Strong and localized recurrence controls dimensionality of neural activity across brain areas

David Dahmen[1,*], Stefano Recanatesi[2,*], Xiaoxuan Jia[3], Gabriel K. Ocker[3,4], Luke Campagnola[3], Tim Jarsky[3], Stephanie Seeman[3], Moritz Helias[1,5,+], and Eric Shea-Brown[2,3,+]

[1]Institute of Neuroscience and Medicine (INM-6 and INM-10) and Institute for Advanced Simulation (IAS-6), Jülich Research Centre, Jülich, Germany
[2]University of Washington Center for Computational Neuroscience and Swartz Center for Theoretical Neuroscience, Seattle, WA, USA
[3]Allen Institute, Seattle, WA, USA
[4]Department of Mathematics and Statistics, Boston University, MA, USA
[5]Department of Physics, Faculty 1, RWTH Aachen University, Aachen, Germany
[*]These authors share first authorship
[+]These authors share senior authorship

The brain contains an astronomical number of neurons, but it is their collective activity that underlies brain function. The number of degrees of freedom that this collective activity explores – its dimensionality – is therefore a fundamental signature of neural dynamics and computation (1–7). However, it is not known what controls this dimensionality in the biological brain – and in particular whether and how local synaptic networks play a role (8–10). Through analysis of high-density Neuropixels recordings (11), we argue that areas across the mouse cortex operate in a *sensitive regime* that gives these synaptic networks a very strong role in controlling dimensionality. Moreover, we show that this control is expressed through highly tractable features of these synaptic networks. We then analyze these key features via a massive synaptic physiology dataset (12). Quantifying these features in terms of cell-type specific network motifs, we find that the synaptic patterns that impact dimensionality are prevalent in both mouse and human brains. Thus local circuitry scales up systematically to help control the degrees of freedom that brain networks may explore and exploit.

Introduction

The complexity of a neural network's activity can be measured by its dimensionality – that is, the number of collective degrees of freedom that its neurons explore. Dimensionality is closely linked to neural computation. Signal classification, for example, benefits from network activities that increase the dimensionality of the incoming signals to be classified (2–4, 13, 14). However, compressing inputs into lower-dimensional activity patterns helps generalization to novel signals (1, 15, 16). Studies have emphasized the comparatively high (3, 6, 7) or low (17, 18) dimensionality of recordings in various experimental settings. Moreover, the dimensionality of neural dynamics can change over time (19), throughout the information processing hierarchy (20), or during learning (15, 21, 22). These findings, taken together, highlight the importance of dimensionality as a property of network activity that will vary depending on the type of computation performed in a circuit. A key question is: how can the connectivity of a network regulate the dimensionality of its activity (8, 9, 23–25)?

This question is of particular interest for cortical networks. The preponderance of inhibitory feedback in these networks leads to a *balanced*, asynchronous state with weak correlations between neurons (26–29). Such asynchronous dynamics at first appear to imply high dimensional dynamics, in which all neurons are roughly independent. However, by analyzing electrophysiological recordings from over 30,000 neurons (11, 30) we show that the opposite is the case: dynamics are constrained to spaces of very low dimension relative to the number of neurons in areas across the brain. This results from the rapid accumulation of many weak but diverse pairwise correlations across the networks (cf. (31)) – as quantified by the variance of these correlations, over and above their average.

To understand the mechanistic origins of this low relative dimensionality of cortical activity, we advance the theory of dynamics in balanced networks, to show how the key variance of correlations results from the network's recurrent connectivity. Moreover, we show that in this strongly recurrent regime, dimensionality is highly sensitive to changes in the structure of recurrent connections. Beyond overall synaptic strength, specific connectivity patterns, or motifs, between pairs and triplets of cells (25, 32–40) can significantly tune the dimensionality of neural activity. To test whether this is broadly the case in biological circuits, we analyze newly released synaptic physiology datasets (12, 41), quantifying connections among more than 22,000 pairs of neurons. We find that the connectivity motifs implicated by our theory were strongly present in both mouse and human brain, that they differ across cortical layers in ways consistent with layer-specific dimensionality of neural activity, and show how previously established patterns of cell-type specific modulation and adaptation can have a new effect: to further regulate connectivity motifs and hence dimensionality across time and brain state.

Low dimensionality across brain regions

We quantify the dimensionality of neural activity via the participation ratio D_{PR}, a widely used measure of dimensionality which, in particular, often corresponds to the number of

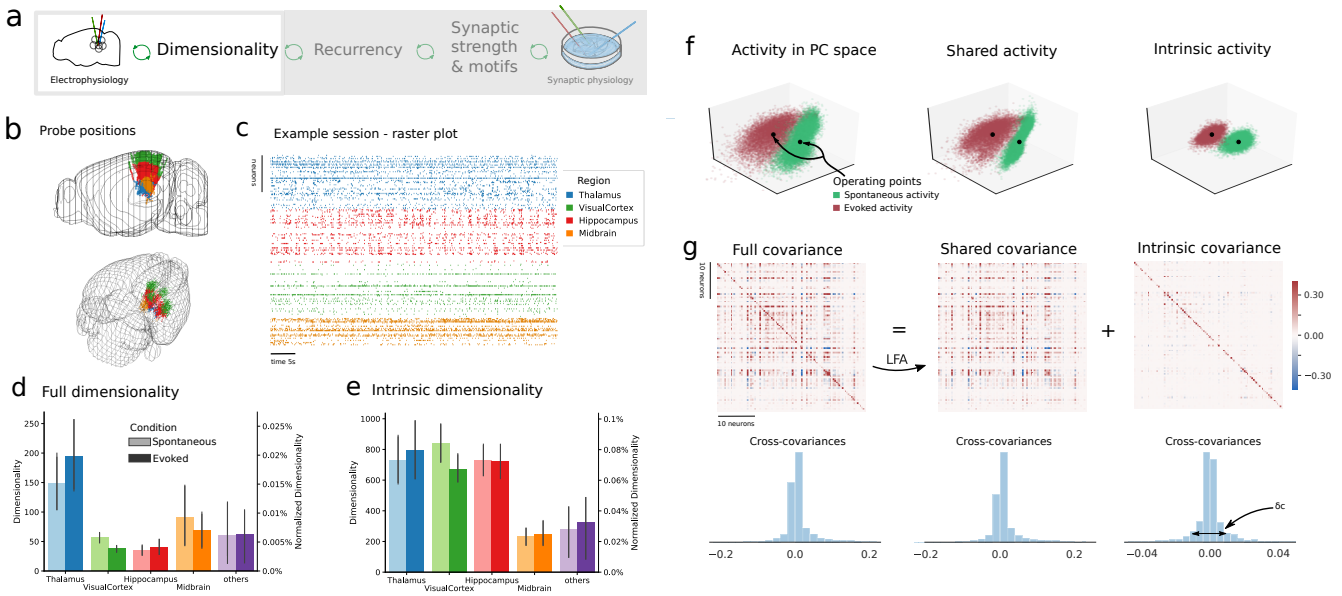

Fig. 1. Intrinsic dimensionality estimation in neural circuits. a) Figure focus: Dimensionality inferred via electrophysiology recordings. b) Sites of Neuropixels recordings colored by brain region. c) Raster plot example of Neuropixels recordings for one experimental session (session id=715093703). d) Dimensionality based on full covariance across brain regions and conditions (evoked and spontaneous) for a network of size $N = 10^6$ neurons (cf. Methods). e) Intrinsic dimensionality based on intrinsic covariance across brain regions and conditions (evoked and spontaneous) for a network of size $N = 10^6$ neurons (cf. Methods). f) Neural activity of example session for spontaneous (green) and evoked (red) condition in the coordinate axes given by the top Principal Components (PC) determined across both conditions. The evoked condition corresponds to drifting grating stimuli with 75 repeats per stimulus orientation. The three panels represent the total, shared, and intrinsic activity, respectively (cf. Methods, Fig. S5). Operating points are defined as the average activity per condition. g) Top panels: Schematic of Latent Factor Analysis decomposition of the full covariance into shared and intrinsic covariances (cf. Fig. S5a). Bottom panels: distribution of cross-covariances for the three matrices.

principal components required to capture roughly 80% of a signal's variability (17) (Fig. S2a).

D_{PR} is defined via the eigenvalues of the covariance matrix and can be rewritten in terms of the statistics of covariances (19) (Fig. S2b) (see Methods). The mean and variance of cross-covariances across neurons define the participation ratio D_{PR} for a large number of recorded neurons N:

$$D_{\mathrm{PR}}(C) \approx \frac{N}{1 + N(m^2 + s^2)}. \qquad \textbf{(1)}$$

Specifically, $m = \frac{\bar{c}}{\bar{a}}$ and $s = \frac{\delta c}{\bar{a}}$ are the ratios between the mean $\bar{c}$ or standard deviation δc of cross-covariances and the average auto-covariance $\bar{a}$, the latter acting as a firing rate normalization (Fig. S2b). The two, m and s, correspond to independent factors influencing dimensionality. For example, a high m can be driven by a small number of "low-rank" behavioral components (42) while s, as we will show, may result from strong recurrent connections between neurons. However, both contribute equally to modulating dimensionality, since even unstructured correlations between neurons, measured by s, can accumulate in large networks to substantially constrain the possible modes of neural activity.

We applied this measure Eq. (1) to large-scale Neuropixels recordings collected across multiple regions of the mouse brain at the Allen Institute for Brain Science, Figs. 1b to 1c (cf. (11, 30)). We analyzed 32,043 neurons across 5 brain regions (Table S1), recorded during sessions lasting on average more than 3 hours (cf. sample of 2 minutes of recorded activity, Fig. 1c). We focused on periods of spontaneous activity (no stimulus was presented to the animal) and evoked activity (where drifting gratings were displayed), cf. Methods and Fig. S1. The results revealed that dimensionality – when extrapolating to realistic cortical network sizes (see Methods) – had values in the order of ~ 100 dimensions (cf. Fig. 1d, Fig. S3a), extremely low when compared to the number of neurons in each brain region (normalized dimensionality, right y-axis) – on the range of 0.01%. Subdividing the data for visual cortex, we further found an increase in dimensionality across visual areas aligned with the underlying functional hierarchy (11) for the evoked condition (Fig. S4), underscoring a likely functional role for dimensionality in the circuits' computations.

However, in electrophysiology recordings there are two factors that could drive such a low dimensionality: low-rank components driven by inputs to the network (7) and intrinsic network connectivity constraining neural dynamics. Distinct from recent studies that have focused on the dimensionality of extrinsic, stimulus-related activity (6, 17), here we focus on intrinsic network connectivity. While sensory inputs may entrain the network's neural responses to be low-dimensional, the question of whether neural activity is inherently low-dimensional captures the effects of recurrent connectivity and therefore indicates the operating regime of the underlying circuits.

To isolate the contribution of intrinsic network connectivity we use a cross-validated Latent Factor Analysis (LFA) (9), which removes low rank components of shared neural variability (Figs. S5a to S5e). We deem the remaining variability intrinsic – inherent to the analyzed brain region – and it may be regarded as an upper bound on the dimensionality of network responses (cf. Fig. 1e, (19)). The resulting dimensionality values were in the hundreds, which was an order of

magnitude greater than the dimensionality of the full covariance; nevertheless, they are extremely low, on the range of 0.1%, when compared to the average number of neurons in each of these brain regions. These values were also consistent across conditions despite being evaluated around distinct operating points (Fig. 1f). We verified our approach by utilizing cross-validated Principal Component Analysis instead of LFA (Fig. S6), as well as simulations in a network model with established ground truth (Fig. S7), which confirmed that our method produced tight upper bounds in dimensionality estimation. In addition, for the spontaneous condition, we utilized a Hidden Markov Model to evaluate the results' robustness with respect to different behaviorally relevant stationary intervals in the activity (Fig. S8); while for the evoked condition, we demonstrated consistency of our findings across drifting grating orientations (Fig. S9).

Strong recurrence as a mechanism for low dimensionality in cortical networks

In the case of cortical areas, we propose a mechanistic theory for the origin of the remarkably low dimensionality ($\lessapprox 0.1\%$) of neural activity, in terms of the reverberation of this activity through the underlying network. Recall that Eq. (1) isolates two factors that can drive low dimensionality, m and s. The theory of balanced cortical networks predicts that there is strong inhibitory feedback which drives nearly asynchronous activity and hence a nearly vanishing average correlation $m \sim 0$ (26, 29), a feature we also find for intrinsic covariances in our cortical data. Consequently the leading factor in determining the dimensionality is the standard deviation of cross-covariances s.

Fig. 2b illustrates how the value of s is determined by the level of recurrency in a balanced network (23, 43): as this recurrency becomes stronger, there is the potential for longer and longer paths that significantly impact the co-variation in activity between each pair of neurons. These longer paths are highly variable from one neuron pair to the next, and this variability drives a wide range in the cross-covariances across neural pairs. This intuition can be formalized through a single number R, derived from the eigenvalues of the connectivity matrix, which characterizes the overall strength of recurrent coupling (see Suppl. Mat. for a formal derivation based on (23, 43) and (25) for an alternative derivation).

Therefore, establishing a three-way link between low-dimensional neural activity (Fig. 2b bottom), large variance of correlations (Fig. 2b center) and strong recurrent connections (Fig. 2b top), we found a direct relationship between dimensionality and recurrency R in the balanced regime:

$$D_{\mathrm{PR}}/N = (1-R^2)^2. \qquad \textbf{(2)}$$

This relationship is extremely robust as shown by our validation in complex nonlinear spiking networks (Fig. 2c and Figs. S10 to S11) and holds for networks with a wide range of topologies, as we will further explore below.

Analyzing the intrinsic dimensionality of activity across cortical layers we found a wide variation across layers, yet consistent across conditions. Intrinsic dimensionality was on the order of 0.1% or less (Fig. 2d), consistent with the hypothesis that cortical circuits operate in a strongly recurrent regime. Layers 2 and 5 had respectively the lowest and highest intrinsic dimensionality, a result consistent with the hypothesis that recurrence in layer 2 is stronger than in layer 5 (44) (Fig. 2d). We then performed the analysis of intrinsic dimensionality for areas along the visual processing hierarchy (Fig. 2f and Fig. S12) (11). Without further subdividing neural activity layerwise, intrinsic dimensionality appeared to be quite constant – consistent with anatomical studies suggesting that connectivity differs less across areas than across layers (45). Dimensionalities in the evoked condition appeared

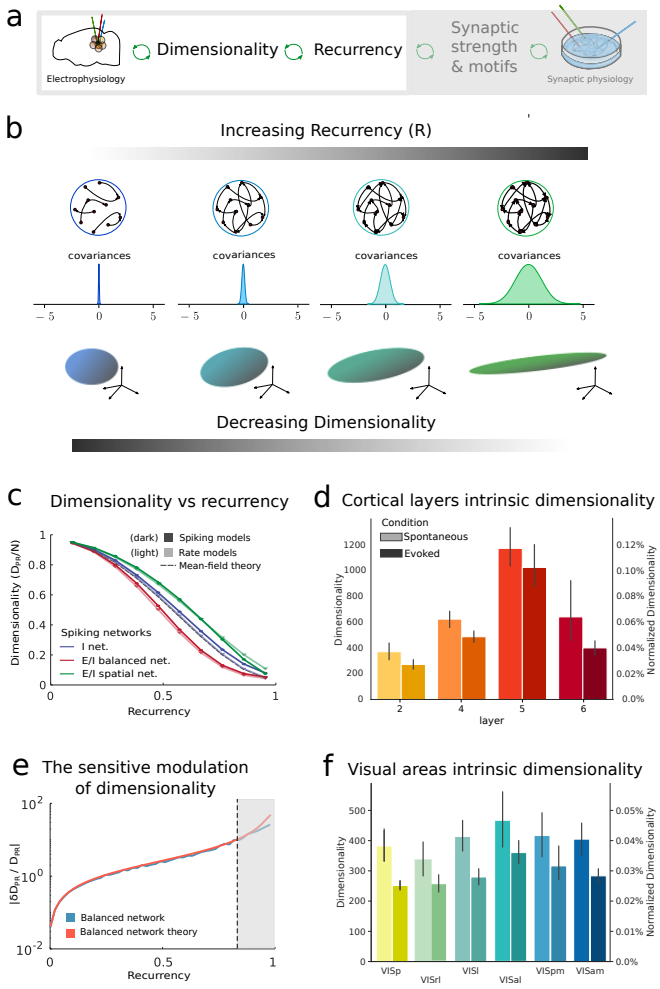

Fig. 2. Dimensionality and recurrency R across the visual cortical circuit. a) Figure focus: Dimensionality is linked to recurrency. b) Three-way connection between recurrency, width of covariances distribution and dimensionality of neural activity. c) Left: Normalized dimensionality D_{PR}/N for a balanced network as a function of the recurrency R. Theoretical predictions for the dimensionality in homogeneous inhibitory networks (gray) are accurate for simulations of rate models (light colors) and spiking models (dark colors) across various network topologies (blue: homogeneous single population inhibitory networks, red: homogeneous two-population excitatory-inhibitory networks, green: spatially organized single-population inhibitory networks). d) Dimensionality of intrinsic covariances across visual cortical layers. Dimensionality values are for networks of size $N = 10^6$ neurons (cf. Methods). e) Relative modulation of dimensionality as a function of the recurrency, Eq. (3). Blue and red curves overlap. The shaded gray area highlights the sensitive regime. f) Dimensionality of intrinsic covariances across visual cortical areas ordered according to the visual cortical hierarchy identified in (11).

to be lower than in the spontaneous condition, suggesting that networks approach more recurrent operating points as they adapt to stimuli.

Overall our analysis shows how trends in the intrinsic dimensionality in strongly coupled, balanced regimes relate to modulations in network recurrency, hypotheses that we revisit in detailed connectivity studies below.

Sensitive control of dimensionality

The connection between D_{PR} and R, coupled with the analysis of cortical areas, suggests that the network's dimensionality in cortical circuits is tightly constrained, being much less than the number of neurons $D_{\mathrm{PR}}/N \lessapprox 0.1\%$. Furthermore, in the same regime, the relative change in dimensionality with respect to the recurrency R is highest (Fig. 2e):

$$\frac{\delta D_{\mathrm{PR}}}{D_{\mathrm{PR}}} = \frac{dD_{\mathrm{PR}}}{dR}\frac{1}{D_{\mathrm{PR}}} = \frac{4R}{R^2-1}. \qquad \textbf{(3)}$$

As a result, with increasing R balanced neural networks gain sensitive control of dimensionality as a function of recurrency. In summary, the low values of D_{PR} obtained for cortical areas of the mouse brain indicate that recurrency for these brain areas is strong (Fig. 2f), suggesting, in turn, that the recurrent strength R has sensitive control over the dimensionality of neural activity (Fig. 2e). We next show that this control can be enacted systematically via the internal structure of recurrent connections.

Local tuning of the global recurrency R

We asked *how* balanced neural networks can regulate their overall recurrency R and hence their dimensionality (Fig. 3a). While many previous studies established how global features of recurrent connectivity affect R (23, 46, 47), here we focus on the impact of local connectivity motifs. These motifs are statistics of the neural connectivity W that involve *pairs* of connections (see Methods), and are the fundamental local building blocks of networks. Second order motifs appear in four types: reciprocal, divergent, convergent, and chain motifs (Fig. 3b), together with the variance (strength) of neural connections already present in purely random models (46). These motifs have been shown to play important roles determining neuron-to-neuron correlations and allied circuit dynamics (32–34, 38, 48–53) and emerge from learning rules consistent with biological STDP mechanisms (54, 55).

We developed a comprehensive theory that takes full account of all second order motifs in networks of excitatory and inhibitory neurons, generalizing allied results developed via distinct theoretical tools (25, 51). Our analysis yields a novel compact analytical quantity that shows how recurrency is modulated by local structure $R = \sigma \cdot R_{\mathrm{motifs}}$ where σ stems from the overall synaptic strength and

$$R_{\mathrm{motifs}} = \frac{1-\tau_{\mathrm{div}}-\tau_{\mathrm{con}}-2\tau_{\mathrm{chn}}+\tau_{\mathrm{rec}}}{\sqrt{1-\tau_{\mathrm{div}}-\tau_{\mathrm{con}}}} \qquad \textbf{(4)}$$

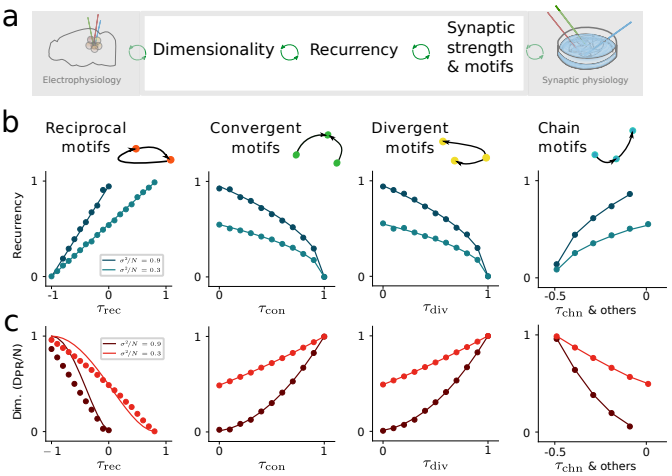

Fig. 3. Theory for recurrency and dimensionality in balanced networks with second-order motifs. a) Figure focus: Modulation of recurrency and dimensionality by local circuit motifs. b) Theoretical dependence of recurrency on motif abundances. c) Theoretical dependence of dimensionality on motif abundances. Solid lines: theory. Markers: simulations.

compactly describes the influence of second order motifs. Here τ_{rec}, τ_{chn}, τ_{div}, τ_{con} denote correlation coefficients between pairs of synapses that capture the abundance of reciprocal, chain, divergent, and convergent motifs, respectively (cf. Methods and Suppl. Mat.). This formula describes how the recurrency R is affected by increasing or decreasing the prevalence of second order motifs (Fig. 3b) and thus links the modulation of auto- and cross-covariances and the dimensionality of neural responses across the global network to the statistics of local circuit connectivity, as shown in Figs. 3b to 3c. While Eq. (4) is exact for the simplest type of balanced networks, which are networks of inhibitory neurons whose recurrent interactions balance the excitatory external input, we show that it generalizes to models of balanced excitatory-inhibitory networks (56). Here, σ and τ combine the corresponding statistics of the excitatory and inhibitory subpopulations (cf. Fig. S13 and Suppl. Mat.). This direct link between quantifiable, local connectivity statistics and the global network property R opened the door to novel functional analyses of very large-scale synaptic physiology datasets in both mouse and human, as we describe next.

Cortical circuits in mouse and human employ local synaptic motifs to modulate their recurrent coupling

We analyzed newly released synaptic physiology datasets from both mouse and human cortex (12, 41) to assess the involvement of synaptic motifs in modulating network recurrency and to probe their possible role in driving the changes in dimensionality seen across layers and conditions in Fig. 2d. This synaptic physiology dataset was based on simultaneous in-vitro recordings of 3-to-8 cell groups (cf. Methods) and consisted of $1,368$ identified synapses from mouse primary visual cortex (out of more than $22,000$ potential connections that were tested) and 363 synapses from human cortex. Recall that the recurrency R as defined

above has an overall scaling term, σ, and a motif contribution term given by Eq. (4). We begin by assessing the probability of occurrence of individual motifs and hence estimating R_{motifs}, cf. Methods. The relationship Eq. (4) defines a specific hypothesis for empirical motif statistics that modulate global circuit dimensionality: if they combine to produce $R_{\text{motifs}} > 1$ then they are tuned to reduce dimensionality, and vice-versa for $R_{\text{motifs}} < 1$.

Beginning with the mouse data, we calculated the statistics of individual motifs, separating those for excitatory (E) and inhibitory (I) synapses (EE, EI, II, Fig. 4b), and found many motifs to be significantly present. We then combined these to compute R_{motifs}. This requires two parameters: one regulating the overall ratio of inhibitory to excitatory neurons (γ), and another the relative strength of the inhibitory synapses (g) (cf. Suppl. Mat.). We found that $R_{\text{motifs}} < 1$ across all choices of these parameters (Fig. 4c).

Many different motifs, distinctly involving E and I cell types, combine to produce this value of R_{motifs}. To study how this occurs, we separated the contribution to R_{motifs} from motifs within the excitatory population (EE type only) by assuming that other motifs occur at chance level. Interestingly, the EE motifs operating alone produced the opposite trend, increasing the radius $R^{\text{EE only}}_{\text{motifs}} > 1$ (Fig. 4d center, one-sided t-test p-value$< 10^{-20}$). The same was true for motifs within the inhibitory population $R^{\text{II only}}_{\text{motifs}} > 1$ (Fig. 4d right, one-sided t-test p-value$< 10^{-20}$), and for motifs within the excitatory population in human cortical circuits (Fig. 4e). We further confirmed that this effect is also predicted for previously published data on excitatory connections in rat visual cortex (36) (cf. Methods). The increased recurrent coupling strengths within both the excitatory and inhibitory populations underscore the prominent role of EI motifs, specifically reciprocal EI motifs, in decreasing and potentially regulating the overall recurrency to be $R_{\text{motifs}} < 1$ (Fig. 4c, Fig. S13).

We found evidence that synaptic motifs contribute to the cross-layer differences in the dimensionality of cortical activity identified above (Fig. 2d). There, activity in mouse cortex layer 2 showed lower dimensionality, corresponding to an increased overall recurrency $R = \sigma R_{\text{motifs}}$ compared to layers 4, 5 or 6. Intriguingly, the corresponding motif contribution R_{motifs} was significantly stronger for layer 2 than for layer 5 (Fig. 4f left), suggesting that motifs play a role in increasing R. Moreover, a similar result held true when performing the analysis on the human dataset for excitatory connections in layers 2 and 3 (Fig. 4g and Figs. S14 to S15). Overall, the distinct roles of motifs among E and I cells types in regulating R_{motifs} point to ways that the recurrency, and hence dimension, may be controlled dynamically in neural circuits.

One pathway for this control is via cell types, which subdivide E and I populations (Fig. 4h) and are separately identified in the synaptic physiology dataset which we analyze. As Table S2e shows, reciprocal EI motifs were prevalent when the inhibitory interneuron was a somatostatin cell (SST) or a parvalbumin cell (PV), but not a VIP cell. Recent findings have shown that VIP interneurons (58, 59) are important regulators of cortical functions, are modulated by arousal and movement (60), and are recruited by reinforcement signals (61). We thus hypothesized that VIP interneurons could adjust the recurrent coupling of cortical circuits by exerting disinhibitory control via the SST population (57, 58) (this could occur while preserving the balanced regime, given the very sparse connectivity within the VIP population and from VIP to pyramidal cells (Fig. 4h)). Under the simplest form of this

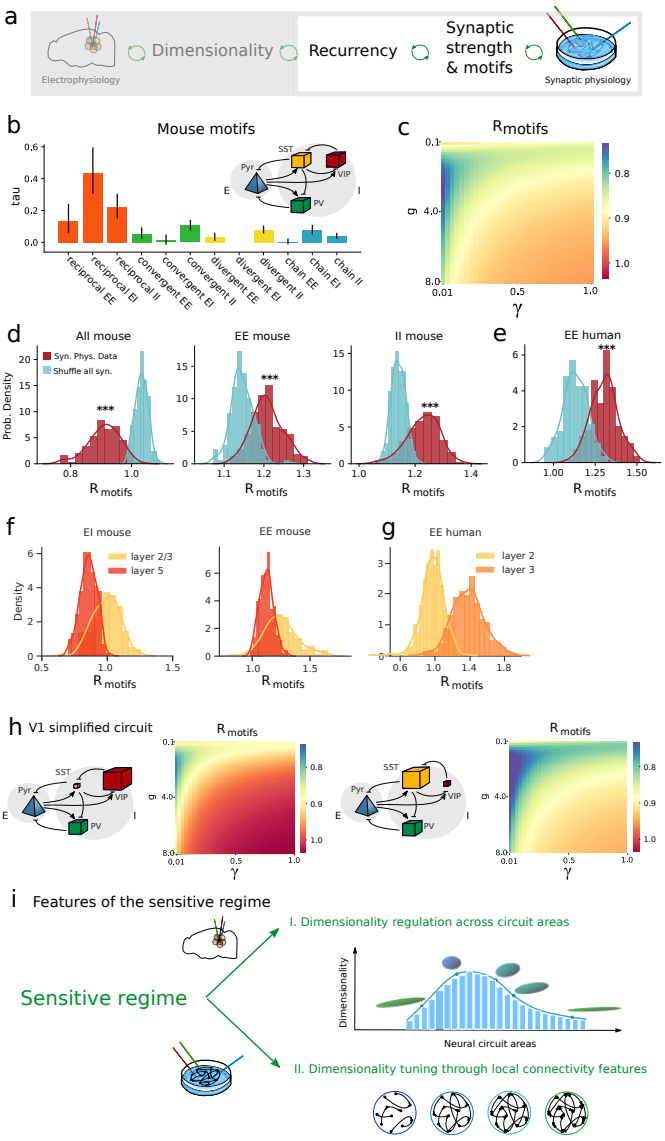

Fig. 4. Motif analysis in synaptic physiology datasets. a) Figure focus: network recurrency inferred via synaptic physiology datasets. b) Motif abundances in mouse V1. Inset: Simplified V1 circuit diagram with only prevalent connections (57). c) Inferred R_{motifs} as a function of relative strength g of inhibitory and excitatory synapses and ratio γ of inhibitory to excitatory population size. d) Estimation of R_{motifs} from mouse data (500 bootstraps based on random subsets of 80% of sessions). Shuffle of synapses within each experimental session preserving EI synapse type (shuffle EI syn.). Effect of all EI motifs ($g = 4$ and $\gamma = 0.25$). e) Same as d for EE motifs in human dataset. f) Layer-wise estimation of R_{motifs} for EI balanced motifs (left), EE-only motifs (right) in mouse. g) Layer-wise estimation of R_{motifs} for EE-only motifs in human. h) Effect of VIP regulation on R_{motifs}. Left: SST is inactive and R_{motifs} is computed over the displayed circuit involving PV and VIP. Right: SST is active and VIP is inhibited resulting in PV and SST balancing the activity of the Pyr population. i) Schematic of sensitive regime and its advantages.

hypothesis, there exists a mutual antagonism between VIP and SST populations that results in only one of these populations being active at a time. We derived the values of R_{motifs} in this case, and found that activation of the VIP pathway substantially increased R_{motifs} (Fig. 4h) and hence decreased predicted dimensionality. This shows how VIP interneurons, which themselves may collect top-down signals from higher cortical areas, can selectively tune the dimensionality of local cortical activity. This adds another channel for population-level control of information processing in cortical circuits, on top of existing hypotheses for on how VIP neurons regulate gain in individual neurons (59). Furthermore, we predicted that a similar trend of increasing R_{motifs} follows from short-term synaptic plasticity (STP) in modulating cell-type specific connections upon stimulus onset, although a detailed analysis awaits future investigation (Figs. S16a to S16c). Finally, we note that our results were robust to inclusion of estimates of the relative synaptic strength of cell type specific connections (Fig. S16d) and the cell type specific prevalence of the three inhibitory subpopulations (Fig. S16d) (see Methods).

In sum, in this section we asked whether the experimentally derived structure of cortical networks – quantified by their motifs – enables the tuning of the recurrency R and hence dimensionality. We found that the answer is yes, and that the VIP disinhibitory pathway and STP both provide examples of how motifs are likely to play a substantial role in this tuning. Specifically, upon accounting for STP modulation, our preliminary analysis suggested that recurrency is increased and, henceforth, dimensionality decreased; this is consistent with the finding that intrinsic dimensionality is lower for evoked, rather than spontaneous, activity across visual cortical areas Fig. 2f. As we reviewed above, high dimensional activity can retain stimulus details, while lower dimensional activity can promote robust and general downstream decoding. Taken together, this points to new functional roles for modulatory and adaptive mechanisms known to take effect across time during stimulus processing and to be engaged across brain states.

Summary and discussion

We showed that neural networks across the mouse cortex operate in a strongly recurrent regime, in which the dimensionality of their activity is much smaller than the number of neurons. A feature of circuits in this regime is the ability to sensitively modulate the relative dimensionality of their activity patterns via their recurrency R, a unifying measure of a network's overall recurrent coupling strength (Fig. 4i). This has potentially important consequences for computation. Indeed, our analyses of large scale Neuropixels recordings from the cortex showed systematic trends in this dimensionality across cortical layers and stimulus conditions. Our theory links these findings to clear predictions for the recurrency in cortical areas: a higher dimensionality suggests a lower recurrency and vice-versa. Moreover, we showed that the critical circuit features that determine a circuit's recurrency R – and hence the dimensionality of its activity patterns – are not just its overall synaptic strength, but also a tractable set of local synaptic motifs. We use theoretical tools to quantify the effect of these motifs via a compact index R_{motifs}. This provides a concrete target quantity that can, as we show, be readily obtained from emerging, large-scale synaptic connectivity datasets and used to check predictions about the role of synaptic structure in controlling dimensionality. Thus theory and brain-wide experimental analyses converge to provide new evidence for an intriguing concept (51, 62, 63): that the connectivity of cortical brain networks exert global control over their activity in a highly local and tractable manner, via the building blocks of their local circuitry (Fig. 4i). This concept may extend beyond cortex: indeed, individual areas in hippocampal and thalamic circuits also show systematic trends in dimensionality (Fig. S4 and Fig. S12) whose mechanistic origins could be similar.

Acknowledgments

We thank Krešimir Josić, Brent Doiron, Luca Mazzucato, Stefan Mihalas, Nicholas Steinmetz, Michael Okun, Byron Yu, Daniel Denman, Yu Hu, Severine Durand and Leenoy Meshulam for helpful discussions. S.R. was supported by a Swartz Fellowship in Theoretical Neuroscience at the University of Washington, and by NIH BRAIN Grant R01EB026908, and E.S.B. by NIH R01EB026908 and NSF DMS Grant 1514743. D.D. and M.H. were supported by the HGF young investigator's group VH-NG-1028, the European Union's Horizon2020 research and innovation program under Grant agreements No. 785907 (Human Brain Project SGA2) and No.945539 (Human Brain Project SGA3), and funded under the Excellence Strategy of the Federal Government and the Länder (G:(DE-82)EXS-PF-JARA-SDS005). We thank the Allen Institute for Brain Science founder, Paul G. Allen, for his vision, encouragement, and support.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Code availability

The code for the numerical simulations and data analyses are immediately available from the corresponding author upon reasonable request, and will be released in a public repository before final publication of the manuscript.

Bibliography

1. S. Fusi, E. K. Miller, M. Rigotti, *Current opinion in neurobiology* **37**, 66 (2016).
2. A. Litwin-Kumar, K. D. Harris, R. Axel, H. Sompolinsky, L. F. Abbott, *Neuron* **93**, 1153 (2017).
3. M. Rigotti, *et al.*, *Nature* **497**, 585 (2013). Number: 7451 Publisher: Nature Publishing Group.
4. N. A. Cayco-Gajic, C. Clopath, R. A. Silver, *Nature Communications* **8**, 1116 (2017).
5. J. P. Cunningham, B. M. Yu, *Nature Neuroscience* **17**, 1500 (2014).
6. C. Stringer, M. Pachitariu, N. Steinmetz, M. Carandini, K. D. Harris, *Nature* **571**, 361 (2019).
7. C. Stringer, *et al.*, *Science* **364** (2019).
8. C. Huang, *et al.*, *Neuron* **101**, 337 (2019).
9. R. C. Williamson, *et al.*, *PLOS Computational Biology* **12**, 1 (2016).
10. T. A. Engel, N. A. Steinmetz, *Current opinion in neurobiology* **58**, 181 (2019).
11. J. H. Siegle, *et al.*, *Nature* **592**, 86 (2021).
12. L. Campagnola, *et al.*, *bioRxiv* (2021).
13. T. M. Cover, *IEEE Transactions on Electronic Computers* **EC-14**, 326 (1965).
14. V. N. Vapnik, *Statistical learning theory* (1998).
15. M. Farrell, S. Recanatesi, G. Lajoie, E. Shea-Brown, *bioRxiv* p. 564476.

16. V. Papyan, X. Han, D. L. Donoho, *Proceedings of the National Academy of Sciences* **117**, 24652 (2020).
17. P. Gao, *et al.*, *BioRxiv* p. 214262 (2017).
18. J. A. Gallego, M. G. Perich, L. E. Miller, S. A. Solla, *Neuron* **94**, 978 (2017).
19. L. Mazzucato, A. Fontanini, G. La Camera, *Frontiers in Systems Neuroscience* **10** (2016).
20. U. Cohen, S. Chung, D. D. Lee, H. Sompolinsky, *Nature communications* **11**, 1 (2020).
21. S. Recanatesi, *et al.*, *bioRxiv* p. 471987 (2019).
22. M. e. a. Stern, In the footsteps of learning: Changes in network dynamics and dimensionality with task acquisition. (2020). COSYNE Conference Abstract.
23. D. Dahmen, S. Grün, M. Diesmann, M. Helias, *Proceedings of the National Academy of Sciences* **116**, 13051 (2019).
24. S. Recanatesi, *et al.*, *arXiv preprint arXiv:1906.00443* .
25. Y. Hu, H. Sompolinsky, *bioRxiv* (2020).
26. T. Tetzlaff, M. Helias, G. T. Einevoll, M. Diesmann, *Plos Computational Biology* **8**, e1002596 (2012). WOS:000308553500003.
27. C. van Vreeswijk, H. Sompolinsky, *Science (New York, N.Y.)* **274**, 1724 (1996).
28. A. S. Ecker, *et al.*, *Science* **327**, 584 (2010). WOS:000274020500039.
29. A. Renart, *et al.*, *Science* **327**, 587 (2010). WOS:000274020500040.
30. V. C. . . S. dev documentation, https://allensdk.readthedocs.io/en/v2.2.0/visual_coding_neuropixels.
31. E. Schneidman, M. J. Berry, R. Segev, W. Bialek, *Nature* **440**, 1007 (2006). WOS:000236906000026.
32. Y. Hu, J. Trousdale, K. Josic, E. Shea-Brown, *Physical Review E* **89**, 032802 (2014). WOS:000332672200015.
33. V. Pernice, B. Staude, S. Cardanobile, S. Rotter, *Plos Computational Biology* **7**, e1002059 (2011). WOS:000291015800032.
34. J. Trousdale, Y. Hu, E. Shea-Brown, K. Josic, *Plos Computational Biology* **8**, e1002408 (2012). WOS:000302244000016.
35. Y. Hu, J. Trousdale, K. Josic, E. Shea-Brown, *Journal of Statistical Mechanics-Theory and Experiment* p. P03012 (2013). WOS:000316056900012.
36. S. Song, P. J. Sjostrom, M. Reigl, S. Nelson, D. B. Chklovskii, *Plos Biology* **3**, 507 (2005). WOS:000227984000018.
37. R. Perin, T. K. Berger, H. Markram, *Proceedings of the National Academy of Sciences* p. 201016051 (2011).
38. L. Zhao, B. Beverlin, T. Netoff, D. Q. Nykamp, *Frontiers in Computational Neuroscience* **5** (2011).
39. O. Sporns, *Discovering the human connectome* (MIT press).
40. F. Mastrogiuseppe, S. Ostojic, *Plos Computational Biology* **13**, e1005498 (2017). WOS:000402542900036.
41. S. P. A. documentation, https://portal.brain-map.org/explore/connectivity/synaptic-physiology.
42. R. Rosenbaum, M. A. Smith, A. Kohn, J. E. Rubin, B. Doiron, *Nature Neuroscience* **20**, 107 (2017). WOS:000391085500018.
43. D. Dahmen, *et al.*, *bioRxiv* (2020).
44. S. Peron, *et al.*, *Nature* **579**, 256 (2020).
45. J. A. Harris, *et al.*, *Nature* **575**, 195 (2019).
46. H. Sompolinsky, A. Crisanti, H.-J. Sommers, *Physical review letters* **61**, 259 (1988).
47. J. Aljadeff, D. Renfrew, M. Vegué, T. O. Sharpee, *Physical Review E* .
48. N. Brunel, *Nature Neuroscience* **19**, 749 (2016).
49. D. Martí, N. Brunel, S. Ostojic, *Physical Review E* **97**, 062314 (2018). Publisher: American Physical Society.
50. G. K. Ocker, *et al.*, *Current Opinion in Neurobiology* **46**, 109 (2017).
51. S. Recanatesi, G. K. Ocker, M. A. Buice, E. Shea-Brown, *PLoS computational biology* **15**, e1006446.
52. D. Zhang, C. Zhang, A. Stepanyants, *Journal of Neuroscience* **39**, 6888. Publisher: Soc Neuroscience.
53. Y. Hu, J. Trousdale, K. Josić, E. Shea-Brown, *Journal of Statistical Mechanics: Theory and Experiment* **2013**, P03012 (2013).
54. M. Gilson, A. N. Burkitt, D. B. Grayden, D. A. Thomas, J. L. van Hemmen, *Biological Cybernetics* **101**, 427 (2009). WOS:000272176000008.
55. G. K. Ocker, A. Litwin-Kumar, B. Doiron, *Plos Computational Biology* **11**, e1004458 (2015). WOS:000360824500049.
56. N. Brunel, *Journal of Computational Neuroscience* **8**, 183 (2000). WOS:000087725300001.
57. D. J. Millman, *et al.*, *Elife* **9**, e55130 (2020).
58. M. M. Karnani, *et al.*, *Journal of Neuroscience* **36**, 3471 (2016).
59. K. A. Ferguson, J. A. Cardin, *Nature Reviews Neuroscience* **21**, 80 (2020).
60. Y. Fu, *et al.*, *Cell* **156**, 1139 (2014).
61. H.-J. Pi, *et al.*, *Nature* **503**, 521 (2013).
62. V. Pernice, B. Staude, S. Cardanobile, S. Rotter, *Physical Review E* **85**, 031916 (2012). WOS:000302117900006.
63. Y. Hu, *et al.*, *Phys. Rev. E* **98**, 062312 (2018).
64. M. M. Churchland, *et al.*, *Nature neuroscience* **13**, 369 (2010).