Symmetries constrain the transition to heterogeneous chaos in balanced networks

Andrea K. Barreiro[1], J. Nathan Kutz[2], Eli Shlizerman[2]

[1] Department of Mathematics, Southern Methodist University, Dallas, TX, 75275, USA
[2] Department of Applied Mathematics, University of Washington, Seattle, WA, 98195, USA
E-mail: abarreiro@smu.edu

Biological neural circuits display both spontaneous asynchronous activity, and complex, yet ordered activity while actively responding to input. Recently, researchers have demonstrated how this spontaneous behavior underlies computational capabilities in large, recurrently connected networks of firing rate [1,2] and spiking [3] units. Yet, not all spontaneous activity is equal: complex computations may require the rich phase-space of *heterogeneous chaos,* in which each neuron has a different time-dependent firing rate [2,3]. Here, we address the question of how network connectivity structure may affect the transition to heterogeneous chaos in echo-state networks.

We choose a family of firing-rate networks in which the neurobiological constraint of Dale's Law — that most neurons are either excitatory or inhibitory — is satisfied, and in which excitation and inhibition are balanced. We first study the transition to heterogeneous chaos in this setting, using principal component analysis (PCA) to provide a lower-dimensional description of network activity. We find that key characteristics of this transition differ in constrained networks, versus unconstrained networks with similar variability: the transition to heterogenous chaos occurs at higher coupling strengths, and is variable across specific networks.

These properties are a consequence of the fact that the constrained system may be described as a perturbation from a system with non-trivial symmetries. These symmetries imply the presence of both fixed points and periodic orbits that are an organizing center for solutions, even for large perturbations. In comparison, spectral characteristics of the network coupling matrix [4-6] are relatively uninformative about the behavior of the constrained system.

We next investigate the impact of this structure on the computational capabilities of constrained vs. unconstrained networks. We first examine the response of networks to time-dependent inputs [3]. In comparison to unconstrained networks, constrained networks with a delayed transition perform better on traditional population coding of firing rate, but have less ability to separate two different time-dependent inputs in phase space.

In contrast, the delayed transition to chaos has little effect on the ability of constrained networks to reproduce a learning task recently investigated in unconstrained networks [7]. We inspect example networks and find that the addition of the feedback loop quickly moves effective network connectivity away from symmetry and into the chaotic regime at the onset of training.

References

[1] Bertschinger N, Natschläger T: **Real-Time Computation at the Edge of Chaos in Recurrent Neural Networks**, *Neural Computation* 2004, **16:**1413-1436.
[2] Sussillo D, Abbott LF: **Generating Coherent Patterns of Activity from Chaotic Neural Networks**, *Neuron* 2009, **63:**544-557.
[3] Ostojic S: **Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons**, *Nature Neuroscience* 2014, **17**:594-600.
[4] Rajan K and Abbott LF: **Eigenvalue Spectra of Random Matrices for Neural Networks**, *Physical Review L4tters* 2006, **97:**188104.
[5] Wei Y: **Eigenvalue spectra of asymmetric random matrices for multicomponent neural networks**, *Physical Review E* 2012, **85:**066116.
[6] Aljadeff J, Stern M, Sharpee T: **Chaos in heterogenous neural networks: I. The critical transition point.** *BMC Neuroscience* 2014, **15(Suppl 1)**:O20.
[7] Sussillo D, Barak O: **Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks**, *Neural Computation* 2013, **25**:1-24.

Symmetries constrain the transition to heterogeneous chaos in balanced networks (Summary)

Andrea K. Barreiro[1], J. Nathan Kutz[2], Eli Shlizerman[2]

[1] Department of Mathematics, Southern Methodist University, Dallas, TX, 75275, USA
[2] Department of Applied Mathematics, University of Washington, Seattle, WA, 98195, USA
E-mail: abarreiro@smu.edu

We study a family of firing rate networks with random connection weights G_{ij}:

$$\dot{x}_i = -x_i + \sum_j G_{ij}\left(\tanh\left(gx_j\right)\right), \quad x_i(0) = x_i^0, \quad i = 1\ldots N$$

(In results here, $N = 1000$). Each weight G_{ij} is a random variable with a column-dependent mean, so that each cell is either excitatory or inhibitory (the parameter $0 < f < 1$ identifies the proportion of E cells), and so that input to each cell is balanced.

$$\mathbf{G}_{ij} \sim N\left(\mu_E/\sqrt{N}, \sigma_E^2/N\right), \quad i \neq j,\ 1 \le j \le fN$$
$$\sim N\left(\mu_I/\sqrt{N}, \sigma_I^2/N\right), \quad i \neq j,\ fN < j \le N$$
$$\left(f\mu_E + (1-f)\mu_I = 0\right)$$

(We refer to these networks as "constrained" because the choice of **G** has the effect of enforcing Dale's Law.) We contrast with Gaussian networks, in which connections strengths are mean zero ("unconstrained"). Previous work has found that the spectrum of **G** (and therefore the stability of the origin) is largely unchanged (exactly, under certain conditions) (Rajan and Abbott, 2006); Gaussian networks will show a sharp transition to chaos at $g = 1$ in the limit of large N (Sompolinsky et al., 1988).

To examine the transition to chaos, we generated a set of networks from each ensemble; for a range of coupling parameters g, solution trajectories were generated for a set of 10 initial conditions. We used the number of *principal components* (Jolliffe, 2002) needed to capture 99% of the trajectory (in terms of the L^2 norm) as a proxy for attractor dimensionality (see Fig. 1A,B). In contrast to unconstrained networks, constrained networks showed large variability over when the number of principal components needed to represent the solution indicated that chaotic solutions had emerged (see Fig. 1C,D for two examples).

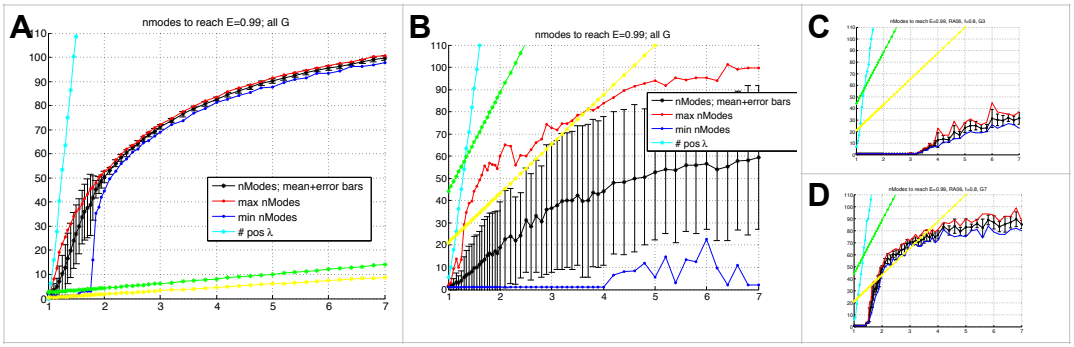

Figure 1: (A-C) Number of principal modes required to represent 99% of the solution energy, as a function of global coupling parameter g. Results are shown as averages (black), maximum (red) and minimum (blue) over both several networks and initial conditions (A-B) or initial conditions only (C-D). (A) Unconstrained networks: all solution trajectories are chaotic for $g > 1.7$. (B) Constrained networks. (C-D) Two example constrained networks; in (C) all trajectories found for $g < 3.5$ are periodic (or fixed points).

To analyze the symmetry underlying the constrained networks, we perform a bifurcation analysis on Eqn. (1,2) for $\sigma_E, \sigma_I = 0$. For N sufficiently large, Eqns. (1,2) supports families of fixed points in which the inhibitory cells split into two clusters, while excitatory cells remain synchronized; a Hopf bifurcation creates a periodic orbit with

the same clustering behavior. As σ_E, σ_I move away from zero, these clusters persist (although "smeared"). In some networks with a delayed transition to heterogeneous chaos, we found that the basins of attraction of the resulting solutions appeared to dominate phase space (they appeared generically with unbiased, asymmetric initial conditions).

We next examined the potential impact of structure on network functionality on computational capacity. First, we gave networks time-dependent inputs and studied their ability to perform two population coding tasks: coding of a single variable through population firing rate, and separating two distinct inputs in phase space. In previous studies, these two tasks have been demonstrated to be better performed in the *classical asynchronous* and *heterogeneous asynchronous* states respectively (Ostojic, 2014). Gaussian networks again showed a robust and stereotyped transition to heterogeneous chaos (one example is shown in Fig. 2A), while balanced E/I networks had significant variability across both individual connectivity matrices an initial conditions (Fig. 2B). Although the classical asynchronous state has no direct analogue in our system (in this firing rate network it corresponds to the fixed point at the origin), we hypothesized that population firing rate will be encoded more accurately for low-mode, periodic solutions; i.e., in the balanced E/I networks in which the transition to heterogeneous chaos has been delayed. Indeed, we found that normalized error in capturing firing rate was about 50% smaller in the constrained networks (Fig. 2C). However, this difference could not be explained simply by the number of principal components required to represent the solutions (Fig. 2D).

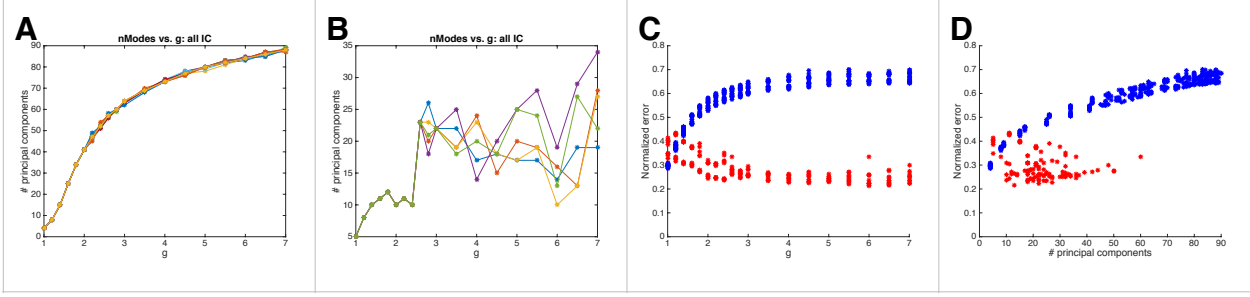

Figure 2: (A-B) Number of principal modes required to represent 99% of the solution energy, as a function of global coupling parameter g. Here all neurons received a time-dependent input (combination of sinusoids). (A) An example unconstrained network (B) An example constrained network. (C-D) Normalized error, for randomly chosen ten-cell subpopulations in representing the true population mean. Blue: Unconstrained networks (data from 2 examples shown); Red: Constrained networks (data from 2 examples shown) (C) Normalized error vs. coupling parameter g. (D) Normalized error vs. number of principal components.

Second, we trained networks to perform a 3-bit memory task (Fig. 3A) using the FORCE learning algorithm (Sussillo and Abbott, 2009). Networks were augmented by a feedback loop, so that the new equations of evolution are:

$$\dot{\mathbf{x}} = -\mathbf{x} + \mathbf{G} \bullet \left(\tanh\left(g\mathbf{x}\right)\right) + \mathbf{W}^{FB}\mathbf{z} + \mathbf{B}\mathbf{u}$$

$$\mathbf{z} = \left(\mathbf{W}^{FF}\right)^{T}\left(\tanh\left(g\mathbf{x}\right)\right)$$

$$\mathbf{W}^{FF}(t) \rightarrow \mathbf{W}^{FF}(t + \Delta t)$$

The matrices $\mathbf{W}^{FF}$ and $\mathbf{W}^{FB}$ are low-rank and therefore sparsely sample the recurrent network: only feedforward weights $\mathbf{W}^{FF}$ are modified during training.

Constrained networks perform this task by construction of a cube of stable fixed points in principal component space, as in the unconstrained case (Sussillo and Barak 2013) (Fig. 3A). Error rates remain similar between the two families of networks (black solid lines in Fig. 3B,C).

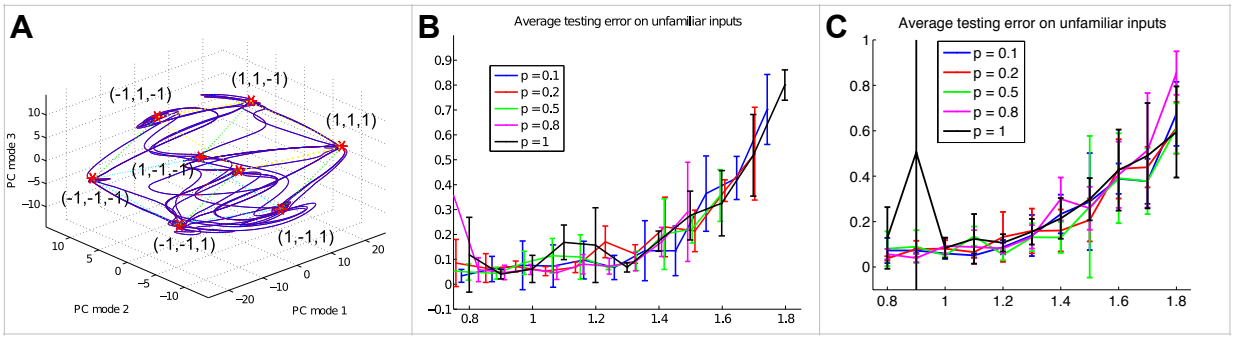

Figure 3: (A) Solution trajectories in the subspace spanned by the first 3 principal components. (B-C) Average error on testing data sets for unconstrained (B) and constrained networks (C), as a function of coupling strength *g* (graphs show averages and error bars over several networks and input data streams).

References

Aljadeff J, Stern M, Sharpee T: **Chaos in heterogenous neural networks: I. The critical transition point.** *BMC Neuroscience* 2014, **15(Suppl 1)**:O20.

Bertschinger N, Natschläger T: **Real-Time Computation at the Edge of Chaos in Recurrent Neural Networks**, *Neural Computation* 2004, **16:**1413-1436.

Jolliffe, I.T.: **Principal Component Analysis**, second edition (Springer), 2002.

Ostojic S: **Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons**, *Nature Neuroscience* 2014, **17**:594-600.

Rajan K and Abbott LF: **Eigenvalue Spectra of Random Matrices for Neural Networks**, *Physical Review L4tters* 2006, **97:**188104.

Sompolinsky H, Crisanti A, Sommers H: **Chaos in random neural networks**, *Physical Review Letters* 1998, **61**:259–262.

Sussillo D, Abbott LF: **Generating Coherent Patterns of Activity from Chaotic Neural Networks**, *Neuron* 2009, **63:**544-557.

Sussillo D, Barak O: **Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks**, *Neural Computation* 2013, **25**:1-24.

Wei Y: **Eigenvalue spectra of asymmetric random matrices for multicomponent neural networks**, *Physical Review E* 2012, **85:**066116.