Slow Dynamics and High Variability in Clustered Networks

Reproduction of Litwin-Kumar and Doiron (2012)

Arjun Puri, Johan Kirikal, and Shuo Wang
Berlin Center for Computational Neuroscience

February 11, 2026

Abstract. *This report reproduces the key findings from Litwin-Kumar and Doiron (2012) on the slow dynamics and variability in clustered spiking networks. We implement their clustered network model and show that potentiated within-cluster connectivity introduces metastable firing states within the clusters. Using analyses from the original paper, we identify slow firing-rate fluctuations (via long-timescale decays in spike-count autocorrelation), higher trial-to-trial spike-count variability (via super-Poisson Fano Factors), and weak average pairwise correlations with a slightly elevated tail for within-cluster neurons. We also reproduce stimulus-quenched spiking variability. These results support the original paper's finding that clustered excitatory connections are sufficient to generate slow firing-rate fluctuations and stimulus-quenched variability in LIF networks—features also seen in cortical recordings.*

1. Introduction

Cortical networks are often modeled as unstructured, randomly connected networks with balanced excitation and inhibition [1, 4]. This balance can place the network in an irregular, asynchronous regime in which neurons exhibit high trial-to-trial variability in spike timing. In this regime, excitation and inhibition are large but cancel in the mean. Thus, spikes are often fluctuation-driven, resulting in high spike-time variability.

However, cortical networks also exhibit slower, population-level changes in firing rate over longer timescales (hundreds of milliseconds to seconds). Baseline, unstructured, balanced excitation-inhibition networks do not typically capture these slower fluctuations in firing rate.

In this study, we find that introducing clustered excitatory connectivity creates slow dynamics. Clusters transition between states of high and low firing, such that neurons exhibit both high spike-time variability and slow firing-rate fluctuations. Introducing stimuli can bias the network toward particular activity states, reducing variability.

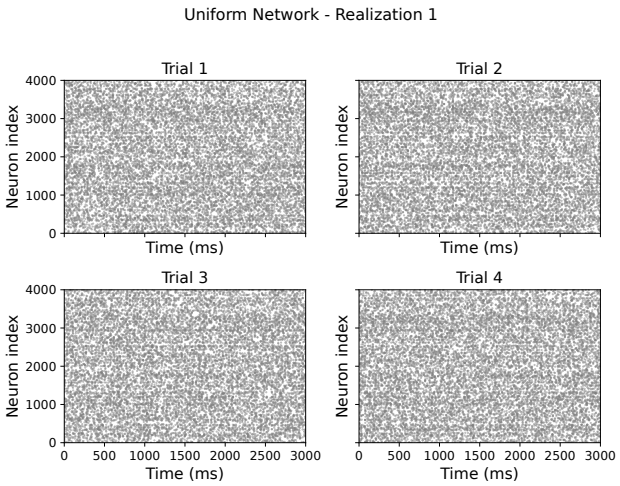

(a) Uniform (unclustered) network.

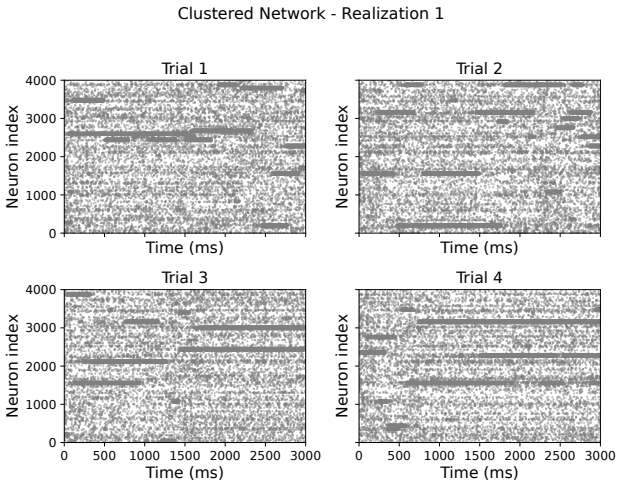

(b) Clustered network.

Figure 1. Spike rasters across multiple trials.

2. Results

We present our reproduction of the key findings from the original study.

2.1. Metastable Dynamics in Clustered Networks

An excitatory-inhibitory (E/I) network with clustered E→E connectivity produces metastable "high" firing activity states within clusters that are stable yet quick to transition. These high activity states persist for hundreds of milliseconds to seconds before switching, leading to slow rate dynamics and elevated trial-to-trial spike count variability (super-Poisson) compared to unclustered networks.

2.2. Mean Rate

Despite the switching between high and low firing states, the clustered network exhibits a low mean fir-

ing rate similar to that of the unstructured network (Figure 2). This is likely due to additional suppression within the network during high firing activity states. The high activity epochs in clustered networks do not alter the overall mean firing rate; however, the firing rate distribution broadens somewhat, likely reflecting the alternation between high and low activity states.

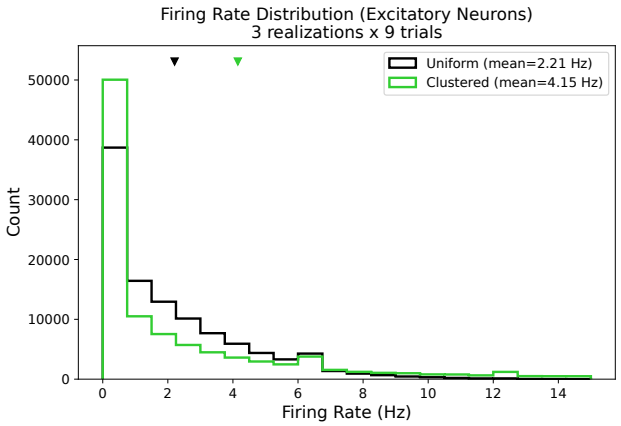

Figure 2. Distribution of firing rates for excitatory neurons in uniform and clustered networks. Both networks maintain low mean firing rates, though the clustered network shows a broader distribution reflecting the alternation between high and low activity states. Triangles indicate mean values.

2.3. Variability

Neurons in the clustered network exhibit Fano factors greater than one ("super-Poisson" spike count variability), as shown in Figure 3. Furthermore, the Fano factor increases with the size of the counting window (Figure 4). This pattern illustrates the presence of slow firing rate variability in clustered networks—larger windows capture more of the slow rate fluctuations.

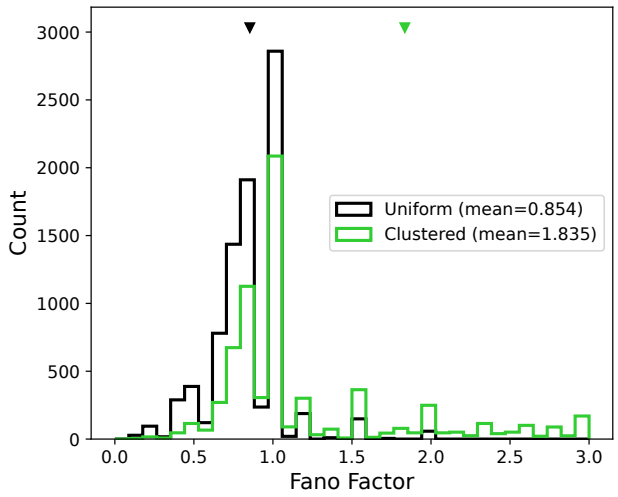

Figure 3. Distribution of Fano factors across neurons for uniform and clustered networks. The clustered network exhibits a broader distribution with a higher mean (triangles), indicating greater spike count variability.

2.4. Correlations

Pairwise spike count correlations remain near zero on average for both network types (Figure 5). However, in the clustered network, within-cluster neuron

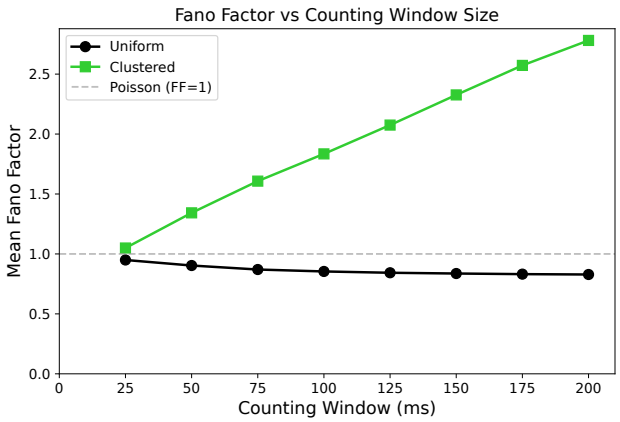

Figure 4. Mean Fano factor as a function of counting window size. In the clustered network, the Fano factor increases with window size, reflecting slow rate fluctuations. The uniform network remains near the Poisson baseline (FF = 1).

pairs show a heavier positive tail in the distribution of correlation coefficients (Figure 6). This observation is consistent with the notion that clusters undergo additional correlated firing rate transitions—switching between high and low activity states—when examined across temporal windows spanning the simulation. Furthermore, this positive tail becomes increasingly pronounced as the counting window size increases (50 ms vs. 100 ms), as expected given that within-cluster neurons share these slow rate fluctuations.

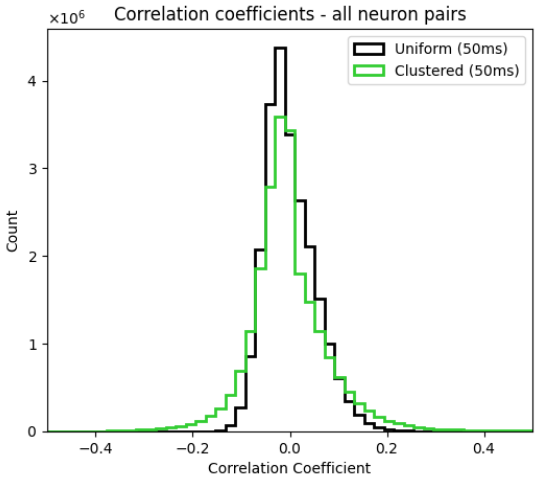

Figure 5. Distribution of pairwise spike count correlations for all neuron pairs. Correlations remain centered near zero on average for both network types.

2.5. Timescales

Clustered networks exhibit substantially longer correlation times than unstructured networks. The autocovariance function decays much more slowly as a function of lag in clustered networks compared to their unstructured counterparts. Similarly, the cross-covariance for within-cluster neuron pairs decays slowly, indicating a shared correlated rate change among neurons belonging to the same cluster (Figures 7 and 8).

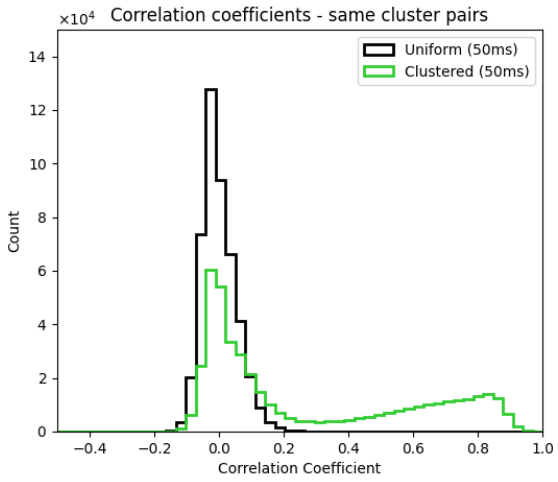

Figure 6. Distribution of pairwise spike count correlations for within-cluster pairs only. The clustered network shows a heavy positive tail that increases with counting window size (50 ms vs. 100 ms).

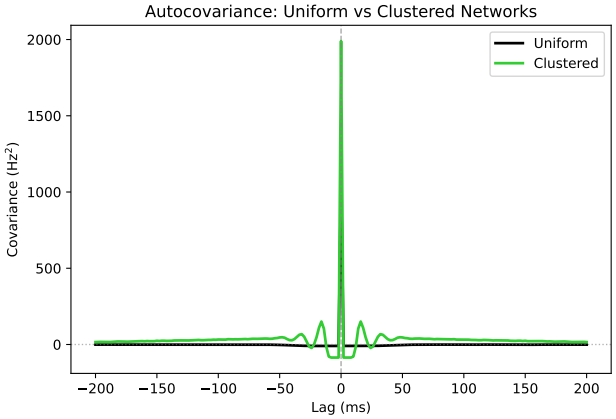

Figure 7. Autocovariance as a function of lag for uniform and clustered networks. The clustered network shows slower decay, indicating longer intrinsic timescales.

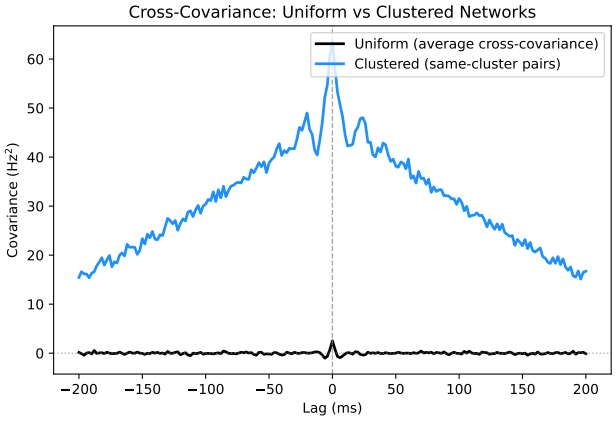

Figure 8. Cross-covariance as a function of lag for uniform and clustered networks. Within-cluster pairs in the clustered network show slower decay, indicating stronger shared rate fluctuations.

2.6. Stimulus-Dependent Variability

To probe how an external stimulus affects variability, we increased the input stimulus for one cluster and tracked the Fano factor over time (Figure 9). The clustered network shows a stronger time-dependent change in variability than the uniform network. During the stimulus, the clustered network shows a dip in Fano Factor compared to the uniform network. During that time, the Fano factor reaches approximately the same level as the uniform network. This is expected because stimulating one cluster inhibits all of the other clusters, removing overall firing rate variability. This finding is consistent with the original paper's finding of stimulus-quenched spiking variability.

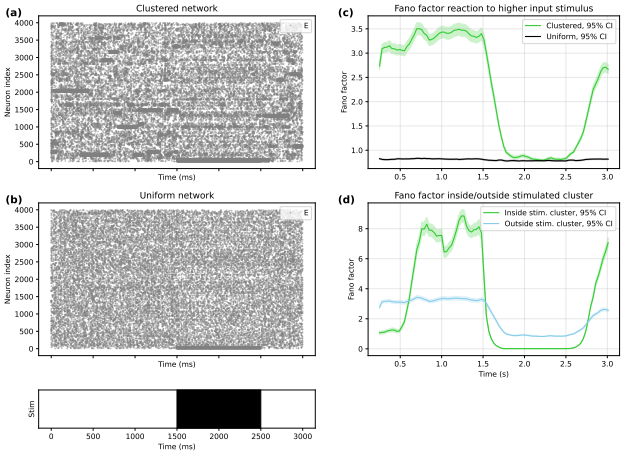

Figure 9. Effect of stimulus on network dynamics. **(a)** Clustered network activity with increased stimulus on the first cluster (first 80 neurons). **(b)** Increased stimulation on the first cluster, uniform network. **(c)** Fano factor as a function of time under input stimulus. The clustered network exhibits larger variability changes over time than the uniform network. **(d)** Fano factor only in the stimulated cluster, only outside the cluster.

2.7. Critical Cluster Strength

We experimented with networks of different connectivity strengths (R_{EE}) and found a critical value above which spike-count trial-to-trial variability (i.e., Fano factor) increases. This is consistent with the original paper's finding. We tracked the Fano factor within a fixed time window (on a similar timescale to the switching dynamics) across several trials (random initial conditions) for a given network realization at each connectivity strength. With increasing connectivity strength, a cluster of neurons remains in a high-activity state for longer. As a result, across many trials, spike-count variability increases, because the analysis window for a cluster is increasingly dominated by a high-activity state (or not).

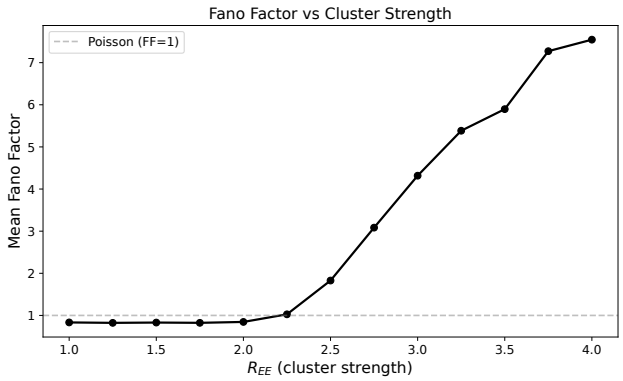

Figure 10. Mean Fano factor as a function of cluster strength (R_{EE}). As within-cluster connection strength increases, the network transitions from Poisson-like variability to super-Poisson variability. The dashed line indicates the Poisson baseline (FF = 1).

3. Discussion

The main goal of this project was to reproduce the paper by Litwin-Kumar and Doiron (2012) [2]. Through replicating their results, we got our first experience in building a computational experimental model while working as a team. We acquired key skills in neural network simulations using Brian2 and in converting the technical descriptions found in the original paper into actual code. In this process we also had to convert some key equations to a format more applicable to the Brian2 framework, since the original study [2] did not use Python.

To understand how successful our replication of the original paper was, Figure 11 and Figure 12 give us a comprehensive comparison between our work and the paper.

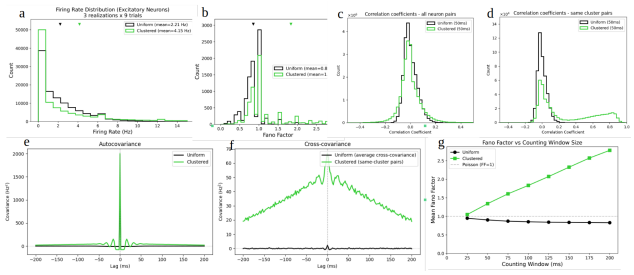

Figure 11. The graphs combined together from our statistical analysis.

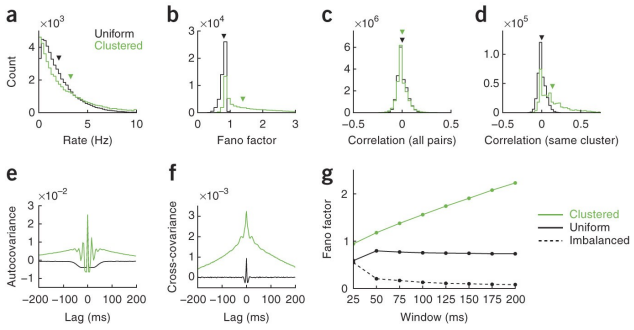

Figure 12. Statistical analysis results from the original paper by Litwin-Kumar and Doiron (2012) [2].

From this comparison, we see that most of the statistical analyses are very similar. The most fundamental difference is between Figure 11 **(d)** and Figure 12 **(d)**. From our experiment, the correlation coefficient figure displays a long heavy tail, with a significant bump around the 0.8 mark, while the original paper displayed a gradually diminishing positive tail, having almost no coefficients at 0.8.

Although it is hard to know exactly the firing dynamics of the network in the original paper, we can still guess the possible explanation for this difference. In our clustered network, neurons that entered the high activity state (high firing-rate), exhibited regular inter-spike intervals. These neurons were also firing very synchronously with the rest of their cluster (Figure 11). As the correlation coefficient calculation used window sizes of 50ms in our analysis, it can be expected that a cluster that is currently in the active state produces a very high correlation.

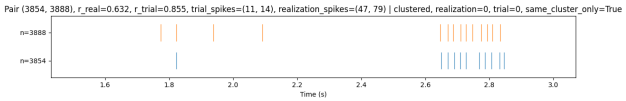

Figure 13. Spiking of two neurons in the same cluster. These neurons were sorted by their high correlation coefficient. We can see that both of the neurons are actively firing around 2.6 - 2.9 second interval, having highly synchronous spikes.

Therefore, the difference in our results and the original could lie in differences in network dynamics. It is possible that clusters in the original paper showed less synchrony when in the active (high firing-rate) state.

Another important difference is the fact, that we had to convert the equations from the original paper to a different form. While our neurons behaved according to Equations (1), (2), and (3), the synaptic current for neuron i in population x in the original paper was calculated as:

$$I^x_{i,\mathrm{syn}}(t) = \sum_{j,y} J^{xy}_{ij} F^y * s^y_j(t)$$

where the index j denotes the presynaptic neuron, y is the presynaptic neuron group, J is the synaptic strength, F is a synaptic filter and s_j is the presynaptic spike train. The synaptic filter $F(t)$ was defined as:

$$F^y(t) = \frac{1}{\tau_2 - \tau_1}\left(e^{-t/\tau_2} - e^{-t/\tau_1}\right)$$

Otherwise, our results coincided with the results from the original paper quite well. With our experiments, we confirmed the phenomenon that clusters can influence the default asynchronous network dynamics by introducing new attractor states and slow dynamics.

4. Methods

Our aim in this project was to study how clusters in an artificial neural network model influence its dynamics. For clarity, we divide this goal into three main parts: model architecture, experimental setup, and statistical analysis.

The code was written using Python 3.13[1] and the artificial neural networks were implemented using Brian2 [3][2].

4.1. Model

A good model architecture provides a framework, which should be biologically plausible, but at the same time simple enough to have as few tunable parameters as possible, computationally inexpensive and easy to grasp. This framework should provide us with a "sandbox" environment, facilitating an easy way to design and test hypotheses.

The uniform model was built with 4000 excitatory neurons and 1000 inhibitory neurons. The connection probability between neurons was $p^{EE} = 0.2$ for excitatory and $p^{EI}, p^{IE}, p^{II} = 0.5$ for all connections involving an inhibitory neuron. The strengths of the connections were $J^{EE} = 0.024$, $J^{IE} = 0.014$, $J^{EI} = -0.045$ and $J^{II} = -0.057$ (J^{YX} denotes a connection from neuron group X to neuron group Y.)

Individual neuronal dynamics were set with the following equations:

$$\dot{V} = \frac{\mu - V}{\tau_m} + I_{\text{syn}} \tag{1}$$

$$\dot{I}_{\text{syn}} = \frac{x - I_{\text{syn}}}{\tau_2} \tag{2}$$

$$\dot{x} = -\frac{x}{\tau_1} \tag{3}$$

If neuron **a** and **b** have a connection $\mathbf{a} \to \mathbf{b}$, then in the event of neuron **a** firing, $x_b = x_b + \frac{J^{BA}}{\tau_1}$.

The model used non-dimensional values in the simulations. The parameters were set as follows: $\tau_m =$ 15 ms for excitatory or 10 ms for inhibitory neurons. $\tau_1 = 1$ ms for both neuron types. $\tau_2 = 3$ ms for excitatory or 2 ms for inhibitory neurons. Spike threshold $V_{th} = 1$. Reset potential $V_{re} = 0$. Refractory period $r = 5$ ms. External input μ was sampled from a uniform distribution in the range $[1.1, 1.2]$ for excitatory and $[1, 1.05]$ for inhibitory neurons. The initial V was sampled from a uniform distribution in the range $[V_{re}, V_{th}]$.

Dimensionality was introduced only when plotting the figures – the parameters were multiplied by 15 mV and transformed so that $V_{th} = -50$ mV and $V_{re} = -65$ mV.

Integration was performed using the Euler method, with timestep $dt = 0.1$ ms.

When constructing the clustered network, clusters were formed only between excitatory neurons. In a network with 4000 excitatory neurons, every cluster had 80 members. Clustering introduced two main changes to the network architecture: (1) the within-cluster synaptic weights were scaled by a factor of 1.9, i.e., $J^{EE}_{\text{in}} = 1.9 \times J^{EE}$; (2) neurons had a higher probability of forming a synapse with a same-cluster counterpart.

The probability for forming excitatory connections inside clusters was set to p^{EE}_{in}, and to neurons in other clusters to p^{EE}_{out}. The probabilities were chosen such that the total average connection probability remained $p^{EE} = 0.2$. The ratio

$$R^{EE} = \frac{p^{EE}_{in}}{p^{EE}_{out}} \tag{4}$$

determines the level of clustering in the network. In our experiments, we used $R^{EE} = 2.5$ in the clustered networks.

4.2. Analysis

For understanding what kind of dynamics clustering introduces to neural networks, we ran parallel simulations on a uniformly connected network and a clustered network. The network was run for 3 s for all of our experiments. After simulations, the network behavior was analyzed using five key metrics: firing rate, Fano factor, correlation coefficients, autocorrelation, and cross-correlation. For the analysis, we discarded the first 1.5 s of the simulation and worked with the rest.

Firing rate for neuron i was calculated as

$$r_i = \frac{N_i(t, t + \Delta t)}{\Delta t} \tag{5}$$

where $N_i(t, t + \Delta t)$ is the spike count of neuron **i** in a time interval $[t, t + \Delta t]$.

The Fano factor is a metric that captures spike variability. A Poisson process with a fixed underlying firing rate has a Fano factor of 1. Higher Fano factor values indicate a change in the neuron's underlying firing rate (for example, a transition from a low to a high firing rate state). The Fano factor for neuron **i** was calculated as

$$F_i = \frac{Var(N_i(t, t + \Delta t))}{\langle N_i(t, t + \Delta t) \rangle}$$

The variance and mean of a neuron were calculated over 9 trials, using $\Delta t = 100$ ms and $t_{start} = 1.5$ s.

The correlation coefficients were calculated between neuron pairs. This metric tells us whether the spiking of two neurons is correlated or not. The correlation coefficient between neurons **i** and **j** was calculated as:

$$\rho_{ij} = \frac{Cov(N_i(t, t + \Delta t), N_j(t, t + \Delta t))}{\sqrt{Var(N_i(t, t + \Delta t)) Var(N_j(t, t + \Delta t))}}$$

Covariance and variance was calculated on overlapping windows within each trial and then averaged across the trials. In our experiments, we used $\Delta t =$ 50 ms and $\Delta t = 100$ ms with a window step of 25 ms.

For autocovariance and cross-covariance analysis, spikes were binned in 2 ms bins over the same post-transient analysis window (1.5 s to 3.0 s). We computed centered spike-count rates and evaluated covariance as a function of lag. Autocovariance was averaged across neurons, while cross-covariance was averaged across sampled neuron pairs (all pairs for the uniform

[1] https://www.python.org
[2] https://brian2.readthedocs.io

network and within-cluster pairs for the clustered network), then averaged across trials and realizations.

To evaluate the effect of cluster strength, we varied R_{EE} and computed the corresponding mean Fano factor for each value from repeated simulations. The same Fano-factor pipeline was used, and results were summarized by averaging across excitatory neurons and trials to produce the Fano-factor-versus-R_{EE} curve.

Finally, we investigated how changing the input stimulus μ influences the network dynamics. For this, we raised the stimulus by a multiplier of 2x for the first cluster (80 neurons). The effect of elevated stimulus on the network was analyzed by calculating the Fano factor over 9 trials as before, but it was calculated for overlapping windows of 250 ms, with a shift of 30 ms, resulting in a time-dependent function for the Fano factor.

References

[1] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3): 183–208, 2000.

[2] Ashok Litwin-Kumar and Brent Doiron. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*, 15(11):1498–1505, 2012. doi: 10.1038/nn.3220.

[3] Marcel Stimberg, Romain Brette, and Dan FM Goodman. Brian 2, an intuitive and efficient neural simulator. *eLife*, 8:e47314, 2019.

[4] Carl van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, 1996.