Random Recurrent Networks and Dynamical Mean-Field Theory

Lecture 6

Roadmap. In Lecture 5, we designed the recurrent matrix $\mathbf{J}$ to produce integration, selective amplification, or attractor memories. What happens if we instead draw its entries at random? We first use the circular law to find when the zero state loses stability. We then use dynamical mean-field theory (DMFT) to describe the fluctuations that replace it. The calculation follows the sequence

$$\text{random connectivity} \longrightarrow \text{Gaussian input} \longrightarrow \text{self-consistent correlations.}$$

The last step is to distinguish irregular activity from chaos by studying sensitivity to initial conditions.

Notation and connection to Lecture 5. We use the current-based convention introduced there:

$$\tau_r \frac{d\mathbf{x}}{dt} = -\mathbf{x} + \mathbf{J}\phi(\mathbf{x}) + \mathbf{h}. \tag{1}$$

Here x_i is a synaptic-current-like variable, $r_i = \phi(x_i)$ is its instantaneous output, J_{ij} connects presynaptic neuron j to postsynaptic neuron i, and $\mathbf{h}$ is external input. For the mathematical model $\phi(x) = \tanh x$, r_i is signed activity, not a nonnegative biological firing rate. The corresponding rate-based model has the nonlinearity outside the recurrent sum; as discussed in Lecture 5, the two descriptions are related by a change of variables under the appropriate assumptions.

From now on, measure time in units of τ_r and write $\tau_r = 1$. A physical decay rate or Lyapunov exponent is obtained by dividing its dimensionless value by τ_r. We reserve $\Delta = t - s$ for a time lag, so it cannot be confused with the neuronal time constant. The symbol g is the random-coupling gain, unrelated to the synaptic conductance $\bar{g}_s$ in Lecture 5.

The linear network: stability without symmetry

Set $\phi(x) = x$ in Equation 1. In dimensionless time,

$$\dot{\mathbf{x}} = (\mathbf{J} - \mathbf{I})\mathbf{x} + \mathbf{h}(t). \tag{2}$$

Its exact solution is

$$\mathbf{x}(t) = e^{(\mathbf{J}-\mathbf{I})t}\mathbf{x}(0) + \int_0^t e^{(\mathbf{J}-\mathbf{I})(t-s)}\mathbf{h}(s)\, ds. \tag{3}$$

This expression remains valid even when $\mathbf{J}$ is not diagonalizable. If it is diagonalizable, write $\mathbf{J} = \mathbf{R}\mathbf{\Lambda}\mathbf{R}^{-1}$ and let $\mathbf{c} = \mathbf{R}^{-1}\mathbf{x}$, $\mathbf{b} = \mathbf{R}^{-1}\mathbf{h}$. For constant input,

$$c_\mu(t) = e^{(\lambda_\mu - 1)t} c_\mu(0) + \frac{b_\mu}{1-\lambda_\mu}\left[1 - e^{(\lambda_\mu - 1)t}\right], \tag{4}$$

when $\lambda_\mu \neq 1$. At $\lambda_\mu = 1$, the expression becomes $c_\mu(t) = c_\mu(0) + b_\mu t$.

Unlike the symmetric matrix of Lecture 5, a real asymmetric matrix can have complex eigenvalues, which occur in conjugate pairs. The real part of $\lambda_\mu - 1$ controls growth or decay; its imaginary part produces oscillation. Consequently,

$$\boxed{\text{Linear asymptotic stability} \quad \Longleftrightarrow \quad \max_\mu \operatorname{Re}\lambda_\mu(\mathbf{J}) < 1.} \tag{5}$$

For a stable network and constant input, $\mathbf{x}_\infty = (\mathbf{I} - \mathbf{J})^{-1}\mathbf{h}$.

Stable eigenvalues do not exclude transient amplification

Asymmetric random matrices are generally non-normal: $\mathbf{J}\mathbf{J}^T \neq \mathbf{J}^T\mathbf{J}$. Their eigenvectors are not generally orthogonal. For example, let

$$\mathbf{J} = \begin{pmatrix} 0 & K \\ 0 & 0 \end{pmatrix}, \qquad \mathbf{x}(0) = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Both eigenvalues are zero, but

$$\mathbf{x}(t) = e^{-t}\begin{pmatrix} Kt \\ 1 \end{pmatrix}.$$

For large K, a perturbation is strongly amplified before it eventually decays. Eigenvalues determine asymptotic linear stability, whereas singular values of $e^{(\mathbf{J}-\mathbf{I})t}$ determine the largest amplification at a specified time. This distinction matters when interpreting random network responses. Asymmetry also removes the particular Hopfield energy construction, but does not imply that every possible Lyapunov function is absent or that the dynamics must be chaotic.

A random recurrent matrix and the circular law

We now choose a fully asymmetric Gaussian ensemble:

$$\overline{J_{ij}} = 0, \qquad \overline{J_{ij}J_{kl}} = \frac{g^2}{N}\delta_{ik}\delta_{jl}. \tag{6}$$

The overbar denotes an average over matrices. Entries are independent, including J_{ij} and J_{ji} for $i \neq j$. Here "fully asymmetric" means uncorrelated reciprocal couplings; it does not mean antisymmetric ($\mathbf{J}^T = -\mathbf{J}$). Setting $J_{ii} = 0$ changes none of the large-N results discussed below.

Why scale individual weights as $N^{-1/2}$? If presynaptic activities have an order-one second moment and no coherent alignment with the incoming weights, the variance of a summed input is of order

$$N \times \frac{g^2}{N} \times O(1) = O(1).$$

Thus the recurrent fluctuations neither disappear nor diverge as N grows. This is a fluctuation argument: summing the absolute values of the weights would give a different estimate. The self-consistent justification for the activity assumptions will be supplied by DMFT.

Statement and interpretation of the circular law

Let λ_μ be the eigenvalues of $\mathbf{J}$ and define their empirical measure in the complex plane by

$$\rho_N(z) = \frac{1}{N} \sum_{\mu=1}^{N} \delta^{(2)}(z - \lambda_\mu).$$

For independent, identically distributed centered entries with variance g^2/N, the limiting eigenvalue measure is uniform in a disk [2]:

$$\boxed{\rho(z) = \frac{1}{\pi g^2} \mathbf{1}_{\{|z|<g\}}.} \tag{7}$$

This is a statement of convergence of measures, not pointwise convergence of the delta functions in ρ_N. The area of the disk is πg^2, so the density integrates to one. Uniformity in area also means that, for $0 \le a \le g$, the limiting fraction of eigenvalues with $|\lambda| < a$ is a^2/g^2; their radii are not uniformly distributed.

The circular law is distinct from the semicircle law for symmetric random matrices, whose eigenvalues are real. Independence of the reciprocal couplings is essential to this distinction. The theorem is much more universal than the Gaussian example, but we retain Gaussian couplings throughout the lecture.

For this Gaussian ensemble the spectral edge also approaches the disk boundary. Combining this fact with Equation 5, the linearized dynamics has a stability boundary at $g = 1$. Strictly speaking, a limiting eigenvalue density alone does not exclude a finite number of outliers; the edge behavior of this ensemble supplies that additional information. At finite N the rightmost eigenvalue fluctuates, so the sample-dependent threshold need not be exactly one.

From loss of stability to nonlinear dynamics

Consider the autonomous network

$$\boxed{\dot{x}_i = -x_i + \sum_{j=1}^{N} J_{ij}\phi(x_j), \qquad \phi(x) = \tanh x.} \tag{8}$$

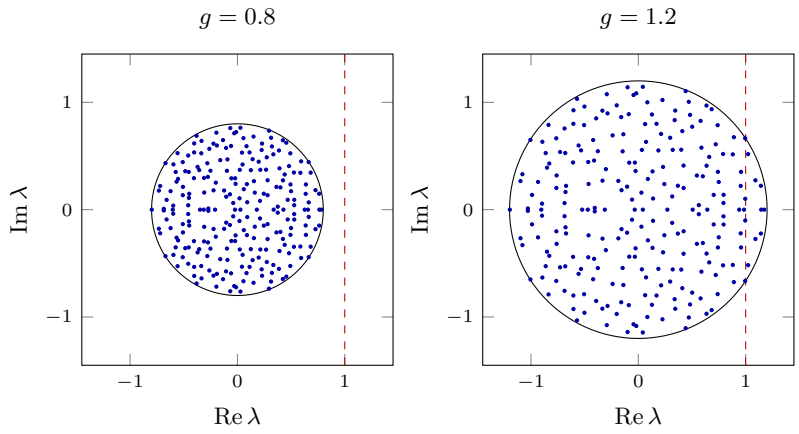


Figure 1: Eigenvalues of Gaussian recurrent matrices with $N = 240$. The same standard-normal matrix is rescaled to the two gains. Black circles show the limiting support $|\lambda| = g$; the dashed line marks $\mathrm{Re}\,\lambda = 1$. The eigenvalues of the linear dynamics $\mathbf{J} - \mathbf{I}$ are obtained by shifting these points one unit to the left. These finite samples illustrate the law; they are not its proof.

Because $\phi(0) = 0$, $\mathbf{x} = 0$ is always a fixed point. Linearization around it gives

$$\delta\dot{\mathbf{x}} = [-\mathbf{I} + \phi'(0)\mathbf{J}]\delta\mathbf{x}.$$

For positive slope $\phi'(0)$ the large-N threshold is $g\phi'(0) = 1$; for tanh, $\phi'(0) = 1$ and hence $g_c = 1$.

For $g < 1$, sufficiently small perturbations decay. For $g > 1$, the origin is unstable, but saturation prevents the unbounded growth predicted by the linear model. For any fixed finite realization, bounded ϕ bounds each recurrent input by $\sum_j |J_{ij}|$, and the leak keeps the trajectory bounded. Boundedness and instability alone do not determine the attractor: other fixed points, periodic orbits, and chaos are different possibilities.

For this fully asymmetric autonomous model, the classical large-N result is a transition to self-sustained chaotic activity for $g > 1$ [1]. We shall see how its correlations are selected by DMFT, and why a further stability calculation establishes chaos. Finite networks can show sample-dependent behavior and long transients. Moreover, a trajectory initialized exactly at zero remains there even when zero is unstable; a simulation must start with a nonzero perturbation.

Quenched disorder and the averages we need

Draw $\mathbf{J}$ once and hold it fixed during the dynamics. This is *quenched disorder*. Redrawing the weights at every time step would produce a different, externally randomized model.

We use $\langle\cdot\rangle$ for averages over the effective stochastic process, or over a specified ensemble of initial conditions in the original network. For a stationary mixing

state, long-time averages can estimate the same quantities. Typical population averages are assumed to concentrate around deterministic limits as $N \to \infty$: this is self-averaging. It is an assertion about appropriate averaged observables, not a claim that all individual neurons have identical trajectories or that time and disorder averages are automatically interchangeable in every state.

The basic observables are the current and output correlations,

$$C(t,s) = \lim_{N\to\infty} \frac{1}{N} \sum_i \overline{\langle x_i(t) x_i(s)\rangle}, \tag{9}$$

$$C_{\phi\phi}(t,s) = \lim_{N\to\infty} \frac{1}{N} \sum_i \overline{\langle \phi(x_i(t))\phi(x_i(s))\rangle}. \tag{10}$$

With zero input and a sign-symmetric initial ensemble the mean vanishes, so these are also covariances. We do not use $r(t,s)$ for a correlation, since r_i already denotes neuronal output in Lecture 5.

The effective Gaussian input

Define the recurrent input

$$\eta_i(t) = \sum_j J_{ij} \phi(x_j(t)). \tag{11}$$

It is tempting to factor a disorder average as

$$\overline{J_{ij}J_{ik}\phi(x_j(t))\phi(x_k(s))} \stackrel{?}{=} \overline{J_{ij}J_{ik}}\,\overline{\phi(x_j(t))\phi(x_k(s))}.$$

This is not an exact identity at finite N: the trajectories themselves depend on the couplings. A cavity argument explains the leading large-N result rather than assuming this factorization.

Remove neuron i and all its incident connections, and denote the remaining trajectories by $x_j^{(i)}(t)$. They are independent of the removed incoming row $\{J_{ij}\}_j$. The cavity input is

$$\eta_i^{\mathrm{cav}}(t) = \sum_{j\neq i} J_{ij} \phi(x_j^{(i)}(t)).$$

Conditioned on the cavity trajectories, this is jointly Gaussian at any finite set of times, because it is a linear combination of Gaussian weights. Its conditional mean is zero and its covariance is exactly

$$\overline{\eta_i^{\mathrm{cav}}(t)\eta_i^{\mathrm{cav}}(s) \mid \{x_j^{(i)}\}} = \frac{g^2}{N} \sum_{j\neq i} \phi(x_j^{(i)}(t))\phi(x_j^{(i)}(s)). \tag{12}$$

Self-averaging turns the random conditional covariance into $g^2 C_{\phi\phi}(t,s)$ in the large-N limit.

Restoring neuron i perturbs the cavity trajectories. For fully asymmetric independent couplings, these perturbations do not produce an order-one systematic feedback term in the limiting single-site equation. At the simplest reciprocal-loop level, the feedback contains sums proportional to $\sum_j J_{ij}J_{ji}$: their mean is zero and their standard deviation is $O(N^{-1/2})$. This illustrates the cancellation; controlling the full trajectory dependence requires the dynamical cavity or generating-functional argument [3]. Gaussian input is therefore a large-N self-consistent description, not a consequence of independence of the finite-network activities from their own couplings.

The dynamical mean-field equations

A representative neuron is described by

$$\boxed{\begin{aligned} \dot{x}(t) &= -x(t) + \eta(t), \\ \langle \eta(t) \rangle &= 0, \\ \langle \eta(t)\eta(s) \rangle &= g^2 \langle \phi(x(t))\phi(x(s)) \rangle. \end{aligned}} \tag{13}$$

The process η is Gaussian but generally temporally correlated. Its covariance is an unknown to be determined, not an externally prescribed white-noise strength. For a fixed $\mathbf{J}$ and initial state, the finite network is deterministic; the stochastic description represents the statistics of typical neurons and realizations.

This is a dynamical mean field rather than a replacement of the input by its mean. The mean input vanishes, but its order-one fluctuations remain. The nonlinear network has become a linear single-site filter coupled to a nonlinear self-consistency condition. For a prescribed Gaussian input, its solution is

$$x(t) = e^{-(t-t_0)}x(t_0) + \int_{t_0}^{t} e^{-(t-u)}\eta(u)\, du.$$

After the initial term has decayed, x is Gaussian as well. The output $\phi(x)$ is generally not Gaussian.

Before imposing stationarity, act with the two filter operators to obtain

$$(1 + \partial_t)(1 + \partial_s)C(t, s) = g^2 C_{\phi\phi}(t, s). \tag{14}$$

For a nonstationary problem one must also supply initial conditions and use the time-dependent variances $C(t, t)$ and $C(s, s)$ in the Gaussian closure. DMFT itself does not require stationarity.

Stationary correlations and Gaussian closure

Assume now a stationary state and write

$$\Delta = t - s, \qquad C(t, s) = C(\Delta), \qquad C_0 = C(0).$$

For a real stationary process $C(-\Delta) = C(\Delta)$ and $|C(\Delta)| \le C_0$. Since $\partial_t = \partial_\Delta$ and $\partial_s = -\partial_\Delta$ when acting on a stationary correlation,

$$\boxed{\left(1 - \frac{d^2}{d\Delta^2}\right) C(\Delta) = g^2 C_{\phi\phi}(\Delta).} \tag{15}$$

The minus sign in the second derivative comes from differentiating with respect to the two different times. If physical time is retained, it is $1 - \tau_r^2 d^2/d\Delta^2$ on the left.

To close this equation, let $X = x(t)$ and $Y = x(s)$ be jointly Gaussian with

$$\langle X^2 \rangle = \langle Y^2 \rangle = q, \qquad \langle XY \rangle = c, \qquad q = C_0,\ c = C(\Delta).$$

For $q > 0$, one convenient representation is

$$X = \sqrt{q - c^2/q}\, z_1 + \frac{c}{\sqrt{q}} z_2, \qquad Y = \sqrt{q}\, z_2, \tag{16}$$

where z_1, z_2 are independent standard normal variables. Define

$$\begin{aligned} F_\phi(q, c) &= \int \mathrm{D}z_1 \mathrm{D}z_2\, \phi\!\left(\sqrt{q - c^2/q}\, z_1 + \frac{c}{\sqrt{q}} z_2\right) \phi(\sqrt{q}\, z_2), \\ \mathrm{D}z &= \frac{e^{-z^2/2}}{\sqrt{2\pi}}\, dz. \end{aligned} \tag{17}$$

Thus $C_{\phi\phi}(\Delta) = F_\phi(C_0, C(\Delta))$, and

$$C''(\Delta) = C(\Delta) - g^2 F_\phi(C_0, C(\Delta)). \tag{18}$$

The case $q = 0$ is handled separately: $x = 0$ almost surely and $C \equiv 0$. For odd ϕ, useful checks are

$$F_\phi(q, 0) = 0, \qquad F_\phi(q, q) = \int \mathrm{D}z\, \phi(\sqrt{q}\, z)^2.$$

At zero covariance the Gaussian variables are independent, whereas at $c = q$ they are identical. These two limits are valuable checks on both algebra and numerical quadrature.

The correlation potential and Price's theorem

Equation 18 has a mechanical analogy. At fixed C_0, define

$$V(c; C_0) = -\frac{1}{2}c^2 + g^2 \int_0^c F_\phi(C_0, u)\, du. \tag{19}$$

Then

$$C'' = -\left.\frac{\partial V}{\partial c}\right|_{c=C}, \qquad \mathcal{E} = \frac{1}{2}(C')^2 + V(C; C_0) = \text{constant}. \tag{20}$$

We have normalized $V(0; C_0) = 0$. The variance C_0 is a parameter of this potential and also the initial position of the fictitious particle; it is not fixed in advance.

This construction must not be confused with the Hopfield energy. The coordinate is the correlation C, and the independent variable is time lag Δ, not physical trajectory time. The mechanical energy is conserved rather than monotonically decreasing. A conservative equation for correlations is entirely consistent with chaotic neuronal dynamics.

A useful Gaussian identity

Let $p_{q,c}(X, Y)$ be the centered bivariate Gaussian density with diagonal variances q and off-diagonal covariance c. For $|c| < q$ it obeys

$$\frac{\partial p_{q,c}}{\partial c} = \frac{\partial^2 p_{q,c}}{\partial X \, \partial Y}.$$

Two integrations by parts therefore give Price's theorem:

$$\boxed{\frac{\partial}{\partial c}\langle A(X)B(Y)\rangle_{q,c} = \langle A'(X)B'(Y)\rangle_{q,c}.} \tag{21}$$

The variances q are held fixed. Gaussian decay eliminates the boundary terms for the smooth functions used here; singular endpoint covariances are understood by continuity.

Choose a scalar primitive $\Phi'(x) = \phi(x)$ and define the distinct bivariate average

$$\mathcal{P}(q, c) = \langle \Phi(X)\Phi(Y)\rangle_{q,c}.$$

Then $\partial_c \mathcal{P} = F_\phi$, and the potential is

$$V(c; C_0) = -\frac{1}{2}c^2 + g^2[\mathcal{P}(C_0, c) - \mathcal{P}(C_0, 0)]. \tag{22}$$

For $\phi(x) = \tanh x$, take $\Phi(x) = \log\cosh x$. Keeping the scalar primitive Φ separate from the two-variable function $\mathcal{P}$ avoids confusing which derivative appears in Price's theorem.

Selecting the fluctuating state: boundary conditions

An arbitrary solution of Equation 18 need not be a physically admissible autocorrelation. In particular, a covariance must be positive semidefinite for every set of times. We seek the stationary mixing branch of the zero-input model, with vanishing long-lag correlation:

$$\boxed{C(0) = C_0, \quad C'(0) = 0, \quad C(\Delta) \to 0, \quad C'(\Delta) \to 0 \quad (\Delta \to \infty).} \tag{23}$$

The zero slope at the origin follows from stationarity and differentiability: $\langle x\dot{x}\rangle = \frac{1}{2}d\langle x^2\rangle/dt = 0$. An externally white-noise-driven model can have a cusp instead, but there is no such forcing in the present model.

Energy conservation requires the same potential at the two endpoints:

$$V(C_0; C_0) = V(0; C_0) = 0. \tag{24}$$

At $c = C_0$, $X = Y$; at $c = 0$, X and Y are independent. Consequently, for the tanh network the unknown variance is selected by a one-dimensional Gaussian integral equation:

$$\boxed{\frac{C_0^2}{2} = g^2 \left\{ \int \mathrm{D}z \, [\log\cosh(\sqrt{C_0}z)]^2 - \left[\int \mathrm{D}z \, \log\cosh(\sqrt{C_0}z) \right]^2 \right\}.} \tag{25}$$

The trivial root $C_0 = 0$ is always present. For $g > 1$, the positive root selects the fluctuating branch. Its correlation starts at rest at C_0 and approaches the unstable mechanical equilibrium $c = 0$ at infinite lag. Along the monotone branch,

$$C'(\Delta) = -\sqrt{-2V(C(\Delta); C_0)}. \tag{26}$$

The potential is negative between the two endpoints. This is a separatrix of the fictitious mechanics, not a periodic orbit of the neuron.

Why the equal-time equation alone is not enough

Evaluating Equation 18 at $\Delta = 0$ gives

$$C''(0) = C_0 - g^2 \int \mathrm{D}z \, \phi(\sqrt{C_0}z)^2.$$

For a differentiable stationary process, $C''(0) = -\langle \dot{x}^2 \rangle$. Thus

$$g^2\langle \phi(x)^2 \rangle = C_0 + \langle \dot{x}^2 \rangle. \tag{27}$$

Setting $C_0 = g^2\langle \phi(x)^2 \rangle$ would silently impose zero velocity variance, describing a static rather than a fluctuating state. The stationary fixed-point ansatz has $C(\Delta) = q$ at all lags and obeys $q = g^2 F_\phi(q, q)$; existence of a solution of that equation does not establish its dynamical stability. The chaotic branch instead uses the boundary conditions in Equation 23.

Near the transition: amplitude and correlation time

We can see how the nonzero solution emerges without solving the full Gaussian integrals. For small variance, use $\tanh x = x - x^3/3 + 2x^5/15 + \cdots$. Gaussian moment identities give

$$F_\phi(q, c) = (1 - 2q + 5q^2)c + \frac{2}{3}c^3 + O(q^4), \qquad |c| \le q. \tag{28}$$

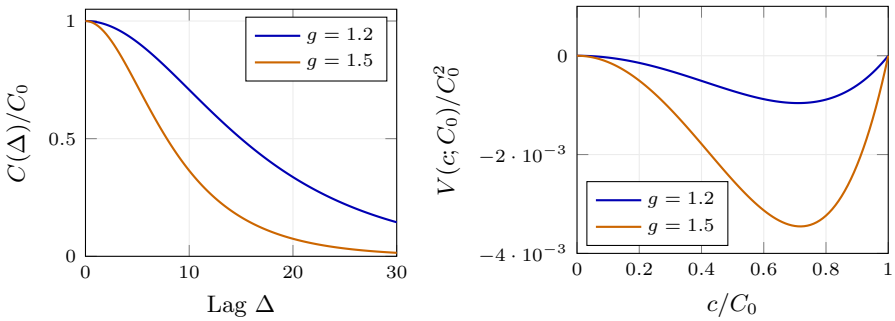


Figure 2: Numerical solutions of the stationary DMFT equations for the tanh model. Left: normalized current autocorrelation. Right: the corresponding correlation potential, with both endpoint energies zero. Gaussian quadrature and Equation 25 give $C_0 \simeq 0.24263$ at $g = 1.2$ and $C_0 \simeq 0.74769$ at $g = 1.5$. Time lag is in units of τ_r. These are mean-field calculations, not finite-network simulations.

For example, $\langle X^3 Y \rangle = 3qc$ and $\langle X^3 Y^3 \rangle = 9q^2 c + 6c^3$. These identities show explicitly where the linear-in-c and cubic-in-c terms come from.

Substitution into the endpoint-energy condition yields

$$1 = g^2 \left(1 - 2C_0 + \frac{16}{3} C_0^2 + O(C_0^3) \right).$$

Hence, with $g = 1 + \epsilon$ and $0 < \epsilon \ll 1$,

$$C_0 = \epsilon + O(\epsilon^2). \tag{29}$$

To leading order the correlation equation becomes

$$C'' = \kappa^2 C - \frac{2g^2}{3} C^3, \qquad \kappa^2 \simeq \frac{g^2 C_0^2}{3},$$

with separatrix solution

$$C(\Delta) \simeq C_0 \operatorname{sech}(\kappa \Delta), \qquad \kappa \simeq \frac{g-1}{\sqrt{3}}. \tag{30}$$

The activity variance vanishes continuously at the transition, while the correlation time grows as $(g-1)^{-1}$. In physical units the long-lag exponential decay time is asymptotically $\sqrt{3}\,\tau_r/(g-1)$. This approximation is for the weakly chaotic state near $g = 1$, not for arbitrary gain.

Is the fluctuating state chaotic?

A decaying autocorrelation does not by itself prove chaos. The relevant test is sensitivity to infinitesimal changes of initial conditions while holding the couplings fixed. A tangent perturbation obeys

$$\delta\dot{x}_i = -\delta x_i + \sum_j J_{ij}\phi'(x_j(t))\delta x_j. \tag{31}$$

Its largest Lyapunov exponent is

$$\lambda_{\mathrm{L}} = \lim_{t\to\infty} \frac{1}{t} \log \frac{\|\delta\mathbf{x}(t)\|}{\|\delta\mathbf{x}(0)\|}, \tag{32}$$

understood in the tangent dynamics, with repeated normalization in numerical calculations. A positive value diagnoses exponential sensitivity. The instantaneous Jacobian is $-\mathbf{I} + \mathbf{J}\,\mathrm{Diag}[\phi'(x_j(t))]$; one cannot infer its long-time perturbation growth simply by applying the circular law to $\mathbf{J}$ again.

Optional derivation: a second eigenvalue problem

For completeness, the tangent DMFT reduces the Lyapunov calculation to a one-dimensional spectral problem [1, 3]. Let $K(t, s)$ be the population-averaged covariance of a small, initially random tangent perturbation. To leading order,

$$(1 + \partial_t)(1 + \partial_s)K(t, s) = g^2 A(t - s)K(t, s), \tag{33}$$

where

$$A(\Delta) = \langle \phi'(x(t))\phi'(x(t - \Delta))\rangle = \partial_c F_\phi(C_0, c)\big|_{c=C(\Delta)}.$$

The last equality follows from Price's theorem. Equation 33 comes from linearizing the two-copy self-consistency problem; it should not be interpreted as an exact finite-N factorization of tangent and activity variables.

Set $T = (t + s)/2$ and $\Delta = t - s$, and seek $K(t, s) = e^{2\lambda T}\psi(\Delta)$. The factor of two reflects growth of a squared perturbation. Since

$$(1 + \partial_t)(1 + \partial_s) = \left(1 + \tfrac{1}{2}\partial_T\right)^2 - \partial_\Delta^2,$$

we obtain

$$\mathcal{H}\psi = E\psi, \qquad \mathcal{H} = -\frac{d^2}{d\Delta^2} + 1 - g^2 A(\Delta), \qquad E = 1 - (1 + \lambda)^2. \tag{34}$$

On the full lag axis, the lowest eigenvalue E_0 therefore determines

$$\boxed{\lambda_{\mathrm{L}} = -1 + \sqrt{1 - E_0}.} \tag{35}$$

As a check, at the zero fixed point $A = 1$ and the bottom of the spectrum is $E_0 = 1 - g^2$, giving $\lambda_{\mathrm{L}} = g - 1$, consistent with the circular-law stability calculation.

For the nontrivial, even, monotonically decaying correlation, differentiating Equation 18 shows that

$$\mathcal{H}C'(\Delta) = 0.$$

This normalizable odd eigenfunction has a node at zero. The ground state of the one-dimensional Schrödinger operator has no nodes and lies below it, so $E_0 < 0$ and $\lambda_{\mathrm{L}} > 0$. Within the large-N tangent DMFT, this establishes that the selected fluctuating branch is chaotic. The fictitious correlation potential and the tangent spectral problem thus answer different questions: the first selects a covariance, and the second tests sensitivity of trajectories.

How to solve and test the theory numerically

The mean-field calculation and the network simulation provide independent checks on one another.

1. **Solve for the variance.** Choose $g > 1$ and evaluate the Gaussian integrals in Equation 25 and find its positive root. The root at zero must be excluded deliberately. Gauss–Hermite quadrature is convenient.

2. **Compute the correlation.** Evaluate $F_\phi(C_0, c)$ and integrate Equation 18 from $C(0) = C_0$, $C'(0) = 0$. Check conservation of $\mathcal{E}$ and the long-lag limit. The separatrix is numerically sensitive at large lag: small errors in C_0 can cause a turn or a spurious zero crossing. Energy quadrature or a boundary-value solver can be more robust than unrestricted forward integration.

3. **Simulate one quenched network.** Draw $J_{ij} = gZ_{ij}/\sqrt{N}$ with independent standard normal Z_{ij}, initialize x_i with small nonzero random values, and integrate Equation 8. Keep $\mathbf{J}$ fixed. Discard the transient and verify results by decreasing the time step and increasing N and the recording duration.

4. **Measure the same observable.** Estimate

$$\widehat{C}(\Delta) = \frac{1}{NT_{\mathrm{obs}}} \sum_i \int_{t_b}^{t_b+T_{\mathrm{obs}}} x_i(t+\Delta)x_i(t)\,dt,$$

   using a recorded interval long enough to include both times. Also measure $C_{\phi\phi}$ separately: it is not interchangeable with C. Near the transition, allow for the long correlation time.

5. **Test chaos.** Integrate the tangent equation alongside the trajectory, periodically normalize its vector, and accumulate logarithmic growth factors. Two full trajectories eventually separate by a finite amount and saturate; that saturation is not the infinitesimal Lyapunov rate.

Scope of the closure and possible extensions

The assumptions of this calculation identify its domain of validity, not a list of conditions required by every version of DMFT.

- **Dense, weak, independent couplings.** The $N^{-1/2}$ scaling and absence of a dominant input support a Gaussian field. Sparse networks of fixed degree, heavy-tailed weights, or a few dominant modes may require a different limit or additional order parameters.
- **Full asymmetry.** Correlated reciprocal couplings generate a systematic retarded feedback involving a response function. Symmetric and partially symmetric networks can have mean-field theories, but not generally the noise-only single-site equation used here.
- **Self-averaging and finite size.** The equations describe typical large networks and specified initial ensembles. Near criticality, finite-size fluctuations and long transients can be substantial. The large-N and long-time limits must be interpreted with care.
- **Stationarity and mixing.** Stationarity reduces a two-time problem to one lag. Mixing with zero persistent mean gives the additional condition $C(\infty) = 0$. Nonstationary states require the two-time equations; states with a nonzero correlation plateau require different boundary conditions. Neither is excluded from DMFT in principle.
- **Biological structure.** Independent zero-mean Gaussian weights do not obey Dale's law. Excitatory–inhibitory populations, nonzero means, low-rank structure, and external drive call for an appropriately extended mean-field description. Zero mean alone is not a complete model of excitatory–inhibitory balance.

A prescribed Gaussian external input with covariance $D(t, s)$, independent of the couplings, adds its covariance to the effective drive. In a stationary zero-mean setting the correlation equation becomes

$$(1 - \partial_\Delta^2)C(\Delta) = g^2 C_{\phi\phi}(\Delta) + D(\Delta).$$

The autonomous boundary conditions and critical behavior derived above must then be reconsidered. A deterministic or nonzero-mean input also requires keeping track of the mean and of how the nonlinearity is shifted.

Questions and exercises

1. Derive Equation 4, including its initial-condition term. How does a complex eigenvalue affect the real-valued trajectory? Why is $|\lambda_\mu| < 1$ not the correct continuous-time stability criterion?

2. For the non-normal two-neuron example, find the maximum of $x_1(t)$ and the initial conditions for transient growth of the Euclidean norm. Compare with asymptotic stability.

3. Draw a Gaussian matrix numerically and verify the disk radius as g changes. Compare with a symmetric Gaussian ensemble having the same leading off-diagonal variance. Why should their thresholds differ?

4. Verify the Gaussian representation in Equation 16. Derive Price's identity by differentiating the density and integrating by parts, keeping both variances fixed.

5. Derive the variance equation from endpoint-energy equality. Explain why replacing it by $C_0 = g^2\langle \tanh^2 x\rangle$ gives a different state.

6. Use Wick's theorem to obtain Equation 28, then derive the leading variance and correlation-time scalings near $g = 1$.

7. At $g = 1.5$, reproduce the positive variance root and the correlation in Figure 2. Compare them with simulations at increasing N.

8. *Optional.* Differentiate the correlation equation to show $\mathcal{H}C' = 0$. Explain why the node of C' implies a positive maximal Lyapunov exponent for the decaying mean-field solution.

References

[1] H. Sompolinsky, A. Crisanti, and H. J. Sommers, *Chaos in Random Neural Networks*, Physical Review Letters **61**, 259–262 (1988). doi:10.1103/PhysRevLett.61.259.

[2] T. Tao and V. Vu, with an appendix by M. Krishnapur, *Random matrices: Universality of ESDs and the circular law*, Annals of Probability **38**, 2023–2065 (2010). arXiv:0807.4898.

[3] A. Crisanti and H. Sompolinsky, *Path integral approach to random neural networks*, Physical Review E **98**, 062120 (2018). arXiv:1809.06042.