



Published in final edited form as:

Phys Rev E Stat Nonlin Soft Matter Phys. 2014 December ; 90(0): 062710.

Dynamics of Random Neural Networks with Bistable Units

M. Stern^{1,4,*}, H. Sompolinsky^{1,3}, and L. F. Abbott^{4,5}

¹The Edmond and Lily Safra Center for Brain Sciences Hebrew University Jerusalem, Israel

²Racah Institute of Physics Hebrew University Jerusalem, Israel

³Center for Brain Science Harvard University Cambridge, MA 02138 USA

⁴Department of Neuroscience Columbia University College of Physicians and Surgeons New York NY 10032-2695 USA

⁵Department of Physiology and Cellular Biophysics Columbia University College of Physicians and Surgeons New York NY 10032-2695 USA

Abstract

We construct and analyze a rate-based neural network model in which self-interacting units represent clusters of neurons with strong local connectivity and random inter-unit connections reflect long-range interactions. When sufficiently strong, the self-interactions make the individual units bistable. Simulation results, mean-field calculations and stability analysis reveal the different dynamic regimes of this network and identify the locations in parameter space of its phase transitions. We identify an interesting dynamical regime exhibiting transient but long-lived chaotic activity that combines features of chaotic and multiple fixed-point attractors.

A substantial fraction of the synaptic input to a cortical neuron comes from nearby neurons within local circuits, while the remaining synapses carry signals from more distal locations. Local connectivity can have a strong effect on network activity [1]. In firing-rate models, a cluster of neurons with similar response properties is grouped together, and their collective activity is described by the output of a single unit [2]. Interactions between the neurons within a cluster are represented in these models by self-coupling, that is, feedback connections from a unit to itself, whereas interactions between clusters are represented by connections between units. Networks consisting of N units with connections chosen randomly and independently have provided a particularly fruitful area of study because they have interesting features and can be analyzed, in the large N limit, using mean-field methods [3]. Self-couplings in the networks that have been studied in this p way to date are either non-existent or weak (of order $1/\sqrt{N}$). If these units represent strongly interacting local clusters of neurons, we should include self-coupling of order 1 in the network model. For this reason, we consider the properties of firing-rate networks with strong self-interactions. The remaining interactions, those between units, are taken to be random in our study, reflecting the fact that we are investigating the properties of generic networks, not networks designed to perform specific tasks.

*merav.stern@mail.huji.ac.il.

The self-coupling we introduce to represent intracluster connectivity, together with the neuronal nonlinearity, can cause the individual units of the network to be bistable. The random inter-unit connectivity promotes chaotic activity, as has been previously established [3]. With both forms of connectivity, the networks we study combine two features normally seen independently, chaotic and multiple-fixed-point attractors. Our goal is to reveal the different types of activity that arise in networks with self-interacting units and to explore how chaotic and multiple-fixed-point dynamics interact. We begin by using network simulations to uncover the different dynamic regimes that the network exhibits, and then we use both static and dynamic mean-field methods to determine, in the limit of large network size, the properties of the activity within these regimes and to compute the phase boundaries between them in the space of network parameters.

The Model and Simulation Results

The networks we study consist of N units described by activation variables x_i , for $i=1, 2, \dots, N$, obeying the equations

$$\frac{dx_i}{dt} = -x_i + s \tanh(x_i) + g \sum_{j \neq i}^N J_{ij} \tanh(x_j). \quad (1)$$

The second term on the right side of this equation describes the within-cluster coupling, which has a strength determined by the parameter s . The last term on the right side reflects the random cluster-to-cluster interactions. The elements of the $N \times N$ connection matrix J are drawn independently from a Gaussian distribution with mean 0 and variance $1/N$, and the parameter g defines the strength of the inter-unit coupling, also known as the network gain. Note that the form of equation 1 implies that time is dimensionless or, equivalently, that it is measured in units of the network time constant.

The self coupling s and the network gain g determine the network dynamics. Before considering the full range of values for s and g , it is instructive to consider two special cases. The first is when the self coupling vanishes, $s=0$. In this case, previous work [3] has shown that, in the limit $N \rightarrow \infty$, the network exhibits activity when $g > 1$ chaotic and activity that decays to 0 when $g < 1$. The second special case is when the network gain vanishes, $g = 0$. In this situation, the units decouple, and each drives its own activity to a fixed point determined by $x = s \tanh(x)$. For $s < 1$, the only solution is $x = 0$, which is stable, and therefore all unit activity decays to zero from any initial state. For $s > 1$, there are two nonzero stable solutions (the zero solution is unstable) that are negatives of each other. Thus, in this case the units show bistability and, because they are independent, there are 2^N possible stable fixed-point configurations of the network. Nonzero values of both s and g can give rise to an interesting interplay between chaos and bistability.

As a preliminary indication of this richness, we investigate the network dynamics over a range of s and g values by computer simulation (figure 1). In the region below the long-dashed line in figure 1, any initial activity in the network decays to zero. Above the solid curve, the network exhibits transient irregular activity that eventually settles into one of a number of possible nonzero fixed points. This settling can take an extremely long time (as

we show below, exponentially long in N). In the region between these two curves, the network activity is persistently irregular. In the following, we will show that both the persistent irregular activity within this region and the transient irregular activity in the region with stable fixed points are chaotic, and we have labelled them as such in figure 1. The region shown with transient chaos and multiple fixed points is a distinctive feature due to the self interaction.

The Mean-Field Approach

To determine the type of activity that the network exhibits in different regions of the space of s and g values, we need to characterize solutions of equation 1 and evaluate their stability. For both of these computations, we take advantage of the random nature of the networks we consider. To analyze stability, we compute the eigenvalues of stability matrices for various solutions using results from the study of eigenvalue spectra of random matrices. To extract solutions of the network equations, we make use of mean-field methods that have been developed to analyze the properties of network models in the limit $N \rightarrow \infty$, averaged over the randomness of their connectivity [3]. In this section, we provide a brief introduction to the mean-field approach.

The basic idea of the mean-field method is to replace the network interaction term in equation 1 (the last term on the right side) by a Gaussian random variable [3] and to compute network properties averaged over realizations of the connectivity matrix J . Because we are averaging over realizations of J , all the units in the network are equivalent, so the N network equations 1 get replaced by the single stochastic differential equation

$$\frac{dx}{dt} = -x + s \tanh(x) + \eta(t). \quad (2)$$

We denote the solutions of this equation by $x(t; \eta)$ to indicate that they depend on the particular realization of the random variable η used in equation 2. If the mean and covariance of the Gaussian distribution that generates $\eta(t)$ are chosen properly and N is sufficiently large, the family of solutions $x(t; \eta)$ across the distribution of η will match the distribution of $x_i(t)$ across $i = 1, 2, \dots, N$ that solve equation 1, averaged over J . The consistency conditions that assure this require that the first and second moments of η match the first and second moments of the interaction term that it replaces. The first moment of η is 0. In the original network model, the average autocorrelation function of the interaction term, averaged over realizations of the random matrix J , is

$$\begin{aligned} C(\tau) &= \frac{1}{N} \sum_{i=1}^N \left\langle \sum_{j=1}^N J_{ij} \tanh(x_j(t)) \sum_{k=1}^N J_{ik} \tanh(x_k(t+\tau)) \right\rangle \\ &= \frac{1}{N} \sum_{j=1}^N \langle \tanh(x_j(t)) \tanh(x_j(t+\tau)) \rangle, \end{aligned} \quad (3)$$

where the square brackets denote an average of realizations of J , the angle brackets denote an average over t , and we have used the identity

$$[J_{ij}J_{kl}] = \frac{1}{N} \delta_{ik} \delta_{jl}. \quad (4)$$

The second moment of η is $g^2 C(\tau)$. As we will show below, $C(\tau)$ is calculated in the mean-field approach by averaging over $x(t; \eta)$ rather than over the different x_i in the network as in equation 3.

In the following sections, we use these results to obtain self-consistent mean-field results for both static and dynamic solutions, that depend on the family of solutions $x(t; \eta)$ of equation 2. When $s=0$, equation 2 is linear and can be solved analytically [3]. With nonzero s , equation 2 must be solved numerically. We describe procedures for doing this in the following sections, first for the simpler static case, when x and η do not depend on time, and then for the more complex dynamic case when they do.

Analysis of the Fixed Points of the Model

The phase plot in figure 1 has 3 regions separated by 2 phase boundaries. As we will see, fixed-point solutions exist in all three of these regions. Their stability defines the 2 phase boundaries. In this section, we identify the fixed points, estimate their number, and analyze their stability.

The Zero Fixed-Point The trivial solution $x_i = 0$ for all i always satisfies equation 1. To determine whether this solution will robustly appear we must compute its stability. The stability matrix for equation 1 expanded around the zero solution is

$$M_{ij} = (-1+s) \delta_{ij} + gJ_{ij}. \quad (5)$$

Because J is a random matrix with variance $1/N$, its eigenvalues, for large N , lie in a circle of unit radius in the complex plane [4–6]. For M , a factor of g scales this radius, and the diagonal terms shift the eigenvalues along the real axis by an amount $-1 + s$. To ensure that all eigenvalues of M have real parts less than zero, so that the zero fixed point is stable, we must therefore require $-1 + s + g < 0$. Thus, the long-dashed line in figure 1 is described by $s = 1 - g$.

Nonzero Fixed Points

In addition to the zero fixed point just discussed, the model exhibits non-zero fixed points. We now use the mean-field approach to find solutions corresponding to these non-zero fixed-points and to determine their stability as a function of g and s . Because we are searching for fixed points, the mean field, η , is a time-independent Gaussian random variable with zero mean and variance σ^2 to be determined. The solutions of the static version of equation 2,

$$x - s \tanh(x) = \eta, \quad (6)$$

are time-independent functions $x(\eta)$. For $s = 0$, equation 6 must be solved numerically, and figure 2 reveals that there are three possible solutions when $s > 1$ and $|\eta| < \eta_m$. These multiple solutions correspond to multiple fixed points in the original network model.

We are interested in determining the parameter range that supports stable non-zero fixed points and in computing their number. The stability matrix from equation 2 for a fixed point with values $x_i, i=1 \dots N$, is

$$M_{ij} = \delta_{ij} \left(-1 + s \left(1 - \tanh^2(x_j) \right) + g J_{ij} \left(1 - \tanh^2(x_j) \right) \right), \quad (7)$$

and stability requires that none of its eigenvalues have real parts greater than 0. We can evaluate stability using the mean-field solutions $x(\eta)$, rather than the networks values x_i that appear in equation 7. In the limit $N \rightarrow \infty$, the matrix 7 has an eigenvalue at the point z in the complex plane if [7]

$$\int_{-\infty}^{\infty} D_{\sigma} \eta \left(\frac{g \left(1 - \tanh^2(x(\eta)) \right)}{\left| z + 1 - s \left(1 - \tanh^2(x(\eta)) \right) \right|} \right)^2 > 1, \quad (8)$$

where we use the notation

$$D_{\sigma} \eta = \frac{d\eta}{\sqrt{2\pi\sigma}} \exp\left(-\frac{\eta^2}{2\sigma^2}\right). \quad (9)$$

If we ask whether there is an eigenvalue at the point $z=0$, this expression simplifies to

$$Q = \int_{-\infty}^{\infty} D_{\sigma} \eta \left(\frac{g}{\cosh^2(x(\eta)) - s} \right)^2 > 1. \quad (10)$$

We use this latter condition below because the inequality 8 does not support isolated eigenvalues [7], so stability can be assessed by determining whether or not there is an eigenvalue at $z=0$. Stability requires $Q > 1$, with the edge of stability defined by $Q=1$.

When $s > 1$, the expressions in equations 8 and 10 are ill-defined as written because $x(\eta)$ is a multivalued function for $|\eta| < \eta_m$. We can eliminate one of the 3 possible solutions of equation 6 in this range by noting that the denominator of the expression in 8 is equal to z plus the slope of the curve drawn in figure 2. Any value of $x(\eta)$ located on a region of this curve with negative slope will cause the denominator to vanish at a positive real value of z , indicated the presence of a positive real eigenvalue and instability. Thus, if we are interested in stable solutions, we can eliminate values of $x(\eta)$ located in the region of negative slope in figure 2, that is, we must require $|x(\eta)| > x_m$. This reduces the number of allowed solutions for $|\eta| < \eta_m$ from 3 to 2, one positive, which we call $x_+(\eta)$ and one negative, which we call $x_-(\eta)$. Because we are interested in evaluating Q of equation 10, we define, in the region where there are two solutions,

$$f_{\pm}(\eta) = \left(\frac{g}{\cosh^2(x_{\pm}(\eta)) - s} \right)^2. \quad (11)$$

In the region where there is no ambiguity, we just use $f(\eta)$ to denote this quantity. Note that, for positive η , $f_-(\eta) > f_+(\eta)$ or, equivalently, f is smaller for the solution with larger $|x(\eta)|$. This means that the solution with larger $|x(\eta)|$, for a given value of η in the range $|\eta| < \eta_m$, contributes less to Q and hence favors stability. On the other hand, restricting the solutions to the one with larger $|x(\eta)|$ could eliminate valid stable solutions.

The number of stable fixed points, when they exist, is exponential in N , which means that it is exponentially dominated by the configuration of $x(\eta)$ with the different solutions corresponding to stable fixed points. We define a weighting factor $m(\eta)$ to be the fraction of solutions that are chosen as $x_+(\eta)$ from the two possible values for $x(\eta)$ in this configuration. Then, $1 - m(\eta)$ is the fraction solutions chosen as $x_-(\eta)$. With this weighting specified,

$$Q = 2 \int_{\eta_m}^{\infty} D_{\sigma} \eta f(\eta) + 2 \int_0^{\eta_m} D_{\sigma} \eta (m(\eta) f_+(\eta) + (1 - m(\eta)) f_-(\eta)). \quad (12)$$

Note that we have used the $\eta \rightarrow -\eta$ symmetry of the system to express Q in terms of integrals only over the positive range of η .

The self-consistency condition that determines σ is also written in terms of the weighting factor $m(\eta)$ as

$$\sigma^2 = 2g^2 \int_{\eta_m}^{\infty} D_{\sigma} \eta \tanh^2(x(\eta)) + 2g^2 \int_0^{\eta_m} D_{\sigma} \eta (m(\eta) \tanh^2(x_+(\eta)) + (1 - m(\eta)) \tanh^2(x_-(\eta))). \quad (13)$$

Finally, the entropy, defined as $1/N$ times the average of the logarithm of the number of stable fixed points, is given by counting the number of combinations of x_+ and x_- solutions,

$$S = \lim_{N \rightarrow \infty} \frac{2}{N} \int_0^{\eta_m} D_{\sigma} \eta \ln \left(\frac{N!}{(m(\eta)N)!((1-m(\eta))N)!} \right) = -2 \int_0^{\eta_m} D_{\sigma} \eta (m(\eta) \ln(m(\eta)) + (1 - m(\eta)) \ln(1 - m(\eta))). \quad (14)$$

To complete the mean-field calculation, we need to determine the weighting function $m(\eta)$. We do this by imposing stability on the solutions being integrated in equation 15. The entropy is exponentially dominated by solutions at the edge of stability, so we constrain Q to the value 1, rather than imposing the inequality $Q \leq 1$. We then chose $m(\eta)$ to be the function that provides the maximum contribution to the entropy subject to the constraint $Q = 1$. Introducing the Lagrange multiplier λ to impose this constraint, we maximize $S + \lambda Q$ (using equations 15 and 12) with respect to $m(\eta)$, obtaining

$$m(\eta) = \frac{\exp(\lambda(f_+(\eta) - f_-(\eta)))}{1 + \exp(\lambda(f_+(\eta) - f_-(\eta)))}. \quad (15)$$

Substituting this expression into equation 15, we see that computing the entropy require the determination of two variables, σ^2 and λ . These are computed numerically by simultaneously solving equation 13 and the condition $Q=1$, using equations 12 and 15.

The transition curve between the nonzero-fixed-point and chaotic regions in figure 1 is the set of s and g values for which $Q = 1$ and the entropy of stable solutions vanishes (i.e. the number of stable solutions approaches 0). Figure 3A reveals three sets of values for which this occurs. Recall, that choosing the solution with larger $|x(\eta)|$ decreases Q , enhancing stability. At the transition, $\lambda \rightarrow -\infty$, and $m(\eta)$ is therefore a step function (figure 3C). This restricts all the solutions for $\eta > 0$ to be positive and all the solutions for $\eta < 0$ to be negative. The set of s and g values for which $S \rightarrow 0$ is the solid curve in (figure 3B), which is given approximately by $s_c(g) \approx 1 + 0.157 \ln(0.443g + 1)$. This is the transition line between the chaotic and transient-fixed-point regions in figure 1. Results for the entropy away from the transition line are indicated by colors in figure 3B.

Analysis of the Dynamics of the Model

We now examine the dynamics of the network model . We begin by studying solutions of the dynamic mean-field equations and using them to compute the average autocorrelation function of the network units. We then examine other properties of the network dynamics.

Autocorrelation

To study network dynamics, we return to the time-dependent mean-field equation, equation 2. The mean-field $\eta(t)$ is a random variable with zero mean and correlation function

$$\langle \eta(t) \eta(t+\tau) \rangle = g^2 C(\tau), \quad (16)$$

where the angle brackets denote averages over the distribution that generates $\eta(t)$, and $C(\tau)$ is to be determined self-consistently [3]. It is easiest to express $\eta(t)$ in terms of its Fourier transform $\tilde{\eta}(\omega)$ and the Fourier transform of $C(\tau)$, $\tilde{C}(\omega)$, as

$$\tilde{\eta}(\omega) = g \tilde{C}^{1/2}(\omega) \xi(\omega). \quad (17)$$

Here, ξ is a complex random variable with real and imaginary parts chosen independently, and independently for each discrete value of ω , from a Gaussian distribution with zero mean and variance 1/2. This assures that equation 16 is satisfied. The self-consistency condition that determines $C(\tau)$, which equates the correlation of the mean field to the average autocorrelation of the network interaction term in equation 1, is then expressed in terms of a functional integral over $\xi(\omega)$ as

$$C(\tau) = \int \frac{\delta \xi(\omega)}{\sqrt{2\pi}} \int \frac{\delta \xi^*(\omega)}{\sqrt{2\pi}} \exp\left(-\int d\omega |\xi(\omega)|^2\right) \times \tanh(x(t;\eta)) \tanh(x(t+\tau;\eta)), \quad (18)$$

with η given by the Fourier transform of equation 17 and $x(t, \eta)$ by equation 2. Equation 18 is a self-consistency condition because its right side depends on $C(\tau)$ through equation 17.

We use an iterative approach to solve the dynamic mean-field equations. We start by making an initial guess for the function $C(\tau)$, perform a discrete Fourier transform $C(\tau) \rightarrow \tilde{C}(\omega)$, and use this in equation 17. We then compute $\eta(t)$ by inverse discrete Fourier transformation and solve equation 2 to obtain $x(t)$. Computing $x(t)$ for many different draws of $\xi(\omega)$, we compile a large set of solutions that allows us to compute $C(\tau)$ as a Monte-Carlo approximation of the integrals in equation 18. Starting with this new $C(\tau)$, instead of our initial guess, we repeat the entire procedure, obtain yet another $C(\tau)$, and iterate until the average across iterations of $C(\tau)$ converges. To check against previous calculation, we have verified that we obtain the same results as in Sompolinsky et al. (1988) for $s=0$.

There are no self-consistent solutions of the dynamic mean-field equations in the region where the zero fixed point is stable, but such solutions exists everywhere above the transition line where the zero fixed point becomes unstable (figure 1). The shape of the autocorrelation function (figure 4A) varies continuously across the phase diagram, with no discontinuity at the transition between the regions that do and do not support stable nonzero fixed points (figure 1). The peak height, $C(0)$, increases steadily as a function of either g or s (figure 4B) until it saturates at 1. This reflects the increased cross- or self-coupling driving the units to saturation. The width of the autocorrelation shows a more interesting non-monotonic dependence (figure 4C). As expected, the width diverges at the phase transition between the chaotic and zero-fixed-point regions (left side of the plot in figure 4C). It also diverges for large s and small g (left side of the plot in figure 4C).

In the region with stable nonzero fixed points (top of figure 1), we have thus obtained two mean-field solution, one static and one dynamic, suggesting the coexistence of stable nonzero fixed points and irregular time-dependent activity in the limit $N \rightarrow \infty$. To understand how this limiting behavior arises, we study numerically the relationship between these two types of solutions for finite N .

Lifetime of the Transient Activity

As shown in figure 1, activity arising from typical initial conditions in the region with stable nonzero fixed points exhibits irregular fluctuations that ultimately decay to one of the stable fixed points. Results for the lifetime of this transient dynamic activity for different values of g and networks of different sizes are presented in figure 5. The lifetime depends on the initial state of the network, which was chosen randomly, and, for small networks, on the realization of J . We ran 10,000 trials with different draws of J and different initial conditions to obtain a distribution of lifetimes for the transient activity in the region with nonzero fixed points. This distribution is log-normal (inset, figure 5). We then computed the average of the logarithm of the lifetime for different s and g values (using 100 trials in each case). As can be seen in figure 5, the average log-lifetime is linear in the size of the network, and it increases more rapidly with N as g is increased. The average log-lifetime divided by N and the entropy follow roughly inverse patterns (not shown). This makes sense because the smaller the number of stable fixed points the longer it should take for the network to find one of them. In conclusion, we find that the coexistence of static and dynamic states in the mean-field analysis corresponds to the $N \rightarrow \infty$ limit of a transient fluctuating state that transitions to a non-zero fixed point after a time that grows exponentially with N .

Maximum Lyapunov Exponents

In this subsection, we show numerical results for the largest Lyapunov exponents over a range of g and s values in regions with both persistent and transient irregular activity (figure 1). The long lifetime of the transient activity in the region with nonzero fixed points for finite N allows us to analyze its properties numerically. In particular, we made sure to use large enough networks so the calculation of the Lyapunov exponent converged before the network reached a fixed point. The largest Lyapunov exponent is positive in both of these regions, indicating the exponential sensitivity to initial condition typical of chaos. The largest Lyapunov exponent increases smoothly as a function of both s and g with no indication of any discontinuity at the transition between the persistent and transient regions (figure 6). This suggests that there is no sharp distinction between these two forms of chaotic activity, other than their long-term stability. Rather, as supported by our mean-field results on the correlation function, characteristics of the chaos change continuously across the phase diagram.

Bimodality

To further characterize the nature of the chaotic activity, we simulated networks exhibiting both transient and persistent chaos and extracted distribution of x values over time and network units. As seen in figure 7, these show bimodality that starts within the persistent chaotic region and become more apparent in the region where stable fixed points exist. Although bistability of individual isolated units requires $s > 1$, bimodality appears for values of s well below 1.

Bimodality, as seen in the histograms of figure 7, arises because individual units, especially in the region where stable nonzero fixed points exist, fluctuate chaotically around the two values where $-x + s \tanh(x) = 0$ figure 2. The resulting quasi-bistable behavior can be seen by plotting $\tanh(x)$ as a function of time (figure 8A). Especially for $s > 1$ and small g , the activity is characterized by relatively infrequent flips between fluctuations about these two x values ($x = \pm x_0$), corresponding to $\tanh(x)$ near 1 or -1 , with a log-normal distribution of inter-flip times (figure 8B). The average time between flips shrinks as a function of g and grows as a function of s (figure 8C). For $s = 0$, the inter-flip times, or equivalently times between zero-crossings, follow an exponential distribution. Between small and large value of s , the inter-flip distribution changes continuously from exponential to log-normal.

The flipping of units between quasi-stable states due to network fluctuations may appear similar to the well-studied problem in which a bistable system is perturbed by noise. A unique feature of this system, however, is that the correlation time of the “noise”, which is actually the result of chaotic fluctuations, is on the same order as the time between flips. Thus, the self-consistency condition relates the increase in the width of the correlation function at large s (figure 4C) to the increase in the average time between flips seen in figure 7C.

Discussion

The network model we have studied interpolates between chaotic networks (for s near zero) and networks with large numbers of stable fixed points (for large s), with an intermediate region in which the activity shares features of both. The intermediate activity ranges from patterns dominated by approximate fixed points that are destabilized by chaotic fluctuations to chaotic activity with bimodal activity distributions. In the former case, the chaotic activity acts as a form of colored noise, the “color” induced by its correlations, and drives sign changes in the baseline around which the chaotic fluctuations occur. This form of activity, dominated by flip-like transitions, ultimately terminates when the network finds a true dynamic fixed point, but this occurs over time periods given by a log-normal distribution with a mean that depends exponentially on the size of the network.

Recordings in cortical brain areas reveal firing-rate fluctuations over time scales longer than those expected to be produced by single-neuron properties [8–10]. Thus, it seems likely that these fluctuations arise from network effects [3]. We have shown that strong self-interactions can extend the time-scales over which fluctuations in model network occur, something that is needed to match biological data [1]. A substantial fraction of the connections made by an excitatory cortical neuron are local, and evidence for local clustering of connections has been reported [11, 12]. However, there is also substantial local inhibition. At present, it is not known whether local cortical excitation is strong enough to induce bistability in local clusters, in our language making $s > 1$, but we view it as a possibility.

We can envision two types of applications of clustered networks. First, the long-time-scale dynamics of the flip-like activity might be harnessed through learning algorithms for tasks requiring processing or coherence over long times. Second, the system could be used as a quasi-stochastic “noise” source with a tunable spectrum, which could drive internal network states producing a realization of a Hidden-Markov model. For example, the flips shown in figure 7A have the characteristics of log-normal distributed random events, although they are, of course, actually deterministic. We leave such applications to future work.

ACKNOWLEDGEMENTS

We are grateful to Omri Barak for helpful discussions and to Yashar Ahmadian for providing results on random matrix spectra. Research was supported by the Gatsby Charitable Foundation (M.S., H.S. and L.F.A.), the James S. McDonnell Foundation (H.S.), and the Swartz Foundation and NIH grant MH093338 (L.F.A.).

References

- [1]. Litwin-Kumar A, Doiron B. *Nature neuroscience*. 2012; 15:1498.
- [2]. Wilson HR, Cowan JD. *Biophys. J.* 1972; 12:1. [PubMed: 4332108]
- [3]. Sompolinsky H, Crisanti A, Sommers H. *Phys Rev Lett.* 1988; 61:259. [PubMed: 10039285]
- [4]. Ginibre J. *Journal of Mathematical Physics.* 1965; 6:440.
- [5]. Girko VL. *Teoriya Veroyatnostei I Ee Primeneniya.* 1984; 29:669.
- [6]. Tao T, Vu V, Krishnapur M, et al. *The Annals of Probability.* 2010; 38:2023.
- [7]. Ahmadian, Y.; Fumarola, F.; Miller, KD. *ArXiv e-prints.* 2013. arXiv:1311.4672 [q-bio.NC]
- [8]. Kohn A, Smith MA. *The Journal of neuroscience.* 2005; 25:3661. [PubMed: 15814797]

- [9]. Churchland AK, Kiani R, Chaudhuri R, Wang X-J, Pouget A, Shadlen MN. *Neuron*. 2011; 69:818. [PubMed: 21338889]
- [10]. Churchland MM, Byron MY, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, et al. *Nature neuroscience*. 2010; 13:369.
- [11]. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB. *PLoS biology*. 2005; 3:e68. [PubMed: 15737062]
- [12]. Yoshimura Y, Dantzker JL, Callaway EM. *Nature*. 2005; 433:868. [PubMed: 15729343]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

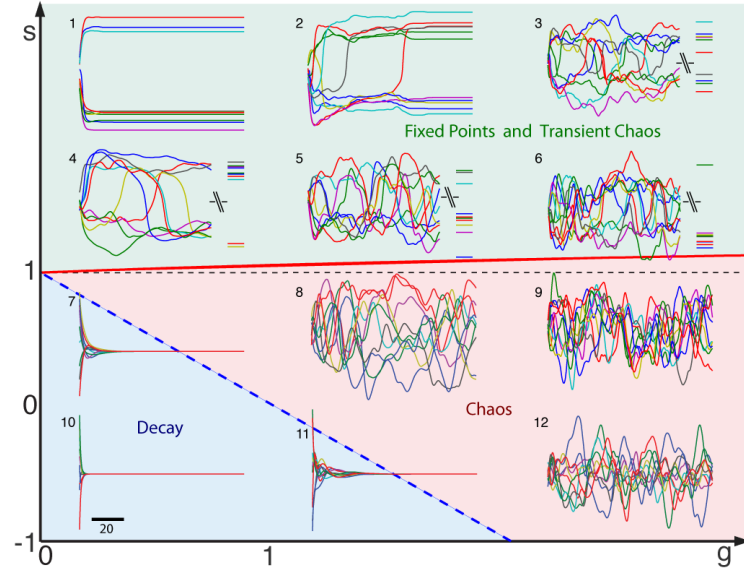




FIG. 1.

Examples of network activity as a function of s and g . Each inset shows $x(t)$ for 6 out of 400 network units as a function of time, with its location indicating the values g and s used: for insets 1-12 (in order): $(g, s) = (0.5, 2.5), (1.3, 2.5), (2.5, 2.5), (0.6, 1.5), (1.5, 1.5), (2.5, 1.5), (0.4, 0.4), (1.5, 0.5), (2.5, 0.5), (0.4, -0.4), (1.2, -0.3), (2.5, -0.5)$. The long-dashed line is the boundary between activity that decays to 0 (inserts 7, 10 & 11) and persistent chaotic activity (inserts 8, 9 & 12). The solid curve is the boundary between persistent chaos and what we will show to be transient chaotic activity that ultimately converges to one of many nonzero fixed points (inserts 1-6). For inserts 3-6, there is a break in the time axis, reflecting the long time required for convergence to a fixed point. The short-dashed line simply indicates $s=1$.

**FIG. 2.**

Graphical solution of the static mean-field equation 6. Solutions $x(\eta)$ are points on the curve $x - s \tanh(x)$ corresponding to a particular value of η . The dashed line shows one such η value and indicates that there are three possible solutions in the region $-\eta_m \leq \eta \leq \eta_m$ (dots). Then open circle indicates that solutions along the portion of the curve with negative slope correspond to unstable fixed points (see text). The arrows show $\pm x_m$, the two local extrema of the function $x - s \tanh(x)$, and $\mp \eta_m$ are its values at these points as indicated by the dotted lines. Along the transition line between the regions with stable fixed point solutions and persistent chaos (figure 1), the unique stable solutions is restricted to the highlighted portion of the curve where $x(\eta)$ has the same sign as η .

**FIG. 3.**

The entropy as a function of s and g , which is $1/N$ times the average (over J) of the logarithm of the number of stable fixed points in the network model. In both plots, the entropy goes to 0 at the values of s and g corresponding to the phase transition between the nonzero fixed-point and persistent chaotic regions in figure 1. A) The entropy as a function of s for, from top to bottom curves, $g = 3, 3.5$ and 4 . B) The entropy over a range of s and g values represented by colors. The green line shows where the entropy reaches 0. The white circles indicate the results obtained from the zero intercepts of the curves in A. C) The weighting function $m(\eta)$ near the transition (solid curve; $s=1.133$ and $g=3$) and away from the transition (dashed curve; $s=1.135$ and $g=3$).




FIG. 4. Autocorrelation functions for different s values and $g = 1.5$. A) $C(\tau)$ for, from the top to the bottom curve, s ranging from 1.6 to -0.4 in steps of 0.2. B) Peak heights of the curves in A. C) Widths at 1/2 peak for the curves in A.

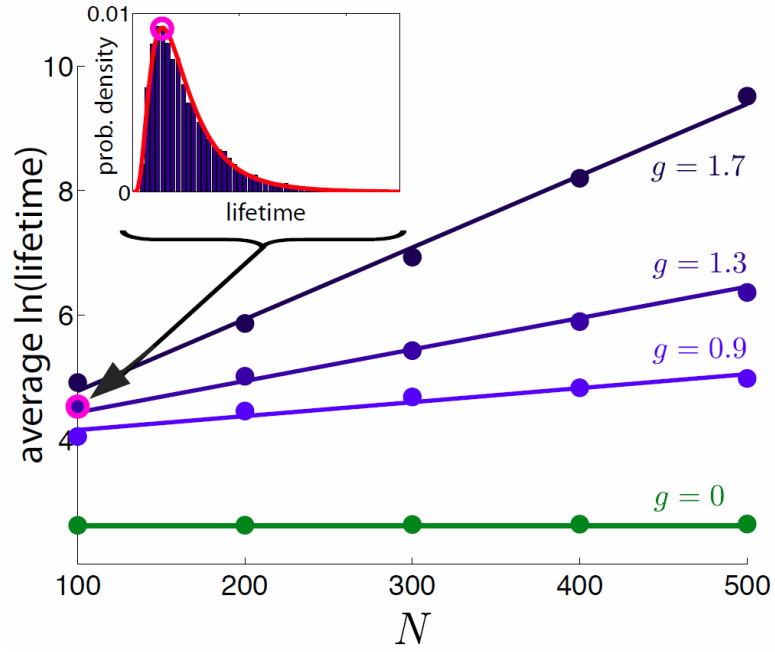


FIG. 5. Exponential dependence on the network size N of the lifetime of the transient activity in the region with stable nonzero fixed points. Inset: Distribution of times to reach a fixed point for $N=100$, $s=2.3$ and $g=1.3$, shown with bars. The curve is a fit to a log-normal distribution. Main figure: The average of the logarithm of the time to reach a fixed point plotted as a function of N , for different g values. In all these examples, $s = 2.3$. For $g = 0$ the units are decoupled and the lifetime is independent of N . For $g > 0$, the lifetime is exponential in N .






FIG. 6. Maximum Lyapunov exponents as a function of s for three different g values. The vertical lines indicate where the transition from persistent chaos (to the left of these lines) to stable fixed points (to the right) occurs. The maximal Lyapunov exponents vary continuously and smoothly through this transition.

**FIG. 7.**

Normalized histograms of unit x values for different s , with $g=3.5$ and $N=1000$. Run time was 20, 000, and x values were sampled at intervals of 50 to avoid temporal correlations. The width of the distributions increases with s , and bimodality first becomes apparent between $s = 0.2$ and $s = 0.3$ (not shown). The transition to the region of transient chaos occurs at $s=1.15$ for this value of g .

**FIG. 8.**

A) Typical activity of one unit in the region where stable fixed points exist ($s = 1.6$, $g = 0.7$). $\tau_{\text{in state}}$ is defined as the time between flips of the unit between states that fluctuate near $\tanh(x) = \pm 1$. B) Distribution of $\tau_{\text{in state}}$ values extracted as in A from all the units. C) The average $\tau_{\text{in state}}$ for different values of g and s , note the logarithmic scale of the vertical axis. From top to bottom, the curves correspond to s ranging from 2 to 1 in steps of 0.2.