Statistical mechanics for networks of real neurons

William Bialek

Joseph Henry Laboratories of Physics, and
Lewis-Sigler Institute for Integrative Genomics
Princeton University

Initiative for the Theoretical Sciences
The Graduate Center, City University of New York

University of Virginia **23 March 2018**

Interesting functions result from coordinated activity among large numbers of neurons

These behaviors are "emergent"

Neural networks and physical systems with emergent collective computational abilities

(associative memory/parallel processing/categorization/content-addressable memory/fail-soft devices)

J. J. HOPFIELD

Emergent phenomena are all around us, even in equilibrium systems.

These phenomena are captured in the language of statistical mechanics.

The first step of equilibrium statistical mechanics is to write the probability distribution over "microscopic" states of the system.

Often we use models which are much simpler than the microscopic reality.

Thanks to the renormalization group, we understand why this works.

Can we write down the (joint!) probability distribution for the activity of many neurons in a network?

(for simplicity, let's focus on one moment in time)

Is there any reason to think that this distribution is simpler than it could be?

$$N \text{ neurons} \Rightarrow 2^N \text{ states}$$

$$N = 10 \qquad 2^N \sim 1000$$

$$N = 20 \qquad 2^N \sim 10^6$$

$$N = 100 \qquad 2^N \sim 10^{30}$$

In principle, every state has a different probability, and there doesn't need to be any pattern. If that's true, we're sunk.

This problem is different because now we can observe the activity of many neurons simultaneously.

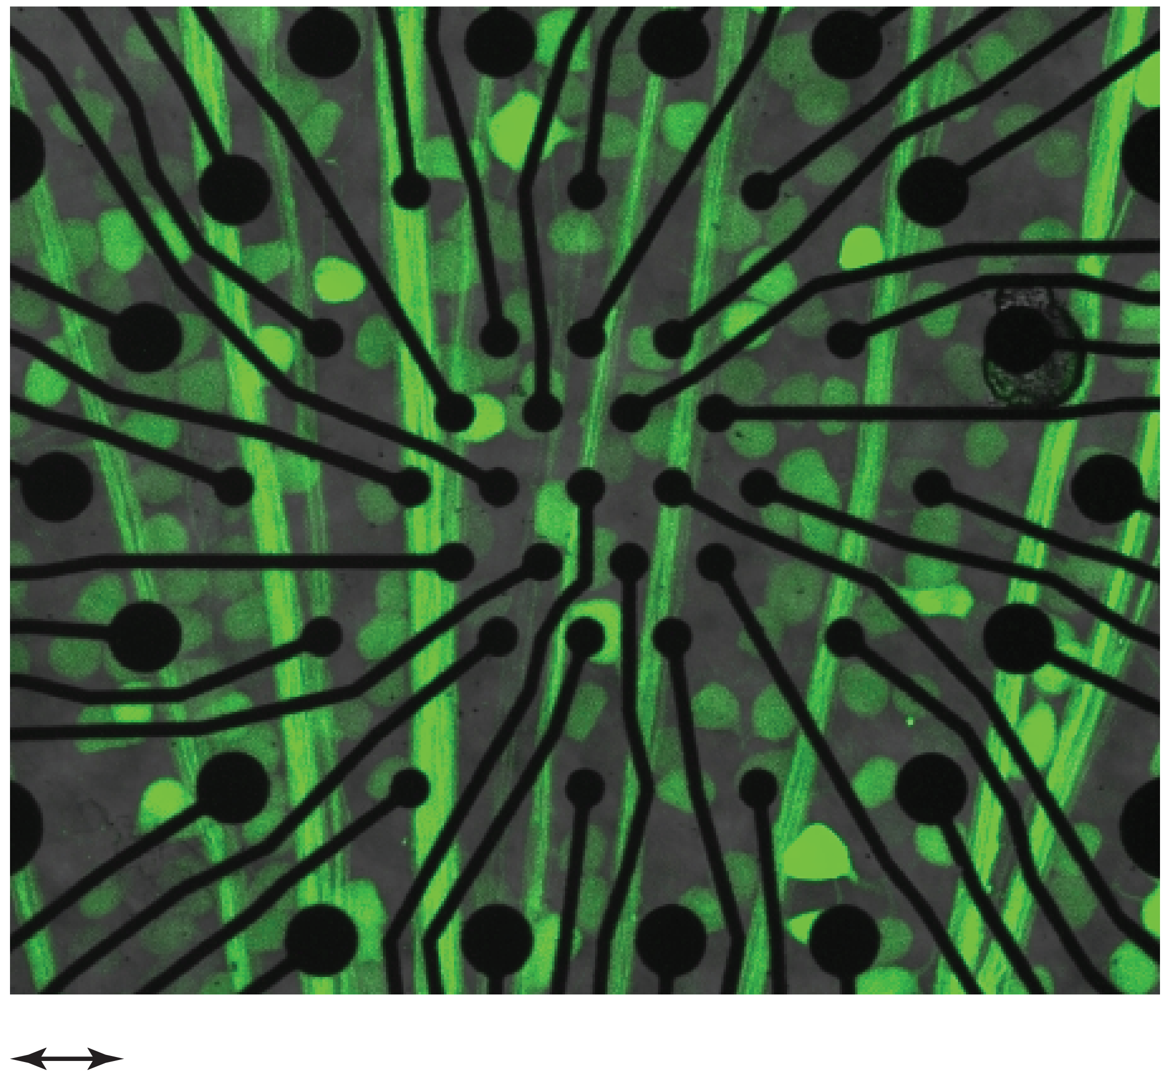

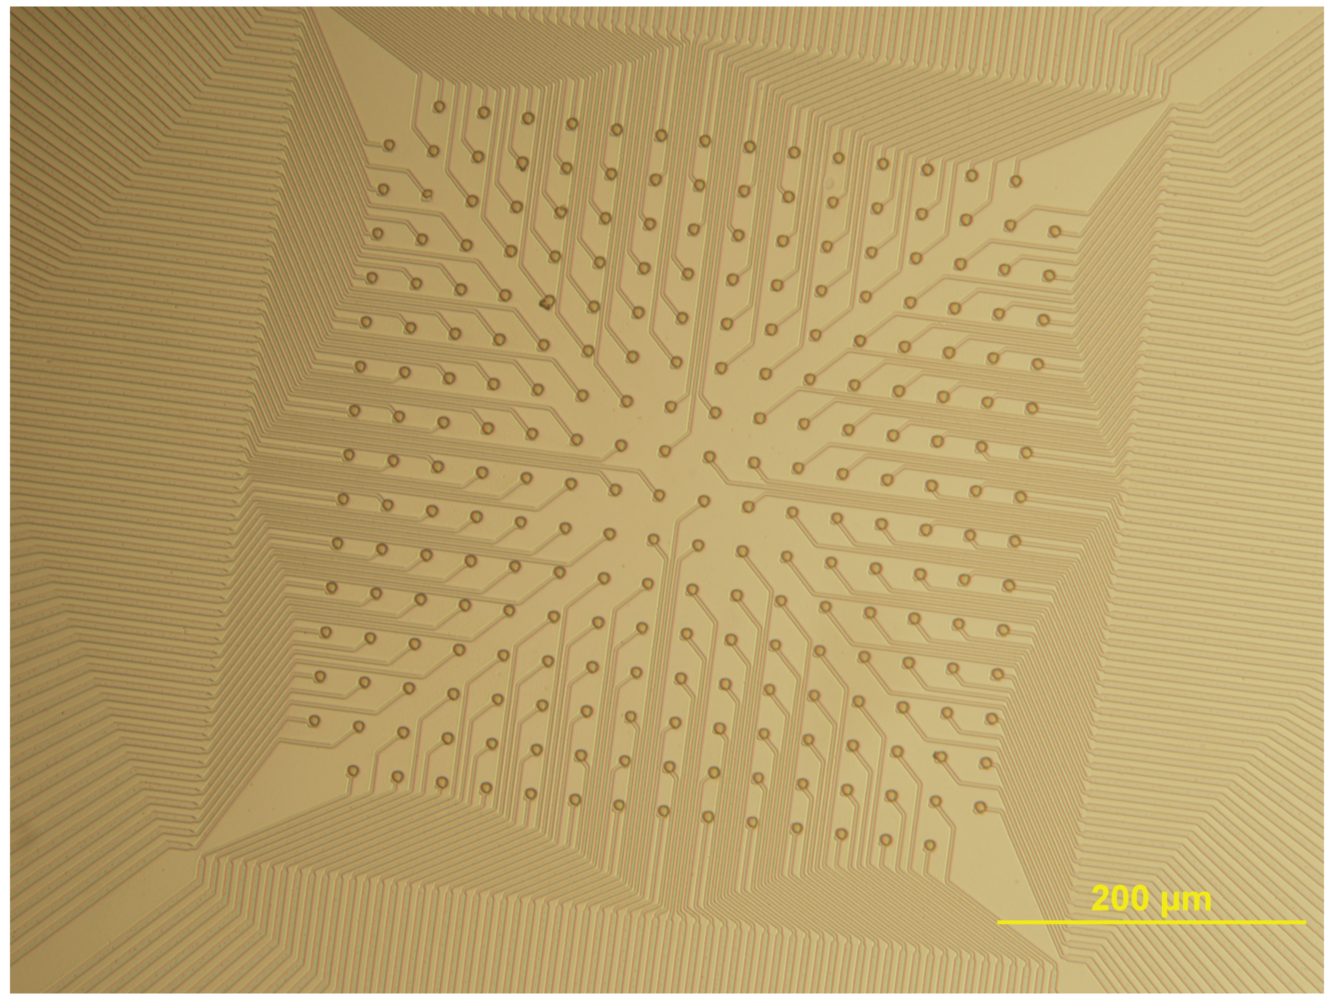

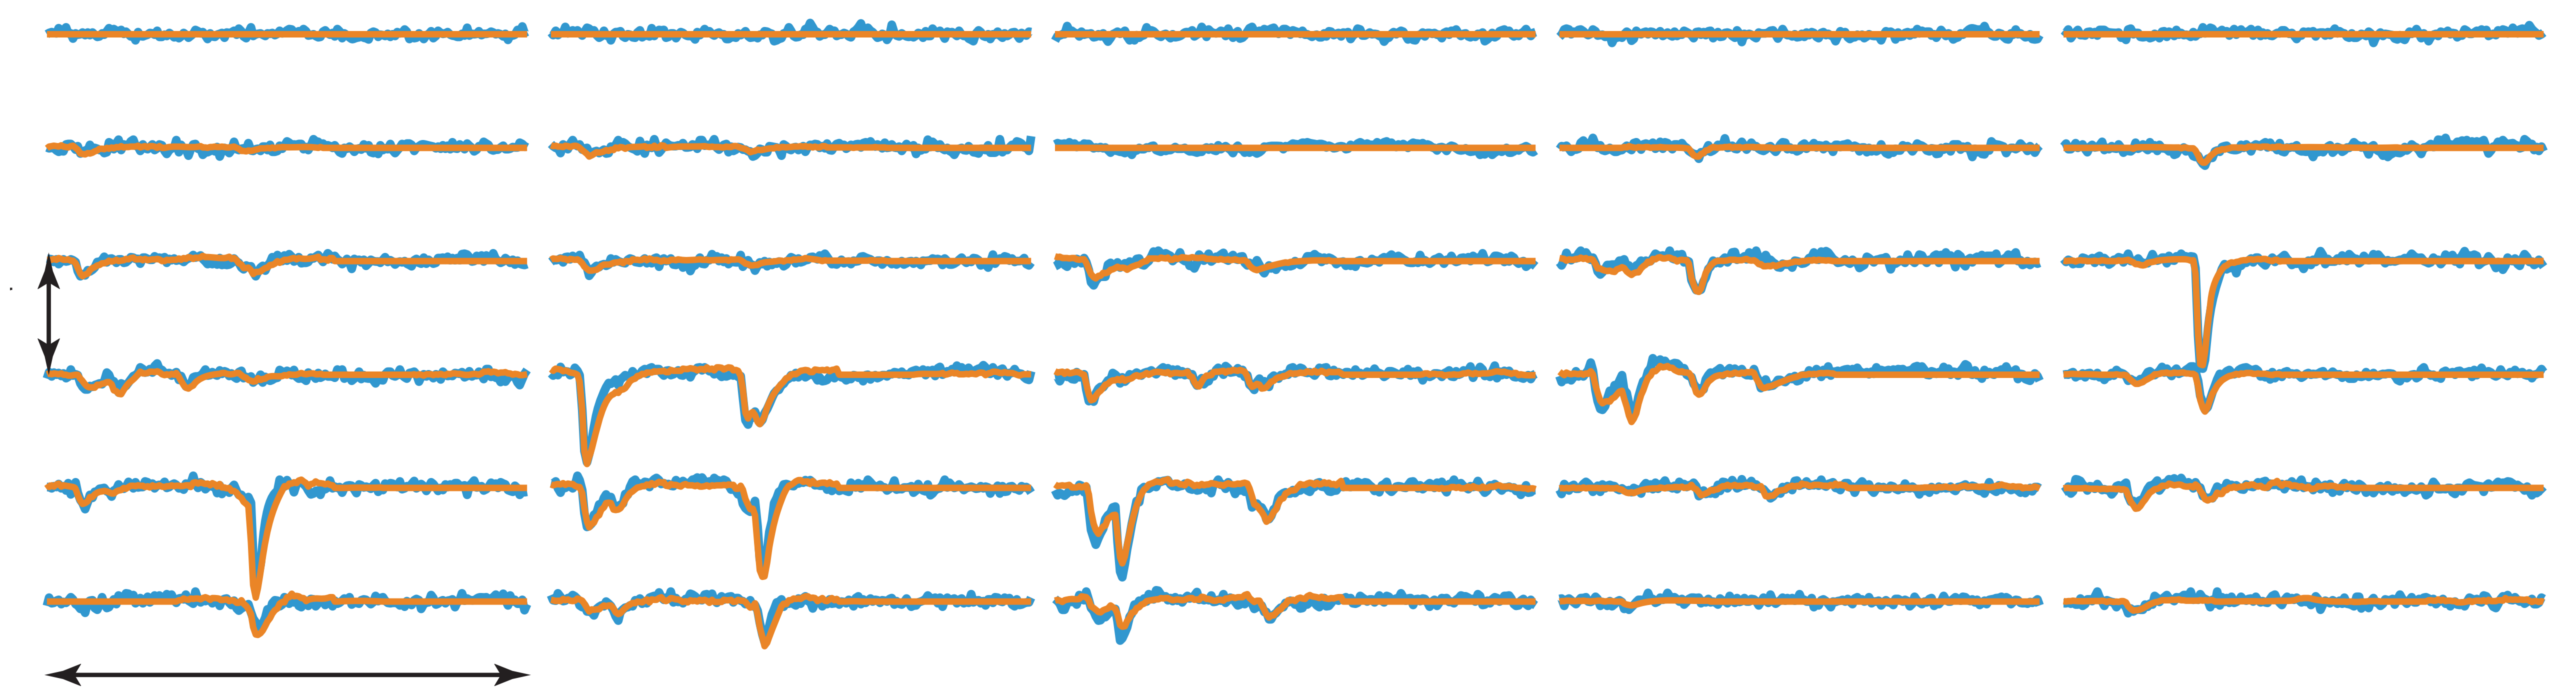

Using arrays of electrodes to record from 100+ neurons in the retina.

R Segev, J Goodhouse, JL Puchalla, and MJ Berry II, *Nat Neurosci* 7:1155 (2004).
O Marre, D Amodei, N Deshmukh, K Sadeghu, F Soo, TE Holy, and MJ Berry II *J Neurosci* 32:14859 (2012).

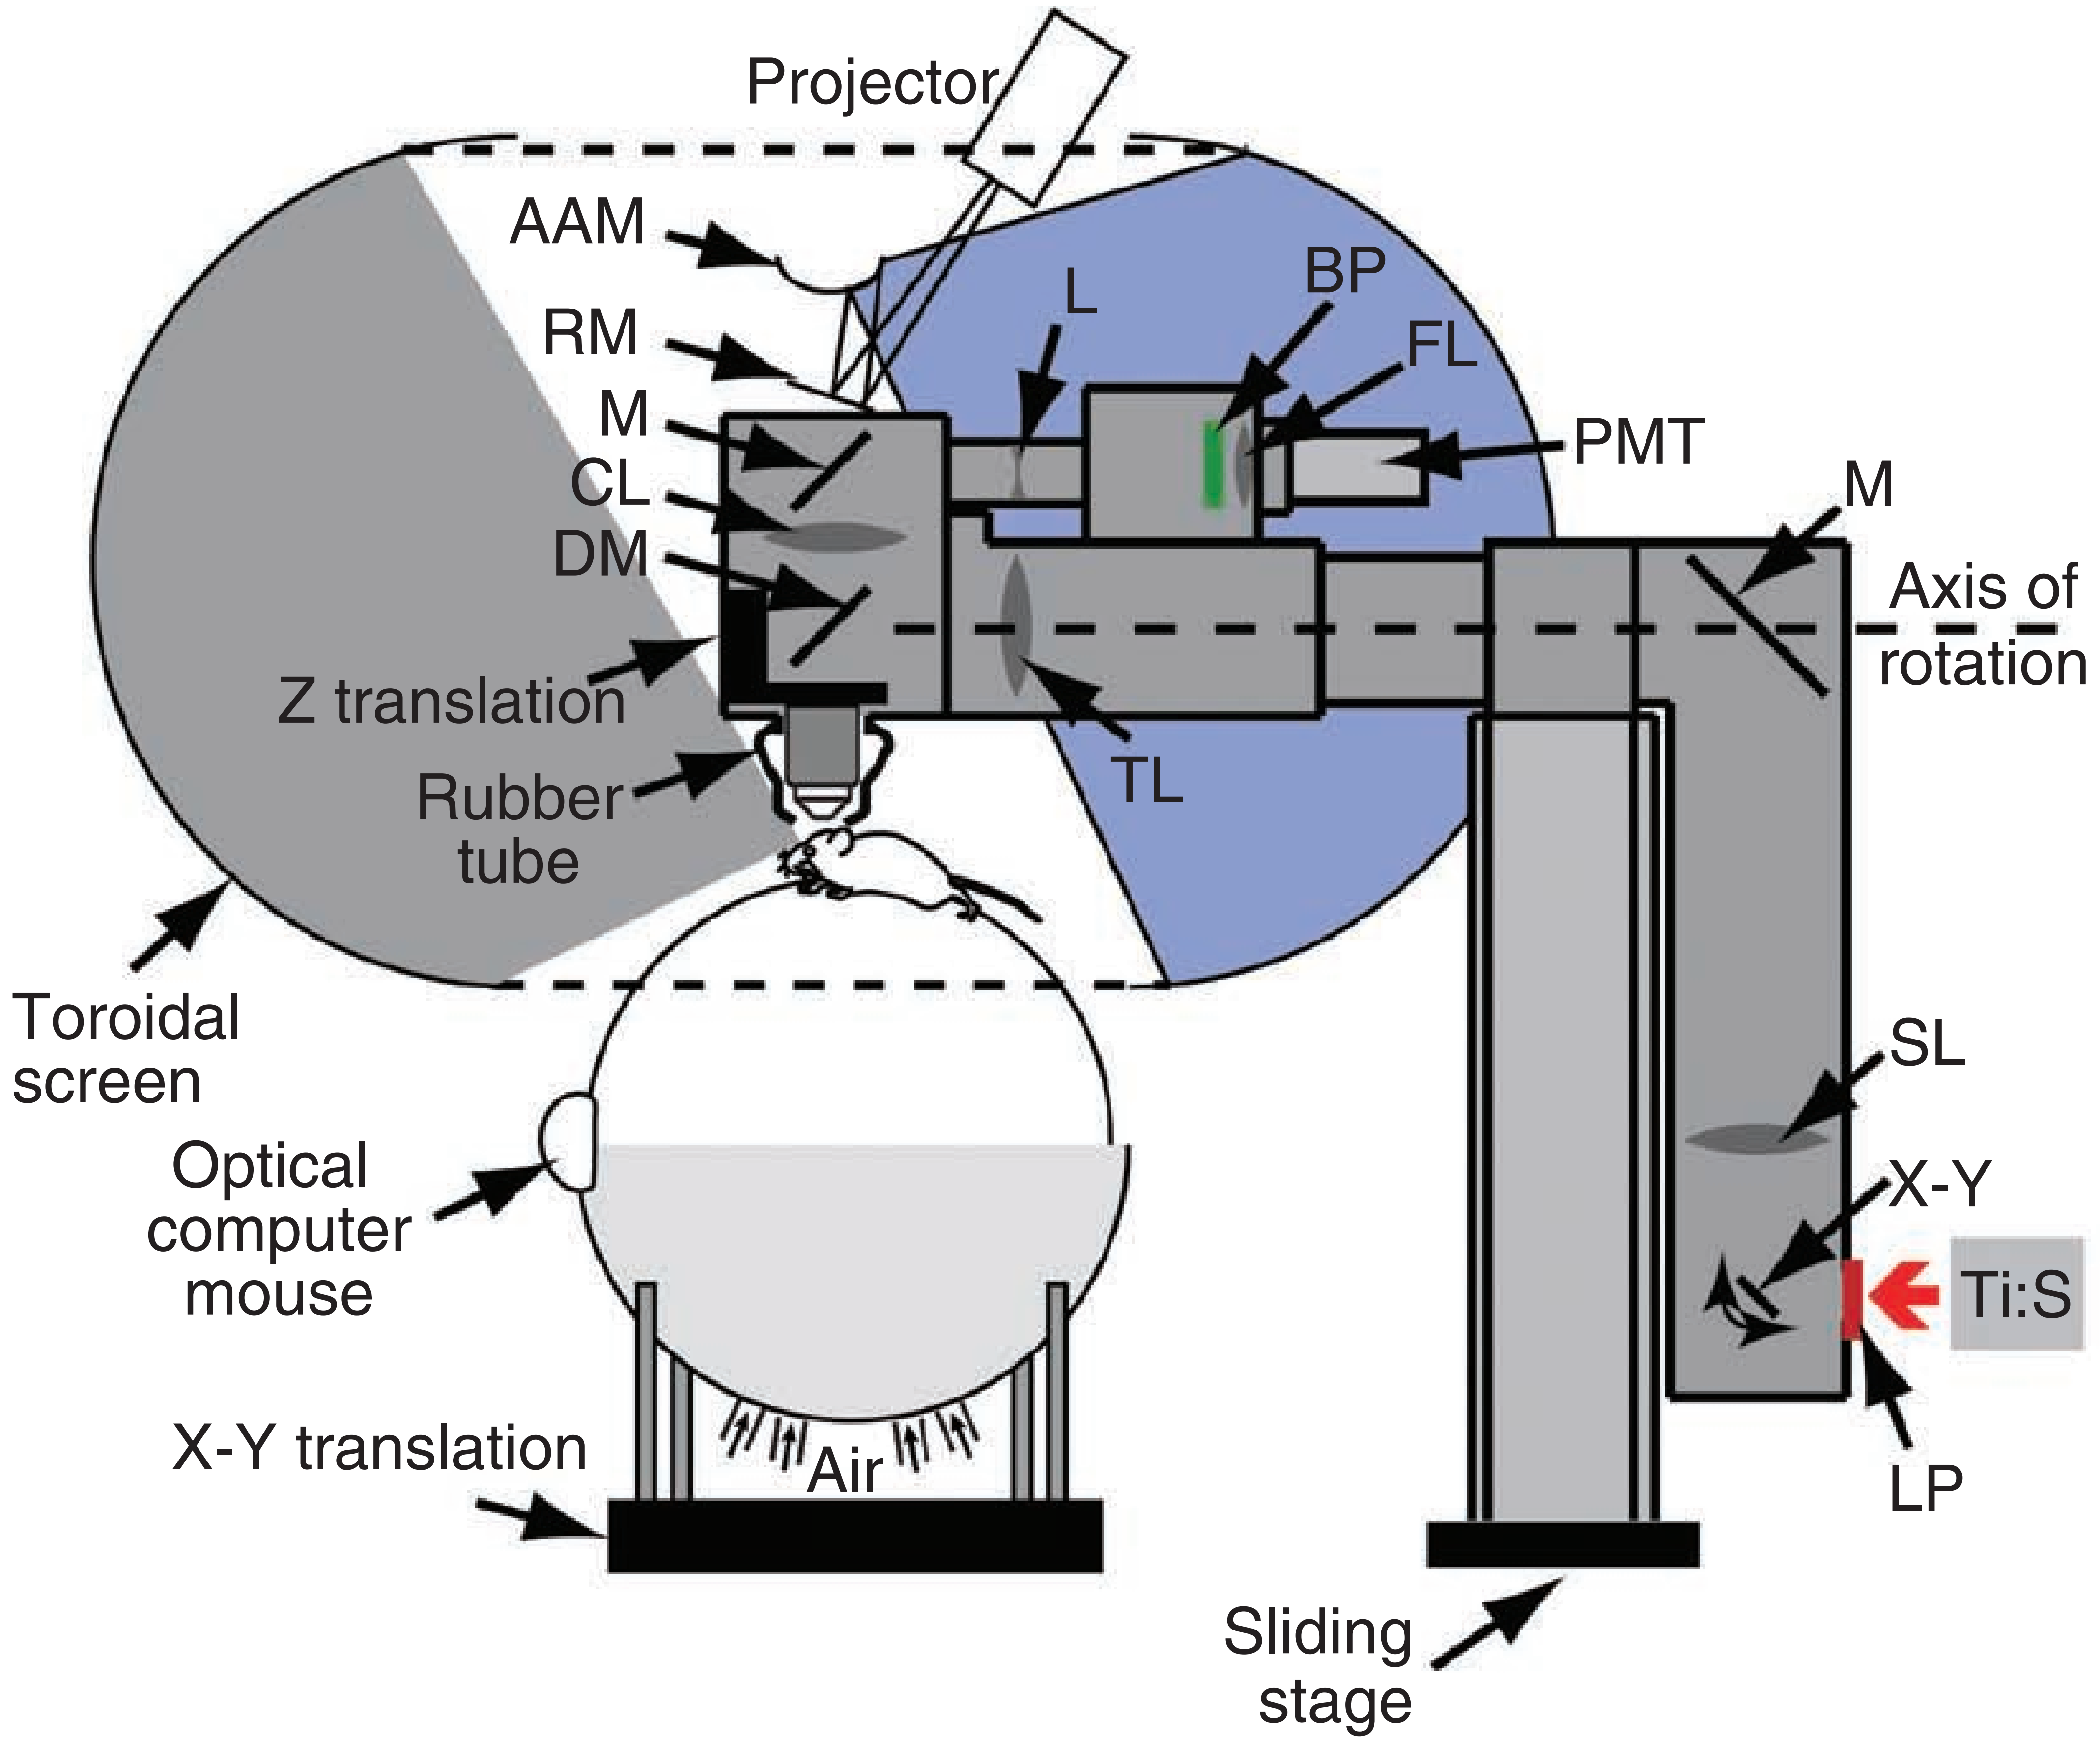

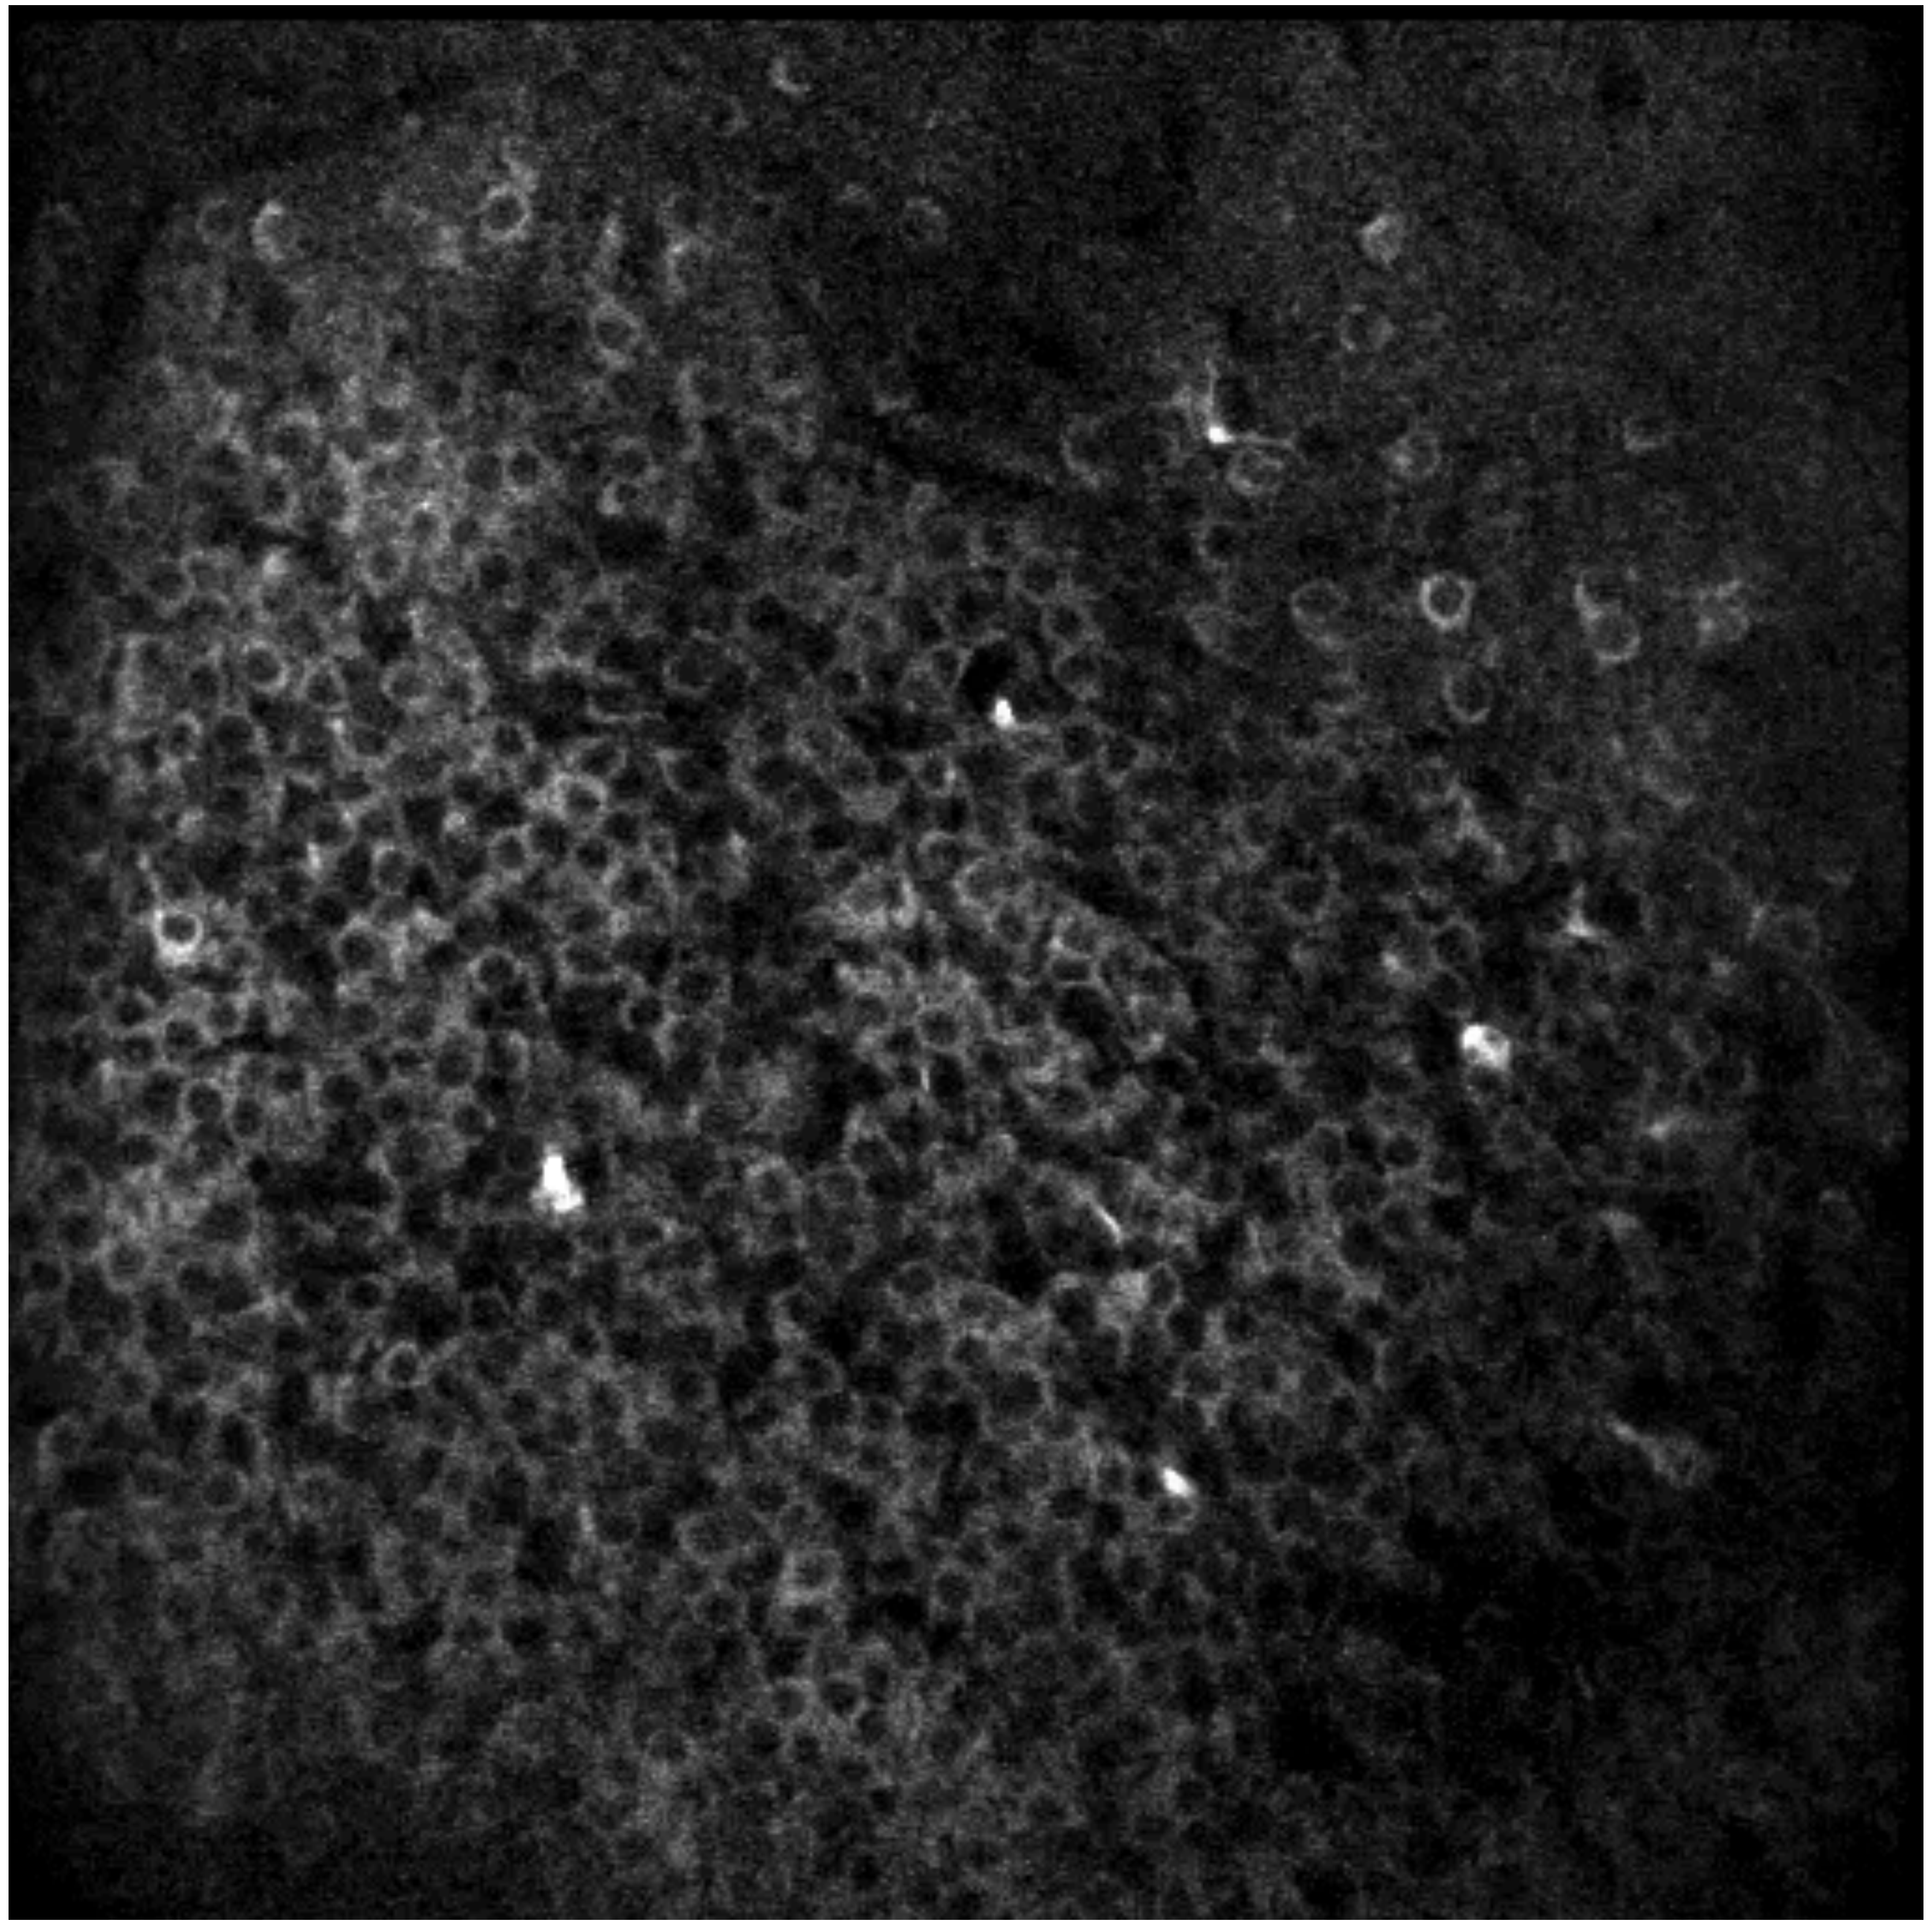

Combining genetic engineering, two-photon microscopy, and virtual reality to record from 1000+ neurons in the hippocampus.

DA Dombeck, CD Harvey, L Tian, LL Looger, and DW Tank, *Nat Neurosci* 13:1433 (2010).

MJ Berry II
R Segev (Ben Gurion University)
D Amodei (Stanford University)
O Marre (Vision Institute, Paris)

CD Brody
DW Tank
JL Gauthier

AM Leifer
JW Shaevitz
JP Nguyen
F Randi

R Ranganathan (UT Southwestern)

A Cavagna & I Giardina (Rome)

E Schneidman (Weizmann Institute)
S Still (University of Hawaii)
G Tkačik (IST Austria)
T Mora (CNRS/ENS Paris)
SE Palmer (U Chicago)
S Bradde (Physical Review)
L Meshulam
X Chen

CG Callan
AM Walczak (CNRS/ENS Paris)

GJ Stephens (VU Amsterdam)

LJ Colwell (Cambridge)
ED Lee (Cornell)

Optical recording from hippocampal neurons as a mouse moves in a virtual environment

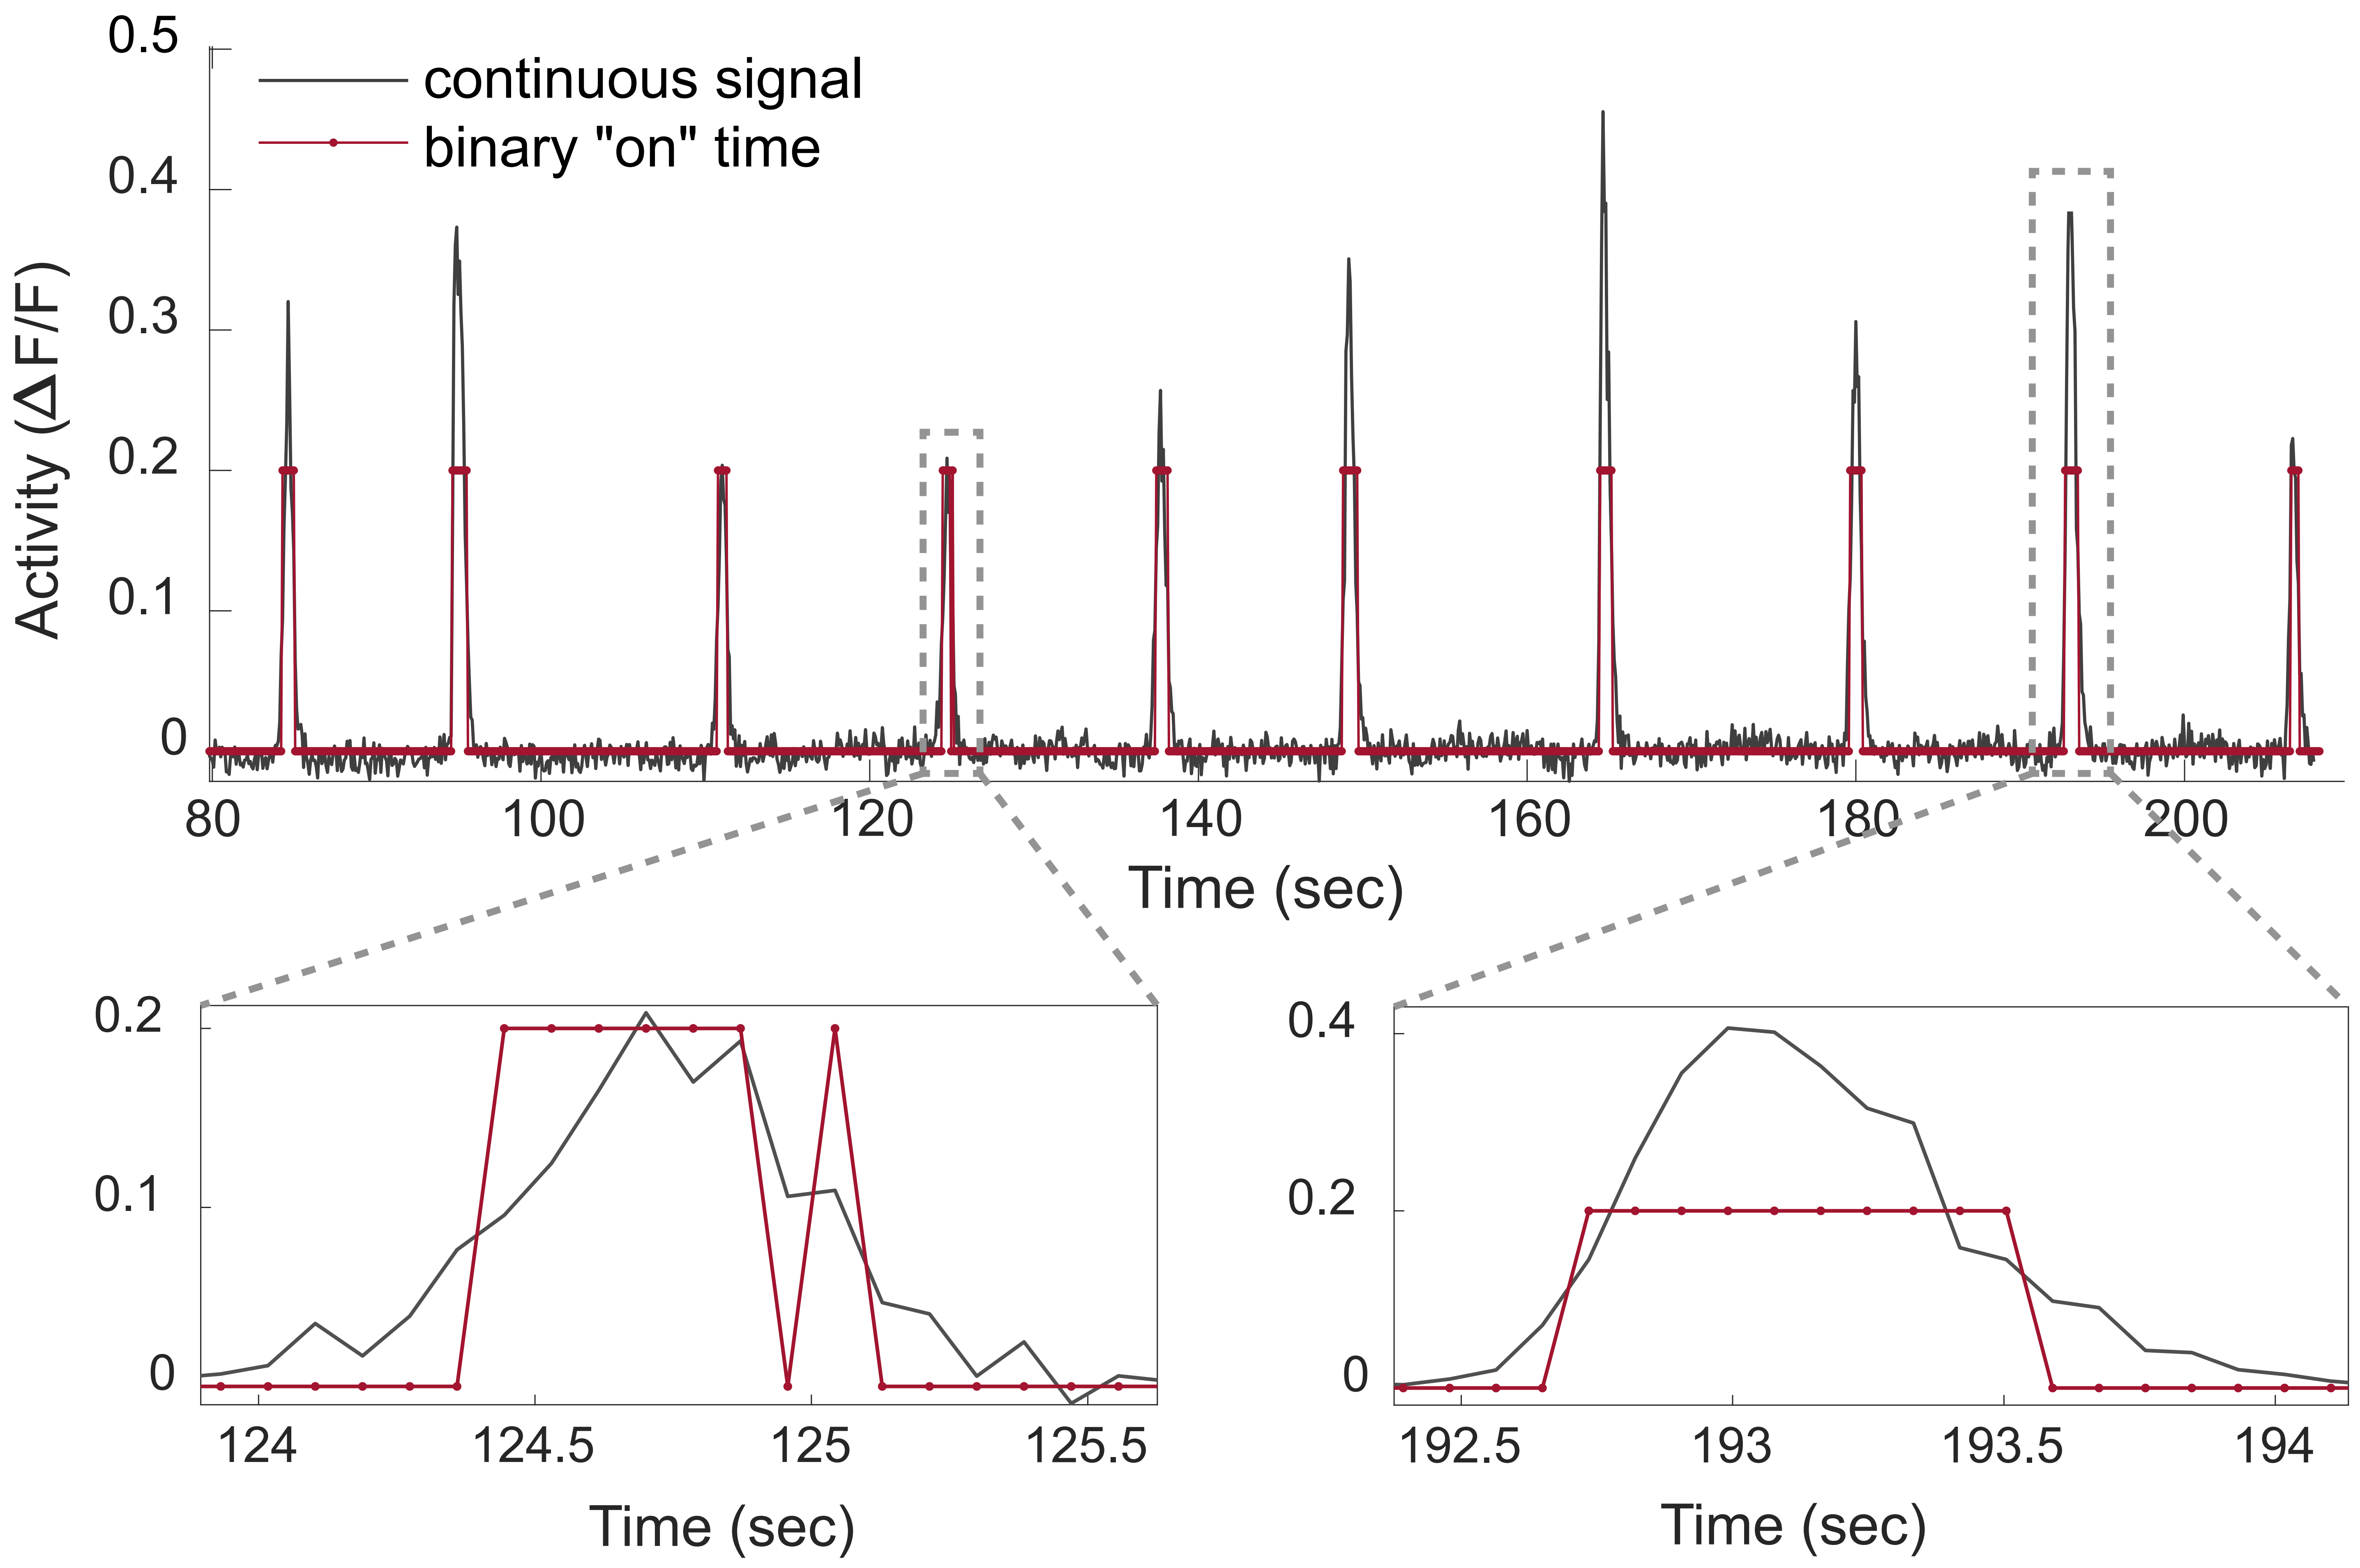

Denoising + discretization leads to a binary activity variable for each neuron

$$\sigma_{\rm i}(t) = \begin{cases} \text{1 (active)} \\ \text{0 (silent)} \end{cases}$$

State of the network $\{\sigma_{\rm i}\}$

What is $P\left(\{\sigma_{\rm i}\}\right)$?

L Meshulam, JL Gauthier, CD Brody, DW Tank, and WB, *Neuron* 96:1 (2017).

What features of the data do we want to capture?

Mean activity of individual neurons
Correlations between pairs of neurons

$$\langle \sigma_{\mathrm{j}} \rangle_{\mathrm{model}} \equiv \sum_{\{\sigma_{\mathrm{i}}\}} P\left(\{\sigma_{\mathrm{i}}\}\right) \sigma_{\mathrm{j}} = \langle \sigma_{\mathrm{j}} \rangle_{\mathrm{data}}$$

$$\langle \sigma_{\mathrm{j}} \sigma_{\mathrm{k}} \rangle_{\mathrm{model}} \equiv \sum_{\{\sigma_{\mathrm{i}}\}} P\left(\{\sigma_{\mathrm{i}}\}\right) \sigma_{\mathrm{j}} \sigma_{\mathrm{k}} = \langle \sigma_{\mathrm{j}} \sigma_{\mathrm{k}} \rangle_{\mathrm{data}}$$

Infinitely many models are consistent with these constraints
Choose the one with the least structure - maximum entropy

$$P\left(\{\sigma_{\mathrm{i}}\}\right) = \frac{1}{Z} \exp\left[\sum_{\mathrm{i}} h_{\mathrm{i}} \sigma_{\mathrm{i}} + \frac{1}{2} \sum_{\mathrm{i} \neq \mathrm{j}} J_{\mathrm{ij}} \sigma_{\mathrm{i}} \sigma_{\mathrm{j}}\right]$$

Since we used the pair correlations to build the model, can we predict correlations among triplets?

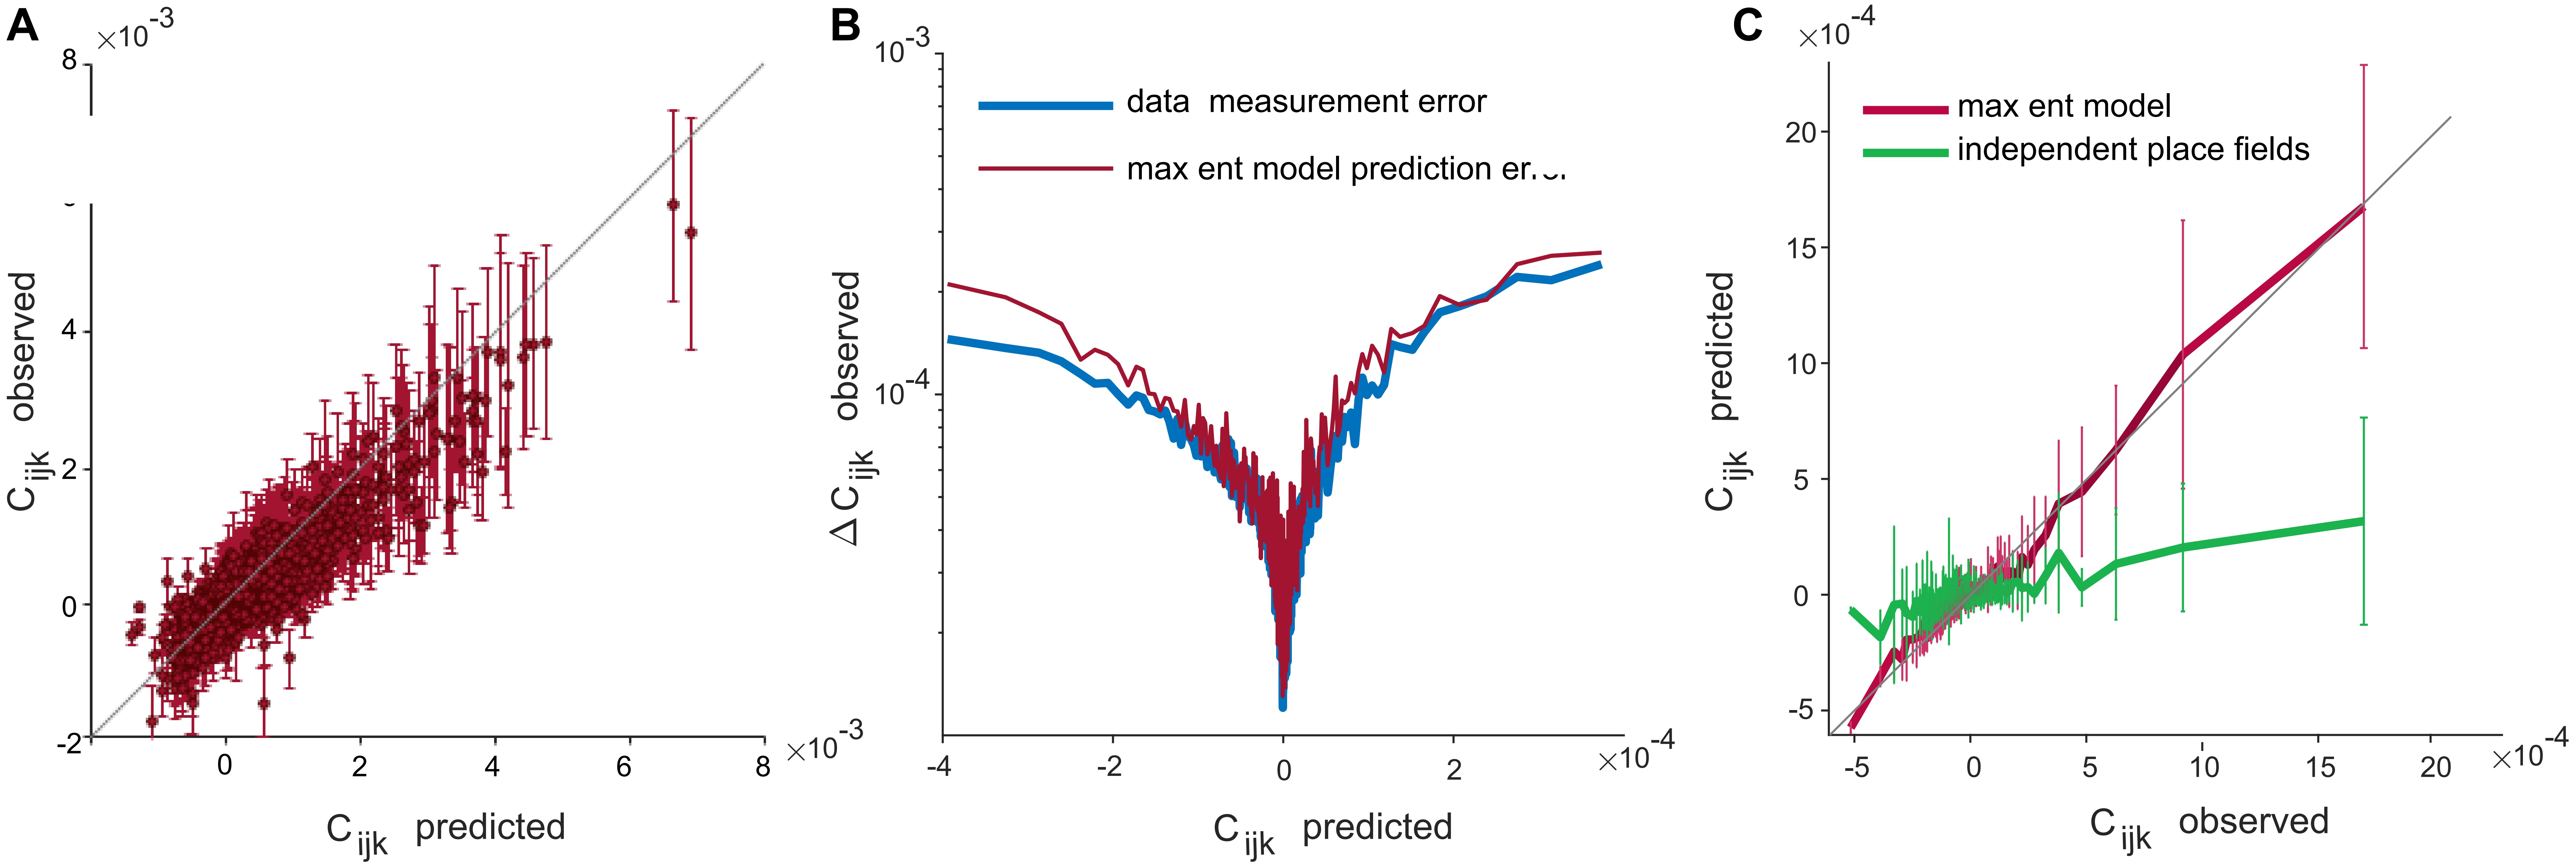

Are correlations inherited from place fields?

If activity really is collective, we can predict the probability of one cell being active from the state of all the other cells in the network.

$$P(\sigma_{\rm i} = 1 | \{\sigma_{\rm j \neq i}\}) = \frac{1}{1 + \exp(-h_{\rm i}^{\rm eff})}$$

$$h_{\rm i}^{\rm eff} = h_{\rm i} + \sum_{\rm j \neq i} J_{\rm ij} \sigma_{\rm j}$$

Let's "unfold" this relationship over time …

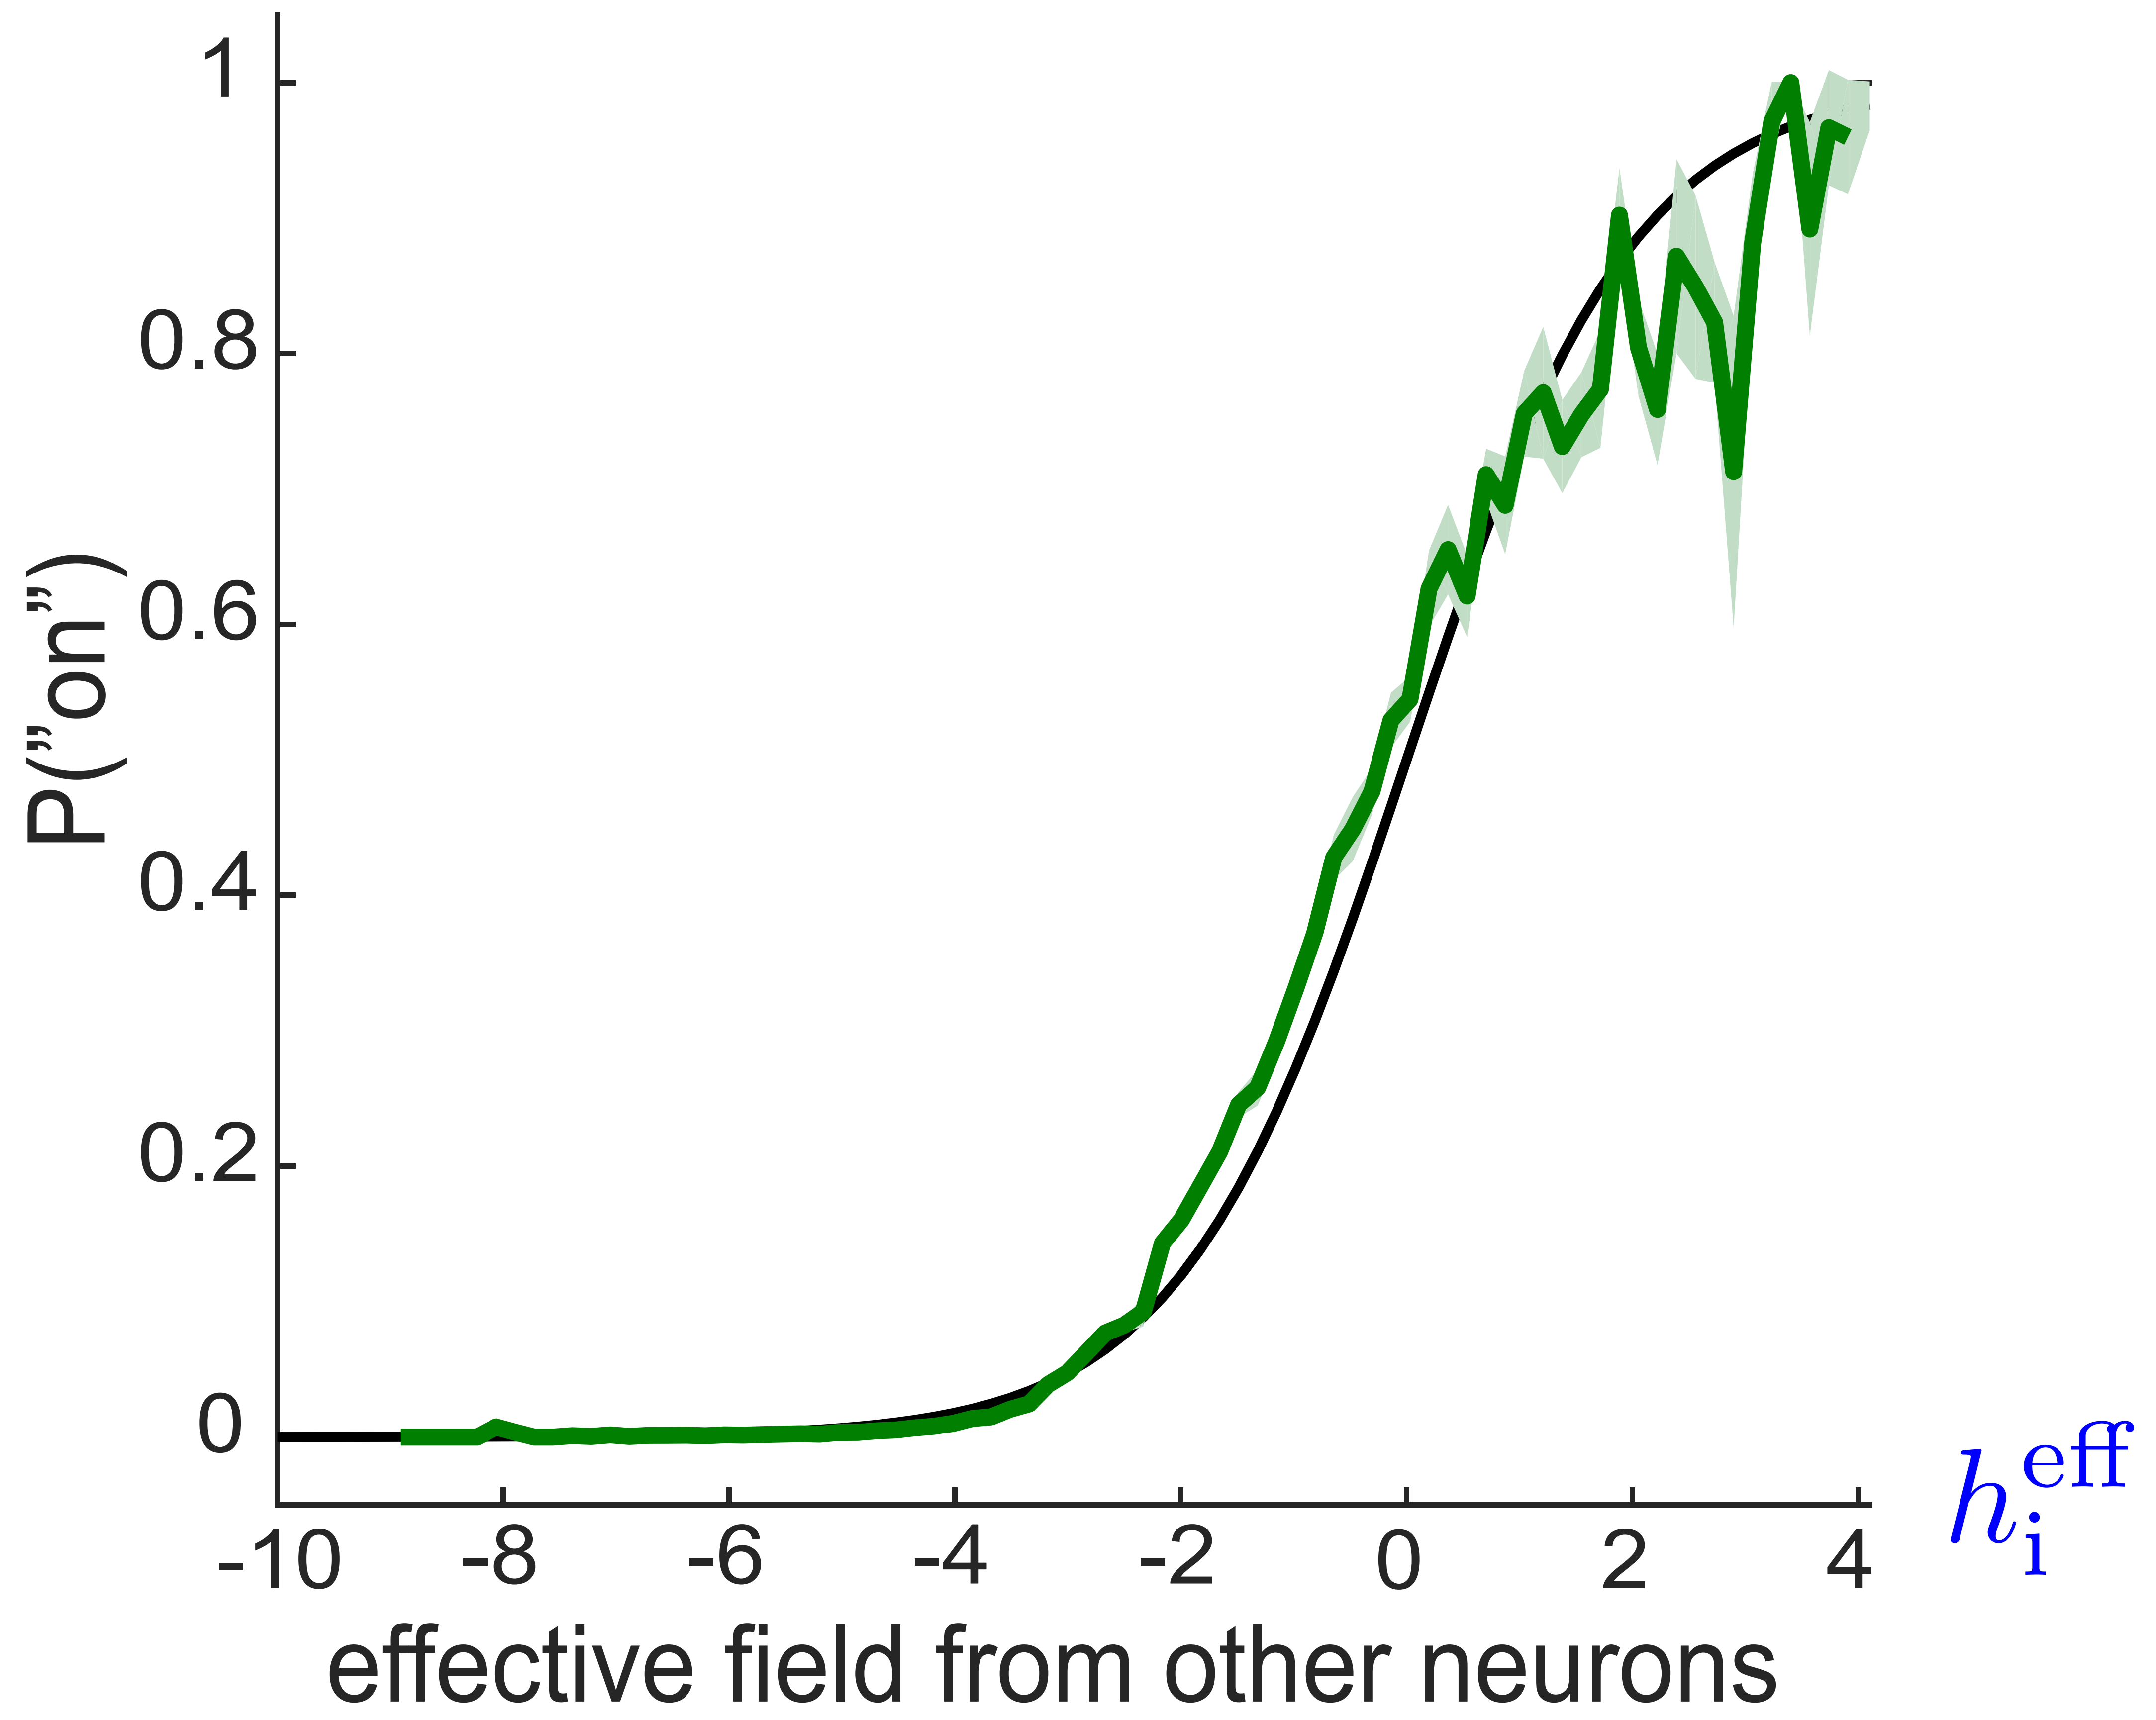

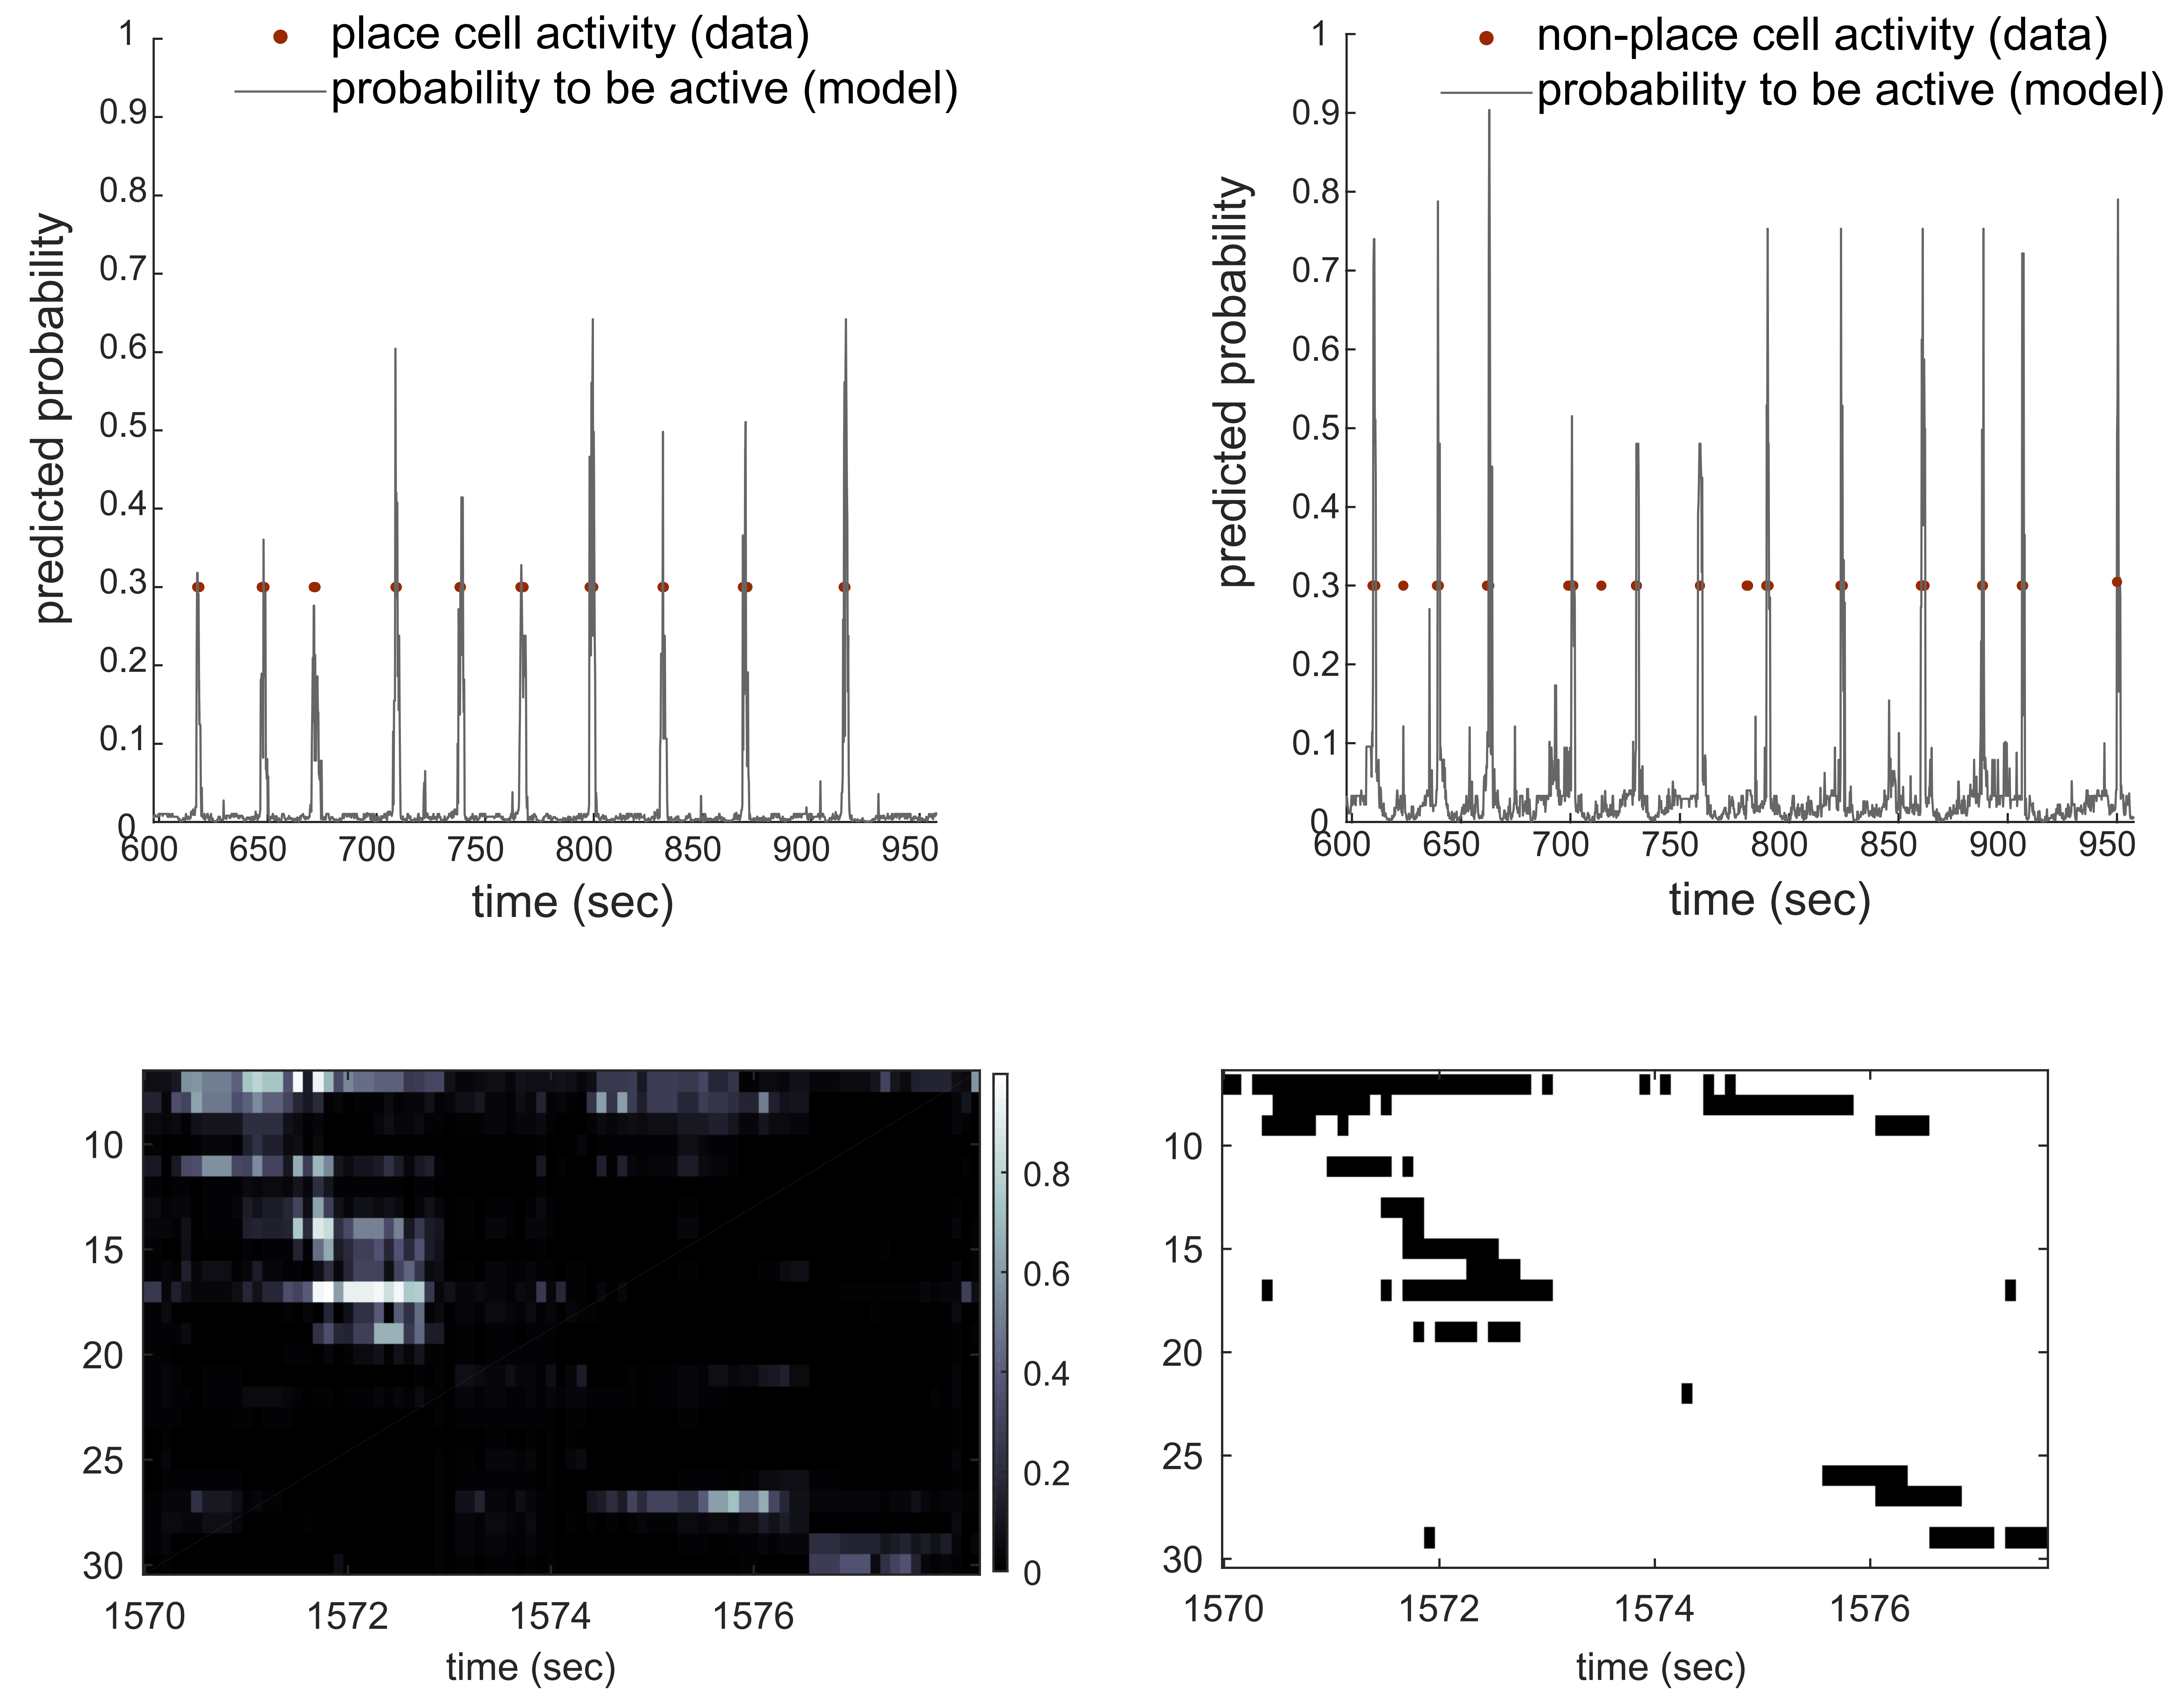
place cell activity (data)
probability to be active (model)
predicted probability
time (sec)
non-place cell activity (data)
probability to be active (model)
predicted probability
time (sec)
time (sec)
time (sec)

Can we write down the (joint!) probability distribution for the activity of many neurons in a network?

Yes. In fact, with ~100 neurons, we can construct models that are surprisingly precise.

Is there any reason to think that this distribution is simpler than it could be?

(a brief reminder about the RG)

S Bradde & WB, *J Stat Phys* 167:462 (2017).
L Mehsulam, JL Gauhtier, CD Brody, DW Tank, and WB (almost done).

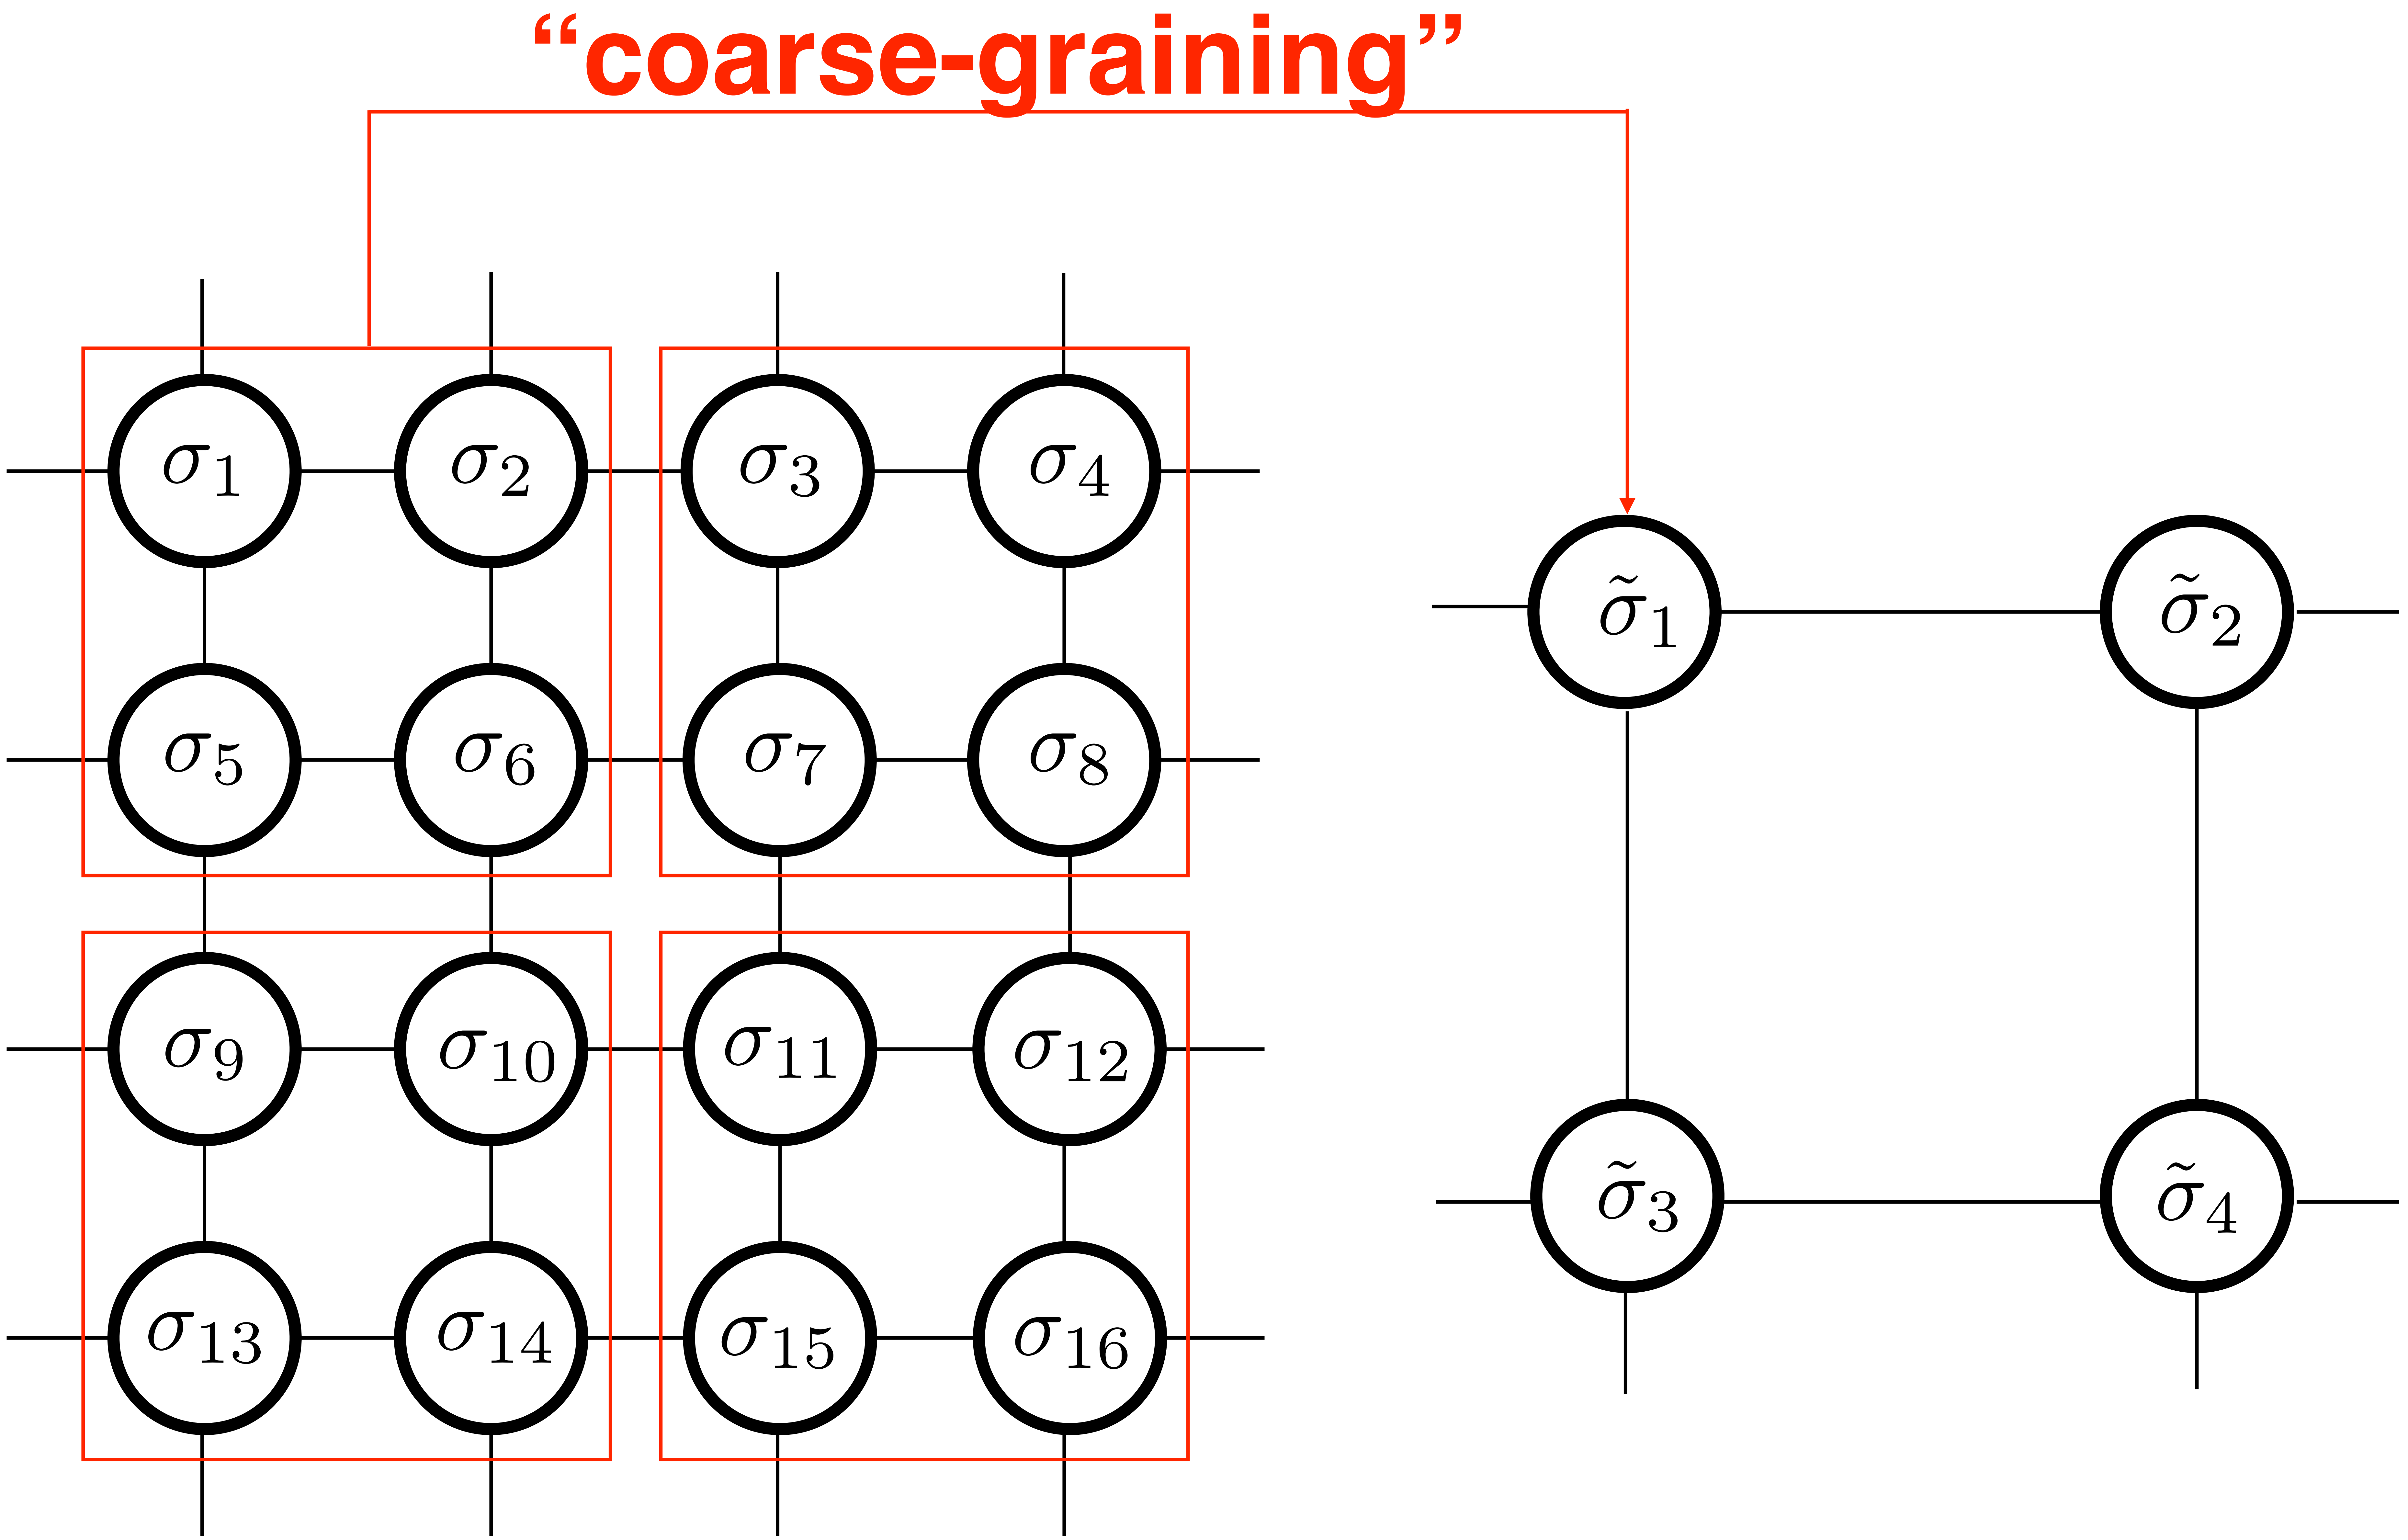

$$P(\{\sigma_{\rm i}\}) \xrightarrow[\textbf{of models}]{\textbf{flow in the space}} P(\{\tilde{\sigma}_{\rm i}\})$$

Instead of spatial neighbors, add together activity of maximally correlated pairs.

$$\tilde{\sigma}_{\mathrm{i}} = \sigma_{\mathrm{i}} + \sigma_{\mathrm{j}_*(\mathrm{i})}$$

Iterate.

Produces clusters of 2, 4, 8, … analogous to spatially contiguous regions.

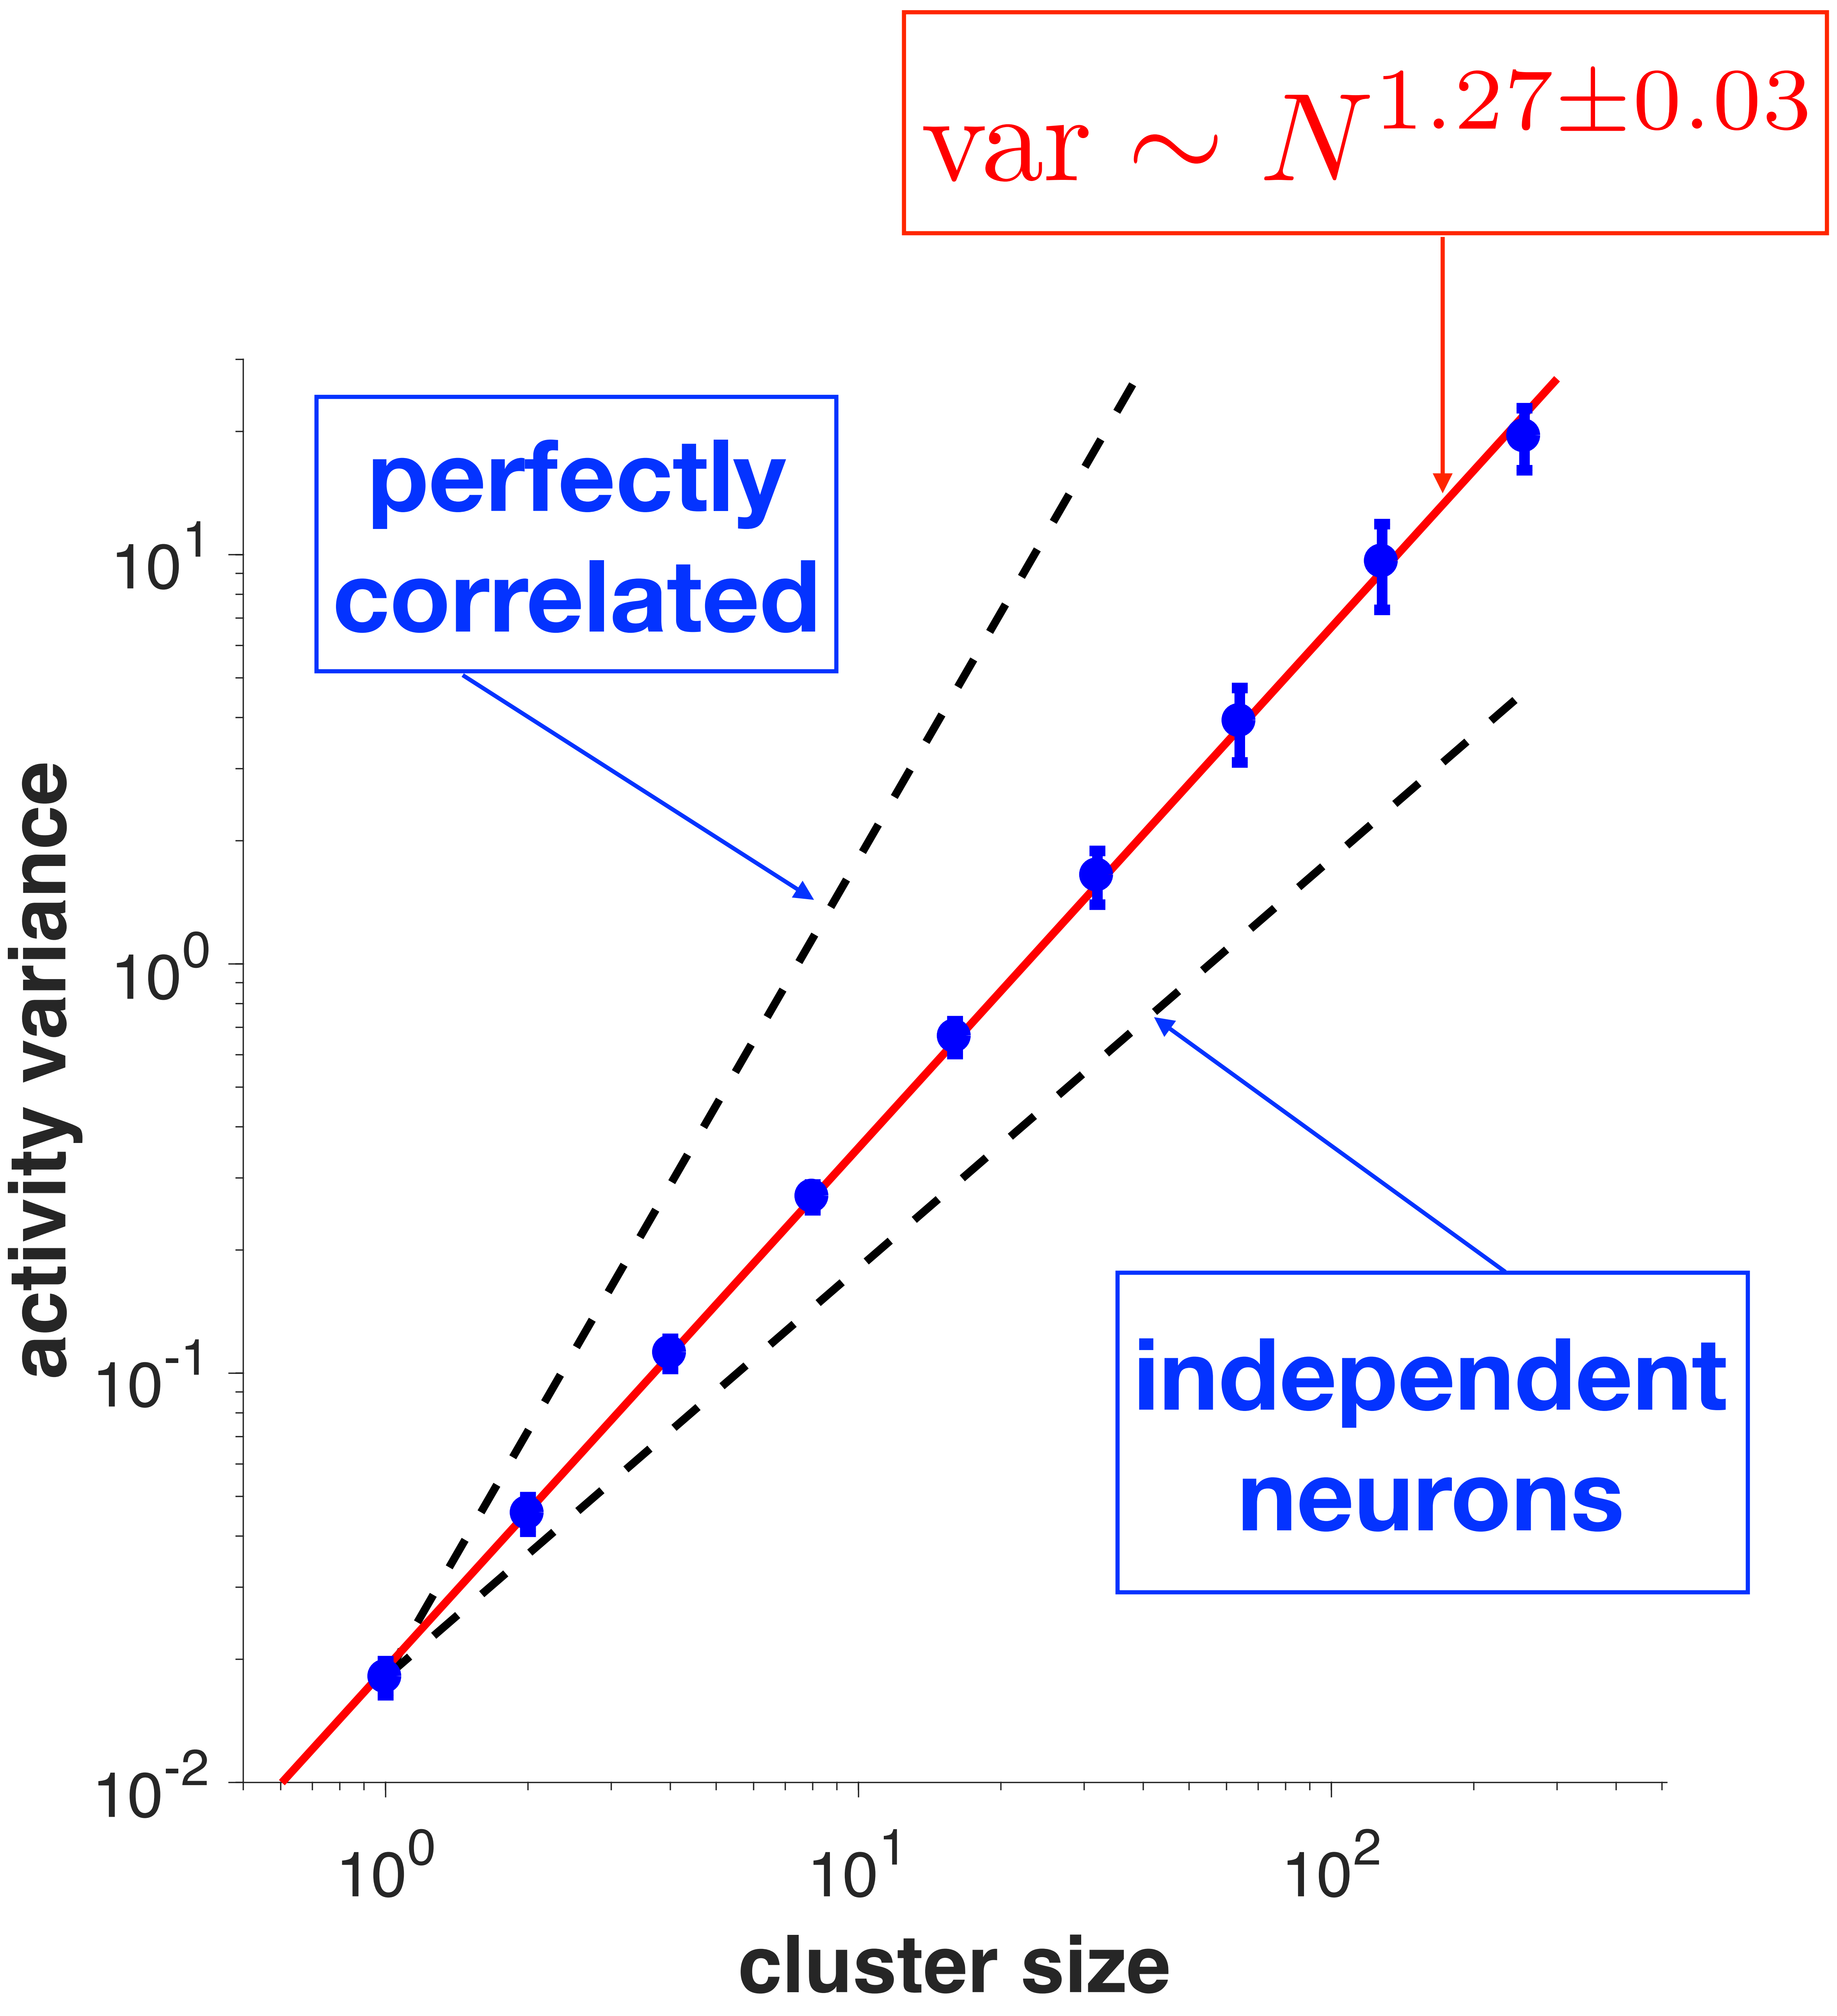

Correlations inside the clusters

$$C_{\mathrm{ij}} = \langle \sigma_{\mathrm{i}} \sigma_{\mathrm{j}} \rangle - \langle \sigma_{\mathrm{i}} \rangle \langle \sigma_{\mathrm{j}} \rangle$$

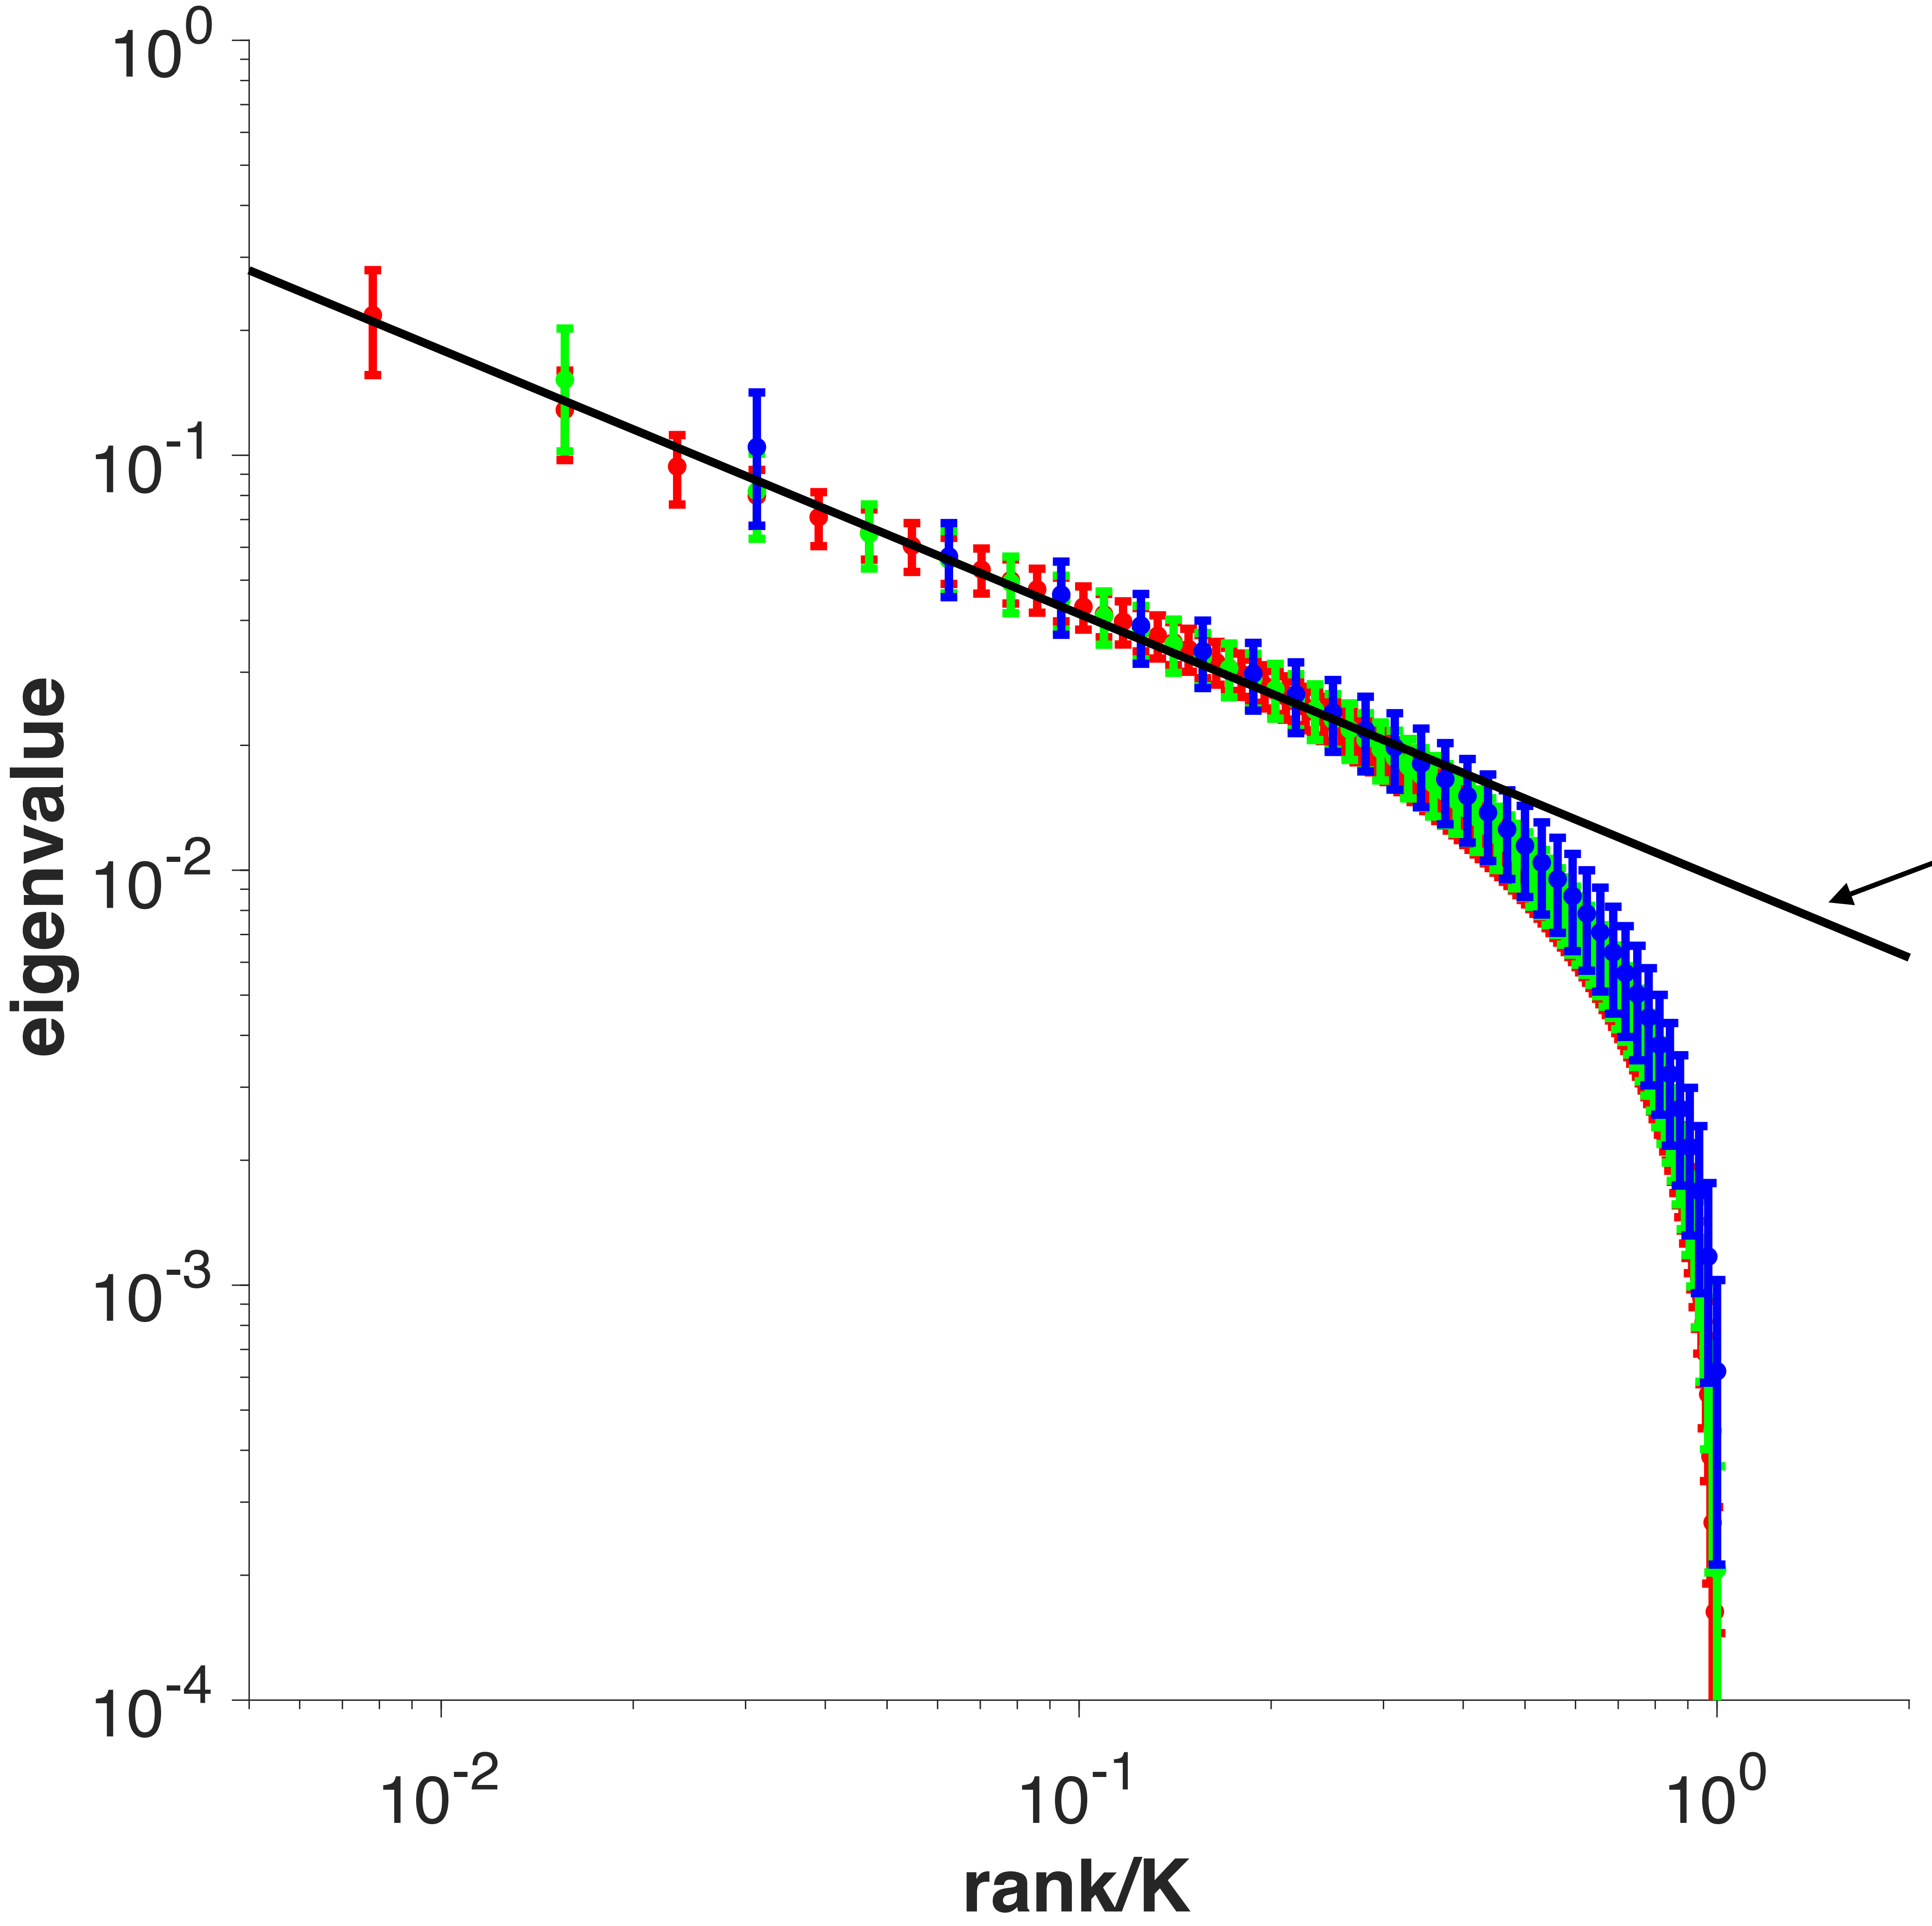

Find the eigenvalues in clusters of different sizes
(be careful about sampling problems!)

$$\lambda = A \left(\frac{\text{cluster size}}{\text{rank}} \right)^{0.64 \pm 0.02}$$

K = 32
K = 64
K = 128

Probability that the entire cluster is silent

$$P\left(\{\sigma_{\rm i}\}\right) = \frac{1}{Z} \exp\left[\sum_{\rm i} h_{\rm i}\sigma_{\rm i} + \frac{1}{2}\sum_{\rm i\neq j} J_{\rm ij}\sigma_{\rm i}\sigma_{\rm j}\right]$$

$$\sigma_{\rm i} = 0, 1$$

$$\Rightarrow P\left(\{\sigma_{\rm i} = 0\}\right) = \frac{1}{Z} = e^{F}$$

So we can estimate the "free energy" as a function of cluster size

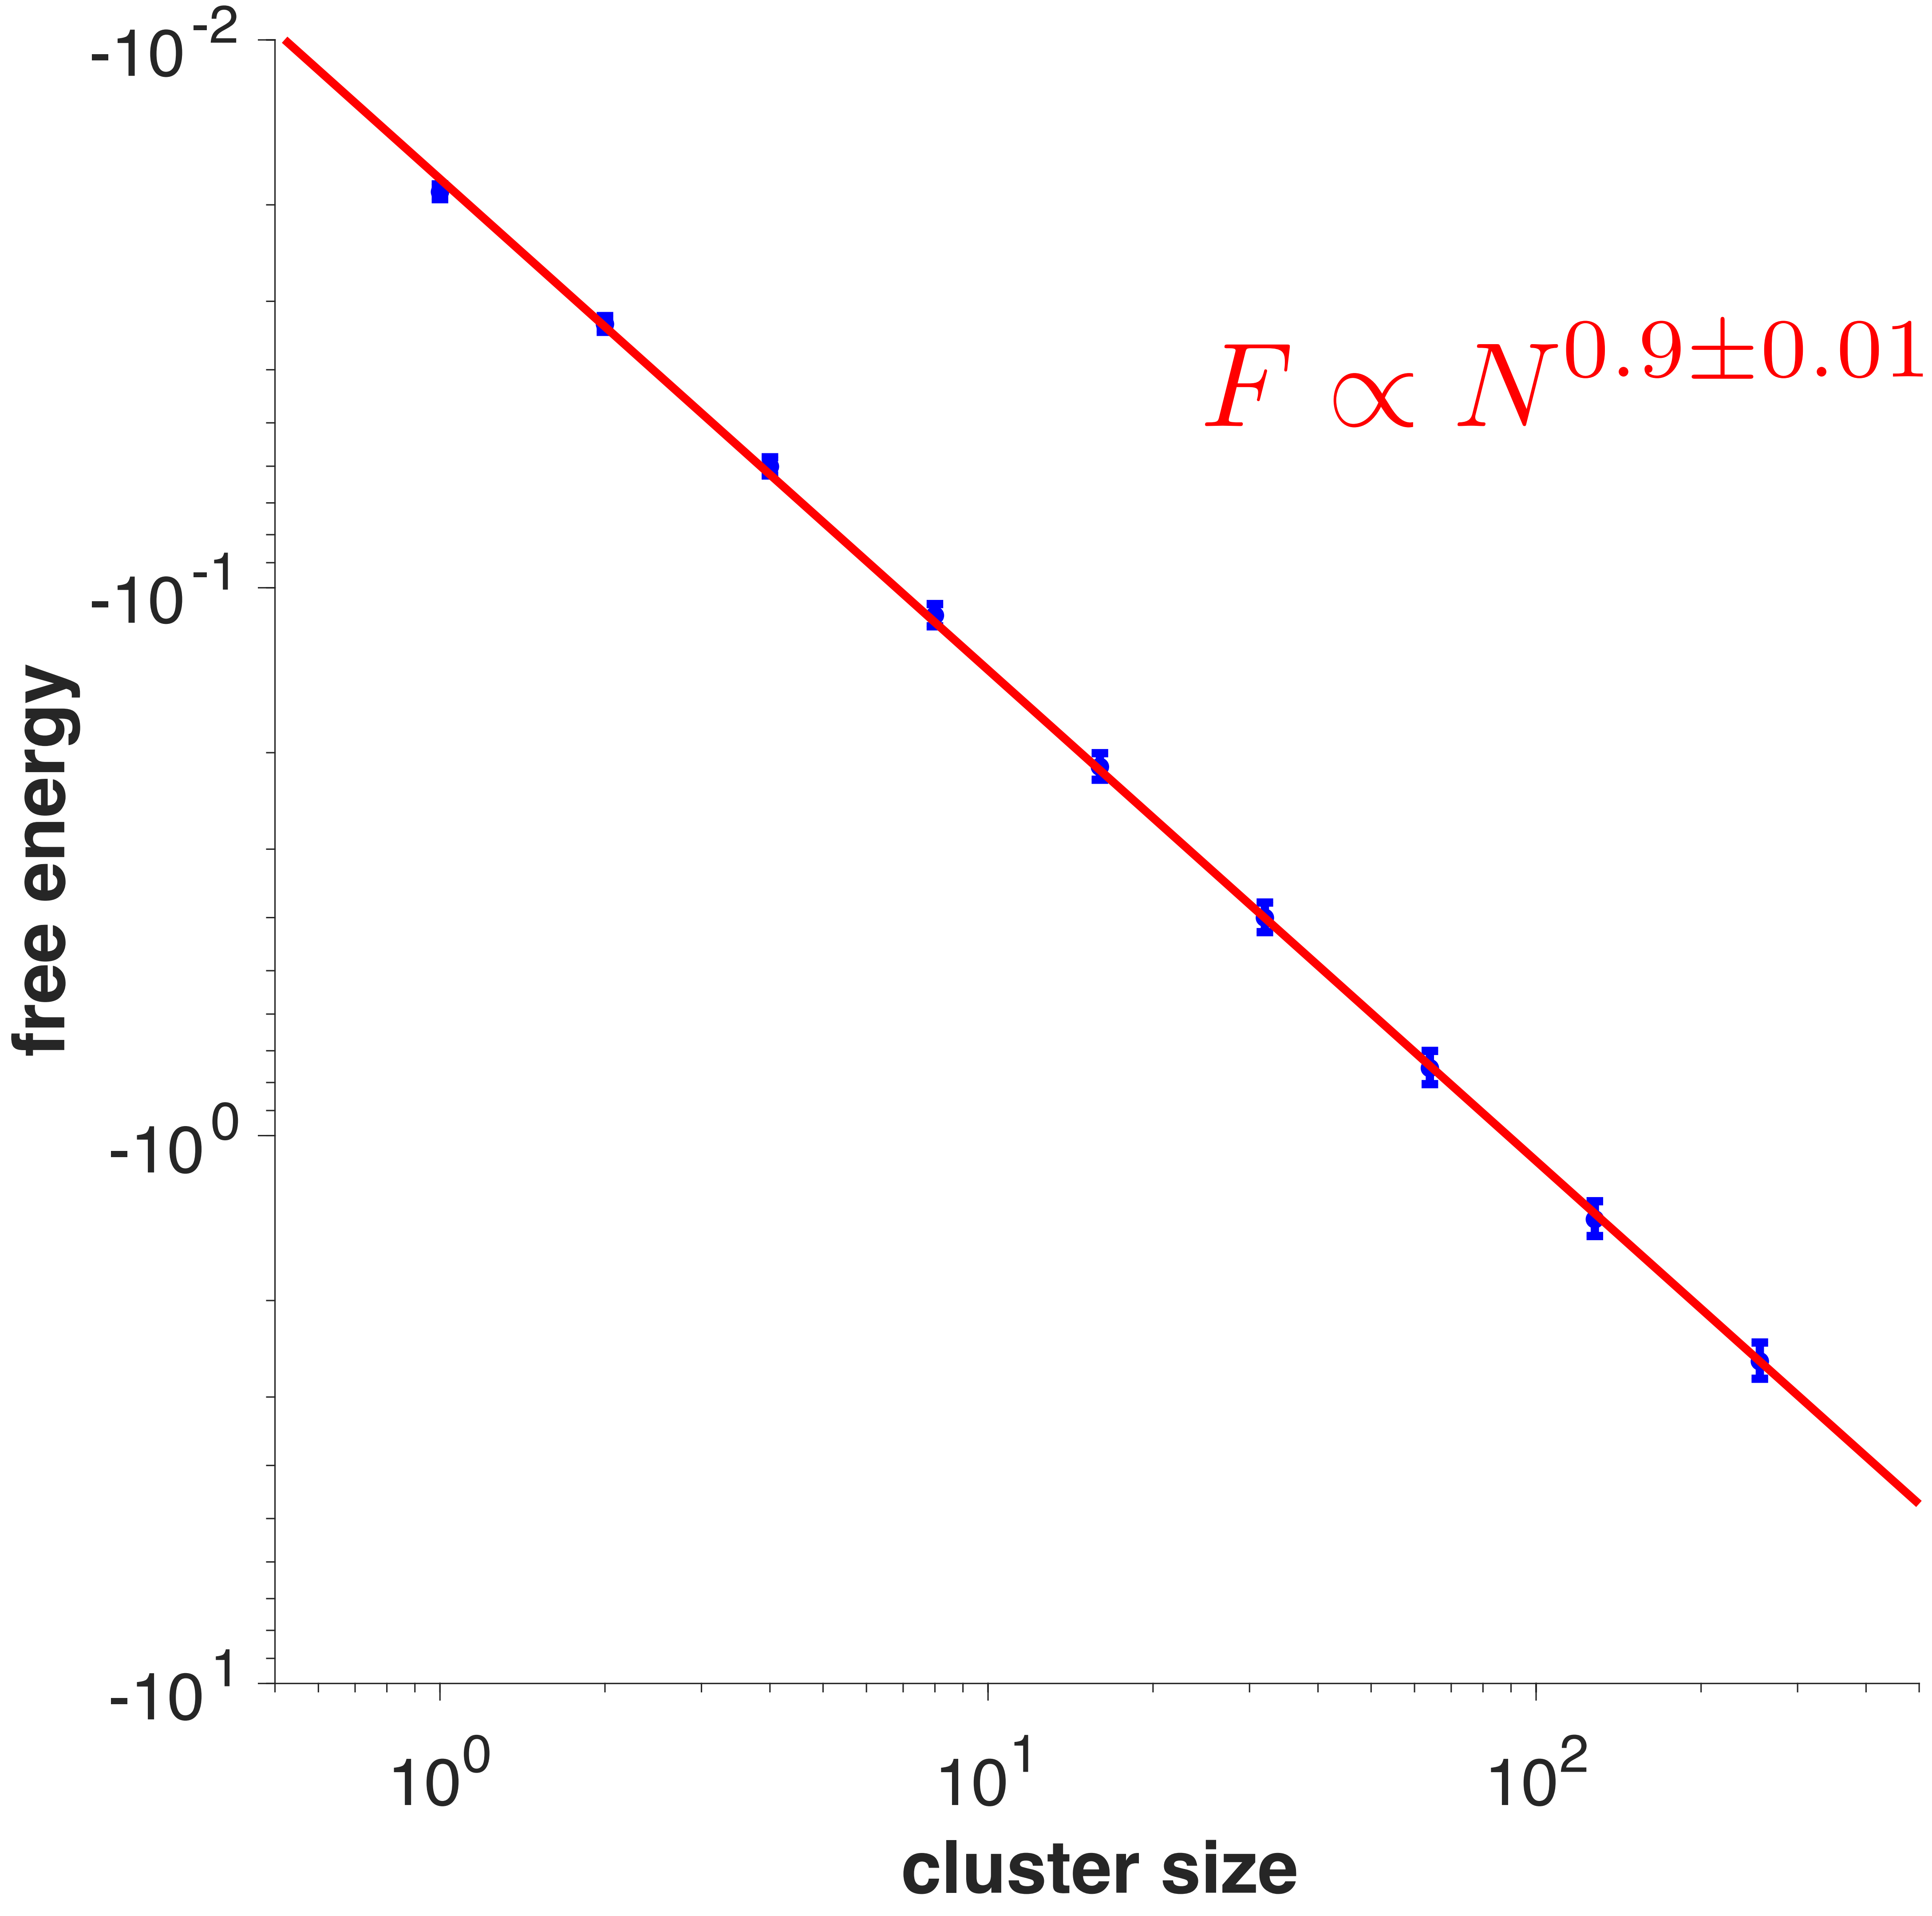

Distribution of nonzero activity

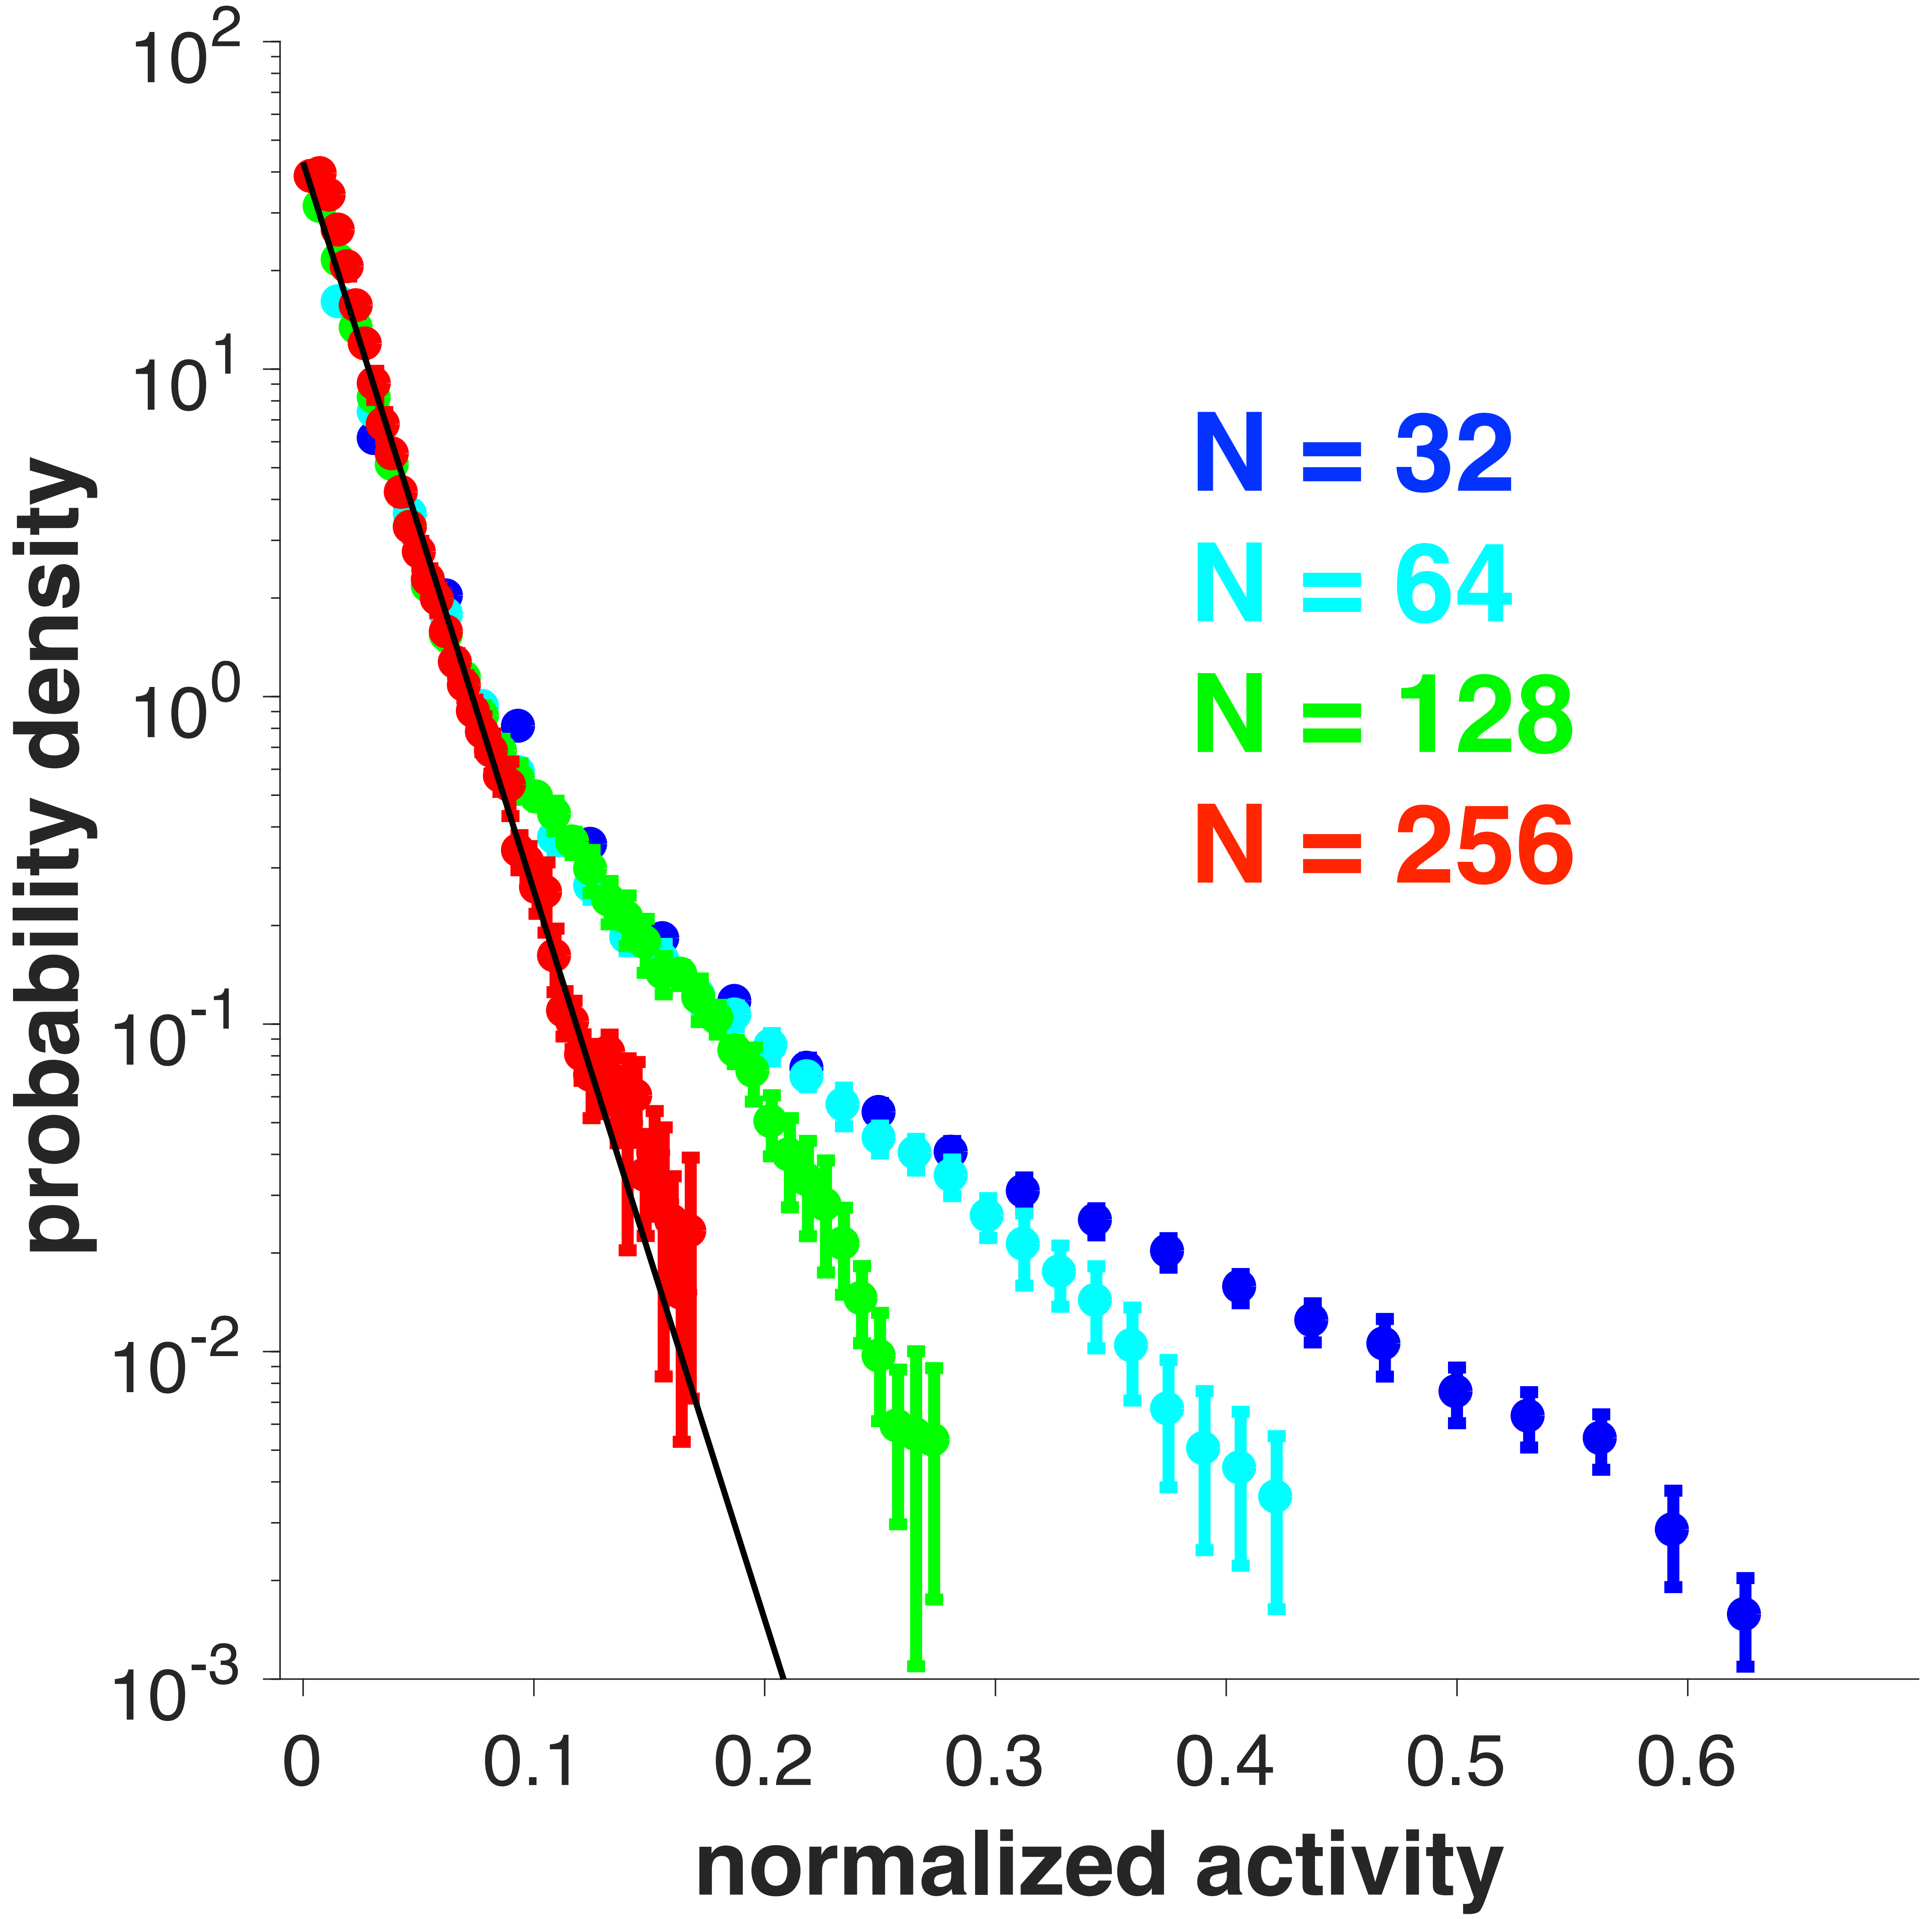

Distribution approaches a fixed form at large scales.

This is also visible in raw fluorescence data.

Larger clusters have slower dynamics …

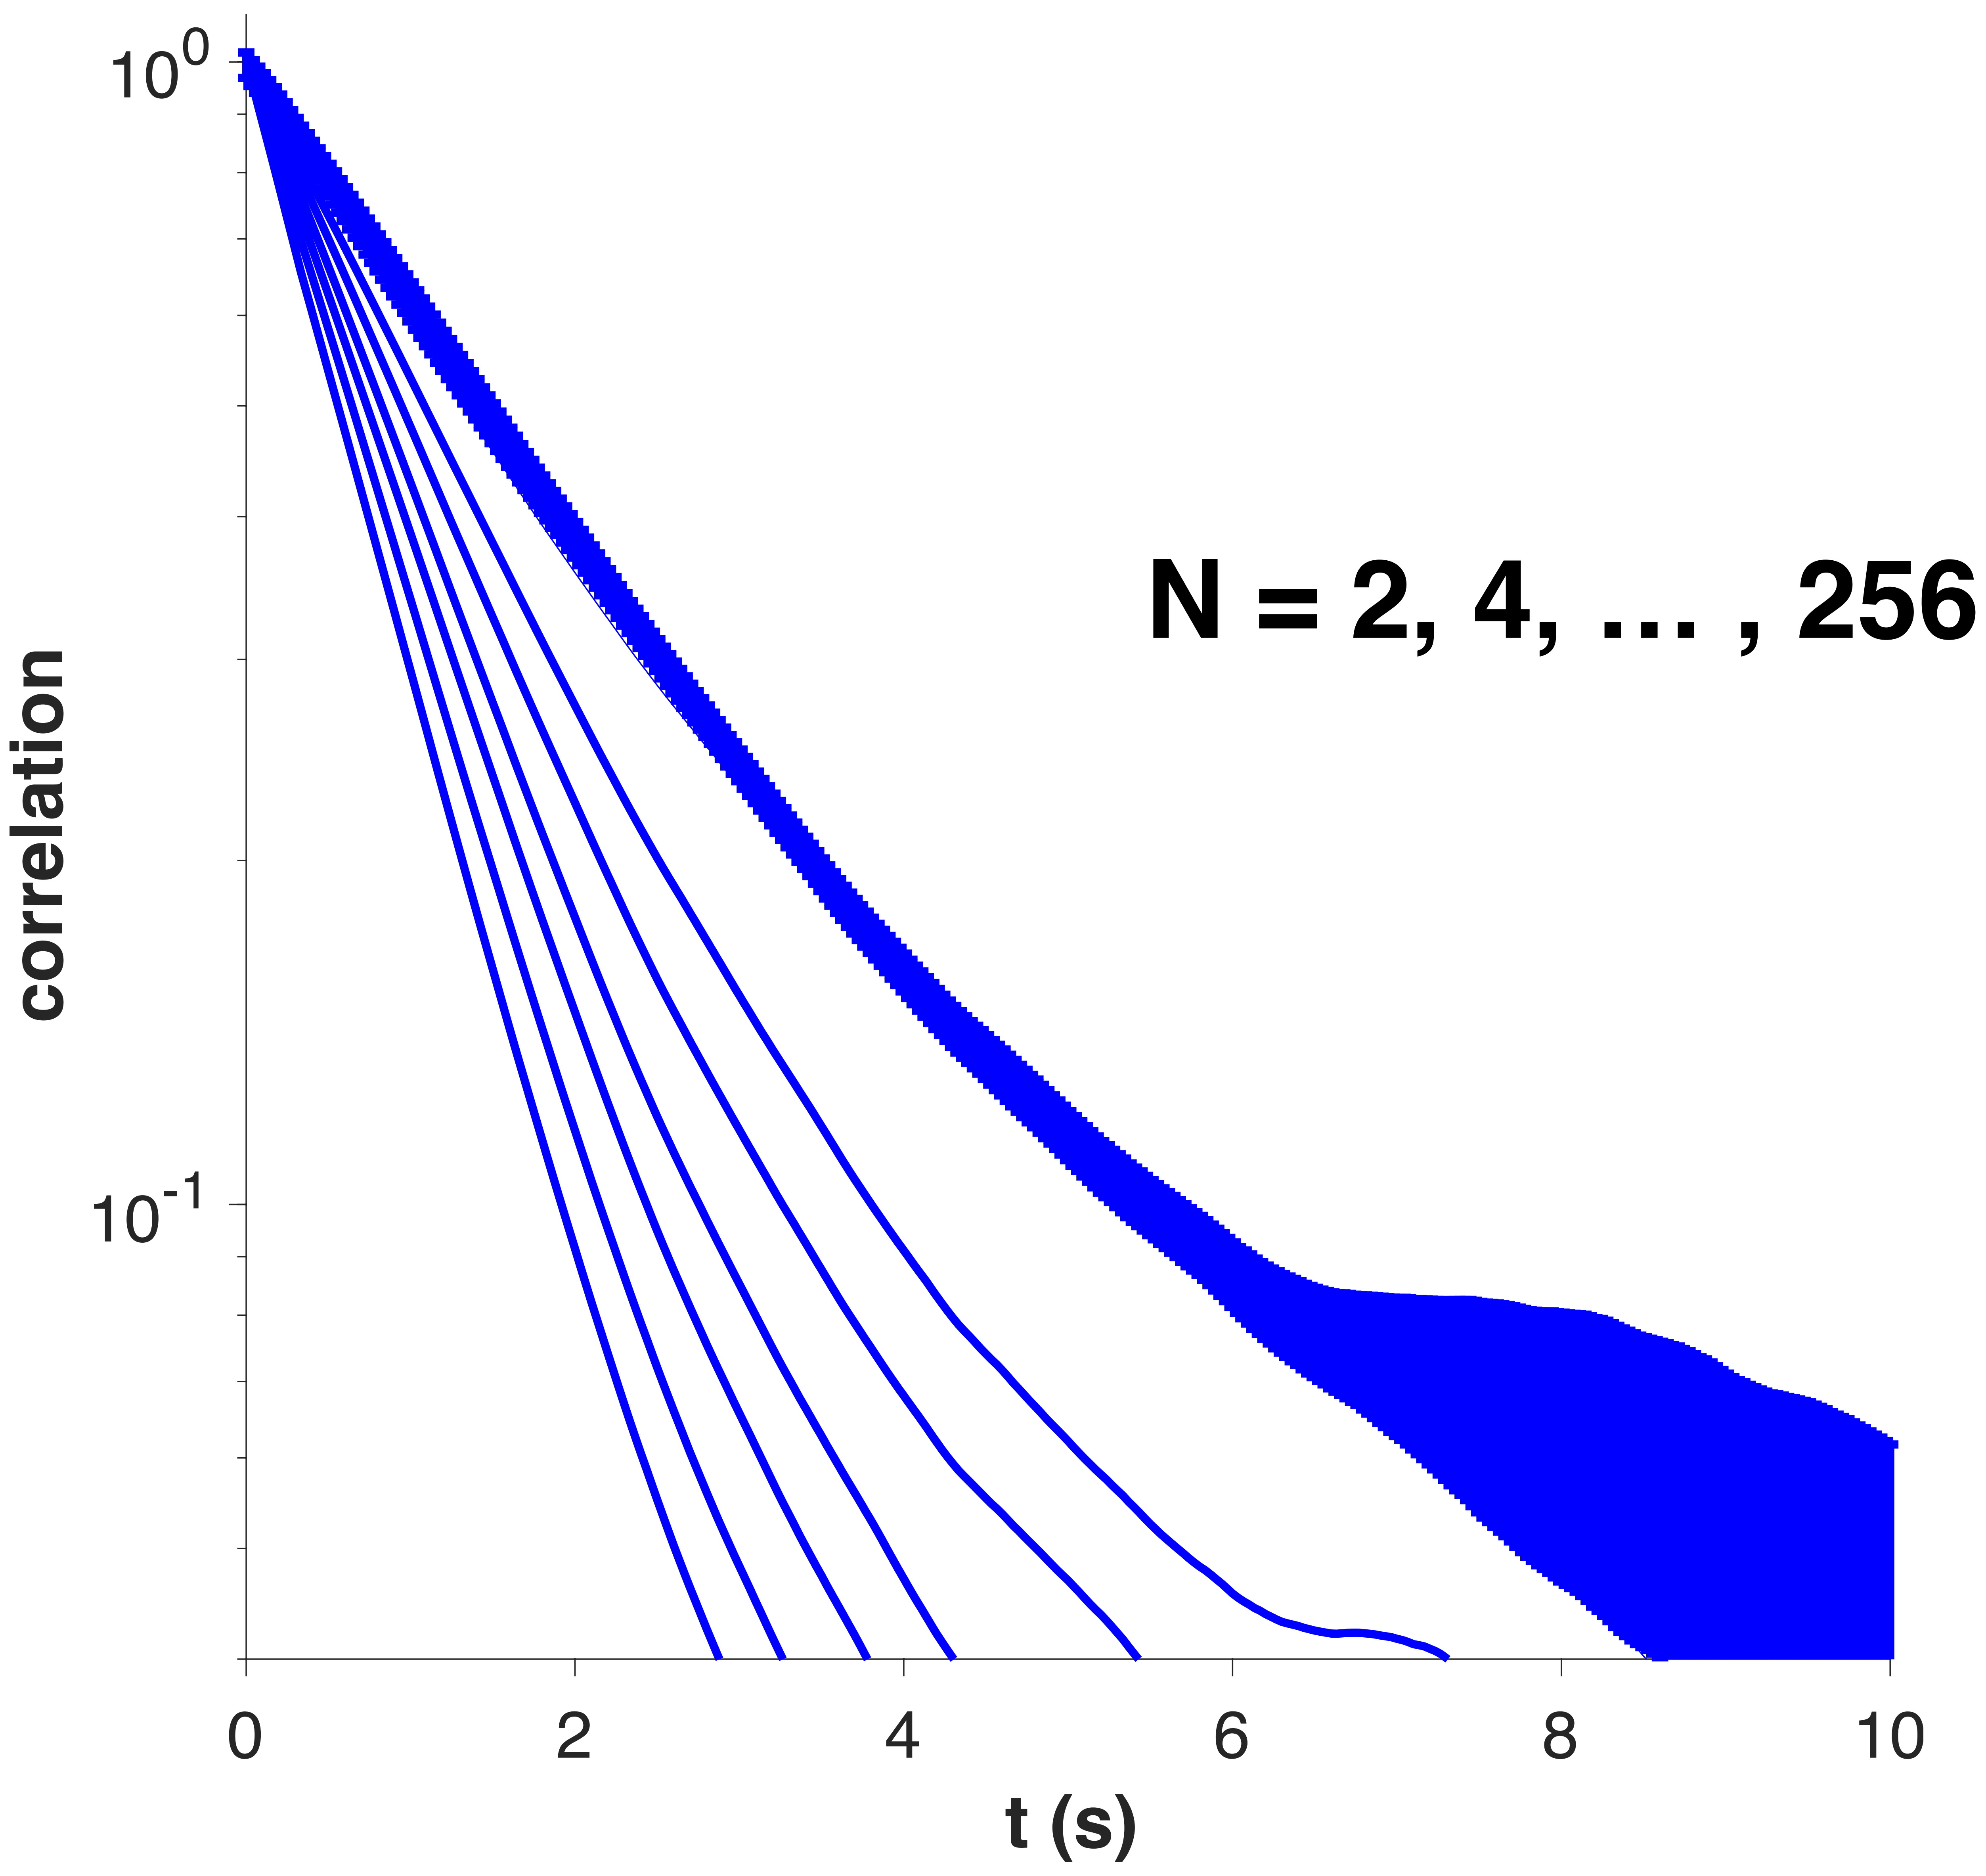

but these dynamics scale

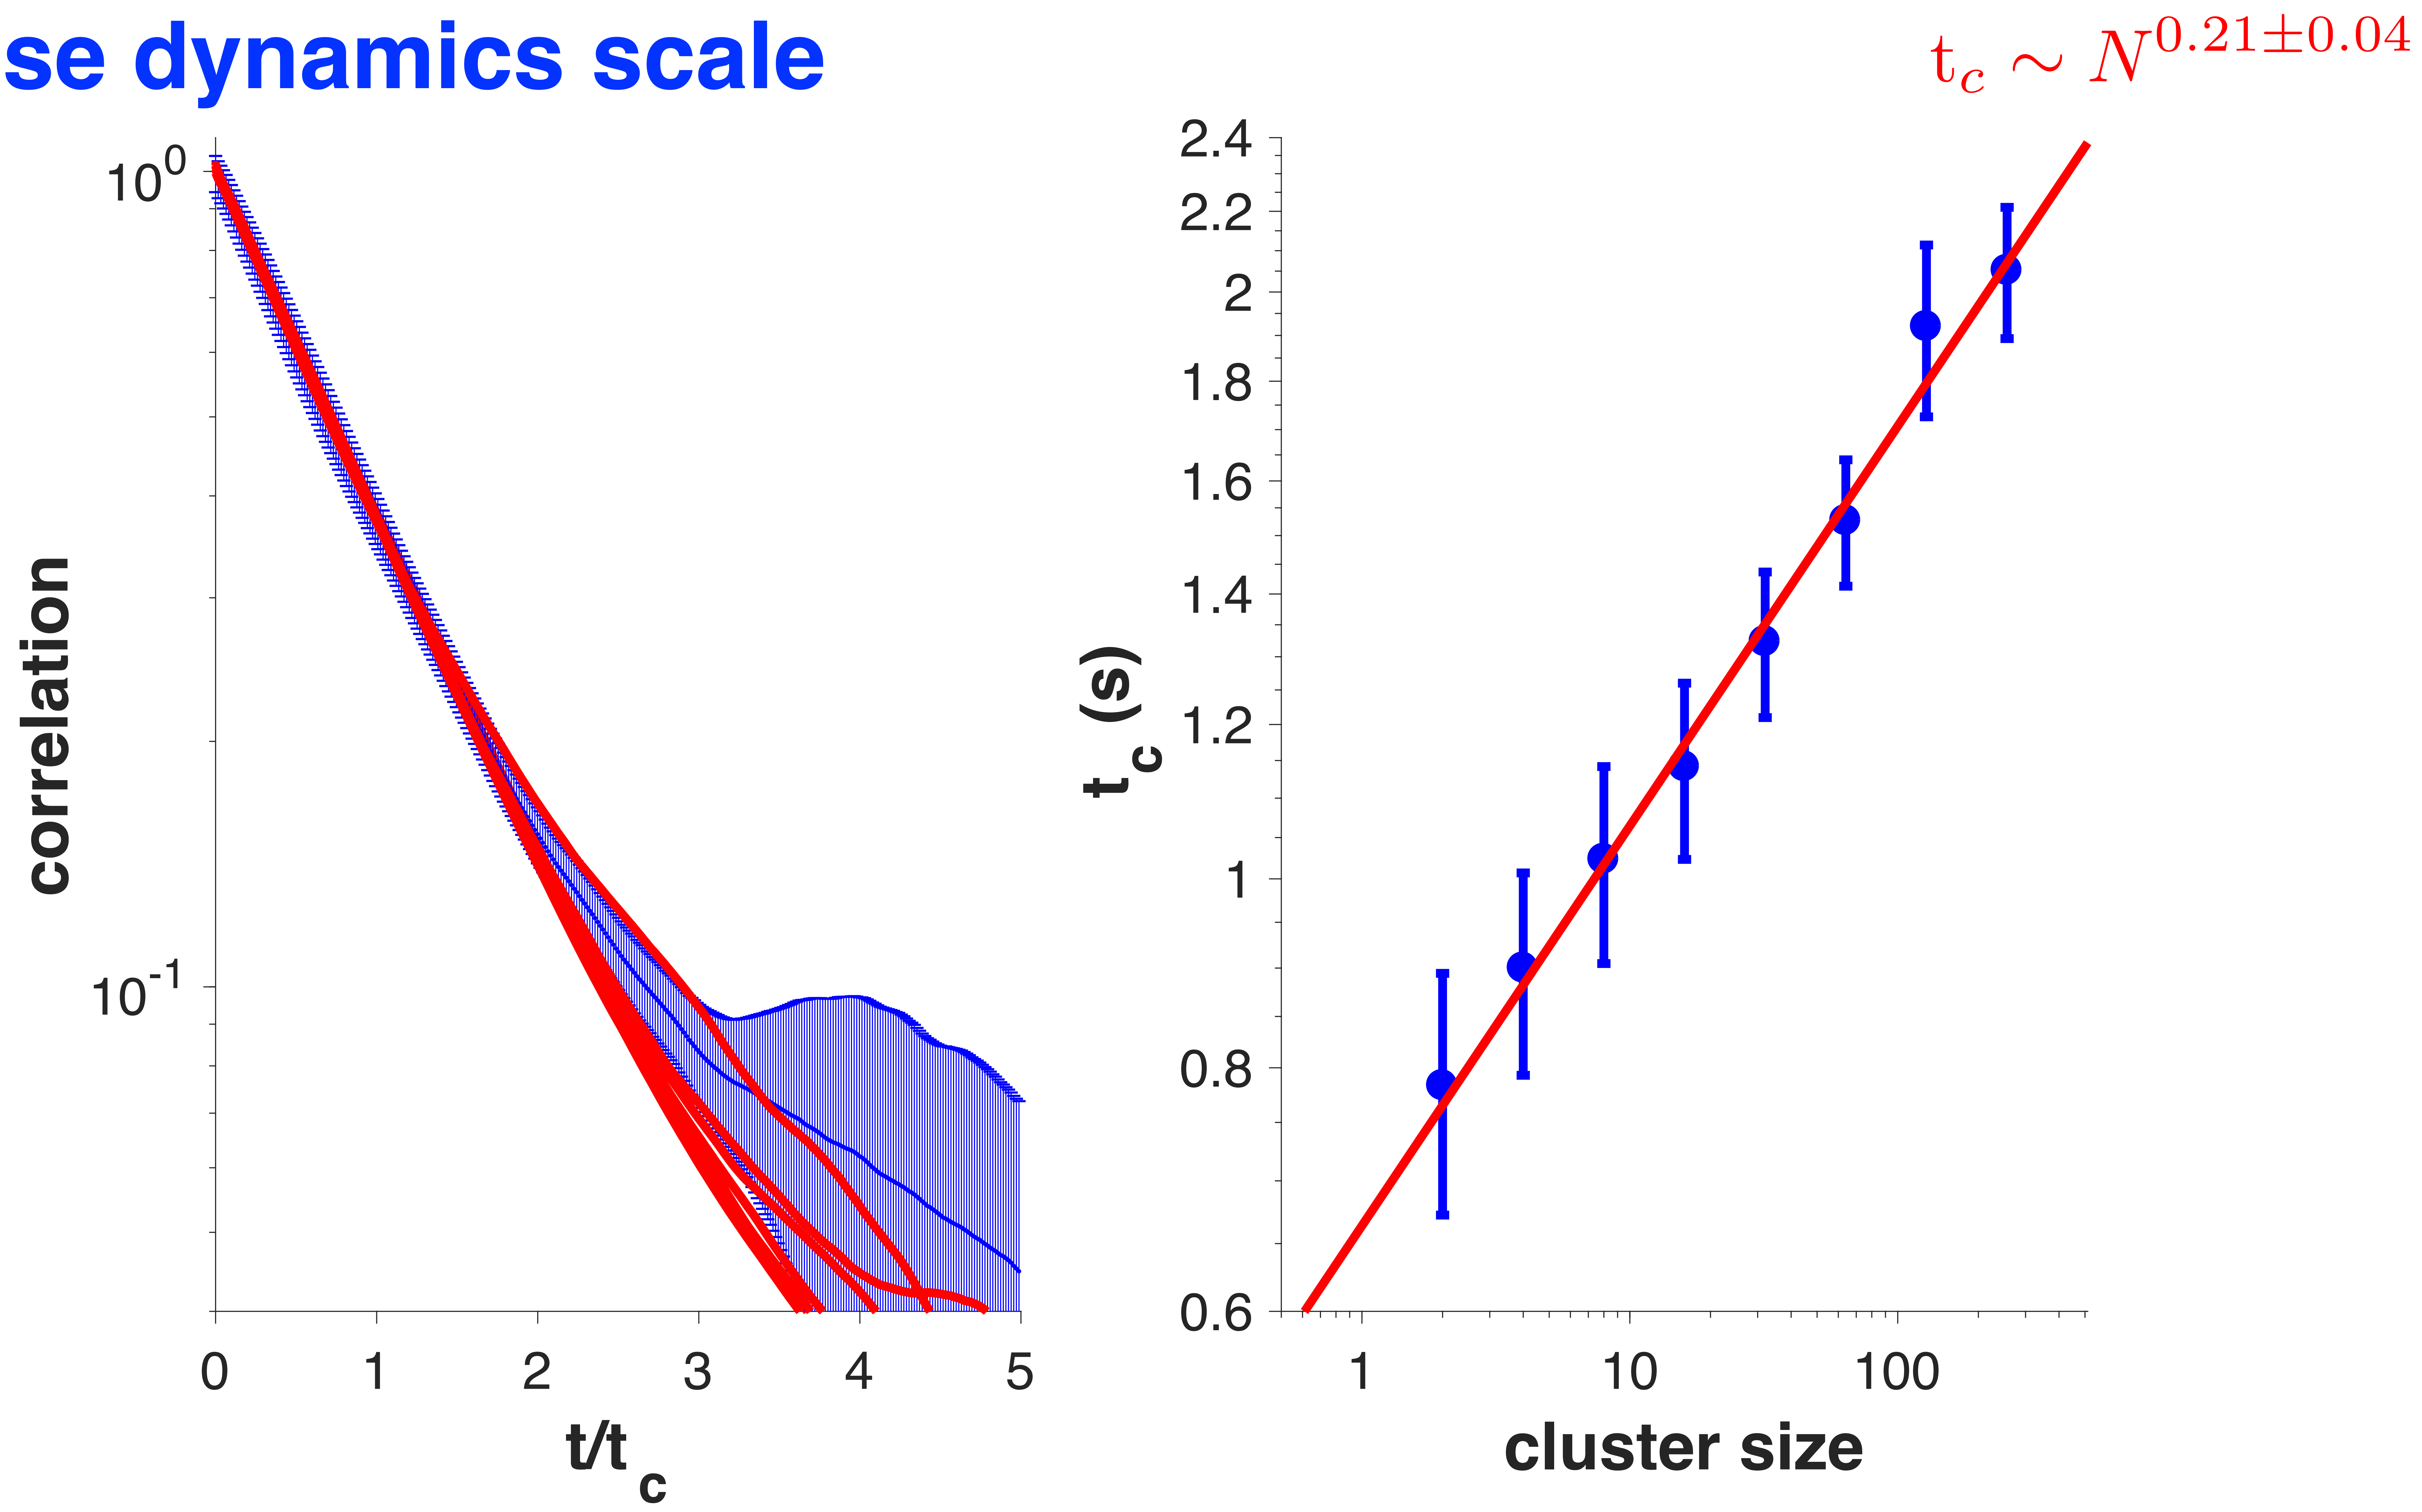

What have we learned?

Coarse-graining the patterns of neural activity leads to simpler, but not trivial, descriptions.

Many characteristics “scale” as a power-law in the number of neurons that we group together.

These results suggest that patterns of neural activity have a surprising self-similarity.

This is not what we expect from “typical” networks.

Path to a fuller theory?

Can we find a model that does

this

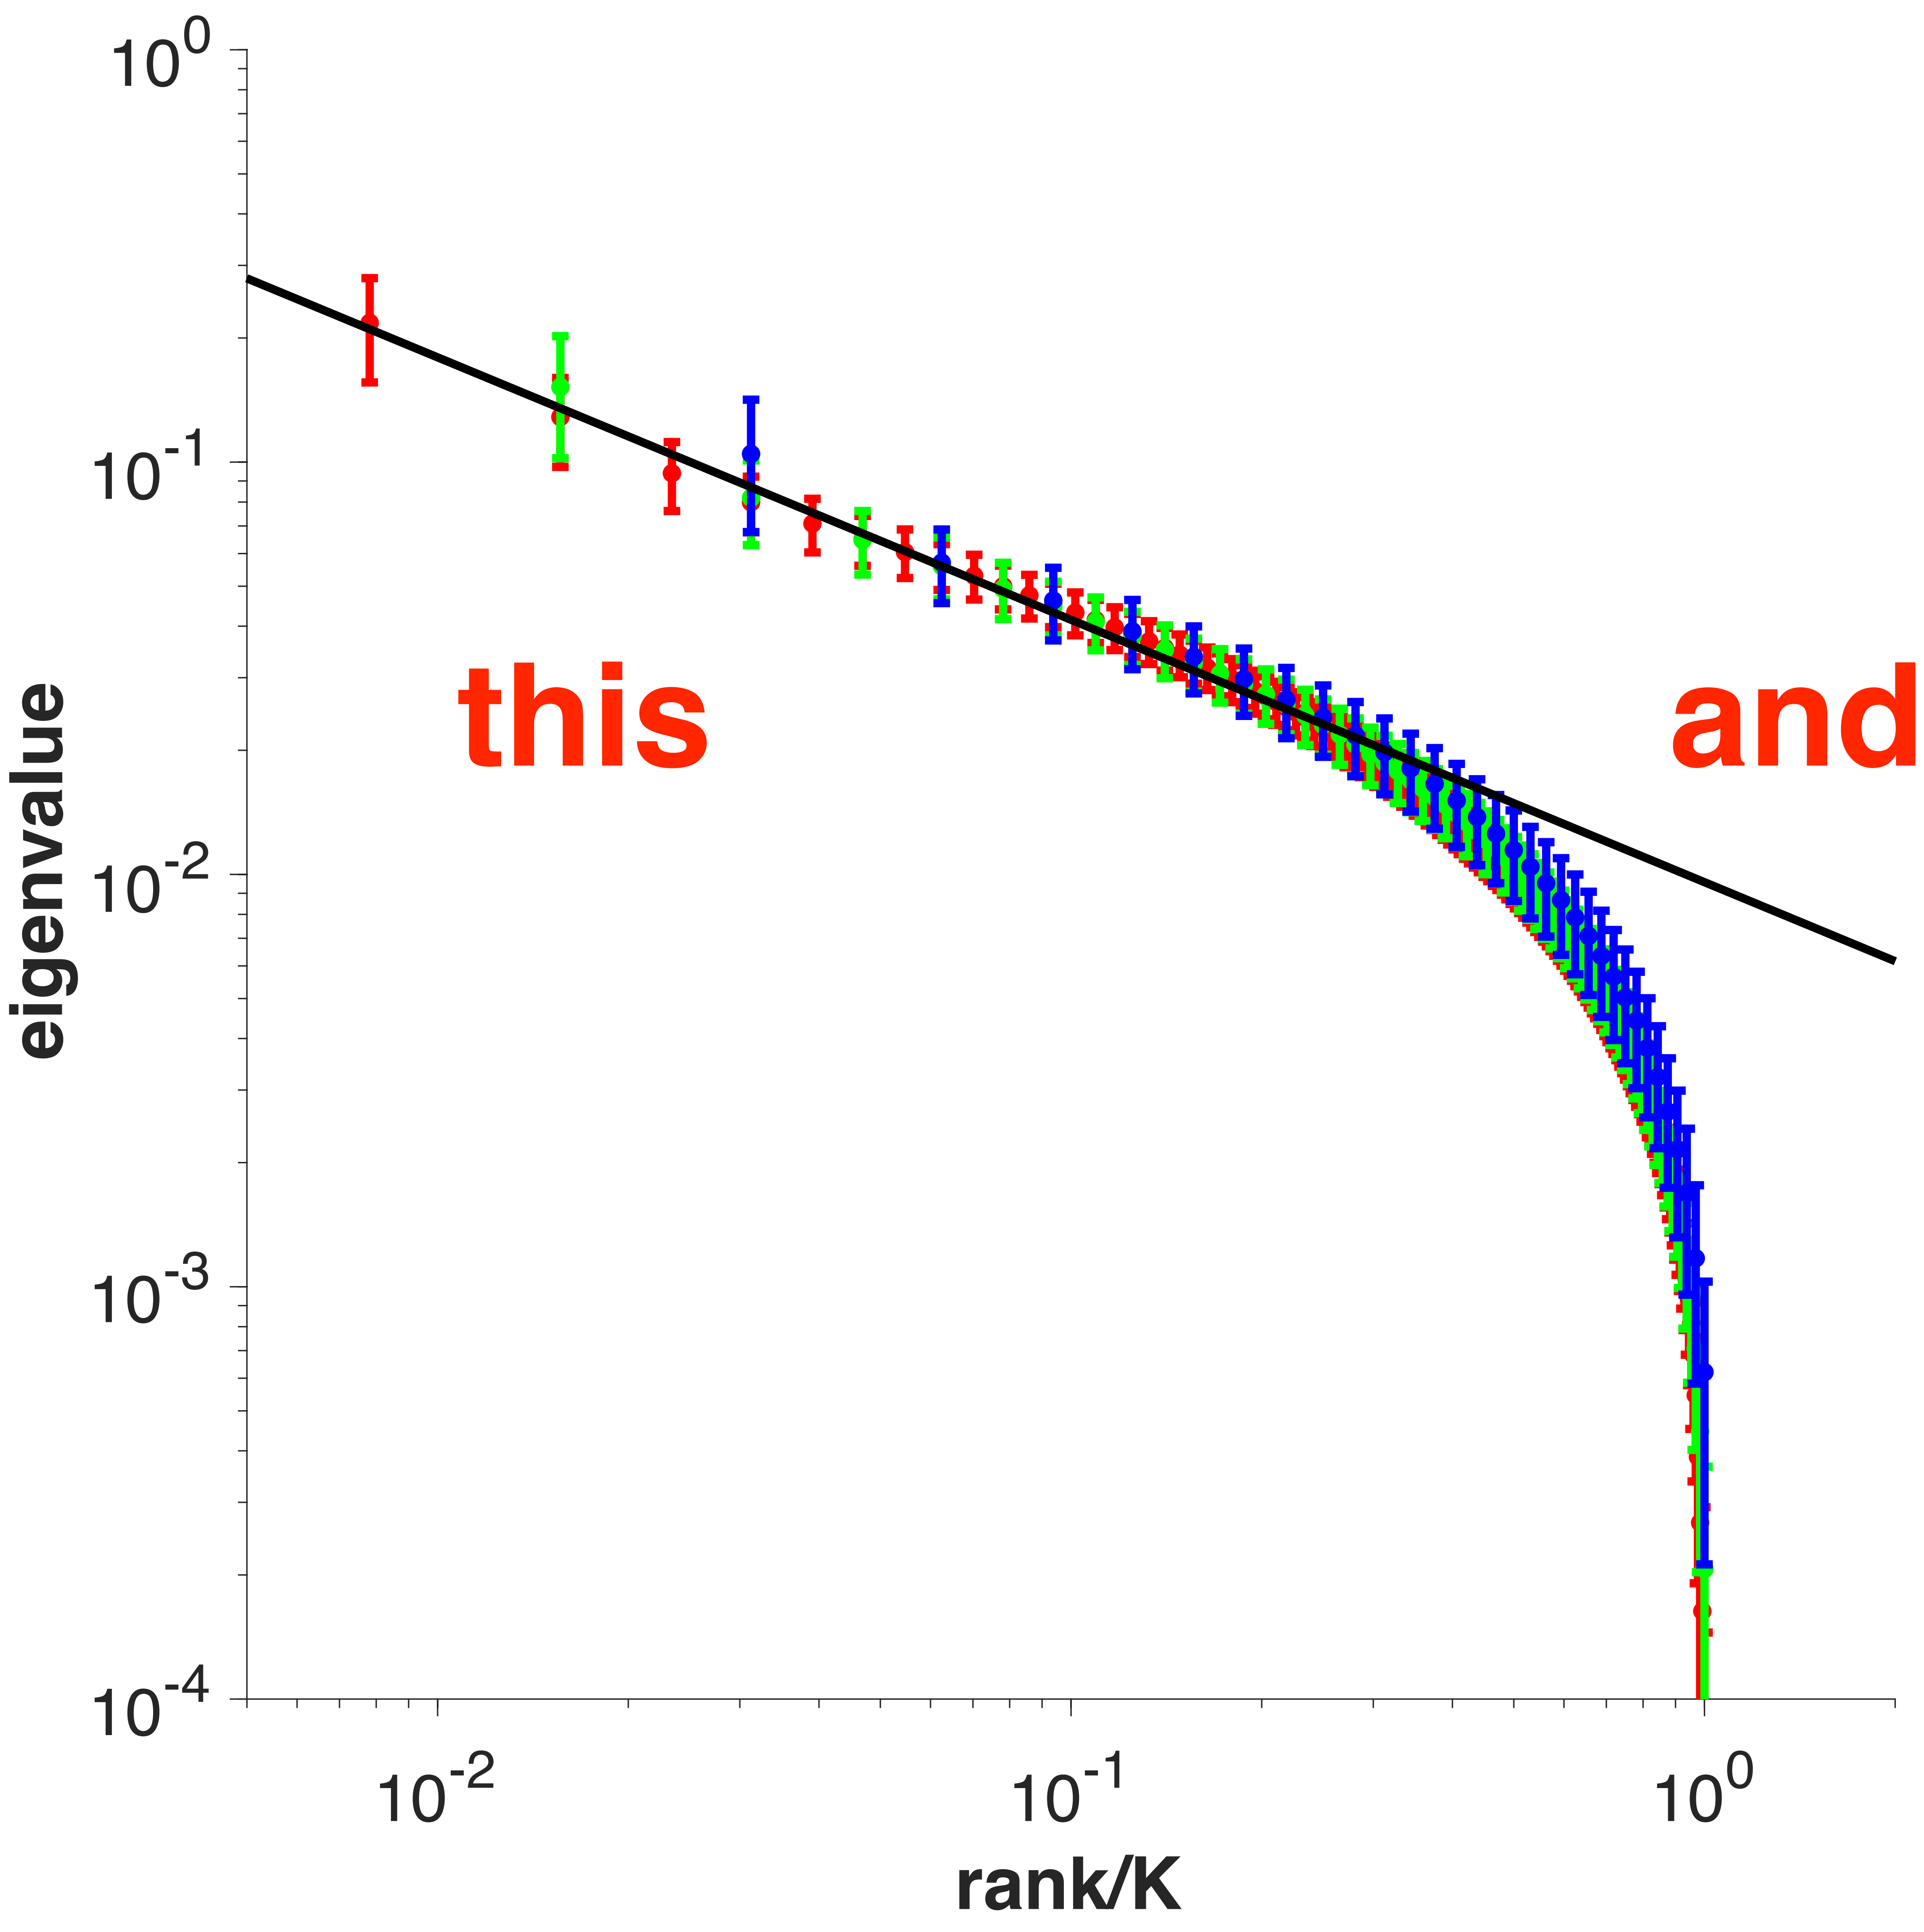

and

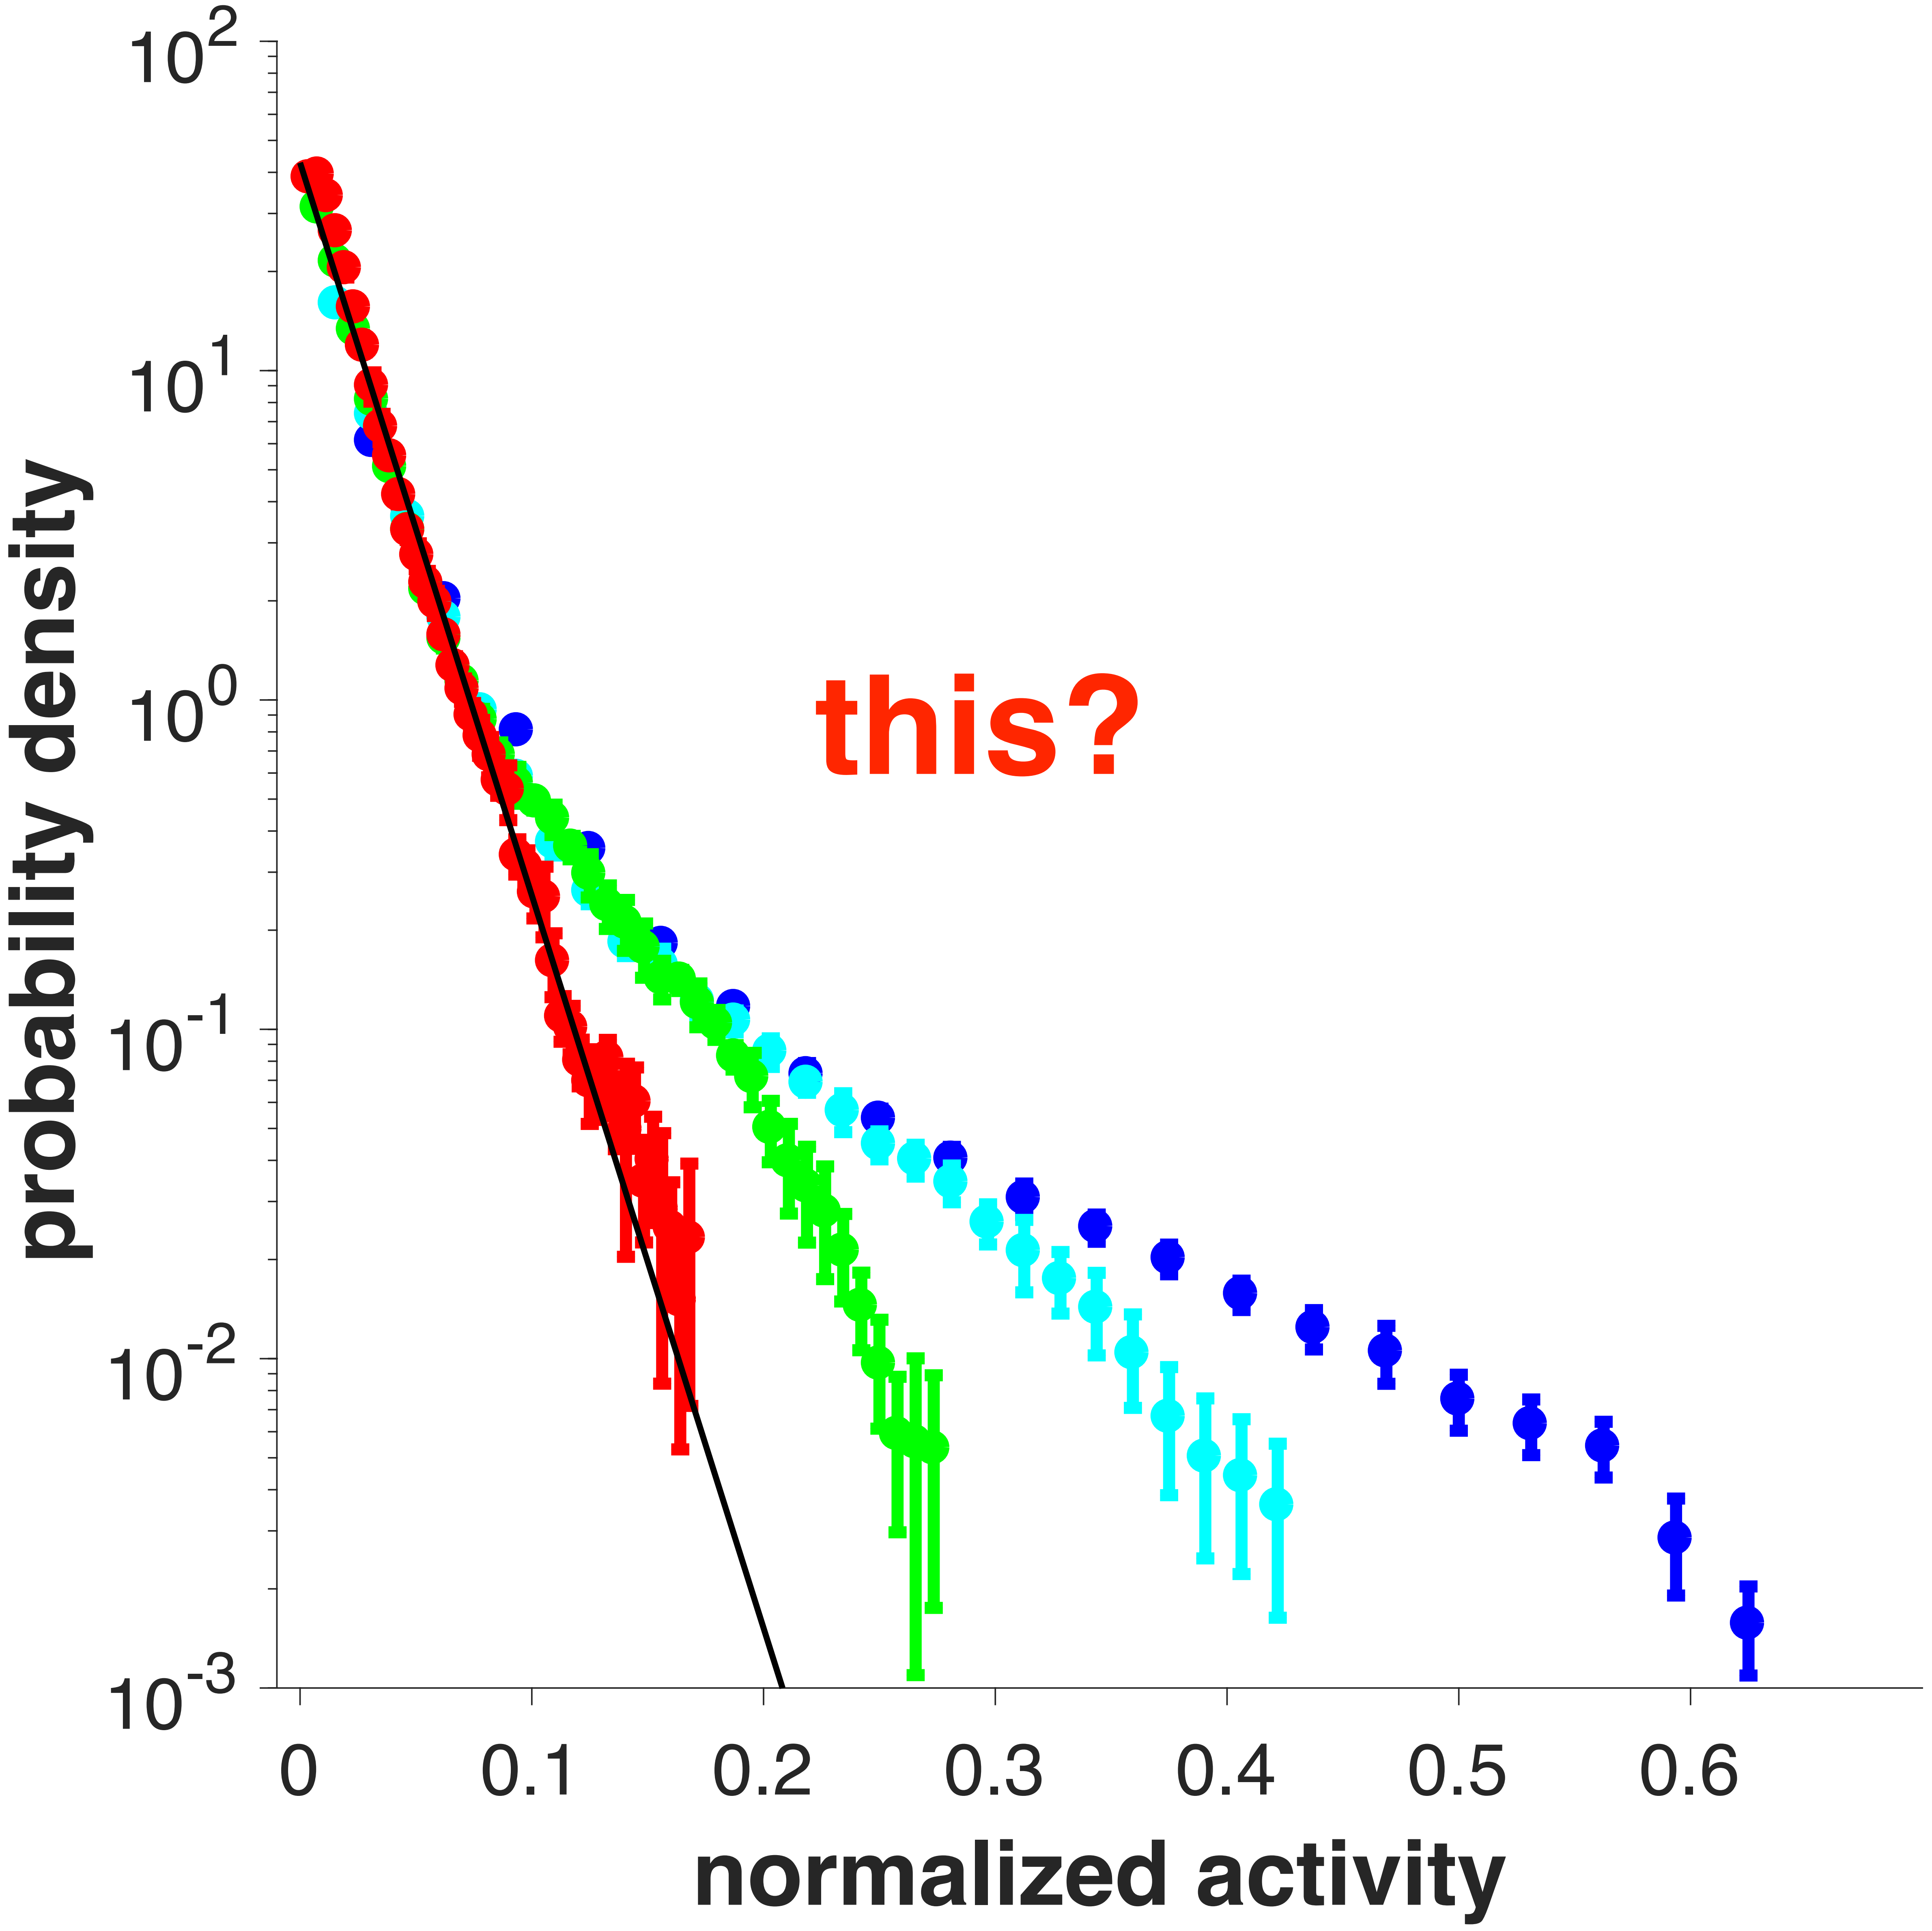

this?