

Theory of Gating in Recurrent Neural Networks

Kamesh Krishnamurthy^{1,*}, Tankut Can^{2,†}, and David J. Schwab³

¹*Joseph Henry Laboratories of Physics and PNI, Princeton University, Princeton, New Jersey 08544, USA*

²*Institute for Advanced Study, Princeton, New Jersey 08540, USA*

³*Initiative for Theoretical Sciences, Graduate Center, CUNY, New York, New York 10016, USA*



(Received 15 March 2021; revised 21 June 2021; accepted 11 November 2021; published 18 January 2022)

Recurrent neural networks (RNNs) are powerful dynamical models, widely used in machine learning (ML) and neuroscience. Prior theoretical work has focused on RNNs with additive interactions. However, gating, i.e., multiplicative, interactions are ubiquitous in real neurons and also *the* central feature of the best-performing RNNs in ML. Here, we show that gating offers flexible control of two salient features of the collective dynamics: (i) timescales and (ii) dimensionality. The gate controlling timescales leads to a novel, marginally stable state, where the network functions as a flexible integrator. Unlike previous approaches, gating permits this important function without parameter fine-tuning or special symmetries. Gates also provide a flexible, context-dependent mechanism to reset the memory trace, thus complementing the memory function. The gate modulating the dimensionality can induce a novel, discontinuous chaotic transition, where inputs push a stable system to strong chaotic activity, in contrast to the typically stabilizing effect of inputs. At this transition, unlike additive RNNs, the proliferation of critical points (topological complexity) is decoupled from the appearance of chaotic dynamics (dynamical complexity). The rich dynamics are summarized in phase diagrams, thus providing a map for principled parameter initialization choices to ML practitioners.

DOI: [10.1103/PhysRevX.12.011011](https://doi.org/10.1103/PhysRevX.12.011011)

Subject Areas: Interdisciplinary Physics
Nonlinear Dynamics, Statistical Physics

I. INTRODUCTION

Recurrent neural networks (RNNs) are powerful dynamical systems that can represent a rich repertoire of trajectories and are popular models in neuroscience and machine learning. In modern machine learning, RNNs are used to learn complex dynamics from data with rich sequential or temporal structure such as speech [1,2], turbulent flows [3–5], or text sequences [6]. RNNs are also influential in neuroscience as models to study the collective behavior of a large network of neurons [7] (and references therein). For instance, they have been used to explain the dynamics and temporally irregular fluctuations observed in cortical networks [8,9] 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 [6,12]. Gating interactions are also ubiquitous in real neurons due to mechanisms such as shunting inhibition [13]. Moreover, when single-neuron models are endowed with more realistic conductance dynamics, the effective interactions at the network level have gating effects, which confer robustness to time-warped inputs [14]. Thus, RNNs with gating interactions not only have superior information processing capabilities, but they also embody a prominent feature found in real neurons.

Prior theoretical work on understanding the dynamics and functional capabilities of RNNs has mostly focused on RNNs with additive interactions. The original work by Sompolinsky, Crisanti, and Sommers [15] identifies a phase transition in the autonomous dynamics of randomly connected RNNs from stability to chaos. Subsequent work extends this analysis to cases where the random connectivity additionally has correlations [16], a low-rank structured component [17,18], strong self-interaction [19], and heterogeneous variance across blocks [20]. The role of sparse connectivity and the single-neuron nonlinearity is studied in Ref. [9]. The effect of a Gaussian noise input is analyzed in Ref. [21].

*Lead and corresponding author.

kameshk@princeton.edu

†Corresponding author.

tankut.can@gmail.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

In this work, we study the consequences of gating interactions on the dynamics of RNNs. We introduce a gated RNN model that naturally extends a 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 output of a neuron. The choice of these forms for gates are motivated by biophysical considerations (e.g., Refs. [14,22]) and retain the most functionally important aspects of the gated RNNs in machine learning. Our gated RNN reduces to the classical RNN [15,23] 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 [24]. We further elaborate on this connection in Sec. VIII.

We develop a theory for the gated RNN based on non-Hermitian random matrix techniques [25,26] and the Martin–Siggia–Rose–De Dominicis–Janssen (MSRDJ) formalism [21,27–32] and use the theory to map out, in a phase diagram, the rich, functionally significant dynamical phenomena produced by gating.

We show that the update gate produces slow modes and a marginally stable critical state. Marginally stable systems are of special interest in the context of biological information processing (e.g., Ref. [33]). Moreover, the network in this marginally stable state can function as a robust integrator—a function that is critical for memory capabilities in biological systems [34–37] but has been hard to achieve without parameter fine-tuning and handcrafted symmetries [38]. Gating permits the network to serve this function without any symmetries or fine-tuning. For a detailed discussion of these issues, we refer the reader to Ref. [39] (pp. 329–350) and Refs. [38,40]. Integratorlike dynamics are also empirically observed in gated machine learning (ML) RNNs successfully trained on complex sequential tasks [41]; our theory shows how gates allow for this robustly.

The output gate allows fine control over the dimensionality of the network activity; control of the dimensionality can be useful during learning tasks [42]. In certain regimes, this gate can mediate an input-driven chaotic transition, where static inputs can push a stable system abruptly to a chaotic state. This behavior with gating is in stark contrast to the typically stabilizing effect of inputs in high-dimensional systems [21,43,44]. The output gate also leads to a novel, discontinuous chaotic transition, where the proliferation of critical points (a static property) is decoupled from the appearance of chaotic transients (a dynamical property); this is in contrast to the tight link between the two properties in additive RNNs as shown by Wainrib and Touboul [45]. This transition is also characterized by a nontrivial state where a stable fixed point coexists with long chaotic transients. Gates also provide a flexible, context-dependent way to reset the state, thus providing a way to selectively erase the memory trace of past inputs.

We summarize these functionally significant phenomena in phase diagrams, which are also practically useful for ML

practitioners—indeed, the choice of parameter initialization is known to be one of the most important factors deciding the success of training [46], with best outcomes occurring near critical lines [10,47–49]. Phase diagrams, thus, allow a principled and exhaustive exploration of dynamically distinct initializations.

II. A RECURRENT NEURAL NETWORK MODEL TO STUDY GATING

We study an extension of a classical RNN [15,23] by augmenting it with multiplicative *gating* interactions. Specifically, we consider two gates: (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

$$\dot{h}_i(t) = \sigma_z(z_i)[-h_i(t) + R_i(t)] + I_i^h(t), \quad (1)$$

where h_i represents the internal state of the i th unit and $\sigma_{(\cdot)}(x) = [1 + \exp(-\alpha_{(\cdot)}x + \beta_{(\cdot)})]^{-1}$ are sigmoidal gating functions. The recurrent input to a neuron is $R_i(t) = \sum_{j=1}^N J_{ij}^h \phi[h_j(t)] \cdot \sigma_r[r_j(t)]$, where J_{ij}^h are the coupling strengths between the units and $\phi(x) = \tanh(g_h x + \beta_h)$ is the activation function. ϕ and $\sigma_{z,r}$ are parametrized by gain parameters ($g_h, \alpha_{z,r}$) and biases ($\beta_{h,z,r}$), which constitute the parameters of the gated RNN. Finally, I^h represents external input to the network. The gating variables $z_i(t)$ and $r_i(t)$ evolve according to dynamics driven by the output $\phi[h(t)]$ of the network:

$$\tau_x \dot{x}_i(t) = -x_i(t) + \sum_{j=1}^N J_{ij}^x \phi[h_j(t)] + I_i^x, \quad (2)$$

where $x \in \{z, r\}$. Note that the coupling matrices $J^{z,r}$ for z, r are distinct from J^h . We also allow for different inputs I^r and I^z being fed to the gates. For instance, they might be zero, or they might be equal up to a scaling factor to I^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 $\sigma_r(r_i)$ modulates the output strength of unit i . In the presence of external input, the r gate can control the relative strengths of the internal (recurrent) activity and the external input I^h . In the limit $\sigma_z, \sigma_r \rightarrow 1$, we recover the dynamics of the classical RNN.

We choose the coupling weights from a Gaussian distribution with variance scaled such that the input to each unit remains $O(1)$. Specifically, $J_{ij}^{h,z,r} \sim \mathcal{N}(0, N^{-1})$. This choice of couplings is a popular initialization scheme for RNNs in machine learning [6,46] and also in models of cortical neural circuits [15,20]. If the gating variables are purely internal, then $(J^{z,r})$ is diagonal; however, we do not consider this case below. In the rest of the paper, we analyze

the various dynamical regimes the gated RNN exhibits and their functional significance.

III. HOW THE GATES SHAPE THE LINEARIZED DYNAMICS

We first study the linearized dynamics of the gated RNN through the lens of the instantaneous 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 timescales of growth and decay of perturbations and the local stability of the dynamics. As we show below, the spectral density of the Jacobian depends on *equal-time* correlation functions, which are the order parameters in the mean-field picture of the dynamics, developed in the Appendix C. We study how the gates shape the support and the density of Jacobian eigenvalues in the steady state, through their influence on the correlation functions.

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

$$\delta \dot{\mathbf{x}} = \mathcal{D}(t) \delta \mathbf{x}, \quad (3)$$

where \mathcal{D} is the $3N \times 3N$ -dimensional instantaneous Jacobian of the full network equations. Linearization of Eqs. (1) and (2) yields

$$\mathcal{D} = \begin{pmatrix} [\sigma_z](-\mathbb{1} + J^h[\phi' \sigma_r]) & \mathbb{D} & [\sigma_z]J^h[\phi \sigma_r'] \\ \tau_z^{-1}J^z[\phi'] & -\tau_z^{-1}\mathbb{1} & 0 \\ \tau_r^{-1}J^r[\phi'] & 0 & -\tau_r^{-1}\mathbb{1} \end{pmatrix}, \quad (4)$$

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 is zero for fixed points. We introduce the additional shorthand $\phi'(t) = \phi'(h(t))$ and $\sigma'_{r/z} = \sigma'_{r/z}(r/z(t))$.

The Jacobian is a block-structured matrix involving random elements ($J^{z,h,r}$) and functions of various state variables. We need additional tools from non-Hermitian random matrix theory (RMT) [26] and dynamical mean-field theory (DMFT) [15] to analyze the spectrum of the Jacobian \mathcal{D} . We provide a detailed, self-contained derivation of the calculations in Appendix C (DMFT) and Appendix A (RMT). Here, we state only the main results derived from these formalisms.

One of the main results is an analytical expression for the spectral curve, which describes the boundary of the Jacobian spectrum, in terms of the moments of the state variables. The most general expression for the spectral curve [Appendix A, Eq. (A34)] involves empirical averages over the $3N$ -dimensional state variables. However, 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 (cf. Appendix C)—i.e., we can replace empirical averages of any function of the state variables $(1/N) \sum_i F(h_i, z_i, r_i)$ with $\langle F[h(t), z(t), r(t)] \rangle$, where the brackets indicate average over the steady-state distribution. The DMFT + RMT prediction for the spectral curve for a generic steady-state point is given in Appendix A, Eq. (A35). Strictly speaking, the analysis of the DMFT around a generic time-dependent steady state is complicated by the fact that the distribution for h is not Gaussian (while r and z are Gaussian). For fixed points, however, the distributions of h , z , and r are all Gaussian, and the expression for the spectral curve reduces simplifies. It is given by the set of $\lambda \in \mathbb{C}$ which satisfy

$$\langle \phi'^2 \rangle \left(\langle \sigma_r'^2 \rangle + \frac{\langle \phi'^2 \rangle \langle \sigma_r'^2 \rangle}{|1 + \tau_r \lambda|^2} \right) \left\langle \frac{\sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle_z = 1. \quad (5)$$

Here, the averages are taken over the Gaussian fixed-point distributions $(h, z, r) \sim \mathcal{N}(0, \Delta_{h,z,r})$ which follow from the MFT [Eq. (C26)]. For example, $\langle \phi'^2 \rangle = \mathbb{E}_{h \sim \mathcal{N}(0, \Delta_h)} [\phi'(h)^2]$.

We make two comments on the Jacobian of a time-varying state: (i) One might wonder if any useful information can be gleaned by studying the Jacobian at a time-varying state where the Hartman-Grobman theorem is not valid. Indeed, as we see below, the limiting form of the Jacobian in steady state crucially informs us about the suppression of unstable directions and the emergence of slow dynamics due to *pinching* and marginal stability in certain parameter regimes (also see Ref. [50]). In other words, the instantaneous Jacobian charts the approach to marginal stability and provides a quantitative justification for the approximate integrator functionality exhibited in Sec. IV B. (ii) Interestingly, the spectral curve calculated using the MFT [Eq. (5)] for a *time-varying* steady state not deep in the chaotic regime is a very good approximation for the true spectral curve (see Fig. 8 in Appendix A).

Figures 1(a)–1(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. (5), we get a condition for the stability of the zero fixed point. The leading edge of the spectral curve for the zero fixed point (FP) crosses the origin when $g_h < 1 + e^{-\beta r}$. So, in the absence of biases, $g_h > 2$ makes the zero FP unstable. More generally, the leading edge of the spectrum crossing the origin gives us the condition for the FP to become unstable:

$$\langle \phi'^2 \rangle (\langle \phi'^2 \rangle \langle \sigma_r'^2 \rangle + \langle \sigma_r'^2 \rangle) > 1 \Rightarrow \text{unstable FP}. \quad (6)$$

We 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 via the equation for the spectral curve.

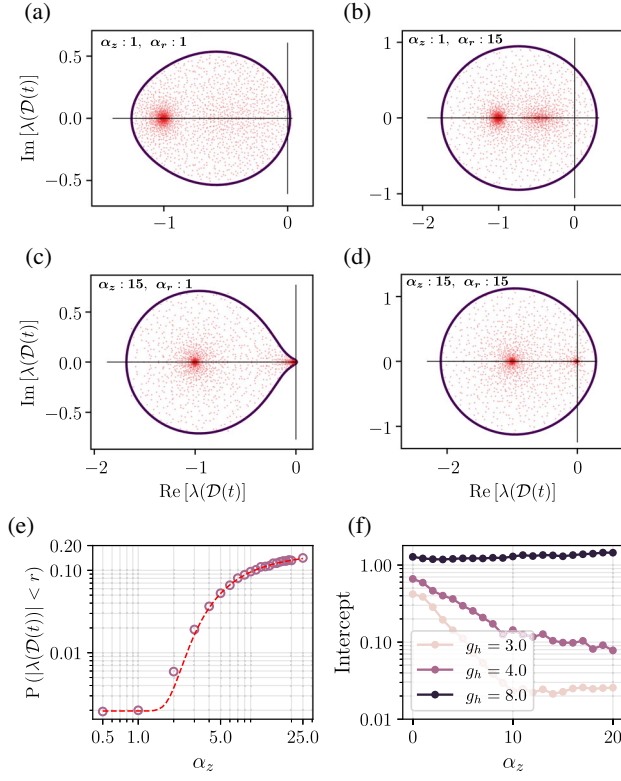


FIG. 1. How gates shape the Jacobian spectrum. (a)–(d) Jacobian eigenvalues (red dots) of the gated RNN in (time-varying) steady state. The dark outline is the spectral support curve predicted by Eq. (5). The bottom row corresponds to larger α_z , and the right column corresponds to large α_r . (e) Cumulative distribution function 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 ten instances), and the dashed line is a fit from Eq. (7). (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, $N = 2000$, $g_h = 3$, and $\tau_r = \tau_z = 1$ unless otherwise stated, and all biases are zero.

A. 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 and the shape of the spectral support curve—the eigenvalues tell us about the rate of growth or decay of small perturbations and, thus, timescales in the local dynamics, and the spectral curve informs us about stability. For ease of exposition, we consider the case without biases in the main text ($\beta_{r,z,h} = 0$); we discuss the role of biases in Appendix H.

Figure 1 shows how the gain parameters of the update and output gates— α_z and α_r , respectively—shape the Jacobian spectrum. In Figs. 1(a)–1(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. 1(f); also see Sec. IV A]. In Figs. 1(a)–1(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 increasing the local dimensionality of the unstable manifold. This behavior of the linearized dynamics is also reflected in the nonlinear dynamics, where, as we show in Sec. V, α_r has the effect of controlling the *dimensionality* of full phase-space dynamics.

The accumulation of eigenvalues near zero with increasing α_z suggests the emergence of a wide spectrum of timescales in the local dynamics. 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 piecewise 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 B):

$$P[|\lambda(\mathcal{D}_x)| < \delta] \sim c_0 \text{erf}\left(\frac{c_1}{\alpha_z}\right), \quad (7)$$

where c_0 and c_1 are constants that, in general, depend on α_r , δ , and g_h . Figure 1(e) 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 B)

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

where $f_z = \langle \sigma_z(z) \rangle$ is the fraction of update gates which are nonzero 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. Analogous phenomena are observed for discrete-time gated RNNs in Ref. [51], using a similar combination of analytical and numerical techniques [52].

In Sec. V A, we show that the slow modes, as seen from linearization, persist asymptotically (i.e., in the nonlinear regime). This can be seen from the Lyapunov spectrum in Fig. 3(a), which for large α_z exhibits an analogous accumulation of Lyapunov exponents near zero.

In the next section, we study the profound *functional* consequences of the combination of spectral pinching and accumulation of eigenvalues near zero.

IV. MARGINAL STABILITY AND ITS CONSEQUENCES

As the update gate becomes more switchlike (higher α_z), we see an accumulation of slow modes and a pinching of the spectral curve which drastically suppresses the unstable

directions. In the limit $\alpha_z \rightarrow \infty$, this can make previously unstable points marginally stable by pinning the leading edge of the spectral curve exactly at zero. Marginally stable systems are of significant interest because of the potential benefits in information processing—for instance, they can generate long timescales in their collective modes [33,39]. Moreover, achieving marginal stability often requires fine-tuning parameters close to a bifurcation point. As we see, gating allows us to achieve a marginally stable critical state over a wide range of parameters; this has been typically highly nontrivial to achieve (e.g., Ref. [39], pp. 329–350, and Ref. [33]). We first investigate the conditions under which marginal stability arises, and then we touch on one of its important functional consequences: the appearance of “line attractors” which allow the system to be used as a robust integrator.

A. Condition for marginal stability

Marginal stability is a consequence of pinching of the spectral curve with increasing α_z , wherein the (positive) leading edge of the spectrum and the intercept of the spectral curve on the imaginary axis both shrink with α_z [e.g., Fig. 1(f) and compare Figs. 1(a) and 1(c)]. However, we see in Fig. 1(f) (via the intercept) that pinching does not happen if g_h is sufficiently large (even as $\alpha_z \rightarrow \infty$). Here, we provide the conditions when pinching can occur and, thus, marginal stability can emerge. For simplicity, let us consider the case where $\tau_r = 1$ and there are no biases.

Marginal stability strictly exists only for $\alpha_z = \infty$. We first examine the conditions under which the system can become marginally stable in this limit, and then we explain the *route* to marginal stability for large but finite α_z , i.e., how a time-varying state ends up as a marginally stable fixed point. For $\alpha_z = \infty$, the spectral density has an extensive number $N[1 - \langle \sigma_z(z) \rangle]$ of zero eigenvalues, and the remaining eigenvalues are distributed in a disk centered at $\lambda = -1$ with radius ρ . If $\rho < 1$, then the spectral density has two topologically disconnected configurations (the disk and the zero modes) and the system is marginally stable. If $\rho > 1$, the zero modes get absorbed in the interior of the disk and the system is unstable with fast, chaotic dynamics. The radius ρ is given by $\rho^2 = \frac{1}{2}a + \frac{1}{2}\sqrt{4b + a^2} < 1$, where $a = \langle \phi'^2 \rangle \langle \sigma_z \rangle \langle \sigma_r^2 \rangle$ and $b = \langle \phi'^2 \rangle \langle \sigma_z \rangle \langle \phi^2 \rangle \langle \sigma_r'^2 \rangle$. This follows from Eq. (5) by evaluating the z -expectation value assuming σ_z is a binary variable. Thus, the system is marginally stable in the limit $\alpha_z = \infty$ as long as

$$\langle \phi'^2 \rangle (\langle \phi^2 \rangle \langle \sigma_r'^2 \rangle + \langle \sigma_r^2 \rangle) < \langle \sigma_z \rangle^{-1}. \quad (8)$$

The crucial difference between this expression and Eq. (6) is that the rhs now has a factor of $\langle \sigma_z \rangle^{-1}$ which can be greater than unity, thus pushing the transition to chaos further out along the g_h and α_r directions, as depicted in the phase diagram (Fig. 7). For concreteness, we report

here how the transition changes at $\alpha_r = 0$. In this setting, the transition to chaos moves from $g_h = 2$ to $g_h \lesssim 6.2$, and the system is marginally stable for $2 < g_h \lesssim 6.2$.

Having identified the region in the phase diagram that can be made marginally stable for $\alpha_z = \infty$, we can now discuss the route to marginal stability for large but finite α_z . In other words, how does an unstable chaotic state become marginally stable with increasing α_z ? Since the marginally stable region is characterized by a disconnected spectral density, evidently increasing α_z must lead to singular behavior in the spectral curve. This takes the form of a pinching at the origin. We show that, for values of g_h supporting marginal stability, the leading edge λ_e of the spectrum for the time-varying state gets pinched exponentially fast with α_z as $\lambda_e \sim e^{-c\alpha_z \sqrt{\Delta_h}}$ (see Appendix B). This accounts for the fact that, already for $\alpha_z = 15$, we observe the pinching in Fig. 1(c). In contrast, the parameters in Fig. 1(d) lie outside the marginally stable region, and, thus, there is no pinching, since the zero modes are asymptotically (in α_z) buried in the bulk of the spectrum.

In summary, as $\alpha_z \rightarrow \infty$ 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 end up sitting exactly at zero, while the rest occupy a finite circular region. If the leading edge of the circular region crosses zero in this limit, then the state remains unstable; otherwise, the state becomes marginally stable. The latter case is achieved through a gradual pinching of the spectrum near zero; there is no pinching in the former case.

We emphasize that marginal stability requires more than just an accumulation of eigenvalues near zero. Indeed, this happens even when g_h is outside the range supporting marginal stability as $\alpha_z \rightarrow \infty$, but there is no pinching and the system remains unstable [e.g., see Fig. 1(d)]. We return to this when we describe the phase diagram for the gated RNN (Sec. VII). There, we see that the marginally stable region occupies a macroscopic volume in the parameter space adjoining the critical lines on one side.

B. Functional consequences of marginal stability

The marginally stable critical state produced by gating can subserve the function of a robust integrator. This integratorlike function is crucial for a variety of computational functions such as motor control [34–36], decision making [37], and auditory processing [53]. However, achieving this function has typically required fine-tuning or special handcrafted architectures [38], but gating permits the integrator function over a range of parameters and without any specific symmetries in $J^{h,z,r}$. Specifically, for large α_z , any perturbation in the span of the eigenvectors corresponding to the eigenvalues with a magnitude close to zero is integrated by the network, and, once the input perturbation ceases, the memory trace of the input is

retained for a duration much longer than the intrinsic time constant of the neurons; perturbations along other directions, however, relax with a spectrum of timescales dictated by the inverse of (the real part of) their eigenvalues. Thus, the manifold of slow directions forms an approximate continuous attractor on which input can effortlessly move the state vector around. These approximate continuous attractor dynamics are illustrated in Fig. 2. At time $t = 0$, an input I^h (with $I^r = I^z = 0$) is applied till $t = 10$ (between dashed vertical lines) along an eigenvector of the Jacobian with an eigenvalue close to zero. Inputs along this 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 one unit. The exponentially fast (in α_z) pinching of the spectral curve (discussed above in Sec. IV A) suggests this slow-manifold behavior should also hold for moderately large α_z (as in Fig. 2). Therefore, even though the state is technically unstable, the local structure of the Jacobian is responsible for giving rise to

extremely long timescales and allows the network to operate as an approximate integrator within relatively long windows of time, as demonstrated in Fig. 2.

Of course, after sufficiently long times, the instability causes the state to evolve and the memory is lost. Exactly how long the memory lasts depends on the asymptotic stability of the network, which is revealed by the Lyapunov spectrum, discussed below in Sec. V A.

V. OUTPUT GATE CONTROLS DIMENSIONALITY AND LEADS TO A NOVEL CHAOTIC TRANSITION

We thus far use insights from local dynamics to study the functional consequences of the gates. To study the salient features of the output gate, it is useful to analyze the effect of inputs and the long-time behavior of the network through the lens of Lyapunov spectra. We see that the output gate controls the *dimensionality* of the dynamics in the phase space; dimensionality is a salient aspect of the dynamics for task function [42]. The output gate also gives rise to a novel discontinuous chaotic transition, near which inputs (even static ones) can abruptly push a stable system into strongly chaotic behavior—contrary to the typically stabilizing effect of inputs. Below, we begin with the Lyapunov analyses of the dynamics and then proceed to study the chaotic transition.

A. Long-time behavior of the network

We study the asymptotic behavior of the network and the nature of the time-varying state through the lens of its Lyapunov spectra. In this section, where we study the effects of α_z , our results are numerical except in cases where $\alpha_z = 0$ [e.g., in Fig. 3(d)]. Lyapunov exponents specify how infinitesimal perturbations $\delta\mathbf{x}(t)$ grow or shrink along the trajectories of the dynamics—in particular, if the growth or decay is exponentially fast, then the rate is dictated by the maximal Lyapunov exponent defined as [54] $\lambda_{\max} := \lim_{T \rightarrow \infty} T^{-1} \lim_{\|\delta\mathbf{x}(0)\| \rightarrow 0} \ln[\|\delta\mathbf{x}(T)\|/\|\delta\mathbf{x}(0)\|]$. 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 behavior. We first numerically study how the gates shape the full Lyapunov spectrum (details in Appendix D) and derive an analytical prediction for the maximum Lyapunov exponent using the DMFT (Sec. V A 1) [55].

Figures 3(a) and 3(b) show 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, generating long timescales. As the output gate becomes more sensitive (larger α_r), all Lyapunov exponents increase, thus increasing the rate of growth in unstable directions.




FIG. 2. Network in the marginally stable state functions as an integrator. (a) Sample traces from a network with switchlike update gates ($\alpha_z = 30, g_h = 3$) show slow evolution (time on x axis is relative to τ_h). (b) An input is applied to the same network in (a) 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). The 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. If the input is along the slow manifold, the trace of the input is retained for a long time after the cessation of input. [The traces in (a) are for the network with an input along the manifold.]

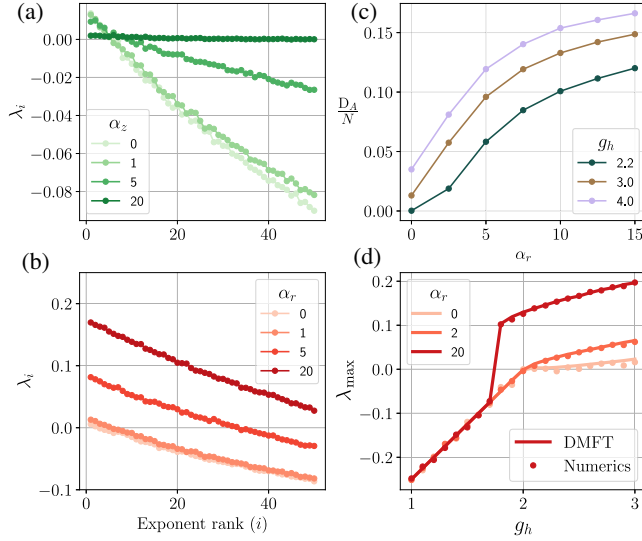


FIG. 3. Lyapunov spectra and dimensionality of the gated RNN. (a),(b) The first 50 ordered Lyapunov exponents for a gated RNN ($N = 2000$) as a function of varying (a) α_z and (b) α_r . The Lyapunov spectrum is calculated as described in Appendix D. (c) The Kaplan-Yorke dimensionality of the dynamics as a function of α_r . (d) The maximal Lyapunov exponent λ_{\max} predicted by the DMFT [solving Eqs. (10) and (11); solid line] and obtained numerically using the *QR* method (circles; $N = 2000$ and $\alpha_z = 0$). Note that the transition for $\alpha_r = 20$ is sharp; also cf. Fig. 5(c). $\tau_z = \tau_r = 2.0$ here.

We can estimate the dimensionality of the activity in the chaotic state by calculating an upper bound D_A on the dimension according to a conjecture by Kaplan and Yorke [54]. The Kaplan-Yorke upper bound for the attractor dimension D_A is given by

$$D_A = M + \frac{\sum_{i=1}^M \lambda_i}{|\lambda_{M+1}|}, \text{ where } M = \max_j \left\{ \sum_{i=1}^j \lambda_i \geq 0 \right\}, \quad (9)$$

where λ_i are the rank-ordered Lyapunov exponents. We see in Fig. 3(c) that the sensitivity of the output gate (α_r) can shape the dimensionality of the dynamics—a more sensitive output gate leads to higher dimensionality. As we see below, this effect of the output gate is different from how the gain g_h shapes dimensionality and can lead to a novel chaotic transition. Even more directly, if the r gate for neurons $i_1 \dots i_K$ is set to zero, then the activity is constrained to evolve in an $N - K$ -dimensional subspace; however, the r gate allows the possibility—i.e., the “inductive bias”—of doing this *dynamically*.

1. DMFT prediction for λ_{\max}

We would also 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 maximum Lyapunov exponent, and, to do this, we make use of a technique suggested by Refs. [56,57] and clearly elucidated in Ref. [21]. The details are provided in Appendix E, and the end result of the calculation is the DMFT prediction for λ_{\max} as the solution to a generalized eigenvalue problem for κ involving the correlation functions of the state variables:

$$\begin{aligned} & [(\langle \sigma_z \rangle + \kappa)^2 - \partial_\tau^2 + C_{\sigma_z}(\tau) - \langle \sigma_z \rangle^2] \chi_h(\tau) \\ &= C_{\sigma_z}(\tau) [C_{\phi \cdot \sigma_r}(\tau) - C_h(\tau)] \chi_z(\tau) \\ &+ C_{\sigma_z}(\tau) \frac{\partial C_{\phi \cdot \sigma_r}(\tau)}{\partial C_h} \chi_h(\tau), \end{aligned} \quad (10)$$

$$[(1 + \tau_{z/r} \kappa)^2 - \tau_{z/r}^2 \partial_\tau^2] \chi_{z/r}(\tau) = \frac{\partial C_\phi(\tau)}{\partial C_h} \chi_h(\tau), \quad (11)$$

where we denote the two-time correlation function $C_x(t, t') \equiv \langle x(t)x(t') \rangle$ for different (functions of) state variables $x(t)$ [see Eq. (C25) for more context]. The largest eigenvalue solution to this problem is the required maximal Lyapunov exponent [58]. Note that this is the analog of the Schrödinger equation for the maximal Lyapunov exponent in the vanilla RNN. 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 rhs of Eqs. (10) and (11) by simple correlation functions, for instance, $\partial C_\phi(\tau)/\partial C_h(\tau) = C_{\phi'}(\tau)$. In this limit, we see good agreement between the numerically calculated maximal Lyapunov exponent [Fig. 3(c), dots] compared to the DMFT prediction [Fig. 3(c), solid line] obtained by solving the eigenvalue problem [Eqs. (10) and (11)]. For large values of α_z , we see quantitative deviations between the DMFT prediction and the true λ_{\max} . Indeed, for large α_z , the distribution of h is strongly non-Gaussian, and there is no reason to expect that variational formulas given by Price’s theorem are even approximately correct. For more on this point, see the discussion toward the end of Appendix C.

2. 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, across a continuous transition to chaos, the point at which λ_{\max} becomes positive is not dependent on α_z (data not shown). 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. (6), and we see that it has no dependence on α_z or the update gate. We stress that this condition [Eq. (6)] 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. 3(c), $\alpha_r = 0, 2$]. However, in the gated RNN, there is another discontinuous

transition to chaos [Fig. 3(c), $\alpha_r = 20$]: 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). To our knowledge, this is a novel type of transition which is not present in the vanilla RNN and not visible from an analysis that considers only the stability of fixed points. We characterize this phenomenon in detail below.

B. Output gate induces a novel chaotic transition

Here, we describe a novel phase, characterized by a proliferation of unstable fixed points and the coexistence of a stable fixed point with chaotic dynamics. It is the appearance of this state that gives rise to the discontinuous transition observed in Fig. 3(c). The appearance of this state is mediated by the output gate becoming more switchlike (i.e., increasing α_r) in the quiescent region for g_h . To our knowledge, no such comparable phenomenon exists in RNNs with additive interactions. The full details of the calculations for this transition are provided in Appendix G. Here, we simply state and describe the salient features. For ease of presentation, the rest of the section assumes that all biases are zero. The results in this section are strictly valid only for $\alpha_z = 0$. In Appendix G 3, we argue that they should also hold for moderate α_z .

This discontinuous transition is characterized by a few noteworthy features.

1. Spontaneous emergence of fixed points

When $g_h < 2.0$, the zero fixed point is stable. Moreover, for $\sqrt{2} < g_h < 2$, when α_r crosses a threshold value $\alpha_{r,FP}^*(g_h)$, unstable fixed points spontaneously appear in the phase space. The only dynamical signature of these unstable FPs are short-lived transients which do not scale with system size (see Fig. 11). Thus, we have a condition for fixed-point transition:

$$\sqrt{2} < g_h \leq 2 \quad \text{and} \quad \alpha_r > \alpha_{r,FP}^*(g_h). \quad (12)$$

These unstable fixed points correspond to the emergence of nontrivial solutions to the time-independent MFT. Figure 4(a) shows the appearance of fixed-point MFT solutions for a fixed g_h , and Fig. 4(b) shows the critical $\alpha_{r,FP}^*(g_h)$ as a function of g_h . As $g_h \rightarrow 2^-$, we see that $\alpha_{r,FP}^* \rightarrow \sqrt{8}$.

These spontaneous MFT fixed-point solutions are unstable according to the criterion Eq. (6) derived from RMT. Moreover, in Appendix J, using a Kac-Rice analysis, we show that in this region the full $3N$ -dimensional system does indeed have a number of unstable fixed points that grows exponentially fast with N . Thus, this transition line $\alpha_{r,FP}^*$ represents a topological trivialization transition as conceived by, e.g., Refs. [59,60]. Our analysis shows that instability is intimately connected to the proliferation of

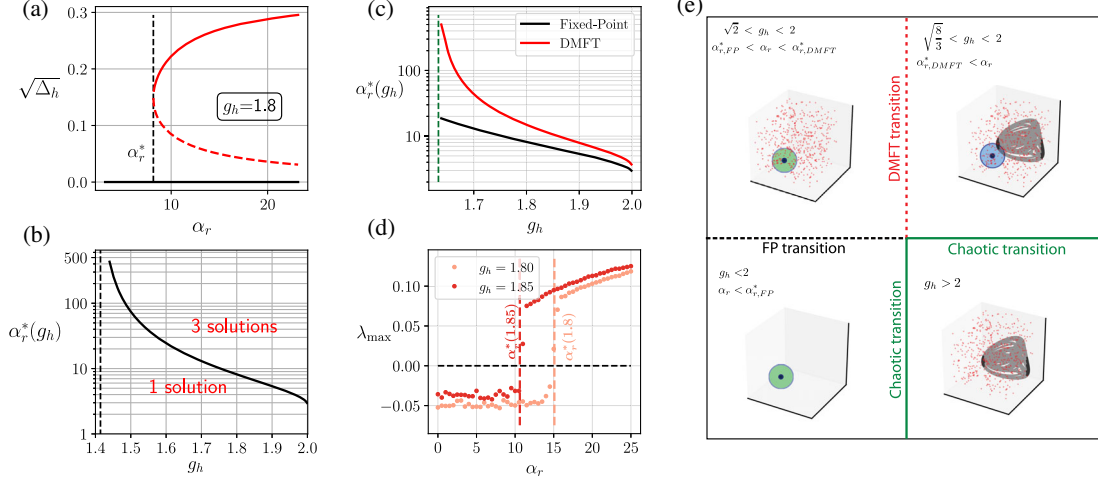


FIG. 4. The discontinuous dynamical transition. (a) Spontaneous appearance of nonzero solutions (dashed and solid red lines) to the FP equations once α_r crosses a critical value $\alpha_{r,FP}^*(g_h)$ at fixed g_h . (b) The critical $\alpha_{r,FP}^*(g_h)$ as a function of g_h . The vertical dashed line represents left critical value $g_c = \sqrt{2}$, below which a bifurcation is not possible. (c) The critical DMFT transition curve $\alpha_{r,DMFT}^*(g_h)$ (red curve) calculated using Eqs. (G8) and (G9). The FP transition curve from (b) is shown in black. The green dashed line corresponds to $g_c = \sqrt{8/3}$, below which the dynamical transition is not possible. (d) 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 (c). (e) Schematic of the bifurcation transition: For $g_h < 2$ and $\alpha_r < \alpha_{r,FP}^*$, the zero FP is the only (stable) solution (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 obvious 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).

fixed points. Remarkably, however, a time-dependent solution to the DMFT does *not* emerge across this transition, and the microscopic dynamics are insensitive to the transition in topological complexity, bringing us to the next point.

2. A delayed dynamical transition that shows a decoupling between topological and dynamical complexity

On increasing α_r beyond $\alpha_{r,FP}^*$, there is a second transition when α_r crosses a critical value $\alpha_{r,DMFT}^*$. This happens when we satisfy the condition for dynamical transition:

$$\sqrt{\frac{8}{3}} < g_h \leq 2 \quad \text{and} \quad \alpha_r > \alpha_{r,DMFT}^*(g_h), \quad (13)$$

derived in Appendix G 2. Figure 4(c) shows how $\alpha_{r,DMFT}^*(g_h)$ varies with g_h . As $g_h \rightarrow 2^-$, we see that $\alpha_{r,DMFT}^* \rightarrow \sqrt{12}$. Across this transition, a dynamical state spontaneously emerges, and the maximum Lyapunov exponent jumps from a negative value to a positive value [Fig. 4(d)]. This state exhibits chaotic dynamics that coexist with the stable zero fixed point. The presence of the stable FP means that the dynamical state is not strictly a chaotic attractor but rather a spontaneously appearing “chaotic set.” On increasing g_h beyond 2.0, for large but fixed α_r , the stable fixed point disappears, and the state smoothly transitions into a full chaotic attractor that is characterized above. This picture is summarized in the schematic in Fig. 4(e). This gap between the proliferation of unstable fixed points and the appearance of the chaotic dynamics differs from the result of Wainrib and Touboul [45] for purely additive RNNs, where the proliferation (topological complexity) is tightly linked to the chaotic dynamics (dynamical complexity). Thus, for gated RNNs, there appears to be another distinct mechanism for the transition to chaos, and the accompanying transition is a discontinuous one.

3. Long chaotic transients

For finite systems, across the transition the dynamics eventually flow into the zero FP after chaotic transients. Moreover, 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 the stable fixed point still exists. This is because the relative volume of the basin of attraction of the fixed point vanishes as $N \rightarrow \infty$. In Appendix G [Figs. 11(c) and 11(d)], we do indeed see that the transient time for a fixed g_h scales with system size [Fig. 11(c)] once α_r is above the second transition (dashed line) and not otherwise [see Figs. 11(a) and 11(e), dashed lines].

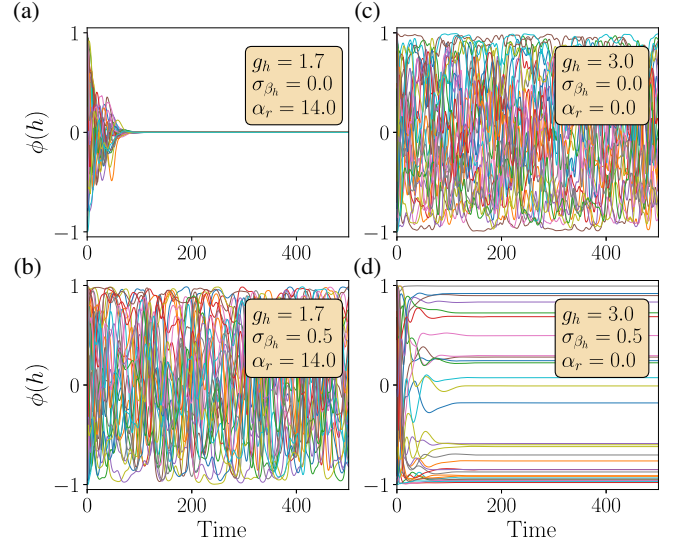


FIG. 5. Input-driven chaos. (a),(b) Near the discontinuous chaotic transition (in region 2 in Fig. 7), static input I^h (with $I^r = I^z = 0$) can push a stable system (a) to chaotic activity (b). (c),(d) In the purely chaotic state [(c), $g_h = 3.0$], input has the familiar effect of stabilizing the dynamics (d). The elements of the input vector I^h are random Gaussian variables with zero mean and variance $\sigma_{\beta_h}^2$.

4. An input-induced chaotic transition

The discontinuous chaotic transition has a functional consequence: Near the transition, static inputs can push a stable system to strong chaotic activity. This is in contrast to the typically stabilizing effects of inputs on the activity of random additive RNNs [21,43,44]. In Figs. 5(a) and 5(b), we see that, when static input with variance σ_{β_h} is applied to a stable system (a) near the discontinuous chaotic transition (in region 2 in Fig. 7), it induces chaotic activity (b); however, for the same input when applied to the system in the chaotic state [Fig. 5(c)], the dynamics are stabilized (d) as reported before.

This phenomenon for static inputs can be understood using the phase diagram with nonzero biases, discussed in Sec. VII. There, we see how the transition curves move when a random bias β_h is included. Near the classic chaotic transition ($\alpha_r = 0$), the bias moves the curve toward larger g_h , thus suppressing chaos. Near the discontinuous chaotic transition $\alpha_{r,DMFT}^*$, the bias pulls the curve toward smaller values of α_r , thus *promoting* chaos. Thus, inputs can have opposite effects of inducing or stabilizing chaos within the same model in different parameter regimes. This phenomenon could, in principle, be leveraged for shaping the interaction between inputs and internal dynamics.

VI. GATES PROVIDE A FLEXIBLE RESET MECHANISM

Here, we discuss how the gates provide another critical function—a mechanism to flexibly reset the memory trace

depending on external input or the internal state. This function complements the memory function; a memory that cannot be erased when needed is not very useful. To build intuition, let us consider a linear network $\dot{\mathbf{h}} = -\mathbf{h} + J\mathbf{h}$, where the matrix $-\mathbb{1} + J$ has a few eigenvalues that are zero, while the rest have a negative real part. The slow modes are good for memory function; however, that fact also makes it hard to forget memory traces along the slow modes. This trade-off is pointed out in Ref. [61]. To be functionally useful, it is critical that the memory trace can be erased flexibly in a context-dependent manner. The r gate allows this function naturally. Consider the same net, but now augmented with an r gate: $\dot{\mathbf{h}} = -\mathbf{h} + J\mathbf{h} \odot \sigma_r$. If the gate is turned off ($\sigma_r = 0$) for a short duration, the state \mathbf{h} is reset to zero. One can actually be more specific: We may choose a $J^r = -\mathbb{1}\mathbf{u}^T$ with $\sigma_r = \sigma[J^r\phi(\mathbf{h})]$, such that the r gate turns off whenever $\phi(\mathbf{h})$ gets aligned with \mathbf{u} , thus providing an internal-context-dependent reset.

Apart from resetting to zero, the z gate also allows the possibility of rapidly scrambling the state to a random value by means of the input-induced chaos. This phenomenon is illustrated in Fig. 6, where the network in the marginally stable state normally functions as a memory (retains traces for long times, as in Fig. 2), but positive inputs I^z (with $I^h = I^r = 0$) to the z gate above a threshold strength even for a short duration can induce chaos, thereby scrambling the state and erasing the previous memory state (Fig. 6, bottom panel). The mechanism for this scrambling can be understood by appealing to Eq. (8). A finite input I^z with nonzero mean is able to change $\langle\sigma(z)\rangle$ and, thus, push the critical line for

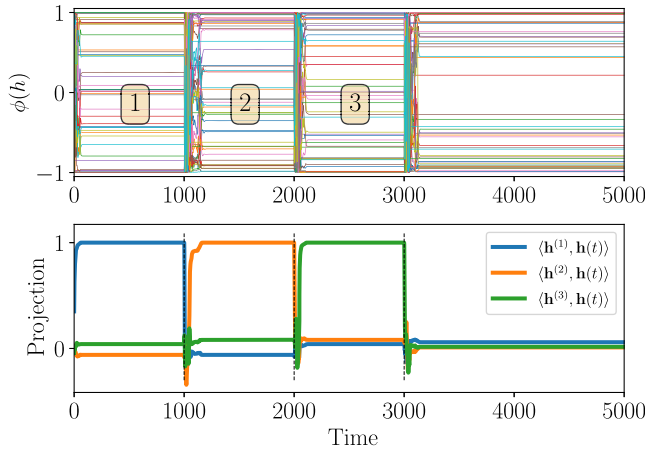


FIG. 6. Gates provide a reset mechanism. Positive static inputs are applied to the z gate when the RNN is in the marginally stable state ($g_h = 3.0$, $\alpha_r = 2.5$, and $\alpha_z = \infty$) for 20 time units at times indicated by dashed lines. The input induces chaos which rapidly scrambles the network state, thus erasing the trace of the previous memory; the bottom panel shows the normalized projection of the state $\mathbf{h}(t)$ on the directions $\mathbf{h}^{(1,2,3)}$ aligned with the state in regions 1, 2, and 3.

marginal stability in one way or the other. For instance, if $\langle I^z \rangle > 0$, $\langle \sigma(z) \rangle > 1/2$, which (for $\alpha_r = 0$) moves the transition to marginal stability to a smaller value of g_h . This implies that a marginally stable state can be made chaotic in the presence of I^z with finite mean. This mechanism for input-induced chaos actually appears to be different from that explored in the previous section, which occurs across the discontinuous chaotic transition. We discuss this more in Sec. VII.

In summary, gating imbues the RNN with the capacity to flexibly reset memory traces, providing an “inductive bias” for context-dependent reset. The specific method of reset depends on the task or function, and this can be selected, e.g., by gradient-based training. This inductive bias for resetting is found to be critical for performance in ML tasks [62].

VII. PHASE DIAGRAMS FOR THE GATED NETWORK

Here, we summarize the rich dynamical phases of the gated RNN 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 continuous or critical chaotic transitions. On the other hand, it influences the discontinuous transition to chaos for sufficiently large values of α_z (see Sec. G 3 for discussion). Furthermore, the update gate has a strong effect on the dynamical aspects of the states near the critical lines. There are macroscopic regions of the parameter space adjacent to the critical lines where the states can be made marginally stable in the limit of $\alpha_z \rightarrow \infty$. The shape of this marginal stability region is influenced by β_z and I^z .

Figure 7(a) shows the dynamical phases for the network with no biases in the (g_h, α_r) plane. When g_h is below 2.0 and $\alpha_r < \alpha_{r,\text{FP}}^*$, the zero fixed point is the only solution (region 1). As discussed in Sec. V B, on crossing the fixed-point bifurcation line [green line, Fig. 7(a)], there is a spontaneous proliferation of unstable fixed points in the phase space (region 2). This can occur only when $g_h > \sqrt{2}$. The 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). When g_h is increased beyond the critical value of 2.0, the stable zero fixed point becomes unstable for all α_r , and we get a chaotic attractor (region 4). All the 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 5 in Fig. 7(a)].




FIG. 7. Phase diagram for the gated RNN. (a) (no biases) In regions 1 and 2, the zero FP is the global attractor of dynamics; however, in region 2, there is a proliferation of unstable FPs without any asymptotic dynamical signatures. In region 3, the (stable) zero FP coexists with chaotic dynamics. Note that the plotted curve separating regions 2 and 3 is computed for $\alpha_z = 0$ and remains valid for sufficiently small values of α_z . In region 4, the zero FP is unstable, and dynamics are chaotic. For all parameter values in region 5, a previously unstable or chaotic state can be made marginally stable when $\alpha_z = \infty$. For any given parameter values in region 5, there are infinitely many marginally stable points in the phase space to which the dynamics converge. The red dashed line indicates the critical transition between a stable fixed point (below the line) and chaos (above the line) in the presence of static random input (to the h variable) with standard deviation $\sigma_h = 0.5$. Note that, while chaos is suppressed for small α_r along the g_h axis, for larger α_r there are regions of stable FPs that become chaotic with finite input. This leads to the phenomenon of input-induced chaos.

A. Role of biases and static inputs

Biases have the effect of generating nontrivial fixed points and controlling stability by moving the edge of the spectral curve. Another key feature of biases is the suppression of the discontinuous bifurcation transition observed without biases. This is explained in detail in Appendix H. A particularly illuminating illustration of the effects of a bias can be inferred from the critical line (red dashed) for finite bias shown in Fig. 7. This curve, computed using the FP stability criterion (6) combined with the MFT equations [(C28)–(C30)], represents the transition between stability and chaos for finite bias with zero mean and nonzero variance. Equivalently, we may think of this as the critical line for a network with static input $I_i^h \sim \mathcal{N}(0, \sigma_h^2)$ (with $I_r = I_z = 0$). Along the g_h axis, we can observe the well-documented phenomena whereby an input suppresses chaos. This corresponds to the region $g_h > 2$ which lies to the left of the red dashed critical line, which is chaotic in the absence of input and flows to a stable fixed point in the presence of input. However, this behavior is reversed for $g_h < 2$. Here, we see a significant swath of phase space which is stable in the absence of input

but which becomes chaotic when input is present. Thus, the stability-to-chaos phase boundary in the presence of biases (or inputs) reveals that the output (r) gate can facilitate an input-induced transition to chaos.

VIII. DISCUSSION

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 prominent feature of biological neurons. Prior theoretical work on RNNs has considered only 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 that have significant functional benefits.

The continuous-time gated RNN (gRNN) we study resembles a popular model used in machine learning applications, the gated recurrent unit (GRU) [see the note below Eq. (C27)]. Previous work [51] looks at the instantaneous Jacobian spectrum for the discrete-time GRU using RMT methods similar to those presented in Appendix A; however, this work does not go beyond time-independent MFT and presents a phase diagram showing only boundaries across which the MFT fixed points become unstable [63]. In the present manuscript, we illuminate the full *dynamical* phase diagram for our gated RNN, uncovering much richer structure. Both the GRU and our gRNN have a gating function which dynamically scales the time constant, which in both cases leads to a marginally stable phase in the limit of a binary gate. However, the dynamical phase diagram presented in Fig. 7 reveals a novel discontinuous transition to chaos. We conjecture that such a phase transition should also be present in the GRU. Also, Ref. [51] lacks any discussion of the influence of inputs or biases. The present paper includes discussion of the functional significance of the gates in the presence of *inputs*. We believe these results, combined with the refined dynamical phase diagram, can shed some light on the role of analogous gates in the GRU and other gated ML architectures. We review the significance of the gates in more detail below.

A. Significance of the update gate

The update gate modulates the rate of integration. In single-neuron models, such a modulation is shown to make the neuron’s responses robust to time-warped inputs [14]. Furthermore, normative approaches, requiring time reparametrization invariance in ML RNNs, naturally imply the existence of a mechanism that modulates the integration rate [64]. We show that, for a wide range of parameters, a more sensitive (or switchlike) update gate leads to marginal stability. Marginally stable models of biological function have long been of interest with regard to their benefits for information processing (cf. Ref. [33] and references therein). In the gated RNN, a functional consequence of

the marginally stable state is the use of the network as a robust integrator—such integratorlike function is shown to be beneficial for a variety of computational functions such as motor control [34–36], decision making [37], and auditory processing [53]. However, previous models of these integrators often require handcrafted symmetries and fine-tuning [38]. We show that gating allows this function without fine-tuning. Signatures of integratorlike behavior are also empirically observed in successfully trained gated ML RNNs on complex tasks [41]. We provide a theoretical basis for how gating produces these. The update gate facilitates accumulation of slow modes and a pinching of the spectral curve which leads to a suppression of unstable directions and overall slowing of the dynamics over a range of parameters. This is a manifestly self-organized slowing down. Other mechanisms for slowing down dynamics have been proposed where the slow timescales of the network dynamics are inherited from other slow internal processes such as synaptic filtering [65,66]; however, such mechanisms differ from the slowing due to gating; they do not seem to display the pinching and clumping, and they also do not rely on self-organized behavior.

B. Significance of the output gate

The output gate dynamically modulates the outputs of individual neurons. Similar shunting mechanisms are widely observed in real neurons and are crucial for performance in ML tasks [62]. We show that the output gate offers fine control over the dimensionality of the dynamics in phase space, and this ability is implicated in task performance in ML RNNs [42]. This gate also gives rise to a novel discontinuous chaotic transition where inputs can abruptly push stable systems to strongly chaotic activity; this is in contrast to the typically stabilizing role of inputs in additive RNNs. In this transition, there is a decoupling between topological and dynamical complexity. The chaotic state across this transition is also characterized by the coexistence of a stable fixed point with chaotic dynamics—in finite size systems, this manifests as long transients that scale with the system size. We note that there are other systems displaying either a discontinuous chaotic transition or the existence of fixed points overlapping with chaotic (pseudo)attractors [19] or apparent chaotic attractors with finite alignment with particular directions [67]; however, we are not aware of a transition in large RNNs where static inputs can induce strong chaos or the topological and dynamical complexity are decoupled across the transition. In this regard, the chaotic transition mediated by the output gated seems to be fundamentally different. More generally, the output gate is likely to have a significant role in controlling the influence of external inputs on the intrinsic dynamics.

We also show how the gates complement the memory function of the update gate by providing an important, context- and input-dependent *reset* mechanism. The ability

to erase a memory trace flexibly is critical for function [62]. Gates also provide a mechanism to avoid the accuracy-flexibility trade-off noted for purely additive networks—where the stability of a memory comes at the cost of the ease with which it is rewritten [61].

We summarize the rich behavior of the gated RNN via phase diagrams indicating the critical surfaces and regions of marginal stability. From a practical perspective, the phase diagram is useful 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 [68–70]. Thus, the phase diagrams provide ML practitioners with a map for principled parameter initialization—one of the most critical choices deciding training success.

ACKNOWLEDGMENTS

K. K. is supported by a C. V. Starr fellowship and a CPBF fellowship (through NSF PHY-1734030). T. C. is supported by a grant from the Simons Foundation (891851, TC). D. J. S. 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 Grant No. R01EB026943. K. K. and D. J. S. thank the Simons Institute for the Theory of Computing at U. C. Berkeley, where part of the research was conducted. T. C. gratefully acknowledges the support of the Initiative for the Theoretical Sciences at the Graduate Center, CUNY, where most of this work was completed. We are most grateful to William Bialek, Jonathan Cohen, Andrea Crisanti, Rainer Engelken, Moritz Helias, Jonathan Kadmon, Jimmy Kim, Itamar Landau, Wave Ngampruetikorn, Katherine Quinn, Friedrich Schuessler, Julia Steinberg, and Merav Stern for fruitful discussions.

K. K. and D. J. S. conceived the overall project. K. K. and T. C. developed the theory and did all the analyses. K. K. and T. C. wrote the paper. All authors discussed the results and the paper.

APPENDIX A: DETAILS OF RANDOM MATRIX THEORY FOR SPECTRUM OF THE JACOBIAN

In this section, we provide details of the calculation of the bounding curve for the Jacobian spectrum for both fixed points and time-varying states. Our approach to the problem utilizes the method of Hermitian reduction [25,26] to deal with non-Hermitian random matrices. The analysis here resembles that in Ref. [51], which also considers Jacobians that are highly structured random matrices arising from discrete-time gated RNNs.

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

approaches 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

$$\mu(\lambda) = \frac{1}{3N} \mathbb{E} \left[\sum_{i=1}^{3N} \delta(\lambda - \lambda_i) \right], \quad (\text{A1})$$

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

$$\mu(\lambda) = \frac{1}{\pi} \frac{\partial}{\partial \bar{\lambda}} \mathbb{E} \left[\frac{1}{3N} \text{Tr}[(\lambda \mathbb{1}_{3N} - \mathcal{D})^{-1}] \right], \quad (\text{A2})$$

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 $G(\lambda, \bar{\lambda}) = (3N)^{-1} \text{Tr} \mathbb{E}[(\lambda \mathbb{1}_{3N} - \mathcal{D})^{-1}]$ by analytic continuation is not applicable, since it is nonholomorphic on the support. Instead, we use the method of Hermitization [25,26] to analyze the resolvent for an expanded $6N \times 6N$ Hermitian matrix H :

$$\mathcal{G}(\eta, \lambda, \bar{\lambda}) = \mathbb{E}[(\eta \mathbb{1}_{6N} - H)^{-1}], \quad (\text{A3})$$

$$H = \begin{pmatrix} 0 & \lambda - \mathcal{D} \\ \bar{\lambda} - \mathcal{D}^T & 0 \end{pmatrix}, \quad (\text{A4})$$

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

$$G(\lambda, \bar{\lambda}) = \lim_{\eta \rightarrow i0^+} \frac{1}{3N} \text{Tr} \mathcal{G}_{21}(\eta, \lambda, \bar{\lambda}). \quad (\text{A5})$$

To make this problem tractable, we invoke an ansatz called the local chaos hypothesis [57,71], 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 Ref. [72]). This implies that the Jacobian [Eq. (4)] has an explicit linear dependence only on $J^{h,z,r}$, and the state variables are governed by their steady-state distribution from the disorder-averaged DMFT (Appendix C). 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 $N \rightarrow \infty$. We detail this procedure below.

The Jacobian itself can be decomposed into structured (A, L, R) and random parts (\mathcal{J}):

$$\mathcal{D} = \underbrace{\begin{pmatrix} -[\sigma_z] & \mathbb{D} & 0 \\ 0 & -\tau_z^{-1} \mathbb{1} & 0 \\ 0 & 0 & -\tau_r^{-1} \mathbb{1} \end{pmatrix}}_A + \underbrace{\begin{pmatrix} [\sigma_z] & 0 & 0 \\ 0 & \tau_z^{-1} \mathbb{1} & 0 \\ 0 & 0 & \tau_r^{-1} \mathbb{1} \end{pmatrix}}_L \\ \times \underbrace{\begin{pmatrix} J^h & 0 & 0 \\ 0 & J^z & 0 \\ 0 & 0 & J^r \end{pmatrix}}_{\mathcal{J}} \underbrace{\begin{pmatrix} [\phi' \sigma_r] & 0 & [\phi \sigma'_r] \\ [\phi'] & 0 & 0 \\ [\phi'] & 0 & 0 \end{pmatrix}}_R. \quad (\text{A6})$$

At this point, we must make a crucial assumption: The structured matrices A , L , and R are *independent* of the random matrices appearing \mathcal{J} . This implies that the dynamics is *self-averaging* and that the state variables reach a steady-state distribution determined by the DMFT and insensitive to the particular quenched disorder \mathcal{J} . This self-averaging assumption leads to theoretical predictions which are in very good agreement with simulations of large networks, as presented in Fig. 1.

This independence assumption renders \mathcal{D} a linear function of the random matrix \mathcal{J} , whose entries are Gaussian random variables. The next steps are to develop an asymptotic series in the random components of H , compute the resulting moments, and perform a resummation of the series. This is conveniently accomplished by the self-consistent Born approximation (SCBA). The output of the SCBA is a self-consistently determined self-energy functional $\Sigma[\mathcal{G}]$ which succinctly encapsulates the resummation of moments. With this, the Dyson equation for \mathcal{G} is given by

$$\mathcal{G}^{-1} = \mathcal{G}_0^{-1} - \Sigma[\mathcal{G}], \quad (\text{A7})$$

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

$$\mathcal{G}_0^{-1} = \begin{pmatrix} \eta \mathbb{1} & \lambda - A \\ \bar{\lambda} - A^T & \eta \mathbb{1} \end{pmatrix}, \quad (\text{A8})$$

$$\Sigma[\mathcal{G}] = \begin{pmatrix} LQ[RG_{22}R^T]L & 0 \\ 0 & R^T Q[L^T \mathcal{G}_{11} L] R \end{pmatrix}, \quad (\text{A9})$$

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

$$Q[M] = \begin{pmatrix} \frac{1}{N} \text{Tr} M_{11} & 0 & 0 \\ 0 & \frac{1}{N} \text{Tr} M_{22} & 0 \\ 0 & 0 & \frac{1}{N} \text{Tr} M_{33} \end{pmatrix}. \quad (\text{A10})$$

Here, we express the self-energy using the $3N \times 3N$ subblocks of the Green's function \mathcal{G} :

$$\mathcal{G} = \begin{pmatrix} \mathcal{G}_{11} & \mathcal{G}_{12} \\ \mathcal{G}_{21} & \mathcal{G}_{22} \end{pmatrix}. \quad (\text{A11})$$

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 (A7), along with the free Green's function (A8) and the self-energy (A9). Most of what is left is complicated linear algebra. However, in the interest of completeness, we 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, it is useful to define the following transformed Green's functions, which can be written in terms of $N \times N$ subblocks:

$$\tilde{\mathcal{G}}_{11} \equiv L^T \mathcal{G}_{11} L = \begin{pmatrix} \tilde{G}_{11} & \tilde{G}_{12} & \tilde{G}_{13} \\ \tilde{G}_{21} & \tilde{G}_{22} & \tilde{G}_{23} \\ \tilde{G}_{31} & \tilde{G}_{32} & \tilde{G}_{33} \end{pmatrix}, \quad (\text{A12})$$

$$\tilde{\mathcal{G}}_{22} \equiv R \mathcal{G}_{22} R^T = \begin{pmatrix} \tilde{G}_{44} & \tilde{G}_{45} & \tilde{G}_{46} \\ \tilde{G}_{54} & \tilde{G}_{55} & \tilde{G}_{56} \\ \tilde{G}_{64} & \tilde{G}_{65} & \tilde{G}_{66} \end{pmatrix}. \quad (\text{A13})$$

Denote also the mean trace of these subblocks as

$$\tilde{g}_{ij} = \frac{1}{N} \text{Tr}[\tilde{G}_{ij}]. \quad (\text{A14})$$

Then the self-energy matrix in Eq. (A9) is block diagonal, i.e., $\Sigma[\mathcal{G}] = \text{bdiag}(\Sigma_{11}, \Sigma_{22})$, with

$$\Sigma_{11} = \begin{pmatrix} [\sigma_z^2] \tilde{g}_{44} & 0 & 0 \\ 0 & \tau_z^{-2} \tilde{g}_{55} & 0 \\ 0 & 0 & \tau_r^{-2} \tilde{g}_{66} \end{pmatrix}, \quad (\text{A15})$$

$$\Sigma_{22} = \begin{pmatrix} [\phi' \sigma_r]^2 \tilde{g}_{11} + [\phi']^2 (\tilde{g}_{22} + \tilde{g}_{33}) & 0 & [\phi' \sigma_r] [\phi \sigma_r'] \tilde{g}_{11} \\ 0 & 0 & 0 \\ [\phi' \sigma_r] [\phi \sigma_r'] \tilde{g}_{11} & 0 & [\phi \sigma_r']^2 \tilde{g}_{11} \end{pmatrix}. \quad (\text{A16})$$

With the self-energy in this form, we are able to solve the Dyson equation for the full Green's function \mathcal{G} by direct matrix inversion:

$$\mathcal{G} = \begin{pmatrix} \eta - \Sigma_{11} & \lambda - A \\ \bar{\lambda} - A^T & \eta - \Sigma_{22} \end{pmatrix}^{-1}, \quad (\text{A17})$$

which can be carried out easily by symbolic manipulation software. The rhs of Eq. (A17) is a function of \tilde{g}_{ii} , whereas the lhs is a function of the Green's function before the transformations (A12) and (A13). Thus, to get a set of

equations we can solve, we apply these same transformations to both sides of Eq. (A17) after solving the Dyson equation. The final step is to take the limit $\eta \rightarrow 0$, recovering the problem we originally wished to solve.

The result of these manipulations is a set of six equations for the mean traces of the transformed Green's function defined in Eq. (A14). In order to write these down, we introduce some additional notation. The self-consistent equations are of the form

$$\tilde{g}_{ii} = \left\langle \frac{\Gamma_i}{\Gamma} \right\rangle, \quad (\text{A18})$$

where we denote $\langle M \rangle \equiv N^{-1} \text{Tr} M$ for shorthand and i runs from 1 to 6. Denote the state-variable-dependent diagonal matrices as

$$p = [\phi'], \quad q = [\phi \sigma_r'], \quad r = [\phi' \sigma_r], \quad (\text{A19})$$

and, because they appear frequently in the resulting equations, define

$$X = \tilde{g}_{11} |\lambda \tau_r + 1|^2 r^2 + (\tilde{g}_{22} + \tilde{g}_{33}) p^2 Z, \quad (\text{A20})$$

$$Y = \mathbb{D}^2 \tilde{g}_{55} + |\lambda \tau_z + 1|^2 [\sigma_z^2] \tilde{g}_{44}, \quad (\text{A21})$$

$$Z = |\lambda \tau_r + 1|^2 - \tilde{g}_{11} \tilde{g}_{66} q^2. \quad (\text{A22})$$

The denominator in Eq. (A18) is then given by

$$\Gamma = |\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 Z - XY, \quad (\text{A23})$$

and the numerators Γ_i are given by

$$\Gamma_1 = \sigma_z^2 |\lambda \tau_z + 1|^2 X, \quad (\text{A24})$$

$$\Gamma_2 = \mathbb{D}^2 X, \quad (\text{A25})$$

$$\Gamma_3 = \tilde{g}_{11} |\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 q^2 - \tilde{g}_{11} (\tilde{g}_{22} + \tilde{g}_{33}) p^2 q^2 Y, \quad (\text{A26})$$

$$\Gamma_4 = \tilde{g}_{66} |\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 q^2 + [|\lambda \tau_r + 1|^2 r^2 - \tilde{g}_{66} (\tilde{g}_{22} + \tilde{g}_{66}) p^2 q^2] Y, \quad (\text{A27})$$

$$\Gamma_5 = \Gamma_6 = p^2 YZ. \quad (\text{A28})$$

The numerators and denominator are all diagonal matrices with real entries, which is why we use the simple notation of a ratio when referring to matrix inversion.

Solving these equations gives us the \tilde{g}_{ii} as implicit functions of λ . They are, in general, complicated and resist exact solution. 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 $\tilde{g}_{ii} \rightarrow 0$.

We must take this limit carefully, since the ratio of these functions can remain constant. For this reason, it is necessary to define

$$x_2 = \tilde{g}_{22}/\tilde{g}_{11}, \quad x_3 = \tilde{g}_{33}/\tilde{g}_{11}. \quad (\text{A29})$$

We may do the same for \tilde{g}_{44} , \tilde{g}_{55} , and \tilde{g}_{66} , but it turns out that x_2 and x_3 are sufficient to compute the spectral curve. Next, divide by \tilde{g}_{11} and send all $\tilde{g}_{ii} \rightarrow 0$, keeping the ratios fixed. Applying this to the equation for \tilde{g}_{11} results in

$$1 = \lim_{\tilde{g}_{ii} \rightarrow 0} \frac{1}{\tilde{g}_{11}} \left\langle \frac{\Gamma_1}{\mathbf{\Gamma}} \right\rangle = \gamma_1 + \gamma_2(x_2 + x_3). \quad (\text{A30})$$

Similarly, for \tilde{g}_{22} and \tilde{g}_{33} , we get

$$x_2 = \gamma_3 + \gamma_4(x_2 + x_3), \quad (\text{A31})$$

$$x_3 = \gamma_5, \quad (\text{A32})$$

where the coefficients γ_i , which are functions of λ , are given by

$$\gamma_1 = \left\langle \frac{\sigma_z^2 r^2}{|\lambda + \sigma_z|^2} \right\rangle, \quad \gamma_2 = \left\langle \frac{p^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle, \quad \gamma_5 = \frac{\langle q^2 \rangle}{|\lambda \tau_r + 1|^2},$$

$$\gamma_3 = \left\langle \frac{\mathbb{D}^2 r^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle, \quad \gamma_4 = \left\langle \frac{\mathbb{D}^2 p^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle.$$

The linear system of equations (A30)–(A32) is consistent iff

$$(1 - \gamma_1)(1 - \gamma_4) = \gamma_2(\gamma_3 + \gamma_5). \quad (\text{A33})$$

In other words, γ_i must satisfy Eq. (A33) when $\tilde{g}_{ii} \rightarrow 0$. This expression depends on λ and implicitly defines a curve in \mathbb{C} , which is the boundary of the support of the spectral density.

Plugging in the explicit expression for γ_i , we get the implicit equation for the spectral curve as all $\lambda \in \mathbb{C}$ that satisfy

$$\left\{ 1 - \left\langle \frac{r^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle \right\} \left\{ 1 - \left\langle \frac{\mathbb{D}^2 p^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle \right\} = \left\langle \frac{\sigma_z^2 p^2}{|\lambda + \sigma_z|^2} \right\rangle \left\{ \left\langle \frac{\mathbb{D}^2 r^2}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2} \right\rangle + \frac{\langle q^2 \rangle}{|\lambda \tau_r + 1|^2} \right\}. \quad (\text{A34})$$

For large systems, we can replace the empirical traces of the state variable by their averages given by the DMFT variances. Then, the equation for the curve for a general steady state is given by




FIG. 8. Jacobian spectrum at a time-varying state. Red dots are the Jacobian eigenvalues for the full network in a (time-varying) steady state, and the spectral curve of the Jacobian is calculated using moments from (i) the full state vectors (blue curve) or using the variances from the fixed-point MFT (green). Surprisingly, the agreement is reasonably good. For network simulations, $N = 1000$, $g_h = 2.5$, $\alpha_r = 1$, $\alpha_z = 15$, and all biases are zero.

$$\left(\langle \sigma_r^2 \rangle + \frac{\langle \phi^2 \sigma_r'^2 \rangle}{|1 + \tau_r \lambda|^2} \right) \left\langle \frac{\phi'^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle + \frac{1}{|1 + \tau_z \lambda|^2} \left\langle \frac{\mathbb{D}^2 \phi'^2}{|\lambda + \sigma_z|^2} \right\rangle = 1. \quad (\text{A35})$$

For fixed points, we have $\mathbb{D} = 0$, which makes $\gamma_3 = \gamma_4 = 0$. The equation for the spectral curve simplifies to that which is quoted in the main text [Eq. (5)]:

$$1 = \left\langle \frac{r^2 \sigma_z^2}{|\lambda + \sigma_z|^2} \right\rangle + \frac{\langle q^2 \rangle}{|\lambda \tau_r + 1|^2} \left\langle \frac{\sigma_z^2 p^2}{|\lambda + \sigma_z|^2} \right\rangle. \quad (\text{A36})$$

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. (A5)] in a simpler form. Recall that

$$G(\lambda, \bar{\lambda}) = \lim_{\eta \rightarrow 0^+} \frac{1}{3N} \text{tr} \mathcal{G}_{21}(\eta, \lambda, \bar{\lambda}). \quad (\text{A37})$$

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

$$G(\lambda, \bar{\lambda}) = \frac{1}{3} \left\langle \frac{|\lambda \tau_z + 1|^2 (\bar{\lambda} + \sigma_z)}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) p^2 \tilde{Y}} \right\rangle + \frac{1}{3} \left\langle \frac{(\bar{\lambda} + \tau_z^{-1})(|\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) p^2 \sigma_z^2)}{|\lambda + \tau_z^{-1}|^2 |\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) p^2 \tilde{Y}} \right\rangle \quad (\text{A38})$$

$$+ \frac{1}{3} \frac{1}{\lambda + \tau_r^{-1}}, \quad (\text{A40})$$

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

$$1 = \left\langle \frac{p^2 \bar{Y}}{|\lambda \tau_z + 1|^2 |\lambda + \sigma_z|^2 - \xi(\lambda, \bar{\lambda}) p^2 \bar{Y}} \right\rangle. \quad (\text{A41})$$

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:

$$G(\lambda, \bar{\lambda}) = \frac{1}{3} \left(\left\langle \frac{1}{\lambda + \sigma_z} \right\rangle + \frac{1}{\lambda + \tau_z^{-1}} + \frac{1}{\lambda + \tau_r^{-1}} \right), \quad (\text{A42})$$

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 nonanalytic 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 nonanalytic solution, at $\xi = 0$. This is the condition used to find the boundary curve above.

APPENDIX B: 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 switchlike ($\alpha_z \rightarrow \infty$). To focus on the key aspects of these phenomena, we consider the case when the reset gate is off and there are no biases ($\alpha_r = 0$ and $\beta_{r,h,z} = 0$). Moreover, we consider a piecewise linear approximation—sometimes called “hard” tanh—to the tanh function given by

$$\phi_{\text{lin}}(x) = \begin{cases} 1 & x > 1/g_h, \\ g_h x & |x| \leq 1/g_h, \\ -1 & x < -1/g_h. \end{cases} \quad (\text{B1})$$

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

$$P(\sigma_z = x) = f_z \delta(x - 1) + (1 - f_z) \delta(x), \quad (\text{B2})$$

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 [Eqs. (A38)–(A42)] to yield

$$G(\lambda, \bar{\lambda}) = \frac{1 - f_z}{\lambda} + f_z(1 - f_h) \frac{1}{\lambda + 1} + \frac{1}{\lambda + \tau_z^{-1}} + \frac{(1 + \bar{\lambda})}{g_h^2 \sigma(\beta_r)^2} \mathbb{I}_{\{|\lambda| < g_h^2 \sigma(\beta_r)^2\}}, \quad (\text{B3})$$

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:

$$\mu(\lambda) = (1 - f_z) \delta(\lambda) + f_z(1 - f_h) \delta(\lambda + 1) + \frac{4}{\pi g_h^2} \mathbb{I}_{\{|\lambda| \leq g_h^2/4\}}. \quad (\text{B4})$$

Thus, we see an extensive number of eigenvalues at zero.

Now, let us study the regime where α_z is large but not infinite. We would like to get the scaling behavior 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

$$\int \mathcal{D}z \frac{\sigma_z(\sqrt{\Delta_z} \cdot z)^2}{|\lambda_0 e^{-c\alpha_z \sqrt{\Delta_h}} + \sigma_z(\sqrt{\Delta_z}^2)|^2} = \frac{\sigma_r(\beta_r)^{-2}}{\langle \phi'(\sqrt{\Delta_h} \cdot h)^2 \rangle}. \quad (\text{B5})$$

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

$$\text{erfc}\left(\frac{c}{\sqrt{2}}\right) \approx \frac{4}{\langle \phi'(\sqrt{\Delta_h} \cdot h)^2 \rangle}. \quad (\text{B6})$$

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 have marginal stability.

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

$$P[|\lambda(\mathcal{D})| < \delta] = \frac{1}{2\pi i} \oint_{\mathcal{C}} dz G(z), \quad (\text{B7})$$

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:

$$P[|\lambda(\mathcal{D})| < \delta] \approx \frac{1}{2} \text{erfc}\left(-\frac{\log(\delta)}{\alpha_z \sqrt{2\Delta_h}}\right). \quad (\text{B8})$$

APPENDIX C: DETAILS OF THE DYNAMICAL MEAN-FIELD THEORY

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 [73,74] and has been later applied to the study of random neural networks [9,15,21,75]. In our case, the DMFT reduces the description of the full $3N$ -dimensional (deterministic) ordinary differential equations (ODEs) describing $(\mathbf{h}, \mathbf{z}, \mathbf{r})$ to a set of three coupled stochastic differential equations for scalar variables (h, z, r) .

Here, we provide a detailed, self-contained description of the dynamical mean-field theory for the gated RNN using the Martin–Siggia–Rose–De Dominicis–Janssen formalism. 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

$$Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = \mathbb{E} \left[\exp \left(i \sum_{j=1}^N \int \hat{\mathbf{b}}_j(t)^T \mathbf{x}_j(t) dt \right) \right], \quad (\text{C1})$$

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 also include 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. The generating functional is then 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

$$\langle h_i(t) h_i(t') \rangle = \frac{\delta^2}{\delta \hat{b}_i^h(t') \delta \hat{b}_i^h(t)} Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] \Big|_{\mathbf{b}=0}. \quad (\text{C2})$$

Up until this point, things are quite general and do 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. Thus, we can study the disorder averaged (over \mathcal{J}), the generating functional $\bar{Z} = \langle Z_{\mathcal{J}} \rangle_{\mathcal{J}}$, and approximate \bar{Z} with its value evaluated

at the saddle point of the action. This approximation gives us the single-site DMFT picture of dynamics described in Eqs. (C19) and (C20).

To see how this all works, we start with the equations of motion (in vector form)

$$\tau_z \dot{z} = -z + J^z \phi_z(h), \quad (\text{C3})$$

$$\tau_r \dot{r} = -r + J^r \phi_r(h), \quad (\text{C4})$$

$$\dot{h} = \sigma_z(z) \odot (-h + \{J^h[\sigma_r(r) \odot \phi_h(h)]\}), \quad (\text{C5})$$

where \odot stands for elementwise multiplication.

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

$$h_i(t+1) - h_i(t) = \sigma_{z,i}(t) \{-h_i(t) + \sum_j J_{ij}^h \sigma_{r,j}(t) \phi_j(t) + b_i^h(t)\} \delta t,$$

$$\tau_z [z_i(t+1) - z_i(t)] = \left\{ -z_i(t) + \sum_j J_{ij}^z \phi_j(t) + b_i^z(t) \right\} \delta t,$$

$$\tau_r [r_i(t+1) - r_i(t)] = \left\{ -r_i(t) + \sum_j J_{ij}^r \phi_j(t) + b_i^r(t) \right\} \delta t,$$

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

$$Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = \mathbb{E} \left[\exp \left(i \sum_{j=1}^N \sum_t \hat{\mathbf{b}}_j(t)^T \mathbf{x}_j(t) \delta t \right) \right], \quad (\text{C6})$$

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:

$$\begin{aligned} Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = & \int \prod_{i,t} \prod_{k,t'} \prod_{m,t''} dh_i(t) dz_k(t') dr_m(t'') \cdot \exp \left(i \left\{ \sum_{i,t} \hat{b}_i^h(t) h_i(t) + \hat{b}_i^z(t) z_i(t) + \hat{b}_i^r(t) r_i(t) \right\} \delta t \right) \\ & \times \delta \left(h_i(t+1) - h_i(t) + \left\{ h_i(t) \sigma_{z,i}(t) - \sigma_{z,i}(t) \left[\sum_j J_{ij}^h \sigma_{r,j}(t) \phi_j(t) \right] - b_i^h(t) \right\} \delta t \right) \\ & \times \delta \left(z_k(t'+1) - z_k(t') + \frac{1}{\tau_z} \left\{ z_k(t') + \sum_l J_{kl}^z \phi_l(t') + b_k^z(t') \right\} \delta t \right) \\ & \times \delta \left(r_m(t''+1) - r_m(t'') + \frac{1}{\tau_r} \left\{ r_m(t'') + \sum_n J_{mn}^r \phi_n(t'') + b_m^r(t'') \right\} \delta t \right). \end{aligned} \quad (\text{C7})$$

Now, let us introduce the Fourier representation for the δ function; this introduces an auxiliary field variable, which as we see allows us to calculate the response function in the MSRDJ formalism. The generating functional can then be expressed as

$$\begin{aligned} Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = & \int \prod_{i,t} \prod_{k,t'} \prod_{m,t''} dh_i(t) \frac{d\hat{h}_i(t)}{2\pi} dz_k(t') \frac{d\hat{z}_k(t')}{2\pi} dr_m(t'') \frac{d\hat{r}_m(t'')}{2\pi} \\ & \times \exp \left[-i \sum_{i,t} \hat{h}_i(t) [h_i(t+1) - h_i(t) - f_h(h_i, z_i, r_i) \delta t - b_i^h(t) \delta t] + i \sum_{i,t} \hat{b}_i^h(t) h_i(t) \delta t \right] \\ & \times \exp \left[-i \sum_{k,t'} \hat{z}_k(t') \left(z_k(t'+1) - z_k(t') - f_z(h_k, z_k) \frac{\delta t}{\tau_z} - b_k^z(t') \frac{\delta t}{\tau_z} \right) + i \sum_{k,t'} \hat{b}_k^z(t') z_k(t') \delta t \right] \\ & \times \exp \left[-i \sum_{m,t''} \hat{r}_m(t'') \left(r_m(t''+1) - r_m(t'') - f_r(h_m, r_m) \frac{\delta t}{\tau_r} - b_m^r(t'') \frac{\delta t}{\tau_r} \right) + i \sum_{m,t''} \hat{b}_m^r(t'') r_m(t'') \delta t \right], \quad (\text{C8}) \end{aligned}$$

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

$$f_h(h_i, z_i, r_i) = \sigma_{z,i}(t) \left(-h_i(t) + \sum_j J_{ij}^h \sigma_{r,j}(t) \phi_j(t) \right),$$

$$f_z(h_k, z_k) = -z_k(t') + \sum_l J_{kl}^z \phi_l(t'),$$

$$f_r(h_m, r_m) = -r_m(t'') + \sum_n J_{mn}^r \phi_n(t'').$$

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:

$$\begin{aligned} Z_{\mathcal{J}}[\hat{\mathbf{b}}, \mathbf{b}] = & \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \mathcal{D}z_i \mathcal{D}\hat{z}_i \mathcal{D}r_i \mathcal{D}\hat{r}_i \exp \left\{ -S[\hat{\mathbf{x}}, \mathbf{x}] \right. \\ & \left. + i \int dt [\hat{\mathbf{b}}(t)^T \mathbf{x}(t) + \mathbf{b}(t)^T \hat{\mathbf{x}}(t)] \right\}, \quad (\text{C9}) \end{aligned}$$

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

$$\begin{aligned} S[\hat{\mathbf{x}}, \mathbf{x}] = & i \sum_i \int dt \hat{h}_i(t) [\partial_t h_i(t) - f_h(h_i, z_i, r_i)] \\ & + i \sum_k \int dt \hat{z}_k(t) \left[\partial_t z_k(t) - \frac{f_z(h_k, z_k)}{\tau_z} \right] \\ & + i \sum_m \int dt \hat{r}_m(t) \left[\partial_t r_m(t) - \frac{f_r(h_m, r_m)}{\tau_r} \right]. \quad (\text{C10}) \end{aligned}$$

The functional is properly normalized, 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.

1. Disorder averaging

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

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

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

$$\begin{aligned} \bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] = & \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \mathcal{D}z_i \mathcal{D}\hat{z}_i \mathcal{D}r_i \mathcal{D}\hat{r}_i \exp \left\{ -\bar{S}[\hat{\mathbf{x}}, \mathbf{x}] \right. \\ & \left. + i \int dt [\hat{\mathbf{b}}(t)^T \mathbf{x}(t) + \mathbf{b}(t)^T \hat{\mathbf{x}}(t)] \right\}, \quad (\text{C11}) \end{aligned}$$

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

$$\begin{aligned}\bar{S}[\hat{\mathbf{x}}, \mathbf{x}] = & i \sum_i \int dt \hat{h}_i(t) [\partial_t h_i(t) + h_i(t) \sigma_{z,i}(t)] + \frac{1}{2N} \sum_{i,j} \left(\int dt \hat{h}_i(t) \sigma_{z,i}(t) \phi_j(t) \sigma_{r,j}(t) \right)^2 \\ & + i \sum_k \int dt \hat{z}_k(t) \left(\partial_t z_k(t) + \frac{z_k(t)}{\tau_z} \right) + \frac{1}{2N} \sum_{k,l} \left(\int dt \frac{\hat{z}_k(t)}{\tau_z} \cdot \phi_l(t) \right)^2 \\ & + i \sum_m \int dt \hat{r}_m(t) \left(\partial_t r_m(t) + \frac{r_m(t)}{\tau_r} \right) + \frac{1}{2N} \sum_{m,n} \left(\int dt \frac{\hat{r}_m(t)}{\tau_r} \cdot \phi_n(t) \right)^2.\end{aligned}\quad (\text{C12})$$

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 nonlocal terms (e.g., involving both i and j), and we can introduce the following auxiliary fields to decouple these nonlocal terms:

$$\begin{aligned}C_{\phi\sigma_r}(t, t') &:= \frac{1}{N} \sum_i \phi_i(t) \phi_i(t') \sigma_{r,i}(t) \sigma_{r,i}(t'), \\ C_\phi(t, t') &:= \frac{1}{N} \sum_k \phi_k(t) \phi_k(t').\end{aligned}\quad (\text{C13})$$

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:

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

This allows us to make the following transformations to decouple the nonlocal terms in the action \bar{S} :

$$\begin{aligned}\frac{1}{2N} \sum_{i,j} [\hat{h}_i(t) \sigma_{z,i}(t) \phi_j(t) \sigma_{r,j}(t)]^2 &\rightarrow \frac{1}{2} \sum_i \int dt dt' \hat{h}_i(t) \sigma_{z,i}(t) C_{\phi\sigma_r}(t, t') \hat{h}_i(t') \sigma_{z,i}(t'), \\ \frac{1}{2N} \sum_{k,l} \left(\int dt \frac{\hat{z}_k(t)}{\tau_z} \cdot \phi_l(t) \right)^2 &\rightarrow \frac{1}{2} \sum_k \int dt dt' \frac{\hat{z}_k(t)}{\tau_z} C_\phi(t, t') \frac{\hat{z}_k(t')}{\tau_z}, \\ \frac{1}{2N} \sum_{m,n} \left(\int dt \frac{\hat{r}_m(t)}{\tau_r} \cdot \phi_n(t) \right)^2 &\rightarrow \frac{1}{2} \sum_m \int dt dt' \frac{\hat{r}_m(t)}{\tau_r} C_\phi(t, t') \frac{\hat{r}_m(t')}{\tau_r}.\end{aligned}$$

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 decouple the sites by summarizing 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

$$\begin{aligned}\bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] &= \int \mathcal{D}\hat{\mathbf{C}} \mathcal{D}\mathbf{C} \exp(-N \cdot \mathcal{L}[\hat{\mathbf{C}}, \mathbf{C}; \hat{\mathbf{b}}, \mathbf{b}]), \\ \mathcal{L} &= \frac{i}{2} \int dt dt' [\mathbf{C}(t, t')^T \hat{\mathbf{C}}(t, t')] - W[\hat{\mathbf{C}}, \mathbf{C}; \hat{\mathbf{b}}, \mathbf{b}], \\ \exp(N \cdot W) &= \int \prod_i \mathcal{D}h_i \mathcal{D}\hat{h}_i \mathcal{D}z_i \mathcal{D}\hat{z}_i \mathcal{D}r_i \mathcal{D}\hat{r}_i \times \exp \left\{ i \int dt [\mathbf{b}(t)^T \hat{\mathbf{h}}(t) + \hat{\mathbf{b}}(t)^T \mathbf{h}(t)] - S_d[\hat{\mathbf{h}}, \mathbf{h}; \{\mathbf{C}, \hat{\mathbf{C}}\}] \right\},\end{aligned}\quad (\text{C14})$$

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

$$S_d[\hat{\mathbf{h}}, \mathbf{h}; \{\mathbf{C}, \hat{\mathbf{C}}\}] = i \int dt [\hat{\mathbf{h}}(t)^T \partial_t \mathbf{h}(t) + \hat{\mathbf{h}}_\tau(t)^T \mathbf{h}(t)] \\ + \frac{1}{2} \int dt dt' \hat{\mathbf{h}}_\tau(t)^T \mathbb{D}\mathbf{C}(t, t') \hat{\mathbf{h}}_\tau(t') \\ - \frac{i}{2} \int dt dt' \mathbf{S}_x(t)^T \mathbb{D}\hat{\mathbf{C}}(t, t') \mathbf{S}_x(t'), \quad (\text{C15})$$

where

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

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

So far, we do not make 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:

$$\bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] \simeq e^{N \cdot \mathcal{L}_0[\hat{\mathbf{b}}, \mathbf{b}; \mathbf{C}^0, \hat{\mathbf{C}}^0]} \int \mathcal{D}\hat{\mathbf{Q}} \mathcal{D}\mathbf{Q} e^{-N \cdot \mathcal{L}_2[\hat{\mathbf{Q}}, \mathbf{Q}, \hat{\mathbf{b}}, \mathbf{b}]}.$$

We approximate the action \mathcal{L} in Eq. (C14) by its saddle-point value plus a Hessian term: $\mathcal{L} \simeq \mathcal{L}_0 + \mathcal{L}_2$ and the \mathbf{Q} and $\hat{\mathbf{Q}}$ fields represent Gaussian fluctuations about the saddle-point values \mathbf{C}^0 and $\hat{\mathbf{C}}^0$, respectively. At the saddle-point the action is stationary with respect to variations; thus, the saddle-point values of C fields satisfy

$$C_{\phi\sigma_r}^0(t, t') = \frac{1}{N} \sum_{i=1}^N \langle \phi_i(t) \sigma_{r,i}(t) \phi_i(t') \sigma_{r,i}(t') \rangle_0, \\ \hat{C}_{\phi\sigma_r}^0(t, t') = \frac{1}{N} \sum_{i=1}^N \langle \hat{h}_i(t) \sigma_{z,i}(t) \hat{h}_i(t') \sigma_{z,i}(t') \rangle_0 \\ = \frac{\delta^2 \langle \sigma_{z,i}(t) \sigma_{z,i}(t') \rangle_0}{\delta b_i(t) \delta b_i(t')} = 0, \\ C_\phi^0(t, t') = \frac{1}{N} \sum_{k=1}^N \langle \phi_k(t) \phi_k(t') \rangle_0, \\ \hat{C}_\phi^0(t, t') = 0. \quad (\text{C16})$$

In evaluating the saddle-point correlation function in the second line, we use 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, $\langle \cdots \rangle_0$ denotes averages with respect to 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 sitewise generating functionals:

$$\bar{Z}[\hat{\mathbf{b}}, \mathbf{b}] = Z_0[\hat{\mathbf{b}}, \mathbf{b}]^N, \quad (\text{C17})$$

where the sitewise functionals are given by

$$Z_0[\hat{\mathbf{b}}, \mathbf{b}] = \int \mathcal{D}h \mathcal{D}\hat{h} \mathcal{D}z \mathcal{D}\hat{z} \mathcal{D}r \mathcal{D}\hat{r} \\ \times e^{i \int dt [\mathbf{b}(t)^T \hat{\mathbf{h}}(t) + \hat{\mathbf{b}}(t)^T \mathbf{h}(t)] - S_d[\hat{\mathbf{h}}, \mathbf{h}; \{\mathbf{C}^0, \hat{\mathbf{C}}^0\}]}, \quad (\text{C18})$$

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

The sitewise 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-Stratonovich transform which makes the action linear at the cost of introducing three auxiliary Gaussian fields η_h , η_z , and η_r with correlation functions $C_{\phi\sigma_r}^0(t, t')$, $C_\phi^0(t, t')$, and $C_\phi^0(t, t')$, respectively. With this transformation, the action for each site corresponds to stochastic dynamics for three scalar variables given by

$$\dot{h}(t) = -\sigma_z(z) \cdot h(t) + \sigma_z(z) \cdot \eta_h(t), \quad (\text{C19})$$

$$\tau_z \dot{z}(t) = -z(t) + \eta_z(t), \quad (\text{C20})$$

$$\tau_r \dot{r}(t) = -r(t) + \eta_r(t), \quad (\text{C21})$$

where the Gaussian noise processes η_h , η_z , and η_r have correlation functions that must be determined self-consistently:

$$\langle \eta_h(t) \cdot \eta_h(t') \rangle = \langle \phi(t) \sigma_r(t) \cdot \phi(t') \sigma_r(t') \rangle, \\ \langle \eta_z(t) \cdot \eta_z(t') \rangle = \langle \phi(t) \cdot \phi(t') \rangle, \\ \langle \eta_r(t) \cdot \eta_r(t') \rangle = \langle \phi(t) \cdot \phi(t') \rangle.$$

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 summarize 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.

We comment here on the unique difficulty that gating presents to an analysis of the DMFT. While $r(t)$ and $z(t)$ are both described by Gaussian processes in the DMFT, the multiplicative $\sigma_z(z)$ interaction in Eq. (C19) spoils the Gaussianity of $h(t)$. Note that $r(t)$ is always Gaussian and uncorrelated to $h(t)$. In order to try solving for the correlation functions, we need to make a factorization assumption, justified numerically in Fig. 10. The story simplifies at a fixed point, where $h = \eta^h$ (since $\sigma_z > 0$), and is, thus, Gaussian and independent of r .

In order to solve the DMFT equations, we use a numerical method described in Ref. [76]. 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. 9, we see that the solution to the mean-field equations agrees well with the true population-averaged correlation function; Fig. 9 also shows the scale of fluctuations around the mean-field solutions (Fig. 9, thin black lines).

The correlation functions in the DMFT picture such as $C_h(t, t') = \langle h(t)h(t') \rangle$ are the order parameters and correspond to the population-averaged correlation functions in the full network. These turn out to be useful in our analysis of the RNN dynamics in some analyses. Qualitative changes in the correlation functions correspond to transitions between dynamical regimes of the RNN.

In general, the non-Gaussian nature of h makes it impossible to get equations governing the correlation functions. However, when α_z is not too large, Eqs. (C19) and (C20) can be extended to get equations of motions for the correlation functions C_h , C_z , and C_r , which proves useful later on. This requires a separation assumption between the h and σ_z correlators, which seems reasonable for moderate α_z (see Fig. 10). “Squaring” Eqs. (C19) and (C20), we get

$$[-\partial_\tau^2 + C_{\sigma_z}(\tau)]C_h(\tau) = C_{\sigma_z}(\tau)C_{\sigma_r}(\tau)C_\phi(\tau), \quad (\text{C22})$$

$$[-\tau_z^2 \partial_\tau^2 + 1]C_z(\tau) = C_\phi(\tau), \quad (\text{C23})$$

$$[-\tau_r^2 \partial_\tau^2 + 1]C_r(\tau) = C_\phi(\tau), \quad (\text{C24})$$

where we use the shorthand $\sigma_z(t) \equiv \sigma_z[z(t)]$, $\phi(t) \equiv \phi[h(t)]$, and denote the two-time correlation functions as

$$C_x(t, t') = \langle x(t)x(t') \rangle, \quad (\text{C25})$$

where $x \in \{h, z, r, \sigma_z, \sigma_r, \phi\}$ and the expectation here is over the random Gaussian fields in Eqs. (C19)–(C21). We

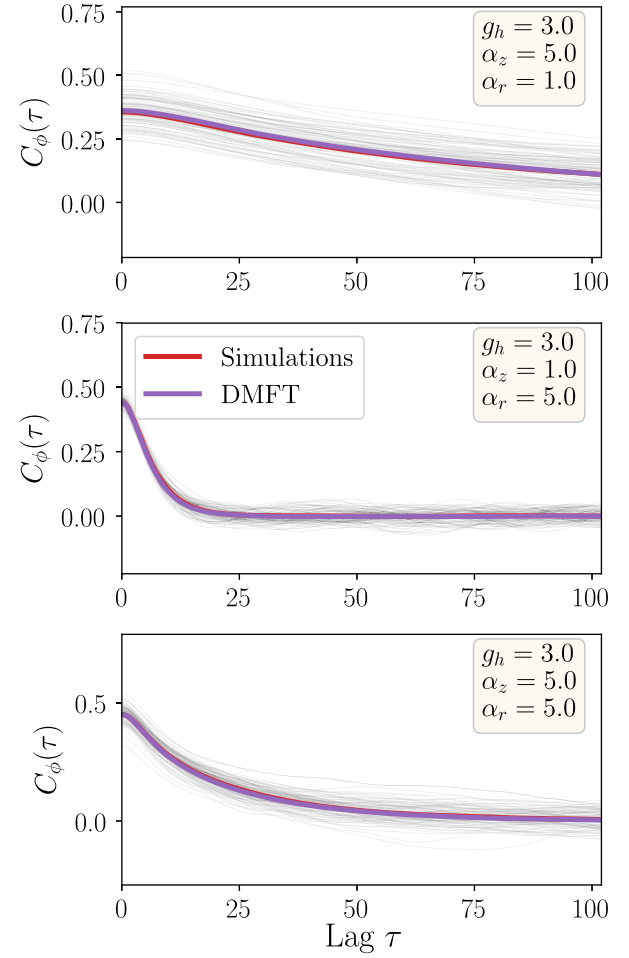


FIG. 9. Validating the DMFT. We show the comparison between the population-averaged correlation functions $C_\phi(\tau) \equiv \langle \phi(t)\phi(t+\tau) \rangle$ obtained from the full network simulations of a single instantiation in steady state (purple line) and from solving the DMFT equations (red line) for three distinct parameter values. The lag τ is relative to τ_h (taken to be unity). Thin black lines are the time-averaged correlation functions for individual neurons sampled from the network, to show the scale of fluctuations around the population-averaged correlation functions. $N = 5000$ for all the panels.




FIG. 10. The validity of the approximation $C_{h\sigma_z} = C_h C_{\sigma_z}$ for two values of α_z . The correlation functions are calculated numerically in a network with $n = 1000$, $g_h = 3.5$, and $\alpha_r = 0$.

assume that the network reaches 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. (C22).

For time-independent solutions, i.e., fixed points, Eqs. (C22)–(C24) simplify to read

$$\Delta_z \equiv \langle z^2 \rangle = \int Dx \phi(\sqrt{\Delta_h}x)^2 = \Delta_r, \quad (\text{C26})$$

$$\Delta_h \equiv \langle h^2 \rangle = \int Dx Dy \phi(\sqrt{\Delta_h}x)^2 \sigma_r(\sqrt{\Delta_r}y)^2, \quad (\text{C27})$$

where we use Δ instead of C to indicate fixed-point variances and Dx is the standard Gaussian measure. It is interesting to note that these mean-field equations can be mapped to those obtained in Ref. [51] for the discrete-time GRU.

We also make use of the MFT with static random inputs. For completeness, we include the resulting equations here. With $I_i^{h,z,r} \sim \mathcal{N}(0, \sigma_{h,z,r}^2)$, the MFT time-independent solution satisfies

$$\Delta_z = \int Dx \phi(\sqrt{\Delta_h}x)^2 + \sigma_z^2, \quad (\text{C28})$$

$$\Delta_r = \int Dx \phi(\sqrt{\Delta_h}x)^2 + \sigma_r^2, \quad (\text{C29})$$

$$\Delta_h = \int Dx Dy \phi(\sqrt{\Delta_h}x)^2 \sigma_r(\sqrt{\Delta_r}y)^2 + \sigma_h^2. \quad (\text{C30})$$

APPENDIX D: DETAILS OF THE NUMERICS FOR THE LYAPUNOV SPECTRUM

The evolution of perturbations $\delta\mathbf{x}(t)$ along a trajectory follow the tangent-space dynamics governed by the Jacobian

$$\partial_t \delta\mathbf{x}(t) = \mathcal{D}(t) \delta\mathbf{x}(t). \quad (\text{D1})$$

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

$$\delta\mathbf{x}(t) = U(t, 0) \delta\mathbf{x}(0), \quad U(t, 0) = \mathcal{T}[e^{\int_0^t ds \mathcal{D}(s)}], \quad (\text{D2})$$

where $\mathcal{T}[\cdot]$ is the time-ordering operator applied to the contents of the bracket. When the infinitesimal perturbations grow (shrink) exponentially, the rate of this exponential growth (decay) is dictated by the maximal Lyapunov exponent defined as [54]

$$\lambda_{\max} := \lim_{T \rightarrow \infty} \frac{1}{T} \lim_{\|\delta\mathbf{x}(0)\| \rightarrow 0} \ln \frac{\|\delta\mathbf{x}(T)\|}{\|\delta\mathbf{x}(0)\|}. \quad (\text{D3})$$

For ergodic systems, this limit is independent of almost all initial conditions, as guaranteed by the Oseledets multiplicative ergodic theorem [54]. 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, provides a fuller characterization of asymptotic behavior. 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 Oseledets matrix, defined as [54]

$$\mathbf{M}(t) = \lim_{t \rightarrow \infty} [U(t, 0)^T U(t, 0)]^{1/2t}. \quad (\text{D4})$$

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 Ref. [77] (also cf. Ref. [78] 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. (D1)] for a short time interval t_0 . Therefore, the new set of vectors is given by

$$\hat{Q} = U(t_0, 0) Q^0. \quad (\text{D5})$$

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

$$\lambda_i = \lim_{N_I \rightarrow \infty} \frac{1}{N_I t_0} \sum_{j=1}^{N_I} \ln R_{ii}^j, \quad i \in \{1, \dots, k\}. \quad (\text{D6})$$

APPENDIX E: DETAILS OF THE DMFT PREDICTION FOR λ_{\max}

The starting point of the method to calculate the DMFT prediction for λ_{\max} 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 or divergence of the trajectories reveals the maximal Lyapunov exponent. To this end, let us define $d(t, s) := N^{-1} \sum_i [x_i^1(t) - x_i^2(s)]^2$ and study the growth rate of $d(t, t)$. In the large N limit, we expect population averages like $C^{12}(t, s) := N^{-1} \sum_i x_i^1(t) x_i^2(s)$ to be self-averaging (like in the DMFT for a single system) [79], and, thus, we can write

$$d(t, s) = C^{11}(t, t) + C^{22}(s, s) - C^{12}(t, s) - C^{21}(t, s). \quad (\text{E1})$$

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 Ref. [21], this can be achieved by considering a joint generating functional for the replicated system:

$$\tilde{Z}_J[\hat{\mathbf{b}}^1, \hat{\mathbf{b}}^2, \mathbf{b}^1, \mathbf{b}^2] = \mathbb{E} \left[\exp \left(i \sum_{\mu=1}^2 \sum_{j=1}^N \int \hat{\mathbf{b}}_j^\mu(t)^T \mathbf{x}_j^\mu(t) dt \right) \right]. \quad (\text{E2})$$

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 (cf. Appendix C) yields a system of coupled stochastic differential equations (SDEs) (one for each replica), similar to Eq. (C20), 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

$$[-\partial_\tau^2 + C_{\sigma_z}^{\mu\nu}(\tau)] C_h^{\mu\nu}(\tau) = C_{\sigma_z}^{\mu\nu}(\tau) C_\phi^{\mu\nu}(\tau) C_{\sigma_r}^{\mu\nu}(\tau), \quad (\text{E3})$$

$$[-\tau_z^2 \partial_\tau^2 + 1] C_z^{\mu\nu}(\tau) = C_\phi^{\mu\nu}(\tau), \quad (\text{E4})$$

$$[-\tau_r^2 \partial_\tau^2 + 1] C_r(\tau) = C_\phi(\tau), \quad (\text{E5})$$

where $\mu, \nu \in \{1, 2\}$ are the replica indices. Note that the single-replica solution clearly is 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. (E1), we see that $d(t, t) \sim \epsilon \tilde{\chi}(t, t)$, and, thus, the growth rate of $\tilde{\chi}$ yields the required Lyapunov exponent. To this end, we make an ansatz $\tilde{\chi}_{h,z,r} = e^{\kappa T} \chi(\tau)$, where $2T = t + s$, $2\tau = t - s$, and κ is the DMFT prediction of the maximum Lyapunov exponent that needs to be solved for. Substituting this back into Eq. (E3), we get a generalized eigenvalue problem for κ as stated in the text [Eqs. (10) and (11)].

APPENDIX F: CALCULATION OF MAXIMAL LYAPUNOV EXPONENT FROM RMT

The DMFT prediction for how gates shape λ_{\max} (via the correlation functions) is somewhat involved; thus, we provide an alternate expression for the maximal Lyapunov exponent λ_{\max} , derived using RMT which relates it to the

relaxation time of the dynamics. The starting point to get λ_{\max} is the Oseledets multiplicative ergodic theorem, which guarantees that [80]

$$\lambda_{\max} = \lim_{t \rightarrow \infty} \frac{1}{2t} \log \frac{\|\chi(t)\|^2}{N} \quad (\text{F1})$$

$$= \lim_{t \rightarrow \infty} \frac{1}{2t} \log \frac{1}{N} \text{Tr}[\chi(t) \chi(t)^T], \quad (\text{F2})$$

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

$$\mathcal{D} = -\mathbb{1} + J[\phi'(t)]. \quad (\text{F3})$$

We expect the maximal Lyapunov exponent to be independent of the random network realization and, thus, equal to its value after disorder averaging. Furthermore, to make any progress, we use a short-time approximation for $\chi(t) \approx e^{\int_0^t dt' \mathcal{D}(t')}$. Defining the diagonal matrix $R(t) = \int^t [\phi'(t')] dt'$, these assumptions give

$$\frac{1}{N} \text{Tr}[\chi(t) \chi(t)^T] \approx e^{-2t} \left\langle \frac{1}{N} \text{Tr} e^{JR(t)} e^{R(t)J^T} \right\rangle \quad (\text{F4})$$

$$= e^{-2t} \sum_{n=0}^{\infty} \frac{1}{(n!)^2} \left(\frac{1}{N} \text{Tr} R(t)^2 \right)^n, \quad (\text{F5})$$

where the second line in Eq. (F5) follows after disorder averaging over J and keeping only terms to leading order in N . Next, we may apply the DMFT to write

$$\frac{1}{N} \text{Tr} R(t)^2 = \int^t dt' dt'' \frac{1}{N} \sum_{i=1}^N \phi_i'(t') \phi_i'(t'') \quad (\text{F6})$$

$$\approx \int dt' dt'' C_{\phi'}(t', t''). \quad (\text{F7})$$

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

$$\int dt' dt'' C_{\phi'}(t', t'') \approx \int_0^{2t} \frac{du}{2} \int_0^t d\tau C_{\phi'}(\tau) \equiv t^2 \tau_R, \quad (\text{F8})$$

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

$$\tau_R \equiv \frac{1}{t} \int_0^t d\tau C_{\phi'}(\tau). \quad (\text{F9})$$

Substituting Eq. (F8) in Eq. (F4), we get

$$\frac{1}{N} \text{Tr}[\chi(t) \chi(t)^T] = e^{-2t} I_0(2t\sqrt{\tau_R}), \quad (\text{F10})$$

which for long times behaves like $\exp[2(\sqrt{\tau_R} - 1)t]$. By inserting this into Eq. (F1), we obtain a bound for the maximal Lyapunov exponent for the vanilla RNN:

$$\lambda_{\max} \geq \sqrt{\tau_R} - 1, \quad (\text{F11})$$

$$\text{where } \tau_R \equiv \frac{1}{t} \int_0^t d\tau C_{\phi'}(\tau). \quad (\text{F12})$$

This formula relates the asymptotic Lyapunov exponent to relaxation time of a local correlation function in steady state. It is interesting to note that the bound also follows by applying the variational theorem to the potential energy obtained from the Schrödinger equation that arises in computing the Lyapunov exponent using DMFT (e.g., see Refs. [15,32]). Specifically, if one uses the potential obtained in these works $V(\tau) = 1 - C_{\phi'}(\tau)$, and assumes a uniform “ground state wave function,” the variational theorem implies that the true ground state energy E_0 is upper bounded $E_0 \leq \lim_{T \rightarrow \infty} (1/T) \int_{-T/2}^{T/2} V(\tau) d\tau \equiv 1 - \tau_R$, which consequently implies the bound (F11).

Now 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

$$\mathcal{D}_t = -\mathbb{1}_{2N} + \begin{pmatrix} J^r & 0 \\ 0 & J^h \end{pmatrix} \begin{pmatrix} 0 & P_t \\ Q_t & R_t \end{pmatrix} \equiv -\mathbb{1}_{2N} + \hat{J} S_t, \quad (\text{F13})$$

$$Q_t = [\phi(\mathbf{h}) \odot \sigma'_r(\mathbf{r})], \quad P_t = [\phi'(\mathbf{h})], \quad (\text{F14})$$

$$R_t = [\phi'(\mathbf{h}) \odot \sigma_r(\mathbf{r})], \quad (\text{F15})$$

where $h = h(t)$ and $r = r(t)$ are time dependent.

Let us define the quantity of interest

$$\sigma_\chi^2 = \left\langle \frac{1}{2N} \text{Tr}[\chi(t) \chi^T(t)] \right\rangle \quad (\text{F16})$$

$$= e^{-2t} \left\langle \frac{1}{2N} \text{Tr} e^{\hat{J} \hat{S}_t} e^{\hat{S}_t^T \hat{J}^T} \right\rangle, \quad (\text{F17})$$

where we additionally define $\hat{S}_t = \int^t dt' S_t$ 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

$$c_n = \left\langle \frac{1}{2N} \text{Tr}[(\hat{J} \hat{S}_t)^n (\hat{S}_t^T \hat{J}^T)^n] \right\rangle. \quad (\text{F18})$$

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 come from

Wick contractions that are “noncrossing”—in the language of diagrams, the moment is 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

$$\hat{c}_n = \langle (\hat{J} \hat{S}_t)^n (\hat{S}_t^T \hat{J}^T)^n \rangle \quad (\text{F19})$$

$$= \langle \hat{J} [\hat{S}_t (\hat{J} \hat{S}_t)^{n-1} (\hat{S}_t^T \hat{J}^T)^{n-1} \hat{S}_t^T] \hat{J}^T \rangle \quad (\text{F20})$$

$$= \begin{pmatrix} a_n \mathbb{1} & 0 \\ 0 & b_n \mathbb{1} \end{pmatrix} + O(N^{-1}). \quad (\text{F21})$$

We wish to determine a_n and b_n . Next, define

$$g_P = \frac{1}{N} \text{Tr} \int^t dt' dt'' P_{t'} P_{t''}, \quad (\text{F22})$$

$$g_Q = \frac{1}{N} \text{Tr} \int^t dt' dt'' Q_{t'} Q_{t''}, \quad (\text{F23})$$

$$g_R = \frac{1}{N} \text{Tr} \int^t dt' dt'' R_{t'} R_{t''}. \quad (\text{F24})$$

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

$$\hat{c}_n = \langle \hat{J} \hat{S}_t (\hat{J} \hat{S}_t)^{n-1} (\hat{S}_t^T \hat{J}^T)^{n-1} \hat{S}_t^T \hat{J}^T \rangle \quad (\text{F25})$$

$$= \langle \hat{J} \hat{S}_t \hat{c}_{n-1} \hat{S}_t^T \hat{J}^T \rangle + O(N^{-1}) \quad (\text{F26})$$

$$= \begin{pmatrix} b_{n-1} g_P \mathbb{1} & 0 \\ 0 & (a_{n-1} g_Q + b_{n-1} g_R) \mathbb{1} \end{pmatrix} + O(N^{-1}). \quad (\text{F27})$$

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

$$a_n = g_P b_{n-1}, \quad b_n = g_R b_{n-1} + g_Q a_{n-1}. \quad (\text{F28})$$

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

$$\begin{pmatrix} a_n \\ b_n \end{pmatrix} = \begin{pmatrix} 0 & g_P \\ g_Q & g_R \end{pmatrix}^n \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad (\text{F29})$$

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

$$c_n = \frac{1}{2} \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} 0 & g_P \\ g_Q & g_R \end{pmatrix}^n \begin{pmatrix} 1 \\ 1 \end{pmatrix}. \quad (\text{F30})$$

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

$$v_{\pm} = \frac{1}{2} \left(g_R \pm \sqrt{g_R^2 + 4g_P g_Q} \right), \quad (\text{F31})$$

which are real valued. The eigenvectors are

$$\mathbf{v}_{\pm} = \left(-\frac{v_{\pm}}{g_Q}, 1 \right). \quad (\text{F32})$$

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

$$c_n = \frac{1}{2} \mathbf{I}^T (v_+^n \mathbf{v}_+ \mathbf{v}_+^T + v_-^n \mathbf{v}_- \mathbf{v}_-^T) \mathbf{I} \quad (\text{F33})$$

$$= \frac{1}{2} \left(1 - \frac{v_-}{g_Q} \right)^2 v_+^n + \frac{1}{2} \left(1 - \frac{v_+}{g_Q} \right)^2 v_-^n. \quad (\text{F34})$$

The final expression for the mean-squared singular value is then

$$\sigma_{\chi}^2 = e^{-2t} \sum_{n=0}^{\infty} \frac{c_n}{(n!)^2}. \quad (\text{F35})$$

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

$$\sigma_{\chi}^2 = \frac{1}{2} e^{-2t} \left[\left(1 - \frac{v_-}{g_Q} \right)^2 I_0(2\sqrt{v_+}) + \left(1 - \frac{v_+}{g_Q} \right)^2 I_0(2\sqrt{v_-}) \right]. \quad (\text{F36})$$

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

$$g_R = \int dt dt' R_t R_{t'} \approx t^2 \frac{1}{t} \int d\tau C_R(\tau) = t^2 \tau_R, \quad (\text{F37})$$

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

$$v_{\pm} = t^2 \frac{1}{2} \left(\tau_R \pm \sqrt{\tau_R^2 + 4\tau_P \tau_Q} \right). \quad (\text{F38})$$

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

$$\sigma_{\chi}^2 \sim \exp(-2t + 2\sqrt{v_+}), \quad (\text{F39})$$

which gives the Lyapunov exponent

$$\lambda_L \geq \left(\frac{\tau_R + \sqrt{\tau_R^2 + 4\tau_P \tau_Q}}{2} \right)^{1/2} - 1, \quad (\text{F40})$$

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

$$\tau_R = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau C_{\phi'}(\tau) C_{\sigma_r}(\tau), \quad (\text{F41})$$

$$\tau_A = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau C_{\phi'}(\tau), \quad (\text{F42})$$

$$\tau_Q = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau C_{\phi}(\tau) C_{\sigma_r}(\tau). \quad (\text{F43})$$

APPENDIX G: DETAILS OF THE DISCONTINUOUS CHAOTIC TRANSITION

In this section, we provide the details for the calculations involved in the discontinuous chaotic transition.

1. 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. (C26)] for Δ_h . However, as we increase α_r for a fixed g_h , two additional nonzero solutions to Δ_h spontaneously appear at a critical value $\alpha_{r,\text{FP}}^*(g_h)$ as shown in Fig. 4(a). Numerical solutions to the fixed-point equations reveal the form of the bifurcation curve $\alpha_{r,\text{FP}}^*(g_h)$ and the associated value of $\Delta_h^*(g_h)$. We see that $\alpha_{r,\text{FP}}^*(g_h)$ increases rapidly with decreasing g_h , dividing the parameter space into regions with either one or three 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 I 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. 4(b). The larger nonzero solution for the variance at the bifurcation point scales as

$$\Delta_h^* \approx (\alpha_r^2 - 8) \cdot \xi_0 + \xi_1 \epsilon \text{ for } \alpha_r \rightarrow \alpha_{r,\text{FP}}^*(2) = \sqrt{8}, \quad (\text{G1})$$

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

At the other extreme, to determine the smallest value of g_h for which a bifurcation is possible, we note from Fig. 4(b) 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 I 2)

$$\Delta_h^*(\sqrt{2}^+) \sim \frac{g_h^2 - 2}{2g_h^4} \quad \text{for } g_h \rightarrow \sqrt{2}^+. \quad (\text{G2})$$

Thus, in the region $g_h \in (\sqrt{2}, 2)$, there exist nonzero 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.

2. 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 $\alpha_r > \alpha_{r,\text{FP}}^*(g_h)$. However, it turns out that the spontaneous appearance of these unstable fixed points is not accompanied by any asymptotic dynamical signatures—as measured by the Lyapunov exponents (see Fig. 4) or by the transient times (see Fig. 11). It is only when α_r is increased further beyond a second critical value $\alpha_{r,\text{DMFT}}^*(g_h)$ that we see the appearance of chaotic and long-lived transients. This is significant in regard to a result by Wainrib and Touboul [45], where they show 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) is the same as the maximal Lyapunov exponent (a dynamical property).

Let us characterize the second dynamical transition curve given by $\alpha_{r,\text{DMFT}}^*(g_h)$ [Fig. 4(c), red curve]. For ease of discussion, we turn off the update gate ($\alpha_z = 0$) and introduce a functional F_ψ for a 2D Gaussian average of a given function $\psi(x)$:

$$F_\psi[C_h(0), C_h(\tau)] = \mathbb{E}[\psi(z_1)\psi(z_2)], \quad (\text{G3})$$

$$\text{where } \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_h), \quad \mathbf{C}_h = \begin{pmatrix} C_h(0) & C_h(\tau) \\ C_h(\tau) & C_h(0) \end{pmatrix}. \quad (\text{G4})$$

The DMFT equations for the correlation functions then become

$$\begin{aligned} \frac{1}{4} C_h(\tau) - \partial_\tau^2 C_h(\tau) &= \frac{1}{4} F_\phi[C_h(0), C_h(\tau)] F_{\sigma_r}[C_r(0), C_r(\tau)], \\ C_r(\tau) - \tau_r^2 \partial_\tau^2 C_r(\tau) &= F_\phi[C_\phi(0), C_\phi(\tau)]. \end{aligned} \quad (\text{G5})$$

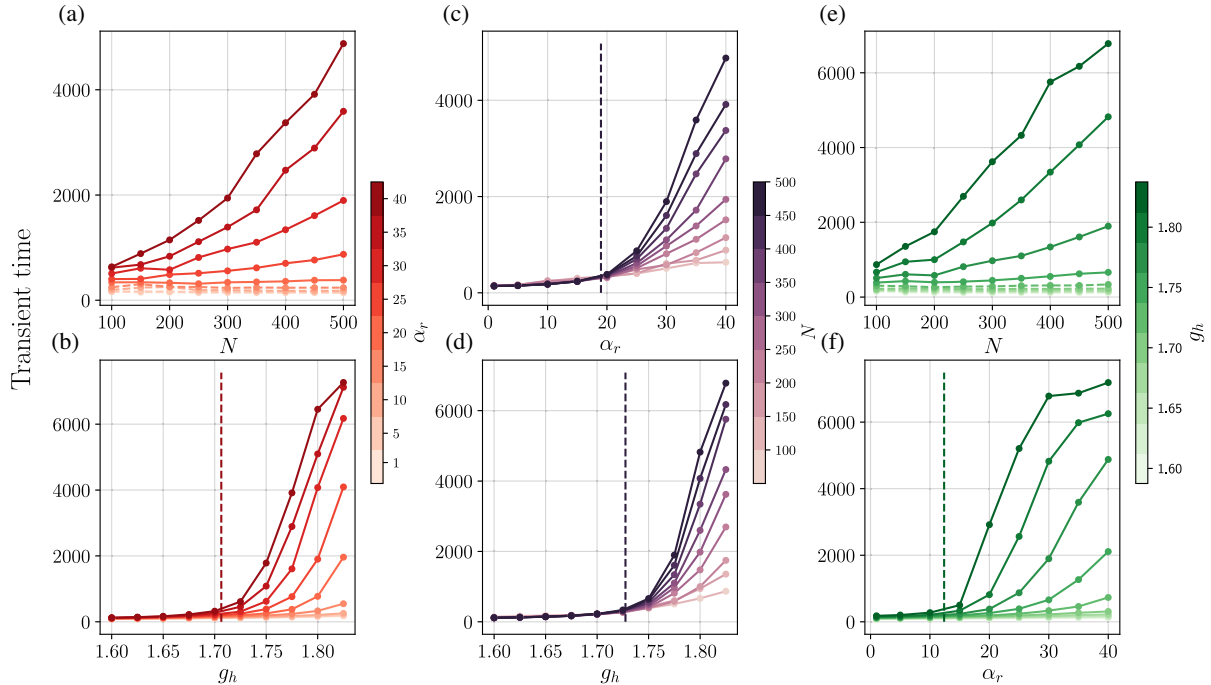


FIG. 11. Transient times at the bifurcation transition. Transient times (τ_T , relative to τ_h) as a function of g_h , α_r , and system size N . Dashed plot lines correspond to situations where $\alpha_r < \alpha_{r,\text{DMFT}}^*(g_h)$. Dashed vertical lines are critical values of α_r or g_h . (a) τ_T vs N for $g_h = 1.775$ and (b) τ_T vs g_h for $N = 500$; the dashed line indicates g_h such that $\alpha_{r,\text{DMFT}}^*(g_h) = 40$. (c) τ_T vs α_r for $g_h = 1.775$; the dashed vertical line is $\alpha_{r,\text{DMFT}}^*(1.775)$. (d) τ_T vs g_h for $\alpha_r = 30$; the dashed line is g_h such that $\alpha_{r,\text{DMFT}}^*(g_h) = 30$. (e) τ_T vs N for $\alpha_r = 30$; dashed plot lines correspond to situations where $30 < \alpha_{r,\text{DMFT}}^*(g_h)$. (f) τ_T vs α_r for $N = 500$; the dashed vertical line is $\alpha_{r,\text{DMFT}}^*(1.85)$. Transient times are averaged over 2000 instances of random networks.

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 . With these approximations, we can integrate the equations for $C_h(\tau)$ to arrive at an equation for the variance $C_h^0 \equiv C_h(0)$. We do this by multiplying by $\partial_\tau C_h(\tau)$ and integrating from τ to ∞ , and we get

$$\frac{1}{2} \dot{C}_h(\tau)^2 = \frac{1}{4} \frac{(C_h)^2}{2} - \frac{1}{4} \int_0^{C_h^0} dC_h F_\phi(C_h, C_h^0) F_{\sigma_r}(C_\phi, C_\phi^0). \quad (\text{G6})$$

Using the boundary condition that $\dot{C}_h(0) = 0$, we get the equation for the variance:

$$\frac{1}{8} C_h(0)^2 - \frac{1}{4} \int_0^{C_h^0} dC_h F_\phi(C_h, C_h^0) F_{\sigma_r}(C_\phi, C_\phi^0) = 0. \quad (\text{G7})$$

Solving this equation gives the DMFT prediction for the variance for any g_h and α_r . Beyond the critical value of α_r , two nonzero solutions for C_h^0 spontaneously emerge. In order to use Eq. (G7) to find a prediction for the DMFT bifurcation curve $\alpha_{r,\text{DMFT}}^*(g_h)$, we need to use the additional fact that at the bifurcation point the two solutions coincide, and there is only one nonzero solution. To proceed, we can view the lhs of Eq. (G7), 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 the following two equations for $C_h^{0,*}$ and α_r^* :

$$\mathcal{F}(g_h, \alpha_r^*, C_h^{0,*}) = 0, \quad (\text{G8})$$

$$\left. \frac{\partial \mathcal{F}(g_h, \alpha_r, C_h^0)}{\partial C_h^0} \right|_{\alpha_r^*, C_h^{0,*}} = 0. \quad (\text{G9})$$

To get the condition for the dynamical bifurcation transition, we need to differentiate the lhs of Eq. (G7) $[\mathcal{F}(g_h, \alpha_r, C_h^0)]$ with respect to C_h^0 and set it to 0. This involves terms like

$$\frac{\partial F_\psi(C_h^0, C_h^0)}{\partial C_h^0}; \quad \frac{\partial F_\psi(C_h^0, 0)}{\partial C_h^0}. \quad (\text{G10})$$

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

$$\begin{aligned} F_\psi(C_h^0, C_h) &= \mathbb{E} \left[\int \frac{dk}{2\pi} \int \frac{dk'}{2\pi} \tilde{\psi}(k) e^{-kz_1} \tilde{\psi}(k') e^{-k'z_2} \right] \\ &= \int \frac{dk}{2\pi} \int \frac{dk'}{2\pi} \tilde{\psi}(k) \tilde{\psi}(k') \exp \left[-\frac{C_h^0}{2} (k^2 + k'^2) - C_h(\tau) k k' \right]. \end{aligned} \quad (\text{G11})$$

This immediately gives us

$$\begin{aligned} \frac{\partial F_\psi(C_h^0, C_h^0)}{\partial C_h^0} &= \int \mathcal{D}x \psi \left(\sqrt{C_h^0} x \right) \psi'' \left(\sqrt{C_h^0} x \right) \\ &\quad + \int \mathcal{D}x \psi' \left(\sqrt{C_h^0} x \right)^2, \\ \frac{\partial F_\psi(C_h^0, 0)}{\partial C_h^0} &= \int \mathcal{D}x \psi \left(\sqrt{C_h^0} x \right) \int \mathcal{D}x \psi'' \left(\sqrt{C_h^0} x \right). \end{aligned} \quad (\text{G12})$$

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 Eqs. (G8) and (G9) to get the bifurcation curve shown in Fig. 4(c). Figure 4(d) shows 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. 4(a)], and (ii) the lower critical value of g_h which permits a dynamical transition [dashed green curve in Figs. 4(a) and 4(b)] is smaller than the corresponding fixed-point critical value of $\sqrt{2}$.

We now calculate the lower critical value of g_h and provide an analytical description of the asymptotic behavior near the lower and higher critical values of g_h . From the red curve in Fig. 4(c), we know that, as g_h tends toward the lower critical value, $\alpha_{r,\text{DMFT}}^* \rightarrow \infty$ 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

$$F_{\sigma_r}(C_\phi^0, C_\phi) \approx \frac{1}{4} + \frac{1}{2\pi} \tan^{-1} \left(\frac{x}{\sqrt{1-x^2}} \right), \quad (\text{G13})$$

$$\text{where } x := \frac{C_h(\tau)}{C_h(0)} \approx \frac{C_\phi(\tau)}{C_\phi(0)}. \quad (\text{G14})$$

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 has $O(C_h(0)^2)$ 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 consistent only to leading order in $C_h(0)$ when g_h is equal to its critical value:

$$g_h^* = \sqrt{\frac{8}{3}}, \quad (\text{G15})$$

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

$$C_h^0 = c_0 + c_1\epsilon + c_2\epsilon^2 + \dots \quad (\text{G16})$$

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

$$F_\phi \approx g_h^2 C_h(\tau) - 2g_h^4 C_h^0 \cdot C_h(\tau) + 5g_h^6 (C_h^0)^2 \cdot C_h(\tau) \quad (\text{G17})$$

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

$$\alpha_{r,\text{DMFT}}^*(2) = \sqrt{12}. \quad (\text{G18})$$

The DMFT prediction for the dynamical bifurcation agrees well with the full network simulations. In Fig. 4(e), 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).

3. Influence of update gate on the discontinuous transition

Here, we comment briefly on the possible influence of the z gate on the discontinuous dynamical phase transition given by the curve $\alpha_{r,\text{DMFT}}^*$. Assuming Eq. (C22) is valid (discussed in more detail toward the end of Appendix C), we may rewrite the DMFT equation for the two-point correlation functions as

$$\frac{1}{2} \dot{C}_h(\tau)^2 = \int_0^{C_h(\tau)} F_{\sigma_z}(C_\phi^0, C_\phi) Q(C_h, C_h^0), \quad (\text{G19})$$

where

$$Q(C_h, C_h^0) = C_h - F_\phi(C_h, C_h^0) F_{\sigma_r}(C_\phi, C_\phi^0). \quad (\text{G20})$$

Noting that a time-dependent solution corresponds to a nonzero solution for $C_h(0)$ and satisfies the boundary condition $\dot{C}_h(0) = 0$ then requires

$$\mathcal{F}_{\alpha_z} \equiv \int_0^{C_h^0} dC_h F_{\sigma_z}(C_\phi, C_\phi^0) Q(C_h, C_h^0) = 0, \quad (\text{G21})$$

where we define a new “potential” function which is related to that defined above by

$$\mathcal{F}_{\alpha_z}|_{\alpha_z=0} = \mathcal{F} = \frac{1}{4} \int_0^{C_h^0} dC_h Q(C_h, C_h^0). \quad (\text{G22})$$

We leave the arguments (g_h, α_r, C_h^0) implicit, for ease of presentation. We proceed to bound the new potential by establishing bounds on F_{σ_z} . To be explicit, we have

$$F_{\sigma_z} = \langle \sigma_z(\tau) \sigma_z(0) \rangle = \langle \sigma_z(\tau) \sigma_z(0) \rangle_c + \langle \sigma_z \rangle^2, \quad (\text{G23})$$

which we express as the sum of a connected component (indicated by a subscript c) and a disconnected component. We can consider two limiting behaviors. When the correlation time tends to zero, the connected component vanishes and (at zero bias $\beta_z = 0$)

$$F_{\sigma_z} \approx \langle \sigma_z \rangle^2 = \frac{1}{4}. \quad (\text{G24})$$

Increasing the correlation time can serve only to increase the two-point function, since $\sigma \geq 0$. In the extreme limit of very long correlation time, we have that

$$F_{\sigma_z} \approx \langle \sigma_z^2(0) \rangle \leq \frac{1}{2}. \quad (\text{G25})$$

The inequality is saturated at $\alpha_z = \infty$, when σ_z becomes a step function of its argument. Therefore, the two-point correlation function of the update gate is bounded above and below:

$$\frac{1}{4} \leq F_{\sigma_z} \leq \frac{1}{2}, \quad (\text{G26})$$

and this bound is uniform in the sense that it holds for all values of the argument $0 \leq C_h \leq C_h^0 < \infty$. Consequently, we are able to bound the potential

$$\frac{1}{4} \mathcal{F} \leq \mathcal{F}_{\alpha_z} \leq \frac{1}{2} \mathcal{F}. \quad (\text{G27})$$

It follows immediately that the derivative is similarly bounded. Consequently, the zeros of \mathcal{F}_{α_z} and $\partial \mathcal{F}_{\alpha_z} / \partial C_h^0$ coincide with the zeros of \mathcal{F} and $\partial \mathcal{F} / \partial C_h^0$, respectively. As a result, the discontinuous transition, determined by Eqs. (G8) and (G9), remain unchanged for values of α_z for which Eq. (C22) is valid. Thus, for moderately large α_z [approximately 10, where Eq. (C22) is valid], the critical line for the discontinuous transition remains unchanged.

APPENDIX H: THE ROLE OF BIASES

We thus far describe 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

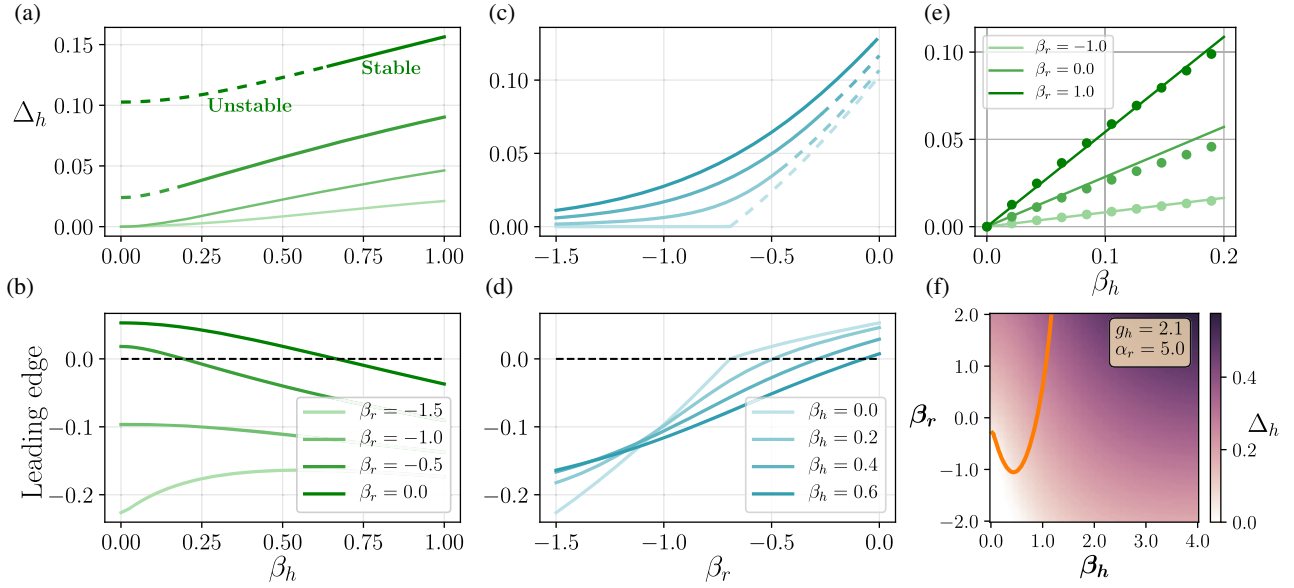


FIG. 12. The role of biases. (a) FP solutions as a function of increasing β_h ; different shades of green correspond to different values of β_r . Dashed lines correspond to FP solutions that are unstable (time-varying states). (b) The leading edge of the spectrum corresponding to the FP solutions calculated 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. (H2) (solid lines). (f) FP solution as a function of β_r and β_h . The orange line indicates the stability line—i.e., regions on top of the orange line correspond to unstable or time-varying states.

always a fixed point of the dynamics, and the zero fixed point is stable provided

$$-1 + \phi'(0)\sigma_r(0) < 0, \quad (\text{H1})$$

where $\phi(x) = \tanh(g_h x + \beta_h)$. This gives the familiar $g_h < 2$ condition when $\beta_r = 0$ [81]. 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 nonzero fixed points, and both β_r and β_h determine the stability of these nonzero fixed points.

To gain some insight into the role of β_h in generating fixed points, we treat the mean-field FP equations [Eq. (C26)] perturbatively around the operating point g_c where the zero fixed point becomes unstable [Eq. (H1)]. 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 I 1)

$$\Delta_h \approx \begin{cases} \frac{\beta_h + \epsilon}{g_c^2(2 - g_c^2 a_1)} & g_c^2 a_1 < 2, \\ (g_c^2 a_1 - 2)f_1 + \epsilon \cdot f_2 & g_c^2 a_1 > 2, \end{cases} \quad (\text{H2})$$

$$\text{where } a_1 = \frac{\alpha_r^2}{16} [\phi_0^{(1)}(\beta_r/2)^2 + \phi_0(\beta_r/2)\phi_0^{(2)}(\beta_r/2)], \quad (\text{H3})$$

where $\phi_0 \equiv \tanh$ and $f_2(\alpha_r, \beta_r)$ and $f_1(\alpha_r, \beta_r)$ are constant functions with respect to ϵ . Therefore, we see that the bias β_h gives rise to nonzero fixed points near the critical point which scale linearly with the bias. In Fig. 12(e), 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. 12(a), 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. 12(b)]. In Fig. 12(c), 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 nonzero FPs. The combined effect of β_h and β_r can be seen in Fig. 12(f), where the nonzero solutions to the left of the orange line indicate unstable (time-varying) solutions. We choose the parameters to illustrate an interesting aspect of the biases: In some cases, increasing β_h can have a nonmonotonic effect on the stability, wherein the solution becomes unstable with increasing β_h and is then eventually stabilized for sufficiently large β_h .

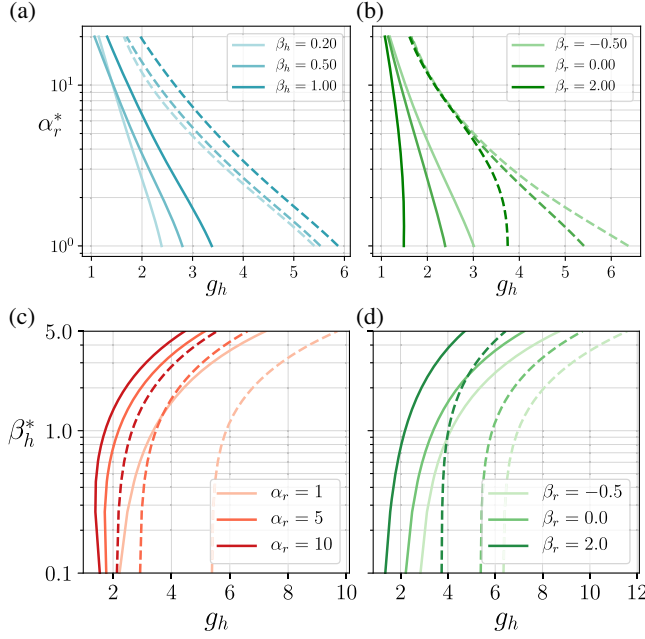


FIG. 13. How the biases alter the transition between stability and chaos. (a) Critical lines indicating boundaries for stability (solid lines) or marginal stability (dashed lines) for different values of β_h . (b) Similar to (a) but for different values of β_r . (c), (d) How the boundaries of stability (solid lines) or marginal stability (dashed lines) change as we vary α_r (c) or β_r (d).

1. Effect of biases on the phase boundaries

In Figs. 13(a) and 13(b), we look at how the critical line for the chaotic transition, in the $\alpha_r - g_h$ plane, changes as we vary β_h (a) or β_r (b). Positive values of β_r (“open” output gate) tend to make the transition line less dependent on α_r [Fig. 13(b)], 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 see above, higher values of β_h have a stabilizing effect, requiring higher g_h and α_r to make the (nonzero) stable fixed point unstable. In both cases, the critical lines for marginal stability [Figs. 13(a) and 13(b), dashed lines] are also influenced in a similar way. In Figs. 13(c) and 13(d), we see how the stability-to-chaos transition is affected by α_r (c) and β_r (d). Consistent with the discussion above, larger α_r and β_r have a destabilizing effect, requiring a larger β_h to make the system stable.

APPENDIX I: DETAILS OF THE 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

$$\Delta_z \equiv \langle z^2 \rangle = \int \mathcal{D}x \phi(\sqrt{\Delta_h} x)^2 = \Delta_r, \quad (\text{I1})$$

$$\Delta_h \equiv \langle h^2 \rangle = \int \mathcal{D}x \mathcal{D}y \phi(\sqrt{\Delta_h} x)^2 \sigma_r(\sqrt{\Delta_r} y)^2. \quad (\text{I2})$$

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 $O(\Delta_h^2)$:

$$\begin{aligned} \langle \sigma_r(\sqrt{\Delta_r} x)^2 \rangle_x &= a_0 + a_1 g_h^2 \Delta_h + a_2 g_h^4 \Delta_h^2, \\ a_0 &= \frac{1}{4} [1 + \phi_0(\beta_r/2)]^2, \end{aligned} \quad (\text{I3})$$

$$a_1 = \frac{\alpha_r^2}{16} \left[\phi_0^{(1)}(\beta_r/2)^2 + \phi_0(\beta_r/2) \phi_0^{(2)}(\beta_r/2) + \phi_0^{(2)}(\beta_r/2) \right], \quad (\text{I4})$$

$$\begin{aligned} a_2 &= \frac{\alpha_r^4}{256} \left[12 \phi_0^{(2)}(\beta_r/2)^2 + 4 \phi_0(\beta_r/2) \phi_0^{(4)}(\beta_r/2) \right. \\ &\quad \left. + 16 \phi_0^{(1)}(\beta_r/2) \phi_0^{(3)}(\beta_r/2) + \phi_0^{(4)}(\beta_r/2) \right], \end{aligned} \quad (\text{I5})$$

where we use the following identities involving the derivatives of tanh:

$$\phi_0(x) = \tanh(x), \quad (\text{I6})$$

$$\phi_0^{(1)}(x) = 1 - \phi_0(x)^2, \quad (\text{I7})$$

$$\phi_0^{(2)}(x) = -2\phi_0(x)[1 - \phi_0(x)^2], \quad (\text{I8})$$

$$\phi_0^{(3)}(x) = 2[1 - \phi_0(x)^2][3\phi_0(x)^2 - 1], \quad (\text{I9})$$

$$\phi_0^{(4)}(x) = -8\phi_0(x)[1 - \phi_0(x)^2][3\phi_0(x)^2 - 2]. \quad (\text{I10})$$

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

$$\Delta_h \approx [c_0 + c_1 \Delta_h + c_2 \Delta_h^2] \langle \sigma_r(\sqrt{\Delta_r} x)^2 \rangle_x, \quad (\text{I11})$$

$$c_0 = \phi_0(\beta_h)^2, \quad (\text{I12})$$

$$c_1 = g_h^2 [\phi_0^{(1)}(\beta_h)^2 + \phi_0^{(2)}(\beta_h) \phi_0(\beta_h)], \quad (\text{I13})$$

$$\begin{aligned} c_2 &= g_h^4 \left[\frac{1}{4} \phi_0(\beta_h) \phi_0^{(4)}(\beta_h) \right. \\ &\quad \left. + \phi_0^{(1)}(\beta_h) \phi_0^{(3)}(\beta_h) + \frac{3}{4} \phi_0^{(2)}(\beta_h)^2 \right], \end{aligned} \quad (\text{I14})$$

and, therefore,

$$\Delta_h \approx (c_0 + c_1 \Delta_h + c_2 \Delta_h^2)(a_0 + a_1 g_h^2 \Delta_h + a_2 g_h^4 \Delta_h^2). \quad (\text{I15})$$

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

$$-1 + \phi'(0)\sigma_r(0) = 0. \quad (\text{I16})$$

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 :

$$\Delta_h = \epsilon^\eta (d_0 + d_1 \epsilon + d_2 \epsilon^2), \quad (\text{I17})$$

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 make sense only for small β_h . How small should β_h be relative to ϵ ? We make the following ansatz:

$$\beta_h = \beta_0 \epsilon^\delta, \quad (\text{I18})$$

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

We now express the coefficients for small β_h :

$$c_0 \approx \beta_0^2 \epsilon^{2\delta}, \quad (\text{I19})$$

$$c_1 \approx g_h^2 (1 - 2\beta_h^2), \quad (\text{I20})$$

$$c_2 \approx g_h^4 (-2 + 17\beta_h^2). \quad (\text{I21})$$

After solving Eqs. (I15)–(I19) self-consistently in terms of the expansion parameter ϵ , we get the following perturbative solution for $\delta \leq 1$:

$$\Delta_h \approx \begin{cases} \frac{2\beta_0 \epsilon^\delta}{g_c^2 (2 - g_c^2 a_1)} & g_c^2 a_1 < 2, \\ (g_c^2 a_1 - 2)f_1 + \epsilon \cdot f_2 & g_c^2 a_1 > 2, \end{cases} \quad (\text{I22})$$

where $a_1 = \frac{\alpha_r^2}{16} [\phi_0^{(1)}(\beta_r/2)^2 + \phi_0(\beta_r/2)\phi_0^{(2)}(\beta_r/2)]$. (I23)

$f_2(\alpha_r, \beta_r)$ and $f_2(\alpha_r, \beta_r)$ are constant functions (with respect to 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 nonzero c_0 corresponding to the bifurcation point. Since we expect, Δ_h to be small in this regime, we can expand Δ_r as

$$\Delta_r \approx g_h^2 \Delta_h - 2g_h^4 \Delta_h^2 + \frac{17}{3} g_h^6 \Delta_h^3 + O(\Delta_h^4), \quad (\text{I24})$$

and, similarly, we can also approximate

$$\langle \sigma_r(\sqrt{\Delta_r} x)^2 \rangle_x \approx \frac{1}{4} \left[1 + \frac{\alpha_r^2}{4} \Delta_r - \frac{\alpha_r^4}{8} \Delta_r^2 \right]. \quad (\text{I25})$$

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

$$c_0 = \frac{3(\alpha_r^2 - 8)}{2(-136 + 24\alpha_r^2 + 3\alpha_r^4)}, \quad (\text{I26})$$

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

$$\Delta_h \approx \frac{1}{2} [g_h^2 \Delta_h - 2g_h^4 \Delta_h^2]. \quad (\text{I27})$$

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

$$\Delta_h^*(\sqrt{2}^+) \sim \frac{g_h^2 - 2}{2g_h^4} \quad \text{for } g_h \rightarrow \sqrt{2}^+. \quad (\text{I28})$$

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

Here, we study the asymptotic behavior 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 $O(C_z^3)$,

$$C_{\sigma_z}(\tau) = g_0 + g_1 C_z(\tau) + g_3 C_z(\tau)^3, \quad (\text{I29})$$

$$\text{where } g_0 = \frac{1}{4}, \quad (\text{I30})$$

$$g_1 = \frac{\alpha_z^2}{16} - \frac{\alpha_z^4}{32} C_z(0) + \frac{5\alpha_z^6}{256} C_z(0)^2, \quad (\text{I31})$$

$$g_3 = \frac{\alpha_z^6}{384} - \frac{\alpha_z^8}{192} C_z(0). \quad (\text{I32})$$

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

$$C_\phi(\tau) = g_0 + g_1 C_h(\tau) + g_3 C_h(\tau)^3, \quad (\text{I33})$$

$$\text{where } g_0 = 0 \quad (\text{for } \beta_h = 0), \quad (\text{I34})$$

$$g_1 = g_h^2 - 2g_h^4 C_h(0) + 5g_h^6 C_h(0)^2, \quad (\text{I35})$$

$$g_3 = \frac{2}{3} g_h^6 - \frac{16}{3} g_h^8 C_h(0). \quad (\text{I36})$$

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

$$[-\partial_\tau^2 + C_{\sigma_z}(\tau)] C_h(\tau) = \frac{1}{4} C_{\sigma_z}(\tau) C_\phi(\tau), \quad (\text{I37})$$

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

$$\partial_\tau^2 C_h(\tau) \simeq a_1 C_h(\tau) + a_2 C_h(\tau)^2 + a_3 C_h(\tau)^3, \quad (\text{I38})$$

$$\text{where } a_1 = \frac{1}{16} (4 - \Gamma), \quad (\text{I39})$$

$$a_2 = \frac{\alpha_z^2}{64} (4 - \Gamma) \Gamma, \quad (\text{I40})$$

$$a_3 = -\frac{g_h^6}{24}, \quad (\text{I41})$$

$$\Gamma = g_h^2 - 2g_h^4 C_h(0) + 5g_h^6 C_h(0)^2. \quad (\text{I42})$$

Integrating with respect to τ gives

$$[\partial_\tau C_h(\tau)]^2 = 2 \left(\frac{a_1}{2} C_h(\tau)^2 + \frac{a_2}{3} C_h(\tau)^3 + \frac{a_3}{4} C_h(\tau)^4 + \text{const} \right). \quad (\text{I43})$$

The boundary conditions are

$$\partial_\tau C_h(0) = 0, \quad \lim_{\tau \rightarrow \infty} \partial_\tau C_h(\tau) = 0. \quad (\text{I44})$$

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

$$\frac{a_1}{2} + \frac{a_2}{3} C_h(0) + \frac{a_3}{4} C_h(0)^2 = 0. \quad (\text{I45})$$

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

$$C_h(0) \approx \frac{1}{8} \epsilon + O(\epsilon^2), \quad (\text{I46})$$

and the α_z enters the timescale-governing term a_1 only at $O(\epsilon^2)$. At first, it might seem counterintuitive 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 J: TOPOLOGICAL COMPLEXITY VIA KAC-RICE FORMULA

The arguments here are similar to those presented in Ref. [82], which use a self-averaging assumption to express the topological complexity (defined below) in terms of a spectral integral. Let us begin.

The goal is to estimate the total number of fixed points for a dynamical system $\dot{\mathbf{x}} = G(\mathbf{x})$. The Kac-Rice analysis proceeds by constructing the integral over the state space \mathbf{x} whose integrand has delta-functional support only on the fixed points:

$$\mathcal{N} = \int d\mathbf{x} \mathbb{E}[\delta(G(\mathbf{x})) |\det \mathcal{D}|], \quad (\text{J1})$$

where $\mathcal{D} = \partial G / \partial \mathbf{x}$ is the instantaneous Jacobian. The expectation value here is over the random coupling matrices. The average number of fixed points is related to the so-called topological complexity \mathcal{C} via the definition

$$\mathcal{N} = \exp(N\mathcal{C}).$$

We seek a saddle-point approximation of this quantity below.

For the gated RNN, the state space $\mathbf{x} = (\mathbf{h}, \mathbf{z}, \mathbf{r})$, and the fixed points satisfy

$$\sigma_z(z_i)(-h_i + \eta_i^h) = 0, \quad (\text{J2})$$

$$-z_i + \eta_i^z = 0, \quad (\text{J3})$$

$$-r_i + \eta_i^r = 0, \quad (\text{J4})$$

where for notational shorthand we introduce $\eta_i^h = \sum_j J_{ij}^h \phi(h_j) \sigma_r(r_j)$ and $\eta_i^{r/z} = \sum_j J_{ij}^{r/z} \phi(h_j)$, anticipating the mean-field approximation to come. Notice that only the first equation for h provides a nontrivial constraint. Once h is found, the second and third equations can be used to determine z and r , respectively. Notice, furthermore, that, since $\sigma(z_i) > 0$, the solutions h_i to the first equation do not depend on z_i . Indeed, the dependence on $\sigma(z)$ can be factorized out of the Kac-Rice integral. This requires noting first that, for the fixed point Jacobian, Eq. (A6) implies that the Jacobian can be written (setting $\tau_r = \tau_z = 1$ for simplicity)

$$\mathcal{D} = A(-1 + \mathcal{J}R) \quad (\text{J5})$$

and that the determinant can be factorized:

$$\det |\mathcal{D}| = \det |A| \times \det |-1 + \mathcal{J}R| \quad (\text{J6})$$

$$= \left(\prod_i \sigma(z_i) \right) \det |-1 + \mathcal{J}R|. \quad (\text{J7})$$

The product of $\sigma(z_i)$ produced by the determinant is canceled by the product of delta functions, using the fact that $\sigma(z_i) > 0$ and the transformation law

$$\prod_i \delta[\sigma(z_i)(-h_i + \eta_i^h)] = \frac{1}{\prod_i \sigma(z_i)} \prod_i \delta[-h_i + \eta_i^h]. \quad (\text{J8})$$

So we see that what evidently matters for the topological complexity is the *fixed-point Jacobian*:

$$\mathcal{D}^{fp} = -1 + \mathcal{J}R, \quad (\text{J9})$$

whose eigenvalues we denote by λ_i for $i = 1, \dots, N$ and with the spectral density

$$\hat{\mu}(z) = \frac{1}{N} \sum_i \delta^{(2)}(z - \lambda_i). \quad (\text{J10})$$

The preceding analysis is all basically to show that we could easily have set $\alpha_z = 0$ and gotten the same answer; i.e., the z gate does not influence the topological properties of the dynamics. For $\alpha_z = \infty$, the situation changes drastically, and the analysis likely needs to be significantly reworked. Indeed, in this limit, we most likely do not have discrete fixed points anymore, so the very notion of counting fixed points no longer makes sense.

Having introduced the spectral density, we can rewrite the Kac-Rice integral as

$$\mathcal{N} = \int \prod_{\mathbf{x} \in \{h, r\}} d\mathbf{x} \mathbb{E} \left[\delta^{(N)}(\mathbf{x} - \boldsymbol{\eta}^{\mathbf{x}}) e^{N \int d^2 z \hat{\mu}(z) \log |z|} \right]. \quad (\text{J11})$$

Note that, since the spectral density of \mathcal{D}^{fp} is independent of z , the integral over z is trivial to perform and leaves only h and r in the integrand.

So far, everything is exact. We begin now to make some approximations. The first crucial approximation is that the spectral density is self-averaging. The RMT analysis in the previous sections shows us furthermore that the spectral density depends only on macroscopic correlation functions of the state variables. Let us denote the spectral integral factor

$$I(\mathbf{x}, \mathcal{J}) = \exp \left(N \int d^2 z \hat{\mu}(z) \log |z| \right), \quad (\text{J12})$$

by which we mean that it depends on the particular realization of the random coupling \mathcal{J} and the state vector \mathbf{x} . The self-averaging assumption implies that

$$I(\mathbf{x}, \mathcal{J}) \approx \mathbb{E}[I(\mathbf{x}, \mathcal{J})] \equiv \bar{I}(\mathbf{x}); \quad (\text{J13})$$

i.e., this factor does not depend on the particular realization of \mathcal{J} but just on the state vector. Equivalently, we are assuming that the spectral density $\hat{\mu}_{\mathbf{h}, \mathbf{r}}(z)$ depends only on the configurations \mathbf{h} and \mathbf{r} and not the particular realization $J^{h, r}$. This allows us to pull this factor outside of the expectation value:

$$\mathcal{N} \approx \int \prod_{\mathbf{x} \in \{h, r\}} \bar{I}(\mathbf{x}) \mathbb{E}[\delta^{(N)}(\mathbf{x} - \boldsymbol{\eta}^{\mathbf{x}})]. \quad (\text{J14})$$

Now we give some nonrigorous arguments for how one might evaluate the remaining expectation value. In order to carry out the average over J^h and J^r , we utilize the Fourier representation of the delta function to write

$$\mathbb{E} \left[\int d\hat{\mathbf{x}} e^{i\hat{\mathbf{x}}(\mathbf{x} - \boldsymbol{\eta}^{\mathbf{x}})} \right] \quad (\text{J15})$$

$$= \int d\mathbf{x} d\hat{\mathbf{x}} \mathbb{E} \left[\exp \sum_{i,j} \{ i\hat{h}_i [h_i - J_{ij}^h \phi(h_j) \sigma_r(r_j)] \right. \quad (\text{J16})$$

$$\left. + i\hat{r}_i [r_i - J_{ij}^r \phi(h_j)] \} \right], \quad (\text{J17})$$

which upon disorder averaging yields

$$\int d\mathbf{x} d\hat{\mathbf{x}} \exp \left\{ \sum_i \left(i\hat{h}_i h_i + i\hat{r}_i r_i - \frac{1}{2} \hat{h}_i^2 \hat{C}_{\phi\sigma_r} - \frac{1}{2} \hat{r}_i^2 \hat{C}_{\phi} \right) \right\}, \quad (\text{J18})$$

where we define

$$\hat{C}_{\phi\sigma_r} = \frac{1}{N} \sum_i \phi(h_i)^2 \sigma_r(r_i)^2, \quad \hat{C}_\phi = \frac{1}{N} \sum_i \phi(h_i)^2. \quad (\text{J19})$$

This is where we make our second crucial assumption: that the empirical averages appearing in Eq. (J19) converge to their average value

$$\hat{C}_{\phi\sigma_r} \rightarrow C_{\phi\sigma_r} = \mathbb{E}_{\mathbf{h}, \mathbf{r}} \left[\frac{1}{N} \sum_i \phi(h_i)^2 \sigma_r(r_i)^2 \right], \quad (\text{J20})$$

$$\hat{C}_\phi \rightarrow C_\phi \equiv \mathbb{E}_{\mathbf{h}} \left[\frac{1}{N} \sum_i \phi(h_i)^2 \right]. \quad (\text{J21})$$

This means we are assuming the strong law of large numbers. With this essential step, the integral in Eq. (J18) evaluates to

$$\frac{1}{\sqrt{2\pi\Delta_h}} \frac{1}{\sqrt{2\pi\Delta_r}} \exp(-\|\mathbf{h}\|^2/2\Delta_h - \|\mathbf{r}\|^2/2\Delta_r) \quad (\text{J22})$$

$$= \prod_{i=1}^N P_h(h_i) P_r(r_i), \quad (\text{J23})$$

where $\Delta_h = C_{\phi\sigma_r}$ and $\Delta_r = C_\phi$ —which are just the time-independent (fixed-point) MFT equations (C26).

Returning to the expression for the complexity, this series of approximations gives us

$$\mathcal{N} \approx \int \prod_{i=1}^N dh_i dr_i P_h(h_i) P_r(r_i) \bar{I}(\mathbf{r}, \mathbf{h}). \quad (\text{J24})$$

Let us now describe our derivation more intuitively. We start with the formal expression for the Kac-Rice formula, which uses the delta functional integrand to find fixed points and counts them with the weighting factor related to the Jacobian. Our first assumption allows us to simplify the calculation involving the Jacobian, since we argue that this term is self-averaging. The second assumption allows us to deal with the remaining expectation value of the delta functions. The expectation value adds a number of delta functions (however many there may be for that $J^{h/r}$) for each configuration of the connectivity. For continuously distributed connectivity, this implies that the expectation value smears out the delta functions and results in a smooth distribution. What should this distribution be? Well, we know from the mean-field analysis that the state vectors are distributed as Gaussians at a fixed point. Furthermore, the mean-field theory becomes exact for large N . Therefore, we should expect that, in this limit, the delta functions are smeared out into the Gaussian distributions determined by the MFT. This is what our derivation shows.

The final step is to recall that the spectral density depends on the state vectors only via empirical averages. For instance, in the absence of an r gate, the spectral

density depends on the empirical average $\hat{C}_{\phi'}$. Again invoking the strong law of large numbers, we may argue that the self-averaging goes a step further and that

$$\mathcal{N} \approx \int \prod_i dP_h(h_i) dP_r(r_i) \exp \left(N \int d^2 z \hat{\mu}_{\mathbf{h}, \mathbf{r}}(z) \log |z| \right) \quad (\text{J25})$$

$$\approx \exp \left(N \int d^2 z \bar{\mu}(z) \log |z| \right), \quad (\text{J26})$$

where

$$\bar{\mu}(z) = \mathbb{E}_{\mathbf{h}, \mathbf{r}} [\hat{\mu}_{\mathbf{h}, \mathbf{r}}(z)] = \int d\mathbf{h} d\mathbf{r} P(\mathbf{h}) P(\mathbf{r}) \hat{\mu}_{\mathbf{h}, \mathbf{r}}(z). \quad (\text{J27})$$

This is precisely the spectral density we study in a preceding Appendix and the one for which we obtain an explicit expression for the spectral curve. These approximations give us the topological complexity

$$\mathcal{C} = \int d^2 z \bar{\mu}(z) \log |z|. \quad (\text{J28})$$

Now we take a closer look at the spectral density. The eigenvalues of \mathcal{D}^{fp} form a circular droplet of finite radius ρ and centered on -1 . Therefore, the eigenvalues have the form $\lambda = -1 + re^{i\theta}$, and the spectral density is a function only of r . The value of the radius is found from Eq. (5) by removing the z gate (i.e., setting $\alpha_z = 0$). After some algebraic steps, we find for the radius

$$\rho^2 = \frac{1}{2} \left(C_1 + \sqrt{C_1^2 + 4C_2} \right), \quad (\text{J29})$$

$$C_1 = C_{\phi'} C_{\sigma_r}, \quad C_2 = C_{\phi'} C_{\sigma_r'} C_\phi. \quad (\text{J30})$$

Using these facts, we can write the topological complexity as

$$\mathcal{C} = \int r dr d\theta \bar{\mu}(r) \mathbb{I}_{\{r < \rho\}} \log |re^{i\theta} - 1| \quad (\text{J31})$$

$$= \begin{cases} 2\pi \int_1^\rho r dr \log r \bar{\mu}(r) \geq 0, & \text{for } \rho > 1, \\ 0, & \text{for } \rho < 1, \end{cases} \quad (\text{J32})$$

where $\mathbb{I}_{\{r < \rho\}}$ is the indicator function which is one for $r < \rho$ and vanishes for $r > \rho$. Thus, we see that the topological complexity is zero for $\rho < 1$. This is *precisely the fixed-point stability condition derived in the main text [Eq. (6)]*. Conversely, the topological complexity is nonzero for $\rho > 1$, which corresponds to unstable fixed points. Thus, we see, under our set of reasonable approximations, unstable MFT fixed points correspond to a finite

topological complexity and, consequently, to a number of “microscopic” fixed points that grows exponentially with N .

The final missing ingredient, necessary to show that region 2 in the phase diagram has an exponentially growing number of fixed points, is to show that the MFT fixed points which appear after the bifurcation are indeed unstable. At the moment, we lack any analytical handle on this. However, we confirm numerically that, along the bifurcation curve, the fixed points are unstable and that increasing the variance Δ_h serves only to increase ρ . However, is it possible for the lower branch, on which Δ_h decreases with α_r ? Evidently not, since Δ_h scales with α_r in such a way that C_{σ_r} ends up growing like α_r^2 , thus once again increasing ρ . Therefore, we conclude that the MFT fixed points appearing after the bifurcation are always unstable, with $\rho > 1$. This concludes our informal proof of the transition in topological complexity between regions 1 and 2 in the phase diagram in Fig. 7.

-
- [1] A. Graves, A.-R. Mohamed, and G. Hinton, in *Proceedings of the 2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (IEEE, New York, 2013), pp. 6645–6649.
 - [2] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, *How to Construct Deep Recurrent Neural Networks*, [arXiv:1312.6026](#).
 - [3] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, *Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach*, *Phys. Rev. Lett.* **120**, 024102 (2018).
 - [4] P. R. Vlachas, J. Pathak, B. R. Hunt, T. P. Sapsis, M. Girvan, E. Ott, and P. Koumoutsakos, *Backpropagation Algorithms and Reservoir Computing in Recurrent Neural Networks for the Forecasting of Complex Spatiotemporal Dynamics*, *Neural Netw.* **126**, 191 (2020).
 - [5] L. Guastoni, P. A. Srinivasan, H. Azizpour, P. Schlatter, and R. Vinuesa, *On the Use of Recurrent Neural Networks for Predictions of Turbulent Flows*, [arXiv:2002.01222](#).
 - [6] R. Jozefowicz, W. Zaremba, and I. Sutskever, *An Empirical Exploration of Recurrent Network Architectures*, *Proc. Mach. Learn. Res.* **37**, 2342 (2015).
 - [7] T. P. Vogels, K. Rajan, and L. F. Abbott, *Neural Network Dynamics*, *Annu. Rev. Neurosci.* **28**, 357 (2005).
 - [8] Y. Ahmadian and K. D. Miller, *What Is the Dynamical Regime of Cerebral Cortex?*, [arXiv:1908.10101](#).
 - [9] J. Kadmon and H. Sompolinsky, *Transition to Chaos in Random Neuronal Networks*, *Phys. Rev. X* **5**, 041030 (2015).
 - [10] D. Sussillo and L. F. Abbott, *Generating Coherent Patterns of Activity from Chaotic Neural Networks*, *Neuron* **63**, 544 (2009).
 - [11] R. Laje and D. V. Buonomano, *Robust Timing and Motor Patterns by Taming Chaos in Recurrent Neural Networks*, *Nat. Neurosci.* **16**, 925 (2013).
 - [12] S. Hochreiter and J. Schmidhuber, *Long Short-Term Memory*, *Neural Comput.* **9**, 1735 (1997).
 - [13] S. J. Mitchell and R. A. Silver, *Shunting Inhibition Modulates Neuronal Gain during Synaptic Excitation*, *Neuron* **38**, 433 (2003).
 - [14] R. Gütig and H. Sompolinsky, *Time-Warp-Invariant Neuronal Processing*, *PLoS Biol.* **7**, e1000141 (2009).
 - [15] H. Sompolinsky, A. Crisanti, and H.-J. Sommers, *Chaos in Random Neural Networks*, *Phys. Rev. Lett.* **61**, 259 (1988).
 - [16] D. Martí, N. Brunel, and S. Ostojic, *Correlations between Synapses in Pairs of Neurons Slow Down Dynamics in Randomly Connected Neural Networks*, *Phys. Rev. E* **97**, 062314 (2018).
 - [17] F. Schuessler, A. Dubreuil, F. Mastrogiuseppe, S. Ostojic, and O. Barak, *Dynamics of Random Recurrent Networks with Correlated Low-Rank Structure*, *Phys. Rev. Research* **2**, 013111 (2020).
 - [18] F. Mastrogiuseppe and S. Ostojic, *Linking Connectivity, Dynamics, and Computations in Low-Rank Recurrent Neural Networks*, *Neuron* **99**, 609 (2018).
 - [19] M. Stern, H. Sompolinsky, and L. F. Abbott, *Dynamics of Random Neural Networks with Bistable Units*, *Phys. Rev. E* **90**, 062710 (2014).
 - [20] J. Aljadeff, M. Stern, and T. Sharpee, *Transition to Chaos in Random Networks with Cell-Type-Specific Connectivity*, *Phys. Rev. Lett.* **114**, 088101 (2015).
 - [21] J. Schuecker, S. Goedeke, and M. Helias, *Optimal Sequence Memory in Driven Random Networks*, *Phys. Rev. X* **8**, 041029 (2018).
 - [22] R. Brette, *Exact Simulation of Integrate-and-Fire Models with Synaptic Conductances*, *Neural Comput.* **18**, 2004 (2006).
 - [23] S.-I. Amari, *Characteristics of Random Nets of Analog Neuron-Like Elements*, *IEEE Trans. Syst. Man Cybernet.* **SMC-2**, 643 (1972).
 - [24] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, *Learning Phrase Representations Using RNN Encoder-Decoder for Statistical Machine Translation*, [arXiv:1406.1078](#).
 - [25] J. T. Chalker and B. Mehlig, *Eigenvector Statistics in Non-Hermitian Random Matrix Ensembles*, *Phys. Rev. Lett.* **81**, 3367 (1998).
 - [26] J. Feinberg and A. Zee, *Non-Hermitian Random Matrix Theory: Method of Hermitian Reduction*, *Nucl. Phys.* **B504**, 579 (1997).
 - [27] P. C. Martin, E. Siggia, and H. Rose, *Statistical Dynamics of Classical Systems*, *Phys. Rev. A* **8**, 423 (1973).
 - [28] C. De Dominicis, *Dynamics as a Substitute for Replicas in Systems with Quenched Random Impurities*, *Phys. Rev. B* **18**, 4913 (1978).
 - [29] J. A. Hertz, Y. Roudi, and P. Sollich, *Path Integral Methods for the Dynamics of Stochastic and Disordered Systems*, *J. Phys. A* **50**, 033001 (2017).
 - [30] H.-K. Janssen, *On a Lagrangean for Classical Field Dynamics and Renormalization Group Calculations of Dynamical Critical Properties*, *Z. Phys. B* **23**, 377 (1976).
 - [31] A. Crisanti and H. Sompolinsky, *Path Integral Approach to Random Neural Networks*, *Phys. Rev. E* **98**, 062120 (2018).

- [32] M. Helias and D. Dahmen, *Statistical Field Theory for Neural Networks* (Springer, New York, 2020).
- [33] T. Mora and W. Bialek, *Are Biological Systems Poised at Criticality?*, *J. Stat. Phys.* **144**, 268 (2011).
- [34] H. S. Seung, *Continuous Attractors and Oculomotor Control*, *Neural Netw.* **11**, 1253 (1998).
- [35] H. S. Seung, *How the Brain Keeps the Eyes Still*, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 13339 (1996).
- [36] H. S. Seung, D. D. Lee, B. Y. Reis, and D. W. Tank, *Stability of the Memory of Eye Position in a Recurrent Network of Conductance-Based Model Neurons*, *Neuron* **26**, 259 (2000).
- [37] C. K. Machens, R. Romo, and C. D. Brody, *Flexible Control of Mutual Inhibition: A Neural Model of Two-Interval Discrimination*, *Science* **307**, 1121 (2005).
- [38] R. Chaudhuri and I. Fiete, *Computational Principles of Memory*, *Nat. Neurosci.* **19**, 394 (2016).
- [39] W. Bialek, *Biophysics: Searching for Principles* (Princeton University Press, Princeton, NJ, 2012).
- [40] M. S. Goldman, *Memory without Feedback in a Neural Network*, *Neuron* **61**, 621 (2009).
- [41] N. Maheswaranathan, A. Williams, M. D. Golub, S. Ganguli, and D. Sussillo, *Reverse Engineering Recurrent Networks for Sentiment Classification Reveals Line Attractor dynamics*, in *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alch'e-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., New York, 2019), Vol 32, p. 15696.
- [42] M. Farrell, S. Recanatesi, T. Moore, G. Lajoie, and E. Shear-Brown, *Recurrent Neural Networks Learn Robust Representations by Dynamically Balancing Compression and Expansion*, *bioRxiv* 10.1101/564476.
- [43] L. Molgedey, J. Schuchhardt, and H. G. Schuster, *Suppressing Chaos in Neural Networks by Noise*, *Phys. Rev. Lett.* **69**, 3717 (1992).
- [44] K. Rajan, L. F. Abbott, and H. Sompolinsky, *Stimulus-Dependent Suppression of Chaos in Recurrent Neural Networks*, *Phys. Rev. E* **82**, 011903 (2010).
- [45] G. Wainrib and J. Touboul, *Topological and Dynamical Complexity of Random Neural Networks*, *Phys. Rev. Lett.* **110**, 118101 (2013).
- [46] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, *On the Importance of Initialization and Momentum in Deep Learning*, *Proc. Mach. Learn. Res.* **28**, 1139 (2013).
- [47] R. Legenstein and W. Maass, *Edge of Chaos and Prediction of Computational Performance for Neural Circuit Models*, *Neural Netw.* **20**, 323 (2007).
- [48] H. Jaeger and H. Haas, *Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication*, *Science* **304**, 78 (2004).
- [49] T. Toyozumi and L. F. Abbott, *Beyond the Edge of Chaos: Amplification and Temporal Integration by Recurrent Networks in the Chaotic Regime*, *Phys. Rev. E* **84**, 051908 (2011).
- [50] Since the Jacobian spectral density depends on correlation functions (see Appendix A), in the dynamical steady state the spectral density becomes time-translation invariant. In other words, the spectral density also reaches a steady-state distribution. As a result, a snapshot of the spectral density at any given time has the same form. Instability then implies that the eigenvectors must evolve over time in order to keep the dynamics bounded. The timescale involved in the evolution of the eigenvectors should correspond roughly with the correlation time implied by the DMFT. Within this window, the spectral analysis of the Jacobian in the steady state gives a meaningful description of the range of time-scales involved. Furthermore, we see empirically that the local structure appears very informative of the true dynamics, in particular, with understanding the emergence of continuous attractors and marginal stability, as we discuss in Sec. IV.
- [51] T. Can, K. Krishnamurthy, and D. J. Schwab, *Gating Creates Slow Modes and Controls Phase-Space Complexity in GRUs and LSTMs*, *Proc. Mach. Learn. Res.* **107**, 476 (2020).
- [52] The continuous-time gated RNN we study in this paper is most closely related to the GRU architecture studied in Ref. [51].
- [53] V. M. Eguíluz, M. Ospeck, Y. Choe, A. J. Hudspeth, and M. O. Magnasco, *Essential Nonlinearities in Hearing*, *Phys. Rev. Lett.* **84**, 5232 (2000).
- [54] J.-P. Eckmann and D. Ruelle, in *The Theory of Chaotic Attractors* (Springer, New York, 1985), pp. 273–312.
- [55] For reference, we also supply a bound on the maximal Lyapunov exponent in Appendix F, showing that the relaxation time of the dynamics enters into an upper bound on λ_{\max} .
- [56] B. Derrida and Y. Pomeau, *Random Networks of Automata: A Simple Annealed Approximation*, *Europhys. Lett.* **1**, 45 (1986).
- [57] B. Cessac, *Increase in Complexity in Random Neural Networks*, *J. Phys. I (France)* **5**, 409 (1995).
- [58] One might worry that the h and $\sigma(z)$ correlators are not separable, in general. However, this issue arises only for large α_z . For moderate α_z , the separability assumption is valid.
- [59] Y. V. Fyodorov, *Complexity of Random Energy Landscapes, Glass Transition, and Absolute Value of the Spectral Determinant of Random Matrices*, *Phys. Rev. Lett.* **92**, 240601 (2004).
- [60] Y. V. Fyodorov and P. Le Doussal, *Topology Trivialization and Large Deviations for the Minimum in the Simplest Random Optimization*, *J. Stat. Phys.* **154**, 466 (2014).
- [61] J. Pereira and X.-J. Wang, *A tradeoff between Accuracy and Flexibility in a Working Memory Circuit Endowed with Slow Feedback Mechanisms*, *Cereb. Cortex* **25**, 3586 (2015).
- [62] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, *LSTM: A Search Space Odyssey*, *IEEE Trans. Neural Netw. Learn. Syst.* **28**, 2222 (2017).
- [63] In fact, the fixed-point phase diagrams for the current model and the GRU are in one-to-one correspondence. What this static phase diagram importantly lacks is region 3 in Fig. 7.
- [64] C. Tallec and Y. Ollivier, *Can Recurrent Neural Networks Warp Time?*, *arXiv:1804.11188*.
- [65] S. P. Muscinelli, W. Gerstner, and T. Schwalger, *How Single Neuron Properties Shape Chaotic Dynamics and Signal Transmission in Random Neural Networks*, *PLoS Comput. Biol.* **15**, e1007122 (2019).
- [66] M. Beiran and S. Ostojic, *Contrasting the Effects of Adaptation and Synaptic Filtering on the Timescales of*

- Dynamics in Recurrent Networks*, *PLoS Comput. Biol.* **15**, e1006893 (2019).
- [67] U. Pereira and N. Brunel, *Attractor Dynamics in Networks with Learning Rules Inferred from In Vivo Data*, *Neuron* **99**, 227 (2018).
- [68] N. Bertschinger and T. Natschl ger, *Real-Time Computation at the Edge of Chaos in Recurrent Neural Networks*, *Neural Comput.* **16**, 1413 (2004).
- [69] R. Legenstein and W. Maass, in *New Directions in Statistical Signal Processing: From Systems to Brain*, edited by Simon Haykin, Jose C. Principe, Terrence J. Sejnowski, and John McWhirter (The MIT Press, 2006), p. 127.
- [70] J. Boedecker, O. Obst, J. T. Lizier, N. M. Mayer, and M. Asada, *Information Processing in Echo State Networks at the Edge of Chaos*, *Theory Biosci.* **131**, 205 (2012).
- [71] S. Geman and C.-R. Hwang, *A Chaos Hypothesis for Some Large Systems of Random Equations*, *Z. Wahrscheinlichkeitstheorie Verwandte Gebiete* **60**, 291 (1982).
- [72] 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.
- [73] H. Sompolinsky and A. Zippelius, *Relaxational Dynamics of the Edwards-Anderson Model and the Mean-Field Theory of Spin-Glasses*, *Phys. Rev. B* **25**, 6860 (1982).
- [74] H. Sompolinsky and A. Zippelius, *Dynamic Theory of the Spin-Glass Phase*, *Phys. Rev. Lett.* **47**, 359 (1981).
- [75] C. C. Chow and M. A. Buice, *Path Integral Methods for Stochastic Differential Equations*, *J. Math. Neurosci.* **5**, 8 (2015).
- [76] F. Roy, G. Biroli, G. Bunin, and C. Cammarota, *Numerical Implementation of Dynamical Mean Field Theory for Disordered Systems: Application to the Lotka–Volterra Model of Ecosystems*, *J. Phys. A* **52**, 484001 (2019).
- [77] K. Geist, U. Parlitz, and W. Lauterborn, *Comparison of Different Methods for Computing Lyapunov Exponents*, *Prog. Theor. Phys.* **83**, 875 (1990).
- [78] R. Engelken, F. Wolf, and L. Abbott, *Lyapunov Spectra of Chaotic Recurrent Neural Networks*, [arXiv:2006.02427](https://arxiv.org/abs/2006.02427).
- [79] The local chaos hypothesis employed by Cessac [57] amounts to the same assumption.
- [80] Strictly speaking, Oseledets theorem guarantees that $\lambda_{\max} = \lim_{t \rightarrow \infty} (1/2t) \log[(\|\chi \mathbf{u}\|^2)/(\|\mathbf{u}\|^2)]$ for almost every \mathbf{u} . In particular, we can take \mathbf{u} to be the all-ones vector. The term inside the log then becomes $\frac{1}{N} \sum_{i,k} \chi_{ik}^2 + (1/N) \sum_{i \neq j} \sum_k \chi_{ik} \chi_{jk}$, and the second term is subleading in N , since the susceptibilities are random functions. This justifies Eq. (F1).
- [81] 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 $\beta_r \rightarrow \infty$.
- [82] J. R. Ipsen and A. D. H. Peterson, *Consequences of Dale's Law on the Stability-Complexity Relationship of Random Neural Networks*, *Phys. Rev. E* **101**, 052412 (2020).