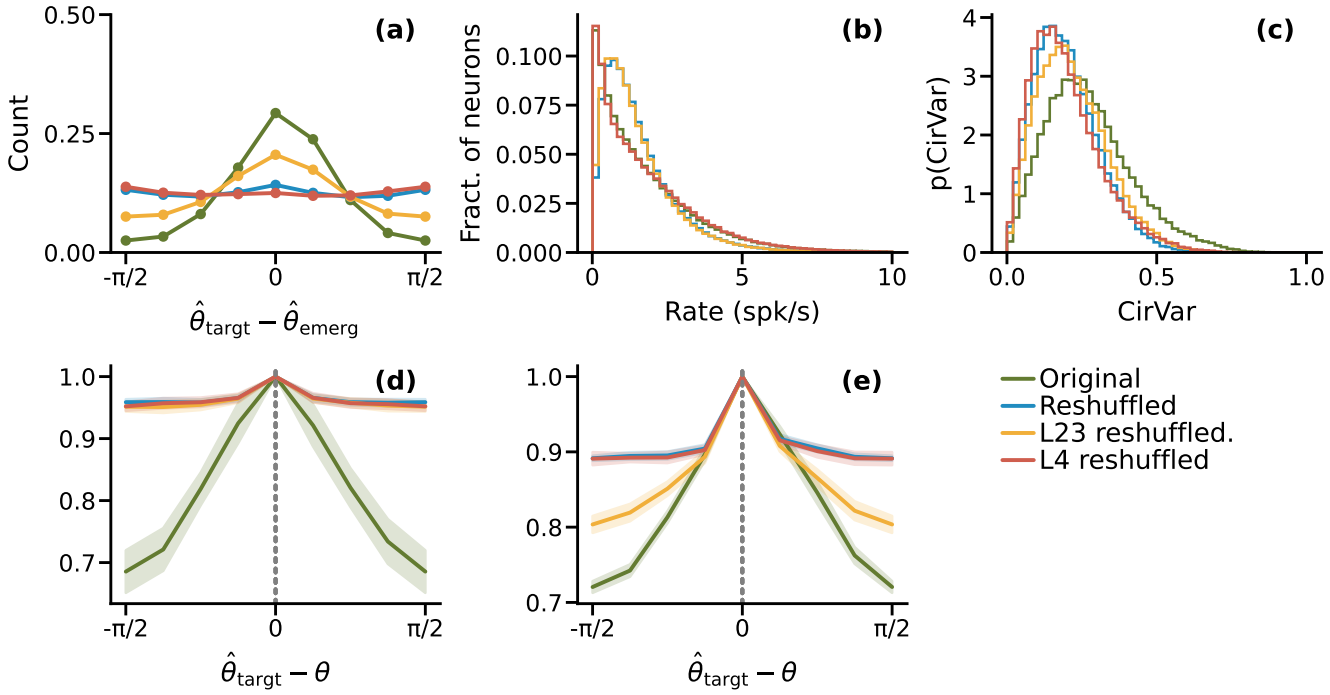

FIG. 4. **Effects of shuffling connections on connectome-constrained recurrent network models.** (**a**) Difference between target and emergent preferred orientation. This and the following panels compare the best-fit model to models with shuffled connectivity (all, recurrent only, feedforward only). (**b**) Firing rate distributions. (**c**) Circular variance distributions. (**d**) Connection probability between layer 2/3 neurons vs. emergent preferred orientation difference. (**e**) Same as **d**, but for connections from layer 4 to layer 2/3.

(Fig. 4c), indicating that while the connectivity structure contributes significantly to selectivity, the majority of selectivity arises from random connections. Changes in the circular variance distribution revealed that this decrease resulted from a collective reduction in single-neuron circular variance, rather than a reduction in the fraction of selective neurons. This suggests that the primary role of structure is to enhance single-neuron selectivity.

How is connectivity structure enhancing selectivity in the circuit? Shuffling connections significantly reduced the dependency of connection probability on emergent orientation differences (Fig. 4d,e), suggesting that, in structured networks, neurons with similar emergent orientations are more likely to excite each other, enhancing selectivity. To test this hypothesis, we examined the specific roles of feedforward and recurrent connectivity by shuffling these connections independently. Shuffling feedforward connections alone disrupted the relationship between emergent and target orientations almost as effectively as shuffling all connections, while shuffling recurrent connections alone was substantially less disruptive (Fig. 4a). This indicates that feedforward connectivity primarily determines the preferred orientations. Aligning target and preferred orientations is essential to strengthen like-to-like connectivity within the network. In fact, shuffling feedforward connections was as effective as, or even more effective than, shuffling recurrent connections in reducing the dependency of connection probability on emergent orientation differences (Fig. 4d,e). Finally, shuffling feedforward connections was as effective as shuffling all connections in reducing circular variance (30%, Fig. 4c), whereas shuffling recurrent connections had an intermediate effect (19%, Fig. 4c). These findings support the idea that feedforward connectivity plays a critical role in aligning emergent preferred orientations with the circuit's structural organization, while recurrent connections further amplify selectivity.

DISCUSSION

Feature selectivity is a widely observed phenomenon across cortical circuits [2–16], and the mechanisms underlying its emergence remain a topic of ongoing debate [2, 17–22]. In this study, we introduced a novel approach to building mechanistic models of cortical function that leverages newly available connectomic data [1], aiming to minimize the assumptions typically required in model construction. We applied this approach to investigate the mechanisms underlying the emergence of orientation selectivity in layer 2/3 of mouse V1, driven by inputs from layer 4. Our analysis of synaptic input statistics reveals that recurrent interactions within layer

2/3 play a more substantial role in driving selectivity than inputs from layer 4, with connection heterogeneity emerging as a primary factor. Network dynamics in models constrained by both functional and structural data further support the idea that heterogeneity is the principal driver of selectivity. Additionally, our results highlight the interplay between structured recurrent and feedforward connectivity in amplifying this selectivity.

The exploration of feature selectivity has long been a central focus in neuroscience, with orientation selectivity being one of the most extensively studied phenomena. Various theoretical models have been proposed to explain this, relying either on structured feedforward [2, 19] or recurrent [20–22] connections, or unstructured networks with strong recurrent interactions in an inhibition-dominated regime [17, 18]. Our findings suggest that the latter class of models is the main contributor to orientation selectivity in layer 2/3 of mouse V1. However, the picture that emerges from our analysis is more nuanced. We observed like-to-like feedforward and recurrent connectivity, consistent with previous experimental data [27, 53, 66–68] and with assumptions of structured models [2, 19–22]. While the majority of orientation selectivity in the circuit does not rely on this structure, it nonetheless plays a significant role in amplifying selectivity. Two key factors contribute to this: first, structured feedforward connectivity biases neurons' preferred orientations to align with the network's recurrent structure, and second, structured recurrent connectivity promotes excitation between neurons with similar preferred orientations, amplifying selectivity.

Our analysis underscores the importance of a quantitative approach to circuit dynamics that incorporates both structured connectivity and heterogeneity in model construction. This data-driven analysis was enabled by the MICrONS dataset [1], which provided joint data on neuronal selectivity and synaptic connection properties across a large number of neurons and synapses. We emphasize that our findings are specific to the problem studied—orientation selectivity in layer 2/3 driven by layer 4 inputs—and recognize that distinct mechanisms may support selectivity for other stimuli within this or other circuits. For example, recent electron microscopy reconstructions in the posterior parietal cortex have revealed evidence of structured recurrent connectivity [34], in which excitatory pyramidal neurons preferentially target inhibitory interneurons with the same selectivity, while inhibitory interneurons target pyramidal neurons with opposite selectivity, forming an opponent inhibition motif. This connectivity motif was subsequently incorporated into a connectome-based model to investigate mechanisms underlying decision-making [34]. It is important to note that these studies did not quantitatively assess heterogeneity, leaving open the possibility of its significant role. Our findings suggest that heterogeneity could meaningfully contribute to selectivity and provide a framework for probing its influence.

Connectomic reconstruction relies on automatic segmentation algorithms, which are prone to errors [54]. These errors are likely the cause of the challenges we encountered when reconstructing synaptic currents to individual neurons in [1]. A common approach to mitigate such errors is manual proofreading, which has been extensively applied, for example, in the reconstruction of the *Drosophila* brain [35, 36]. However, this strategy is impractical for larger circuits, such as those in the mouse cortex, due to the significantly larger neuron size and far greater number of synapses. This raises the question of how connectomic data can be effectively leveraged to build mechanistic models of cortical functions. To address this, we developed a method that focuses on the statistical properties of neural responses and connectivity, rather than individual neurons and synapses. Our approach is based on two assumptions: first, that population-level cortical responses are driven by connection statistics, not individual connections, as shown by mean-field approximations [57–61]; and second, that while reconstruction errors affect absolute connection probabilities, relative probabilities are expected to remain robust. This latter observation allowed us to infer certain quantities from partially proofread data, as proofreading errors act as independent noise that averages out, reducing reliance on neurons with extensive proofreading. This approach provided greater statistical power compared to other methods [78], revealing that the connection probability for inputs from both layers 2/3 and 4 decreases as the difference in preferred orientation between pre- and postsynaptic neurons increases, consistent with previous studies [27, 53, 66–68].

The advent of large-scale reconstruction of neural circuits represents a significant milestone in neuroscience, poised to shape the field for years to come. For species with smaller brains, such as zebrafish and *Drosophila*, wiring diagrams of entire circuits have been available for several years and successfully integrated into theoretical models [37–46]. In contrast, wiring diagrams of cortical circuits are only now beginning to be developed [1, 27, 34], and we currently lack a systematic approach to integrate this information into models. In this study, we identified critical obstacles to the direct use of this data for studying cortical computations; we also introduced a novel approach that allowed us to effectively leverage this data to draw conclusions about the mechanisms underlying orientation selectivity in layer 2/3 of mouse V1. We anticipate that this method will be crucial for studying other computations within mouse visual cortex and will gain increasing relevance as the field progresses toward reconstructing larger parts of mouse brains as well as cortical circuits of other species.

MATERIALS AND METHODS

Analysis of the MICrONS dataset. We relied on the MICrONS dataset [1] v.661 to extract functional and

structural properties of neurons in layer 2/3 and 4 of mouse V1.

Functional properties. We extracted responses from neurons in layers 2/3 and 4 of V1 from the functional dataset. Neurons in these areas were identified using table `coregistration_manual_v3`, and layer assignment was based on soma depth, following the boundaries described in [78] (layer 2/3 from 98 µm to 283 µm, and layer 4 from 283 µm to 371 µm). We identified 6034 functionally matched excitatory neurons, of which 4167 were in layer 2/3 and 1867 in layer 4.

We analyzed trials with 'Monet' stimuli—parametric textures that vary in orientation and spatial frequency—each trial presenting one of 8 orientations from 0 to 180 degrees, with drift perpendicular to the stimulus's orientation (details in [1]). To characterized the selectivity of each neuron to stimulus orientation, we first identified the frames in each recording session during which the animal was exposed to a specific stimulus orientation. We then calculated the average response across these repeated trials for all orientations using the deconvolved spike traces. These responses were fitted with a von Mises function [79] utilizing least square optimization through the `curve_fit` function implemented as part of `scipy` library in Python,

$$f(\theta) = Ae^{k\left(\cos(2(\theta-\hat{\theta}))-1\right)}. \tag{4}$$

Here, A is the amplitude of the cosine wave, k is a width parameter, and $\hat{\theta}$ is the preferred orientation. Note that we considered $\hat{\theta}$ to be a discrete variable with 8 possible values that coincide with those of the stimulus presentation. We calculated the orientation at which the function's minimum occurred relative to the estimated $\hat{\theta}$ and retrieved the corresponding average spike traces for both orientations from the deconvolved spike traces. We then performed a Wilcoxon signed-rank test on the two sets; a p-value <0.05 indicated significantly higher activity at $\hat{\theta}$, classifying the neuron as selective. If the test was not significant, we classified the neuron as non-selective.

We characterized the tuning of selective neurons through circular variance

$$\text{CircVar} = \frac{\sum_\theta r(\theta)e^{i\theta}}{\sum_\theta r(\theta)} \tag{5}$$

where i is the imaginary constant, and the sum goes over the potential discrete values of the preferred orientation. This parameter is close to 0 for untuned neurons and close to 1 for highly selective ones.

Structural properties. We characterized connectivity within layers 2/3 and 4 of V1 using version v661 of the structural dataset. To assess the area each neuron belonged to, we used coregistered neurons, i.e., neurons identified in both the functional and structural datasets. We utilized table `coregistration_manual_v3` from the functional data to determine the area of these neurons. This information was then used to train a random forest classifier (from `scikit-learn` library, default parameters, 99% accuracy on cross-validated tests) to predict the area of a neuron based on its position in the reconstructed volume. Neuronal layers were identified based on depth using the same distances as for the functional data (see above).

In order to identify the cell type in the connectomic data, we levaraged infrormation from 3 different tables, `baylor_log_reg_cell_type_coarse_v1`, `baylor_gnn_cell_type_fine_model_v2` [80] and `aibs_soma_nuc_metamodel_preds_v117` [54], which performed a classification of cell types. We restrained our analysis only to neurons which were present in all the tables, and decide if a neuron was excitatory or inhibitory by majority agreement among the three tables. We compared the prediction of this method with neurons manually classified by experts, determining that combining the three tables lead to a slightly better accuracy in determining the cell type. Through this approach, we identified a total of 17412 neurons, of which 16580 are excitatory (5629 layer 4 and 10951 in layer 2/3) and 832 inhibitory neurons, of which 568 are in layer 2/3. Synapses between the neurons were obtained through API calls to the synapse table to query connections of the neuron subset of interest.

As discussed in the main text and in [54], automatic synaptic connection reconstructions are subject to error, which can be addressed through manual proofreading. In the MICrONS dataset, dendritic and axonal reconstructions were proofread independently, with three levels of precision: `none`, `clean` and `extended`. In our analysis, neurons were selected based on their reconstruction status, applying stricter criteria to axonal reconstructions due to the higher error rates associated with their smaller size. Dendrites, being thicker, are generally well-segmented, with most merge errors involving small axons [81]. For most analyses, we extracted connectivity properties from synaptically connected pairs with clean axonal proofreading and no dendritic proofreading. Although connectivity may be biased by limited axonal extent, we focused on relative measures, which are less sensitive to reconstruction errors and expected to average out noise. Supp. Fig. 3 supports this approach, demonstrating that several relative measures show weak dependency on the proofreading status of presynaptic and postsynaptic targets. In other analyses, where we were explicitly interested in absolute quantities (discussed in the main text), reconstruction errors had a more systematic effect. For these analyses, we used pairs where the axons were proofread to the `extended` level and the dendrites to the `clean` level. We estimated connection strength using synapse size from the table of synapses; when multiple synapses were present, we combined their sizes to measure the total connection strength. Synaptic sizes were then normalized such that the average size was set to 1.

The data analysis pipeline is fully automated, documented and easy to use, and will be released as open

source code upon the publication of the final version of the paper. The codebase also contains several utilities to work with the processed data, to easily filter neurons and synapses for common criteria, work with neuronal activity, etc.

Inference of network statistics.

Estimation of the in-degree. We estimated the in-degree of excitatory neurons by combining the probability of connection between excitatory neuron pairs and the total number of excitatory neurons. The connection probability was determined from neuron pairs with axonal reconstructions proofread to the `extended` level and dendritic reconstructions to the `clean` level in order to have an accurate estimate. The total number of excitatory neurons was obtained by counting the identified E neurons in the dataset. This approach yielded an expected in-degree of 598.

Connection probability and strength. To determine the statistics of connectivity between different neuron types, we cataloged all excitatory and inhibitory neurons in layer 2/3, as well as excitatory neurons in layer 4 (see above for determination of excitatory and inhibitory neurons). From this list, we extracted all recorded synapses between these neurons to construct a connectivity matrix. We quantified connection probability between each pair of presynaptic and postsynaptic neuron types and selectivity properties by calculating the fraction of non-zero entries in this matrix relative to the total number of possible connections. Due to potential reconstruction errors, some entries in the matrix are either missing or may represent connections that do not actually exist. To address this, we approached the inference in two stages. First, we estimated connection probability using only neuron pairs with at least `extended` axonal proofreading and `clean` dendritic proofreading, which reduced the number of pairs included. Then, to examine how neuron type and selectivity modulates this probability, we computed the relative magnitude of connection probability using pairs with at least `clean` axonal proofreading and no dendritic proofreading. This step allowed us to increase the sample size while expecting that reconstruction errors would be unbiased with respect to selectivity, thus averaging out in our calculations. Finally, we extracted connection strength distributions using all pairs with at least `clean` axonal proofreading and no dendritic proofreading.

Analyses of selectivity of synaptic currents.

Direct derivation of synaptic currents from connectome. To validate the relationship between preferred orientation derived from firing rate and synaptic currents predicted by Eq. (3), we restricted our analysis to neuron pairs with `clean` or `extended`-level proofreading. For each selective postsynaptic neuron, we identified all its functionally matched presynaptic partners with at least `clean` level of proofreading in axons. We calculated the individual synaptic current contributions by multiplying the tuning curve of each presynaptic neuron (i.e., its firing rate as a function of stimulus orientation) by the size of the corresponding synapse. For non-selective neurons, we assumed orientation-independent tuning curves, using the average firing rate as the constant value. Summing across all synapses provided an estimate of the total synaptic current for the postsynaptic neuron, where the peak of this estimate represented the predicted preferred orientation. This procedure yielded the results shown in Fig. 1d,e. To assess the model's efficacy at the population level, we repeated this procedure for all neurons in the dataset, measuring the difference between each neuron's observed preferred orientation and the orientation predicted by synaptic currents. Fig. 1f presents a histogram of these differences. As a control, we shuffled the dataset by randomly assigning each postsynaptic neuron a set of presynaptic partners with randomly sized synapses, then computed synaptic currents as before. Fig. 1f shows the results obtained by repeating this procedure 1000 times, with error bars representing the standard deviation over bootstrap samples.

Statistical derivation of synaptic currents from connectome. To investigate the statistics of synaptic currents in selective neurons, we compiled a list of all synapses to selective neurons along with the tuning curves of presynaptic neurons. Connections are filtered so presynaptic neurons had a `clean` level of proofreading at least. We adjusted each tuning curve so that the postsynaptic neuron was aligned to a zero orientation preference. Using this table, we generated model neurons by sampling presynaptic inputs according to the observed statistics in the data. We then summed these inputs to calculate the synaptic current for each model neuron, following the same method as in the direct derivation of synaptic currents. For each model neuron, we extracted the contributions from both layer 2/3 and layer 4 components, repeating the procedure 1000 times. Fig. 2e,f displays the mean synaptic current across these repetitions, with error bars representing the standard error of the mean. Fig. 2g shows a histogram of the peak positions of these currents across repetitions.

Simulation and analysis of recurrent network models. *Mathematical description of the network model.* Dynamics of neural activity in the network model is Eq. (1). The single-neuron response function was taken to be the f-I curve of leaky integrate-and-fire neurons driven by white noise [70]. For a neuron belonging to population $A \in [E, I]$ in response to an input μ, this was given by

$$\phi^A(\mu) = \left[\tau_{rp} + \tau_A\sqrt{\pi}\int_{u^A_{\min}}^{u^A_{\max}} e^{u^2}\left(1+\mathrm{erf}(u)\right)du\right]^{-1}, \tag{6}$$

with

$$u^A_{\max} = \frac{V_{th}-\mu-I_A}{\sigma_A}, \quad u^A_{\min} = \frac{V_r-\mu-I_A}{\sigma_A}. \tag{7}$$

The parameters σ_A and V_{th} control the slope of the response function and were assumed to be independent of neuron type. The parameter I_A controls the offset of

the transfer function and implicitly models feedforward inhibition from layer 4; this parameter was assumed to differ between E and I neurons. We fit σ_A, V_{th}, and I_A to the data, while standard values were assumed for the refractory period $\tau_{rp} = 2$ ms, resting potential $V_r = 10$ mV, and $\sigma_{E,I} = 10$ mV.

We modeled networks comprising a total of $N = 8000$ neurons, with 4883 excitatory (E) neurons, 409 inhibitory (I) neurons, and 2708 additional (X) neurons. The orientation-tuning responses of X neurons were sampled from the functional data, preserving the observed statistics for the fraction of selective neurons and the distribution of preferred orientations among selective neurons. Neurons were connected using random connectivity matrices, with connection probabilities and synaptic strength distributions directly derived from experimental data. The connection probability between E neurons, as well as from X to E neurons, was dependent on each neuron's assigned orientation selectivity. Each E neuron was assigned a target preferred orientation based on the distribution observed in the functional dataset, with connectivity structured according to this distribution. Specifically, the probability of a connection as a function of tuning was given by

$$p\left(\hat{\theta}^t_{\text{post}}, \hat{\theta}^t_{\text{pre}}\right) = p_{AB}\left[1 + \beta_{AB}\cos\left(\hat{\theta}^t_{\text{post}} - \hat{\theta}^t_{\text{pre}}\right)\right], \tag{8}$$

where $\hat{\theta}^t_{\text{post}}$ and $\hat{\theta}^t_{\text{pre}}$ are the target preferred orientations of the post- and pre-synaptic neurons, respectively. Here, p_{AB} represents the baseline connection probability between neurons of types A and B, as obtained from the structural dataset, and β_{AB} modulates this probability according to selectivity of the neurons.

Simulation procedure. We conducted simulations of the network dynamics in Python. For each parameter set and stimulus orientation θ, we generated one realization of the network structure, including the connectivity matrices J^{ij}_{AB} and the firing rates of the external population $r^i_X(\theta)$. We simulated the network dynamics according to Eq. (1) using an explicit Euler method with timestep $\Delta t = 5 \cdot 10^{-3}$ ms. We ran the simulations for a total duration of $200\tau_E = 4$ s and stored computed firing rates every $\tau_I/3 = 3.33$ ms. We determined the firing rate of each cell by calculating its average rate over the simulation period, excluding the initial $10\tau_E = 0.2$ s. We repeated this procedure for each fixed value of θ and obtained tuning curves for the neurons by combining results across different values of θ.

Fitting procedure. We identified the set of parameters that best fit the data using a simulation-based inference (SBI) approach employing the `sbi` package [77]. For a given value of $\beta_{E,X}$, we conducted a total of 10^5 simulations by first sampling 10^4 parameter sets from a uniform prior and running 10 repetitions for each set, using different realizations of the connectivity matrix and layer 4 responses. These simulations were used to train a deep network via the SNPE algorithm [82] using the average tuning curves of excitatory neurons in layer 2/3 as summary statistics. The trained network can estimate a posterior distribution for the observed average tuning curves. These parameters were used to run another 10^5 simulations as the first time. This approach allowed us to generate several hundreds of simulations with reasonable parameters. Furthermore, we selected the parameter sets that produced average excitatory responses and selectivity within 25% of those observed in layer 2/3 excitatory neurons among all the possibilities. This procedure was repeated for different values of $\beta_{E,X}$. The results shown in Fig. 3 and Fig. 4 presents data from the simulations the provided the best overall simulation.

Analysis of model responses. For network models with parameter sets that produced responses matching experimental data, we performed further analysis to measure selectivity and determine the preferred orientation of each neuron. This analysis generated Fig. 4a, comparing each neuron's selectivity with its preferred orientation. We then used the same network model and scrambled the columns of the E-E subnetwork, the E-X subnetwork, or both, running simulations on each scrambled configuration. In these simulations, layer 4 responses were identical, and only the connectivities were changed. These scrambles removed the tuning-dependent structure in connectivity. In each scrambled network, we remeasured neuronal selectivity to estimate the resulting circular variance. We also determined the preferred orientations and measured the connection probability as a function of differences in emergent preferred orientation, following the same procedure applied to experimental data.

[1] MICrONS Consortium, J Alexander Bae, Mahaly Baptiste, Caitlyn A Bishop, Agnes L Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J Bumbarger, Manuel A Castro, Brendan Celii, et al. Functional connectomics spanning multiple areas of mouse visual cortex. BioRxiv, pages 2021–07, 2021.

[2] David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. The Journal of physiology, 160(1):106, 1962.

[3] Michael M Merzenich, Paul L Knight, and G Linn Roth. Representation of cochlea within primary auditory cortex in the cat. Journal of neurophysiology, 38(2):231–249, 1975.

[4] Jay A Gottfried. Central mechanisms of odour object perception. Nature Reviews Neuroscience, 11(9):628–641, 2010.

[5] JUN Tanji and Edward V Evarts. Anticipatory ac-

tivity of motor cortex neurons in relation to direction of an intended movement. Journal of neurophysiology, 39(5):1062–1068, 1976.

[6] Apostolos P Georgopoulos, John F Kalaska, Roberto Caminiti, and Joe T Massey. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. Journal of Neuroscience, 2(11):1527–1537, 1982.

[7] Doug P Hanes and Jeffrey D Schall. Neural control of voluntary movement initiation. Science, 274(5286):427–430, 1996.

[8] John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. Brain research, 1971.

[9] David B Omer, Shir R Maimon, Liora Las, and Nachum Ulanovsky. Social place-cells in the bat hippocampus. Science, 359(6372):218–224, 2018.

[10] Jay Hegdé and David C Van Essen. Selectivity for complex shapes in primate visual area v2. The Journal of Neuroscience, 20(5):RC61, 2000.

[11] Minami Ito and Hidehiko Komatsu. Representation of angles embedded within contour stimuli in area v2 of macaque monkeys. Journal of Neuroscience, 24(13):3313–3324, 2004.

[12] Eucaly Kobatake and Keiji Tanaka. Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. Journal of neurophysiology, 71(3):856–867, 1994.

[13] Keiji Tanaka. Inferotemporal cortex and object vision. Annual review of neuroscience, 19(1):109–139, 1996.

[14] Chou P Hung, Gabriel Kreiman, Tomaso Poggio, and James J DiCarlo. Fast readout of object identity from macaque inferior temporal cortex. Science, 310(5749):863–866, 2005.

[15] James J DiCarlo, Davide Zoccolan, and Nicole C Rust. How does the brain solve visual object recognition? Neuron, 73(3):415–434, 2012.

[16] Ha Hong, Daniel LK Yamins, Najib J Majaj, and James J DiCarlo. Explicit information for category-orthogonal object properties increases along the ventral stream. Nature neuroscience, 19(4):613–622, 2016.

[17] David Hansel and Carl van Vreeswijk. The mechanism of orientation selectivity in primary visual cortex without a functional map. Journal of Neuroscience, 32(12):4049–4064, 2012.

[18] Cengiz Pehlevan and Haim Sompolinsky. Selectivity and sparseness in randomly connected balanced networks. PloS one, 9(2):e89992, 2014.

[19] David Ferster and Kenneth D Miller. Neural mechanisms of orientation selectivity in the visual cortex. Annual review of neuroscience, 23(1):441–471, 2000.

[20] Rani Ben-Yishai, R Lev Bar-Or, and Haim Sompolinsky. Theory of orientation tuning in visual cortex. Proceedings of the National Academy of Sciences, 92(9):3844–3848, 1995.

[21] Haim Sompolinsky and Robert Shapley. New perspectives on the mechanisms for orientation selectivity. Current opinion in neurobiology, 7(4):514–522, 1997.

[22] David Hansel and Haim Sompolinsky. Modeling feature selectivity in local cortical circuits, 1998.

[23] John G White, Eileen Southgate, J Nichol Thomson, Sydney Brenner, et al. The structure of the nervous system of the nematode caenorhabditis elegans. Philos Trans R Soc Lond B Biol Sci, 314(1165):1–340, 1986.

[24] Moritz Helmstaedter, Kevin L Briggman, Srinivas C Turaga, Viren Jain, H Sebastian Seung, and Winfried Denk. Connectomic reconstruction of the inner plexiform layer in the mouse retina. Nature, 500(7461):168–174, 2013.

[25] Narayanan Kasthuri, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, Richard Lee Schalek, José Angel Conchello, Seymour Knowles-Barley, Dongil Lee, Amelio Vázquez-Reina, Verena Kaynig, Thouis Raymond Jones, et al. Saturated reconstruction of a volume of neocortex. Cell, 162(3):648–661, 2015.

[26] Matthew E Berck, Avinash Khandelwal, Lindsey Claus, Luis Hernandez-Nunez, Guangwei Si, Christopher J Tabone, Feng Li, James W Truman, Rick D Fetter, Matthieu Louis, et al. The wiring diagram of a glomerular olfactory system. Elife, 5:e14859, 2016.

[27] Wei-Chung Allen Lee, Vincent Bonin, Michael Reed, Brett J Graham, Greg Hood, Katie Glattfelder, and R Clay Reid. Anatomy and function of an excitatory network in the visual cortex. Nature, 532(7599):370–374, 2016.

[28] David Grant Colburn Hildebrand, Marcelo Cicconet, Russel Miguel Torres, Woohyuk Choi, Tran Minh Quan, Jungmin Moon, Arthur Willis Wetzel, Andrew Scott Champion, Brett Jesse Graham, Owen Randlett, et al. Whole-brain serial-section electron microscopy in larval zebrafish. Nature, 545(7654):345–349, 2017.

[29] Katharina Eichler, Feng Li, Ashok Litwin-Kumar, Youngser Park, Ingrid Andrade, Casey M Schneider-Mizell, Timo Saumweber, Annina Huser, Claire Eschbach, Bertram Gerber, et al. The complete connectome of a learning and memory centre in an insect brain. Nature, 548(7666):175–182, 2017.

[30] Zhihao Zheng, J Scott Lauritzen, Eric Perlman, Camenzind G Robinson, Matthew Nichols, Daniel Milkie, Omar Torrens, John Price, Corey B Fisher, Nadiya Sharifi, et al. A complete electron microscopy volume of the brain of adult drosophila melanogaster. Cell, 174(3):730–743, 2018.

[31] Louis K Scheffer, C Shan Xu, Michal Januszewski, Zhiyuan Lu, Shin-ya Takemura, Kenneth J Hayworth, Gary B Huang, Kazunori Shinomiya, Jeremy Maitlin-Shepard, Stuart Berg, et al. A connectome and analysis of the adult drosophila central brain. elife, 9:e57443, 2020.

[32] Benjamin Scholl, Connon I Thomas, Melissa A Ryan, Naomi Kamasawa, and David Fitzpatrick. Cortical response selectivity derives from strength in numbers of synapses. Nature, 590(7844):111–114, 2021.

[33] Alexander Shapson-Coe, Michał Januszewski, Daniel R Berger, Art Pope, Yuelong Wu, Tim Blakely, Richard L Schalek, Peter H Li, Shuohong Wang, Jeremy Maitin-Shepard, et al. A connectomic study of a petascale fragment of human cerebral cortex. BioRxiv, pages 2021–05, 2021.

[34] Aaron T Kuan, Giulio Bondanelli, Laura N Driscoll, Julie Han, Minsu Kim, David GC Hildebrand, Brett J Graham, Daniel E Wilson, Logan A Thomas, Stefano Panzeri, et al. Synaptic wiring motifs in posterior parietal cortex support decision-making. Nature, 627(8003):367–373, 2024.

[35] Sven Dorkenwald, Arie Matsliah, Amy R Sterling, Philipp Schlegel, Szi-Chieh Yu, Claire E McKellar, Albert Lin, Marta Costa, Katharina Eichler, Yijie Yin,

et al. Neuronal wiring diagram of an adult brain. Nature, 634(8032):124–138, 2024.

[36] Philipp Schlegel, Yijie Yin, Alexander S Bates, Sven Dorkenwald, Katharina Eichler, Paul Brooks, Daniel S Han, Marina Gkantia, Marcia Dos Santos, Eva J Munnelly, et al. Whole-brain annotation and multi-connectome cell typing of drosophila. Nature, 634(8032):139–152, 2024.

[37] Jinseop S Kim, Matthew J Greene, Aleksandar Zlateski, Kisuk Lee, Mark Richardson, Srinivas C Turaga, Michael Purcaro, Matthew Balkam, Amy Robinson, Bardia F Behabadi, et al. Space–time wiring specificity supports direction selectivity in the retina. Nature, 509(7500):331–336, 2014.

[38] Horace B Barlow and Richard M Hill. Selective sensitivity to direction of movement in ganglion cells of the rabbit retina. Science, 139(3553):412–414, 1963.

[39] Adrian A Wanner and Rainer W Friedrich. Whitening of odor representations by the wiring diagram of the olfactory bulb. Nature neuroscience, 23(3):433–442, 2020.

[40] Shin-ya Takemura, Arjun Bharioke, Zhiyuan Lu, Aljoscha Nern, Shiv Vitaladevuni, Patricia K Rivlin, William T Katz, Donald J Olbris, Stephen M Plaza, Philip Winston, et al. A visual motion detection circuit suggested by drosophila connectomics. Nature, 500(7461):175–181, 2013.

[41] Kazunori Shinomiya, Thangavel Karuppudurai, Tzu-Yang Lin, Zhiyuan Lu, Chi-Hon Lee, and Ian A Meinertzhagen. Candidate neural substrates for off-edge motion detection in drosophila. Current Biology, 24(10):1062–1070, 2014.

[42] Elena A Westeinde, Emily Kellogg, Paul M Dawson, Jenny Lu, Lydia Hamburg, Benjamin Midler, Shaul Druckmann, and Rachel I Wilson. Transforming a head direction signal into a goal-oriented steering command. Nature, 626(8000):819–826, 2024.

[43] Peter Mussells Pires, Lingwei Zhang, Victoria Parache, LF Abbott, and Gaby Maimon. Converting an allocentric goal into an egocentric steering signal. Nature, 626(8000):808–818, 2024.

[44] Janne K Lappalainen, Fabian D Tschopp, Sridhama Prakhya, Mason McGill, Aljoscha Nern, Kazunori Shinomiya, Shin-ya Takemura, Eyal Gruntman, Jakob H Macke, and Srinivas C Turaga. Connectome-constrained networks predict neural activity across the fly visual system. Nature, pages 1–9, 2024.

[45] Dean A Pospisil, Max J Aragon, Sven Dorkenwald, Arie Matsliah, Amy R Sterling, Philipp Schlegel, Szi-chieh Yu, Claire E McKellar, Marta Costa, Katharina Eichler, et al. The fly connectome reveals a path to the effectome. Nature, 634(8032):201–209, 2024.

[46] Ashok Litwin-Kumar and Srinivas C Turaga. Constraining computational models using electron microscopy wiring diagrams. Current opinion in neurobiology, 58:94–100, 2019.

[47] Peter Dayan and Laurence F Abbott. Theoretical neuroscience: computational and mathematical modeling of neural systems. MIT press, 2005.

[48] Cristopher M Niell and Michael P Stryker. Highly selective receptive fields in mouse visual cortex. Journal of Neuroscience, 28(30):7520–7536, 2008.

[49] Saskia EJ de Vries, Jerome A Lecoq, Michael A Buice, Peter A Groblewski, Gabriel K Ocker, Michael Oliver, David Feng, Nicholas Cain, Peter Ledochowitsch, Daniel Millman, et al. A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. Nature neuroscience, 23(1):138–151, 2020.

[50] Simone Holler, German Köstinger, Kevan AC Martin, Gregor FP Schuhknecht, and Ken J Stratford. Structure and function of a neocortical synapse. Nature, 591(7848):111–116, 2021.

[51] Bao-hua Liu, Ya-tang Li, Wen-pei Ma, Chen-jie Pan, Li I Zhang, and Huizhong Whit Tao. Broad inhibition sharpens orientation selectivity by expanding input dynamic range in mouse simple cells. Neuron, 71(3):542–554, 2011.

[52] Andrew YY Tan, Brandon D Brown, Benjamin Scholl, Deepankar Mohanty, and Nicholas J Priebe. Orientation selectivity of synaptic input to neurons in mouse and cat primary visual cortex. Journal of Neuroscience, 31(34):12339–12350, 2011.

[53] L Federico Rossi, Kenneth D Harris, and Matteo Carandini. Spatial connectivity matches direction selectivity in visual cortex. Nature, 588(7839):648–652, 2020.

[54] Brendan Celii, Stelios Papadopoulos, Zhuokun Ding, Paul G. Fahey, Eric Wang, Christos Papadopoulos, Alexander B. Kunin, Saumil Patel, J. Alexander Bae, Agnes L. Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J. Bumbarger, Manuel A. Castro, Erick Cobos, Sven Dorkenwald, Leila Elabbady, Akhilesh Halageri, Zhen Jia, Chris Jordan, Dan Kapner, Nico Kemnitz, Sam Kinn, Kisuk Lee, Kai Li, Ran Lu, Thomas Macrina, Gayathri Mahalingam, Eric Mitchell, Shanka Subhra Mondal, Shang Mu, Barak Nehoran, Sergiy Popovych, Casey M. Schneider-Mizell, William Silversmith, Marc Takeno, Russel Torres, Nicholas L. Turner, William Wong, Jingpeng Wu, Szi-chieh Yu, Wenjing Yin, Daniel Xenes, Lindsey M. Kitchell, Patricia K. Rivlin, Victoria A. Rose, Caitlyn A. Bishop, Brock Wester, Emmanouil Froudarakis, Edgar Y. Walker, Fabian Sinz, H. Sebastian Seung, Forrest Collman, Nuno Maçarico da Costa, R. Clay Reid, Xaq Pitkow, Andreas S. Tolias, and Jacob Reimer. Neurd offers automated proofreading and feature extraction for connectomics. bioRxiv, 2024.

[55] Yashar Ahmadian and Kenneth D Miller. What is the dynamical regime of cerebral cortex? Neuron, 109(21):3373–3391, 2021.

[56] Larry F Abbott, Davi D Bock, Edward M Callaway, Winfried Denk, Catherine Dulac, Adrienne L Fairhall, Ila Fiete, Kristen M Harris, Moritz Helmstaedter, Viren Jain, et al. The mind of a mouse. Cell, 182(6):1372–1376, 2020.

[57] Daniel J Amit and Nicolas Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cerebral cortex (New York, NY: 1991), 7(3):237–252, 1997.

[58] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. Journal of computational neuroscience, 8:183–208, 2000.

[59] Alex Roxin, Nicolas Brunel, David Hansel, Gianluigi Mongillo, and Carl van Vreeswijk. On the distribution of firing rates in networks of cortical neurons. Journal of Neuroscience, 31(45):16217–16226, 2011.

[60] Itamar D Landau, Robert Egger, Vincent J Dercksen, Marcel Oberlaender, and Haim Sompolinsky. The impact of structural heterogeneity on excitation-inhibition balance in cortical networks. Neuron, 92(5):1106–1121,

2016.

[61] Alessandro Sanzeni, Mark H Histed, and Nicolas Brunel. Emergence of irregular activity in networks of strongly coupled conductance-based neurons. Physical Review X, 12(1):011044, 2022.

[62] Paul G Fahey, Taliah Muhammad, Cameron Smith, Emmanouil Froudarakis, Erick Cobos, Jiakun Fu, Edgar Y Walker, Dimitri Yatsenko, Fabian H Sinz, Jacob Reimer, et al. A global map of orientation tuning in mouse visual cortex. BioRXiv, page 745323, 2019.

[63] Ahna R Girshick, Michael S Landy, and Eero P Simoncelli. Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. Nature neuroscience, 14(7):926–932, 2011.

[64] Sen Song, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. PLoS biology, 3(3):e68, 2005.

[65] Christina You Chien Chou, Hovy Ho-Wai Wong, Connie Guo, Kiminou Elizabeth Boukoulou, Cleo Huang, Javid Jannat, Tal Klimenko, Vivian Yuwei Li, Tasha Aurora Liang, Vivian Chenfei Wu, et al. Principles of visual cortex excitatory microcircuit organization. bioRxiv, pages 2023–12, 2023.

[66] Ho Ko, Sonja B Hofer, Bruno Pichler, Katherine A Buchanan, P Jesper Sjöström, and Thomas D Mrsic-Flogel. Functional specificity of local synaptic connections in neocortical networks. Nature, 473(7345):87–91, 2011.

[67] Lee Cossell, Maria Florencia Iacaruso, Dylan R Muir, Rachael Houlton, Elie N Sader, Ho Ko, Sonja B Hofer, and Thomas D Mrsic-Flogel. Functional organization of excitatory synaptic strength in primary visual cortex. Nature, 518(7539):399–403, 2015.

[68] Adrian Wertz, Stuart Trenholm, Keisuke Yonehara, Daniel Hillier, Zoltan Raics, Marcus Leinweber, Gergely Szalay, Alexander Ghanem, Georg Keller, Balázs Rózsa, et al. Single-cell–initiated monosynaptic tracing reveals layer-specific cortical network modules. Science, 349(6243):70–74, 2015.

[69] M Florencia Iacaruso, Ioana T Gasler, and Sonja B Hofer. Synaptic organization of visual space in primary visual cortex. Nature, 547(7664):449–452, 2017.

[70] A. J. F. Siegert. On the first passage time probability problem. Phys. Rev., 81:617–623, 1951.

[71] D. J. Amit and M. V. Tsodyks. Quantitative study of attractor neural network retrieving at low spike rates I: Substrate – spikes, rates and neuronal gain. Network, 2:259–274, 1991.

[72] D. J. Amit and N. Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cerebral Cortex, 7:237–252, 1997.

[73] Alexander Rauch, Giancarlo La Camera, Hans-Rudolf Luscher, Walter Senn, and Stefano Fusi. Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo–like input currents. Journal of neurophysiology, 90(3):1598–1612, 2003.

[74] Alessandro Sanzeni, Mark H Histed, and Nicolas Brunel. Response nonlinearities in networks of spiking neurons. PLoS computational biology, 16(9):e1008165, 2020.

[75] Alessandro Sanzeni, Agostina Palmigiano, Tuan H Nguyen, Junxiang Luo, Jonathan J Nassi, John H Reynolds, Mark H Histed, Kenneth D Miller, and Nicolas Brunel. Mechanisms underlying reshuffling of visual responses by optogenetic stimulation in mice and monkeys. Neuron, 111(24):4102–4115, 2023.

[76] Paul K LaFosse, Zhishang Zhou, Jonathan F O'Rawe, Nina G Friedman, Victoria M Scott, Yanting Deng, and Mark H Histed. Cellular-resolution optogenetics reveals attenuation-by-suppression in visual cortical neurons. Proceedings of the National Academy of Sciences, 121(45):e2318837121, 2024.

[77] Alvaro Tejero-Cantero, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann, Conor Durkan, Pedro J. Gonçalves, David S. Greenberg, and Jakob H. Macke. sbi: A toolkit for simulation-based inference. Journal of Open Source Software, 5(52):2505, 2020.

[78] Zhuokun Ding, Paul G Fahey, Stelios Papadopoulos, Eric Y Wang, Brendan Celii, Christos Papadopoulos, Andersen Chang, Alexander B Kunin, Dat Tran, Jiakun Fu, et al. Functional connectomics reveals general wiring rule in mouse visual cortex. bioRxiv, pages 2023–03, 2023.

[79] N.V. Swindale. Orientation tuning curves: Empirical description and estimation of parameters. Biological Cybernetics, pages 45–56, 1998.

[80] Leila Elabbady, Sharmishtaa Seshamani, Shang Mu, Gayathri Mahalingam, Casey Schneider-Mizell, Agnes Bodor, J. Alexander Bae, Derrick Brittain, JoAnn Buchanan, Daniel J. Bumbarger, Manuel A. Castro, Erick Cobos, Sven Dorkenwald, Paul G. Fahey, Emmanouil Froudarakis, Akhilesh Halageri, Zhen Jia, Chris Jordan, Dan Kapner, Nico Kemnitz, Sam Kinn, Kisuk Lee, Kai Li, Ran Lu, Thomas Macrina, Eric Mitchell, Shanka Subhra Mondal, Barak Nehoran, Stelios Papadopoulos, Saumil Patel, Xaq Pitkow, Sergiy Popovych, Jacob Reimer, William Silversmith, Fabian H. Sinz, Marc Takeno, Russel Torres, Nicholas Turner, William Wong, Jingpeng Wu, Wenjing Yin, Szichieh Yu, Andreas Tolias, H. Sebastian Seung, R. Clay Reid, Nuno Maçarico Da Costa, and Forrest Collman. Quantitative census of local somatic features in mouse visual cortex. bioRxiv, 2022.

[81] Allen Institute. Microns tutorial. `https://alleninstitute.github.io/microns_tutorial/`, 2024.

[82] Michael Deistler, Pedro J Goncalves, and Jakob H Macke. Truncated proposals for scalable and hassle-free simulation-based inference. arXiv, 2022.