Neural Manifolds and Low-dimensional dynamics: What are Neural Manifolds?

1. What are Neural Manifolds?
- experimental observations

2. Two views of Neural Activity
- computing (Hopfield model)
- neural circuits (field model)

3. Low-dimensional dynamics
- formalism and assumption
- dynamics

4. Examples of low-dim dynamics
- context-dependent decision making

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Cambridge Univ. Press

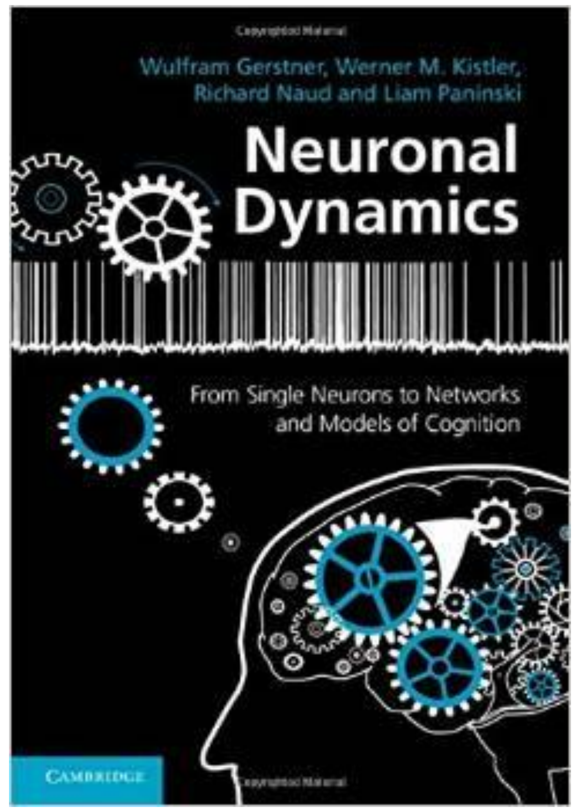

Introduction: What are Neural Manifolds?

Mante, V., Sussillo, D., Shenoy, K.V., Newsome, W.T.: Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503(7474), 78–84 (2013)

Shenoy, K.V., Sahani, M., Churchland, M.M.: Cortical control of arm movements: A dynamical systems perspective. Annual Review of Neuroscience 36(1), 337–359 (2013)

Mastrogiuseppe, F., Ostojic, S.: Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. Neuron 99(3), 609–62329 (2018)

Chaudhuri, R., Gercek, B., Pandey, B., Peyrache, A., Fiete, I.: The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. Nature Neuroscience 22(9), 1512–1520 (2019)

Vyas, S., Golub, M.D., Sussillo, D., Shenoy, K.V.: Computation through neural population dynamics. Annual Review of Neuroscience 43(1), 249–275 (2020)

Barack, D.L., Krakauer, J.W.: Two views on the cognitive brain. Nature Reviews Neuroscience 22(6), 359–371 (2021)

Langdon, C., Genkin, M., Engel, T.A.: A unifying perspective on neural manifolds and circuits for cognition. Nature Reviews Neuroscience 24(6), 363–377 (2023)

DePasquale, B., Sussillo, D., Abbott, L.F., Churchland, M.M.: The centrality of population-level factors to network computation is demonstrated by a versatile approach for training spiking networks. Neuron 111(5), 631–64910 (2023)

Pezon, L., Schmutz, V, Gerstner, W. (2024), Linking Neural Manifolds to Principles of Circuit Structure in Recurrent Networks bioRxiv doi: https://doi.org/10.1101/2024.02.28.582565

Introduction: low-dimensional response manifold

Neural Manifold: The activity does not fill the N-dimensional space

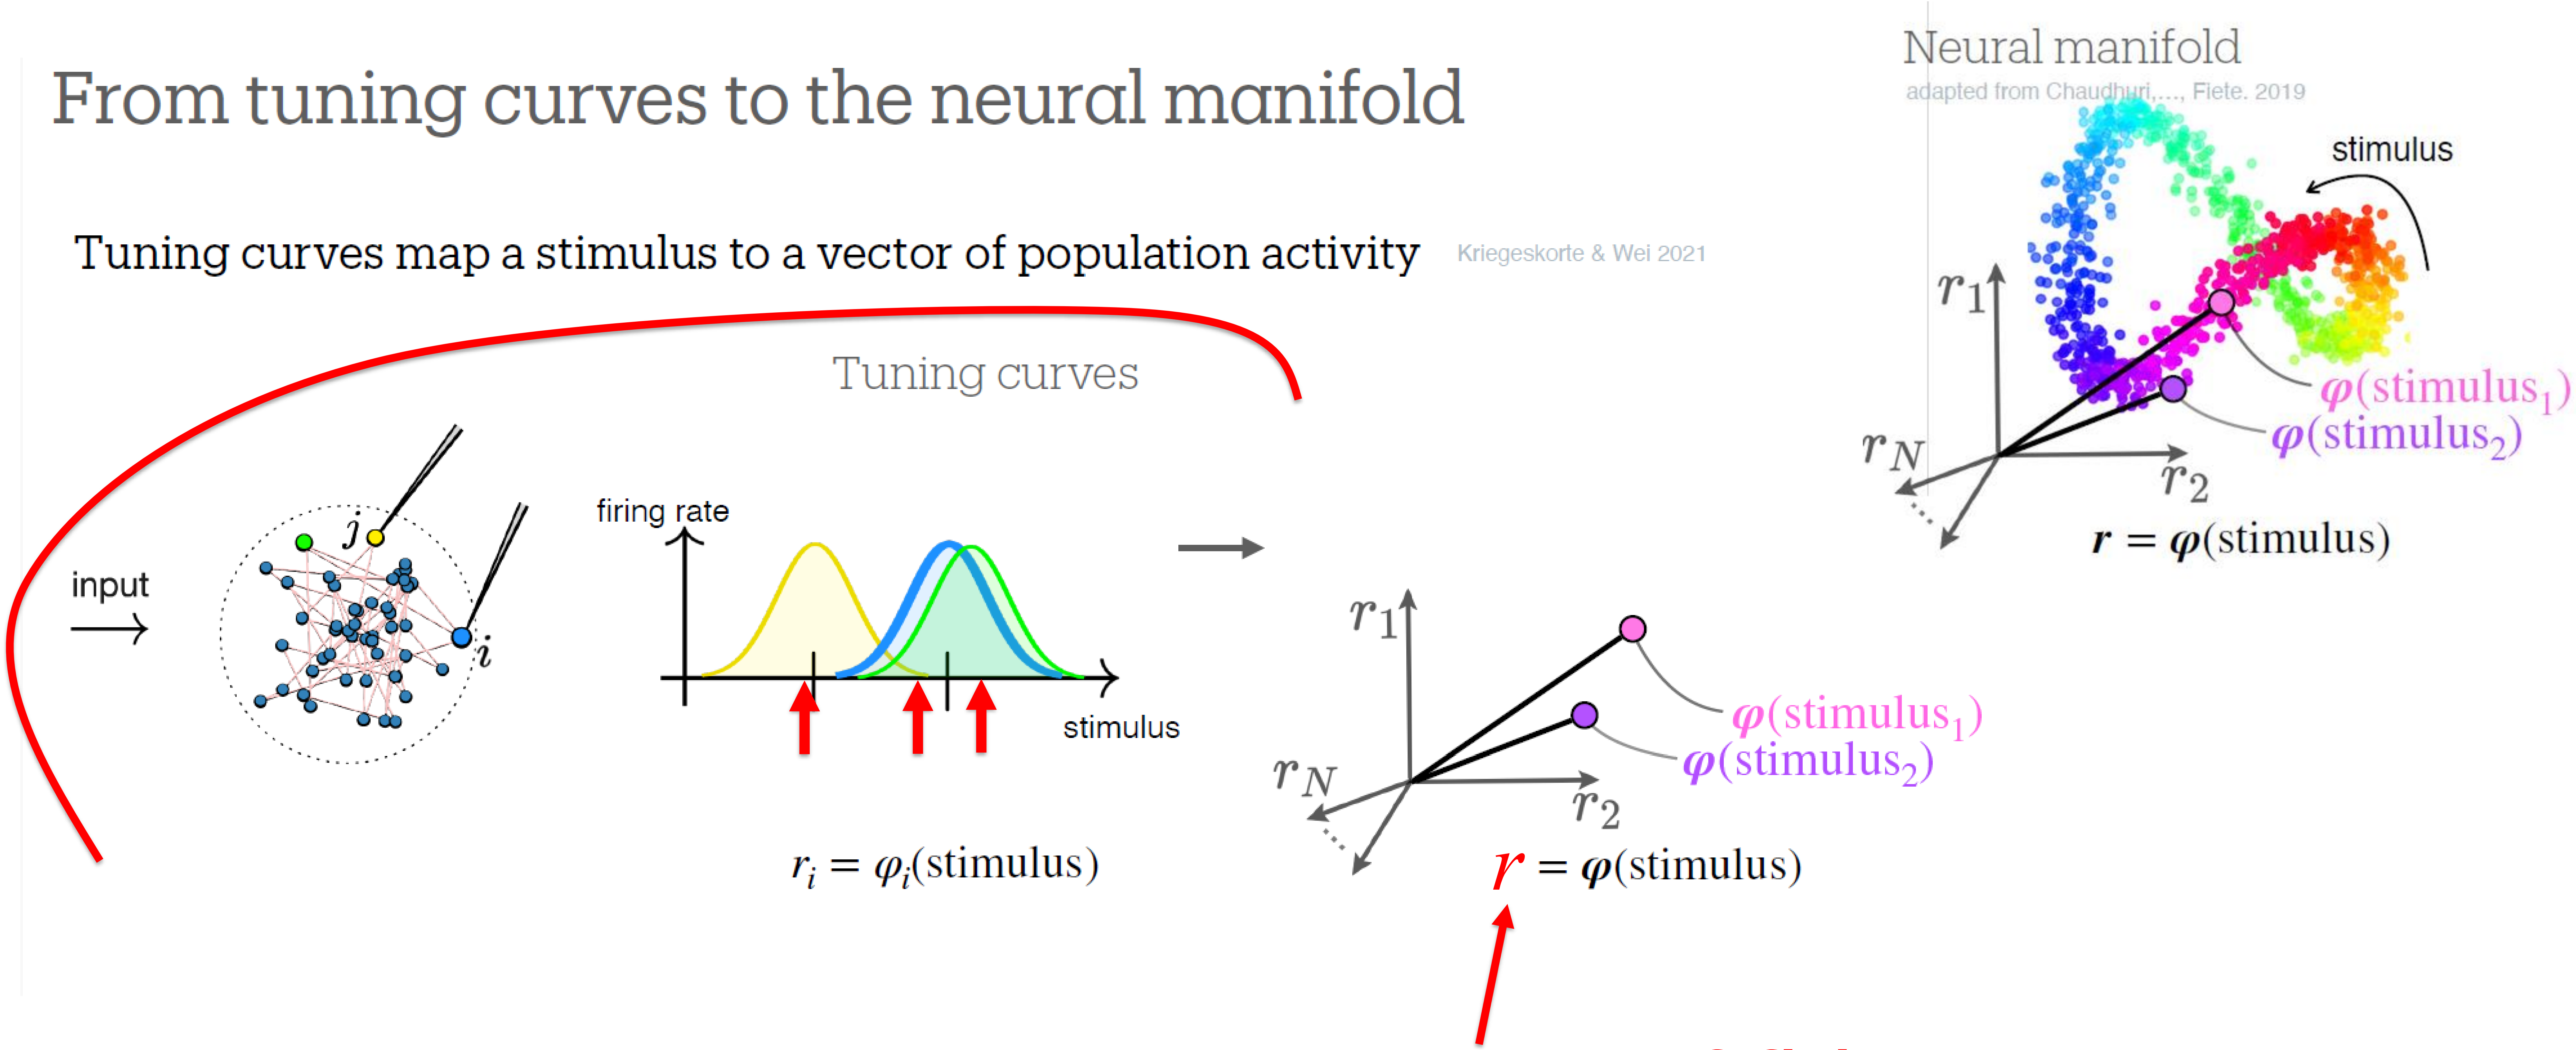

vector of firing rates

$r=(r_1, r_2, r_3, \ldots r_N)$

Introduction: low-dimensional dynamics

How can we think about neural activity?

low-dimensional dynamics

$r_i(t)$= sliding estimate of firing rate of neuron ***i***

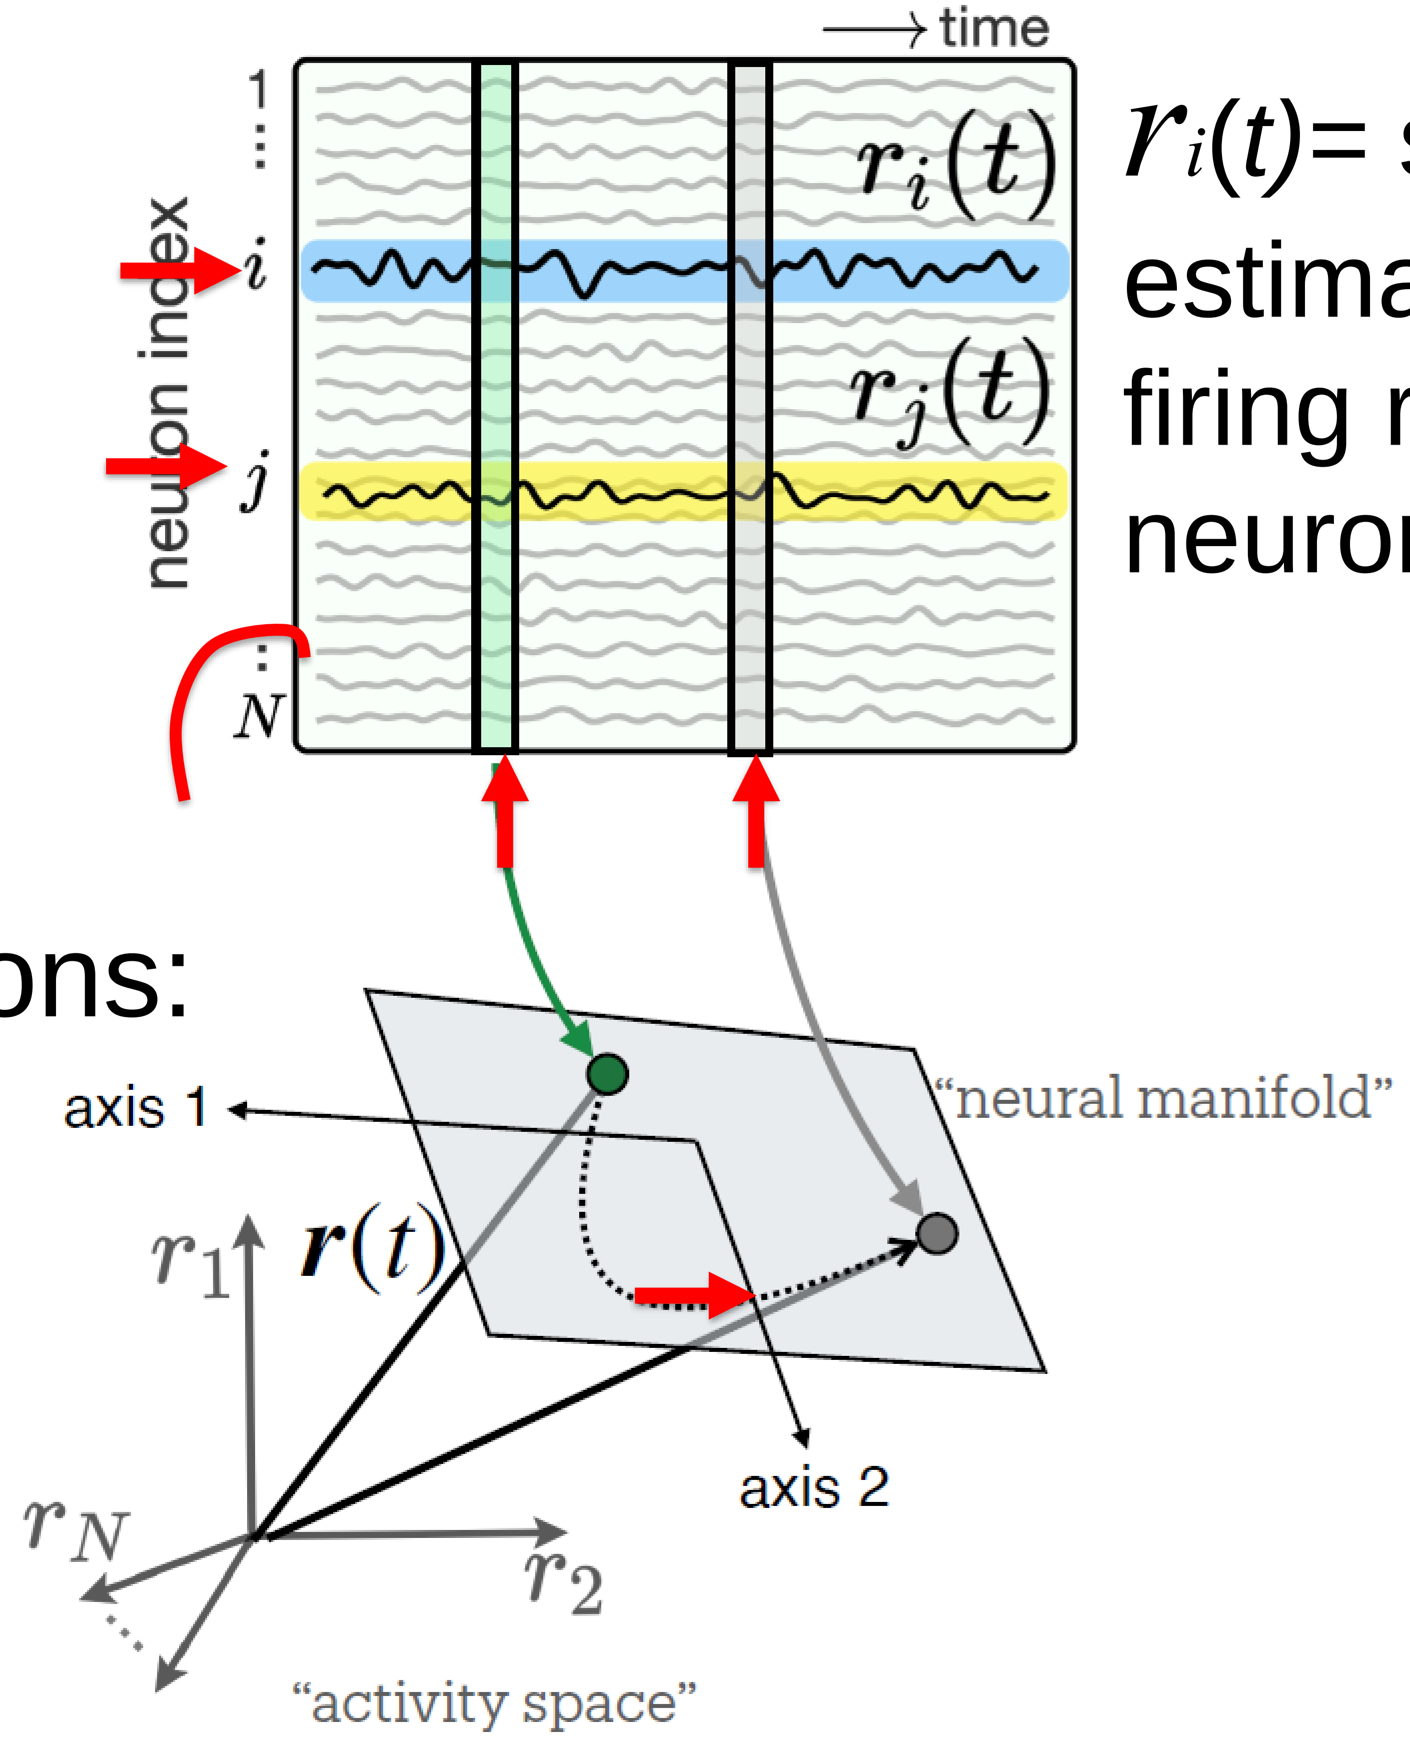

Simultaneous recordings from hundreds of neurons:

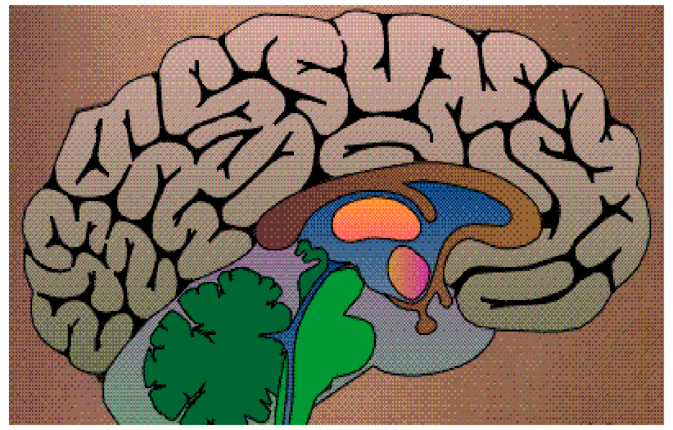

neuropixel probe

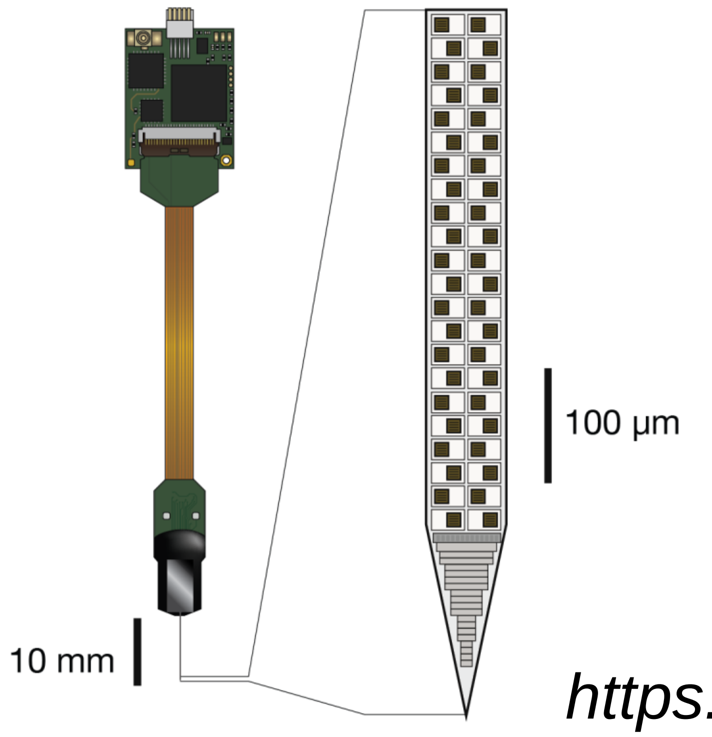

https://www.neuropixels.org/

Image: Pezon et al. 2024

Introduction: low-dimensional dynamics

Brain computation = dynamics in manifold

- **observation:** in many brain areas, the high-dimensional activity lies in a low-dimensional "***manifold***"

ex: HD cells & grid cells, prefrontal cortex, motor control

Ex: HD system (mouse)

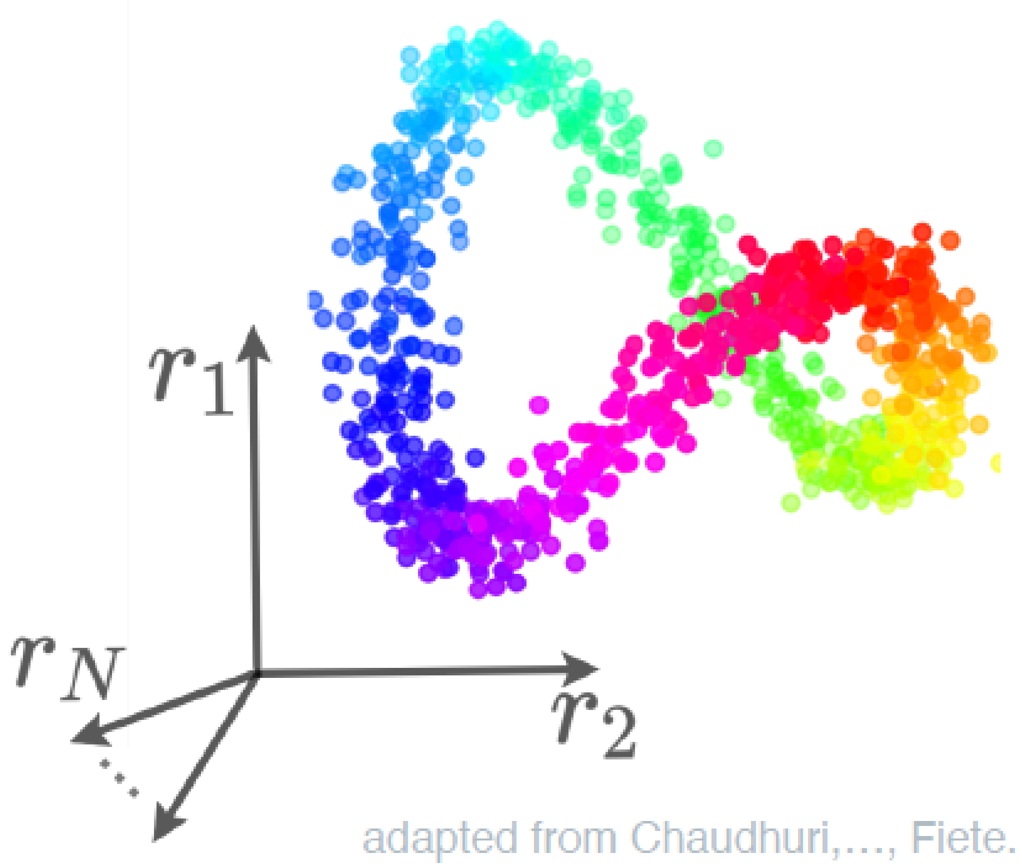

adapted from Chaudhuri,..., Fiete. 2019

Ex: motor cortex (monkey)

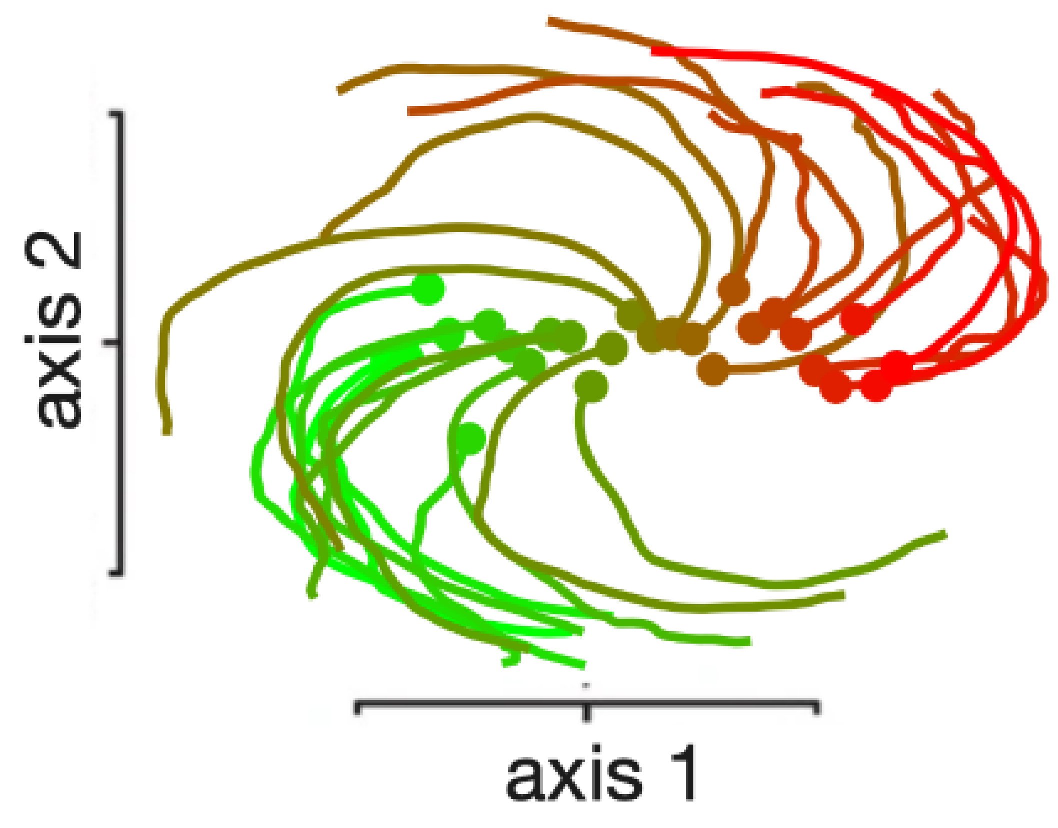

adapted from Churchland,...,Shenc

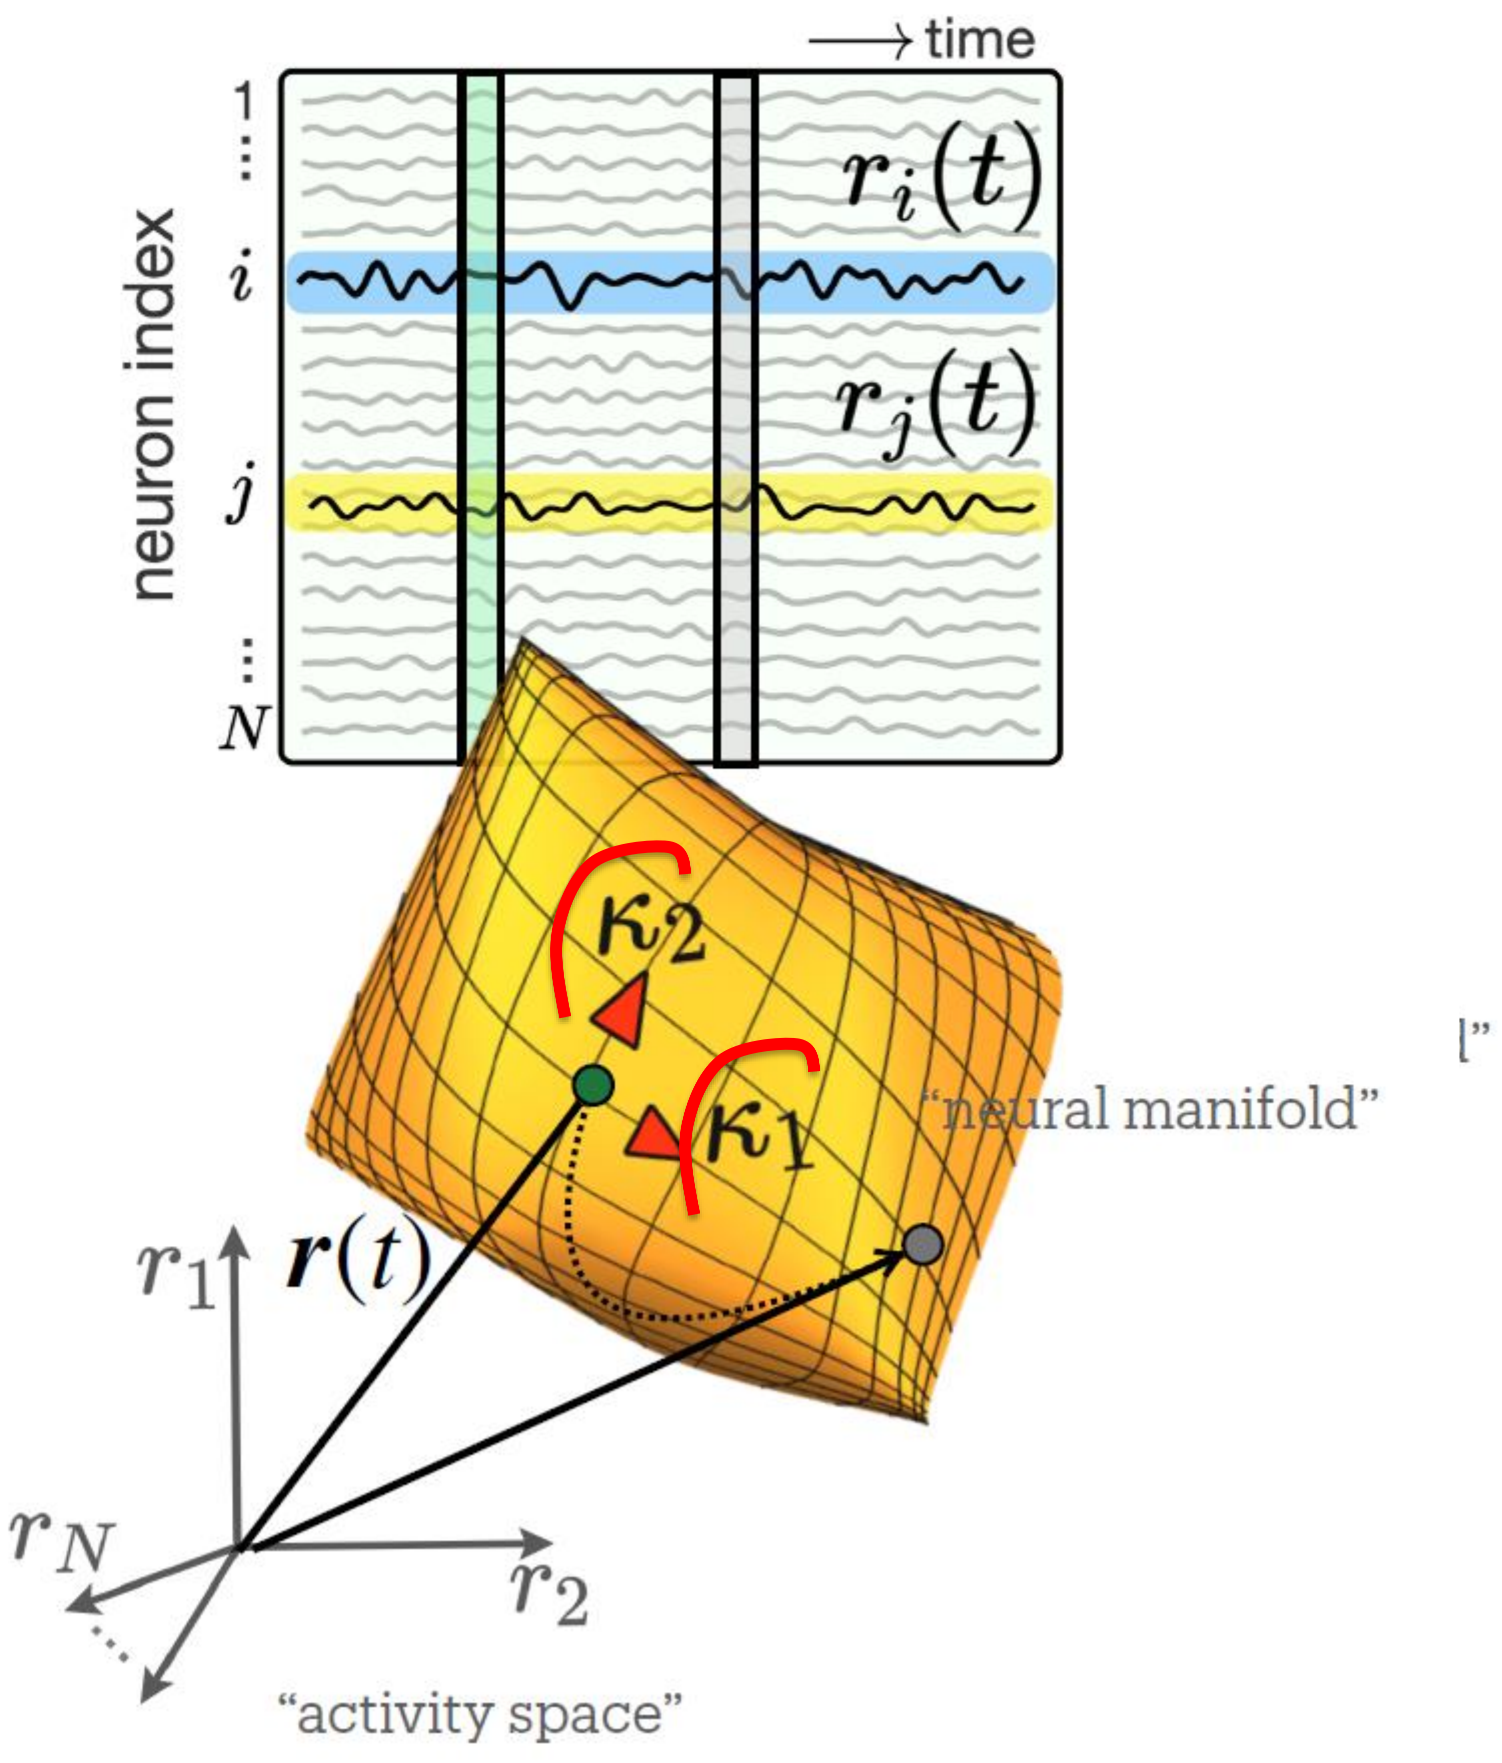

Image: Pezon et al. 2024

Introduction: low-dimensional dynamics

computations are described by **collective dynamics** in the manifold
Vyas, S., Golub, M.D., Sussillo, D., Shenoy, 2020

Ex: motor cortex (monkey)

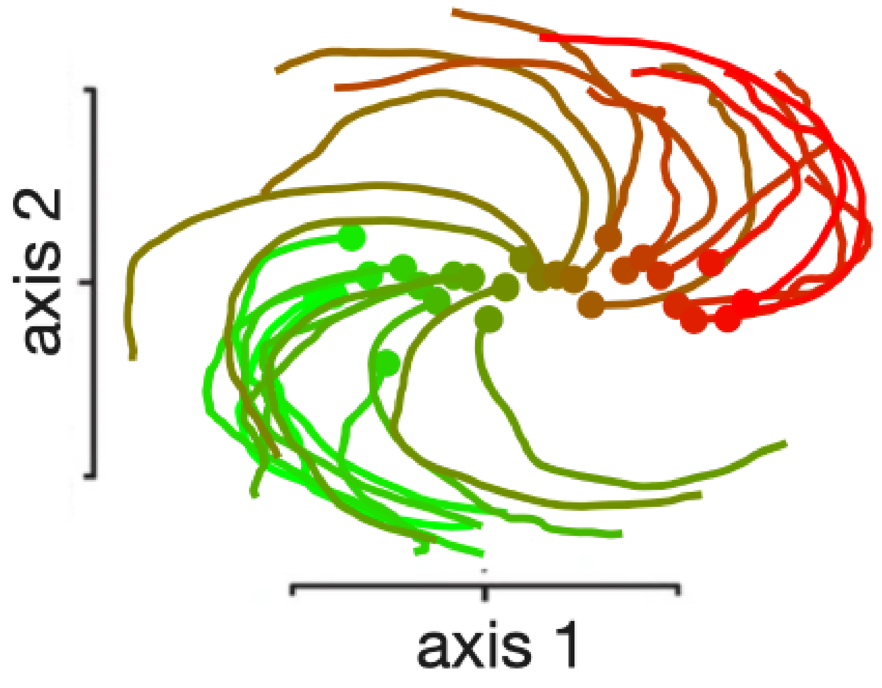

adapted from Churchland,...,Shenoy. 2012

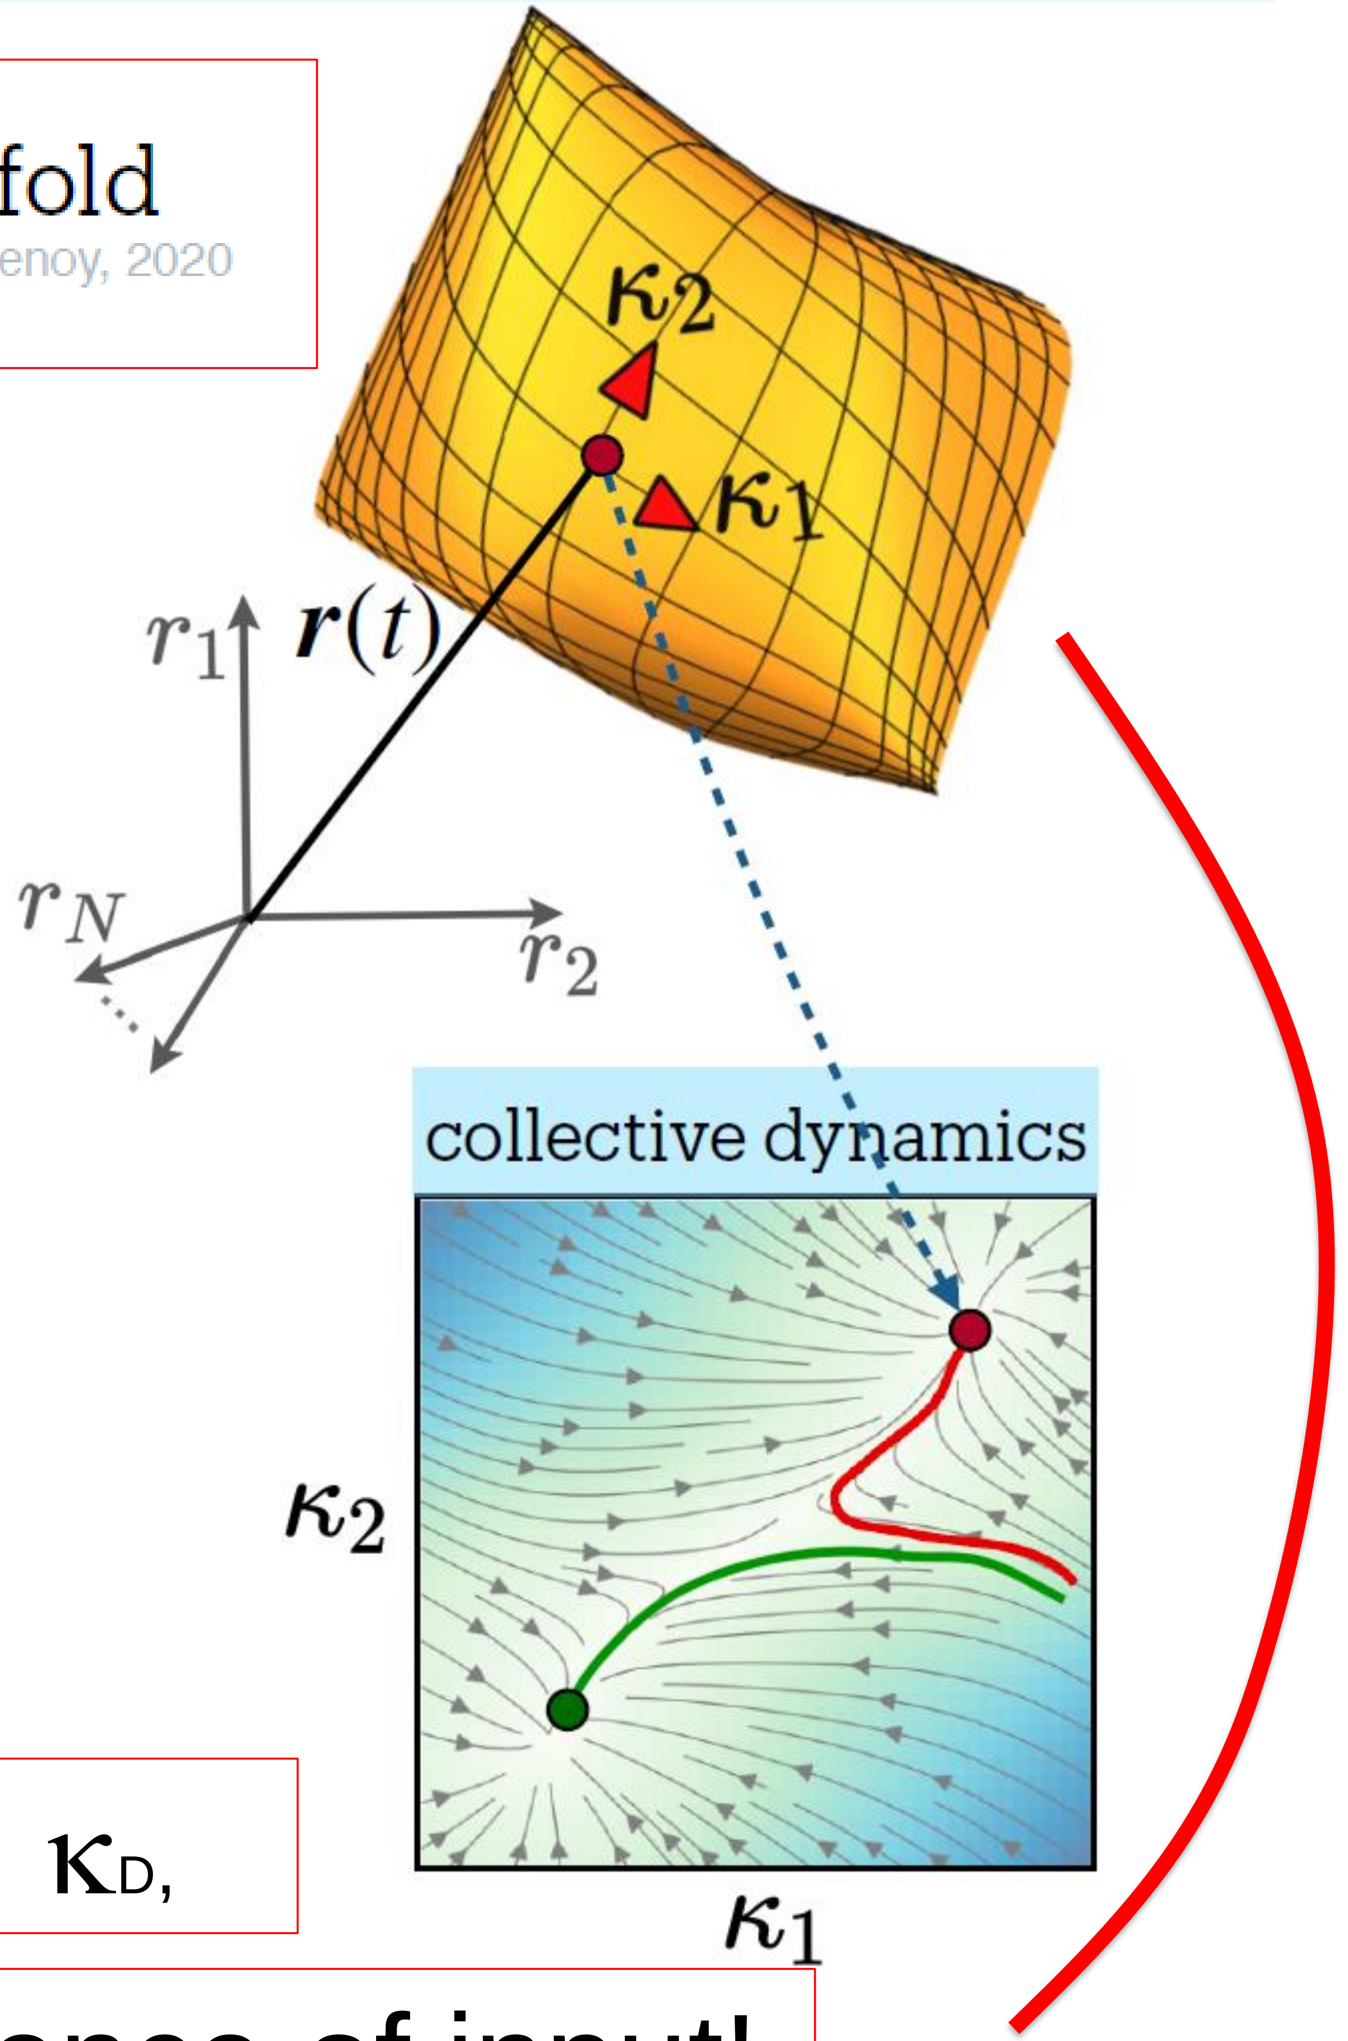

Flow described by small number of variables $\kappa_1, \ldots \kappa_D$,

Low-dimensional dynamics even during sleep/absence of input!

Chaudhuri et al, 2019

Image: Pezon et al. 2024

Quiz: low-dimensional dynamics

[ ] If an experimentalist records simultaneously from 287 neurons, then the momentary 'rate vector' of observed activity represents the network state as a point in 287 dimensions.

[ ] Different time points of the rate vector do not 'fill' the 287-dimensional space. Rather they live on a low-dimensional manifold.

[ ] A black-and-white camera of 1024 pixels mounted stably on a robot moving in an indoor environment of 2mX2m, generates measurement values that live on a 3-dimensional manifold in 1024-dimensional space

view manifold: Arleo&Gerstner (2000) https://doi.org/10.1007/s004220000171

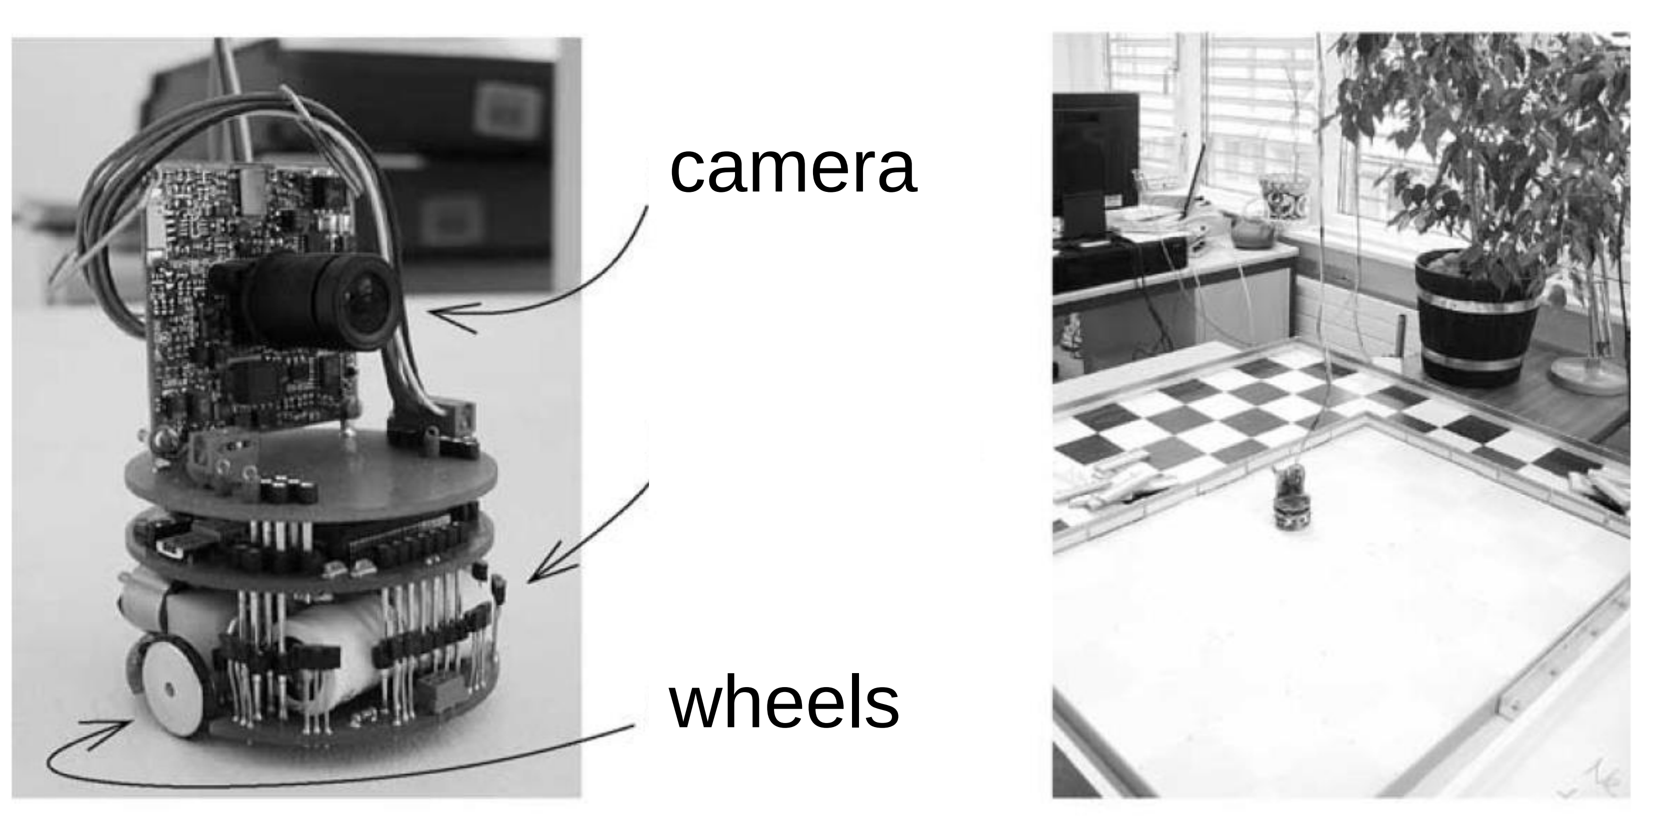

Even during dreaming, the neural activity lives on a low-dim. manifold

waking and sleep: Chaudhuri et al, 2019, https://doi.org/10.1038/s41593-019-0460-x

Neural Manifolds and low-dimensional dynamics

List of video lectures on Computational Neuroscience, organized by topics:
https://lcnwww.epfl.ch/gerstner/NeuronalDynamics-MOOCall.html

YouTube Channel:
https://www.youtube.com/@gerstnerlab

References:

Kriegeskorte and Wei (2021), Neural Tuning and Representational Geometry, Nat. Rev. Neuroscience 22:703-718

Churchland et al. (2012), Neural Population Dynamics During Reaching, Nature 487:51-56

Shenoy, K.V., Sahani, M., Churchland, M.M.: Cortical control of arm movements: A dynamical systems perspective. Annual Review of Neuroscience 36(1), 337–359 (2013)

Vyas, S., Golub, M.D., Sussillo, D., Shenoy, K.V.: Computation through neural population dynamics. Annual Review of Neuroscience 43(1), 249–275 (2020)

Pezon, L., Schmutz, V, Gerstner, W. (2024), Linking Neural Manifolds to Principles of Circuit Structure in Recurrent Networks bioRxiv doi: https://doi.org/10.1101/2024.02.28.582565

Chaudhuri, R., et al. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. Nat. Neurosci. 22(9), 1512–1520 (2019)

Arleo, A, Gerstner W. (2000), Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity *Biol. Cybern* 83, 287–299 (2000)

EPFL

Neural Manifolds and Low-dimensional dynamics: How can we interpret neural activity?

1. **What are Neural Manifolds?**
   - experimental observations

2. **Two views of Neural Activity**
   2.1 first view: computing (Hopfield model)
   2.2 2^{nd} view: neural circuits (field model)

3. **Low-dimensional dynamics**
   - formalism and assumption
   - dynamics

4. **Examples of low-dim dynamics**
   - context-dependent decision making

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Cambridge Univ. Press

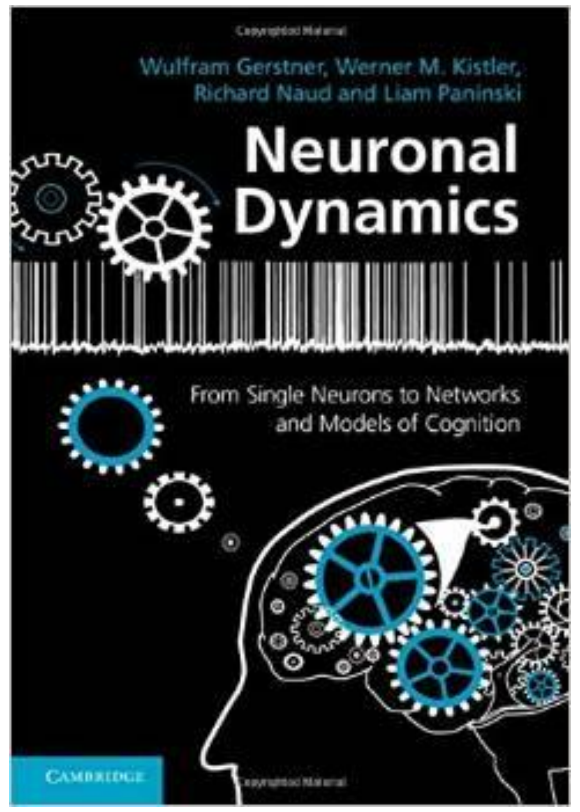

Introduction: How can we understand neuronal activity?

How can we understand principles of neuronal activity?

D. Barack and J. Krakauer, 2021
C. Langdon and T. Engel, 2023

Two different perspectives

- **Low-dimensional dynamics:**
  e.g., flow towards fixed point/attractor dynamics:
  → Hopfield model

→ computation as flow!

- **Receptive fields** and wiring/circuits:
  neurons can be classified according to
  functional similarity
  → field model

→ function from wiring/ circuit structure!

Review: Hopfield Model of Associative Memory

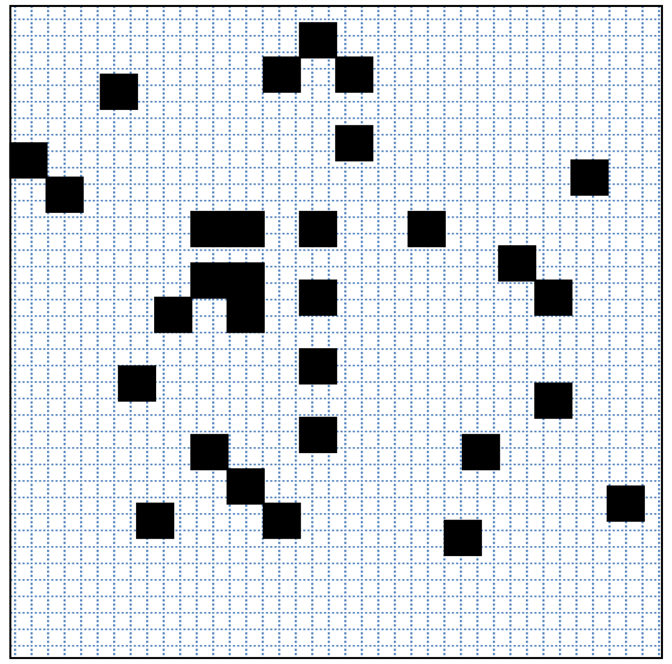

Prototype

$\vec{p}^{\,1}$ (random pattern)

interactions

$$w_{ij} = \frac{1}{N} \sum_{\mu} p_i^{\mu} p_j^{\mu}$$

Sum over all prototypes

This rule is very good for **random** patterns

It does not work well for correlated patters

dynamics

$$S_i(t+1) = \mathrm{sgn}\left[\sum_j w_{ij} S_j(t)\right]$$

all interactions with i

Random patterns, fully connected:

Hopfield model

J. Hopfield, 1982

Review: Hopfield Model of Associative Memory

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

weights: $w_{ij} = \frac{1}{N} \sum_\mu p_i^\mu p_j^\mu$

$$S_i(t+1) = \mathrm{sgn}[\sum_\mu p_i^\mu m^\mu(t)]$$

overlap (similarity) $m^\mu(t) = \frac{1}{N} \sum_j p_j^\mu S_j(t)$

Review Hopfield model: memory retrieval

$$\Pr\{S_i(t+1)=+1 \mid h_i\} = g[h_i] = g\left[\sum_j w_{ij} S_j(t)\right] = g\left(\sum_{\mu} p_i^{\mu} m^{\mu}(t)\right)$$

overlap (similarity) $m^{\mu}(t) = \frac{1}{N}\sum_j p_j^{\mu} S_j(t)$

If we start close to pattern ν, **1-dimensional dynamics**

$$m^{\nu}(t+1) = \tilde{F}(m^{\nu}(t))$$

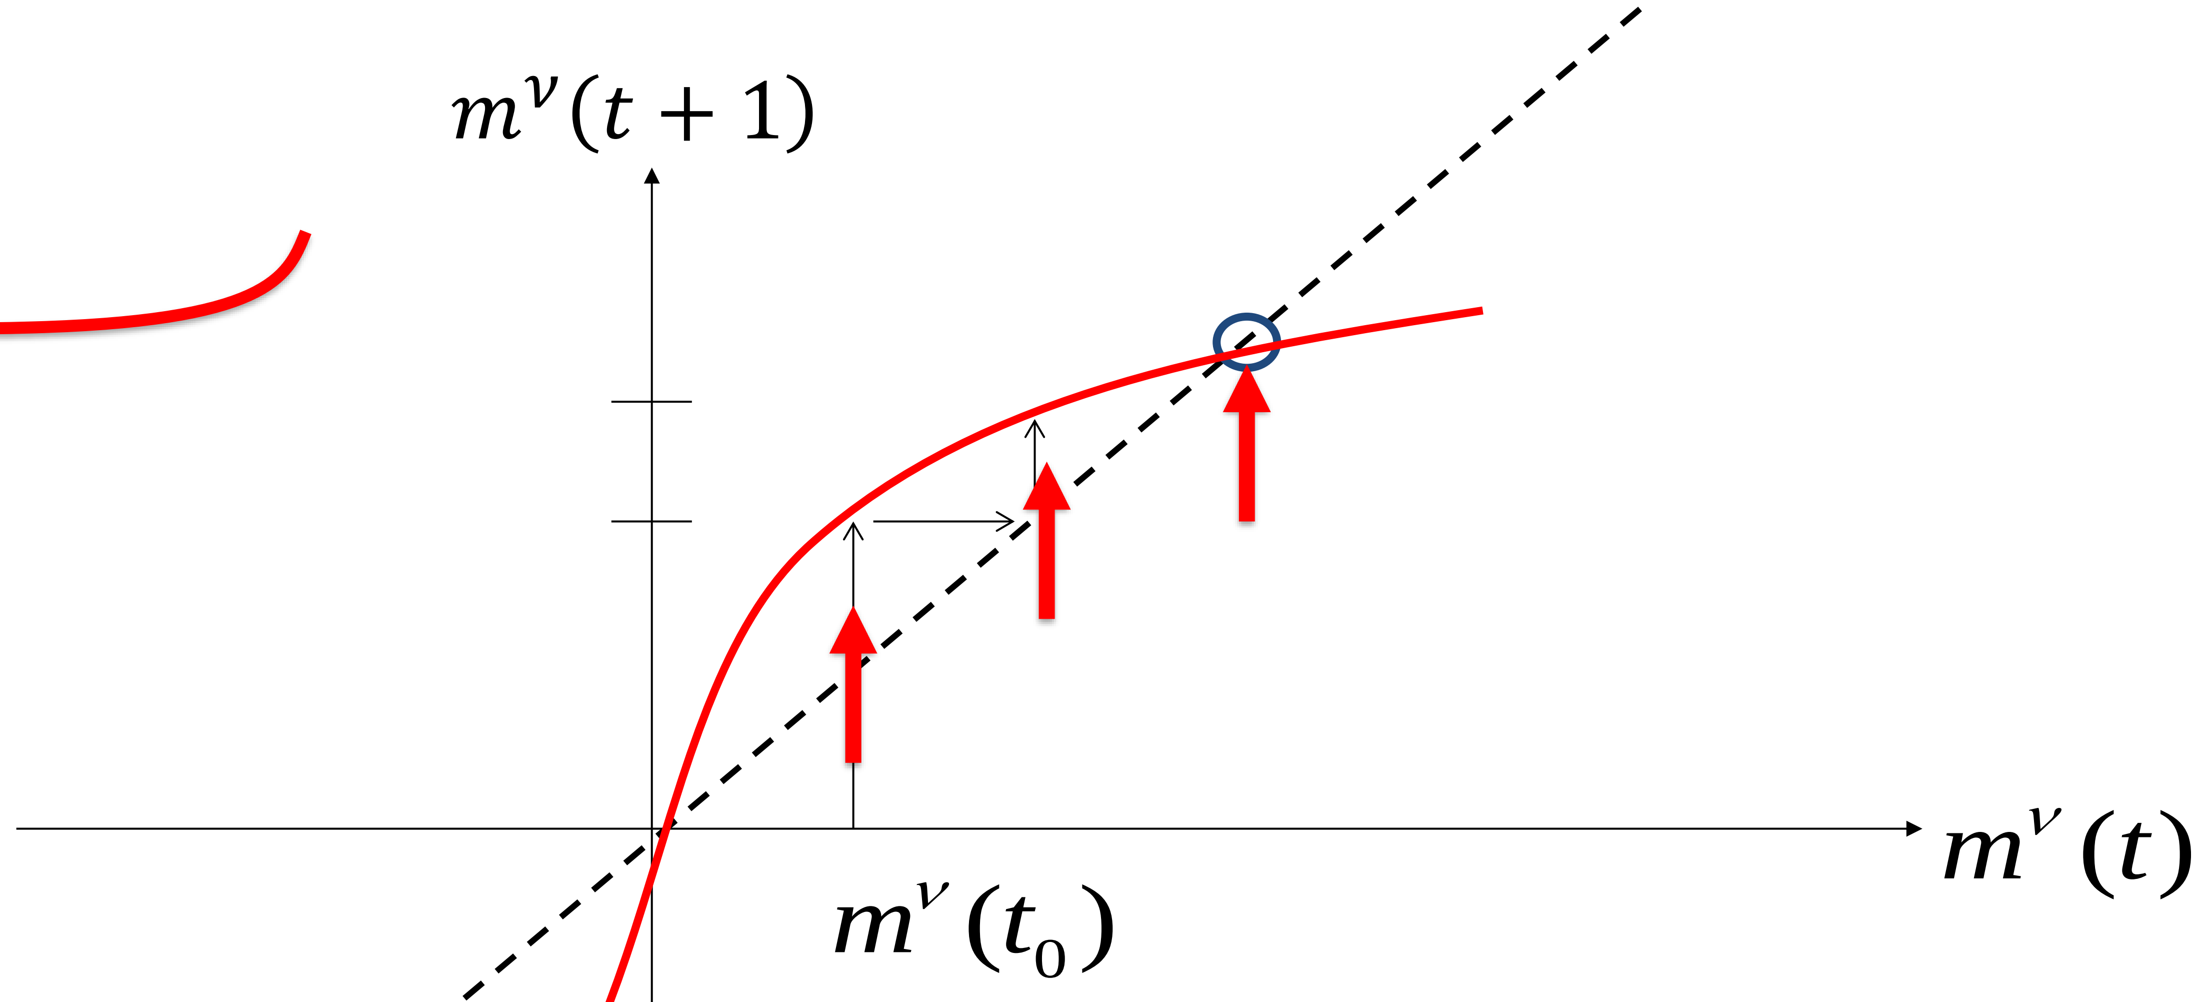

Review: Hopfield model: attractor dynamics

Overlap (definition)

$$m^3(t+1) = \sum_j p_j^3 S_j$$

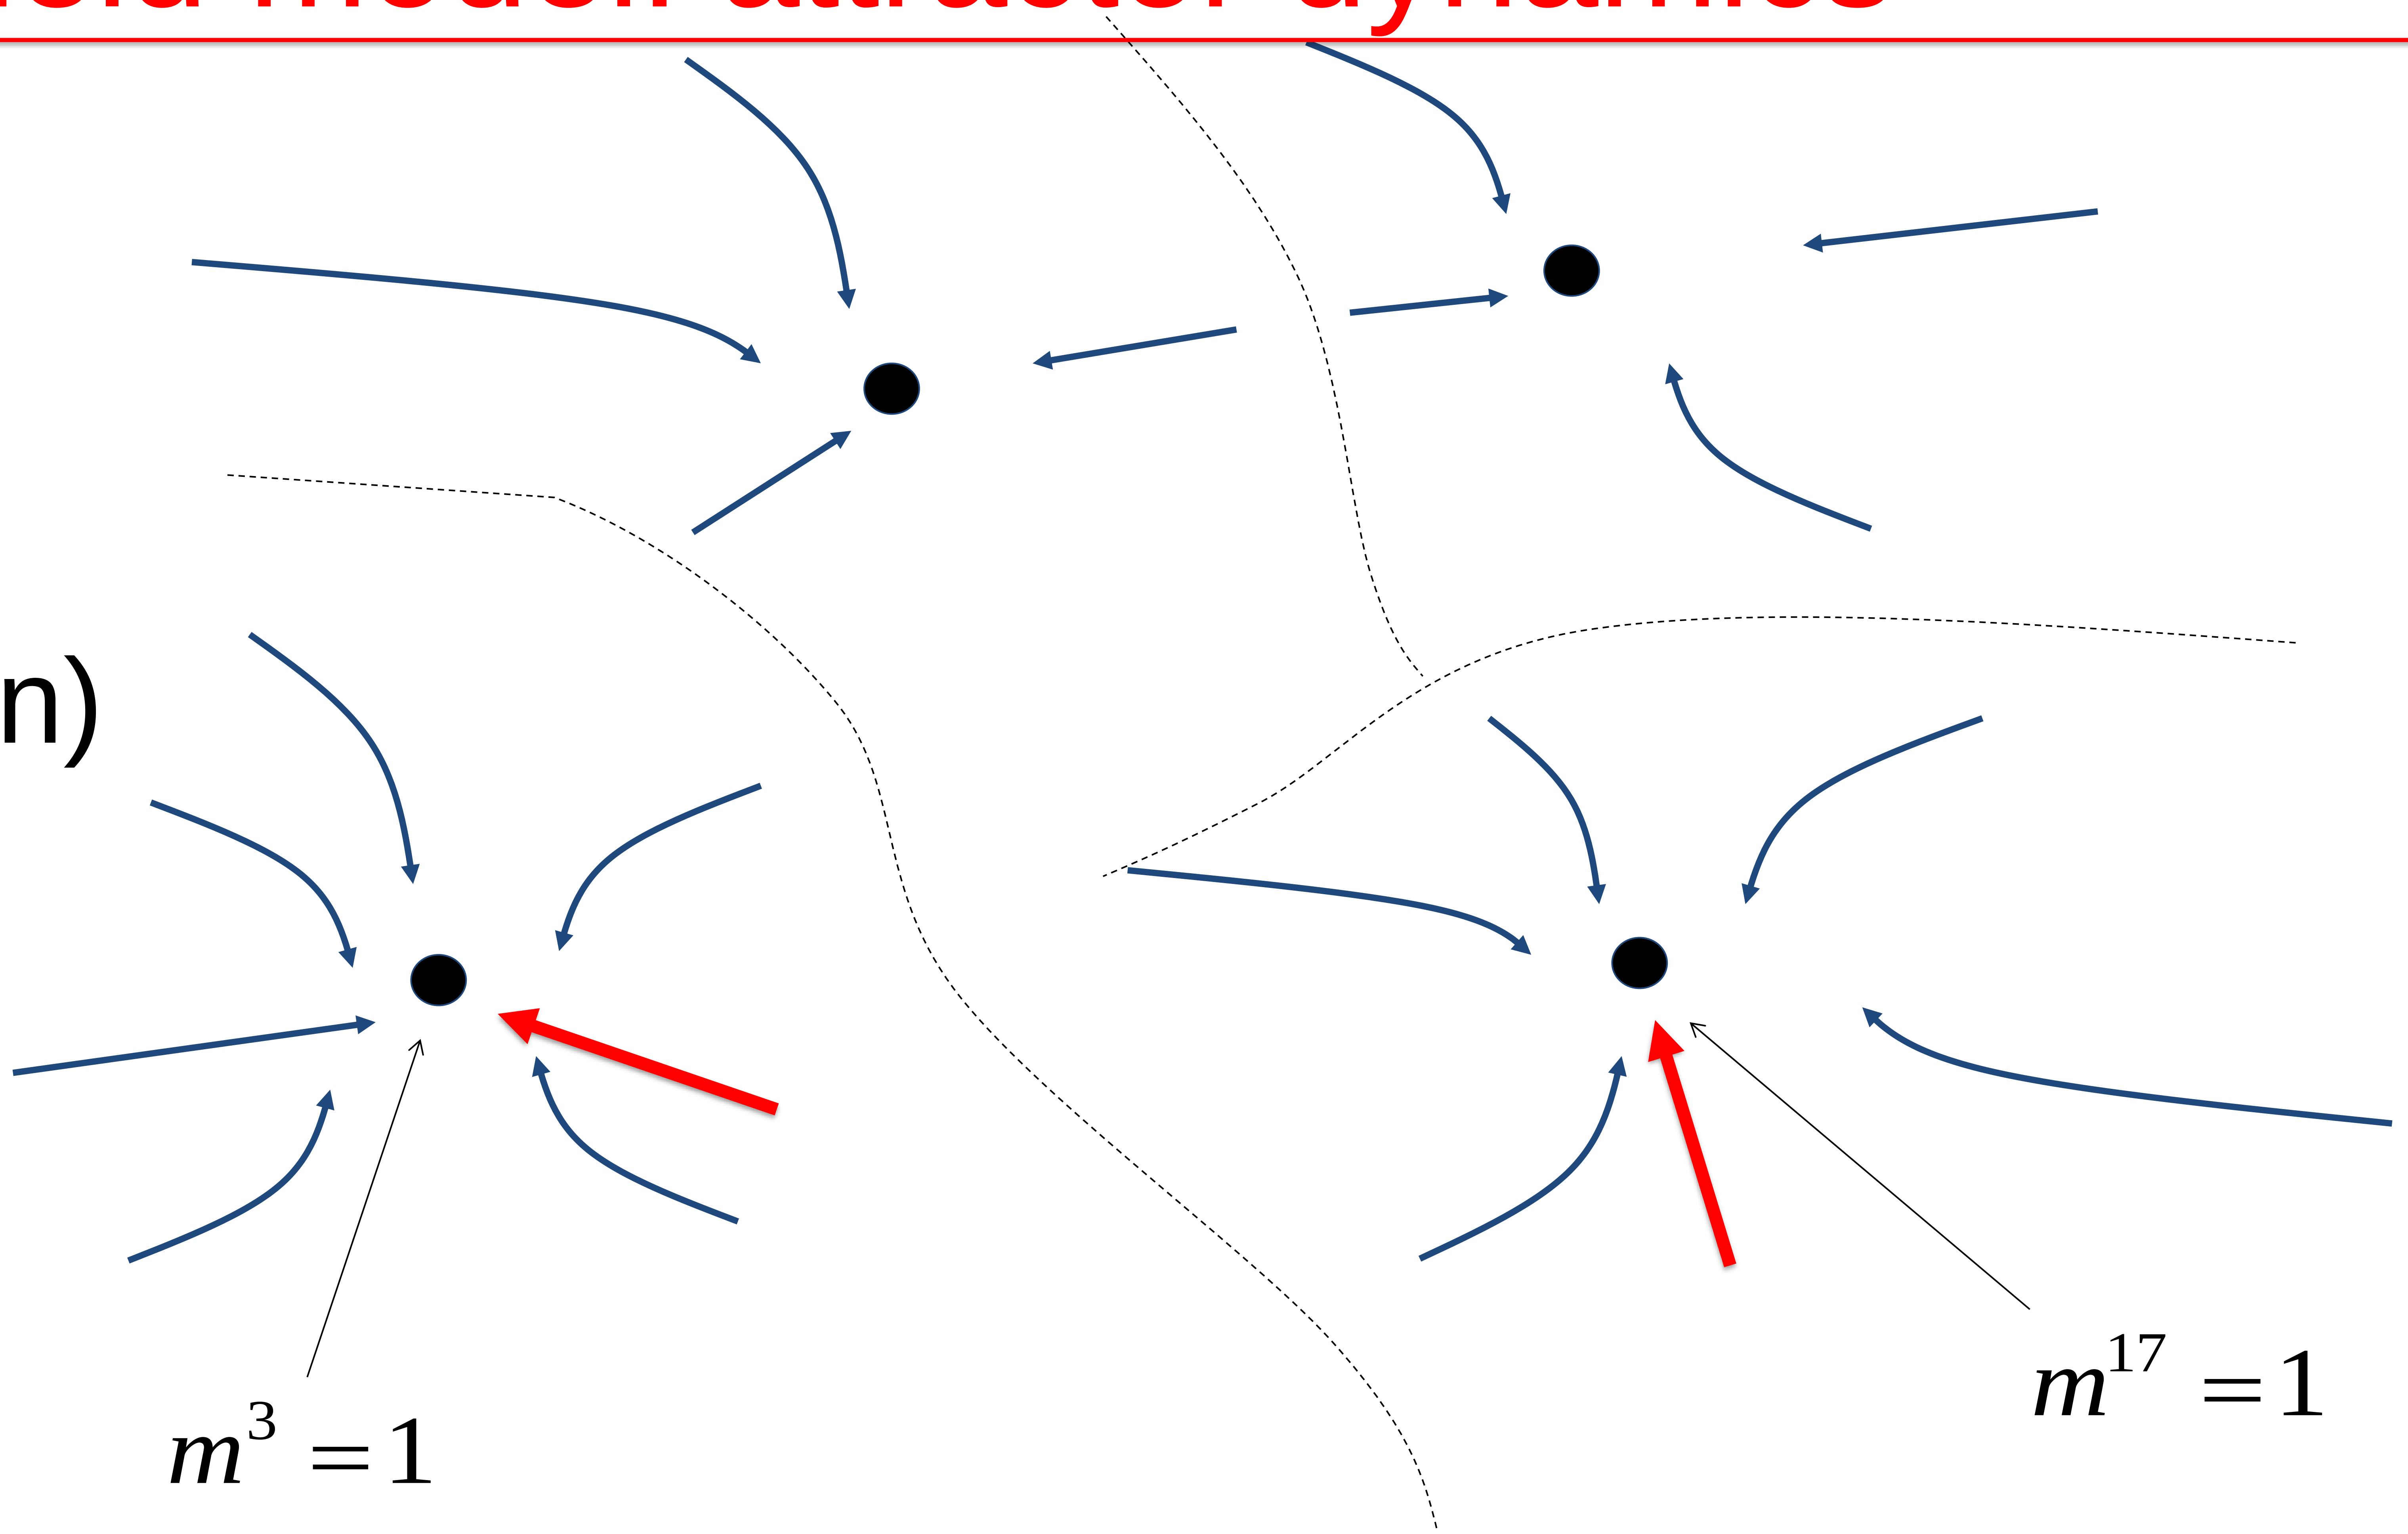

Attractor Dynamics: *https://www.youtube.com/watch?v=9LjFqc2JRk4*

Review: Stochastic Hopfield model: memory retrieval

- Memory retrieval possible with stochastic dynamics

- Fixed point at value with large overlap (e.g., 0.95)

- Random patterns: nearly orthogonal

- Pattern retrieval yields **low-dimensional dynamics,**
  even if '*state*'=N variables (i.e. configuration of all neurons)

Question: are overlap variables m^{ν} 'somehow related'
to the low-dimensional variables κ in experiments?

Neural Manifolds and Low-dimensional dynamics: How can we interpret neural activity?

1. **What are Neural Manifolds?**
   - experimental observations

2. **Two views of Neural Activity**
   - **2.1** first view: computing (Hopfield model)
   - **2.2 2nd view: neural circuits (field model)**

3. **Low-dimensional dynamics**
   - formalism and assumption
   - dynamics

4. **Examples of low-dim dynamics**
   - context-dependent decision making

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Cambridge Univ. Press

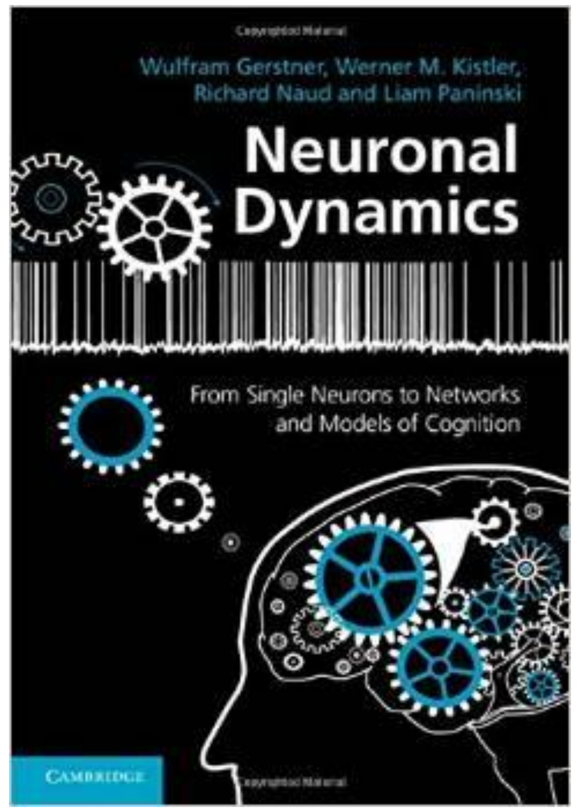

Introduction: How can we understand neuronal activity?

How can we understand principles of neuronal activity?

D. Barack and J. Krakauer, 2021
C. Langdon and T. Engel, 2023

Two different perspectives

- Hopfield model: **low dimensional dynamics**
  (e.g., flow towards fixed point/attractor dynamics)

- Field models for perception:
  each neuron has a receptive field
  (neurons can be classified according to
  functional similarity)

Review: receptive fields and cortical maps

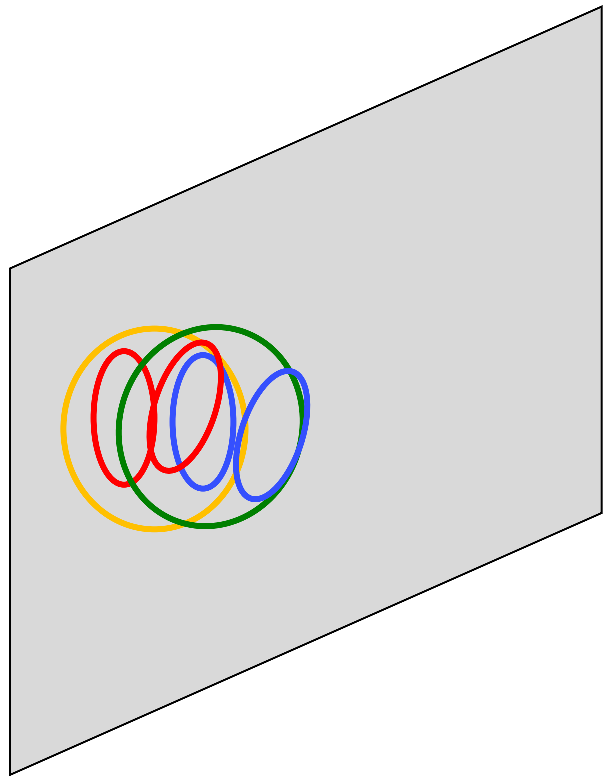

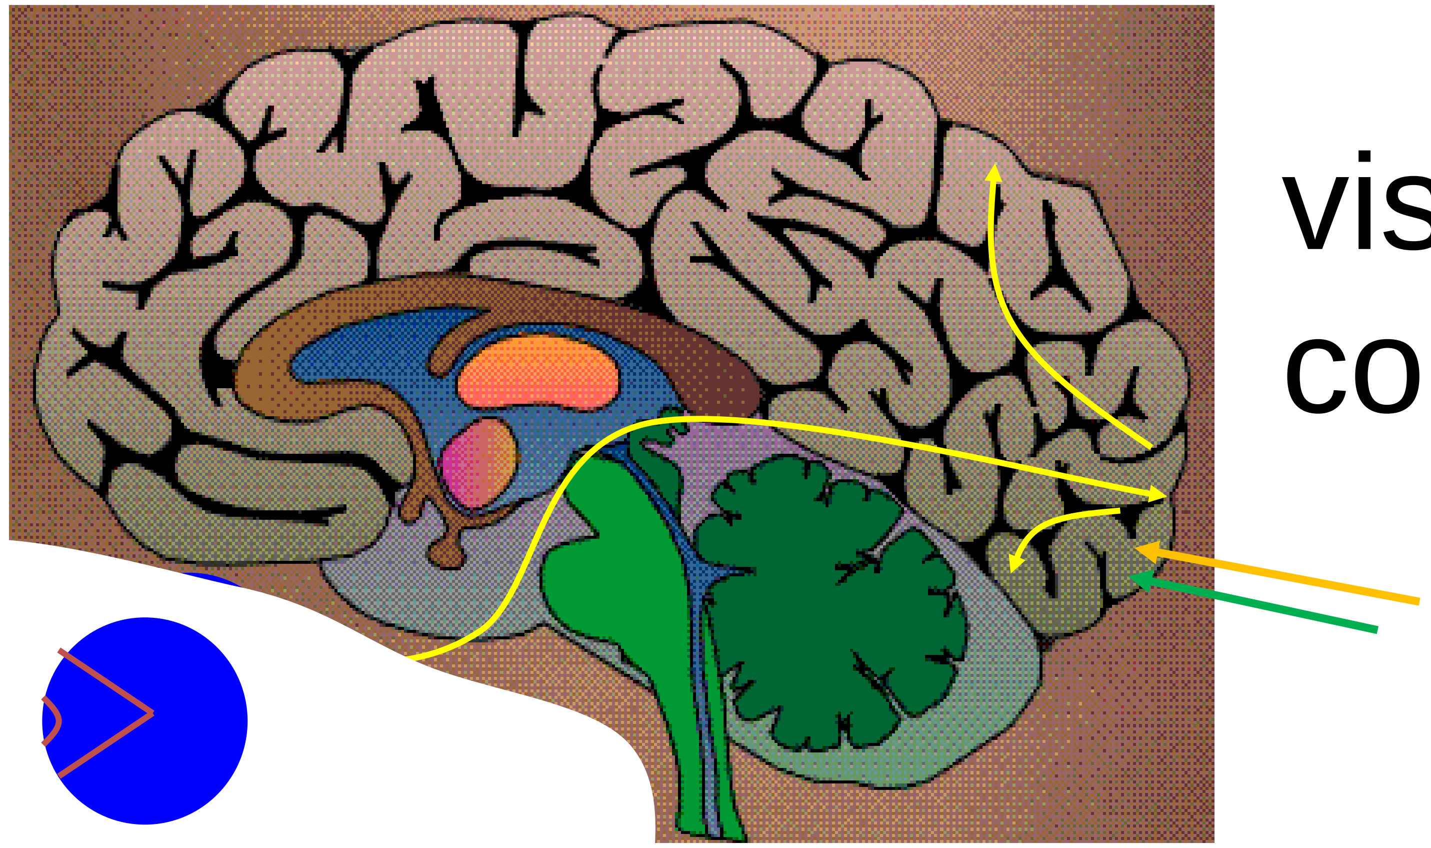

visual cortex

Neighboring cells in visual cortex
- have similar center of receptive field
  → **spatial map of visual field**
- have similar preferred orientation:
  → **cortical orientation map**
- **connectivity stronger between cells with similar orientation**

Hubel and Wiesel 1968; Bonhoeffer&Grinvald, 1991; Bressloff&Cowan, 2002; Kaschube et al. 2010

Review: receptive fields and cortical maps

neighboring neurons: similar orientation and similar RF center
along cortical surface: orientation AND RF center change

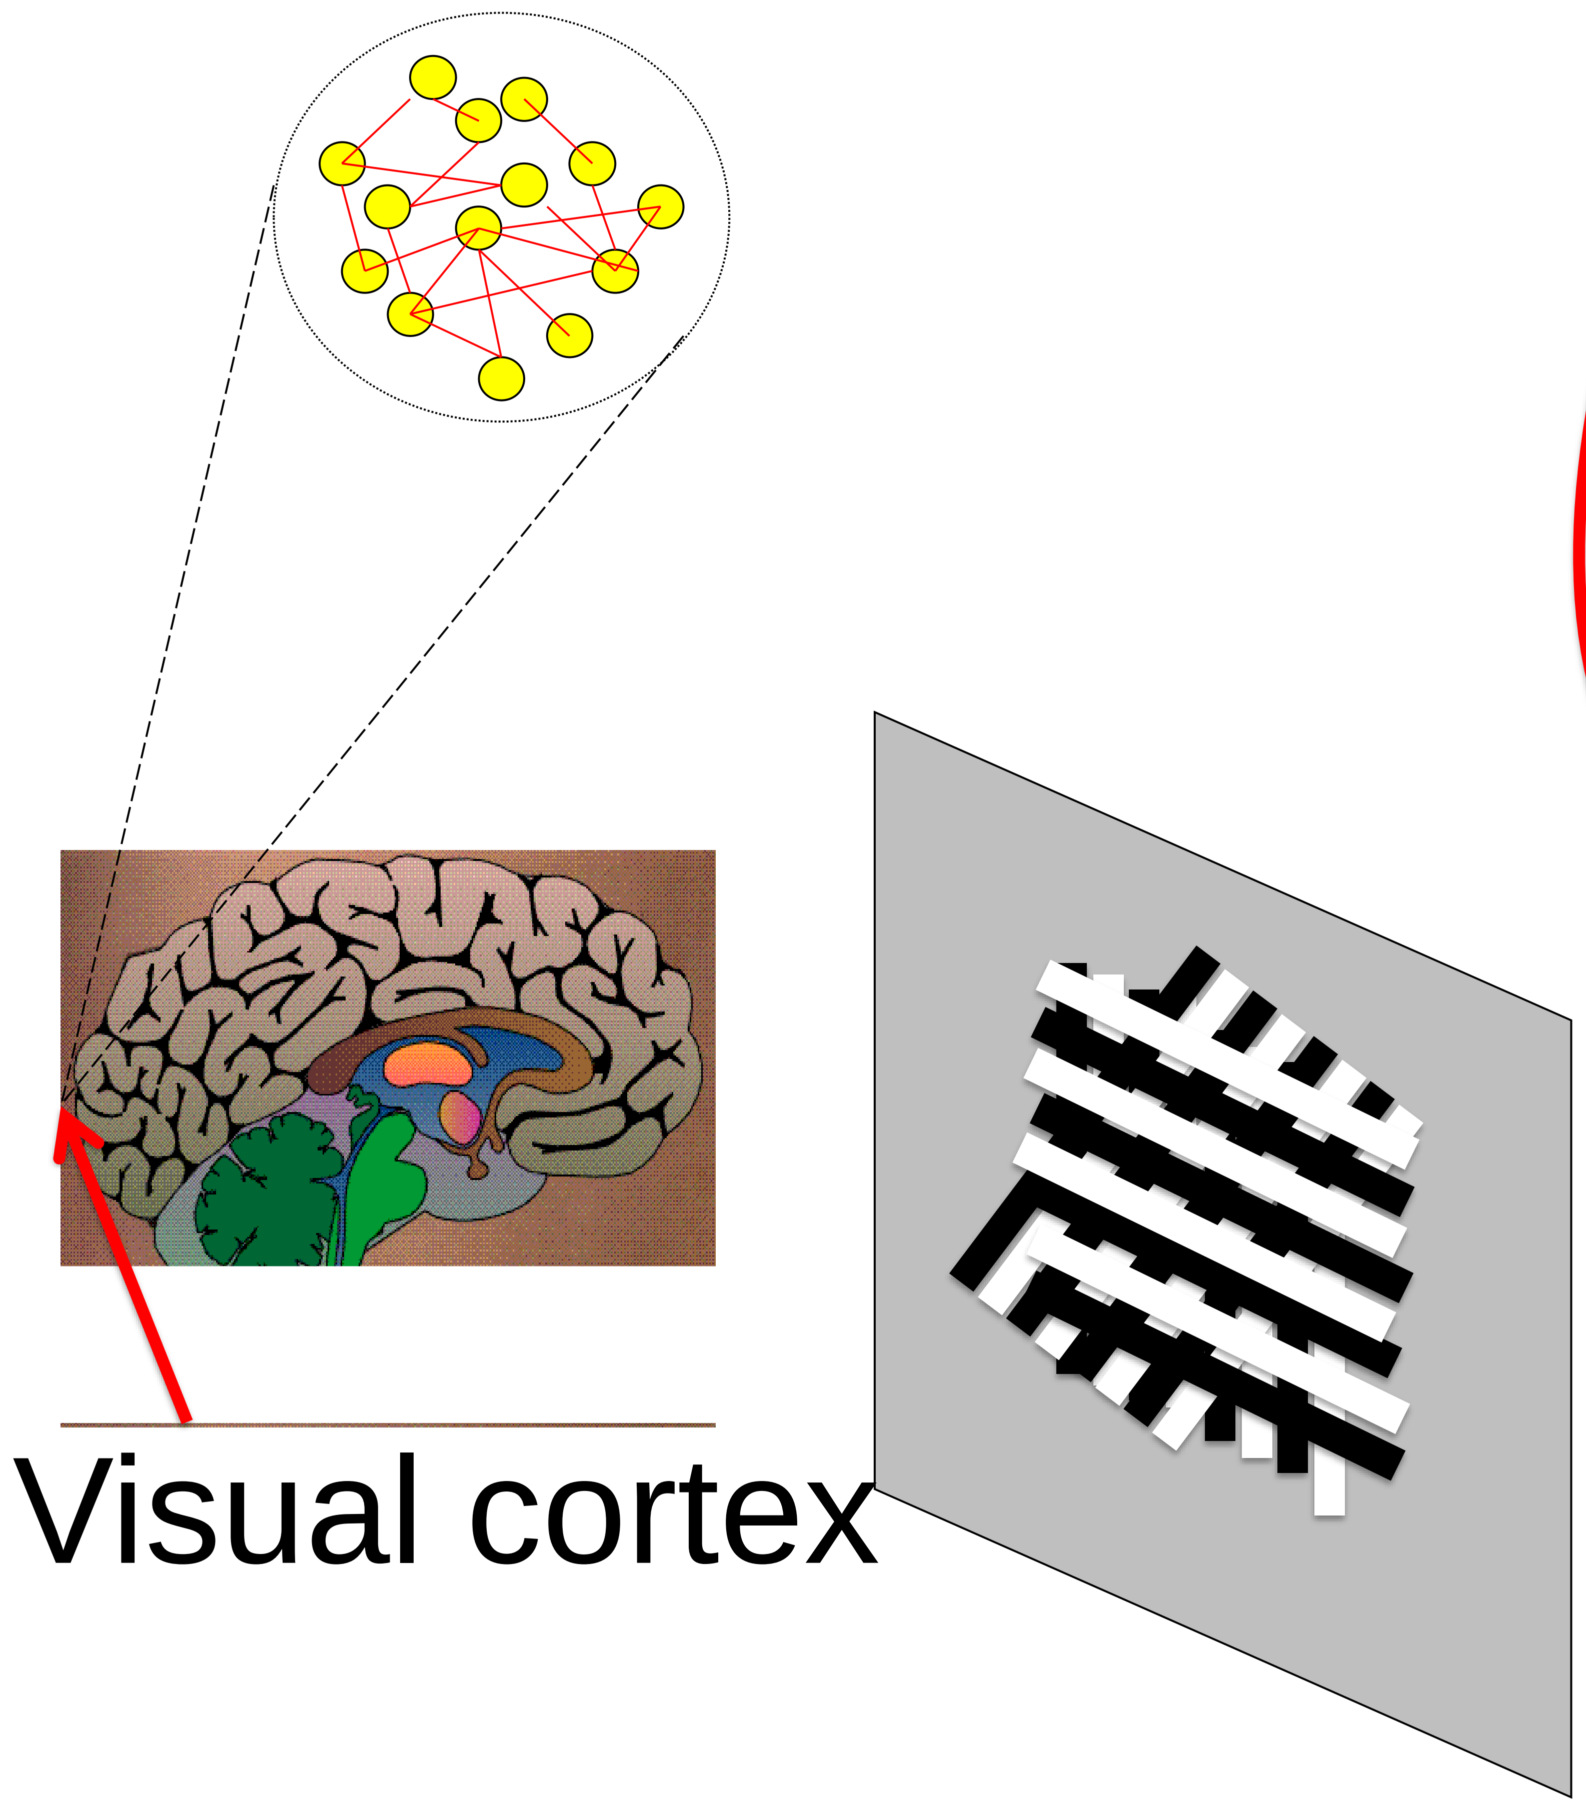

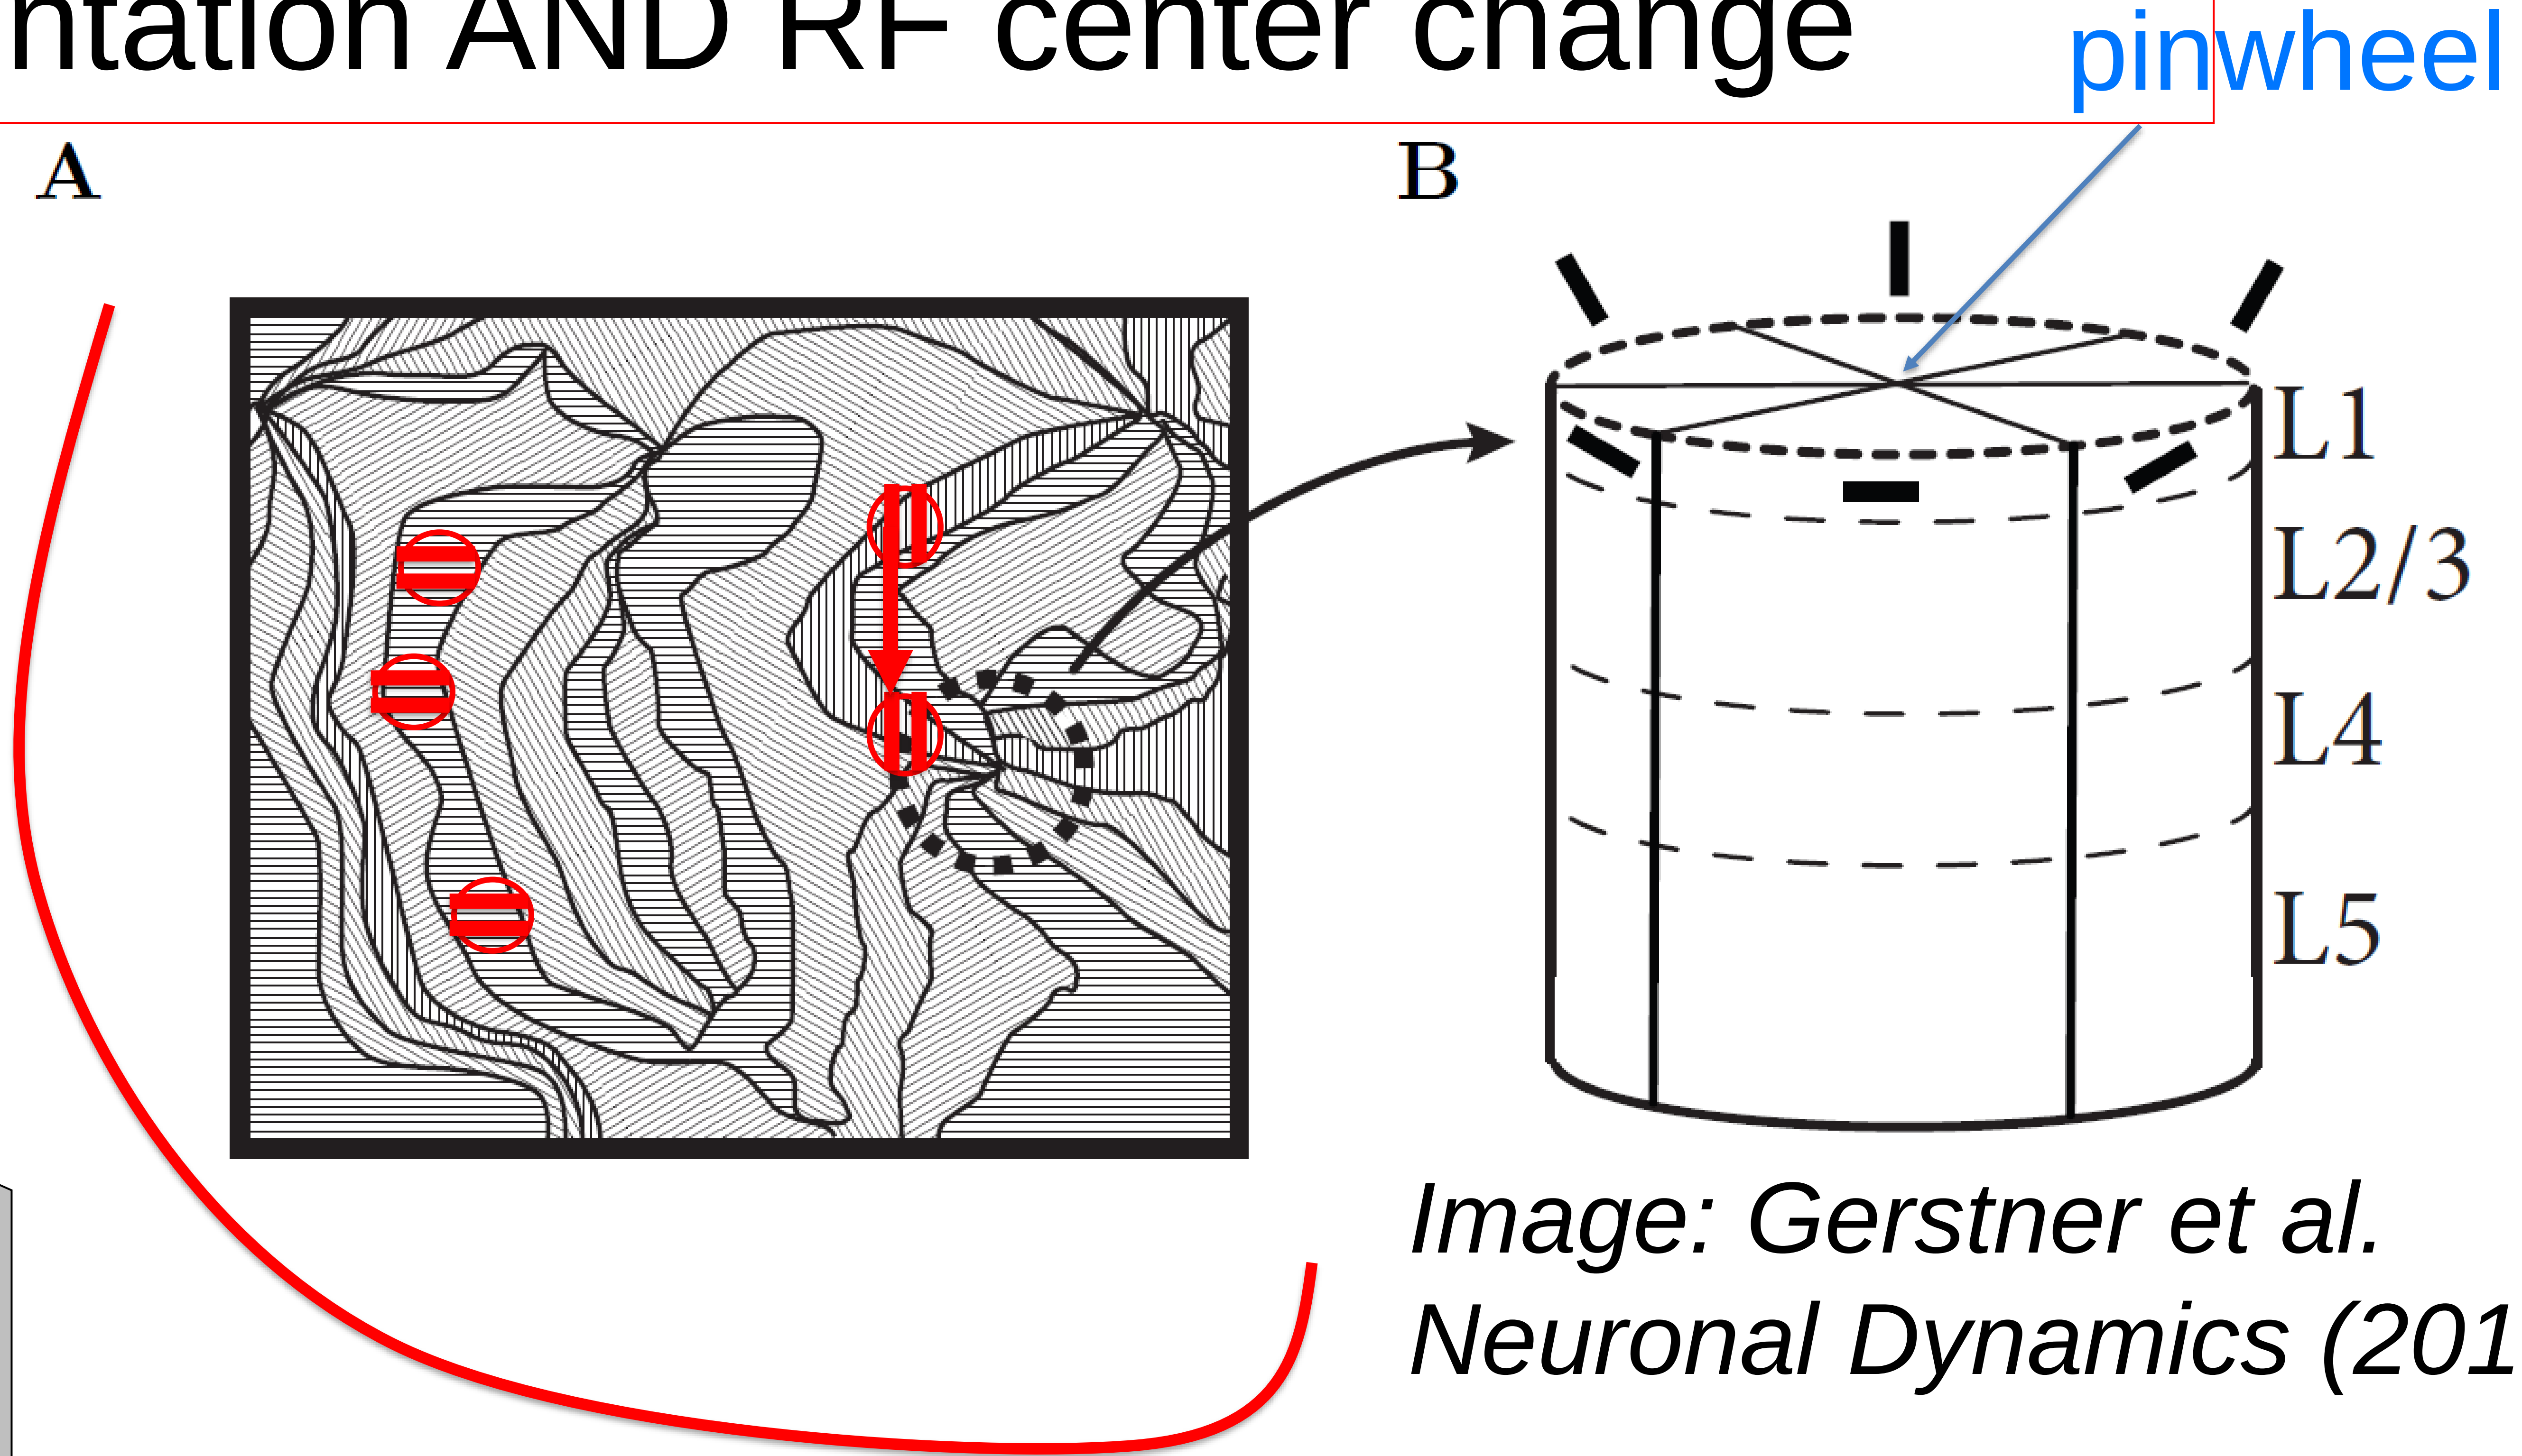

Image: Gerstner et al. Neuronal Dynamics (2014)

Bonhoeffer&Grinvald, 1991; Bressloff&Cowan, 2002; Kaschube et al. 2010

Review: receptive fields and cortical maps

neighboring neurons: similar orientation and similar RF center
along cortical surface: orientation AND RF center change

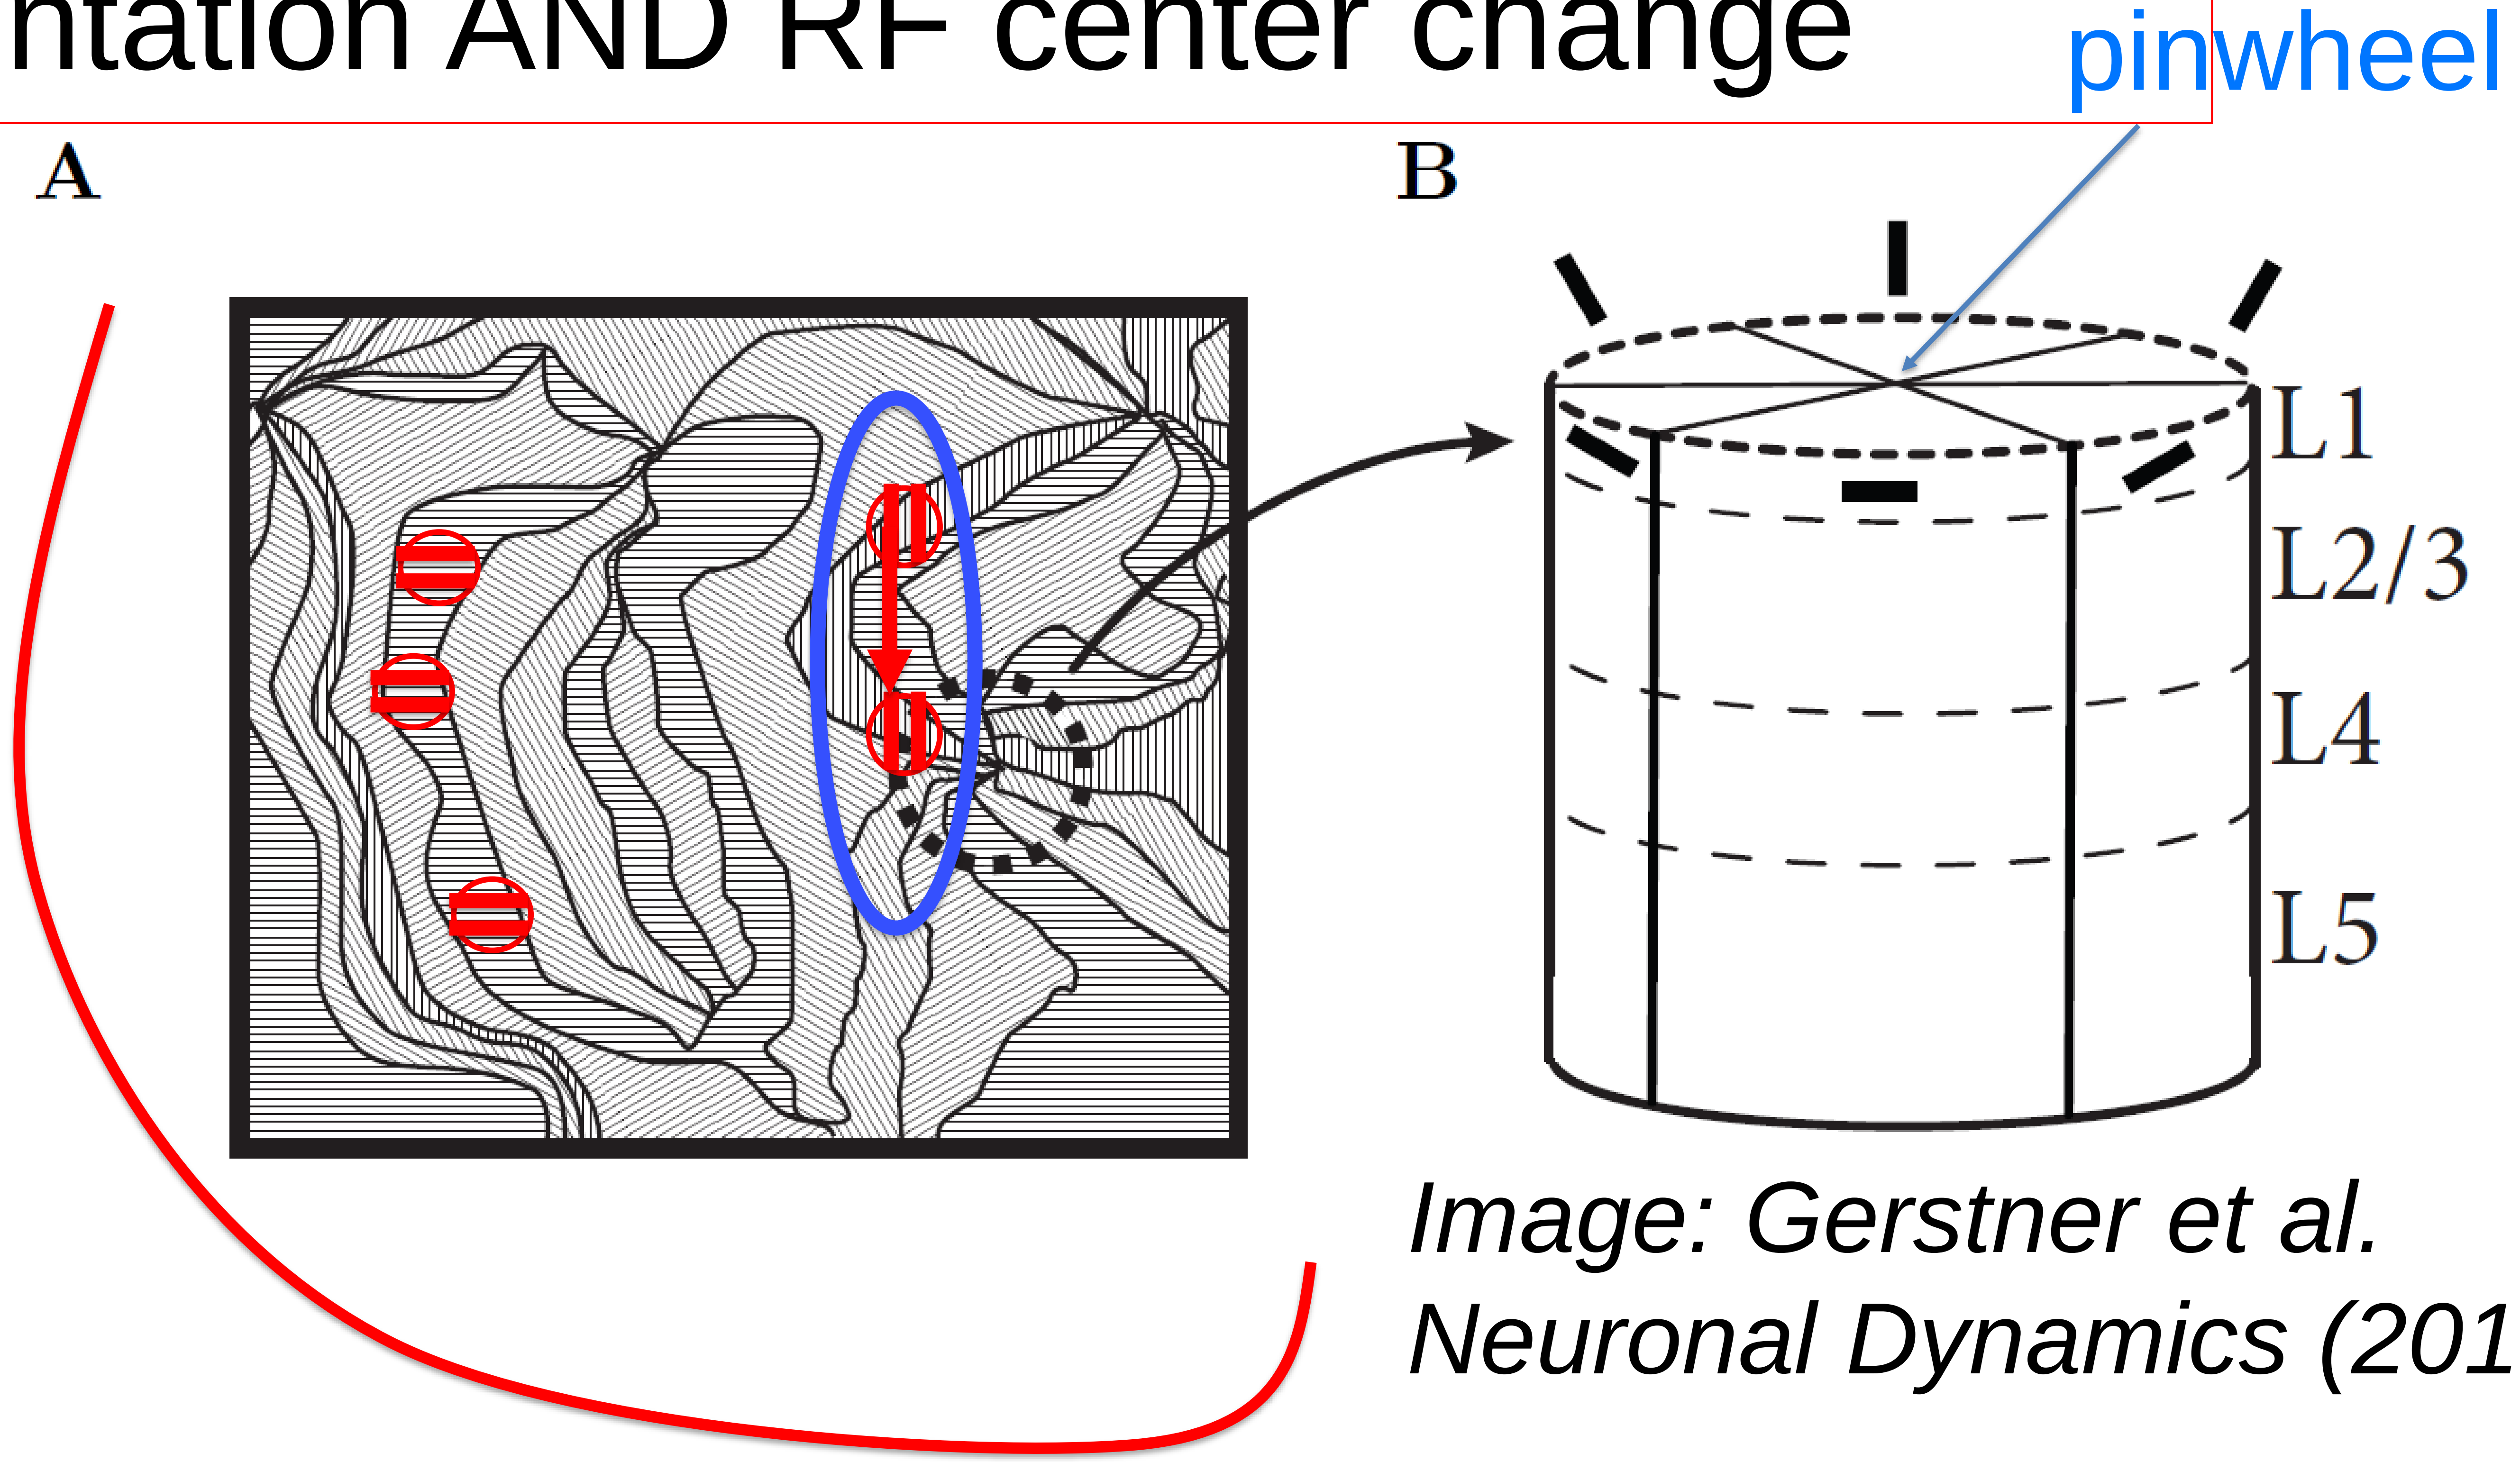

Image: Gerstner et al. Neuronal Dynamics (2014)

Strong connections
- between neurons of similar orientation
- between neurons of similar RF Center

'patchy connectivity'

Bonhoeffer&Grinvald, 1991; Bressloff&Cowan, 2002; Kaschube et al. 2010

Review: functional similarity of neurons

functional characterization of neuron
- orientation of rec. field: z_1
- horizontal placement of rec. field: z_2
- vertical placement of rec. field: z_3

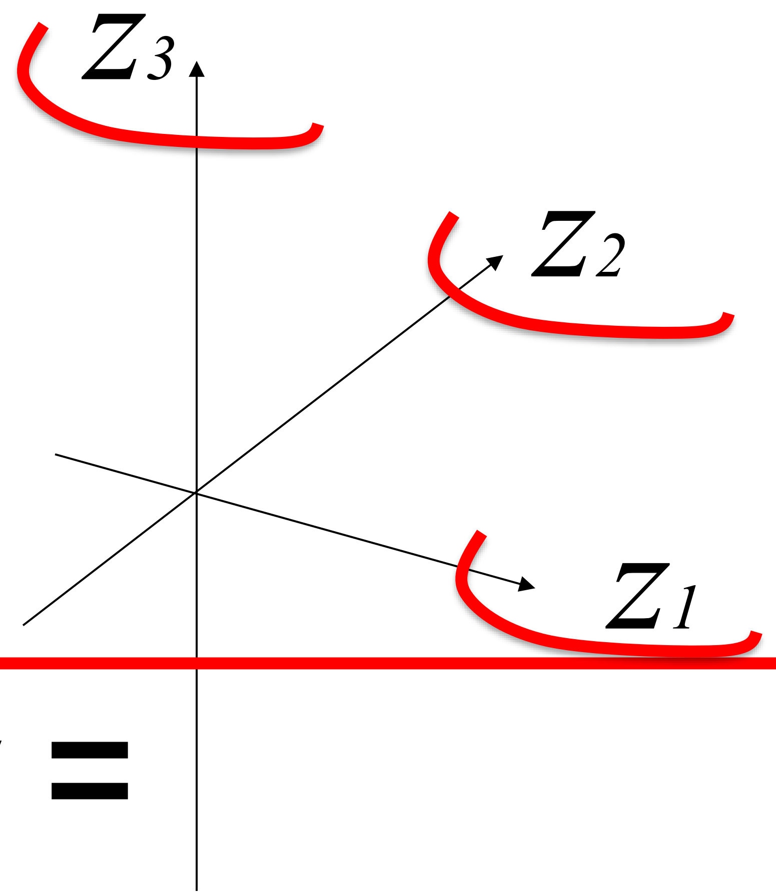

functional similarity = neighborhood in abstract space

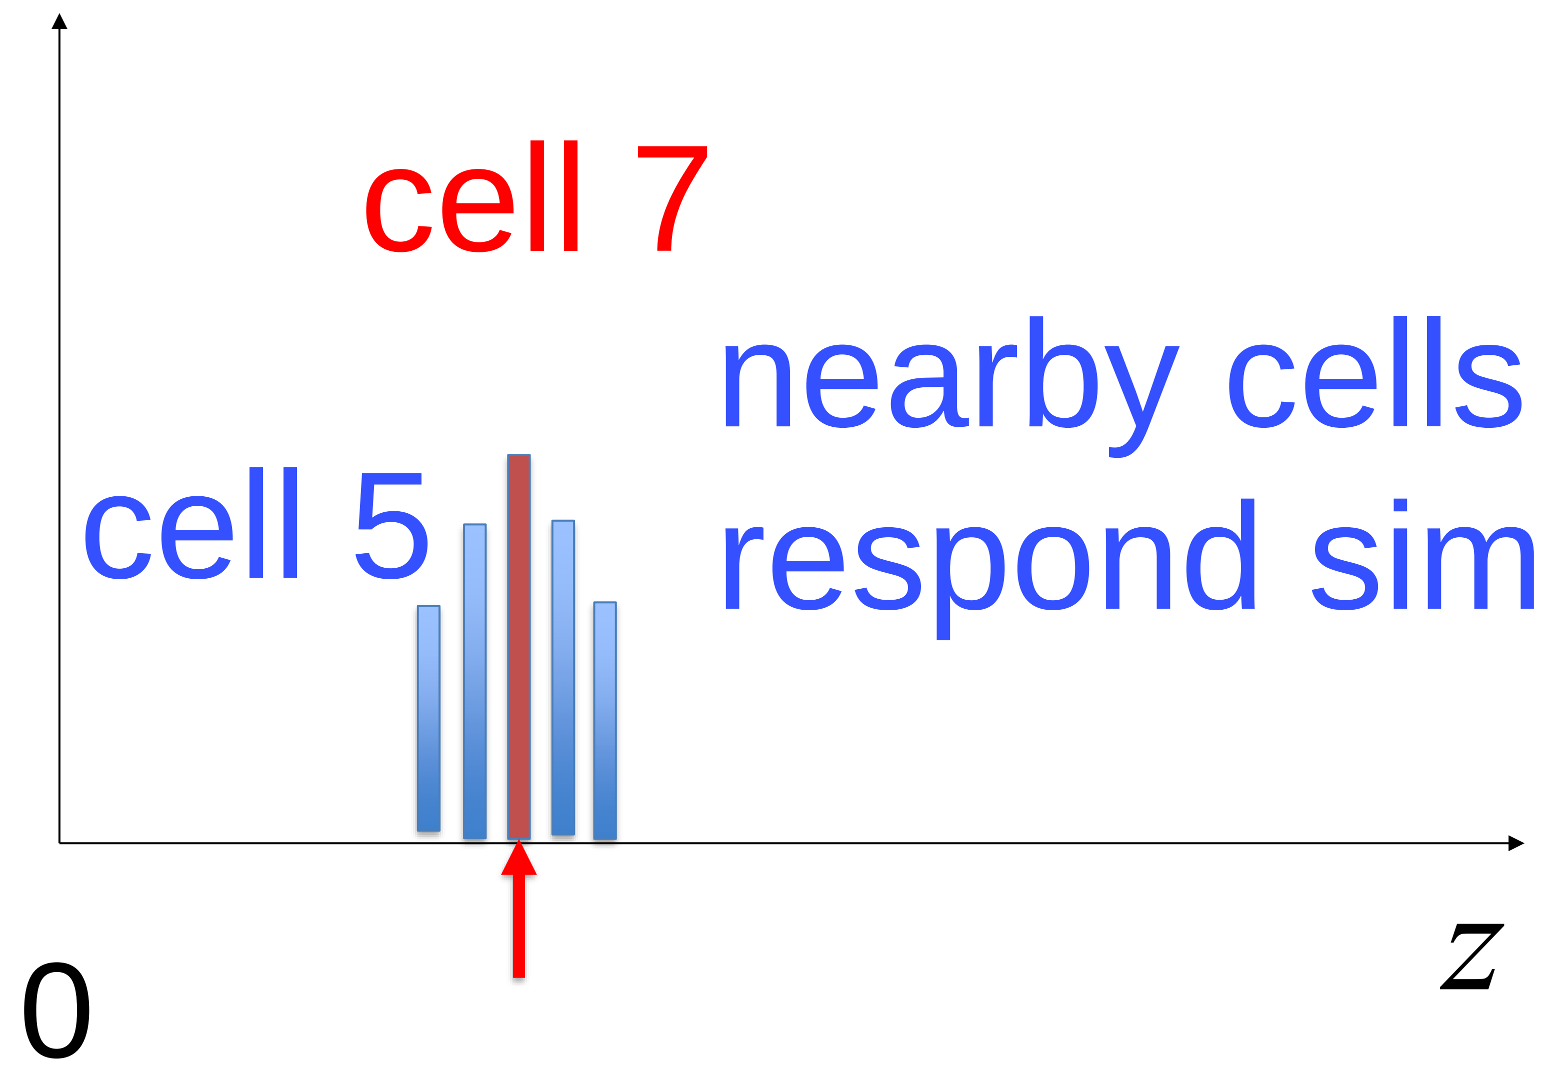

nearby cells (along abstract axis) respond similarly

abstract axis: - a feature of receptive field

a stimulus that maximally excites cell 7

Review: functional similarity of neurons

functional similarity =
neighborhood in abstract space

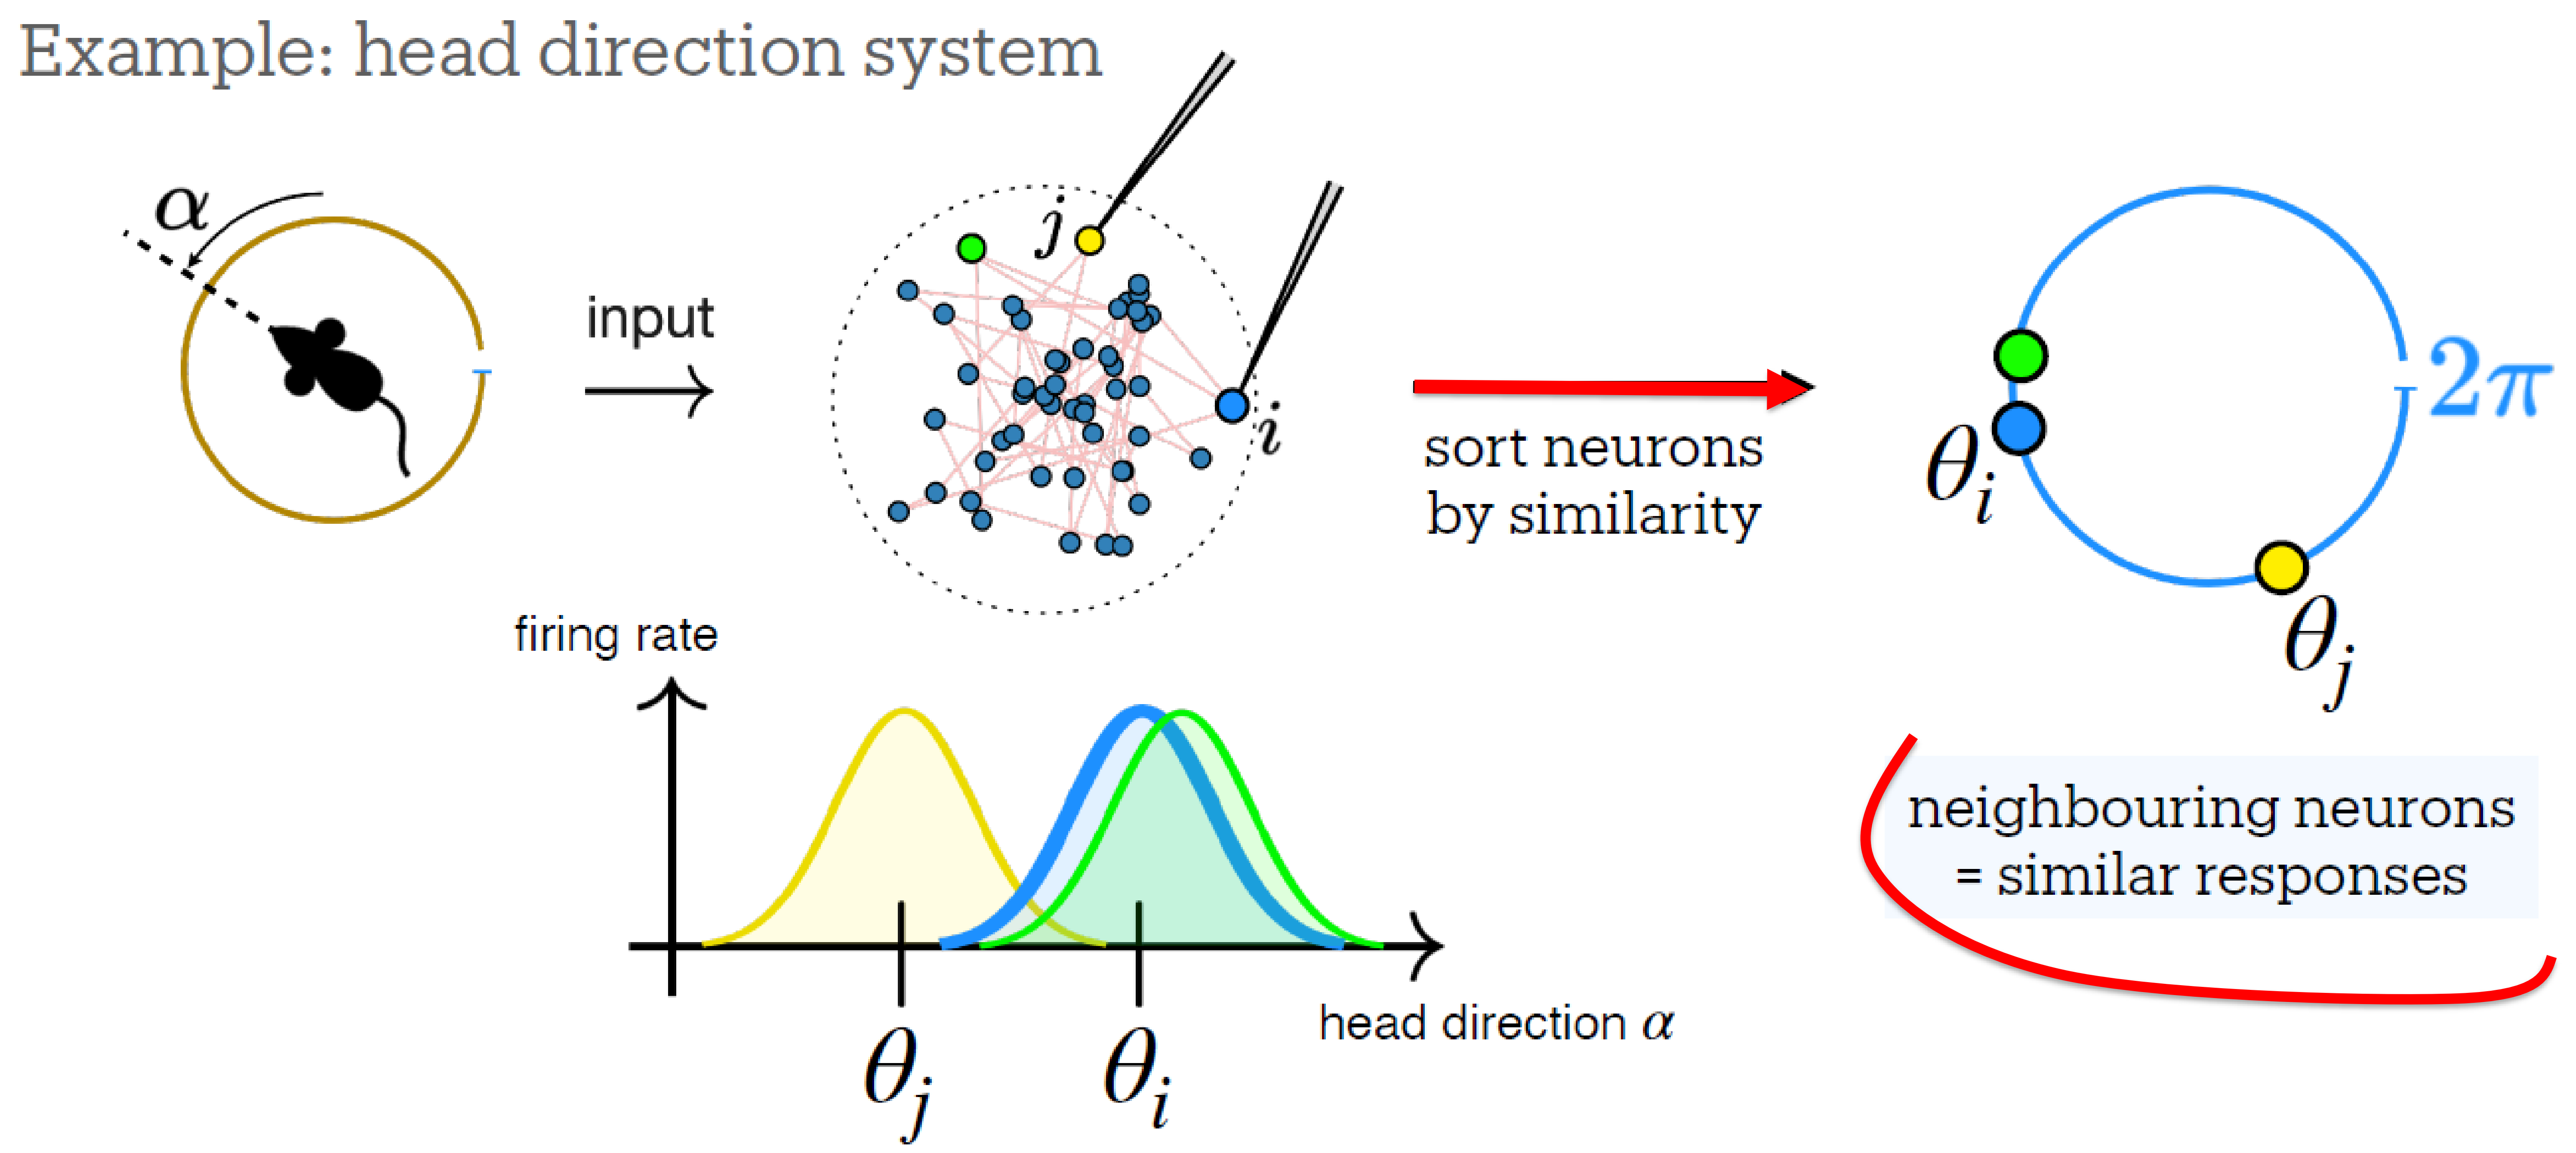

variable z:
position on ring

Review: multiple populations → continuum

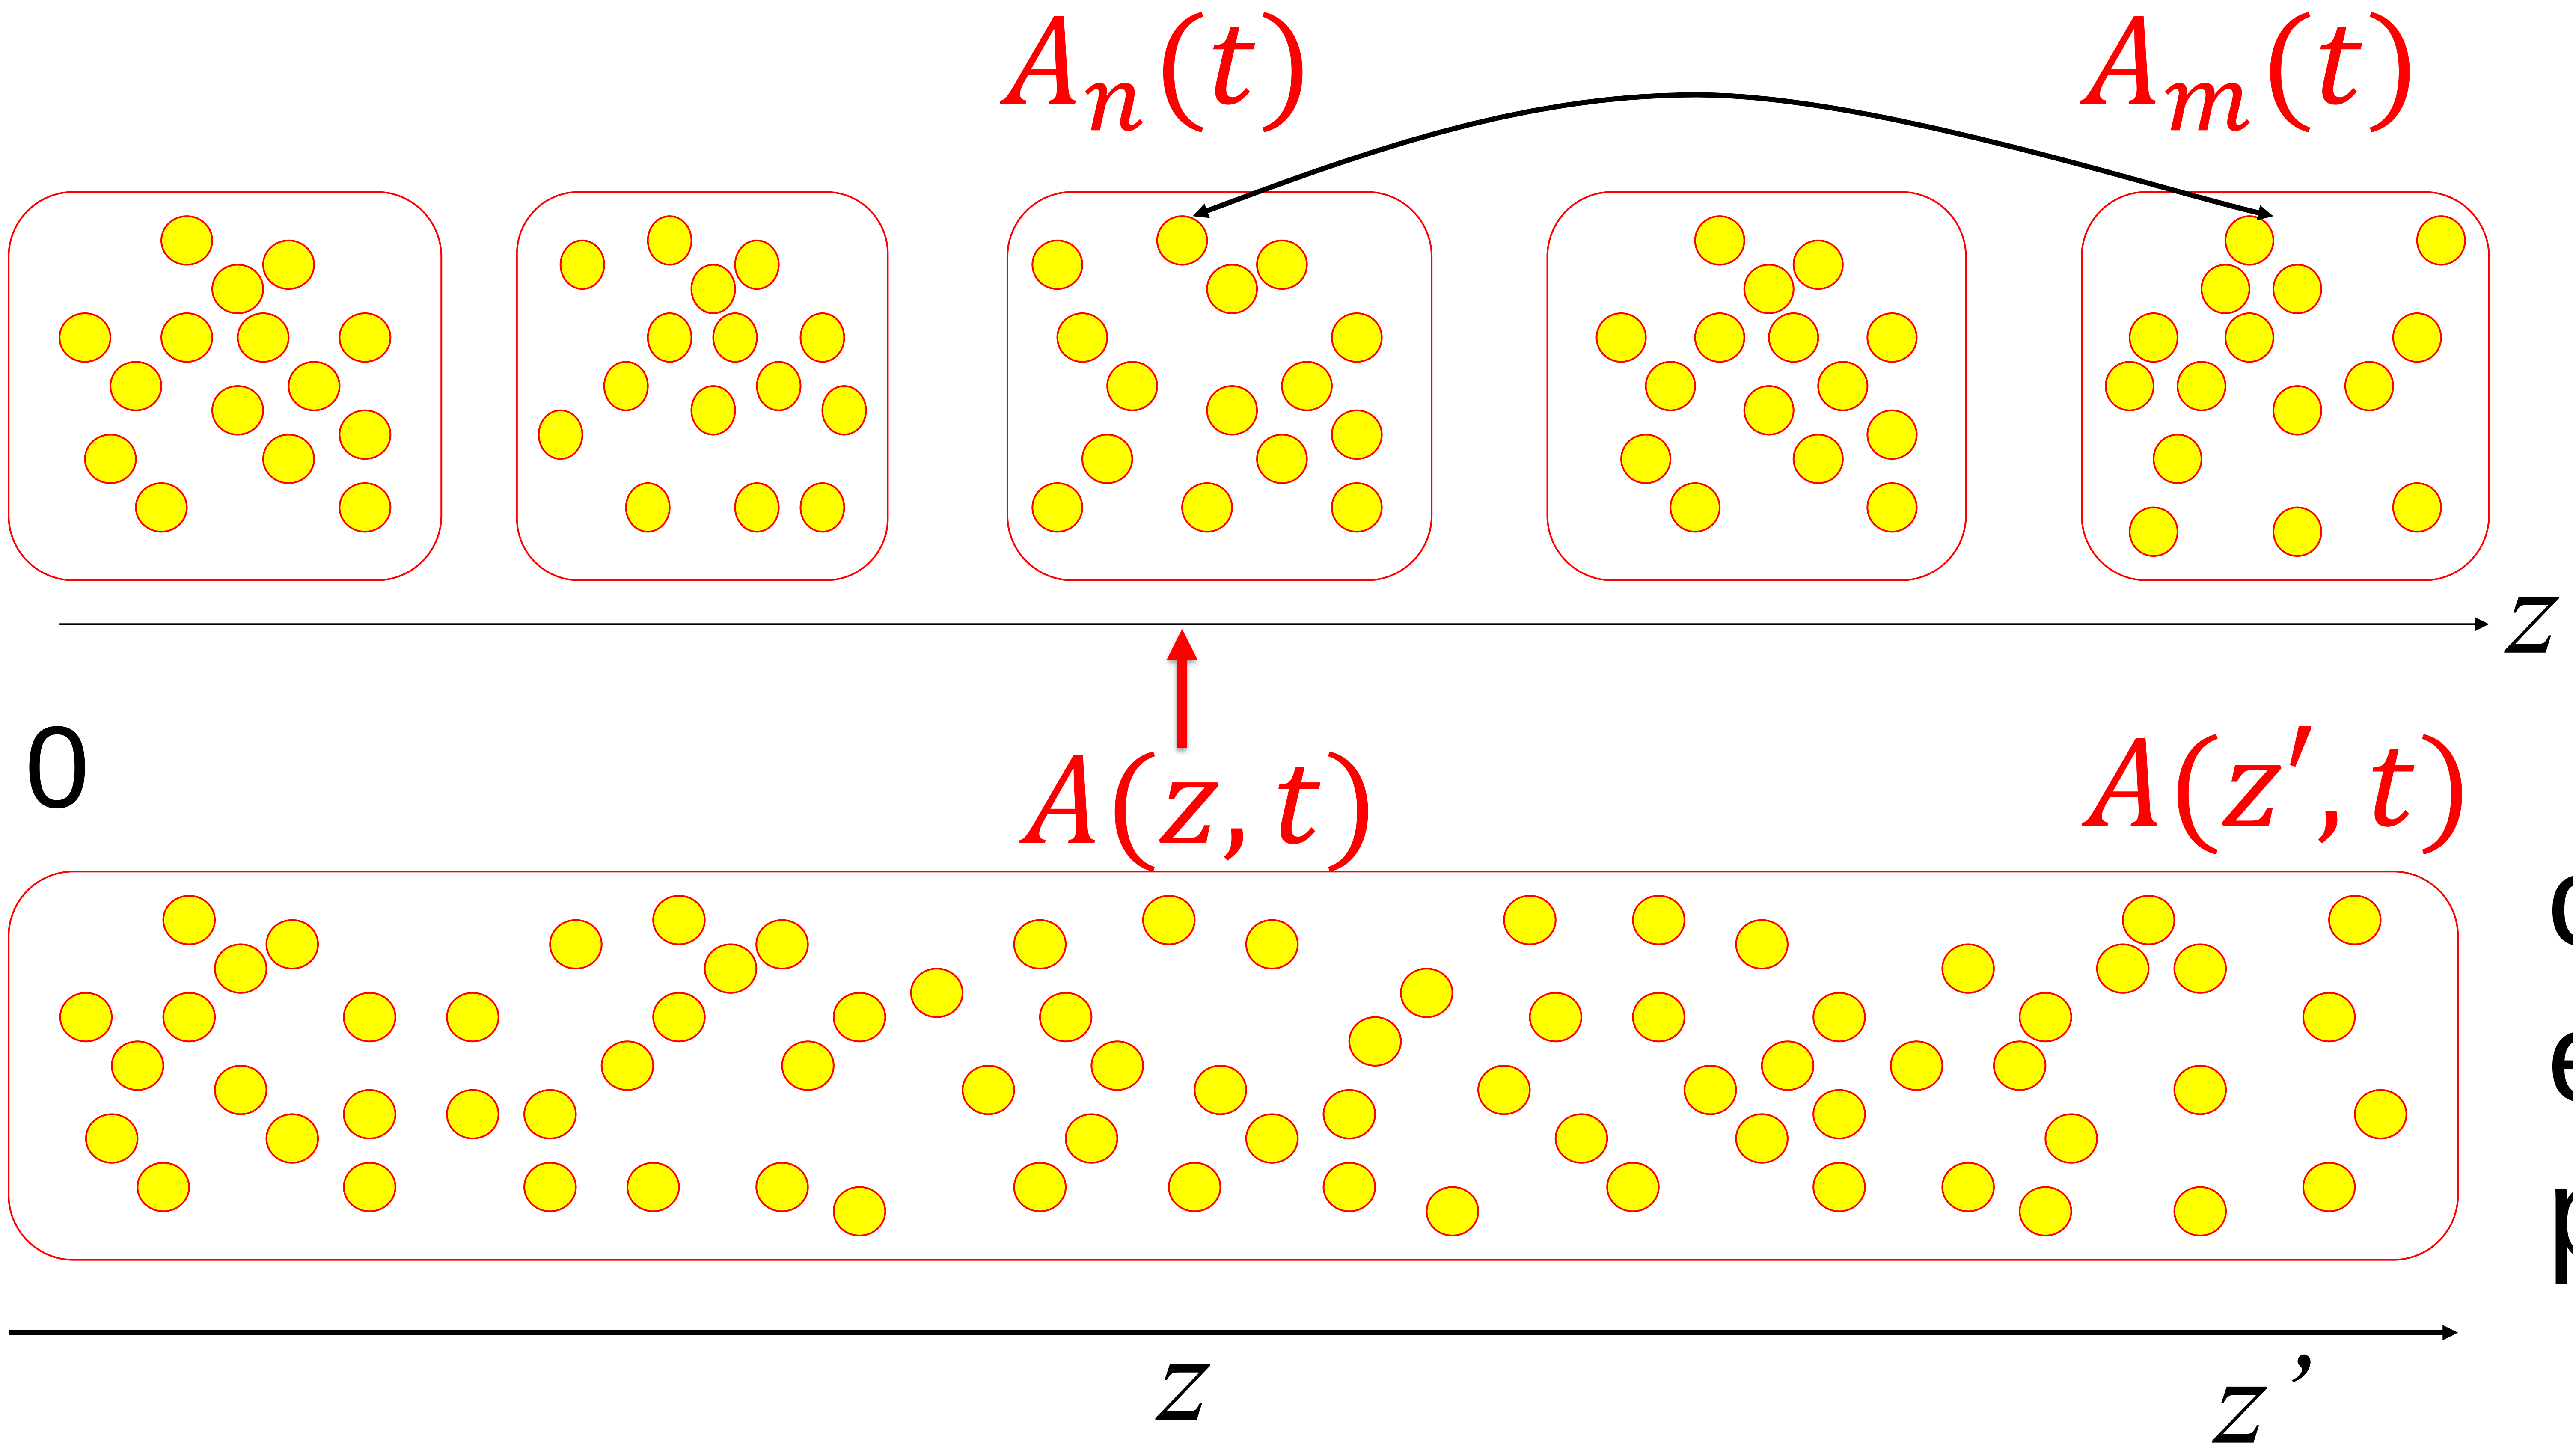

discrete populations

continuous/ extended population

Continuum models/Field Models *https://www.youtube.com/watch?v=f2tadlvuhcg*

Review: Field equation (continuum model)

Wilson and Cowan, 1973

Membrane potential

$$\tau \frac{d}{dt} h(z,t) = -h(z,t) + RI^{ext}(z,t) + \int w(z,z') F(h(z',t)) dz'$$

- field equation = population activity model in continuum
- position **z = abstract variable** (functional similarity)
- **coupling weight depends on functional similarity**:
- neurons with ‘similar function’ strongly connected

- population activity (rate)

$$A(z',t) = F(h(z',t))$$

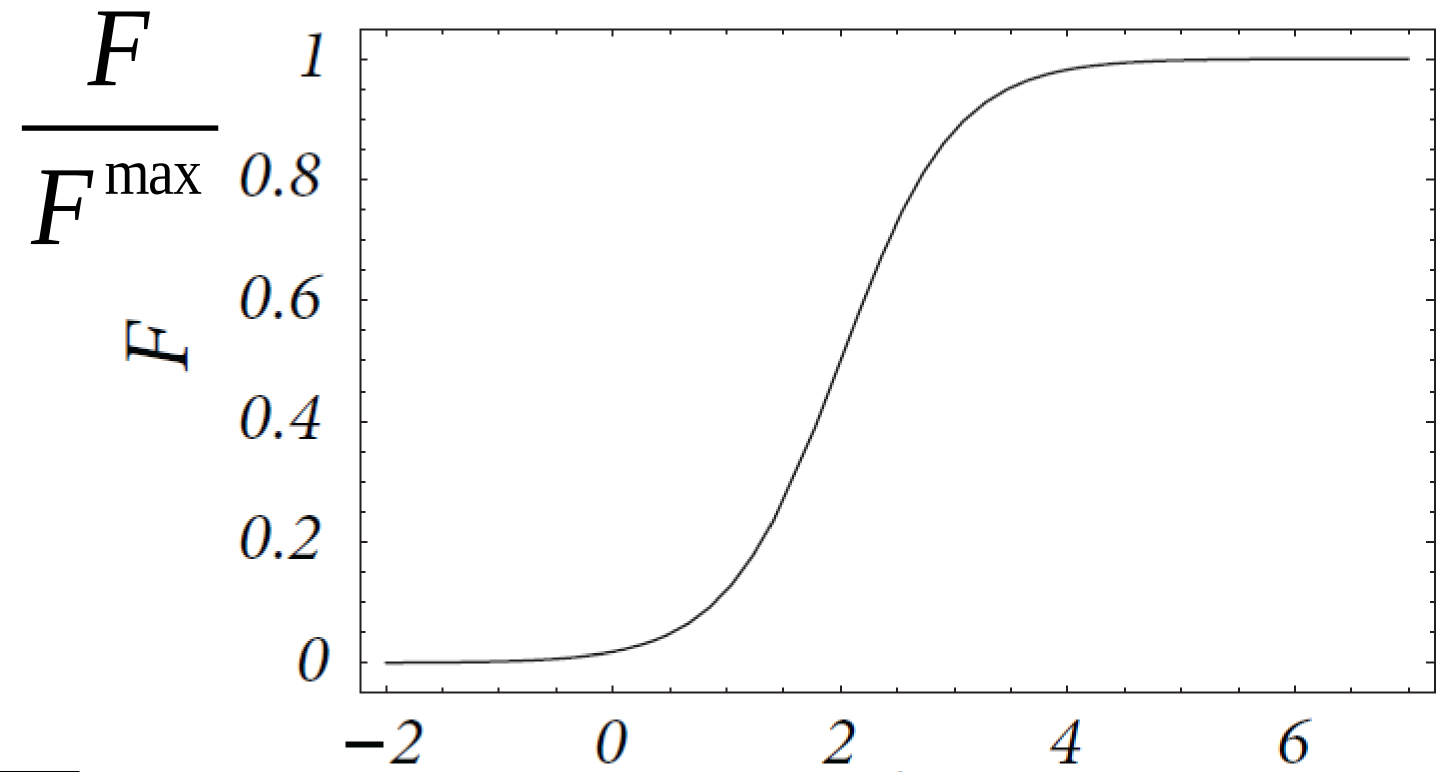

Summary/review: Field equation

A population rate model in continuous space is also called a field equation.

$$\tau \frac{d}{dt} h(z,t) = -h(z,t) + RI^{ext}(z,t) + \int w(z,z')F(h(z',t))dz'$$

Here the variable z can be interpreted as an **abstract quantity**, such as the orientiation of the preferred visual stimulus: **Functional similarity**

In the general model w(z,z') could be an arbitrary function; but in most field equations it is taken as a distance-dependent function w(z-z'). **Connectivity is stronger between cells with similar 'functional role'.**

A classic choice is the Mexican-Hat function with long-range inhibition and short-range excitation. Note that in real neural networks, inhibition involves a separate class of neurons.

Summary: How can we interpret neural activity?

How can we understand principles of neuronal activity?

D. Barack and J. Krakauer, 2021
C. Langdon and T. Engel, 2023

Two different perspectives

- **low dimensional dynamics** → Hopfield model
  (e.g., flow towards fixed point/attractor dynamics)

→ continuum model

- **neurons and functional similarity**
  (functional similarity reflected in wiring,
  wiring causes dynamics)

→Relation between the two views? Relation to known models?

Summary:
There are **two different perspectives** on how to interpret neuronal activity:

- The classic view since Hubel and Wiesel was to start with **receptive fields.** We can then define **functional similarity** between neurons as neurons with similar receptive fields. On the theory side, this view has led to **field models** where neurons are organized along one or several abstract axis. Functionally similar neurons have typically stronger (more positive) connections to each other than to functionally different neurons. Hence **wiring** reflects functional similarity.

- The modern view is that neurons perform computational and that these computations can be described by a **flow or dynamics in low-dimensional manifolds**: Even though modern experiment probe the activity of hundreds of neurons simultaneously, we do not need 100 variables to describe the activity but only a few. On the theory this is similar to mean-field models or the Hopfield model. In the Hopfield model, we have encountered **effective variables** (‘overlap’) that describe the **collective dynamics.**

The question of the following videos is how the two views are connected to each other and to standard models of computational neuroscience

Neural Manifolds and low-dimensional dynamics

List of video lectures on Computational Neuroscience, organized by topics:
https://lcnwww.epfl.ch/gerstner/NeuronalDynamics-MOOCall.html

YouTube Channel:
https://www.youtube.com/@gerstnerlab

References:

Mastrogiuseppe, F., Ostojic, S.: Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. Neuron 99(3), 609–62329 (2018)

→ **Barack, D.L., Krakauer, J.W.: Two views on the cognitive brain.** Nature Reviews Neuroscience 22(6), 359–371 (2021)

→ **Langdon, C., Genkin, M., Engel, T.A.: A unifying perspective on neural manifolds and circuits for cognition.** Nature Reviews Neuroscience 24(6), 363–377 (2023)

Pezon, L., Schmutz, V, Gerstner, W. (2024), Linking Neural Manifolds to Principles of Circuit Structure in Recurrent Networks bioRxiv doi: https://doi.org/10.1101/2024.02.28.582565

EPFL

Neural Manifolds and Low-dimensional dynamics: Low-Rank Recurrent Neural Networks

1. What are Neural Manifolds?
   - experimental observations

2. Two views of Neural Activity
   - computing (Hopfield model)
   - neural circuits (field model)

3. Low-rank recurrent networks
   - formalism of low-rank networks
   - dynamics

4. Examples of low-dim dynamics
   - context-dependent decision making

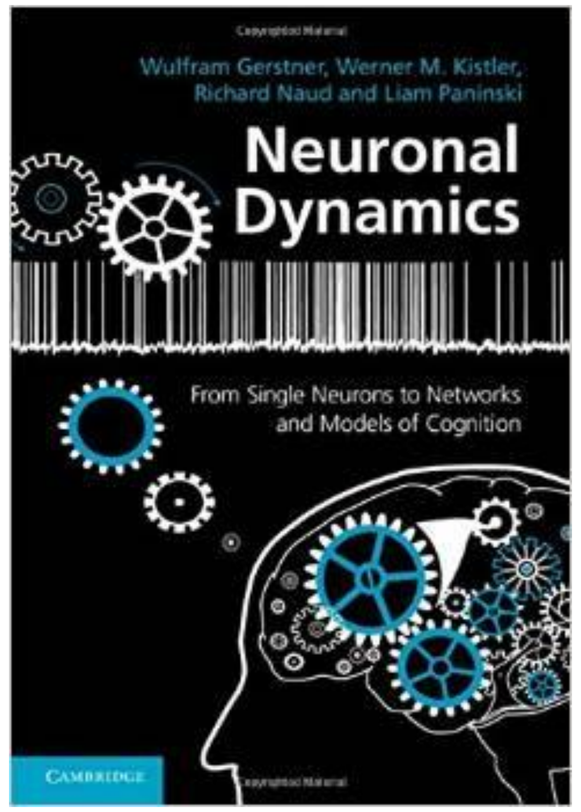

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Cambridge Univ. Press

Recurrent Neural Network (RNN)

Recurrent network of N neurons.

Membrane potential of neuron i:

$$\frac{d}{dt}h_i(t) = -\frac{1}{\tau}h_i(t) + \sum_j W_{ij}\,\phi(h_j(t))$$

Firing rate:

$\phi(h)$

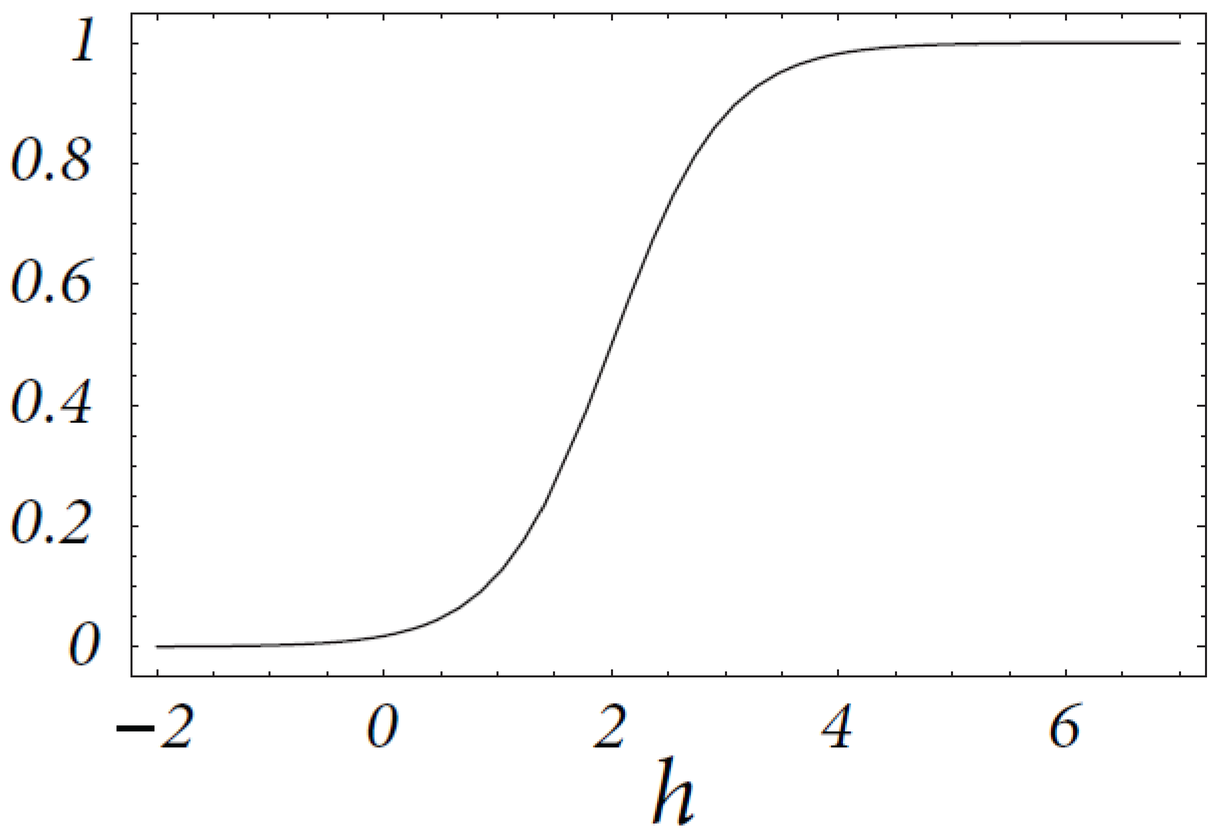

Three assumptions

Assumption 1: neurons are functionally characterized by features

functional characterization of neuron
- orientation of rec. field: z_1
- horizontal placement of rec. field: z_2
- vertical placement of rec. field: z_3

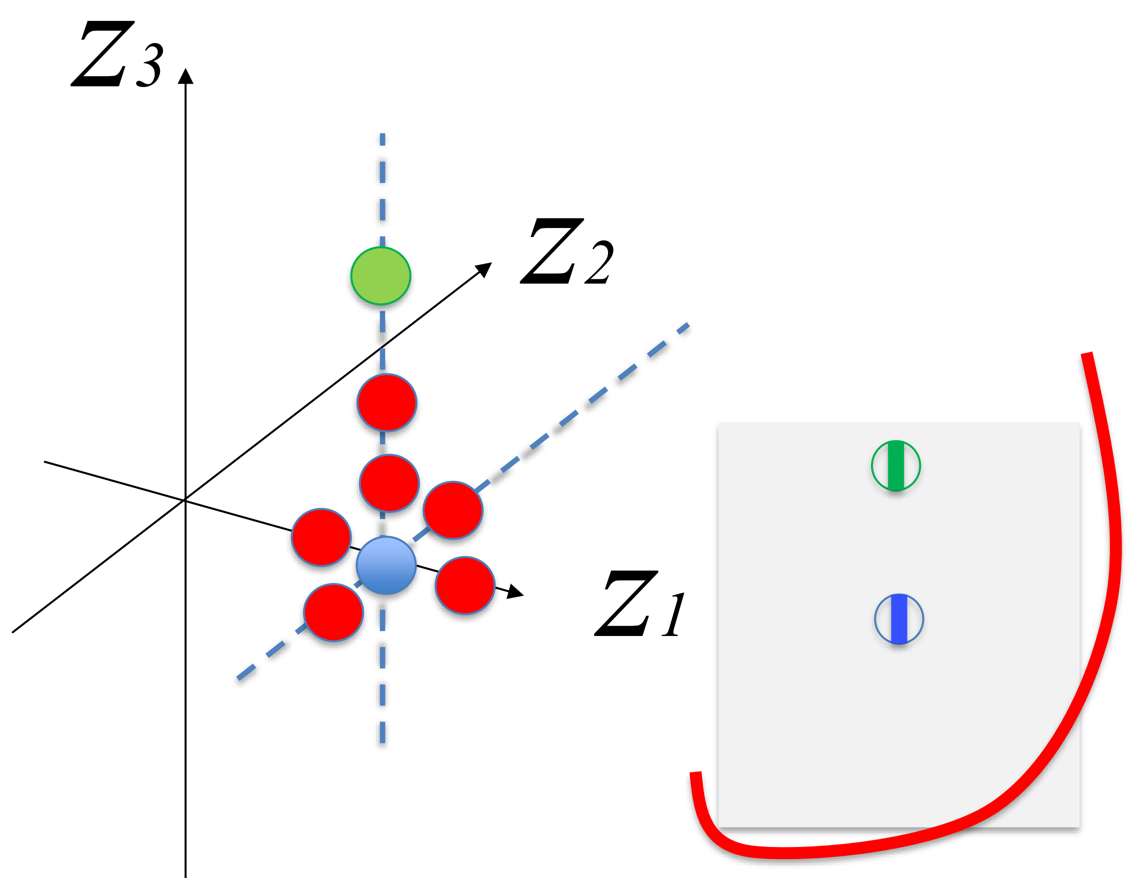

Each abstract axis: a feature of receptive field

functional similarity = neighborhood in abstract space

Three assumptions

Assumption 2: Similar neurons have similar connectivity: functionally similar neurons are strongly connected

functional characterization of neuron
- orientation of rec. field: z_1
- horizontal placement of rec. field: z_2
- vertical placement of rec. field: z_3

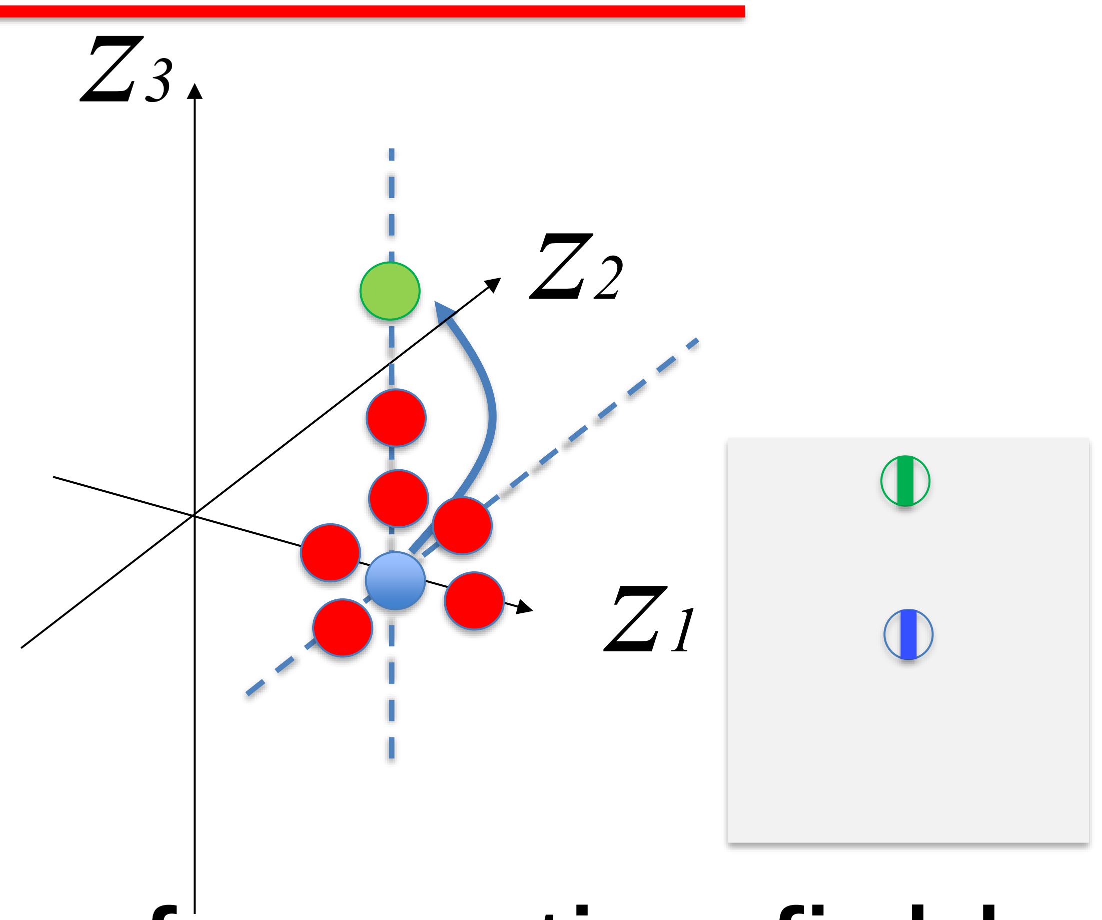

Each abstract axis: a feature of receptive field

Example of 'patchy connectivity': neurons with similar orientation are strongly connected (even if far distance on cortical surface)

Assumption 3: Connectivity is of 'low rank' (outer product)

$$W_{ij} = \sum_{\mu}^{D} F_i^{\mu} \, G_j^{\mu}$$

for example: **D=1**
→ all columns of matrix W_{ij} are linearly dependent
→ **rank 1 (not 2!)**

Example of low-rank: connectivity in Hopfield model

$$W_{ij} = \sum_{\mu}^{D} p_i^{\mu} \, p_j^{\mu}$$

with p_i^{μ}=+/-1 the target value of neuron i in pattern μ

Functional similarities and 'wiring'

functional similarity = neighborhood in abstract space

Assumption 1:

Position of neuron i in abstract space: $\boldsymbol{z_i} = (z_1, z_2, z_3, \dots)(\boldsymbol{i})$

Assumption 2:

Weight of connection from j to i depends on the positions $\boldsymbol{z_i}$, $\boldsymbol{z_j}$:

$$W_{ij} = w(\boldsymbol{z_i}, \boldsymbol{z_j})$$

Assumption 3:

Specific choice of weight from j to i:

$$W_{ij} = \sum_{\mu}^{D} F_i^{\mu} G_j^{\mu} = \sum_{\mu}^{D} f_{\mu}(\boldsymbol{z_i}) g_{\mu}(\boldsymbol{z_j})$$

$$f_{\mu}(\boldsymbol{z_i}) = F_i^{\mu}$$

$$g_{\mu}(\boldsymbol{z_j}) = G_j^{\mu}$$

Field equation in functional similarity space

$$\frac{d}{dt} h_i(t) = -\frac{1}{\tau} h_i(t) + \sum_j W_{ij}\, \phi(h_j(t))$$

with neuron i at position $\mathbf{z_i}$

use weights:

$$W_{ij} = \sum_{\mu}^{D} f_\mu(\mathbf{z_i}) g_\mu(\mathbf{z_j})$$

$$\frac{d}{dt} h(\mathbf{z_i}, t) = -\frac{1}{\tau} h(\mathbf{z_i}, t) + \sum_j \sum_{\mu}^{D} f_\mu(\mathbf{z_i}) g_\mu(\mathbf{z_j})\, \phi(h(\mathbf{z_j}, t))$$

$$\frac{d}{dt} h(\mathbf{z}, t) = -\frac{1}{\tau} h(\mathbf{z}, t) + \int d\mathbf{z}' \rho(\mathbf{z}') \sum_{\mu}^{D} f_\mu(\mathbf{z}) g_\mu(\mathbf{z}')\, \phi(h(\mathbf{z}', t))$$

generalized field equation (large number of neurons)

Field equation and low-dimensional dynamics

$$\frac{d}{dt}h(\mathbf{z},t) = -\frac{1}{\tau}h(\mathbf{z},t) + \underbrace{\int d\mathbf{z}'\rho(\mathbf{z}')\; \sum_{\mu}^{D} f_\mu(\mathbf{z})\, g_\mu(\mathbf{z}')\,\phi(h(\mathbf{z}',t))}$$

$\alpha_\mu(t)$

$$\frac{d}{dt}h(\mathbf{z},t) = -\frac{1}{\tau}h(\mathbf{z},t) + \sum_{\mu}^{D} f_\mu(\mathbf{z})\alpha_\mu(t)$$

coefficient

D ‘basis functions’

Idea: write

$$h(\mathbf{z},t) = \sum_{\mu}^{D} f_\mu(\mathbf{z})\kappa_\mu(t)$$

projection onto basis f.

Field equation and low-dimensional dynamics

(X) $$\frac{d}{dt}h(\mathbf{z},t) = -\frac{1}{\tau}h(\mathbf{z},t) + \sum_{\mu}^{D} f_\mu(\mathbf{z})\alpha^\mu(t) \quad +I(\mathbf{z},t)$$

coefficient

D 'basis functions'

Idea: write

$$h(\mathbf{z},t) = \sum_{\mu}^{D} f_\mu(\mathbf{z})\kappa_\mu(t)$$

projection onto basis f.

Example:

$f_1(z) = sin(z)$

$cos(z)$

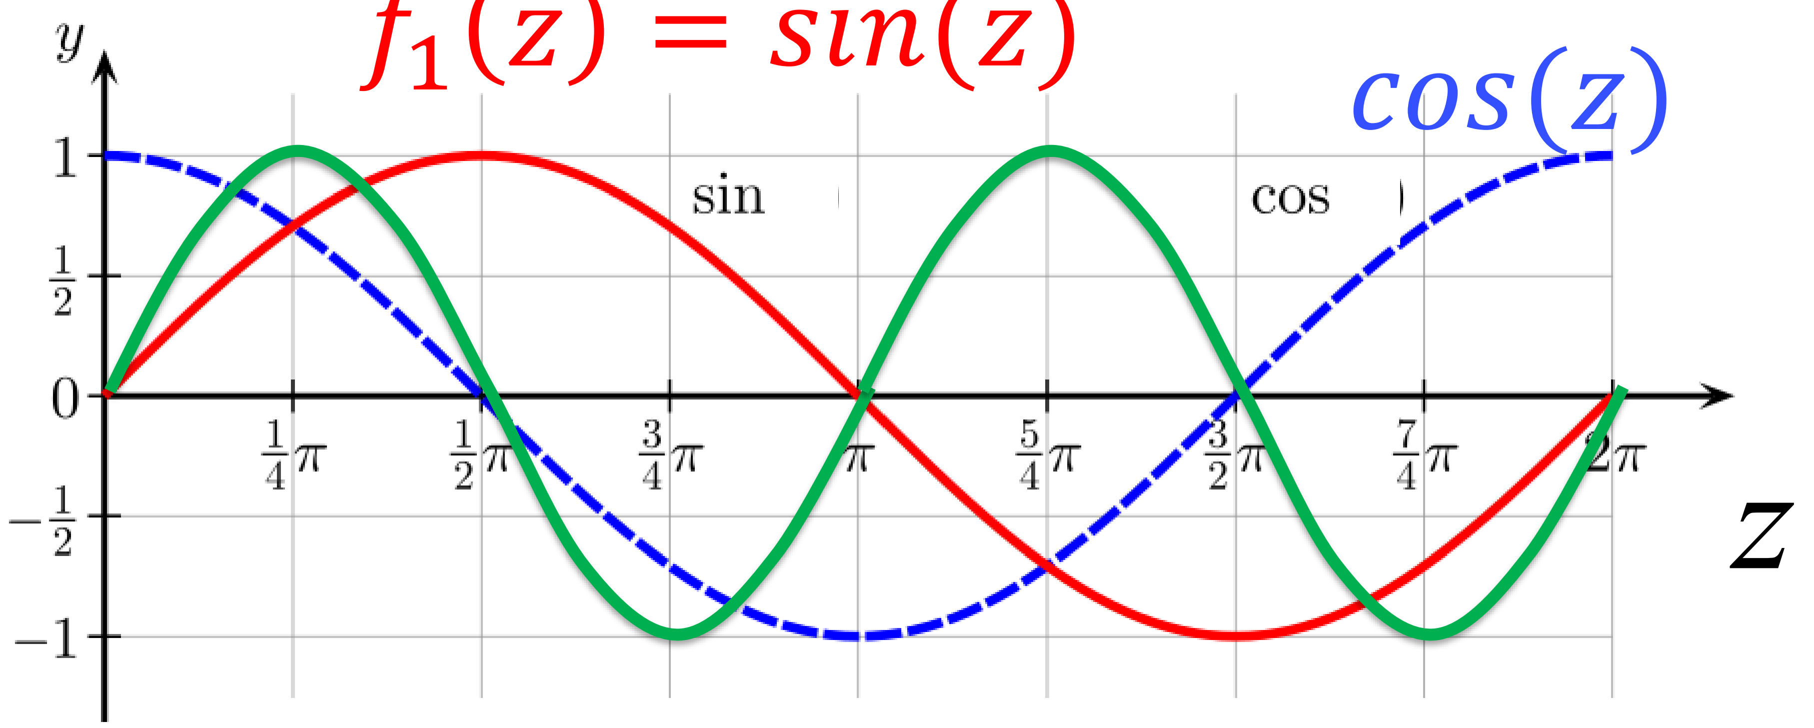

$$I(\mathbf{z},t) = \sum_{\mu}^{D+3} f_\mu(\mathbf{z}) I_\mu(t)$$

external input in same basis

Field equation and low-dimensional dynamics

Idea: write

$$h(\mathbf{z},t) = \sum_{\mu}^{D} f_\mu(\mathbf{z})\kappa_\mu(t)$$

→yields D coupled equations

$$\frac{d}{dt}\kappa_\mu(t) = -\frac{1}{\tau}\kappa_\mu(t) + \int d\mathbf{z}\rho(\mathbf{z})g_\mu(\mathbf{z})\phi(\underbrace{\sum_{\nu}^{D} f_\nu(\mathbf{z})\kappa_\nu(t)}_{\phi(h(\mathbf{z},t))})$$

→ activity of all N neurons (N>>1) is described by D equations in recurrent network (without external input)

Summary: low-dimensional dynamics

-What is relation between functional similarity and manifold?

functional similarity reflected in wiring, wiring causes dynamics

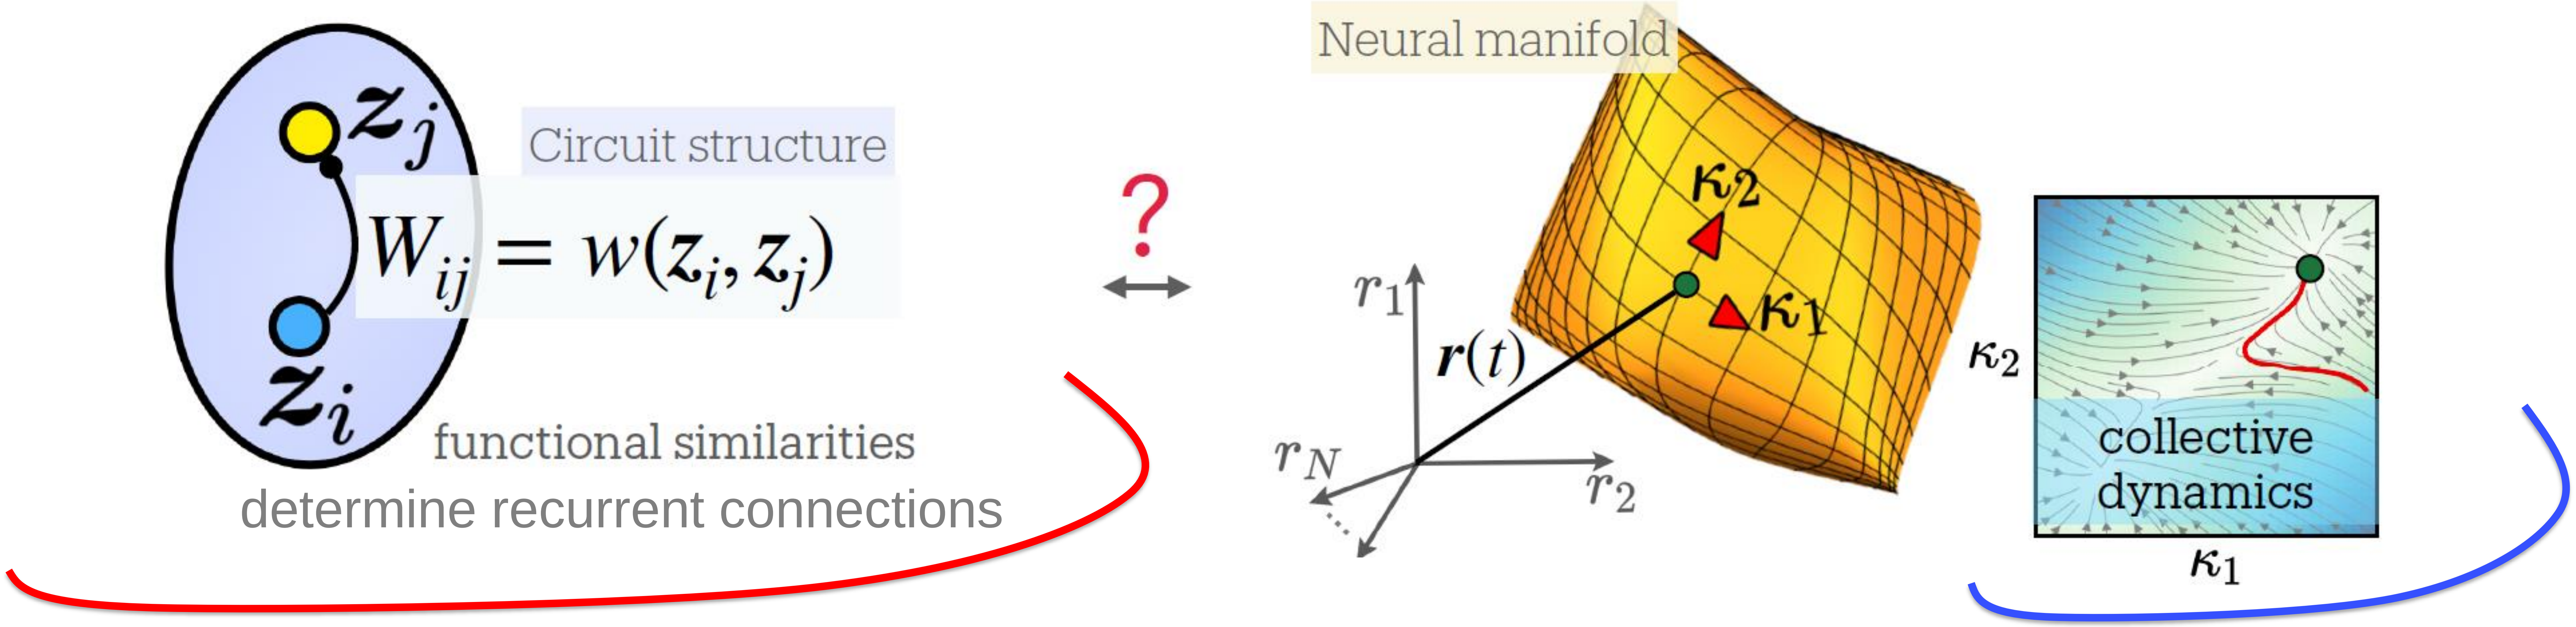

with weights $W_{ij} = \sum_{\mu}^{D} f_\mu(\boldsymbol{z_i}) g_\mu(\boldsymbol{z_j})$, dynamics evolves in D dim.

→ flow described by small number of variables $\kappa_1, \ldots \kappa_D,$

Summary: low-dimensional dynamics

To generate **low-dimensional dynamics** in **heterogeneous** networks of N neurons, three ingredients are important:

(i) neurons characterized by abstract positions z representing functional similarity

(ii) weight matrix depends on z and z'

(iii) weight matrix is of low rank: outer-product with D terms

→ field model for large network (N to infinity)
→ collective dynamics evolves in D dimensions
→ external input can also be included in formalism

References:

Mastrogiuseppe, F., Ostojic, S. (2018), Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. Neuron 99(3), 609–62329

Pezon, L., Schmutz, V, Gerstner, W. (2024), Linking Neural Manifolds to Principles of Circuit Structure in Recurrent Networks bioRxiv doi: https://doi.org/10.1101/2024.02.28.582565

List of video lectures on Computational Neuroscience, organized by topics:

https://lcnwww.epfl.ch/gerstner/NeuronalDynamics-MOOCall.html

YouTube Channel:

https://www.youtube.com/@gerstnerlab

Textbook (online):

https://neuronaldynamics.epfl.ch/

Part B:
The following slides correspond to the video here:

https://youtu.be/eO4F-j0Z6RA

EPFL

From Spiking Neurons to Rate Units: Emergent Rate-based Dynamics in Spiking Neural Networks

Valentin Schmutz, Johanni Brea, Wulfram Gerstner

EPFL, Lausanne, Switzerland

1. The problem of Firing Rates
- textbook introduction

2. Firing rates without duplicates

V. Schmutz, J. Brea, W. Gerstner (2025) **Emergent rate-based dynamics in duplicate-free populations of spiking neurons**
Physical Review Letters, 134:018401
DOI 10.1103/PhysRevLett.134.018401

What is the Firing Rate? 1. spike count (temporal average)

spikes in response to stimulus

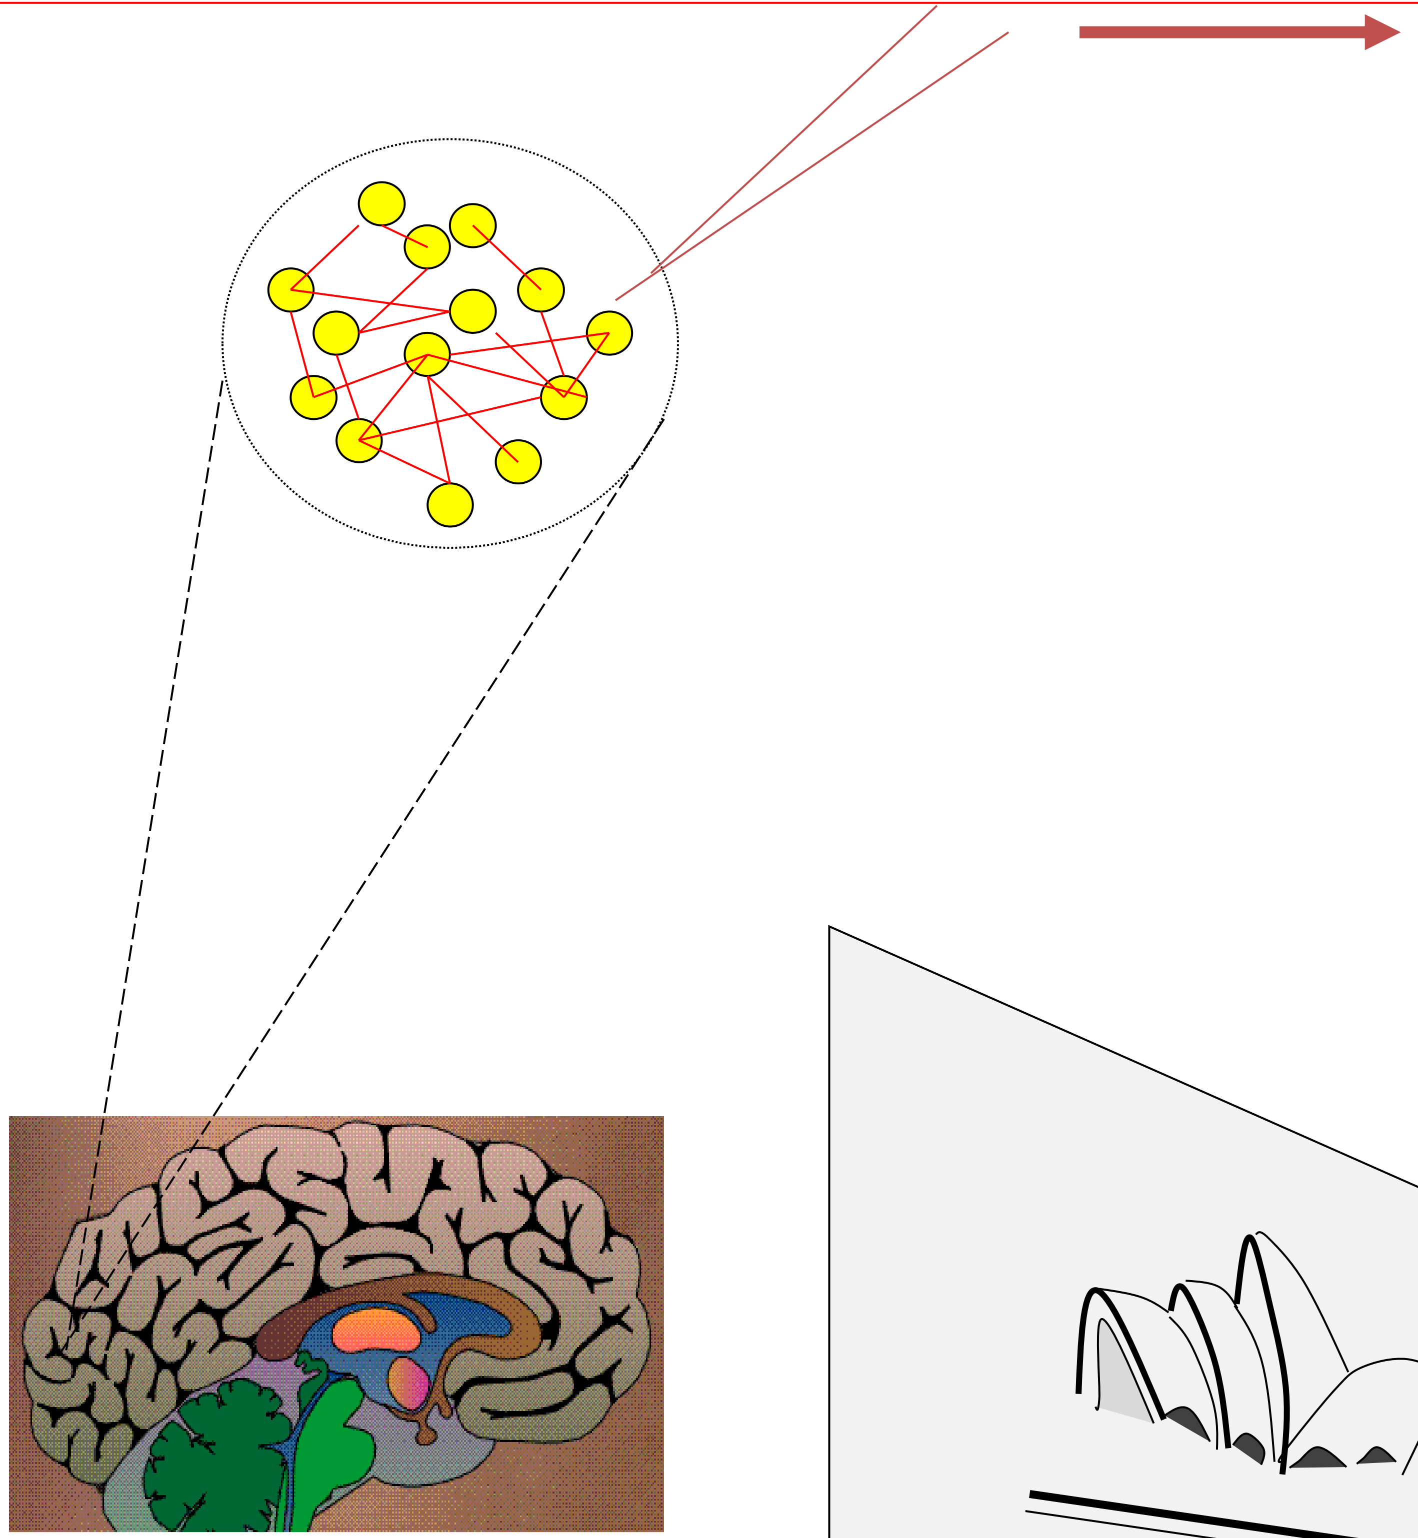

brain

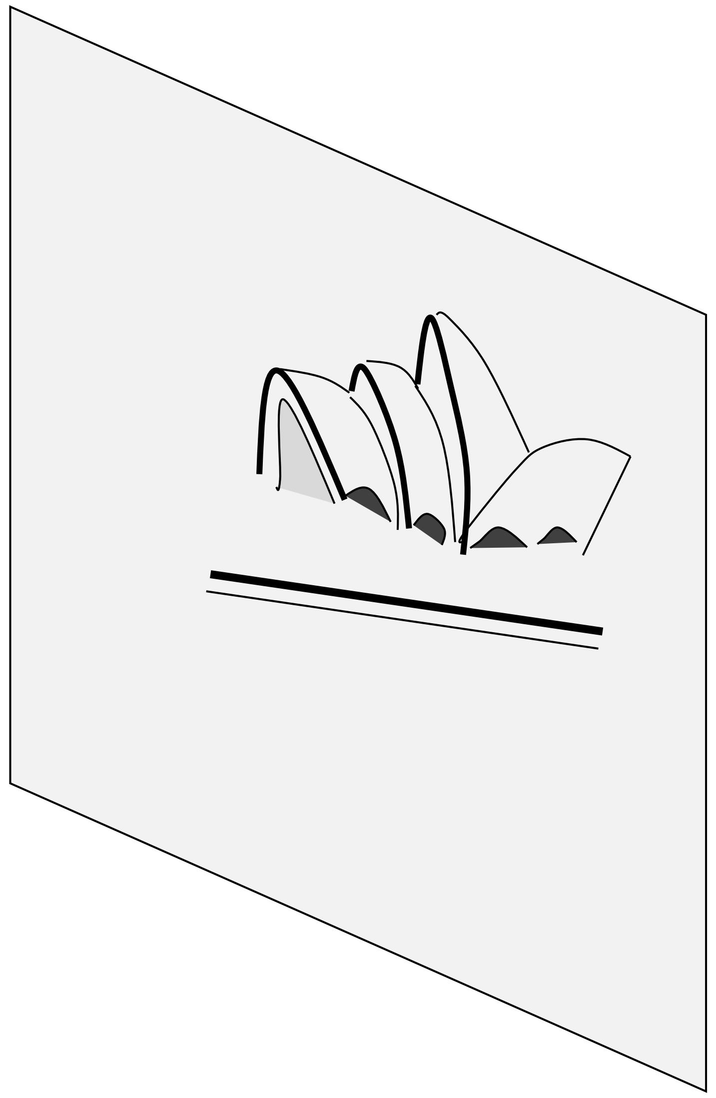

trial 1

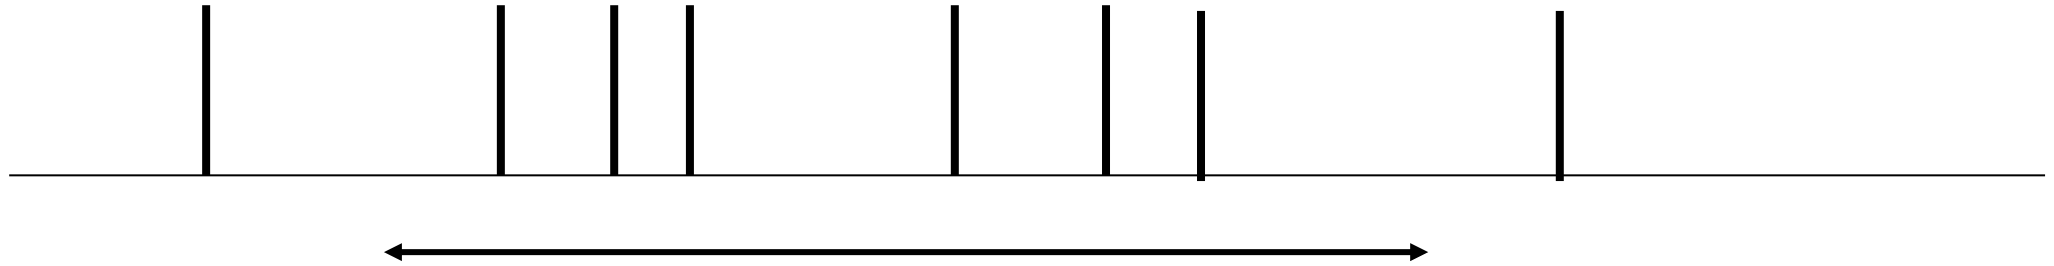

rate as a (normalized) spike count:

$$\nu(t) = \frac{n^{sp}}{T}$$

single neuron/single trial: temporal average

T=1s

stim

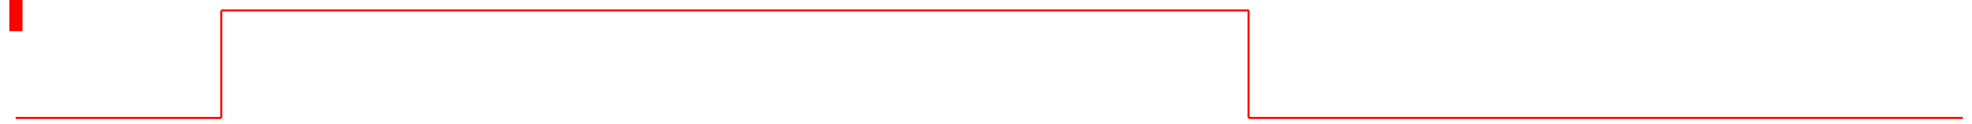

What is the firing rate? 2. population activity (spatial average)

population of neurons
with similar properties

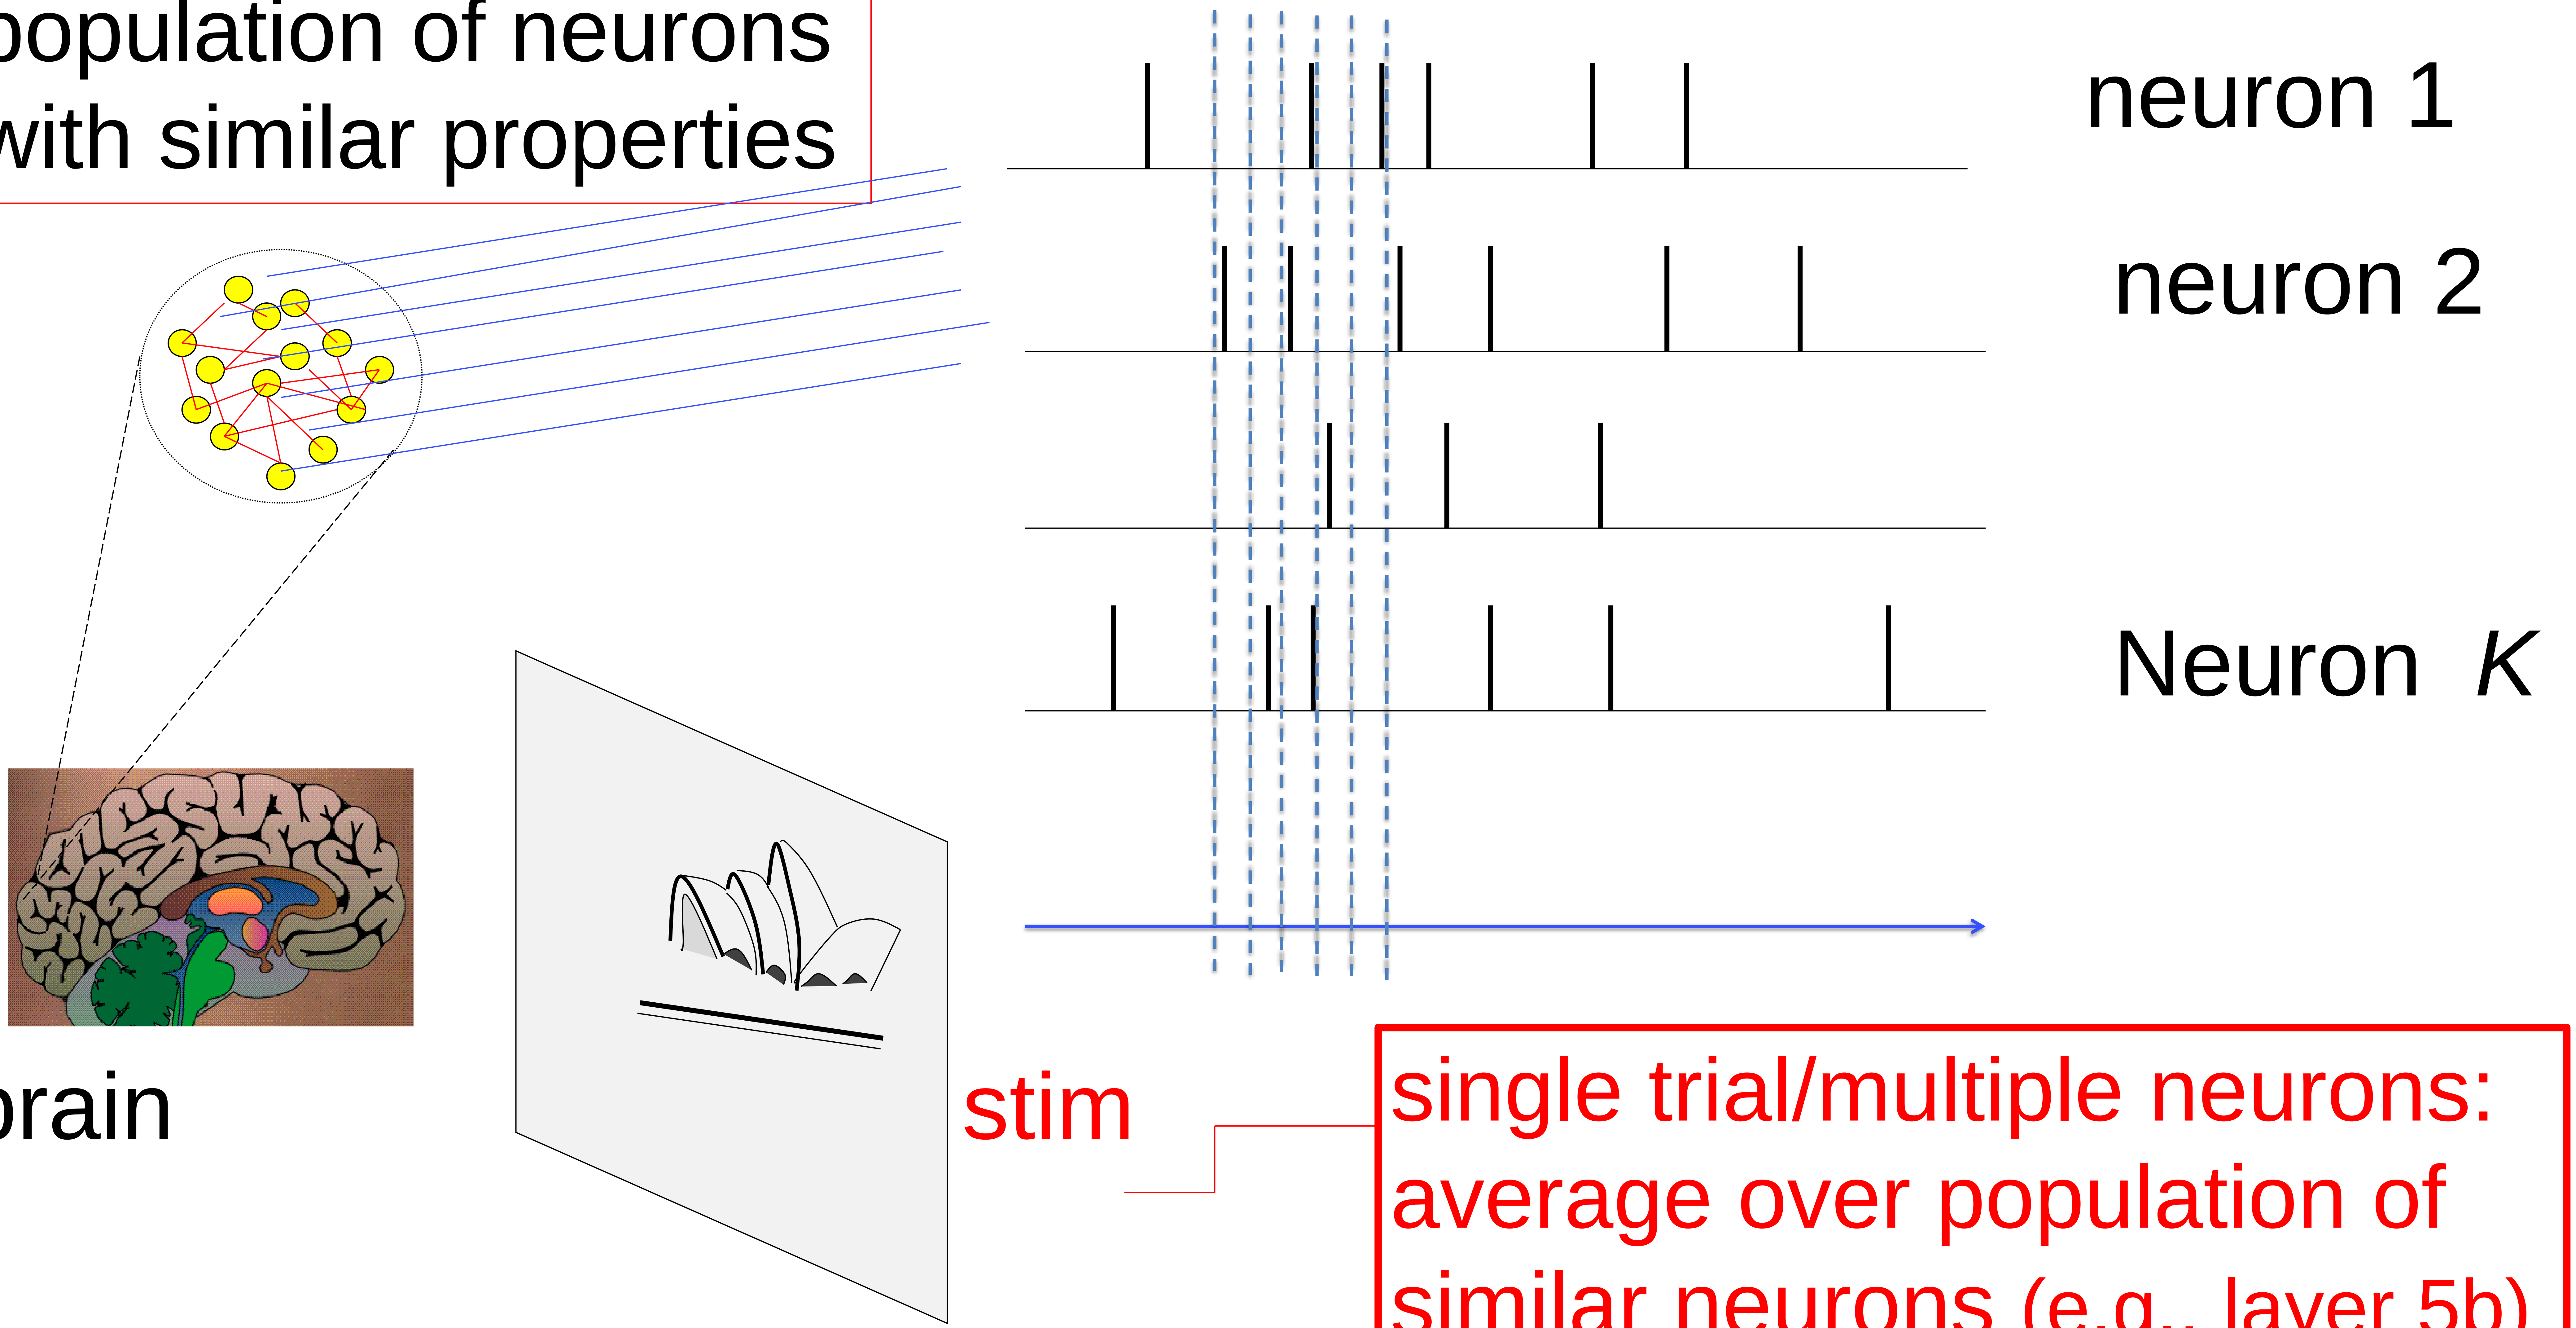

single trial/multiple neurons:
average over population of
similar neurons (e.g., layer 5b)

Rate model: interacting populations (duplicate neurons)

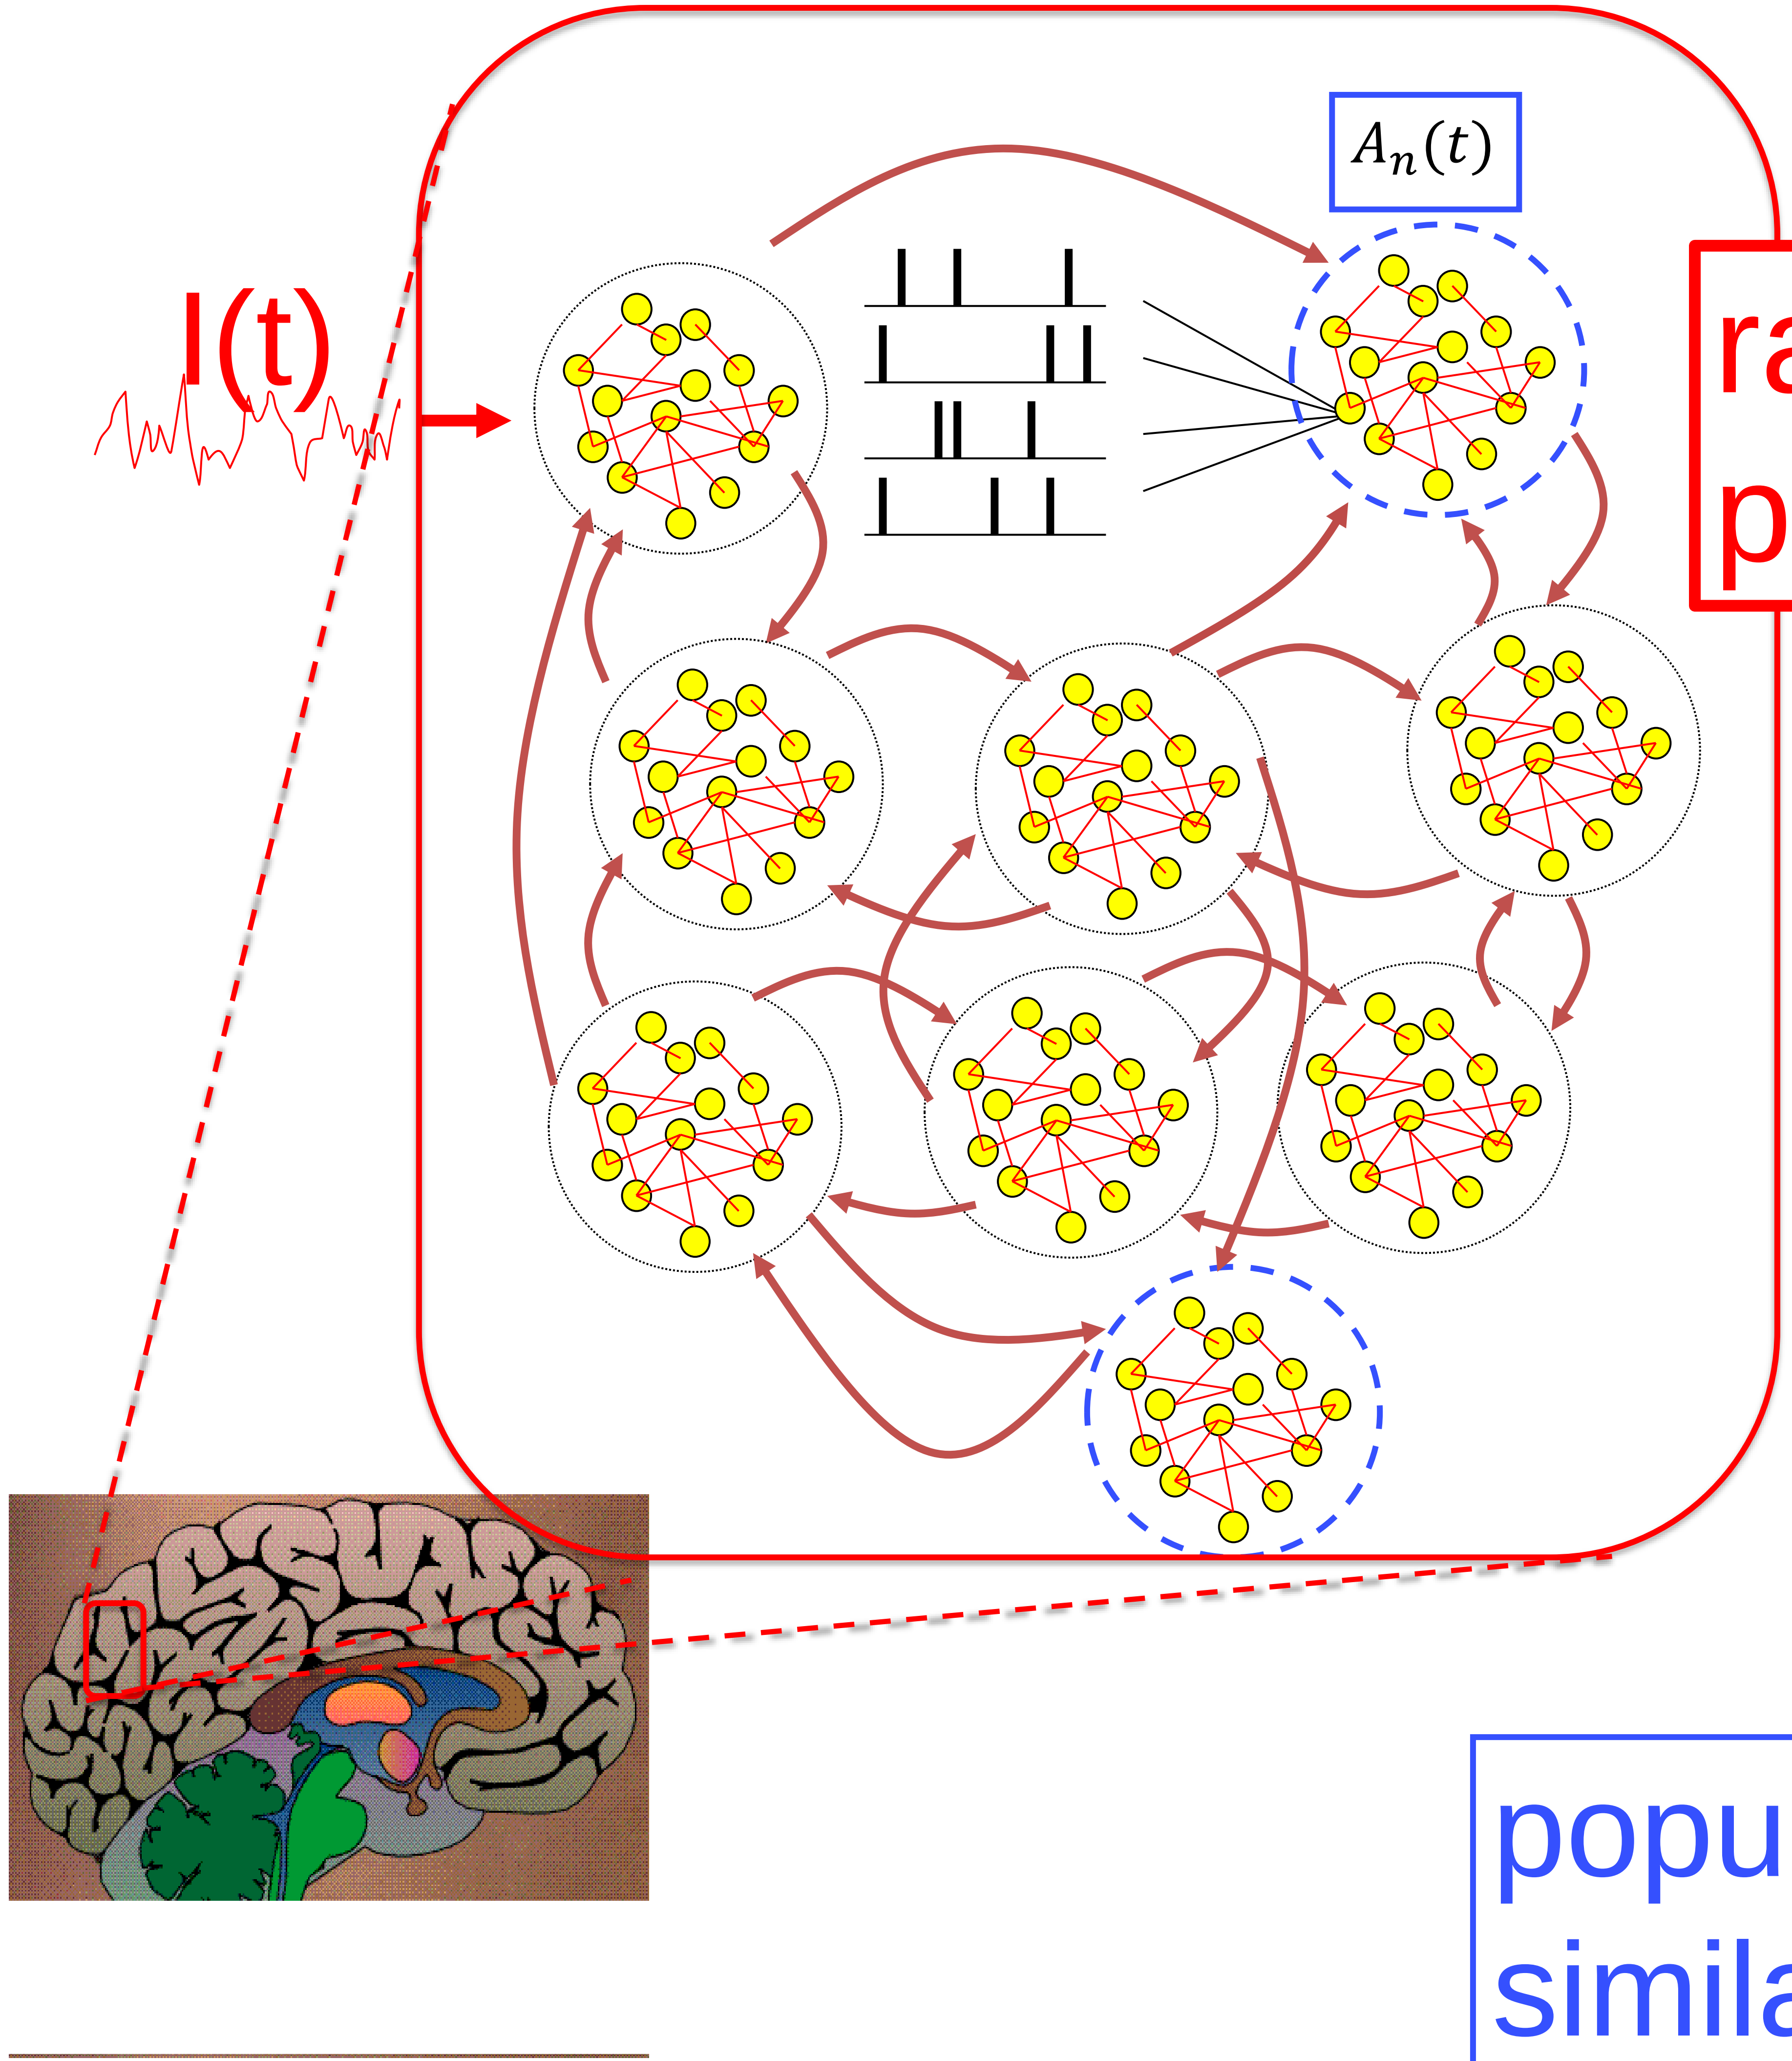

rate network:
populations talk to each other

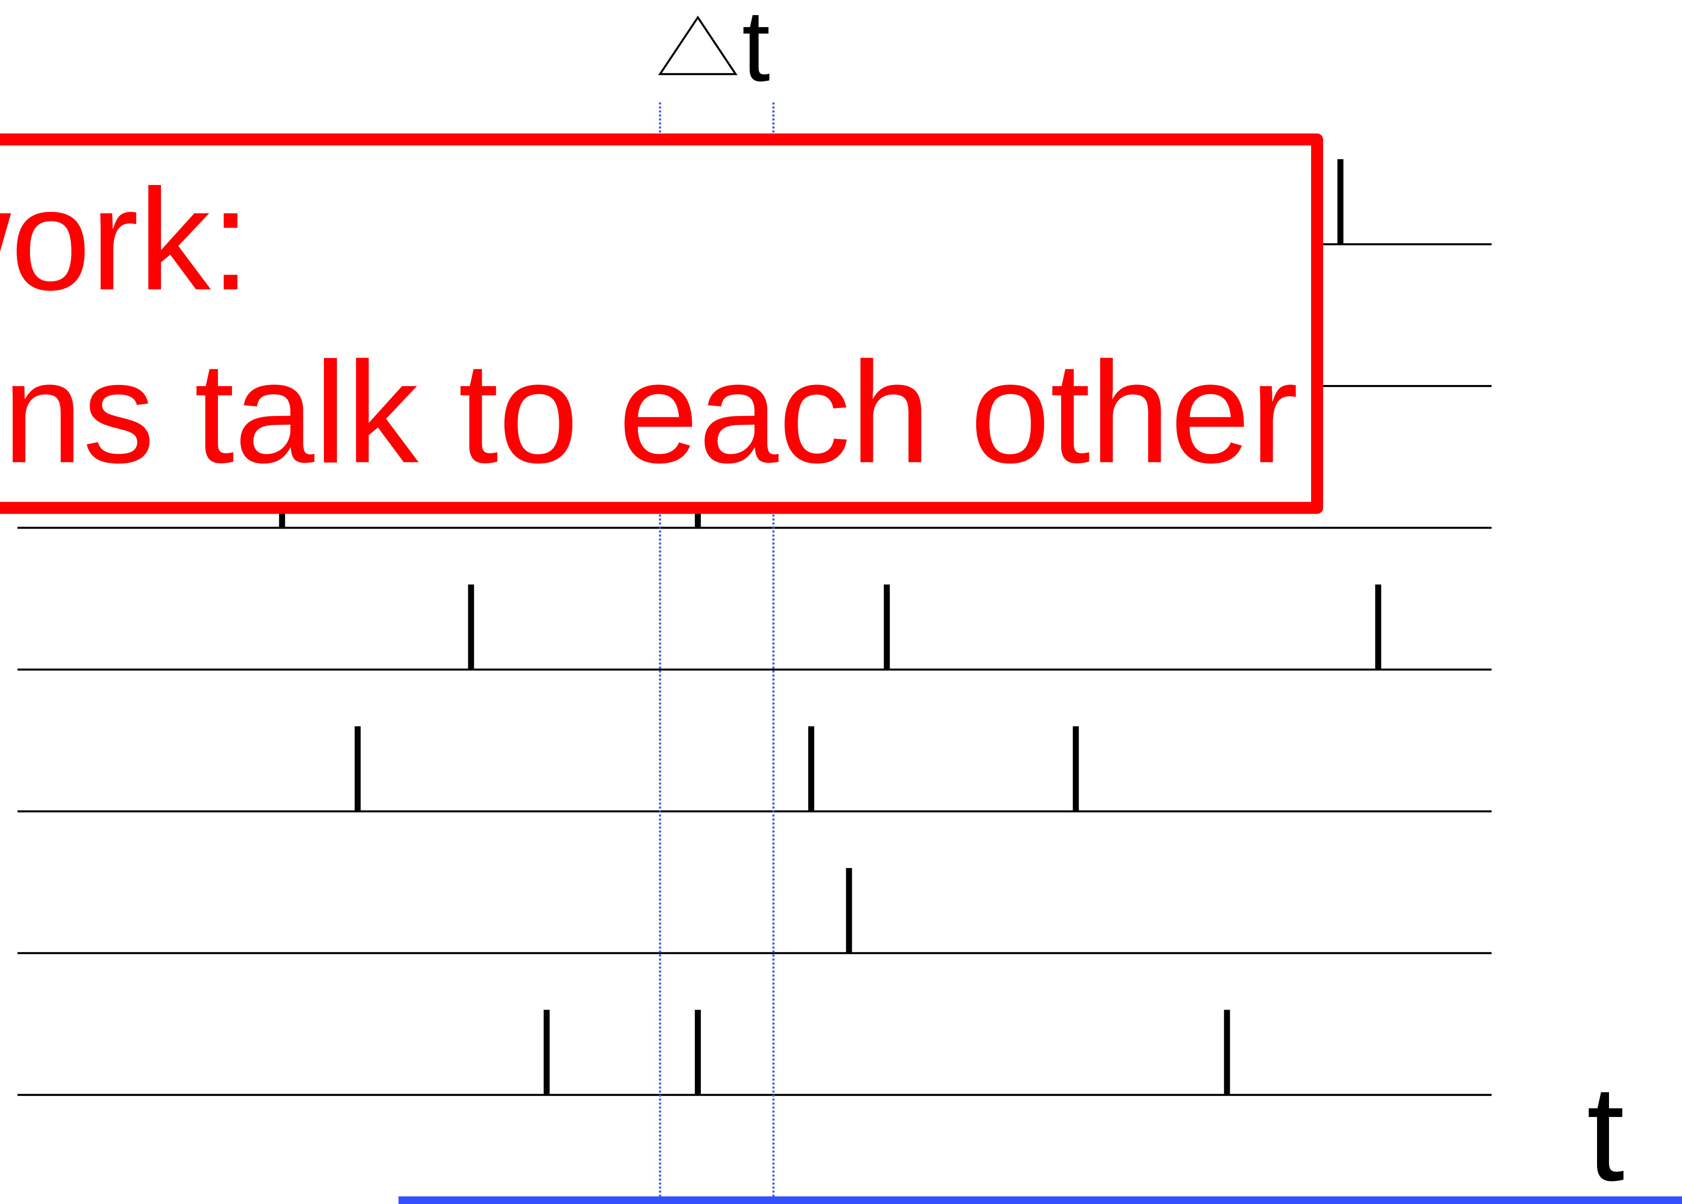

population activity (rate)

$$A_n(t) = \frac{n(t;t+\Delta t)}{N\Delta t}$$

population of similar neurons

rate unit n

Rate codes: population activity

population activity - rate defined by population average

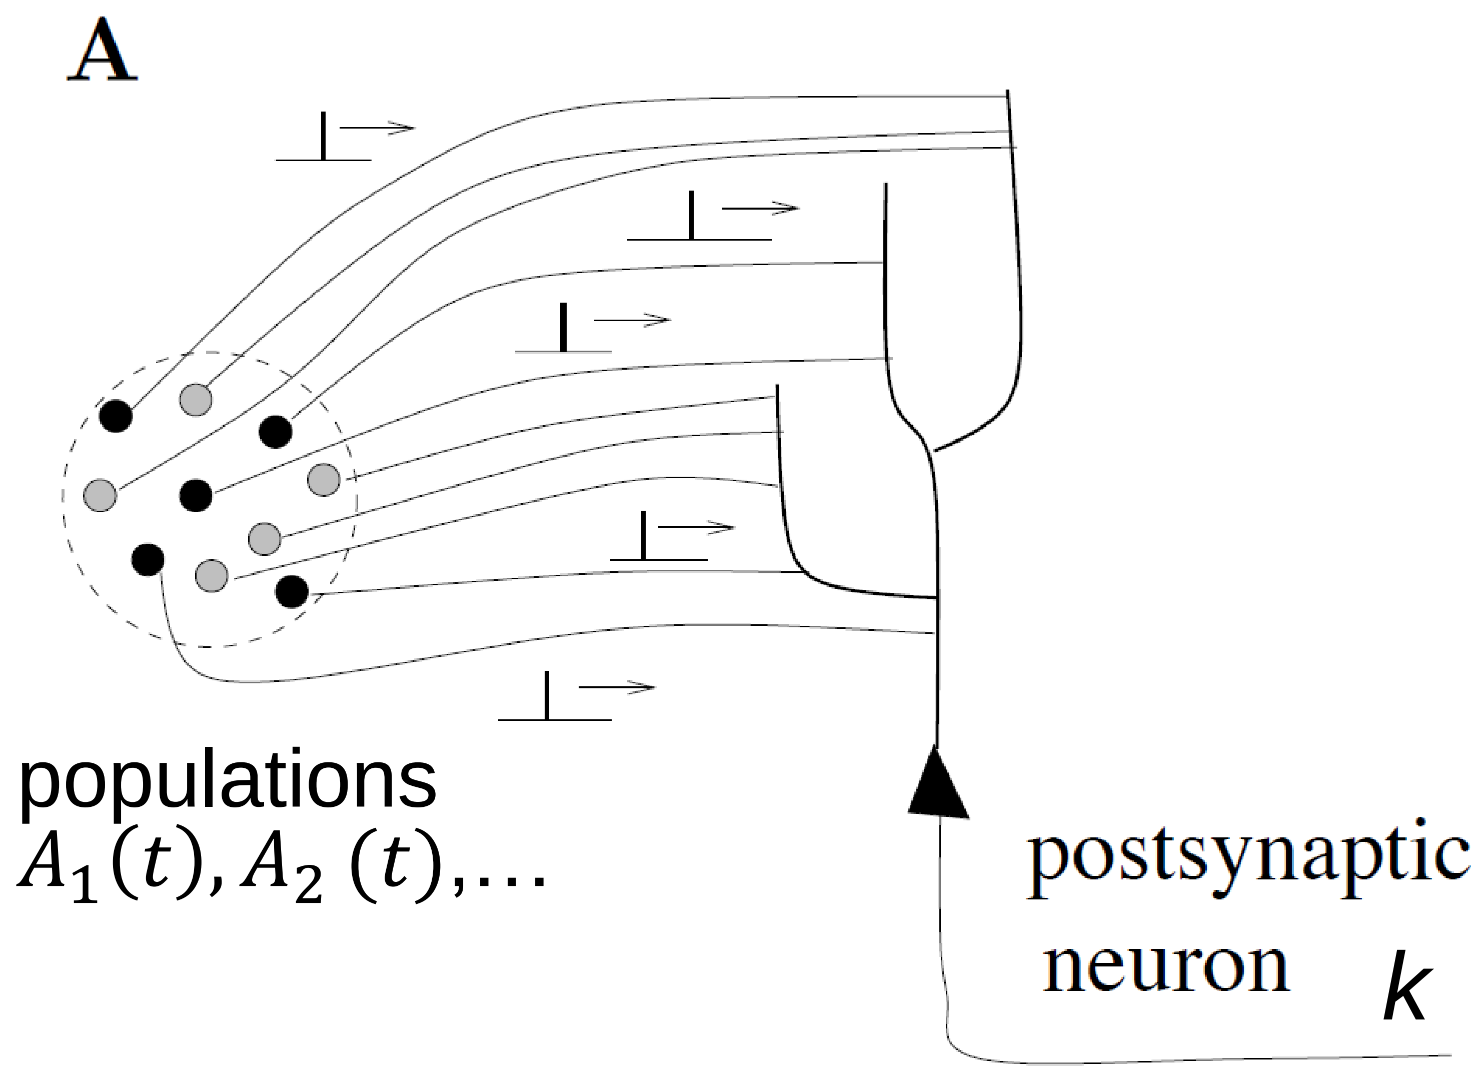

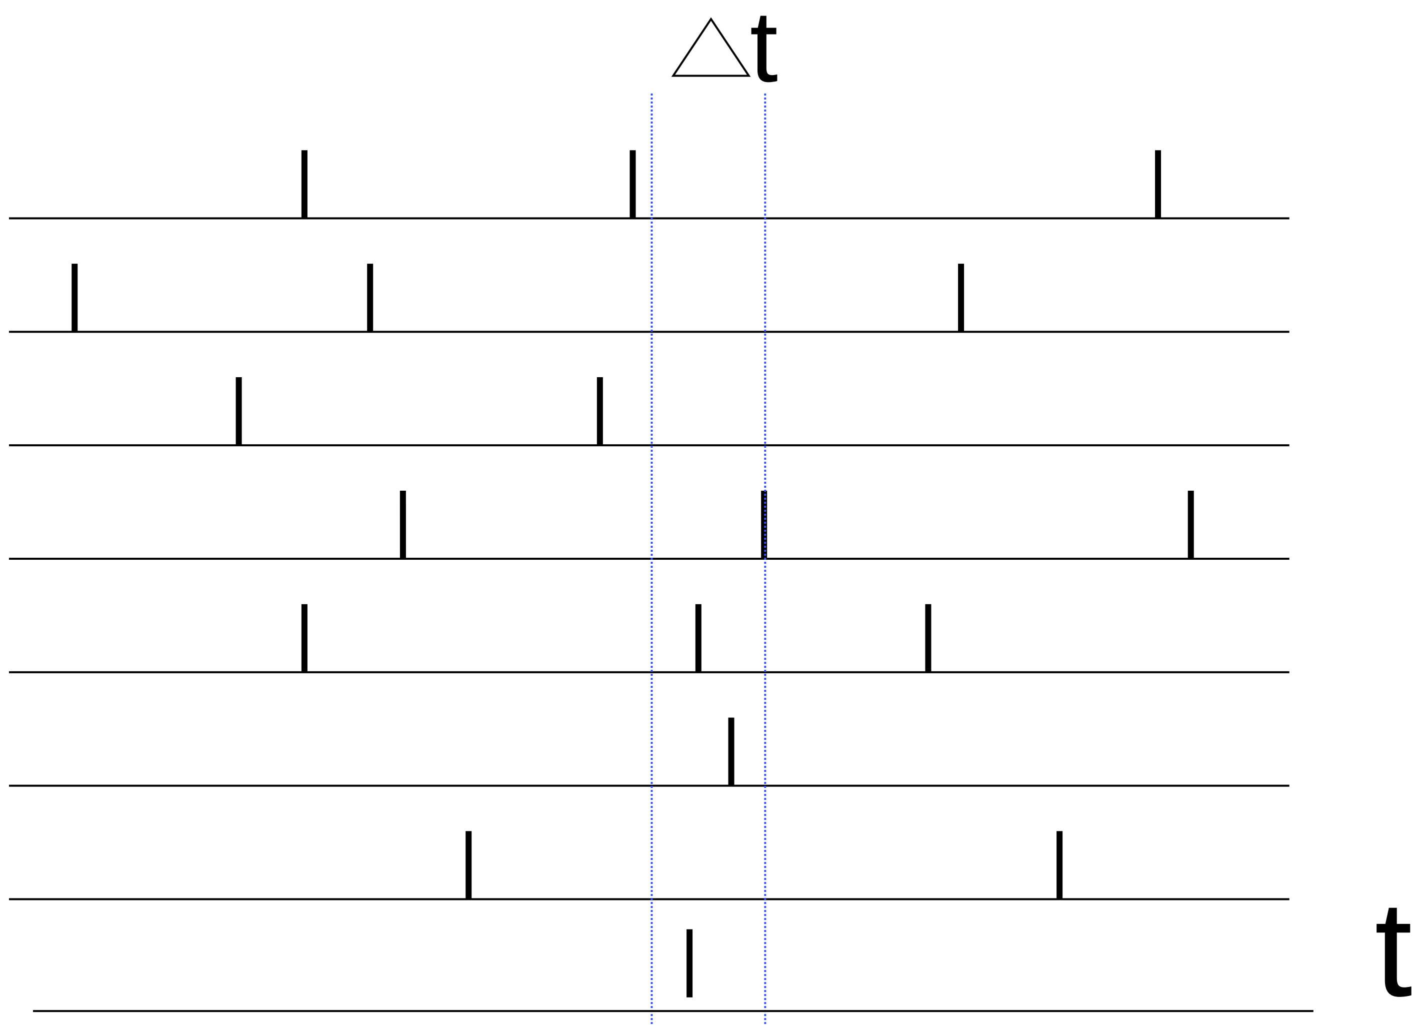

'natural readout'

population activity (rate)

$$A_n(t) = \frac{n(t; t + \Delta t)}{N \Delta t}$$

Textbooks: e.g.

but are the presynaptic pools really **homogenous populations of duplicate neurons/similar neurons**?
→ weighted average over very heterogeneous group!

Definitions of Rate codes: summary

Two averaging methods

- single trial, average over time

single neuron

too slow
for animal!!!

- single trial, average over population

many neurons

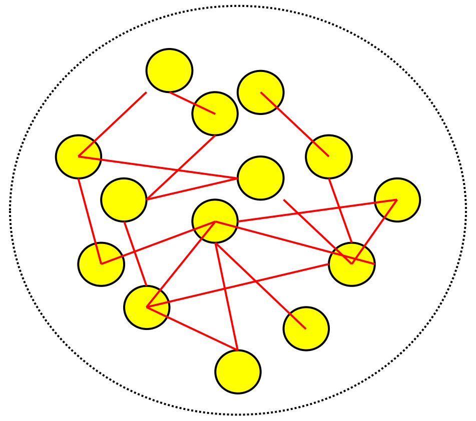

'natural', but do we have enough **duplicate/similar neurons**?

Textbooks: e.g. - *Neuronal Dynamics, Gerstner et al., (Cambridge Univ. Press, 2014)*
- *Theoretical Neuroscience, Dayan and Abbott (MIT Press, 2001)*

Big question:

Is a rate description meaningful in spiking neurons, if

- temporal averaging is impossible because **signals are fast**

- there are **no duplicate neurons (no similar neurons)**

?

[ ] intuitively plausible
[ ] intuitively not plausible
[ ] may be, but if yes, then under
very **strict** conditions

From Spiking Neurons to Rate Units
Emergent Rate-based Dynamics in Spiking Neural Networks:

Valentin Schmutz, Johanni Brea,
Wulfram Gerstner

EPFL, Lausanne, Switzerland

1. **The problem of Firing Rates**
   - textbook introduction

2. **Firing rates without duplicates**

V. Schmutz, J. Brea, W. Gerstner (2025) **Emergent rate-based dynamics in duplicate-free populations of spiking neurons**
Physical Review Letters, 134:018401
[DOI 10.1103/PhysRevLett.134.018401](https://doi.org/10.1103/PhysRevLett.134.018401)

Review: functional similarity of neurons

**functional similarity =
neighborhood in abstract space**

Example: head direction system

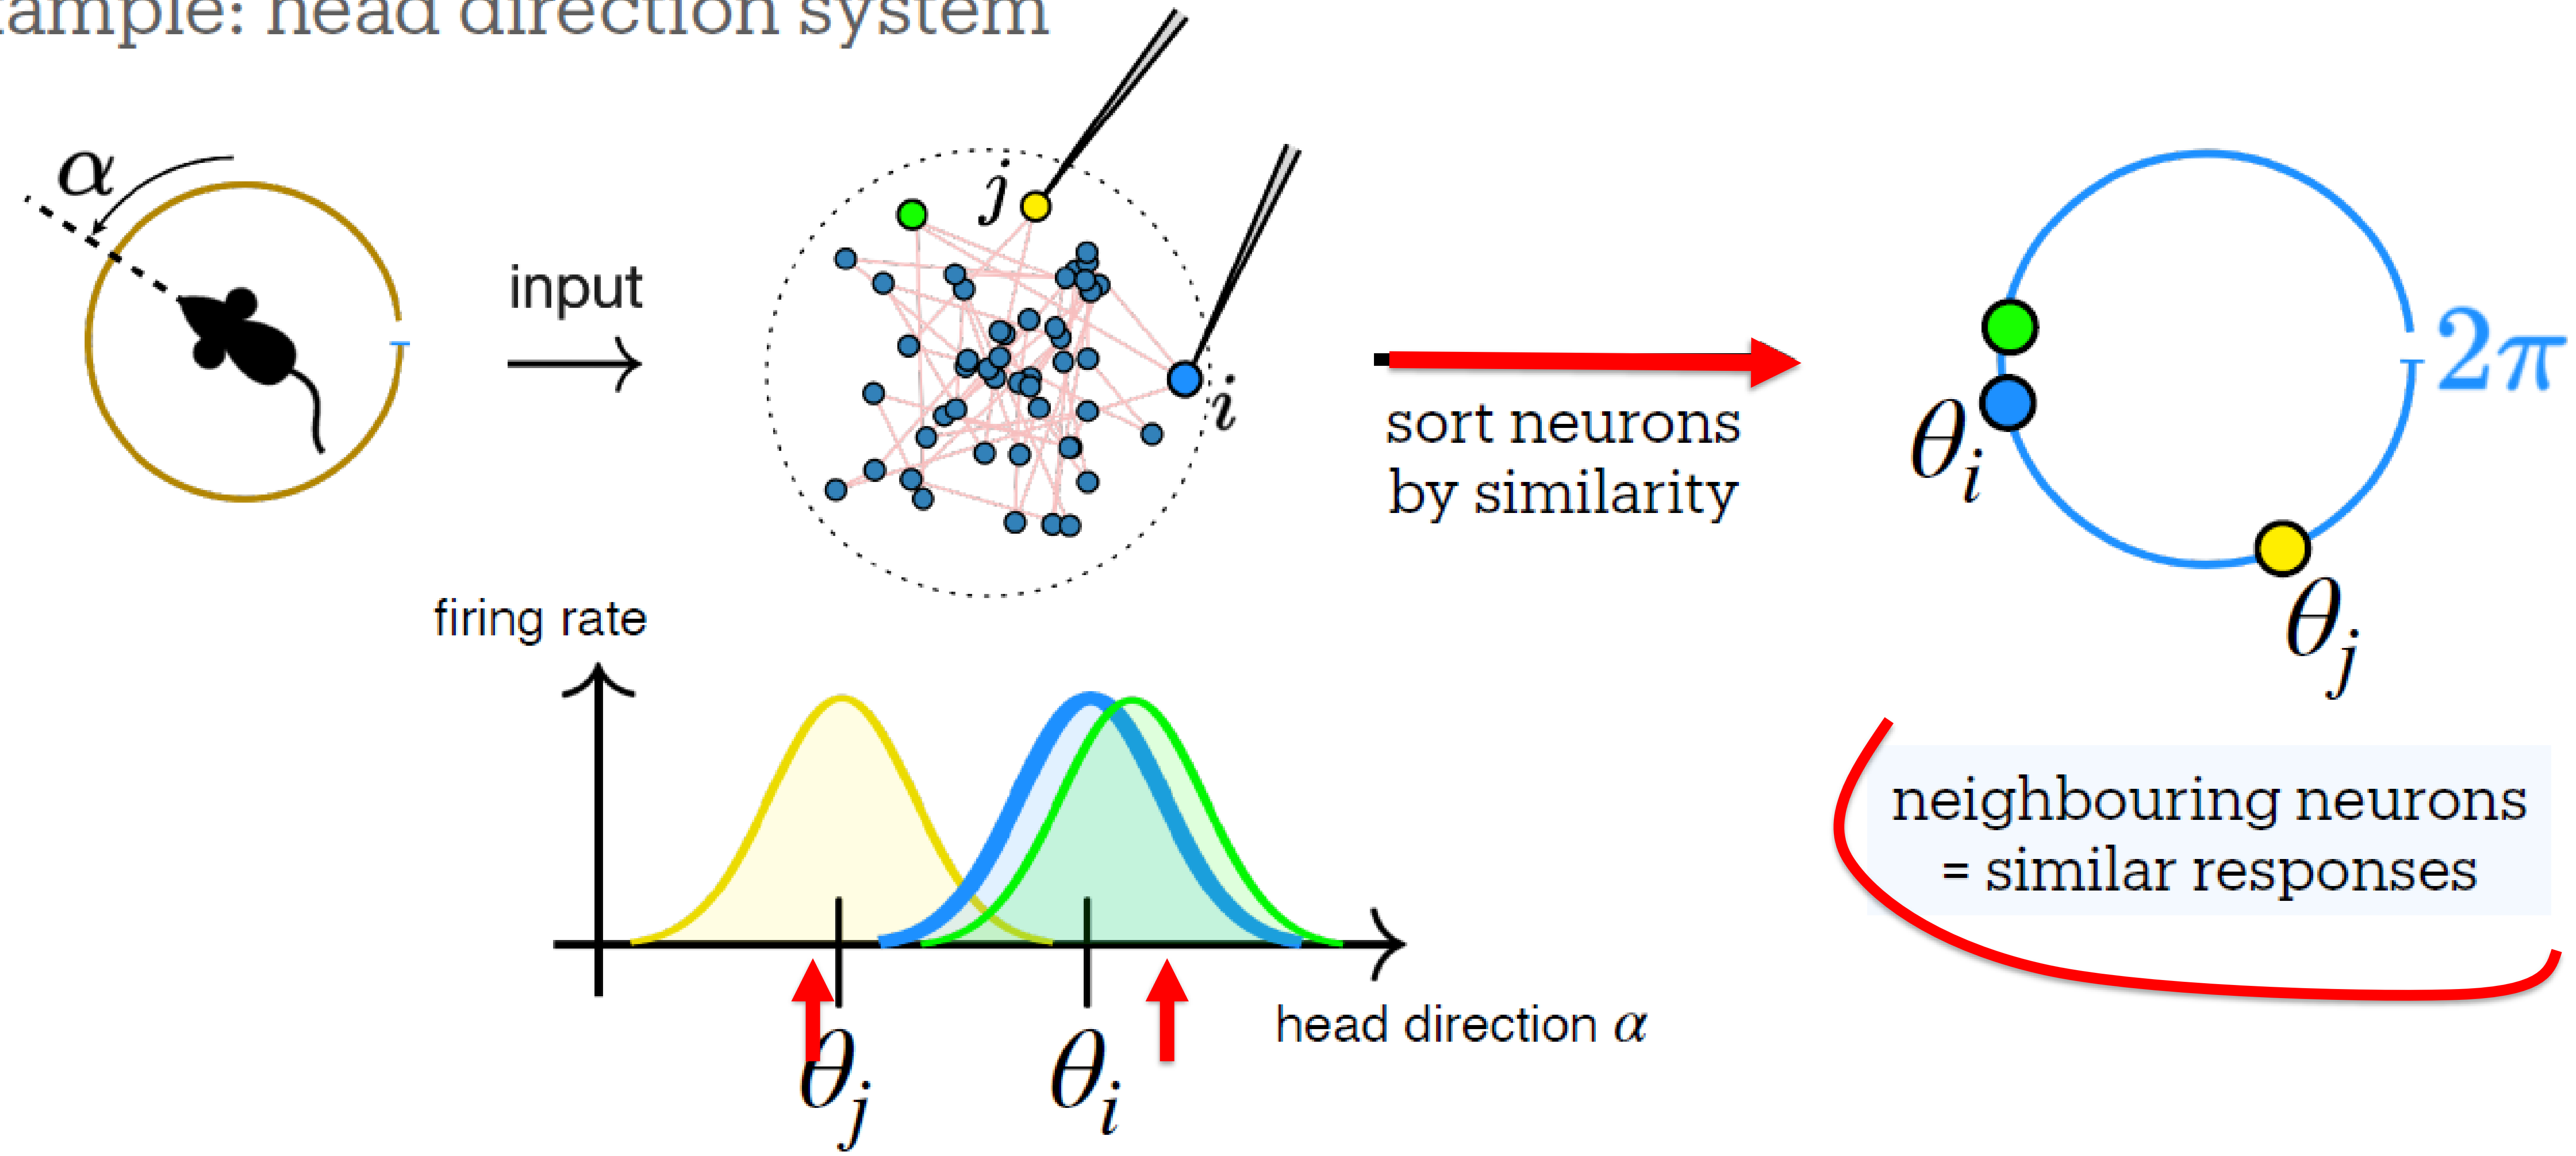

variable z:
position on ring

'ring model'

Functionally similar neurons do not always sit next to each other in cortex.

Correlations between two neurons ($N = 10^6$ neurons total)

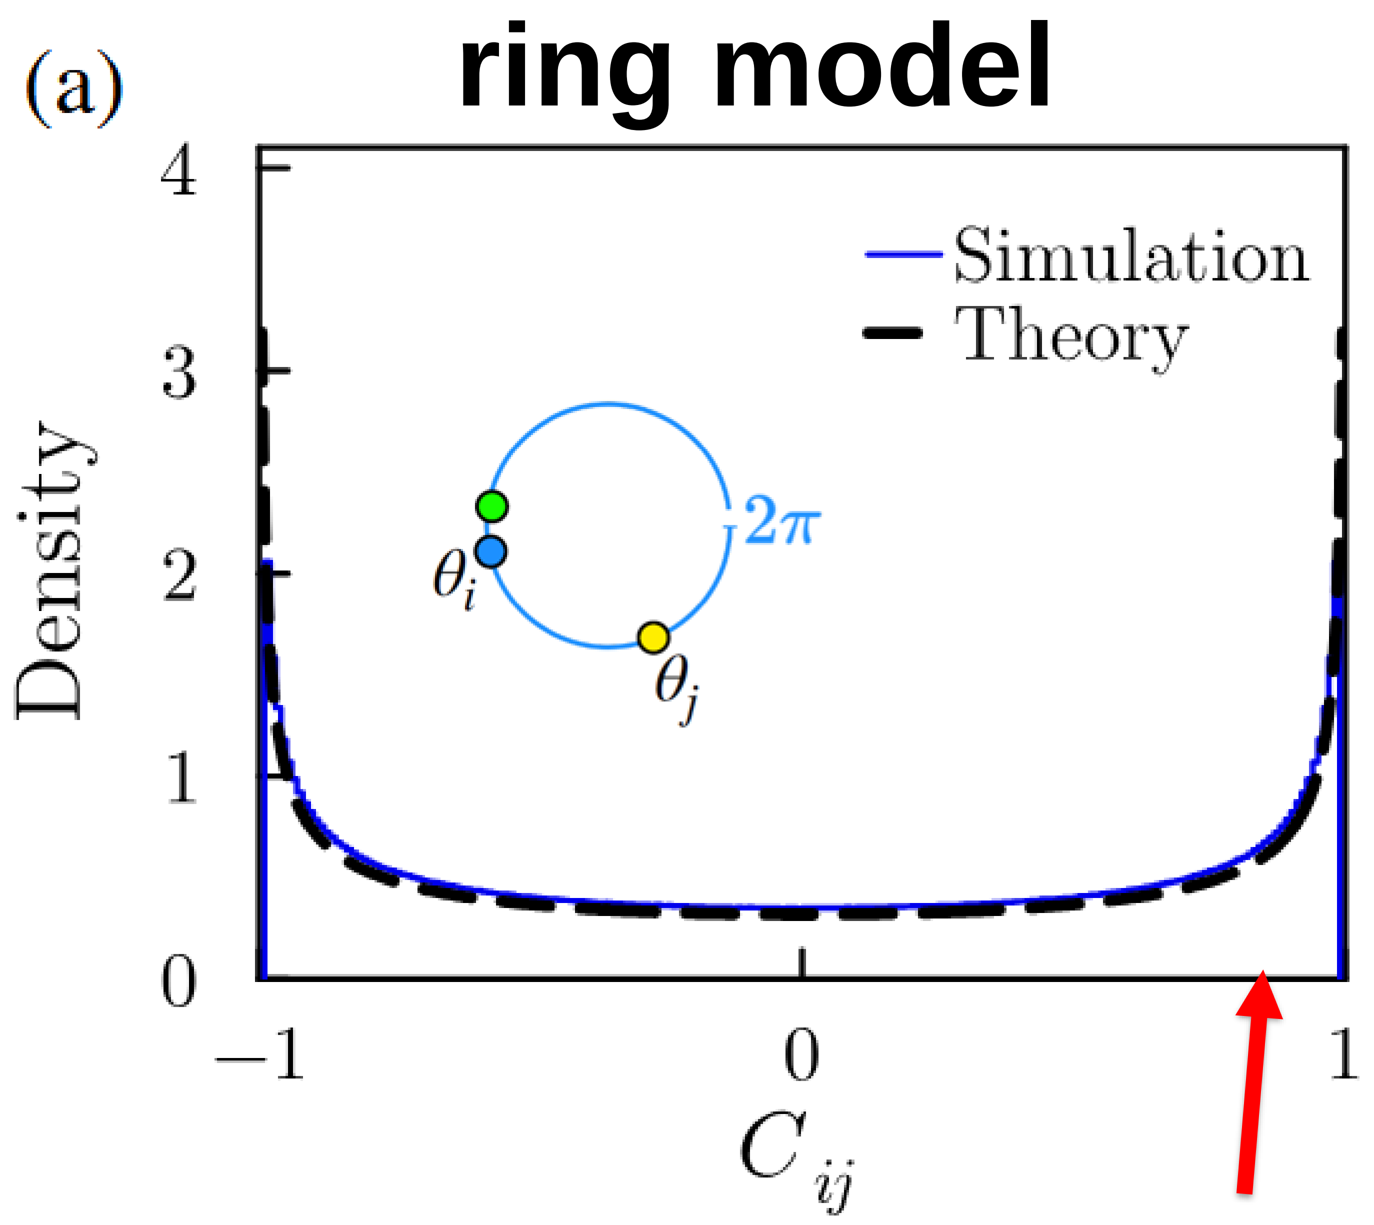

-ring model with Poisson neurons
-stimulation at random location

Horizontal axis:
amount of correlation between 2 neurons

high correlation,
caused by pairs of neighboring neurons

Ring model: many pairs of neurons are strongly correlated.

→ **“duplicate neurons”: identical or strongly correlated**

→ duplicate neurons respond ‘nearly the same’

Correlations between two neurons: low-rank weight matrix

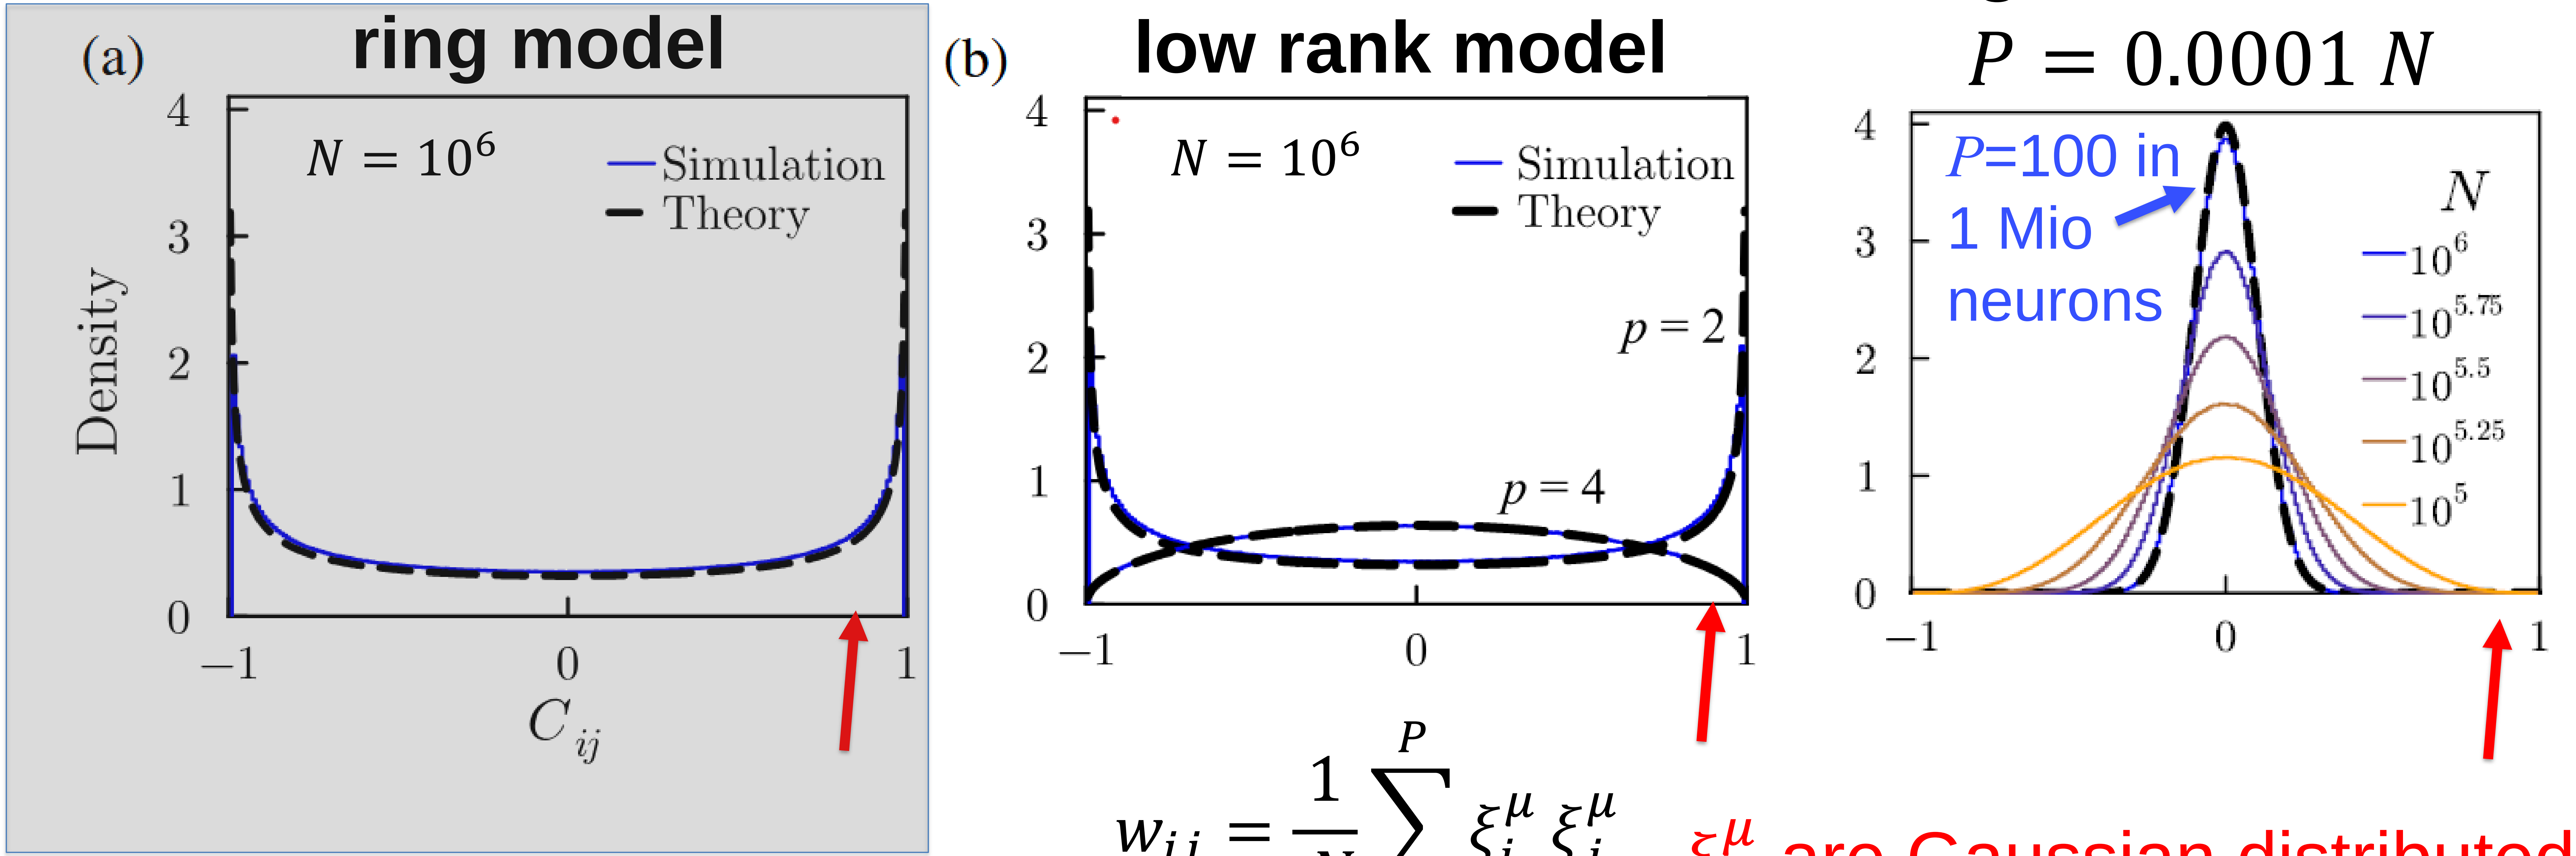

$$w_{ij} = \frac{1}{N} \sum_{\mu}^{P} \xi_i^\mu \, \xi_j^\mu$$

ξ_j^μ are Gaussian distributed

Neurons become uncorrelated for $P \to \infty;\ N \to \infty;\ \frac{P}{N} \to 0$

$e.g.\ P = N^{1/3}$

→ **no duplicate neurons**

V. Schmutz, J. Brea, W. Gerstner (2025) ***Emergent rate-based dynamics in duplicate-free populations of spiking neurons***
Physical Review Letters, 134:018401

1st important finding:

For low-rank weights

$$w_{ij} = \frac{1}{N} \sum_{\mu}^{P} \xi_i^\mu \, \xi_j^\mu$$

Proof:
Concentration of Measure

ξ_j^μ are Gaussian distributed

neurons become uncorrelated for $P \to \infty;\ N \to \infty;\ \frac{P}{N} \to 0$

→ **no duplicate neurons**

$e.g.\ P = N^{1/3}$

V. Schmutz, J. Brea, W. Gerstner (2025) ***Emergent rate-based dynamics in duplicate-free populations of spiking neurons*** *Physical Review Letters, 134:018401*

With low-rank weights and $P = N^{1/3}$,
we can exclude duplicate neurons: **is rate coding possible**?

Recurrent Neural Network (RNN)

Recurrent network of N neurons.

Membrane potential of neuron i:

$$\frac{d}{dt}x_i(t) = -\frac{1}{\tau}x_i(t) + \sum_j w_{ij}\,\phi\left(x_j(t)\right) + I_i^{ext}(t)$$

"rate model"

firing rate (rate variable):

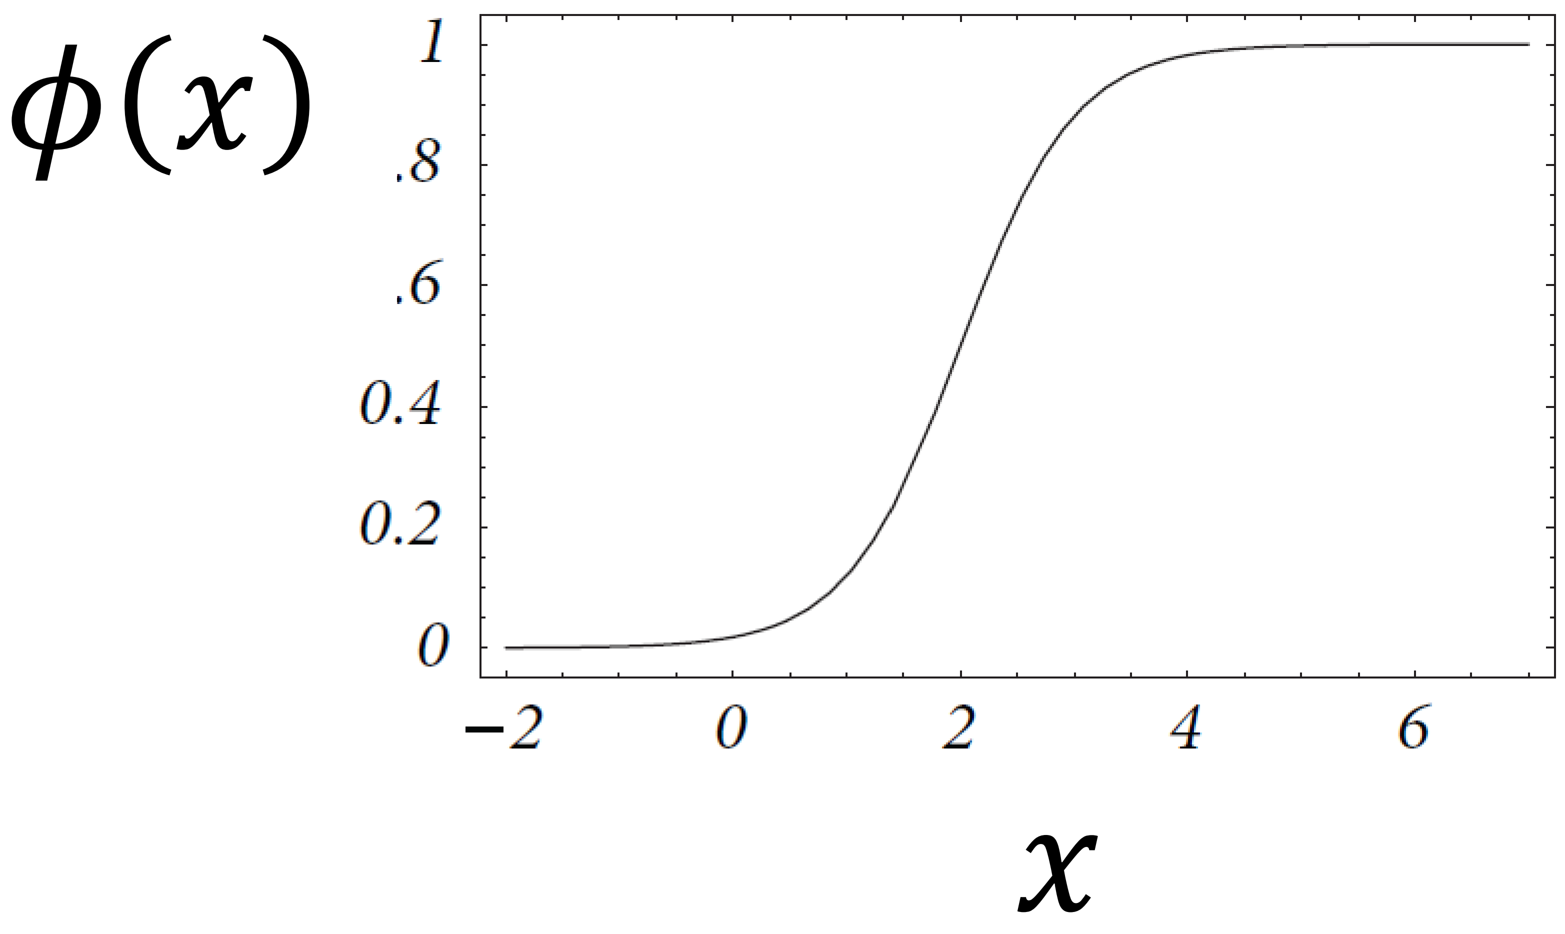

Spiking Neural Network (SNN)

Recurrent network of N neurons.

Membrane potential of neuron i:

$$\frac{d}{dt}h_i(t) = -\frac{1}{\tau}h_i(t) + \sum_j w_{ij}\,S_j(t) + I_i^{ext}(t)$$

"spiking model": spike causes jump

spike generated by inhomogeneous Poisson pr. with stochastic intensity

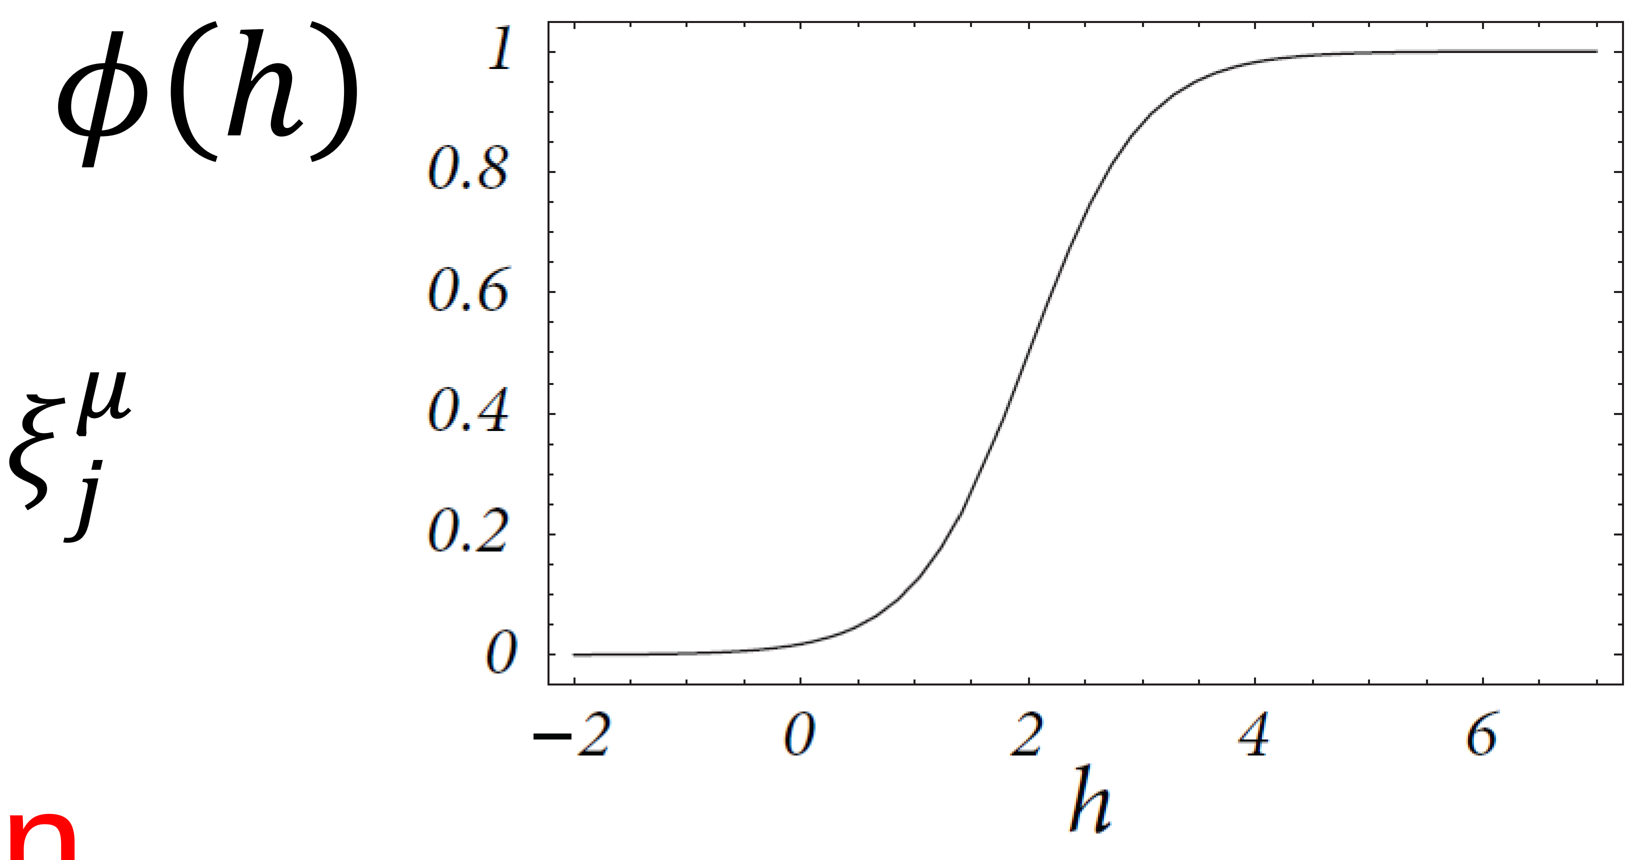

$$w_{ij} = \frac{1}{N}\sum_{\mu}^{P} \xi_i^{\mu}\,\xi_j^{\mu}$$

Gaussian

Compare SNN and RNN for same input, same connections

$$w_{ij} = \frac{1}{N}\sum_{\mu}^{P} \xi_i^{\mu}\,\xi_j^{\mu}$$

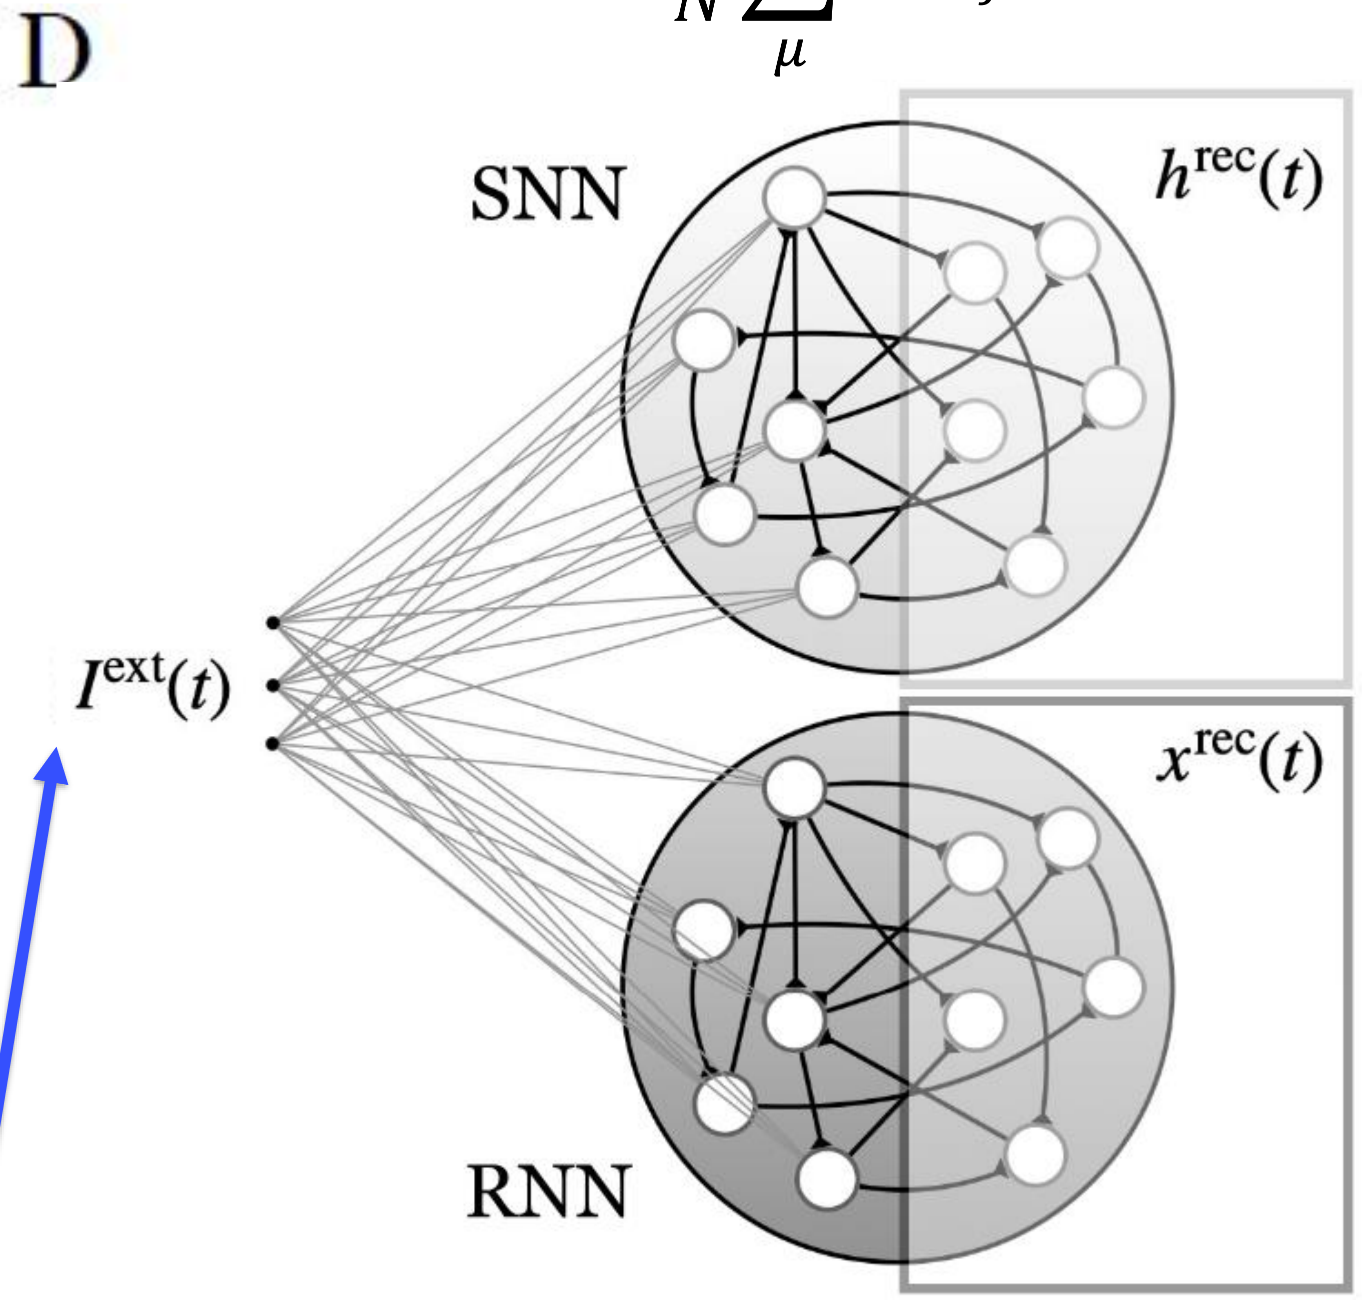

E

external input

$$I_i^{ext}(t) = \frac{1}{\sqrt{p}}\sum_{\mu}^{P} \xi_i^{\mu}\,\eta^{\mu}(t) \quad \text{if } i<N/2$$

N/2 neurons do not receive external input

V. Schmutz, J. Brea, W. Gerstner (2025) ***Emergent rate-based dynamics in duplicate-free populations of spiking neurons***
Physical Review Letters, 134:018401

Distance between potential in SNN (spikes) and RNN (rates)

Simulation for large N,

$$\Delta_N^{\text{rec}}(\alpha) := \frac{2}{N} \sum_{i=N/2+1}^{N} \lim_{T\to\infty} \frac{1}{T} \int_0^T \left| h_i^{\text{rec}}(t) - x_i^{\text{rec}}(t) \right| \mathrm{d}t$$

momentary distance

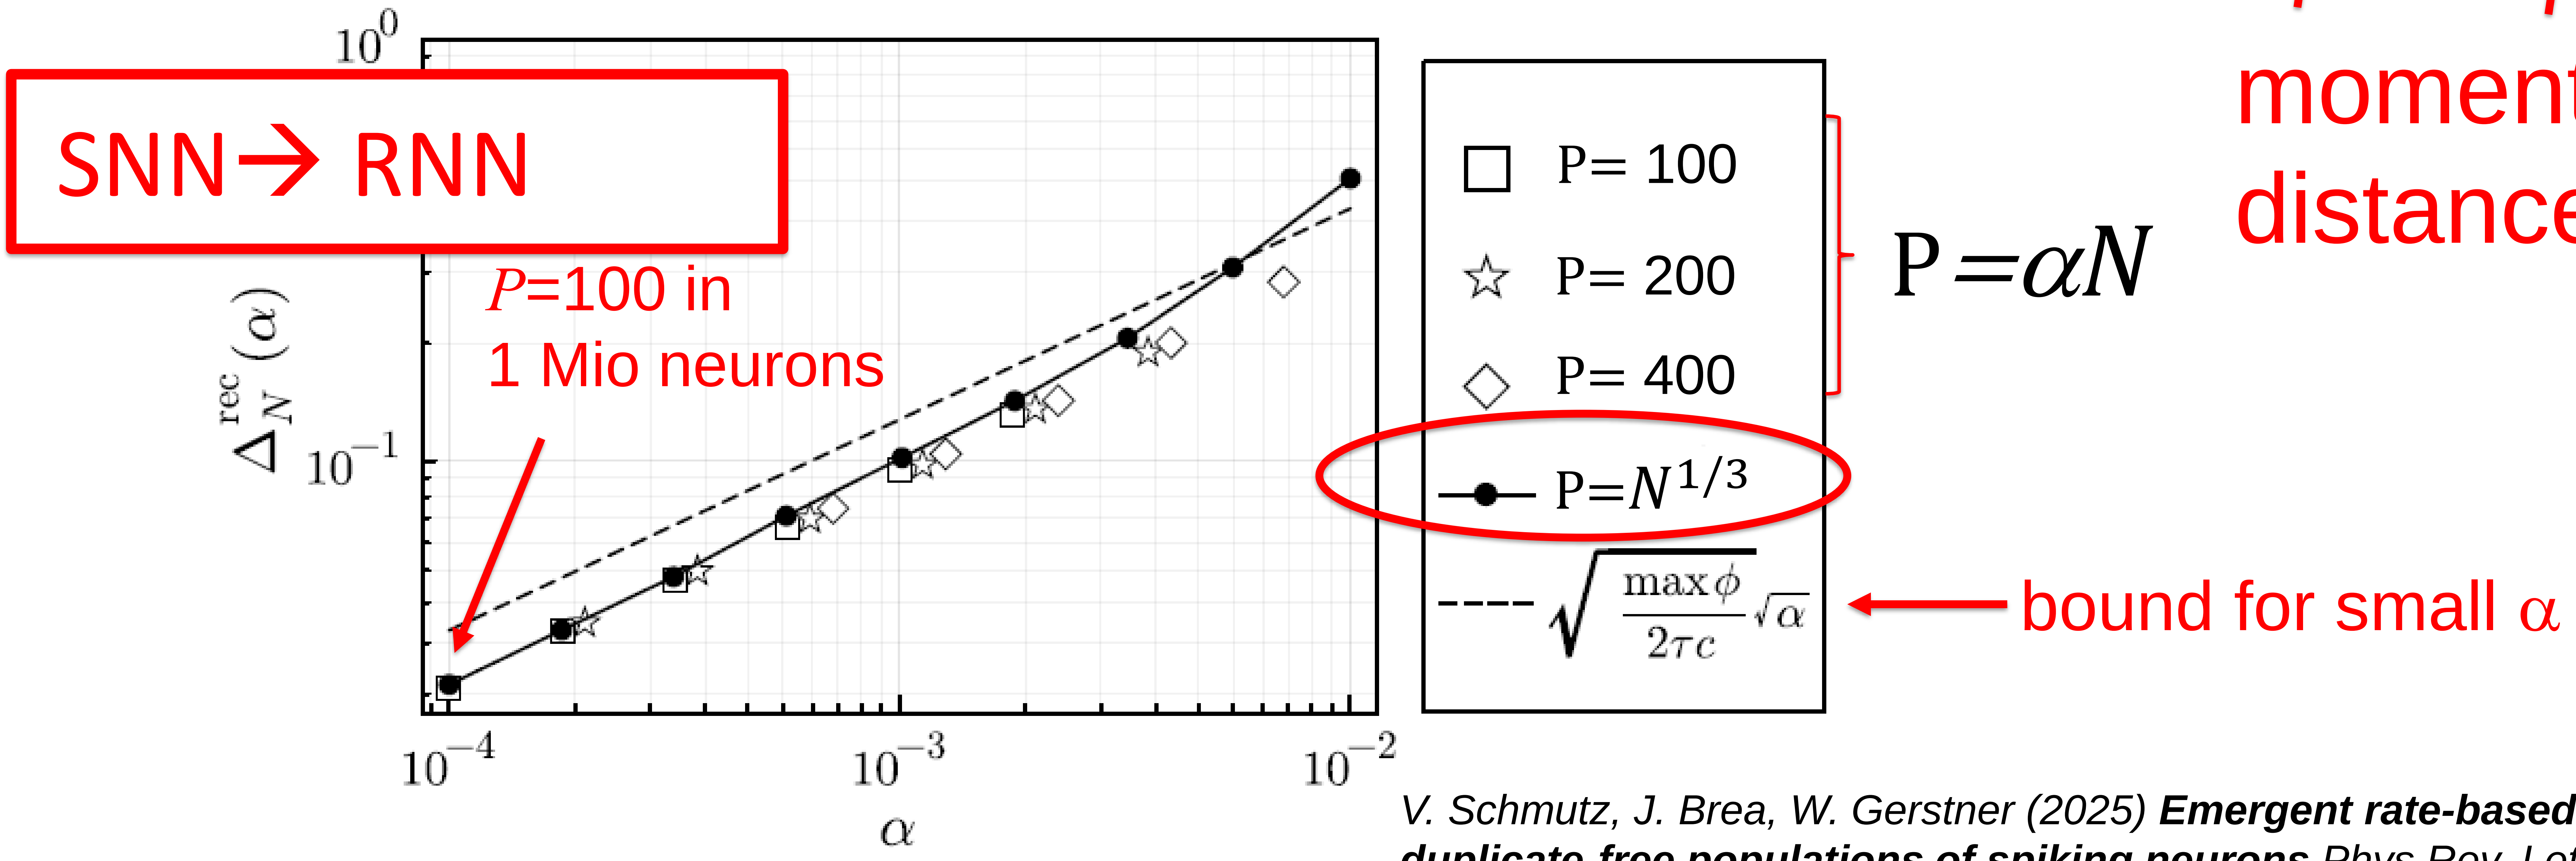

V. Schmutz, J. Brea, W. Gerstner (2025) ***Emergent rate-based dynamics in duplicate-free populations of spiking neurons*** *Phys.Rev. Lett. 134:018401*

effect of spikes disappears for large networks,
no need to average over time or neuron duplicates!

Summary Rate coding with **instantaneous time-dependent rates** is possible in network of spiking neurons even though not a single pair of neurons is correlated (**no duplicates**)

→ completely heterogeneous population
→ no spatial averaging
→ no temporal averaging

SNN → RNN

Rather: low-rank weight matrix

→ **low-dimensional** network-input to each neuron
→ neural activity lives in a $\boldsymbol{P}$**-dimensional manifold**
→ $e.g.\ P = N^{1/3}$
→ P=100-dimensional activity in 1 Mio neurons

V. Schmutz, J. Brea, W. Gerstner (2025) ***Emergent rate-based dynamics in duplicate-free populations of spiking neurons*** *Phys.Rev. Lett. 134:018401*

Answer to our big question:

Is rate coding meaningful in spiking neurons, even if

- **signals are fast** (no temporal averaging possible)
- there are **no duplicate neurons**

YES!!!!, if 'low-rank' connectivity

NxN matrix

$$w_{ij} = \frac{1}{N} \sum_{\mu}^{P} \xi_i^{\mu} \, \xi_j^{\mu}$$

→ **rank *P***

Gaussian (or binary +/-1)

V. Schmutz, J. Brea, W. Gerstner (2025) ***Emergent rate-based dynamics in duplicate-free populations of spiking neurons*** *Phys.Rev. Lett. 134:018401*

Is low-rank connectivity a strange assumption?

1) *"Neurons have receptive fields and wiring patterns: is a low-rank model realistic AT ALL?"*

Barack, D.L., Krakauer, J.W.: Two views on the cognitive brain. Nat. Rev. Neurosc. (2021)
Langdon, C., Genkin, M., Engel, T.A.: A unifying perspective on neural manifolds and circuits for cognition. Nat. Rev. Neurosci. (2023)

Answer: All standard models of cortex are dominated by a low-rank connectivity matrix

Pezon, L., Schmutz, V, Gerstner, W. (2024), Linking Neural Manifolds to Principles of Circuit Structure in Recurrent Networks bioRxiv doi: https://doi.org/10.1101/2024.02.28.582565

2) *"How are low-rank networks related to low-dim. dynamics?"*

Answer: rank *P* weight matrix (outer product matrix) **always** generate *P*-dimensional dynamics (→neural manifolds)

Mastrogiuseppe, F., Ostojic, S.: Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. Neuron 99(3), 609–62329 (2018)

Conclusions

- SNN → RNN without averaging!
- rather ‘loose’ conditions
- rank *P* can be ‘relatively large’

V. Schmutz, J. Brea, W. Gerstner (2025) ***Emergent rate-based dynamics in duplicate-free populations of spiking neurons***
Phys.Rev. Lett. 134:018401