

Network structure of neural systems supporting cascading dynamics predicts stimulus propagation and recovery

Harang Ju (harang@upenn.edu)

Neuroscience Graduate Group, 240 S 33rd St
Philadelphia, PA 19104 USA

Jason Z. Kim (jinsu1@upenn.edu)

Department of Bioengineering, 240 S 33rd St
Philadelphia, PA 19104 USA

Danielle S. Bassett (dsb@upenn.edu)

Departments of Bioengineering, Electrical & Systems Engineering,
Physics & Astronomy, Neurology, and Psychiatry, 240 S 33rd St
Philadelphia, PA 19104 USA

Abstract

Many neural systems display cascading behavior characterized by uninterrupted sequences of neuronal firing. When the distributions of cascade size and duration follow a power law, theoretical models suggest that such dynamics support optimal information transmission and storage. However, the unknown role of network structure on neural dynamics precludes an understanding of how variations in network structure either support or impinge upon information processing. Here, we develop a theoretical understanding of how network structure supports information processing through network dynamics and validate our theory with empirical data. Using a generalized spiking model and mathematical intuitions from linear systems theory, network control theory, and information theory, we show how network structure can be designed to temporally extend the propagation and recovery of certain stimulus patterns. Moreover, we observe cycles as structural and dynamic motifs that are prevalent in such networks. Broadly, our results demonstrate how network structure constrains cascading dynamics and supports persistent activation that could potentially contribute to cognitive faculties, such as working memory or attention.

Keywords: linear systems; network control; mutual information

In many neural systems, neurons display spontaneous cascades of activity, in which neurons spike in patterns across consecutive time windows. These cascades often form heavy-tailed distributions of size and duration. When such distributions follow a power law, those cascades are called avalanches and have been linked to optimal information processing in critical systems (Beggs & Plenz, 2003; Kinouchi & Copelli, 2006). However, often left implicit in the analysis of this phenomenon are (i) the computational function of avalanches without the assumption of self-organized criticality (Priesemann et al., 2014; Touboul & Destexhe, 2017) and (ii) the network structure that supports various cascading dynamics (Perin, Berger, & Markram, 2011; Brunel, 2016). Thus, it is

not yet clear how network structures shape cascading dynamics, which in turn perform computations (Larremore, Shew, & Restrepo, 2011; Chambers & MacLean, 2016).

Here, we address this gap in knowledge through a series of analyses and numerical simulations of a stochastic neural network model instantiated upon various network structures. Using linear systems theory, network control theory, and information theory, we demonstrate that network structure constrains cascade duration, identify topological features that extend cascade duration, and show that long cascade duration allows recovery of stimuli. Importantly, we empirically validate these theoretical results with multielectrode array (MEA) recordings from neurons in the mouse somatosensory cortex (Ito et al., 2016). Collectively, our findings show that the network topology reported extensively in the empirical literature supports the persistent activation of a cluster of neurons, which in turn allows for the recovery of stimulus patterns implicated in working memory (Durstewitz, Seamans, & Sejnowski, 2000; Eriksson, Vogel, Lansner, Bergström, & Nyberg, 2015).

Network structure constrains cascade duration

To investigate the role of network structure, we first formalize a network as a directed graph of nodes $\mathcal{V} = \{1, \dots, n\}$ and edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$, and we represent this graph as a weighted adjacency matrix $A = [a_{ij}]$ (Figure 1a). We model the activity of an n -neuron network as a binary vector $\mathbf{y}(t) \in \{0, 1\}^n$ that evolves stochastically as $y_i(t) \sim B(\mathbf{a}_i \cdot \mathbf{y}(t-1))$, where $B(p)$ is a Bernoulli process with probability $\min(0, \max(1, p))$. The average state of the stochastic model can be written as $\mathbb{E}[\mathbf{y}(t)] = \mathbf{x}(t)$, and given equal initial states $\mathbf{x}(0) = \mathbf{y}(0)$ and $\forall i \in \mathcal{V} : \sum_j a_{ij} \leq 1$, this average neural activity evolves as a linear system $\mathbf{x}(t) = A\mathbf{x}(t-1)$ (Ju, Kim, & Bassett, 2019).

From this formulation, we can analytically represent any pattern of neural activity $\mathbf{y}(t)$ at time t as a state $\mathbf{s}^i \in \{0, 1\}^n$. Then, we can write the evolution of activity as a Markov process, where the probability that the network is in any state is given by the probability vector $\mathbf{p}(t) = [P(\mathbf{y}(t) = \mathbf{s}^1); \dots; P(\mathbf{y}(t) = \mathbf{s}^{2^n})]$. As a Markov process, we can write the evolution of this probability vector as a linear map according



to a transition matrix T such that $\mathbf{p}(t) = T\mathbf{p}(t-1) = T^t\mathbf{p}(0)$. Then, the fraction of cascades that terminate by time t is simply given by the first entry of $\mathbf{p}(t)$ where $\mathbf{s}^1 = \mathbf{0}$ (Figure 1b). Because the transition matrix is an explicit and deterministic function of the network $T = f(A)$ (Ju et al., 2019), the cascade duration is analytically constrained by the network A .

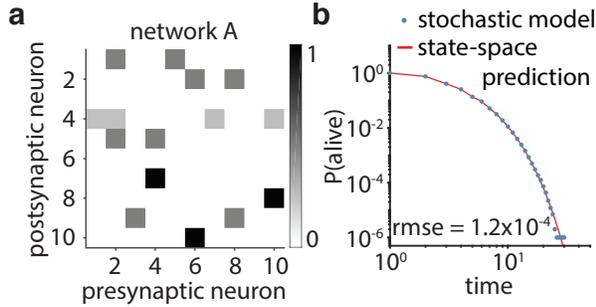


Figure 1: **Network structure constrain cascade duration.** **a**, An example of a network adjacency matrix. **b**, The estimation of the fraction of neurons alive at time t in simulations of the network in panel **a**.

As a more general and tractable description, we then numerically demonstrated the role of network structure on cascade duration by using eigenvalue analysis. Because the average of the stochastic model $\mathbb{E}[\mathbf{y}(t)] = \mathbf{x}(t)$ evolves linearly as $\mathbf{x}(t) = A\mathbf{x}(t-1)$, we used the dominant eigenvalue of the network matrix $\lambda_1 = \max_{\lambda_i \in \text{eig}(A)} |\lambda_i|$ to predict the distribution of cascade duration. We fitted such a distribution using maximum likelihood estimation (MLE) as a truncated power law $p(x) \sim x^{-\alpha}e^{-x/\tau}$ with parameters α , the log-log slope of the power law, and τ , the scale of the exponential truncation (Clauset, Shalizi, & Newman, 2009; Alstott, Bullmore, & Plenz, 2014). To avoid overgeneralizing, we bounded τ by the max cascade duration and denoted this new value as τ' . In simulations of cascades on 192 networks of 2^8 nodes, the parameters α and τ' were monotonically and linearly correlated with λ_1 , as reflected in a Spearman's ρ of 0.97 ($p \approx 0$) and a Pearson's r of 0.99 ($p = 3.9 \times 10^{-31}$; $\lambda_1 > 0.8$), respectively (Figure 2a,b).

To empirically validate this theory, we used λ_1 to predict cascade duration in 25 multielectrode (MEA) recordings of spiking neurons in the mouse somatosensory cortex (Ito et al., 2016). To measure λ_1 , we derived directed networks from the recordings using vector autoregression (Neumaier & Schneider, 2001; Schneider & Neumaier, 2001). Then, we binned the spikes with 2ms bins, identified cascades of consecutively active bins, and fitted distributions of cascade duration to truncated power laws. In these recordings, the parameters α and τ' are monotonically and linearly correlated with λ_1 , as reflected in a Spearman's ρ of 0.78 ($p = 7.4 \times 10^{-6}$) and Pearson's r of 0.62 ($p = 9.9 \times 10^{-4}$), respectively (Figure 2c,d). Moreover, in simulations of cascades on these empirical networks, we find a significant positive correlation between τ' and

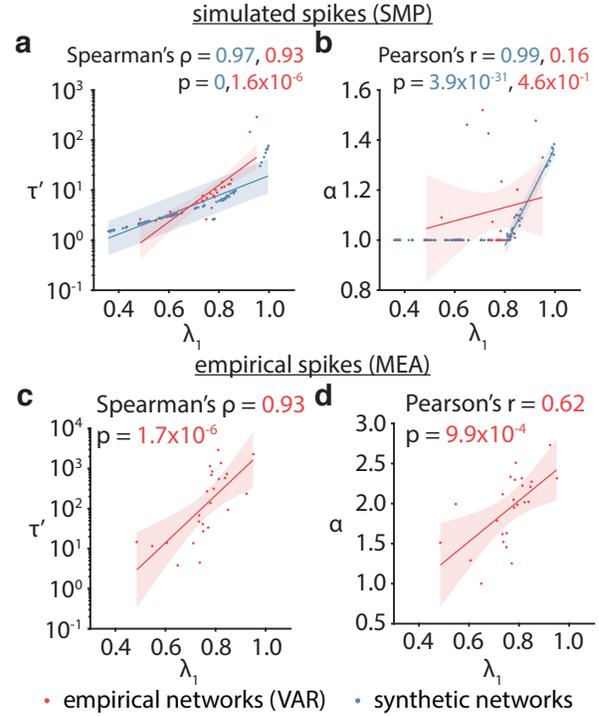


Figure 2: **Eigenvalue analysis captures distributions of cascade duration.** The dominant eigenvalue of a network monotonically scales both parameters τ' (**a**, **c**) and α (**b**, **d**). **a-b**, This relation holds in simulations of the stochastic model on synthetic (red) and empirically derived (blue) networks. **c-d**, We validate these relations in spiking neurons from MEA recordings.

λ_1 , with a Spearman's ρ of 0.93 ($p = 1.7 \times 10^{-6}$; Figure 2a,b).

Cyclical network topology extends cascades

Previous empirical studies have found cycles in the topology of cortical connectivity (Wang et al., 2006; Ko et al., 2011). Here, we hypothesized that these cycles allow networks to display heavy-tailed distributions of cascade duration. To test this hypothesis, we simulated cascades in networks with varying degrees of cycle density, defined as the number of simple cycles divided by the number of edges. We found that as cycle density increases with edge rewiring, cascade duration also increases with a Pearson's correlation coefficient of $r = 0.82$ ($p = 1.1 \times 10^{-27}$; Figure 3a).

Depending on the extent of refractoriness, these cycles, even if they exist structurally, may not support cyclical propagation of activity. Thus, to ensure that activity can propagate through cycles in neuronal systems, we measured the number of n -cycles in the MEA data, where an n -cycle occurs when a neuron spikes again after n time bins of its previous spike. We found that 1-, 2-, 3-, and 4-cycles occur an average of 1.9 times per cascade, with an average of $2.7 \times 10^5 \pm 1.6 \times 10^4$ cascades per recording (standard error; Figure 3b).

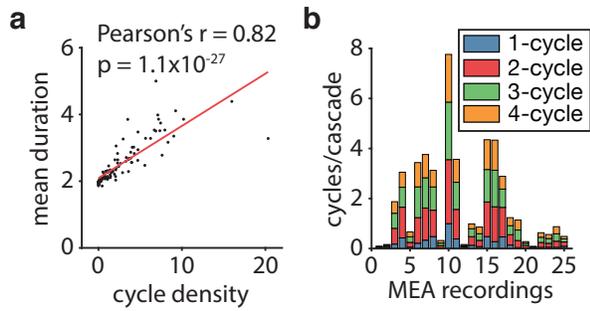


Figure 3: Cyclical network topology extends cascades. **a**, Increasing cycle density through rewiring edges extends cascade duration. **b**, In MEA recordings, spikes propagate frequently through cycles.

Highly controllable neurons extend cascades

Even within a single network, the duration of cascades can vary depending on the nodes that are stimulated, either spontaneously or exogenously. Recently, studies have used network control theory to determine the capacity of a node in controlling the state of a network (Yan et al., 2017). In particular, one metric called *average controllability* measures the impulse response of a node (Pasqualetti, Zampieri, & Bullo, 2014), and thus, we hypothesized that it can predict cascade duration.

To test this hypothesis, we simulated cascades on 100-node random networks by stimulating individual nodes. We observed that the mean cascade duration and finite average controllability (a finite version of average controllability) were significantly positively correlated with a Pearson's correlation coefficient of $r = 0.93$ ($p = 1.0 \times 10^{-43}$; Figure 4a). In expanded simulations of 30 random instantiations of the network, we observed consistent effects with a median Pearson's correlation coefficient of $r = 0.90$.

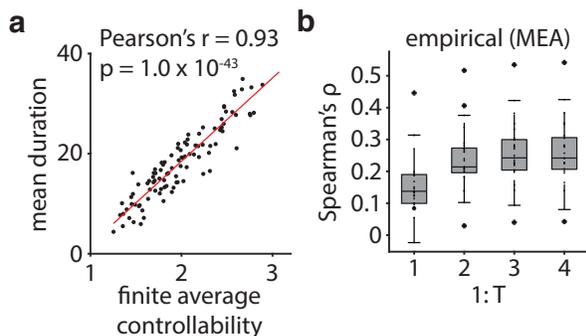


Figure 4: Stimulation in neurons with high controllability extends cascade duration. **a**, Finite average controllability of a stimulated neuron is highly correlated with mean duration of the resulting cascades. **b**, In MEA recordings, controllability of neurons in the first T time bins predicts cascade duration.

To empirically validate this theory, we tested these predic-

tions in the same MEA recordings used previously. We find that cascade duration is correlated with mean average controllability of active neurons in the initial T states with an average Spearman's ρ of 0.16 for $T = 1$ and 0.24 for $T = 2$ ($p < 0.001$, Bonferroni-corrected; Figure 4b). Considering that the cascades are stochastic and cannot be predicted deterministically, we find it notable to observe any correlation between controllability and cascade duration in empirical data. These results suggest that stimulus patterns must be tailored for a network to produce the desired neural dynamics and lead to the question of how stimulation, either endogenous or exogenous, can be used for information processing.

Lasting cascades allow information recovery

If certain networks and stimulus patterns can produce long-lasting cascades, how can these long-lasting cascades contribute to information processing? Intuitively, one cannot recover information about stimuli from cascades that have already terminated, but for lasting cascades, network states can be discriminated and provide information about stimuli. Such delayed recovery of stimuli can allow the associative learning of stimuli across temporal delays (Durstewitz et al., 2000; Eriksson et al., 2015).

To test this intuition, we first show in our Markov formulation that lasting cascades allow information recovery. We define stimulus recoverability as the mutual information $I(S; Y_t)$ between initial states $\mathbf{y}(0) \in S$ and states $\mathbf{y}(t) \in Y_t$ at time t . Given two initial states $\mathbf{y}^i(0)$ and $\mathbf{y}^j(0)$, the probability vectors of the two cascades evolve as $\mathbf{p}^i(t) = T^t \mathbf{y}^i(0)$ and $\mathbf{p}^j(t) = T^t \mathbf{y}^j(0)$. For quickly decaying systems, $\mathbf{p}^i(t)$ and $\mathbf{p}^j(t)$ will both have high probabilities of the zero state \mathbf{s}^1 , inherently reducing recoverability.

To numerically demonstrate this relationship, we simulated cascades on four types of networks (weighted random, random geometric, modular, Watts-Strogatz) and measured the mutual information $I(S; Y_t)$ at each t and cascade duration. Consistent with our intuition, we observe that mutual information is maintained longer when cascades last longer on average. To quantify the relationship between cascade duration and the decay in mutual information, we calculated the Pearson's correlation between average cascade duration and the slope of the linear regression between mutual information and time (Figure 5a). For all network types, we found mean correlation coefficients greater than 0.9 (Figure 5b). Collectively, these findings demonstrate the interplay among network structure, dynamics, and information processing.

Conclusion

Neural systems display cascading dynamics that harbor the marks of a complex underlying network structure and support diverse range of computations (Beggs & Plenz, 2003; Kinouchi & Copelli, 2006). Yet, precisely how network structure supports computations through cascading dynamics remains unclear. Here, using the rich mathematical properties of linear systems, we describe how network structure and stimulus patterns together determine the manner in which a stimulus

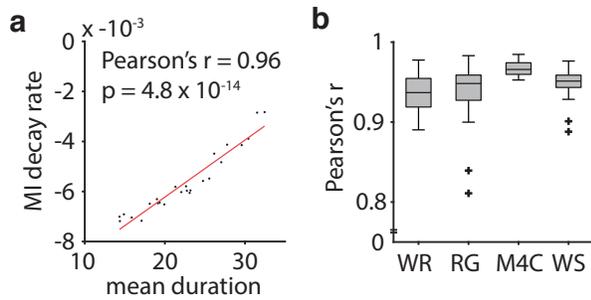


Figure 5: **Lasting cascade allow stimulus recovery.** **a**, In an example network, mean duration of a stimulus pattern is correlated with the rate of decay in mutual information (MI) between the stimulus pattern $y(0)$ and a later network state $y(t)$. **b**, The mean correlation between MI decay rate and mean duration is above 0.9 in simulations of four graph types (WR: weighted random; RG: random geometric; M4C: modular with 4 communities; WS: Watts-Strogatz, small world).

propagates through the network. We then demonstrate that long-lasting cascades can allow for temporally delayed recovery of desired patterns of stimulation. Importantly, we validate these results in empirical data. Broadly, our work blends dynamical systems theory, network control theory, information theory, and computational neuroscience to address the wide gap in the field's current understanding of the relations between architecture, dynamics, and computation.

Acknowledgments We acknowledge support from the Paul G. Allen Family Foundation.

References

- Alstott, J., Bullmore, E., & Plenz, D. (2014). power-law: A python package for analysis of heavy-tailed distributions. *PLOS ONE*, *9*(1), 1-11. doi: 10.1371/journal.pone.0085777
- Beggs, J. M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of Neuroscience*, *23*(35), 11167–11177. doi: 10.1523/JNEUROSCI.23-35-11167.2003
- Brunel, N. (2016, 04). Is cortical connectivity optimized for storing information? *Nature Neuroscience*, *19*, 749 EP. doi: 10.1038/nn.4286
- Chambers, B., & MacLean, J. N. (2016). Higher-order synaptic interactions coordinate dynamics in recurrent networks. *PLOS Computational Biology*, *12*(8), 1-23. doi: 10.1371/journal.pcbi.1005078
- Clauset, A., Shalizi, C., & Newman, M. (2009). Power-law distributions in empirical data. *SIAM Review*, *51*(4), 661-703. doi: 10.1137/070710111
- Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000, 11). Neurocomputational models of working memory. *Nature Neuroscience*, *3*, 1184 EP. doi: 10.1038/81460
- Eriksson, J., Vogel, E. K., Lansner, A., Bergström, F., & Nyberg, L. (2015). Neurocognitive architecture

- of working memory. *Neuron*, *88*(1), 33 - 46. doi: 10.1016/j.neuron.2015.09.020
- Ito, S., Yeh, F.-C., Timme, N. M., Hottowy, P., Litke, A. M., & Beggs, J. M. (2016). Spontaneous spiking activity of hundreds of neurons in mouse somatosensory cortex slice cultures recorded using a dense 512 electrode array. *CR-CNS.org*. doi: 10.6080/K07D2S2F
- Ju, H., Kim, J. Z., & Bassett, D. S. (2019). Network topology of neural systems supporting avalanche dynamics predicts stimulus propagation and recovery. *arXiv*, 1812.09361.
- Kinouchi, O., & Copelli, M. (2006). Optimal dynamical range of excitable networks at criticality. *Nature Physics*, *2*, 348 EP. doi: 10.1038/nphys289
- Ko, H., Hofer, S. B., Pichler, B., Buchanan, K. A., Sjöström, P. J., & Mrcic-Flogel, T. D. (2011). Functional specificity of local synaptic connections in neocortical networks. *Nature*, *473*, 87-91. doi: 10.1038/nature09880
- Larremore, D. B., Shew, W. L., & Restrepo, J. G. (2011). Predicting criticality and dynamic range in complex networks: Effects of topology. *Phys. Rev. Lett.*, *106*, 058101. doi: 10.1103/PhysRevLett.106.058101
- Neumaier, A., & Schneider, T. (2001). Estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Trans. Math. Softw.*, *27*(1), 27-57. doi: 10.1145/382043.382304
- Pasqualetti, F., Zampieri, S., & Bullo, F. (2014). Controllability metrics, limitations and algorithms for complex networks. *IEEE Transactions on Control of Network Systems*, *1*(1), 40-52.
- Perin, R., Berger, T. K., & Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *Proceedings of the National Academy of Sciences*, *108*(13), 5419–5424. doi: 10.1073/pnas.1016051108
- Priesemann, V., Wibral, M., Valderrama, M., Pröpper, R., Le Van Quyen, M., Geisel, T., ... Munk, M. H. J. (2014). Spike avalanches in vivo suggest a driven, slightly subcritical brain state. *Frontiers in Systems Neuroscience*, *8*, 108. doi: 10.3389/fnsys.2014.00108
- Schneider, T., & Neumaier, A. (2001). Algorithm 808: Arfit—a matlab package for the estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Trans. Math. Softw.*, *27*(1), 58–65. doi: 10.1145/382043.382316
- Touboul, J., & Destexhe, A. (2017). Power-law statistics and universal scaling in the absence of criticality. *Phys. Rev. E*, *95*, 012413. doi: 10.1103/PhysRevE.95.012413
- Wang, Y., Markram, H., Goodman, P. H., Berger, T. K., Ma, J., & Goldman-Rakic, P. S. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nature Neuroscience*, *9*, 534-42. doi: 10.1038/nn1670
- Yan, G., Vértes, P. E., Towson, E. K., Chew, Y. L., Walker, D. S., Schafer, W. R., & Barabási, A.-L. (2017, 10 18). Network control principles predict neuron function in the *Caenorhabditis elegans* connectome. *Nature*, *550*, 519-523. doi: 10.1038/nature24056