Weight statistics controls dynamics in recurrent neural networks

Patrick Krauss[1,2]*, **Marc Schuster**[2], **Verena Dietrich**[2], **Achim Schilling**[1,2], **Holger Schulze**[2], **Claus Metzner**[2,3]

1 Cognitive Computational Neuroscience Group at the Chair of English Philology and Linguistics, Department of English and American Studies, Friedrich-Alexander University Erlangen-Nürnberg (FAU), Erlangen, Germany, **2** Experimental Otolaryngology, Neuroscience Group, University Hospital Erlangen, Friedrich-Alexander University Erlangen-Nürnberg (FAU), Erlangen, Germany, **3** Biophysics Group, Department of Physics, Friedrich-Alexander University Erlangen-Nürnberg (FAU), Erlangen, Germany

* patrick.krauss@uk-erlangen.de

Abstract

Recurrent neural networks are complex non-linear systems, capable of ongoing activity in the absence of driving inputs. The dynamical properties of these systems, in particular their long-time attractor states, are determined on the microscopic level by the connection strengths w_{ij} between the individual neurons. However, little is known to which extent network dynamics is tunable on a more coarse-grained level by the *statistical* features of the weight matrix. In this work, we investigate the dynamics of recurrent networks of Boltzmann neurons. In particular we study the impact of three statistical parameters: *density* (the fraction of non-zero connections), *balance* (the ratio of excitatory to inhibitory connections), and *symmetry* (the fraction of neuron pairs with $w_{ij} = w_{ji}$). By computing a 'phase diagram' of network dynamics, we find that balance is the essential control parameter: Its gradual increase from negative to positive values drives the system from oscillatory behavior into a chaotic regime, and eventually into stationary fixed points. Only directly at the border of the chaotic regime do the neural networks display rich but regular dynamics, thus enabling actual information processing. These results suggest that the brain, too, is fine-tuned to the 'edge of chaos' by assuring a proper balance between excitatory and inhibitory neural connections.

Introduction

In contrast to the artificial neural networks used in deep learning, which typically have a strict feed-forward structure, the networks of the brain contain many loops and are therefore recurrent in nature. This feature allows the cortex to maintain dynamical activity even without incoming external stimuli [1] and may therefore underlie such diverse operations as short-term memory [2–4], the modulation of neuronal excitability with attention [2, 5, 6], or the generation of spontaneous activity during sleep [7–9].

The relation between structure and function in recurrent neural networks is a topic of considerable interest in the neurosciences and has already been addressed in several

important publications. In 1988, a transition from a stationary phase to a chaotic phase was predicted [10], occurring at a critical value of a 'gain' parameter that controls the width of the neuron's sigmoidal function. In 2007, it was shown that Network structure of the cerebral cortex shapes functional connectivity on multiple time scales [11]. A paper from 2011 studied the effect of broad degree distributions on network dynamics by interpolating between a binomial and a truncated power-law distribution for the in-degree and out-degree of the neurons independently [12]. In 2014, it was demonstrated that (1) for homogeneous external input, the structure of pairwise correlations between neuron states is mainly determined by the local recurrent connectivity, that (2) homogeneous external inputs provide an additive, unspecific contribution to the correlations, that (3) inhibitory feedback effectively decorrelates neuronal activity, even if neurons receive identical external inputs, and, finally, that (4) identical synaptic input statistics to excitatory and to inhibitory cells increases intrinsically generated fluctuations and pairwise correlations [13]. In 2016, a paper analyzed the anatomical origins of oscillations in the cortical microcircuit. Based on a theoretical reduction of network dynamics, a sensitivity measure was derived, resulting in a frequency-dependent connectivity map that revealed connections crucial for the peak amplitude and frequency of the observed oscillations and identifies the minimal circuit generating a given frequency [14].

Recently, more sophisticated statistical properties of weight matrices were explored. For instance, a model introduced a structured component of connectivity, in addition to random connections, which effectively embeds a feed-forward structure via unidirectional coupling between a pair of orthogonal modes [15]. Another approach studied a class of recurrent network models in which the connectivity is a sum of a random part and a minimal, low-dimensional structure. It was shown that, in such networks, the dynamics is low dimensional and can be directly inferred from connectivity using a geometrical approach [16]. Still another approach focused on the eigenvalue structure of the weight matrix and identified structural properties of networks that are associated with non-normality [17]. A more coarse-grained approach showed that coupling among cortical modules is central. The highest dynamical richness of the network emerges at a critical connectivity at the verge of physical disconnection. Stronger coupling leads to a persistently coherent activity among the modules, while weaker coupling precipitates the activity to be localized solely within the modules [18]. Finally, the effect of external inputs on network dynamics was explored. A paper identified a general criterion that distinguishes two classes of networks depending on properties of the connectivity matrix: networks in which all inputs lead to weak, decaying transients, and networks in which specific inputs elicit strongly amplified transient responses and are mapped onto orthogonal output states during the dynamics [19].

Recent micro-anatomical studies of the brain revealed that neural connectivity in the mammalian cortex has unique statistical properties. In particular, it was found that connections are sparse (low density), so that only a small fraction of possible connections are realized. The distribution of connection strengths is close to log-normal, and thus highly skewed, with a fat tail towards large magnitudes [20, 21]. Although the total number of non-zero connections can vary strongly between neurons, the ratio of excitatory to inhibitory connections is relatively constant [22]. Moreover, cortical networks contain a 'skeleton' of strongly connected neurons, linked pairwise in a bidirectional, symmetric way. This skeleton is embedded in a 'sea' of more weakly, non-symmetrically connected neurons [20].

Whereas the role of this peculiar connection structure is still poorly understood, certain features seem to affect whether the brain can properly act as an information processor. For example, it has been shown that recurrent neural networks can show chaotic behavior for certain ratios between excitatory and inhibitory connections [1, 23]. It has even been speculated that

certain social dysfunctions, such as autism and schizophrenia, are related to an elevated cortical excitation/inhibition balance [24]. Moreover, the discovered skeleton of neurons with strong bi-directional links may help to optimize information storage [25].

In a recent paper [26], we have investigated the relation between connectivity and system dynamics in small motifs of probabilistic neurons with binary outputs, assuming discrete, ternary connection strengths. We found that the balance between excitatory and inhibitory connections has a strong effect on the transition probabilities between successive motif states, whereas the total density of non-zero connections is less important.

Here, we extent our study to larger recurrent networks that consist of deterministic neurons with continuous outputs. Connection strength follow a random, log-normal weight distribution, but have prescribed values of the three control parameters density, balance, and symmetry. We analyze how these parameters affect the dynamical properties of the networks, in particular the Lyapunov exponent of the system trajectory in state space, the period length of cyclic attractors, and the cross correlation between individual neuron states.

As has been previously shown by Hopfield [27], networks with a very large fraction of symmetric bidirectional connections (symmetry parameter close to one) tend to end up in stationary fixed points. We therefore focus on moderate and small symmetry parameters, and explore the two-dimensional phase diagram of system dynamics as a function of balance and density.

We find that this two-dimensional phase plane consists of three basic regions, corresponding to the possible attractors in deterministic and autonomous dynamical systems: periodic state cycles, chaos, and stationary fixed point behavior. Strikingly, it is almost exclusively the balance parameter that controls in which of these three regimes a neural network is located, while the overall density of connections has a much weaker influence. In particular, the networks behave in a way that is suitable for information processing purposes only in a narrow range of balance parameters, located at the edge of the chaotic phase. This theoretical result is in line with the experimental finding that neural networks in the mammalian cortex have moderate degrees of symmetry and are tuned to rather specific values of balance, whereas connection density can vary widely between neurons and over time.

Methods

Neural network model

Our neural networks are based on simple deterministic neurons with zero bias (zero threshold). The total input $z_i(t)$ of neuron i at time t is calculated as:

$$z_i(t) = \sum_{j=1}^{n} w_{ij} \, y_j(t-1) \tag{1}$$

where $y_j(t-1)$ is the state of neuron j at time $t-1$ and w_{ij} is the connection weight from neuron j to neuron i. The new state $y_i(t)$ of neuron i is computed as

$$y_i(t) = \frac{1}{1 + e^{-z_i(t)}}. \tag{2}$$

When simulating the dynamics of the networks, all neurons are updated simultaneously. The total state of a neural network at time step t can be summarized by the n-dimensional vector $\vec{y}(t) = (y_1(t), \ldots, y_n(t))$, where $y_i(t)$ is the output state of neuron i at this time. In each particular case, we simulated $T_{max} = 10^6$ time steps.

Random weight matrix

The structure of a given neural network is defined by its weight matrix $W = \{w_{ij}\}$. Here, we consider networks in which self-connections are forbidden, so that $w_{ii} = 0$. For all non-zero matrix elements, the magnitudes of the weights are distributed according to a log-normal distribution,

$$p(w_{ij}) \propto \mathrm{lognormal}(\mu, \sigma) \tag{3}$$

with location μ and scale σ.

Statistical control parameters $\hat{d}, \hat{b}, \hat{s}$

For a network with n neurons, the weight matrix has dimensions $n \times n$. Due to the excluded self-connections, the diagonal elements of this matrix are zero, leaving a maximum possible number $n(n-1)$ of non-zero matrix elements. We denote the actual number of non-zero weights by $m = m_+ + m_-$, where m_+ and m_- are the numbers of positive and negative weights, respectively. Furthermore, we denote the number of non-zero matrix elements w_{ij} for which a symmetric reverse connection $w_{ji} = w_{ij}$ exists by m_s. Based on these numbers, we define the *density* parameter $\hat{d}$, which varies between 0 for an unconnected and 1 for a fully connected network, by

$$\hat{d} = \frac{m}{n(n-1)}. \tag{4}$$

The *balance* parameter $\hat{b}$, which varies between -1 for a purely inhibitory and $+1$ for a purely excitatory connection matrix, is defined by

$$\hat{b} = \frac{m_+ - m_-}{m}. \tag{5}$$

The *symmetry* parameter $\hat{s}$, which varies between 0 for a completely non-symmetric and $+1$ for a completely symmetric (Hopfield-like [27]) network, is defined by

$$\hat{s} = \frac{m_s}{m}. \tag{6}$$

The meaning of these three control parameters is visualized in Fig 1.

Generation of weight matrices

Random weight matrices with prescribed values of the parameters $\hat{d}$, $\hat{b}$, and $\hat{s}$ were generated in a series of steps. First, a fraction $\hat{d}$ of the weights were drawn independently from a log-normal distribution with location $\mu = 0$ and scale $\sigma = 1$, whereas all remaining weights were set to zero. Second, in order to introduce inhibitory connections to the network, a fraction $(1 - \hat{b})/2$ of the non-zero-weights were randomly selected and multiplied by -1. For all cases with symmetry $\hat{s} = 0$, this produced already the final weight matrix.

For the case $\hat{s} > 0$, further steps were required: the weights above the diagonal of the weight matrix were copied to below the diagonal, thereby creating a perfectly symmetric matrix. Finally, pairs of matrix elements below the diagonal were randomly selected and swapped iteratively, until the desired degree of symmetry $\hat{s}$ was achieved.
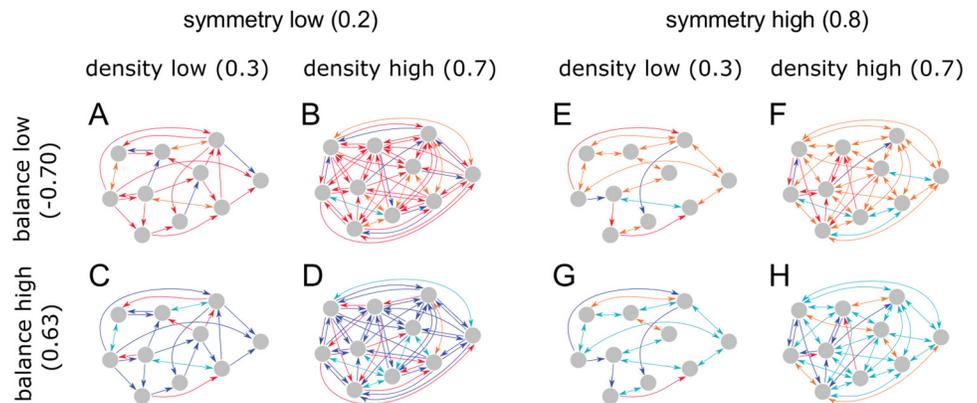
Fig 1. Visualization of the control parameters density $\hat{d}$, balance $\hat{b}$, and symmetry $\hat{s}$ for eight example networks (A-H). Neurons are represented by gray circles, non-zero connections between neurons by arrows. One-headed arrows stand for uni-directional, two-headed arrows for bi-directional connections. Blue/magenta connections are excitatory ($w_{ij} > 0$), red/orange connections inhibitory ($w_{ij} < 0$).

https://doi.org/10.1371/journal.pone.0214541.g001

Fraction of positive Lyapunov exponents $f_{\lambda>0}$

Computing the new network state $\vec{y}(t)$ from the previous state $\vec{y}(t-1)$ can be formally described by a vectorial update function

$$\vec{y}(t) = \vec{F}(\vec{y}(t-1)). \tag{7}$$

Starting from an infinitesimally close initial state $\vec{y}_*(t-1) = \vec{y}(t-1) + \vec{\epsilon}$ will lead to a different final state $\vec{y}_*(t) = \vec{F}(\vec{y}_*(t-1))$. The sensitivity of the update function to this infinitesimal perturbation can be measured by the differential quotient

$$\vec{F}\,'(\vec{y}(t)) \approx \frac{\vec{y}_*(t) - \vec{y}(t)}{|\vec{\epsilon}|}. \tag{8}$$

The *maximum Lyapunov coefficient* λ of the update function is defined as

$$\lambda = \left\langle \ln \left| \vec{F}\,'(\vec{y}(t)) \right| \right\rangle_t, \tag{9}$$

where $\langle\rangle_t$ denotes the time average over all successive states of the system. It can be computed using well-established algorithms [28, 29]. A positive Lyapunov coefficient $\lambda > 0$ indicates that two nearby points in state space diverge exponentially, thus leading to irregular (chaotic) behavior. A zero or negative $\lambda \leq 0$ indicates regular behavior. In general, within an ensemble of networks that are all characterized by the same set of control parameters $(\hat{d}, \hat{b}, \hat{s})$, one finds λ-values of both signs. We therefore counted the fraction of networks (within the statistical ensemble of the considered parameter combination) that had a positive Lyapunov exponent. This fraction $f_{\lambda>0}$ is represented as a color code and shown for all paramater combinations in the 'phase diagrams' below. Note that, the so-defined quantity $f_{\lambda>0}$ reveals a transition at around 50%.

Average period length T_{av}

Our recurrent networks are deterministic and autonomous dynamical systems. Thus, their trajectory $\vec{y}(t)$ through n-dimensional state space is eventually governed by one of three possible attractors: a stationary fixed point, a cycle of period T, or chaotic behavior. For each investigated network, we characterize the type of attractor by the measured period length T, that is, the number of time steps before the system state repeats itself for the first time

$(\vec{y}(t) = \vec{y}(t + T))$. This includes the special cases $T = 1$, corresponding to a stationary fixed point, and $T = \infty$, corresponding to a chaotic attractor. To identify repeating system states, we make use of a hash table. Since period lengths fluctuate for different networks from the same ensemble $(\hat{d}, \hat{b}, \hat{s})$, we compute the ensemble average T_{av} and use this average for color coding the phase diagrams.

Root mean square of cross correlations ρ_{rms}

The Lyapunov coefficient λ and the period length T characterize the long-time behavior of the neural networks. Another property that is relevant for a network's information processing ability is the degree of correlation between individual neuron states $y_i(t)$ at the same time step t. For each pair i, j of neurons, it can be quantified by the Pearson cross correlation coefficient, defined as

$$\rho_{i,j} = \frac{\langle (y_i(t) - \bar{y}_i)(y_j(t) - \bar{y}_j) \rangle_t}{\sigma_i \sigma_j},$$ (10)

where $\bar{y}_k$ is the temporal mean of the time series $y_k(t)$ and σ_k its standard deviation. In cases where σ_i or σ_j were zero, $\rho_{i,j}$ was set to 1. To characterize the global degree of correlation in a given neural network (without caring about the sign of the individual ρ_{ij}), we computed the root mean square (RMS) over all neuron pairs

$$\rho_{\mathrm{rms}} = \sqrt{\frac{1}{n^2} \sum_{i,j} (\rho_{i,j})^2}.$$ (11)

This quantity was additionally averaged over all members of a given $(\hat{d}, \hat{b}, \hat{s})$ ensemble and then used for color coding the phase diagrams.

Results

We first consider non-symmetric networks ($\hat{s} = 0$), that is, networks without any bidirectional links of exactly the same strength. For each combination of balance $\hat{b}$ and density $\hat{d}$ on a 11×10 regular grid, we generate an ensemble of 100 random networks. We then simulate the temporal dynamics of these networks, starting from random initial states. For each ensemble, we compute the fraction of positive Lyapunov coefficients $f_{\lambda > 0}$, the average period length T_{av}, and the RMS of cross correlations ρ_{rms}. The dependence of these dynamical quantities on the statistical control parameters is presented in the form of heat maps, which can be interpreted as dynamical 'phase diagrams' of these recurrent neural networks.

We initially focus on small networks of 100 neurons. When keeping the density close to one and gradually increasing the balance from negative to positive values, we find that the fraction of positive Lyapunov coefficients $f_{\lambda > 0}$, indicating chaotic behavior, is close to zero, except for a narrow interval of balance values around $\hat{b} \approx 0.1$. This chaotic interval broadens as the density parameter is reduced (Fig 2A). In the 2D phase diagram, the chaotic regime therefore has an approximately triangular shape.

Inspecting the temporal output signals of selected neurons in the investigated networks (Fig 2B–2G), it turns out that the two regimes with $f_{\lambda > 0} \approx 0$ at the 'left' and 'right' side of the chaotic regime correspond to periodic attractors (Fig 2B) and fixed point attractors (Fig 2D), respectively. The most interesting dynamics is found at the edge of the chaotic regime (Fig 2C and 2E), where one finds cases of periodic behavior with large period length $T > 2$,
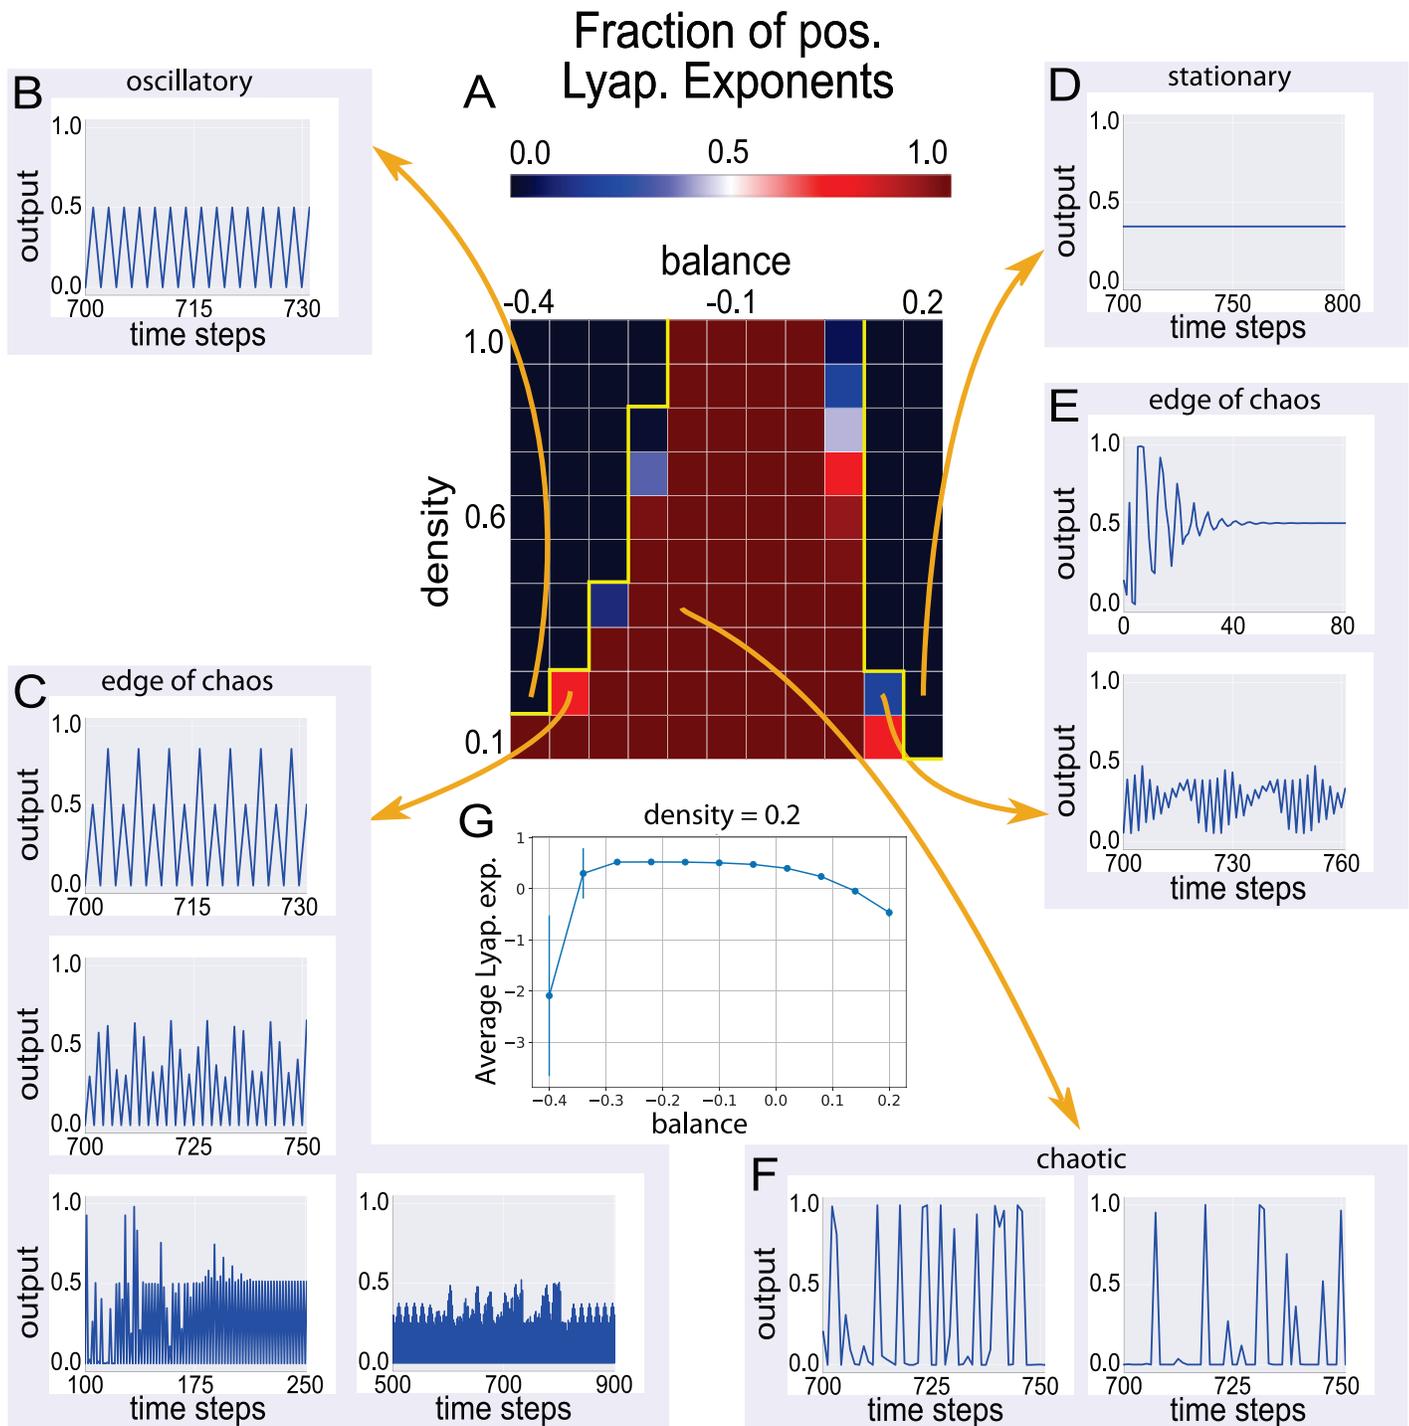
Fig 2. Dynamical phases in recurrent neural networks and characteristic output signals of individual neurons. (A) Two-dimensional phase diagram, showing the fraction of positive Lyapunov exponents $f_{\lambda>0}$ ($\hat{b}, \hat{d}$ | $\hat{s}=0$) as a function of the control parameters balance and density, for a constant symmetry parameter $\hat{s}=0$ (Note that in part G, we show the average Lyapunov exponent along a a one-dimensional cut through the 2D phase space of balance and density, for constant density 0.2). In the heat map, dark blue colors indicate $f_{\lambda>0} \approx 0$, dark red colors $f_{\lambda>0} \approx 1$. The red region in the center of the phase diagram is the chaotic regime, consistent with the irregular outputs of selected neurons (F). The 'left' blue region at negative balance values is the regime of cyclic attractors, often with small period lengths $T \approx 2$, as demonstrated with the neuron output (B). The 'right' blue region at positive balance values is the regime of fixed points, as exemplified with the constant neuron output (D). Note that, in both cases (B, D) the fraction of positive Lyapunov exponents is zero since the dynamics is non-chaotic, and hence the color coding is identical in both cases. The most interesting dynamics is found at the edges of the chaotic regime (C, E), where one finds cases of periodic behavior with large period length $T > 2$, periodic behavior with intermittent bursts, decaying oscillatory behavior, and 'beating' oscillatory behavior. Note that, the sampled time traces depicted in the figure are from selected neurons, not necessarily from within the same network.

periodic behavior with intermittent bursts, decaying oscillatory behavior, and 'beating' oscillatory behavior.

In a next step, we compare the phase distribution of $f_{\lambda>0}$ with that of the other two dynamical quantities (middle and right column in Fig 3). At the same time, we investigate the effect of system size (rows in Fig 3, with different numbers of neurons N).
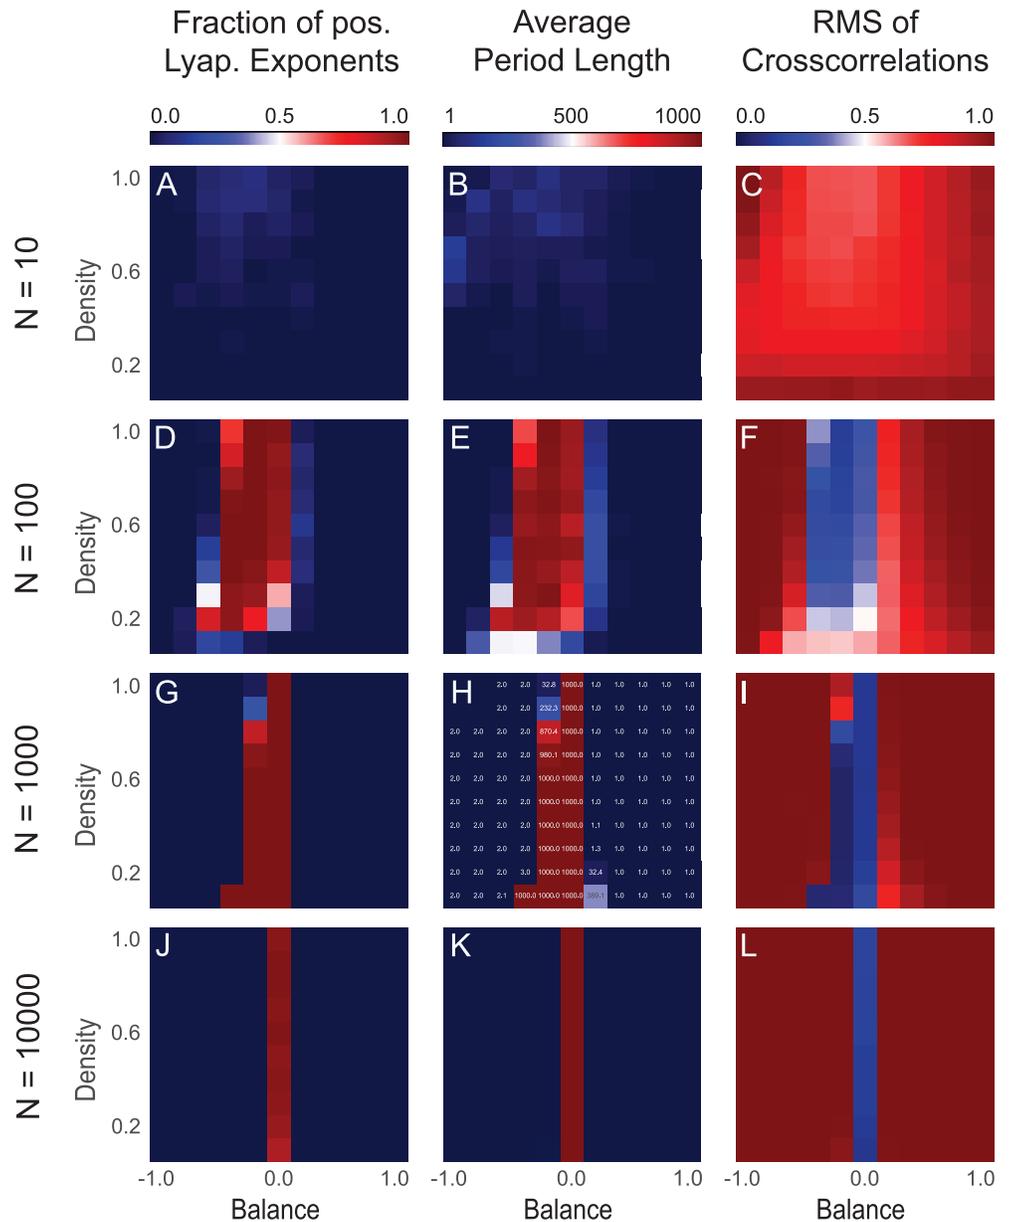

Fig 3. Comparing different dynamical measures, and the effect of system size. The columns correspond to the quantities $f_{\lambda>0}$ (left), T_{av} (middle) and ρ_{rms} (right), as defined in the methods section. The rows from top to bottom correspond to increasing system sizes, characterized by the number of neurons N in the neural networks. For each of the 12 cases, a two-dimensional phase diagram is shown as a function of balance and density, keeping a constant symmetry parameter of $\hat{s} = 0$. The three dynamic phases become apparent only for systems with a minimum size of $N \geq 100$. The three different dynamical measures are mutually consistent. In particular, the chaotic regime is characterized by a $f_{\lambda>0}$ close to one, by a diverging T_{av}, and by a vanishing ρ_{rms}. For large systems with $N \geq 10000$, the density parameter has no more effect on the system dynamics, which is then controlled by the balance only.
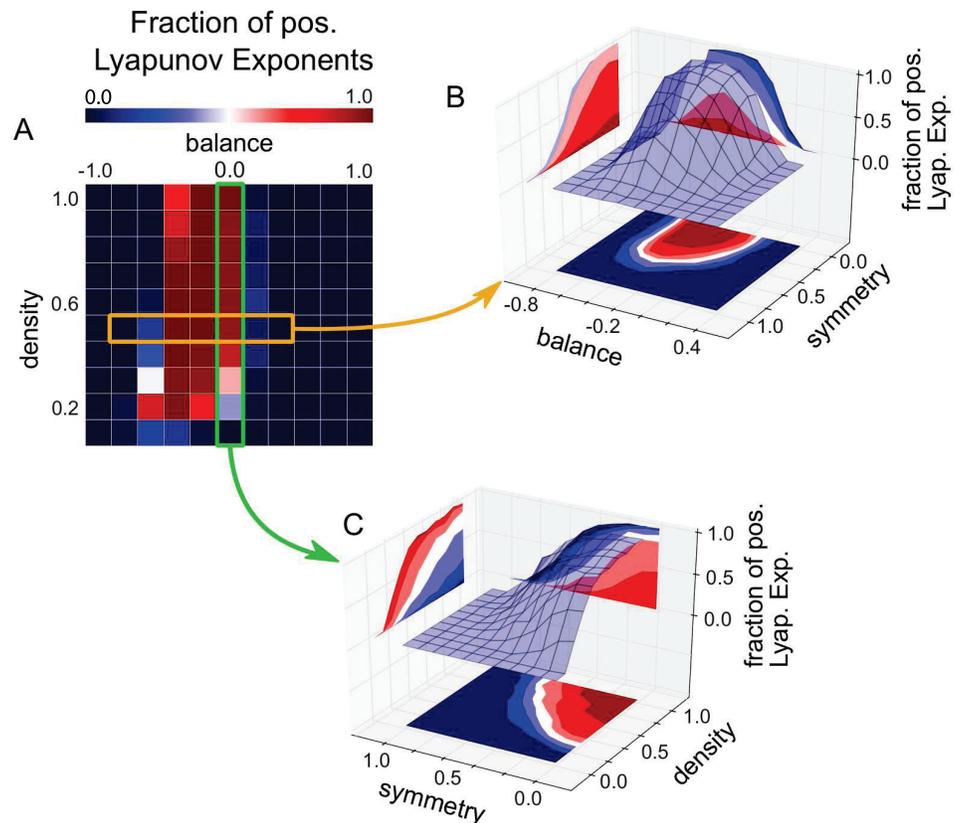
Fig 4. Effect of symmetry $\hat{s}$ on system dynamics. (A): Standard plot of $f_{\lambda>0}$ as a function of balance and density, for constant symmetry $\hat{s} = 0$. (B): Plot of $f_{\lambda>0}$ as a function of balance and symmetry, for constant density $\hat{d} = 0.5$ (see orange box in (A)). (C): Plot of $f_{\lambda>0}$ as a function of symmetry and density, for constant balance $\hat{b} = 0$ (see green box in (A)). For too large symmetry $\hat{s} \approx 1$, the system ends up in fixed point attractors, irrespective of balance and density. Note that the phase diagram shown in (A) is the same as shown in Fig 3D.

https://doi.org/10.1371/journal.pone.0214541.g004

We find that the three different dynamical quantities are mutually consistent. In particular, the chaotic regime is characterized by $f_{\lambda>0} \approx 1$, by a diverging T_{av}, and by a vanishing ρ_{RMS}. The periodic regime is characterized by $T_{\mathrm{av}} \approx 2$ and by a relatively large ρ_{rms}. The fixed point regime is characterized by $T_{\mathrm{av}} = 1$ and, again, by a relatively large ρ_{rms}. Approaching the chaotic regime from either side by changing the balance parameter, T_{av} is rapidly increasing in the border region.

With increasing system size, the influence of the density parameter on the dynamical phase of the networks is diminishing. For large networks with $N \geq 1000$ neurons, the network dynamics is exclusively controlled by the balance parameter.

Finally, we investigate the effect of the symmetry parameter on the network dynamics (Fig 4). By computing a complete 3D phase diagram of $f_{\lambda>0}$ as a function of all three statistical control parameters, we find that balance and density have only an effect on the system dynamics when the symmetry is smaller than one, that is, when there are sufficiently many non-symmetric connections between the neurons. For a too large symmetry $\hat{s} \approx 1$, the system ends up in fixed point attractors, irrespective of balance and density.

Conclusion

One of the earliest studies [10] dealing with dynamical regimes in recurrent neural networks found that a transition from regular to chaotic behavior can be induced by increasing the nonlinearity of the neuron's sigmoidal function. In this work, we have demonstrated that the dynamical behaviour of recurrent neural networks can be effectively tuned by certain statistical properties of the network's connection weight matrix.

In particular, a large fraction of symmetric, bi-directional neural connections ($\hat{s} \approx 1$) favors fixed point attractors, and may therefore be useful for pattern completion tasks, as in the Hopfield model [27]. However, rich dynamical behavior is only possible for moderate or small degrees of symmetry. We point out, that besides our definition of the symmetry parameter, there are alternative definitions which are more fine grained. For instance, the symmetry definition by Esposito et al. [30] takes into account the magnitudes of the forward and backward connections, and vanishes in the case when only uni-directional connections exist. In addition, future extensions of our work could improve our present method to generate weight matrices with pre-defined symmetry, as the present method may create spurious correlations of weight magnitudes between the upper and lower triangle.

For non-symmetric networks, the statistical parameter with the largest impact on system dynamics is the balance $\hat{b}$ [26]. This ratio between excitatory and inhibitory connections controls, with high fidelity, whether a free-running neural network will behave stationary, oscillatory, or irregularly. Moreover, fine tuning of the balance parameter can bring the system to the edge of the chaotic regime, where the outputs of the neurons produce complex wave forms, and where the system may depend sensibly, but still regularly, on external inputs. We speculate that this regime is most suitable for purposes of neural information processing [31–35], and that biological brains may therefore control the parameter $\hat{b}$ in a homeostatic way [1, 36, 37].

By contrast, the impact of the overall connection density $\hat{d}$ on network dynamics, at least in realistically large systems with many neurons, is much smaller than that of the balance $\hat{b}$. In principle, a recurrent neural network can gain or loose a large random fraction of neural connections without changing its dynamical attractor state, as long as the balance $\hat{b}$ remains unchanged. This surprising robustness, for which the term graceful degradation has been coined [38], may help to keep the cortex functional during periods of growth and decay.

In this work we abstracted from biological detail in that each given neuron could have both, positive and negative output weights. By contrast, in the human brain the vast majority of neurons is either purely excitatory or purely inhibitory (Dale's principle), although there are prominent exceptions to this rule, such as the dopaminergic transmission within the basal ganglia [39]. Nevertheless, it might be worthwhile to explore the impact of Dale's principle on networks dynamics.

Future work will also need to clarify how recurrent neural networks, statistically tuned into specific attractor states, react to external inputs. A particularly interesting question will be whether the edge of chaos is also marked by a large mutual information between input signals and the internal sequence of states within the recurrent neural network. Furthermore, Wernecke et al. [40] proposed a method to test for partially predictable chaos, which might be applied to derive a more fine-grained description of the chaotic regime, including the edges of chaos.

Acknowledgments

Author Contributions

Conceptualization: Patrick Krauss, Claus Metzner.

Software: Marc Schuster, Verena Dietrich.

Supervision: Patrick Krauss, Achim Schilling, Holger Schulze, Claus Metzner.

Validation: Achim Schilling, Claus Metzner.

Visualization: Marc Schuster, Verena Dietrich.

Writing – original draft: Patrick Krauss, Claus Metzner.

References

1. Shu Y, Hasenstaub A, McCormick DA. Turning on and off recurrent balanced cortical activity. Nature. 2003; 423(6937):288–293. https://doi.org/10.1038/nature01616 PMID: 12748642

2. Grossberg S, Stone G. Neural dynamics of attention switching and temporal-order information in short-term memory. Memory & Cognition. 1986; 14(6):451–468. https://doi.org/10.3758/BF03202517

3. Stopfer M, Laurent G. Short-term memory in olfactory network dynamics. Nature. 1999; 402(6762):664. https://doi.org/10.1038/45244 PMID: 10604472

4. Kopell N, Whittington M, Kramer M. Neuronal assembly dynamics in the beta1 frequency range permits short-term memory. Proceedings of the National Academy of Sciences. 2011; p. 201019676. https://doi.org/10.1073/pnas.1019676108

5. Carpenter GA, Grossberg S. Neural dynamics of category learning and recognition: Attention, memory consolidation, and amnesia. In: Advances in psychology. vol. 42. Elsevier; 1987. p. 239–286.

6. Buschman TJ, Kastner S. From behavior to neural dynamics: an integrated theory of attention. Neuron. 2015; 88(1):127–144. https://doi.org/10.1016/j.neuron.2015.09.017 PMID: 26447577

7. Babloyantz A, Salazar J, Nicolis C. Evidence of chaotic dynamics of brain activity during the sleep cycle. Physics letters A. 1985; 111(3):152–156. https://doi.org/10.1016/0375-9601(85)90444-X

8. Kavanau JL. Sleep and dynamic stabilization of neural circuitry: a review and synthesis. Behavioural brain research. 1994; 63(2):111–126. https://doi.org/10.1016/0166-4328(94)90082-5 PMID: 7999294

9. Watson BO, Levenstein D, Greene JP, Gelinas JN, Buzsáki G. Network homeostasis and state dynamics of neocortical sleep. Neuron. 2016; 90(4):839–852. https://doi.org/10.1016/j.neuron.2016.03.036 PMID: 27133462

10. Sompolinsky H, Crisanti A, Sommers HJ. Chaos in random neural networks. Physical review letters. 1988; 61(3):259. https://doi.org/10.1103/PhysRevLett.61.259 PMID: 10039285

11. Honey CJ, Kötter R, Breakspear M, Sporns O. Network structure of cerebral cortex shapes functional connectivity on multiple time scales. Proceedings of the National Academy of Sciences. 2007; 104 (24):10240–10245. https://doi.org/10.1073/pnas.0701519104

12. Roxin A. The role of degree distribution in shaping the dynamics in networks of sparsely connected spiking neurons. Frontiers in computational neuroscience. 2011; 5:8. https://doi.org/10.3389/fncom.2011.00008 PMID: 21556129

13. Helias M, Tetzlaff T, Diesmann M. The correlation structure of local neuronal networks intrinsically results from recurrent dynamics. PLoS computational biology. 2014; 10(1):e1003428. https://doi.org/10.1371/journal.pcbi.1003428 PMID: 24453955

14. Bos H, Diesmann M, Helias M. Identifying anatomical origins of coexisting oscillations in the cortical microcircuit. PLoS computational biology. 2016; 12(10):e1005132. https://doi.org/10.1371/journal.pcbi.1005132 PMID: 27736873

15. Landau ID, Sompolinsky H. Coherent chaos in a recurrent neural network with structured connectivity. bioRxiv. 2018; p. 350801.

16. Mastrogiuseppe F, Ostojic S. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. Neuron. 2018; 99(3):609–623. https://doi.org/10.1016/j.neuron.2018.07.003 PMID: 30057201

17. Asllani M, Lambiotte R, Carletti T. Structure and dynamical behavior of non-normal networks. Science advances. 2018; 4(12):eaau9403. https://doi.org/10.1126/sciadv.aau9403 PMID: 30547090

18. Yamamoto H, Moriya S, Ide K, Hayakawa T, Akima H, Sato S, et al. Impact of modular organization on dynamical richness in cortical networks. Science advances. 2018; 4(11):eaau4914. https://doi.org/10.1126/sciadv.aau4914 PMID: 30443598

19. Bondanelli G, Ostojic S. Coding with transient trajectories in recurrent neural networks. arXiv preprint arXiv:181107592. 2018.

20. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB. Highly nonrandom features of synaptic connectivity in local cortical circuits. PLoS biology. 2005; 3(3):0507–0519. https://doi.org/10.1371/journal.pbio.0030068

21. Buzsáki G, Mizuseki K. The log-dynamic brain: how skewed distributions affect network operations. Nature reviews Neuroscience. 2014; 15(4):264–78. https://doi.org/10.1038/nrn3687 PMID: 24569488

22. Gal E, London M, Globerson A, Ramaswamy S, Reimann MW, Muller E, et al. Rich cell-type-specific network topology in neocortical microcircuitry. 2017;(June).

23. van Vreeswijk C, Sompolinsky H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science (New York, NY). 1996; 274(5293):1724–6. https://doi.org/10.1126/science.274.5293.1724

24. Yizhar O, Fenno LE, Prigge M, Schneider F, Davidson TJ, O'shea DJ, et al. Neocortical excitation/inhibition balance in information processing and social dysfunction. Nature. 2011; 477(7363):171. https://doi.org/10.1038/nature10360 PMID: 21796121

25. Brunel N. Is cortical connectivity optimized for storing information? Nature Neuroscience. 2016; 19(5):749–755. https://doi.org/10.1038/nn.4286 PMID: 27065365

26. Krauss P, Zankl A, Schilling A, Schulze H, Metzner C. Analysis of structure and dynamics in three-neuron motifs. Frontiers in Computational Neuroscience. 2019; 13:5. https://doi.org/10.3389/fncom.2019.00005 PMID: 30792635

27. Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. Proceedings of the national academy of sciences. 1982; 79(8):2554–2558. https://doi.org/10.1073/pnas.79.8.2554

28. Wolf A, Swift JB, Swinney HL, Vastano JA. Determining Lyapunov exponents from a time series. Physica D: Nonlinear Phenomena. 1985; 16(3):285–317. https://doi.org/10.1016/0167-2789(85)90011-9

29. Rosenstein MT, Collins JJ, De Luca CJ. A practical method for calculating largest Lyapunov exponents from small data sets. Physica D: Nonlinear Phenomena. 1993; 65(1-2):117–134. https://doi.org/10.1016/0167-2789(93)90009-P

30. Esposito U, Giugliano M, Van Rossum M, Vasilaki E. Measuring symmetry, asymmetry and randomness in neural network connectivity. PloS one. 2014; 9(7):e100805. https://doi.org/10.1371/journal.pone.0100805 PMID: 25006663

31. Skarda CA, Freeman WJ. How brains make chaos in order to make sense of the world. Behavioral and brain sciences. 1987; 10(2):161–173. https://doi.org/10.1017/S0140525X00047336

32. Schiff SJ, Jerger K, Duong DH, Chang T, Spano ML, Ditto WL. Controlling chaos in the brain. Nature. 1994; 370(6491):615. https://doi.org/10.1038/370615a0 PMID: 8065447

33. Birbaumer N, Flor H, Lutzenberger W, Elbert T. Chaos and order in the human brain. Electroencephalography and Clinical Neurophysiology/Supplement. 1995; 44:450–459.

34. Chialvo DR. Critical brain networks. Physica A: Statistical Mechanics and its Applications. 2004; 340(4):756–765. https://doi.org/10.1016/j.physa.2004.05.064

35. Chialvo DR. Emergent complex neural dynamics. Nature physics. 2010; 6(10):744. https://doi.org/10.1038/nphys1803

36. Effenberger F, Jost J, Levina A. Self-organization in balanced state networks by STDP and homeostatic plasticity. PLoS computational biology. 2015; 11(9):e1004420. https://doi.org/10.1371/journal.pcbi.1004420 PMID: 26335425

37. Wilting J, Dehning J, Pinheiro Neto J, Rudelt L, Wibral M, Zierenberg J, et al. Operating in a reverberating regime enables rapid tuning of network states to task requirements. Frontiers in Systems Neuroscience. 2018; 12:55. https://doi.org/10.3389/fnsys.2018.00055 PMID: 30459567

38. Rolls ET, Treves A. The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. Network: computation in neural systems. 1990; 1(4):407–421. https://doi.org/10.1088/0954-898X_1_4_002

39. Kandel ER, Schwartz JH, Jessell TM, of Biochemistry D, Jessell MBT, Siegelbaum S, et al. Principles of neural science. vol. 4. McGraw-hill New York; 2000.

40. Wernecke H, Sándor B, Gros C. How to test for partially predictable chaos. Scientific reports. 2017; 7 (1):1087. https://doi.org/10.1038/s41598-017-01083-x PMID: 28439074