Slow Transition to Low-Dimensional Chaos in Heavy-Tailed Recurrent Neural Networks

1 ALLEN INSTITUTE
2 PRINCETON Neuroscience INSTITUTE

Eva Yi Xie[1,2,*], Stefan Mihalas[1], Lukasz Kusmierz[1,*]

*evayixie@princeton.edu, lukasz.kusmierz@alleninstitute.org

Full paper

Code

Model Setup

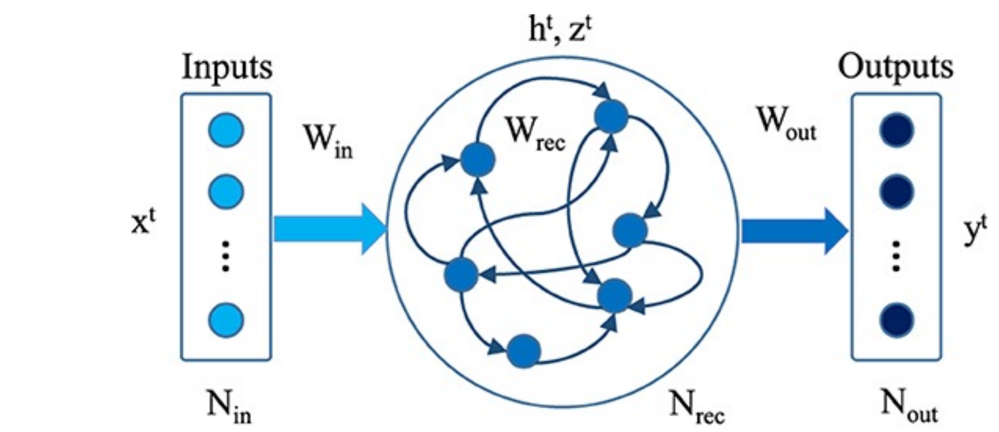

Schematic of a Recurrent Neural Network (RNN) as a model for brain network dynamics, adopted from Rostami et al., 2022.

$$y_t = W_{out} h_t,$$
$$h_t = \phi(z_t),$$
$$z_t = W_{rec} h_{t-1} + W_{in} x_t$$

where $(W_{\rm rec})_{ij} \overset{\rm i.i.d.}{\sim} L_\alpha(\sigma)$, with $\sigma = \frac{g_{\rm rec}}{N^{1/\alpha}}$

Characteristic function $\phi_{L_\alpha(\sigma)}(k) = \exp(-|\sigma k|^\alpha)$

- Smaller the **α**, heavier the tail.
- Gaussian is a special case when **α** = 2.

RNNs as common simplified models for brain dynamics

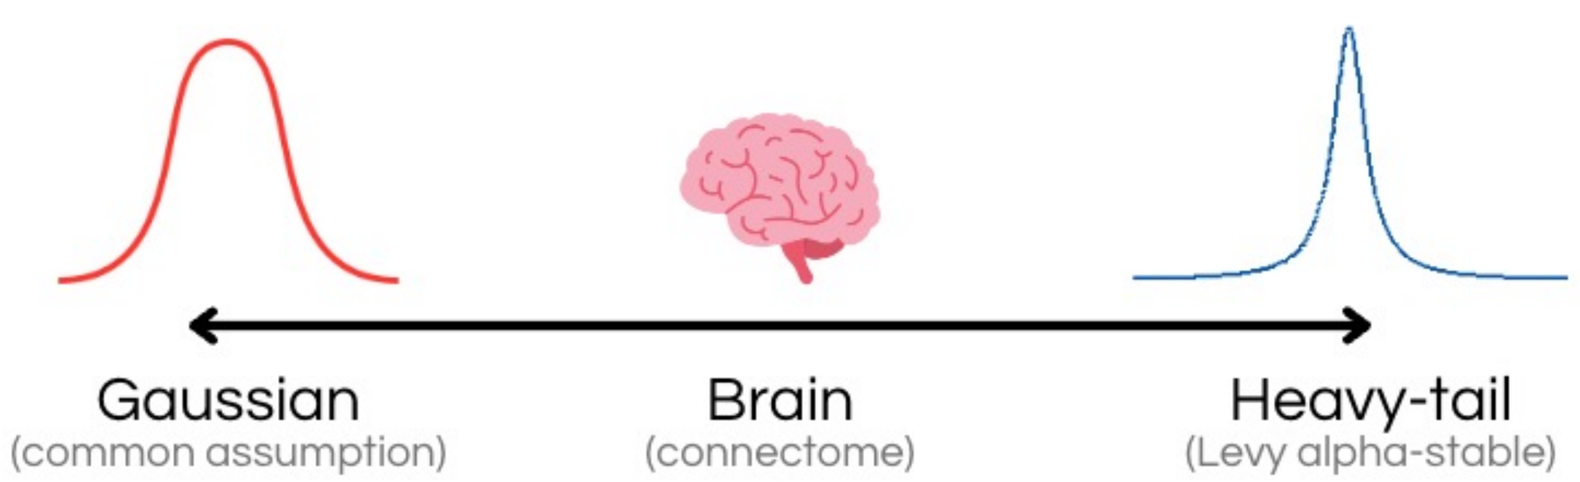

What Structural Principles Govern Neural Computation?

Most theory *assumes* a brain network of **Gaussian** weight distribution,
BUT the actual brain synaptic weights follow a more **heavy-tailed** distribution!

We show heavy-tailed weights can **strongly affect** dynamics in **finite-size** networks, revealing a **biologically aligned tradeoff** between robustness ↔️ attractor dimension.

1

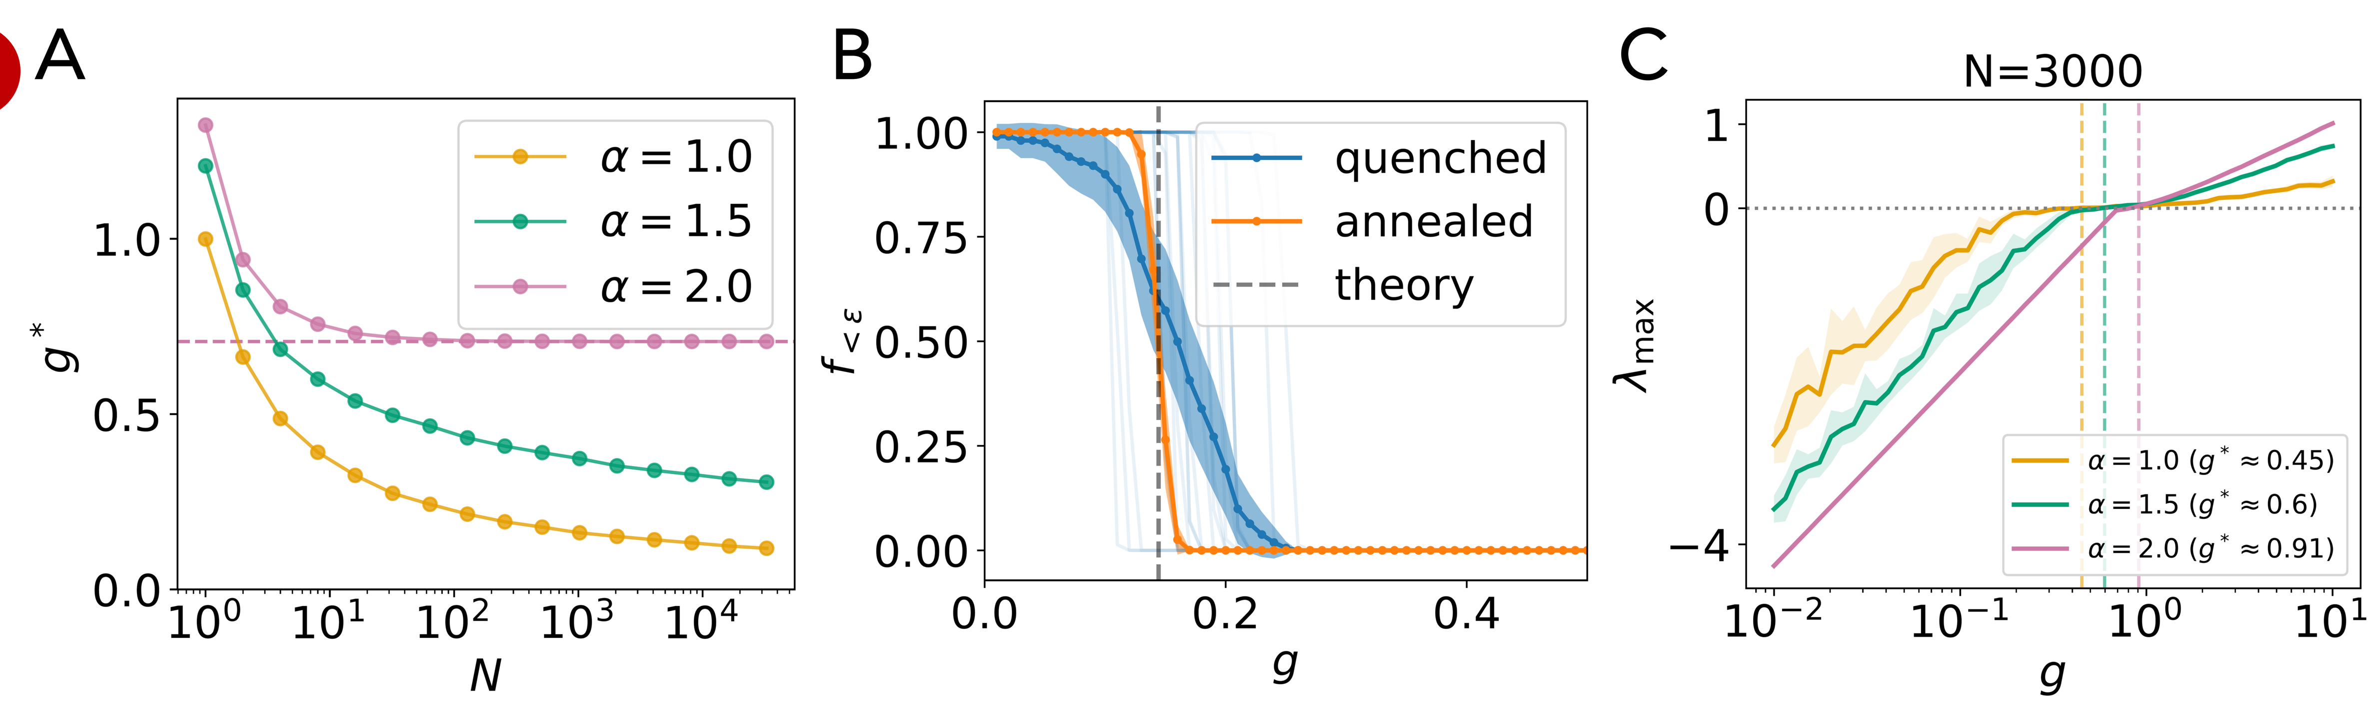

- Finite heavy-tailed networks exhibit a predictable, **slow** 🐢, quiescent-to-chaotic transition, **in contrary** to the mean-field prediction of ubiquitous chaos.
- We theoretically predict the gain (g) at which the transition happens [Eqn. 7 in paper].

2

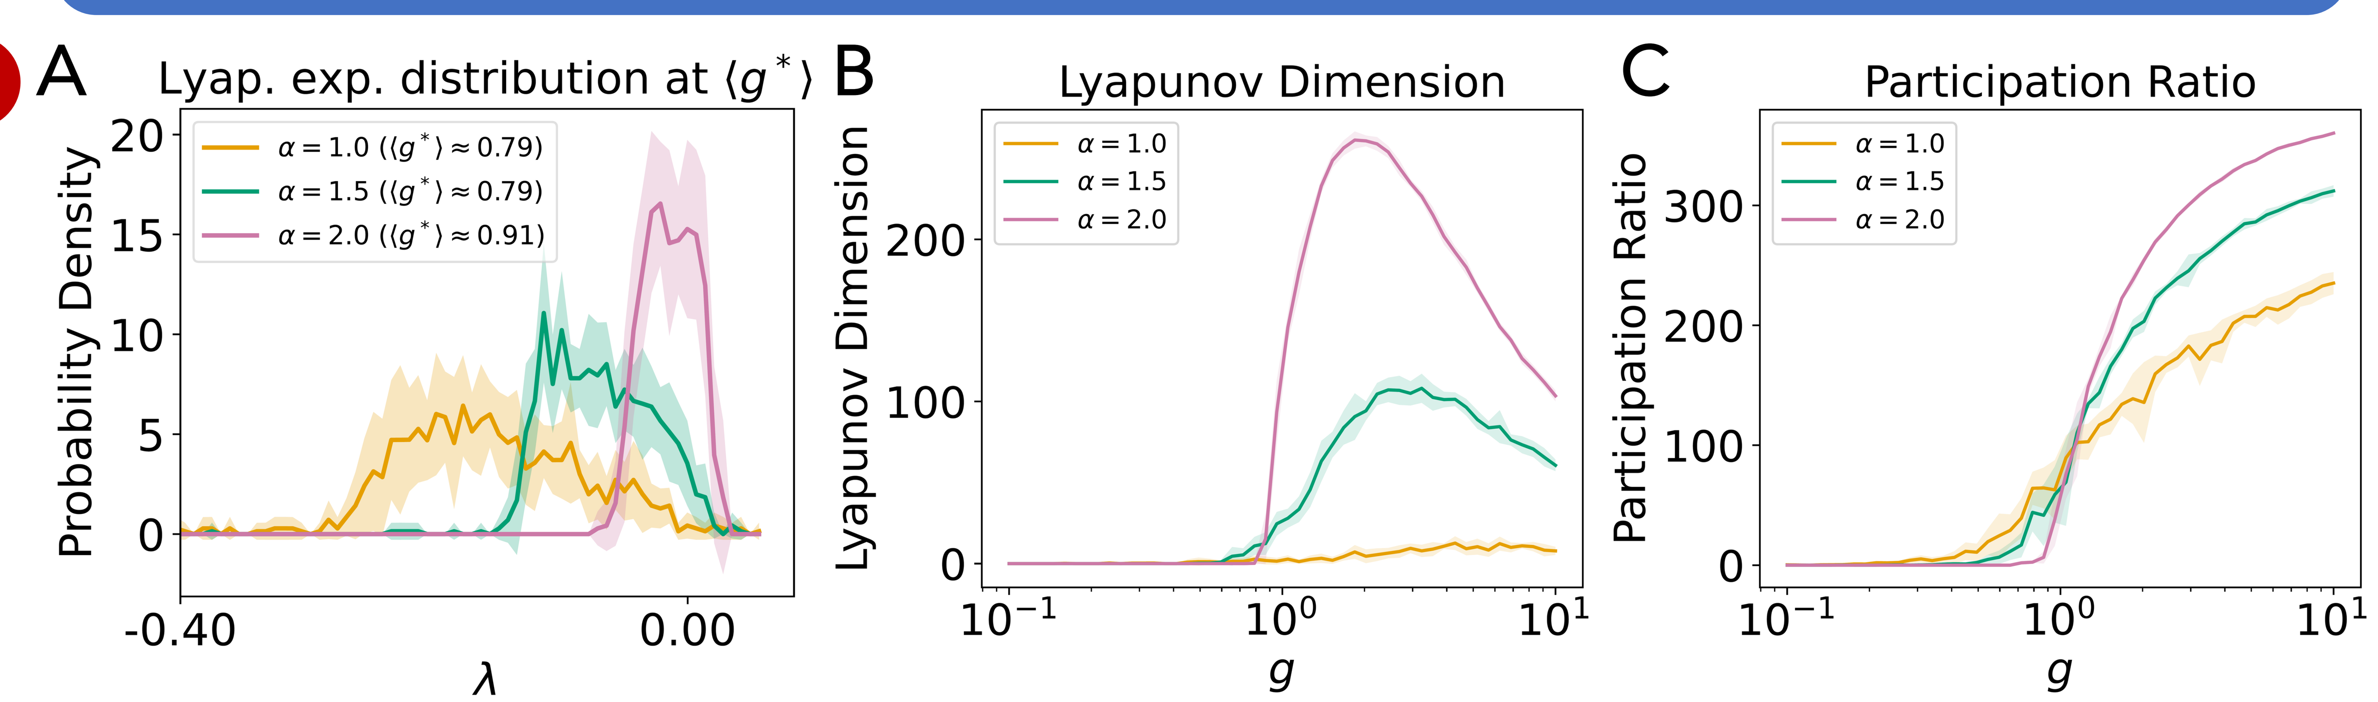

However, the robustness from slow transition (improved info-processing in simple tasks) comes with **low attractor dimensionality**, commonly observed in neural data.

Metrics

Quantify system dynamical stability

Maximum Lyapunov Exponent (λ_{max})

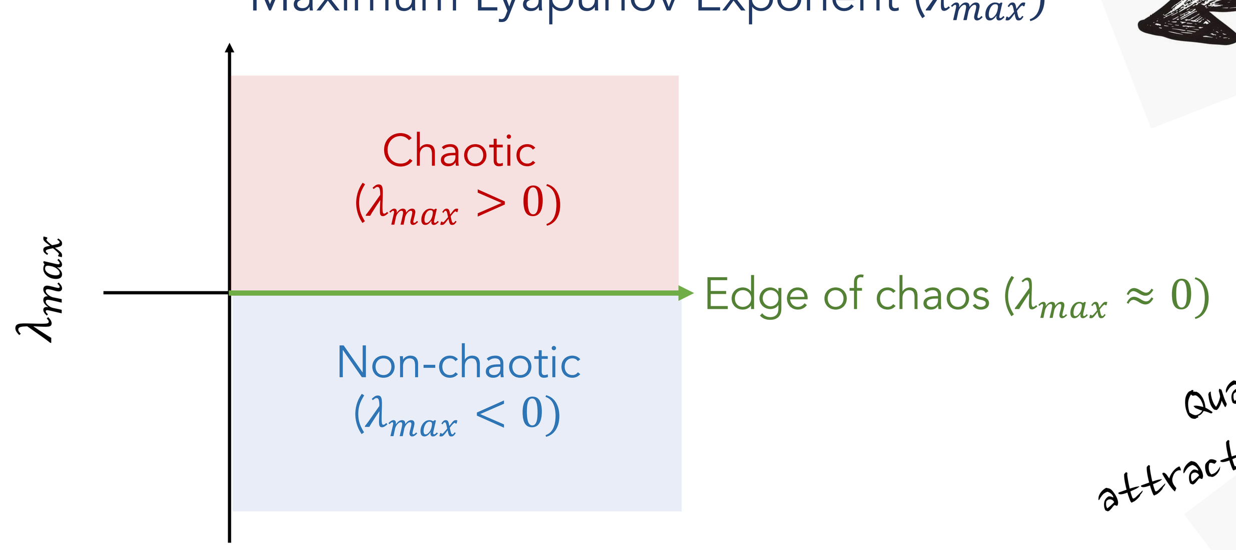

Quantify attractor dimension

Lyapunov Dimension

$$D_{KY} := k + \frac{\sum_{i=1}^{k} \lambda_i}{|\lambda_{k+1}|}$$

- $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ are ordered Lyapunov exponents.
- Define k as the largest index s.t. $\sum_{i=1}^{k} \lambda_i \geq 0$.
- Higher D_{KY} means more modes contribute to long-term variability ➔ greater capacity to support rich, temporally extended computations.

Participation Ratio

$$\mathrm{PR} := \frac{\left(\sum_i \tilde{\lambda}_i\right)^2}{\sum_i \tilde{\lambda}_i^2}$$

- Let $x^{(t)} \in \mathrm{R}^N$ be the hidden state of the RNN at time t.
- Empirical covariance matrix
$$S = \frac{1}{T-1}\sum_{t=1}^{T}\left(x^{(t)} - \bar{x}\right)\left(x^{(t)} - \bar{x}\right)^\top, \bar{x} = \frac{1}{T}\sum_{t=1}^{T} x^{(t)}.$$
- $\tilde{\lambda}$ denote the eigenvalues of S.
- 1 (all variance in one mode) to N (uniform variance).

More on Theoretical Results

- Any α–stable network of finite size N exhibits a predictable transition location when
$$g^* = \exp(-\langle \Xi_{N,\alpha} \rangle)$$
$$\Xi_{N,\alpha} = \frac{1}{\alpha}\ln\left(\frac{1}{N}\sum_{j=1}^{N}|z_j|^\alpha\right), \text{ where } z_j \sim L_\alpha(1)$$
- ✅ **Finite-size correction:** our theory generalizes mean-field result to finite-size networks, with g^* decreases with network size N.
- Decay in heavy-tailed networks follows
$$g^* \propto 1/(\ln N)^{1/\alpha}.$$

A simple case where a broader edge-of-chaos regime translates to benefits in information processing.

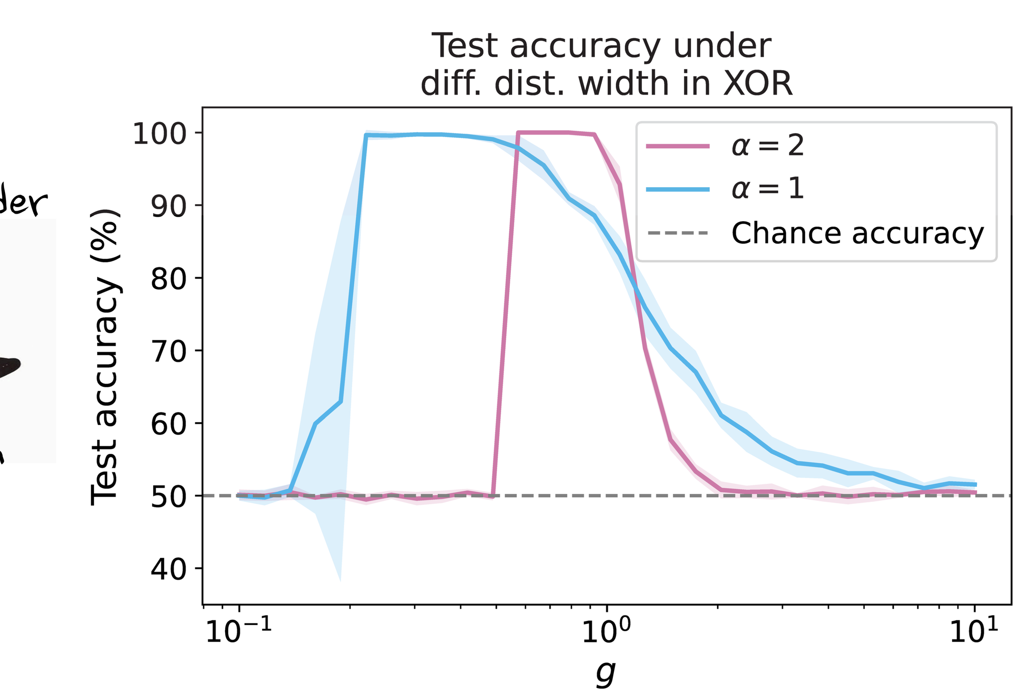

Conclusion

🧠 For Neuroscience, we show
- A biologically-aligned tradeoff between **robustness** + **richness** of neural dynamics;
- A **tractable framework** for finite-size, heavy-tailed neural circuits.

🤖 For Machine Learning, we show
- Robustness benefits of brain-like **heavy-tailed weights;**
- A theoretical perspective on dynamical consequences of heavy-tailed structures that can naturally arise from brain-like learning. [Cornford et al., 2024]

Bluesky/X: @EvaYiXie