<u>The Maps Inside Your Head</u>

Vijay Balasubramanian

Department of Physics
&
Computational Neuroscience Initiative
University of Pennsylvania

Small brain, big world

The brain gathers information from the world, makes decisions for future actions, learns from experience, and tries to remember.

The Challenge: It's a very big, complex, and often unpredictable world, and the brain has very limited resources

The Question: What are the organizational principles that allow neural circuits to meet this enormous challenge?

Specialization of Circuits in The Brain

Neurons & Neural Communication

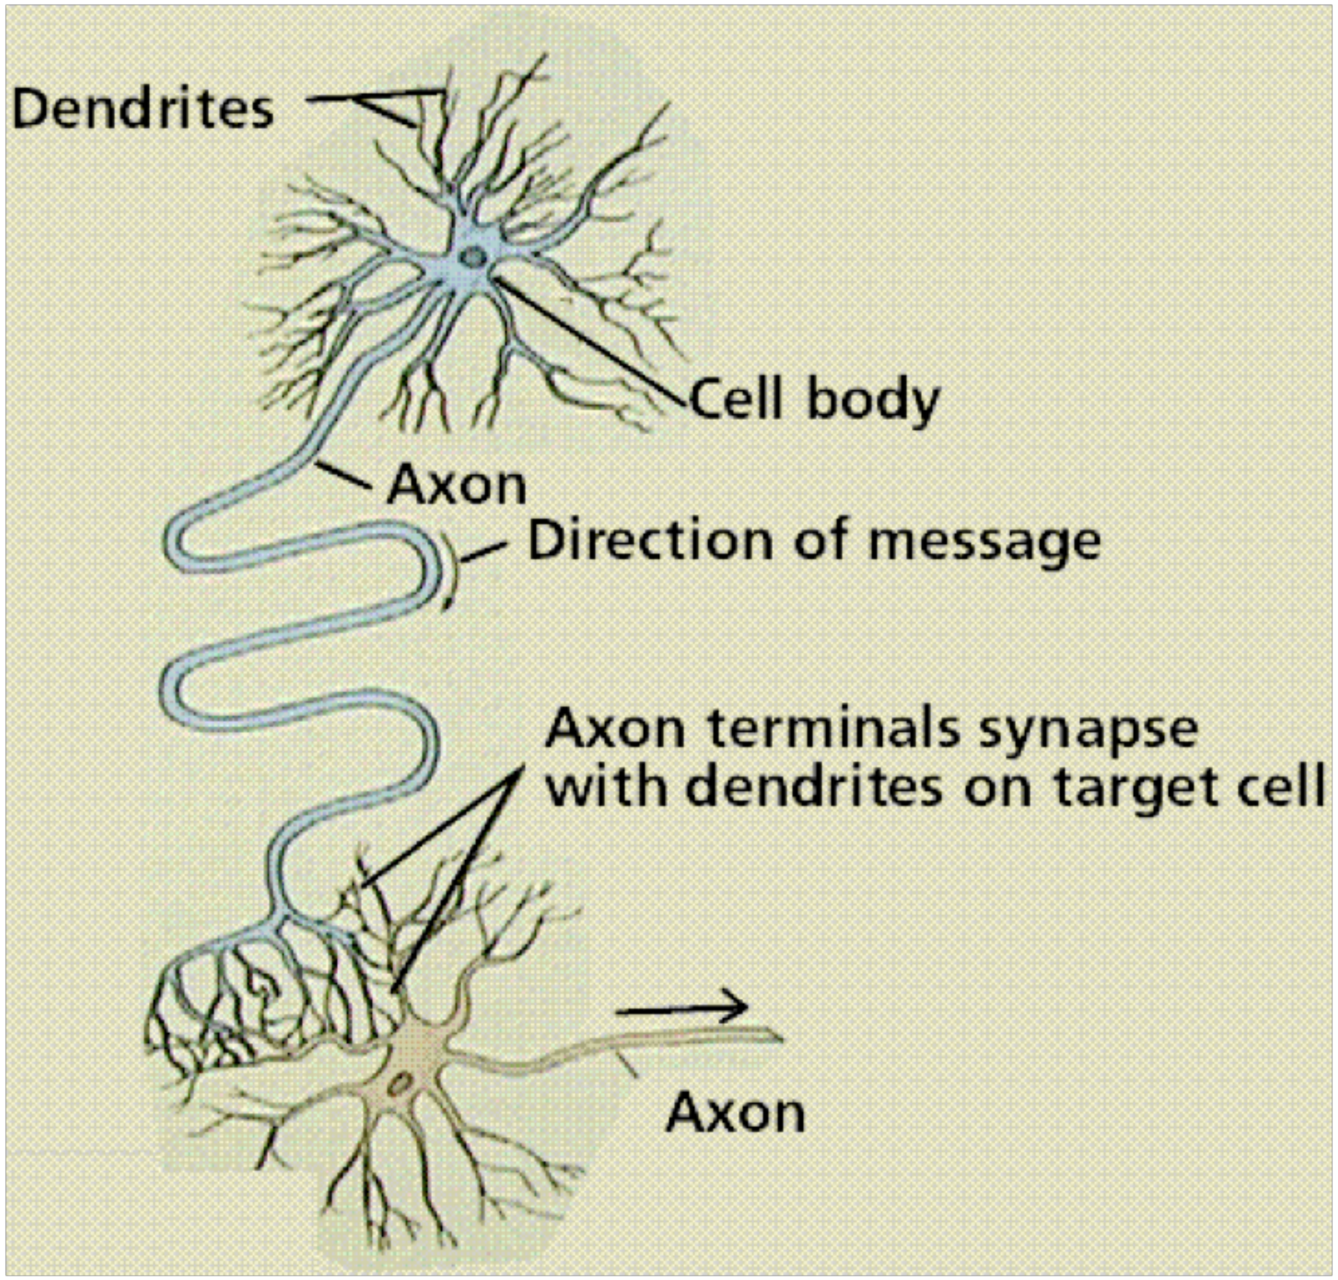

From Purves et al., *Life: The Science of Biology*

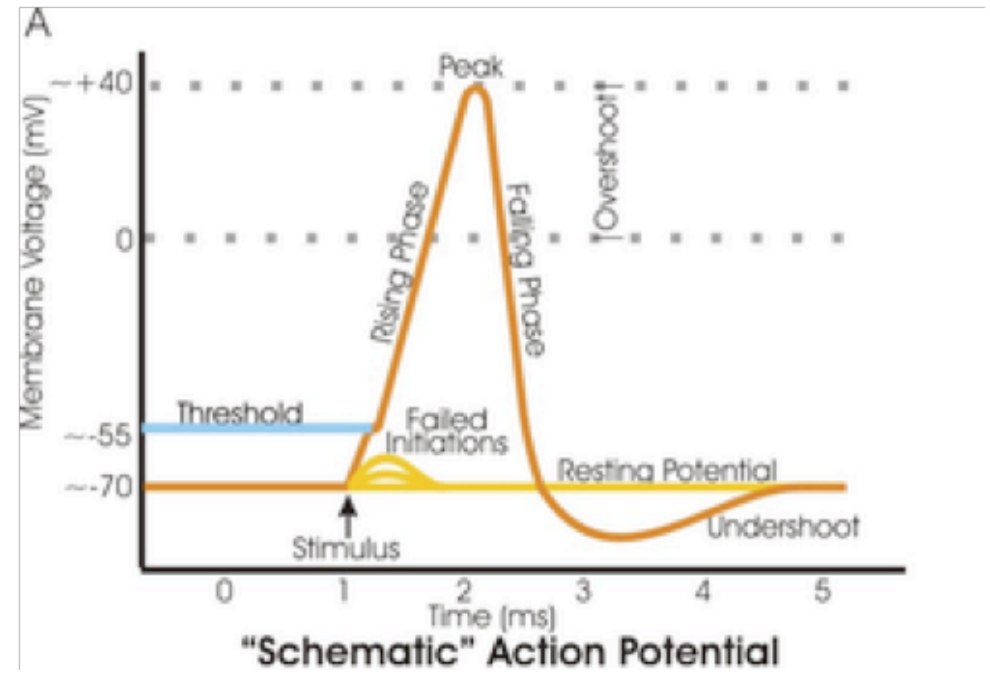

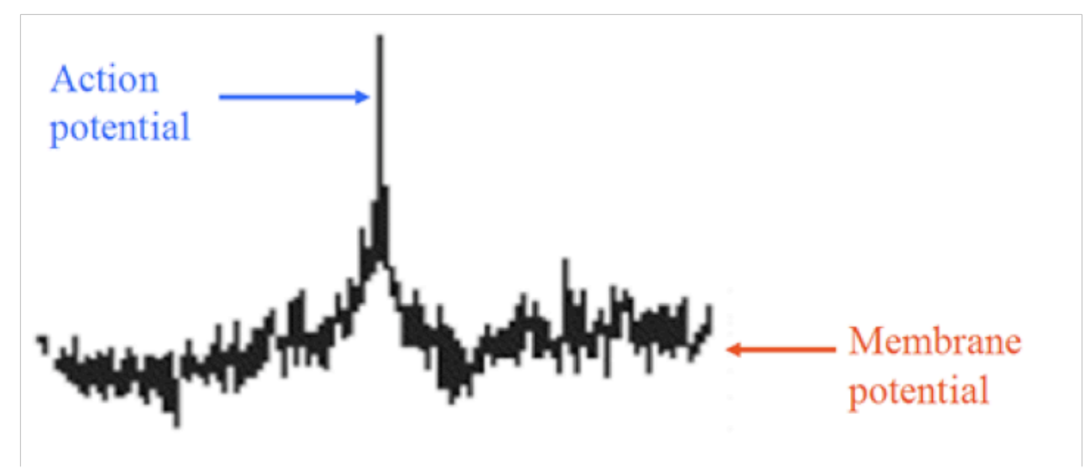

From Wikipedia article on Action Potentials

Specialization of function: brain areas

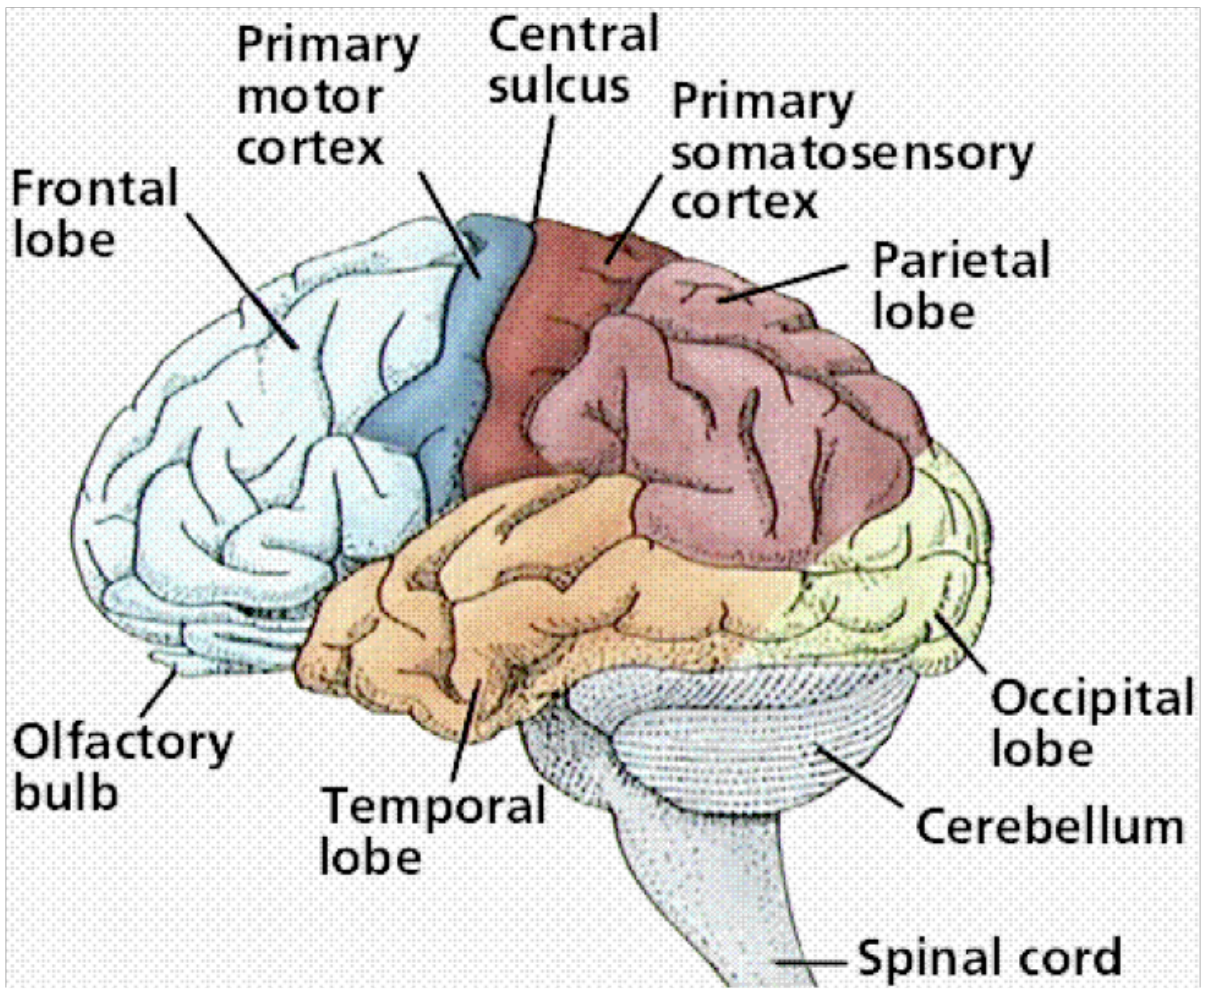

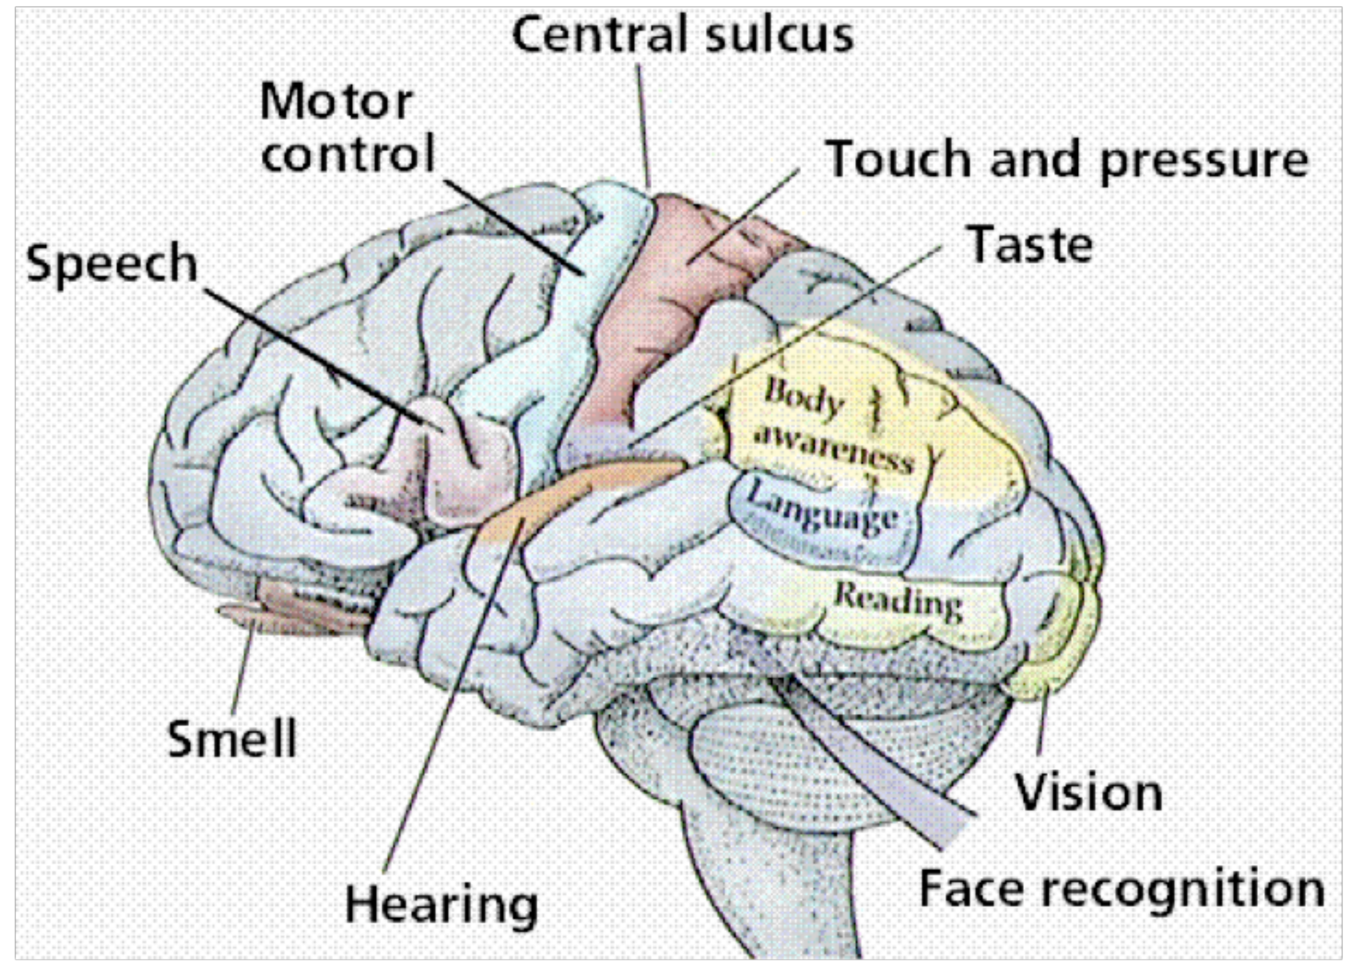

From Purves et al., *Life: The Science of Biology*

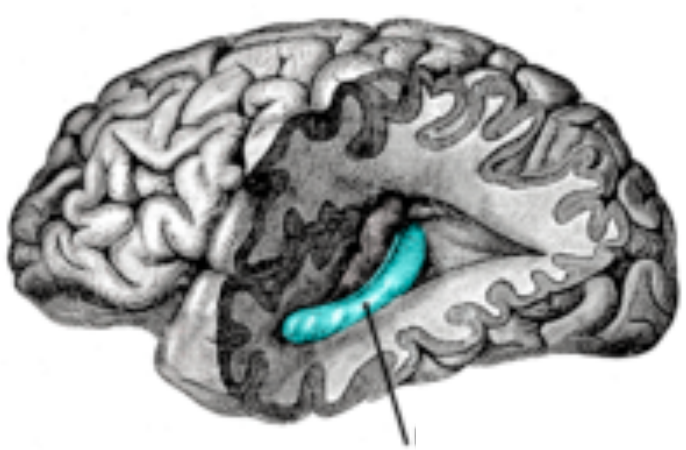

From Wikipedia article on HIPPOCAMPUS

memory formation

Specialization of function: layers within brain areas

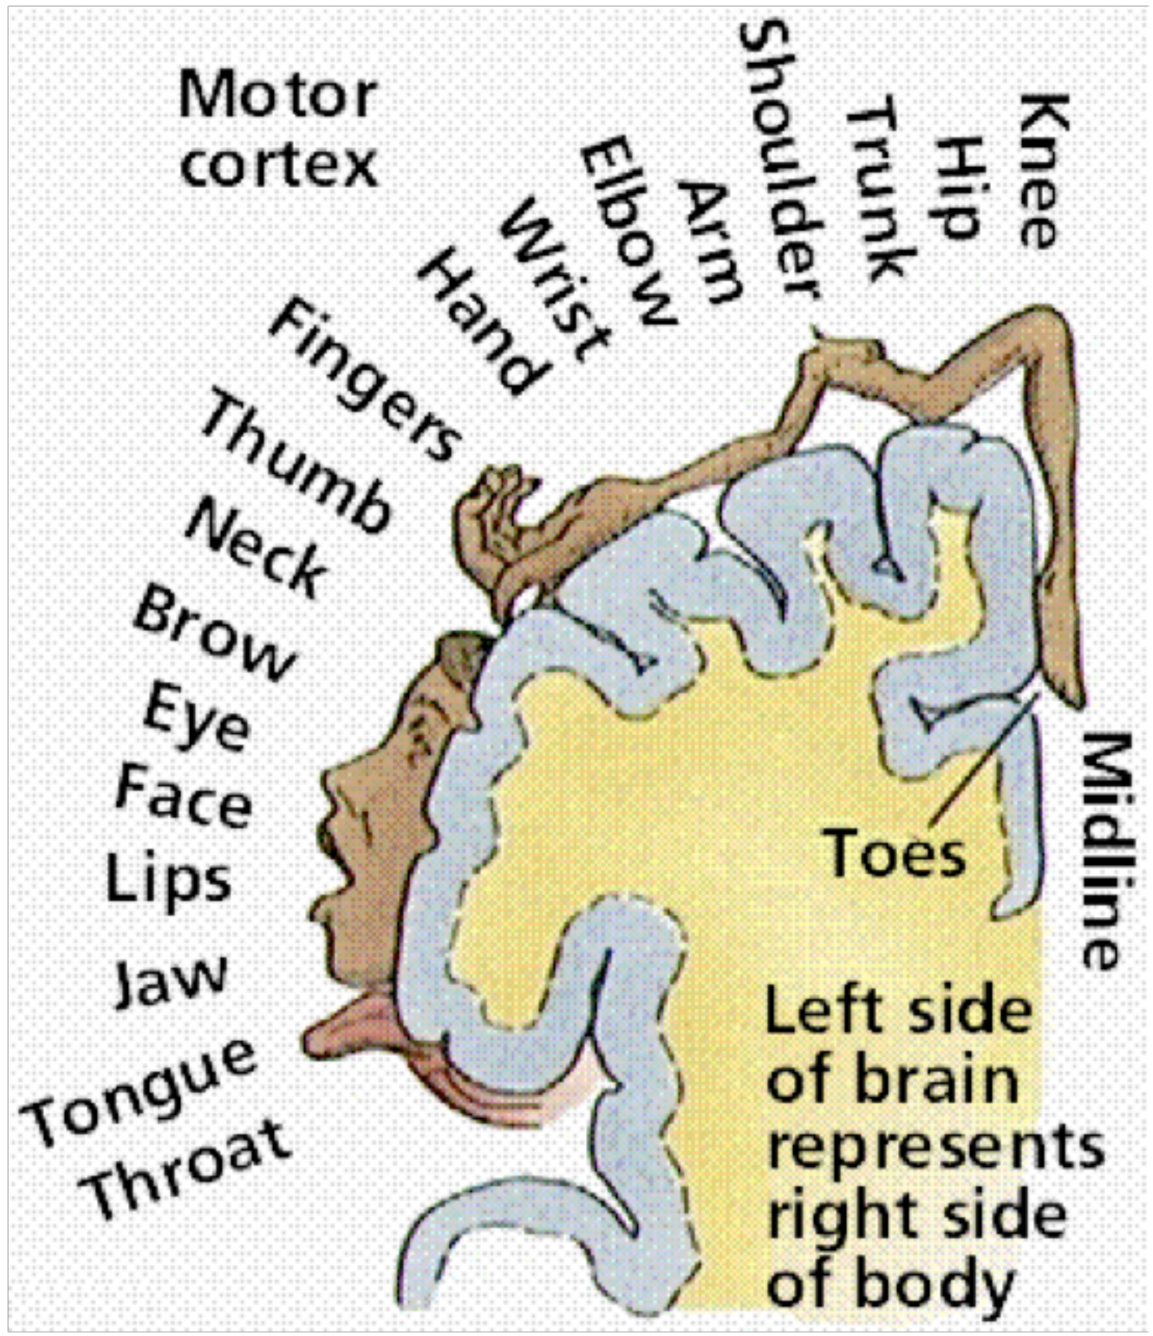

From Purves et al., *Life: The Science of Biology*

Specialization of circuits: the retina

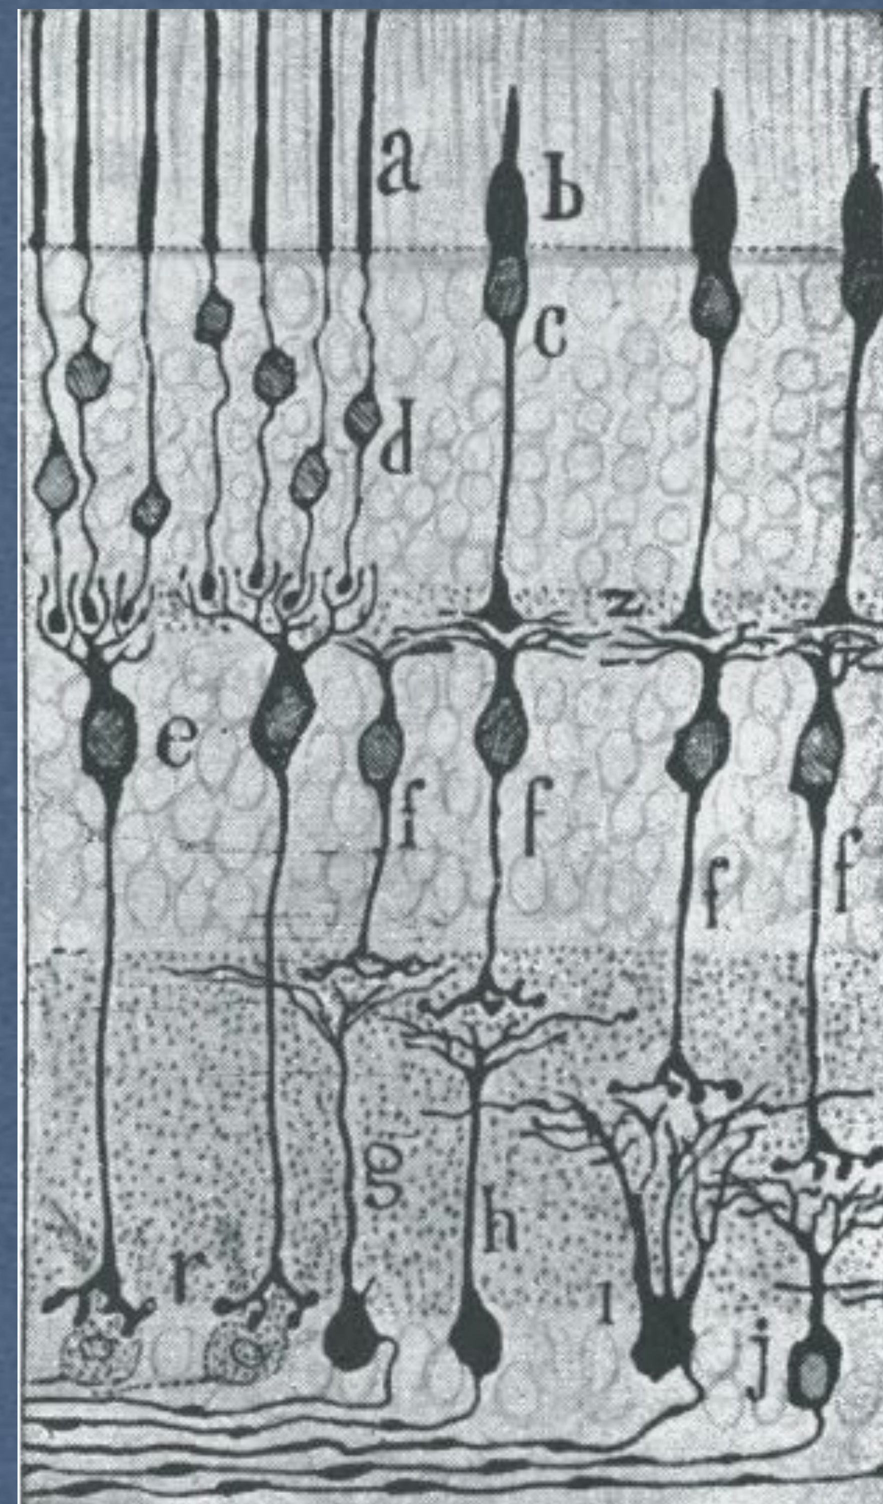

Cajal 1917

Specialization: cell types in the retina

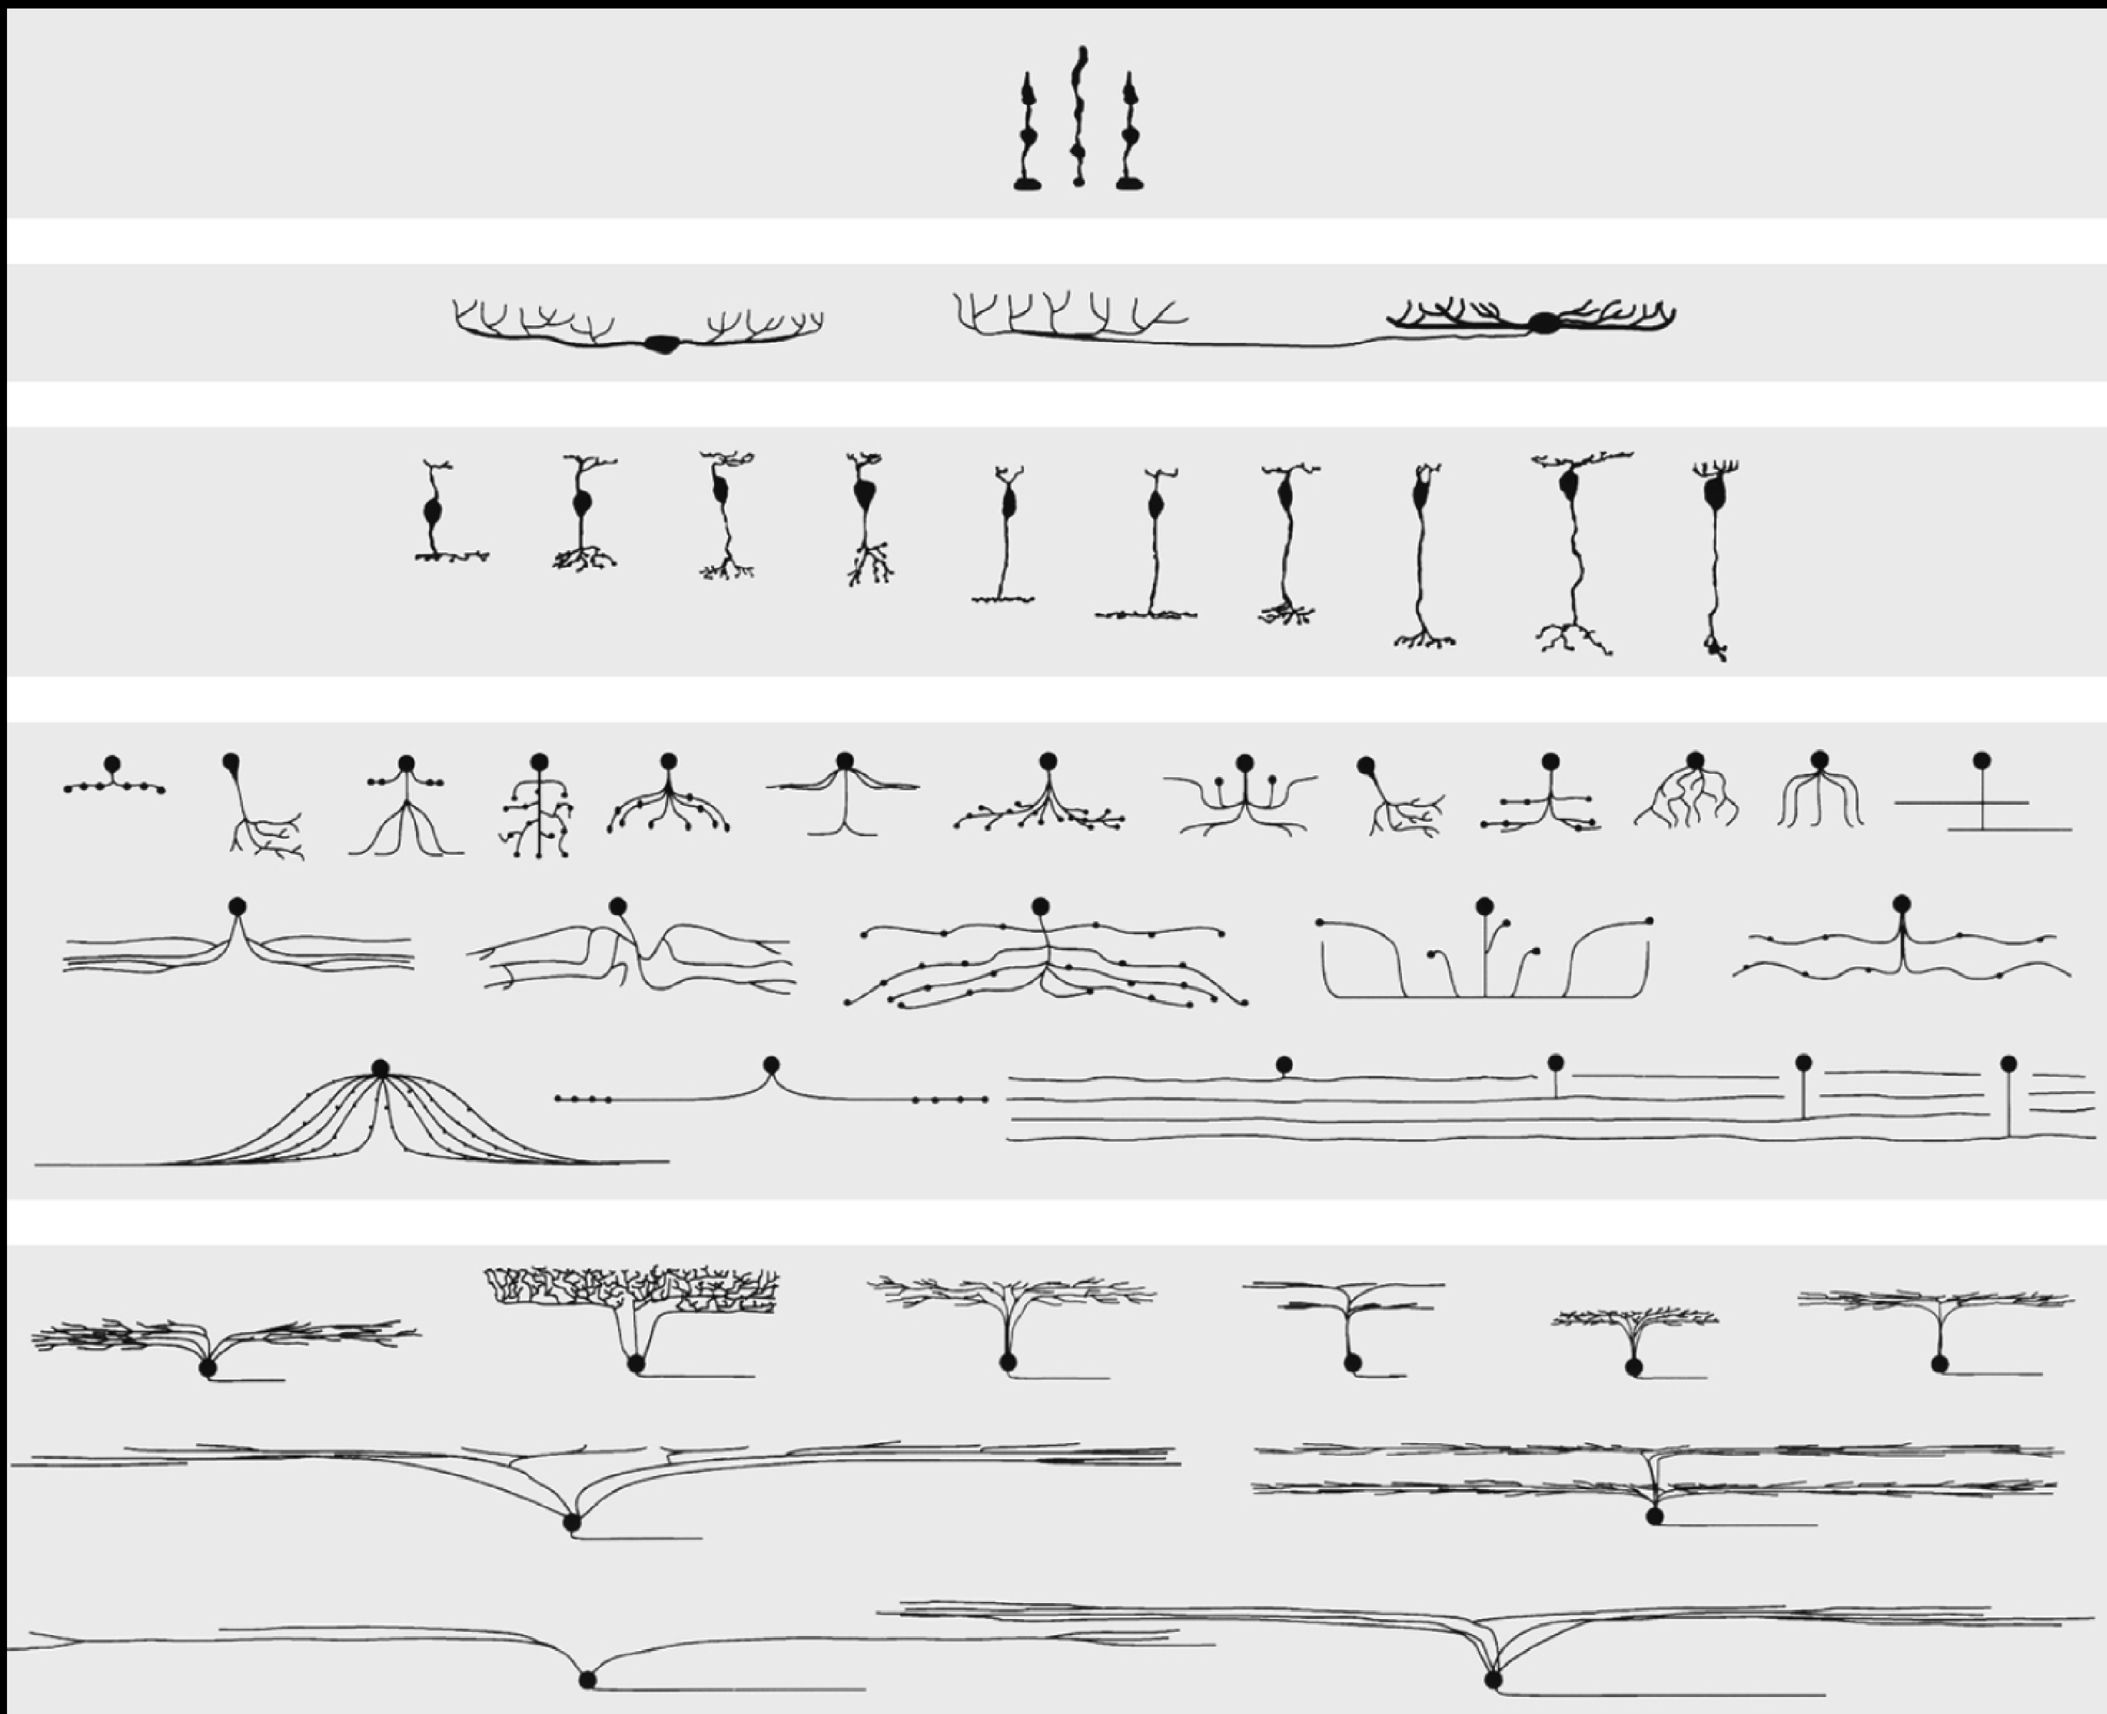

3 cones
↓
2 horizontal cells
↓
10 bipolars
↓
30 interneurons
↓
15 ganglion cells

Masland 2001

Specialization: Precise microcircuitry

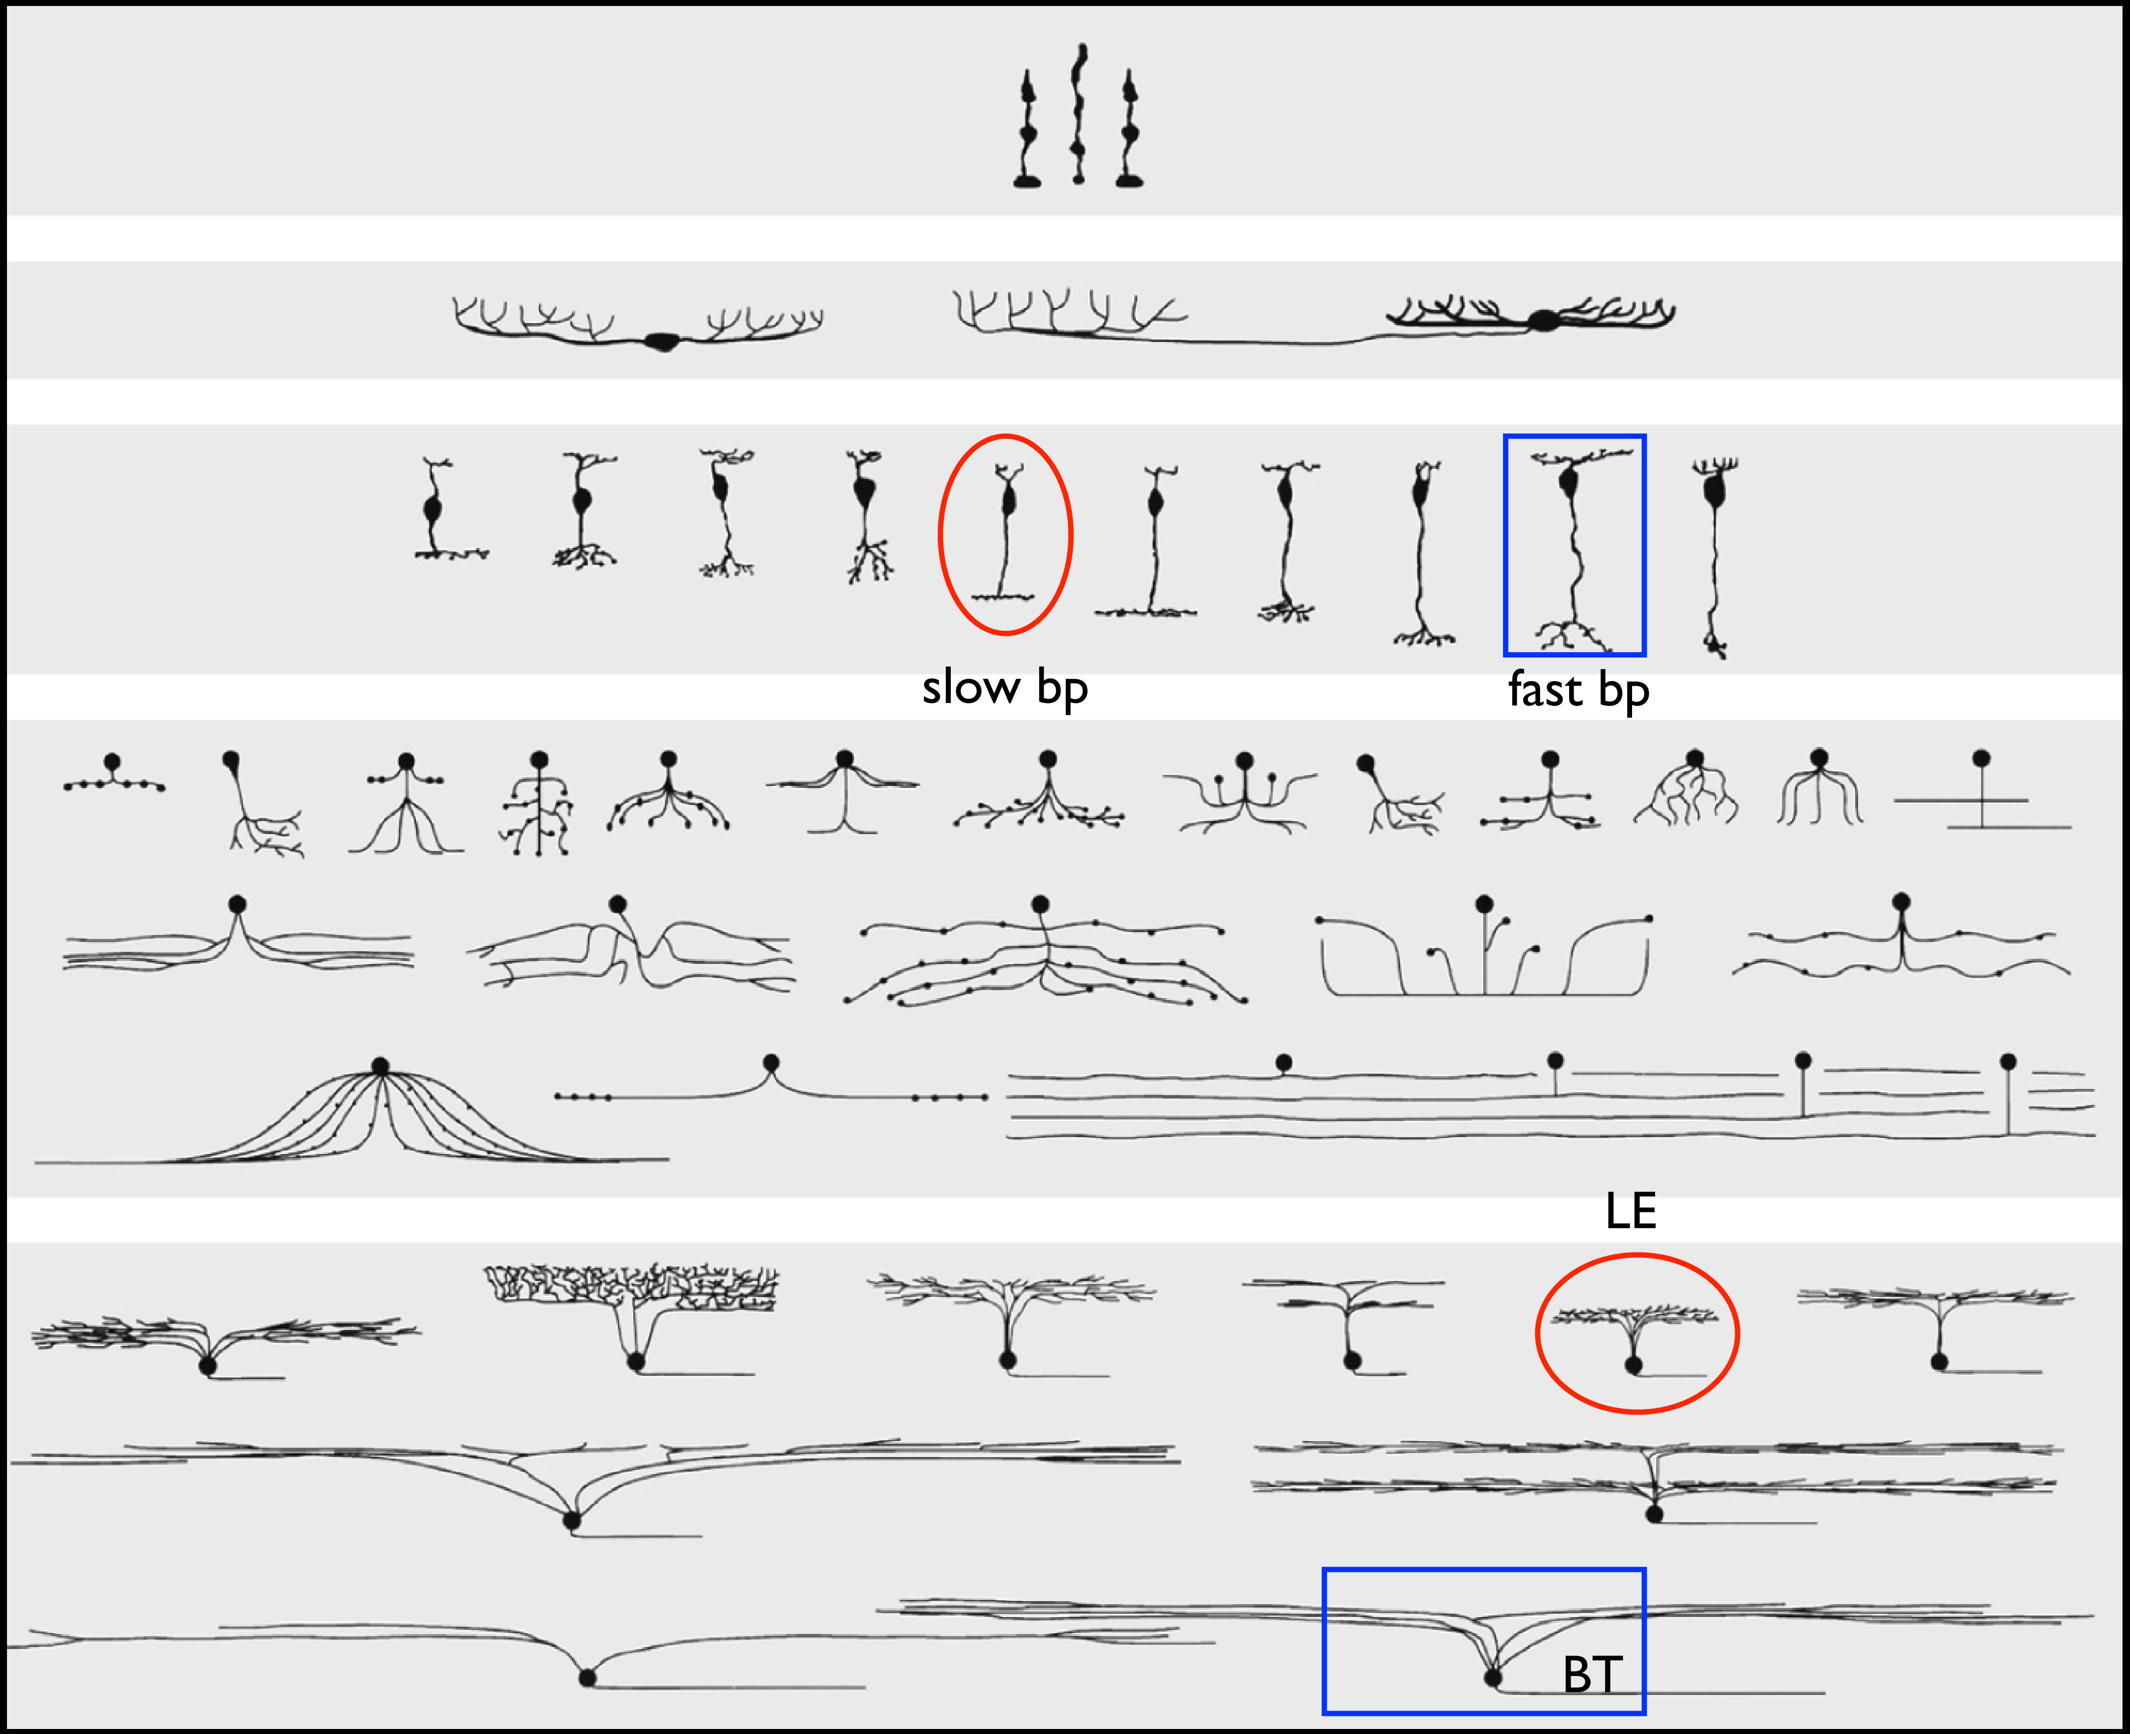

slow bp → LE

fast bp → BT

Maps in the brain

At every scale of organization, the brain has a diverse repertoire of functional units whose coordinated activity produces the desired overall function.

e.g. Whole Brain
(everything you do)

e.g. Entorhinal Cortex
(the sense of place)

e.g Retina
(vision)

e.g Olfactory Cortex
(smell)

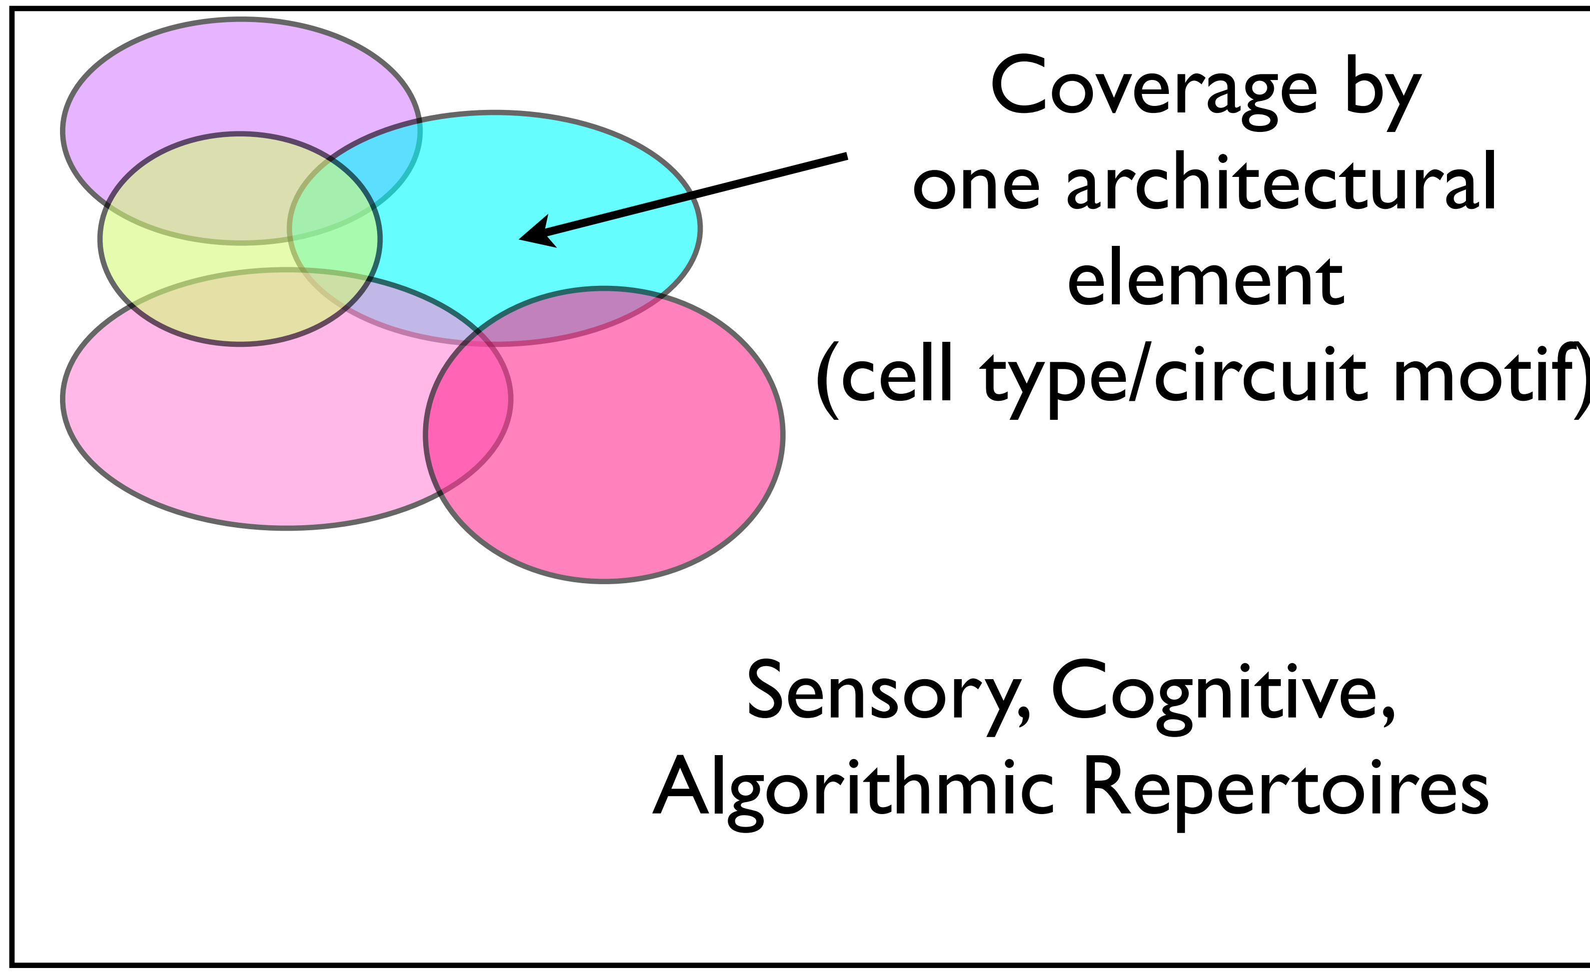

A view of the repertoire: A memory of computations that have predictive value for behavior, learned over evolutionary time, encoded in the genome and the developmental program.

What organizational principles (“laws”) control the computational & information-processing repertoires of the brain?

The Costs of Computation

vs.

- Brain: 2% of body weight, but 20% of metabolic load.
- Brain: Every mm^3 contains 4 km of wire

Power and space are major constraints -- (Attwell & Laughlin; Wen & Chklovskii)

- Brain consumes ~12W of power (refrigerator lightbulb)
- Packing seems to minimize wire length

How does the brain achieve its efficiency?

Idea: Specialization in the circuit repertoire &
Adaptation to structure in the world

A theory of maps in the brain?

HYPOTHESIS: Brains exploit structure in the world to efficiently allocate limited computational resources to maximize gain for the organism

- Retina
(visual features)
- Entorhinal Cortex
(the sense of place)
- Olfactory Cortex
(complex smells)
- Whole Brain
(everything you do)

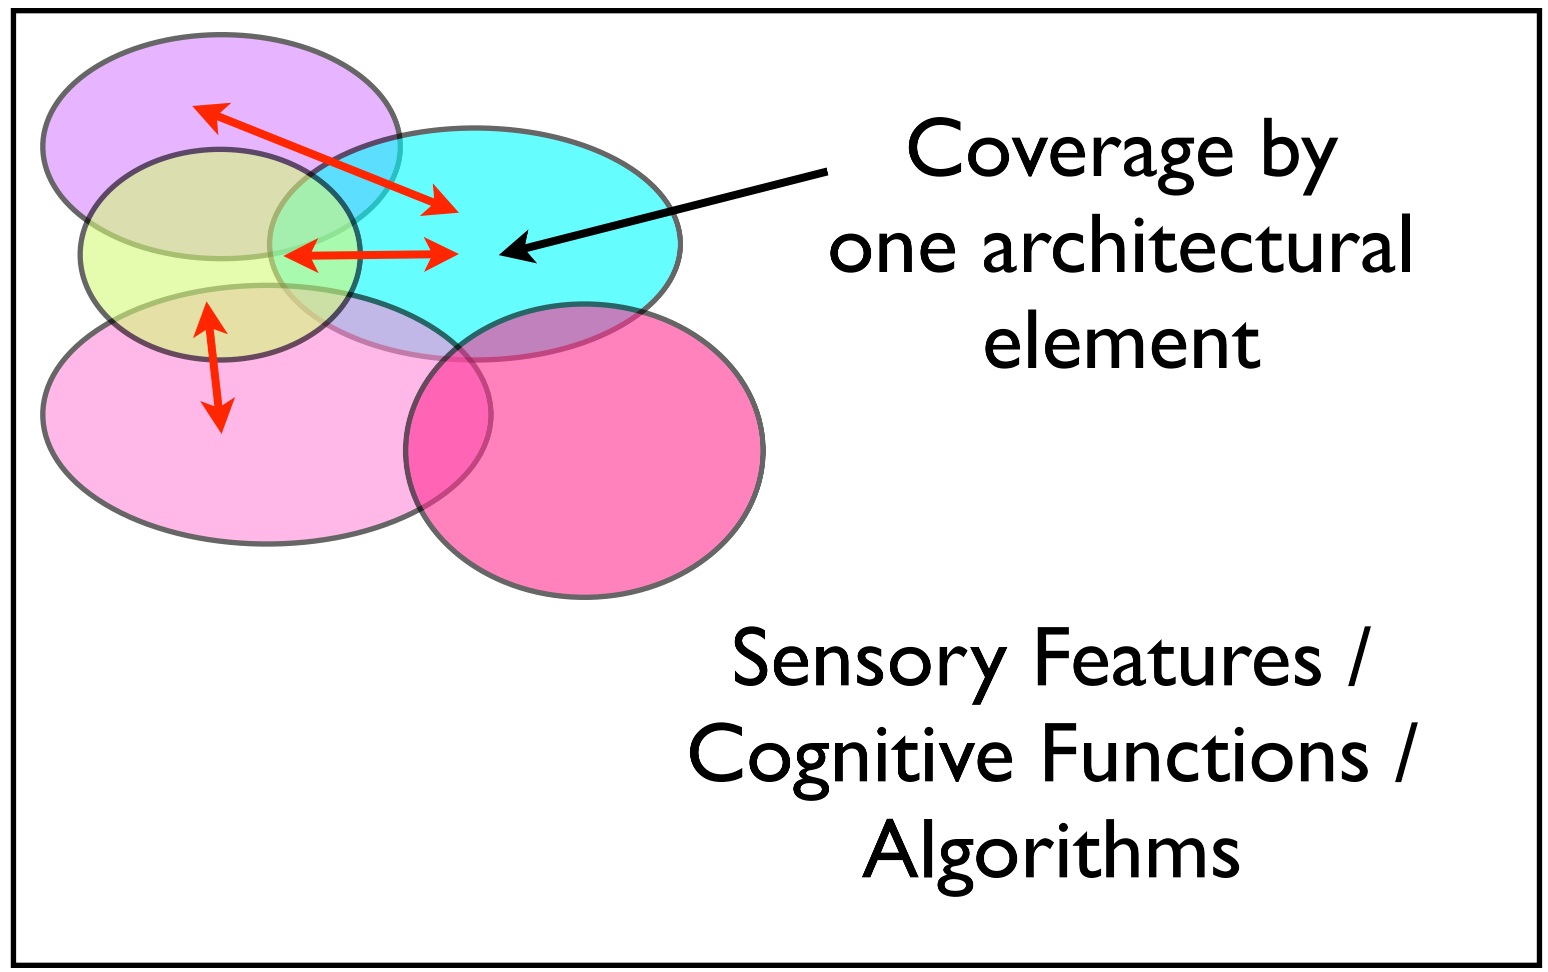

Roadmap

- Example 1: Vision (the sense of sight)
- Example 2: Spatial cognition (the sense of place)
- Example 3: Olfaction (the sense of smell)

In each case we will see that evolution seems to have exploited sophisticated mathematical principles of information processing that have only recently been discovered.

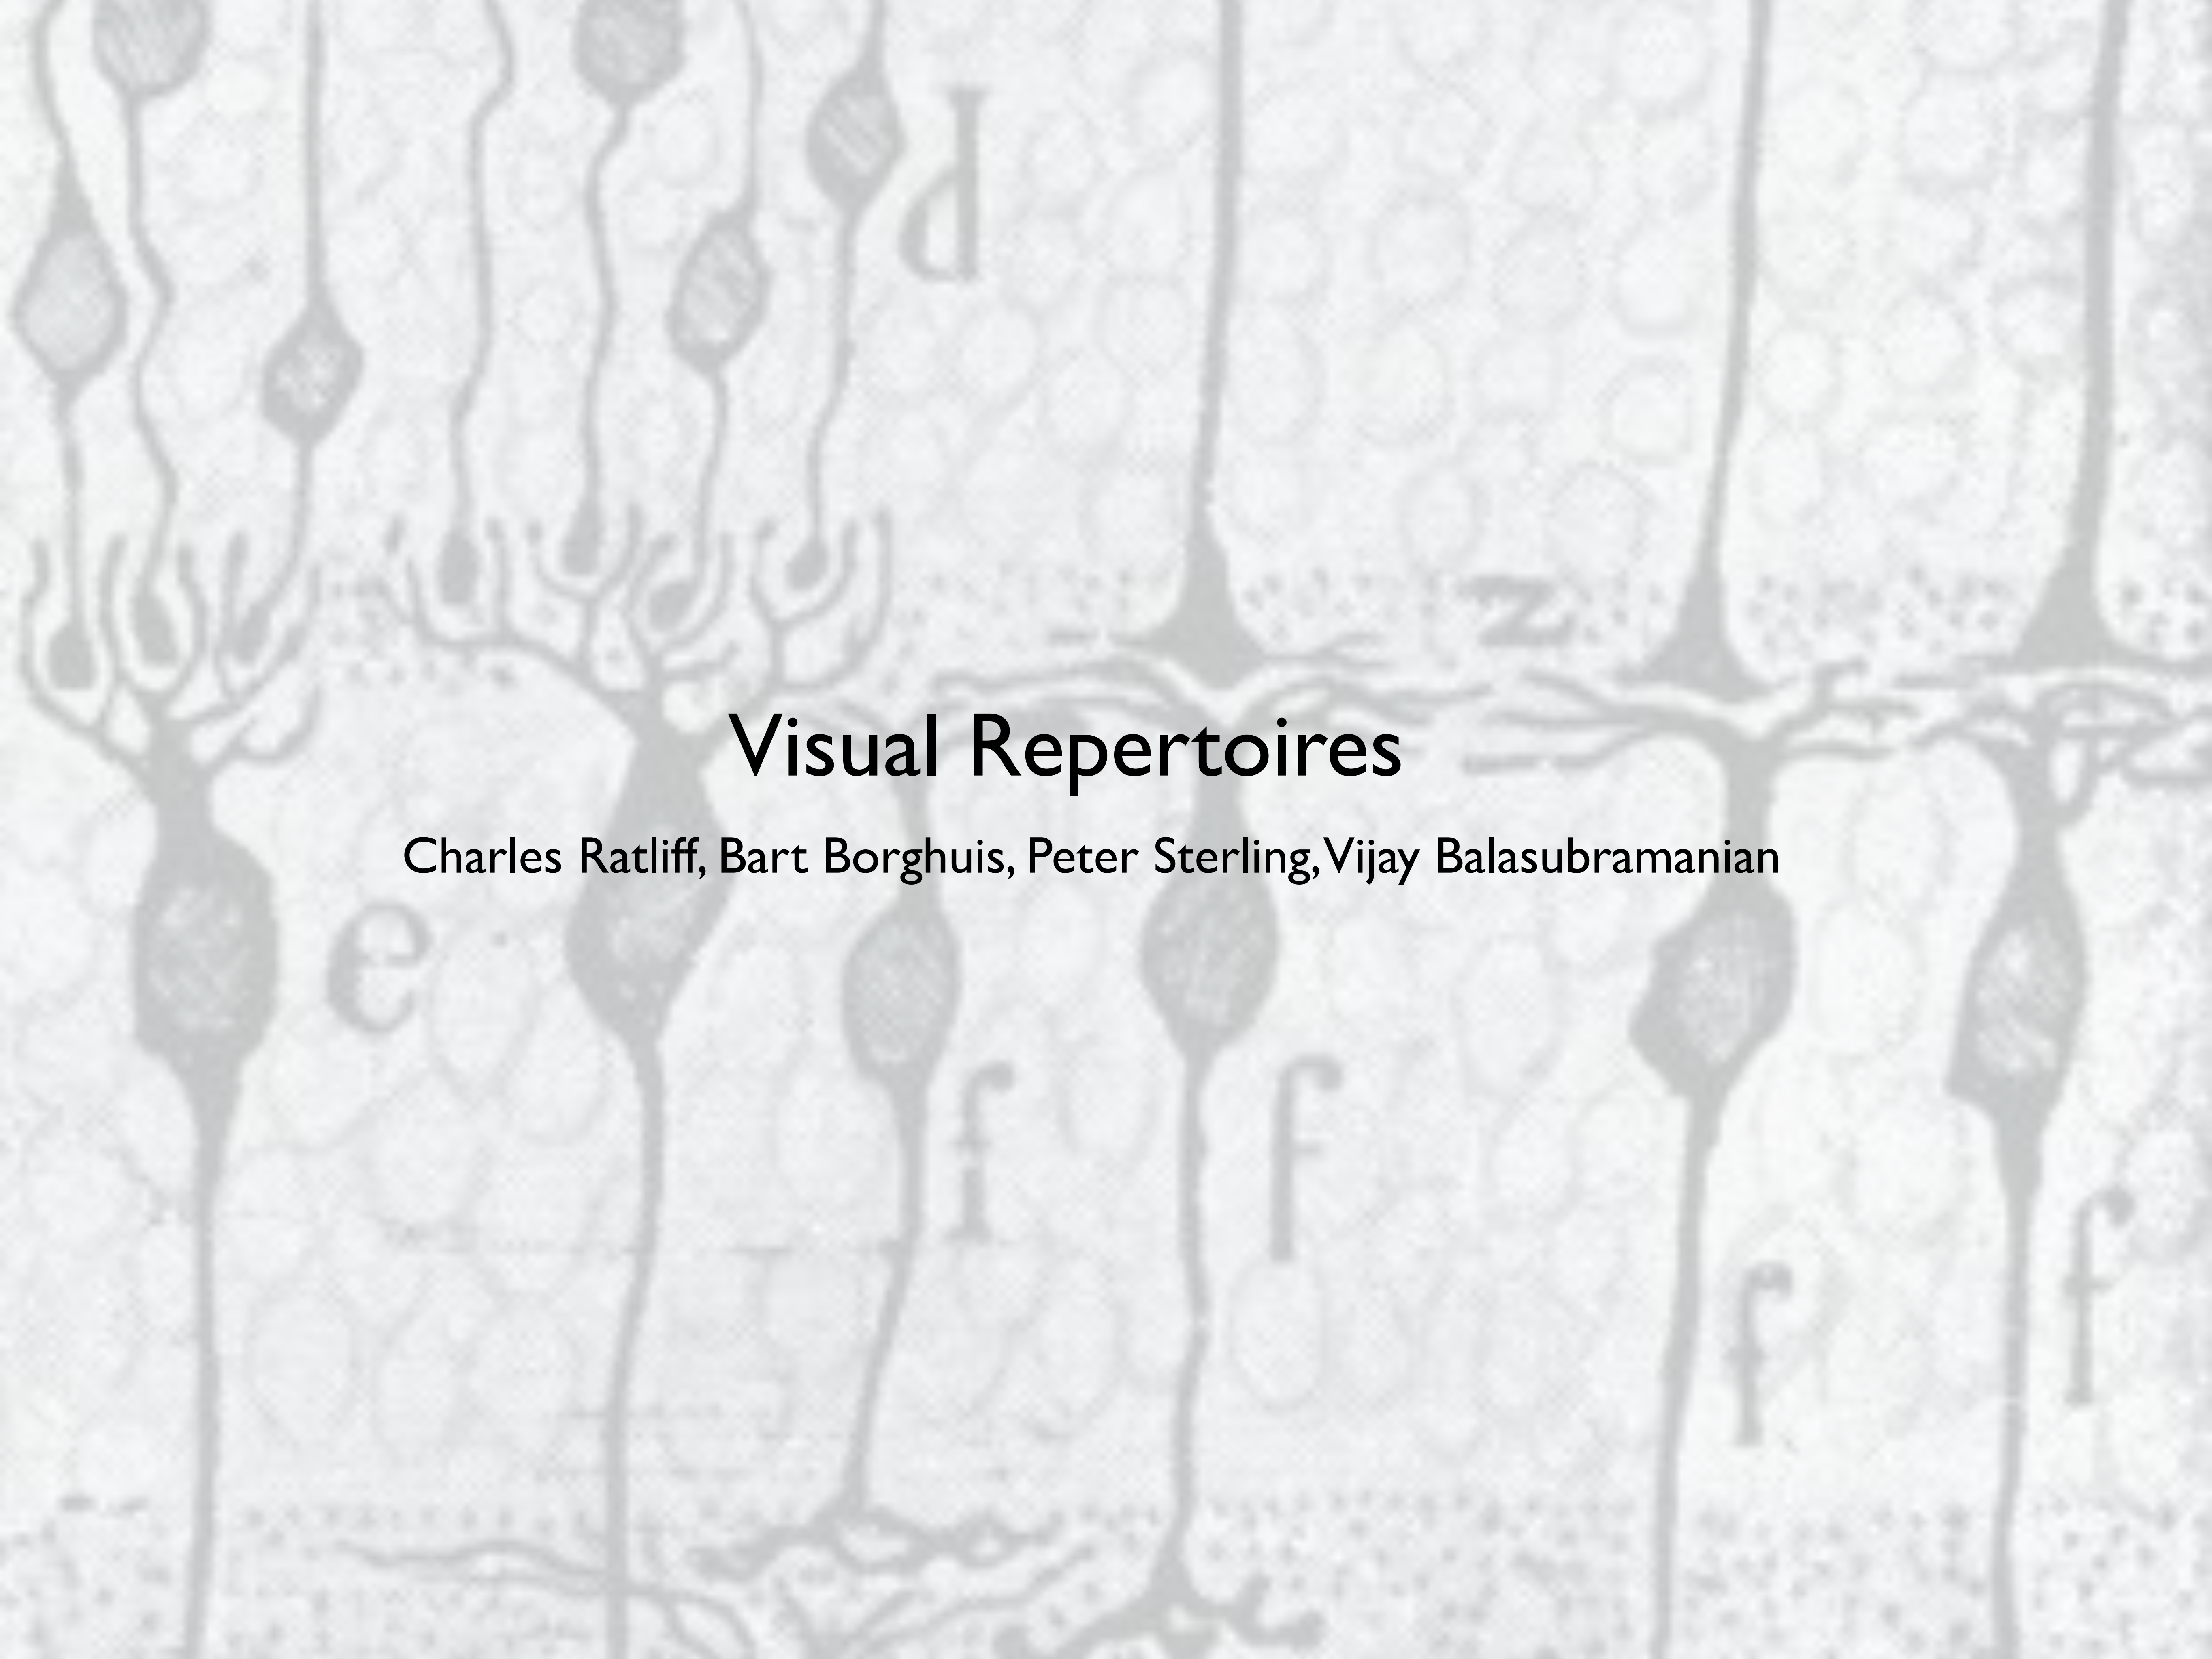

Visual Repertoires

Charles Ratliff, Bart Borghuis, Peter Sterling, Vijay Balasubramanian

The retinal repertoire

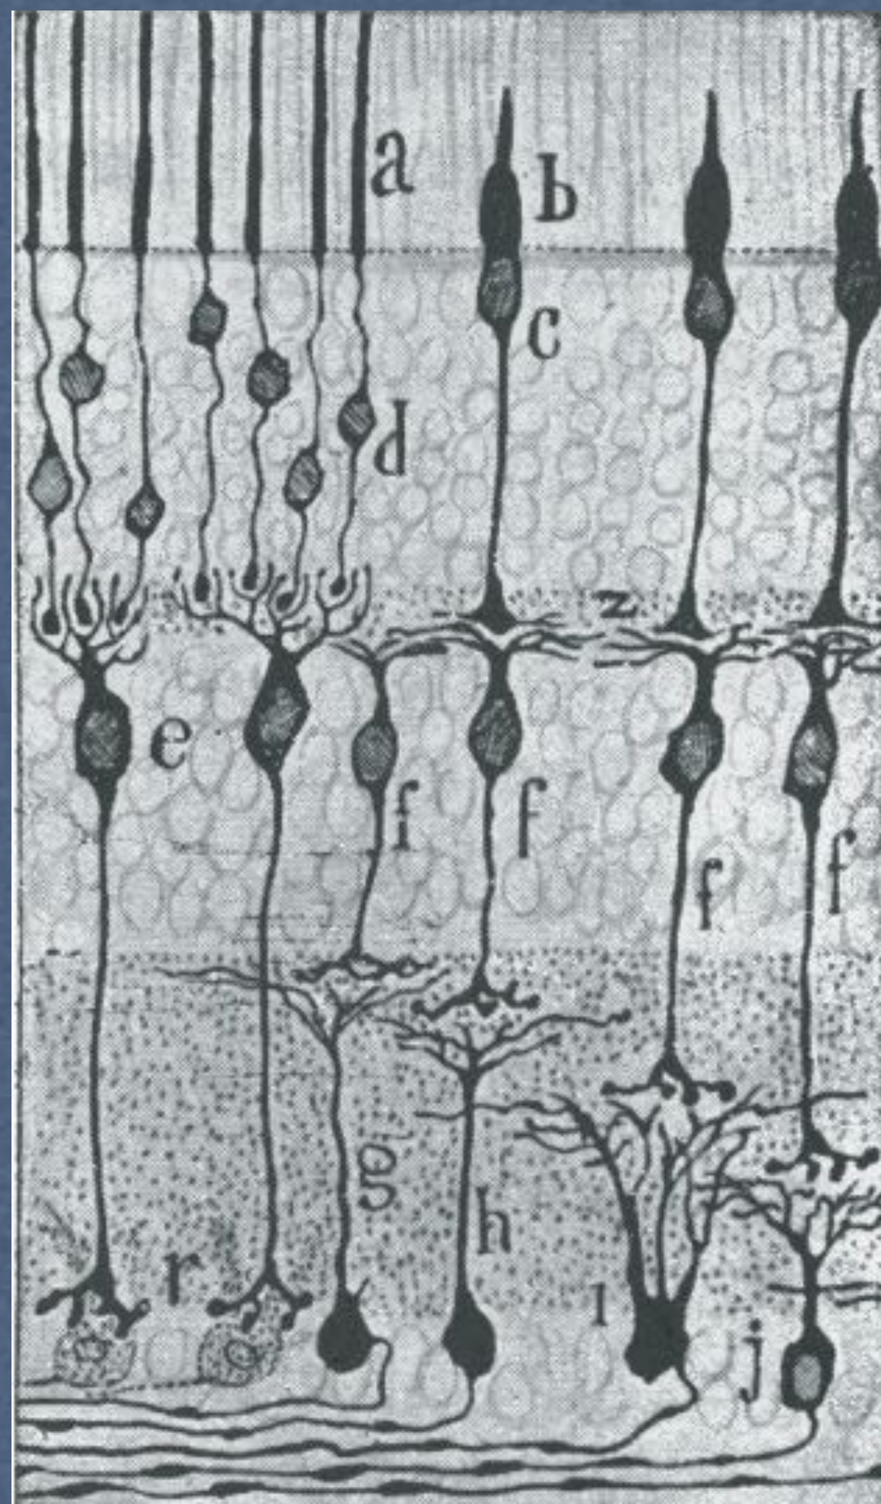

Cajal 1917

- Retinal ganglion cells (the output cells of the retina) detect "features" of the world (bright spots / dark spots / color / local motion) and report them to the brain.

- How should the repertoire of ganglion cells (1,000,000 in humans) be divided into types responding to these different features.

- Let's consider the example of bright and dark spot detectors (ON and OFF cells)

Statistical structure of natural scenes

Garrigan, Ratliff, Sterling, Brainard, Balasubramanian (PLoS Comp Bio, 2010)

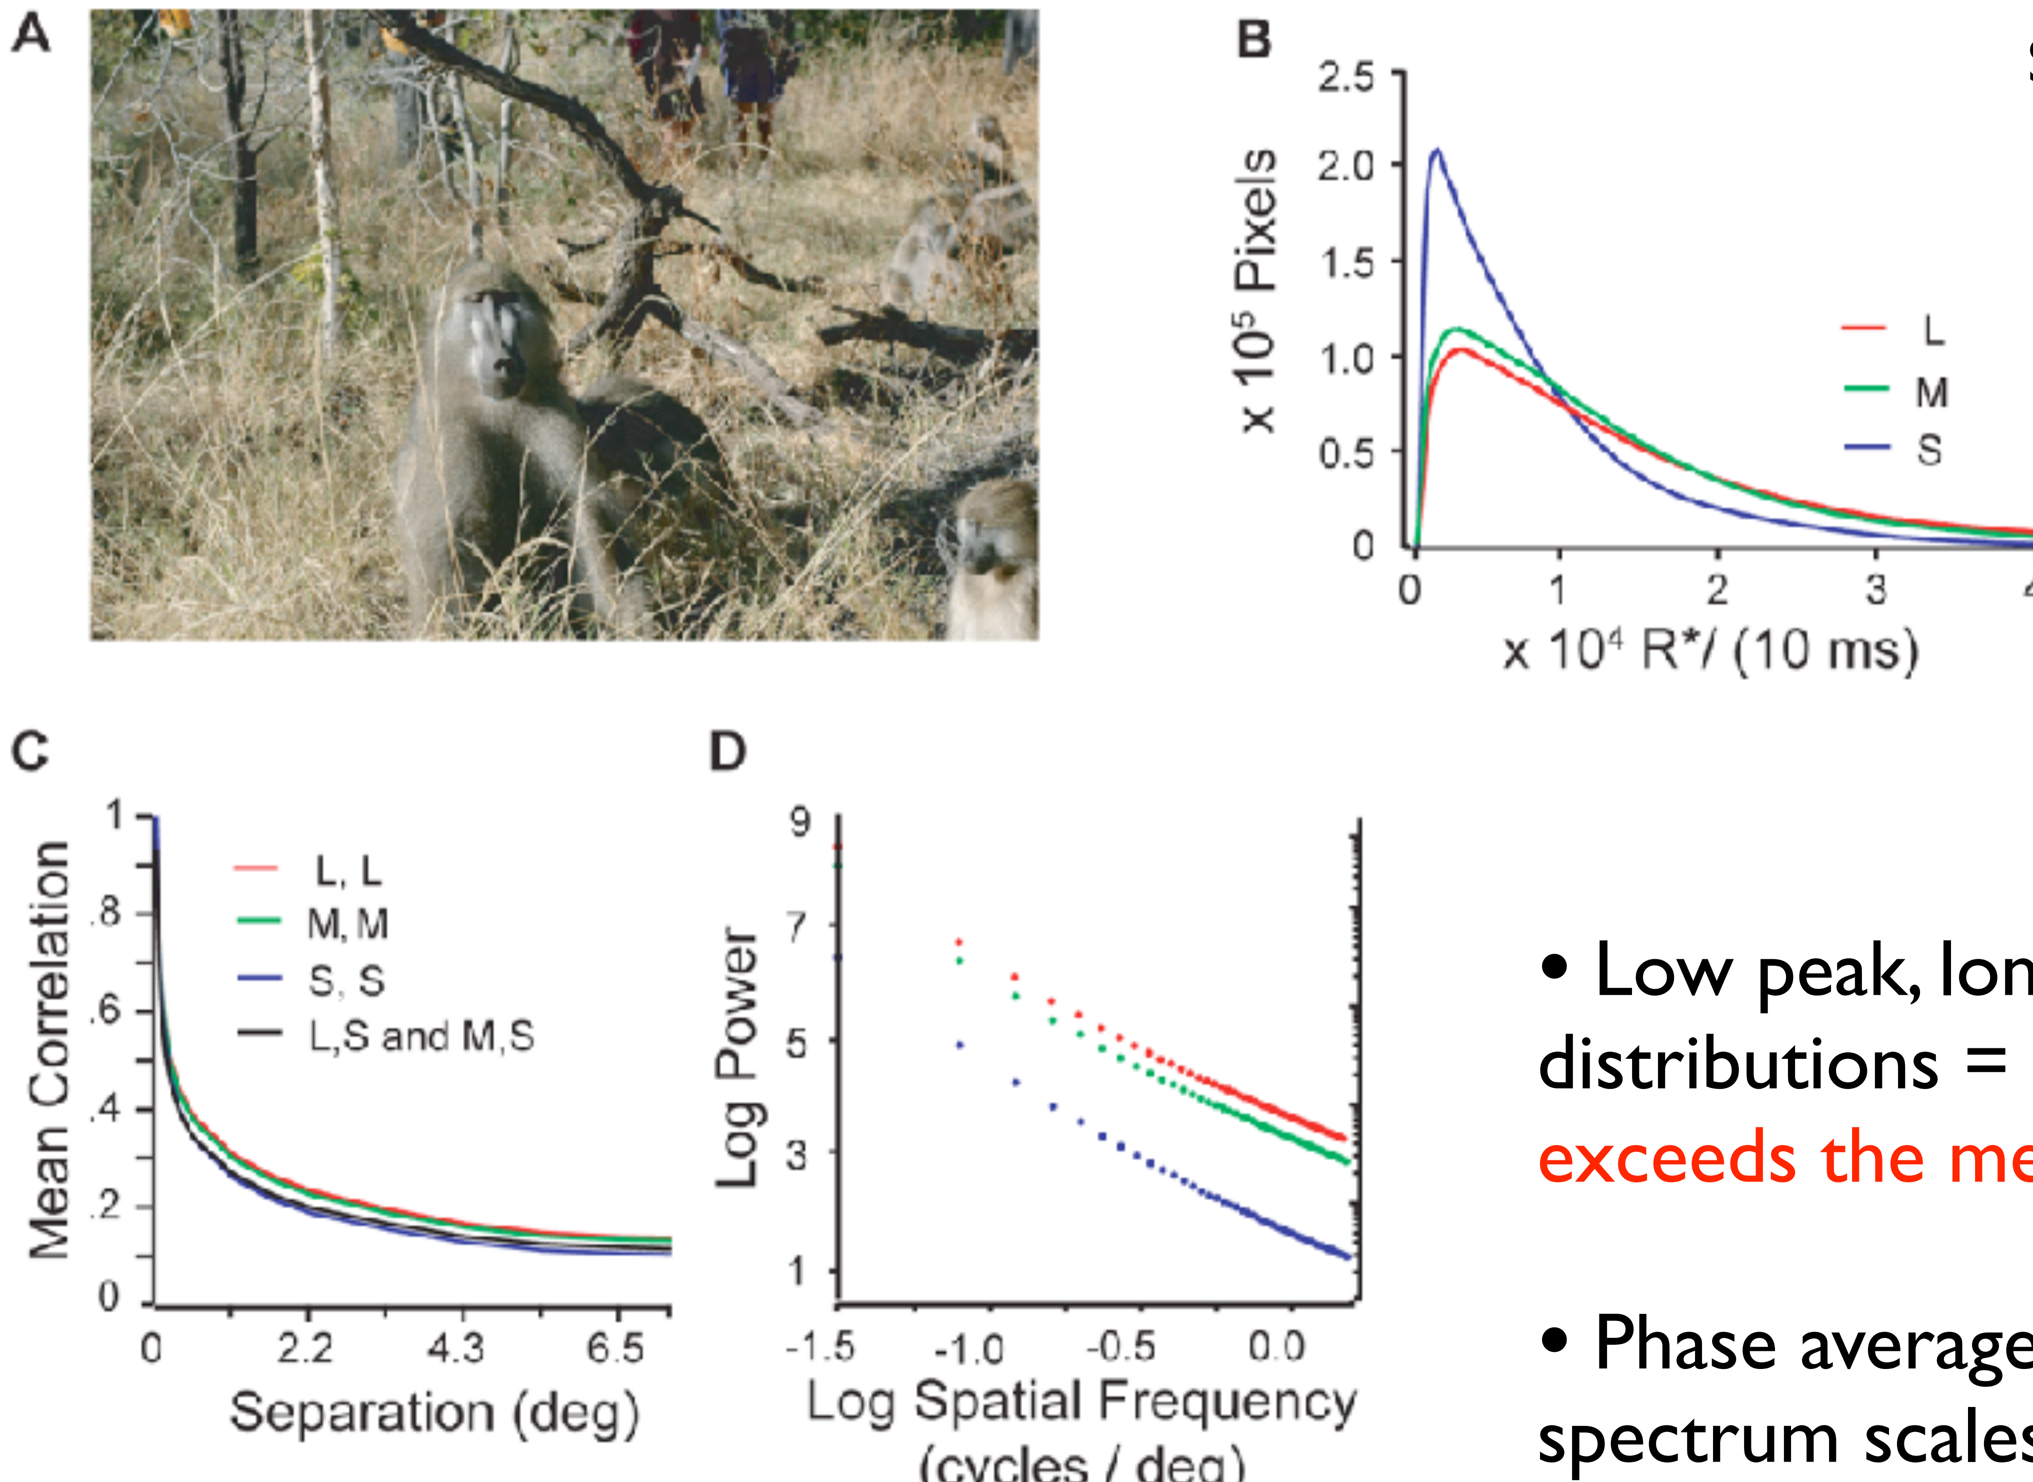

- Low peak, long tail distributions = mean exceeds the median
- Phase averaged power spectrum scales as $\sim 1/k^2$

Scale invariant

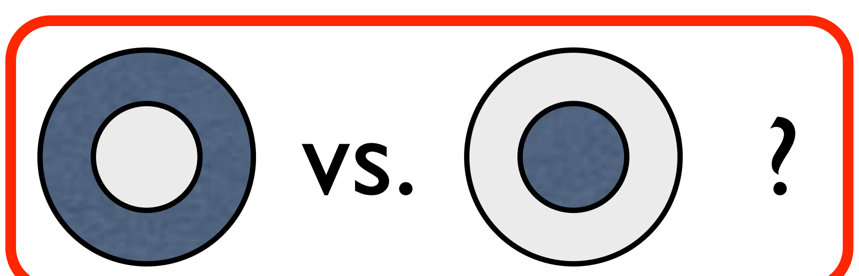

Natural images contain more dark spots

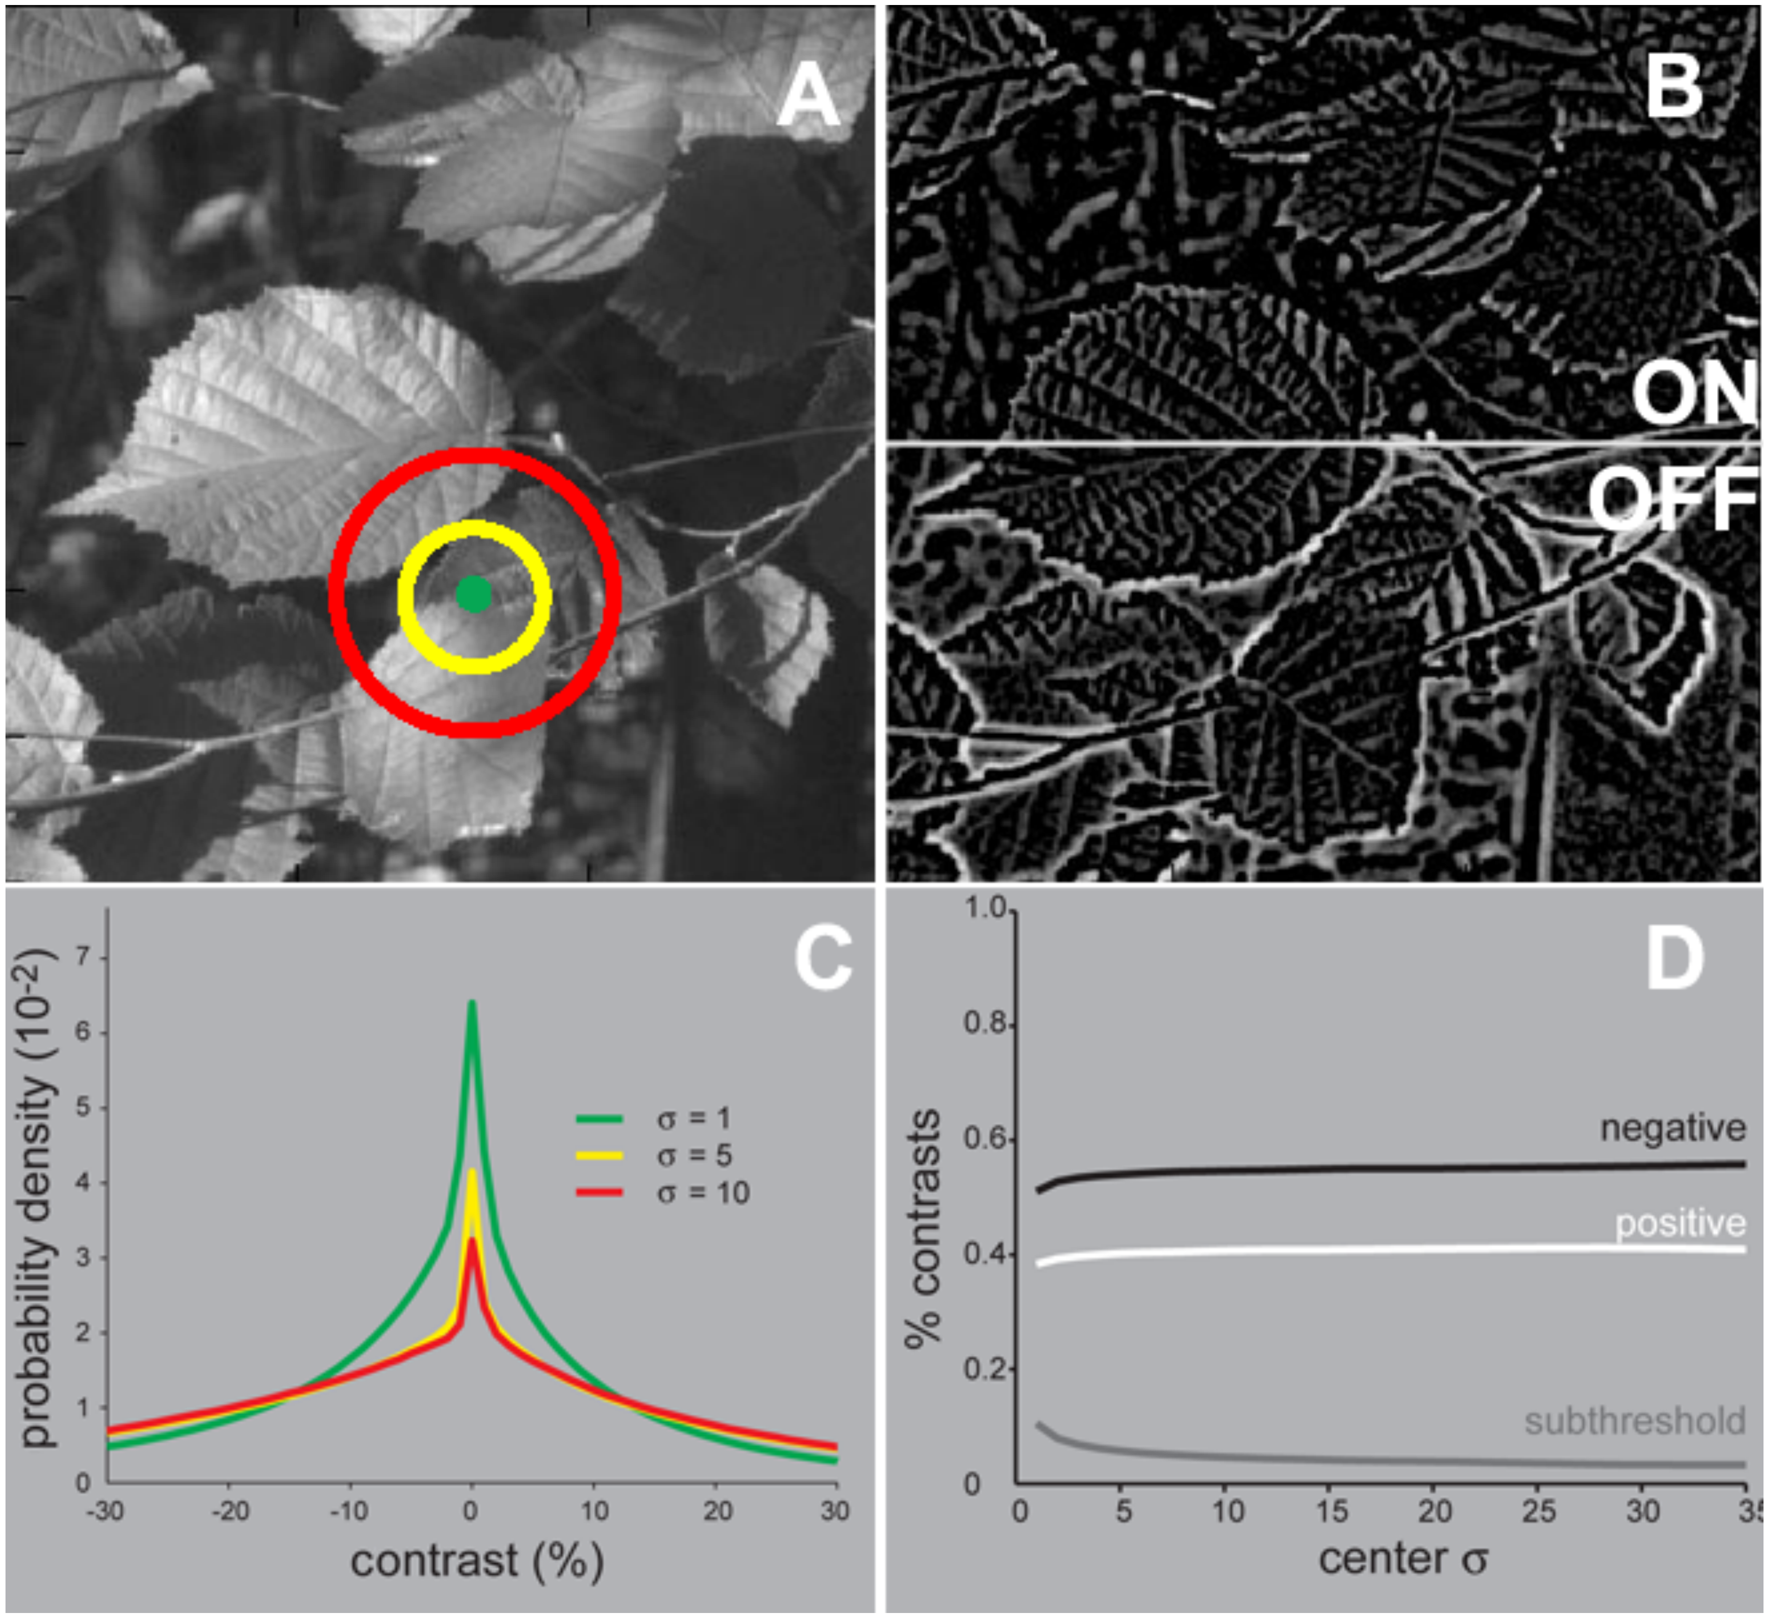

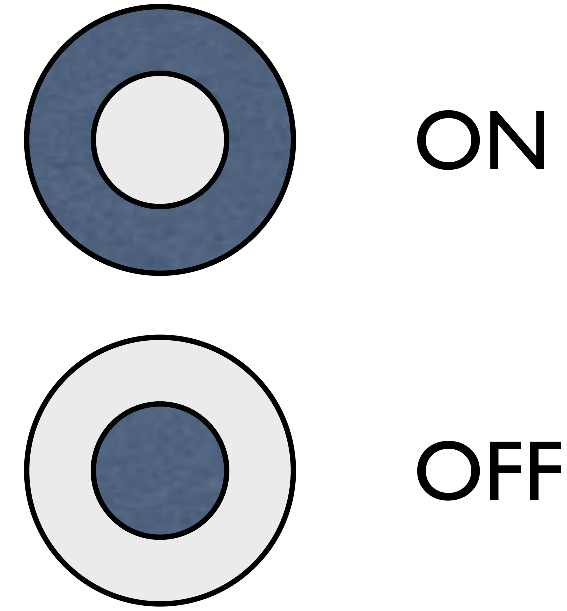

Simple difference of Gaussians model:

$$\frac{I_c(x,y) - I_s(x,y)}{I_s(x,y)}$$

What should be the relative proportion of ON and OFF cells?

How to design the best detector array

Assume that resource constraints require that a particular ON/OFF channel contains N cells. Let $N = N_{OFF} + N_{ON}$.

Given N, find the OFF:ON ratio that maximizes total information.

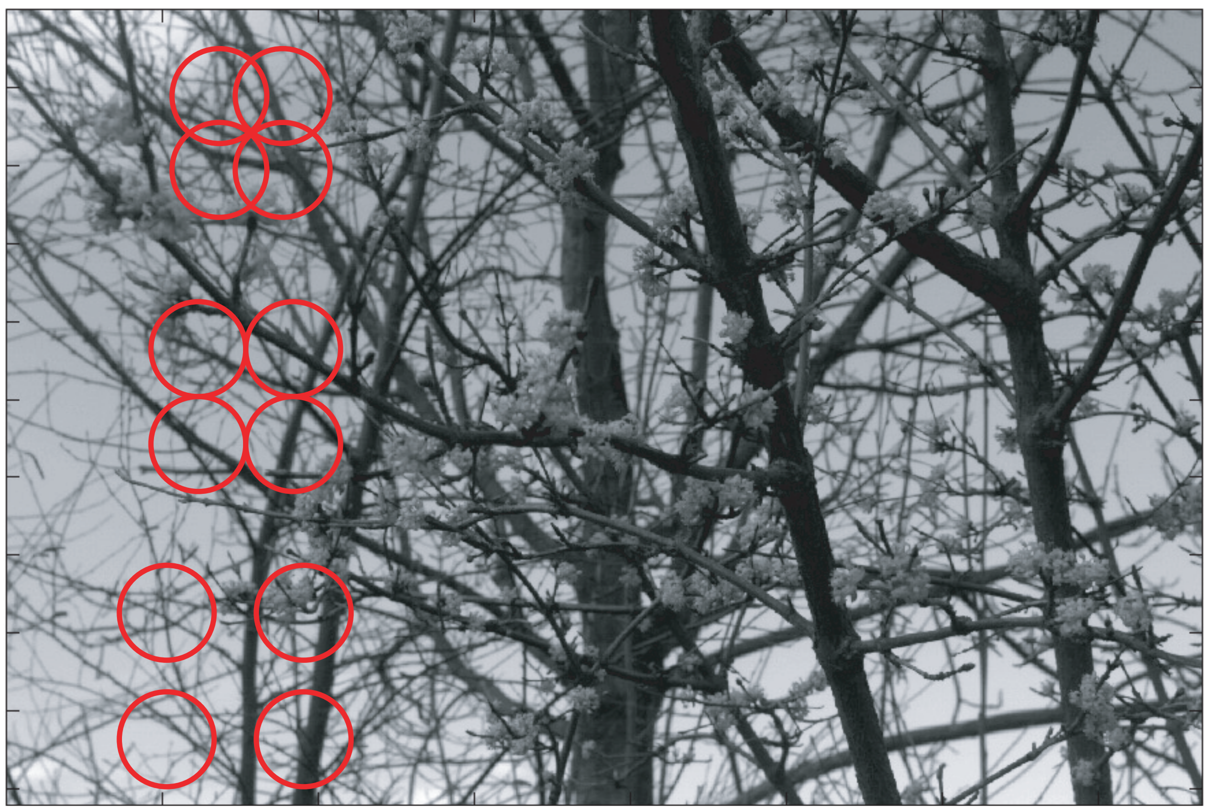

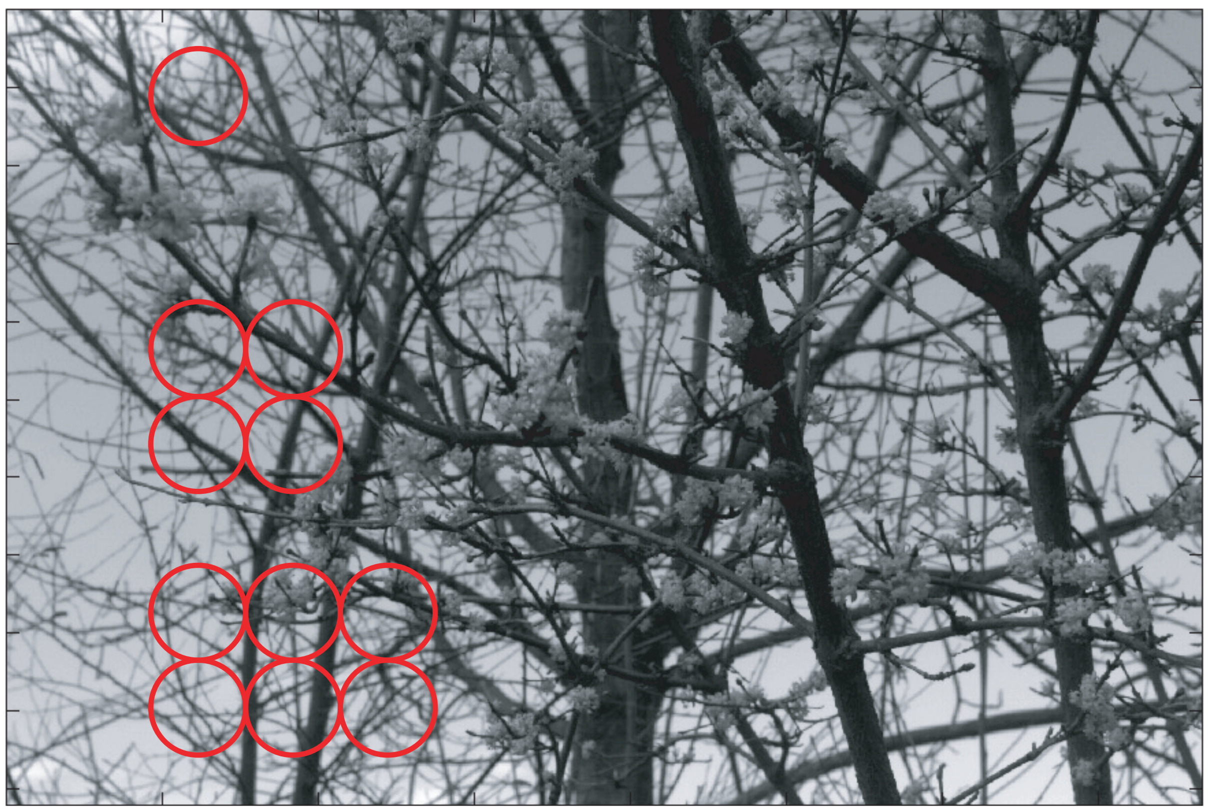

For N=1 the answer is clear: choose an OFF cell -- it is more likely to respond.

Characterizing the optimal mosaic: simplest model

Total information in the array:

$$I = \rho_{ON} N_{ON} I^1_{ON} + (N - N_{ON})\rho_{OFF} I^1_{OFF}$$

$$\boxed{\frac{\partial I}{\partial N_{on}} = 0}$$

ON-ON redundancy, depends only on Receptive Field overlap.
Assume fixed Receptive Field overlap like real cells so that redundancy is constant and equal for ON and OFF types.

Simple SNR + redundancy approximation of each mosaic:

$$I^1_{ON} = \frac{1}{2}\log(1 + f^2_{ON}\,\mathrm{SNR}_{\mathrm{cone}})$$

Receptive field SNR improves with area of receptive field:

$$f^2_{on} = \beta_{on}\, A_{rc} = \beta_{on}\, \frac{A}{N_{on}} \quad \Longrightarrow \quad \frac{\partial I^1_{on}}{N_{on}} \to 0 \text{ for large A}$$

$$\text{Thus } \frac{\partial I}{\partial N_{on}} \quad \Longrightarrow \quad I^1_{on} = I^1_{off}$$

Information equality in the optimal mosaic.

Characterizing the optimal mosaic

Total information in the array:

$$I = \rho_{ON}\, N_{ON}\, I^1_{ON} + \rho_{OFF}\,(N - N_{ON})\, I^1_{OFF} - M$$

ON-ON redundancy

ON-OFF redundancy

$$\boxed{\frac{\partial I}{\partial N_{on}} = 0}$$

Simple model (equally used signaling levels; p_{OFF} = 1 - p_{ON}):

$$I^1_{ON} = -(1 - p_{ON})\log(1 - p_{ON}) - \sum_{i=1}^{l_{ON}} \frac{p_{ON}}{l_{ON}} \log \frac{p_{ON}}{l_{ON}}$$

info. of non-response

Number of signaling levels (SNR) improves with area of receptive field:

$$l_{ON} = \beta_{ON} \left(\frac{A}{N_{ON}}\right)^{1/2}$$

Realistic model: optimal ratio is ~1.7 times as many OFF cells.

Mutual information due to anti-correlation between ON and OFF cells - if an ON cell fails to fire, overlapping OFF cells do fire. Thus the entropy of *non-response* of ON cells is redundant with OFF responses ➡ drop it.

Replace I^1 by: $\tilde{I}^1_{ON} = -\sum_{i=1}^{l_{ON}} \frac{p_{ON}}{l_{ON}} \log \frac{p_{ON}}{l_{ON}}$

The brain separates light from dark unequally

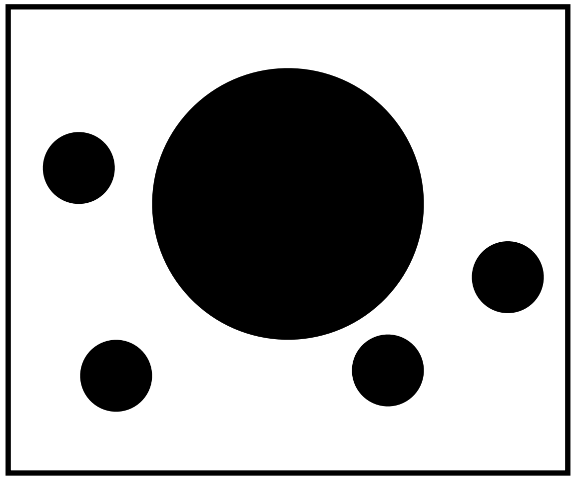
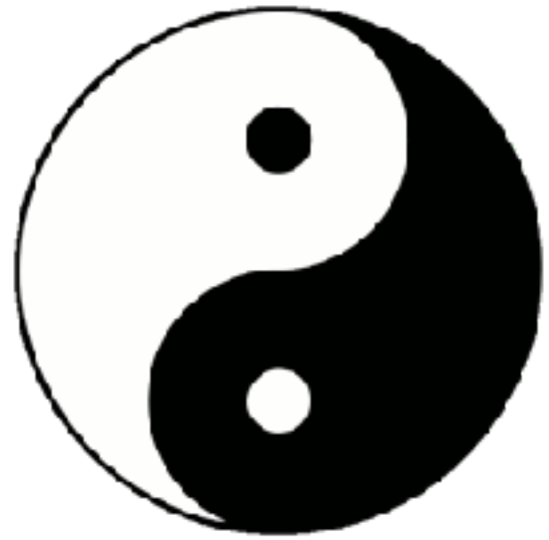
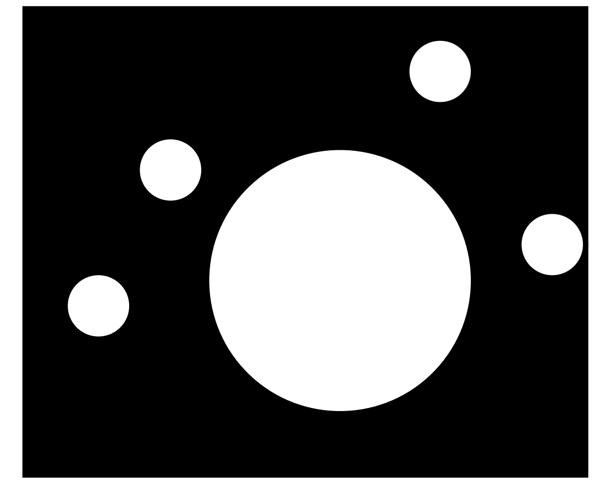

- Behavioral measurements show greater sensitivity to light decrements and dark spots in images (Zemon et al., '88; Chubb et al., 2004)

- More cortical cells respond to negative (dark) than to positive (bright) contrasts (Jin et al., 2008)

- Retinal OFF cells are ~1.3-2 times as numerous as ON. Conserved across types and species: guinea pig (Ratliff et al, 2010), rabbit (de Vries & Baylor, 1997), rat (Morigiwa 1989), monkey (Chichilnisky & Kalmar 2002), human (Dacey and Petersen 1992).

PREDICTED BECAUSE: There are more dark regions in natural scenes and information is more densely packed in them.

Can this approach be applied generally?

Q. What is optimized?
A. Information is an approximation of the "objective" of the early visual system.

Q. Shouldn't the forms and function be determined by evolutionary history?
A. Within the lineage, which is constrained by its history, better adapted forms and functions, are selected over time.

Q. Nothing is ever optimal -- life is a work in progress. Why should anything be optimal?
A. It isn't. But the optimal solution guides us to the principles underlying circuit organization.

There may not be any order in nature, but those of us who look for it have a better chance of finding it if it is there.

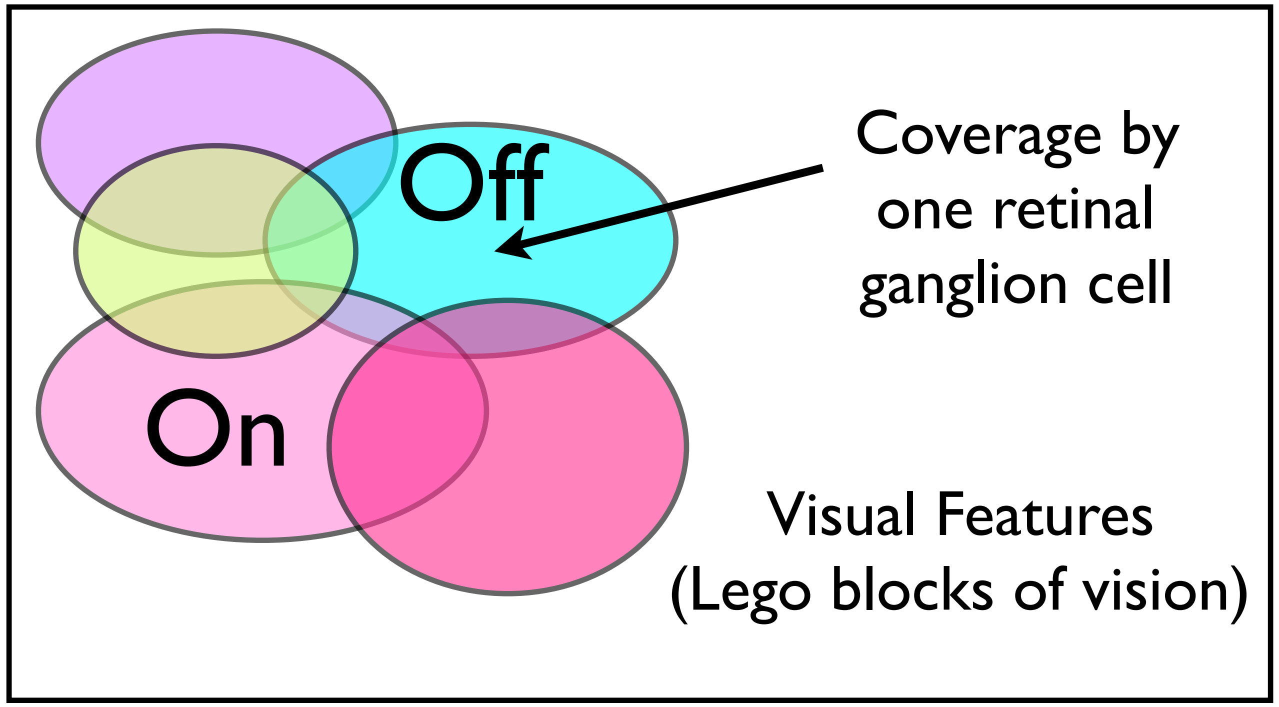

CHALLENGE: Explain the relative proportions of different elements of the visual repertoire in terms of the value they have for vision

Cognitive repertoires: The sense of place

Xuexin Wei, Jason Prentice, Vijay Balasubramanian

The sense of place

What is place?   How do you know where you are?

Inside your head, "here" is an abstract pattern of neurons firing. These patterns maintain a map of your location.

A simple representation of one dimensional space

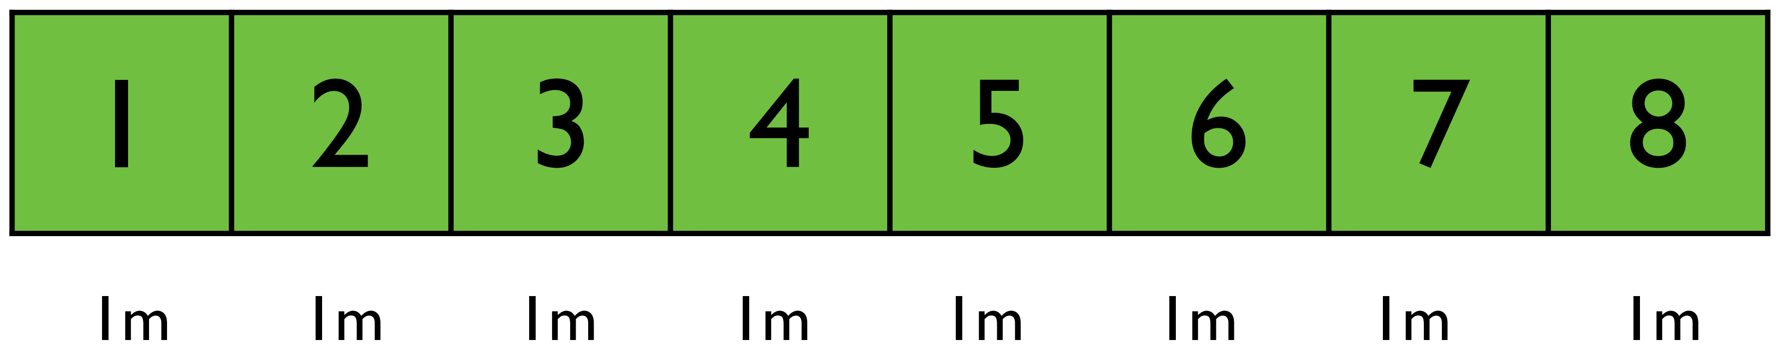

8m linear track

To achieve 1m resolution on an 8m track can have 8 “place neurons”, each of which fire when you are in a particular 1m wide location. **This requires 8 neurons.**

A more efficient representation: binary numbers

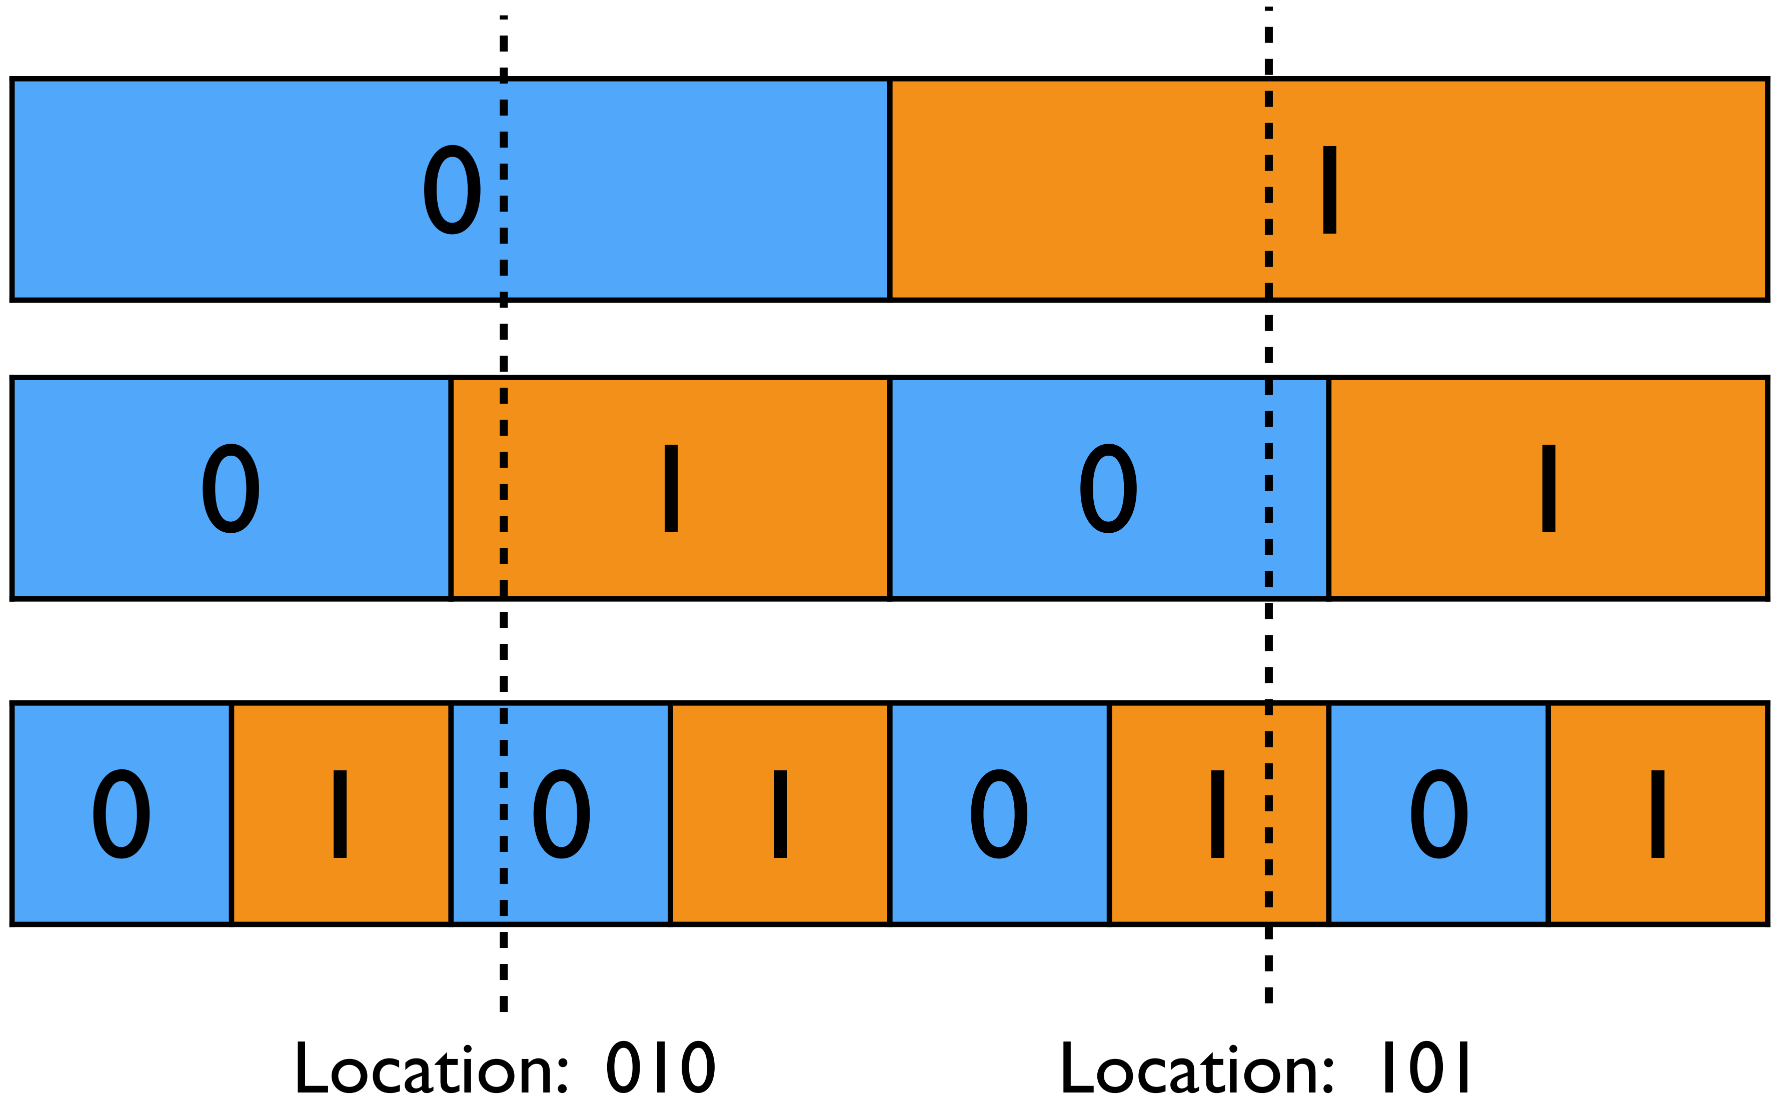

This is a "binary" representation of space (i.e. a *base 2* number system) and requires only 6 neurons — it is more efficient.

You can use other bases, like decimal, i.e. base 10.

A simple representation of two dimensional space

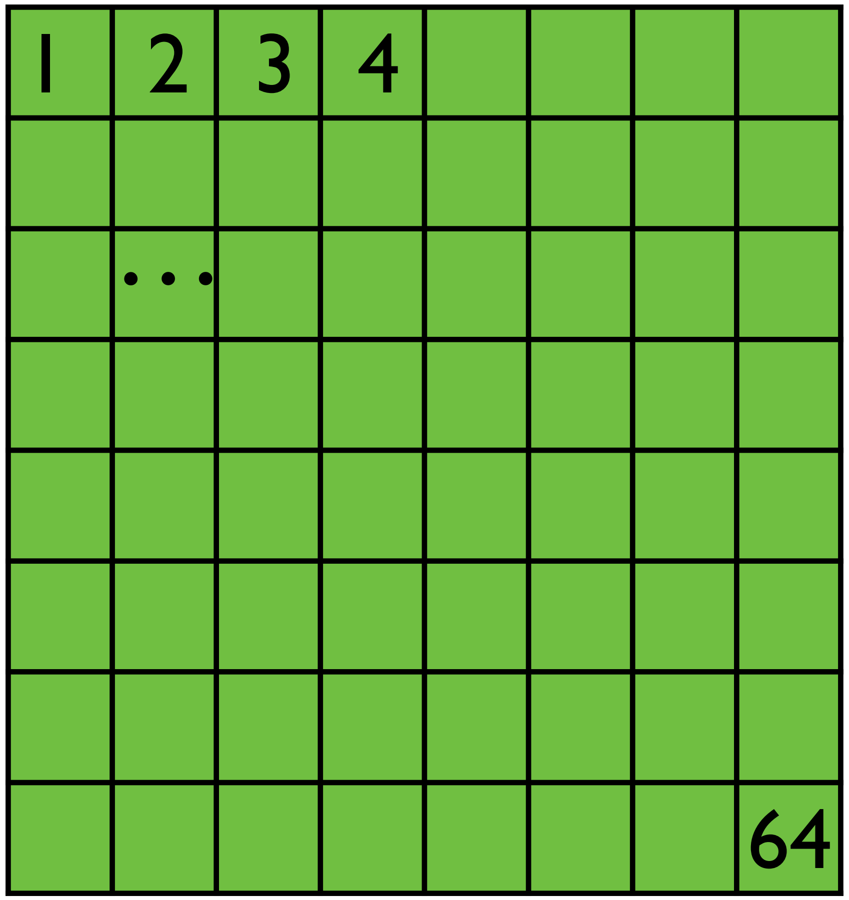

- Square: 8m on each side
- Resolution: 1m on each side
- Need: 64 neurons

In two dimensions you could imagine different neurons responding when the animal is in different locations

A two dimensional analog of binary numbers

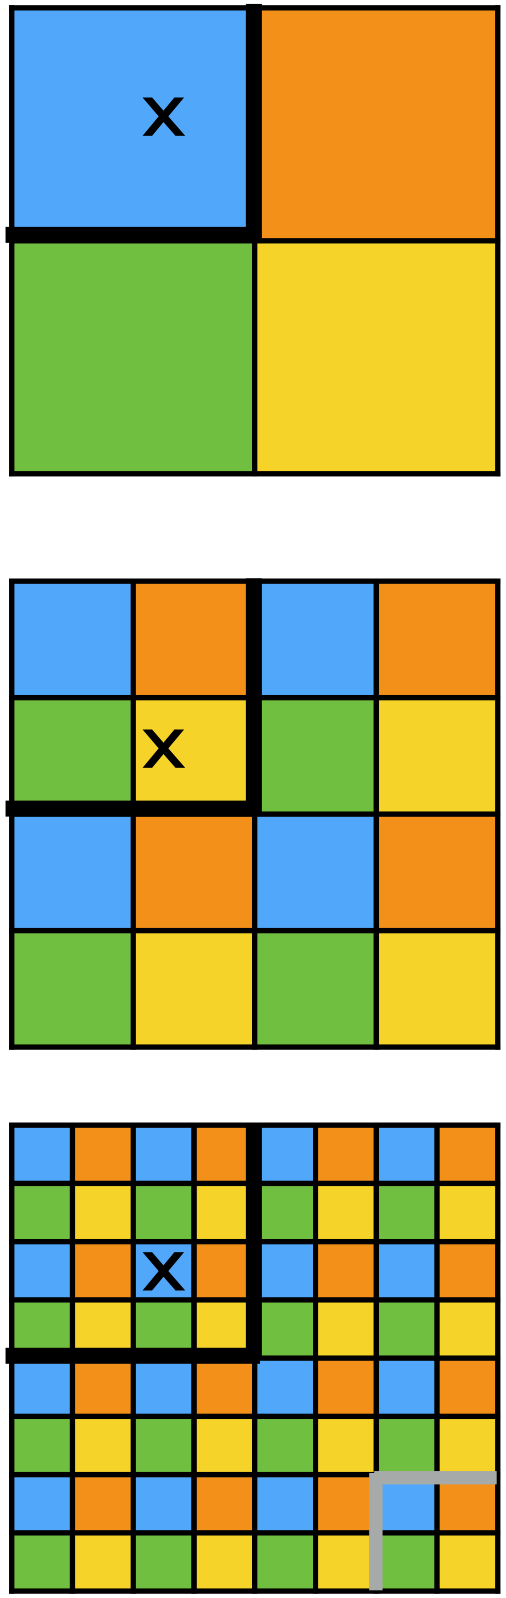

- Square: 8m on each side
- Resolution: 1m on each side
- Need: 12 neurons

You can use other bases, like decimal, i.e. base 10.

Grid Cells: a numbering system for location?

- Grid cells in the ***entorhinal cortex*** respond when an animal is physically in locations lying on a triangular lattice

- The grids increase in size along the axis of the entorhinal cortex

- Different cells have randomly varying offsets (phases) for their grids

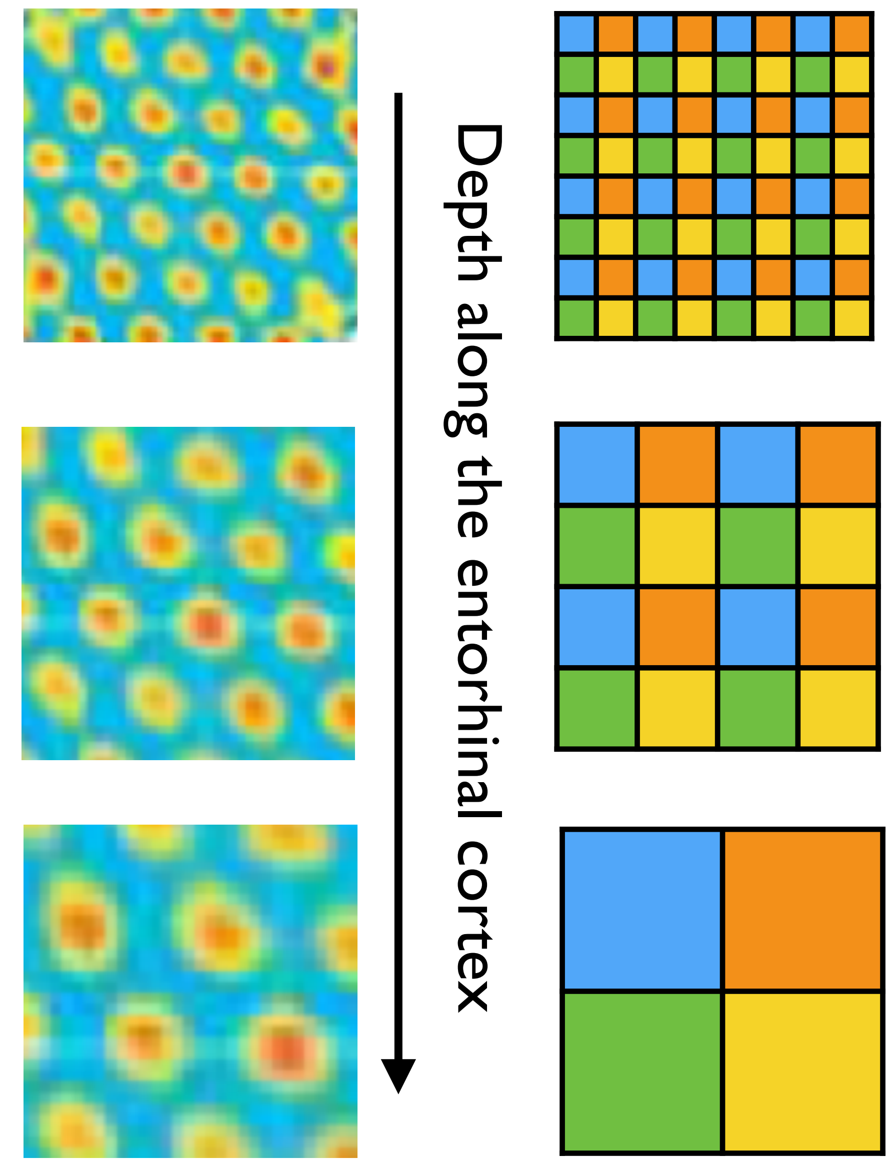

Hafting et al Nature 2005

➢ Like a two-dimensional, fuzzy, neural number system

Which number system (binary, decimal, etc.) should the brain pick to represent space?

The ratio of adjacent grid sizes is the "base" of the number system (e.g. binary or base 2; decimal or base 10)

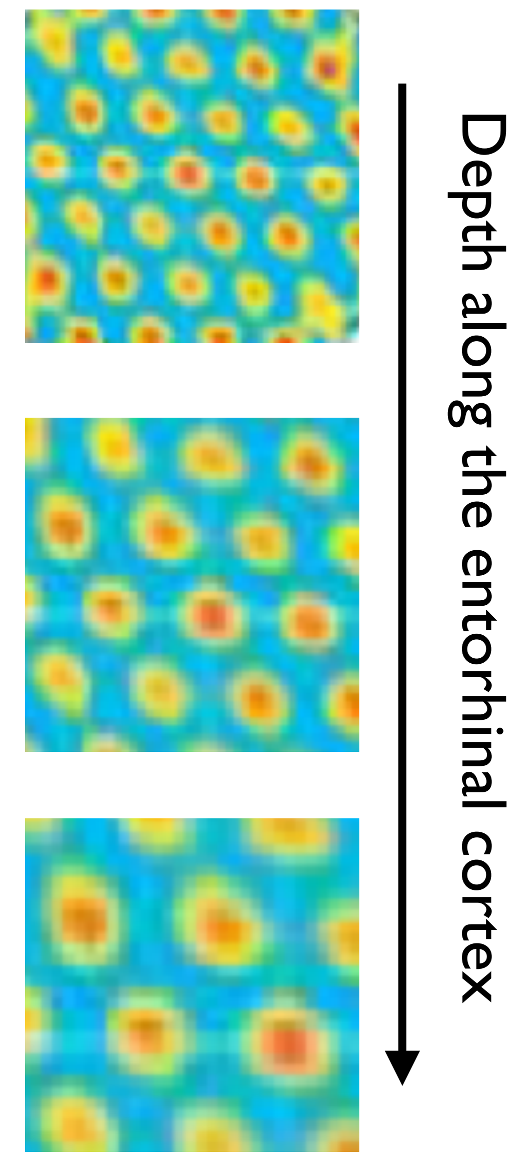

Spatial characteristics of the grid system

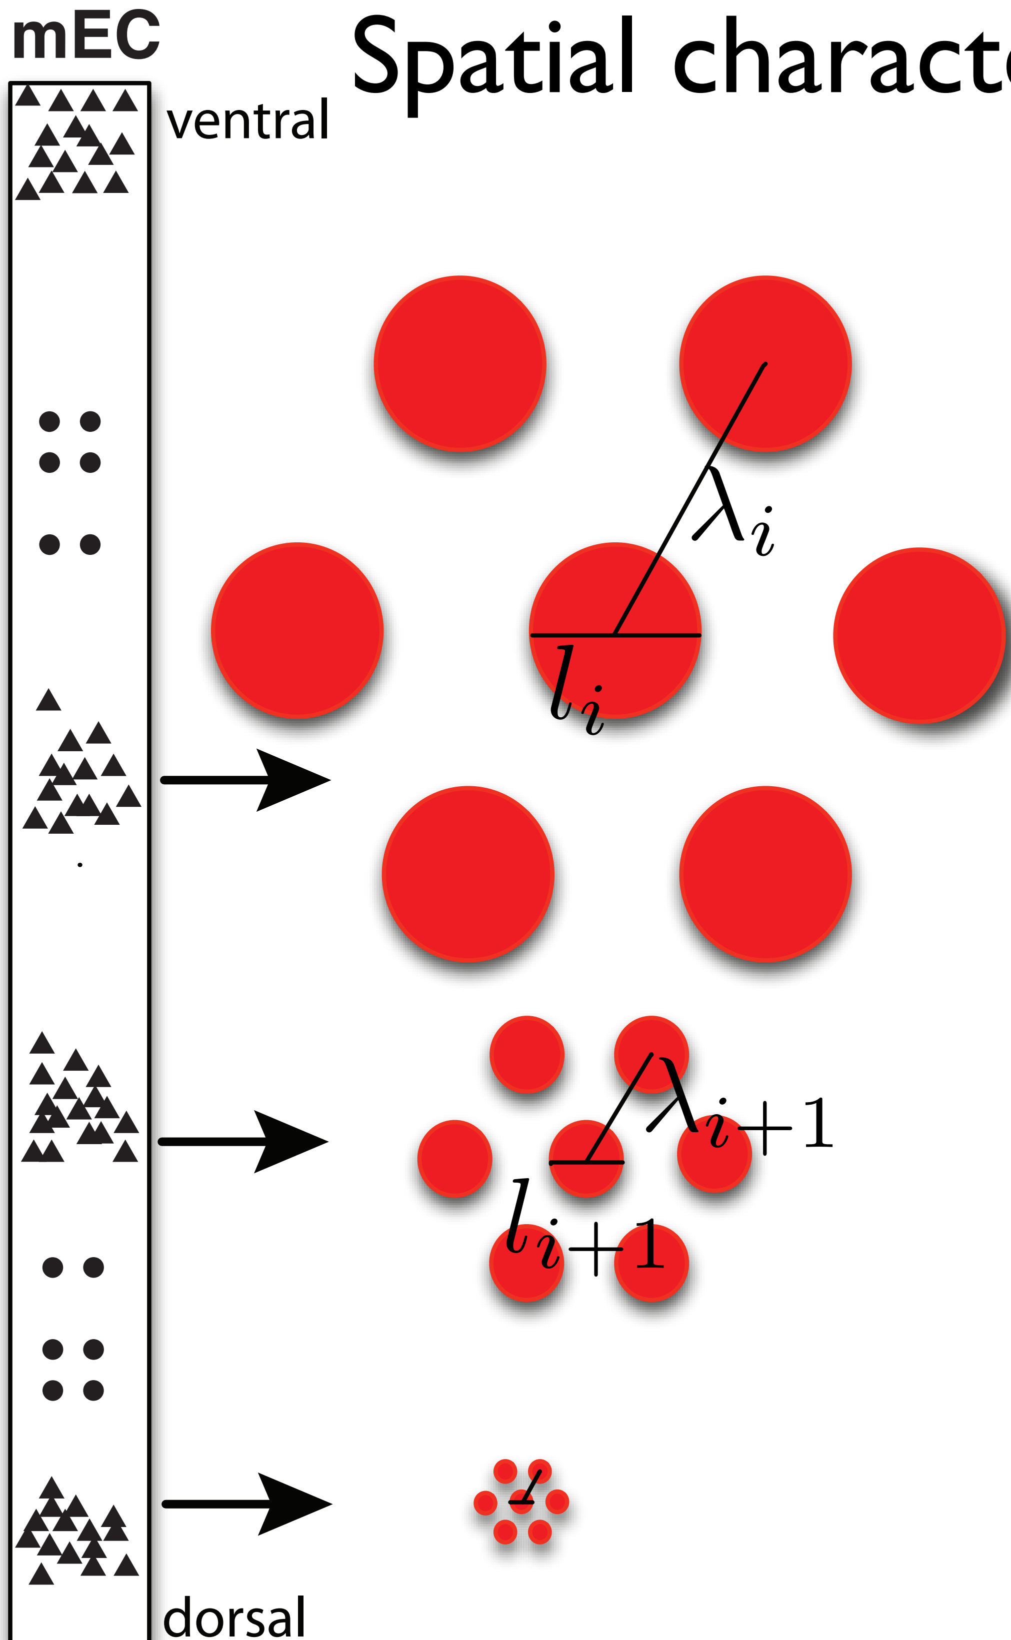

λ_i = Scale of grid (assume largest scale matched to environment)

l_i = Diameter of single grid field

<u>Assume</u>: Uniform spatial phases, and constant coverage at each scale

$r_i = \frac{\lambda_i}{\lambda_{i+1}}$ = Ratio of grid scales

$R = \frac{\lambda_1}{\lambda_n} = \prod_i r_i$ = Resolution

$N \propto \sum_i \frac{\lambda_i}{l_i}$ = Number of grid cells

Cost of grid = f(N)

formulae written for 1d grids

What ratio between scales minimizes the number of cells required achieve a given spatial resolution?

One dimensional grids

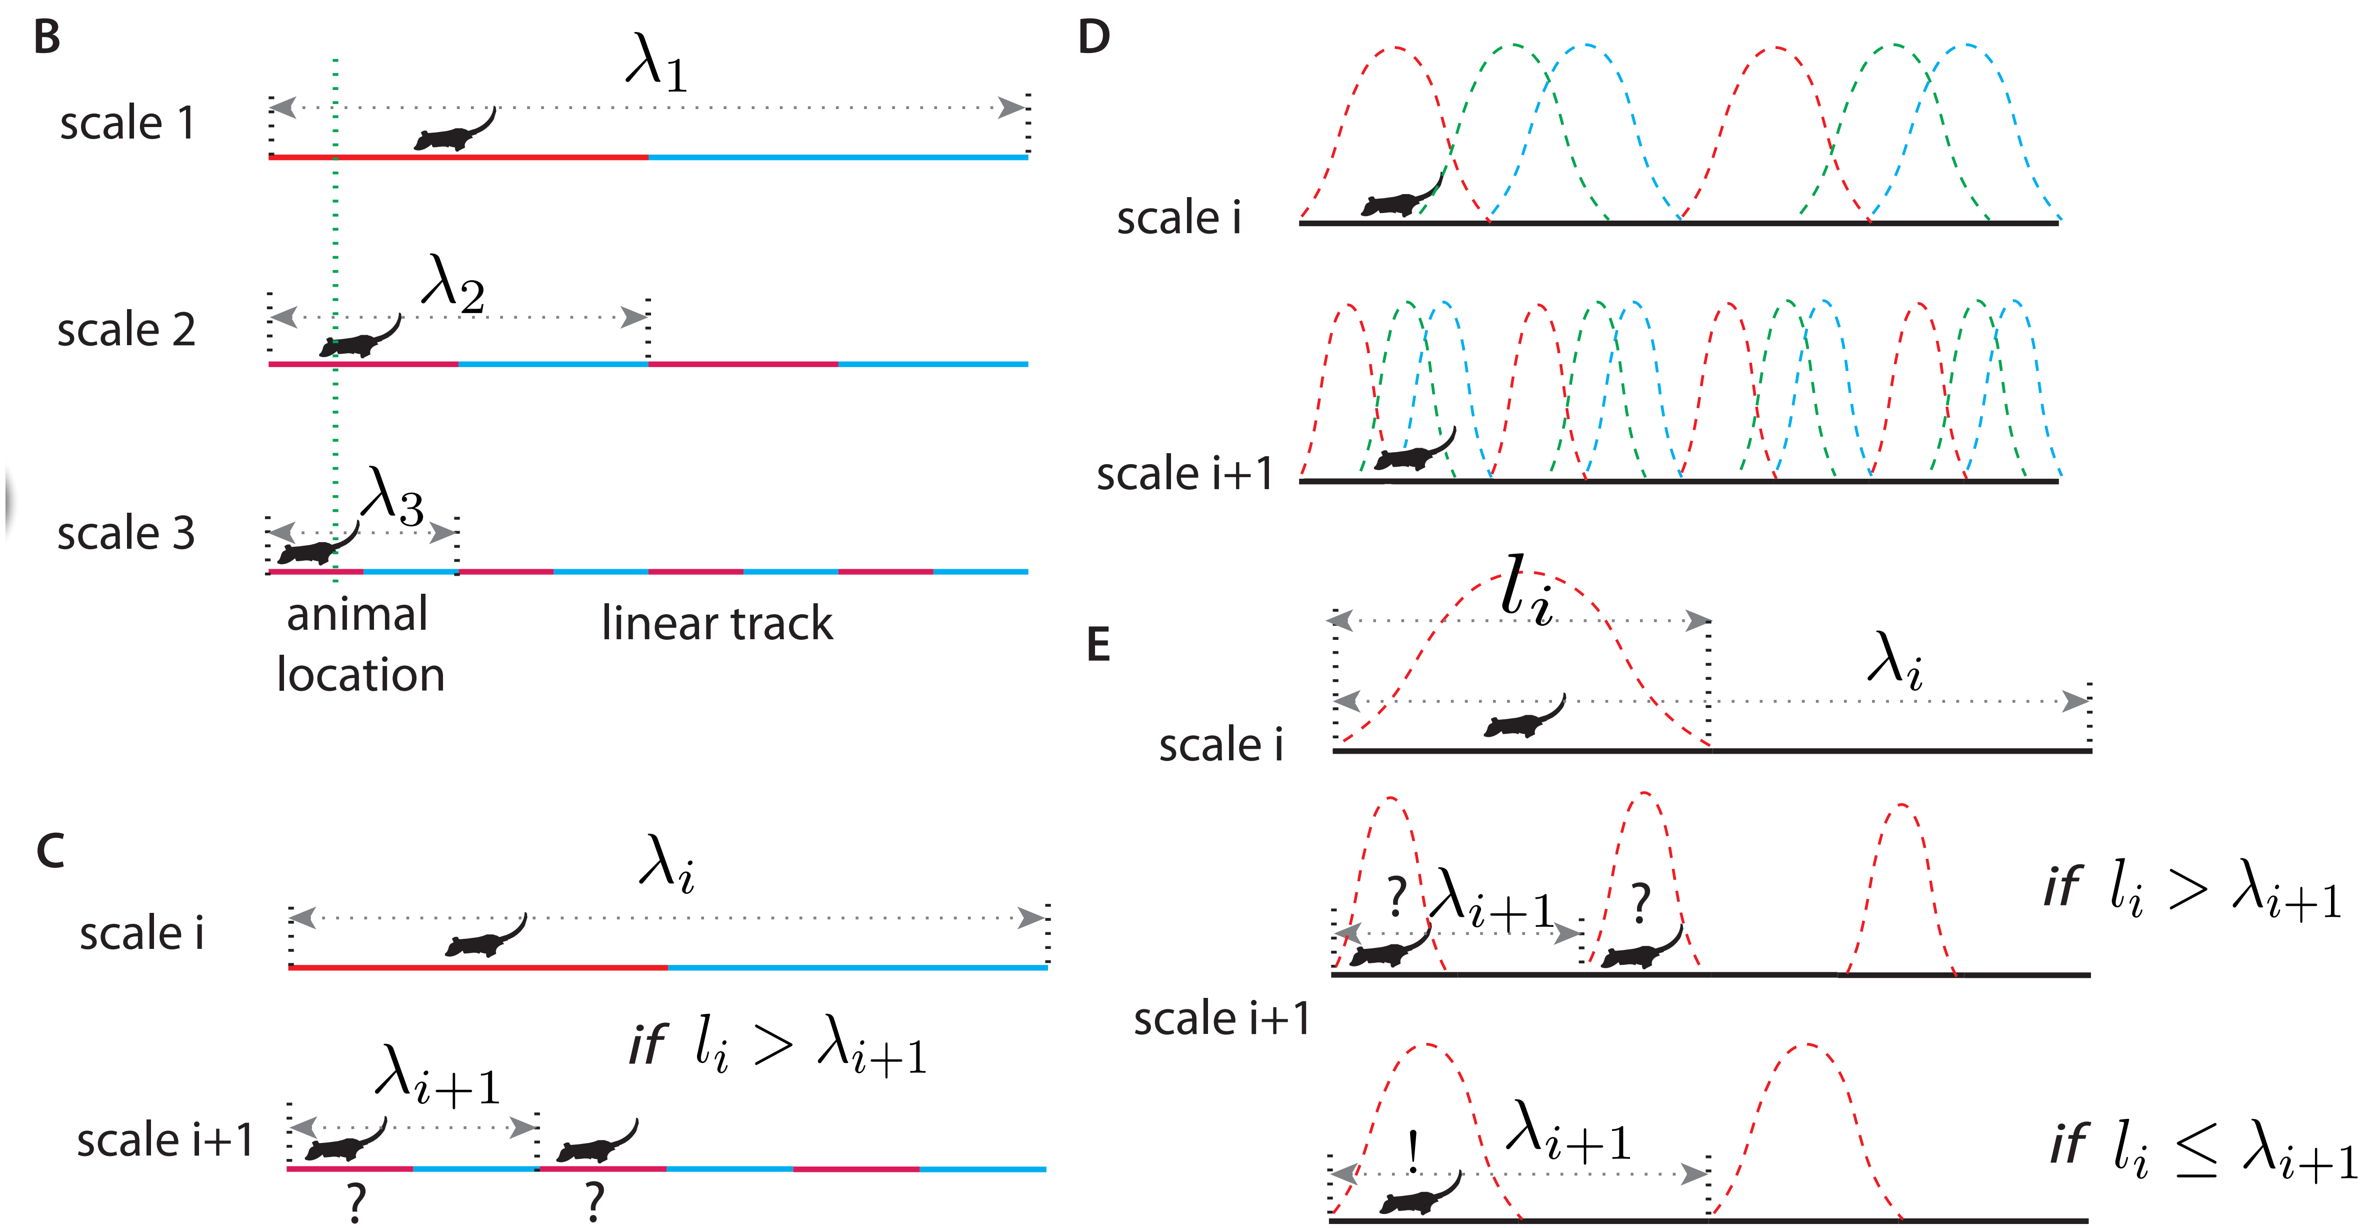

Ambiguities arise if the grid field width is too large compared to the next scale

Optimizing the grid: a simple model

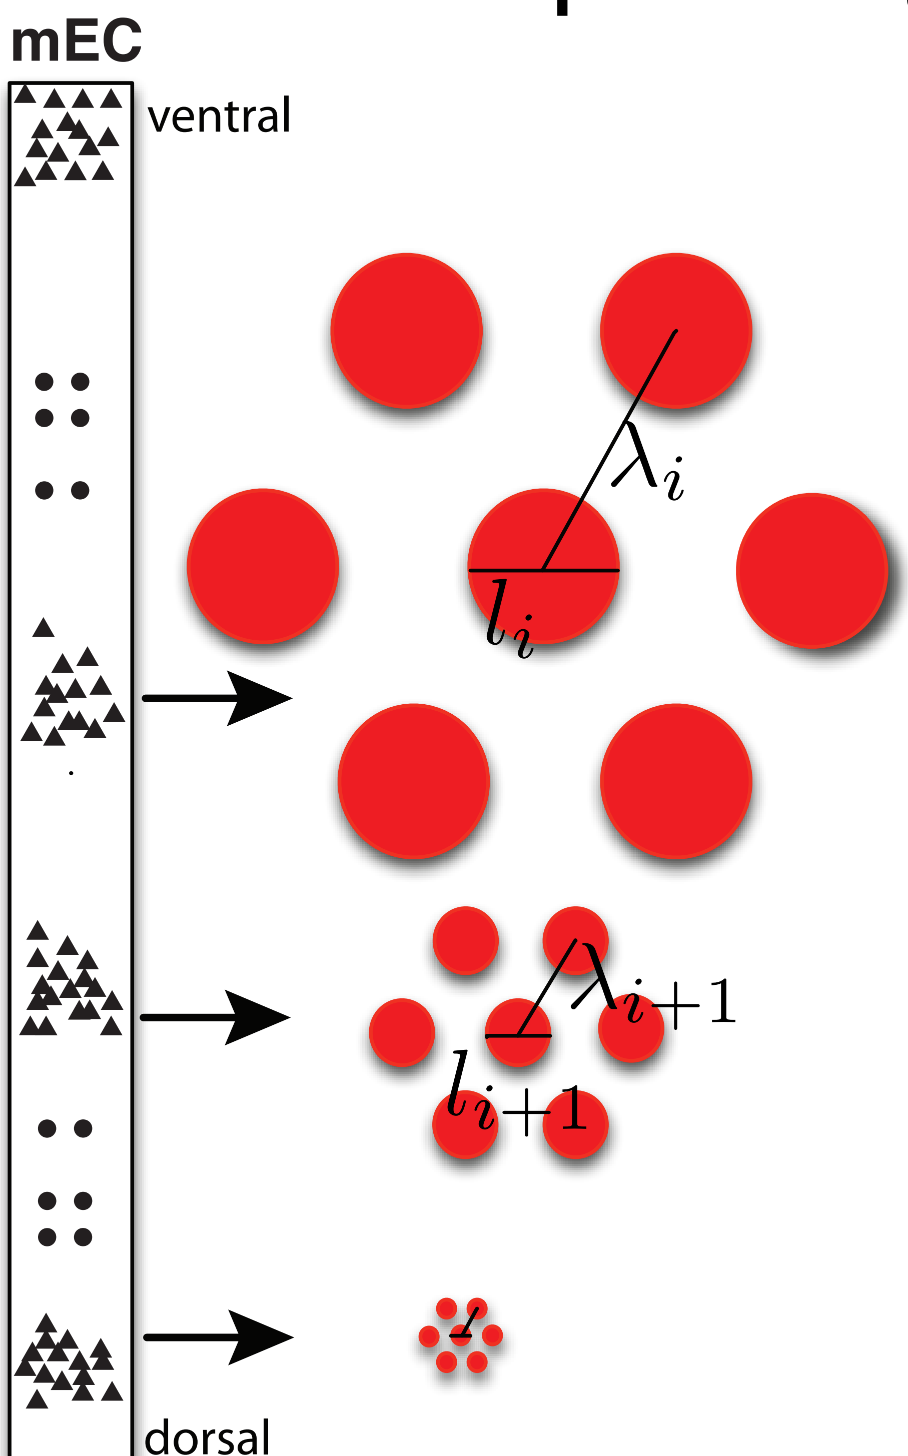

$$r_i = \frac{\lambda_i}{\lambda_{i+1}}$$ = Ratio of grid scales

$$R = \frac{\lambda_1}{\lambda_n} = \prod_i r_i$$ = Resolution

$$N \propto \sum_i \frac{\lambda_i}{l_i}$$ = Number of grid cells

formulae written for 1d grids

$$\lambda_1 > \lambda_2 > \cdots \lambda_m$$

Minimize number of cells (N) for fixed resolution (R)

Unambiguous decoding:

$$\lambda_{i+1} = \frac{\lambda_i}{r_i} \geq l_i \implies \frac{\lambda_i}{l_i} \geq r_i$$

Optimizing the grid

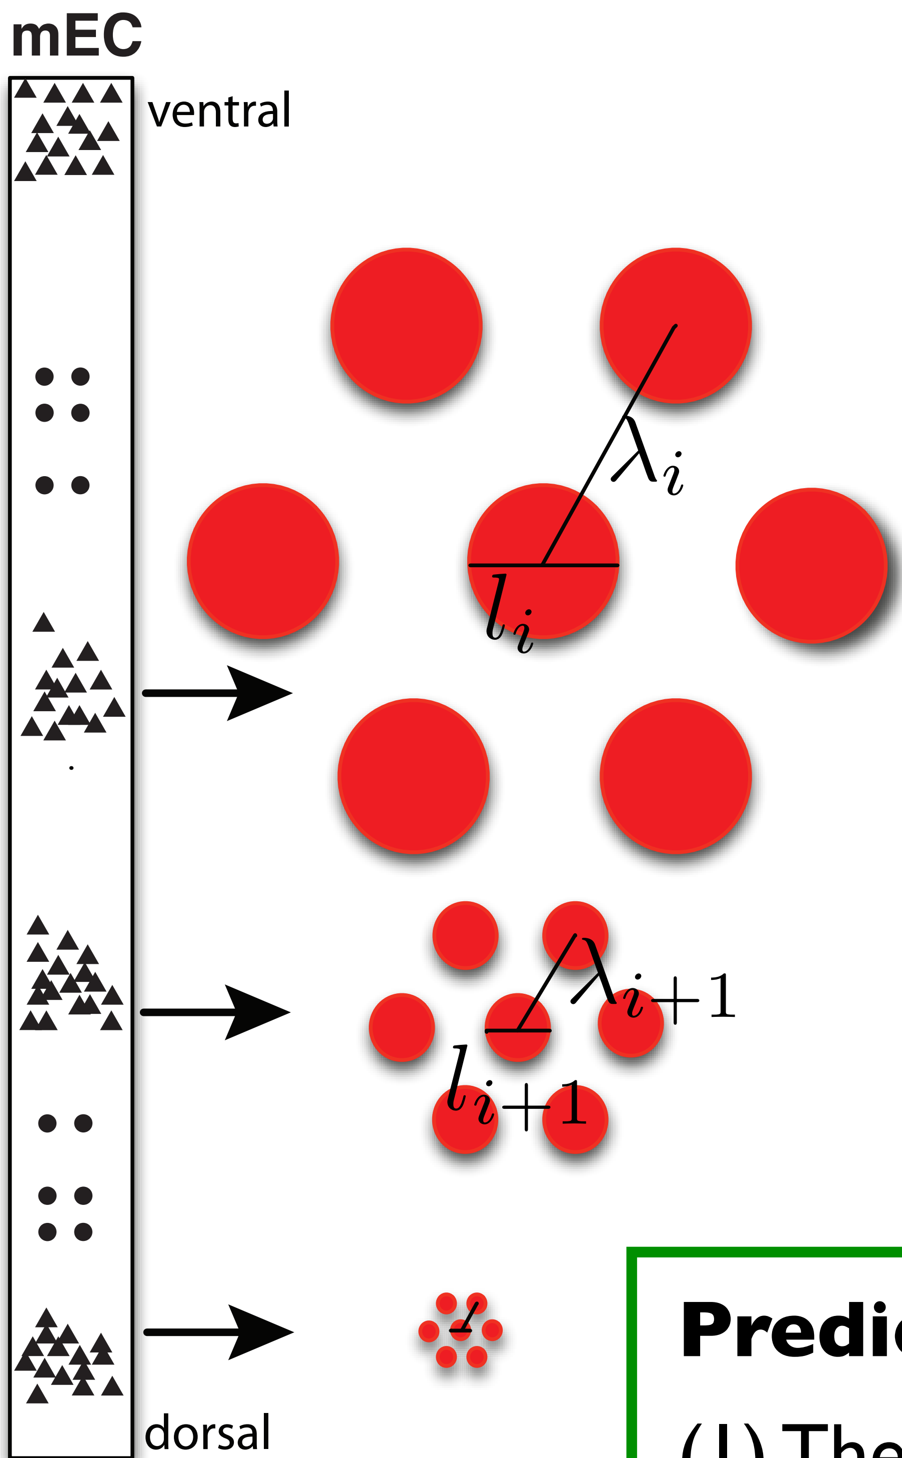

$r_i = \dfrac{\lambda_i}{\lambda_{i+1}}$ = Ratio of grid scales

$R = \dfrac{\lambda_1}{\lambda_n} = \prod_i r_i$ = Resolution

$N \propto \sum_i r_i$ = Number of grid cells

(formulae written for 1d grids)

Minimize number of cells (N) for fixed resolution (R)

Predictions:

(1) The ratios of adjacent periods will be equal

(2) The constant ratio is $r = (e)^{1/d}$ in d dimensions.

(3) $\lambda_i / l_i = r$

Optimizing the grid: probabilistic decoding

- Asymptotically, the posterior distribution over position of each module may be approximated by a periodic series of Gaussian bumps.

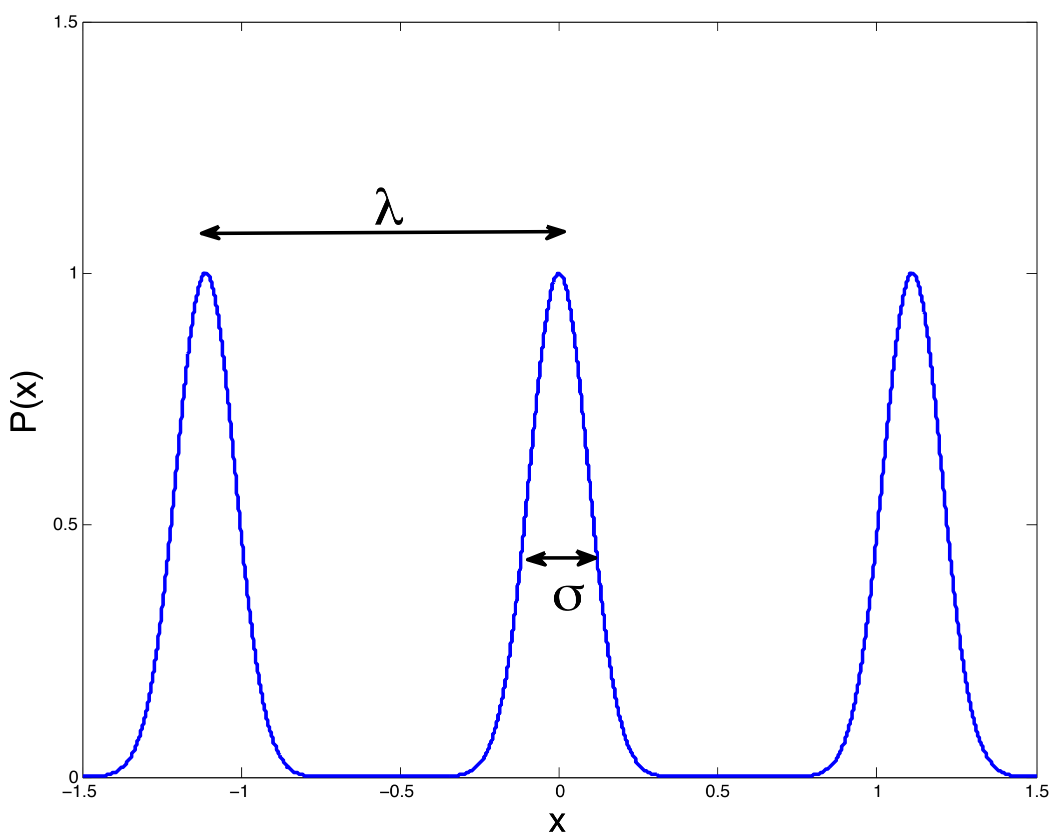

- Combine information from modules (scales) by multiplying the posterior distributions.

Resolution vs. Ambiguity

Shrinking the scale improves resolution

$$R \sim \prod_i \frac{\delta_{i-1}}{\delta_i}$$

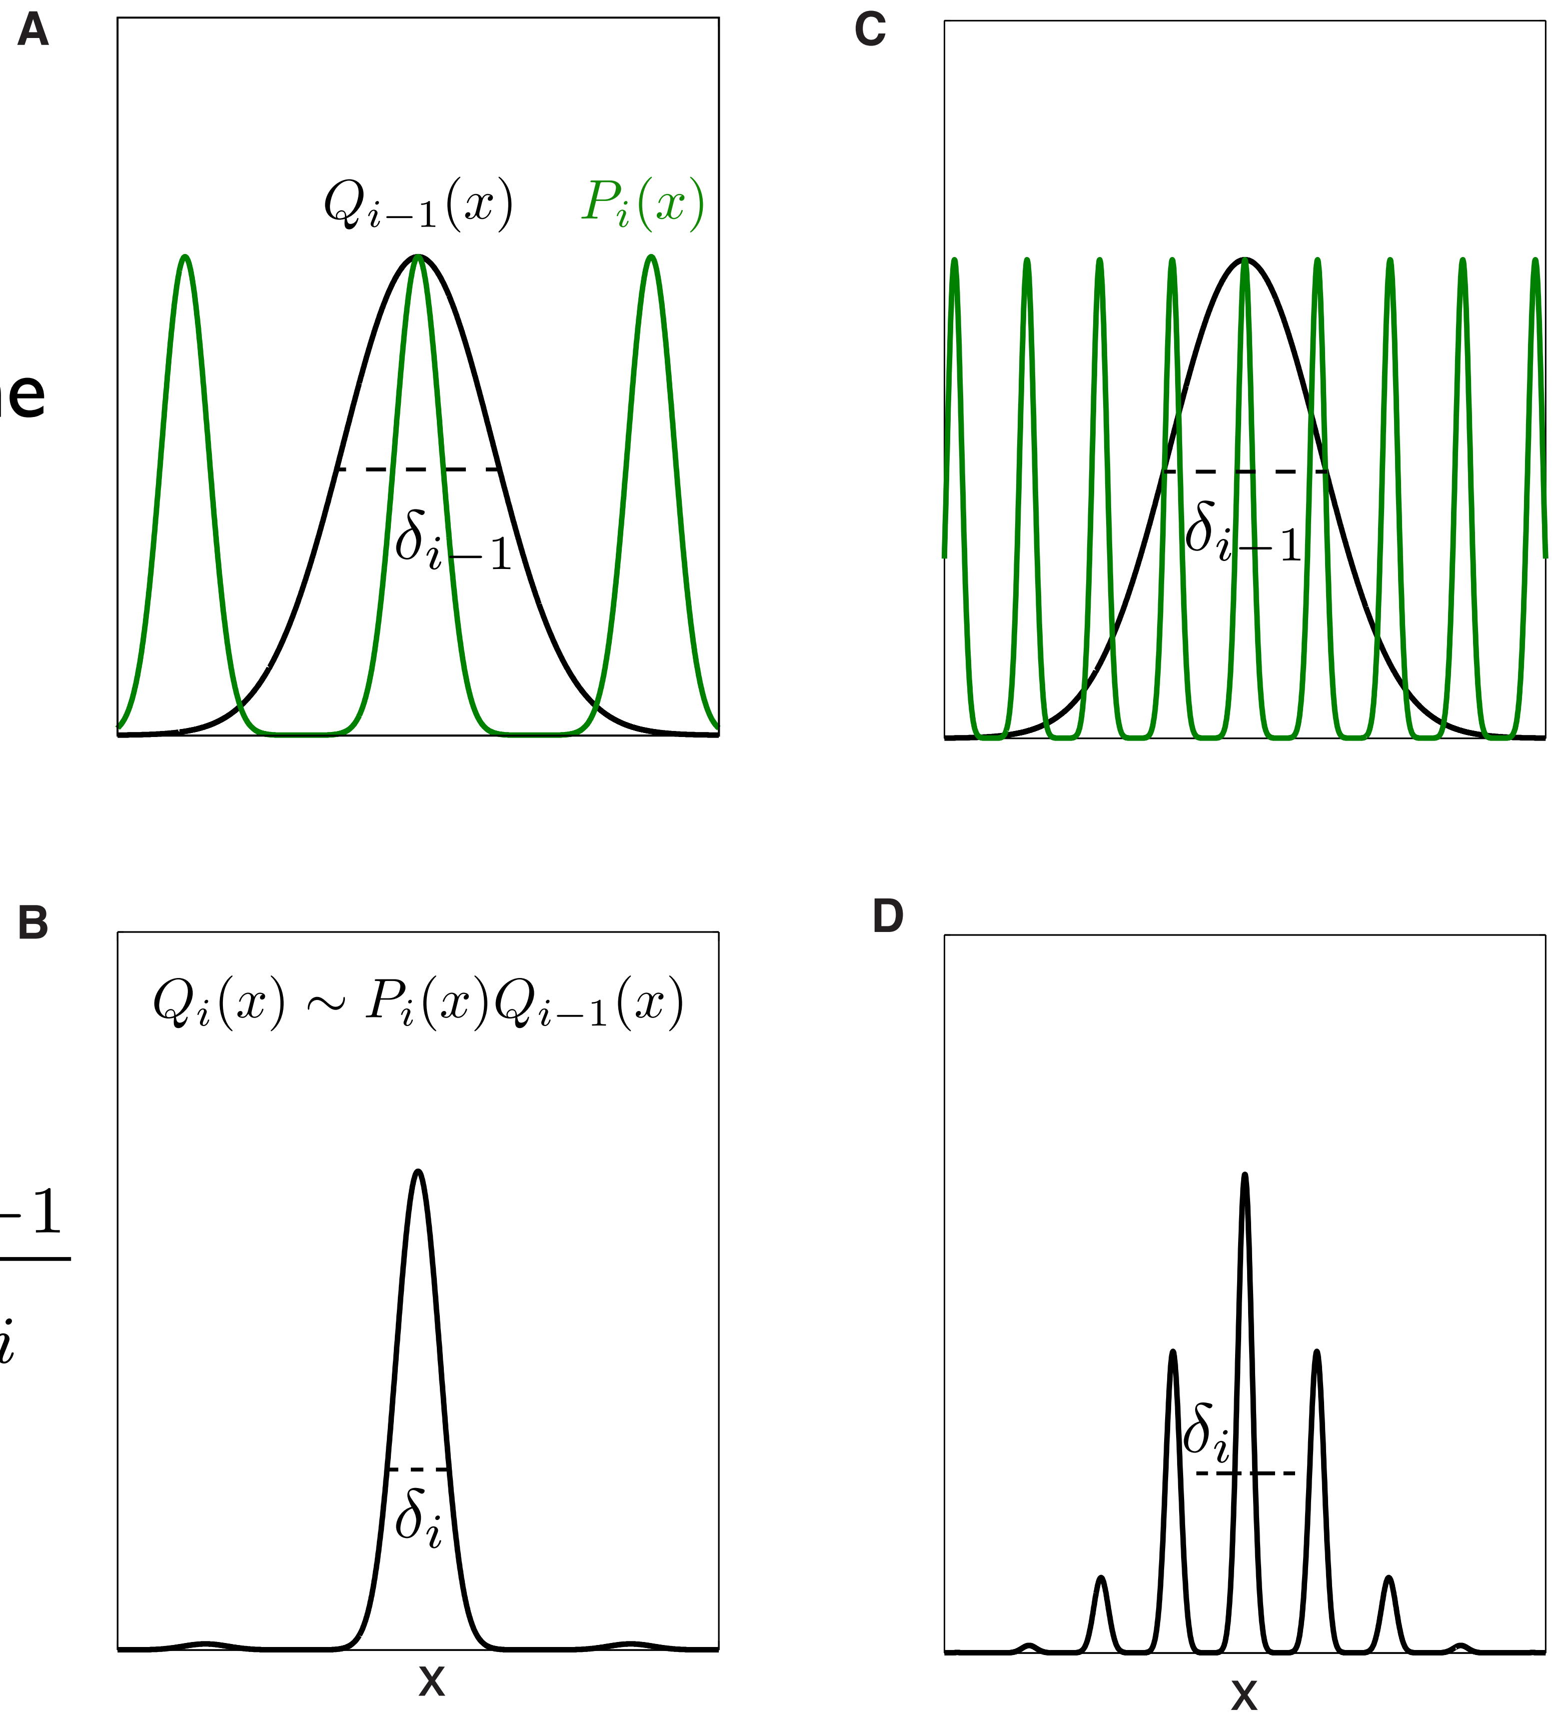

Ambiguities arise if the scales shrink too quickly

Two dimensional grids

Firing maps are doubly-periodic, with period vectors $\lambda_i \mathbf{u}$ and $\lambda_i \mathbf{v}$ $(|\mathbf{u}| = 1)$

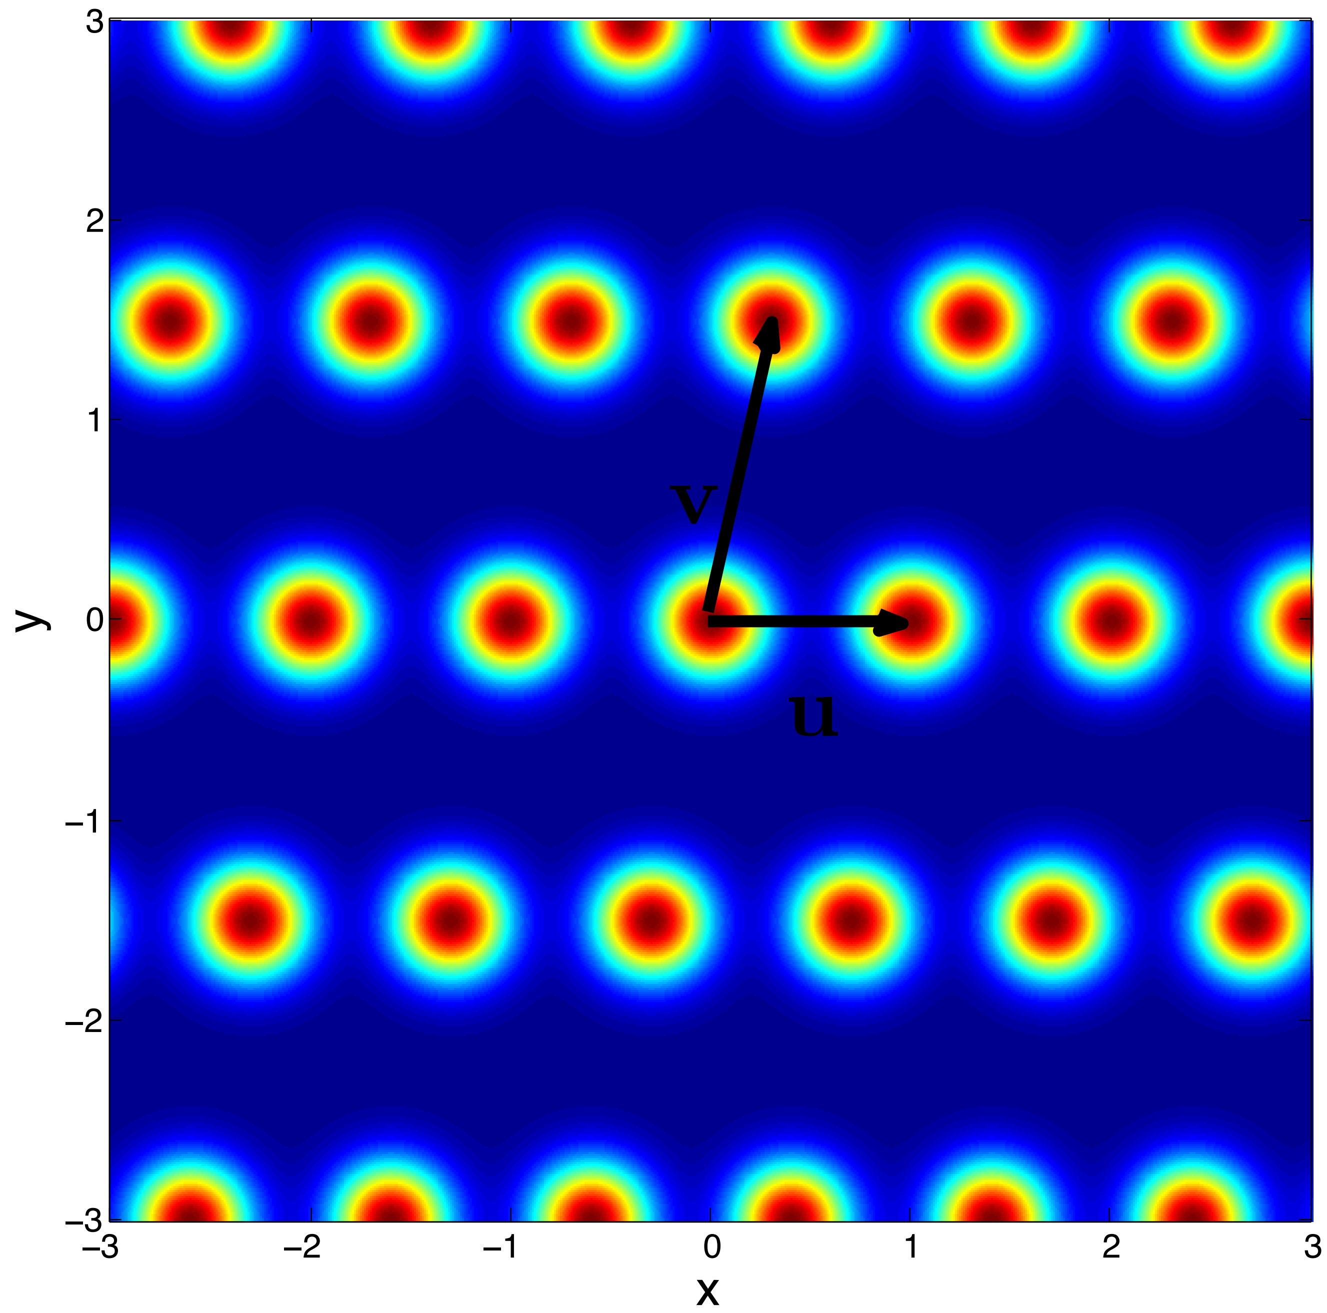

Theory matches experiment

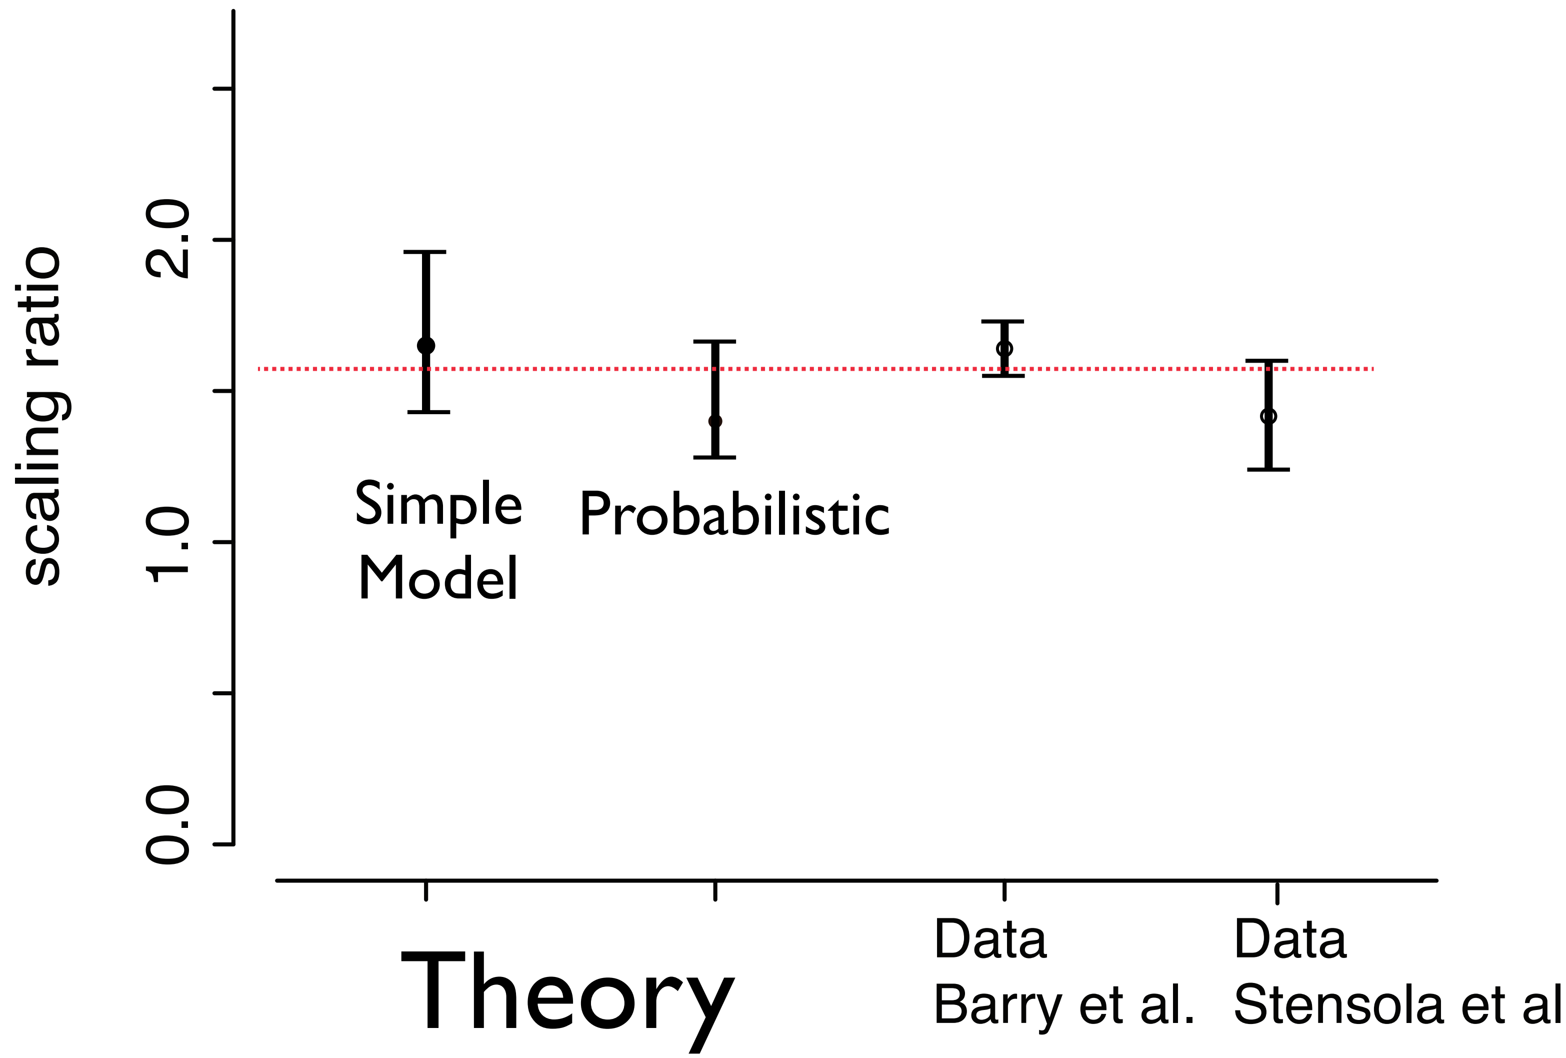

Excellent match with our theory! Evolution seems to have invented base-n number systems, and optimized them for neural hardware!

An efficiency principle seems to explain the organization of complex circuits supporting a cognitive function.

CHALLENGES

- Dynamical mechanism for self-organization of grid module repertoire via an attractor mechanism (Louis Kang, VB)

- Dynamical mechanism for explaining deformations of grids with sudden changes of the environment (Alex Keinath, VB)

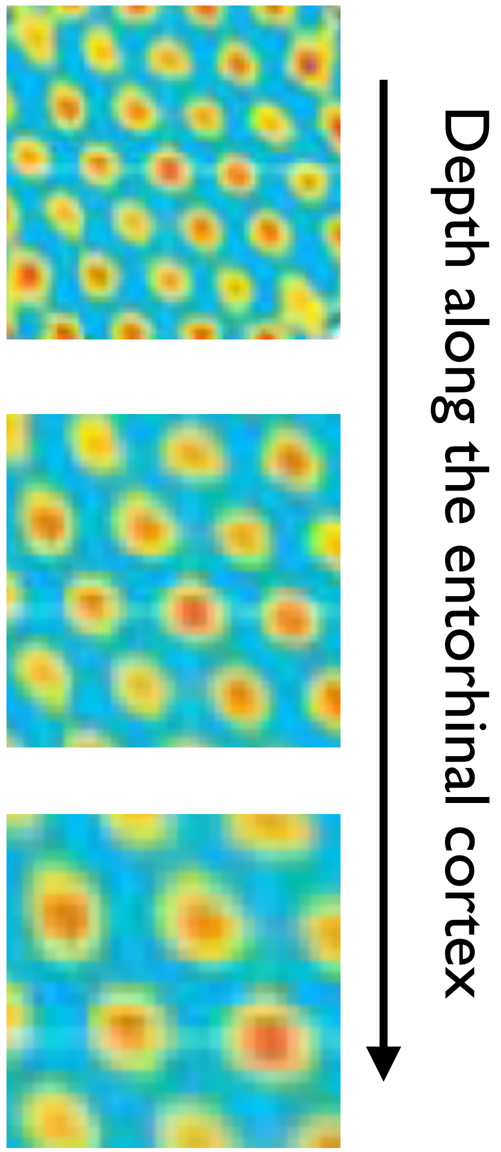

<u>Olfactory Repertoires</u>
Disordered Sensing and the Sense of Smell

Kamesh Krishnamurthy
Ann Hermundstad
Thierry Mora
Aleksandra Walczak
Vijay Balasubramanian

The intimidatingly diverse space of smells

- There are a very large number of volatile molecules, maybe 1,000,000

- Complex odors contain 100 or more molecules ➔ $(1{,}000{,}000)^{100}$ = bazillions of odor types, in each of which you can vary all the concentrations

- Odors change with seasons, and as new opportunities and threats come to light.

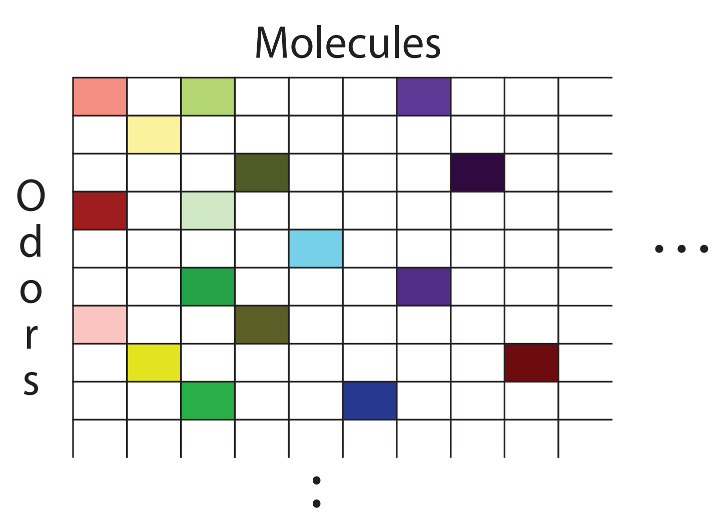

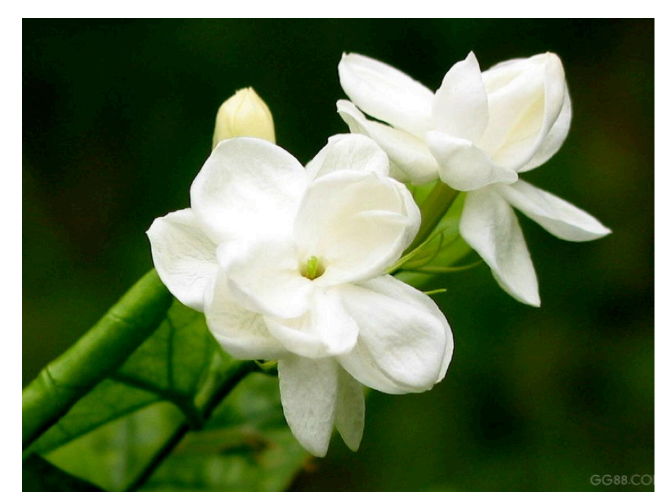

The challenge of identifying odors

- Odors are sensed when molecules bind to Olfactory Receptors in the nose

- Every receptor needs a separate gene. Flies have ~100, humans have ~ 300, mice have ~1000.

- How can you possibly represent so many odors with ~1000 sensors?

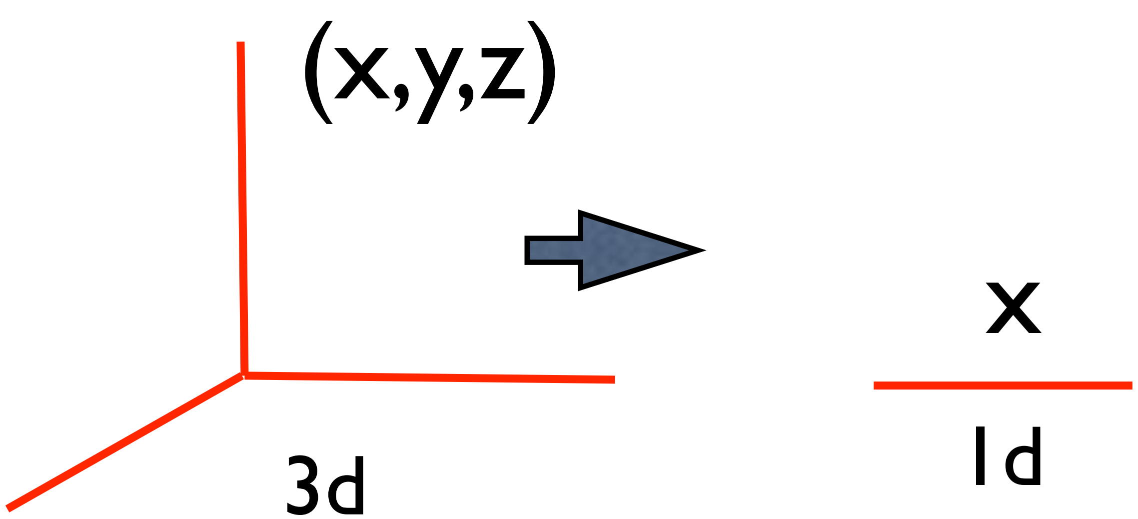

Imagine representing three dimensional positions using a single number. Also want to preserve proximity relations.

The space of odors is perhaps 1,000,000 dimensional and we have maybe 100-1000 numbers to describe it. Can this be done?

Sensory processing with limited resources

- **Strategy**: adaptation to the environment
- Exploit stable structure in the world to produce compact and easily manipulated information architectures

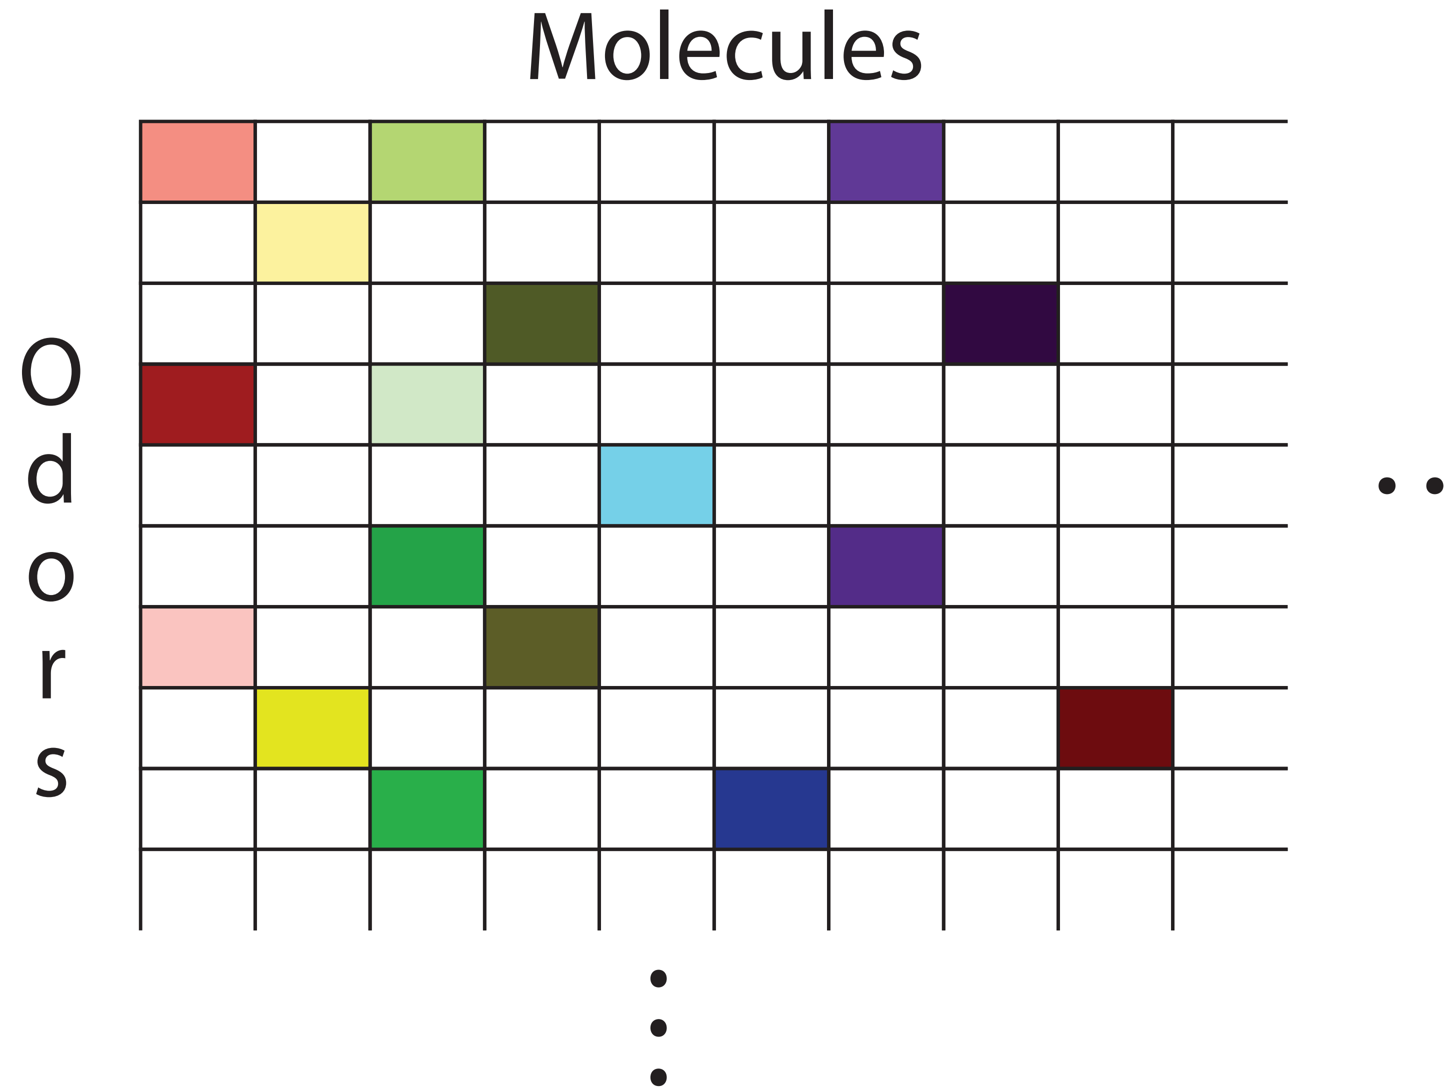

Natural odours are sparse in “chemotopic” space

Efficient packing of odors by disordered sensing

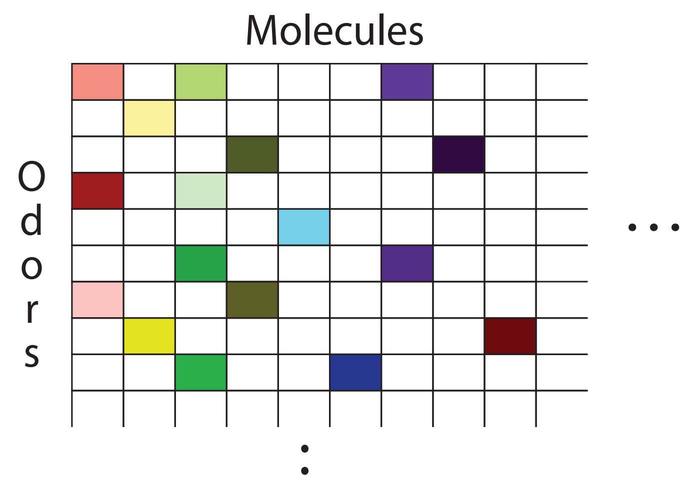

Natural odours are sparse in "chemotopic" space

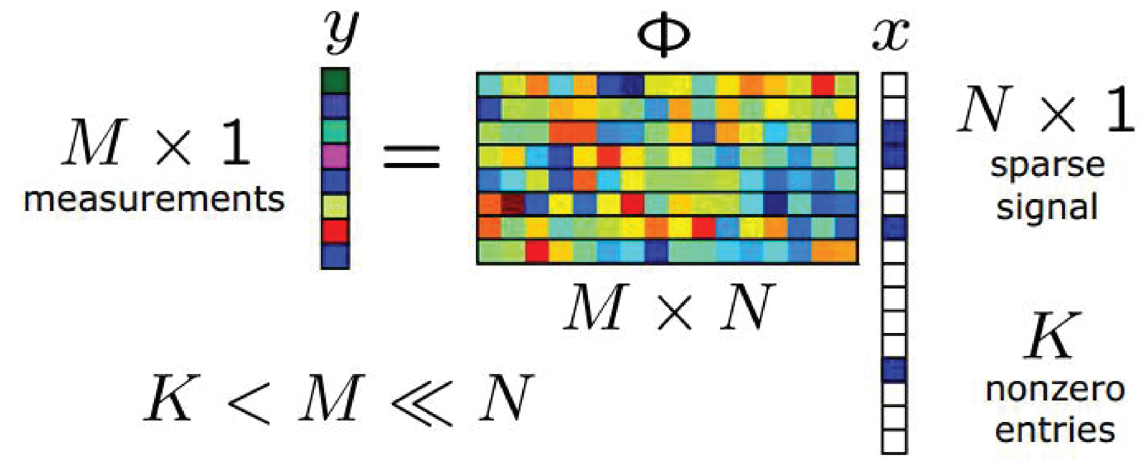

Can recover x from y as long as x is known to be sparse

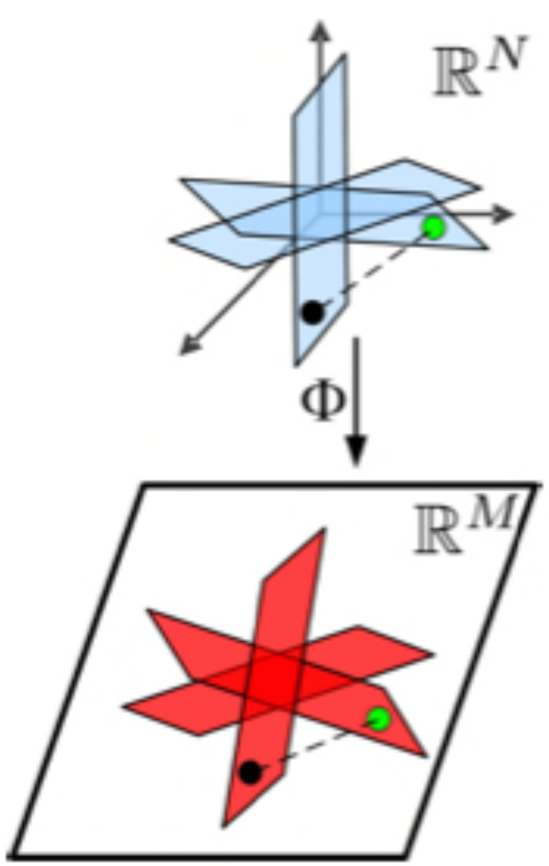

Sparse vectors (odors) that are nearby in the high-dimensional odor space will be nearby after the projection into neural space

<u>Decoding</u>

Find a vector **x** such that **y**= A **x** (where **y** is the measured output), and x minimizes the L1 norm.

$$||\mathbf{x}||_1 = \sum_{i=1}^{N} |x_i|$$

Olfaction & Disordered Sensing?

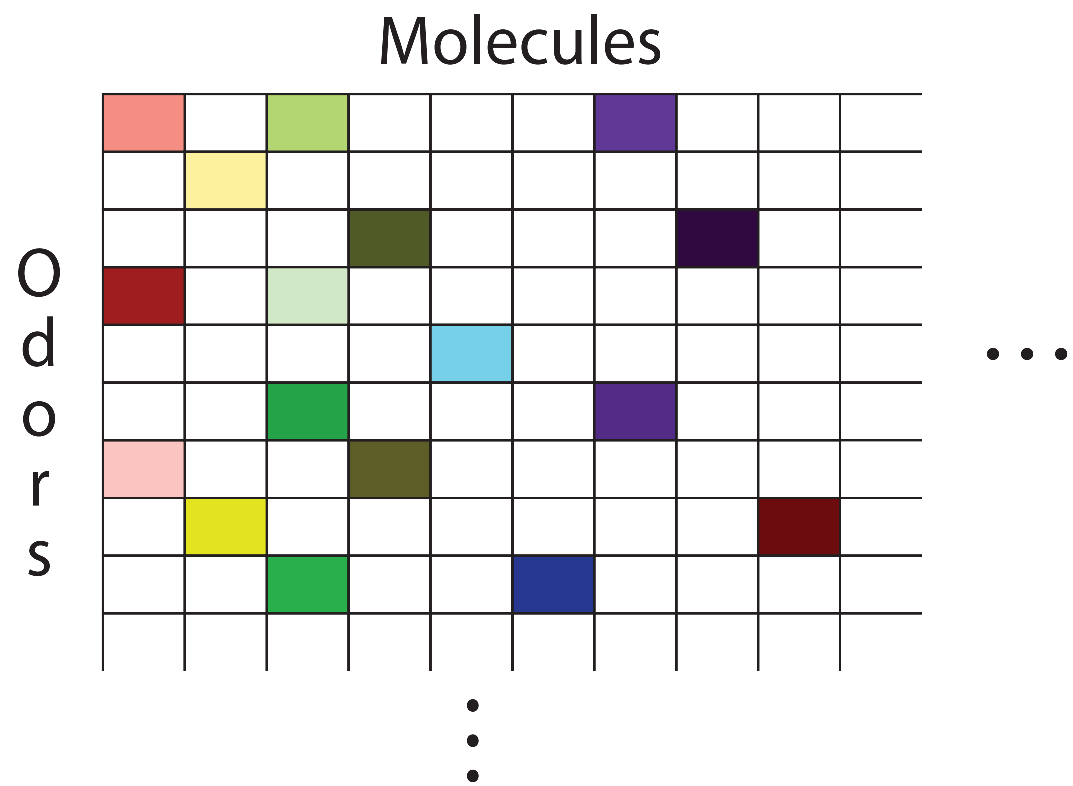

The typical odor contains maybe ~50 of the millions of possible molecules

If each receptor binds with "random" affinities to volatile molecules then you only need ~O(100) sensors to represent all odors.

Does the olfactory system use this approach?

Key Property: Diffuse randomized sensing

The disordered structure of olfactory circuits

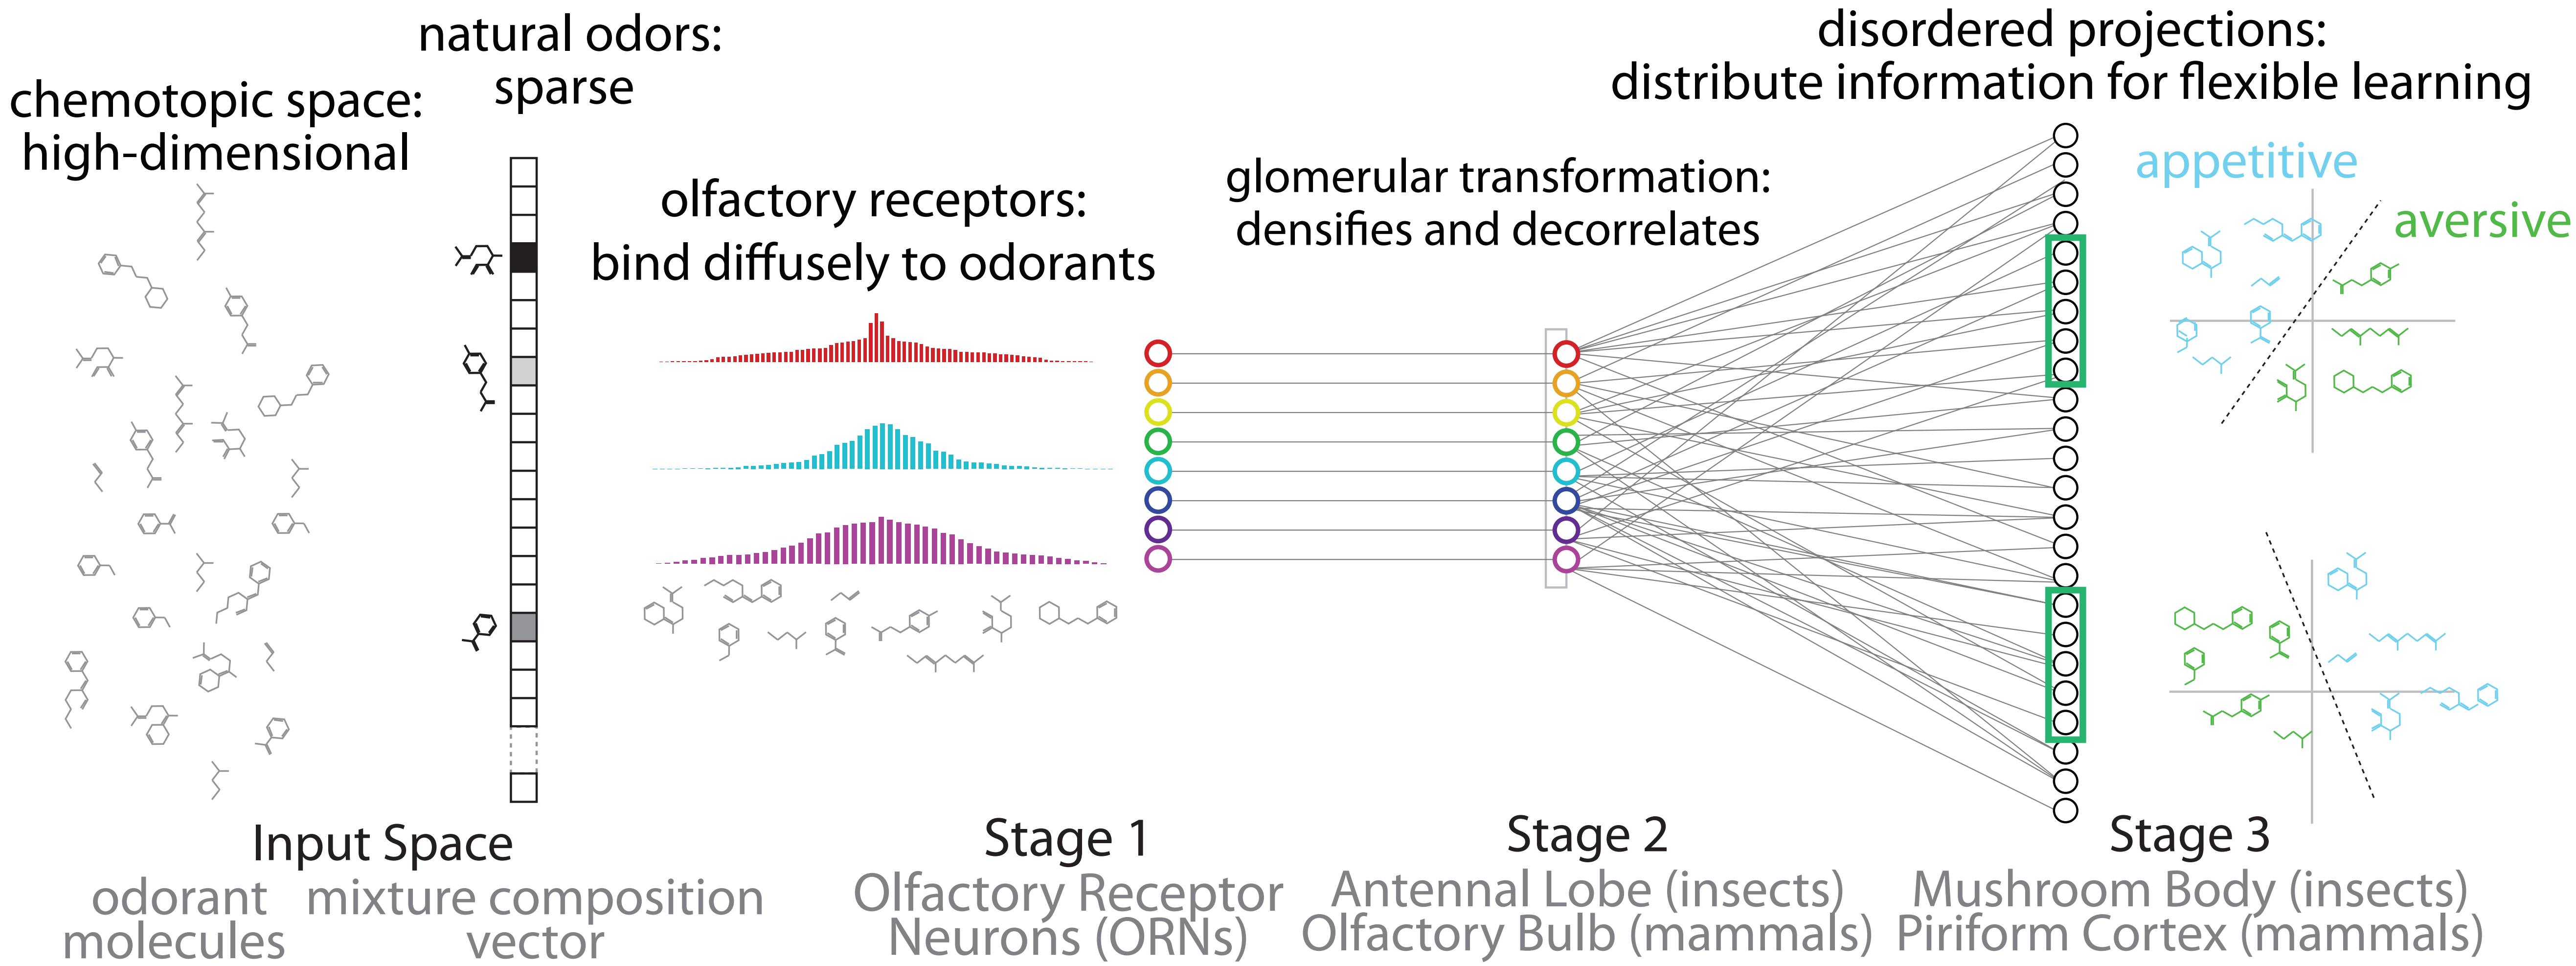

Maybe the brain disorders odor information as mathematicians would recommend

Stage 1: Receptors and random sensing

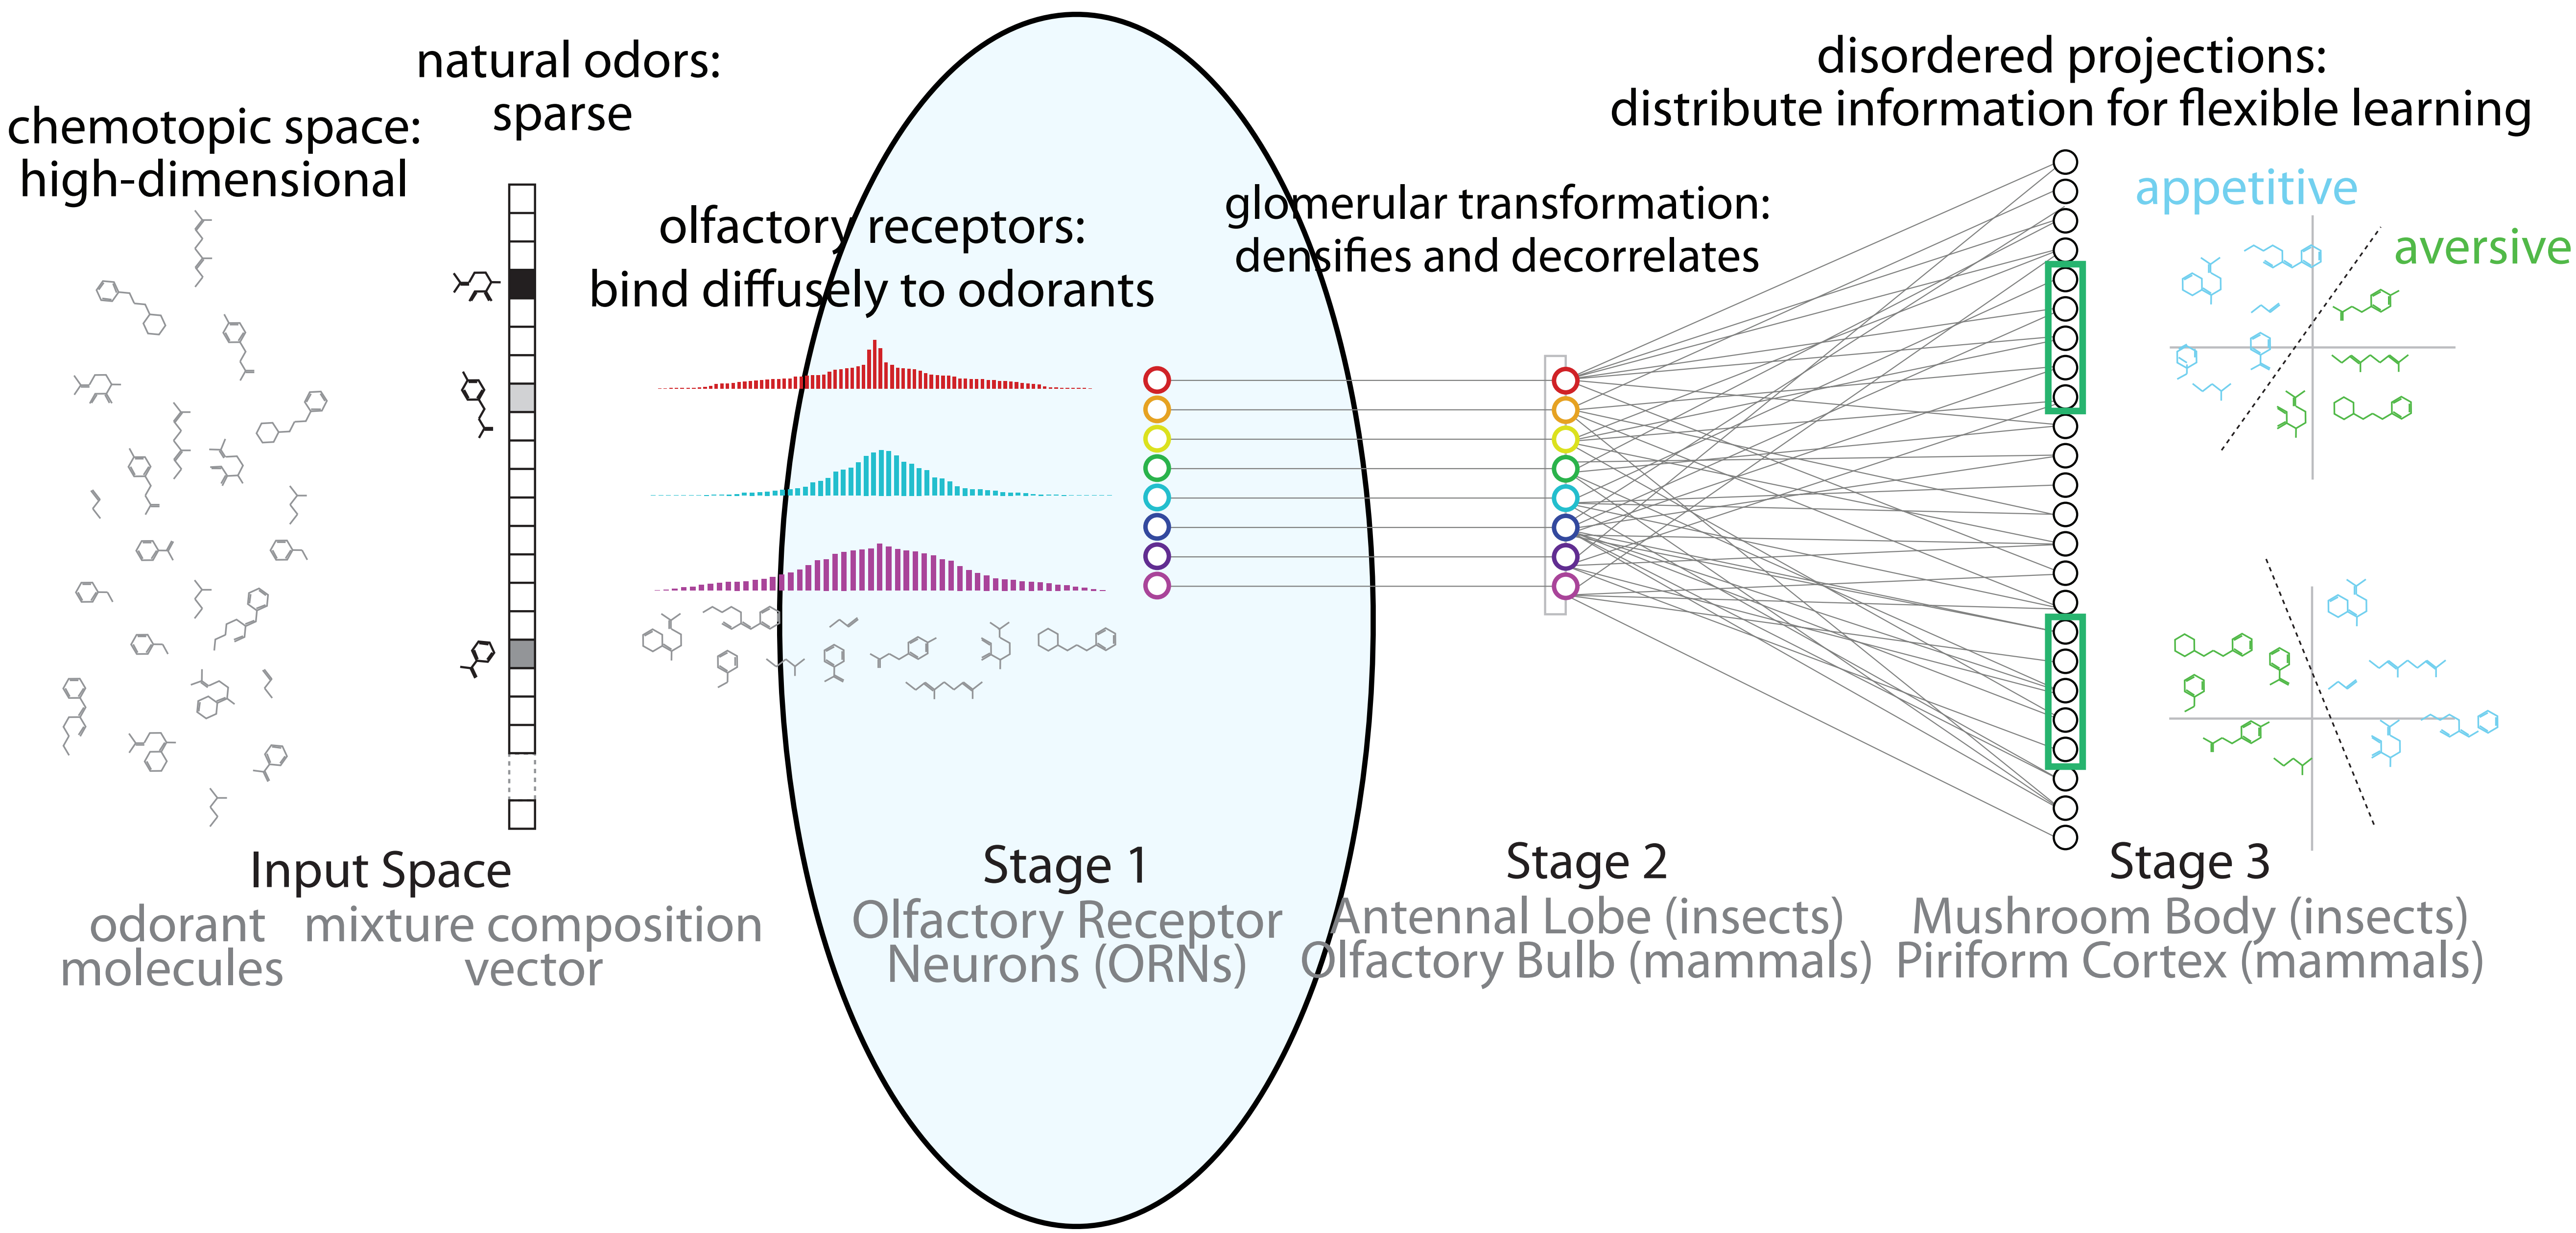

Stage 1: Decoding from receptor responses

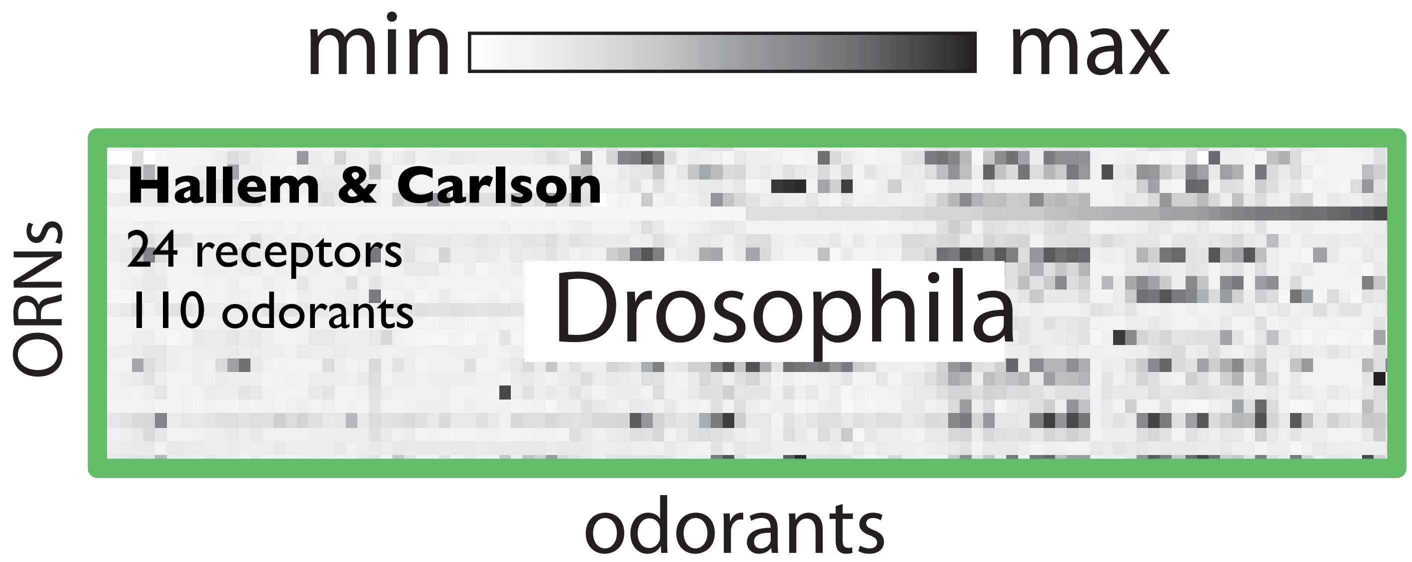

R_{ij} = response of receptor i to odor j

Linear sensing model: $\vec{y} = R\vec{x}$

Decoding error = fraction of decoded odors that differ from original by more than 1% in norm

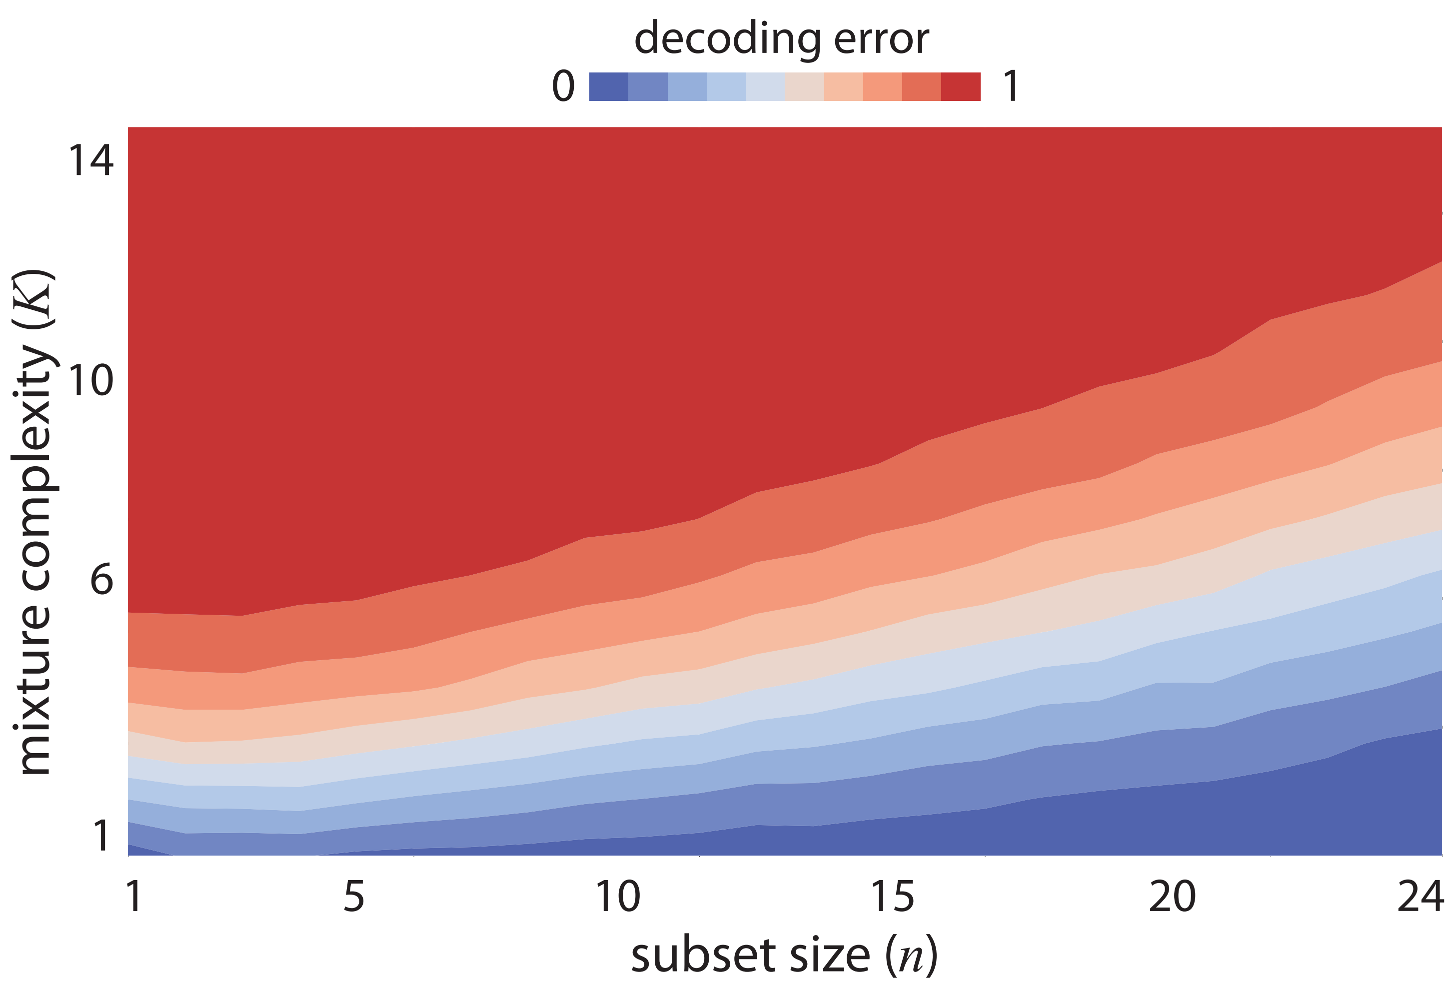

Mixture complexity = number of odor components in a mixture

- 67% of odors with 5 or fewer components drawn from 110 odorants can be accurately decoded from responses of 24 receptors.
- There are ~100 million such mixtures

Decoding odor mixtures from receptor responses

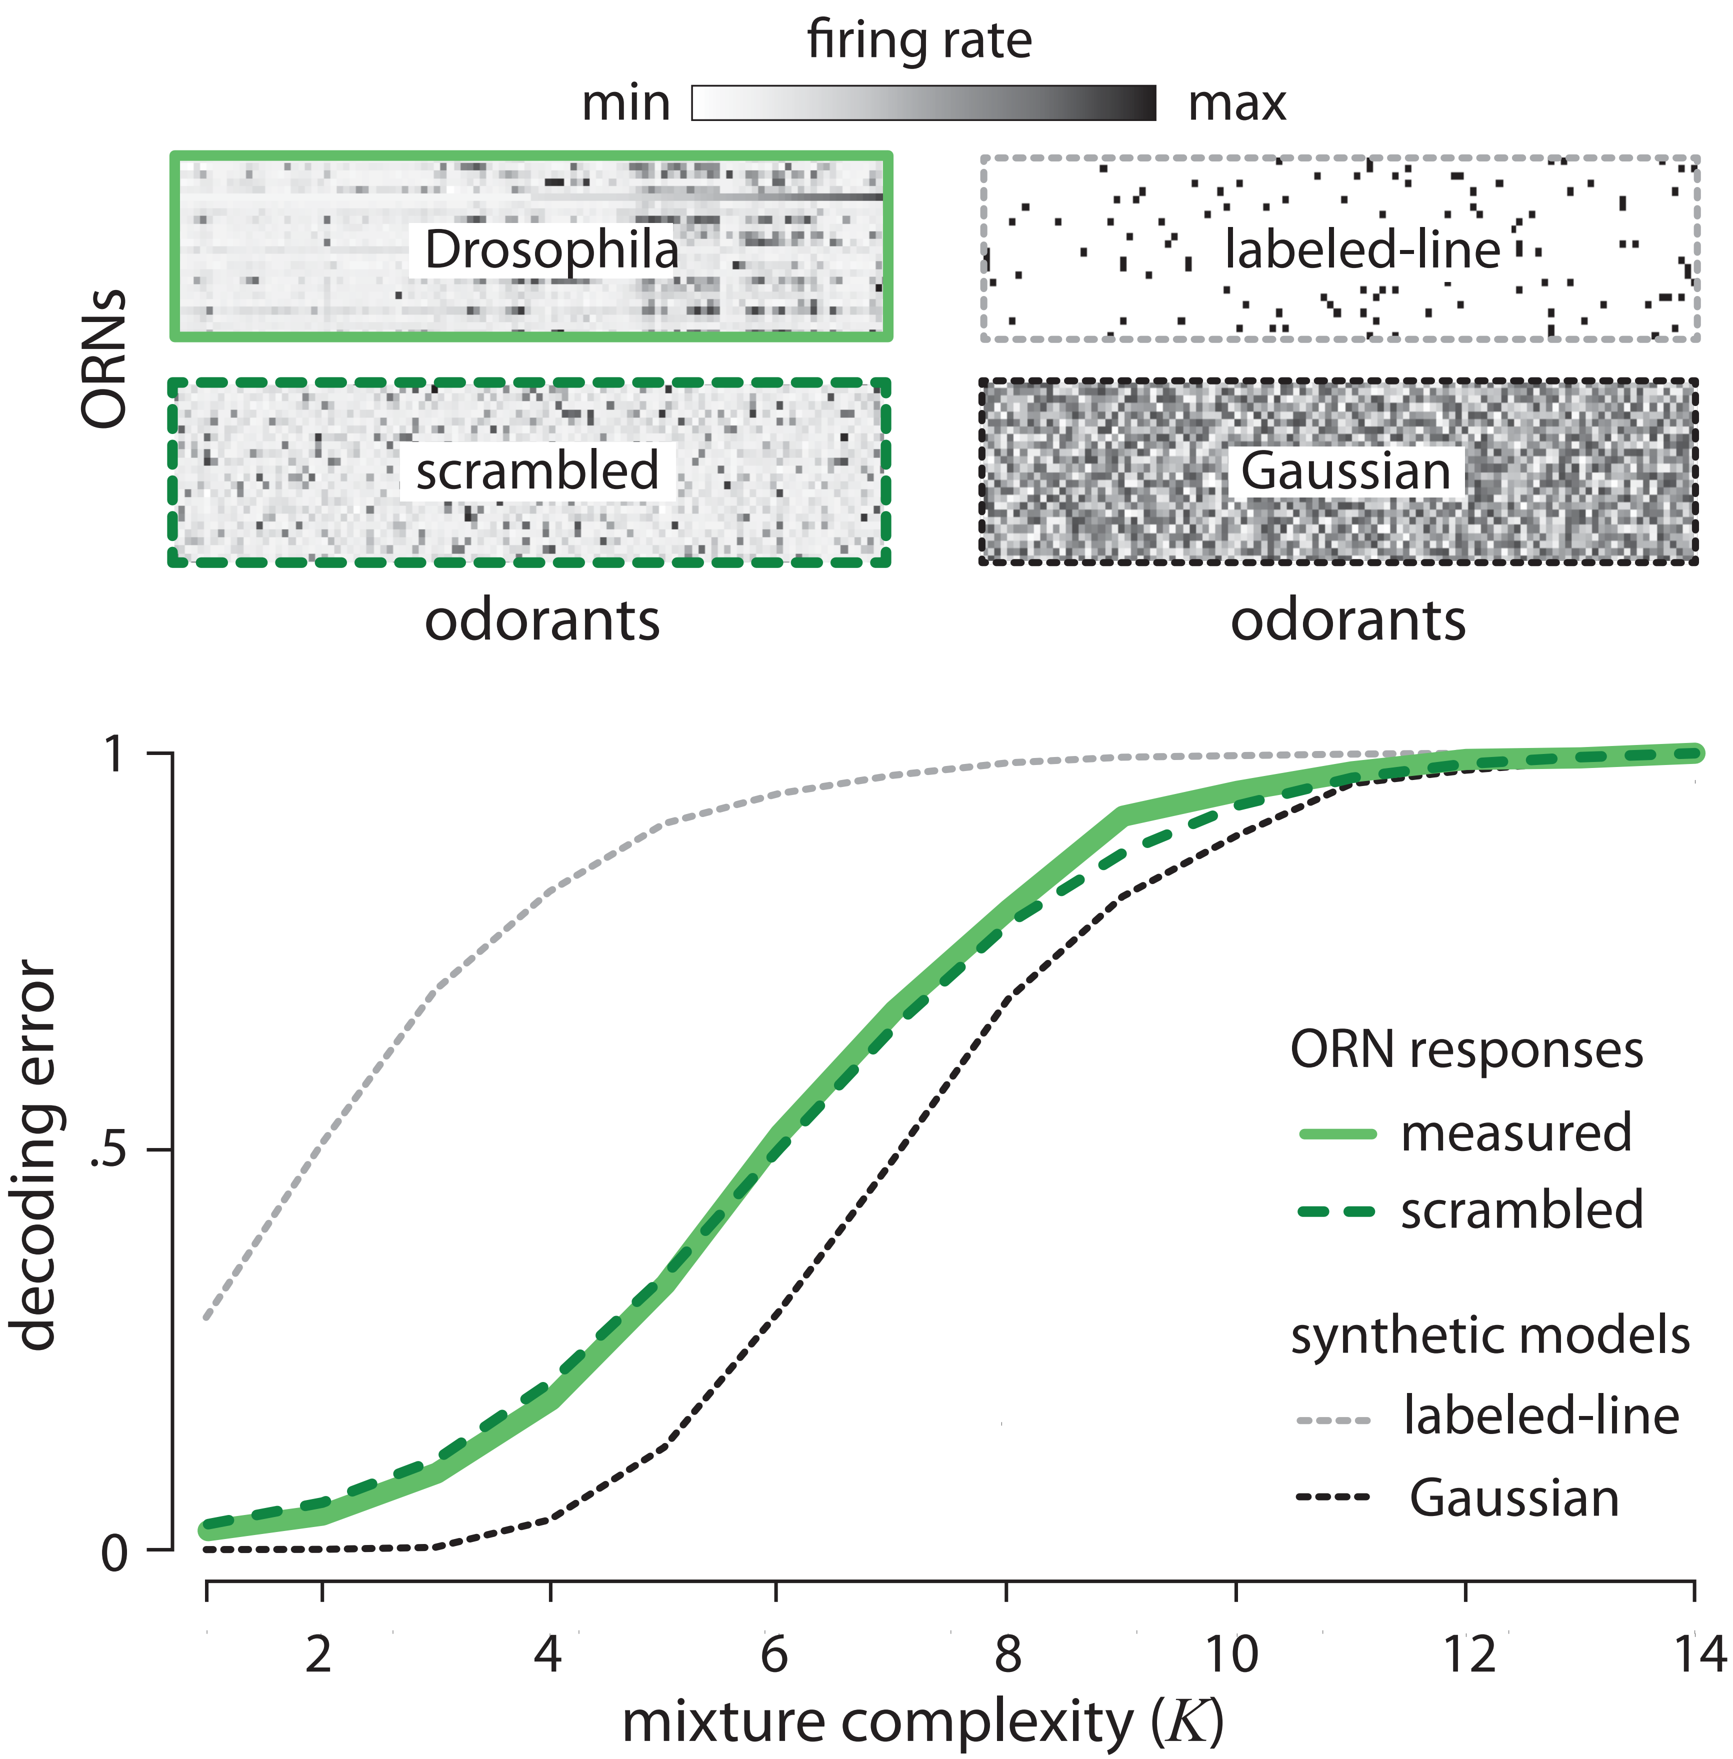

- scrambled = random permutation of Drosophila sensing matrix
- labeled-line = threshold to keep the k largest responses (k = 5)
- Gaussian = ideal random sensing model

Receptors approach ideal performance

Stage 2: Decorrelation in the olfactory bulb

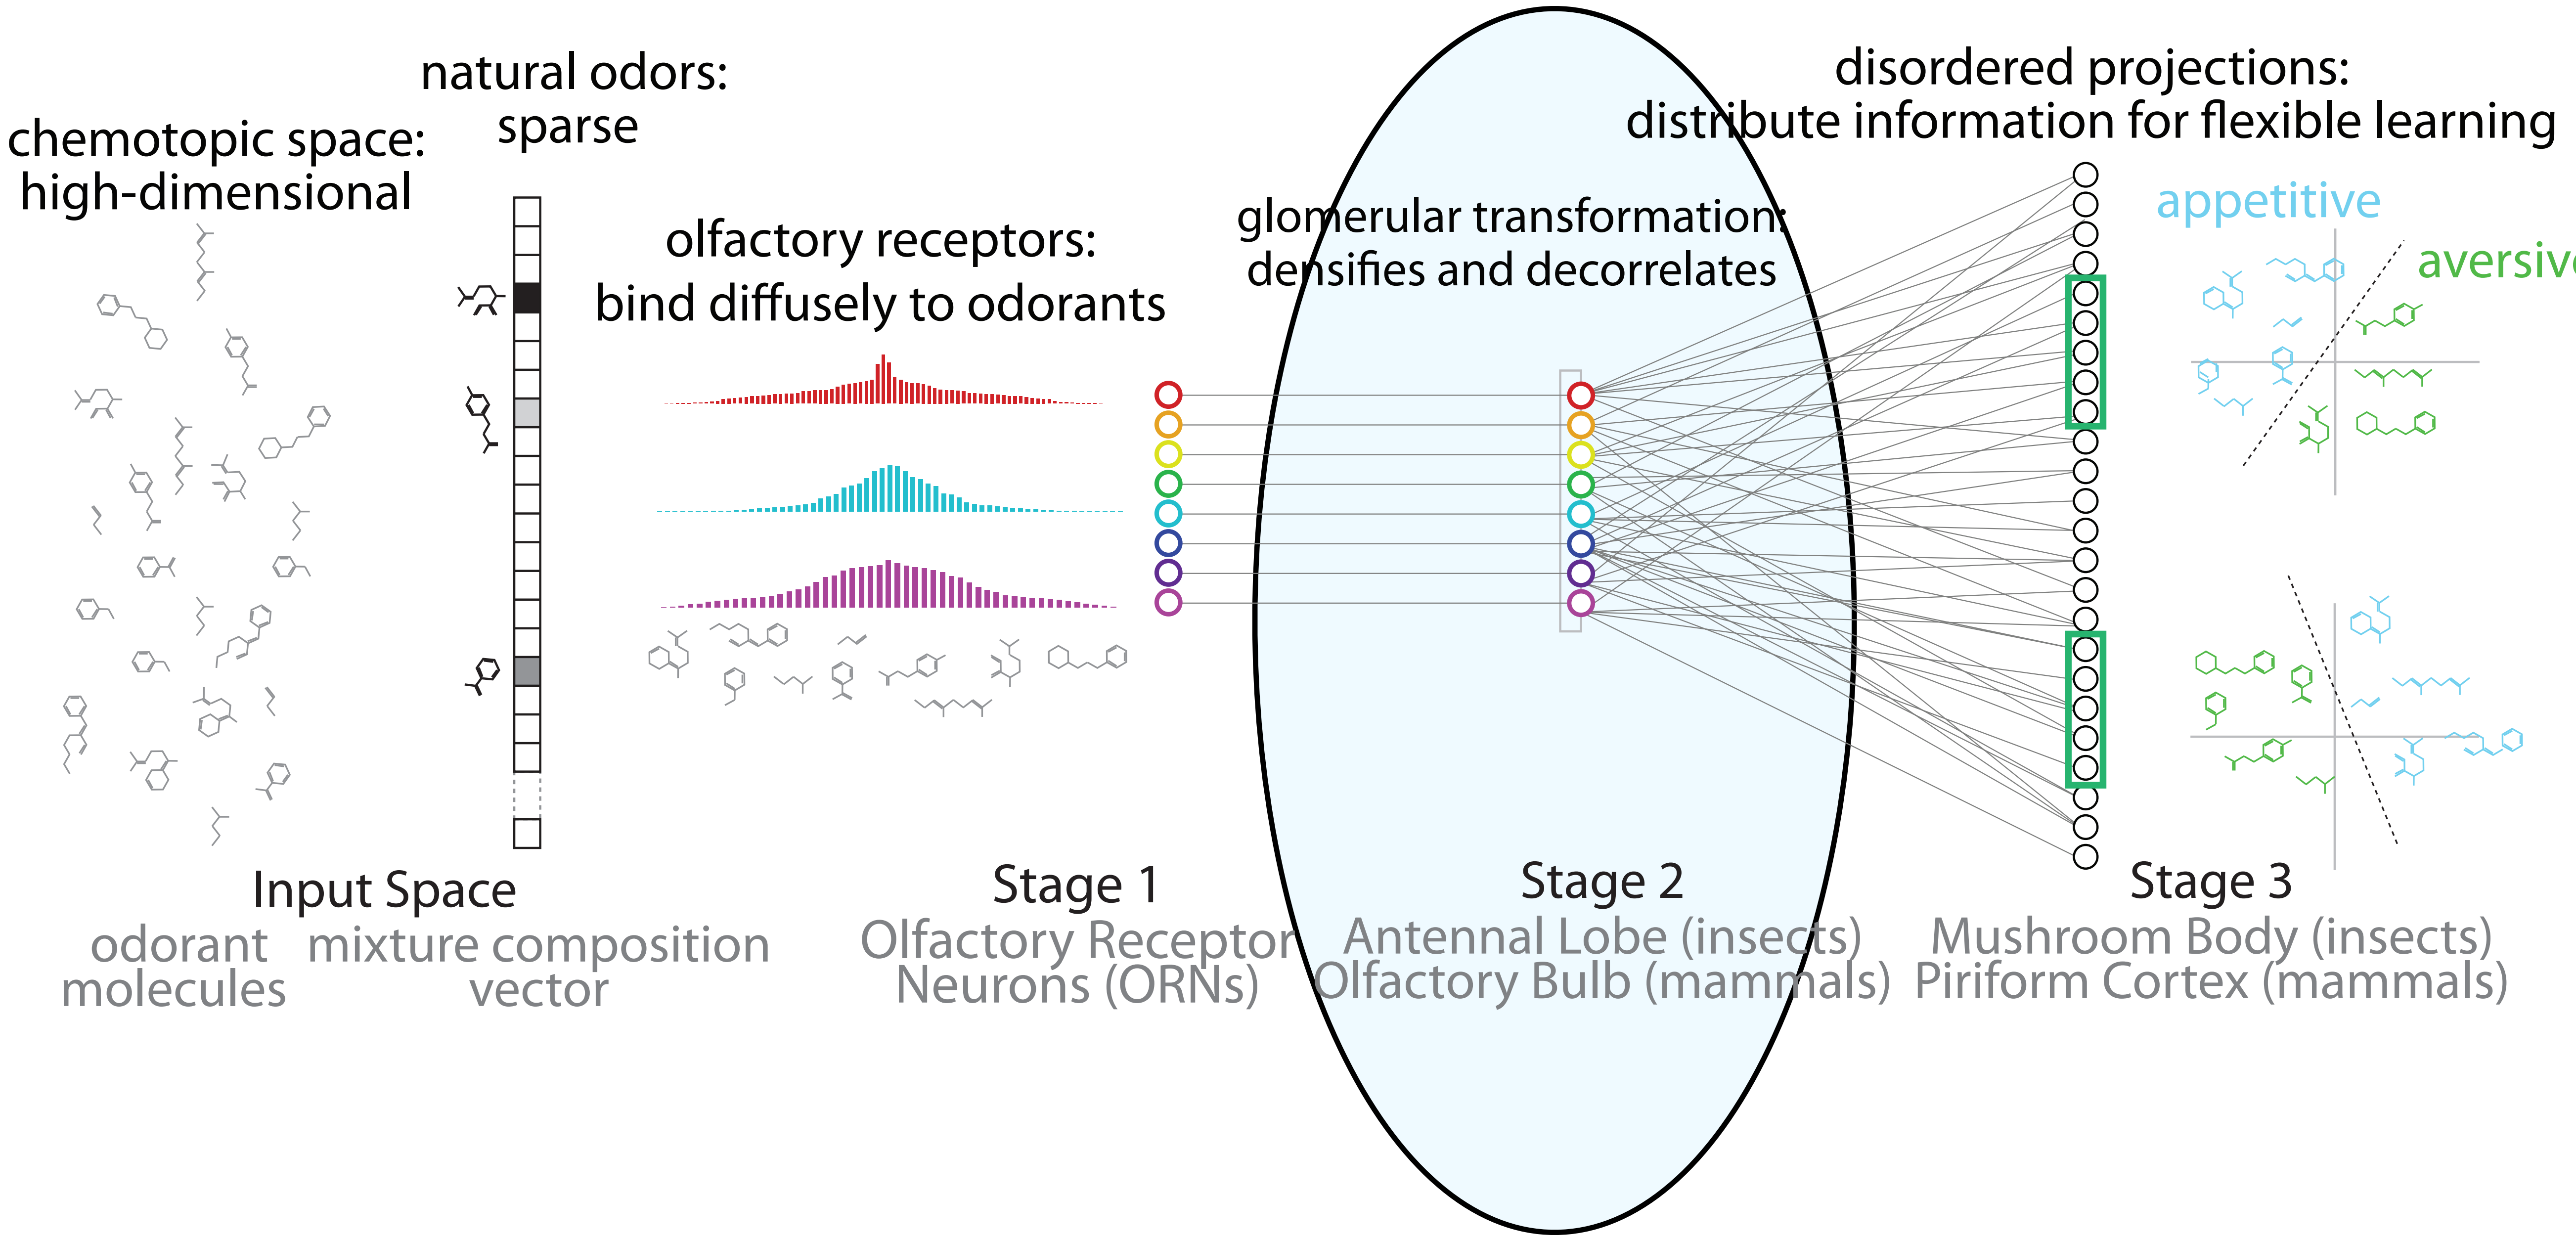

Stage 3: Random expansion to cortex

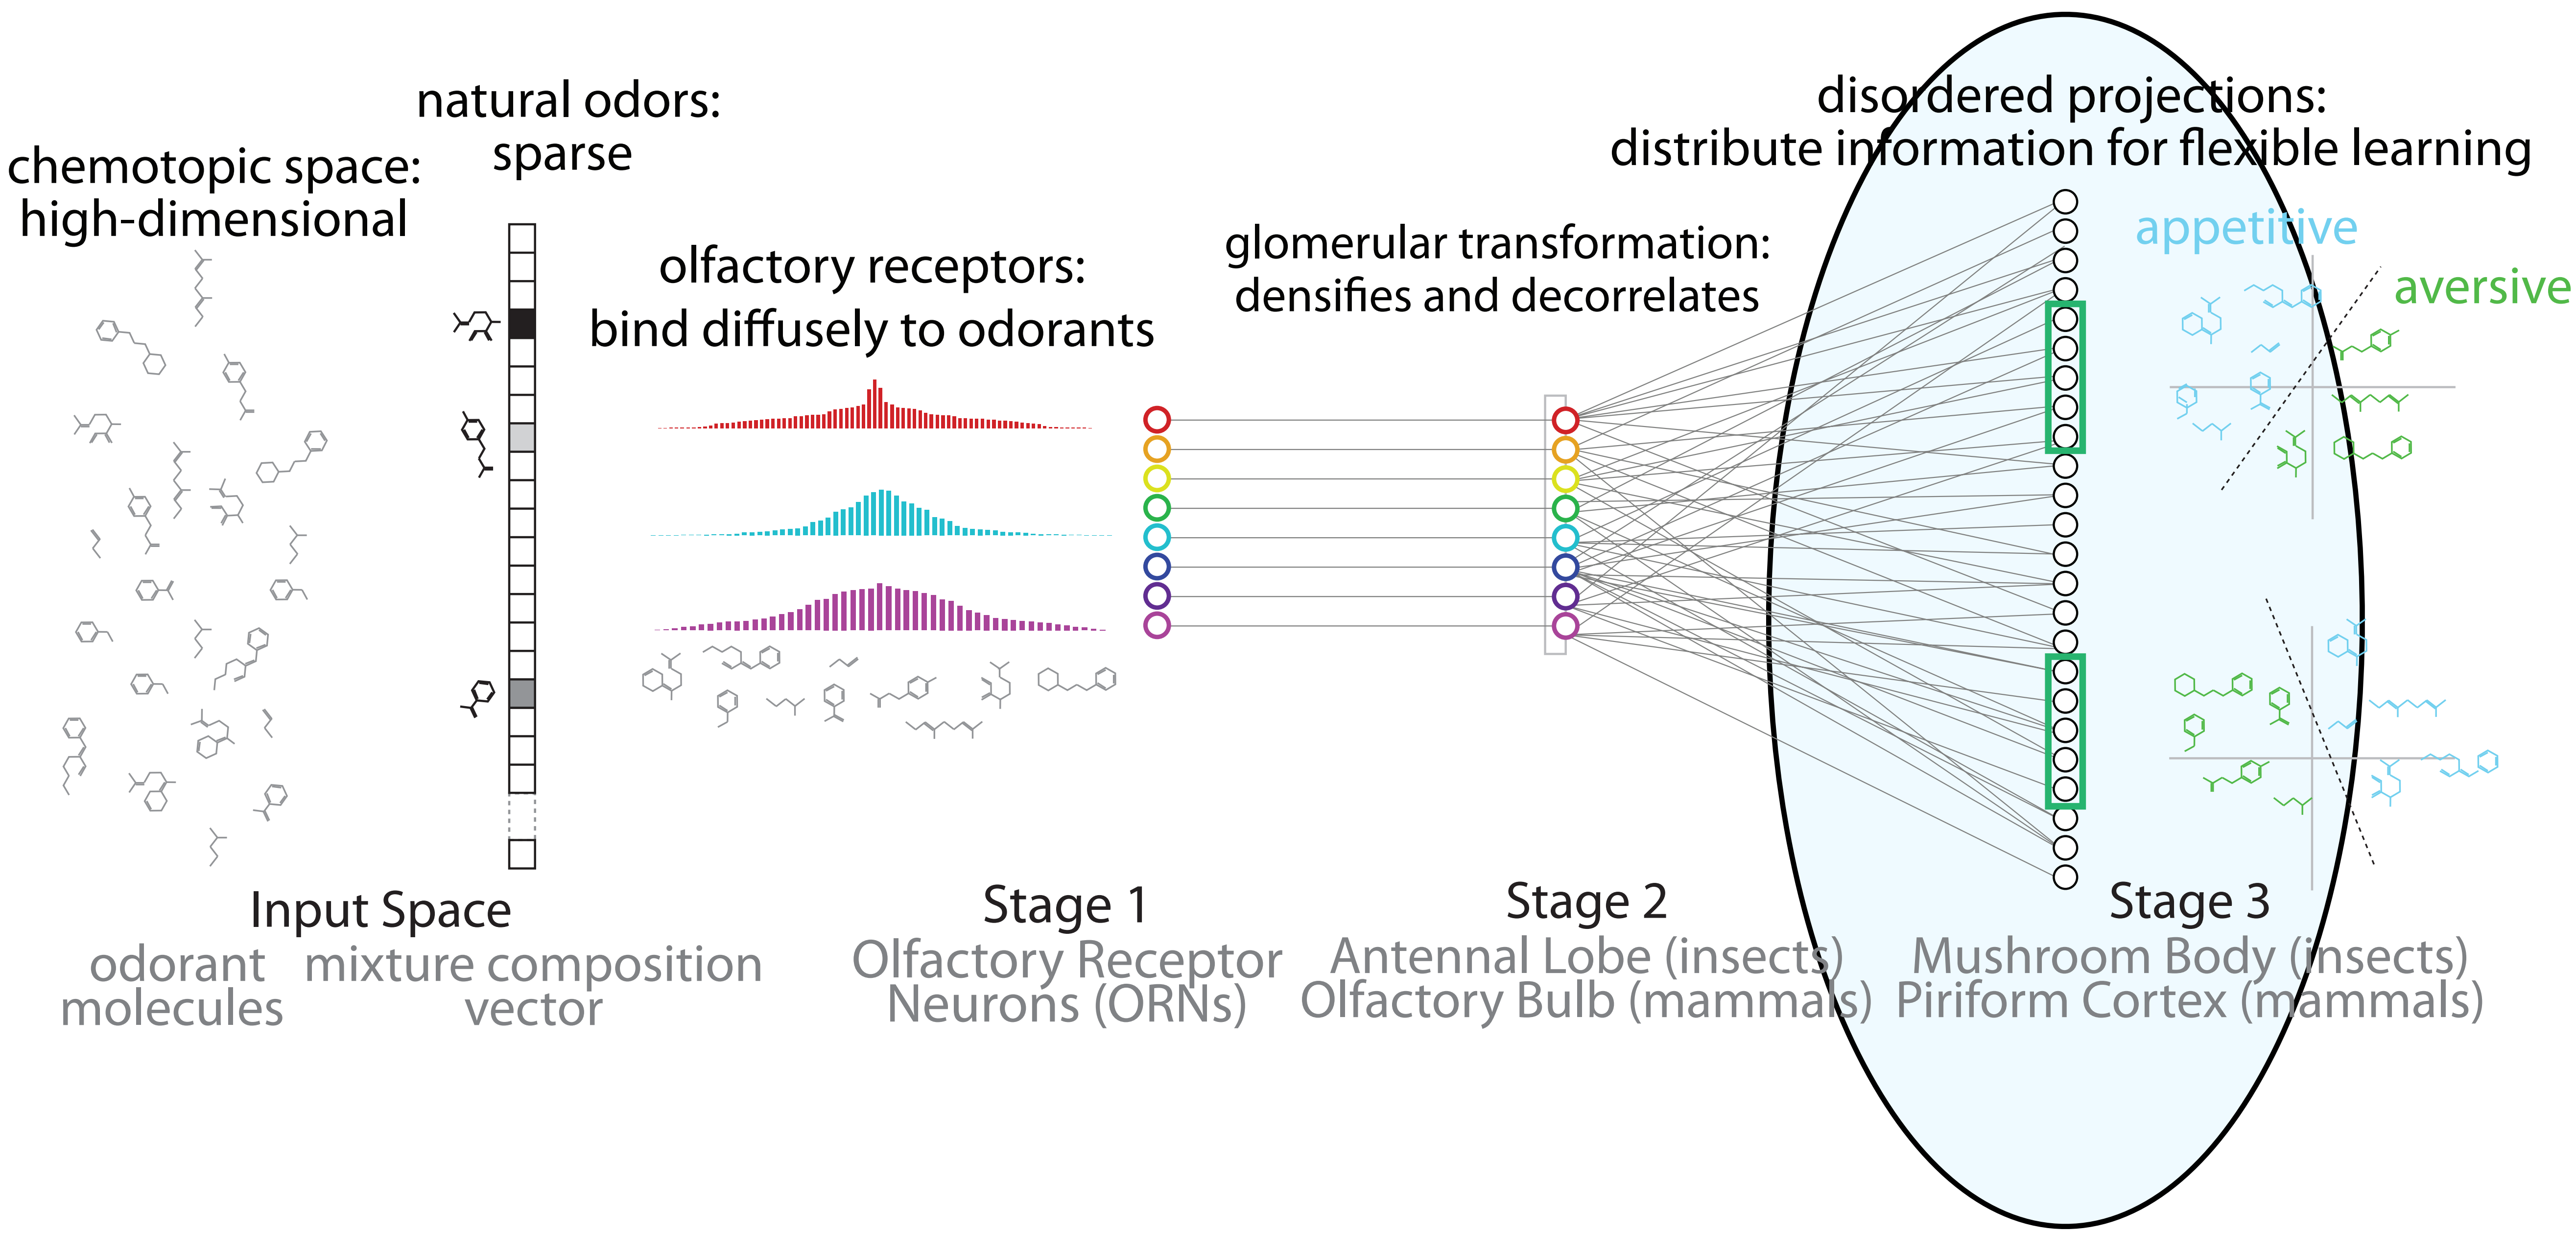

The maps inside your head

Towards a theory of *functional maps* and *computational repertoires* in the brain — today, examples from vision, olfaction & spatial cognition

- Adaptation to the environment and to the task
- Constraints of neural computation
- Efficiency and parsimony

An attempt to explain the immense diversity and complexity of computational architecture in the brain

For the future

More generally, the same ideas turn out to be relevant for understanding "functional repertoires" in e.g. the immune system, biochemical circuits in cells.

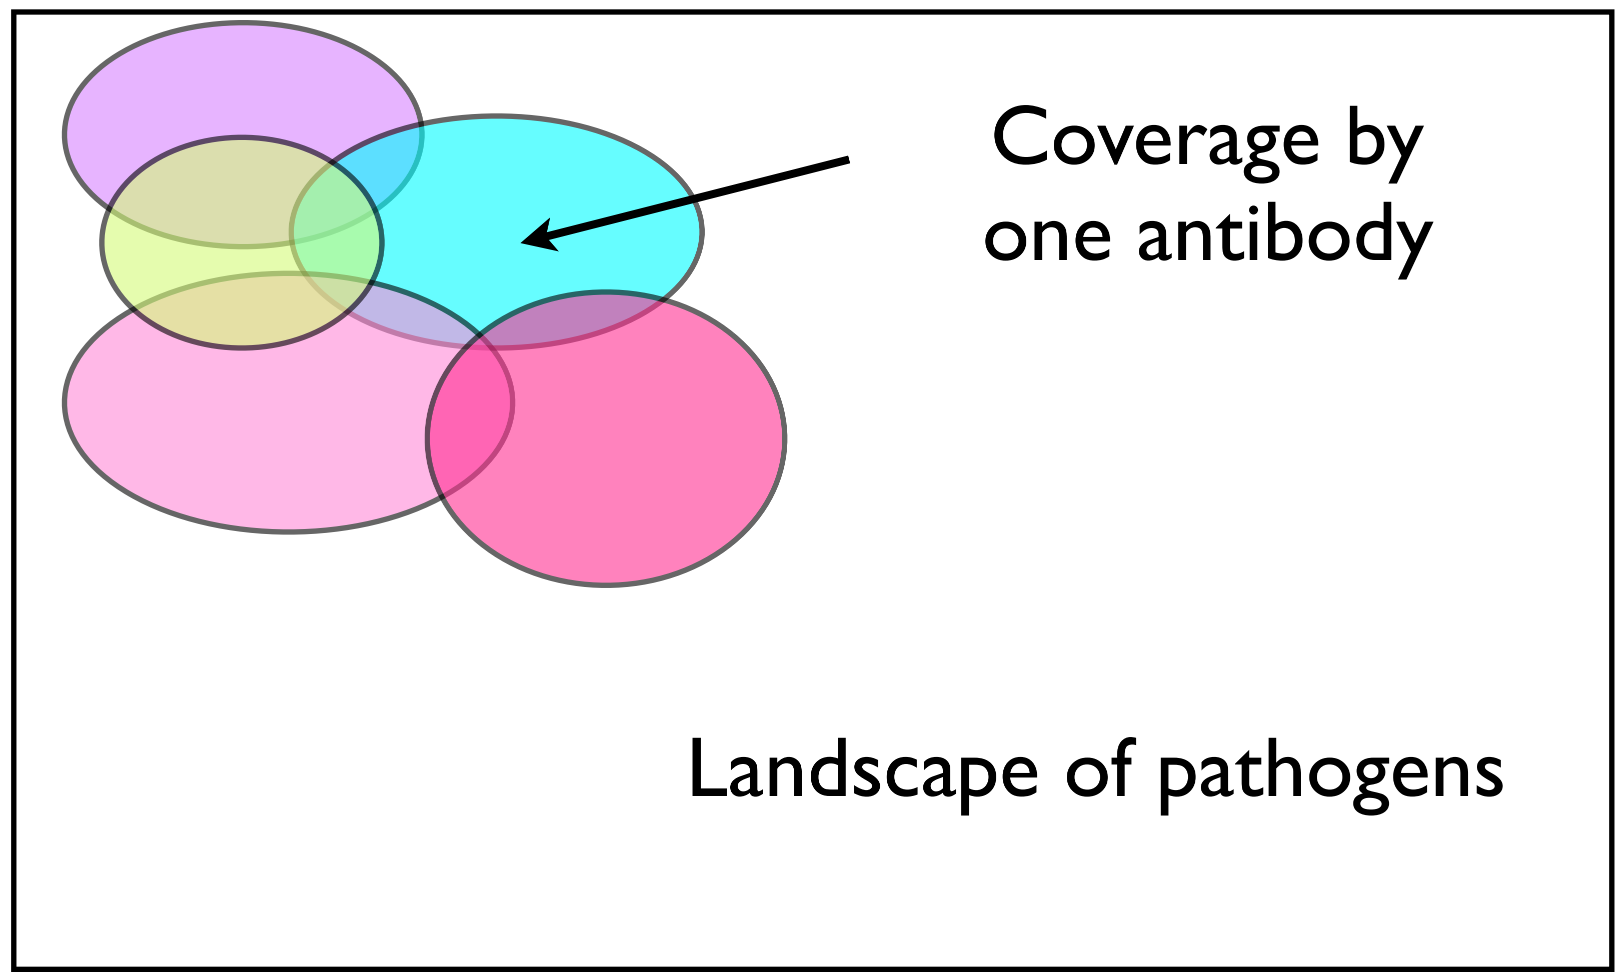

The End

Stage 3: Densification and disorder increase robustness

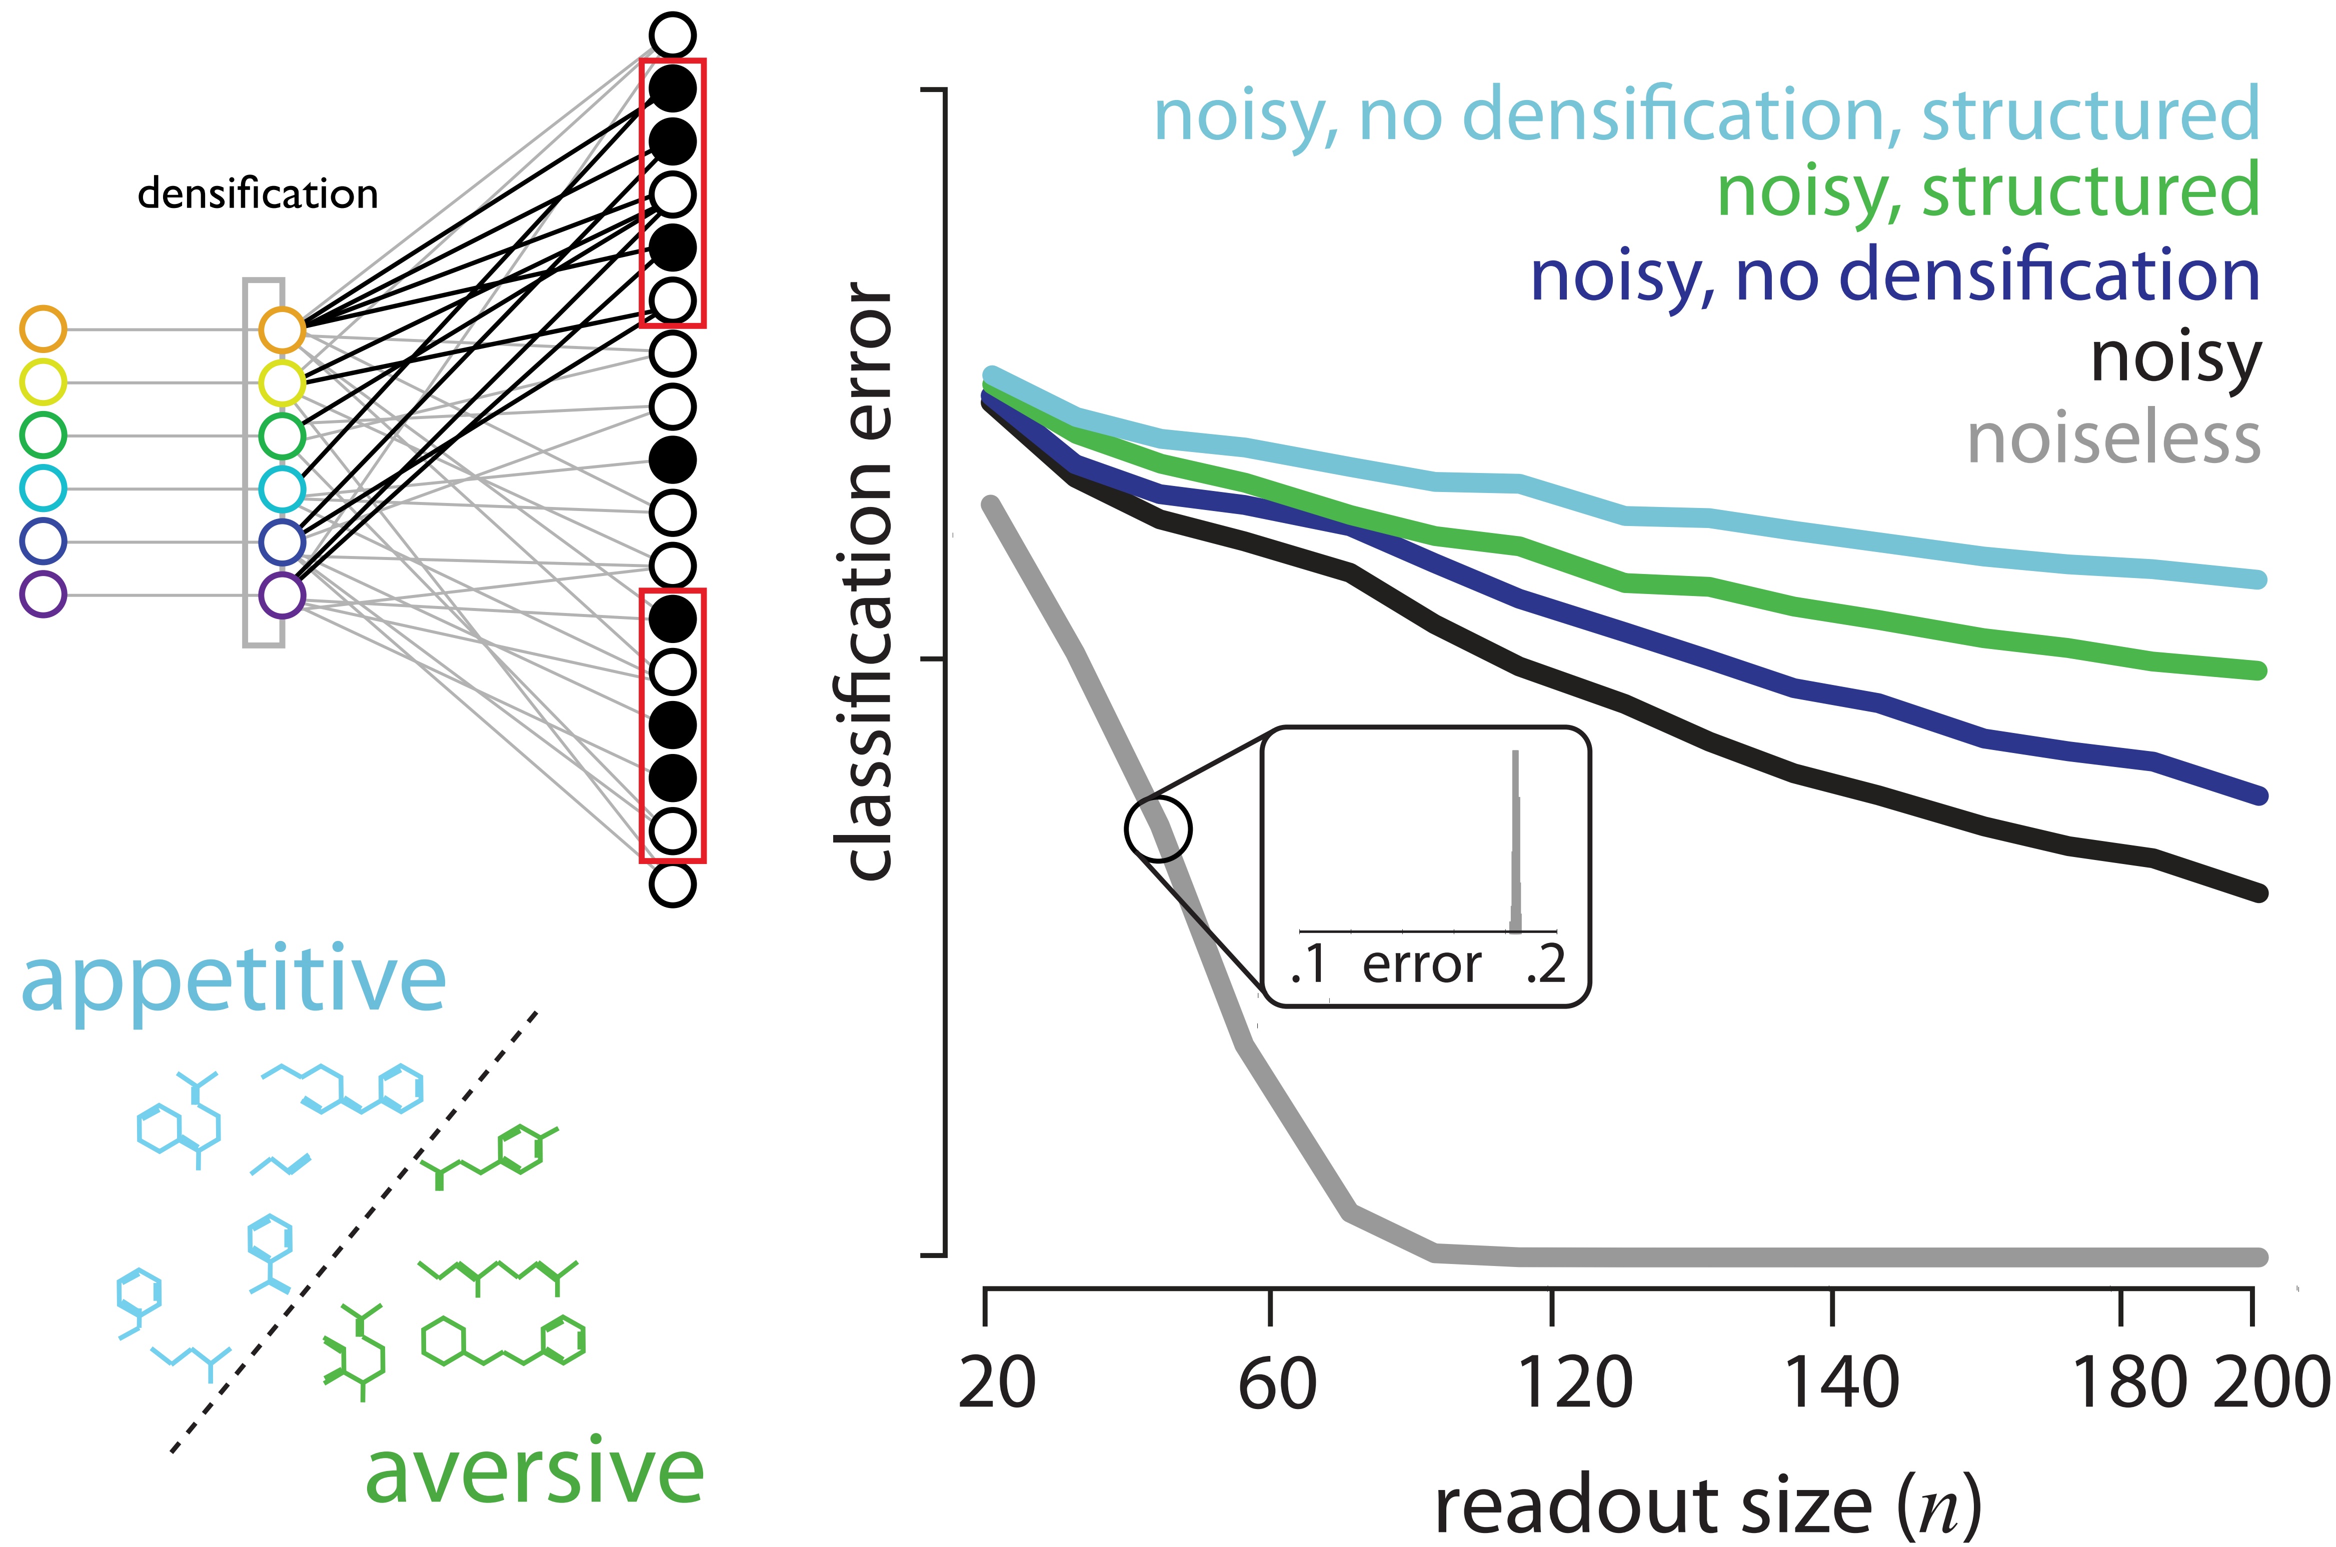

Stage 2: Decorrelation

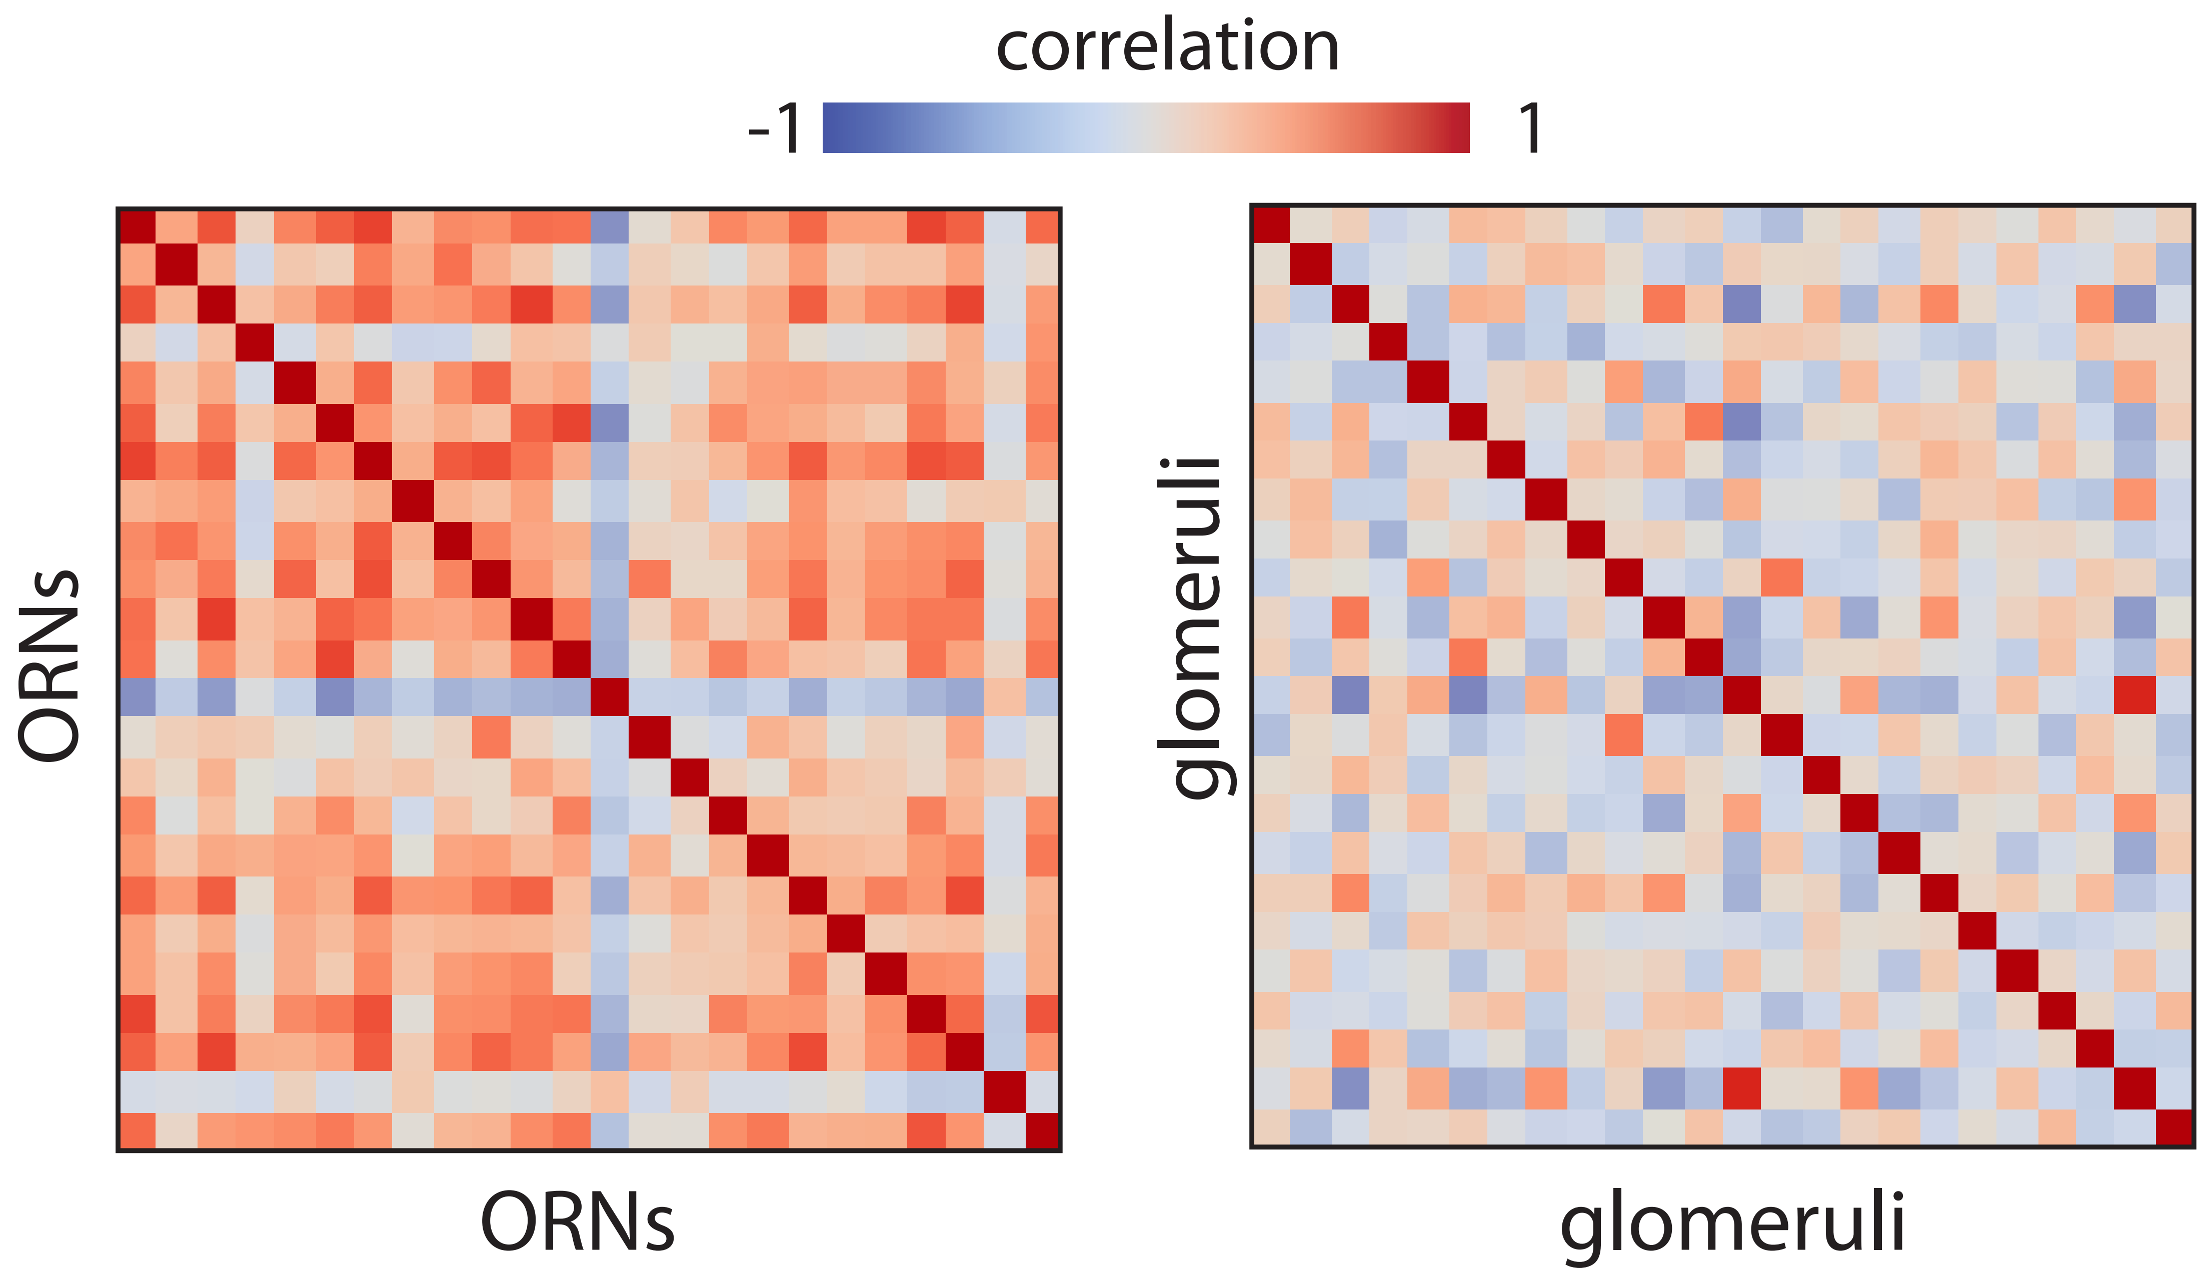

The nonlinear transformation decorrelates the receptor responses.

Animals can learn arbitrary associations from cortex

Choi et al.

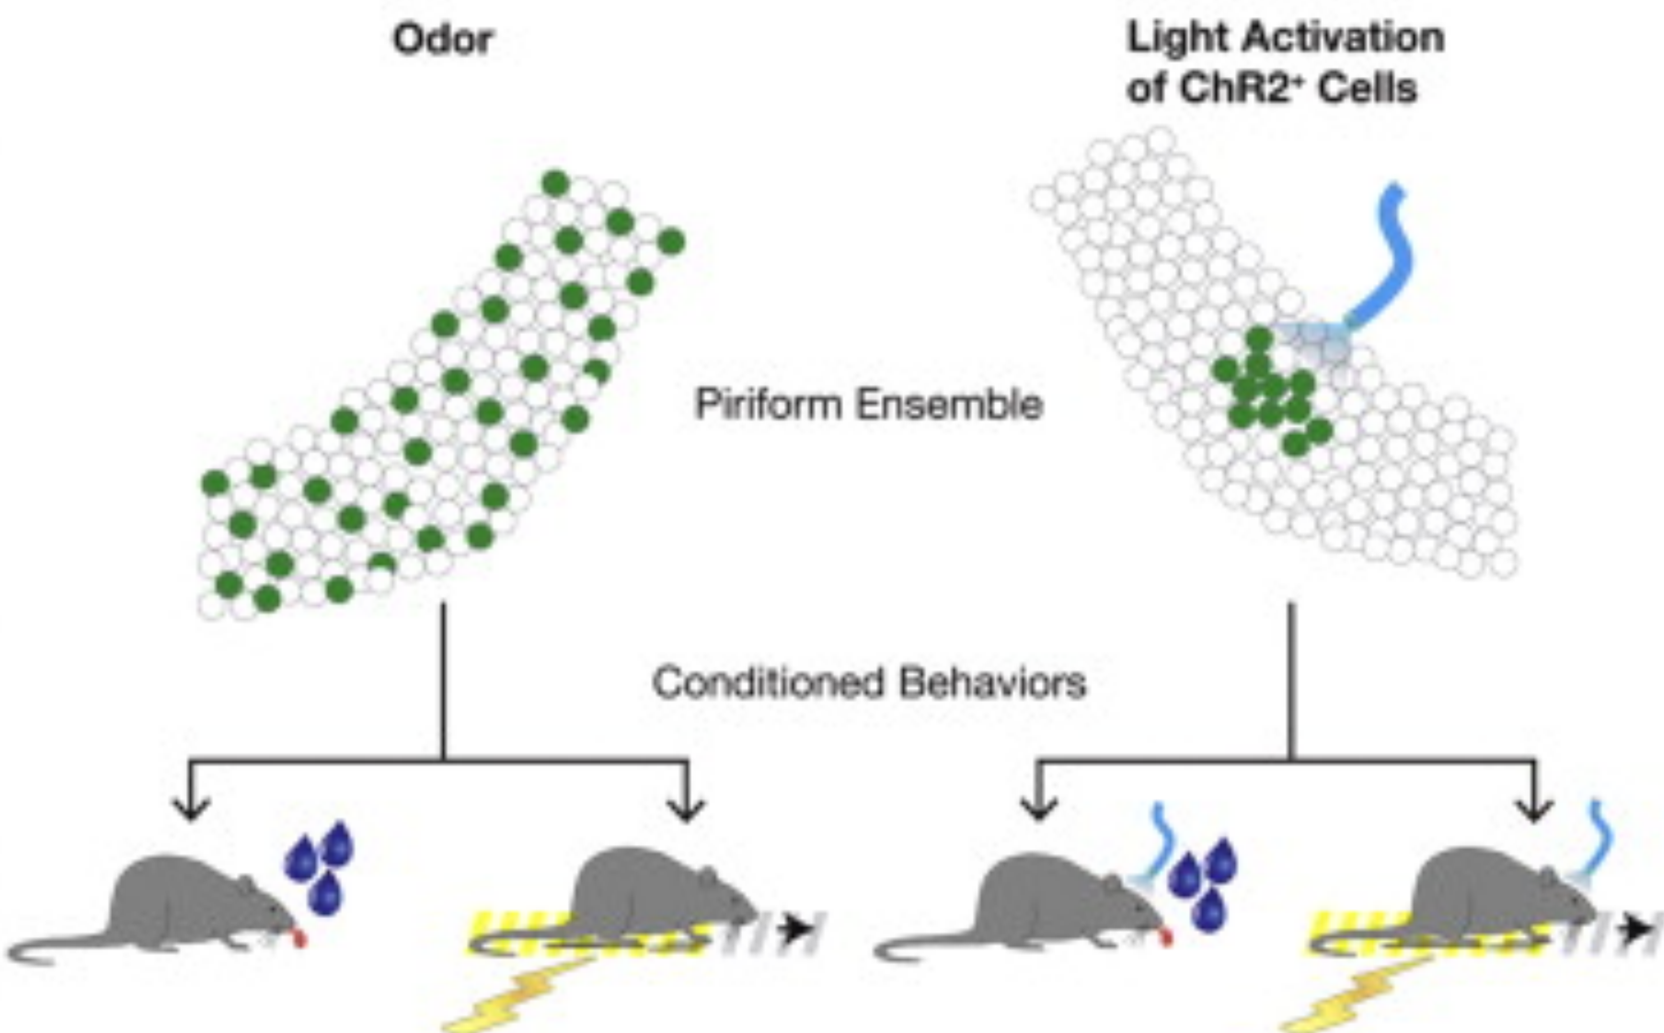

- Random ensembles of cortical neurons can be entrained to elicit appetitive and aversive behaviours
- A random ensemble can go from *appetitive* to *aversive* to *appetitive*
- Location of the ensemble in the piriform does not seem to matter
- Similar results in fly

Alternative models of connectivity

- **Random**: a reasonable interpretation of data
- **Secretly Structured**: a possible alternative

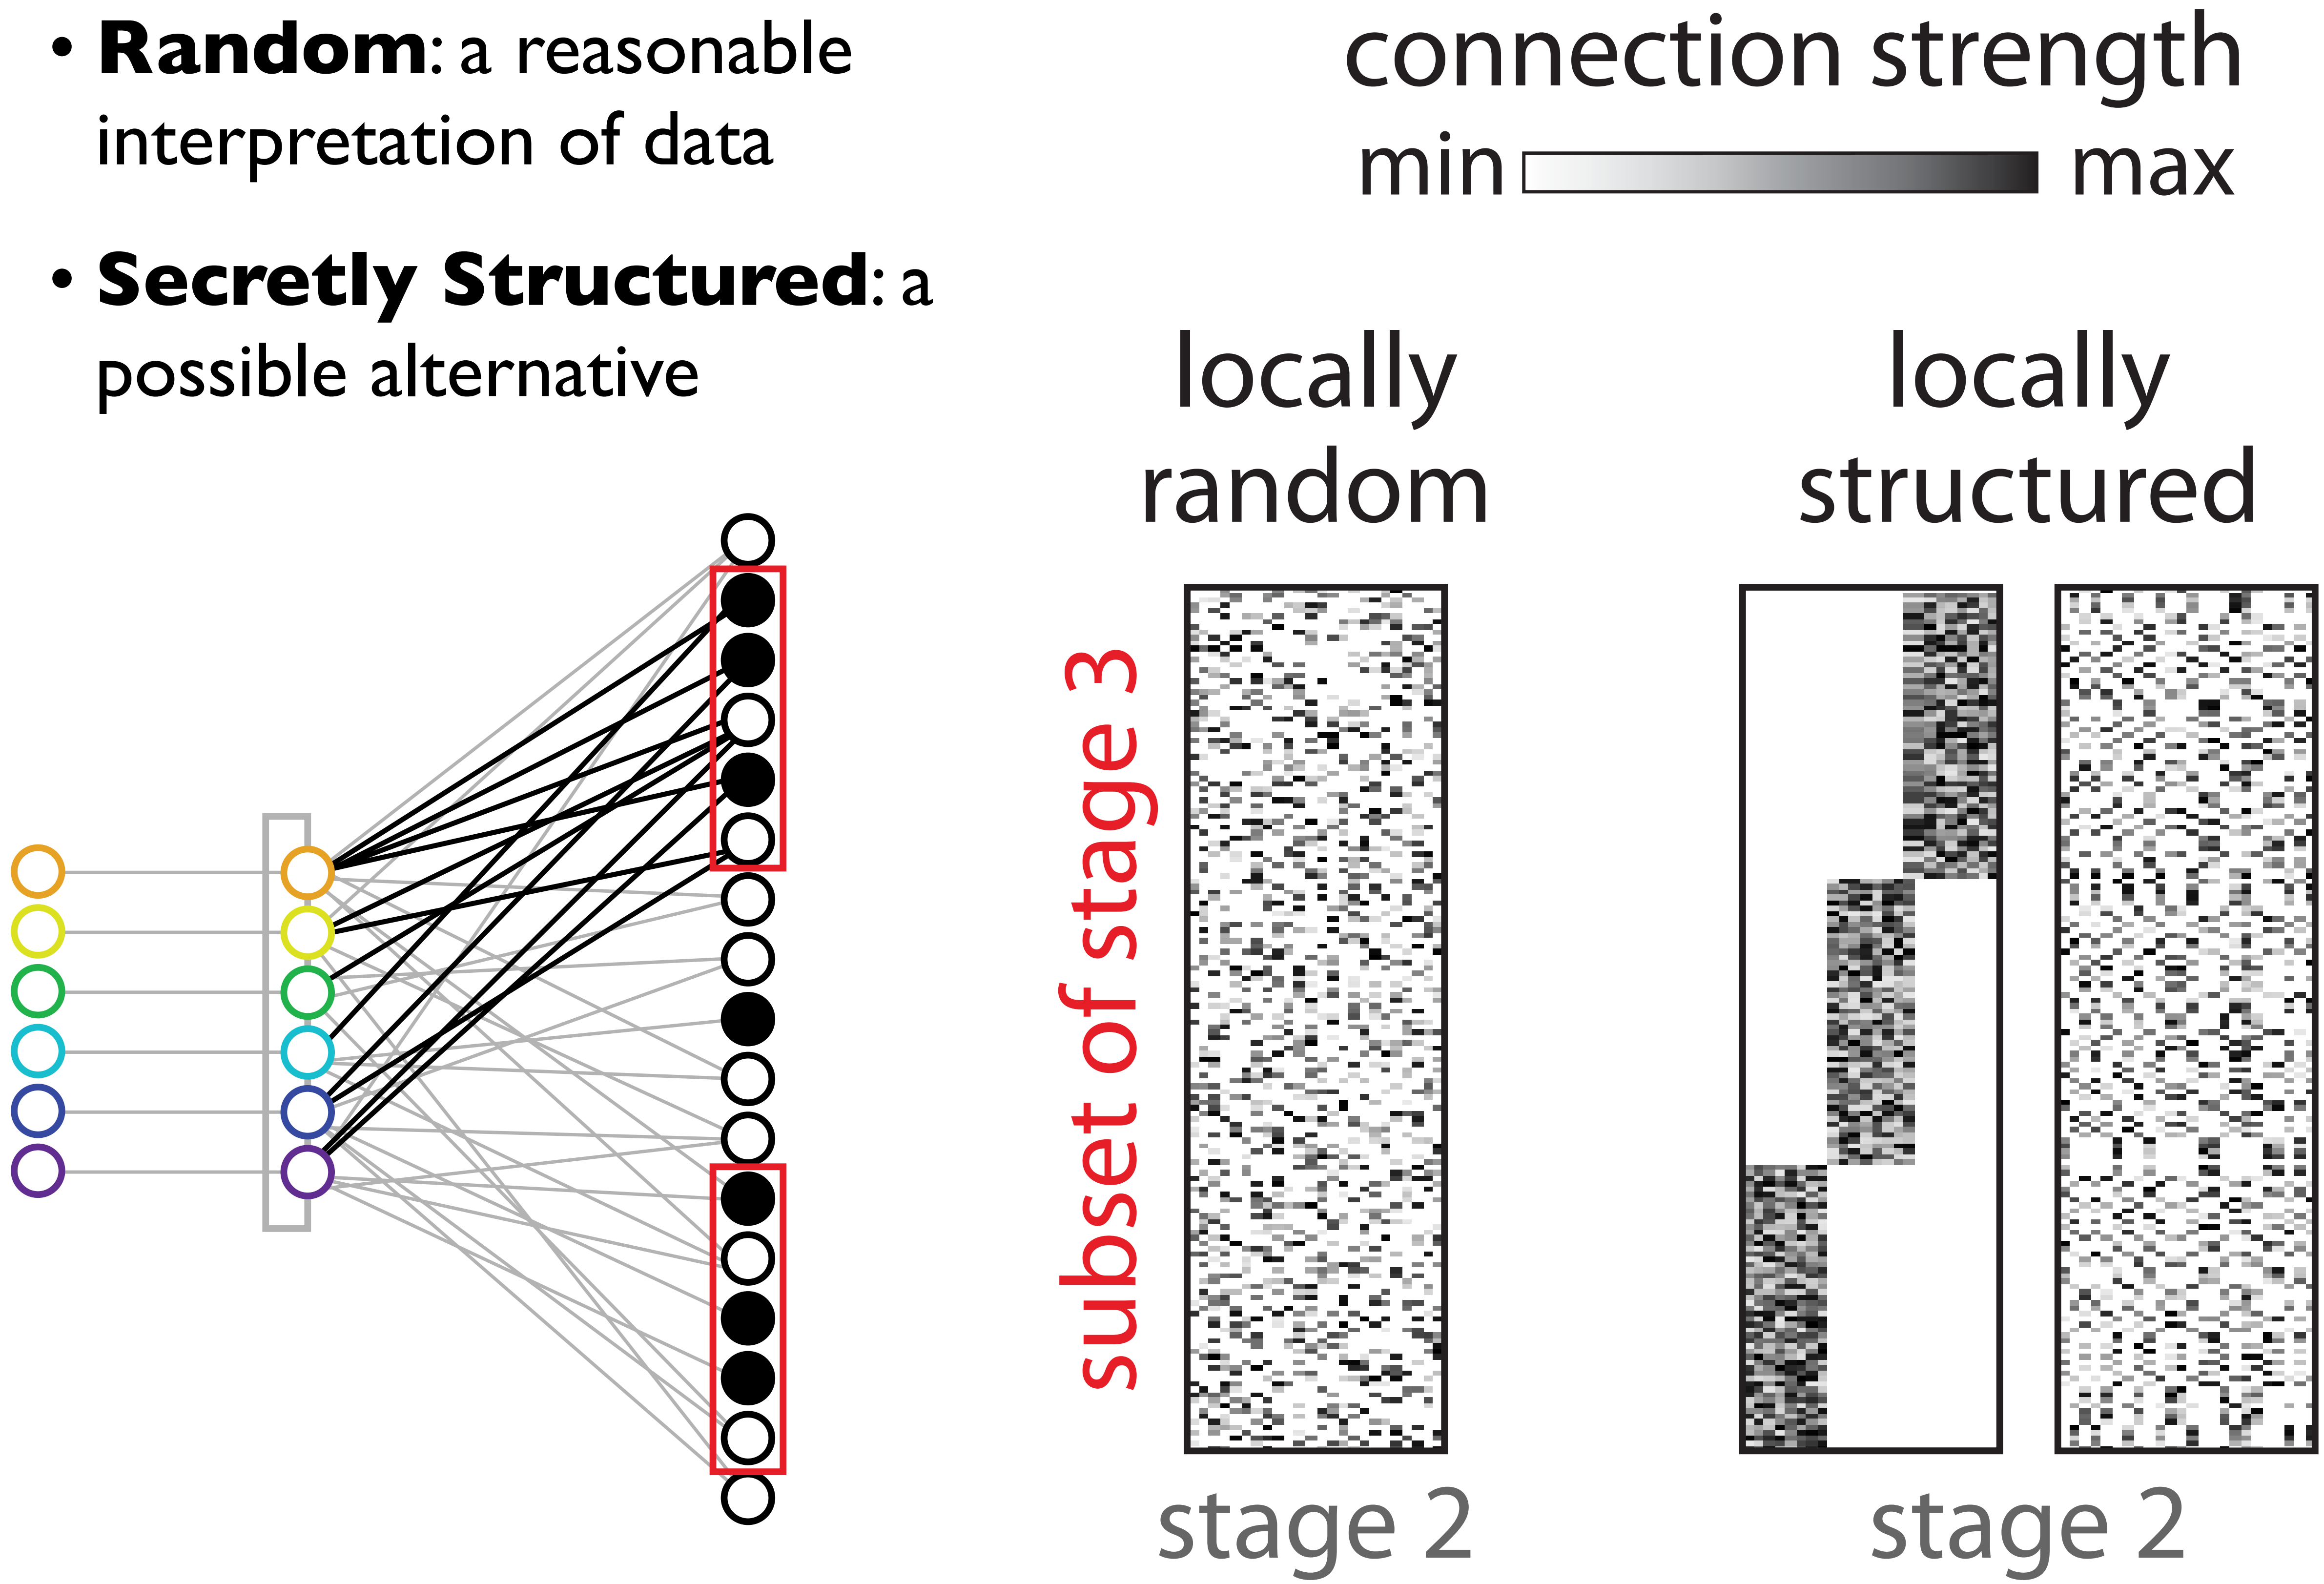

Stage 2: Decorrelation and diffusion of responses

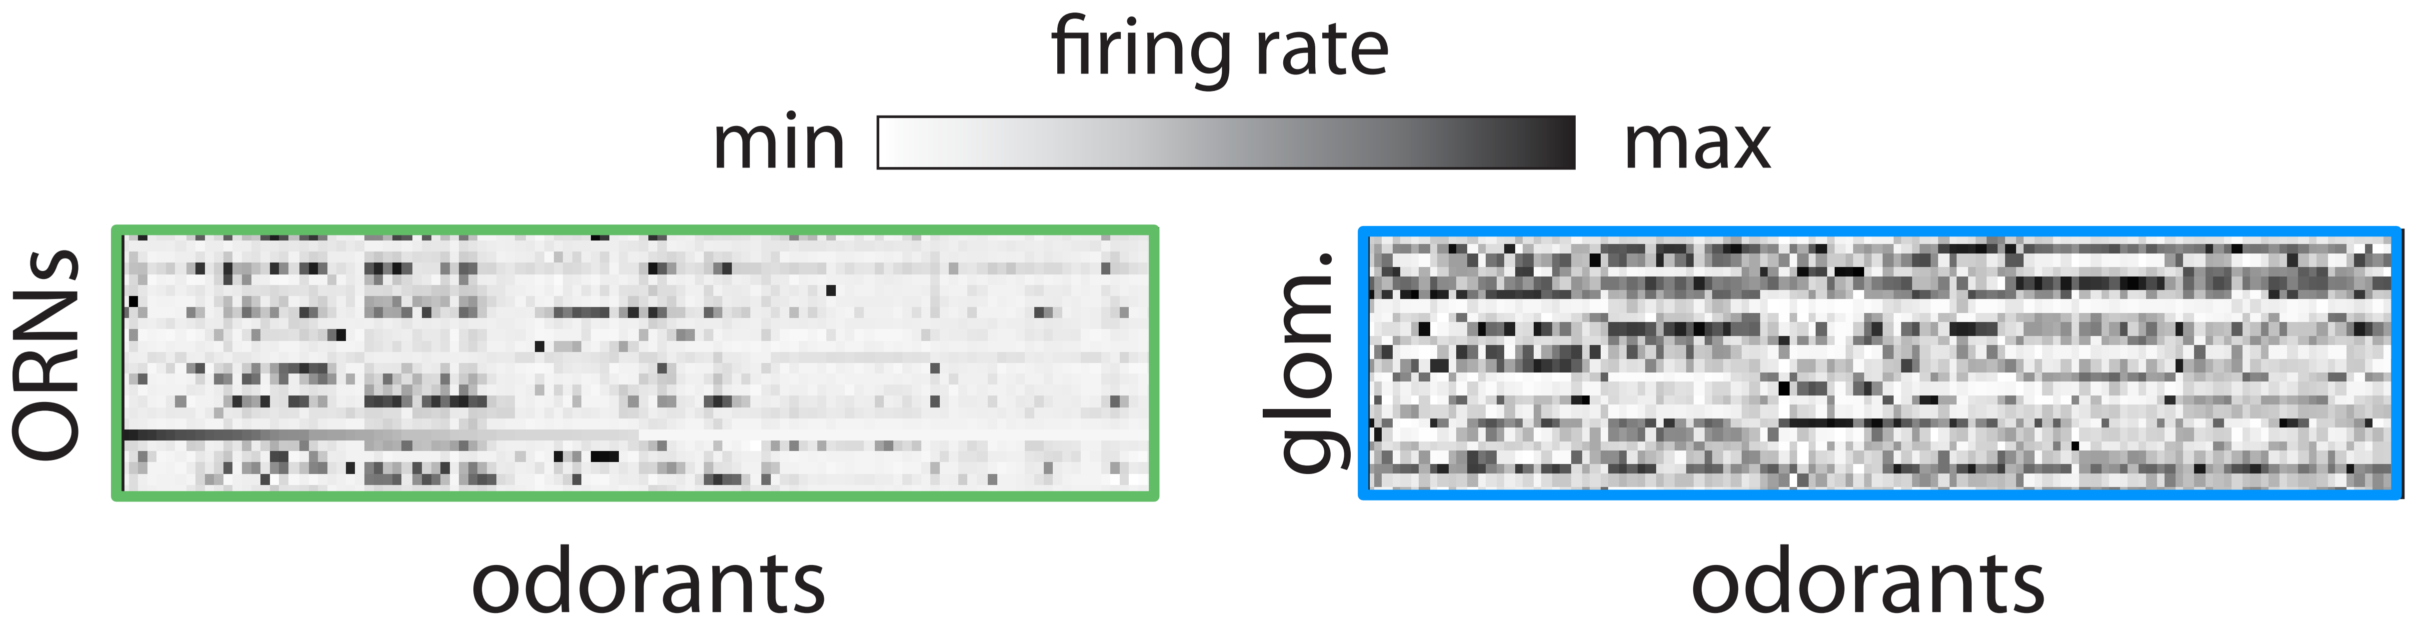

Odor receptor responses are nonlinearly transformed by circuitry in the *glomeruli* of the second stage. Empirical model (divisive normalization)

$$R_i^{AL} = \frac{R_{max} \cdot \left(R_i^{ORN}\right)^{1.5}}{\sigma^{1.5} + \left(R_i^{ORN}\right)^{1.5} + \left(m \cdot \sum_i R_i^{ORN}\right)^{1.5}}$$

Responses are spread out more evenly and broadly

Transformed responses are easier to decode

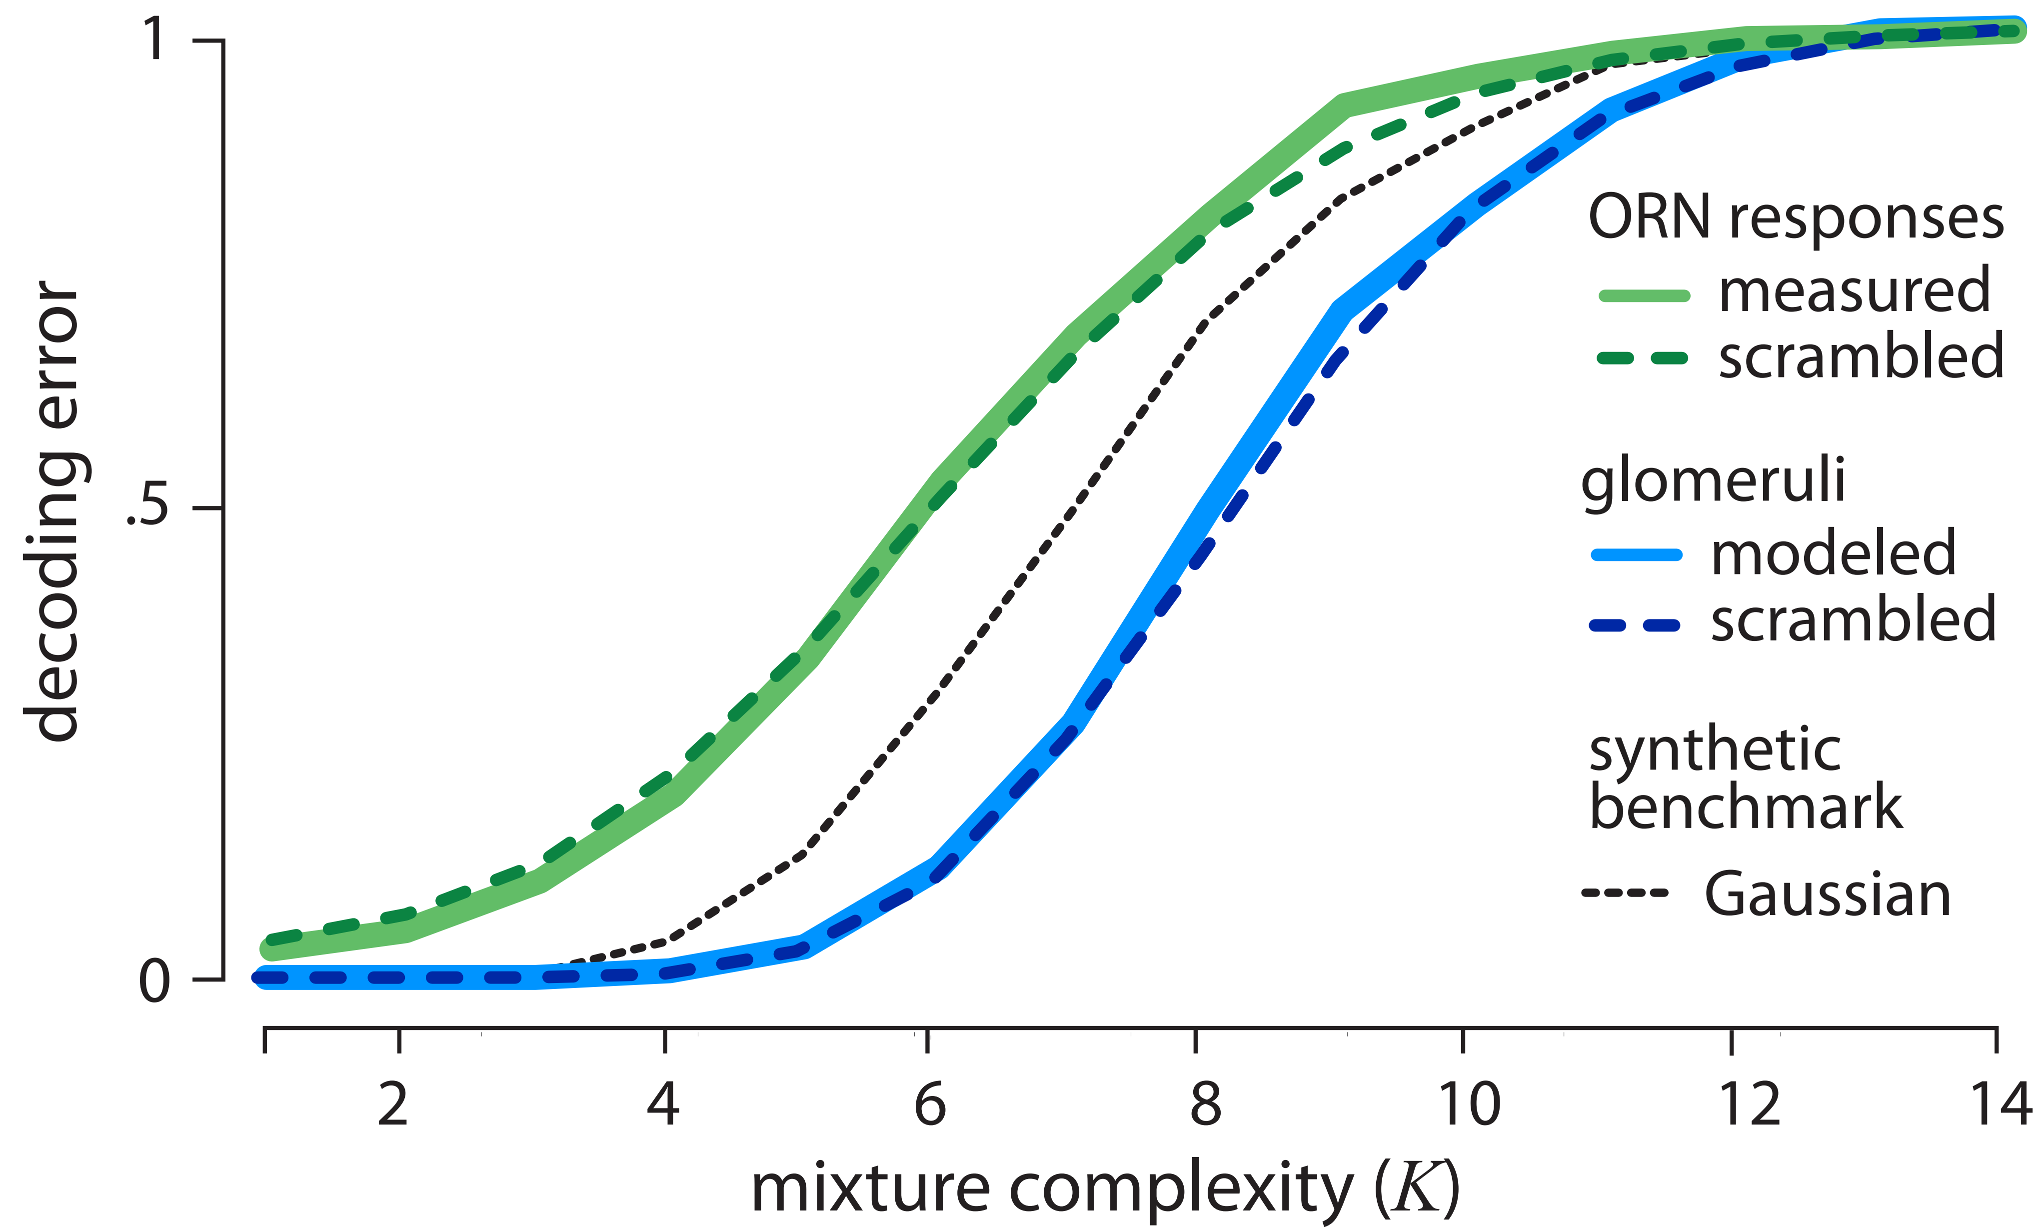

Scrambling leaves results unchanged — only the overall distribution of responses matters

Animals can reversibly learn arbitrary odor associations

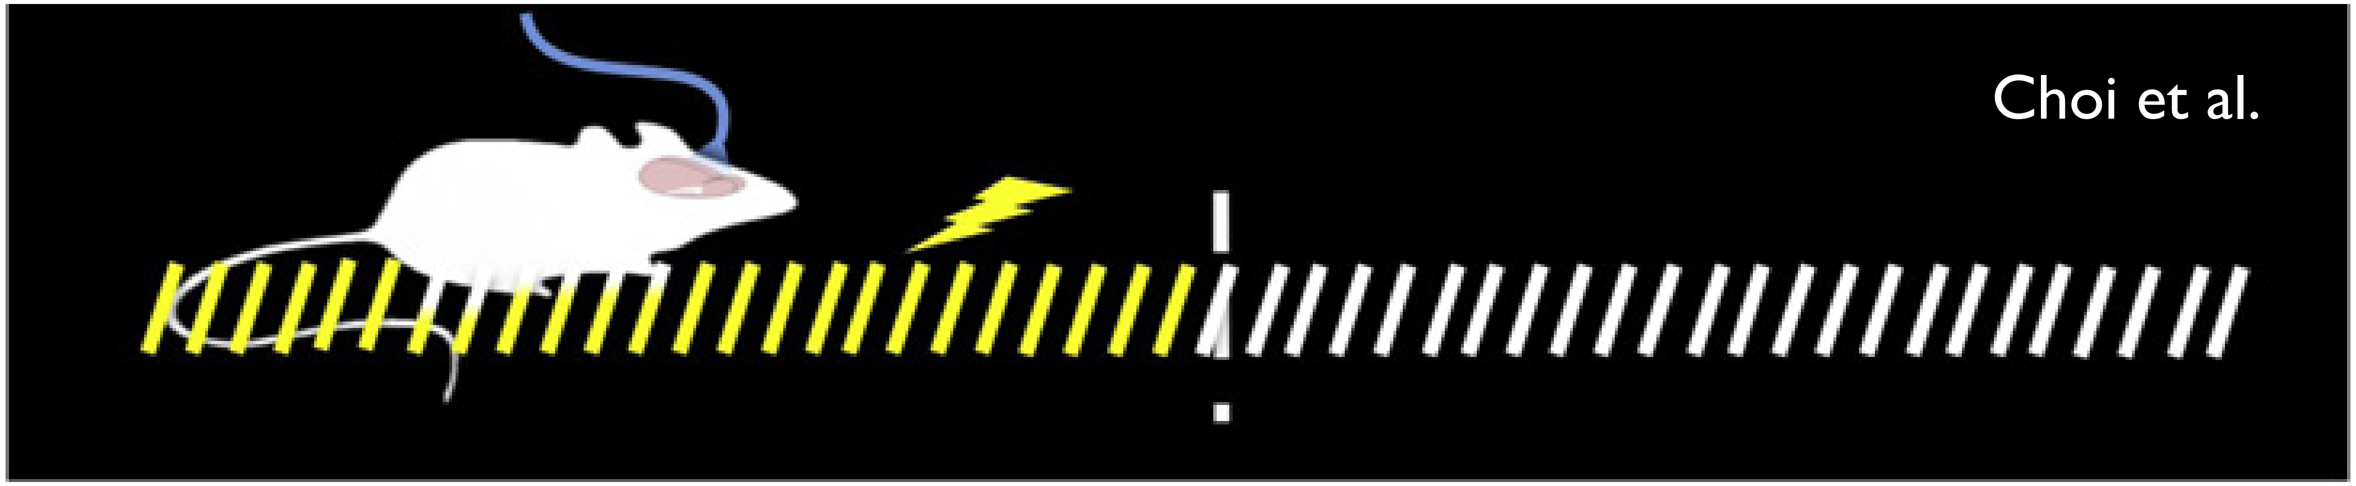

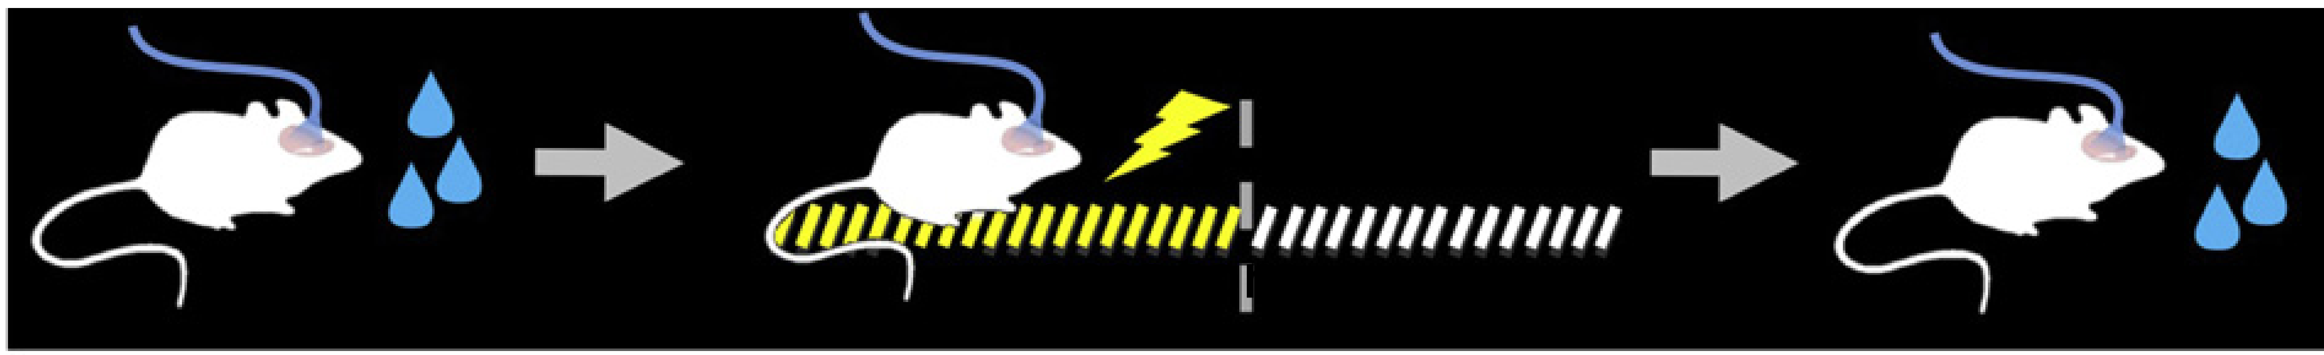

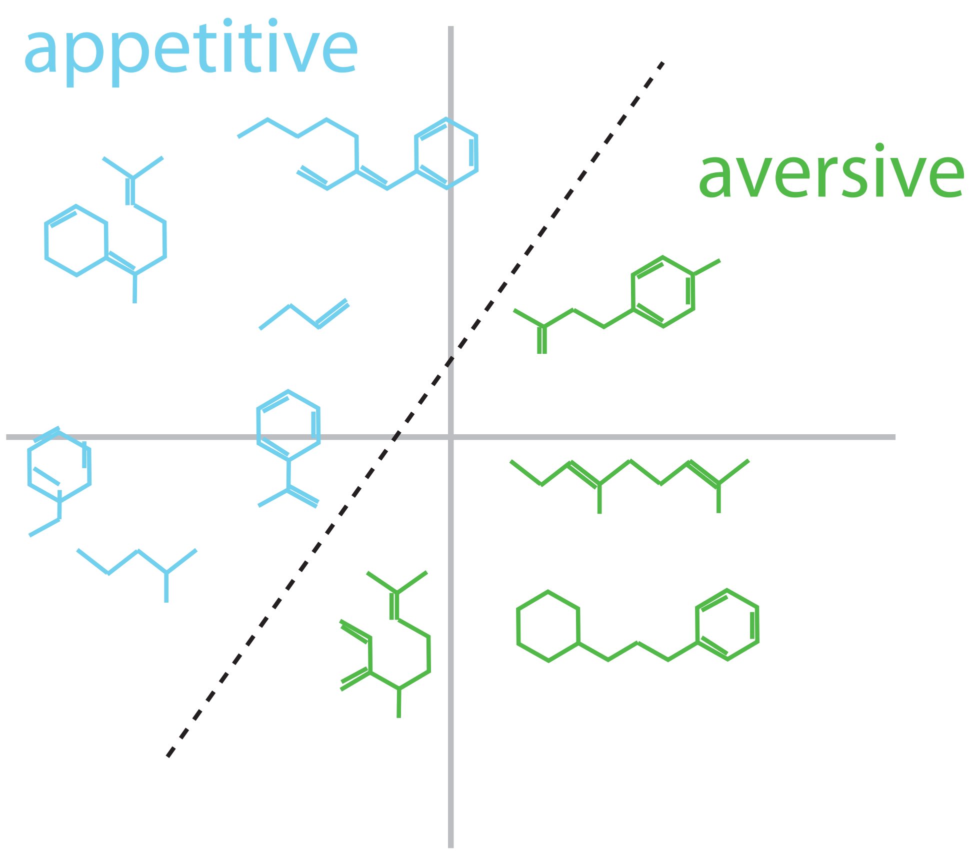

Can a linear classifier (model of a simple readout neuron) learn arbitrary assignments of "appetitive" and "aversive" classes from the neural responses?

Paradox!? Poor linear classification performance

Assign labels "appetitive" or "aversive" to 50% of 5-component odors and classify linearly.

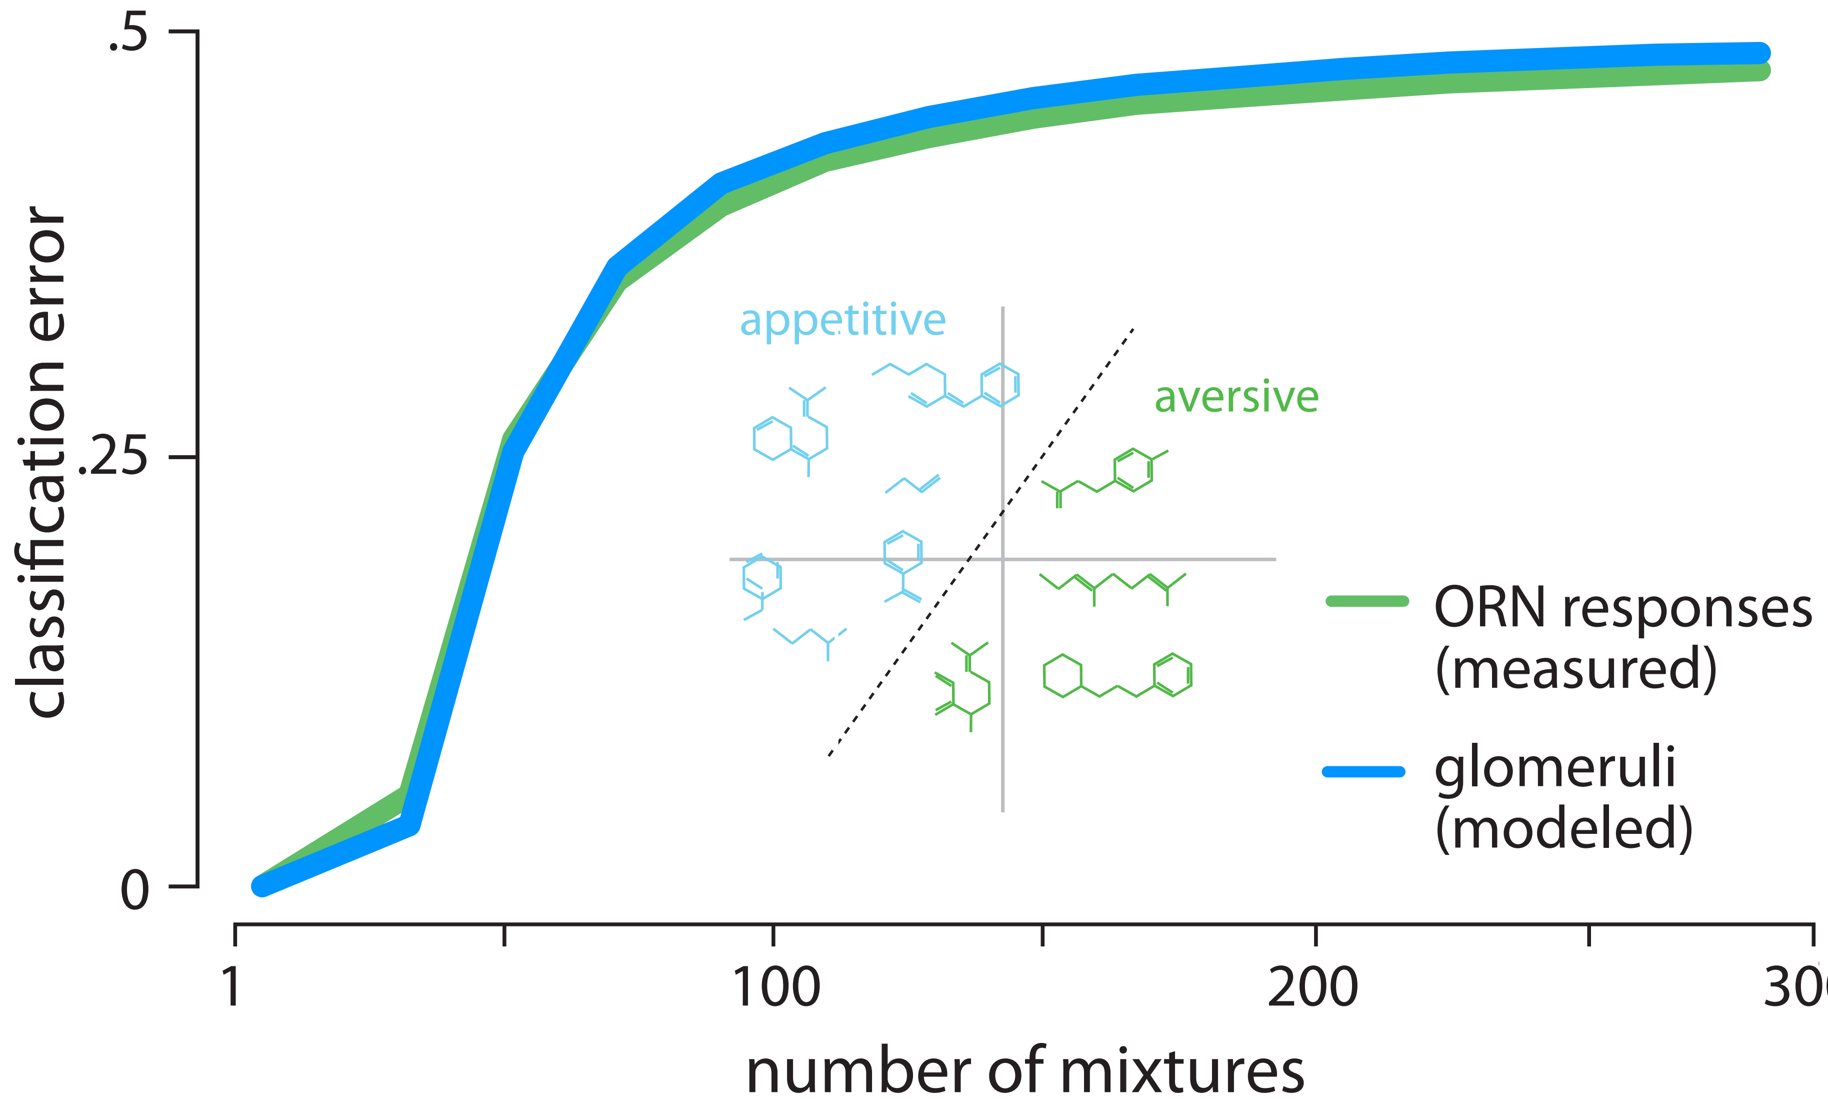

Classification performance with 24 receptors is close to chance with just 100 mixtures. Thus, information about odors needs to be "untangled" for linear classification.

Model task: two-way classification of odor classes

Assign labels "appetitive" or "aversive" to 50% of mixtures and classify linearly.

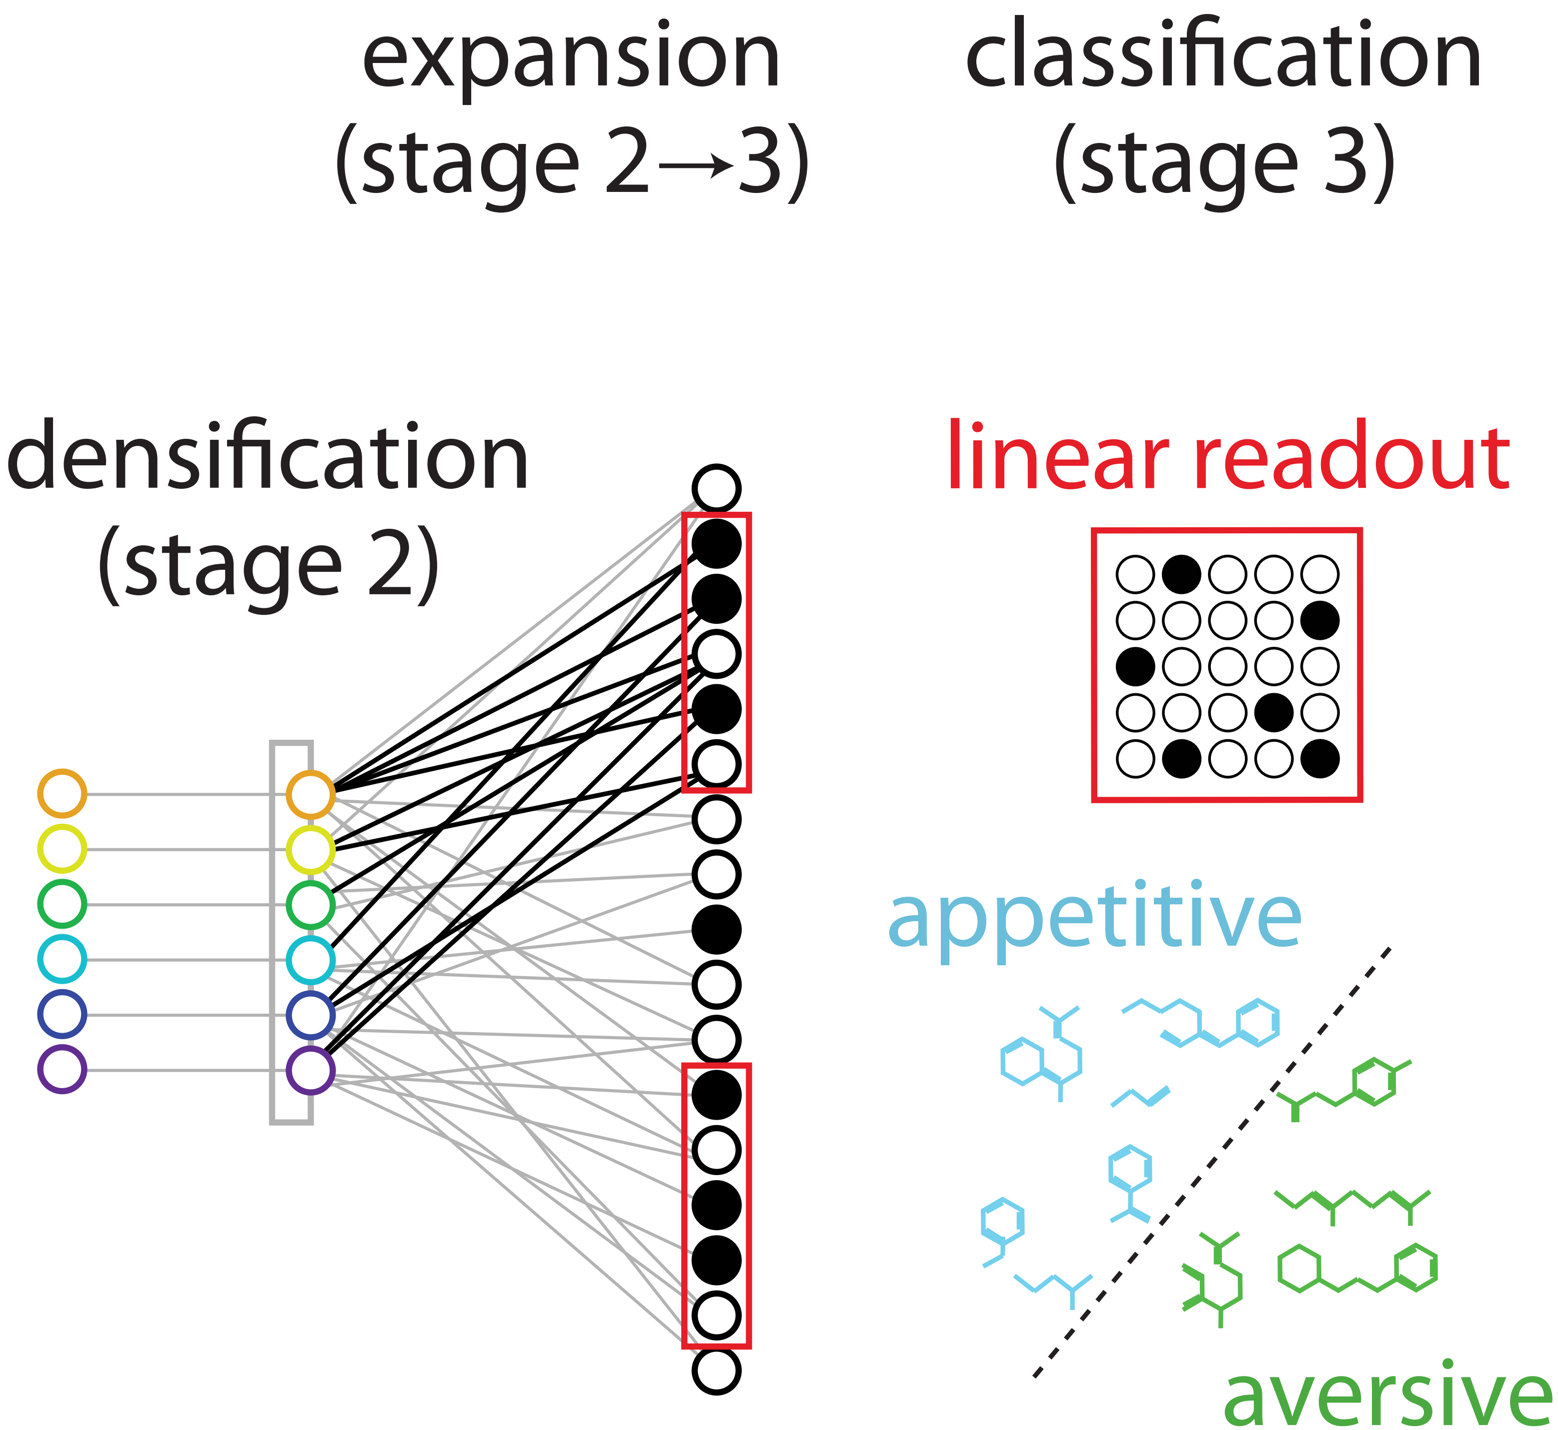

IDEA: A disordered (random) projection into a higher dimension should make a simple readout (linear classification) easier

Stage 3: Excellent classification

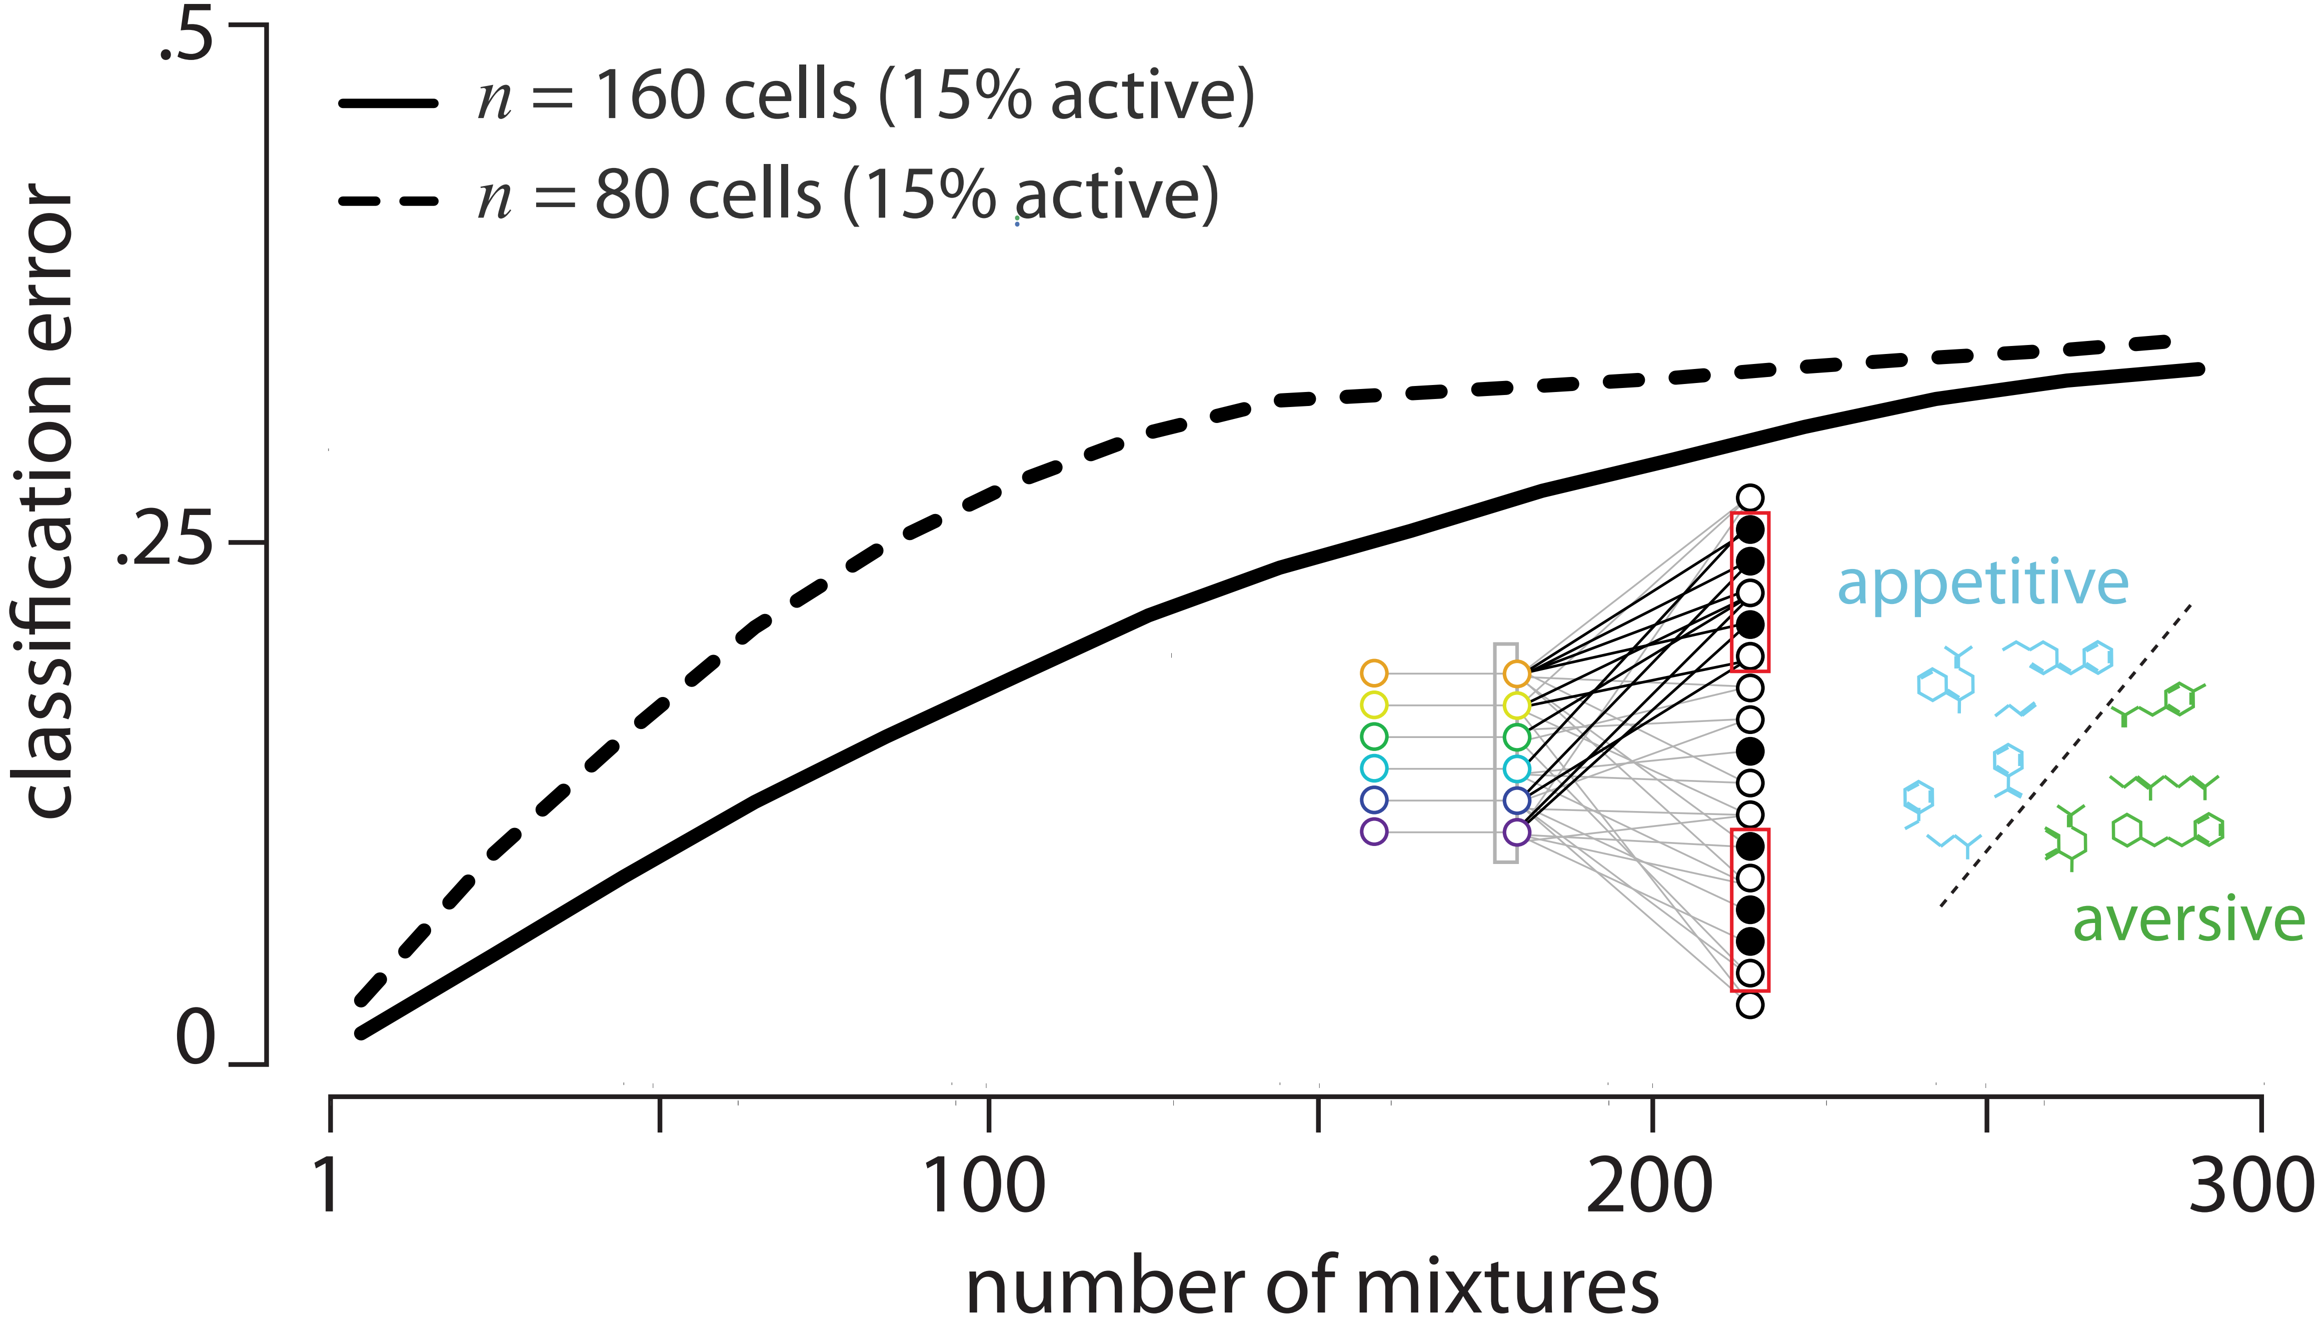

Disorder and the sense of smell

Olfactory circuits employ two kinds of disorder to smell in the real world

- disordered sensing compresses chemical space into receptor space
- decorrelation and disordered expansion reformat information for flexible learning
- Disorder as an adaptation to the olfactory world