

Theory of gating in recurrent neural networks

Kamesh Krishnamurthy,^{1,*} Tankut Can,^{2,†} and David J. Schwab²

¹*Joseph Henry Laboratories of Physics and PNI, Princeton University, Princeton NJ*

²*Initiative for Theoretical Sciences, Graduate Center, CUNY, New York*

(Dated: September 7, 2020)

Recurrent neural networks (RNNs) are powerful dynamical models, widely used in machine learning (ML) for processing sequential data, and also in neuroscience, to understand the emergent properties of networks of real neurons. Prior theoretical work in understanding the properties of RNNs has focused on models with additive interactions. However, real neurons can have gating i.e. multiplicative interactions, and gating is also a central feature of the best performing RNNs in machine learning. Here, we develop a dynamical mean-field theory (DMFT) to study the consequences of gating in RNNs. We use random matrix theory to show how gating robustly produces marginal stability and *line attractors* – important mechanisms for biologically-relevant computations requiring long memory. The long-time behavior of the gated network is studied using its Lyapunov spectrum, and the DMFT is used to provide a novel analytical expression for the maximum Lyapunov exponent demonstrating its close relation to relaxation time of the dynamics. Gating is also shown to give rise to a novel, discontinuous transition to chaos, where the proliferation of critical points (topological complexity) is decoupled from the appearance of chaotic dynamics (dynamical complexity), contrary to a seminal result for additive RNNs. Critical surfaces and regions of marginal stability in the parameter space are indicated in phase diagrams, thus providing a map for principled parameter choices for ML practitioners. Finally, we develop a field-theory for gradients that arise in training, by incorporating the adjoint sensitivity framework from control theory in the DMFT. This paves the way for the use of powerful field-theoretic techniques to study training/gradients in large RNNs.

I. INTRODUCTION

Recurrent neural networks (RNNs) are powerful dynamical systems that can represent a rich repertoire of trajectories; in fact, they are known to be Turing complete [1]. In modern machine learning, RNNs are used to learn complex dynamics from data with rich sequential/temporal structure such as speech [2, 3], turbulent flows [4–6] or text sequences [7]. RNNs are also influential in neuroscience as models to study the collective behavior of a large network of neurons. For instance, they have been used to explain the temporally-irregular fluctuations observed in cortical networks [8, 9] (and references therein), and how the motor-cortex network generates movement sequences [10, 11].

Classical RNN models typically involve units that interact with each other in an additive fashion – i.e. each unit integrates a weighted sum of the output of the rest of the network. However, researchers in machine learning have empirically found that RNNs with *gating* – a form of multiplicative interaction – can be trained to perform significantly more complex tasks than classical RNNs [7]. This is attributed to the superior ability of gated RNNs to learn long-time dependencies in data [12, 13]. Gating interactions between neurons are also more realistic from a biophysical perspective due to mechanisms such as shunting inhibition [14]. When single-neuron models are endowed with more realistic conductance dynamics,

the effective interactions at the network level have gating effects, and this confers robustness to time-warped inputs [15]. Thus, RNNs with gating interactions not only have superior information processing capabilities, but they also embody a salient feature found in real neurons.

Prior theoretical work on understanding the dynamics of RNNs has mostly focused on RNNs with additive interactions. The original work by Crisanti et al. [16] identified a phase transition in the autonomous dynamics of randomly connected RNNs from quiescence to chaos, when the variance of the coupling weights exceeded a critical value. Subsequent work has extended this analysis to cases where the random connectivity additionally has correlations [17], a low-rank structured component [18, 19], strong self-interaction [20] and heterogeneous variance across blocks [21]. The role of sparse connectivity, spiking dynamics and the single-neuron nonlinearity was studied in [9]. The effect of a Gaussian noise input was analysed in [22]. With regard to training on tasks, recent theoretical work has focused on the correspondence between classical additive RNNs and kernel methods [23, 24].

In this work, we study the consequences of gating interactions on the dynamics of RNNs. We introduce a gated RNN model that naturally extends the classical RNN by augmenting it with two kinds of gating interactions: i) an update gate that acts like an adaptive time-constant and ii) an output gate which modulates the relative weight of the output of a neuron. The choice of these forms for gates are inspired from single-neuron conductance dynamics [15, 25] and retain the most salient aspects of the gated RNNs in machine learning. Our gated RNN re-

* lead author; kameshk@princeton.edu

† corresponding authors: tankut.can@gmail.com, kameshk@princeton.edu

duces to the classical RNN ([16, 26]) when the gates are open, and is closely related to the state-of-the-art gated RNNs in machine learning when the dynamics are discretized.

We develop a dynamical mean-field theory (DMFT) for the gated RNN based on the Martin-Siggia-Rose-De Dominicis-Jansen (MSRDJ) formalism [8, 22, 27–31] and use it to show how gating can produce substantially richer dynamical phases compared to additive RNNs. Identifying the different dynamical phases is not only useful in characterizing the expressivity of RNNs, but it has significant implications for training success in machine learning tasks – indeed, the regime chosen for initialization is known to be one of the most important factors deciding the outcome of training [32]. Furthermore, for very large networks there are regimes in which parameters barely change from their initial values during training [33, 34]. Thus the phase diagram is of both theoretical and practical interest (in choosing initial parameters).

In mapping out the phase diagram, we identify critical surfaces in the parameter space where the dynamics transition from stationary to chaotic activity. Studying the transition to chaos is not only interesting in understanding the dynamics of large complex networks, but it is also of significant practical interest since training RNNs works best near critical lines – an empirical phenomenon often referred to as “edge of chaos” hypothesis [10, 35–37]. This is presumably due to the presence of long timescales in the dynamics close to chaotic transitions. However, we find that not all critical surfaces might be beneficial – indeed, we identify a new, discontinuous transition to chaos produced by gating, where long timescales are absent, and the proliferation of critical points is decoupled with the appearance of a chaotic attractor; this is in contrast to the interesting result of Wainrib & Touboul [38], which showed a tight link between the two phenomena in classical RNNs. We study this new chaotic transition in detail using numerical and analytical methods.

We combine tools from random matrix theory with the DMFT to study how the gating shapes the timescales in local dynamics. We show that gating produces slow modes, marginal stability and line attractors. These line attractors are widely-studied mechanisms that are critical for computations involving long memory [39–42], and to date, models generating line attractors require fine-tuning. Here, we show that gating can robustly generate line attractors.

Finally, we develop a DMFT for gradients that arise in training RNNs by incorporating the adjoint formalism from control theory [43] in the MSRDJ field theory formalism. This allows using the powerful field-theoretic techniques developed for the DMFT – typically used to study forward dynamics – to investigate gradients and training in large RNNs. This potentially opens a novel direction to incorporate statistical physics tools in the theoretical analysis of machine learning models and biological systems that organize based on gradient informa-

tion. The techniques introduced here naturally generalize to other forms of gating [44], and thus provide a principled way to assess the effect that architectural choices have on dynamics and gradients. We provide a brief list of the main results below:

A. Summary of main results

- We develop a DMFT for the gated RNN based on the MSRDJ formalism [Sec. III] and use it to show how gating can produce substantially richer dynamical regimes compared to additive RNNs.
- We show how gates shape the local dynamics through their effect on the Jacobian spectrum in [Sec. IV B]. To do this, we combine random matrix theory with the DMFT to derive an expression for the spectral support of the Jacobian explicitly showing its dependence on the gates [Sec. IV A].
- We show how the update gate produces slow modes in the dynamics [Sec. IV B] and in the limit, it can poise the system at a marginally stable point.
- We study marginal stability and the resulting line attractor dynamics in [Sec. V]. We show that gating can robustly generate line attractors, where inputs along the line-attractor manifold are retained for long times and inputs not aligned with the manifold decay with a spectrum of timescales.
- We study how gates shape the long-time dynamics via the Lyapunov spectrum [Sec. VI A] and derive a DMFT prediction for the maximum Lyapunov exponent showing the role of the gates [Sec. VI B]. This corroborates the local picture of the dynamics that emerges from Sec. IV B.
- We provide an alternate derivation of the maximum Lyapunov exponent using random matrix theory [Sec. VI C], which explicitly shows the intimate link between the Lyapunov exponent and the relaxation time of the dynamics.
- Gating is shown to give rise to a novel, discontinuous transition to chaos where the appearance of chaotic dynamics is decoupled from the proliferation of unstable fixed-points [Sec. VII]. We show that this transition is characterized by a state where a stable fixed-point coexists with chaotic dynamics. This state displays long-lived chaotic transients that grow with system size [Sec. VII C].
- Biases in the nonlinearities are shown to produce non-trivial fixed-points – useful for memory functions – and control the transition to chaos via their effect on the Jacobian spectrum [Sec. VIII].
- In [Sec. IX] we summarize the dynamical regimes, critical lines, and the regions of marginal stability by constructing phase diagrams for the gated RNN.

- In [Sec. X] we extend the adjoint sensitivity formalism to develop a DMFT for gradients that arise while training RNNs, thus enabling a way for field-theoretic analysis of gradients. As an example application, the DMFT for the gradients is used to provide predictions for the variance of the gradients in large RNNs at initialization.

II. A RECURRENT NEURAL NETWORK MODEL TO STUDY GATING

Classical RNN models that are used to model temporal data or study collective dynamics of neural circuits typically consist of units linked by additive interactions. We aim to study an extension of a classical RNN that has multiplicative *gating* interactions. This gated RNN reduces to the classical RNN in certain limits (when the gates are open) and is closely related to gated RNN models popular in machine learning. First we describe the classical RNN, which is defined by a system of N coupled differential equations of the form [16, 26]:

$$\dot{h}_i(t) = -h_i(t) + \sum_{j=1}^N J_{ij}^h \phi(h_j(t)) \quad (1)$$

where h_i represents the synaptic current of the i^{th} unit and is an exponentially-averaged (with time constant of one) version of the input from the rest of the network, J_{ij}^h are the coupling strengths between the units, and ϕ is a nonlinear activation function such as \tanh with $\phi(h_i(t))$ representing the output or activity of the i^{th} unit. This model has been used to explain the temporally-irregular rate fluctuations in circuits of cortical neurons [16], and variants of this model are a popular choice to learn sequence structure in machine learning applications [7, 10, 11]

We augment the classical RNN model with two gates that facilitate multiplicative interactions: (i) an *update* (or z -) gate which controls the rate of integration, and (ii) an *output* (or r -) gate which modulates the strength of the output. The equations describing the gated RNN are given by:

$$\dot{h}_i(t) = \sigma_z(z_i) \left[-h_i(t) + \sum_{j=1}^N J_{ij}^h \phi(h_j(t)) \cdot \sigma_r(r_j) \right] \quad (2)$$

where $\phi(x) = \tanh(g_h x + \beta_h)$ is the activation function and $\sigma_{z,r}(x) = [1 + \exp(-\alpha_{z,r} x + \beta_{z,r})]^{-1}$ are sigmoidal gating functions. These functions are parametrized by gain parameters ($g_h, \alpha_{z,r}$) and biases ($\beta_{h,z,r}$), which constitute the parameters of the gated RNN. The input to the gating functions – $z_i(t), r_i(t)$ – evolves according to

dynamics driven by the output $\phi(h(t))$ of the network:

$$\tau_z \dot{z}_i(t) = -z_i(t) + \sum_{j=1}^N J_{ij}^z \phi(h_j(t)) \quad (3)$$

$$\tau_r \dot{r}_i(t) = -r_i(t) + \sum_{j=1}^N J_{ij}^r \phi(h_j(t)) \quad (4)$$

Note that the coupling matrices $J^{z,r}$ for z, r are distinct from J^h .

The value of $\sigma_z(z_i)$ can be viewed as a dynamical time-constant for the i^{th} unit, while the output gate controls $\sigma_r(r_i)$ and modulates the output strength of unit i , which in the presence of external input can modulate the relative strength between internal activity and external input. In the limit $\sigma_z, \sigma_r \rightarrow 1$, we recover the dynamics of the classical model (eq. 1).

We choose the coupling weights from a Gaussian distribution with variance scaled such that the input to each unit will remain $O(1)$. Specifically,

$$J_{ij}^{h,z,r} \sim \mathcal{N}\left(0, \frac{1}{N}\right) \quad (5)$$

This choice of couplings is a popular initialization scheme for RNNs in machine learning [7, 32], and also in models of cortical neural circuits [16, 21]. It is possible to consider another alternative for the coupling matrix for the gating variables: we can choose ($J^{z,r}$) to be a diagonal matrix, e.g. $J^{z,r} = \mathbb{1}$. In this case, gating is *internal* – i.e., z_i, r_i only depend on h_i . In the rest of the paper, we analyze the various dynamical regimes the gated RNN exhibits and how the gates shape these dynamics.

III. DYNAMICAL MEAN-FIELD THEORY FOR GATED MODEL

In this section, we develop a Dynamical Mean-Field Theory (DMFT) for the gated RNN using the field theory formalism developed in [27, 28, 30, 45] (see [29, 31] for a recent review). The DMFT is a powerful analytical framework used to study the dynamics of disordered systems, and it traces its origins to the study of dynamical aspects of spin glasses [46, 47] and has been later applied to the study of random neural networks [9, 16, 22, 48]. We provide a detailed, self-contained derivation of the DMFT for the gated RNN in Appendix A, and explain only the key concepts in this section. The starting point is a generating functional – akin to the generating function of a random variable – which takes an expectation over the paths generated by the dynamics. The generating functional is defined as

$$Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = \mathbb{E} \left[\exp \left(i \sum_{j=1}^N \int \hat{\mathbf{b}}_j(t)^T \mathbf{x}_j(t) dt \right) \right] \quad (6)$$

where $\mathbf{x}_j(t) \equiv (h_j(t), z_j(t), r_j(t))$ is the trajectory and $\hat{\mathbf{b}}_j(t) = (\hat{b}_j^h, \hat{b}_j^z, \hat{b}_j^r)$ is the argument of the generating

functional. We have also included external fields $\mathbf{b}_j = (b_j^h, b_j^z, b_j^r)$ which are used to calculate the response functions. The measure in the expectation is a path integral over the dynamics (c.f. Appendix A). The generating functional formalism is typically used to describe stochastic dynamics driven by a noise process, and the expectation is taken with respect to the measure over paths induced by the noise. In our case, the dynamics are deterministic and one can think of the expectation formally as the limiting measure when the noise variance goes to zero [49]. The subscript, \mathcal{J} in $Z_{\mathcal{J}}$ is meant to indicate the dependence on the random matrices $J^{h,z,r}$.

The generating functional is used to calculate correlation and response functions using the appropriate (variational) derivatives. For instance, the two-point function for the h field is given by:

$$\langle h_i(t)h_i(t') \rangle = \frac{\delta^2}{\delta \hat{b}_i^h(t') \delta \hat{b}_i^h(t)} Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] \Big|_{\mathbf{b}=0} \quad (7)$$

Up until this point, the formalism is quite general, and does not rely on the specific form of the dynamics. However, for large random networks, we expect certain quantities such as the population averaged correlation function $C_h \equiv N^{-1} \sum_i \langle h_i(t)h_i(t') \rangle$ to be self-averaging, and thus not vary much across realizations. Indeed, when we perform a disorder average (over \mathcal{J}), the generating functional $\bar{Z} = \langle Z_{\mathcal{J}} \rangle_{\mathcal{J}}$ is given by an integral of an exponential term (the action) which is extensive in N (c.f. Appendix A), and thus we can approximate \bar{Z} with it's value evaluated at the saddle-point of the action. The saddle-point approximation of \bar{Z} is a Gaussian path-integral of the form

$$\bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] \simeq e^{-N \cdot \mathcal{L}_0[\hat{\mathbf{b}}, \mathbf{b}]} \cdot \int \mathcal{D}\hat{Q} \mathcal{D}Q e^{-N \cdot \mathcal{L}_2[\hat{Q}, Q, \hat{\mathbf{b}}, \mathbf{b}]} \quad (8)$$

where \mathcal{L}_0 is the saddle-point value of the action and the integral over \mathcal{L}_2 account for the Gaussian fluctuations about the saddle-point. The generating functional formalism, in principle, allows studying the scale of fluctuations, which is an interesting direction but beyond the scope of this work (also c.f. [8]). If we neglect the fluctuations, then we see that in the saddle-point approximation the generating functional is a product of N identical generating functionals – i.e. *the sites get decoupled*:

$$\bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] \simeq e^{-N \cdot \mathcal{L}_0[\hat{\mathbf{b}}, \mathbf{b}]} = Z_0[\hat{\mathbf{b}}, \mathbf{b}]^N \quad (9)$$

So the problem reduces to studying a single-site picture, which corresponds to three coupled scalar variables (h, z, r) driven by Gaussian noise processes whose correlation functions must be computed self-consistently (Appendix A). The (stochastic) dynamics of the scalar variables takes the form

$$\dot{h}(t) = -\sigma_z(z) \cdot h(t) + \sigma_z(z) \cdot \eta_h(t) \quad (10)$$

$$\tau_z \dot{z}(t) = -z(t) + \eta_z(t) \quad (11)$$

$$\tau_r \dot{r}(t) = -r(t) + \eta_r(t) \quad (12)$$

and the Gaussian processes, η_h, η_z , and η_r have correlation functions given by:

$$\langle \eta_h(t) \cdot \eta_h(t') \rangle = \langle \phi(h(t)) \sigma_r(r(t)) \cdot \phi(h(t')) \sigma_r(r(t')) \rangle \quad (13)$$

$$\langle \eta_z(t) \cdot \eta_z(t') \rangle = \langle \phi(h(t)) \cdot \phi(h(t')) \rangle \quad (14)$$

$$\langle \eta_r(t) \cdot \eta_r(t') \rangle = \langle \phi(h(t)) \cdot \phi(h(t')) \rangle \quad (15)$$

Eqns. 11-10 are the DMFT description of the dynamics, in which the $3N$ deterministic ODEs have been reduced to 3 SDEs. The correlation functions in the single-site picture such as $C_h(t, t') = \langle h(t)h(t') \rangle$ are the order parameters in the DMFT, and correspond to the population-averaged correlation functions in the full network. Qualitative changes in the correlation functions correspond to transitions between dynamical regimes of the neural network.

The DMFT equations 11-10 can be extended to get equations of motions for the correlation functions C_h, C_z, C_r , which will prove useful later on. ‘Squareing’ eqs. 11-10 we get

$$[-\partial_\tau^2 + C_{\sigma_z}(\tau)] C_h(\tau) = C_{\sigma_z}(\tau) C_{\sigma_r}(\tau) C_\phi(\tau) \quad (16)$$

$$[-\tau_z^2 \partial_\tau^2 + 1] C_z(\tau) = C_\phi(\tau) \quad (17)$$

$$[-\tau_r^2 \partial_\tau^2 + 1] C_r(\tau) = C_\phi(\tau) \quad (18)$$

where we have used the shorthand $\sigma_z(t) \equiv \sigma_z(z(t))$; $\phi(t) \equiv \phi(h(t))$, and assumed that the network has reached steady-state, so that the correlation functions are only a function of the time difference $\tau = t - t'$. The role of the z -gate as an adaptive time constant is evident in eq. 16. For time-independent solutions, i.e. fixed points, eqns. 16 - 18 simplify to read

$$\Delta_z \equiv \langle z^2 \rangle = \int Dx \phi(\sqrt{\Delta_h x})^2 = \Delta_r \quad (19)$$

$$\Delta_h \equiv \langle h^2 \rangle = \int Dx Dy \phi(\sqrt{\Delta_h x})^2 \sigma_r(\sqrt{\Delta_r y})^2 \quad (20)$$

where we have used Δ instead of C to indicate fixed-point variances, and Dx is the standard Gaussian measure.

In order to solve the DMFT equations, we use a numerical method described in [50]. Specifically, we generate noise paths $\eta_{h,z,r}$ starting with an initial guess for the correlation functions, and then iteratively update the correlation functions using the mean-field equations till convergence. The classical method of solving the DMFT by mapping the DMFT equations to a second-order ODE describing the motion of a particle in a potential cannot be used in the presence of multiplicative gates. In Fig. 1, we see that the solution to the mean-field equations agrees well with the true population-averaged correlation function; Fig. 1 also shows the scale of fluctuations around the mean-field solutions (Fig. 1 thin black lines). The order parameters from the DMFT (the correlation functions) provide a description of the macroscopic dynamical state of the network in the different parameter

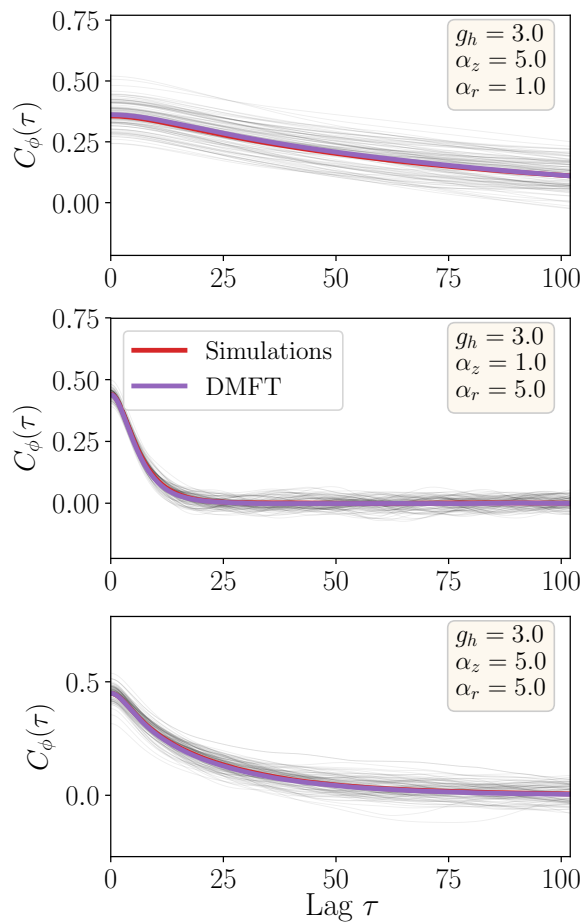


FIG. 1. *Validating the DMFT*: panels show the comparison between the population averaged correlation functions $C_\phi(\tau) \equiv \langle \phi(t)\phi(t+\tau) \rangle$ obtained from the full network simulations in steady-state (purple line) and from solving the DMFT equations (red line) for three distinct parameter values. The lag τ is relative to τ_h (taken to be unity). Thin black lines show the fluctuations around the average correlation functions. $N = 5000$ for all the panels.

regimes. We next use the DMFT to uncover phase transitions and study how the gates affect the local and long-time dynamics of the network.

IV. HOW THE GATES SHAPE THE LOCAL DYNAMICS

We first study the local dynamics of the gated RNN through the lens of the state-to-state Jacobian, and describe how these dynamics are shaped by the gates. The instantaneous Jacobian describes the linearized dynamics about an operating point, and the eigenvalues of the Jacobian inform us about the range of timescales present and the local stability of the dynamics. As we show below, the spectrum of the Jacobian depends on the correlation functions, which are the order parameters in the mean-field picture of the dynamics. We study how the

gates shape the support and the local density of eigenvalues, through their influence on the correlation functions.

The linearized dynamics in the tangent space at an operating point $\mathbf{x} = (\mathbf{h}, \mathbf{z}, \mathbf{r})$ are given by

$$\delta \dot{\mathbf{x}} = \mathcal{D}(t)\delta \mathbf{x} \quad (21)$$

where \mathcal{D} is the $3N \times 3N$ dimensional instantaneous Jacobian of the full network equations. Linearization of eq. 2 yields the following expression for \mathcal{D}

$$\mathcal{D} = \begin{pmatrix} [\sigma_z](-\mathbb{1} + J^h[\phi'\sigma_r]) & \mathbb{D} & [\sigma_z]J^h[\phi'\sigma_r'] \\ \tau_z^{-1}J^z[\phi'] & -\tau_z^{-1}\mathbb{1} & 0 \\ \tau_r^{-1}J^r[\phi'] & 0 & -\tau_r^{-1}\mathbb{1} \end{pmatrix} \quad (22)$$

where $[x]$ denotes a diagonal matrix with the diagonal entries given by the vector x . The term $\mathbb{D}_{ij} = \delta_{ij}\sigma'_z(z_i)(-h_i + \sum_j J_{ij}^h\phi(h_j)\sigma_r(r_j)) = [-\sigma'_z(z)h] + [\sigma'_z \odot J^h(\phi \odot \sigma_r)]$ arises when we linearize about a time-varying state and will be zero for fixed-points. We introduce the additional shorthand $\phi'(t) = \phi'(h(t))$ and $\sigma'_{r/z} = \sigma'_{r/z}(r/z(t))$. We will need additional tools from random matrix theory, along with the mean-field theory developed above, to analyse the spectrum of the Jacobian \mathcal{D} . We develop the random matrix formalism for the gated RNN below.

A. Spectral support of the Jacobian

The Jacobian \mathcal{D} is a block-structured matrix constructed from the random coupling matrices $J^{h,z,r}$ and diagonal matrices of the state variables. In the limit of large N , we expect the spectrum to be self-averaging – i.e. the distribution of eigenvalues for a random instance of the network will approach the ensemble-averaged distribution. We can thus gain insight about typical dynamical behavior by studying the ensemble (or disorder) averaged spectrum of the Jacobian. Our starting point is the disorder-averaged spectral density $\mu(\lambda)$ defined as

$$\mu(\lambda) = \frac{1}{3N} \mathbb{E} \left[\sum_{i=1}^{3N} \delta(\lambda - \lambda_i) \right] \quad (23)$$

where the λ_i are the eigenvalues of \mathcal{D} for a given realization of $J^{h,z,r}$ and the expectation is taken over the distribution of real Ginibre random matrices from which $J^{h,z,r}$ are drawn. Using an alternate representation for the Dirac delta function in the complex plane ($\delta(\lambda) = \pi^{-1}\partial_{\bar{\lambda}}\lambda^{-1}$), we can write the average spectral density as

$$\mu(\lambda) = \frac{1}{\pi} \frac{\partial}{\partial \bar{\lambda}} \mathbb{E} \left[\frac{1}{3N} \text{Tr} \left[(\lambda \mathbb{1}_{3N} - \mathcal{D})^{-1} \right] \right] \quad (24)$$

where $\mathbb{1}_{3N}$ is the $3N$ -dimensional identity matrix. \mathcal{D} is in general non-Hermitian, so the support of the spectrum is not limited to the real line, and the standard procedure of studying the Green's function $G(\lambda, \bar{\lambda}) =$

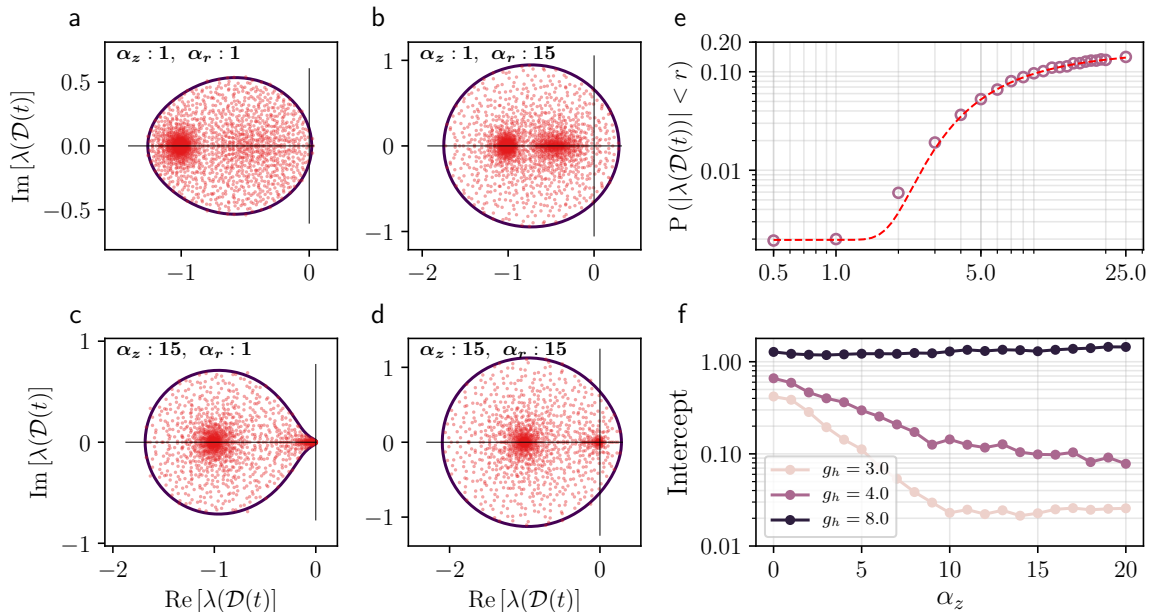


FIG. 2. *How gates shape the Jacobian spectrum:* (a-d) Jacobian eigenvalues (red dots) of the gated RNN in (time-varying) steady-state. The dark outline is the spectral support curve predicted by eq. 29. Bottom row corresponds to larger α_z and right column corresponds to large α_r . (e) Density of Jacobian eigenvalues in a disk of radius $r = 0.05$ centered at the origin plotted against α_z . Circles are numerical density calculated from the true network Jacobian (averaged over 10 instances), and the dashed line is a fit from eq. 33. (f) Intercept of the spectral curve on the imaginary axis, plotted against α_z for three different values of g_h ($\alpha_r = 0$). For network simulations $N = 2000$, $g_h = 3$, $\tau_r = \tau_z = 1$ unless otherwise stated, and all biases are zero.

$(3N)^{-1} \text{Tr} \mathbb{E} \left[(\lambda \mathbb{1}_{3N} - \mathcal{D})^{-1} \right]$ by analytic continuation is not applicable since it is non-holomorphic on the support. Instead, we use the method of Hermitization [51, 52] to analyse the resolvent for an expanded $6N \times 6N$ Hermitian matrix H

$$\mathcal{G}(\eta, \lambda, \bar{\lambda}) = \mathbb{E} \left[(\eta \mathbb{1}_{6N} - H)^{-1} \right] \quad (25)$$

$$H = \begin{pmatrix} 0 & \lambda - \mathcal{D} \\ \bar{\lambda} - \mathcal{D}^T & 0 \end{pmatrix} \quad (26)$$

and the Green's function for the original problem is obtained by considering the lower-left block of \mathcal{G}

$$G(\lambda, \bar{\lambda}) = \lim_{\eta \rightarrow i0^+} \frac{1}{3N} \text{Tr} \mathcal{G}_{21}(\eta, \lambda, \bar{\lambda}) \quad (27)$$

In order to proceed further, we invoke an ansatz called the local chaos hypothesis [53, 54], which posits that for large random networks in steady-state, the state variables are statistically independent of the random coupling matrices $J^{z,h,r}$ (also see [55]). The local chaos hypothesis combined with the assumption of self-averaging for large networks implies that the Jacobian (eq. 22) only has an explicit linear dependence on $J^{h,z,r}$, and the state variables are governed by their steady-state distribution from the disorder-averaged DMFT. These assumptions make the random matrix problem tractable, and we can evaluate the Green's function by using the self-consistent Born approximation, which is exact as $N \rightarrow \infty$ (details

are given in Appendix B; also see [56] for a similar approach in discrete-time RNNs). The result of this calculation is the spectral curve which is the boundary of the eigenvalue support in the complex plane. We note that the most general expression for the spectral curve (Appendix B), eq. B40) involves empirical averages (traces) over the $3N$ dimensional state variables. For large N , we can appeal to a concentration of measure argument to replace these discrete sums with averages over the steady-state distribution from the DMFT – i.e. we can replace $\frac{1}{N} \sum_i F(h_i, z_i, r_i)$ with $\langle F(h(t), z(t), r(t)) \rangle$, where the brackets indicate average over the steady-state distribution.

Thus, by combining the RMT with DMFT, we arrive at the following expression for the spectral boundary curve:

$$\left(\langle \sigma_r^2 \rangle + \frac{\langle \phi^2 \sigma_r'^2 \rangle}{|1 + \tau_r \lambda|^2} \right) \left\langle \frac{\phi'^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle + \frac{1}{|1 + \tau_z \lambda|^2} \left\langle \frac{\mathbb{D}^2 \phi'^2}{|\lambda + \sigma_z|^2} \right\rangle = 1. \quad (28)$$

Two major simplifications occur at fixed-points: first, $\mathbb{D} = 0$ and the entire second term is dropped, and second the random variables h and z become independent and Gaussian distributed (otherwise, the distribution of h will depend on the z -gate through $\sigma_z(z)$). As a result, the equation for the spectral curve at fixed-points is given by

$$\langle \phi'^2 \rangle \left(\langle \sigma_r^2 \rangle + \frac{\langle \phi^2 \rangle \langle \sigma_r'^2 \rangle}{|1 + \tau_r \lambda|^2} \right) \left\langle \frac{\sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle = 1, \quad (29)$$

As alluded to above, the averages in eq. 29 are taken with respect to the steady-state distribution given by the DMFT. For fixed-points, we can thus express eq. 29 in terms of Gaussian averages involving the variances $\Delta_h, \Delta_z, \Delta_r$ given in eq. 19 as

$$\begin{aligned} & \frac{\langle \phi(\sqrt{\Delta_h}x)^2 \rangle_x \langle \sigma_r'(\sqrt{\Delta_r}x)^2 \rangle_x + \langle \sigma_r(\sqrt{\Delta_r}x)^2 \rangle_x}{|1 + \tau_r \lambda|^2} \\ &= \frac{1}{\langle \phi'(\sqrt{\Delta_h}x)^2 \rangle_x} \left\langle \frac{(\sigma_z(\sqrt{\Delta_z}x))^2}{|\sigma_z(\sqrt{\Delta_z}x) + \lambda|^2} \right\rangle_x^{-1} \end{aligned} \quad (30)$$

The notation $\langle \cdot \rangle_x$ in eq. 30 denotes an average w.r.t the standard Gaussian measure on x . Fig. 2 [a-d] show that the RMT prediction of the spectral support (dark outline) agrees well with the numerically calculated spectrum (red dots) in different dynamical regimes. As a consequence of eq. 30, we get a condition for the stability of zero fixed-point. The leading edge of the spectral curve for the zero fixed-point will cross zero when

$$\text{zero-FP stability condition: } g_h < 1 + e^{-\beta r} \quad (31)$$

So, in the absence of biases, $g_h > 2$ will make the zero FP unstable. More generally, a FP will become unstable when the leading edge of the spectrum crosses zero, and the condition for this is given by

condition for FP to become unstable:

$$\begin{aligned} & \langle \phi(\sqrt{\Delta_h}x)^2 \rangle_x \langle \sigma_r'(\sqrt{\Delta_r}x)^2 \rangle_x + \langle \sigma_r(\sqrt{\Delta_r}x)^2 \rangle_x \\ & > \frac{1}{\langle \phi'(\sqrt{\Delta_h}x)^2 \rangle_x} \end{aligned} \quad (32)$$

We will see later on that the time-varying state corresponding to this regime is chaotic. We now proceed to analyze how the two gates shape the Jacobian spectrum.

B. Update gate facilitates slow modes and Output gate causes instability

To understand how each gate shapes the local dynamics, we study their effect on the density of Jacobian eigenvalues near zero and the shape of the spectral support curve – the former tells us about timescales present in the dynamics, and the latter about stability. For ease of exposition, we first consider how these quantities are shaped by the gates in the absence of biases ($\beta_{r,z,h} = 0$); we will consider the role of biases later on.

Fig. 2 shows how the gain parameters of the update and output gates – α_z and α_r respectively – shape the Jacobian spectrum. In Fig. 2 [a-d], we see that α_z has two salient effects on the spectrum: increasing α_z leads

to i) an accumulation of eigenvalues near zero; and ii) a *pinching* of the spectral curve for certain values of g_h wherein the intercept on the imaginary axis gets smaller (Fig. 2b; also see Sec. V). In Fig. 2 [a-d], we also see that increasing the value of α_r leads to an increase in the spectral radius, thus pushing the leading edge ($\max \text{Re} \lambda_i$) to the right and thereby facilitating instability.

The accumulation of eigenvalues near zero with increasing α_z implies a wide spectrum of timescales in the local dynamics since timescales are given by the inverse of the (real part of) the eigenvalues. To understand this accumulation quantitatively, it is helpful to consider the scenario where α_z is large and we replace the tanh activation functions with a piece-wise linear approximation. In this limit, the density of eigenvalues within a radius δ of the origin is well approximated by the following functional form (details in Appendix C):

$$P(|\lambda(\mathcal{D}_x)| < \delta) \sim c_0 \text{erf}\left(\frac{c_1}{\alpha_z}\right) \quad (33)$$

where c_0, c_1 are constants that, in general, depend on α_r, δ , and g_h . Fig. 2e, shows this scaling for a specific value of δ : the dashed line shows the predicted curve and the circles indicate the actual eigenvalue density calculated using the full Jacobian. In the limit of $\alpha_z \rightarrow \infty$ we get an extensive number of eigenvalues at zero, and the eigenvalue density converges to (see Appendix C):

$$\mu(\lambda) = (1 - f_z)\delta(\lambda) + f_z(1 - f_h)\delta(\lambda + 1) + \frac{4}{\pi g_h^2} \mathbb{I}_{\{|\lambda| \leq g_h^2/4\}} \quad (34)$$

where f_z is the fraction of update gates which are non-zero, and f_h is the fraction of unsaturated activation functions $\phi(h)$. For other choices of saturating nonlinearities, the extensive number of eigenvalues at zero remains; however, the expressions are more complicated.

A second interesting aspect of increasing α_z is the pinching of the spectral curve for certain values of g_h . As is evident from Fig. 2f, increasing the value of α_z point of intersection of the spectral curve with the imaginary axis. This pinching further accentuates the accumulation of eigenvalues near zero, and in the limit $\alpha_z \rightarrow \infty$ the leading edge of the spectrum can become pinched exactly at zero thus making an unstable system *marginally stable*. However, as shown in 2f (black line), pinching does not occur for all g_h ; when g_h is large enough, increasing α_z does not lead to pinching. We analyse the phenomenon of pinching and marginal stability below.

V. MARGINAL STABILITY AND LINE ATTRACTORS

We see an accumulation of Jacobian eigenvalues near zero and a pinching of the spectral curve near zero as the update becomes more switch-like (i.e. for larger values of α_z). In the limit $\alpha_z \rightarrow \infty$, this can convert previously

unstable points marginally stable by resting the leading edge of the spectral curve exactly at zero. In the limit $\alpha_z \rightarrow \infty$, the Jacobian spectrum undergoes a topological transition from a single simply connected domain to two domains, both containing an extensive number of eigenvalues. A finite fraction of eigenvalues will end up sitting exactly at zero, while the rest occupy a finite circular region: if the leading edge of the latter crosses zero, then the FP remains unstable. The separation of these domains is accomplished by a gradual *pinching* of the spectrum near zero. However, as seen in Fig. 2f, the pinching doesn't happen if g_h is sufficiently large. Here, we provide the conditions when pinching can occur and thus marginal stability can emerge. We then show how this phenomenon can provide a mechanism for computations using line attractors.

A. Conditions for marginal stability

Let us again start with the case where $\alpha_r = 0$ and there are no biases. To show the emergence of marginal stability, we can focus on the FP spectrum and show how it evolves with increasing α_z . In the limit of large α_z , it can be shown that the leading edge of the spectrum scales as (see Appendix C)

$$\lambda_e \sim e^{-c\alpha_z\sqrt{\Delta_h}} \quad \text{where } c \text{ satisfies} \quad \text{erfc}\left(\frac{c}{\sqrt{2}}\right) = \frac{8}{\langle\phi'^2\rangle}, \quad (35)$$

where $\langle\phi'^2\rangle \equiv \int Dx\phi'(\sqrt{\Delta_h}x)^2$, i.e. the average is taken over the Gaussian fixed point distribution. As long as eq. 35 has a positive solution for c , we expect pinching of the spectral edge with increasing α_z , and in the limit $\alpha_z \rightarrow \infty$, the spectral edge will be poised at zero making the unstable points marginally stable. Thus, marginal stability is possible for all parameters for which

$$\langle\phi'^2\rangle < 8 \quad (36)$$

It's easy to see from eq. 29 that unstable points satisfy $\langle\phi'^2\rangle > 4$, so we see that not all unstable points can be made marginally stable; we require eq. 36 as well. More generally, for finite α_r , in the limit $\alpha_z \rightarrow \infty$, σ_z will become binary and the condition for marginal stability is:

$$\langle\phi'^2\rangle \left(\langle\phi^2\rangle \langle\sigma_r'^2\rangle + \langle\sigma_r^2\rangle \right) < \frac{1}{\langle\sigma_z\rangle} \quad (37)$$

We will return to this when we describe the phase diagram for the gated RNN (Sec. IX). There we will see that the marginally stable region occupies a volume in the parameter space adjoining the critical lines on one side. Next, we turn to implications of pinching and clumping for implementing computations with the dynamics.




FIG. 3. *Gating produces line attractors* Top row: sample traces from a network with switch-like update gates ($\alpha_z = 30$) show slow evolution (time on x -axis is relative to τ_h). Bottom row: An input is applied from $t = 0$ till $t = 10$, either aligned with a slow eigenvector \mathbf{u}_λ (red traces) or unaligned with slow modes (black dashed trace). Plot shows the excess projection of the network state on the left eigenvector \mathbf{u}_λ . Different shades of red correspond to different input strengths. ($N = 2000, g_h = 3.0, \alpha_r = 0.0$)

B. Line attractors

Spectral pinching and clumping of eigenvalues near zero lead to a significant reduction in the number of unstable directions, and a manifold of slow directions emerges. In particular, for large α_z , any perturbation in the span of the eigenvectors corresponding to the eigenvalues with magnitude close to zero will be integrated by the network, and once the input perturbation ceases the memory trace of the input will be retained for a duration much longer than the intrinsic time-constant of the neurons; perturbations along other eigenvectors, however, will relax with a timescale dictated by the inverse of (the real part of) their eigenvalues. Thus the manifold of slow directions will form an approximate continuous attractor on which input can effortlessly move the state vector around. These approximate continuous attractor dynamics are illustrated in Fig. 3. At time $t = 0$, an input is applied for 10 seconds along an eigenvector of the Jacobian with an eigenvalue close to zero. Inputs along the slow manifold with varying strengths (different shades of red) are integrated by the network as evidenced by the excess projection of the network activity on the left-eigenvector (\mathbf{u}_λ) corresponding to the slow mode; on the

other hand, inputs not aligned with the slow modes decay away quickly (dashed black line). Recall that the intrinsic time-constant of the neurons here is set to 1 second. The picture that emerges is the following: there is a manifold spanned by the slowest eigenvectors of the Jacobian where the flow is almost stationary, and any input lying within the span of these eigenvectors will effortlessly move the activity along the direction of the input. Once the input ceases, the memory will be retained for a long time; on the other hand, components of the input that are not in the span of the slowest modes will be integrated in a leaky fashion, decaying with a spectrum of timescales, and the perturbed activity will flow back to the slow manifold spanned by the slow modes. Of course, this is a local picture and network dynamics will slowly drift, so the slow manifold will evolve; however, as we can see from Fig. 3, this picture is accurate over a reasonably long time.

VI. LONG-TIME DYNAMICS AND TRANSITION TO CHAOS

We have thus far studied how the gates shape the local dynamics by looking at their effect on the Jacobian. The nature of these effects could change when we consider longer timescales. Moreover, we would like to know what is the nature of the time-varying state when the dynamics are locally unstable – in particular, is it chaotic? If so, are there many unstable directions? To address these questions, we characterize the asymptotic dynamics of the network using Lyapunov exponents. We study how the gates shape the numerically calculated Lyapunov spectrum of the network, and derive an analytical prediction for the maximum Lyapunov exponent using two different approaches: one based on the DMFT and the other using RMT. The RMT based approach is novel and gives a relation between the Lyapunov exponent and the relaxation time.

A. How gates shape the Lyapunov spectrum

To characterize the asymptotic behavior of the network, we study how infinitesimal perturbations grow or shrink along the trajectories of the dynamics. The evolution of perturbations $\delta\mathbf{x}(t)$ along a trajectory follow the tangent-space dynamics governed by the Jacobian

$$\partial_t \delta\mathbf{x}(t) = \mathcal{D}(t) \delta\mathbf{x}(t) \quad (38)$$

So after a time T , the initial perturbation $\delta\mathbf{x}(0)$ will be given by

$$\delta\mathbf{x}(T) = \left[e^{\int_0^T dt \mathcal{D}(t)} \right] \delta\mathbf{x}(0) \quad (39)$$

When the infinitesimal perturbations grow/shrink exponentially, the rate of this exponential growth/decay will

be dictated by the maximal Lyapunov exponent defined as [57]:

$$\lambda_{max} := \lim_{T \rightarrow \infty} \frac{1}{T} \lim_{\|\delta\mathbf{x}(0)\| \rightarrow 0} \ln \frac{\|\delta\mathbf{x}(T)\|}{\|\delta\mathbf{x}(0)\|} \quad (40)$$

For ergodic systems, this limit is independent of almost all initial conditions, as guaranteed by Oseledec's multiplicative ergodic theorem [57]. Positive values of λ_{max} imply that the nearby trajectories diverge exponentially fast, and the system is chaotic. More generally, the set of all Lyapunov exponents – the Lyapunov spectrum – yields the rates at which perturbations along different directions shrink or diverge, and thus provide a fuller characterization of asymptotic behaviour. The first k ordered Lyapunov exponents are given by the growth rates of k linearly independent perturbations. These can be obtained as the logarithms of the eigenvalues of the Oseledec's matrix, defined as [57]

$$\mathbf{M}(t) = \lim_{t \rightarrow \infty} \left(\left[e^{\int_0^t dt' \mathcal{D}(t')} \right]^T \left[e^{\int_0^t dt' \mathcal{D}(t')} \right] \right)^{\frac{1}{2t}} \quad (41)$$

However, this expression cannot be directly used to calculate the Lyapunov spectra in practice since $\mathbf{M}(t)$ rapidly becomes ill-conditioned. We instead employ a method suggested by [58] (also c.f. [59] for Lyapunov spectra of RNNs). We start with k orthogonal vectors $Q^0 = [q_1, \dots, q_k]$ and evolve them using the tangent-space dynamics, eq. 38, for a short time interval t_0 . Therefore, the new set of vectors are given by

$$\widehat{Q} = \left[e^{\int_0^{t_0} dt' \mathcal{D}(t')} \right] Q^0 \quad (42)$$

We now decompose $\widehat{Q} = Q^1 R^1$ using a QR -decomposition, into an orthonormal matrix Q^1 and a upper-diagonal matrix R^1 with positive diagonal elements, which give the rate of shrinkage/expansion of the volume element along the different directions. We iterate this procedure for a long time, $t_0 \times N_l$, and the first k ordered Lyapunov exponents are given by

$$\lambda_i = \lim_{N_l \rightarrow \infty} \frac{1}{N_l t_0} \sum_{j=1}^{N_l} \ln R_{ii}^j \quad i \in \{1, \dots, k\} \quad (43)$$

Fig. 4a,b shows how the update (z -) and output (r -) gates shape the Lyapunov spectrum. We see that as the update gets more sensitive (larger α_z), the Lyapunov spectrum flattens pushing more exponents closer to zero. Moreover, a more sensitive output gate (larger α_r) makes the Lyapunov exponents larger thus increasing the rate of growth in unstable directions. We note that these effects are largely consistent with the results on how the gates shape the local dynamics. However, a prediction of the Lyapunov exponents based purely on the local estimates of the Jacobian eigenvalues will be incorrect due to potentially large correlations between the Jacobian matrices across time.

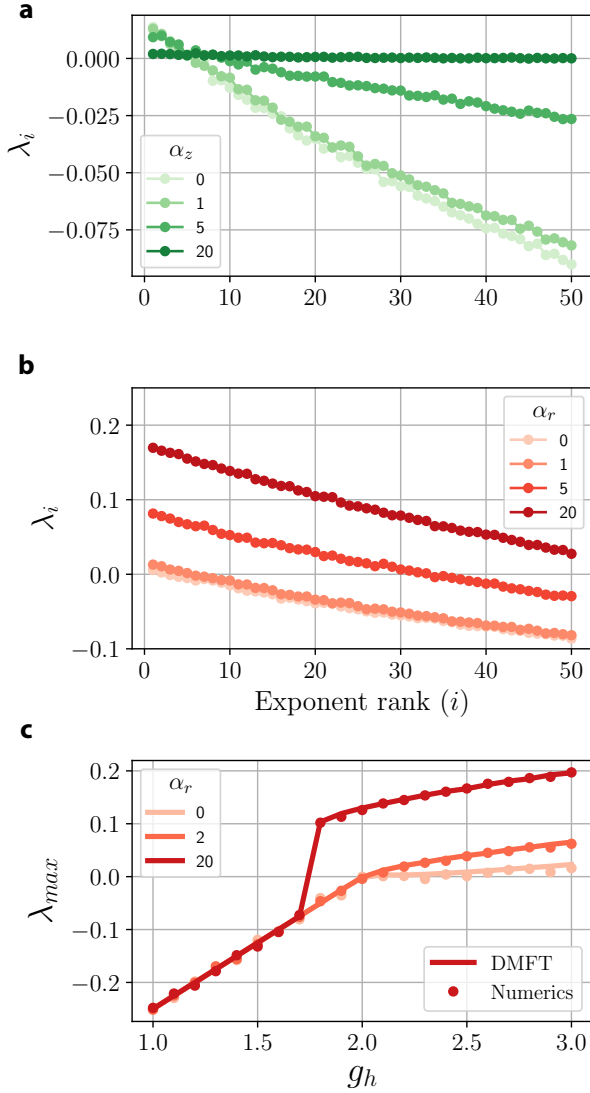


FIG. 4. *Asymptotic behavior of the gated network* (a,b): The first 50 ordered Lyapunov exponents for a gated RNN ($N = 2000$) as a function of varying (a) α_z and (b) α_r . The Lyapunov spectrum was calculated using the *QR*-method described in the main text with total simulation time $T = 1000s$. (c) The maximal Lyapunov exponent λ_{max} predicted by the DMFT (solving eqns. 49-51; solid line) and obtained numerically using the *QR*-method (circles; $N = 2000, \alpha_z = 0$). Note that the transition for $\alpha_r = 20$ is sharp, and the sloped line is a result of discrete sampling of g_h ; also c.f. Fig. 6c. $\tau_z = \tau_r = 2.0$ here.

B. DMFT prediction for the maximal Lyapunov exponent

In addition to shaping of the Lyapunov spectrum, we would like to study the chaotic nature of the time-varying phase by means of the maximal Lyapunov exponent, and characterize when the transition to chaos occurs. We extend the DMFT for the gated RNN to calculate the max-

imum Lyapunov exponent, and to do this, we make use of a technique suggested by [53, 60] and clearly elucidated in [22]. The starting point of the method is two replicas of the system $\mathbf{x}^1(t)$ and $\mathbf{x}^2(t)$ with the same coupling matrices $J^{h,z,r}$ and the same parameters. If the two systems are started with initial conditions which are close, then the rate of convergence/divergence of the trajectories reveals the maximal Lyapunov exponent. To this end, let's define $d(t, s) := N^{-1} \sum (x_i^1(t) - x_i^2(s))^2$, and study the growth rate of $d(t, t)$. In the large N limit, we expect population averages like $C^{12}(t, s) := N^{-1} \sum x_i^1(t) x_i^2(s)$ to be self-averaging (like in the DMFT for a single system) [61], and thus we can write

$$d(t, s) = C^{11}(t, t) + C^{22}(s, s) - C^{12}(t, s) - C^{21}(t, s) \quad (44)$$

For trajectories that start nearby, the asymptotic growth rate of $d(t)$ is the maximal Lyapunov exponent. In order to calculate this using the DMFT, we need a way to calculate C^{12} – the correlation *between* replicas – for a typical instantiation of systems in the large N limit. As suggested by [22], this can be achieved by considering a joint generating functional for the replicated system:

$$\tilde{Z}_{\mathcal{J}}[\hat{\mathbf{b}}^1, \hat{\mathbf{b}}^2, \mathbf{b}^1, \mathbf{b}^2] = \mathbb{E} \left[\exp \left(i \sum_{\mu=1}^2 \sum_{j=1}^N \int \hat{\mathbf{b}}_j^\mu(t)^T \mathbf{x}_j^\mu(t) dt \right) \right] \quad (45)$$

We then proceed to take the disorder average of this generating functional – in much the same way as a single system – and this introduces correlations between the state vectors of the two replicas. A saddle-point approximation as in the single system case (c.f. Appendix A), yields a system of coupled SDEs (one for each replica), similar to eq. 11, but now the noise processes in the two replicas are coupled, so that terms like $\langle \eta_h^1(t) \eta_h^2(t') \rangle$ need to be considered. As before, the SDEs imply the equations of motion for the correlation functions

$$[-\partial_\tau^2 + C_{\sigma_z}^{\mu\nu}(\tau)] C_h^{\mu\nu}(\tau) = C_{\sigma_z}^{\mu\nu}(\tau) C_\phi^{\mu\nu}(\tau) C_{\sigma_r}^{\mu\nu}(\tau) \quad (46)$$

$$[-\tau_z^2 \partial_\tau^2 + 1] C_z^{\mu\nu}(\tau) = C_\phi^{\mu\nu}(\tau) \quad (47)$$

$$[-\tau_r^2 \partial_\tau^2 + 1] C_r(\tau) = C_\phi^{\mu\nu}(\tau) \quad (48)$$

where $\mu, \nu \in \{1, 2\}$ are the replica indices. Note that the single-replica solution will clearly be a solution to this system, reflecting the fact that marginal statistics of each replica is the same as before. When the replicas are started with initial conditions that are ϵ -close, we expect the inter-replica correlation function to diverge from the single replica steady-state solution, so we expand C^{12} to linear order as $C_{h,z,r}^{12}(t, s) \approx C_{h,z,r}(t-s) + \epsilon \tilde{\chi}_{h,z,r}(t, s)$. From eq. 44 we see that $d(t, t) \sim \epsilon \tilde{\chi}(t, t)$, and thus the growth rate of $\tilde{\chi}$ will yield the required Lyapunov exponent. To this end we make an ansatz $\tilde{\chi}_{h,z,r} = e^{\kappa T} \chi(\tau)$ where $2T = t + s$ and $2\tau = t - s$, and κ is the DMFT prediction of the maximum Lyapunov exponent that needs to be solved for. Substituting this back in eq. 46, we get a generalized eigenvalue problem for κ :

$$\begin{aligned} & [(\langle \sigma_z \rangle + \kappa)^2 - \partial_\tau^2 + C_{\sigma_z}(\tau) - \langle \sigma_z \rangle^2] \chi_h(\tau) = \\ & C_{\sigma'_z}(\tau) [C_{\phi \cdot \sigma_r}(\tau) - C_h(\tau)] \chi_z(\tau) \\ & + C_{\sigma_z}(\tau) \frac{\partial C_{\phi \cdot \sigma_r}(\tau)}{\partial C_h} \chi_h(\tau), \end{aligned} \quad (49)$$

$$[(1 + \tau_z \kappa)^2 - \tau_z^2 \partial_\tau^2] \chi_z(\tau) = \frac{\partial C'_\phi(\tau)}{\partial C_h} \chi_h(\tau), \quad (50)$$

$$[(1 + \tau_r \kappa)^2 - \tau_r^2 \partial_\tau^2] \chi_r(\tau) = \frac{\partial C_\phi(\tau)}{\partial C_h} \chi_h(\tau). \quad (51)$$

The largest eigenvalue solution to this problem is the required maximal Lyapunov exponent. Note, this is the analogue of the Schrödinger equation for the maximal Lyapunov exponent in the vanilla RNN. In this case, when $\alpha_z = 0$ (or small), the h field is Gaussian and we can use Price's theorem for Gaussian integrals to replace the variational derivatives on the r.h.s of eq. 49-51 by simple correlation functions. For instance:

$$\frac{\partial C_\phi(\tau)}{\partial C_h(\tau)} = C_{\phi'}(\tau) \quad (52)$$

In this limit, we see good agreement between the numerically calculated maximal Lyapunov exponent (Fig. 4c dots) compared to the DMFT prediction (Fig. 4c solid line) obtained by solving the eigenvalue problem (eq. 49-51). For large values of α_z we see quantitative deviations between the DMFT prediction and the true λ_{max} .

C. RMT prediction for λ_{max} and its relation to relaxation time

We now provide an alternate expression for the maximal Lyapunov exponent, λ_{max} , which relates it to the relaxation time of the dynamics. We first derive the result for the classical RNN, and simply state the result for the gated RNN while leaving the details to Appendix G. Oseledet's multiplicative ergodic theorem guarantees that

$$\lambda_{max} = \lim_{t \rightarrow \infty} \frac{1}{2t} \log \frac{\|\chi(t)\|^2}{N}, \quad (53)$$

$$= \lim_{t \rightarrow \infty} \frac{1}{2t} \log \frac{1}{N} \text{Tr} [\chi(t) \chi(t)^T], \quad (54)$$

where $\chi(t) = e^{\int_0^t dt' \mathcal{D}(t')}$ and \mathcal{D} is the Jacobian. For the vanilla RNN, the Jacobian is given by

$$\mathcal{D} = -\mathbb{1} + J[\phi'(t)]. \quad (55)$$

We expect the maximal Lyapunov exponent to be independent for the random network realization, and thus equal to its value after disorder-averaging. Defining the diagonal matrix $R(t) = \int^t [\phi'(t')] dt'$, this assumption

gives

$$\frac{1}{N} \text{Tr} [\chi(t) \chi(t)^T] \approx e^{-2t} \left\langle \frac{1}{N} \text{Tr} e^{JR(t)} e^{R(t)J^T} \right\rangle, \quad (56)$$

$$= e^{-2t} \sum_{n=0}^{\infty} \frac{1}{(n!)^2} \left(\frac{1}{N} \text{Tr} R(t)^2 \right)^n, \quad (57)$$

where the second line in eq. 56 follows after disorder averaging over J and keeping only terms to leading order in N . Next, we may apply the DMFT to write

$$\frac{1}{N} \text{Tr} R(t)^2 = \int^t dt' dt'' \frac{1}{N} \sum_{i=1}^N \phi'_i(t'') \phi'_i(t'), \quad (58)$$

$$\approx \int dt' dt'' C_{\phi'}(t', t''). \quad (59)$$

In steady-state, the correlation function depends only on the difference of the two times, and thus we can write

$$\int dt' dt'' C_{\phi'}(t', t'') \approx \int_0^{2t} \frac{du}{2} \int_0^t d\tau C_{\phi'}(\tau) \equiv t^2 \tau_R, \quad (60)$$

where we have defined the relaxation time for the $C_{\phi'}$ correlation function

$$\tau_R \equiv \frac{1}{t} \int_0^t d\tau C_{\phi'}(\tau). \quad (61)$$

Substituting eq. 60 in eq. 56 we get

$$\frac{1}{N} \text{Tr} [\chi(t) \chi(t)^T] = e^{-2t} I_0(2t\sqrt{\tau_R}), \quad (62)$$

which for long times behaves like $\exp(2(\sqrt{\tau_R} - 1)t)$. Inserting this into eq. 53 gives us the RMT prediction for the maximal Lyapunov exponent for the vanilla RNN

$$\lambda_{max} = \sqrt{\tau_R} - 1. \quad (63)$$

This formula relates the asymptotic Lyapunov exponent to relaxation time of a local correlation function in steady state.

Now, we state the analogous result for the gated RNN with only the r -gate and relegate details to Appendix G. The maximal Lyapunov exponent for the gated RNN is given by

$$\lambda_{max} = \left(\frac{\tau_R + \sqrt{\tau_R^2 + 4\tau_A\tau_Q}}{2} \right)^{1/2} - 1, \quad (64)$$

where the relaxation times τ_A, τ_R, τ_Q are defined as

$$\tau_R = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau C_{\phi'}(\tau) C_{\sigma_r}(\tau), \quad (65)$$

$$\tau_A = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau C_{\phi'}(\tau), \quad (66)$$

$$\tau_Q = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau C_\phi(\tau) C_{\sigma'_r}(\tau). \quad (67)$$

D. Condition for continuous transition to chaos

The value of α_z affects the precise value of the maximal Lyapunov exponent λ_{max} ; however, numerics suggest that the transition to chaos – the point at which λ_{max} becomes positive – is not dependent on α_z (data not shown). This makes intuitive sense since the update gate acts as a dynamical time-constant, and this should not make a stable system chaotic. We can see this more clearly by calculating the transition to chaos when the leading edge of the spectral curve (for a FP) crosses zero. This condition is given by eq. 32, and we see that it has no dependence on α_z or the update gate.

We note that this condition, eq. 32, for the transition to chaos – when the stable fixed-point becomes unstable – is valid when the chaotic attractor emerges continuously from the fixed-point (Fig. 4c, $\alpha_r = 0, 2$). However, in the gated RNN, there is another another interesting – and discontinuous – transition to chaos (Fig.4c, $\alpha_r = 20$): for sufficiently large α_r , the transition to chaos is discontinuous and occurs at a value of g_h where the zero FP is still stable ($g_h < 2$ with no biases). This is a novel type of transition which is not present in the vanilla RNN, and we characterize this phenomenon in detail below.

VII. A NOVEL, DISCONTINUOUS TRANSITION TO CHAOS

In this section, we describe in detail a novel state, characterized by a proliferation of unstable fixed-points, and the coexistence of a stable fixed-point with a chaotic pseudo-attractor. It is the appearance of this state that gives rise to the discontinuous transition observed in Fig. 4c. The appearance of this state is mediated by the output gate becoming more switch-like (i.e. increasing α_r) in the quiescent region for g_h , and to our knowledge no such comparable phenomenon exists in RNNs with additive interactions. This phenomenon has a few salient aspects: i) we see the spontaneous appearance of unstable fixed-points with increasing α_r , for certain values of g_h in the quiescent region; this transition has no obvious dynamical signatures; ii) if we further increase α_r beyond a second threshold, we observe a dynamical transition characterized by a state where the stable fixed-point coexists with a chaotic dynamics; this state is characterized by long, chaotic transients whose lifetimes scale with system size. This gap between the appearance of fixed-points and the chaotic attractor differs from the result of Wainrib & Touboul [38] for purely additive RNNs, where the proliferation of unstable fixed-points (topological complexity) is tightly linked to the chaotic dynamics (dynamical complexity); and iii) finally, on increasing g_h for large but fixed α_r , the stable fixed-point disappears and the state smoothly transitions into a full chaotic attractor that was characterized above (for $g_h > 2.0$ with no biases). We elaborate on these points below. For ease of presentation, the rest of the section will assume that

all biases are zero.

A. Spontaneous emergence of fixed-points

For $g_h < 2.0$ and small α_r , the zero fixed-point is the globally stable state for the dynamics and the only solution to the fixed-point equations, eq. 19, for Δ_h . However, as we increase α_r for a fixed g_h , two additional non-zero solutions to Δ_h spontaneously appear at a critical value $\alpha_{r,FP}^*(g_h)$ as shown in Fig. 5a. The appearance of these solutions corresponds to a proliferation of unstable fixed-points in the phase space. However, the zero fixed-point continues to be the globally stable state for the dynamics for α_r near $\alpha_{r,FP}^*(g_h)$.

Numerical solutions to the fixed-point equations reveal the form of the bifurcation curve $\alpha_{r,FP}^*(g_h)$ and the associated value of $\Delta_h^*(g_h)$. We see that $\alpha_{r,FP}^*(g_h)$ increases rapidly with decreasing g_h , dividing the parameter space into regions with either 1 or 3 solutions for Δ_h . The corresponding $\Delta_h^*(g_h)$ vanishes at two boundary values of g_h – one at 2.0 and another, g_c , below 1.5 where $\alpha_r^* \rightarrow \infty$. This naturally leads to the question of whether the fixed-point bifurcation exists for all values of g_h below 2.0.

To answer this, we perturbatively solve the fixed-point equations in two asymptotic regimes: i) $g_h \rightarrow 2^-$ and ii) $g_h \rightarrow g_c^+$. Details of the perturbative treatment are in Appendix D 2. For $g_h = 2 - \epsilon$, we see that the perturbative problem undergoes a bifurcation from one solution ($\Delta_h = 0$) to three when α_r crosses the bifurcation threshold $\alpha_r^*(2.0) = \sqrt{8}$, and this is the left limit of the bifurcation curve in Fig. 5b. The larger non-zero solution for the variance at the bifurcation point scales as

$$\Delta_h^* \approx (\alpha_r^2 - 8) \cdot \xi_0 + \xi_1 \epsilon \quad \text{for } \alpha_r \rightarrow \alpha_{r,FP}^*(2) = \sqrt{8}, \quad (68)$$

where ξ_0 and ξ_1 are positive constants (see Appendix D 2).

At the other extreme, to determine the smallest value of g_h for which a bifurcation is possible, we note from Fig. 5b that in this limit $\alpha_r \rightarrow \infty$, and thus we can look for solutions to Δ_h in the limit: $\Delta_h \ll 1$ and $\alpha_r \rightarrow \infty$ and $\alpha_r \sqrt{\Delta_h} \gg 1$. In this limit, there is a bifurcation in the perturbative solution when $g_h > g_h^* = \sqrt{2}$, and close to the critical point, the fixed-point solution is given by (see Appendix D 2):

$$\Delta_h^*(\sqrt{2}^+) \sim \frac{g_h^2 - 2}{2g_h^4} \quad \text{for } g_h \rightarrow \sqrt{2}^+. \quad (69)$$

Thus in the region $g_h \in (\sqrt{2}, 2)$ there exist non-zero solutions to the fixed-point equations once α_r is above a critical value $\alpha_r^*(g_h)$. These solutions correspond to unstable fixed-points appearing in the phase space. Interestingly, as discussed in detail below, there are no obvious dynamical signatures of this transition in the transients of the network, and it is only when α_r crosses another critical value do we see a change in dynamics. The zero fixed-point is still the global attractor of the dynamics.

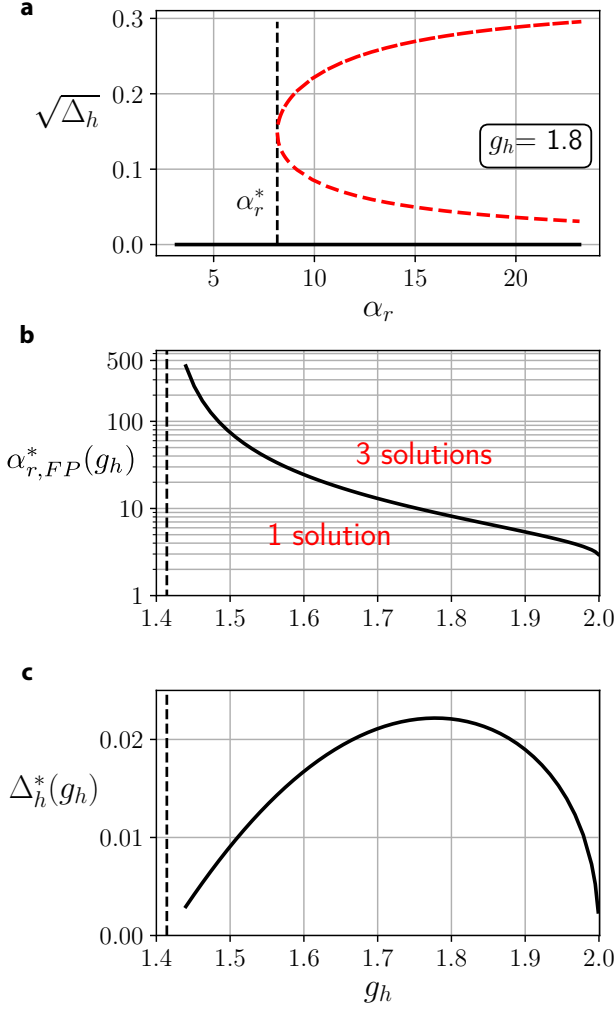


FIG. 5. *Bifurcation transition for fixed-point solutions* (a) spontaneous appearance of non-zero solutions (dashed red lines) to the FP equations with increasing α_r at fixed g_h . Both the FP solutions are unstable with respect to the dynamics. For a given g_h , there is a critical $\alpha_{r,FP}^*(g_h)$ at which the non-zero solution first appears. (b) The critical $\alpha_{r,FP}^*(g_h)$ as a function of g_h dividing the parameter space into regions with just one FP solution (the zero FP) or 3 FP solutions. The dashed line represents left critical value $g_c = \sqrt{2}$ below which a bifurcation is not possible. (c) The corresponding solution to the FP equations, $\Delta_h^*(g_h)$, for $\alpha_r = \alpha_{r,FP}^*(g_h)$.

B. Delayed dynamical transition shows a decoupling between topological and dynamical complexity

The picture from the fixed-point transition above is that when g_h is in the interval $(\sqrt{2}, 2)$, there is a proliferation of unstable fixed-points in the phase space provided $\alpha_r > \alpha_{r,FP}^*(g_h)$. However, it turns out that the spontaneous appearance of these unstable fixed-points is not accompanied by any dynamical signatures – as measured by the Lyapunov exponents (see Fig. 6) or by the transient

times (see Fig. 7). It is only when α_r is increased further beyond a second critical value $\alpha_{r,DMFT}^*(g_h)$, that we see the appearance of chaotic and long-lived transients. This is significant in regard to a result by Wainrib & Touboul [38], where they showed that the transition to chaotic dynamics (dynamical complexity) in random RNNs is tightly linked to the proliferation of critical points (topological complexity), and in their case, the exponential rate of growth of critical points (a topological property) was the same as the maximal Lyapunov exponent (a dynamical property). Here, we see that the topological and dynamical complexity are decoupled with an intermediate regime where critical points proliferate without affecting global dynamics significantly. Thus, for gated RNNs, there appears to be another distinct mechanism for the transition to chaos, and the accompanying transition is a discontinuous one.

Before we elaborate on the nature of the time-varying activity across the dynamical transition, let us characterize the second dynamical transition curve given by $\alpha_{r,DMFT}^*(g_h)$ (Fig. 6a, red curve). For ease of discussion, we turn off the update gate ($\alpha_z = 0$) and introduce a functional F_ψ for a 2-D Gaussian average of a given function $\psi(x)$

$$F_\psi(C_h(0), C_h(\tau)) = \mathbb{E}[\psi(z_1)\psi(z_2)], \quad (70)$$

$$\text{where } \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbb{C}) \quad \mathbb{C} = \begin{pmatrix} C_h(0) & C_h(\tau) \\ C_h(\tau) & C_h(0) \end{pmatrix}. \quad (71)$$

The DMFT equations for the correlation functions then become

$$\begin{aligned} C_h(\tau) - 4\partial_\tau^2 C_h(\tau) &= F_\phi(C_h(0), C_h(\tau))F_{\sigma_r}(C_r(0), C_r(\tau)), \\ C_r(\tau) - \tau_r^2 \partial_\tau^2 C_r(\tau) &= F_\phi(C_\phi(0), C_\phi(\tau)). \end{aligned} \quad (72)$$

We further make an approximation that $\tau_r \ll 1$, which in turn implies $C_r(\tau) \approx C_\phi(\tau)$. This approximation turns out to hold even for moderately large τ_r as will be seen from the numerics in Fig. 6c. With these approximations, we can integrate the equations for $C_h(\tau)$ to arrive at an equation for the variance $C_h^0 \equiv C_h(0)$ by solving the following equation for C_h^0 (details in Appendix E):

$$C_h(0)^2 + 2 \int_0^{C_h^0} dC_h F_\phi(C_h, C_h^0) F_{\sigma_r}(C_\phi, C_\phi^0) = 0. \quad (73)$$

Solving this equation will give the DMFT prediction for the variance for any g_h and α_r . Beyond the critical value of α_r , two non-zero solutions for C_h^0 spontaneously emerge. In order to use eq. 73 to find a prediction for the DMFT bifurcation curve $\alpha_{r,DMFT}^*(g_h)$, we need to use the additional fact that at the bifurcation point the two solutions coincide, and there is only one non-zero solution. To proceed, we can view the L.H.S of eq. 73, as a function of α_r , g_h and C_h^0 : $\mathcal{F}(g_h, \alpha_r, C_h^0)$. Then the equation for the bifurcation curve is obtained by solving

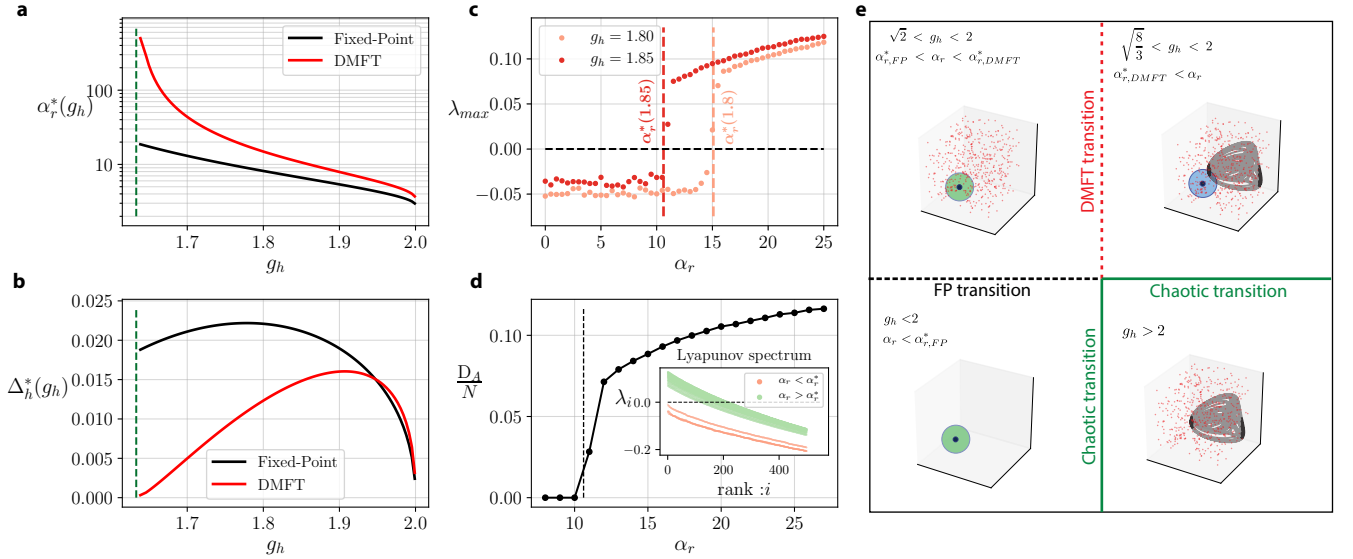


FIG. 6. *The discontinuous dynamical transition* (a) The critical DMFT transition curve $\alpha_{r,DMFT}^*(g_h)$ (red curve) calculated using eqns. 74, 75. The FP transition curve is shown in black. Green dashed line corresponds to $g_c = \sqrt{8/3}$ below which the dynamical transition is not possible. (b) The corresponding DMFT solution to the variance (eq. 73) for $\alpha_r = \alpha_{r,DMFT}^*(g_h)$. (c) Numerically calculated maximum Lyapunov exponent λ_{max} as a function of α_r for two different values of g_h . The dashed lines correspond to the DMFT prediction for the discontinuous transition from (a). (d) The upper bound on the dimension of the chaotic attractor D_A according to the Kaplan-Yorke conjecture (eq. 78) as a function of α_r for $g_h = 1.85$. (Inset) The first 450 ordered Lyapunov exponents for $\alpha_r > \alpha_{r,DMFT}^*(g_h)$ (green) and $\alpha_r < \alpha_{r,DMFT}^*(g_h)$ (red). $N = 2000$; $\alpha_z = 0$ for the numerical simulations. (e) Schematic of the bifurcation transition: for $g_h < 2$ and $\alpha_r < \alpha_{r,FP}^*$ the zero FP is stable and is the only solution to the FP equations (bottom left box); for $\sqrt{2} < g_h < 2$ and $\alpha_{r,FP}^* < \alpha_r < \alpha_{r,DMFT}^*$ the zero FP is still stable, but there is a proliferation of unstable FPs without any clear dynamical signature (top left); for $\sqrt{8/3} < g_h < 2$ and $\alpha_r > \alpha_{r,DMFT}^*$ chaotic dynamics coexist with the stable FP and this transition is discontinuous (top right); finally for $g_h > 2.0$ the stable FP becomes unstable, and only the chaotic attractor remains; this transition is continuous (bottom right).

the following two equations for $C_h^{0,*}$ and α_r^*

$$\mathcal{F}(g_h, \alpha_r^*, C_h^{0,*}) = 0, \quad (74)$$

$$\left. \frac{\partial \mathcal{F}(g_h, \alpha_r, C_h^0)}{\partial C_h^0} \right|_{\alpha_r^*, C_h^{0,*}} = 0. \quad (75)$$

The details of solving these equations are provided in Appendix E.

Fig. 6a,b shows the DMFT prediction for the bifurcation curve $\alpha_{r,DMFT}^*(g_h)$, and the corresponding variance at the bifurcation point $C_h^{0,*}$ (red curves). We note two salient points: i) the DMFT bifurcation curve is always above the fixed-point bifurcation curve (black, in Fig. 6a) and ii) the lower critical value of g_h which permits a dynamical transition (dashed green curve in Fig. 6a,b) is smaller than the corresponding fixed-point critical value of $\sqrt{2}$. Indeed, in Appendix E, we show the following condition for the dynamical transition

$$\text{condition for dynamical transition:} \\ \sqrt{\frac{8}{3}} < g_h \leq 2 \quad \text{and} \quad \alpha_r > \alpha_{r,DMFT}^*(g_h), \quad (76)$$

and at the right limit $g_h \rightarrow 2^-$, we find

$$\sqrt{12} = \alpha_{r,DMFT}^*(2) > \alpha_{r,FP}^*(2) = \sqrt{8}. \quad (77)$$

The DMFT prediction for the dynamical bifurcation agrees well with the full network simulations. In Fig. 6c we see that the maximum Lyapunov exponent experiences a discontinuous transition from a negative value (network activity decays to fixed-point) to a positive value (activity is chaotic) at the critical value of α_r predicted by the DMFT (dashed vertical lines).

C. Nature of the dynamic transition

We now address the question of what happens to the network activity across the dynamic transition. We already alluded to the discontinuous appearance of a time-varying state in Fig. 6c, where the maximum Lyapunov exponent jumps from a negative value to a positive value. This is indeed the case. A dynamical state spontaneously emerges across the transition, which has the characteristics a chaotic attractor with a punctured hole (a stable FP) that is a sink. The presence of the stable FP means that the dynamical state is not strictly a chaotic attractor but rather a ‘‘chaotic set’’; however, when we increase

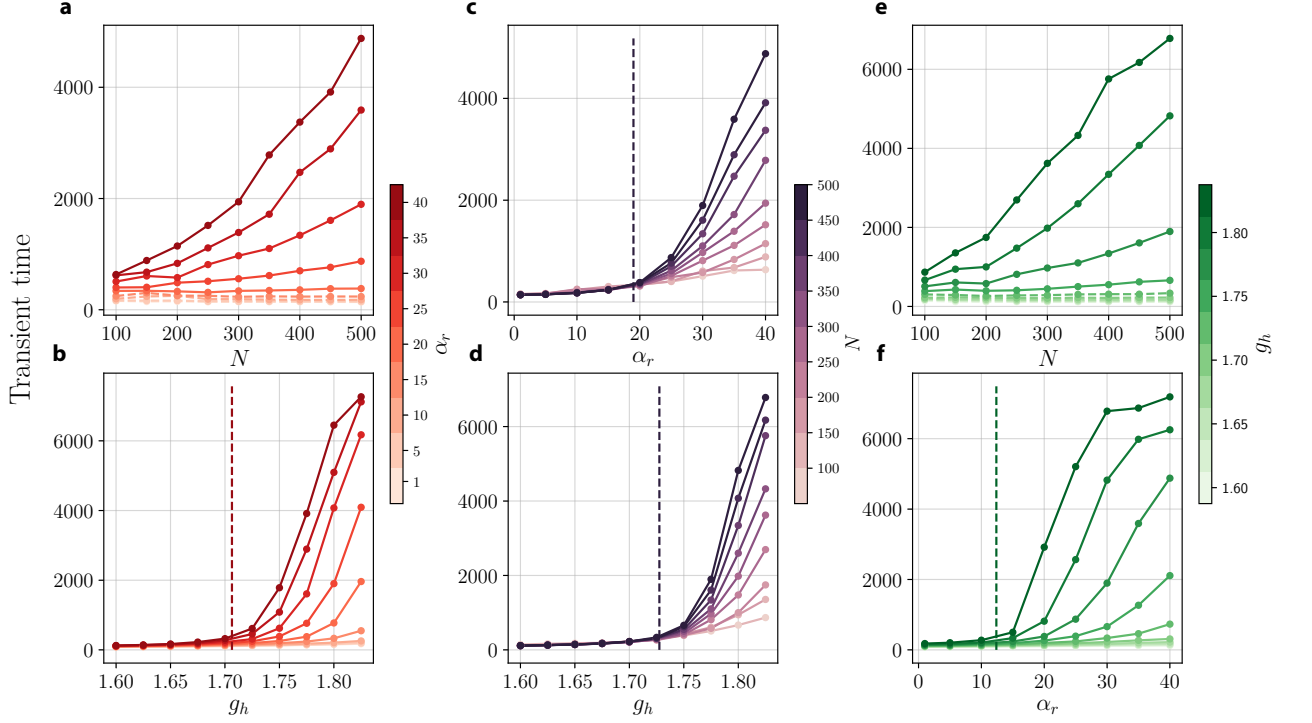


FIG. 7. *Transient times at the bifurcation transition* Transient times (τ_T , relative to τ_h) as a function of g_h , α_r and system size N . Dashed plot lines correspond to situations where $\alpha_r < \alpha_{r,DMFT}^*(g_h)$. Dashed vertical lines are critical values of α_r or g_h (a) τ_T vs. N for $g_h = 1.775$ (b) τ_T vs. g_h for $N = 500$; dashed line indicated g_h such that $\alpha_{r,DMFT}^*(g_h) = 40$. (c) τ_T vs. α_r for $g_h = 1.775$; dashed vertical line is $\alpha_{r,DMFT}^*(1.775)$. (d) τ_T vs. g_h for $\alpha_r = 30$; dashed line is g_h such that $\alpha_{r,DMFT}^*(g_h) = 30$. (e) τ_T vs. N for $\alpha_r = 30$; dashed plot lines correspond to situations where $30 < \alpha_{r,DMFT}^*(g_h)$. (f) τ_T vs. α_r for $N = 500$; dashed vertical line is $\alpha_{r,DMFT}^*(1.85)$. Transient times are averaged over 2000 instances of random networks.

g_h above 2.0, the chaotic set smoothly transitions into a chaotic attractor once the stable FP disappears. We can also see the chaotic nature of the dynamical state by looking at the full Lyapunov spectrum across the transition (Fig. 6d inset), and we see that the full spectrum undergoes a discontinuous jump across the transition (red to green). Moreover, we can estimate the dimension of the chaotic attractor by calculating an upper-bound D_A on the dimension according to a conjecture by Kaplan & Yorke [57]. The Kaplan-Yorke upper bound for the attractor dimension D_A is given by

$$D_A = M + \frac{\sum_{i=1}^M \lambda_i}{|\lambda_{M+1}|} \quad \text{where } M = \max \left\{ \sum_{i=1}^j \lambda_i \geq 0 \right\}, \quad (78)$$

where λ_i are the rank-ordered Lyapunov exponents. We see in Fig. 6d that the attractor dimension jumps from zero to a finite value across the transition (dashed line). The attractor dimension also appears to be extensive in the system size (data not shown). Thus, the dynamical transition is characterized by the spontaneous appearance of a chaotic set.

As we mentioned above, the zero FP is still stable across the dynamical transition and coexists with the chaotic set. For finite systems, the dynamics will eventually flow into the zero FP after chaotic transients. To see this, we study the behavior of transient activity as a function of system size. If the picture of a chaotic attractor with a punctured hole is valid, we expect the transient activity to be chaotic and eventually flow into the fixed-point. Moreover, above the dynamical transition, we expect this transient time to scale with the system size, and in the infinite system size limit, the transient time should diverge in spite of the fact that stable the fixed-point still exists. This is because the volume of the basin of attraction of the fixed-point relative to that of the chaotic set will vanish as $N \rightarrow \infty$.

In Fig. 7c,d we do indeed see that the transient time for a fixed g_h scales with system size (Fig. 7c) once α_r is above the transition (dashed line). Conversely, for a fixed α_r , the transients scale with system size once g_h crosses the value for which the (fixed) α_r is critical (Fig. 7d). If α_r is below the critical value then there is no scaling of transient time with system size (see Fig. 7a,e dashed

lines). Above the transition for a fixed g_h the transient time also increases with increasing α_r (Fig. 7a,f) consistent with the observation in Fig. 6d that the attractor dimension increases with α_r beyond the critical value. Once g_h becomes larger than $g_c = 2.0$, the stable fixed-point becomes unstable; however, the chaotic attractor when it exists, transitions smoothly across the $g_h = 2.0$ line, as can be seen in Fig. 4c ($\alpha_r = 20$).

The picture of the transition that emerges is summarized in the schematic in Fig. 6e: when $g_h < 2.0$ the zero fixed-point is stable. Moreover, when $\sqrt{2} < g_h \leq 2$, and α_r crosses $\alpha_{r,FP}^*$, unstable fixed-points spontaneously appear in the phase space (the FP transition) without any salient dynamical signatures. On increasing α_r further, a dynamical state with chaotic transients and stable FP, discontinuously appears once α_r crosses $\alpha_{r,DMFT}^*$ (the DMFT transition), provided $\sqrt{8/3} < g_h < 2.0$. This coexistence phase – with the chaotic activity and the stable fixed-point – exhibits long chaotic transients whose lifetime scales with the system size. On increasing g_h , with fixed α_r , the chaotic set smoothly varies across the $g_h = 2.0$ transition and becomes a chaotic attractor once the stable FP disappears (if $\alpha_r < \alpha_{r,DMFT}^*$, then a chaotic attractor appears continuously).

VIII. BIASES PRODUCE NON-TRIVIAL FIXED-POINTS AND CONTROL THEIR STABILITY

We have thus far described the salient dynamical aspects for the gated RNN in the absence of biases. Here we describe the role of the biases β_h (bias of the activation ϕ) and β_r (bias of the output gate σ_r). We first note that when $\beta_h = 0$, zero is always a fixed-point of the dynamics, and the zero fixed-point is stable provided

$$-1 + \phi'(0)\sigma_r(0) < 0 \quad (79)$$

where $\phi(x) = \tanh(g_h x + \beta_h)$. This gives the familiar $g_h < 2$ condition when $\beta_r = 0$ [62]. Thus, in this case, there is an interplay between g_h and β_r in determining the leading edge of the Jacobian around the zero fixed-point, and thus its stability. In the limit $\beta_r \rightarrow -\infty$ the leading edge retreats to $-\tau_r^{-1}$. When $\beta_h > 0$, zero cannot be a fixed-point of the dynamics. Therefore, β_h facilitates the appearance of non-zero fixed-points, and both β_r and β_h will determine the stability of these non-zero fixed-points.

To gain some insight into the role of β_h in generating fixed-points, we treat the mean-field FP equations (eq. 19) perturbatively around the operating point g_c where the zero fixed-point becomes unstable (eq. 79). For small β_h and $\epsilon = g_h - g_c$, we can express the solution Δ_h as a power series in ϵ , and we see that to leading order the

fixed-point variance behaves as (details in Appendix D 1):

$$\Delta_h \approx \begin{cases} \frac{\beta_h + \epsilon}{g_c^2(2 - g_c^2 a_1)} & g_c^2 a_1 < 2 \\ (g_c^2 a_1 - 2) f_1 + \epsilon \cdot f_2 & g_c^2 a_1 > 2 \end{cases} \quad (80)$$

$$\text{where } a_1 = \frac{\alpha_r^2}{16} \left[\phi_0^{(1)}(\beta_r/2)^2 + \phi_0(\beta_r/2)\phi_0^{(2)}(\beta_r/2) \right] \quad (81)$$

where $\phi_0 \equiv \tanh$ and $f_2(\alpha_r, \beta_r)$ and $f_2(\alpha_r, \beta_r)$ are constant functions w.r.t ϵ . Therefore, we see that the bias β_h gives rise to non-zero fixed-points near the critical point which scale linearly with the bias. In Fig. 8e, we show this linear scaling of the solution for the case when $\beta_h = \epsilon$, and we see that the prediction (lines) matches the true solution (circles) over a reasonably wide range.

More generally, away from the critical g_c , an increasing β_h gives rise to fixed-point solutions with increasing variance, and this can arise continuously from zero, or it can arise by stabilizing an unstable, time-varying state depending on the value of β_r . In Fig. 8a we see how the Δ_h behaves for increasing β_h for different β_r , and we can see the stabilizing effect of β_h on unstable solutions by looking at its effect on the leading spectral edge (Fig. 8b). In Fig. 8c, we see that an increasing β_r also gives rise to increasing Δ_h . However in this case, it has a destabilizing effect by shifting the leading spectral edge to the right. In particular, when $\beta_h = 0$, increasing β_r destabilizes the zero fixed-point and give rise to a time-varying solution. We note that when $\beta_h = 0$, varying β_r cannot yield stable non-zero FPs. The combined effect of β_h and β_r can be seen in Fig. 8f where the non-zero solutions to the left of the orange line indicate unstable (time-varying) solutions. We have chosen the parameters to illustrate an interesting aspect of the biases: in some cases, increasing β_h can have a non-monotonic effect on the stability, wherein the solution becomes unstable with increasing β_h and is then eventually stabilized for sufficiently large β_h .

IX. PHASE DIAGRAM FOR THE GATED RNN

As we have seen above, the gated RNN exhibits a rich array of qualitatively different states. Here, we summarize the key aspects of these different phases and the critical lines separating them. The key parameters determining the critical lines and the phase diagram are the activation and output-gate gains and the associated biases: $(g_h, \beta_h, \alpha_r, \beta_r)$. The update gate does not play a role in determining the critical lines; however, it can have a strong effect on the dynamical aspects of the states near the critical lines. As we will see below, there are macroscopic regions of the parameter space adjacent to the critical lines (on the unstable/chaotic side) where the states can be made marginally stable in the limit of $\alpha_z \rightarrow \infty$. We first begin with the case when there are no biases, and subsequently describe the role of biases.

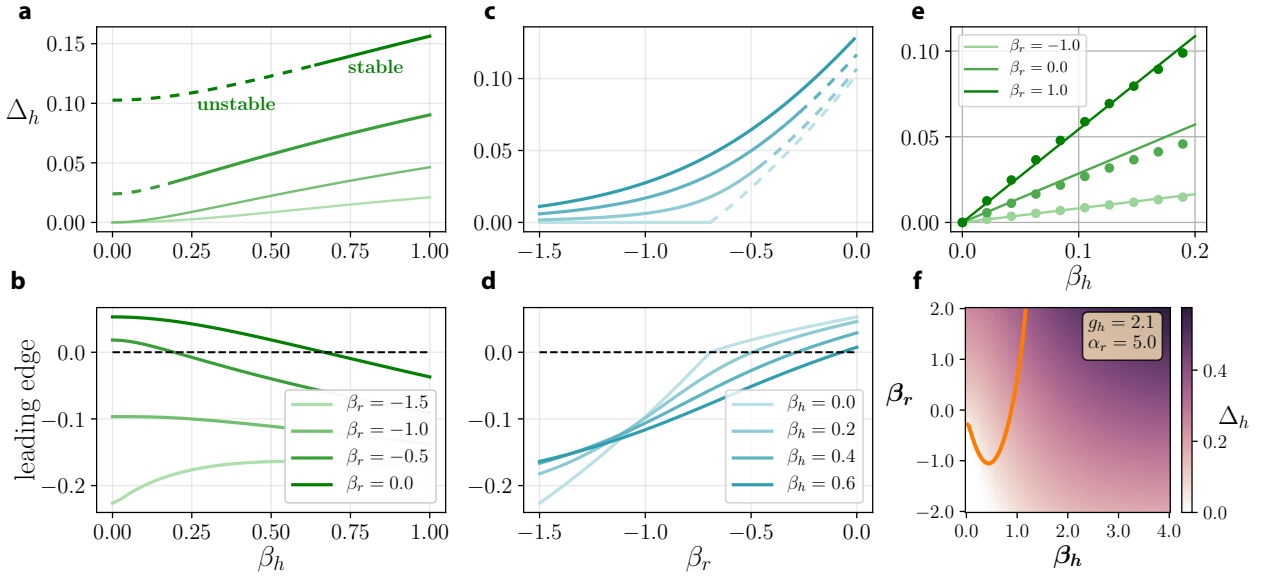


FIG. 8. *The role of biases* (a) FP solutions as a function of increasing β_h ; different shades of green correspond to different values of β_r . Dashed lines correspond to FP solutions that are unstable (time-varying states). (b) The leading edge of the spectrum corresponding to the FP solutions calculated in (a); the FP solution is unstable when the leading edge is positive. (c) similar to (a) but for β_r ; different shades of blue correspond to different values of β_h . (d) similar to (b) but for β_r . (e) FP solutions near critical g_c where the zero FP becomes unstable (circles) compared with the perturbative solution predicted by eq. 80 (solid lines). (f) FP solution as a function of β_r and β_h . Orange line indicates stability line – i.e. regions on top of the orange line correspond to unstable/time-varying states.

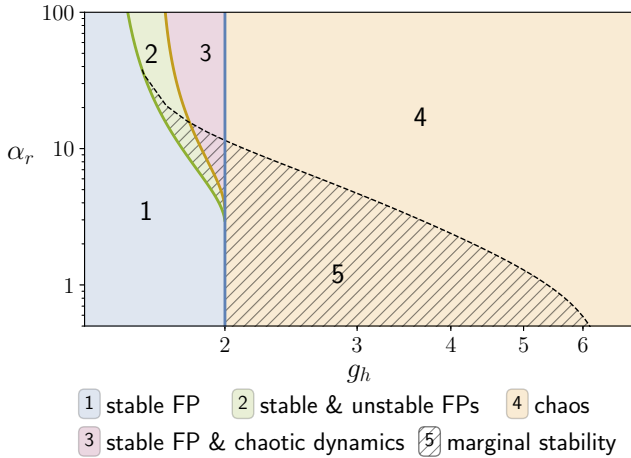


FIG. 9. *Phase diagram for the gated RNN without biases*

A. Phases and critical lines without biases

We can summarize the phase diagram without biases in the 2-D parameter space characterized by the gains for the activation and the output-gate (g_h, α_r). A salient feature of the network without biases are the discontinuous bifurcation transitions described in Sec. VII. In Fig. 9, we see that when g_h is below 2.0 and $\alpha_r < \alpha_{r,FP}^*$, the zero-fixed point is the only solution (region 1). On crossing

the fixed-point bifurcation line (green line, Fig. 9), there is a spontaneous proliferation of unstable fixed-points in the phase space (region 2). This is, however, possible only when $g_h > \sqrt{2}$. This proliferation of fixed-points is not accompanied by any obvious dynamical signatures. However, if $\sqrt{8/3} < g_h < 2$, we can increase α_r further to cross a second discontinuous transition where a dynamical state spontaneously appears featuring the coexistence of chaotic activity and a stable fixed-point (region 3). This phase is characterized by long, chaotic transient activity which eventually flows into the fixed-point at a time that scales with the system size. When g_h is increased beyond the critical value of 2.0, the stable fixed-point becomes unstable for all α_r , and the nonattracting chaotic set, responsible for the chaotic transients, varies continuously across this transition and becomes a chaotic attractor – if there was no chaotic set for $g_h = 2^-$, then a chaotic attractor appears continuously from 0 for $g_h = 2^+$. This region is characterized by chaotic activity with an extensive attractor dimension. All these critical lines are determined by g_h and α_r , and α_z has no explicit role; however, for large α_z there is a large region of the parameter space on the chaotic side of the chaotic transition that can be made marginally stable (thatched region in Fig. 9; determined from eq. 36).




FIG. 10. How the biases alter the transition between stability and chaos (left) Critical lines indicating boundaries for stability (solid lines) or marginal stability (dashed lines) for different values of β_h . (right) Similar to the left panel but for different values of β_r .

B. Role of biases

As discussed in Sec. VIII, the main feature of the biases is facilitating non-zero fixed-point solutions and determining the leading edge of the spectral curve, and thus controlling the stable-to-chaos transition. Another key feature of biases is the suppression of the discontinuous bifurcation transition observed without biases. We note that even in the presence of biases it is possible to have small regions of the parameter space with bifurcations; however, we leave a detailed study of biases and bifurcations for future work. In the following, we consider regimes without a bifurcation and thus the transition to chaos occurs (continuously) when the leading spectral edge becomes unstable.

In Fig. 10, we look at how the critical line for the chaotic transition, in the $\alpha_r - g_h$ plane, changes as we vary β_h (left panel) or β_r (right panel). Positive values of β_r (“open” output gate) tend to make the transition line less dependent on α_r (Fig. 10, right panel), and negative values of β_r have a stabilizing effect by requiring larger values of g_h and α_r to transition to chaos. As we have seen in Sec. VIII, higher values of β_h have a stabilizing effect, requiring higher g_h and α_r to make the (non-zero) stable fixed-point unstable. In both case, the critical lines for marginal stability (Fig. 10, dashed lines) are also influenced in a similar way.

We can also study the role of β_h by plotting the phase diagram in the $\beta_h - g_h$ plane for a fixed α_r, β_r (Fig. 11 top). We see that larger values of β_h require higher activation gains to transition to chaos, and on the chaotic side of the critical line, there is a region of the parameter space that can become marginally stable for $\alpha_z \rightarrow \infty$. Such a stability-to-chaos transition (without marginal stability) for the biases and gains of the activation has also been observed in feed-forward networks [63]. In Fig. 11, bottom panels, we see how the stability-to-chaos transition is affected by α_r (left panel) and β_r (right panel). Consistent with the discussion above, larger α_r and β_r

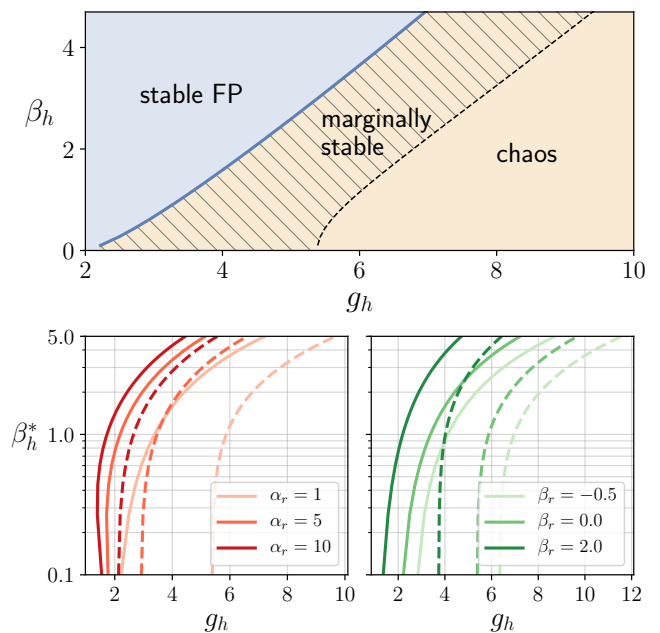


FIG. 11. How the bias β_h controls the transition between stability and chaos (top) stability-to-chaos transition in the $g_h - \beta_h$ plane. Hatched region indicates unstable points that can be made marginally stable in the limit $\alpha_z \rightarrow \infty$. Here $\alpha_r = 1.0, \beta_r = 0.0$. (bottom) how the boundaries of stability (solid lines) or marginal stability (dashed lines) change as we vary α_r (left) or β_r (right).

have a destabilizing effect, requiring a larger β_h to make the system stable.

X. TRAINING RNNs: A DYNAMICAL MEAN-FIELD THEORY OF GRADIENTS

In the previous sections, we discussed in detail the rich nature of the autonomous dynamics of the gated RNN with randomly distributed weights. Here we address the issue of *training* these RNNs on a task – e.g. producing a desired trajectory – using gradient signals. The problem of understanding how variations in the parameters affect the trajectories of a dynamical system has its roots in the control theory literature, and in particular, the Adjoint Sensitivity method introduced by Pontryagin and colleagues [43] provides an elegant framework to study this. This method makes explicit the intimate link between the forward/autonomous dynamics of a dynamical system and the characteristics of the gradients of the trajectories with respect to the parameters. We first describe the Adjoint formalism and how to calculate gradients with it, and then show how to incorporate the Adjoint formalism in the DMFT framework by using the vanilla RNN as a model. Finally, we develop the Adjoint DMFT for the gated RNN and derive predictions for the statistics of gradients in the gated network.

A. The Adjoint formalism

Let us suppose that the dynamics of the system are described by the equation

$$\dot{\mathbf{x}}(t) - G(\mathbf{x}, \boldsymbol{\theta}, t) = 0 \quad (82)$$

where $\boldsymbol{\theta}$ denotes all the parameters – e.g., weights, gains and biases for the gated RNN. \mathbf{x} is the combined vector of all the dynamical state variables – in the case of the gated network, it will be the vector: $\mathbf{x} = (\mathbf{h}, \mathbf{z}, \mathbf{r})$. Furthermore, let us denote by \mathcal{F} a cost-functional that assigns scalar scores to the trajectories of the network:

$$\mathcal{F}(\mathbf{x}(t), \boldsymbol{\theta}) = \int_0^T dt f(\mathbf{x}, \boldsymbol{\theta}, t) \quad (83)$$

For instance, the cost might be the distance from some desired trajectory \mathbf{x}^* in which case an appropriate cost functional would be $\mathcal{F}(\mathbf{x}(t), \boldsymbol{\theta}) = T^{-1} \int_0^T dt \|\mathbf{x}(t) - \mathbf{x}^*(t)\|^2$. We are interested in the gradient of the cost with respect to the parameters: $d\mathcal{F}/d\boldsymbol{\theta}$. However, calculating this can be challenging due to the $d\mathbf{x}/d\boldsymbol{\theta}$ term. The Adjoint method introduced by Pontryagin et al. circumvents this problem by introducing an auxiliary variable governed by another, dual dynamics. The starting point is the Lagrangian

$$\mathcal{L} = \int_0^T dt f(\mathbf{x}, \boldsymbol{\theta}, t) + \int_0^T dt \boldsymbol{\lambda}(t)^T (\dot{\mathbf{x}}(t) - G(\mathbf{x}, \boldsymbol{\theta}, t)) \quad (84)$$

where $\boldsymbol{\lambda}(t)$ is a Lagrange multiplier for enforcing the dynamics. Along a trajectory of the dynamics we can express the gradient of the cost as

$$\frac{d\mathcal{F}}{d\boldsymbol{\theta}} = \frac{d\mathcal{L}}{d\boldsymbol{\theta}} = \int_0^T dt \left(\frac{\partial f}{\partial \boldsymbol{\theta}} + \frac{\partial f}{\partial \mathbf{x}} \frac{d\mathbf{x}}{d\boldsymbol{\theta}} \right) + \int_0^T dt \boldsymbol{\lambda}(t)^T \left(-\frac{\partial G}{\partial \boldsymbol{\theta}} - \frac{\partial G}{\partial \mathbf{x}} \frac{d\mathbf{x}}{d\boldsymbol{\theta}} + \frac{d\dot{\mathbf{x}}}{d\boldsymbol{\theta}} \right) \quad (85)$$

The $d\mathbf{x}/d\boldsymbol{\theta}$ terms in the expression for the gradient can be eliminated by the following choice of $\boldsymbol{\lambda}(t)$: i) $\boldsymbol{\lambda}(T) = 0$ and ii) for $t \neq T$, $\boldsymbol{\lambda}(t)$ satisfies

$$\frac{d\boldsymbol{\lambda}(t)}{dt} = -\left(\frac{\partial G}{\partial \mathbf{x}} \right)^T \boldsymbol{\lambda}(t) + \frac{\partial f}{\partial \mathbf{x}} \quad (86)$$

where $\partial_{\mathbf{x}}G^T$ is the transpose (Adjoint) of the instantaneous Jacobian. The gradient is then given by

$$\frac{d\mathcal{F}}{d\boldsymbol{\theta}} = \frac{d\mathcal{L}}{d\boldsymbol{\theta}} = \int_0^T dt \left(\frac{\partial f}{\partial \boldsymbol{\theta}} - \boldsymbol{\lambda}(t)^T \frac{\partial G}{\partial \boldsymbol{\theta}} \right) \quad (87)$$

The steps to calculate the gradient of the cost with respect to the parameters can be summarised as follows: i) run the forward dynamics (eq. 82) to get the trajectory $\mathbf{x}(t)$; ii) run the dynamics for $\boldsymbol{\lambda}$ (eq. 86) *backwards* in time with the initial condition $\boldsymbol{\lambda}(T) = 0$ and the state-to-state Jacobian $\partial_{\mathbf{x}}G$ evaluated along the trajectory $\mathbf{x}(t)$ and iii) with $\mathbf{x}(t)$ and $\boldsymbol{\lambda}(t)$ the gradient is given by eq. 87.

We now make a few observations:

- The equation for $\boldsymbol{\lambda}(t)$ (eq. 86) is linear and is governed by the same state-to-state Jacobian $\partial_{\mathbf{x}}G$ that characterizes the behaviour of the forward dynamics.
- Empirically, training RNNs for machine learning applications can be difficult because gradients vanish or explode rapidly with the training time T . From eq. 87, we see that in typical cases where the cost f does not explicitly depend on the parameters, the rapid growth/decay of the gradients is dictated by the rapid growth/decay of $\boldsymbol{\lambda}(t)$, which in turn is dictated by the asymptotic behaviour of the linear ODE governed by $\partial_{\mathbf{x}}G$.
- Our Lyapunov analysis of the gated RNN shows that the asymptotic behavior agrees well with the local behavior predicted by the Jacobian spectrum, thus if we want finite (and $O(1)$) gradients we should initialize networks where the Lyapunov spectrum is flat; i.e., near critical lines and especially in the marginally stable regions. This provides a more fine-grained perspective on the empirical observation that it is easier to train RNNs initialized “at the edge of chaos” [10, 35–37].

Having highlighted the role of $\boldsymbol{\lambda}(t)$ in governing the behavior of gradients, we show below how to incorporate the Adjoint dynamics of $\boldsymbol{\lambda}(t)$ into the DMFT, a tool typically used to study the forward dynamics.

B. Adjoint DMFT for the vanilla RNN

We first outline how to incorporate gradients in the MSRDJ formalism for the vanilla RNN, as it is simpler than the gated RNN. The starting point for the Adjoint DMFT is, again, the generating functional in the MSRDJ formalism:

$$Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = \mathbb{E} \left[\exp \left(i \sum_{j=1}^N \int \hat{\mathbf{b}}_j(t)^T \mathbf{x}_j(t) dt \right) \right] \quad (88)$$

where $\mathbf{x}_j(t) \equiv (h_j(t), \lambda_j(t))$ and $\hat{\mathbf{b}}_j(t) = (\hat{b}_j^h, \hat{b}_j^\lambda)$. The measure in the expectation is a path integral over the joint dynamics of $h(t)$ and $\lambda(t)$, and this can be expressed by means of an action by a procedure similar to that used for the forward dynamics:

$$Z_{\mathcal{J}} = \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \prod_k \mathcal{D}\lambda_k \mathcal{D}\hat{\lambda}_k e^{-i\mathbf{S}} \quad (89)$$

$$\mathbf{S} = i \sum_{j=1}^N \int dt \left(\hat{\mathbf{x}}_j(t)^T \mathbf{b}_j(t) dt + \hat{\mathbf{b}}_j(t)^T \mathbf{x}_j(t) \right) + i\mathbf{S}_{\mathcal{J}} \quad (90)$$

$$\mathbf{S}_{\mathcal{J}} = i \int dt \left[\hat{\mathbf{h}}(t)^T (\dot{\mathbf{h}}(t) - G(\mathbf{h}(t), J)) + \hat{\boldsymbol{\lambda}}(t)^T (\dot{\boldsymbol{\lambda}}(t) + \partial_{\mathbf{h}}G(t)^T \boldsymbol{\lambda}(t) - \partial_{\mathbf{h}}f(\mathbf{h}(t))) \right] \quad (91)$$

where, for the vanilla RNN, the function G describing the e.o.m and the state-to-state Jacobian $\partial_{\mathbf{h}}G$ are given by

$$G(\mathbf{h}(t), J) = -\mathbf{h}(t) + J\phi(\mathbf{h}(t)) \quad (92)$$

$$\partial_{\mathbf{h}}G(\mathbf{h}(t), J) = -\mathbb{I} + J[\phi'] \quad (93)$$

We are interested in the typical behavior of gradients for large networks, so as before we disorder average over J to get a joint description of the forward dynamics and the adjoint dynamics as two, coupled stochastic differential equations (details in Appendix F):

$$\dot{h}(t) = -h(t) + \eta(t) \quad (94)$$

$$\dot{\lambda}(t) = \lambda(t) - \phi'(h(t))\xi(t) + \partial_h f(h(t)) \quad (95)$$

$$\langle \eta(t)\eta(t') \rangle = C_\phi(t, t'), \quad \langle \xi(t)\xi(t') \rangle = C_\lambda(t, t') \quad (96)$$

where η and ξ are independent Gaussian processes with correlation functions C_ϕ and C_λ , that represent the population-averaged correlation functions

$$C_\lambda(t, t') = \frac{1}{N} \boldsymbol{\lambda}(t)^T \boldsymbol{\lambda}(t'), \quad C_\phi(t, t') = \frac{1}{N} \boldsymbol{\phi}(t)^T \boldsymbol{\phi}(t') \quad (97)$$

From eq. 94 we find an ODE for the correlation function:

$$(\partial_t - 1)(\partial_s - 1)C_\lambda(t, s) = C_{\phi'}(t, s)C_\lambda(t, s) + \langle \partial_h f(t) \partial_h f(t') \rangle \quad (98)$$

This is the reduced DMFT description of the forward + adjoint dynamics. We would like to use this to calculate the DMFT prediction for the gradient statistics. To do this, we see from eq. 87 that the mean-squared gradient is given by

$$\left\| \frac{d\mathcal{F}}{dJ} \right\|^2 = \frac{1}{N^2} \sum_{i,j} \left\langle \left(\frac{d\mathcal{F}}{dJ_{ij}} \right)^2 \right\rangle \rightarrow \int_0^T dt ds C_\lambda(t, s) C_\phi(t, s) \quad (99)$$

Thus, for instance, we can use the DMFT prediction for the gradient norm to examine conditions under which the gradients will exceed some threshold magnitude.

C. Adjoint DMFT for the gated RNN

The Adjoint DMFT for the gated RNN follows along similar lines as the vanilla RNN, but now we have a state variable which is composed of three variables: $\mathbf{x}_j(t) \equiv (h_j(t), z_j(t), r_j(t))$, and correspondingly we will need to consider more correlation and response functions in the DMFT. In this case, the dynamics for the Adjoint variable $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_h, \boldsymbol{\lambda}_z, \boldsymbol{\lambda}_r)$ will be

$$\dot{\boldsymbol{\lambda}} = \mathcal{D}(t)^T \boldsymbol{\lambda} + \partial_x \mathbf{f} \quad (100)$$

where $\partial_x \mathbf{f} = (\partial_h \mathbf{f}, \partial_z \mathbf{f}, \partial_r \mathbf{f})$, and $\mathcal{D}(t)$ is the state-to-state Jacobian we considered in eq. 22.

After the saddle-point approximation, the single-site mean-field equations look the same as eq. 11 for the state variables (h, z, r) , and the equations for the adjoint variables are given by

$$\dot{\lambda}_h - \lambda_h \sigma_z(z) + \phi'(h) \sigma_r(r) \xi_h + \tau_z^{-1} \phi'(h) \xi_z + \tau_r^{-1} \phi'(h) \xi_r = \partial_h f, \quad (101)$$

$$\dot{\lambda}_z - \lambda_h \sigma'_z(z) h - \tau_z^{-1} \lambda_z + \lambda_h \sigma'_z(z) \zeta_z = \partial_z f, \quad (102)$$

$$\dot{\lambda}_r - \tau_r^{-1} \lambda_r + \phi(h) \sigma'_r(r) \zeta_r = \partial_r f, \quad (103)$$

where $\zeta_{z,r}$ and $\xi_{h,z,r}$ are Gaussian processes whose correlation functions are determined self-consistently:

$$\langle \xi_h(t) \xi_h(t') \rangle = \langle \lambda_h(t) \sigma_z(z_t) \lambda_h(t') \sigma_z(z_{t'}) \rangle, \quad (104)$$

$$\langle \xi_z(t) \xi_z(t') \rangle = \langle \lambda_z(t) \lambda_z(t') \rangle, \quad (105)$$

$$\langle \xi_r(t) \xi_r(t') \rangle = \langle \lambda_r(t) \lambda_r(t') \rangle, \quad (106)$$

$$\langle \zeta_r(t) \zeta_r(t') \rangle = \langle \lambda_h(t) \sigma_z(z_t) \lambda_h(t') \sigma_z(z_{t'}) \rangle, \quad (107)$$

$$\langle \zeta_z(t) \zeta_z(t') \rangle = \langle \phi(h_t) \sigma_r(r_t) \phi(h_{t'}) \sigma_r(r_{t'}) \rangle, \quad (108)$$

$$\langle \zeta_r(t) \zeta_r(t') \rangle = \langle \lambda_h(t) \sigma_z(z_t) \lambda_h(t') \sigma_z(z_{t'}) \rangle. \quad (109)$$

Using eq. 87, we see that the gradients of the loss in the full network can be written as

$$\frac{d\mathcal{F}}{dJ_{ij}^h} = \int dt \lambda_i^h \sigma_z(z_i) \phi(h_j) \sigma_r(r_j), \quad (110)$$

$$\frac{d\mathcal{F}}{dJ_{ij}^z} = \int dt \lambda_i^z \phi(h_j), \quad (111)$$

$$\frac{d\mathcal{F}}{dJ_{ij}^r} = \int dt \lambda_i^r \phi(h_j), \quad (112)$$

which leads to the following DMFT prediction for the mean-squared gradients (i.e. the squared Frobenius norm of the matrix of gradients):

$$\left\langle \left\| \frac{d\mathcal{F}}{dJ^h} \right\|^2 \right\rangle = \int dt dt' C_{\lambda_h}(t, t') C_\phi(t, t') C_{\sigma_r}(t, t'), \quad (113)$$

$$\left\langle \left\| \frac{d\mathcal{F}}{dJ^z} \right\|^2 \right\rangle = \int dt dt' C_{\lambda_z}(t, t') C_\phi(t, t'), \quad (114)$$

$$\left\langle \left\| \frac{d\mathcal{F}}{dJ^r} \right\|^2 \right\rangle = \int dt dt' C_{\lambda_r}(t, t') C_\phi(t, t'). \quad (115)$$

XI. DISCUSSION

Significance of gating interactions: We introduced a gated RNN that naturally extends a canonical continuous-time RNN model to include gating interactions – a salient feature of the best-performing machine learning RNNs [7] and single-neuron models with realistic biophysics [15, 25]. Gating improves performance in machine learning tasks involving language modeling [64], speech recognition [2], neural machine translation

[65–67]. Prior work on gating has studied input-output Jacobian in discrete-time RNNs assuming independent weights at each time-step [68, 69]; this effectively makes the model a feed-forward model. Our gated RNN properly deals with the highly recurrent nature of the dynamics and reduces to the classical RNN [16, 26] when the gates are open. Moreover, when the dynamics are discretized, the gated RNN closely resembles a widely-used machine learning RNN [65]. Specifically, the main differences from the machine learning RNN are that the individual neurons integrate their inputs in a leaky fashion instead of being directly affected by the input at each time-step, and the single-neuron nonlinearity acts on the state variable of each neuron instead of the total summed inputs; these differences are not only more biologically realistic but also endow the individual neurons with richer dynamics [22]. We use our gated RNN to study the consequences of gating on collective dynamics and gradients that arise during training tasks.

Generation of slow modes by the update gate: The update gate in our gated RNN functions as a dynamical time-constant, which controls the rate at which a neuron integrates its inputs. In single neuron models, such a dynamical time-constant has been shown to make the neuron’s responses robust to time-warped inputs [15]. Furthermore, normative approaches, requiring time-warp invariance in machine learning RNNs naturally imply the existence of a gate that effectively acts as an adaptive time-constant [70]. In our gated RNN, the sensitivity of the update gate is a key control parameter, and more sensitive (or switch-like) update gates lead to an accumulation of slow modes in the dynamics. Indeed, the emergence of slow modes with a spectrum of timescales is likely useful for processing inputs which have dependencies over a wide range of timescales [40, 71]. We show that accumulation of these slow modes is a generic feature of gating controlling the rate of integration; moreover, our explicit choice of the gate sensitivity as a parameter, suggests that this could be a potentially useful trainable parameter for gated machine learning RNNs.

Marginal stability and line attractors: When the activation gain g_h is not too large, increasing the update gate sensitivity (α_z) leads to a drastic reduction in the number of unstable directions of the local dynamics, and in the limit it can make a system that was previously unstable marginally stable. Thus, as discussed above, there are extensive volumes in the parameter space adjacent to the critical surfaces (on the chaotic side) which can be made marginally stable by making the update gate switch-like. Marginally stable models of biological function have long been of interest with regard to their benefits for information processing (c.f. [71] and references therein); however, most models to achieve marginal stability have required fine-tuning. Here, we show that gating interactions can be another method to achieve marginal stability in a robust way. One consequence of marginal stability is the emergence of line attractors – such line attractor dynamics have been shown to be ben-

eficial for a variety of computational functions such as motor control [41, 42, 72], decision making [73] and auditory processing [74]. Signatures of line attractors are also empirically observed in successfully trained gated RNNs on classifying text sentences [39]. We provide a theoretical basis for how gating could produce line attractors right at initialization.

Role of the output gate: The output gate has the effect of increasing the spectral radius when it becomes more sensitive (larger α_r). Thus, for chaotic states, more sensitive output gates will cause them to become more chaotic. For stable fixed-points, we show that making the output gate more sensitive can lead to a novel, discontinuous first-order transition to chaos. In addition to its effect on autonomous dynamics, the output gate is likely to have a significant role in controlling the influence of external inputs on the intrinsic dynamics. This can be seen by noting that when the output gate is close to zero, the internal state will mostly integrate the input – as intuited in the work introducing the discrete-time gated RNN for machine learning applications [65]. This can be functionally useful as a means of resetting the dynamics, and erasing the dynamical traces from the past.

Transition to chaos in the gated RNN: In the classical RNN, there is a single transition from quiescence to chaos when the variance of the couplings exceeds a threshold value [16]. This is a continuous transition where the correlation timescale of the fluctuations diverges as we approach the transition from above. In the gated RNN, we have two types of transitions to chaos. When α_r is smaller than $\sqrt{8}$ (no biases), there is a continuous transition to chaos when g_h crosses the critical value 2.0. This transition shares characteristics of the transition in the classical RNN. However, there is a second discontinuous transition when α_r is increased in the quiescent region, which is characterized by the spontaneous emergence of chaotic dynamics, along with the existence of a stable fixed-point. Moreover, this transition also has the characteristic that unstable fixed-points proliferate in the phase space before the chaotic dynamics appear, and thus it differs from the interesting result of Wainrib and Touboul for additive RNNs, where these two phenomena are tightly linked. This suggests that gating could facilitate a distinct mechanism for the transition to chaos and could have important implications for the stability of biological neural networks where changes in effective gating parameters could abruptly push the collective dynamics from stable to strongly irregular.

From a practical perspective, studying the transition to chaos is of interest in light of the observation that it is often easier to train RNNs initialized in the chaotic regime but close to the critical points. This is often referred to as the “edge of chaos” hypothesis [75–77]. This could be a result of the long timescales present in the dynamics at these operating points. Thus a theoretical characterization of the critical surfaces indicating the transition to chaos in the parameter space is helpful for practitioners. We also note that the strong chaotic nature of the

dynamics at the discontinuous chaotic transition seems unlikely to be beneficial for training on natural tasks, and it seems more appropriate to initialize near critical lines corresponding to continuous transitions.

The role of biases: Initialization choices for biases have a significant impact on task performance in machine learning architectures [7, 32, 70, 78]. Here we studied how biases affect the dynamics and qualitatively change the phase diagram. The bias β_h in the nonlinear activation function has the effect of producing stable non-zero fixed point solutions; indeed, for non-zero β_h zero is no longer a stationary point of the fixed-point equations. Moreover, the biases β_h and β_r also control the leading edge of the spectral curve, and thus affect the stability of the fixed-points. Higher values of the β_h, β_r require higher values of g_h for a transition to chaos. Another interesting aspect of β_h is that it suppresses the discontinuous chaotic transition and thus the transition from the stable fixed-point to the chaotic time-varying state is a continuous one. From a functional perspective, biases could be used to stabilize regions of the parameter space or engineer stable fixed-points corresponding to memory states.

DMFT for gradient dynamics: In machine learning applications, the dynamics of RNNs at initialization is critical for training success. Thus studying the autonomous dynamics of random RNNs, provides insights on training performance. Indeed, there is a deeper connection between the forward dynamics and the behaviour of gradients. The adjoint formalism introduced by Pontryagin et al. [43] explicitly shows that the (exponential) growth or decay of gradients is governed by the same state-to-state Jacobian which governs the local behavior in the forward dynamics. It is also known from empirical studies that one of the key difficulties in training RNNs is the problem of exploding or vanishing gradients [3]; this suggests that for efficient training we need gradients to be stable and $O(1)$. The theory suggests a way to achieve this – train at parameter regimes where the Jacobian eigenvalues are close to zero i.e. near critical lines. This is a more fine-grained perspective on why it might be easier to train at the “edge of chaos”. In the case of our gated RNN, in addition to critical lines, we also know that the update gate can make many parameter combinations marginally stable. The arguments above would then suggest that initializing in the marginally stable regions should be beneficial for taming the gradients during training.

We have combined the MSRDJ field theory formalism to incorporate the adjoint framework for gradients, to formulate a DMFT for gradients. The DMFT is typically, a tool typically used to study forward evolution, and our extension allows the powerful techniques of the DMFT to be used in the analysis of gradients. Fully exploring this connection will be undertaken in future work.

Generality of techniques: We provide a self-contained treatment of developing the DMFT for a RNN with multiplicative interactions, and we show how the mean-field theory can be combined with techniques from random

matrix theory to calculate the spectral properties of the state-to-state Jacobian. These theoretical frameworks are general and powerful, and we expect that the self-contained treatment provided here should enable machine learning researchers to extend these analyses to other architectures and forms of gating, thereby providing a principled way to assess the impact of architectural choices on the richness of dynamics and gradients.

A. Conclusions and outlook

Gating is a form of multiplicative interaction that is a central feature of the best performing RNNs in machine learning, and it is also a salient feature of biological neurons. Prior theoretical work on RNNs has only considered RNNs with additive interactions. Here, we present the first detailed study on the consequences of gating for RNNs and show that gating can produce dramatically richer behavior than classical RNNs.

Line attractors are widely-studied mechanisms for biological and computational functions involving long memory. To date, most of the models proposed for generating line attractors require fine-tuning. We show how gating can robustly generate line attractors without fine-tuning.

We provide the first detailed phase diagram for a RNN with gating, thus delineating the dynamical phases in the parameter space. This is highly relevant to practitioners in machine learning, since choosing the initial parameters for training is one of the critical determinants of training outcome. The phase diagram provides a principled map for doing so.

We identify a novel, discontinuous transition from a stable to a chaotic state caused by gating. Such a phenomenon has not been reported in classical RNNs. As a practical consequence, this result suggests that the usual heuristic of initializing RNN parameters at the edge-of-chaos in machine learning needs to be reconsidered since not all transition boundaries have beneficial properties – especially with gating. This novel transition to chaos is also relevant for stability of other large complex systems with multiplicative interactions.

We have set down the first complete DMFT treatment of gradients in large RNNs, by combining the adjoint sensitivity method, widely used in control theory, with field-theoretic methods. This paves the way for applying the powerful techniques of DMFT – mostly used to study the forward evolution of dynamics – to analyze gradient signals at initialization and during training. Our framework also highlights the close link (via the Jacobian) between forward dynamics and the behavior of gradients, and is thus relevant to the broader set of questions surrounding dynamical systems in biology that reorganize based on gradient information.

Lastly, we would like to point out that the analytical methods have been presented in a self-contained manner, and it should be possible to adapt them to study other large complex systems with gating interactions.

ACKNOWLEDGMENTS

We are most grateful to William Bialek, Giulio Biroli, Jonathan Cohen, Andrea Crisanti, Rainer Engelken, Moritz Helias, Jonathan Kadmon, Louis Kang, Jimmy Kim, Itamar Landau, Wave Ngampruetikorn, Jeffrey Pennington, Katherine Quinn, Friedrich Schuessler, James Sethna, Julia Steinberg and Merav Stern for fruitful discussions. KK is supported by a C.V. Starr Fellowship and a CPBF Fellowship (through NSF PHY-1734030). DJS was supported by the NSF through the CPBF (PHY-1734030) and by a Simons Foundation fellowship for the MMLS. This work was partially supported by the NIH under award number R01EB026943. KK & DJS thank the Simons Institute for the Theory of Computing at U.C. Berkeley, where part of the research was conducted.

Appendix A: Details of the Dynamical Mean-Field Theory

In this section, we provide a detailed, self-contained description of the dynamical mean-field theory for the gated RNN using the Martin-Siggia-Rose-DeDominicis-Jansen formalism. We start with the equations of motion

(in vector form)

$$\tau_z \dot{z} = -z + J^z \phi_z(h), \quad (\text{A1})$$

$$\tau_r \dot{r} = -r + J^r \phi_r(h), \quad (\text{A2})$$

$$\dot{h} = \sigma_z(z) \odot (-h + [J^h(\sigma_r(r) \odot \phi_h(h))]), \quad (\text{A3})$$

where \odot stands for element-wise multiplication.

To write down the MSRDJ generating functional, let us discretise the dynamics (in the Itô convention). Note that in this convention the Jacobian is unity.

$$\begin{aligned} h_i(t+1) - h_i(t) &= \sigma_{z,i}(t) \left\{ -h_i(t) + \sum_j J_{ij}^h \sigma_{r,j}(t) \phi_j(t) + b_i^h(t) \right\} \delta t, \\ \tau_z (z_i(t+1) - z_i(t)) &= \left\{ -z_i(t) + \sum_j J_{ij}^z \phi_j(t) + b_i^z(t) \right\} \delta t, \\ \tau_r (r_i(t+1) - r_i(t)) &= \left\{ -r_i(t) + \sum_j J_{ij}^r \phi_j(t) + b_i^r(t) \right\} \delta t, \end{aligned}$$

where we have introduced external fields in the dynamics $\{b_i^h(t)\}$, $\{b_i^z(t)\}$ and $\{b_i^r(t)\}$. The generating functional is given by

$$Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = \mathbb{E} \left[\exp \left(i \sum_{j=1}^N \sum_t \hat{\mathbf{b}}_j(t)^T \mathbf{x}_j(t) \delta t \right) \right], \quad (\text{A4})$$

where $\hat{\mathbf{b}} = (\hat{b}_j^h, \hat{b}_j^z, \hat{b}_j^r)$; $\mathbf{b} = (b_j^h, b_j^z, b_j^r)$ and $\mathbf{x}_j(t) \equiv (h_j(t), z_j(t), r_j(t))$; also, the expectation is over the dynamics generated by the network. Writing this out explicitly, with δ -functions enforcing the dynamics, we get the following integral for the generating functional

$$\begin{aligned} Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] &= \int \prod_{i,t} \prod_{k,t'} \prod_{m,t''} dh_i(t) dz_k(t') dr_m(t'') \cdot \exp \left(i \left\{ \sum_{i,t} \hat{b}_i^h(t) h_i(t) + \hat{b}_i^z(t) z_i(t) + \hat{b}_i^r(t) r_i(t) \right\} \delta t \right), \\ &\times \delta \left(h_i(t+1) - h_i(t) + \left\{ h_i(t) \sigma_{z,i}(t) - \sigma_{z,i}(t) \left[\sum_j J_{ij}^h \sigma_{r,j}(t) \phi_j(t) \right] - b_i^h(t) \right\} \delta t \right), \\ &\times \delta \left(z_k(t'+1) - z_k(t') + \frac{1}{\tau_z} \left\{ z_k(t') + \sum_l J_{kl}^z \phi_l(t') + b_k^z(t') \right\} \delta t \right), \\ &\times \delta \left(r_m(t''+1) - r_m(t'') + \frac{1}{\tau_r} \left\{ r_m(t'') + \sum_n J_{mn}^r \phi_n(t'') + b_m^r(t'') \right\} \delta t \right). \end{aligned} \quad (\text{A5})$$

Now, let us introduce the Fourier representation for the δ -function; this introduces an auxiliary field variable,

which as we will see allows us to calculate the response function in the MSRDJ formalism. The generating functional can then be expressed as

$$\begin{aligned}
Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] &= \int \prod_{i,t} \prod_{k,t'} \prod_{m,t''} dh_i(t) \frac{d\hat{h}_i(t)}{2\pi} dz_k(t') \frac{d\hat{z}_k(t')}{2\pi} dr_m(t'') \frac{d\hat{r}_m(t'')}{2\pi}, \\
&\times \exp \left[-i \sum_{i,t} \hat{h}_i(t) \left(h_i(t+1) - h_i(t) - f_h(h_i, z_i, r_i) \delta t - b_i^h(t) \delta t \right) + i \sum_{i,t} \hat{b}_i^h(t) h_i(t) \delta t \right], \\
&\times \exp \left[-i \sum_{k,t'} \hat{z}_k(t') \left(z_k(t'+1) - z_k(t') - f_z(h_k, z_k) \frac{\delta t}{\tau_z} - b_k^z(t') \frac{\delta t}{\tau_z} \right) + i \sum_{k,t'} \hat{b}_k^z(t') z_k(t') \delta t \right], \\
&\times \exp \left[-i \sum_{m,t''} \hat{r}_m(t'') \left(r_m(t''+1) - r_m(t'') - f_r(h_m, r_m) \frac{\delta t}{\tau_r} - b_m^r(t'') \frac{\delta t}{\tau_r} \right) + i \sum_{m,t''} \hat{b}_m^r(t'') r_m(t'') \delta t \right],
\end{aligned} \tag{A6}$$

where the functions $f_{h,z,r}$ summarise the gated RNN dynamics

$$\begin{aligned}
f_h(h_i, z_i, r_i) &= \sigma_{z,i}(t) \left(-h_i(t) + \sum_j J_{ij}^h \sigma_{r,j}(t) \phi_j(t) \right), \\
f_z(h_k, z_k) &= -z_k(t') + \sum_l J_{kl}^z \phi_l(t'), \\
f_r(h_m, r_m) &= -r_m(t'') + \sum_n J_{mn}^r \phi_n(t'').
\end{aligned}$$

Let us now take the continuum limit $\delta t \rightarrow 0$, and formally define the measures $\mathcal{D}h_i = \lim_{\delta t \rightarrow 0} \prod_t dh_i(t)$. We can then write the generating functional as a path integral

$$\begin{aligned}
Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] &= \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \mathcal{D}z_i \mathcal{D}\hat{z}_i \mathcal{D}r_i \mathcal{D}\hat{r}_i \exp \left\{ -S[\hat{\mathbf{x}}, \mathbf{x}] \right. \\
&\quad \left. + i \int dt [\hat{\mathbf{b}}(t)^T \mathbf{x}(t) + \mathbf{b}(t)^T \hat{\mathbf{x}}(t)] \right\}
\end{aligned} \tag{A7}$$

Where $\hat{\mathbf{b}} = (\hat{b}_i^h, \hat{b}_i^z, \hat{b}_i^r)$; $\mathbf{x} = (h_i, z_i, r_i)$ and $\hat{\mathbf{x}} = (\hat{h}_i, \hat{z}_i/\tau_z, \hat{r}_i/\tau_r)$, and the action S which gives weights to the paths is given by

$$\begin{aligned}
S[\hat{\mathbf{x}}, \mathbf{x}] &= i \sum_i \int dt \hat{h}_i(t) \left[\partial_t h_i(t) - f_h(h_i, z_i, r_i) \right] \\
&\quad + i \sum_k \int dt \hat{z}_k(t) \left[\partial_t z_k(t) - \frac{f_z(h_k, z_k)}{\tau_z} \right] \\
&\quad + i \sum_m \int dt \hat{r}_m(t) \left[\partial_t r_m(t) - \frac{f_r(h_m, r_m)}{\tau_r} \right]
\end{aligned} \tag{A8}$$

The functional is properly normalised, so $Z_{\mathcal{J}}[\mathbf{0}, \mathbf{b}] = 1$. We can calculate correlation functions and response functions by taking appropriate variational derivatives of the

generating functional Z , but first we address the role of the random couplings.

Disorder Averaging:

We are interested in the typical behaviour of ensembles of the networks, so we work with the disorder-averaged generating functional \bar{Z} ; $Z_{\mathcal{J}}$ is properly normalised, so we are allowed to do this averaging on $Z_{\mathcal{J}}$. Averaging over J_{ij}^h involves the following integral

$$\begin{aligned}
&\int dJ_{ij}^h \sqrt{\frac{N}{2\pi}} \exp \left\{ -\frac{N (J_{ij}^h)^2}{2} + \right. \\
&\quad \left. i \cdot J_{ij}^h \int dt \hat{h}_i(t) \sigma_{z,i}(t) \phi_j(t) \sigma_{r,j}(t) \right\},
\end{aligned}$$

which evaluates to

$\exp \left\{ -(1/2N) \cdot \left(\int dt \hat{h}_i(t) \sigma_{z,i}(t) \phi_j(t) \sigma_{r,j}(t) \right)^2 \right\}$ and similarly for J^z and J^r we get terms

$$\begin{aligned}
&\exp \left\{ -(1/2N) \cdot \left(\int dt \hat{z}_k(t) \phi_l(t) \right)^2 \tau_z^{-2} \right\}, \\
&\exp \left\{ -(1/2N) \cdot \left(\int dt \hat{r}_m(t) \phi_n(t) \right)^2 \tau_r^{-2} \right\}.
\end{aligned}$$

The disorder-averaged generating functional is then given by

$$\begin{aligned}
\bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] &= \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \mathcal{D}z_i \mathcal{D}\hat{z}_i \mathcal{D}r_i \mathcal{D}\hat{r}_i \exp \left\{ -\bar{S}[\hat{\mathbf{x}}, \mathbf{x}] \right. \\
&\quad \left. + i \int dt [\hat{\mathbf{b}}(t)^T \mathbf{x}(t) + \mathbf{b}(t)^T \hat{\mathbf{x}}(t)] \right\}
\end{aligned} \tag{A9}$$

where the disorder-averaged action \bar{S} is given by

$$\begin{aligned}
\bar{S}[\hat{\mathbf{x}}, \mathbf{x}] &= i \sum_i \int dt \hat{h}_i(t) \left(\partial_t h_i(t) + h_i(t) \sigma_{z,i}(t) \right) + \frac{1}{2N} \sum_{i,j} \left(\int dt \hat{h}_i(t) \sigma_{z,i}(t) \phi_j(t) \sigma_{r,j}(t) \right)^2 \\
&+ i \sum_k \int dt \hat{z}_k(t) \left(\partial_t z_k(t) + \frac{z_k(t)}{\tau_z} \right) + \frac{1}{2N} \sum_{k,l} \left(\int dt \frac{\hat{z}_k(t)}{\tau_z} \cdot \phi_l(t) \right)^2 \\
&+ i \sum_m \int dt \hat{r}_m(t) \left(\partial_t r_m(t) + \frac{r_m(t)}{\tau_r} \right) + \frac{1}{2N} \sum_{m,n} \left(\int dt \frac{\hat{r}_m(t)}{\tau_r} \cdot \phi_n(t) \right)^2
\end{aligned} \tag{A10}$$

With some foresight, we see the action is extensive in the system size, and we can try to reduce it to a single-site description. However, the issue now is that we have non-local terms (e.g. involving both i and j), and we can introduce the following auxiliary fields to decouple these non-local terms

$$\begin{aligned}
C_{\phi\sigma_r}(t, t') &:= \frac{1}{N} \sum_i \phi_i(t) \phi_i(t') \sigma_{r,i}(t) \sigma_{r,i}(t'), \\
C_\phi(t, t') &:= \frac{1}{N} \sum_k \phi_k(t) \phi_k(t').
\end{aligned} \tag{A11}$$

To make the C 's free fields that we integrate over, we enforce these relations using the Fourier representation of δ functions with additional auxiliary fields:

$$\begin{aligned}
\delta\left(NC_{\phi\sigma_r}(t, t') - \sum_i \phi_i(t) \phi_i(t') \sigma_{r,i}(t) \sigma_{r,i}(t') \right) &= \\
&\int \frac{N}{\pi} d\widehat{C}_{\phi\sigma_r}(t, t') \exp\left[-\frac{i}{2} \widehat{C}_{\phi\sigma_r}(t, t') \left(N \cdot C_{\phi\sigma_r}(t, t') \right. \right. \\
&\quad \left. \left. - \sum_i \phi_i(t) \phi_i(t') \sigma_{r,i}(t) \sigma_{r,i}(t') \right) \right] \\
\delta\left(NC_\phi(t, t') - \sum_k \phi_k(t) \phi_k(t') \right) &= \int \frac{N}{\pi} d\widehat{C}_\phi(t, t') \\
&\exp\left[-\frac{i}{2} \widehat{C}_\phi(t, t') \left(N \cdot C_\phi(t, t') - \sum_k \phi_k(t) \phi_k(t') \right) \right]
\end{aligned}$$

this allows us to make the following transformations to decouple the non-local terms in the action \bar{S}

$$\begin{aligned}
\frac{1}{2N} \sum_{i,j} \left(\hat{h}_i(t) \sigma_{z,i}(t) \phi_j(t) \sigma_{r,j}(t) \right)^2 &\longrightarrow \\
\frac{1}{2} \sum_i \int dt dt' \hat{h}_i(t) \sigma_{z,i}(t) C_{\phi\sigma_r}(t, t') \hat{h}_i(t') \sigma_{z,i}(t') & \\
\frac{1}{2N} \sum_{k,l} \left(\int dt \frac{\hat{z}_k(t)}{\tau_z} \cdot \phi_l(t) \right)^2 &\longrightarrow \\
\frac{1}{2} \sum_k \int dt dt' \frac{\hat{z}_k(t)}{\tau_z} C_\phi(t, t') \frac{\hat{z}_k(t')}{\tau_z} &
\end{aligned}$$

$$\begin{aligned}
\frac{1}{2N} \sum_{m,n} \left(\int dt \frac{\hat{r}_m(t)}{\tau_r} \cdot \phi_n(t) \right)^2 &\longrightarrow \\
\frac{1}{2} \sum_m \int dt dt' \frac{\hat{r}_m(t)}{\tau_r} C_\phi(t, t') \frac{\hat{r}_m(t')}{\tau_r} &
\end{aligned}$$

We see clearly that the $C_{\phi\sigma_r}$ and C_ϕ auxiliary fields which represent the (population averaged) $\phi\sigma_r - \phi\sigma_r$ and $\phi - \phi$ correlation functions have decoupled the sites by summarising all the information present in the rest of the network in terms of two-point functions; two different sites interact only by means of the correlation functions. The disorder-averaged generating functional can now be written as

$$\begin{aligned}
\bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] &= \int \mathcal{D}\widehat{\mathbf{C}} \mathcal{D}\mathbf{C} \exp\left(-N \cdot \mathcal{L}[\widehat{\mathbf{C}}, \mathbf{C}; \hat{\mathbf{b}}, \mathbf{b}] \right) \tag{A12} \\
\mathcal{L} &= \frac{i}{2} \int dt dt' \left[\mathbf{C}(t, t')^T \widehat{\mathbf{C}}(t, t') \right] - W[\widehat{\mathbf{C}}, \mathbf{C}; \hat{\mathbf{b}}, \mathbf{b}] \\
\exp(N \cdot W) &= \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \mathcal{D}z_i \mathcal{D}\hat{z}_i \mathcal{D}r_i \mathcal{D}\hat{r}_i \times \\
&\exp\left\{ i \int dt \left[\mathbf{b}(t)^T \hat{\mathbf{h}}(t) + \hat{\mathbf{b}}(t)^T \mathbf{h}(t) \right] - S_d[\hat{\mathbf{h}}, \mathbf{h}; \{\mathbf{C}, \widehat{\mathbf{C}}\}] \right\}
\end{aligned}$$

where $\mathbf{C} = (C_h, C_z, C_r)$ and $\widehat{\mathbf{C}} = (\widehat{C}_h, \widehat{C}_z, \widehat{C}_r)$. The site-wise decoupled action, S_d only contains terms involving a single site (and the C fields). So, for a given value of $\widehat{\mathbf{C}}$ and \mathbf{C} , the different sites are decoupled and driven by the site-wise action

$$\begin{aligned}
S_d[\hat{\mathbf{h}}, \mathbf{h}; \{\mathbf{C}, \widehat{\mathbf{C}}\}] &= i \int dt \left[\hat{\mathbf{h}}(t)^T \partial_t \mathbf{h}(t) + \hat{\mathbf{h}}_\tau(t)^T \mathbf{h}(t) \right] \\
&+ \frac{1}{2} \int dt dt' \hat{\mathbf{h}}_\tau(t)^T \mathbb{D}\mathbf{C}(t, t') \hat{\mathbf{h}}_\tau(t') \\
&- \frac{i}{2} \int dt dt' \mathbf{S}_x(t)^T \mathbb{D}\widehat{\mathbf{C}}(t, t') \mathbf{S}_x(t')
\end{aligned} \tag{A13}$$

where

$$\begin{aligned}
\hat{\mathbf{h}}_\tau(t) &= (\hat{h}_i \sigma_{z,i}, \hat{z}_i / \tau_z, \hat{r}_i / \tau_r) \\
\hat{\mathbf{h}}(t) &= (\hat{h}_i, \hat{z}_i, \hat{r}_i) \\
\mathbf{S}_x &= (\phi_i \sigma_{r,i}, \phi_i, \phi_i) \\
\mathbb{D}\mathbf{C}(t, t') &= \text{Diag}[C_{\phi\sigma_r}(t, t'), C_\phi(t, t'), C_\phi(t, t')] \\
\mathbb{D}\widehat{\mathbf{C}}(t, t') &= \text{Diag}[\widehat{C}_{\phi\sigma_r}(t, t'), \widehat{C}_\phi(t, t'), \widehat{C}_\phi(t, t')]
\end{aligned}$$

Saddle-point approximation for $N \rightarrow \infty$

So far, we have not made any use of the fact that we are considering large networks. However, noting that N appears in the exponent in the expression for the disorder-averaged generating functional, we can approximate it using a saddle-point approximation:

$$\overline{Z[\hat{\mathbf{b}}, \mathbf{b}]} \simeq e^{N \cdot \mathcal{L}_0[\hat{\mathbf{b}}, \mathbf{b}; \mathbf{C}_0, \mathbf{C}_0]} \int \mathcal{D}\mathbf{Q} \mathcal{D}\mathbf{Q} e^{-N \cdot \mathcal{L}_2[\mathbf{Q}, \mathbf{Q}, \mathbf{b}, \mathbf{b}]}$$

we have approximated the action \mathcal{L} in eq. A12 by its saddle-point value plus a Hessian term: $\mathcal{L} \simeq \mathcal{L}_0 + \mathcal{L}_2$ and the \mathbf{Q}, \mathbf{Q} fields represent Gaussian fluctuations about the saddle-point values $\mathbf{C}_0, \widehat{\mathbf{C}}_0$, respectively. At the saddle-point the action is stationary w.r.t variations, thus, the saddle-point values of C fields satisfy

$$\begin{aligned} C_{0_{\phi\sigma_r}}(t, t') &= \frac{1}{N} \sum_{i=1}^N \langle \phi_i(t) \sigma_{r,i}(t) \phi_i(t') \sigma_{r,i}(t') \rangle_0 \\ \widehat{C}_{0_{\phi\sigma_r}}(t, t') &= \frac{1}{N} \sum_{i=1}^N \langle \widehat{h}_i(t) \sigma_{z,i}(t) \widehat{h}_i(t') \sigma_{z,i}(t') \rangle_0 \\ &= \frac{\delta^2 \langle \sigma_{z,i}(t) \sigma_{z,i}(t') \rangle_0}{\delta b_i(t) \delta b_i(t')} = 0 \\ C_{0_\phi}(t, t') &= \frac{1}{N} \sum_{k=1}^N \langle \phi_k(t) \phi_k(t') \rangle_0 \\ \widehat{C}_{0_\phi}(t, t') &= 0 \end{aligned} \quad (\text{A14})$$

In evaluating the saddle-point correlation function in the second line, we have used the fact that equal-time response functions in the Itô convention are zero [29]. This is perhaps the first significant point of departure from previous studies of disordered neural networks, and forces us to confront the multiplicative nature of the z -gate. Here $\langle \dots \rangle_0$ denotes averages w.r.t paths generated by the saddle-point action, thus these equations are a self-consistency constraint. With the correlation fields fixed at their saddle-point values, if we neglect the contribution of the fluctuations (i.e. ignore \mathcal{L}_2) then the generating functional is given by a product of identical site-wise generating functionals (c.f. eq. (8) in main text).

$$\overline{Z[\hat{\mathbf{b}}, \mathbf{b}]} = Z_0[\hat{\mathbf{b}}, \mathbf{b}]^N \quad (\text{A15})$$

where the site-wise functionals are given by

$$\begin{aligned} Z_0[\hat{\mathbf{b}}, \mathbf{b}] &= \int \mathcal{D}h \mathcal{D}\hat{h} \mathcal{D}z \mathcal{D}\hat{z} \mathcal{D}r \mathcal{D}\hat{r} \times \\ &e^{\left(i \int dt [\mathbf{b}(t)^T \hat{\mathbf{h}}(t) + \hat{\mathbf{b}}(t)^T \mathbf{h}(t)] - S_d[\mathbf{h}, \mathbf{h}; \{\mathbf{C}_0, \mathbf{0}\}] \right)} \end{aligned} \quad (\text{A16})$$

where $\mathbf{C}_0 = (C_{0_{\phi\sigma_r}}, C_{0_\phi})$.

The site-wise decoupled action is now quadratic with the correlation functions taking on their saddle-point values. This corresponds to an action for each site containing three scalar variables driven by Gaussian processes. This can be seen explicitly by using a Hubbard-Stratonovic transform which will make the action linear at the cost of introducing three auxiliary Gaussian fields

η_h, η_z and η_r with correlation functions $C_{0_{\phi\sigma_r}}(t, t')$, $C_{0_\phi}(t, t')$ and $C_{0_\phi}(t, t')$, respectively. With this transformation, the action for each site corresponds to stochastic dynamics for three scalar variables given by (same as eq. (11) in the main text):

$$\begin{aligned} \tau_z \dot{z}(t) &= -z(t) + \eta_z(t), \\ \tau_r \dot{r}(t) &= -r(t) + \eta_r(t), \\ \dot{h}(t) &= -\sigma_z(z) \cdot h(t) + \sigma_z(z) \cdot \eta_h(t). \end{aligned}$$

The Gaussian processes, η_h, η_z, η_r have correlation functions given by:

$$\begin{aligned} \langle \eta_h(t) \cdot \eta_h(t') \rangle &= \langle \phi(t) \sigma_r(t) \cdot \phi(t') \sigma_r(t') \rangle, \\ \langle \eta_z(t) \cdot \eta_z(t') \rangle &= \langle \phi(t) \cdot \phi(t') \rangle, \\ \langle \eta_r(t) \cdot \eta_r(t') \rangle &= \langle \phi(t) \cdot \phi(t') \rangle. \end{aligned}$$

The intuitive picture of the saddle-point approximation is as follows: the sites of the full network become decoupled, and they are each driven by a Gaussian processes whose correlation functions summarise the activity of the rest of the network ‘felt’ by each site. It is possible to argue about the final result heuristically, but one does not have access to the systematic corrections that a field theory formulation affords.

Appendix B: Random matrix theory for spectrum of the Jacobian

In this section, we provide details of calculating the bounding curve for the Jacobian spectrum for both fixed-points and time-varying states. Our starting point is the matrix Green’s function for the Hermitian problem:

$$\mathcal{G}(\eta, \lambda, \bar{\lambda}) = \mathbb{E} \left[(\eta \mathbb{1}_{6N} + H)^{-1} \right], \quad (\text{B1})$$

$$H = \begin{pmatrix} 0 & \lambda - \mathcal{D} \\ \bar{\lambda} - \mathcal{D}^T & 0 \end{pmatrix}. \quad (\text{B2})$$

We can now use standard diagrammatic techniques in random matrix theory to evaluate \mathcal{G} using the self-consistent Born approximation (SCBA), which is exact in the limit $N \rightarrow \infty$. We first decompose the Jacobian into structured (A, L, R) and random parts (\mathcal{J}):

$$\begin{aligned} \mathcal{D} &= \underbrace{\begin{pmatrix} -[\sigma_z] & \mathbb{D} & 0 \\ 0 & -\tau_z^{-1} \mathbb{1} & 0 \\ 0 & 0 & -\tau_r^{-1} \mathbb{1} \end{pmatrix}}_A + \underbrace{\begin{pmatrix} [\sigma_z] & 0 & 0 \\ 0 & \tau_z^{-1} \mathbb{1} & 0 \\ 0 & 0 & \tau_r^{-1} \mathbb{1} \end{pmatrix}}_L \\ &\times \underbrace{\begin{pmatrix} J^h & 0 & 0 \\ 0 & J^z & 0 \\ 0 & 0 & J^r \end{pmatrix}}_{\mathcal{J}} \underbrace{\begin{pmatrix} [\phi' \sigma_r] & 0 & [\phi \sigma_r'] \\ [\phi'] & 0 & 0 \\ [\phi'] & 0 & 0 \end{pmatrix}}_R. \end{aligned} \quad (\text{B3})$$

In the SCBA, \mathcal{G} is given by a Dyson equation with the self-energy functional $\Sigma[\mathcal{G}]$

$$\mathcal{G}^{-1} = \mathcal{G}_0^{-1} - \Sigma[\mathcal{G}], \quad (\text{B4})$$

where the matrices on the right are defined in terms of $3N \times 3N$ blocks:

$$\mathcal{G}_0^{-1} = \begin{pmatrix} \eta \mathbb{1} & \lambda - A \\ \bar{\lambda} - A^T & \eta \mathbb{1} \end{pmatrix}, \quad (\text{B5})$$

$$\Sigma[\mathcal{G}] = \begin{pmatrix} LQ[RG_{22}R^T]L & 0 \\ 0 & R^TQ[L^T\mathcal{G}_{11}L]R \end{pmatrix}, \quad (\text{B6})$$

and Q is a superoperator which acts on its argument as follows:

$$Q[M] = \begin{pmatrix} \frac{1}{N} \text{Tr} M_{11} & 0 & 0 \\ 0 & \frac{1}{N} \text{Tr} M_{22} & 0 \\ 0 & 0 & \frac{1}{N} \text{Tr} M_{33} \end{pmatrix}. \quad (\text{B7})$$

At this point, we have presented all of the necessary ingredients for computing the Green's function, and thus determining the spectral properties of the Jacobian. These are the Dyson equation (B4), along with the free Green's function (B5) and the self-energy (B6). Most of what is left is complicated linear algebra. However, in the interest of completeness, we will proceed to unpack these equations and give a detailed derivation of the main equation of interest, the bounding curve of the spectral density.

To proceed further, let us define a finer block structure in terms of $N \times N$ blocks as follows:

$$\tilde{\mathcal{G}}_{11} \equiv L^T \mathcal{G}_{11} L = \begin{pmatrix} \tilde{G}_{11} & \tilde{G}_{12} & \tilde{G}_{13} \\ \tilde{G}_{21} & \tilde{G}_{22} & \tilde{G}_{23} \\ \tilde{G}_{31} & \tilde{G}_{32} & \tilde{G}_{33} \end{pmatrix}, \quad (\text{B8})$$

$$\tilde{\mathcal{G}}_{22} \equiv R \mathcal{G}_{22} R^T = \begin{pmatrix} \tilde{G}_{44} & \tilde{G}_{45} & \tilde{G}_{46} \\ \tilde{G}_{54} & \tilde{G}_{55} & \tilde{G}_{56} \\ \tilde{G}_{64} & \tilde{G}_{65} & \tilde{G}_{66} \end{pmatrix}, \quad (\text{B9})$$

and the traces of these block matrices as

$$\tilde{g}_{ij} = \frac{1}{N} \text{Tr}[\tilde{G}_{ij}]. \quad (\text{B10})$$

Then the blocks of the self-energy matrix in eq. B6 are given by

$$\Sigma_{11} = \begin{pmatrix} [\sigma_z^2] \tilde{g}_{44} & 0 & 0 \\ 0 & \tau_z^{-2} \tilde{g}_{55} & 0 \\ 0 & 0 & \tau_r^{-2} \tilde{g}_{66} \end{pmatrix}, \quad (\text{B11})$$

$$\Sigma_{22} = \begin{pmatrix} [\phi' \sigma_r]^2 \tilde{g}_{11} + [\phi']^2 (\tilde{g}_{22} + \tilde{g}_{33}) & 0 & [\phi' \sigma_r][\phi \sigma_r'] \tilde{g}_{11} \\ 0 & 0 & 0 \\ [\phi' \sigma_r][\phi \sigma_r'] \tilde{g}_{11} & 0 & [\phi \sigma_r']^2 \tilde{g}_{11} \end{pmatrix}. \quad (\text{B12})$$

These equations allow us to define self-consistent relations between the various \tilde{g}_{ii} . First, using the Dyson equation we invert the block matrix to get

$$\mathcal{G} = \begin{pmatrix} \eta - \Sigma_{11} & \lambda - A \\ \bar{\lambda} - A^T & \eta - \Sigma_{22} \end{pmatrix}^{-1}. \quad (\text{B13})$$

The R.H.S of eq. B13 is a function of \tilde{g}_{ii} , and then we can read off the diagonal blocks to impose the self consistent relations using

$$\tilde{\mathcal{G}}_{11} = L^T \mathcal{G}_{11} L \quad \tilde{\mathcal{G}}_{22} = R \mathcal{G}_{22} R^T, \quad (\text{B14})$$

$$\tilde{g}_{ii} = \frac{1}{N} \text{Tr}[\tilde{G}_{ii}] \quad i \in 1, \dots, 6. \quad (\text{B15})$$

To simplify things we write the blocks of the Greens functions, \tilde{G}_{ii} , as the ratio of a numerator and a denominator that is common to all the blocks. Let

$$r = [\phi' \sigma_r], \quad q = [\phi \sigma_r'], \quad a = [\phi'], \quad (\text{B16})$$

and define

$$A = \tilde{g}_{11} |\lambda \tau_r + 1|^2 r^2 + (\tilde{g}_{22} + \tilde{g}_{33}) a^2 C, \quad (\text{B17})$$

$$B = \mathbb{D}^2 \tilde{g}_{55} + |\lambda \tau_z + 1|^2 [\sigma_z^2] \tilde{g}_{44}, \quad (\text{B18})$$

$$C = |\lambda \tau_r + 1|^2 - \tilde{g}_{11} \tilde{g}_{66} q^2. \quad (\text{B19})$$

Then the denominator can be written

$$\Gamma = |\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 C - AB. \quad (\text{B20})$$

Now we look at the numerator terms Γ_i for all the Green's function blocks \tilde{g}_{ii} :

$$\Gamma_1 = \sigma_z^2 |\lambda \tau_z + 1|^2 A, \quad (\text{B21})$$

$$\Gamma_2 = \mathbb{D}^2 A, \quad (\text{B22})$$

$$\Gamma_3 = \tilde{g}_{11} |\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 q^2 - \tilde{g}_{11} (\tilde{g}_{22} + \tilde{g}_{33}) a^2 q^2 B, \quad (\text{B23})$$

$$\Gamma_4 = \tilde{g}_{66} |\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 q^2 + (|\lambda \tau_r + 1|^2 r^2 - \tilde{g}_{66} (\tilde{g}_{22} + \tilde{g}_{66}) a^2 q^2) B, \quad (\text{B24})$$

$$\Gamma_5 = \Gamma_6 = a^2 C B. \quad (\text{B25})$$

The numerators and denominator are all diagonal matrices with real entries. The self-consistent equations implied by eqns. B14,B15 for the six scalar variables \tilde{g}_{ii} are then given by

$$\tilde{g}_{ii} = \left\langle \frac{\Gamma_i}{\Gamma} \right\rangle, \quad (\text{B26})$$

where we denote $\langle M \rangle \equiv N^{-1} \text{Tr} M$ for shorthand, and i runs from 1 to 6. Solving these equations give us the \tilde{g}_{ii} as implicit functions of λ and are in general complicated. However, the situation simplifies considerably when we are looking for the spectral curve. In this case, we are looking for all $\lambda \in \mathbb{C}$ that satisfy the self-consistent equations with $\tilde{g}_{ii} = 0$.

It is important to note that on the spectral curve even though the \tilde{g}_{ii} vanish, we encounter ratios like $\tilde{g}_{22}/\tilde{g}_{11} \equiv x_2$ and $\tilde{g}_{33}/\tilde{g}_{11} \equiv x_3$ which remain finite. Therefore, to get the equation for the spectral curve we look at the numerator terms in the limit of small \tilde{g}_{ii} and eliminating the ratios mentioned above to get an implicit equation in the state variables and λ . The denominator simplifies to

$$\mathbf{\Gamma} = |1 + \tau_r \lambda|^2 |1 + \tau_z \lambda|^2 |\sigma_z + \lambda|^2. \quad (\text{B27})$$

In this limit, we may take the trace and keeping non-vanishing. In particular, the expression for \tilde{g}_{11} gives

$$\begin{aligned} 1 &= \frac{1}{\tilde{g}_{11}} \frac{1}{N} \text{Tr} \frac{\mathbf{\Gamma}_1}{\mathbf{\Gamma}} \\ &= \left\langle \frac{\sigma_z^2 |\lambda \tau_z + 1|^2 (\mathbb{D}^2 (|\lambda \tau_r + 1|^2 r^2 + (x_2 + x_3) a^2 C))}{\mathbf{\Gamma}} \right\rangle \end{aligned} \quad (\text{B28})$$

And similarly for \tilde{g}_{22} and \tilde{g}_{33}

$$x_2 = \left\langle \frac{\mathbb{D}^2 (|\lambda \tau_r + 1|^2 r^2 + (x_2 + x_3) a^2 C)}{\mathbf{\Gamma}} \right\rangle, \quad (\text{B29})$$

$$x_3 = \left\langle \frac{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 q^2}{\mathbf{\Gamma}} \right\rangle. \quad (\text{B30})$$

Let us define now

$$\gamma_1 = \left\langle \frac{\sigma_z^2 r^2}{|\lambda + \sigma_z|^2} \right\rangle, \quad (\text{B31})$$

$$\gamma_2 = \left\langle \frac{a^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle, \quad (\text{B32})$$

$$\gamma_3 = \left\langle \frac{\mathbb{D}^2 r^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle, \quad (\text{B33})$$

$$\gamma_4 = \left\langle \frac{\mathbb{D}^2 a^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle, \quad (\text{B34})$$

$$\gamma_5 = \frac{\langle q^2 \rangle}{|\lambda \tau_r + 1|^2}. \quad (\text{B35})$$

Then we must solve the system of equations

$$1 = \gamma_1 + (x_2 + x_3) \gamma_2, \quad (\text{B36})$$

$$x_2 = \gamma_3 + \gamma_4 (x_2 + x_3), \quad (\text{B37})$$

$$x_3 = \gamma_5. \quad (\text{B38})$$

This system only has a solution when

$$(1 - \gamma_1)(1 - \gamma_4) = \gamma_2 (\gamma_3 + \gamma_5). \quad (\text{B39})$$

This expression depends on λ . Thus, the $\lambda \in \mathbb{C}$ for which it is satisfied describes the boundary of the support of the eigenvalue spectrum.

We get the spectral curve using eq. B26 by eliminating \tilde{g}_{ii} , and it is given by all $\lambda \in \mathbb{C}$ that satisfy

$$\begin{aligned} &\left\{ 1 - \left\langle \frac{r^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle \right\} \left\{ 1 - \left\langle \frac{\mathbb{D}^2 a^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle \right\} \\ &= \left\langle \frac{\sigma_z^2 a^2}{|\lambda + \sigma_z|^2} \right\rangle \left\{ \left\langle \frac{\mathbb{D}^2 r^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle + \frac{\langle q^2 \rangle}{|\lambda \tau_r + 1|^2} \right\}. \end{aligned} \quad (\text{B40})$$

For fixed-points, we have $\mathbb{D} = 0$, which makes $\gamma_3 = \gamma_4 = 0$, thus we get the simpler equation for the spectral curve given in the main text (eq. 29):

$$1 = \left\langle \frac{r^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle + \frac{\langle q^2 \rangle}{|\lambda \tau_r + 1|^2} \left\langle \frac{\sigma_z^2 a^2}{|\lambda + \sigma_z|^2} \right\rangle. \quad (\text{B41})$$

1. Jacobian spectrum for the case $\alpha_r = 0$

In the case when $\alpha_r = 0$, it is possible to express the Green's function (eq. 27) in a simpler form. Recall that,

$$G(\lambda, \bar{\lambda}) = \lim_{\eta \rightarrow i0^+} \frac{1}{3N} \text{tr} \mathcal{G}_{21}(\eta, \lambda, \bar{\lambda}). \quad (\text{B42})$$

Let $\tilde{B} = \mathbb{D}^2 + \sigma_r^2 \sigma_z^2 |\lambda \tau_z + 1|^2$. Then the Green's function is given by

$$G(\lambda, \bar{\lambda}) = \frac{1}{3} \left\langle \frac{|\lambda \tau_z + 1|^2 (\bar{\lambda} + \sigma_z)}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) a^2 \tilde{B}} \right\rangle, \quad (\text{B43})$$

$$+ \frac{1}{3} \left\langle \frac{(\bar{\lambda} + \tau_z^{-1}) (|\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) a^2 \sigma_z^2)}{|\lambda + \tau_z^{-1}|^2 |\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) a^2 \tilde{B}} \right\rangle, \quad (\text{B44})$$

$$+ \frac{1}{3} \frac{1}{\lambda + \tau_r^{-1}}, \quad (\text{B45})$$

where $\xi(\lambda, \bar{\lambda})$ is defined implicitly to satisfy the equation

$$1 = \left\langle \frac{a^2 \tilde{B}}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) a^2 \tilde{B}} \right\rangle. \quad (\text{B46})$$

The function $\xi(\lambda, \bar{\lambda})$ acts as a sort of order parameter for the spectral density, indicating the transition on the complex plane between zero and finite density μ . Outside the spectral support, $\lambda \in \Sigma^c$, this order parameter vanishes $\xi = 0$ and the Green's function is holomorphic

$$G(\lambda, \bar{\lambda}) = \frac{1}{3} \left(\left\langle \frac{1}{|\lambda + \sigma_z|} \right\rangle + \frac{1}{\lambda + \tau_z^{-1}} + \frac{1}{\lambda + \tau_r^{-1}} \right), \quad (\text{B47})$$

which of course indicates that the density is zero since $\mu(\lambda) = \partial_{\bar{\lambda}} G(\lambda, \bar{\lambda})$. Inside the support $\lambda \in \Sigma$, the order parameter $\xi \neq 0$, and the Green's function consequently picks up non-analytic contributions, proportional to $\bar{\lambda}$. Since the Green's function is continuous on the complex plane, it must be continuous across the boundary of the spectral support. This must then occur precisely when the holomorphic solution meets the non-analytic solution, at $\xi = 0$. This is the condition used to find the boundary curve above.

Appendix C: Spectral clumping and pinching in the limit $\alpha_z \rightarrow \infty$

In this section we provide details on the accumulation of eigenvalues near zero and the *pinching* of the leading

spectral curve (for certain values of g_h) as the update gate becomes switch-like ($\alpha_z \rightarrow \infty$). To focus on the key aspects of these phenomena, we consider the case when the reset gate is off and there are no biases ($\alpha_r = 0, \beta_{r,h,z} = 0$). Moreover, we consider a piece-wise linear approximation – sometimes called “hard” tanh – to the tanh function given by

$$\phi_{lin}(x) = \begin{cases} 1 & x > 1/g_h, \\ g_h x & |x| \leq 1/g_h, \\ -1 & x < -1/g_h. \end{cases} \quad (\text{C1})$$

This approximation does not qualitatively change the nature of the clumping.

In the limit, $\alpha_z \rightarrow \infty$ the update gate σ_z becomes binary with a distribution given by

$$P(\sigma_z = x) = f_z \delta(x - 1) + (1 - f_z) \delta(x), \quad (\text{C2})$$

where $f_z = \langle \sigma_z \rangle$ is the fraction of update gates that are open (i.e. equal to one). Using this, along with the assumption that $\mathbb{D} \approx 0$ – which is valid in this regime – we can simplify the expression for the Green’s function (eq. B43 - B47) to yield:

$$G(\lambda, \bar{\lambda}) = \frac{1 - f_z}{\lambda} + f_z (1 - f_h) \frac{1}{\lambda + 1} + \frac{1}{\lambda + \tau_z^{-1}} + \frac{(1 + \bar{\lambda})}{g_h^2 \sigma(\beta_r)^2} \mathbb{I}_{\{|\lambda| < g_h^2 \sigma(\beta_r)^2\}}, \quad (\text{C3})$$

where f_h is the fraction of hard tanh activations that are not saturated. In the limit of small τ_z and $\beta_r = 0$, we get the expression for the density given in the text:

$$\mu(\lambda) = (1 - f_z) \delta(\lambda) + f_z (1 - f_h) \delta(\lambda + 1) + \frac{4}{\pi g_h^2} \mathbb{I}_{\{|\lambda| \leq g_h^2/4\}}. \quad (\text{C4})$$

Thus we see an extensive number of eigenvalues near zero.

Now, let us study the regime where α_z is large but not infinite. We would like to get the scaling behaviour of the leading edge of the spectrum and the density of eigenvalues contained in a radius δ around the origin. We make an ansatz for the spectral edge close to zero $\lambda \sim e^{-c\alpha_z \sqrt{\Delta_h}}$, where c is a positive constant. With this ansatz, the equation for the spectral curve reads

$$\int \mathcal{D}z \frac{\sigma_z(\sqrt{\Delta_z} \cdot z)^2}{\left| \lambda_0 e^{-c\alpha_z \sqrt{\Delta_h}} + \sigma_z(\sqrt{\Delta_z}^2) \right|^2} = \frac{\sigma_r(\beta_r)^{-2}}{\langle \phi'(\sqrt{\Delta_h} \cdot h)^2 \rangle}. \quad (\text{C5})$$

In the limit of large α_z and $\beta_r = 0$ this implies

$$\text{erfc}\left(\frac{c}{\sqrt{2}}\right) \approx \frac{4}{\langle \phi'(\sqrt{\Delta_h} \cdot h)^2 \rangle}. \quad (\text{C6})$$

If this has a positive solution for c , then the scaling of the spectral edge as $\lambda \sim e^{-c\alpha_z \sqrt{\Delta_h}}$ holds. Moreover, whenever

there is a positive solution for c we also expect pinching of the spectral curve and in the limit $\alpha_z \rightarrow \infty$ we will have marginal stability.

Under the same approximation, we can approximate the eigenvalue density in a radius δ around zero as

$$P(|\lambda(\mathcal{D})| < \delta) = \frac{1}{2\pi i} \oint_{\mathcal{C}} dz G(1 + z), \quad (\text{C7})$$

where we choose the contour along $z = e^{-c\alpha_z \sqrt{\Delta_h} + i\theta}$ for $\theta \in [0, 2\pi)$ and $\delta = e^{-c\alpha_z \sqrt{\Delta_h}}$. In the limit of large α_z (thus $\delta \ll 1$) we get the scaling form described in the main text :

$$P(|\lambda(\mathcal{D})| < \delta) \approx \frac{1}{2} \text{erfc}\left(-\frac{\log(\delta)}{\alpha_z \sqrt{2\Delta_h}}\right). \quad (\text{C8})$$

Appendix D: Perturbative solutions to the mean-field equations

1. Perturbative solutions for the fixed-point variance Δ_h with biases

In this section, we derive the perturbative solutions for the fixed-point variance Δ_h with finite biases, near the critical point where the zero fixed-point becomes unstable. Recall, that fixed-point variances are obtained by solving

$$\Delta_z \equiv \langle z^2 \rangle = \int \mathcal{D}x \phi\left(\sqrt{\Delta_h} x\right)^2 = \Delta_r \quad (\text{D1})$$

$$\Delta_h \equiv \langle h^2 \rangle = \int \mathcal{D}x \mathcal{D}y \phi\left(\sqrt{\Delta_h} x\right)^2 \sigma_r(\sqrt{\Delta_r} y)^2 \quad (\text{D2})$$

The expansion we seek is perturbative in Δ_h . So, expanding the gating and activating functions about their biases under the assumption $\Delta_r \approx g_h^2 \Delta_h$, we have a series expansion to $O(\Delta_h^2)$

$$\left\langle \sigma_r(\sqrt{\Delta_r} x)^2 \right\rangle_x = a_0 + a_1 g_h^2 \Delta_h + a_2 g_h^4 \Delta_h^2$$

$$a_0 = \frac{1}{4} [1 + \phi_0(\beta_r/2)]^2 \quad (\text{D3})$$

$$a_1 = \frac{\alpha_r^2}{16} \left[\phi_0^{(1)}(\beta_r/2)^2 + \phi_0(\beta_r/2) \phi_0^{(2)}(\beta_r/2) + \phi_0^{(2)}(\beta_r/2) \right] \quad (\text{D4})$$

$$a_2 = \frac{\alpha_r^4}{256} \left[12 \phi_0^{(2)}(\beta_r/2)^2 + 4 \phi_0(\beta_r/2) \phi_0^{(4)}(\beta_r/2) + 16 \phi_0^{(1)}(\beta_r/2) \phi_0^{(3)}(\beta_r/2) + \phi_0^{(4)}(\beta_r/2) \right] \quad (\text{D5})$$

where we have used the following identities involving

the derivatives of \tanh :

$$\phi_0(x) = \tanh(x) \quad (\text{D6})$$

$$\phi_0^{(1)}(x) = 1 - \phi_0(x)^2 \quad (\text{D7})$$

$$\phi_0^{(2)}(x) = -2\phi_0(x)(1 - \phi_0(x)^2) \quad (\text{D8})$$

$$\phi_0^{(3)}(x) = 2(1 - \phi_0(x)^2)(3\phi_0(x)^2 - 1) \quad (\text{D9})$$

$$\phi_0^{(4)}(x) = -8\phi_0(x)(1 - \phi_0(x)^2)(3\phi_0(x)^2 - 2) \quad (\text{D10})$$

This gives us to $O(\Delta_h^2)$

$$\Delta_h \approx [c_0 + c_1\Delta_h + c_2\Delta_h^2] \left\langle \sigma_r(\sqrt{\Delta_r}x)^2 \right\rangle_x \quad (\text{D11})$$

$$c_0 = \phi_0(\beta_h)^2 \quad (\text{D12})$$

$$c_1 = g_h^2 \left[\phi_0^{(1)}(\beta_h)^2 + \phi_0^{(2)}(\beta_h)\phi_0(\beta_h) \right] \quad (\text{D13})$$

$$c_2 = g_h^4 \left[\frac{1}{4}\phi_0(\beta_h)\phi_0^{(4)}(\beta_h) + \phi_0^{(1)}(\beta_h)\phi_0^{(3)}(\beta_h) + \frac{3}{4}\phi_0^{(2)}(\beta_h)^2 \right] \quad (\text{D14})$$

and therefore,

$$\Delta_h \approx (c_0 + c_1\Delta_h + c_2\Delta_h^2) (a_0 + a_1g_h^2\Delta_h + a_2g_h^4\Delta_h^2) \quad (\text{D15})$$

To proceed further, we study the solutions to this equation for small deviations for a critical value of g_h . Which critical value should we use? Recall, that the zero fixed-point becomes unstable when

$$-1 + \phi'(0)\sigma_r(0) = 0 \quad (\text{D16})$$

Therefore, we expand around this operating point and our small parameter $\epsilon = g_h - g_c$ where $g_c = \sigma_r(0)^{-1}$. We make an ansatz that we can express Δ_h as a power series in ϵ ,

$$\Delta_h = \epsilon^\eta (d_0 + d_1\epsilon + d_2\epsilon^2) \quad (\text{D17})$$

where η is the exponent for the prefactor scaling, and needs to be determined self-consistently. To get the scaling relations for Δ_h we need to expand the coefficients in the Taylor series for Δ_h in terms of ϵ . We note that $c_0 = \tanh(\beta_h)^2$, and therefore, these approximations only make sense for small β_h . How small should β_h be relative to ϵ ? We make the following ansatz:

$$\beta_h = \beta_0\epsilon^\delta \quad (\text{D18})$$

and thus if $\delta > 1/2$ then $c_0 \sim \beta_0^2\epsilon^{2\delta}$ will increase slower than ϵ .

We now express the coefficients for small β_h :

$$c_0 \approx \beta_0^2\epsilon^{2\delta} \quad (\text{D19})$$

$$c_1 \approx g_h^2(1 - 2\beta_h^2) \quad (\text{D20})$$

$$c_2 \approx g_h^4(-2 + 17\beta_h^2) \quad (\text{D21})$$

After solving Eqns. [D15-D19] self-consistently in terms of the expansion parameter ϵ , we get the following perturbative solution for $\delta \leq 1$:

$$\Delta_h \approx \begin{cases} \frac{2\beta_0\epsilon^\delta}{g_c^2(2 - g_c^2a_1)} & g_c^2a_1 < 2 \\ (g_c^2a_1 - 2)f_1 + \epsilon \cdot f_2 & g_c^2a_1 > 2 \end{cases} \quad (\text{D22})$$

$$\text{where } a_1 = \frac{\alpha_r^2}{16} \left[\phi_0^{(1)}(\beta_r/2)^2 + \phi_0(\beta_r/2)\phi_0^{(2)}(\beta_r/2) \right] \quad (\text{D23})$$

$f_2(\alpha_r, \beta_r)$ and $f_2(\alpha_r, \beta_r)$ are constant functions (w.r.t ϵ). Therefore, we see a linear scaling with the bias β_h .

2. Perturbative solutions for the fixed-point variance Δ_h in the bifurcation region with no biases

The perturbative treatment of the fixed-point solutions in this case closely follows that described above. For $g_h = 2 - \epsilon$, we can express Δ_h as a power-series in ϵ : $\Delta_h = c_0 + c_1\epsilon + c_2\epsilon^2$, and look for a condition that allows for a non-zero c_0 corresponding to the bifurcation point. Since we expect, Δ_h to be small in this regime, we can expand Δ_r as:

$$\Delta_r \approx g_h^2\Delta_h - 2g_h^4\Delta_h^2 + \frac{17}{3}g_h^6\Delta_h^3 + O(\Delta_h^4) \quad (\text{D24})$$

and similarly, we can also approximate

$$\left\langle \sigma_r(\sqrt{\Delta_r}x)^2 \right\rangle_x \approx \frac{1}{4} \left[1 + \frac{\alpha_r^2}{4}\Delta_r - \frac{\alpha_r^4}{8}\Delta_r^2 \right] \quad (\text{D25})$$

Now, equating coefficient of powers of ϵ , we get that either $c_0 = 0$ or

$$c_0 = \frac{3(\alpha_r^2 - 8)}{2(-136 + 24\alpha_r^2 + 3\alpha_r^4)} \quad (\text{D26})$$

which is a valid solution when $\alpha_r \geq \sqrt{8}$. This is the bifurcation curve limit near $g_h = 2^-$.

In the other limit, $\alpha_r^* \rightarrow \infty$ and $\Delta_h^* \rightarrow 0$. We can work in the regime where $\alpha_r\sqrt{\Delta_h} \gg 1$ to see what values of g_h admit a bifurcation in the perturbative solutions. The equation (to $O(\Delta_h^2)$) is given by:

$$\Delta_h \approx \frac{1}{2} [g_h^2\Delta_h - 2g_h^4\Delta_h^2] \quad (\text{D27})$$

Thus, we get a positive solution for Δ_h , when $g_h > \sqrt{2}$ and to the leading order, the solution scales as

$$\Delta_h^*(\sqrt{2}^+) \sim \frac{g_h^2 - 2}{2g_h^4} \quad \text{for } g_h \rightarrow \sqrt{2}^+ \quad (\text{D28})$$

3. $C_h(\tau)$ near critical point

Here we study the asymptotic behaviour of $C_h(\tau)$ near the critical point $g_h = 2.0$ for small α_z . For simplicity, we set the biases to be zero. In this limit we can assume that $C_h(\tau)$ and $C_\phi(\tau)$ are small. Let us begin by approximating $C_{\sigma_z}(\tau)$.

We get up to $O(C_z^3)$,

$$C_{\sigma_z}(\tau) = g_0 + g_1 C_z(\tau) + g_3 C_z(\tau)^3, \quad (\text{D29})$$

$$\text{where } g_0 = \frac{1}{4}, \quad (\text{D30})$$

$$g_1 = \frac{\alpha_z^2}{16} - \frac{\alpha_z^4}{32} C_z(0) + \frac{5\alpha_z^6}{256} C_z(0)^2, \quad (\text{D31})$$

$$g_3 = \frac{\alpha_z^6}{384} - \frac{\alpha_z^8}{192} C_z(0). \quad (\text{D32})$$

This can be obtained, for instance, by expanding $\sigma_z(z(t))$ and taking the Gaussian averages over the argument $z(t)$ in the steady-state. The relation between $C_\phi(\tau)$ and $C_z(\tau)$, in general, does not have a simple form; however, when $g_h \sim 2$, we expect the relaxation time $\tau_R \gg 1$, and therefore, we can approximate $C_z(\tau) \approx C_\phi(\tau)$. We can then approximate C_ϕ as

$$C_\phi(\tau) = g_0 + g_1 C_h(\tau) + g_3 C_h(\tau)^3, \quad (\text{D33})$$

$$\text{where } g_0 = 0, \quad (\text{for } \beta_h = 0) \quad (\text{D34})$$

$$g_1 = g_h^2 - 2g_h^4 C_h(0) + 5g_h^6 C_h(0)^2, \quad (\text{D35})$$

$$g_3 = \frac{2}{3}g_h^6 - \frac{16}{3}g_h^8 C_h(0). \quad (\text{D36})$$

Note that this also gives us an approximation for $C_\phi(0)$. Putting all this together, the equation governing $C_h(\tau)$:

$$[-\partial_\tau^2 + C_{\sigma_z}(\tau)] C_h(\tau) = \frac{1}{4} C_{\sigma_z}(\tau) C_\phi(\tau), \quad (\text{D37})$$

becomes (up to $O(C_h^3)$)

$$\partial_\tau^2 C_h(\tau) \simeq a_1 C_h(\tau) + a_2 C_h(\tau)^2 + a_3 C_h(\tau)^3, \quad (\text{D38})$$

$$\text{where } a_1 = \frac{1}{16} (4 - \Gamma), \quad (\text{D39})$$

$$a_2 = \frac{\alpha_z^2}{64} (4 - \Gamma) \Gamma, \quad (\text{D40})$$

$$a_3 = -\frac{g_h^6}{24} \quad (\text{D41})$$

$$\Gamma = g_h^2 - 2g_h^4 C_h(0) + 5g_h^6 C_h(0)^2. \quad (\text{D42})$$

Integrating w.r.t τ gives

$$(\partial_\tau C_h(\tau))^2 = 2 \left(\frac{a_1}{2} C_h(\tau)^2 + \frac{a_2}{3} C_h(\tau)^3 + \frac{a_3}{4} C_h(\tau)^4 + \text{const.} \right). \quad (\text{D43})$$

The boundary conditions are

$$\partial_\tau C_h(0) = 0 \quad \lim_{\tau \rightarrow \infty} \partial_\tau C_h(\tau) = 0. \quad (\text{D44})$$

The second condition implies the constant is 0. And, the first condition implies

$$\frac{a_1}{2} + \frac{a_2}{3} C_h(0) + \frac{a_3}{4} C_h(0)^2 = 0. \quad (\text{D45})$$

From this, we can solve for $C_h(0)$ (neglecting terms higher than quadratic) to get a solution that is perturbative in the deviation ϵ from the critical point ($g_h = 2 + \epsilon$). To the leading order the variance grows as

$$C_h(0) \approx \frac{1}{8} \epsilon + O(\epsilon^2), \quad (\text{D46})$$

and the α_z only enters the timescale-governing term a_1 at $O(\epsilon^2)$. At first, it might seem counter-intuitive that α_z , which effectively controls the dynamical time constant in the equations of motion, should not influence the relaxation rate to leading order. However, this result is for the dynamical behavior close to the critical point, where the relaxation time is a scaling function of ϵ . Moving away from this critical point, the relaxation time becomes finite, and the z -gate, and thus α_z , should have a more visible effect.

Appendix E: The DMFT bifurcation transition

In this section, we provide the details for calculating the DMFT prediction for the bifurcation transition. As in the main text, we assume $\alpha_z = 0$, and $\tau_r \ll 1$. The DMFT equation is then given by

$$4\partial_\tau^2 C_h(\tau) = C_h(\tau) - F_\phi(C_h^0, C_h(\tau)) F_{\sigma_r}(C_\phi^0, C_\phi(\tau)), \quad (\text{E1})$$

multiplying by $\partial_\tau C_h(\tau)$ and integrating from τ to ∞ , we get

$$2\dot{C}_h(\tau)^2 = \frac{(C_h)^2}{2} + \int_0^{C_h^0} dC_h F_\phi(C_h, C_h^0) F_{\sigma_r}(C_\phi, C_\phi^0). \quad (\text{E2})$$

Using the boundary condition that $\dot{C}_h(0) = 0$, we get the equation for the variance stated in the main text (eq. 73).

To get the condition for the dynamical bifurcation transition, we need to differentiate the L.H.S of eq. 73 ($\mathcal{F}(g_h, \alpha_r, C_h^0)$) w.r.t C_h^0 and set it to 0. This involves terms like

$$\frac{\partial F_\psi(C_h^0, C_h^0)}{\partial C_h^0}; \quad \frac{\partial F_\psi(C_h^0, 0)}{\partial C_h^0}. \quad (\text{E3})$$

We give a brief outline of calculating the first term. It's easier to work in the Fourier domain:

$$\begin{aligned} F_\psi(C_h^0, C_h) &= \mathbb{E} \left[\int \frac{dk}{2\pi} \int \frac{dk'}{2\pi} \tilde{\psi}(k) e^{-kz_1} \tilde{\psi}(k') e^{-k'z_2} \right], \\ &= \int \frac{dk}{2\pi} \int \frac{dk'}{2\pi} \tilde{\psi}(k) \tilde{\psi}(k') \exp \left[-\frac{C_h^0}{2} (k^2 + k'^2) - C_h(\tau) k k' \right]. \end{aligned} \quad (\text{E4})$$

This immediately gives us,

$$\begin{aligned} \frac{\partial F_\psi(C_h^0, C_h^0)}{\partial C_h^0} &= \int \mathcal{D}x \psi(\sqrt{c_h^0}x) \psi''(\sqrt{c_h^0}x) + \\ &\quad \int \mathcal{D}x \psi'(\sqrt{c_h^0}x)^2, \\ \frac{\partial F_\psi(C_h^0, 0)}{\partial C_h^0} &= \int \mathcal{D}x \psi(\sqrt{c_h^0}x) \int \mathcal{D}x \psi''(\sqrt{c_h^0}x). \end{aligned} \quad (\text{E5})$$

Using this fact, we can calculate the derivative of $\mathcal{F}(g_h, \alpha_r, C_h^0)$ as a straightforward (but long!) sum of Gaussian integrals. We then numerically solve eqns. 74,75 to get the bifurcation curve shown in Fig. 6a.

Next, we provide an analytical description of the asymptotic behaviour near the lower and higher critical values of g_h . From the red curve in Fig. 6a we know that as g_h tends towards the lower critical value, $\alpha_{r,DMFT}^* \rightarrow \infty$ and $C_h^0 \rightarrow 0$. So, we can approximate σ_r as a step function in this limit, and F_{σ_r} is approximated as

$$F_{\sigma_r}(C_\phi^0, C_\phi) \approx \frac{1}{4} + \frac{1}{2\pi} \tan^{-1}\left(\frac{x}{\sqrt{1-x^2}}\right), \quad (\text{E6})$$

$$\text{where } x := \frac{C_h(\tau)}{C_h(0)} \approx \frac{C_\phi(\tau)}{C_\phi(0)}. \quad (\text{E7})$$

The DMFT equation then reads:

$$4\ddot{x} = x - g_h^2 x \left(\frac{1}{4} + \frac{1}{2\pi} \tan^{-1}\left(\frac{x}{\sqrt{1-x^2}}\right) \right) + O(C_h(0)^2).$$

Integrating, this equation we get

$$2\dot{x}^2 = \frac{x^2}{2} \left(1 - \frac{g_h^2}{4} \right) + \frac{g_h^2}{8\pi} \left[(1-2x^2) \sin^{-1}(x) - x\sqrt{1-x^2} \right],$$

which will have $O(C_h(0)^2)$ corrections. From the boundary condition $\dot{C}_h(0) = 0$, we know that as $x \rightarrow 1$ then $\dot{x} \rightarrow 0$. We thus find that these boundary conditions are only consistent to leading order in $C_h(0)$ when g_h is equal to its critical value:

$$g_h^* = \sqrt{\frac{8}{3}}. \quad (\text{E8})$$

Which indicates that $C_h(0)$ must vanish as $g_h \rightarrow \sqrt{8/3}^+$.

In the other limit when $g_h \rightarrow 2^-$, we see that α_r^* remains finite and $C_h^{0,*} \rightarrow 0$. We assume that for $g_h = 2 - \epsilon$, C_h^0 has a power-series expansion

$$C_h^0 = c_0 + c_1\epsilon + c_2\epsilon^2 + \dots \quad (\text{E9})$$

We also expand F_ϕ and F_{σ_r} to $O(C_h(0)^2)$

$$F_\phi \approx g_h^2 C_h(\tau) - 2g_h^4 C_h^0 \cdot C_h(\tau) + 5g_h^6 (C_h^0)^2 \cdot C_h(\tau), \quad (\text{E10})$$

and look for values of α_r which permit a non-zero value for c_0 in the leading order solutions to the DMFT. We find that critical value of α_r from the perturbative solution is given by

$$\alpha_{r,DMFT}^*(2) = \sqrt{12}. \quad (\text{E11})$$

Appendix F: Adjoint DMFT details

1. Adjoint DMFT for the vanilla RNN

In this section, we provide details for deriving the Dynamical Mean-Field equations for the combined Adjoint+forward dynamics for the vanilla RNN (eq. 94-96). We start with the generating function expressed using the action \mathbf{S} :

$$Z_{\mathcal{J}} = \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \prod_k \mathcal{D}\lambda_k \mathcal{D}\hat{\lambda}_k e^{-i\mathbf{S}}, \quad (\text{F1})$$

$$\mathbf{S} = i \sum_{j=1}^N \int dt \left(\hat{\mathbf{x}}_j(t)^T \mathbf{b}_j(t) dt + \hat{\mathbf{b}}_j(t)^T \mathbf{x}_j(t) \right) + i\mathbf{S}_{\mathcal{J}}, \quad (\text{F2})$$

$$\begin{aligned} \mathbf{S}_{\mathcal{J}} = i \int dt &\left[\hat{\mathbf{h}}(t)^t (\dot{\mathbf{h}}(t) - G(\mathbf{h}(t), J)) + \right. \\ &\left. \hat{\boldsymbol{\lambda}}(t)^T (\dot{\boldsymbol{\lambda}}(t) + \partial_{\mathbf{h}} G(t)^T \boldsymbol{\lambda}(t) - \partial_{\mathbf{h}} f(\mathbf{h}(t))) \right], \end{aligned} \quad (\text{F3})$$

where, for the vanilla RNN, the function G describing the e.o.m and the state-to-state Jacobian $\partial_{\mathbf{h}} G$ are given by

$$G(\mathbf{h}(t), J) = -\mathbf{h}(t) + J\phi(\mathbf{h}(t)), \quad (\text{F4})$$

$$\partial_{\mathbf{h}} G(\mathbf{h}(t), J) = -\mathbf{1} + J[\phi']. \quad (\text{F5})$$

When we evaluate the disorder-averaged generating functional \bar{Z} , the terms in the action involving J_{ij} are averaged out as follows:

$$\begin{aligned} iJ_{ij} \int dt &(\lambda_i \hat{\lambda}_j \phi'_j - \hat{h}_i \phi_j) \rightarrow \\ &- \frac{g^2}{2N} \int dt dt' (\lambda_i \hat{\lambda}_j \phi'_j - \hat{h}_i \phi_j)_t (\lambda_i \hat{\lambda}_j \phi'_j - \hat{h}_i \phi_j)_{t'} \end{aligned} \quad (\text{F6})$$

As before, we introduce population averaged fields to make the expression local in the spatial indices. Let us define,

$$C_\lambda(t, t') = \frac{1}{N} \boldsymbol{\lambda}(t)^T \boldsymbol{\lambda}(t'), \quad (\text{F7})$$

$$C_\phi(t, t') = \frac{1}{N} \phi(t)^T \phi(t'), \quad (\text{F8})$$

$$R_\lambda(t, t') = \frac{1}{N} \boldsymbol{\lambda}(t)^T \hat{\mathbf{h}}(t'), \quad (\text{F9})$$

$$R_\phi(t, t') = \frac{1}{N} \phi(t)^T \hat{\mathbf{L}}(t'), \quad \hat{\mathbf{L}}(t) = (\hat{\boldsymbol{\lambda}} \odot \phi')(t). \quad (\text{F10})$$

We enforce these constraints using delta functions and the associated auxiliary fields $\hat{C}_\lambda, \hat{C}_\phi, \hat{R}_\lambda, \hat{R}_\phi$. The

disorder-averaged generating functional becomes

$$\begin{aligned} \bar{Z} = & \int [dM] \exp \left(iN \operatorname{tr} \hat{R}_\lambda^T R_\lambda + iN \operatorname{tr} \hat{R}_S^T R_S + i \frac{N}{2} \operatorname{tr} \hat{C}_\lambda C_\lambda \right. \\ & \left. + i \frac{N}{2} \hat{C}_S C_S - 2 \operatorname{tr} R_\lambda R_S + iW[\{M\}] \right) \end{aligned} \quad (\text{F11})$$

$$\begin{aligned} e^{iW[\{M\}]} = & \int [dX] \exp \left(i\mathbf{S}_0 - iN \hat{R}_\lambda(t, t') \frac{1}{N} \boldsymbol{\lambda}(t)^T \hat{\mathbf{h}}(t') \right. \\ & - iN \hat{R}_\phi(t, t') \frac{1}{N} \phi(t)^T \hat{\mathbf{L}}(t') - \frac{iN}{2} \hat{C}_\lambda(t, t') \frac{1}{N} \boldsymbol{\lambda}(t)^T \boldsymbol{\lambda}(t') \\ & - \frac{iN}{2} \hat{C}_\phi(t, t') \frac{1}{N} \phi(t)^T \phi(t') - \frac{N}{2} C_\lambda(t, t') \frac{1}{N} \hat{\mathbf{L}}(t)^T \hat{\mathbf{L}}(t') \\ & \left. - \frac{N}{2} C_\phi(t, t') \frac{1}{N} \hat{\mathbf{h}}(t) \cdot \hat{\mathbf{h}}(t') \right), \end{aligned} \quad (\text{F12})$$

where $\{M\}$ stands for the set $\{C_\phi, \hat{C}_\phi, C_\lambda, \hat{C}_\lambda, R_\phi, \hat{R}_\phi, R_\lambda, \hat{R}_\lambda\}$, $\{X\} = \{h, \hat{h}, \lambda, \hat{\lambda}\}$, and the action \mathbf{S}_0 is given by

$$\begin{aligned} \mathbf{S}_0 = & i \int dt \left[\hat{\mathbf{h}}(t)^t (\dot{\mathbf{h}}(t) + \mathbf{h}(t)), \right. \\ & \left. + \hat{\boldsymbol{\lambda}}(t)^T (\dot{\boldsymbol{\lambda}}(t) - \boldsymbol{\lambda}(t) - \partial_{\mathbf{h}} f(\mathbf{h}(t))) \right]. \end{aligned} \quad (\text{F13})$$

We see that the total action is extensive, and thus we can perform a saddle-point approximation in the limit $N \rightarrow \infty$. Noting that in this limit $\hat{C}_\phi = \hat{C}_\lambda = 0$ and that the response functions R have a vanishing contribution, we get a generating functional that is a product of single-site generating functionals, where each site is driven by two Gaussian noise processes:

$$\begin{aligned} \bar{Z}_0 = & \int [dh d\lambda] \exp \left(iS_0 - \frac{1}{2} C_\lambda(t, t') \hat{L}(t) \cdot \hat{L}(t') \right. \\ & \left. - \frac{1}{2} C_\phi(t, t') \hat{h}(t) \cdot \hat{h}(t') \right). \end{aligned} \quad (\text{F14})$$

This single-site generating functional corresponds to the coupled stochastic differential equations stated in the main text (eq. 94-96).

Appendix G: Calculation of maximal Lyapunov exponent from RMT

Here we present the derivation for the mean squared singular value of the susceptibility matrix for the gated RNN with $\alpha_z = 0$, and $\beta_z = -\infty$. In this limit, $\sigma_z = 1$, and the instantaneous Jacobian becomes the $2N \times 2N$ matrix

$$\mathcal{D}_t = -\mathbb{1}_{2N} + \begin{pmatrix} J^r & 0 \\ 0 & J^h \end{pmatrix} \begin{pmatrix} 0 & A_t \\ Q_t & R_t \end{pmatrix} \equiv -\mathbb{1}_{2N} + \hat{J} S_t, \quad (\text{G1})$$

$$Q_t = [\phi(\mathbf{h}(t)) \odot \sigma'_r(\mathbf{r}(t))], \quad A_t = [\phi'(\mathbf{h}(t))], \quad (\text{G2})$$

$$R_t = [\phi'(\mathbf{h}(t)) \odot \sigma_r(\mathbf{r}(t))]. \quad (\text{G3})$$

Let us define the quantity of interest

$$\sigma_\chi^2 = \left\langle \frac{1}{2N} \operatorname{Tr} (\chi(t) \chi^T(t)) \right\rangle \quad (\text{G4})$$

$$= e^{-2t} \left\langle \frac{1}{2N} \operatorname{Tr} e^{\hat{J} \hat{S}_t} e^{\hat{S}_t^T \hat{J}^T} \right\rangle, \quad (\text{G5})$$

where we have additionally defined $\hat{S}_t = \int^t dt' S_{t'}$, and the integration is performed elementwise. Expanding the exponentiated matrices and computing moments directly, one finds that the leading order in N moments must have an equal number of \hat{J} and \hat{J}^T . Thus, we must evaluate

$$c_n = \left\langle \frac{1}{2N} \operatorname{Tr} \left[(\hat{J} \hat{S}_t)^n (\hat{S}_t^T \hat{J}^T)^n \right] \right\rangle. \quad (\text{G6})$$

The ordering of the matrices is important in this expression. Since all of the \hat{J} appear to the left of \hat{J}^T , the leading order contributions to the moment will come from Wick contractions that are ‘‘non-crossing’’- in the language of diagrams, the moment will be given by a ‘‘rainbow’’ diagram. Consequently, we may evaluate c_n by induction. First, the induction step. Define the expected value of the matrix moment

$$\hat{c}_n = \left\langle (\hat{J} \hat{S}_t)^n (\hat{S}_t^T \hat{J}^T)^n \right\rangle. \quad (\text{G7})$$

$$= \left\langle \hat{J} \left(\hat{S}_t (\hat{J} \hat{S}_t)^{n-1} (\hat{S}_t^T \hat{J}^T)^{n-1} \hat{S}_t^T \right) \hat{J}^T \right\rangle \quad (\text{G8})$$

$$= \begin{pmatrix} a_n \mathbb{1} & 0 \\ 0 & b_n \mathbb{1} \end{pmatrix} + O(N^{-1}). \quad (\text{G9})$$

We wish to determine a_n and b_n . Next, define

$$g_A = \frac{1}{N} \operatorname{Tr} \int^t dt' dt'' A_{t'} A_{t''}, \quad (\text{G10})$$

$$g_Q = \frac{1}{N} \operatorname{Tr} \int^t dt' dt'' Q_{t'} Q_{t''}, \quad (\text{G11})$$

$$g_R = \frac{1}{N} \operatorname{Tr} \int^t dt' dt'' R_{t'} R_{t''}. \quad (\text{G12})$$

Now we can directly determine the induction step at the level of matrix moments by Wick contraction of the rainbow diagram

$$\hat{c}_n = \left\langle \hat{J} \hat{S}_t (\hat{J} \hat{S}_t)^{n-1} (\hat{S}_t^T \hat{J}^T)^{n-1} \hat{S}_t^T \hat{J}^T \right\rangle, \quad (\text{G13})$$

$$= \left\langle \hat{J} \hat{S}_t \hat{c}_{n-1} \hat{S}_t^T \hat{J}^T \right\rangle + O(N^{-1}), \quad (\text{G14})$$

$$= \begin{pmatrix} b_{n-1} g_A \mathbb{1} & 0 \\ 0 & (a_{n-1} g_Q + b_{n-1} g_R) \mathbb{1} \end{pmatrix} + O(N^{-1}). \quad (\text{G15})$$

This implies the following recursion for the diagonal elements of \hat{c}_n

$$a_n = g_A b_{n-1}, \quad b_n = g_R b_{n-1} + g_Q a_{n-1}. \quad (\text{G16})$$

The initial condition is given by observing that $\hat{c}_0 = \mathbb{1}$, which implies $a_0 = b_0 = 1$. The solution to this recursion relation can be written in terms of a transfer matrix

$$\begin{pmatrix} a_n \\ b_n \end{pmatrix} = \begin{pmatrix} 0 & g_A \\ g_Q & g_R \end{pmatrix}^n \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad (\text{G17})$$

which implies the moment $c_n = \frac{1}{2}(a_n + b_n)$ is given by

$$c_n = \frac{1}{2} \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & g_A \\ g_Q & g_R \end{pmatrix}^n \begin{pmatrix} 1 \\ 1 \end{pmatrix}. \quad (\text{G18})$$

To evaluate this, we use the fact that the eigenvalues of the transfer matrix are

$$v_{\pm} = \frac{1}{2} \left(g_R \pm \sqrt{g_R^2 + 4g_A g_Q} \right), \quad (\text{G19})$$

which are real valued. The eigenvectors are

$$\mathbf{v}_{\pm} = \begin{pmatrix} -\frac{v_{\mp}}{g_Q} \\ 1 \end{pmatrix}. \quad (\text{G20})$$

Then, defining $\mathbf{l} = (1, 1)$, the moment can be written

$$c_n = \frac{1}{2} \mathbf{l}^T (v_+^n \mathbf{v}_+ \mathbf{v}_+^T + v_-^n \mathbf{v}_- \mathbf{v}_-^T) \mathbf{l}, \quad (\text{G21})$$

$$= \frac{1}{2} \left(1 - \frac{v_-}{g_Q} \right)^2 v_+^n + \frac{1}{2} \left(1 - \frac{v_+}{g_Q} \right)^2 v_-^n. \quad (\text{G22})$$

The final expression for the mean squared singular value will then be

$$\sigma_{\chi}^2 = e^{-2t} \sum_{n=0}^{\infty} \frac{c_n}{(n!)^2}. \quad (\text{G23})$$

After resumming this infinite series we wind up with an expression in terms of the modified Bessel function

$$\sigma_{\chi}^2 = \frac{1}{2} e^{-2t} \left[\left(1 - \frac{v_-}{g_Q} \right)^2 I_0(2\sqrt{v_+}) + \left(1 - \frac{v_+}{g_Q} \right)^2 I_0(2\sqrt{v_-}) \right]. \quad (\text{G24})$$

In the steady-state, we approximate these expressions by assuming the correlation functions are time-translation invariant. Thus, we may write

$$g_R = \int dt dt' R_t R_{t'} \approx t^2 \frac{1}{t} \int d\tau C_R(\tau) = t^2 \tau_R, \quad (\text{G25})$$

and similarly for g_Q and g_A . Then the eigenvalues of the transfer matrix become

$$v_{\pm} = t^2 \frac{1}{2} \left(\tau_R \pm \sqrt{\tau_R^2 + 4\tau_A \tau_Q} \right). \quad (\text{G26})$$

At late times, using the asymptotic behavior of the modified Bessel function, the moment becomes

$$\sigma_{\chi}^2 \sim \exp(-2t + 2\sqrt{v_+}), \quad (\text{G27})$$

which gives the Lyapunov exponent

$$\lambda_L = \left(\frac{\tau_R + \sqrt{\tau_R^2 + 4\tau_A \tau_Q}}{2} \right)^{1/2} - 1. \quad (\text{G28})$$

-
- [1] H. T. Siegelmann, *Science* **268**, 545 (1995).
[2] A. Graves, A.-r. Mohamed, and G. Hinton, in *2013 IEEE international conference on acoustics, speech and signal processing* (IEEE, 2013) pp. 6645–6649.
[3] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, arXiv preprint arXiv:1312.6026 (2013).
[4] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, *Physical review letters* **120**, 024102 (2018).
[5] P. R. Vlachas, J. Pathak, B. R. Hunt, T. P. Sapsis, M. Girvan, E. Ott, and P. Koumoutsakos, *Neural Networks* (2020).
[6] L. Guastoni, P. A. Srinivasan, H. Azizpour, P. Schlatter, and R. Vinuesa, arXiv preprint arXiv:2002.01222 (2020).
[7] R. Jozefowicz, W. Zaremba, and I. Sutskever, in *International conference on machine learning* (2015) pp. 2342–2350.
[8] A. Crisanti and H. Sompolinsky, *Physical Review E* **98**, 062120 (2018).
[9] J. Kadmon and H. Sompolinsky, *Physical Review X* **5**, 041030 (2015).
[10] D. Sussillo and L. F. Abbott, *Neuron* **63**, 544 (2009).
[11] R. Laje and D. V. Buonomano, *Nature neuroscience* **16**, 925 (2013).
[12] S. Hochreiter and J. Schmidhuber, *Neural computation* **9**, 1735 (1997).
[13] S. Hochreiter, Y. Bengio, P. Frasconi, J. Schmidhuber, *et al.*, “Gradient flow in recurrent nets: the difficulty of learning long-term dependencies,” (2001).
[14] S. J. Mitchell and R. A. Silver, *Neuron* **38**, 433 (2003).
[15] R. Gütig and H. Sompolinsky, *PLoS biology* **7** (2009).
[16] H. Sompolinsky, A. Crisanti, and H.-J. Sommers, *Physical review letters* **61**, 259 (1988).
[17] D. Martí, N. Brunel, and S. Ostojic, *Physical Review E* **97**, 062314 (2018).
[18] F. Schuessler, A. Dubreuil, F. Mastrogiuseppe, S. Ostojic, and O. Barak, *Physical Review Research* **2**, 013111 (2020).
[19] F. Mastrogiuseppe and S. Ostojic, *Neuron* **99**, 609 (2018).
[20] M. Stern, H. Sompolinsky, and L. F. Abbott, *Phys. Rev. E* **90**, 062710 (2014).

- [21] J. Aljadeff, M. Stern, and T. Sharpee, *Physical review letters* **114**, 088101 (2015).
- [22] J. Schuecker, S. Goedeke, and M. Helias, *Physical Review X* **8**, 041029 (2018).
- [23] S. Alemohammad, Z. Wang, R. Balestriero, and R. Baraniuk, arXiv preprint arXiv:2006.10246 (2020).
- [24] J. Dong, R. Ohana, M. Rafayelyan, and F. Krzakala, arXiv preprint arXiv:2006.07310 (2020).
- [25] R. Brette, *Neural Computation* **18**, 2004 (2006).
- [26] S.-I. Amari, *IEEE Transactions on systems, man, and cybernetics*, 643 (1972).
- [27] P. C. Martin, E. Siggia, and H. Rose, *Physical Review A* **8**, 423 (1973).
- [28] C. De Dominicis, *Physical Review B* **18**, 4913 (1978).
- [29] J. A. Hertz, Y. Roudi, and P. Sollich, arXiv preprint arXiv:1604.05775 (2016).
- [30] H.-K. Janssen, *Zeitschrift für Physik B Condensed Matter* **23**, 377 (1976).
- [31] M. Helias and D. Dahmen, arXiv preprint arXiv:1901.10416 (2019).
- [32] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, in *International conference on machine learning* (2013) pp. 1139–1147.
- [33] J. Lee, L. Xiao, S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington, in *Advances in neural information processing systems* (2019) pp. 8572–8583.
- [34] A. Jacot, F. Gabriel, and C. Hongler, in *Advances in neural information processing systems* (2018) pp. 8571–8580.
- [35] R. Legenstein and W. Maass, *Neural networks* **20**, 323 (2007).
- [36] H. Jaeger and H. Haas, *science* **304**, 78 (2004).
- [37] T. Toyozumi and L. Abbott, *Physical Review E* **84**, 051908 (2011).
- [38] G. Wainrib and J. Touboul, *Physical review letters* **110**, 118101 (2013).
- [39] N. Maheswaranathan, A. Williams, M. Golub, S. Ganguli, and D. Sussillo, in *Advances in Neural Information Processing Systems* (2019) pp. 15696–15705.
- [40] D. Haviv, A. Rivkind, and O. Barak, arXiv preprint arXiv:1902.07275 (2019).
- [41] H. S. Seung, *Proceedings of the National Academy of Sciences* **93**, 13339 (1996).
- [42] H. S. Seung, *Neural Networks* **11**, 1253 (1998).
- [43] R. Gamkrelidze, L. S. Pontrjagin, and V. G. Boltjanskij, *The mathematical theory of optimal processes* (Macmillan Company, 1964).
- [44] See [56] for an example of how to extend the local random matrix analysis to a discrete-time model without leak in which the gate controlling the rate of updating is decoupled from the gate modulating the input strength.
- [45] K. Fischer and J. Hertz, “Spin glasses, cambridge university press,” (1991).
- [46] H. Sompolinsky and A. Zippelius, *Physical Review B* **25**, 6860 (1982).
- [47] H. Sompolinsky and A. Zippelius, *Physical Review Letters* **47**, 359 (1981).
- [48] C. C. Chow and M. A. Buice, *The Journal of Mathematical Neuroscience (JMN)* **5**, 8 (2015).
- [49] J. Schücker, S. Goedeke, D. Dahmen, and M. Helias, arXiv preprint arXiv:1605.06758 (2016).
- [50] F. Roy, G. Biroli, G. Bunin, and C. Cammarota, *Journal of Physics A: Mathematical and Theoretical* **52**, 484001 (2019).
- [51] J. Feinberg and A. Zee, *Nuclear Physics B* **504**, 579 (1997).
- [52] J. T. Chalker and B. Mehlig, *Physical review letters* **81**, 3367 (1998).
- [53] B. Cessac, *Journal de Physique I* **5**, 409 (1995).
- [54] S. Geman and C.-R. Hwang, *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **60**, 291 (1982).
- [55] Strictly speaking, the state variables evolve according to dynamics governed by (and thus dependent on) the J_s . However, the local chaos hypothesis states that large random networks approach a steady-state where the state variables are independent of J_s and are distributed according to their steady-state distribution.
- [56] T. Can, K. Krishnamurthy, and D. J. Schwab, arXiv preprint arXiv:2002.00025 (2020).
- [57] J.-P. Eckmann and D. Ruelle, in *The theory of chaotic attractors* (Springer, 1985) pp. 273–312.
- [58] K. Geist, U. Parlitz, and W. Lauterborn, *Progress of theoretical physics* **83**, 875 (1990).
- [59] R. Engelken, F. Wolf, and L. Abbott, arXiv preprint arXiv:2006.02427 (2020).
- [60] B. Derrida and Y. Pomeau, *EPL (Europhysics Letters)* **1**, 45 (1986).
- [61] The local chaos hypothesis employed by Cessac amounts to the same assumption.
- [62] In previous work, $g = 1$ sets the critical value. The difference is simply due to the factor $\sigma_r(0) = 1/2$. The vanilla RNN result is recovered by sending $\beta_r \rightarrow \infty$.
- [63] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein, arXiv preprint arXiv:1611.01232 (2016).
- [64] R. Kirov, Y. Zhu, R. R. Salakhutdinov, R. Zemel, R. Urtasun, A. Torralba, and S. Fidler, in *Advances in neural information processing systems* (2015) pp. 3294–3302.
- [65] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, arXiv preprint arXiv:1406.1078 (2014).
- [66] I. Sutskever, O. Vinyals, and Q. V. Le, in *Advances in neural information processing systems* (2014) pp. 3104–3112.
- [67] D. Bahdanau, K. Cho, and Y. Bengio, arXiv preprint arXiv:1409.0473 (2014).
- [68] M. Chen, J. Pennington, and S. S. Schoenholz, arXiv preprint arXiv:1806.05394 (2018).
- [69] D. Gilboa, B. Chang, M. Chen, G. Yang, S. S. Schoenholz, E. H. Chi, and J. Pennington, arXiv preprint arXiv:1901.08987 (2019).
- [70] C. Tallec and Y. Ollivier, arXiv preprint arXiv:1804.11188 (2018).
- [71] T. Mora and W. Bialek, *Journal of Statistical Physics* **144**, 268 (2011).
- [72] H. S. Seung, D. D. Lee, B. Y. Reis, and D. W. Tank, *Neuron* **26**, 259 (2000).
- [73] C. K. Machens, R. Romo, and C. D. Brody, *Science* **307**, 1121 (2005).
- [74] V. M. Eguíluz, M. Ospeck, Y. Choe, A. Hudspeth, and M. O. Magnasco, *Physical review letters* **84**, 5232 (2000).
- [75] N. Bertschinger and T. Natschläger, *Neural computation* **16**, 1413 (2004).
- [76] R. Legenstein and W. Maass, *New directions in statistical signal processing: From systems to brain*, 127 (2007).
- [77] J. Boedecker, O. Obst, J. T. Lizier, N. M. Mayer, and M. Asada, *Theory in Biosciences* **131**, 205 (2012).
- [78] D. Mishkin and J. Matas, arXiv preprint

arXiv:1511.06422 (2015).