Large-scale algorithmic search identifies stiff and sloppy dimensions in synaptic architectures consistent with murine neocortical wiring

Tarek Jabri[1], Jason N. MacLean[1,2,3,*]

[1]Department of Neurobiology, The University of Chicago, Chicago, Illinois 60637, United States of America; [2]Committee on Computational Neuroscience, The University of Chicago, Chicago, Illinois 60637, United States of America; [3]Neuroscience Institute, Chicago, Illinois 60637, United States of America; *Corresponding author: jmaclean@uchicago.edu

Abstract

Complex systems can be defined by “sloppy” dimensions, meaning that their behavior is unmodified by large changes to specific parameter combinations, and “stiff” dimensions whose changes result in considerable modifications. In the case of the neocortex, sloppiness in synaptic architectures would be crucial to allow for the maintenance of spiking dynamics in the normal range despite a diversity of inputs and both short- and long-term changes to connectivity. Using simulations on neural networks with spiking matched to murine visual cortex, we determined the stiff and sloppy parameters of synaptic architectures across three classes of input (brief, continuous, and cyclical). Large-scale algorithmically-generated connectivity parameter values revealed that specific combinations of excitatory and inhibitory connectivity are stiff and that all other architectural details are sloppy. Stiff dimensions are consistent across a range of different input classes with self-sustaining synaptic architectures occupying a smaller subspace as compared to the other input classes. We also find that experimentally estimated connectivity probabilities from mouse visual cortex are similarly stiff and sloppy when compared to the architectures that we identified algorithmically. This suggests that simple statistical descriptions of spiking dynamics are a sufficient and parsimonious description of neocortical activity when examining structure-function relationships at the mesoscopic scale. Moreover, this study provides further evidence of the importance of the interrelationship of excitatory and inhibitory connectivity to establish and maintain stable spiking dynamical regimes in neocortex.

Significance Statement

Connections between neurons are continuously changing to allow learning and adaptation to new stimuli. However, the ability of neural networks to vary these connections while avoiding excessively high- or low-activity states is still not well understood. We tackled this question by studying how changes in the parameters of connectivity within and between different neuronal populations impacted network activity in computational models. We identified specific combinations of parameters, deemed “stiff”, that must be maintained to observe activity consistent with recordings from

murine visual cortex, while the rest of the parameters can be varied freely with minimal effects on activity. Our results agree with experimentally measured connectivity statistics demonstrating the importance of balancing opposing forces to maintain activity in a natural regime.

Conflicts of interest: None.

Acknowledgements: JNM was granted research funds from the National Institute of Health (NIH Grant R01EY022338) and from the National Science Foundation (NSF CAREER Grant 0952686). TJ was granted research support by the University of Chicago (Dean's Scholarship and Neuroscience Research Metcalf Fellowship). We thank former and current MacLean lab members, Yuqing Zhu, Isabel Garon, Maayan Levy, and Gabriella Wheeler Fox, and Dr. Emil Sidky for assistance with initial simulations, accurate estimations of Fisher Information Matrices, and helpful comments on our manuscript.

Authors' contributions: TJ and JNM conceived the project. TJ performed the research and analyzed the data. JNM supervised the project. TJ and JNM wrote the paper and edited the manuscript.

Introduction

Local synaptic connectivity in neocortex is fundamental to the generation and stipulation of the spiking dynamics (Koulakov et al., 2009; Cossell et al., 2015) that underlie the formation of percepts, decisions, and the generation of appropriate behavioral responses. However, the rules that govern the mesoscopic-scale relationships between local synaptic architectures and spiking activity remain unclear. On one hand, synaptic architectures must be highly dynamic since they underlie, at least in part, the storage of information (Chklovskii et al., 2004) and generate the variety of spiking activity corresponding to distinct brain states (Doiron et al., 2016). On the other hand, it is clear that aberrant synaptic wiring can give rise to detrimental spiking behaviors and pathophysiology such as epilepsy (Engel et al., 2013).

Modeling and theoretical analysis are essential complements to experimental investigations of structure-function relationships since they enable precise manipulation of simulated connectivity (Transtrum et al., 2015; Churchland and Abbott, 2016). Any model of a biological system involves a large number of free parameters that cannot always be determined from experimental data or that vary greatly from one observation to the other. Nevertheless, numerous models of neural systems have successfully replicated key aspects of neuronal network dynamics including different activity patterns (Chambers and MacLean, 2016; Ocker et al., 2017; Vegué and Roxin, 2019) and receptive field properties (Zylberberg et al., 2011; Kerr et al., 2015; Hopkins et al., 2018). Given the difficulty of estimating the exact values of connectivity parameters, and given the variability in these values observed *in-vivo*, a reasonable hypothesis for why these models work is that multiple combinations of parameters can result in similar network activity (Brown and Sethna, 2003). Together, these observations are indicative of sloppy systems, whose behavior depends only on a few stiff combinations of parameters while the majority of parameters are not critical for accurate predictions of the system's behavior (Gutenkunst et al., 2007).

Sloppiness is a universal feature of models in systems biology (Gutenkunst et al., 2007; Daniels et al., 2008; Transtrum et al., 2015). For example, ionic conductances within individual neurons have consistently been found to vary greatly across neurons and between individuals despite regularity in spiking activity (Prinz et al., 2004; Schulz et al., 2006; Ransdell et al., 2013). At the neuronal circuit level, stability and state changes are mediated by a subset of neurons described by a small number of stiff parameter combinations while the parameters of the remainder of the neurons are sloppy (Panas et al., 2015; Ponce-Alvarez et al., 2020).

What remains unclear, however, are the stiff and sloppy parameter combinations that define synaptic architectures capable of producing spiking that is low rate, irregular, and asynchronous—the dominant dynamical regime of the neocortex (Brunel, 2000;

Hromádka et al., 2008; Koulakov et al., 2009; Roxin et al., 2011). The conservation of both connectivity and wiring cost across different species (Assaf et al., 2020) is a strong incentive to find the stiff and sloppy dimensions of synaptic architectures. Moreover, delineating these dimensions of neocortical wiring can better constrain cortical models. Understanding the stiff and sloppy dimensions also carries implications beyond the realm of organic neural systems into artificial neural networks (ANNs). It has been shown that connectivity patterns are directly related to the dimensionality of the activity in recurrent spiking neural networks (SNNs), with the latter decreasing as overall connectivity increases (Recanatesi et al., 2019). Hence, delineating the small number of parameters that describe stiff dimensions of network connectivity will allow further studies to capture meaningful functional and computational principles that define those networks.

Here, we survey a large portion of the synaptic connectivity parameter space to create wiring diagrams and identify parameter combinations capable of producing activity matched to murine visual cortex (V1). We identify the stiff and sloppy parameter combinations of synaptic architectures responsible for producing naturalistic activity and compare these algorithmically identified connectivity parameters to recent experimental values finding them to be largely in agreement.

Methods

Model architecture

Networks consisted of 5000 neurons: 4000 excitatory (e) and 1000 inhibitory (i). Neurons were modeled as adaptive exponential leaky integrate-and-fire (AdEx) (Brette and Gerstner, 2005) units. Membrane potential of each neuron was governed by the following equations:

$$C\frac{dV}{dt} = g_L\left(E_L - V\right) + g_L\Delta_T \cdot exp\left(\frac{V-V_T}{\Delta_T}\right) + g_e\left(E_e - V\right) + g_i\left(E_i - V\right) + g_p\left(E_e - V\right) - w \quad (1)$$

With the decaying adaptation current:

$$\tau_w \frac{dw}{dt} = a(V - E_L) - w \quad (2)$$

Neurons were connected according to four probabilities of connectivity that were varied: p_{ee}, p_{ie}, p_{ei}, and p_{ii}; where the first subscript index represents the presynaptic neuron type and the second represents the postsynaptic neuron type.

Synaptic conductances decayed exponentially as per the following equations:

$$\tau_e \frac{dg_e}{dt} = - g_e \quad (3)$$

$$\tau_i \frac{dg_i}{dt} = - g_i \quad (4)$$

$$\tau_p \frac{dg_p}{dt} = - g_p \quad (5)$$

When a neuron fires, the membrane potential is reset, the adaptation current is increased by a value of b, and the corresponding synaptic conductance at the synapses receiving the signal is increased by the weight of the connection. The weights of the connections were randomly drawn from a log-normal distribution with the parameters of the corresponding normal being $\mu = - 0.64$ and $\sigma = 0.51$. Connections from inhibitory neurons were enhanced by an order of magnitude. The parameters fixed across all simulations are defined and summarized in Table 1.

Initial voltages of all neurons were randomly drawn from a normal distribution with mean $\mu = - 65mV$ and standard deviation $\sigma = 5mV$.

All simulations were implemented in Python 3 using the Brian Simulator (version 2.2.1) (Stimberg et al., 2019).

Table 1. Fixed neuronal parameters.

Parameter Name	Notation	Value
Capacitance	C	281 pF
Leak conductance	g_L	30 nS
Leak reversal potential	E_L	-70.6 mV
Slope factor	Δ_T	2 mV
Firing threshold	V_T	-40.4 mV
Excitatory reversal potential	E_e	0 mV
Inhibitory reversal potential	E_i	-75 mV
Excitatory synaptic time constant	τ_e	10 ms
Inhibitory synaptic time constant	τ_i	3 ms
Input synaptic time constant	τ_p	10 ms
Adaptation time constant	τ_w	144 ms
Subthreshold adaptation	a	4 nS
Spike-triggered adaptation	b	80.5 pA

Network input

Network activity was initiated using a population of 3000 Poisson units which were connected to both the excitatory and inhibitory neurons in the main network with a connection probability of 0.1 unless otherwise specified, and the weights of the connections were randomly drawn from a log-normal distribution with the parameters of the corresponding normal being $\mu = -0.64$ and $\sigma = 0.51$. Three types of input were considered:

- Brief input: the firing rates of the Poisson units were drawn from a log-normal distribution with the parameters of the corresponding normal being $\mu = 2.8$ and $\sigma = 0.3$, resulting in a mode at around 15 Hz. Their activity was halted after 300 ms.
- Continuous input: similar to the brief input but the activity of the input units was maintained for the duration of the simulation.

- Cyclical input: the maximum firing rates (ν_{max}) of the input units were drawn from the same log-normal distribution as that of the brief input, but the instantaneous firing rates varied between 0 and ν_{max} according to a sinusoidal function with a period of 600ms.

Connectivity parameters

To minimize sampling bias, the four connectivity parameters were algorithmically determined using a low resolution four-dimensional grid search in which p_{ee}, p_{ei}, and p_{ie} varied between 0.10 and 0.30 and p_{ii} varied between 0.20 and 0.40 with an increment of 0.02. This resulted in 14641 parameter combinations which were each simulated 5 times with different initial membrane potentials, yielding a total of 73205 simulations.

We also tested values based on the experimental data in V1 published by the Allen Institute for Brain Science at an intersomatic distance of 75μm (Billeh et al., 2020). The connectivity statistics reported in the paper are for pyramidal excitatory neurons and the three major classes of inhibitory neurons determined by the markers parvalbumin, somatostatin, and the ionotropic serotonin receptor 5HT3a in all 6 layers of V1. We calculated the four summary connectivity parameters based on the total number of each neuron type also reported in the paper using the following equation:

$$p_{k \to h} = \sum_m \left(f_{k_m} \cdot \sum_n \left(f_{h_n} \cdot p_{k_m \to h_n} \right) \right) \quad (6)$$

Where:

- $k, h \in \{e, i\}$;
- $p_{k \to h}$ and $p_{k_m \to h_n}$ are respectively the connection probability from neurons of type k to neurons of type h, and the connection probability from neurons of subtype k_m to neurons of subtype h_n;
- f_{k_m} and f_{h_n} are the fraction of neurons of subtype k_m out of all the k neurons, and the fraction of neurons of subtype h_n out of the h neurons respectively.

For simulations using the probabilities calculated from the statistics of either L2/3 or L5, the connectivity probability of the input units to the network was computed based on the connections of L4 and L2/3 excitatory neurons respectively to each layer. In the case of L2/3, the resulting value was decreased to account for the numerous inhibitory connections from L4 to excitatory neurons in L2/3.

Matching spiking statistics to murine visual cortex

The fit of network activity to naturalistic spiking statistics was evaluated based on the following criteria:

- Firing rate of excitatory (ν_e) and inhibitory (ν_i) neurons.
- Synchrony (s): On runs which were tested individually, synchrony was determined using the Van Rossum distance between excitatory spike trains with 10ms time constant using the "elephant" package (van Rossum, 2001; Denker et al., 2018). For the Fisher Information Matrix calculations (see next subsection), we used the Van Rossum distance on the excitatory spikes during the last 150ms only, since these are representative of the activity of the network after cessation of input and it allows faster computation. However, because computing the Van Rossum distance is computationally expensive, we devised a fast method of estimating the synchrony of the excitatory neurons to use during the grid search: activity was divided into 10ms time bins, and for each time bin we calculated the variance of the number of spikes per neuron divided by the mean number of spikes per neuron and then averaged the result over all time bins. To evaluate the validity of this score, we calculated both this new synchrony measure and the Van Rossum distance for 45 networks that showed sustained activity. The two scores had a 94% correlation.
- Proportion of runs that resulted in sustained activity (p_{sus}): activity is considered sustained if excitatory neurons are firing until the end of the simulation with no more than 150ms of inactivity across the network. In general, networks that showed sustained activity for 1 second sustained their activity for the duration of the simulation regardless of simulation time. However, for runs initiated using brief input, some networks had activity that truncated after 1 second of sustained activity. To ensure that all networks that we evaluated would sustain activity for any simulation duration, we developed an additional check. We selected 50 parameter combinations that resulted in activity sustained for 1 second in at least one trial. We then ran 10 additional trials—each with a new adjacency matrix, initial voltages, and input units—on all 50 parameter combinations. Out of these 500 runs, 435 showed sustained activity for more than 1 second. Of these 435, 85 runs had activity that truncated past one second. We then used this data to train a Support Vector Machine (SVM) classifier with a Radial Basis Function (RBF) kernel to predict which of runs that sustained for 1 second will truncate later. This prediction was based on the excitatory and inhibitory firing rates and the fast synchrony measure calculated on the spikes during the first second of activity only. The best hyperparameters were determined using a cross-validation grid search carried out across different methods of scaling the three features

considered. The highest accuracy (96%) was obtained with a power transformation of the scores for a more Gaussian-like distribution and using $\Gamma = 0.1$ and $C = 1000$. This classifier was then used during the grid search on the connectivity parameters. The SVM was implemented using scikit-learn version 0.22.1 (Pedregosa et al., 2011). In this way we ensured that all models considered would spike for the full duration of simulation regardless of the duration.

Fisher Information Matrix estimation

The likelihood model is a multivariate normal distribution defined by:

$$L(x) = N \cdot exp\left[-\frac{1}{2}(y(x) - \mu)^T K_y^{-1}(y(x) - \mu)\right] \quad (7)$$

where N is a normalization factor; $x = (p_{ee}, p_{ei}, p_{ie}, p_{ii})$ is the vector of model parameters; $y(x) = (\bar{\nu}_e, \bar{\nu}_i, \bar{s}, p_{sus})$ is the vector of simulation scores' averages from 5 simulations using the same parameter vector x; K_y is the (estimated) covariance matrix of the scores vector y (estimated using the grid search results); and μ is the mean score vector of the distribution. μ was set to be equal to the average score vector $y(x)$ obtained at the optimal parameter combination x_{opt}. As such, x_{opt} results in the maximum likelihood.

The Fisher Information Matrix (FIM) is equal to the negative of the expected value of the Hessian matrix of the log-likelihood. Consequently the observed FIM can be estimated as the negative of the Hessian of the log-likelihood (Spall, 2005). However, computing the Hessian matrix numerically may result in errors, especially if the estimated covariance matrix has bad conditioning. Instead, we derived the expression of the components of the Hessian matrix:

$$\frac{\partial^2 log(L(x))}{\partial x_m \partial x_n} = -\left(\sum_{ij} \frac{\partial^2 y_i(x)}{\partial x_m \partial x_n}\left[K_y^{-1}\right]_{ij}\left(y_j(x) - \mu_j\right) + \frac{\partial y_i(x)}{\partial x_m}\left[K_y^{-1}\right]_{ij}\frac{\partial y_j(x)}{\partial x_n}\right) \quad (8)$$

The FIM was computed at the optimal parameter combination which nullifies all the second partial-derivative terms since $(y(x) - \mu) = 0$. We are left with the gradients of the scores which were estimated by calculating all four scores for networks that have 3 of the connectivity parameters fixed and one of them varied between $x_{optimal} - 0.01$ and $x_{optimal} + 0.01$, then fitting a line to these results due to the strong linear relationships around the optimal values (fig. 1).

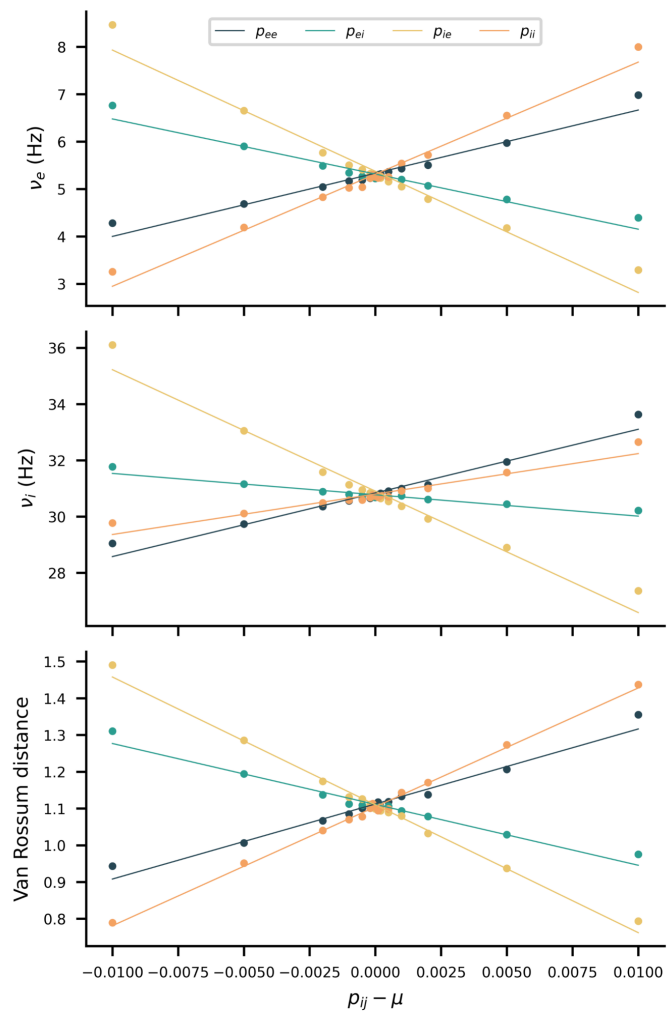

Figure 1. Variation of each spiking activity score with the change in each probability of connectivity around the optimal points.

Gini Coefficient

Sparsity of matrices was evaluated with the Gini coefficient (Hurley and Rickard, 2009) using the code provided by Panas et al. (2015):

$$G = 1 - 2\sum_{i=1}^{N} \frac{f_i}{\sum f_i}\left(\frac{N-i+1/2}{N}\right) \quad (9)$$

Where f_i are the values of the N elements in the matrix.

Statistical testing

Variances in the projections of the parameter vectors onto the eigenvectors were compared using the Levene test.

Results

Grid search for synaptic architectures producing naturalistic spiking

In order to evaluate the impact of specific synaptic connectivity parameter combinations on the statistics of spiking, we carried out large scale simulations of spiking neural network (SNN) models. SNNs were composed of both excitatory and inhibitory adapting exponential leaky integrate-and-fire neurons (AdEx) (Brette and Gerstner, 2005) connected with conductance-based synapses (fig. 2A; see methods). In previous work, we have shown that conductance based synapses are crucial to accurately simulate neuronal integration of synaptic inputs—a critical consideration when evaluating structure-function hypotheses (Chambers and MacLean, 2016; Bojanek et al., 2020). We conducted a grid search over a range of synaptic connectivity parameters defined by both excitatory and inhibitory connectivity and then quantified the outcome in SNN model behavior as connectivity parameter values changed (fig. 2B). Specifically, we determined which connectivity parameter combinations were capable of producing sustained naturalistic spiking activity matched to *in vivo* spiking of murine visual cortex (Niell and Stryker, 2010; Dechery and MacLean, 2018; Steinmetz et al., 2019; Billeh et al., 2020; Siegle et al., 2021). The performance of the networks was quantified using a set of objective functions, each of which corresponded to individual descriptors of spiking activity: firing rates, synchrony, and fraction of trials with sustained activity (fig. 2C; see methods).

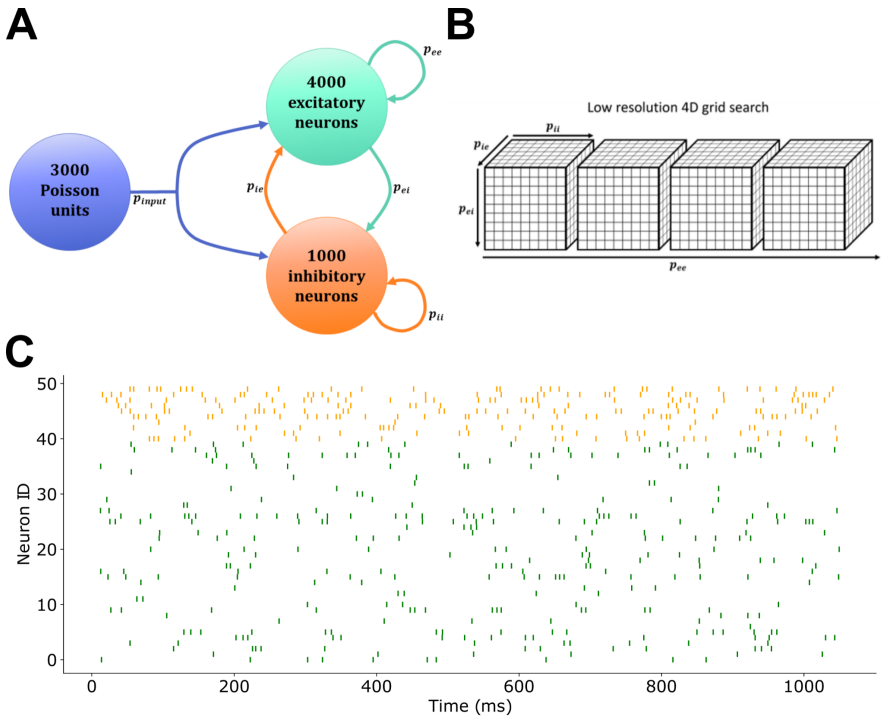

Figure 2. Network structure and activity. A. Simulated network composition. B. Design of grid search of synaptic architectures. C. Raster plot showing the timing of spikes of 50 example neurons, for ease of presentation, in the network with connectivity probabilities that gave the lowest firing rates in the target range using continuous input. Neurons 0-39 (green) are excitatory and neurons 40-49 (orange) are inhibitory.

We began by identifying combinations of parameters which produced firing rates between 8 and 15 Hz (Siegle et al., 2021) and synchrony scores corresponding to a Van Rossum distance greater than 4 in response to the three classes of input (constant **continuous**, **cyclical** continuous, **brief**). Notably, the input classes fall along a continuum of brevity and consequently place increased emphasis on network architectures capable of producing self-sustaining activity following input. The networks were first tested using continuous input from Poisson units firing with rates drawn from a log-normal distribution with a mean of about 17±5.3 Hz. Out of the 14461 unique parameters combinations tested, 5080 resulted in sustained activity for the length of the simulation more than 50% of the time. Of those, 579 synaptic wiring diagrams showed average excitatory firing rates in the desired range ($\overline{\nu_e} = 10.8 \pm 1.9\,Hz$). Those networks had a mean synchrony score of 1.03±0.06 ($\overline{\nu_i} = 27.2 \pm 8.1\,Hz$; $\overline{p_{sus}} = 0.998 \pm 0.029$). Networks receiving cyclical input (1.67Hz with maximal firing rates of units drawn from the same log-normal distribution) showed a lower number of successful parameter combinations at 199, including 47 networks with rates also between 8 and 15 Hz ($\overline{\nu_e} = 10.8 \pm 2.1\,Hz$; $\overline{\nu_i} = 25.6 \pm 7.6\,Hz$; $\overline{s} = 1.05 \pm 0.11$; $\overline{p_{sus}} = 0.97 \pm 0.10$). Finally, we evaluated architectures capable of producing self-sustaining activity in response to brief (300ms) excitatory Poisson input. We similarly simulated 73205 trials for 14461 synaptic architectures which corresponded to a range of different parameter combinations. This specific type of activity, being the hardest to achieve, resulted in the lowest number of successful networks: 241 trials resulted in self-sustained activity, which in turn corresponded to 44 unique parameter combinations that produced self-sustained activity in a minimum of 50% of the simulations. Of these 44 networks, 25 also scored well on the objective functions and exhibited an average firing rate of excitatory neurons between 8 and 15 Hz ($\overline{\nu_i} = 9.9 \pm 1.6\,Hz$; $\overline{\nu_i} = 20.6 \pm 3.6\,Hz$; $\overline{s} = 0.98 \pm 0.02$; $\overline{p_{sus}} = 0.888 \pm 0.161$; fig. 3B). Grid search resolution was low, in order to evaluate large ranges of parameter combinations for synaptic architectures. Thus the low number of viable synaptic wiring diagrams should not be interpreted as indicative of a scarcity of viable architectures. In fact, subsequent calculations of the FIM (discussed below) show that within a narrow range (±0.01) around the parameter combinations found, all networks exhibit sustained activity with low firing rates and low synchrony levels.

This limited number of possible networks spread out over the entire range of parameters tested (fig. 3D) is, by definition, indicative of both sloppiness and stiffness of spiking neural networks. For the rest of the analyses, we will only be considering those networks that achieved sustained activity using each type of input.

Correlated components of naturally spiking synaptic architectures

As an initial investigation of potentially stiff parameter combinations, we examined the differences and ratios between pairs of connectivity probabilities for the parameter combinations that resulted in sustained activity. We observed that certain differences in the connection likelihoods were highly correlated with each other (either positively or negatively), while others were uncorrelated (fig 3E&G). Specifically, $p_{ee} - p_{ie}$ and $p_{ei} - p_{ie}$ had one of the strongest positive correlations at 0.96 in the case of brief input, while $p_{ee} - p_{ie}$ and $p_{ii} - p_{ei}$ consistently showed the strongest negative correlations (-0.99 for brief and cyclical inputs, -0.95 for continuous inputs; fig. 3E&I&K). The same pattern of correlations was observed among pairs of ratios (e.g. p_{ee}/p_{ie} and p_{ei}/p_{ie} very strongly correlated; fig. 3G&J&L) indicating that networks that result in sustained activity are more likely to have connectivity parameters that follow these linear relationships between certain differences and ratios of connectivity. In other words, certain combinations of differences or combinations of ratios may constitute stiff parameters, while the others may be sloppy. In fact, especially in the case of networks receiving brief input, when looking at the pairs of differences that are highly correlated, it appeared that the majority of the networks had the same values for these differences despite having very different probabilities of connectivity (fig. 3F&H). Similar strong correlations were found among differences and ratios of wiring parameters in networks sustaining their activity with low firing rates using both brief and cyclical inputs (fig. 3I&J). When not restricting the firing rate to the target range the correlations remained but with lower magnitudes. Networks that exhibited sustained low rate activity in response to continuous input (n=579) also exhibited the majority of the same significant correlations (fig. 3K&L). We note that if we did not control for firing rate, these correlations were less apparent in networks that spiked in response to continuous inputs. To confirm that these correlations that appear for rate-matched networks are not statistical artifacts due to the relatively small number of networks considered (~11% of of successful networks), we matched the sample size (n=579) in different sets of networks randomly selected from those that showed sustained activity in more than 50% of the trials and evaluated the correlations among the pairs of parameters. The lack of strong correlations in the case of randomly selected networks confirms that the results that we observed are related to the rates of those networks being in the target range.

These correlated parameter combinations are additional indicators of the presence of stiff dimensions within synaptic architectural parameter combinations when matching stable spiking activity in the network to *in vivo* recordings. It is notable that despite the fact that different networks showed sustained spiking with each type of input, the same pairs of connectivity parameter combinations were crucial to the production of sustained and biologically-realistic activity regardless of input type.

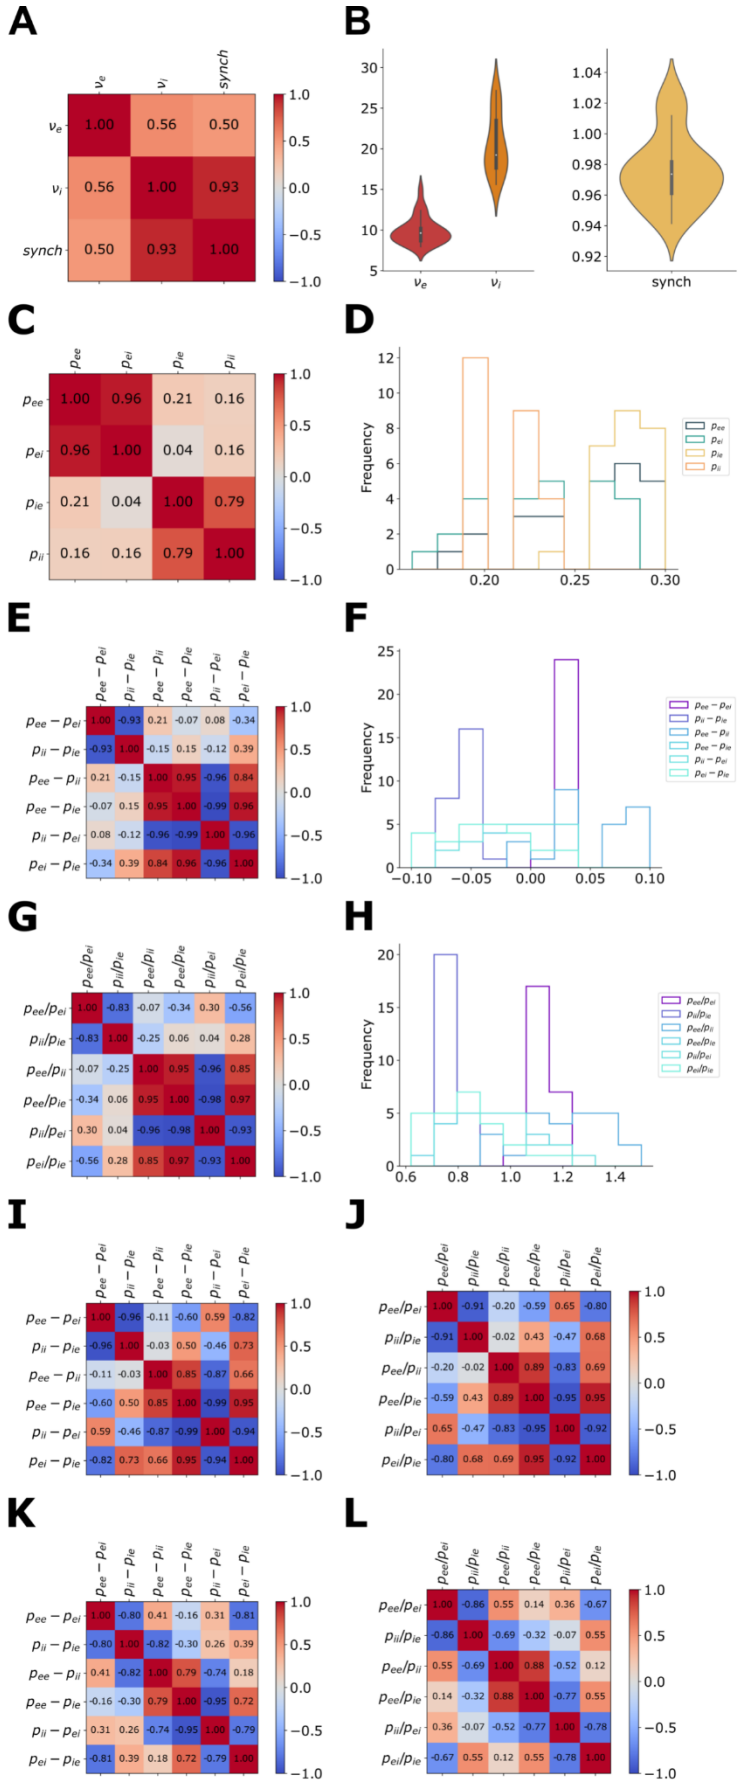
A
B
C
D
E
F
G
H
I
J
K
L
Frequency
synch

Figure 3. Analysis of networks from grid search with rate-matched sustained activity shows correlations between classes of connection. A. Correlation values between spiking measures of networks using brief input. B. Distribution of firing rates and synchrony measures of networks following brief input. C. Correlations between the four connectivity parameters of those successful networks following brief input. D. Distribution of the four parameters of connectivity in successful networks using brief input. E&I&K. Correlations between differences of connectivity probabilities using brief, cyclical, and continuous inputs respectively. F. Distribution of differences between connectivity probabilities from networks using brief input. G&J&L. Correlations between ratios of connectivity probabilities using brief, cyclical, and continuous inputs respectively. H. Distribution of differences between connectivity probabilities from networks using brief input. All networks considered here sustained their activity in more than 50% of the trials and had excitatory firing rates between 8 and 15 Hz.

Experimentally measured synaptic wiring in mouse visual cortex is highly similar to algorithmically identified synaptic architectures

To evaluate the biological plausibility of the networks that we identified using grid search, we compared the algorithmically generated wiring of the SNNs to the values of measured connectivity probabilities recently reported by the Allen Institute for Brain Science (Billeh et al., 2020). To do so, it was necessary to coarse grain the reported connection likelihoods from all classes of inhibitory interneurons and excitatory neurons regardless of laminar location by summarizing all of the connectivity measures as four probability values for excitatory and inhibitory neurons. This allowed us to match the reported values to those that we had varied in the grid search (fig. 4A; see methods). Additionally, we coarse grained the reported connectivities according to laminar position allowing us to separately evaluate individual layers 2/3, 4, and 5. For these simulations, input connectivity probability was maintained at 10% for the entire cortex and for L4, but otherwise was based on the reported connectivity probabilities from L4 for layers 2/3 and layers 2/3 to L5. Despite the fact that after coarse graining some of the resulting connectivity probabilities were not in the range tested in the grid search and despite significant differences in connectivity probabilities between the different lamina, all of the networks exhibited sustained activity in response to continuous input consistent with the initial study (fig. 4B) (Billeh et al., 2020).

Grid search in the continuous input condition found that the ratios $\frac{p_{ee}}{p_{ei}}$ and $\frac{p_{ii}}{p_{ie}}$ have a strong negative correlation despite not sharing any parameter, while $\frac{p_{ee}}{p_{ii}}$ and $\frac{p_{ei}}{p_{ie}}$, which also have distinct pairs, are not correlated (fig. 3L). We found that these correlations, and lack thereof, are also reflected in the coarse grained experimental connectivity values. We tested the importance of each of these pairs of ratios in sustaining network activity using the connectivity parameters derived from Billeh et al. (2020) and found

that if the ratios $\frac{p_{ee}}{p_{ei}}$ and $\frac{p_{ii}}{p_{ie}}$ were maintained while changing the actual parameter values, the majority of the resulting networks continued to exhibit sustained activity regardless of how large the values of the probabilities became (up to 0.99; fig. 4C). However, it should not come as a surprise that not all the resulting networks exhibited sustained activity—this is due to the role of other potentially critical parameter combinations that were being varied simultaneously. In contrast, using those same networks, maintaining the ratios $\frac{p_{ei}}{p_{ie}}$ and $\frac{p_{ee}}{p_{ii}}$ (which are very weakly correlated) did not result in networks that can sustain activity even with continuous input in most cases (fig. 4D).

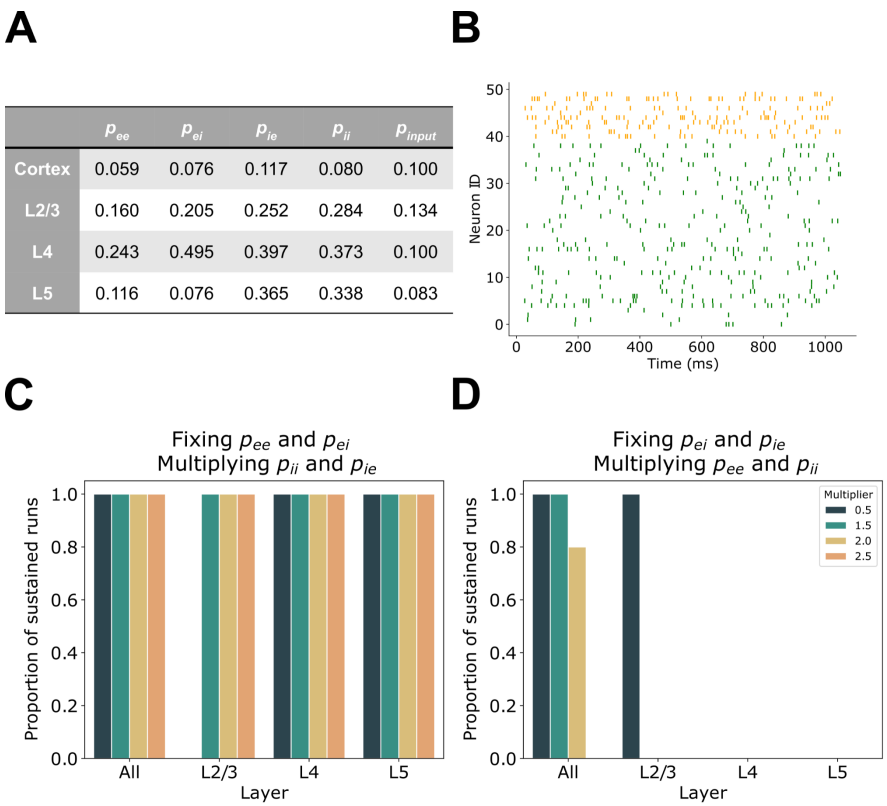

	p_{ee}	p_{ei}	p_{ie}	p_{ii}	p_{input}
Cortex	0.059	0.076	0.117	0.080	0.100
L2/3	0.160	0.205	0.252	0.284	0.134
L4	0.243	0.495	0.397	0.373	0.100
L5	0.116	0.076	0.365	0.338	0.083

Figure 4. Testing correlated ratios using experimental parameters. A. Coarse grained connectivity probabilities calculated from Billeh et al. (2020). B. Raster plot showing the timing of spikes of 50 example neurons in the network with connectivity probabilities experimentally derived from L2/3 using continuous input. Neurons 0-39 (green) are excitatory and neurons 40-49 (orange) are inhibitory. C. Networks that maintained highly correlated ratios regardless of the value of two parameters showed sustained activity. D. Same as C but instead maintaining weakly correlated ratios which did not result in sustained activity in most trials.

All four classes of connectivity contribute to stiff dimensions

To determine the contribution of individual parameters and their combinations to the stiff dimensions of synaptic architectures, as well as the impact of the input on the model's stiffness, we used the Fisher Information Matrix (FIM) (Gutenkunst et al., 2007). By its relation to the Hessian matrix, the FIM evaluated at a specific point in the parameter space examines how the likelihood of matching the spiking activity of murine visual cortex changes along the different dimensions around that point. Large changes in the likelihood along certain dimensions indicates that those dimensions are stiff whereas the others are sloppy. For each class of input, the FIM was computed at the parameter combinations that resulted in the lowest firing rates among the rate-matched networks (8-15 Hz), which we defined as the "optimal" combinations (fig. 5A). In the grid search, we ran five trials for each parameter combination and found negligible differences in the average firing rates of the networks at the lower end of the target range (8-15 Hz). We therefore estimated the FIM at the five parameter combinations with the lowest average firing rates for each class of inputs.

To identify the parameter combinations that had the greatest impact on spiking activity, we decomposed the FIM into eigenvectors and identified the corresponding eigenvalues. The eigenvalues of all FIMs extended over several orders of magnitude, consistent with many synaptic architecture parameter combinations being sloppy (fig. 5B&G&I). The eigenvectors that corresponded to the largest eigenvalues define the stiffest dimensions and indicated that those specific parameter combinations have the greatest impact on SNN spiking activity. Indeed, when projecting all the parameter vectors with sustained activity onto the different eigenvectors, we find that the first eigenvector—corresponding to the largest eigenvalue—had the lowest variance in projections (fig. 5E&F). This confirms that the first eigenvector defines the stiffest dimension of parameter space.

We then evaluated the sparsity of the FIM using the Gini coefficient to establish the complexity of the stiff dimensions (Panas et al., 2015). In the majority of networks the Gini coefficient was low (0.21±0.07), indicating that the FIMs were not sparse (Panas et al., 2015). The lack of sparsity indicates that the stiff dimensions depend on complex combinations of several parameters and not only on a few critical parameters (Gutenkunst et al., 2007; Panas et al., 2015), consistent with the results of the grid search. Indeed, the contribution of each parameter to the first eigenvectors reveals that those dimensions are equally sensitive to changes in all of the parameters (fig. 5D&H&J). These results were consistent across the three types of inputs.

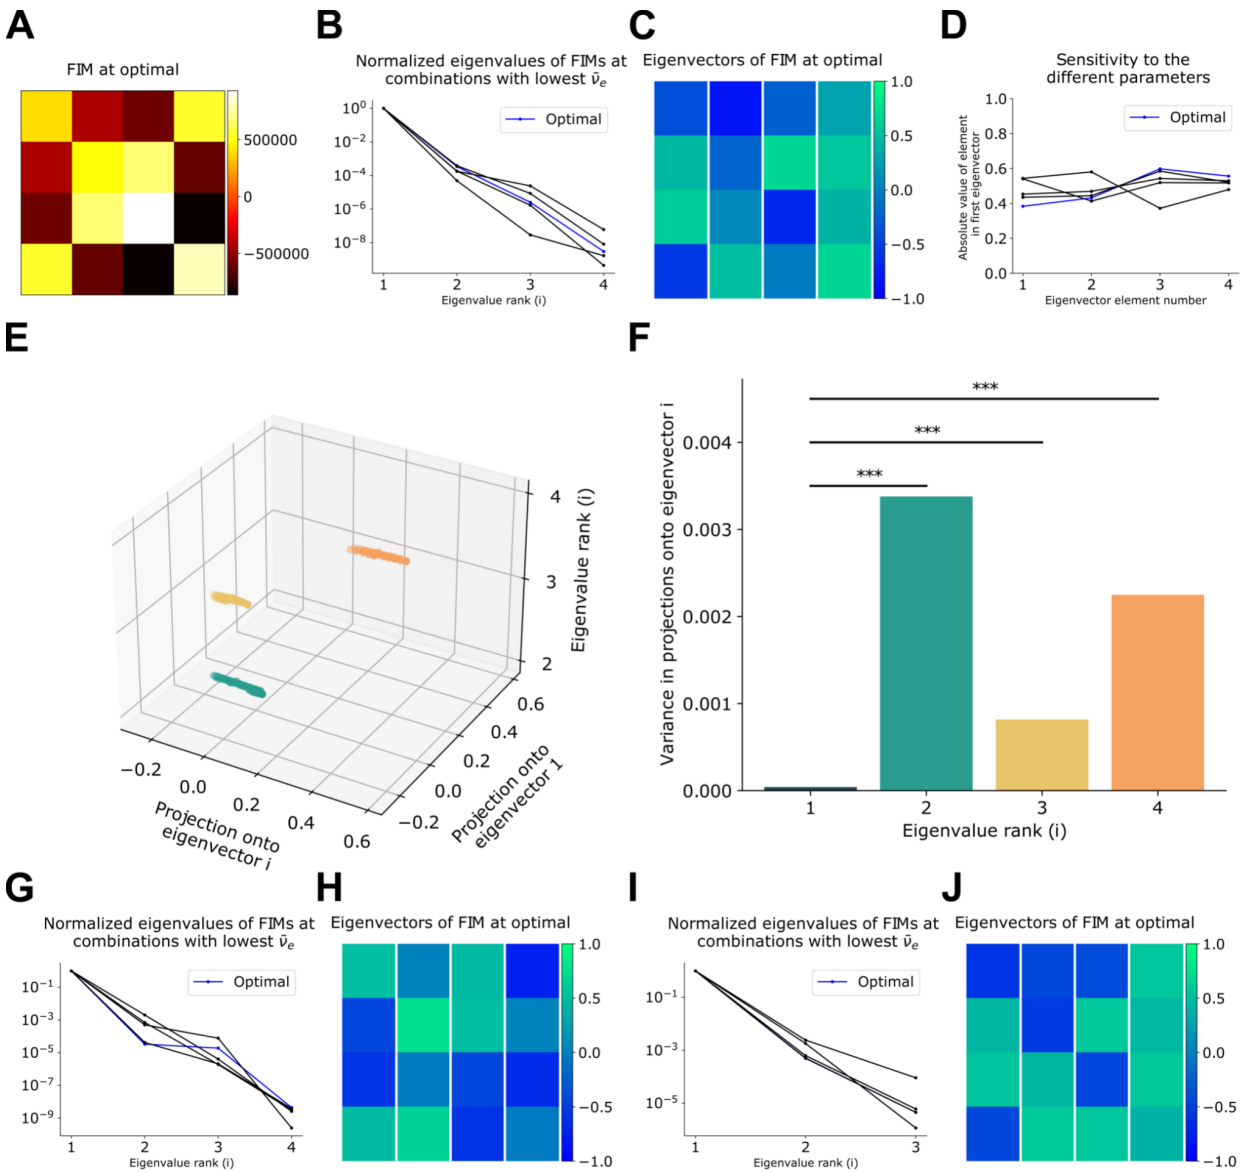

Figure 5. Fisher Information Matrix analysis. A. FIM computed at the optimal parameter combination for brief input. B. Eigenvalues of the FIMs computed at the five parameter combinations with the lowest firing rates in the target range using brief input. C. Eigenvectors of the FIM in A. D. Sensitivity of the eigenvectors in C to each of the parameters based on the absolute value of each vector element. E. Projections of parameter vectors that resulted in sustained activity following brief input onto the eigenvectors of the optimal FIM (green: projection onto the eigenvectors 1 and 2; yellow: projection onto the eigenvectors 1 and 3; orange: projection onto the eigenvectors 1 and 4). F. Variances in the projections in E $\left(p_{1-2} \approx 1.1 \cdot 10^{-14}, p_{1-3} \approx 6.2 \cdot 10^{-10}, p_{1-4} \approx 4.8 \cdot 10^{-14}\right)$. G&H and I&J. As B&C but for cyclical and continuous input respectively.

Input brevity increasingly constrained the viable wiring parameter space

All the networks that had sustained activity in response to brief input also showed sustained activity when receiving continuous input. In addition, similar patterns of correlations were observed among pairs of parameters for all three types of inputs. For this reason, we hypothesized that dimensions deemed stiff for one case may also be informative of stiff dimensions for other inputs. Indeed, when considering the projections of parameter combinations that resulted in sustained activity using continuous input onto the stiffest dimensions, we found that they are grouped in half the parameter space with little overlap with those networks that didn't sustain their activity. Importantly, SNNs that demonstrated self-sustaining activity following brief input correspond to a limited region of the parameter space described by those same eigenvectors (fig. 6A). Interestingly, networks that show sustained activity when receiving cyclical input define an intermediate region of this parameter space that encompasses those sustained with brief input. Both of those smaller regions are included in the larger space of networks sustained using continuous input (fig. 6A).

Parameter combinations derived from Billeh et al. (2020) for L2/3 and for the laminar-agnostic primary visual cortex fall inside the smallest region. However, connectivity parameters from L4 and L5 fall outside of this region and exhibit sustained activity in response to continuous input (fig. 6A). Notably for L5 parameters, the projection onto the first eigenvector fell in the same very narrow range but the projection onto the second eigenvector differed.

Dimensions used for this analysis define the stiffest dimensions of the system—and consistently, the same type of analysis on the less informative dimensions revealed no structure in the data (fig. 6B).

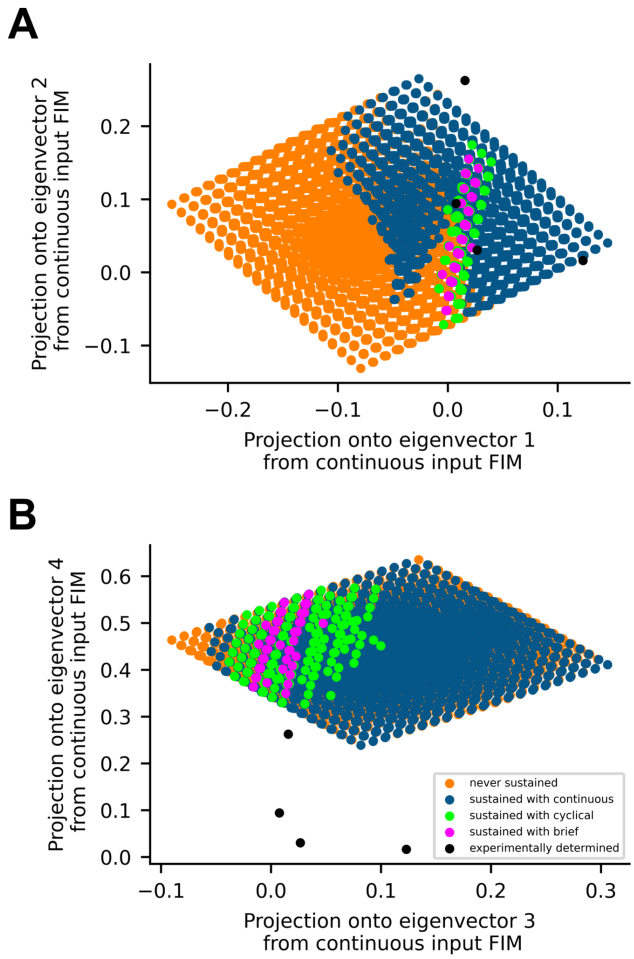

Figure 6. Restrictions of the parameter space along the stiffest dimensions. A&B. Projections of all parameter combinations tested onto the stiffest and sloppiest dimensions respectively, color-coded depending on the resulting spiking activity using each type of input.

Discussion

Using large scale algorithmic grid searches we have shown that the parameter space of synaptic architectures is highly anisotropic: large ranges of parameter values produce spiking that matches that of murine visual cortex, while a specific subset of parameter value combinations dramatically change network activity. These are the sloppy and stiff dimensions respectively. Stiff parameter combinations generalize across three broad classes of input into the network. Notably, the region of viable parameter combinations constricted as the requirement for architectures being capable of self-sustaining activity increased.

A recurring theme in neuroscience is that stiff parameter combinations encompass opposing forces. For example, most combinations of ion channel conductances are sloppy with the exception of a maintained ratio between a hyperpolarizing conductance and a depolarizing conductance (MacLean et al., 2003; Prinz et al., 2004; Schulz et al., 2006; Ransdell et al., 2013). Here we show that the connectivity statistics between and within excitatory and inhibitory neurons comprise the stiff parameter combinations of synaptic architectures. Indeed, maintaining a balance between excitation and inhibition is critical for normal network activity and has been resolved in synaptic conductances at the single cell level *in vivo* (Haider et al., 2006). We show that this balance is achieved in synaptic architectures through inter- and intra-population neuronal connections and involves excitation, inhibition, and disinhibition. It is for this reason that we found that all four connectivity parameters that we implemented contribute to the stiff dimensions and why pairs of connection likelihoods that don't share any parameters were highly correlated.

Surprisingly, despite coarse graining, experimentally estimated connectivity probabilities from the entirety of V1 of mice as well as connectivity probabilities from L2/3 of V1 (Billeh et al., 2020) fall inside the small space of topologies that we identified via algorithmic search as potentially capable of sustaining their activity using highly constrained input. However, there were notable differences between the algorithmically identified and experimentally measured synaptic architectures in other lamina. Experimentally estimated L5 connectivity parameters themselves did not fall within the restricted space of viable synaptic architectures, although the projection onto the first eigenvector does, indicating congruence with the stiffest dimension. The difference between L5 and algorithmically identified topologies is in the projection onto the second eigenvector, a less stiff dimension, that may reflect differences in the proportions of inhibitory neurons between the lamina. In fact, while the probabilities of connections from subtypes of inhibitory neurons to specific subtypes of neurons do not vary substantially between the different cortical layers (Billeh et al., 2020), the difference we observed in the generalized probabilities is likely due to the predominance of

parvalbumin- and somatostatin-positive inhibitory neurons in L5 (48% and 43% of inhibitory neurons respectively) when compared to the percentage of Htr3a-positive neurons (9% vs. 50% in L2/3). As a result, the coarse grained i-i and i-e probabilities in L5 are much higher than the excitatory probabilities in contrast to the other lamina. This could theoretically be compensated for by the much stronger e-e synapses that they obtained in L5; however we did not take these parameters into consideration in our analysis. Connectivity parameters of L4 also fall outside the region of the parameter space identified by the stiffest dimensions. This difference lacks an easy explanation but may simply reflect the different roles that layer 4 is hypothesized to play as compared to other lamina and the fact that thalamocortical connectivity is particularly important and heterogeneous into L4 (Landau et al., 2016).

Machine learning techniques demonstrate that the structure of neural networks can be modified to achieve a specific task and that, following training, these models are capable of accurately modeling neocortical neuronal activity at different stages of the visual processing hierarchy (Yamins et al., 2014). Notably, we found synaptic network architectures that matched experimentally measured connectivity by focusing on the spiking statistics of the network rather than training on a task. It will be of great interest to evaluate whether training similar networks (Bellec et al., 2020) will result in convergence to the same set of synaptic architectures.

Our approach to identify synaptic architectures based solely on spiking statistics and the correspondence with experimental measures demonstrates the utility of spiking statistics as a parsimonious way of studying structure-function relations. Indeed, previous work using maximum entropy models fit to spiking data was similarly successful at identifying cellular-level stiff and sloppy dimensions (Panas et al., 2015; Ponce-Alvarez et al., 2020). These results are consistent with a previous experimental study that identified neurons of varying levels of correlation with the network activity (Okun et al., 2015; Ponce-Alvarez et al., 2020). In sum, these studies along with this work form a compelling argument to study mesoscale connectivity as well as network wide correlations by fitting models to spiking statistics.

References

Assaf Y, Bouznach A, Zomet O, Marom A, Yovel Y (2020) Conservation of brain connectivity and wiring across the mammalian class. Nat Neurosci 23:805–808.

Bellec G, Scherr F, Subramoney A, Hajek E, Salaj D, Legenstein R, Maass W (2020) A solution to the learning dilemma for recurrent networks of spiking neurons. Nat Commun 11:3625.

Billeh YN, Cai B, Gratiy SL, Dai K, Iyer R, Gouwens NW, Abbasi-Asl R, Jia X, Siegle JH, Olsen SR, Koch C, Mihalas S, Arkhipov A (2020) Systematic Integration of Structural and Functional Data into Multi-scale Models of Mouse Primary Visual Cortex. Neuron 106:388-403.e18.

Bojanek K, Zhu Y, MacLean J (2020) Cyclic transitions between higher order motifs underlie sustained asynchronous spiking in sparse recurrent networks. PLOS Comput Biol 16:e1007409.

Brette R, Gerstner W (2005) Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity. J Neurophysiol 94:3637–3642.

Brown KS, Sethna JP (2003) Statistical mechanical approaches to models with many poorly known parameters. Phys Rev E 68:021904.

Brunel N (2000) Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons. J Comput Neurosci 8:183–208.

Chambers B, MacLean JN (2016) Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks. PLOS Comput Biol 12:e1005078.

Chklovskii DB, Mel BW, Svoboda K (2004) Cortical rewiring and information storage. Nature 431:782–788.

Churchland AK, Abbott LF (2016) Conceptual and technical advances define a key moment for theoretical neuroscience. Nat Neurosci 19:348–349.

Cossell L, Iacaruso MF, Muir DR, Houlton R, Sader EN, Ko H, Hofer SB, Mrsic-Flogel TD (2015) Functional organization of excitatory synaptic strength in primary visual cortex. Nature 518:399–403.

Daniels BC, Chen Y-J, Sethna JP, Gutenkunst RN, Myers CR (2008) Sloppiness, robustness, and evolvability in systems biology. Curr Opin Biotechnol 19:389–395.

Dechery JB, MacLean JN (2018) Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons. PLOS Comput Biol 14:e1006153.

Denker M, Yegenoglu A, Grün S (2018) Collaborative HPC-enabled workflows on the HBP Collaboratory using the Elephant framework. In: Neuroinformatics 2018, pp

P19 Available at: https://abstracts.g-node.org/conference/NI2018/abstracts#/uuid/023bec4e-0c35-4563-81ce-2c6fac282abd.

Doiron B, Litwin-Kumar A, Rosenbaum R, Ocker GK, Josić K (2016) The mechanics of state-dependent neural correlations. Nat Neurosci 19:383–393.

Engel J, Thompson PM, Stern JM, Staba RJ, Bragin A, Mody I (2013) Connectomics and epilepsy. Curr Opin Neurol 26:186–194.

Gutenkunst RN, Waterfall JJ, Casey FP, Brown KS, Myers CR, Sethna JP (2007) Universally Sloppy Parameter Sensitivities in Systems Biology Models. PLOS Comput Biol 3:e189.

Haider B, Duque A, Hasenstaub AR, McCormick DA (2006) Neocortical Network Activity In Vivo Is Generated through a Dynamic Balance of Excitation and Inhibition. J Neurosci 26:4535–4545.

Hopkins M, Pineda-García G, Bogdan PA, Furber SB (2018) Spiking neural networks for computer vision. Interface Focus 8:20180007.

Hromádka T, DeWeese MR, Zador AM (2008) Sparse Representation of Sounds in the Unanesthetized Auditory Cortex. PLOS Biol 6:e16.

Hurley N, Rickard S (2009) Comparing Measures of Sparsity. IEEE Trans Inf Theory 55:4723–4741.

Kerr D, McGinnity TM, Coleman S, Clogenson M (2015) A biologically inspired spiking model of visual processing for image feature detection. Neurocomputing 158:268–280.

Koulakov AA, Hromádka T, Zador AM (2009) Correlated Connectivity and the Distribution of Firing Rates in the Neocortex. J Neurosci 29:3685.

Landau ID, Egger R, Dercksen VJ, Oberlaender M, Sompolinsky H (2016) The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks. Neuron 92:1106–1121.

MacLean JN, Zhang Y, Johnson BR, Harris-Warrick RM (2003) Activity-Independent Homeostasis in Rhythmically Active Neurons. Neuron 37:109–120.

Niell CM, Stryker MP (2010) Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. Neuron 65:472–479.

Ocker GK, Hu Y, Buice MA, Doiron B, Josić K, Rosenbaum R, Shea-Brown E (2017) From the statistics of connectivity to the statistics of spike times in neuronal networks. Curr Opin Neurobiol 46:109–119.

Okun M, Steinmetz NA, Cossell L, Iacaruso MF, Ko H, Barthó P, Moore T, Hofer SB, Mrsic-Flogel TD, Carandini M, Harris KD (2015) Diverse coupling of neurons to

populations in sensory cortex. Nature 521:511–515.

Panas D, Amin H, Maccione A, Muthmann O, van Rossum M, Berdondini L, Hennig MH (2015) Sloppiness in Spontaneously Active Neuronal Networks. J Neurosci 35:8480–8492.

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011) Scikit-learn: Machine Learning in Python. J Mach Learn Res 12:2825–2830.

Ponce-Alvarez A, Mochol G, Hermoso-Mendizabal A, de la Rocha J, Deco G (2020) Cortical state transitions and stimulus response evolve along stiff and sloppy parameter dimensions, respectively. eLife 9:e53268.

Prinz AA, Bucher D, Marder E (2004) Similar network activity from disparate circuit parameters. Nat Neurosci 7:1345–1352.

Ransdell JL, Nair SS, Schulz DJ (2013) Neurons within the Same Network Independently Achieve Conserved Output by Differentially Balancing Variable Conductance Magnitudes. J Neurosci 33:9950–9956.

Recanatesi S, Ocker GK, Buice MA, Shea-Brown E (2019) Dimensionality in recurrent spiking networks: Global trends in activity and local origins in connectivity. PLOS Comput Biol 15:e1006446.

Roxin A, Brunel N, Hansel D, Mongillo G, van Vreeswijk C (2011) On the Distribution of Firing Rates in Networks of Cortical Neurons. J Neurosci 31:16217–16226.

Schulz DJ, Goaillard J-M, Marder E (2006) Variable channel expression in identified single and electrically coupled neurons in different animals. Nat Neurosci 9:356–362.

Siegle JH et al. (2021) Survey of spiking in the mouse visual system reveals functional hierarchy. Nature 592:86–92.

Spall JC (2005) Monte Carlo Computation of the Fisher Information Matrix in Nonstandard Settings. J Comput Graph Stat 14:889–909.

Steinmetz NA, Zatka-Haas P, Carandini M, Harris KD (2019) Distributed coding of choice, action and engagement across the mouse brain. Nature 576:266–273.

Stimberg M, Brette R, Goodman DF (2019) Brian 2, an intuitive and efficient neural simulator. eLife 8:e47314.

Transtrum MK, Machta BB, Brown KS, Daniels BC, Myers CR, Sethna JP (2015) Perspective: Sloppiness and emergent theories in physics, biology, and beyond. J Chem Phys 143:010901.

van Rossum MCW (2001) A Novel Spike Distance. Neural Comput 13:751–763.

Vegué M, Roxin A (2019) Firing rate distributions in spiking networks with heterogeneous connectivity. Phys Rev E 100:022208.

Yamins DLK, Hong H, Cadieu CF, Solomon EA, Seibert D, DiCarlo JJ (2014) Performance-optimized hierarchical models predict neural responses in higher visual cortex. Proc Natl Acad Sci 111:8619–8624.

Zylberberg J, Murphy JT, DeWeese MR (2011) A Sparse Coding Model with Synaptically Local Plasticity and Spiking Neurons Can Account for the Diverse Shapes of V1 Simple Cell Receptive Fields. PLOS Comput Biol 7:e1002250.