

# Metastable Dynamics and Learning in Balanced Cortical Networks

Praveen Venkatesh, Majid Mahzoon

## 1 Balanced Cortical Networks

Consider a balanced cortical network: a model of a bunch of neurons, e.g., 4000 excitatory neurons and 1000 inhibitory neurons, which are wired up without any structure, i.e., like an Erdos-Renyi network (with connection probability of 0.2). These neurons are in a balanced state, i.e., there is a lot of recurrence in the network and a lot of soft-coupling between neurons; there is lots of excitation and lots of inhibition, but both of these things cancel each other out because they are of opposite polarities. As a result, even though there is lots of excitation and lots of inhibition, the system has a zero-mean input and it is really driven by fluctuations. Thus, the spike trains happen very radically in time and this is all self-generated in this network. There is no true randomness in this network. From a dynamical systems point of view, this situation is effectively chaos and this is a high-dimensional chaotic system. Chaos at these system-size levels is not distinguishable from noise.

“We studied a simple network of  $N_E$  excitatory and  $N_I$ , inhibitory neurons. An additional  $N_O$  neurons from outside the network provide external excitatory inputs to the network neurons. On average,  $K$  excitatory,  $K$  inhibitory, and  $K$  external neurons project to each neuron in the network. An important aspect of the architecture is the random, sparse connectivity. Although the average number of projections,  $K$ , is large, it is still much smaller than the total number of neurons in the subpopulations  $N_E$  and  $N_I$ . Also important is the assumption that the individual connections are relatively strong, namely, that only  $\sqrt{K}$  excitatory inputs are needed to cross the firing threshold. Because the threshold is only of the order of  $\sqrt{K}$  synaptic inputs, the total synaptic input to a cell will be overwhelmingly depolarizing or hyperpolarizing unless the activity of the excitatory and inhibitory populations dynamically adjust themselves so that the large total inhibitory input nearly cancels the large excitatory one. This balance between excitation and inhibition generates a net synaptic input whose mean and fluctuations are both on the order of the threshold. The precise timing of the crossing of the threshold is determined by the fluctuations of the synaptic input, yielding a strongly irregular pattern of activity. This disorder is a signature of deterministic chaos. Its presence does not require external sources of noise. The resulting state is characterized by strongly chaotic dynamics, even when the external inputs to the network are constant in time. Such a network exhibits a linear response, despite the highly nonlinear dynamics of single

neurons, and reacts to changing external stimuli on time scales much smaller than the integration time constant of a single neuron.” [1]

## 1.1 Spike Train Statistics

Why do people study these balanced networks? Balanced networks have been popular to study because they produce variability that is consistent with cortex. As said above, this network has chaotic internal dynamics. We can form the raster plot of the network: assuming that the  $x$ -axis represents time and the  $y$ -axis represents neuron index, if we look at the spike train from any one of the neurons, any time the neuron crosses a threshold, i.e., fires, we put a little dot in the appropriate coordinates and that is a spike. We can immediately see kind of salt-and-pepper behavior. The individual neurons are quite variable and they are relatively asynchronous from one another. If we look at the distribution of the firing rates in the network, we see that it has a heavy-tailed distribution: most neurons do not fire very much and some neurons fire a little bit more. If we look at their Fano factor which is the ratio of spike count variance to the spike count mean, it is quite high. A Poisson process would have a Fano factor of one and in our model, it is a little less than one.

“The balanced state is characterized also by a broad distribution of (time-averaged) firing rates across the network. This distribution is the outcome of the different synaptic projections on each neuron as well as their different threshold levels. A prominent feature is the skewness of the distribution at low mean rates. Because in our model the temporal fluctuations are generated by the network activity itself, they are relatively small when the mean network activity is low. Hence, only those neurons that have relatively low thresholds fire vigorously; the rest show activity levels much lower than the mean rate. This feature agrees well with experimentally obtained histograms of firing rates of neurons in the monkey prefrontal cortex.” [1] See Figure 1.

*The internally-generated variability produces these kinds of behaviors.*




Figure 1: (Left) The distribution of the rates of neurons in the excitatory population for different population-averaged rates. The rate distribution is shown in terms of the local rates divided by the mean rate. (Right) The distribution of firing rates of neurons in the right prefrontal cortex of a monkey attending to a variety of stimuli (light source and sound) and executing simple reaching movements. The rates were averaged over the duration of events (stimuli or movements) that showed a significant response. The average rate was 15.8 Hz. Most cells fire at a lower rate, whereas a small fraction of the cells fire at much higher rates. (Figure 3 in [1])

### 1.1.1 Firing Rate Distribution

If we look at cortex, we see that the lognormal distribution of firing rates matches what people see in an unanesthetized auditory cortex. This is work from Anthony Zador’s lab where they have measured spontaneous activity and have seen that “the distribution of spontaneous firing rates across the population was remarkably well fit with a lognormal distribution - that is, the logarithm of the firing rates was well fit with a Gaussian distribution. Because the lognormal distribution has a *heavy tail*, most spikes were generated by just a few neurons.” [2] See Figure 2.

*Balanced networks produce this.*




Figure 2: The distribution of firing rates follows a lognormal distribution: Firing rates of most neurons were low and followed a lognormal distribution. (Left) Frequency histogram of non-zero spontaneous firing rates. The filled arrow shows the position of the median spontaneous firing rate, and the open arrow shows the position of the mean spontaneous firing rate. (Right) The distribution of spontaneous firing rates (dots) was fit with a lognormal distribution (gray line), the mean and variance of which were given by the mean and variance of the original firing rate distribution. The lognormal distribution appears as a normal distribution on a (semi-) logarithmic scale. The error bars show 95% confidence intervals determined by bootstrapping. (Figure 3B,C in [2])

### 1.1.2 Spike Train Variability

We can also ask how variable our network is. “We have measured the responses of neurons in area MT of the alert monkey while we varied the strength and direction of the motion signal in such displays [i.e., dynamic random-dot stimuli]” [3]. This is work done in Anthony Movshon’s lab where they have looked at the relationship between the mean spike count and the variance of the spike count and have seen that they are kind of huddled around the unit line: as the mean increases, the variance increases proportionally or more precisely, equal to it which will yield a Fano factor of one. “We also explored the relationship between response magnitude and response variance for these cells and found, in general agreement with other investigators, that this relationship conforms to a power law [ $Variance = k(mean)^b$ ] with an exponent slightly greater than 1. We used a maximum-likelihood technique that takes account of the variability of the estimates of both mean and variance to fit this relationship to our data.” [3] See Figure 3.

*Balanced networks produce this as well.*




Figure 3: Relationship between mean response level and response variance. The dashed line represents the average slope of the relationship estimated using maximum-likelihood fitting to individual cells' data. Therefore, balanced networks without too many assumptions, just that excitation and inhibition are large but balanced so that they cancel (meaning that fluctuations are the only things left), give us responses that are like cortex. (Figure 4 in [3])

### 1.1.3 Asynchronous Dynamics

Recently, Alfonso Renart and Jaime de la Rocha showed that in these balanced networks we get asynchronous behavior. If we look at the pairwise correlation between any pair of neurons in our network, it looks roughly asynchronous in the raster plot and it indeed is asynchronous. Previous to this work, theorists used to think that the asynchrony came from the sparseness condition that the likelihood of any two neurons wire together was  $\epsilon$ ; so, the likelihood of joint connections or common input to a pair of neurons is  $\epsilon^2$  which is negligible. Alfonso and Jaime pushed this further and said that we can get this asynchronous activity even for heavily-coupled networks: their connection probability was 0.2 which is not negligible and they got asynchronous behavior.

“To investigate whether such decorrelation can arise spontaneously from the dynamics of a recurrent network, we analytically characterized the behavior of correlations in a simple recurrent circuit of binary neurons. The network consists of two populations (of size  $N$ ) of  $E$  and  $I$  neurons connected randomly, both receiving excitatory projections from an external ( $X$ ) population of  $N$  cells. The network has two key properties: first, the connectivity is dense so that the connection probability  $p$  is fixed independently of the network size (e.g.,  $p = 0.2$ ). Second, the synaptic couplings are strong such that

only a small fraction of a cell’s excitatory inputs is enough to evoke firing; in the model, although the average number of inputs is proportional to  $N$ , the number of excitatory inputs needed to induce firing is only proportional to  $\sqrt{N}$ . Our analysis showed that, even in the presence of shared input, the network settles into a stationary state in which the population-averaged firing correlation is very weak, if inhibition is sufficiently strong and fast. In fact, in networks of different sizes, correlation decreases inversely proportional to  $N$ , a signature of asynchronous networks.” [4]

In another work from Andreas Tolias’s group, they were looking at the very careful recordings of the visual system of primates. They actually made a big deal about the spike sorting algorithms to really make sure that they were looking at the multiunit activity. Noise correlations in the visual systems of primates seem to be asynchronous.

“Correlated response fluctuations among simultaneously recorded neurons have been observed in a number of cortical areas (referred to as noise correlations). We developed chronically implanted multi-tetrode arrays offering unprecedented recording quality to re-examine this question in the primary visual cortex of awake macaques. We found that even nearby neurons with similar orientation tuning show virtually no correlated variability.” [5] See Figure 4.

*The balanced networks capture this asynchronous activity in cortex as well, even though there is a lot of wiring in the system.*




Figure 4: Distribution of noise correlation. (Figure 3B in [5])

**Therefore, balanced networks capture many key features of cortical dynamics, but there is still stuff they do not capture.**

## 1.2 Super-Poisson Variability on Long Time Scales

Anthony Movshon’s group in the 90’s showed that neurons were Poisson. They looked at the same data recently and they said they are not as Poisson as we thought: they are super-Poisson. What they did is that they looked at lateral geniculate nucleus (LGN), primary visual cortex (V1), secondary visual cortex (V2) and the area of the brain involved in sensing motions (MT) and they saw that when the number of spikes is small, things look relatively Poisson, but if we let the windows over which we observe the system be really large so that the average number of events in that window will grow large, we start to deviate from Poisson. As we look at tens or hundreds of spikes in a window, the variability is sometimes five, six or seven times that for a Poisson process. See Figure 5.

*Therefore, things are super-Poisson which is not captured by balanced networks.*




Figure 5: Comparison of neural response variability for cells in different visual areas. Variance-to-mean relation for 63 LGN cells (orange), 396 V1 cells (dark green), 189 V2 cells (blue) and 137 MT cells (violet). Each data point illustrates the mean and variance of the spike count in a 1,000-ms window of one cell for one stimulus condition. (Figure 2a in [6])

### 1.2.1 Spontaneous and Evoked Dynamics: Doubly Stochastic Process

This super-Poisson activity has also been prevalent in motor cortex where Mark Churchland did some nice work. He looked at variability before and during a task. Suppose that an animal’s task is to reach to a certain area (target) and it is going to do this over and over again. He looked at a single neuron spike train over many trials of this identical task where the target appears midway through the experiment.

“We measured the across-trial variability of neural responses in dorsal premotor cortex of three monkeys performing a delayed-reach task. Such variability was initially high, but declined after target onset, and was maintained at a rough plateau during the delay. An additional decline was observed after the go cue. Between target onset and movement onset, variability declined by an average of 34%. This decline in variability was observed even when mean firing rate changed little. We hypothesize that this effect is related to the progress of motor preparation. In this interpretation, firing rates are

initially variable across trials but are brought, over time, to their appropriate values, becoming consistent in the process.” [7]

Before the target appears, the average firing rate is quite low. When the target appears, since this neuron is responsible for encoding that muscle movement, the average rate goes up. But if we look at the estimation of the firing rate within a specific trial (gray lines) not the trial average, we see that during the spontaneous activity, there is lots of variability in terms of the rate. On some spontaneous trials, the firing rate is a little higher compared to the average and on the other trials, it is lower. But during the task, all these gray lines sort of collapse onto a single line. If we think of this from a stochastic process point of view, they are all Poisson processes, but in the spontaneous state, the firing rate of that Poisson process is itself random. That is why they call it a “doubly stochastic” process. Then, during the trial, that firing rate variability collapses and we will just have the classic Poisson. Thus, this super-Poisson behavior seems to be characterized in the spontaneous state quite heavily. Fano factor is quite high in the spontaneous state, i.e., the variance is higher than the mean, but then it drops during the task. See Figure 6.

Mark Churchland, Krishna Shenoy, and Byron Yu published a paper a few years back where they said that the stimulus- or motion-induced quenching a variability is a generic feature of cortex.

“We measured neural variability in 13 extracellularly recorded datasets and one intracellularly recorded dataset from seven areas spanning the four cortical lobes in monkeys and cats. In every case, stimulus onset caused a decline in neural variability. This occurred even when the stimulus produced little change in mean firing rate. The variability decline was observed in membrane potential recordings, in the spiking of individual neurons and in correlated spiking variability measured with implanted 96-electrode arrays. The variability decline was observed for all stimuli tested, regardless of whether the animal was awake, behaving or anaesthetized. This widespread variability decline suggests a rather general property of cortex, that its state is stabilized by an input.” [8]

In all these cases, the arrow denotes the beginning of a trial or a stimulus. The Fano factor tends to drop pretty reliably when the stimulus comes in or when the task starts. Therefore, this stimulus-induced quenching a variability due to this doubly stochastic process seems to be something that we think cortex has. See Figure 7.

*Balanced networks will not produce these kinds of behaviors.*

**If we put attractors into our system, we can capture these behaviors.**

## 2 Clustered excitatory networks

Balanced cortical networks by themselves show many important properties of biological neural networks, but as mentioned in the previous section, they do not capture all of their properties. Specifically, they do not capture super-Poisson firing rate behaviour and stimulus-dependent decrease in trial-to-trial variability. We wish to explore, therefore,

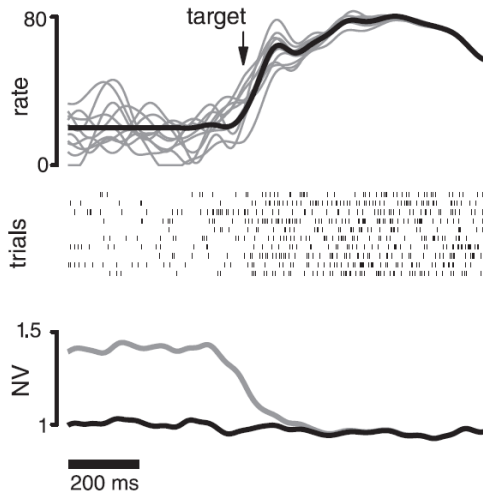


Figure 6: Simulations illustrating how an increasing consistency in across-trial firing rate could be detected using the normalized variance (NV) metric. Simulations were based on the mean firing rate of one recorded neuron (solid black trace at top). Baseline activity was artificially extended (to the left) to allow longer simulations. For each of 10,000 simulated trials, spike trains were generated using Poisson statistics. Two versions of the simulation were run. For the first version, the underlying firing rate was identical (black trace at top) on all simulated trials. The resulting NV is shown by the black trace at the bottom. For the second version, each trial had a different underlying firing rate, generated by adding noise, filtered with a 30 ms SD Gaussian, to the mean. The magnitude of this noise decayed with an exponential time constant of 200 ms after target onset. Ten examples of the resulting underlying firing rates are shown in gray at top, and the resulting spike trains (computed with Poisson statistics, with the time-varying mean taken from the gray traces) are shown in the rasters. The NV computed from 10,000 such spike trains is shown by the gray trace at the bottom. (Figure 2 in [7])

the possibility that introducing clustering into balanced cortical networks can reveal some of these signatures of attractor dynamics in spiking activity.

## 2.1 Introducing clustering

Clustering can be introduced into an Erdos-Renyi model by artificially breaking the network up into groups. Edges are once again assigned randomly and independently of each other, however, the probability of creating an edge within each neuron group ( $p_{in}^{EE}$ ) is different from the probability of creating an edge between two neurons in different groups ( $p_{out}^{EE}$ ). Setting  $p_{in}^{EE} > p_{out}^{EE}$  results in dense clusters with few edges going between them.

It should be noted that this clustering is only present for the excitatory neurons. There has been no evidence of inhibitory clustering. Fino and Yuste used a two-photon



photostimulation technique in order to explore the structure of cortical microcircuits in mouse frontal cortex and found that “most, and sometimes all, inhibitory neurons are locally connected to every sampled pyramidal cell”. Furthermore, “inhibitory innervation of neighboring pyramidal cells is similar, regardless of whether they are connected among themselves or not”, from which they conclude that inhibitory connectivity “can approach the theoretical limit of a completely connected synaptic matrix” [10].

## 2.2 Firing rate transitions

The introduction of clusters creates a new kind of spiking dynamic in the network, as depicted by the neural spike trains [9]. We observe that clusters of neurons tend to increase and decrease their overall firing together. That is, the clusters have spontaneous and *synchronous* changes in their firing rates, against a background average cluster-independent firing rate for all neurons. Due to the presence of denser clusters, neurons within a cluster provide positive feedback to each other as a result of which they have an increased propensity to fire with a higher rate synchronously. This leads to the kind of raster plot depicted in figure 8.

### 2.2.1 Spike train statistics

The clustered network still gives a lognormal distribution of firing rates. The average Fano factor of the clustered network is found to be closer to 1.5, indicating super-Poisson behaviour. The autocovariance of single-neuron variability shows long time-scale dynamics that lasts 100s of milliseconds, despite the fact that the intrinsic time scales in the network (synapses, membrane potentials, etc.) are of the order of 10ms, indicating that the long time-scale dynamics is a feature of the architecture. We also see long time-scale correlations between members of the same cluster.

### 2.2.2 An energy-landscape explanation of synchronous in-cluster firing rate fluctuations

A simple explanation of why clustering causes the kind of dynamics we just saw is provided by an energy-landscape model, shown in figure 9. When the network consists of just a single cluster, it has a unique stable state, in which all neurons have the same average firing rate. But when we introduce clustering into the network, positive feedback between neurons within a cluster creates new stable states or “attractors” within this energy landscape. Chaotic fluctuations can now cause transitions from one stable state to another, resulting in two different average firing rates for a given cluster.

Clusters are not entirely independent of each other due to the absence of inhibitory clustering. So the excess in activation of one cluster will cause inhibition and hence reduced activity of other clusters.

## 2.3 Stimulus-dependent reduction in trial-to-trial variability

The natural question to ask next is whether or not this variability in firing rate is reduced as a result of stimulation. This problem was explored by Doiron and Litwin-Kumar [11], who found that stimulation causes a change in the energy landscape by breaking the symmetry between the low-activation and high-activation states of each cluster. In the absence of stimulation, clusters have an equal likelihood of being in either state (since their “potentials” are equal). High activation is not restricted to any particular cluster, and clusters “compete” for increased firing rate.

But on stimulation of one cluster, the potential of its high-activation state *reduces* and thus it is more likely to fire at a higher rate. Correspondingly, other clusters are likely to fire at a lower rate (due to the inhibitory effect of the stimulated cluster on the others), suggesting an *increase* in the “potential” of their high-activation states. This is shown in figure 10.

*Stimulation, therefore, breaks the symmetry between clusters and makes one of the clusters “more attractive” in the energy landscape than the others. This is essentially a decrease in the variability of the firing rates of clusters, and hence of neurons.*

### 2.3.1 Fano factor reduction

This reduction in variability is shown in simulation by a reduction in the Fano factor. The bias towards firing of the stimulated clusters is evident from the spike train raster plots. The reduction in Fano factor matches what is expected from investigations of actual neurons in motor cortex by Churchland and Yu [8].

*We have therefore shown that the single neuron signature of cortical activity (decrease in variability) supports an architectural hypothesis of clustered networks.*

## 3 Can stable attractors be learned?

Tsodyks et. al. [12] showed that spontaneous dynamics often in cortical networks often resemble or mimic dynamics produced on stimulation. Specifically, they examined the dynamics of a large area of cortical network of a cat, evoking certain patterns by repeatedly showing the cat an image. Later, they found that in the absence of input (eyes closed), the spontaneous patterns produced were similar to the evoked patterns. This is depicted in [figure]. Luczak et. al. [13] have more recently shown that in rat auditory and somatosensory cortex, “vectors [i.e., patterns] evoked by individual stimuli [occupied] subspaces of a larger but still constrained space outlined by the set of spontaneous events”. Han et. al. [14] have also shown, using Calcium waves as identifiers of cortical dynamical patterns, that “visually evoked cortical activity reverberates in subsequent spontaneous waves” in rat visual cortex.

These results indicate that stimulation creates an imprint of patterns in biological neural networks, glimpses of which are seen during subsequent spontaneous activity. This therefore begs the question: is it possible to “learn” such patterns in a balanced cortical network?

### 3.1 Introducing weight update for learning: Spike-timing-dependent plasticity

In order to examine this, we need to create a system of learning, wherein the synaptic weights between different neurons can change. So far, we have assumed that the network structure and synaptic weights themselves are static. Dynamical activity in the network does not cause any change in network connections. But if we want evoked patterns to get imprinted on the network, we need to create a mechanism for learning, by allowing synaptic weights to change.

To this end, we use a simple learning rule called spike-timing-dependent plasticity (STDP). If there is a connection from neuron A to neuron B, we increase the strength of the connection every time A fires just before B fires, and we decrease the strength of the connection every time A fires just after B fires. In other words, this scheme is causality-rewarding, and strengthens connections whenever it finds it likely that A played a role in making B fire. This is depicted in figure 13.

### 3.2 Unstable network activity and STDP of inhibition

Using STDP only for excitation does not work, however, as it leads to large positive feedbacks: rewarding causality strengthens connections, which makes it more likely that those connections will produce causal spike trains, which in turn will continue to increase their strength. The end result of this is that the network starts seizing - the whole network fires synchronously at a constant high rate. This is very pathological behaviour, and is not something observed in normal, healthy biological neural networks. This had been shown earlier by Morrison et. al. [16].

One way to overcome this problem is to introduce plasticity in the inhibitory neurons as well, as suggested by Vogels et. al. [17]. This ensures that excitation and inhibition keep each other in check, and prevent explosive positive feedback. The inhibitory weight-update is depicted in figure 14.

### 3.3 Learning metastable activity

With the introduction of STDP of inhibition, we can train the network to show spontaneous metastable activity. This is done by simultaneously and repeatedly stimulating groups of neurons that we want to make into clusters. The stimulation causes those neurons to fire more, which strengthens inter-neuron connections through STDP. Over multiple stimulation cycles, the neurons form clusters and start exhibiting the same spontaneous metastable dynamics that we saw earlier in clustered balanced cortical networks.

## References

- [1] C. van Vreeswijk and H. Sompolinsky, "Chaos in neuronal networks with balanced excitatory and inhibitory activity," *Science*, vol. 274, no. 5293, pp. 1724–1726, 1996.

- [2] T. Hromádka, M. R. DeWeese, and A. M. Zador, “Sparse representation of sounds in the unanesthetized auditory cortex,” *PLoS biology*, vol. 6, no. 1, p. e16, 2008.
- [3] K. H. Britten, M. N. Shadlen, W. T. Newsome, and J. A. Movshon, “Responses of neurons in macaque mt to stochastic motion signals,” *Visual neuroscience*, vol. 10, no. 06, pp. 1157–1169, 1993.
- [4] A. Renart, J. de la Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K. D. Harris, “The asynchronous state in cortical circuits,” *science*, vol. 327, no. 5965, pp. 587–590, 2010.
- [5] A. S. Ecker, P. Berens, G. A. Keliris, M. Bethge, N. K. Logothetis, and A. S. Tolias, “Decorrelated neuronal firing in cortical microcircuits,” *Science*, vol. 327, no. 5965, pp. 584–587, 2010.
- [6] R. L. Goris, J. A. Movshon, and E. P. Simoncelli, “Partitioning neuronal variability,” *Nature neuroscience*, 2014.
- [7] M. M. Churchland, M. Y. Byron, S. I. Ryu, G. Santhanam, and K. V. Shenoy, “Neural variability in premotor cortex provides a signature of motor preparation,” *The Journal of neuroscience*, vol. 26, no. 14, pp. 3697–3712, 2006.
- [8] M. M. Churchland, M. Y. Byron, J. P. Cunningham, L. P. Sugrue, M. R. Cohen, G. S. Corrado, W. T. Newsome, A. M. Clark, P. Hosseini, B. B. Scott *et al.*, “Stimulus onset quenches neural variability: a widespread cortical phenomenon,” *Nature neuroscience*, vol. 13, no. 3, pp. 369–378, 2010.
- [9] A. Litwin-Kumar and B. Doiron, “Slow dynamics and high variability in balanced cortical networks with clustered connections,” *Nature neuroscience*, vol. 15, pp. 1498–1505, 2012.
- [10] R. Y. Elodie Fino, “Dense inhibitory connectivity in neocortex,” *Neuron*, vol. 69, no. 6, pp. 1188–1203, 2011.
- [11] B. Doiron and A. Litwin-Kumar, “Balanced neural architecture and the idling brain,” *Frontiers in Computational Neuroscience*, vol. 8, no. 56, 2014.
- [12] M. Tsodyks, T. Kenet, A. Grinvald, and A. Arieli, “Linking spontaneous activity of single cortical neurons and the underlying functional architecture,” *Science*, vol. 286, no. 5446, pp. 1943–1946, 1999.
- [13] A. Luczak, P. Bartho, and K. D. Harris, “Spontaneous events outline the realm of possible sensory responses in neocortical populations,” *Neuron*, vol. 62, no. 3, pp. 413–425, 2009.
- [14] D. Y. Han F., Caporale N., “Reverberation of recent visual experience in spontaneous cortical waves,” *Neuron*, vol. 60, no. 2, pp. 321–327, 2008.

- [15] G. qiang Bi and M. ming Poo, “Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type,” *J. Neuroscience*, vol. 18, no. 24, pp. 10 464–10 472, 1998.
- [16] M. D. Abigail Morrison, Ad Aertsen, “Spike-timing-dependent plasticity in balanced random networks,” *Neural Computation*, vol. 19, pp. 1437–1467, 2007.
- [17] T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner, “Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks,” *Science*, vol. 334, pp. 1569–1573, 2011.

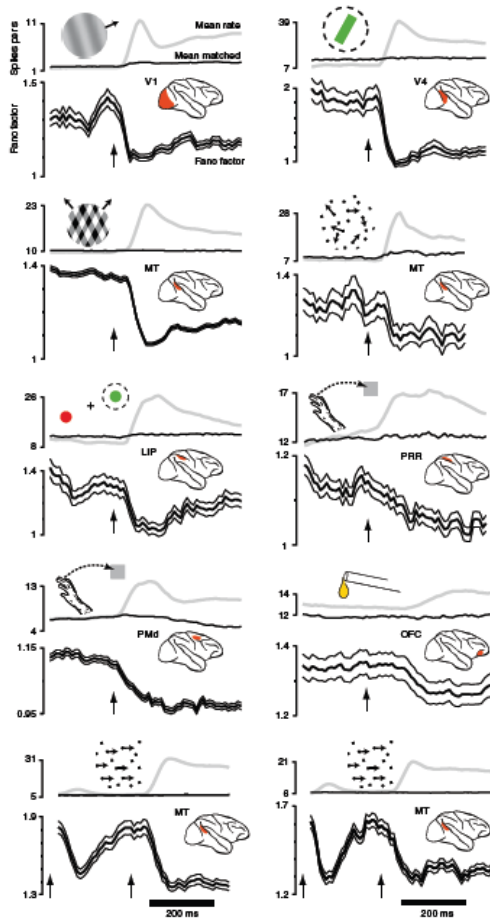


Figure 7: Changes in firing-rate variability for ten datasets (one per panel). Insets indicate stimulus type. Data are aligned on stimulus onset (arrow). For the two bottom panels (MT area/direction and MT speed), the dot pattern appeared at time zero (first arrow) and began moving at the second arrow. The mean rate (gray) and the Fano factor (black) were computed using a 50-ms sliding window. For OFC, where response amplitudes were small, a 100-ms window was used to gain statistical power. The resulting stabilized means are shown in black. The mean number of trials per condition was 100 (V1), 24 (V4), 15 (MT plaids), 88 (MT dots), 35 (LIP), 10 (PRR), 31 (PMd), 106 (OFC), 125 (MT direction and area) and 14 (MT speed). (Figure 1 in [8])

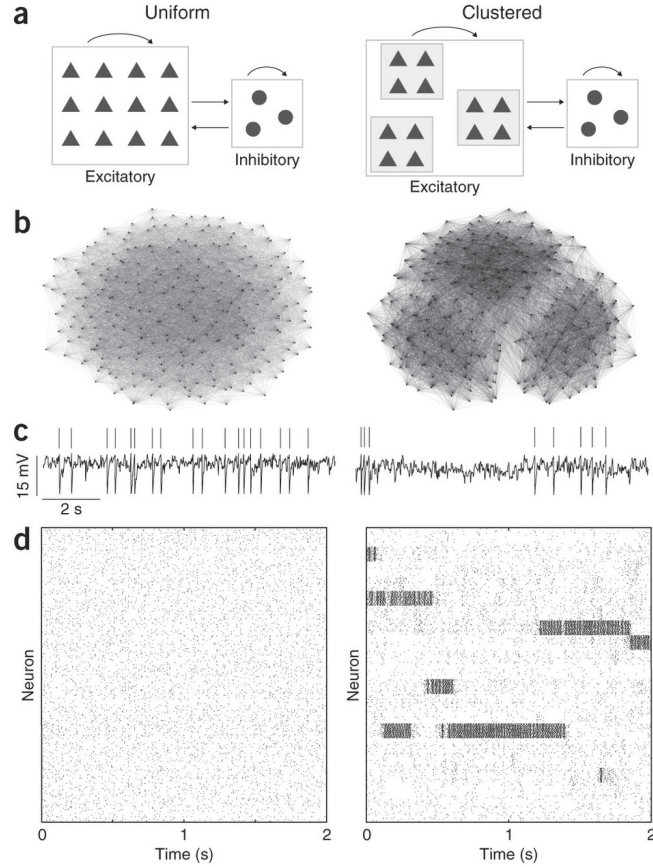


Figure 8: (a) Schematic of recurrent network, showing excitatory (triangles) and inhibitory (circles) neurons. Connections within and between populations were nonspecific for the uniform network. Shaded regions in the clustered network schematic indicate subpopulations with increased connection probability and strength for neurons belonging to the same cluster. The full network contained 4,000 excitatory neurons in 50 clusters of 80 neurons each and 1,000 inhibitory neurons. (b) Visualization of connectivity for a subpopulation of 240 excitatory neurons in three clusters. Nodes correspond to excitatory neurons; edges, synaptic connections. Nodes are positioned according to the Fruchterman-Reingold force algorithm (see Online Methods). Edge widths reflect synaptic strength. (c) Example voltage trace for an excitatory neuron. (d) Spike raster showing the spike times of a subpopulation of 1,600 excitatory neurons. Figure taken from Figure 1 in [9].




Figure 9: (a) Potential governing dynamics of average cluster activity  $\langle s_{in}^E \rangle$ , for different values of clustering  $R^{EE}$ . As clustering increases beyond a critical value, bistability emerges (black curve) and switches (arrow) can occur between the two stable states. (b) Potential (top) and corresponding histogram for average cluster activity (bottom). (c) Top: example activity raster for the parameters in b. Transitions are evident between the high activity (red) and low activity (blue) states. Bottom: histogram of input currents for neurons in the cluster conditioned on the cluster being in the low or high activity state (defined as average activity less than or greater than 0.5, respectively) and normalized in peak height. (d) Average timescale of transitions from high activity to low activity state as a function of well depth  $\Delta U$ . Each point corresponds to an average over four networks simulated for 200,000 time steps. The line corresponds to a linear regression of well depth against the logarithm of the transition time (coefficient of determination  $R^2 = 0.89$ ). Figure taken from Figure 5 in [9].

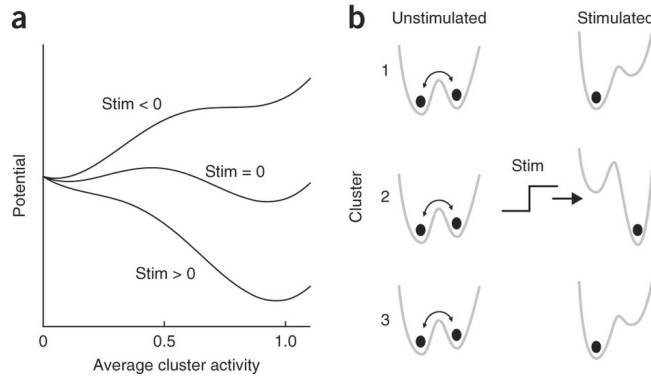


Figure 10: (a) Potentials for  $R^{EE} = 2$  and different values of stimulus (stim) applied to the cluster. A stimulus can bias activity toward the low or high activity state depending on its sign. (b) Consequences of stimulation and recurrent inhibition. Left: when all clusters are symmetric, any cluster can transition to the high activity state. Right: when cluster 2 receives a stimulus to bias it to the high activity state, other clusters are suppressed by recurrent inhibition. Hence the network remains in one particular activity configuration. Figure taken from Figure 6 in [9].



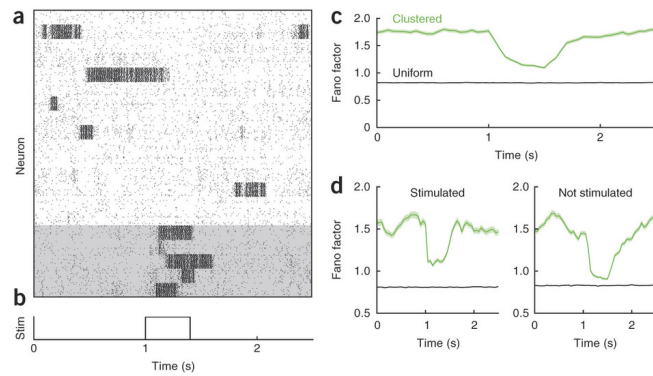


Figure 11: (a) Spike raster showing the activity of 1,600 excitatory neurons on one trial. Neurons in the shaded region received the stimulus when applied. (b) Time course of stimulus, a depolarizing step of current that lasted 400 ms. The stimulus caused the clusters that received it (a, shaded region) to transition into a high activity state. (c) Fano factor computed over 100-ms windows as a function of time for excitatory neurons. For comparison, the average Fano factor for uniform networks receiving an identical stimulus is shown. Only clustered networks exhibited a noticeable decrease in variability when stimulated. Shaded regions denote 95% confidence intervals. (d) Left: same as c, but restricted to neurons in stimulated clusters (corresponding to a, shaded region). Right: same as left, but restricted to neurons not in stimulated clusters. Figure taken from Figure 7 in [9].

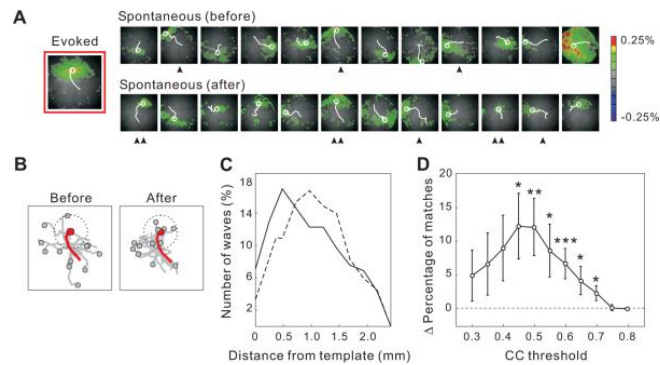


Figure 12: Effect of repeated visual stimulation at a given position on spontaneous waves. (A) Spontaneous waves immediately before and after training (10.24 s/session). Each image shows the initial frame of a wave. Initiation site and propagation path are indicated by a white circle and a line. Left, Evoked wave in response to the training stimulus. Spontaneous waves well matched to the template are indicated;  $\Delta$ ,  $CC > 0.6$ ;  $\Delta\Delta$ ,  $CC > 0.7$ . (B) Superposition of the initiation sites and propagation paths for all spontaneous waves before and after training (gray) and for the training-evoked wave (red). Dotted circle indicates a 0.6 mm radius from the initiation site of the training-evoked wave. (C) Distribution of the distance between the initiation sites of the spontaneous waves and the evoked template before (dashed) and after (solid) training (30 training experiments, 9 rats). (D) Difference in the percentage of matched waves before and after training plotted as a function of the  $CC$  threshold (\*,  $p < 0.05$ ; \*\*,  $p < 0.02$ ; \*\*\*,  $p < 0.01$ , same data as in C). Error bar,  $\pm$ SEM. Figure taken from Figure 2 in [14].




Figure 13: This graph shows the change in excitatory post-synaptic current amplitude, after repeated correlated spiking, plotted against spike timing. The same kind of function is used as a weight-update rule, where the y-axis now represents the percentage change in synaptic strength that is made, for a given spike-timing difference. Figure taken from figure 7 in [15].

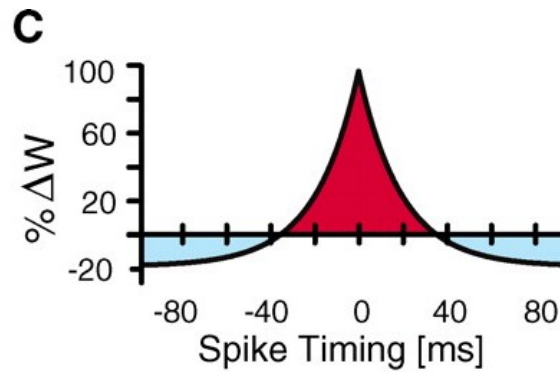


Figure 14: Spike-timing-dependent learning rule: Near-coincident pre- and post-synaptic spikes potentiate inhibitory synapses, whereas every presynaptic spike causes synaptic depression. Figure taken from figure 1 in [17].