Noise driven synaptic pruning

Rishidev Chaudhuri
Depts. of Mathematics and Neurobiology, Physiology & Behavior
University of California, Davis

MADD seminar
November 12, 2019

Overview

Networks in the brain are sparse

Neural activity is noisy

Brain might use noise to detect and prune redundant connections

Proof: graph sparsification + dynamical systems

Graphs and networks in neuroscience

Observation 1: Connections in the brain are sparse

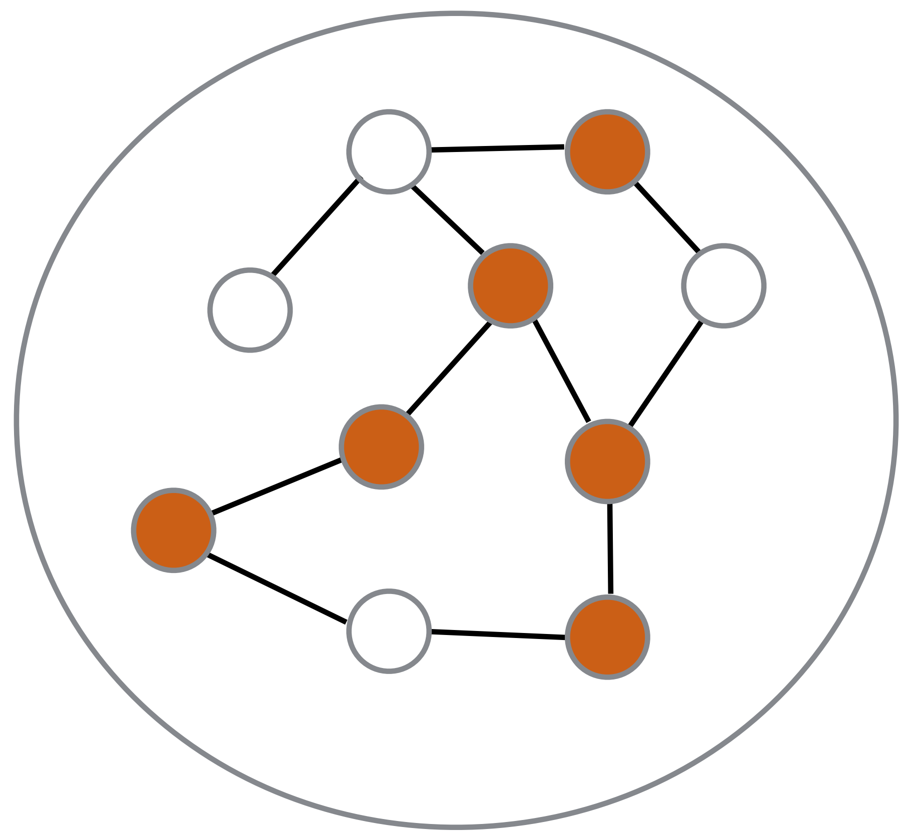

Millions of neurons, but most do not talk to each other

Why study sparse networks?

Sparse = efficient (compare carbon footprint of dense networks in AI)

Connection density disrupted in disease

Many networks need to be coordinated with a small number of interactions

Brains are cool

Possible source of pretty math

Questions about sparse networks?

What principles underlie information flow in these sparse networks?

(one idea, sparse expander graphs)

What are the computational advantages of sparse networks in the brain?

(one idea, credit assignment. e.g., Chaudhuri & Fiete NeurIPS 2019)

How does the brain find a sparse network that performs a given task?

(one idea, noise-driven pruning)

Brain seeks out sparsity

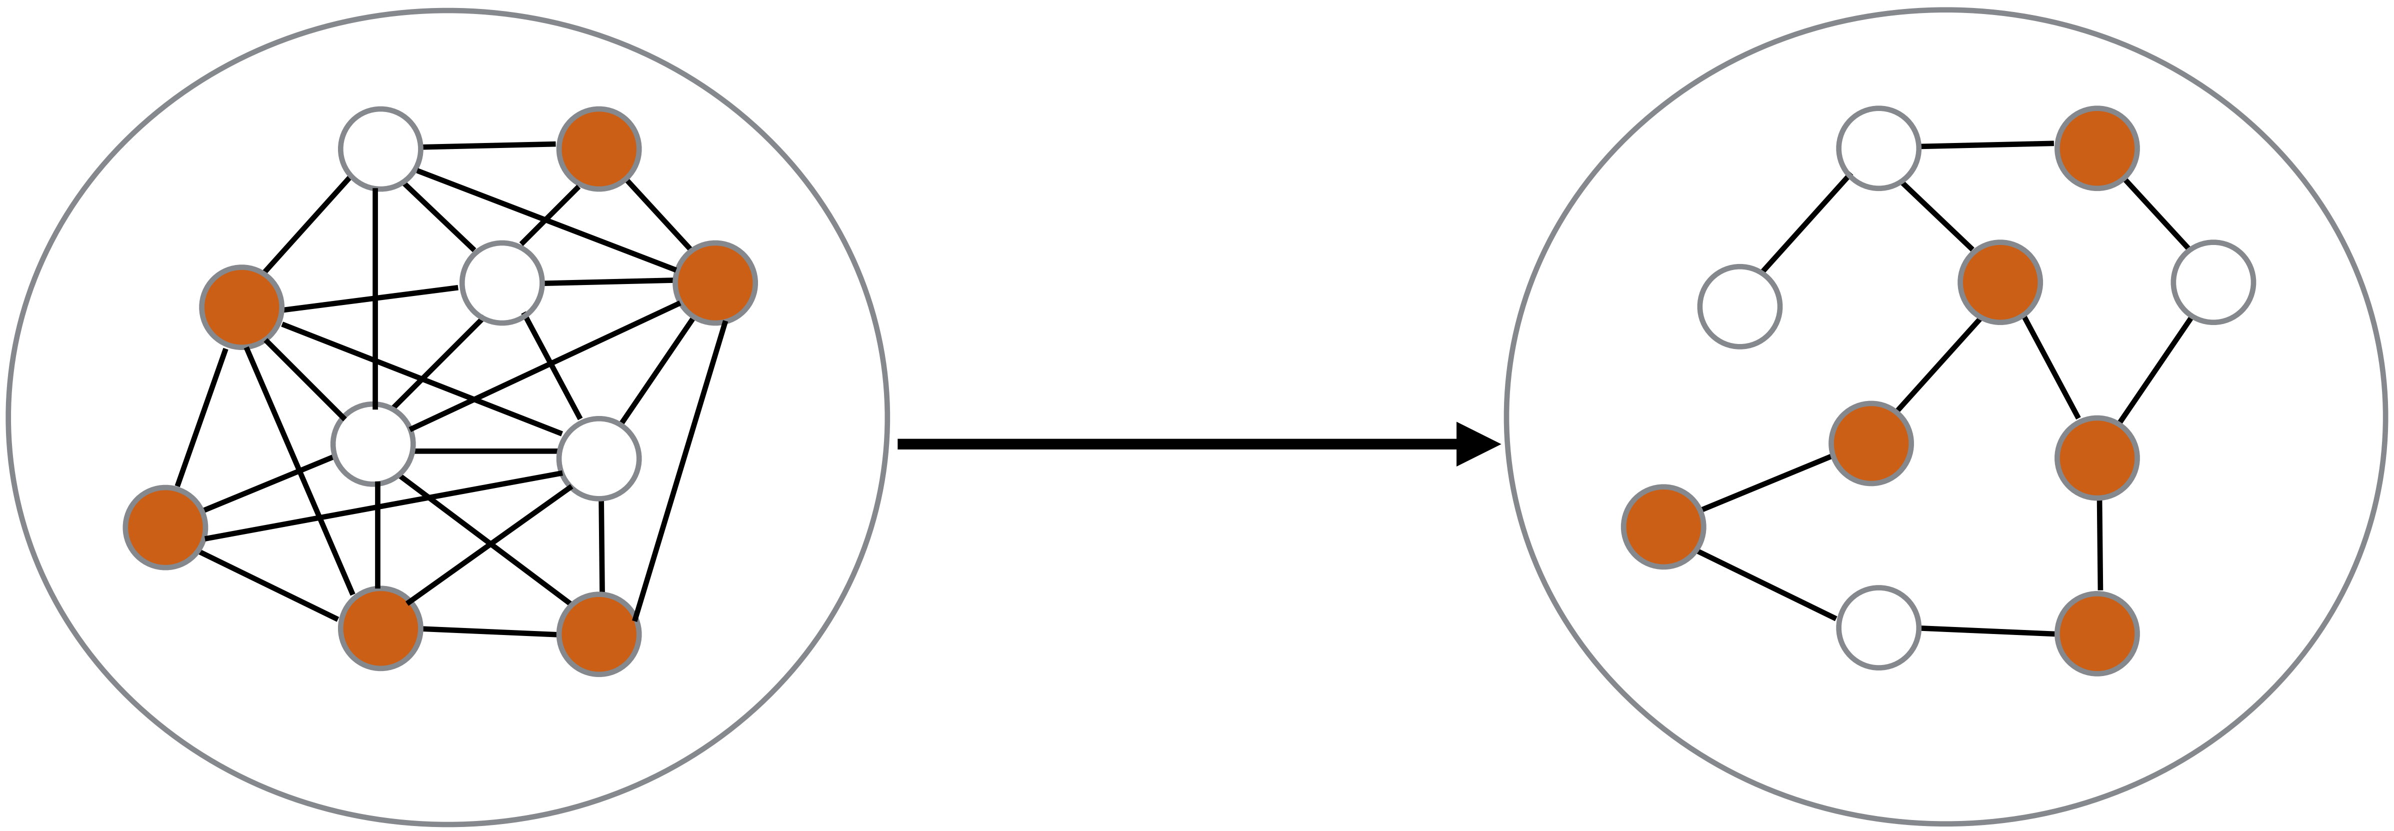

Synapses pruned during development, learning and possibly sleep

Q: How could the brain prune synapses?

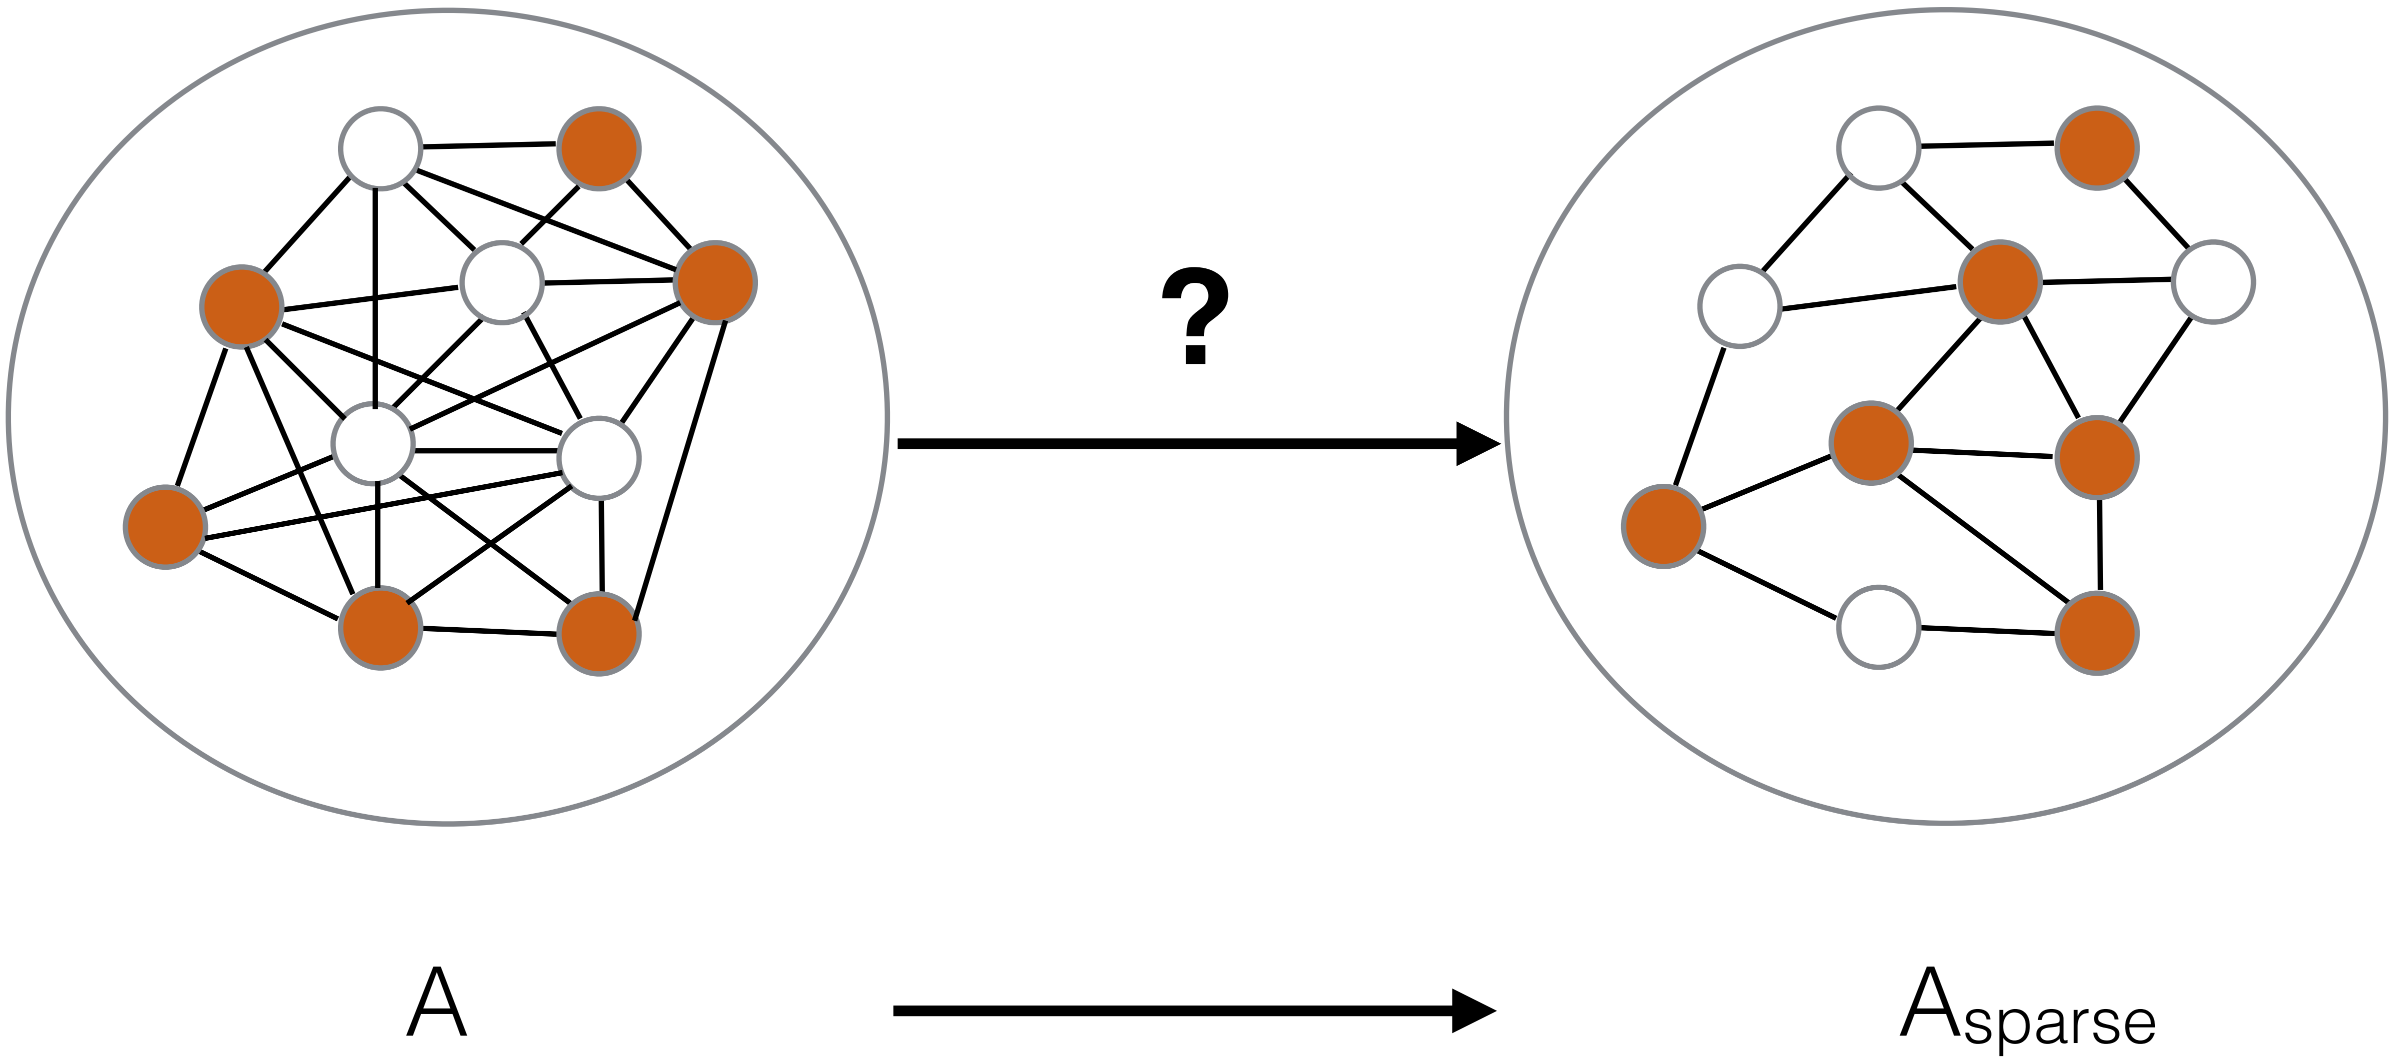

Preserve properties of A or of dynamical system on A

Need to figure out which synapses are redundant

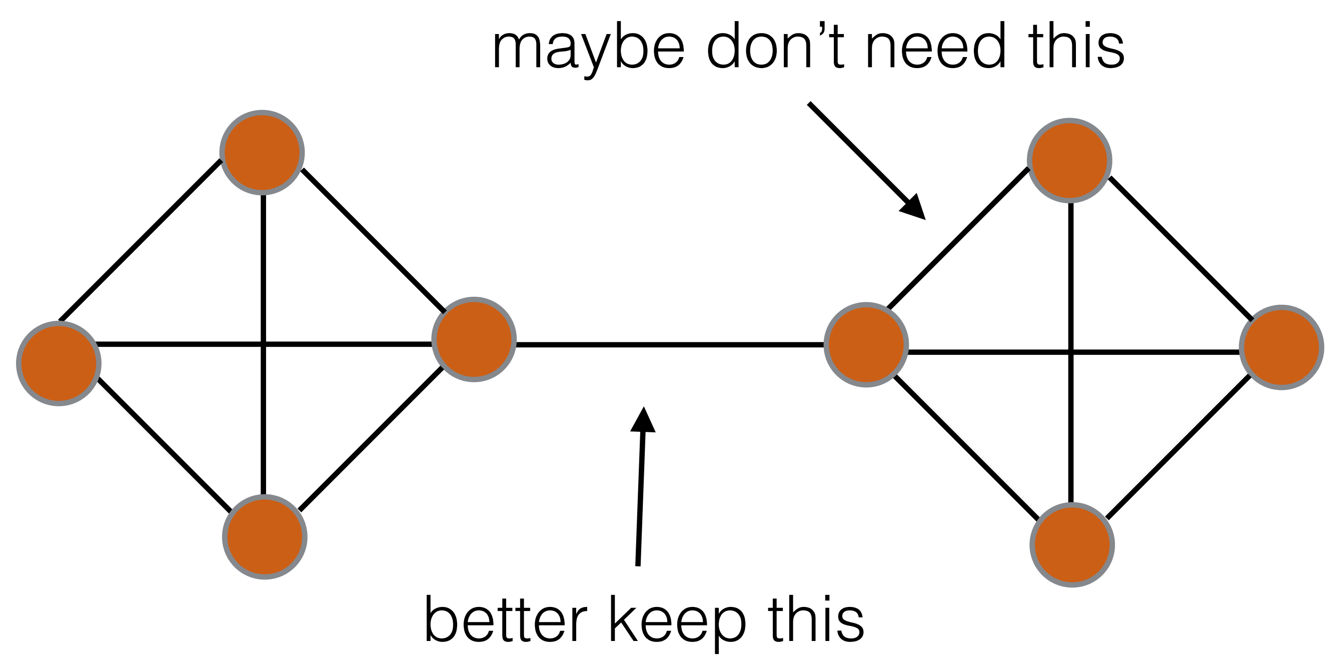

Hard problem: depends on every possible route through the network

How do two neurons, somewhere deep in the brain, figure out their indirect connections?

Must use local information

Observation 2: brains are (seemingly) noisy

Noisy activity

Noisy state

Random connectivity

Why are brains noisy?

Biology can be extraordinarily robust and precise when it wants (though at a price)

Possible answers:

1) Brains are not noisy

2) Randomness averaged away

3) Randomness repurposed for computation

4) Brains actively generate randomness

Various questions related to randomized computation in the brain

Idea: neural networks could use noise to probe and prune connections

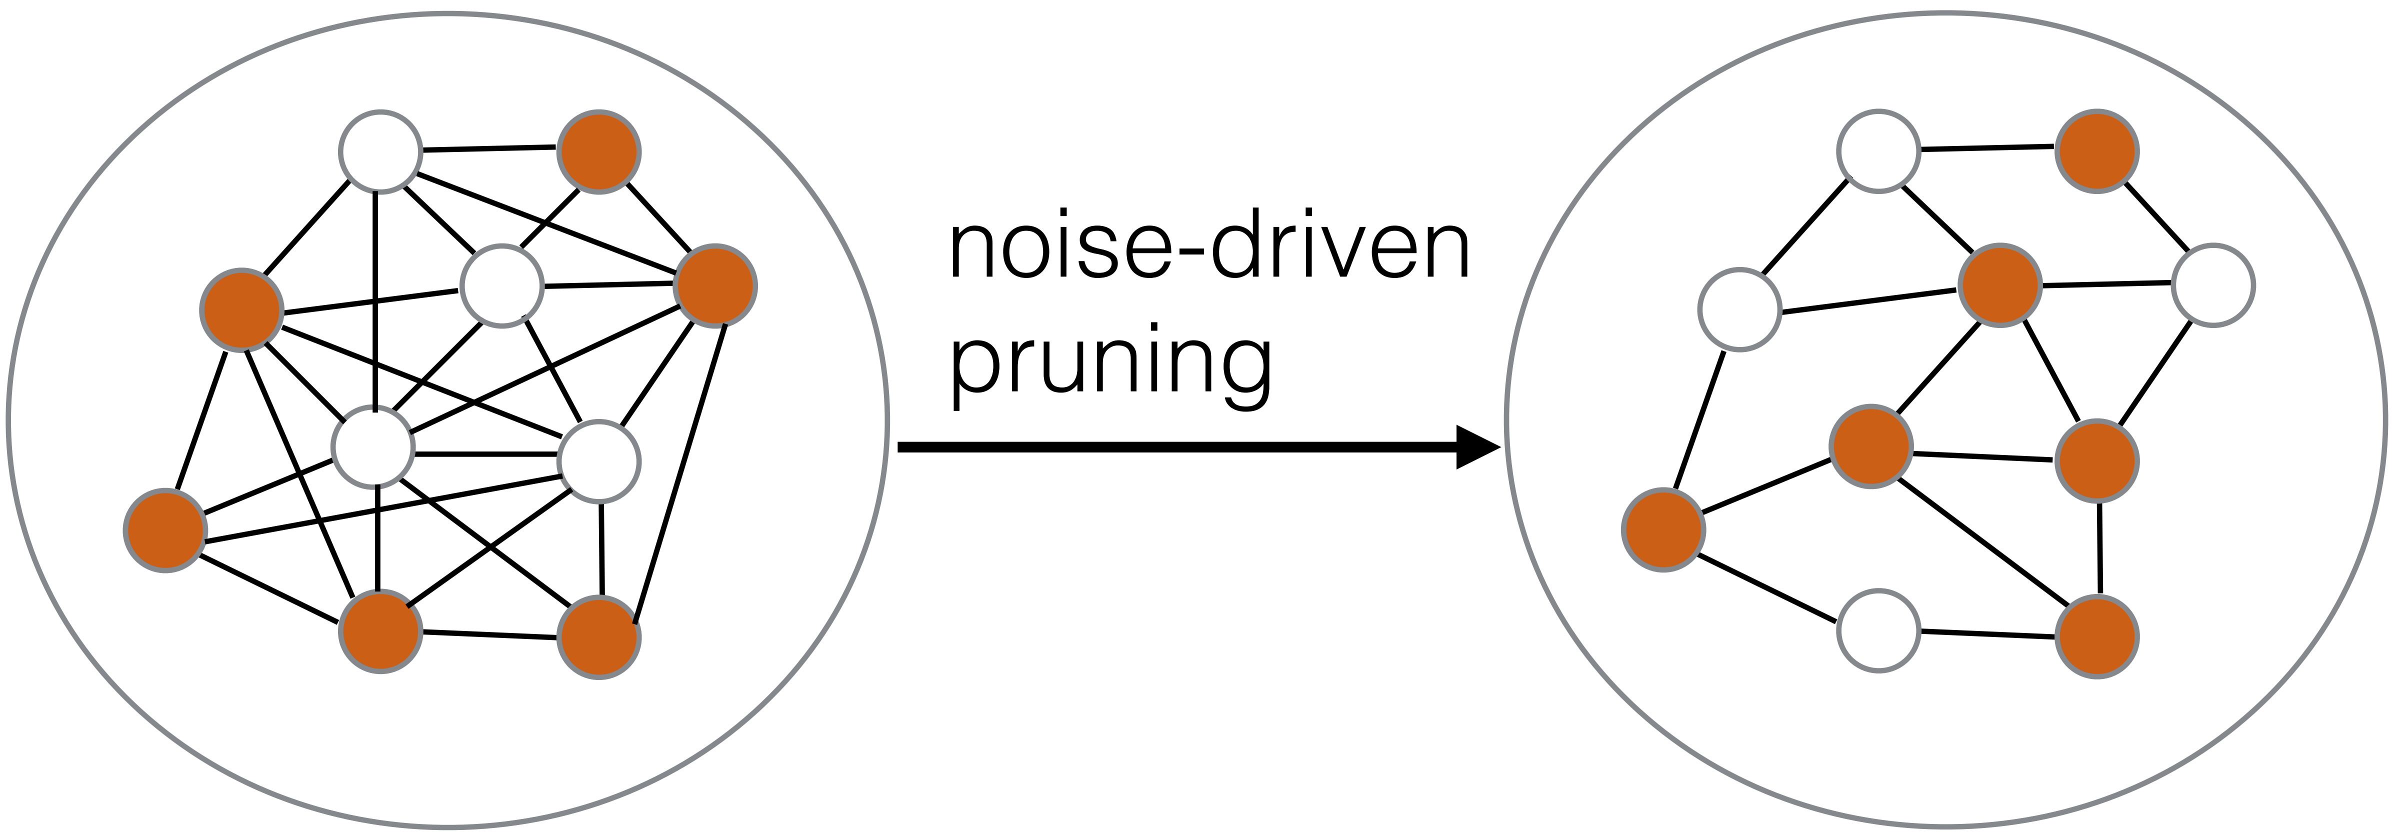

How? Neurons play “telephone”
(or spread rumors about each other)

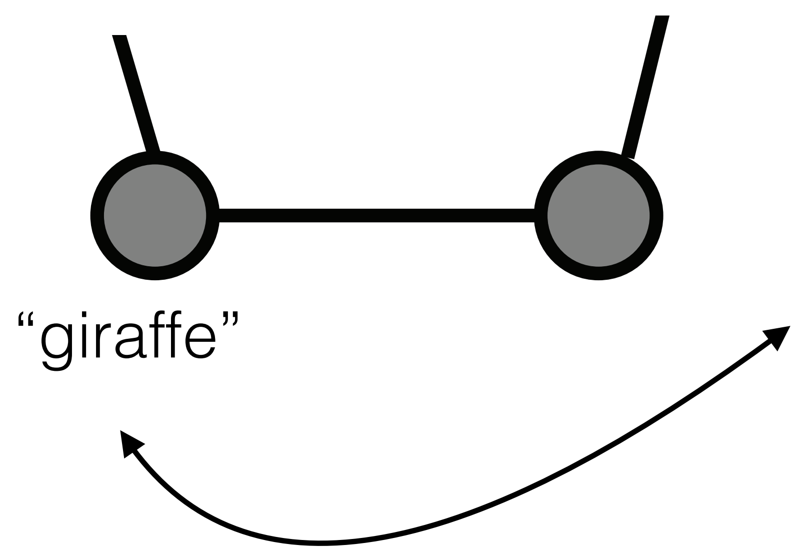

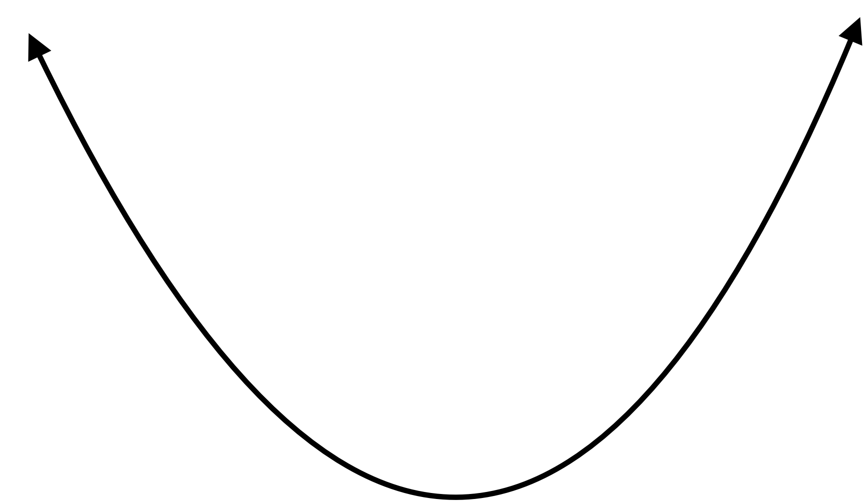

connection important

connection redundant

Use noise to probe and prune network

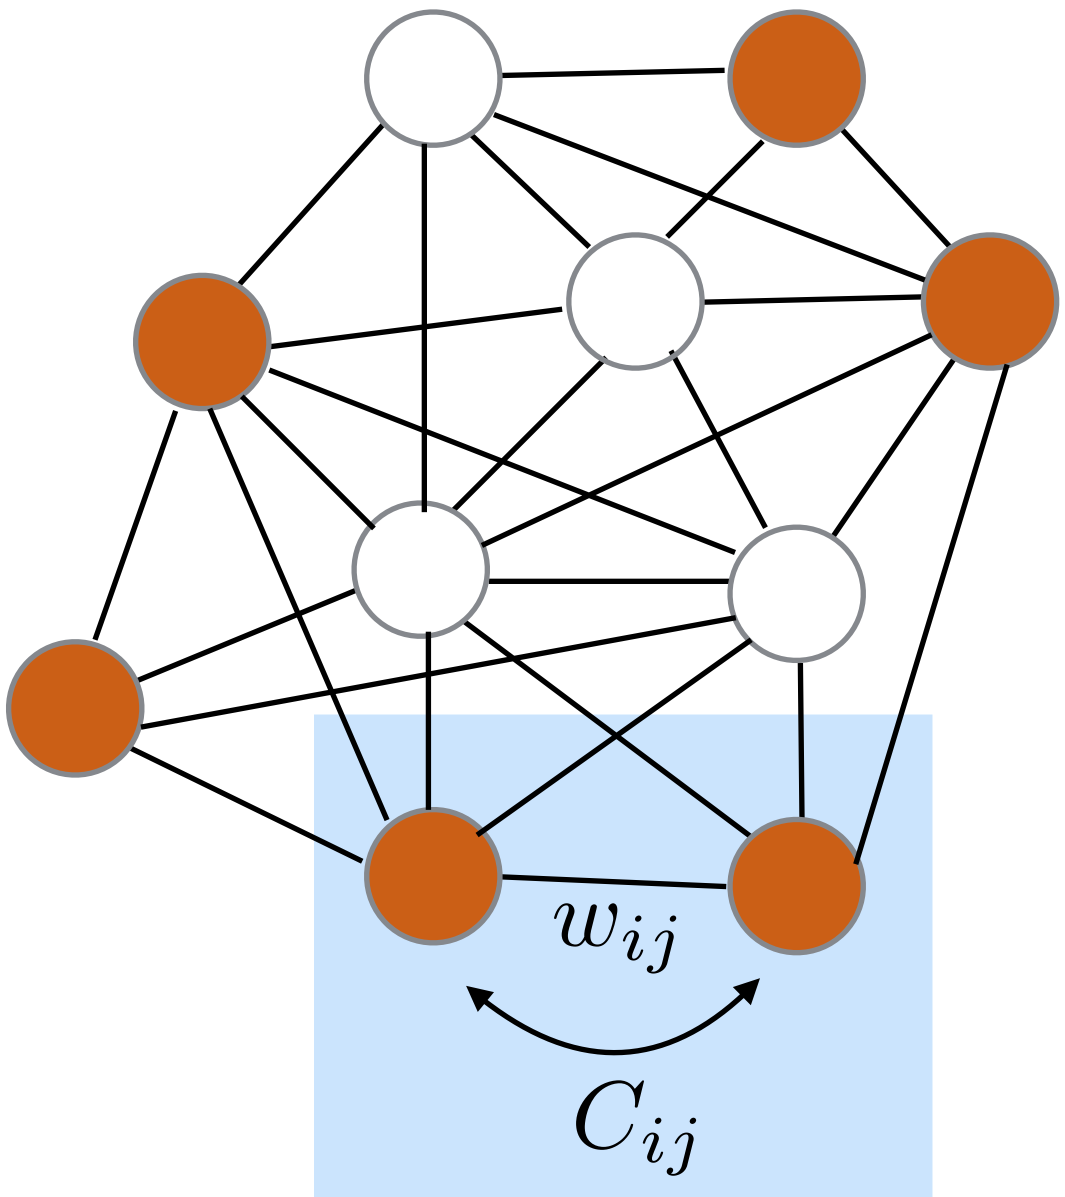

1. Drive network with noise
2. Calculate correlations
3. Compare correlations to connection weight
4. More correlated than weight suggests? Probably redundant

Learning rule:

preserve synapses with probability

excitatory: $p_{ij} \propto w_{ij} \left(C_{ii} + C_{jj} - 2C_{ij}\right)$

inhibitory: $p_{ij} \propto |w_{ij}| \left(C_{ii} + C_{jj} + 2C_{ij}\right)$

Learning on weights + correlations = Hebbian, ubiquitous

Unsupervised, task-agnostic

Linear model of neural dynamics

$$\frac{dx}{dt} = (-\mathbb{I} + W)x + \text{input}$$
$$= Ax + \text{input}$$

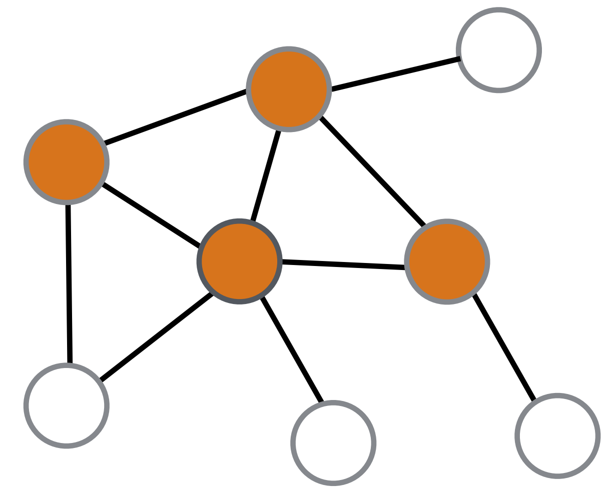

$$x = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{pmatrix} \in \mathbb{R}^N$$

activity

W = connection weights

-I to reflect activity leak

for stability:

$\lambda_{max}(-I + W) < 0$

Pruning rule

$$\frac{dx}{dt} = Ax + \xi(t)$$

white noise or OU with
short correlation time

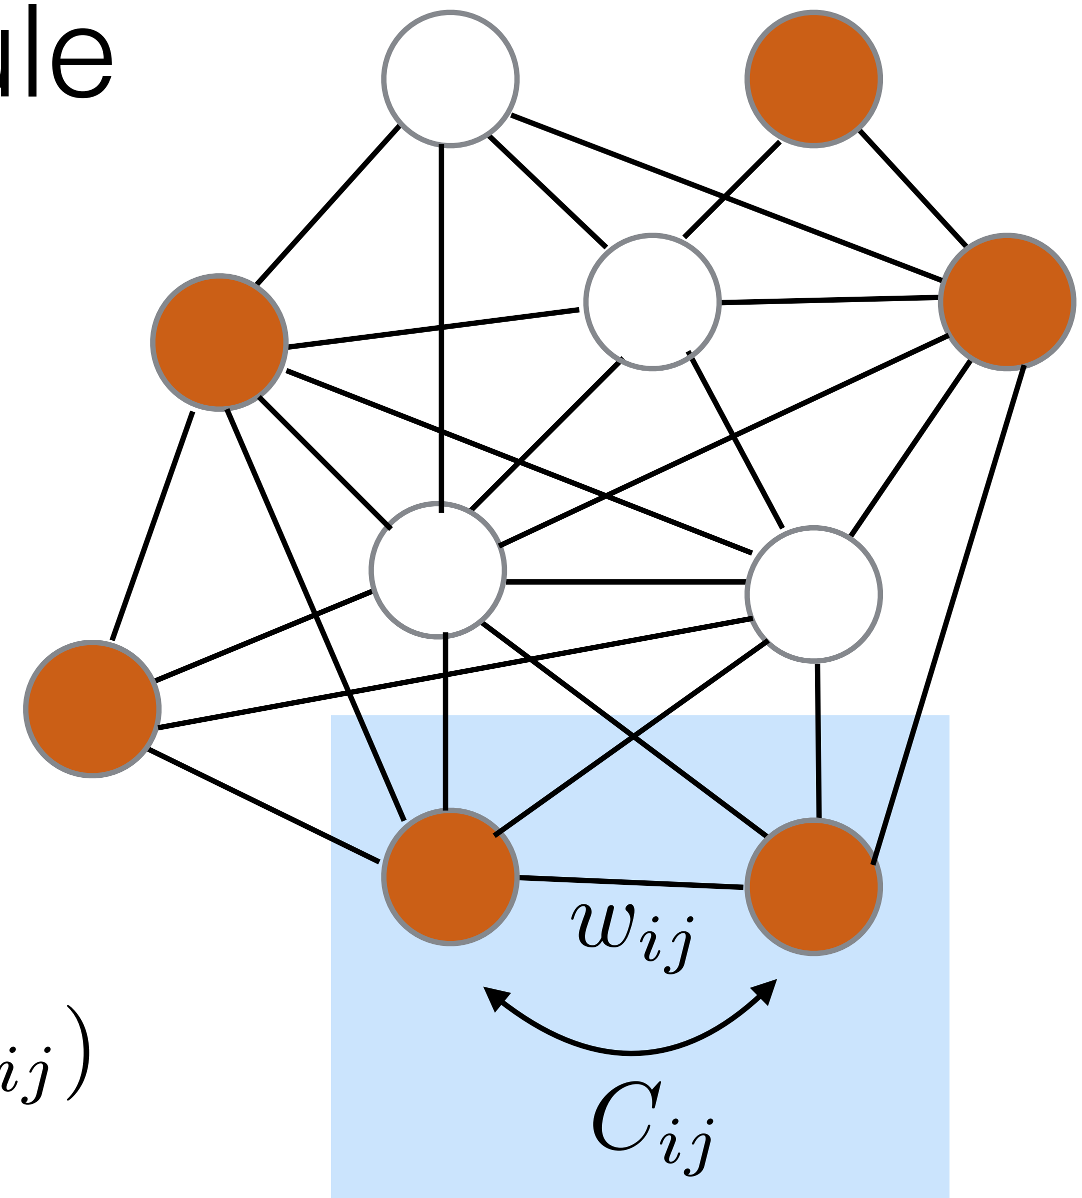

$$p_{ij} \propto |a_{ij}| \left(C_{ii} + C_{jj} - 2\text{sign}(a_{ij}) C_{ij}\right)$$

$$A_{ij}^{\text{sparse}} = \begin{cases} A_{ij}/p_{ij} & \text{with probability } p_{ij} \\ 0 & \text{otherwise} \end{cases}$$

More generally: synapses transition between stable and labile state

Pruning rule identifies important synapses

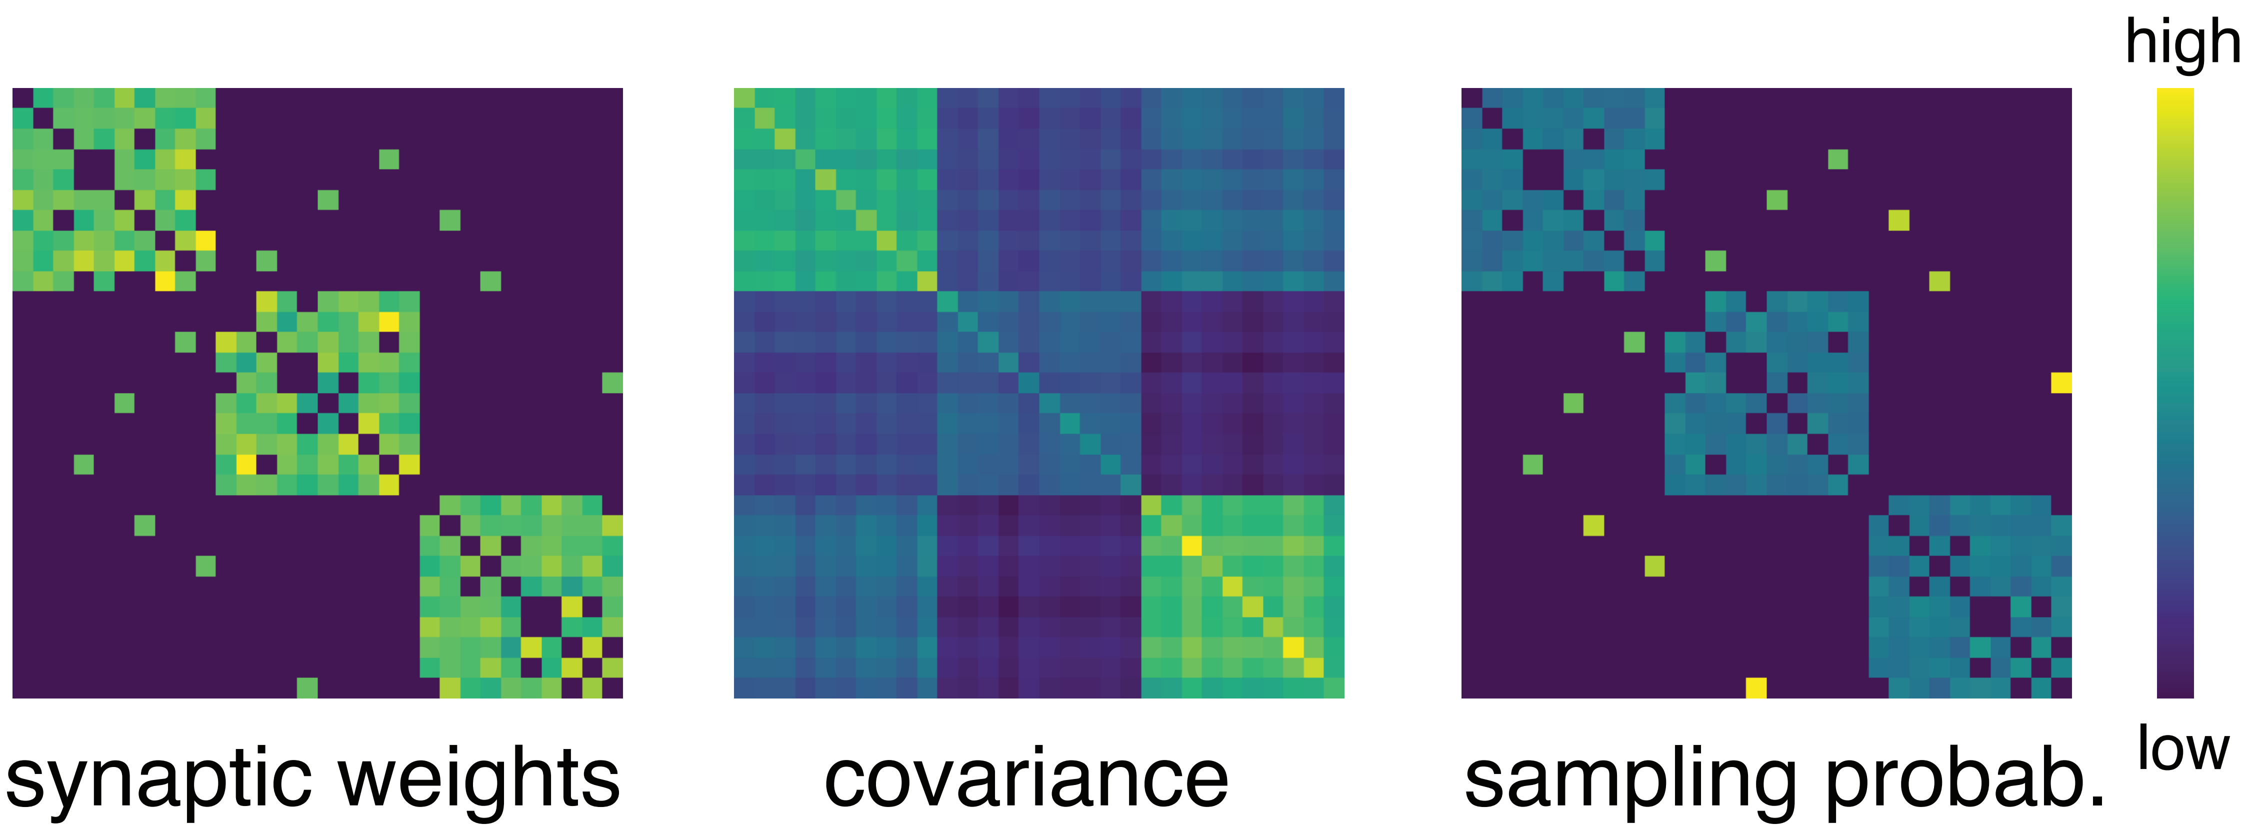

Pruning rule provably preserves important features for a subset of linear systems

Combine ideas from graph sparsification and dynamical systems

Sparsify graph by sampling and reweighing a subset of edges (Spielman & Srivastava, 2011)

Key quantity is set of sampling probabilities

Show that covariance of noise-driven dynamical system encodes sampling probabilities

Sparsification of neural network matrices

A is N x N symmetric, diagonally-dominant, positive diagonal $|a_{ii}| > \sum_i |a_{ij}|$

Let $\epsilon \in (0,1)$ and $k > 3$ Set $K = \dfrac{k \log(N)}{\epsilon^2}$

For each edge, define $p_{ij} = K|a_{ij}|\left(C_{ii} + C_{jj} - 2\mathrm{sign}(a_{ij})C_{ij}\right)$

$$A_{ij}^{\text{sparse}} = \begin{cases} A_{ij}/p_{ij} & \text{with probability } p_{ij} \\ 0 & \text{otherwise} \end{cases}$$

(skipping some details)

Then, with high probability

$$\forall x \in \mathbb{R}^N \quad (1-\epsilon)x^T A x \le x^T A^{sparse} x \le (1+\epsilon)x^T A x$$

(eigenvalues, eigenvectors, linear equation close to original)

Sparsification by sampling

see Spielman & Srivastava 2011 for more on this strategy

$$A = \begin{pmatrix} * & a_{12} & a_{13} \\ a_{21} & * & a_{23} \\ a_{31} & a_{32} & * \end{pmatrix} = \begin{pmatrix} * & a_{12} & 0 \\ a_{21} & * & 0 \\ 0 & 0 & * \end{pmatrix} + \begin{pmatrix} * & 0 & a_{13} \\ 0 & * & 0 \\ a_{31} & 0 & * \end{pmatrix} + \begin{pmatrix} * & 0 & 0 \\ 0 & * & a_{23} \\ 0 & a_{32} & * \end{pmatrix}$$

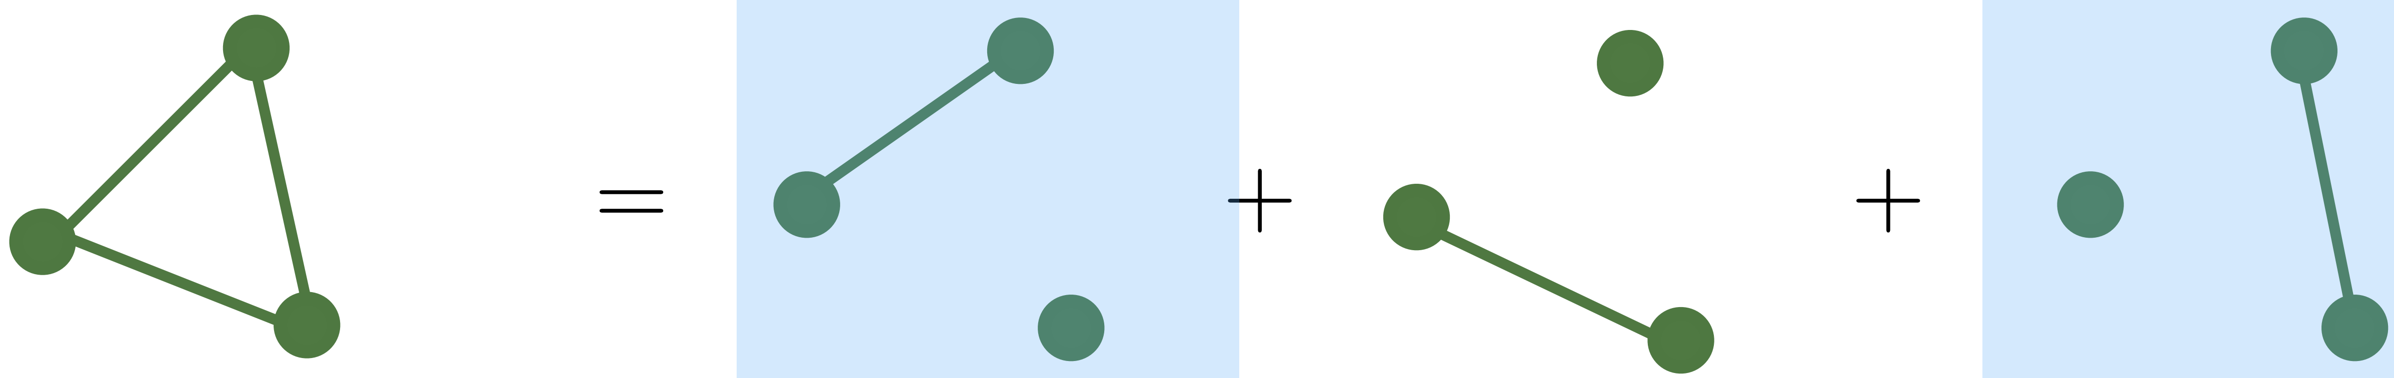

p1 p2 p3

$A_{\text{sparse}} =$

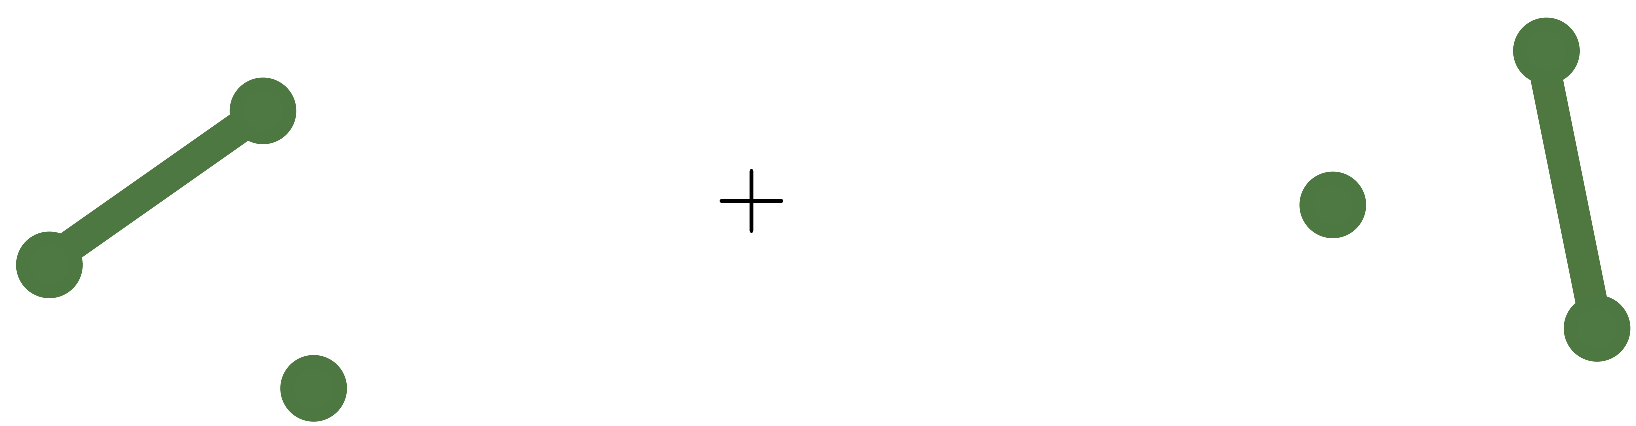

Key ideas in sparsification proof:

1) Write A as sum of positive definite pieces (= edges)

$$A = \sum_{(i,j)\in E} \vec{u}_{ij}\vec{u}_{ij}^T$$

2) Convert to and approximate identity

$$\vec{v}_{ij} = A^{-1/2}\vec{u}_{ij} \qquad I = \sum_{(i,j)\in E} v_{ij}v_{ij}^T$$

3) Sample pieces to yield I_{sparse}. Bound $||I - I_{sparse}||_2$ using matrix Chernoff bound (Tropp 2012)

4) Choose probabilities to optimize concentration bound. Probabilities depend on A, A^{-1}.

$$p_{ij} \propto |a_{ij}| \left(A_{ii}^{-1} + A_{jj}^{-1} - 2\text{sign}(a_{ij})A_{ij}^{-1}\right)$$

How do we compute sampling probabilities?

$$\frac{d\boldsymbol{x}}{d\boldsymbol{t}} = \boldsymbol{A}\boldsymbol{x} + \xi(t)$$

input

white noise

covariance matrix $C_{ij} = \mathbb{E}\left[x_i x_j\right]$

C is solution to Lyapunov equation

$$AC + CA^T = -I$$

(when dynamical system is stable)

If A is normal, then $C \propto A_{symm}^{-1}$ (use ev basis of A)

If A is symmetric, then $C \propto A^{-1}$

$$p_{ij} \propto |a_{ij}| \left(C_{ii}^{-1} + C_{jj}^{-1} \pm 2C_{ij}^{-1}\right)$$

Proof summary

$$A = \begin{pmatrix} * & a_{12} & a_{13} \\ a_{21} & * & a_{23} \\ a_{31} & a_{32} & * \end{pmatrix} = \begin{pmatrix} * & a_{12} & 0 \\ a_{21} & * & 0 \\ 0 & 0 & * \end{pmatrix} + \begin{pmatrix} * & 0 & a_{13} \\ 0 & * & 0 \\ a_{31} & 0 & * \end{pmatrix} + \begin{pmatrix} * & 0 & 0 \\ 0 & * & a_{23} \\ 0 & a_{32} & * \end{pmatrix}$$

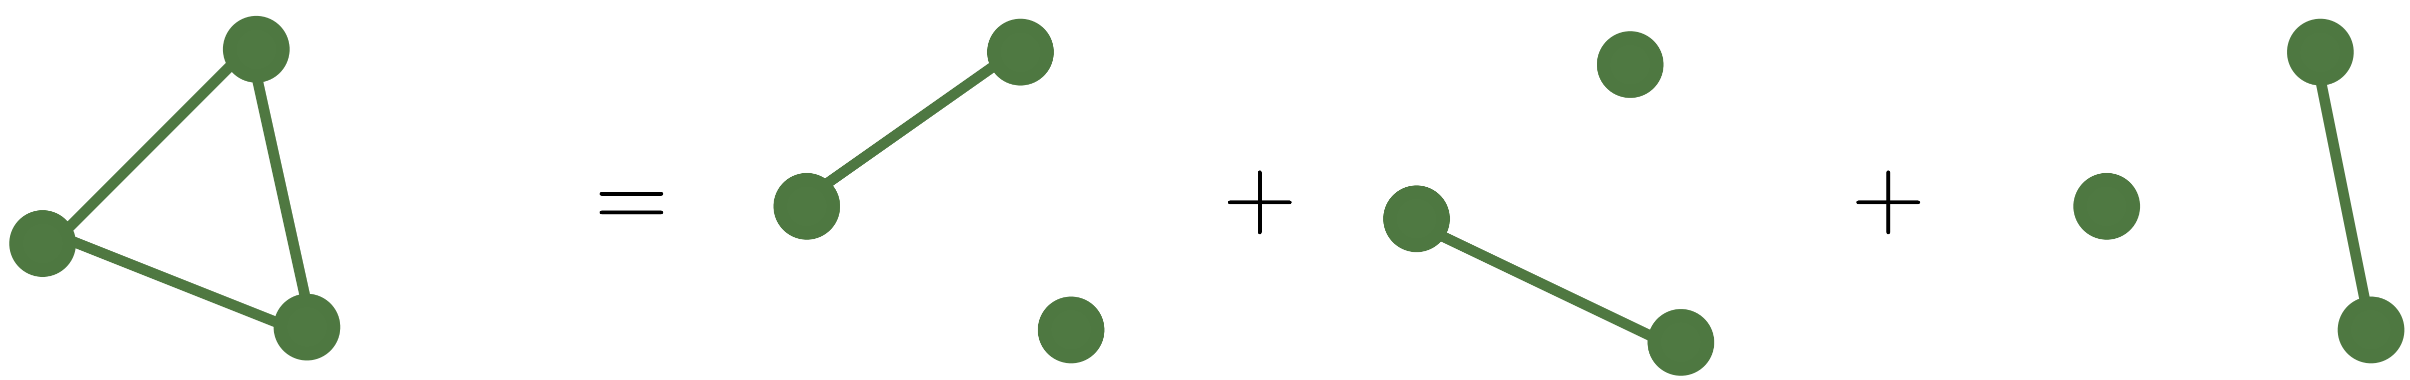

sample $\quad p_1 \quad p_2 \quad p_3$

For good p_i, A_{sparse} close to A with only O(Nlog(N)) edges

Compute p_i from dynamical system on graph

Idea summary

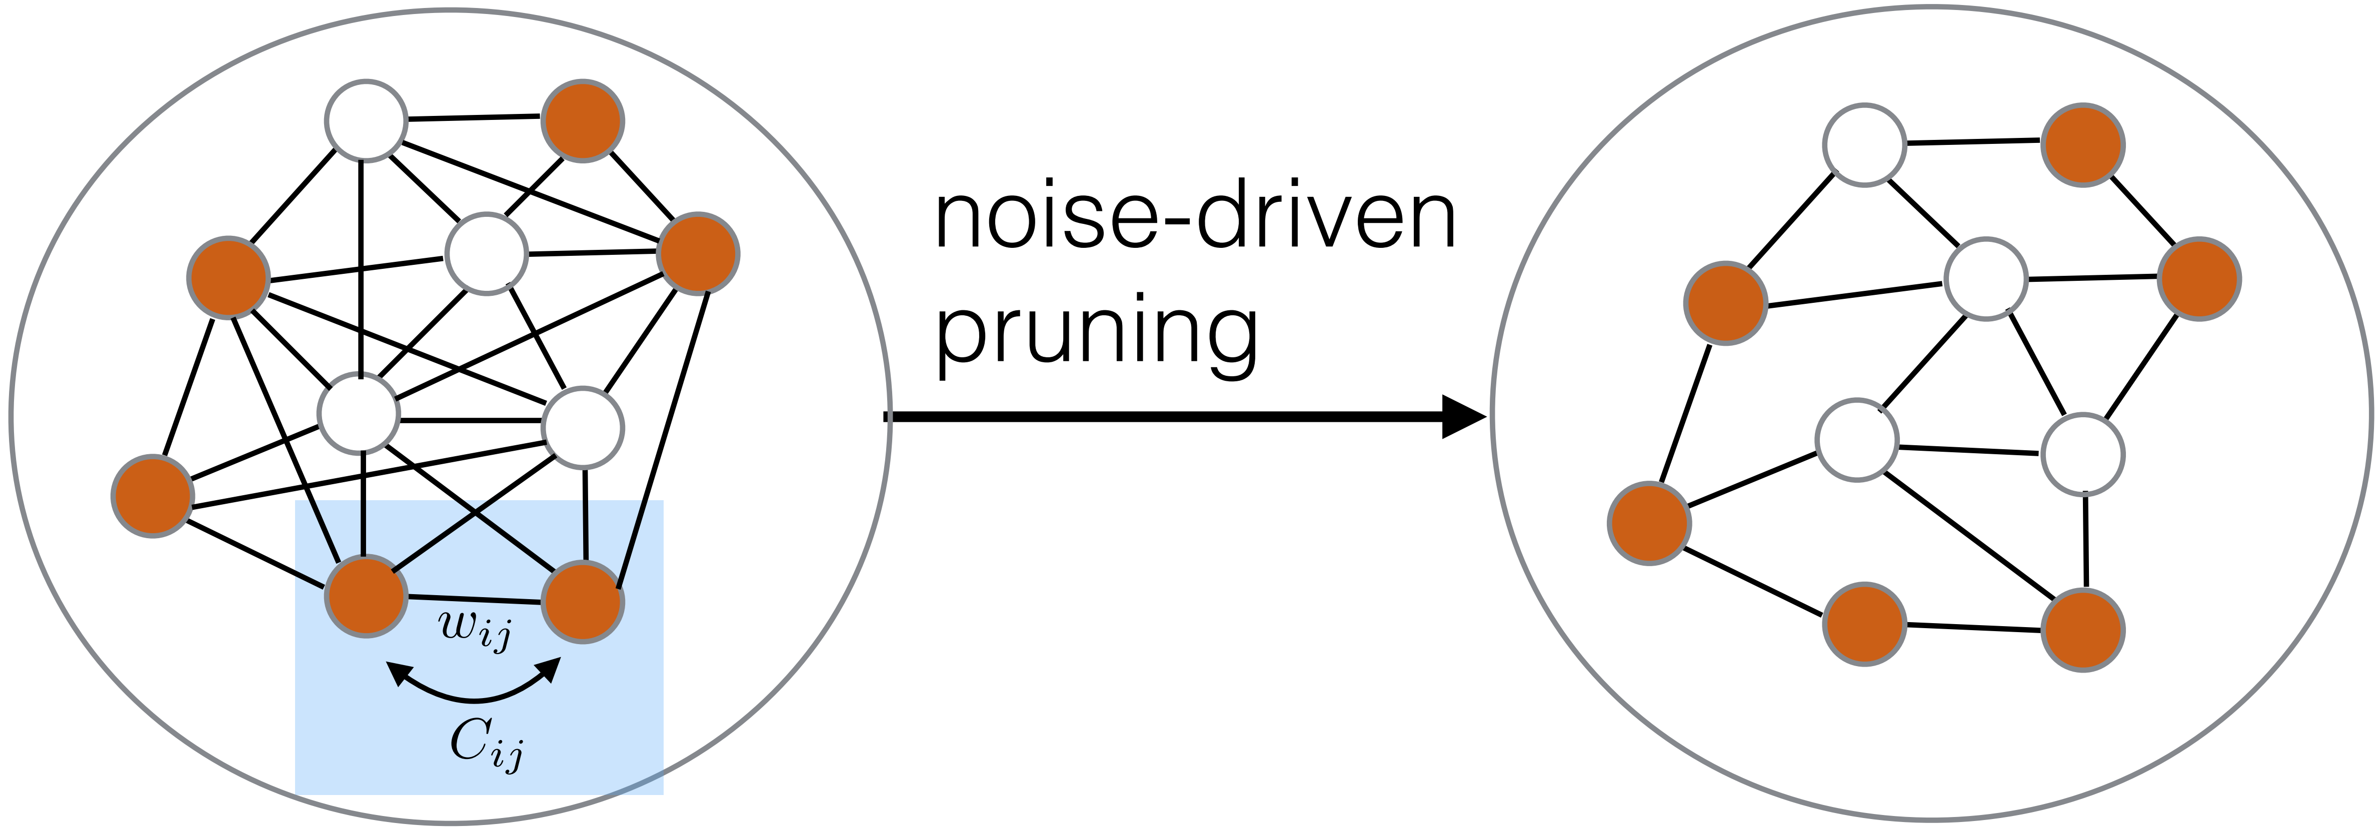

Use noise to probe and prune network (uses randomness in two ways)

Prune synapses based on weight and covariance

Graph sparsification + dynamical systems

Properties of algorithm

Computing sparsification probabilities/matrix inverse is non-trivial

Parallel

Randomized

Dynamical system

Based on local, biologically-accessible information

Network-level homeostatic mechanism

Separation of learning and pruning

Q: How do we learn a sparse network solution?

perhaps separate into:

learning	**pruning**
task-driven	task agnostic, homeostatic
greedily add synapses where might be useful	remove synapses while trying to preserve existing dynamics
Hebbian	anti-Hebbian
may be reward or error driven	unsupervised

could be happening simultaneously or in separate epochs

Some features of the brain this rule might explain

Noise, esp. during sleep

Seemingly random drift: multiple synapses b/w same 2 neurons are weakly correlated

Various kinds of plasticity during sleep. Synaptic strengths seem to go both up and down.

Synapses in different states: labile vs. fixed

Sweeps of excitability during slow-wave sleep: shift network into diagonally-dominant regime?

Conclusions

Networks in the brain are sparse

Neural activity is noisy

Brain might use noise to detect and prune redundant connections

Proof: graph sparsification + dynamical systems

Separate learning and pruning rules

Preliminary results! Working on extending along multiple directions

Graphs and networks in neuroscience

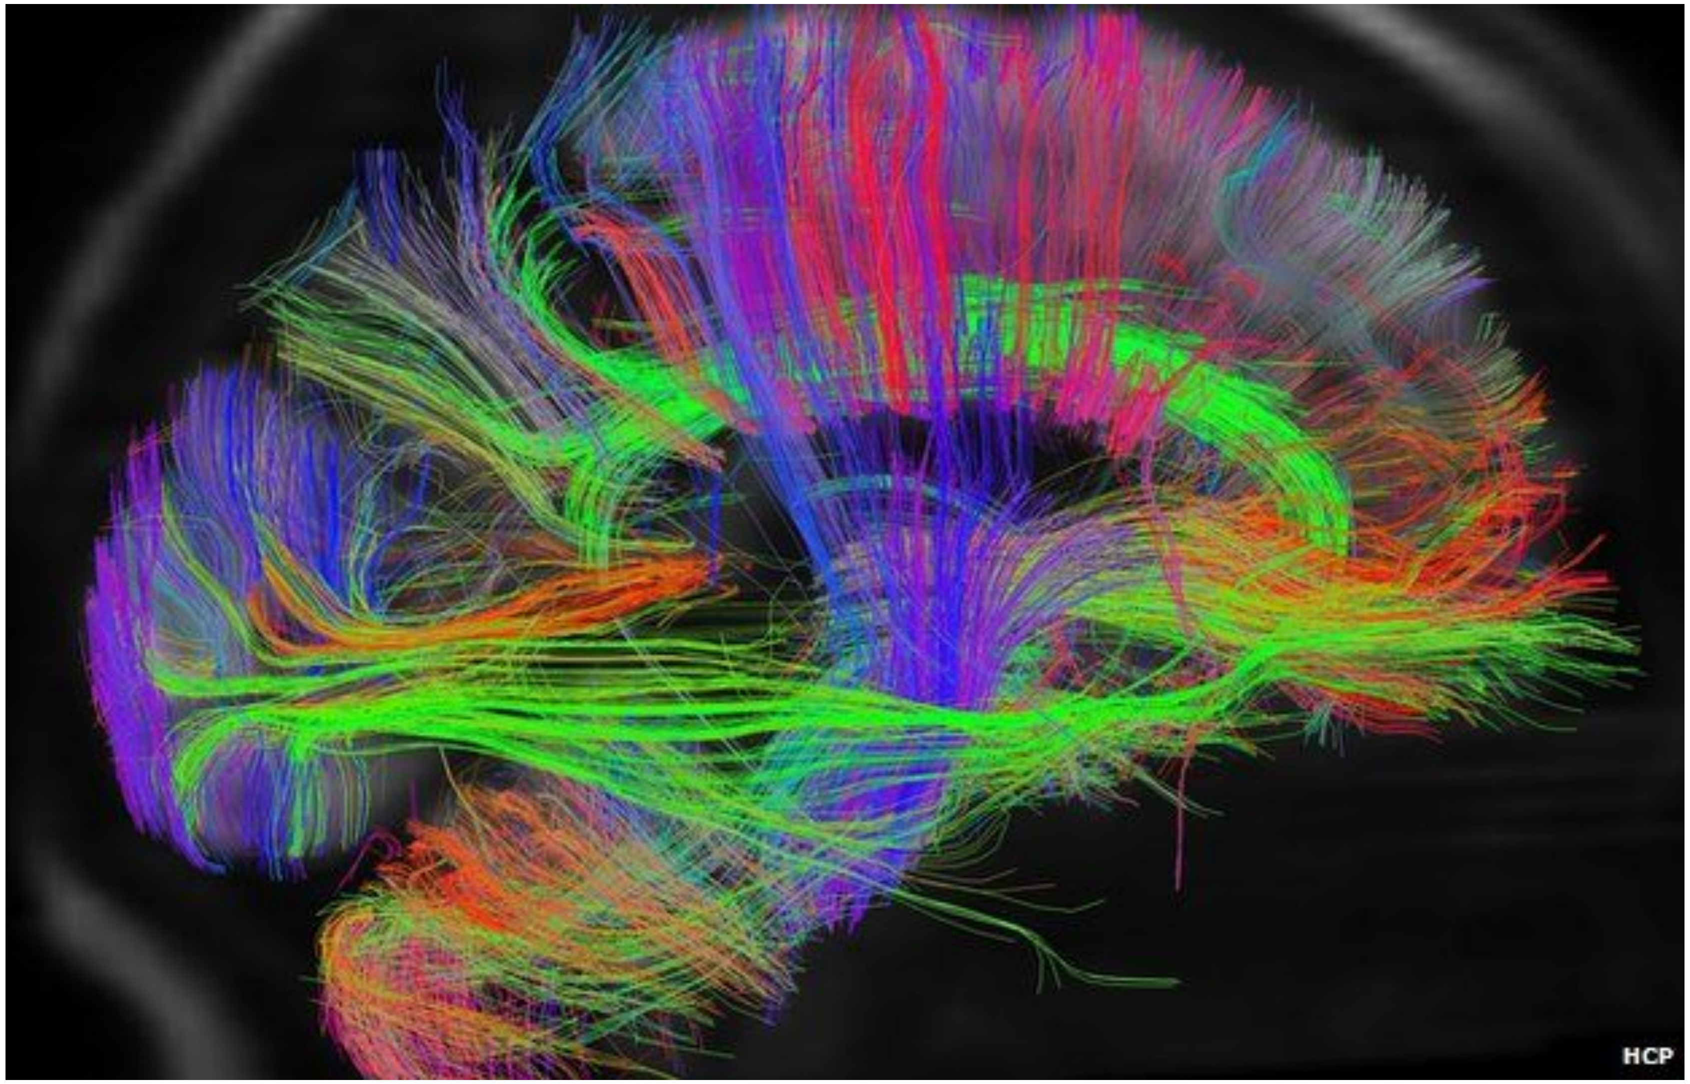

Image: Human Connectome Project

Show up in all sorts of places. I'm going to briefly mention a few interesting directions (with a math bias).

Same network exists at multiple levels

Structural networks (anatomy)

Dynamical system with network structure (physiology)

Distributed computation (function)

How do we fuse these frameworks/levels of understanding?

Networks in the brain are plastic

1) Change due to learning

2) Change due to homeostatic plasticity

3) Drift

4) Development and evolution

How do we understand these changing networks?

Bigger question: what in the brain is stable?

Which features of a network matter? Which can we ignore?

What are the right invariants?

What are the right macroscopic variables?

Very partial observations of a network

Observe 10–1000 neurons out of millions (or more).

What can we conclude from these partial observations?

How to stitch together data from multiple observations, sessions, modalities, levels of description?

Computation on abstract networks

Abstract graphs have been a productive way of thinking about error-correcting codes, conditional independence in probability distributions, etc.

Do these abstract graphs have counterparts in the brain?

Challenges and opportunities:

Neuro as source of interesting mathematical questions

Increasing amounts of data; shortage of theories, models and analysis tools

Challenge: Domain specific knowledge

Idea: neural networks could use noise to probe and prune connections

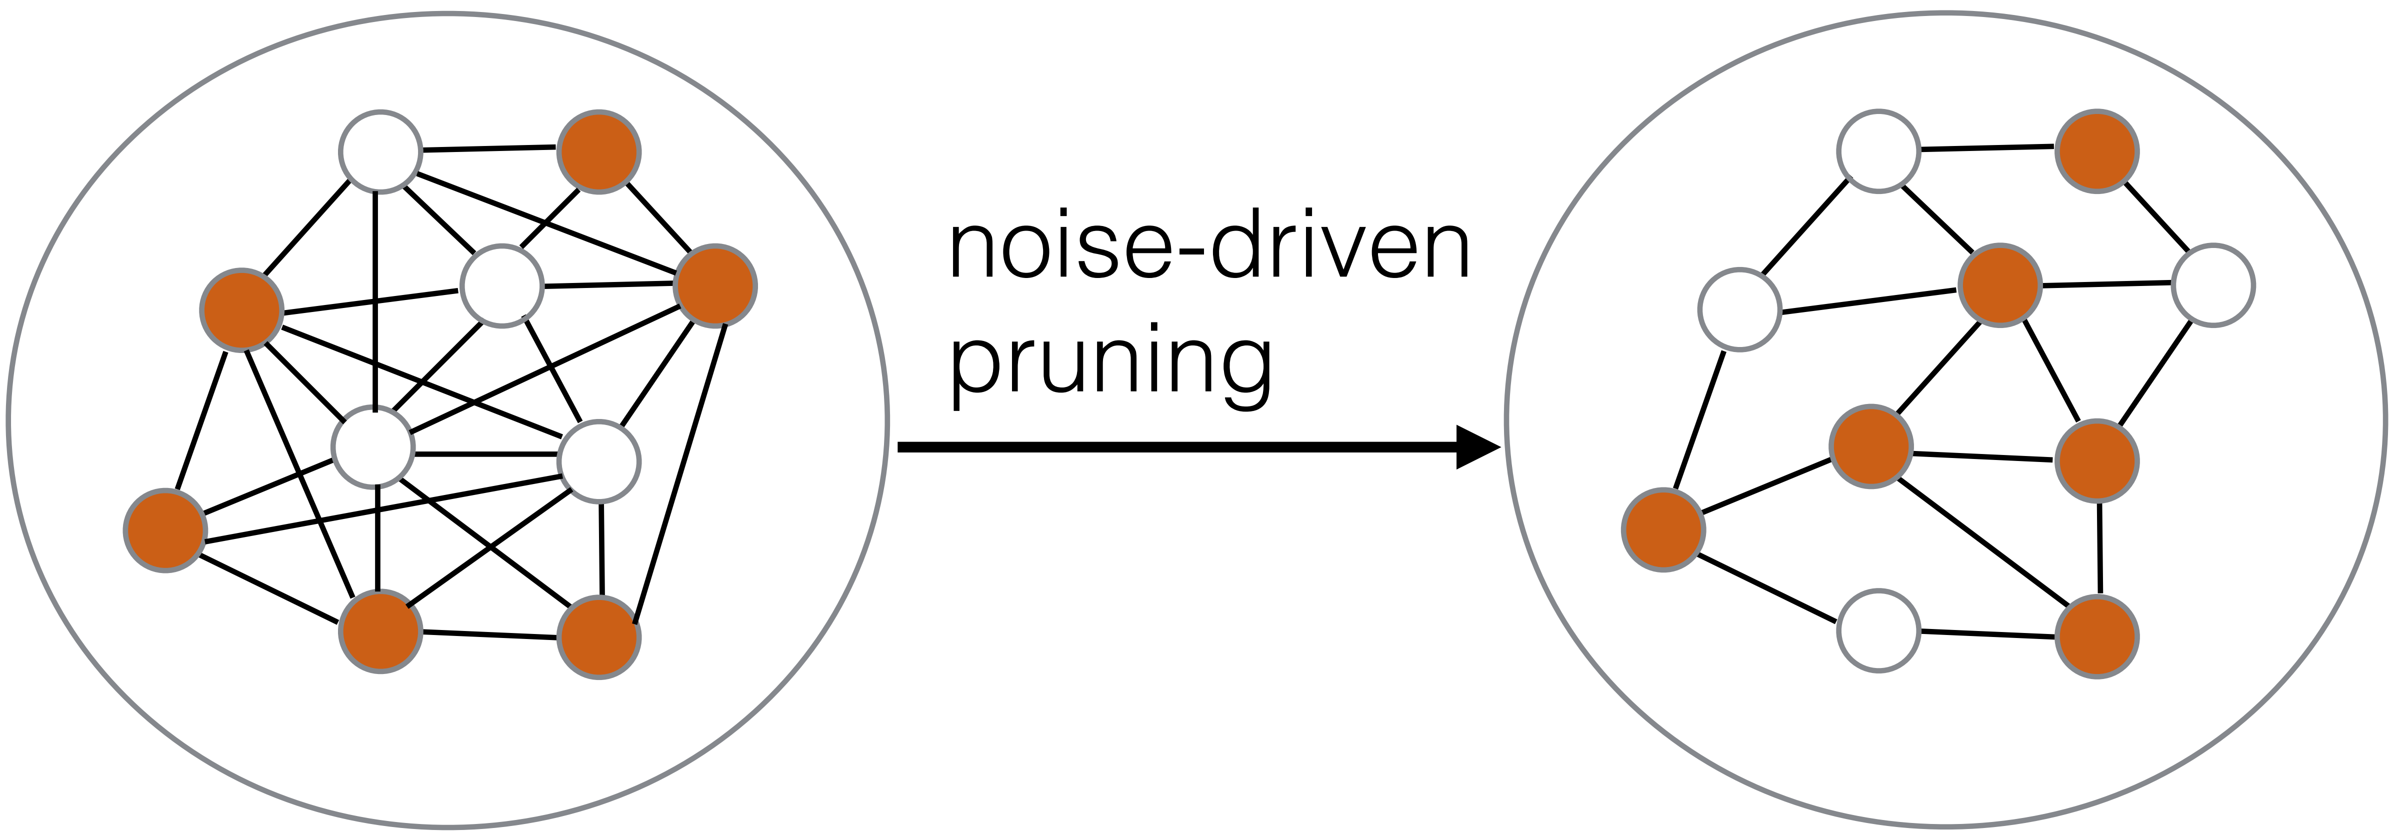

THANK YOU!