AMS Western Sectional 2019

NEURAL REPRESENTATIONS, TOPOLOGY, AND CATEGORY THEORY

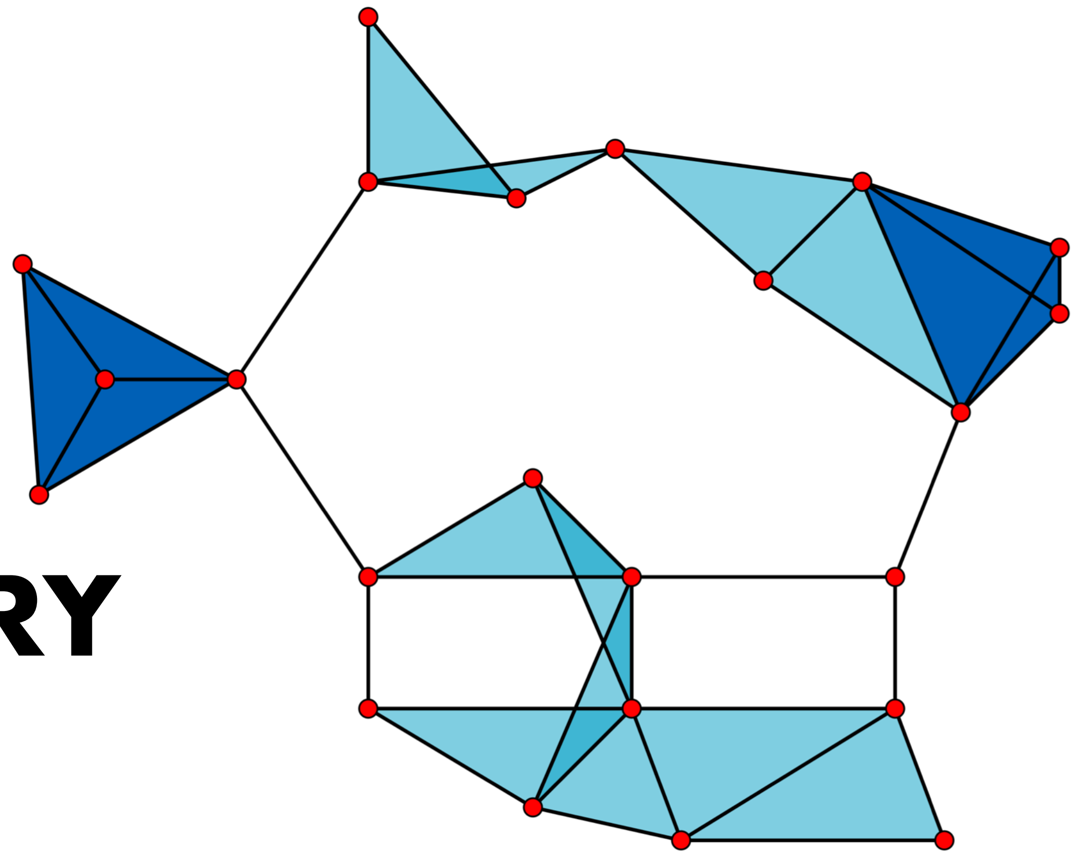

Brad Theilman | Gentner Lab | UC San Diego

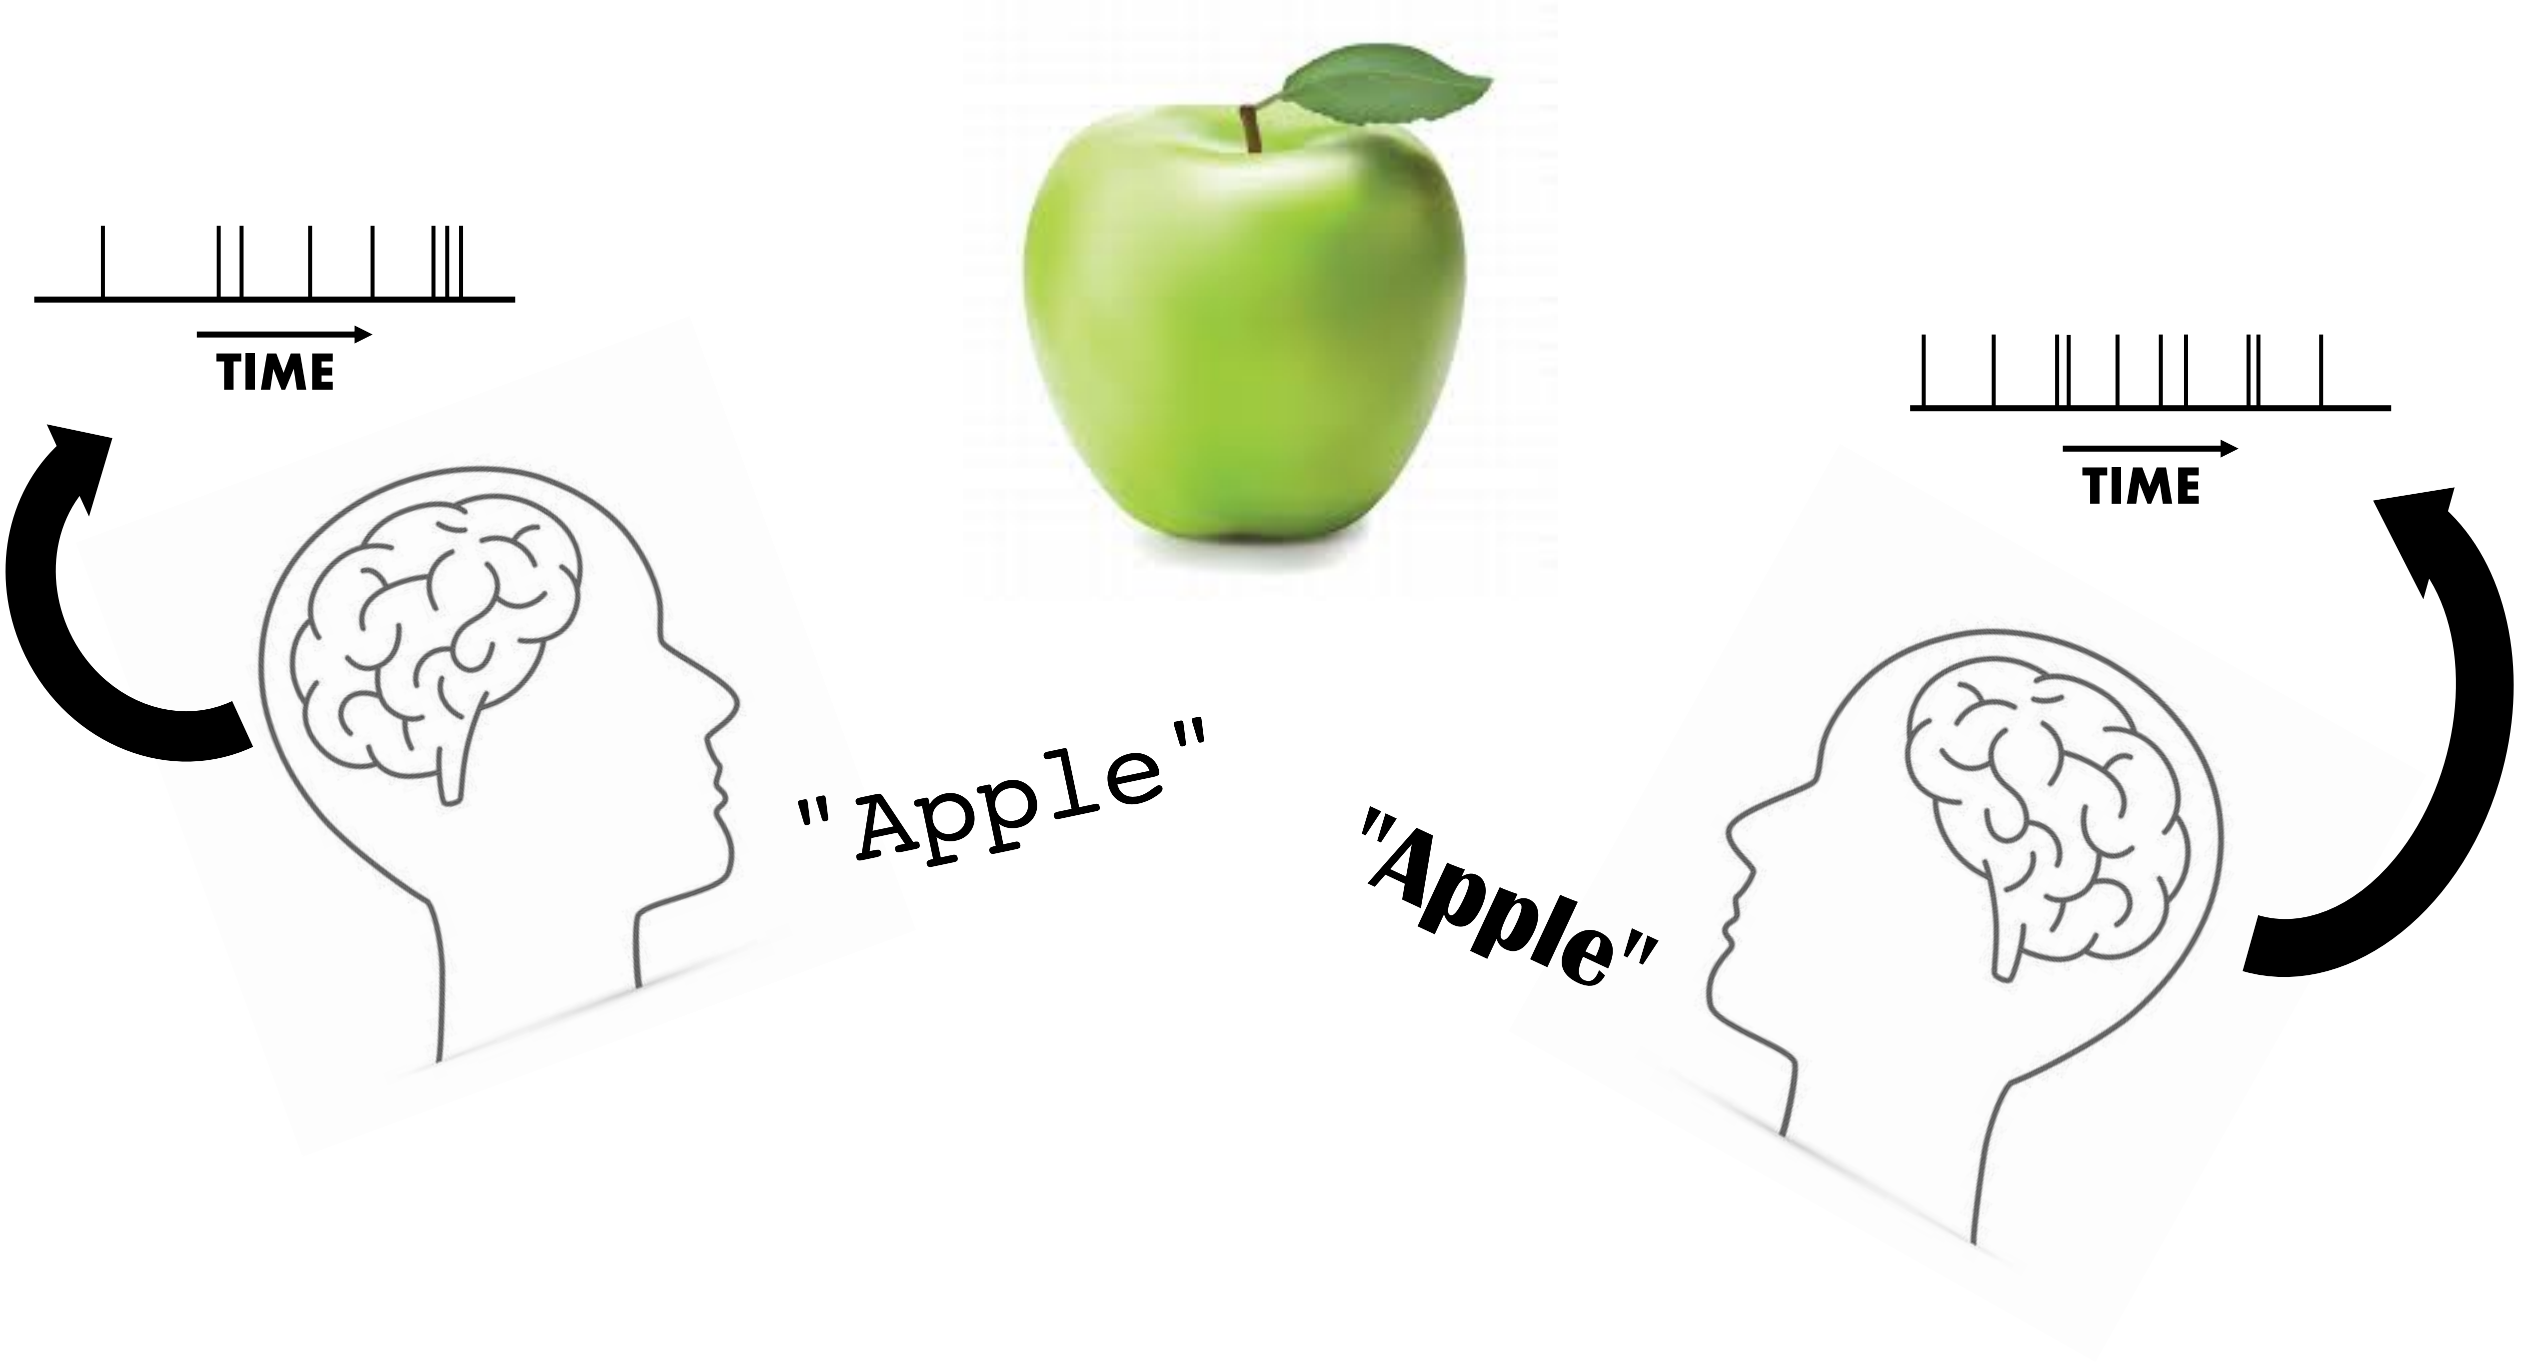
TIME
"Apple"
"Apple"
TIME

NEURAL REPRESENTATIONS

"A neural representation is a pattern of neural activity that stands for some environmental feature in the internal workings of the brain."

Neural Representation. A Survey-Based Analysis of the Notion

 Oscar Vilarroya[1,2*]

[1]Departament de Psiquiatria i Medicina Legal, Universitat Autònoma de Barcelona, Barcelona, Spain

[2]Institut Hospital del Mar d'Investigacions Mèdiques (IMIM), Barcelona, Spain

EUROPEAN STARLING

Complex, natural, acoustic signals

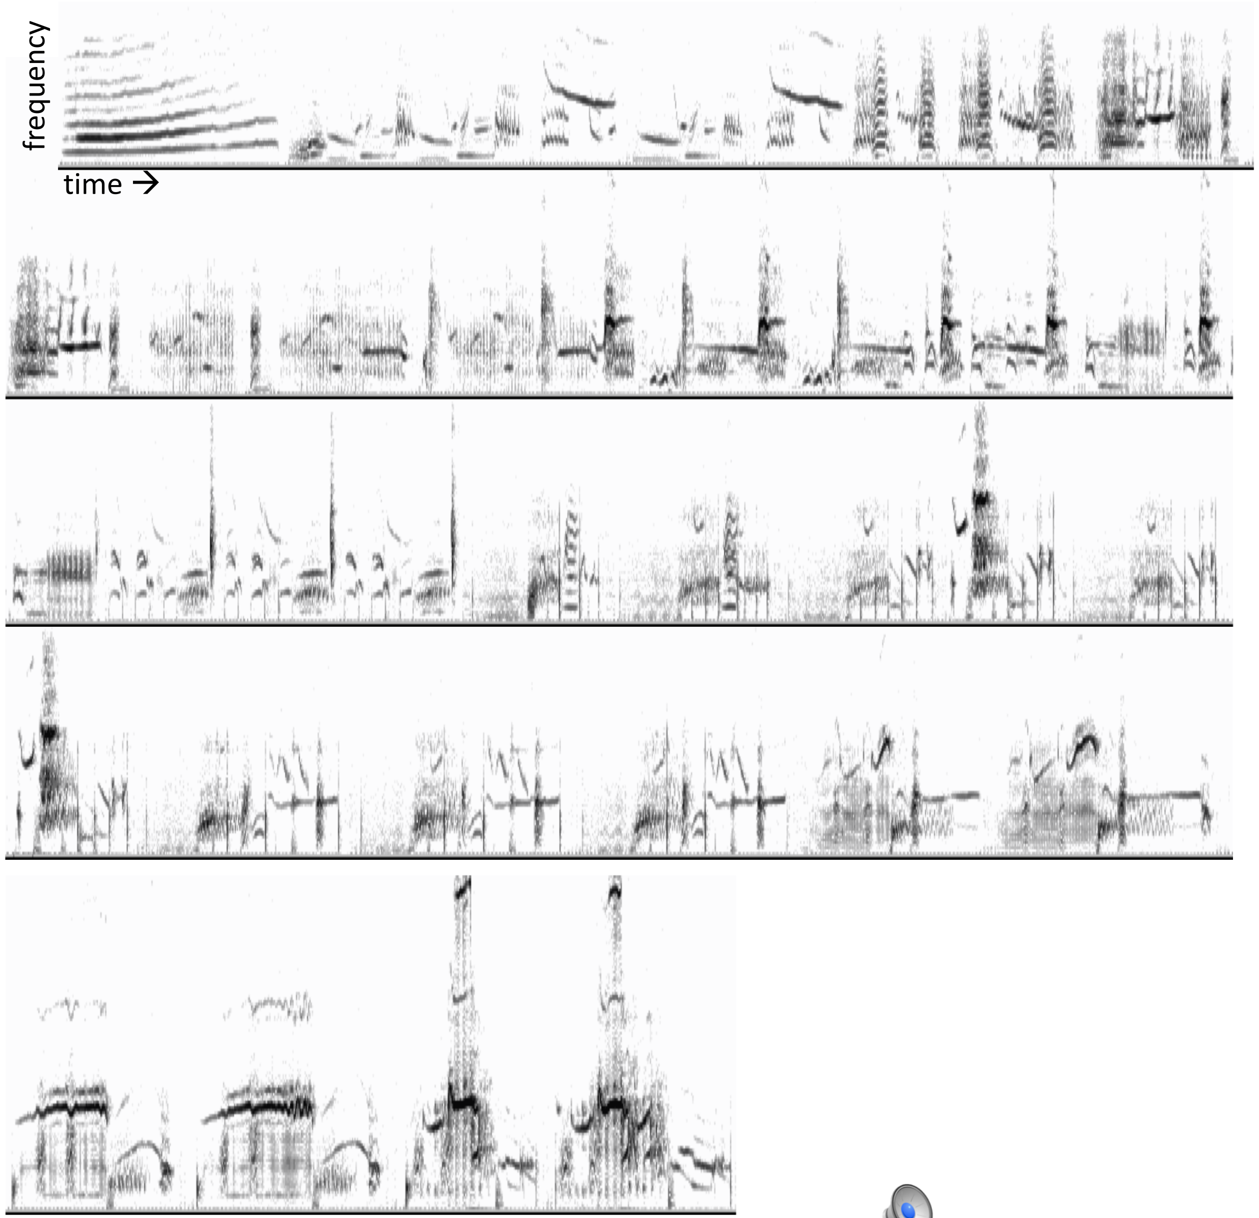

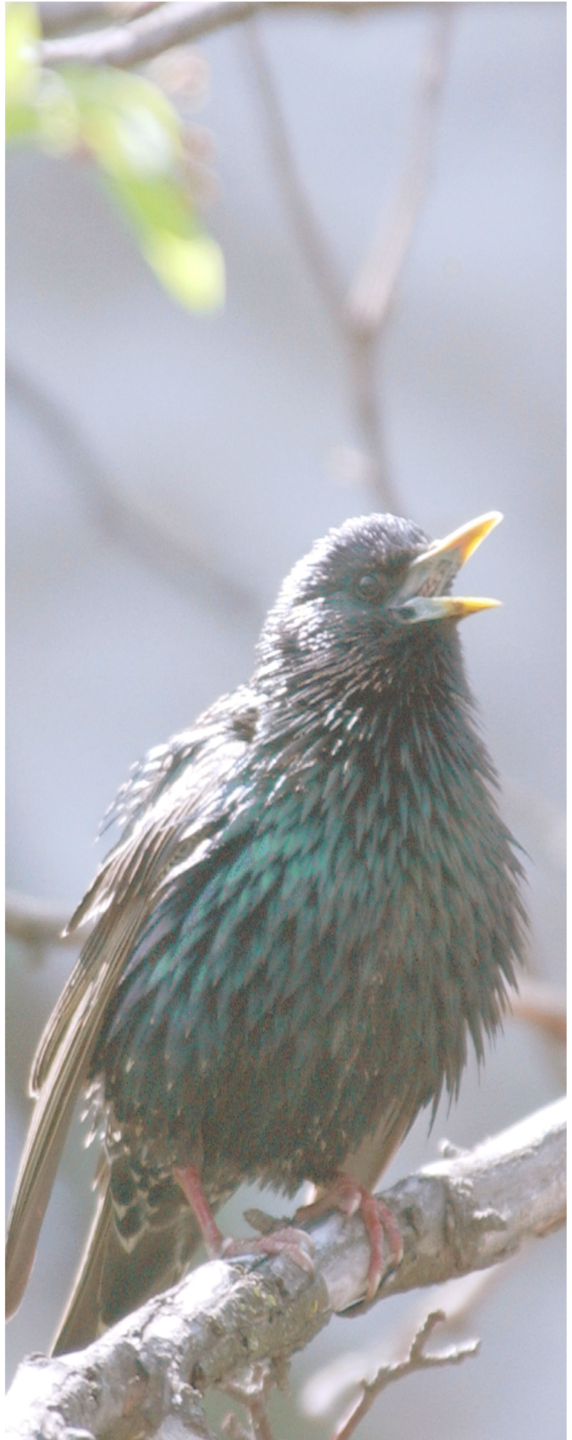

Electrophysiological Recording

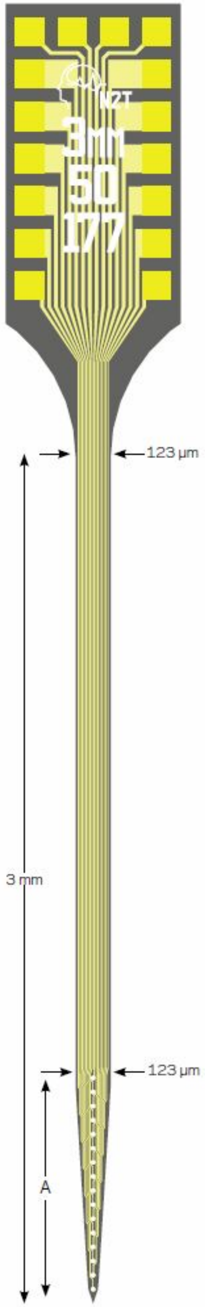

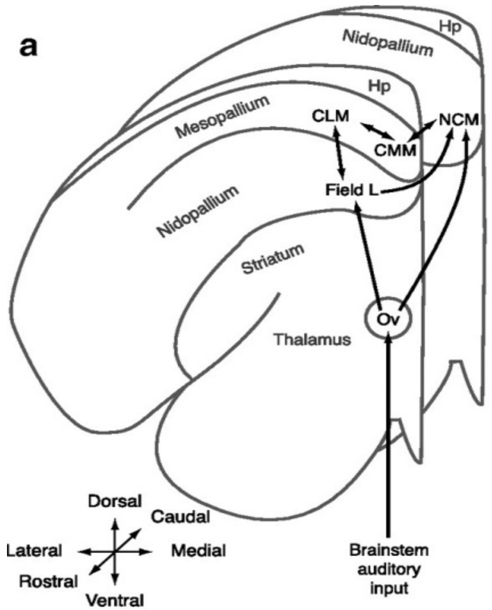

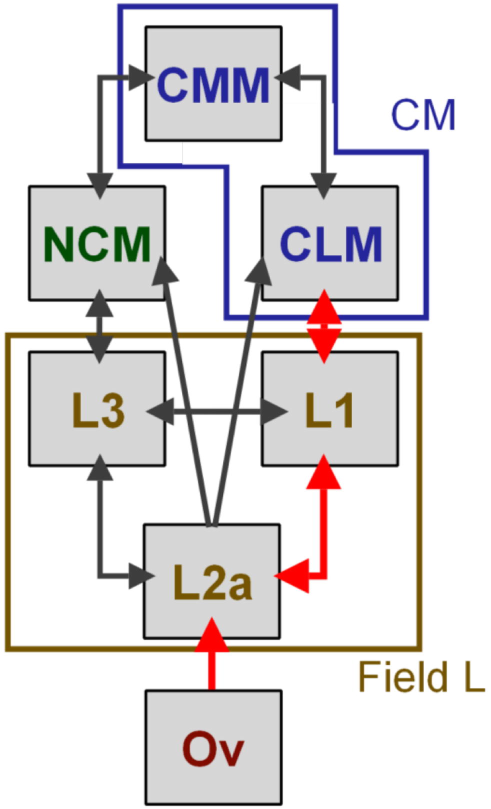

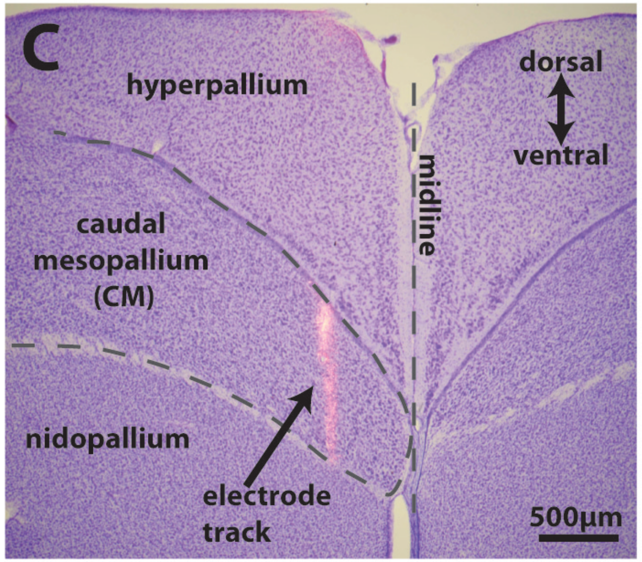

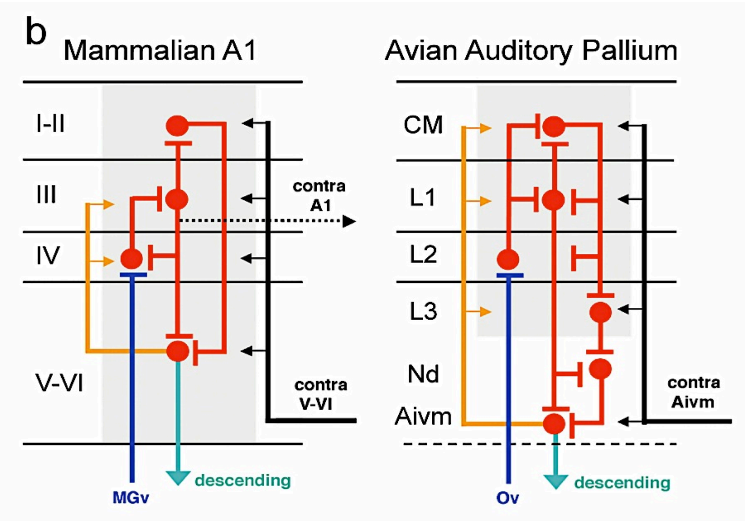

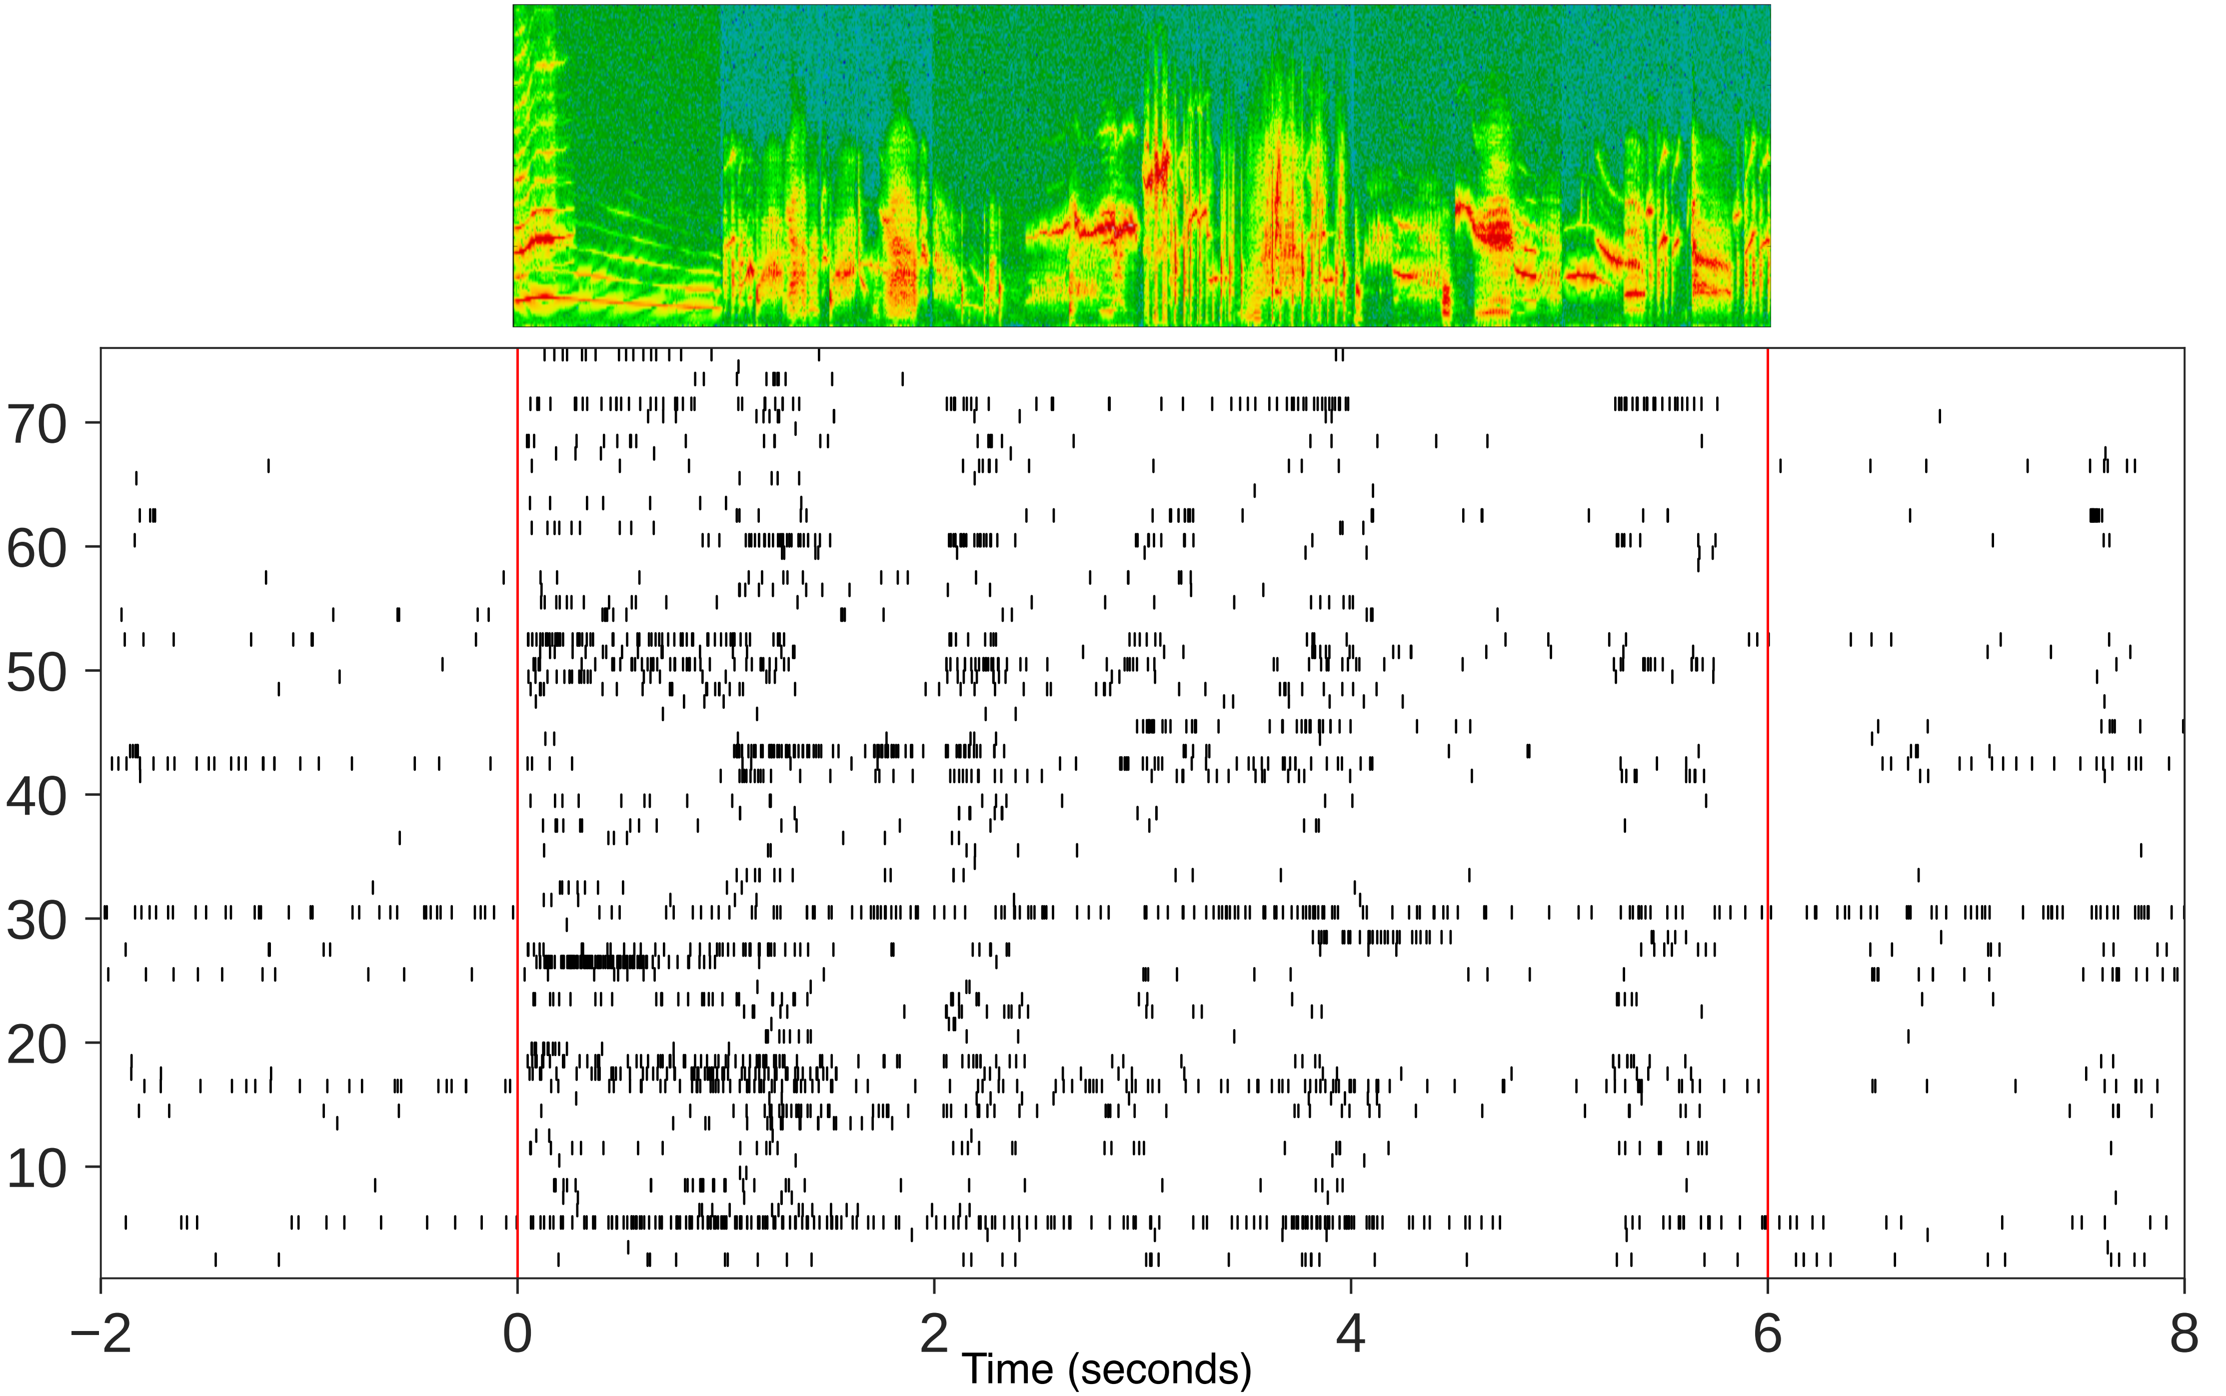

70
60
50
40
30
20
10
−2
0
2
4
6
8
Time (seconds)

SINGLE NEURON, MULTIPLE PRESENTATIONS OF FOUR STIMULI

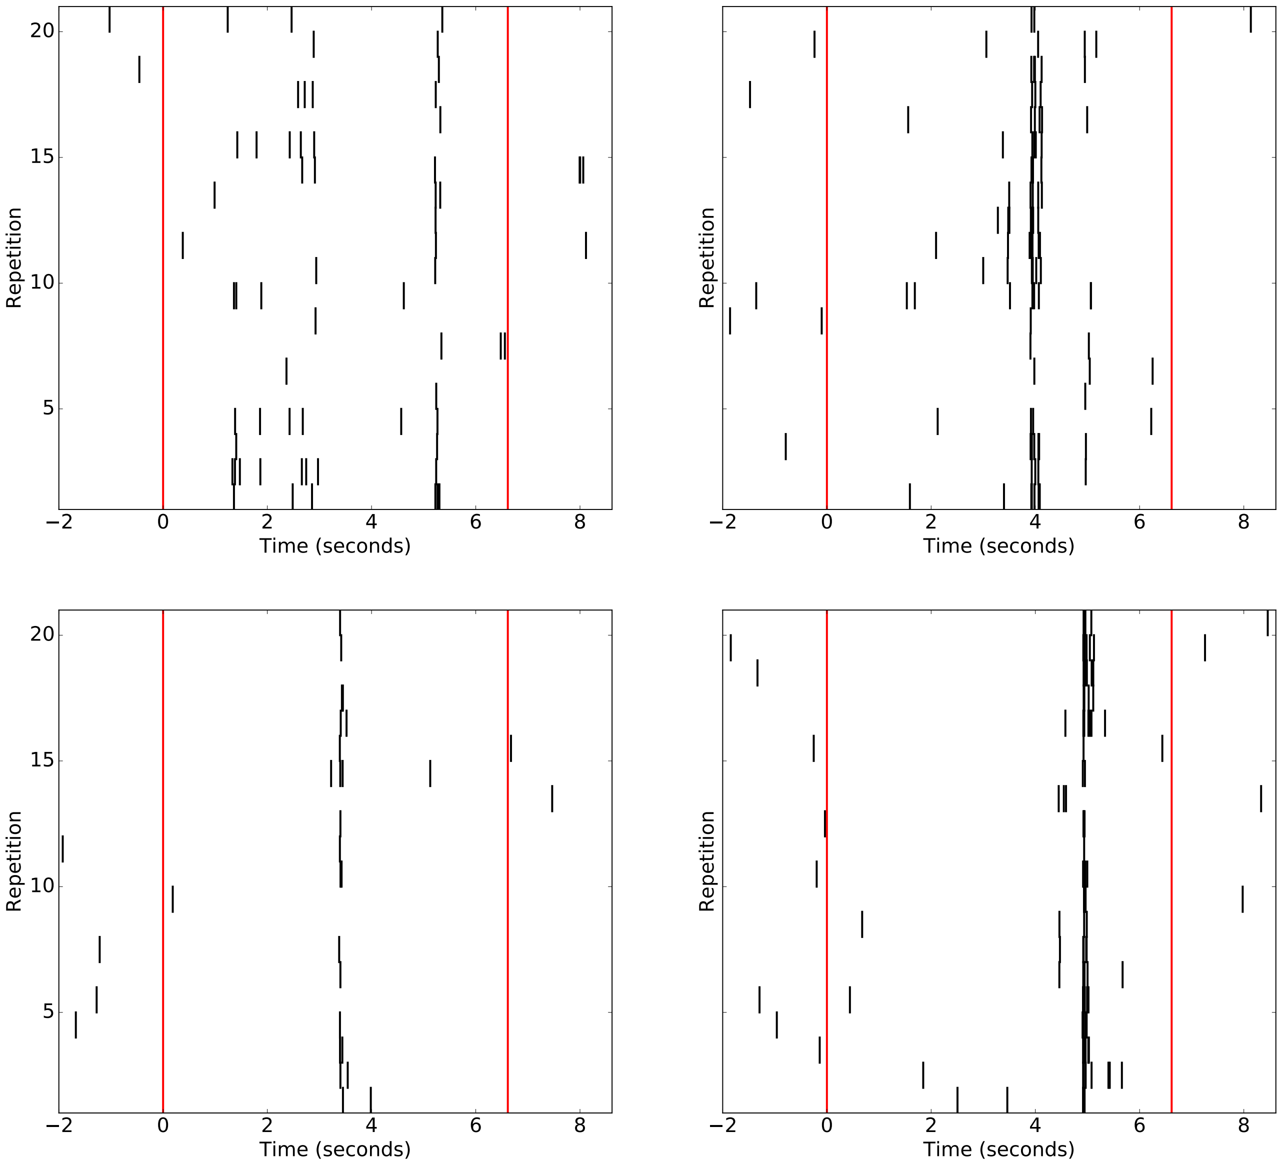

RECEPTIVE FIELDS

Which components of a stimulus cause a neuron to spike?

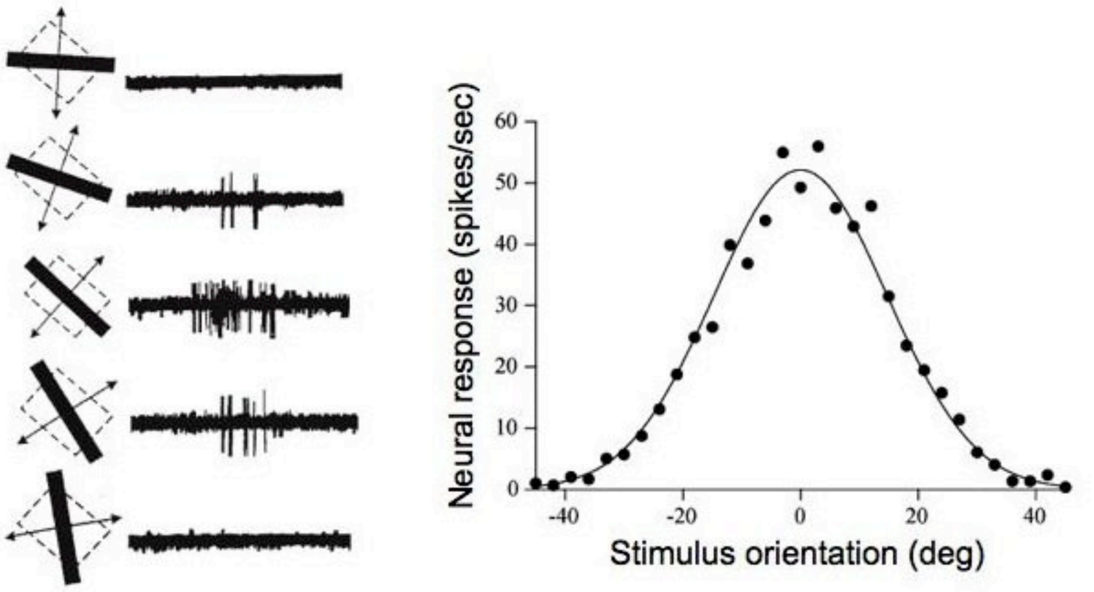

Hubel & Wiesel, 1968

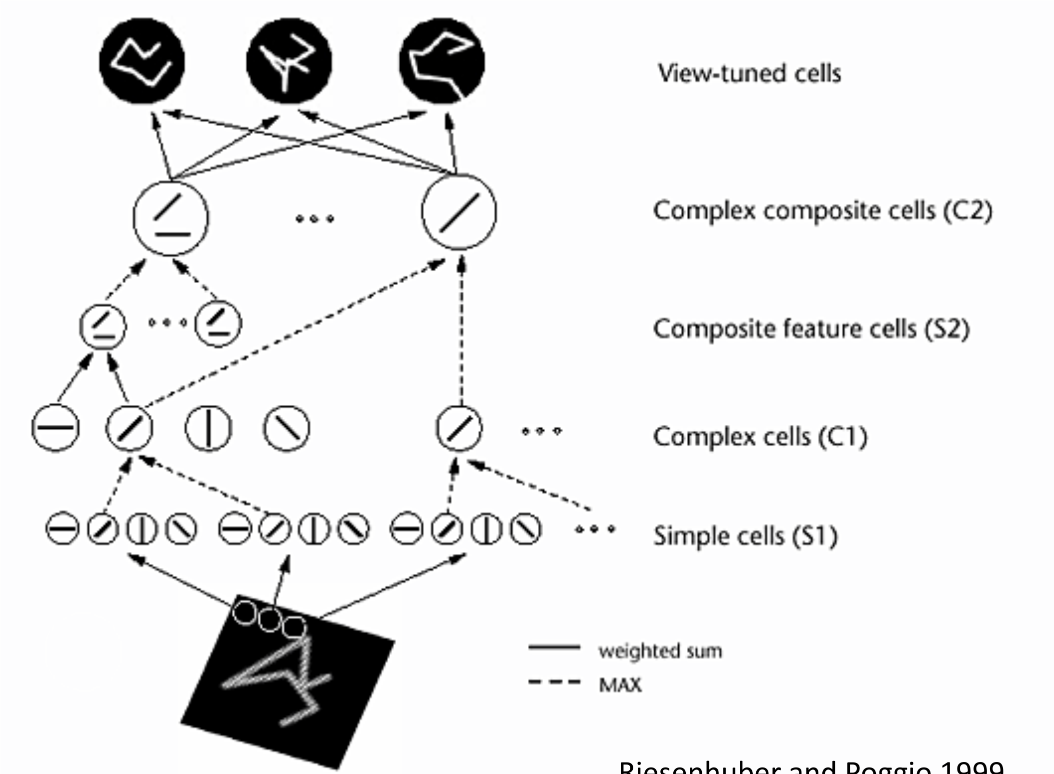

Riesenhuber and Poggio 1999

RECEPTIVE FIELDS

Which components of a stimulus cause a neuron to spike?

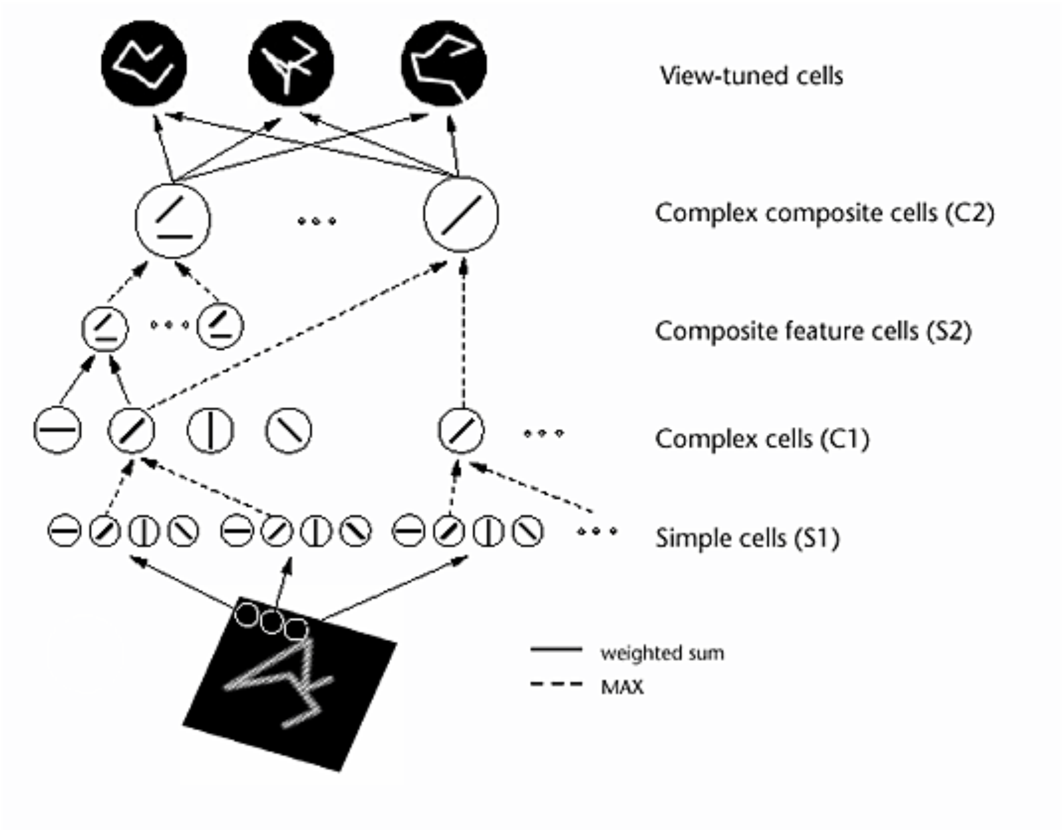

Field L ('A1')

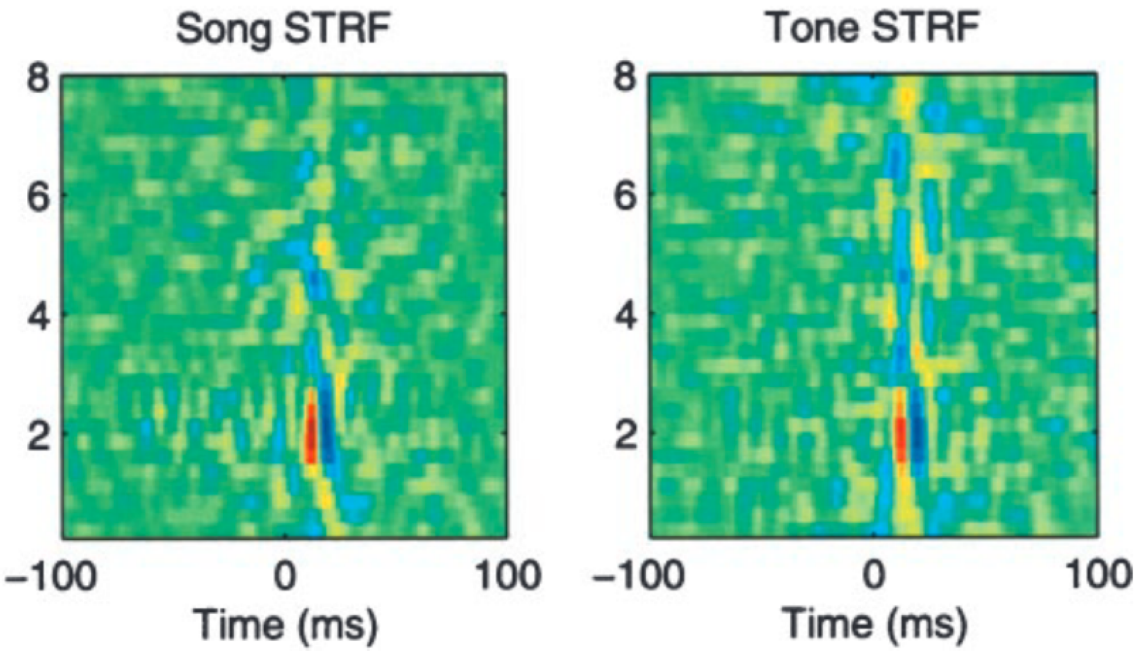

Riesenhuber and Poggio 1999

RECEPTIVE FIELDS

Which components of a stimulus cause a neuron to spike?

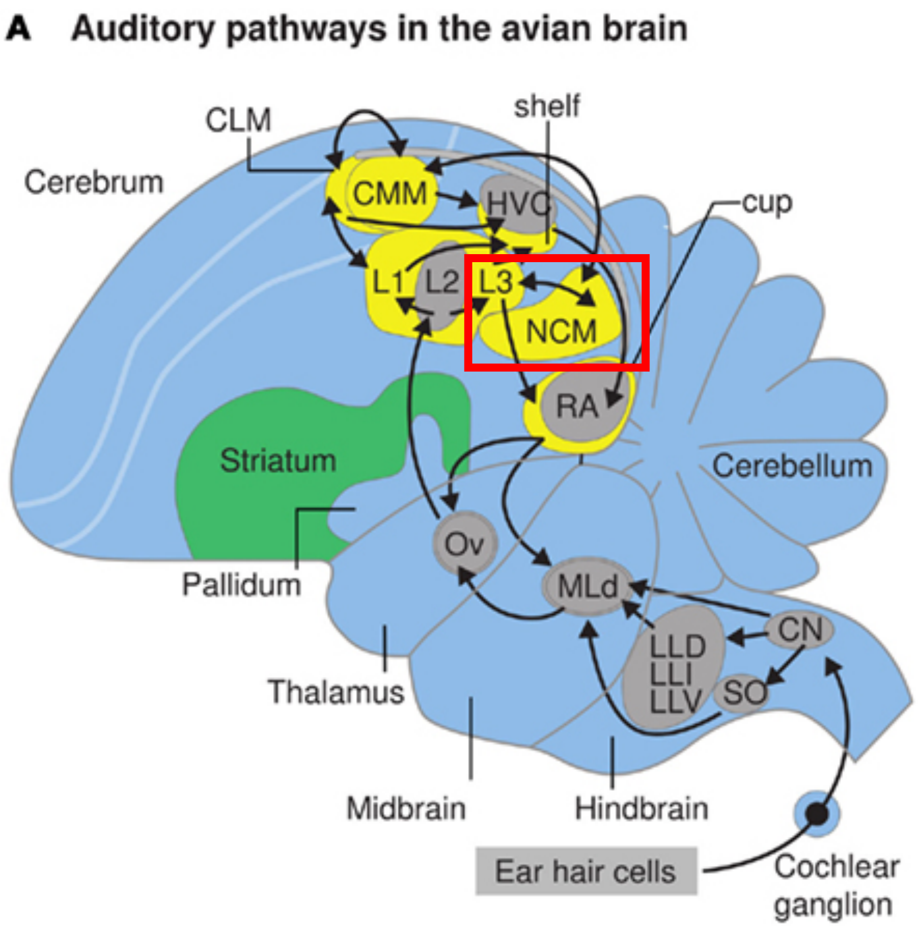

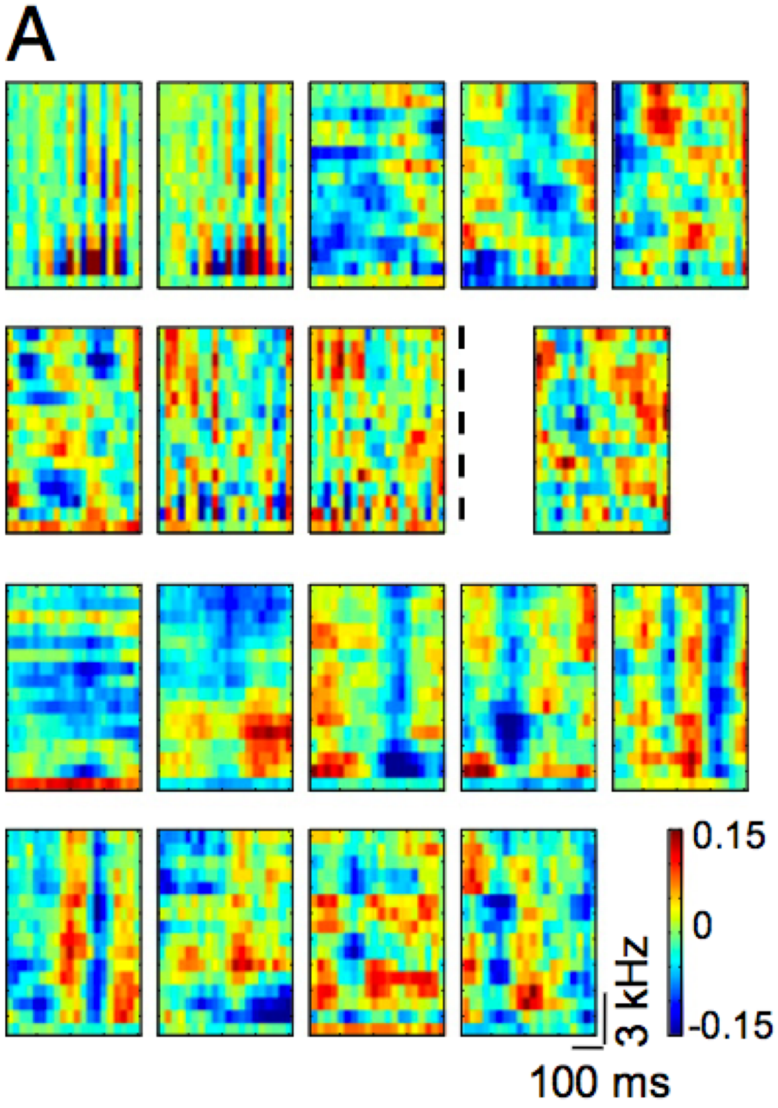

NCM

Kozlov & Gentner

RECEPTIVE FIELDS

Which components of a stimulus cause a neuron to spike?

A single neuron can be selective for multiple features.

Disambiguation must come from population.

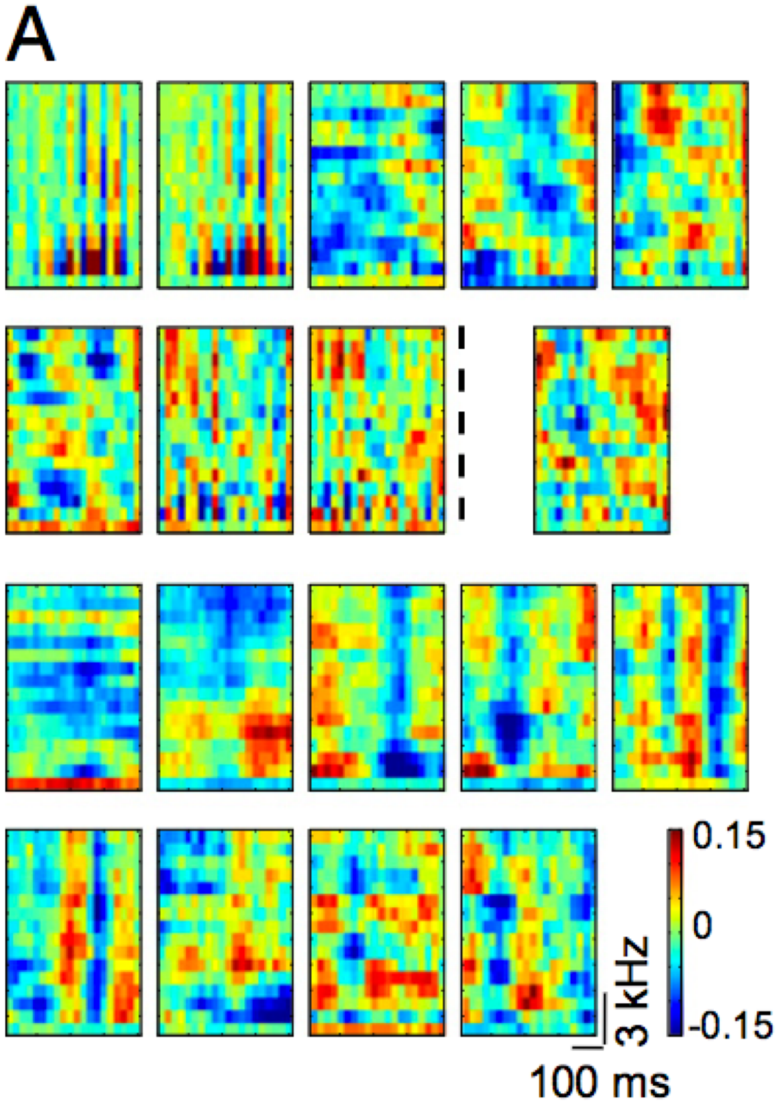

NCM

Kozlov & Gentner

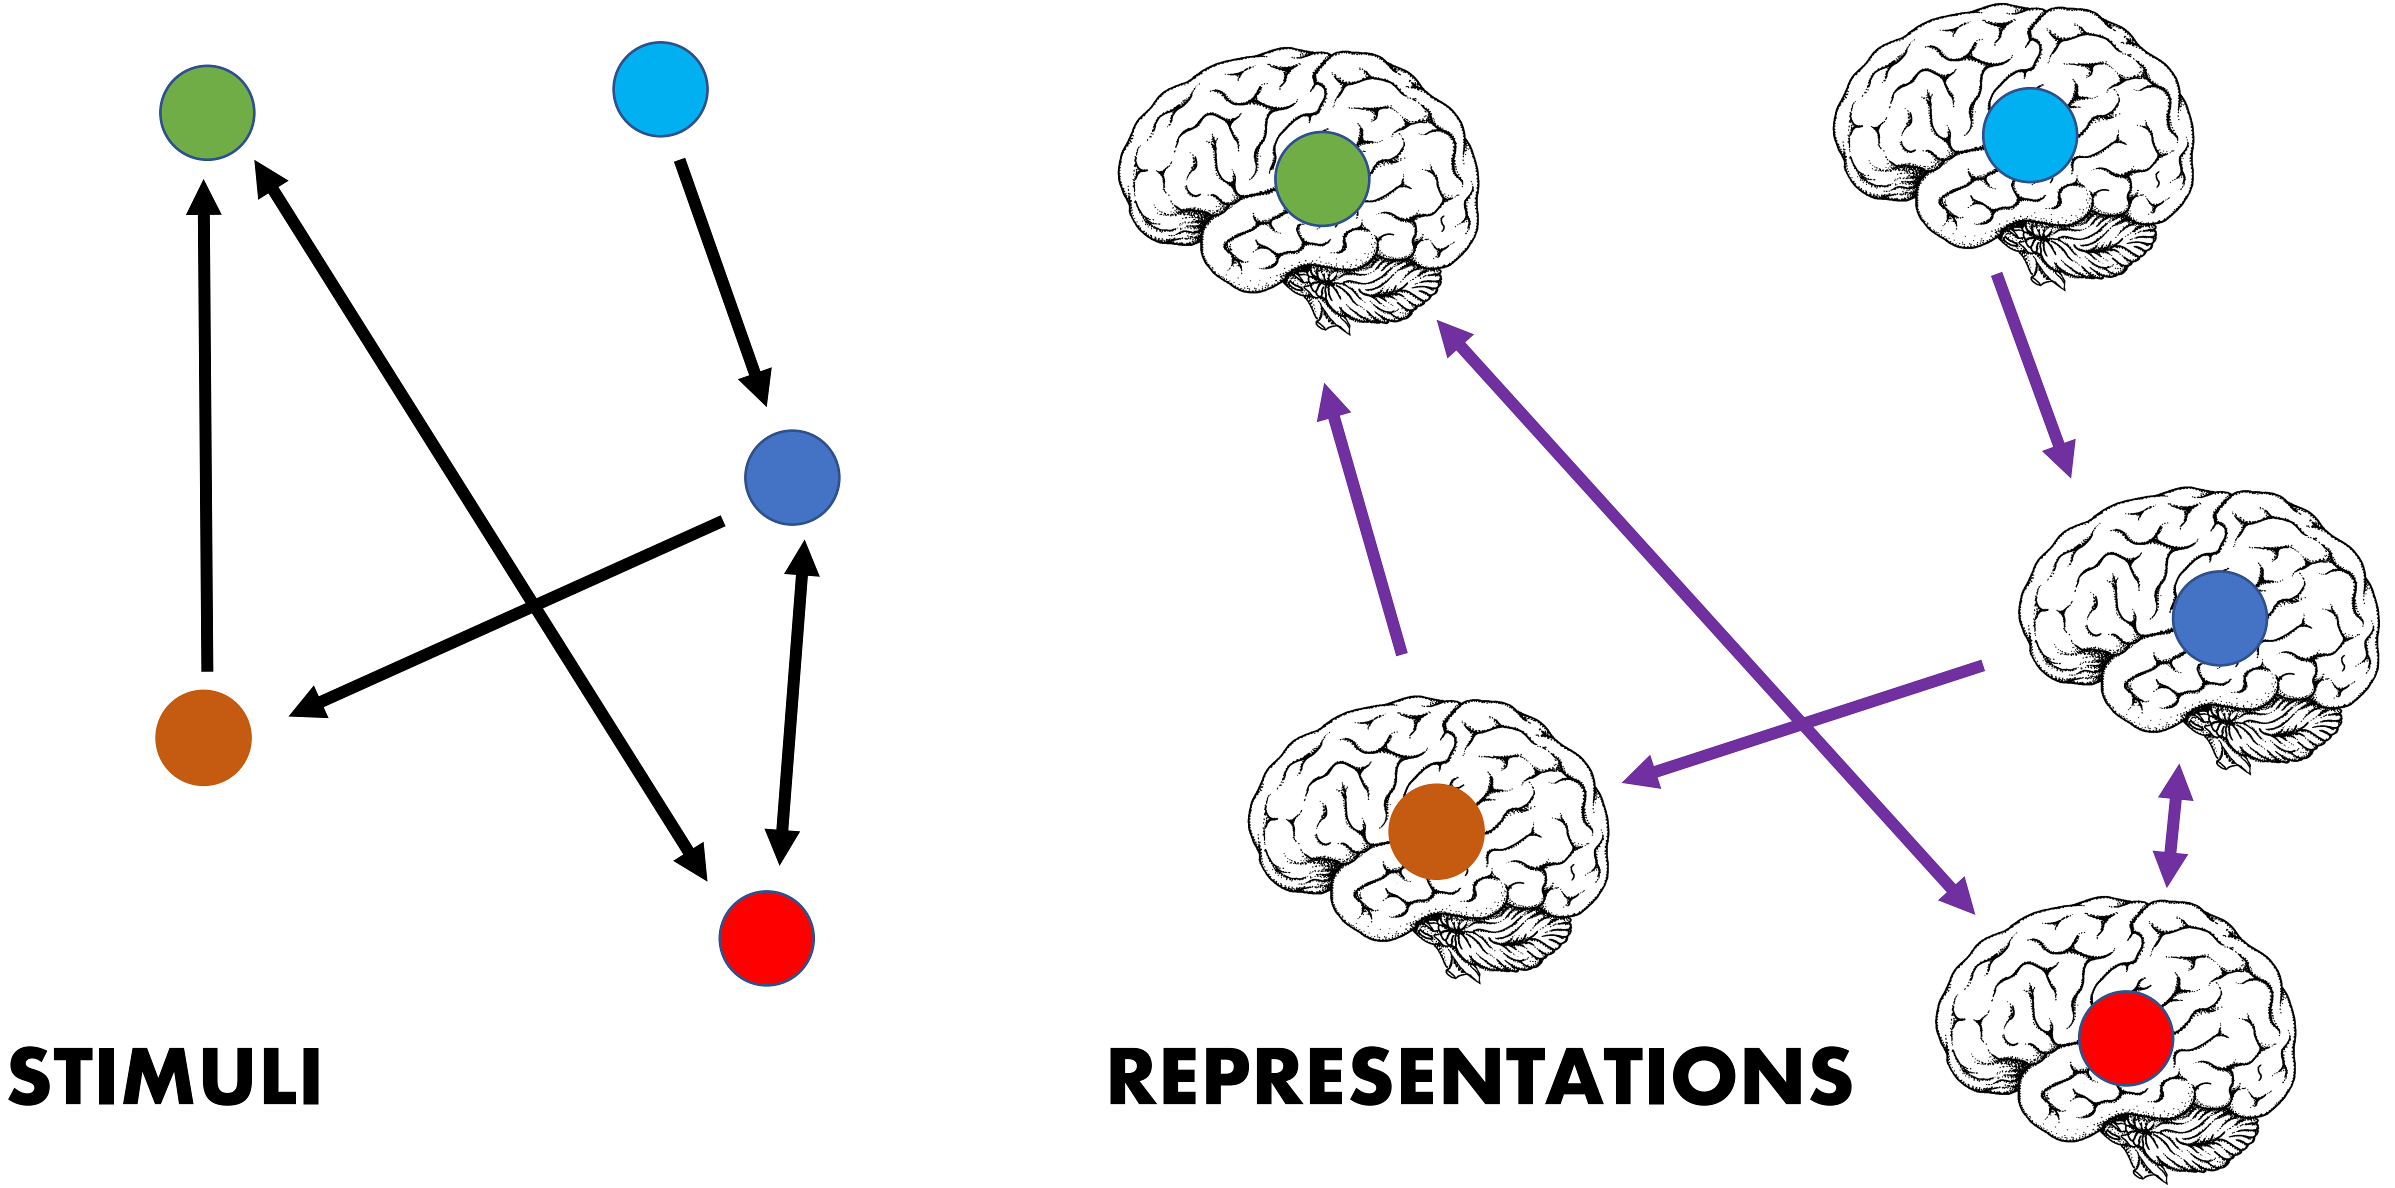

If neural activity is to represent stimuli,
it must respect the relationships between stimuli

(See Gibson, 1979)

COINCIDENCE CODES

The simplest relationship between two neurons available to the brain is temporal coincidence.

Some kind of coincidence code may instantiate invariant relationships.

How to quantify coincidence codes?

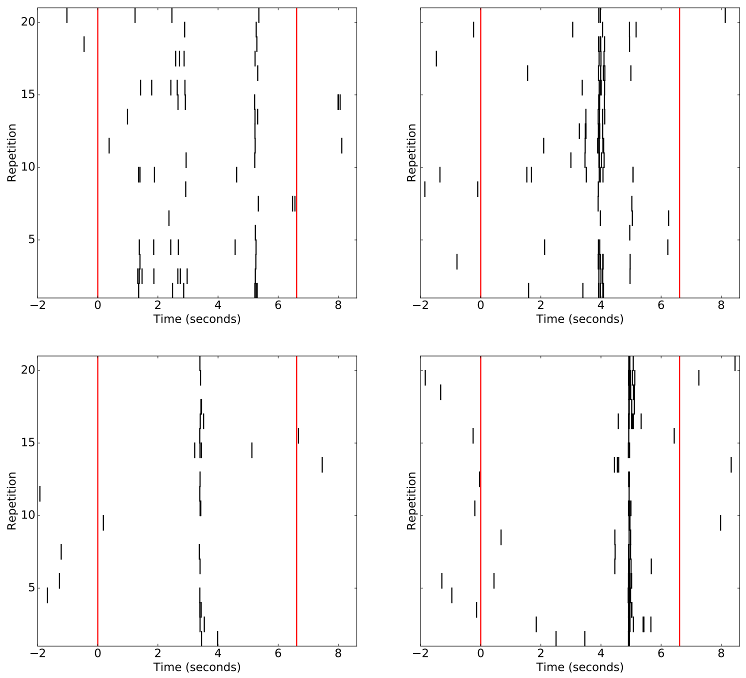

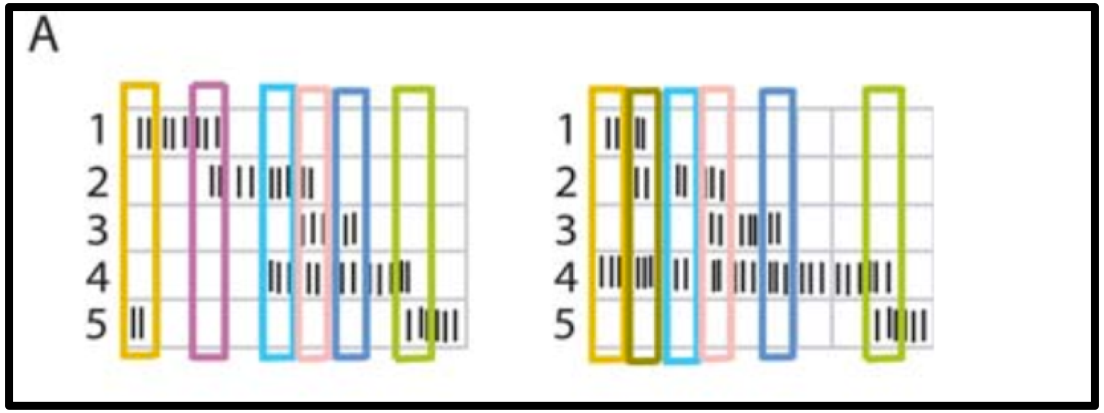

Curto + Itskov 2008

NEURAL TOPOLOGY

The Curto-Itskov construction

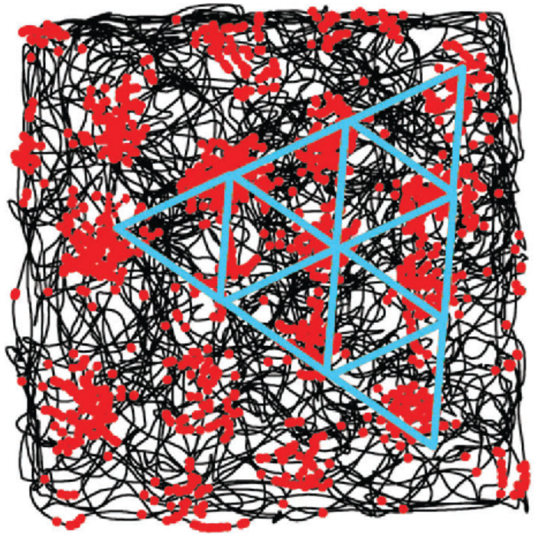

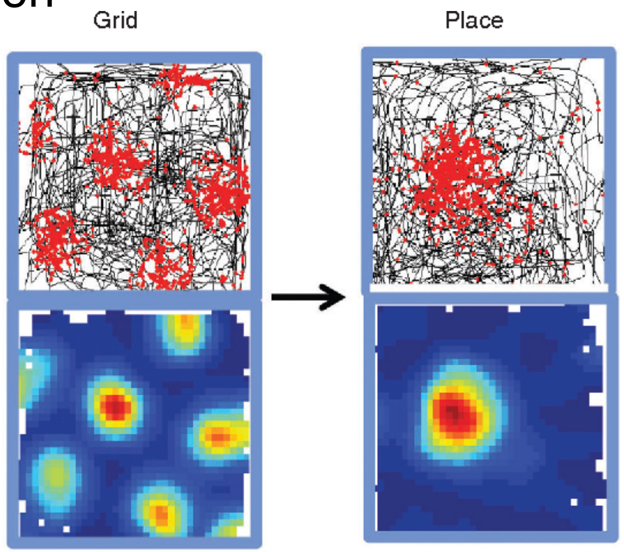

Moser, Rowland, Moser, 2015

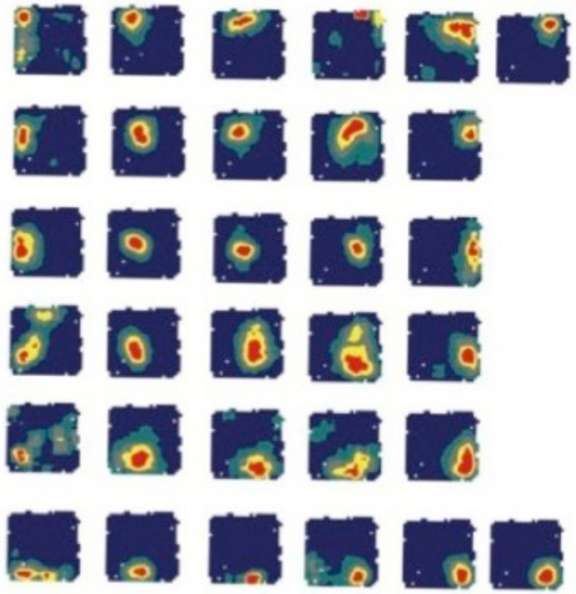

O'Keefe, Nobel Lecture, 2014

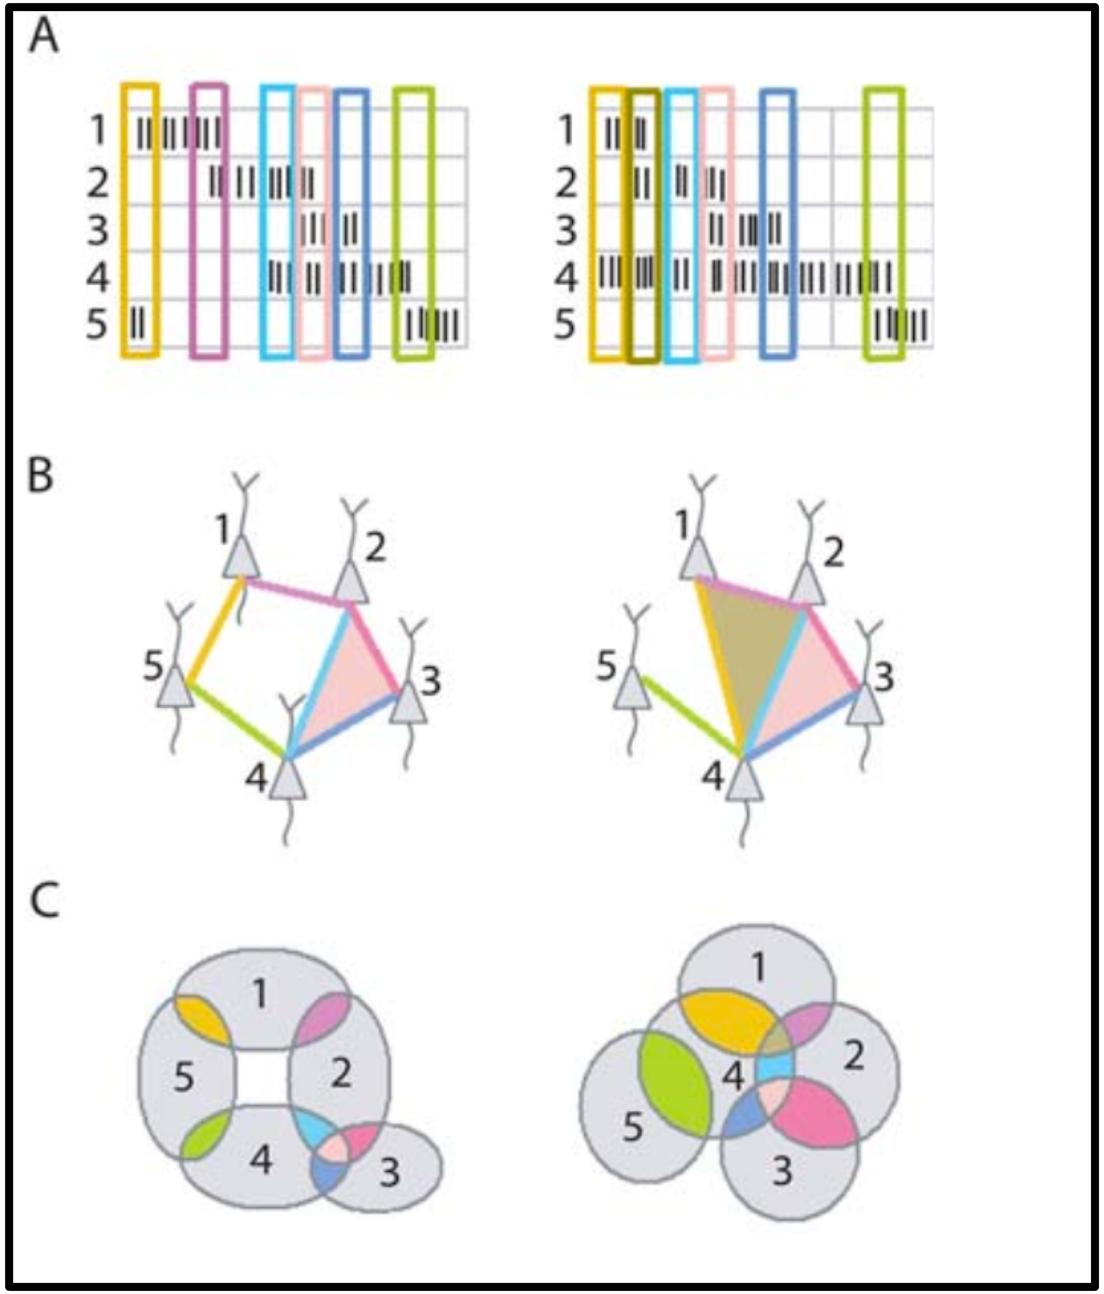

Curto + Itskov 2008

RECONSTRUCTION OF PHYSICAL SPACE FROM NEURAL TOPOLOGY

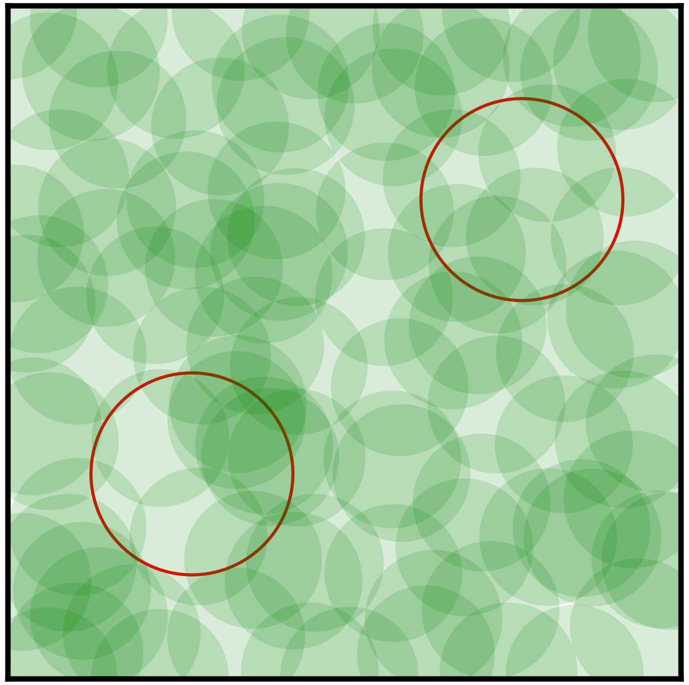

Original Environment

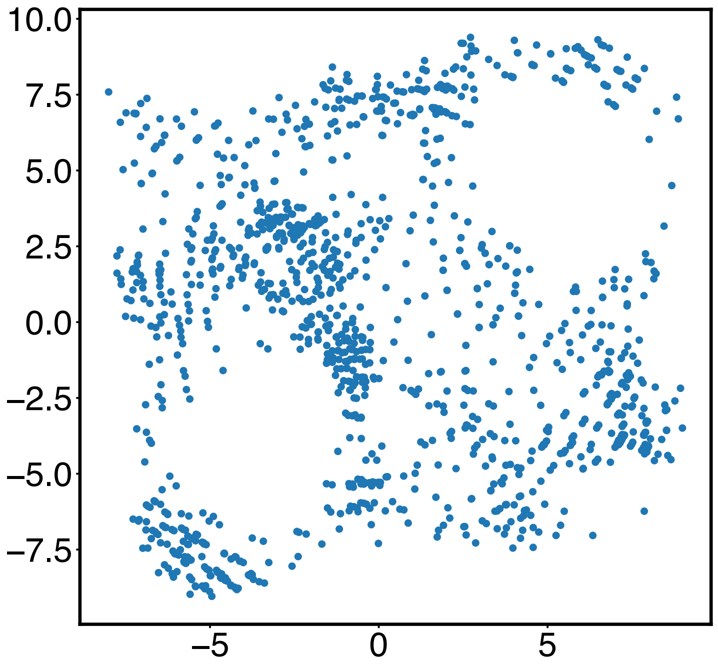

Reconstruction

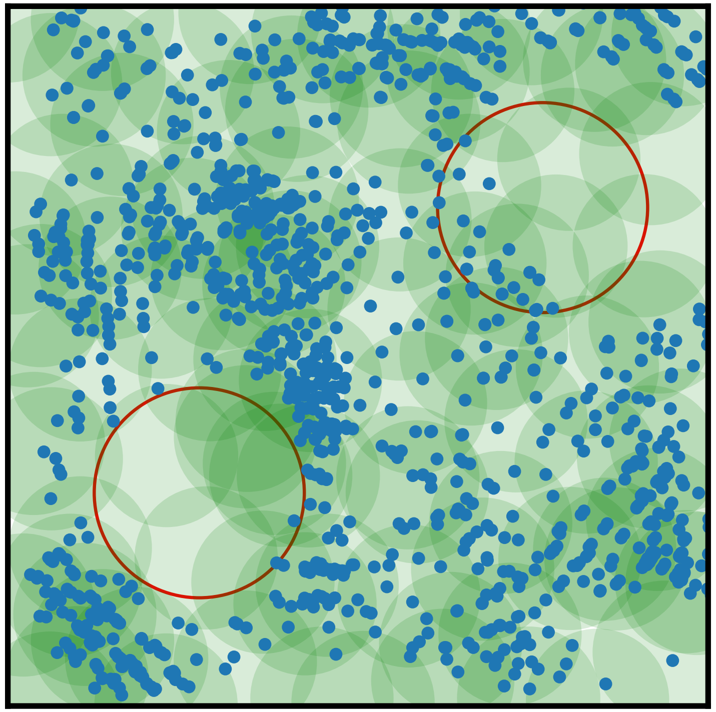

Affine Transformation

THE TOPOLOGICAL STRUCTURE OF PLACE CELL ACTIVITY

↓

INVARIANT RELATIONSHIPS BETWEEN POINTS IN SPACE

THE TOPOLOGICAL STRUCTURE OF AUDITORY NEURON ACTIVITY

↓

INVARIANT RELATIONSHIPS BETWEEN SOUNDS

TWO ALTERNATIVE CHOICE (WITH KRISTA PERKS)

Behavior:

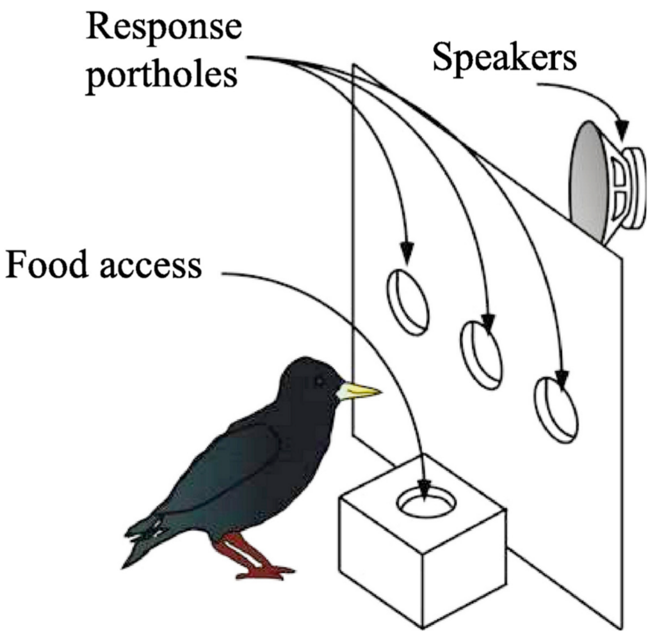

2-alternative choice task

Stimuli A , B: Left
Stimuli C, D: Right
Stimuli E - P: Unfamiliar

Stimuli:

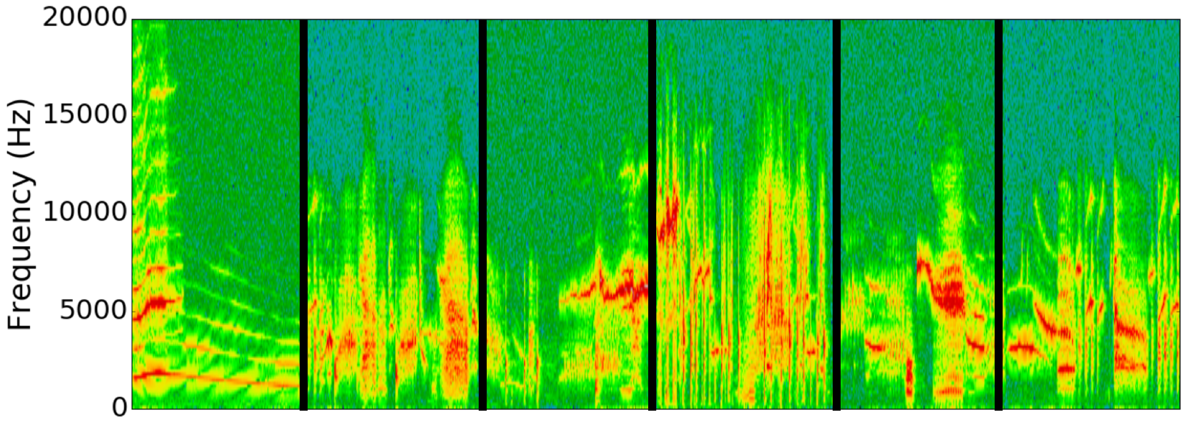

- 6 second long "pseudosongs"
- **common** 1-second introductory motif
- 5 **unique** 1-second motifs

Neural Data:

- Anesthetized, acute recordings
- 16 or 32 channel silicon microelectrode
- Region NCM

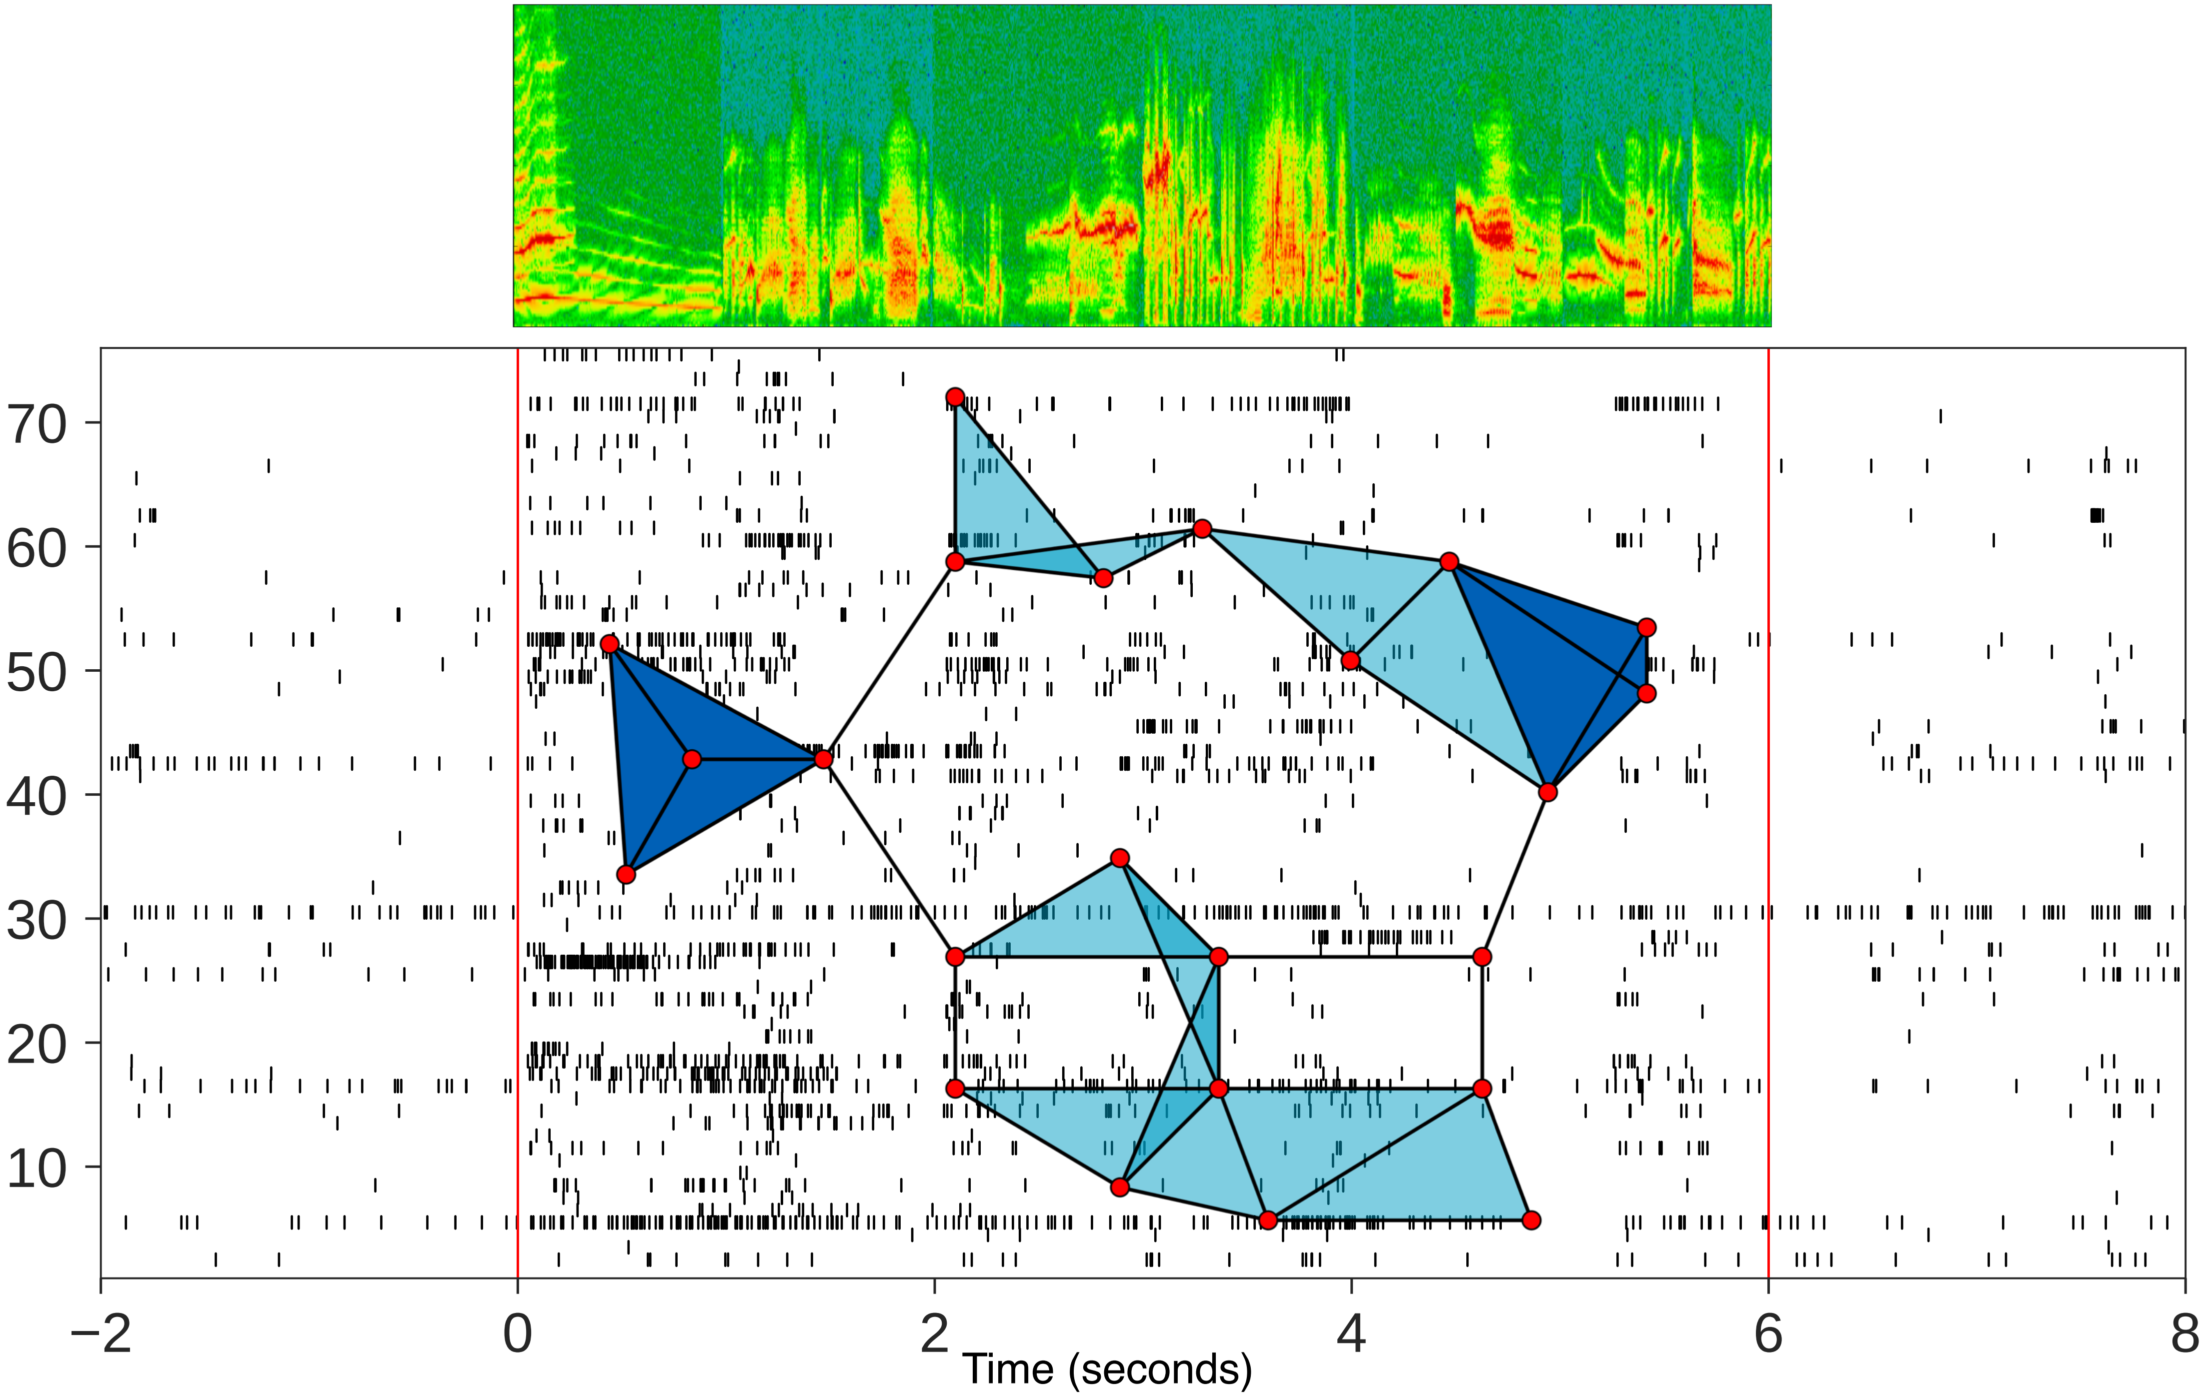

70
60
50
40
30
20
10
−2
0
2
4
6
8
Time (seconds)

SIMPLICIAL COMPLEXES THROUGH TIME

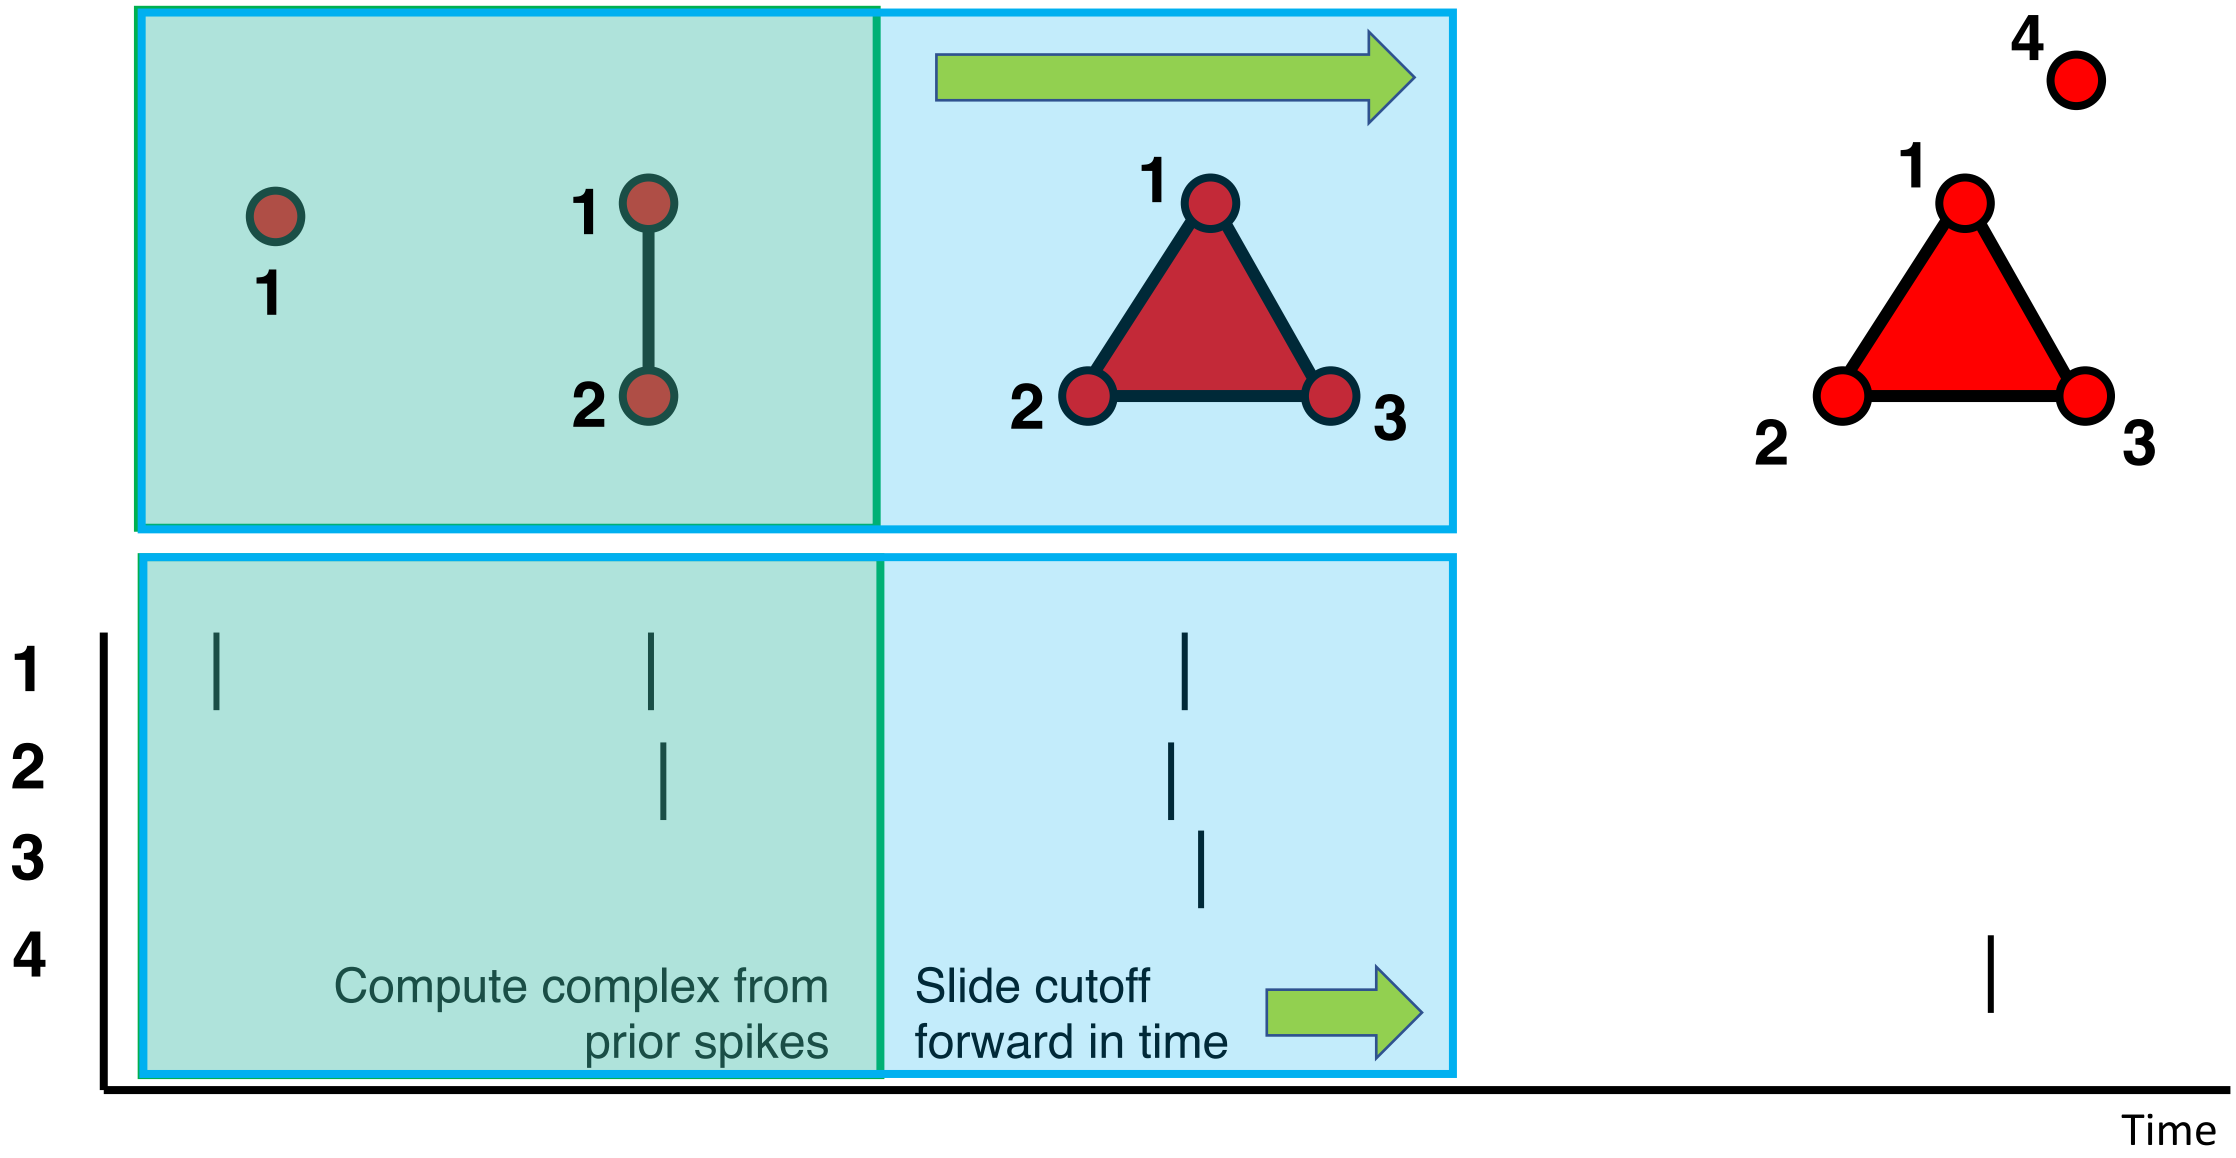

NCM NEURAL TOPOLOGY IS NON-TRIVIAL

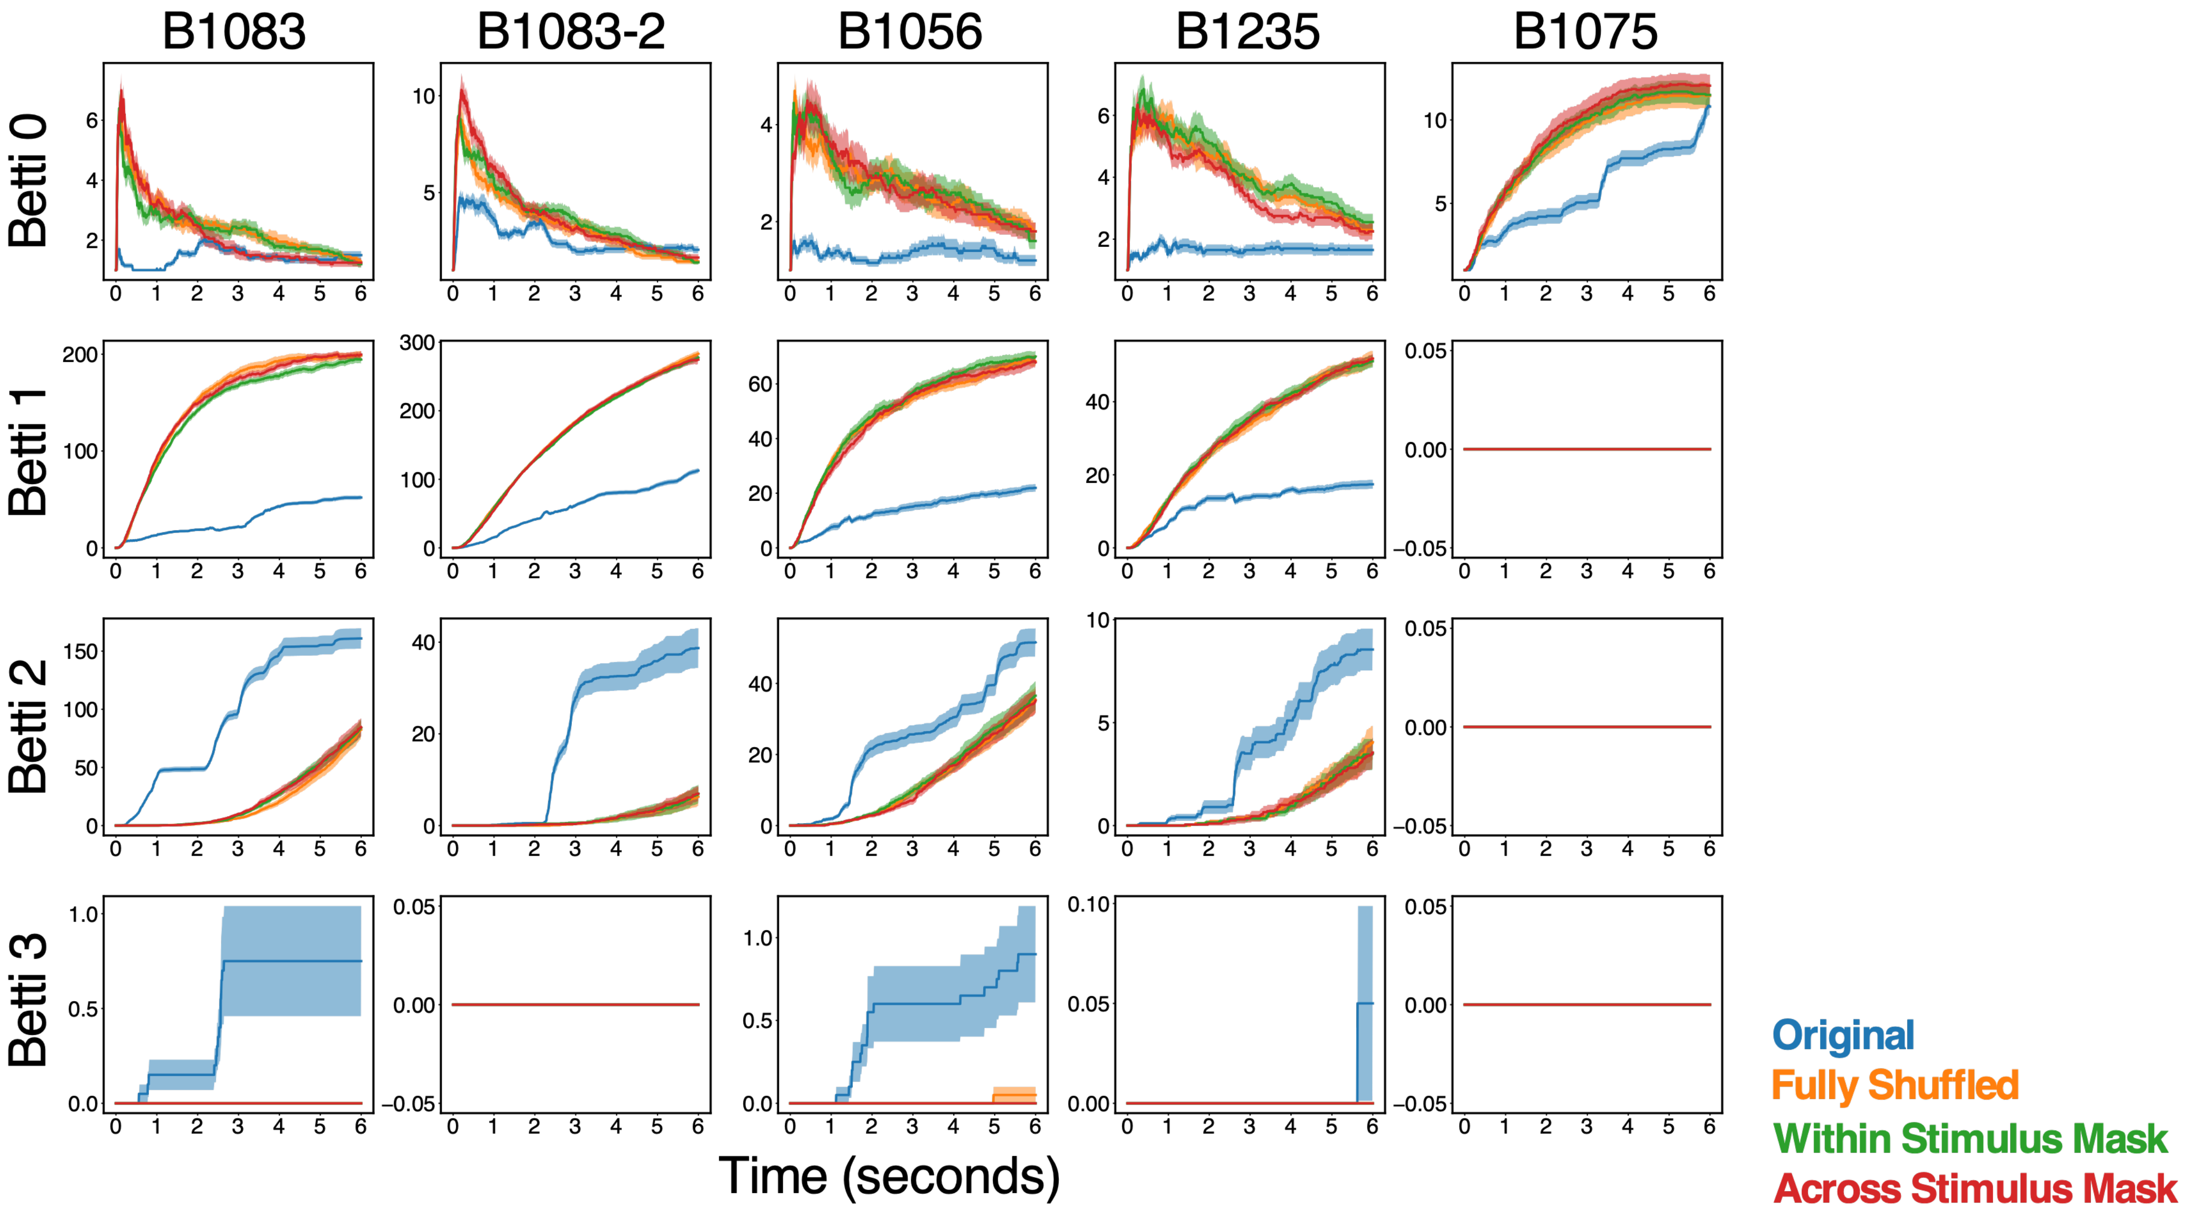

DYNAMICS OF NEURAL TOPOLOGY

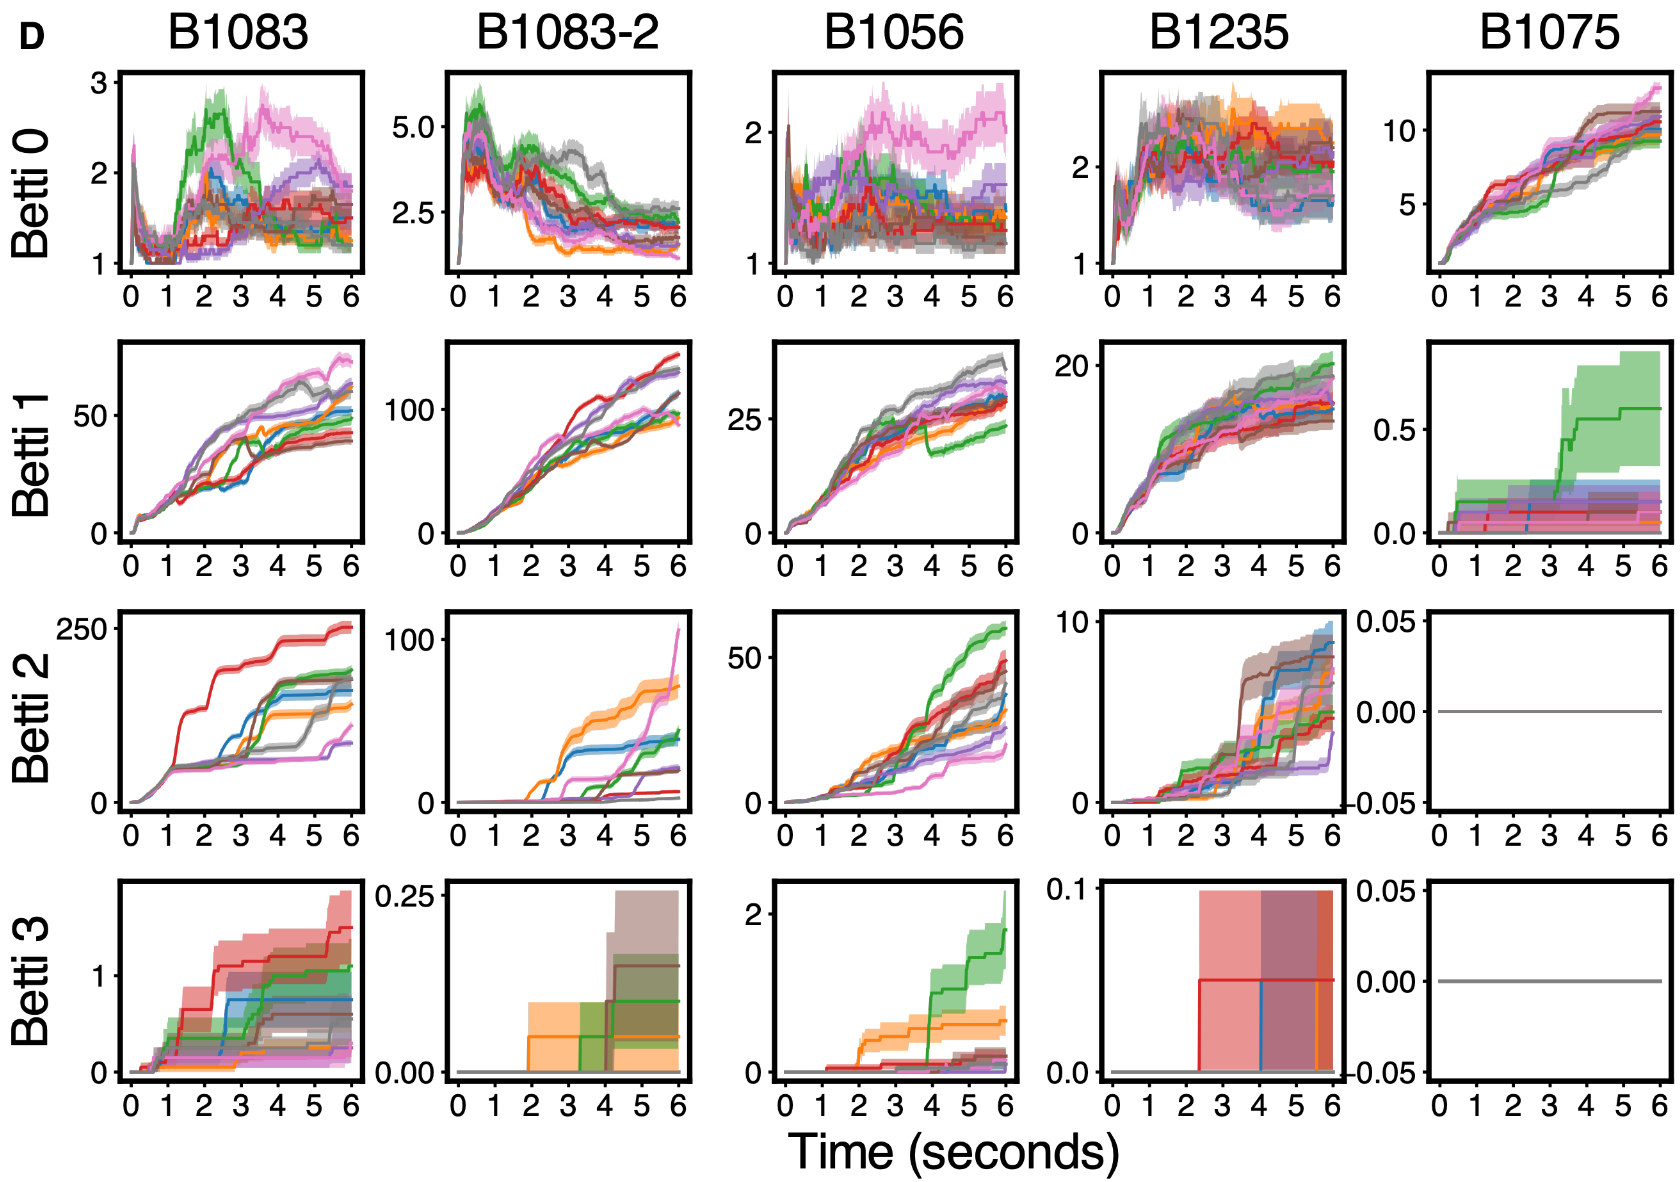

QUANTIFYING "TOPOLOGICAL DISTANCE" BETWEEN POPULATION SPIKE TRAINS

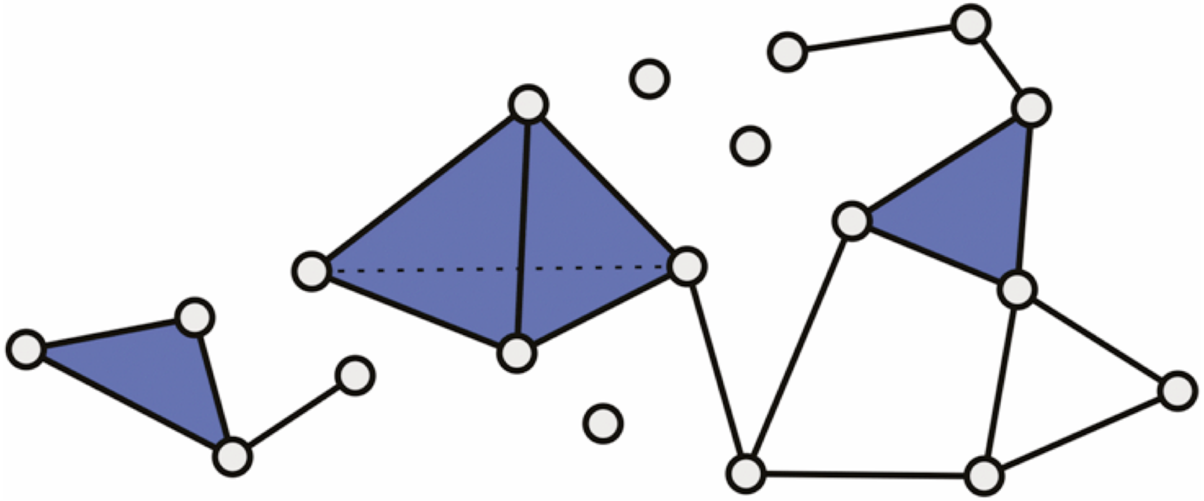

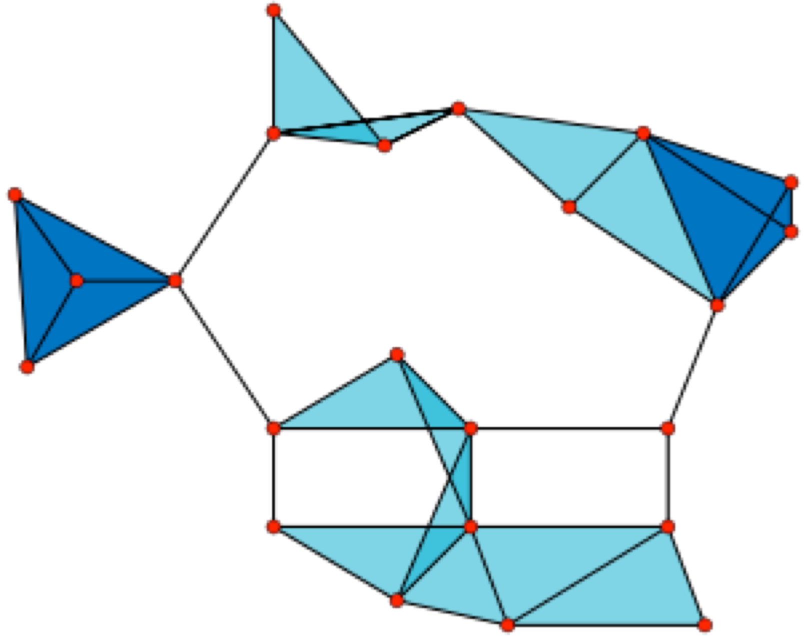

HOW FAR APART ARE THESE SIMPLICIAL COMPLEXES?

"SPECTRAL ENTROPIES FOR COMPLEX NETWORK COMPARISON"

DeDomenico and Biamonte, 2016

Idea: build density matrix from graph Laplacian

$$\rho = \frac{e^{-\beta L}}{\mathrm{Tr}\, e^{-\beta L}}$$

Discrete versions of information-theoretic quantities yield measures of similarity for graphs

$$S = -\mathrm{Tr}\left(\rho \log_2 \rho\right)$$

ENTROPY

RELATIVE ENTROPY

$$D_{KL}(\rho||\sigma) = -\mathrm{Tr}\left(\rho(\log_2 \rho - \log_2 \sigma)\right)$$

$$S = -\int p(x) \log_2 p(x)\, dx$$

ENTROPY

RELATIVE ENTROPY

$$D_{KL}(p||q) = -\int p(x) \log_2 \frac{p(x)}{q(x)}\, dx$$

SIMPLICIAL LAPLACIAN

There exist "simplicial Laplacians"

$$\boxed{L_i = \partial_i^* \partial_i + \partial_{i+1} \partial_{i+1}^*}$$

There is one Laplacian for every dimension in the complex.

Equivalent to the graph Laplacian if the complex is purely one-dimensional

SIMPLICIAL LAPLACIAN SPECTRAL ENTROPY (SLSE) ALGORITHM

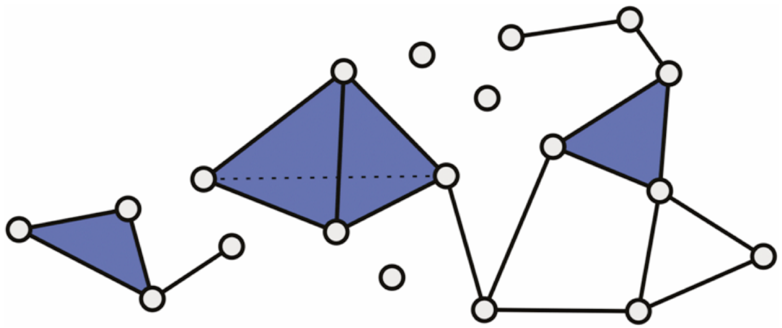

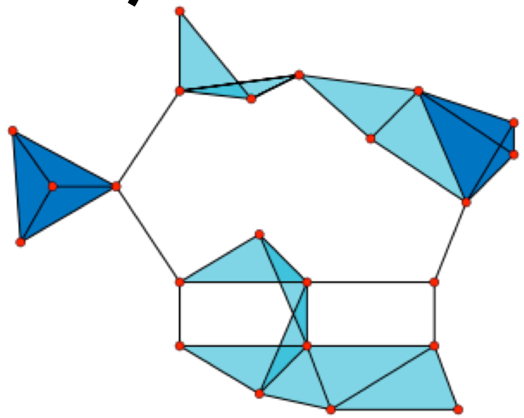

$$\boxed{L_i = \partial_i^* \partial_i + \partial_{i+1} \partial_{i+1}^*}$$

$$\boxed{L_i = \partial_i^* \partial_i + \partial_{i+1} \partial_{i+1}^*}$$

$$\rho = \frac{e^{-\beta L}}{\operatorname{Tr} e^{-\beta L}}$$

$$\sigma = \frac{e^{-\beta L}}{\operatorname{Tr} e^{-\beta L}}$$

$$D_{KL}(\rho || \sigma) = -\operatorname{Tr}\left(\rho(\log_2 \rho - \log_2 \sigma)\right)$$

SLSE RECOVERS RELATIONSHIPS BETWEEN STIMULI

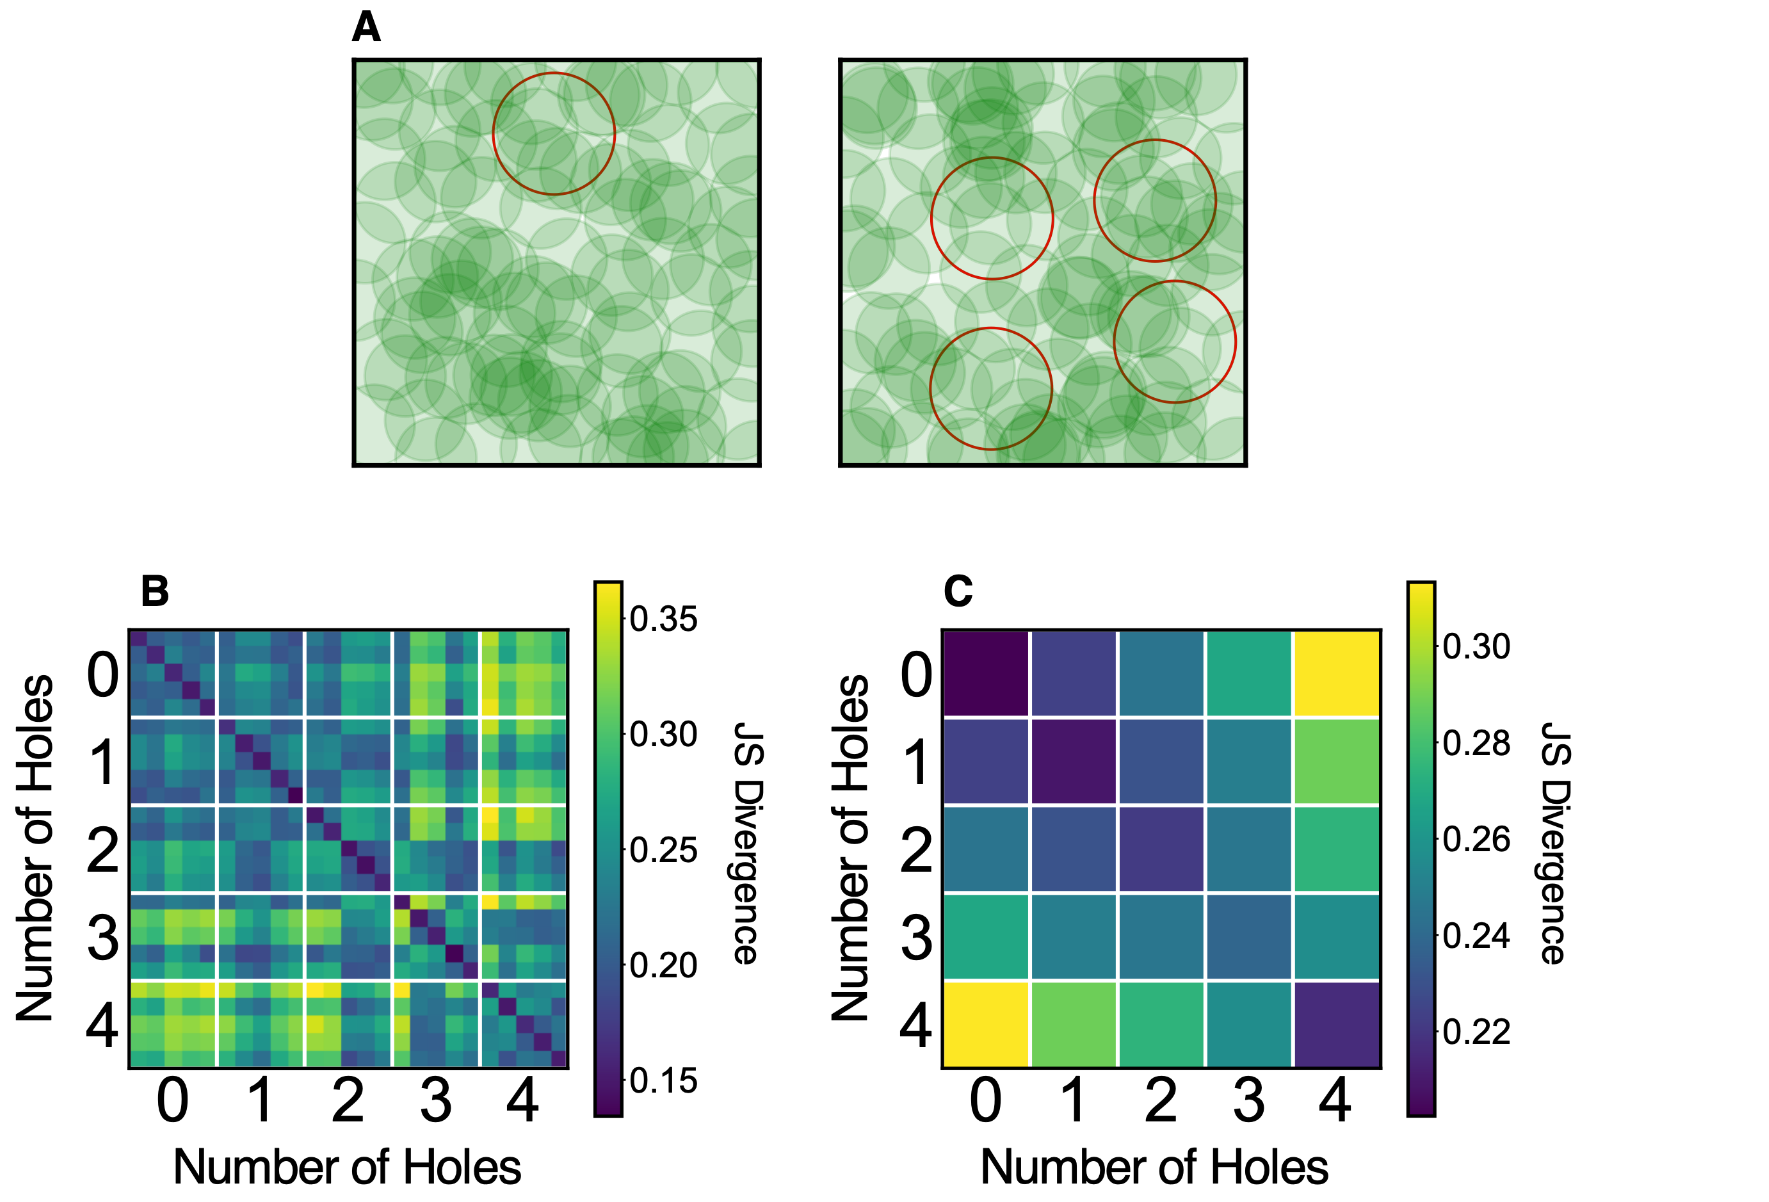

SLSE RECOVERS LEARNED BEHAVIORAL RELATIONSHIPS

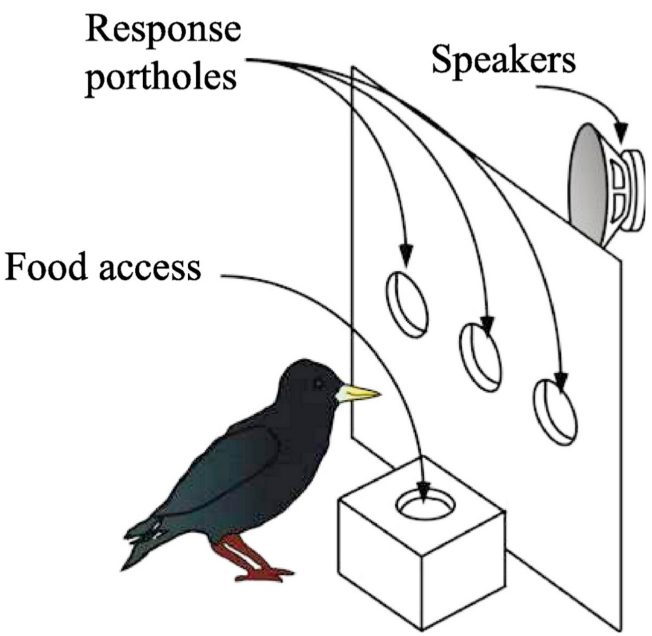

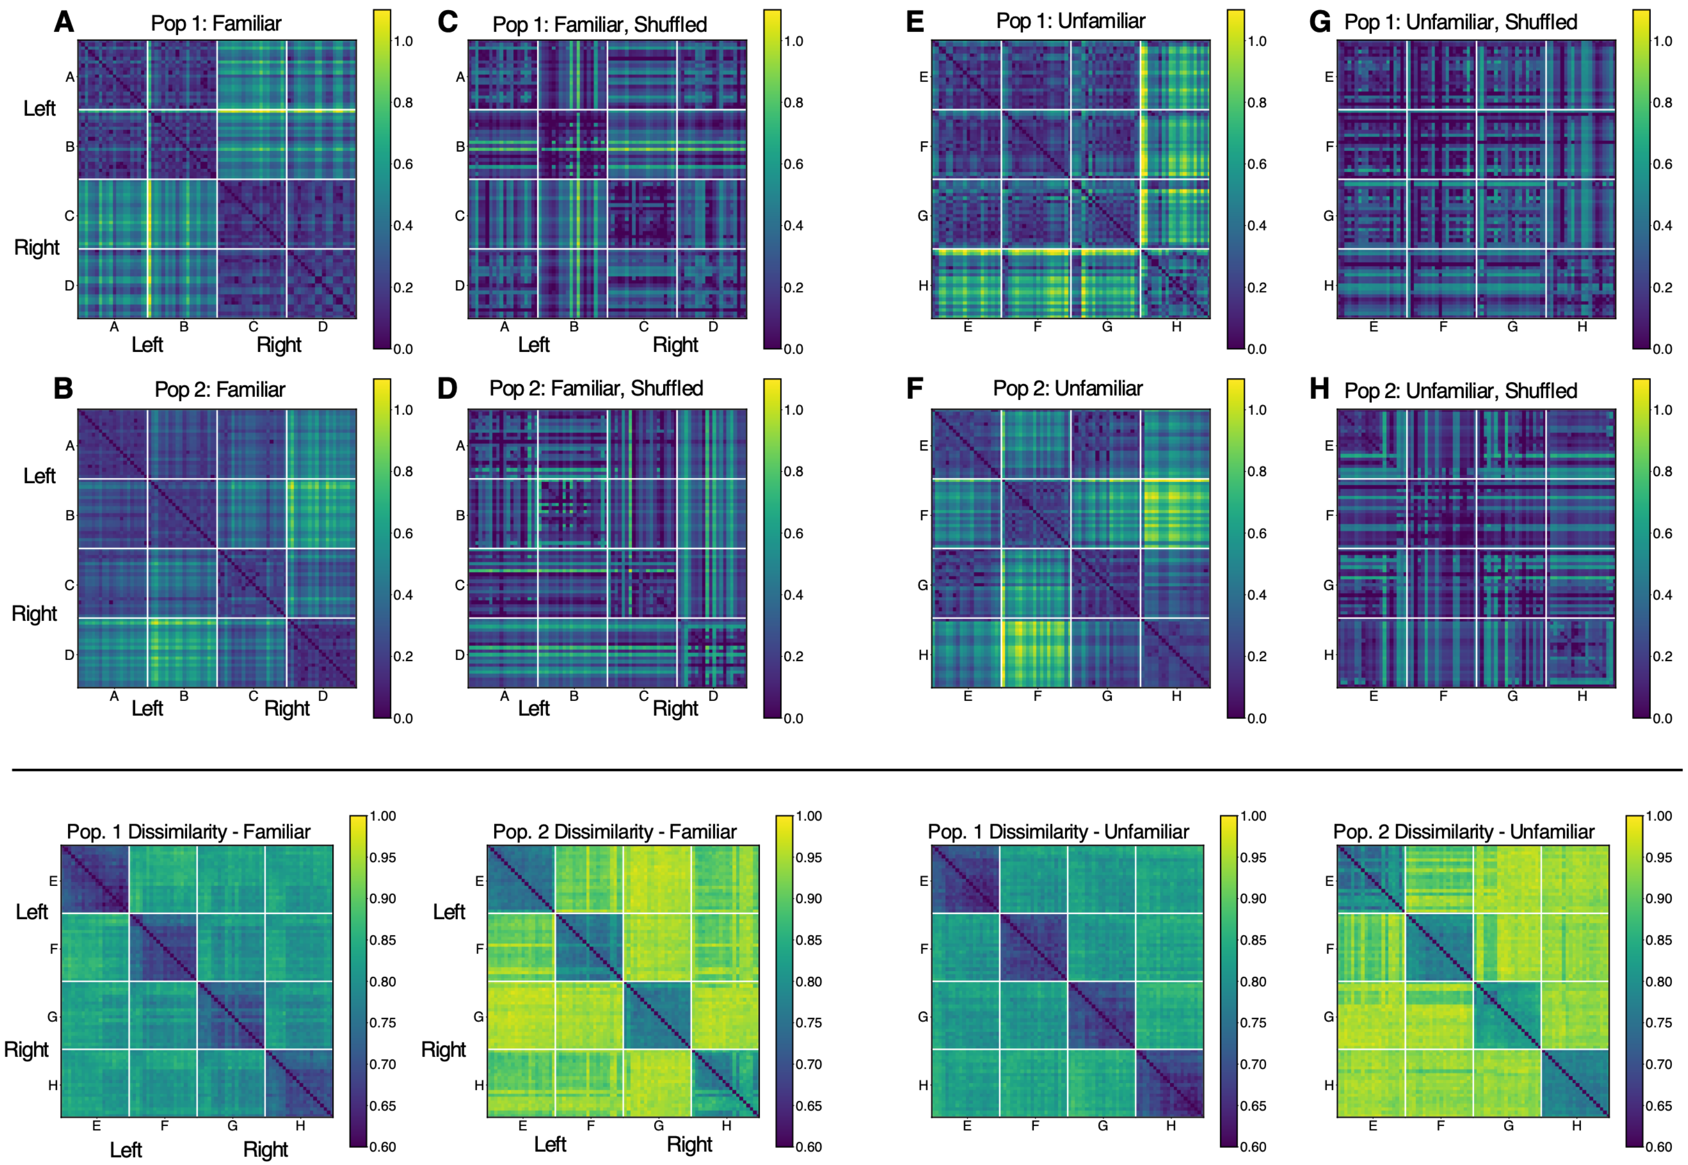

2-alternative choice task

Stimuli A, B: Left
Stimuli C, D: Right
Stimuli E - H: Unfamiliar

RECONSTRUCTIONG "CONCEPTUAL SPACES"

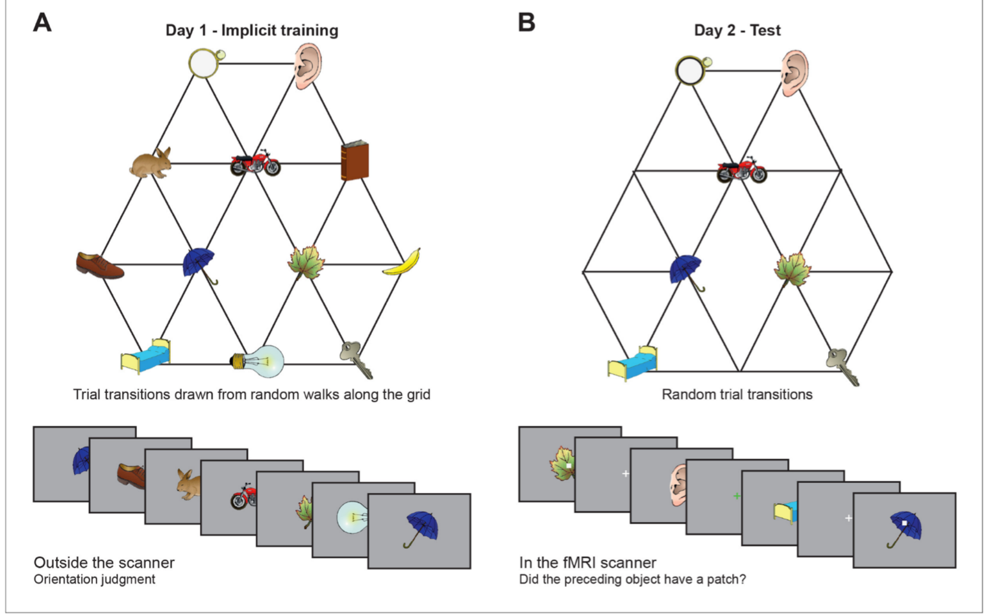

There is increasing evidence that the brain organizes all information spatially - "Cognitive Maps"

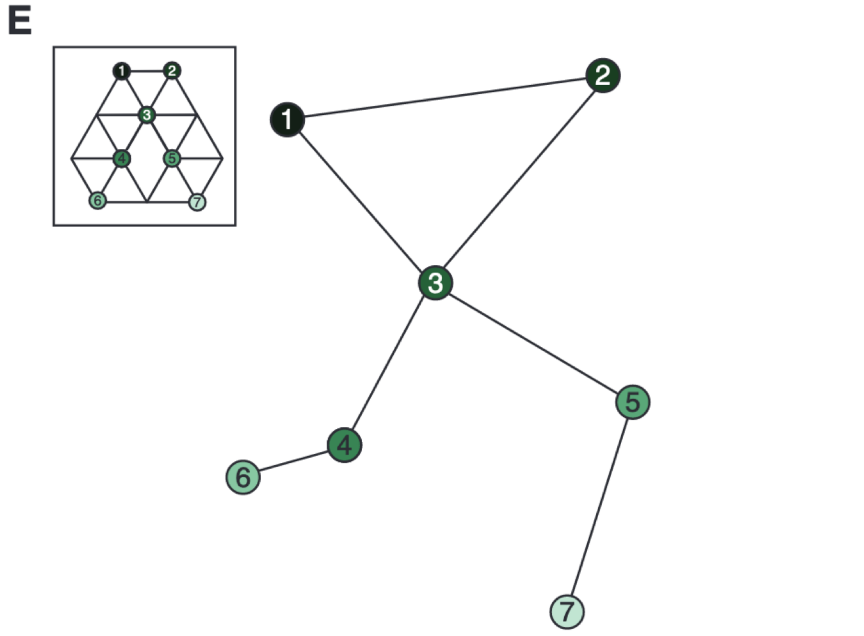

We can reconstruct the Cech nerve of physical space from neural activity.

Can we reconstruct the Cech nerve of more abstract categories?

How can network dynamics support this reconstruction?

A map of abstract relational knowledge in the human hippocampal–entorhinal cortex

Mona M Garvert[1,2*], Raymond J Dolan[1,3], Timothy EJ Behrens[1,2]

COMPETING THEORIES OF VISUAL PERCEPTION

David Marr

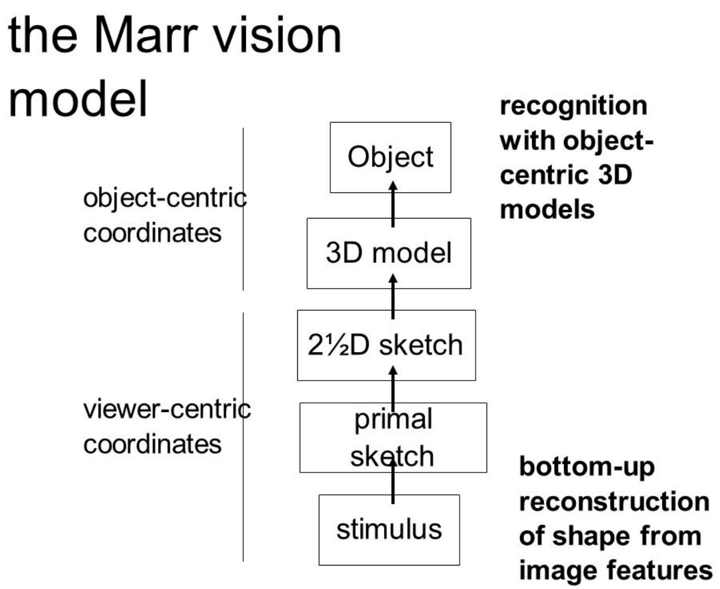

"Vision is a process that produces from images of the external world a description that is useful to the viewer and not cluttered with irrelevant information"

James Gibson

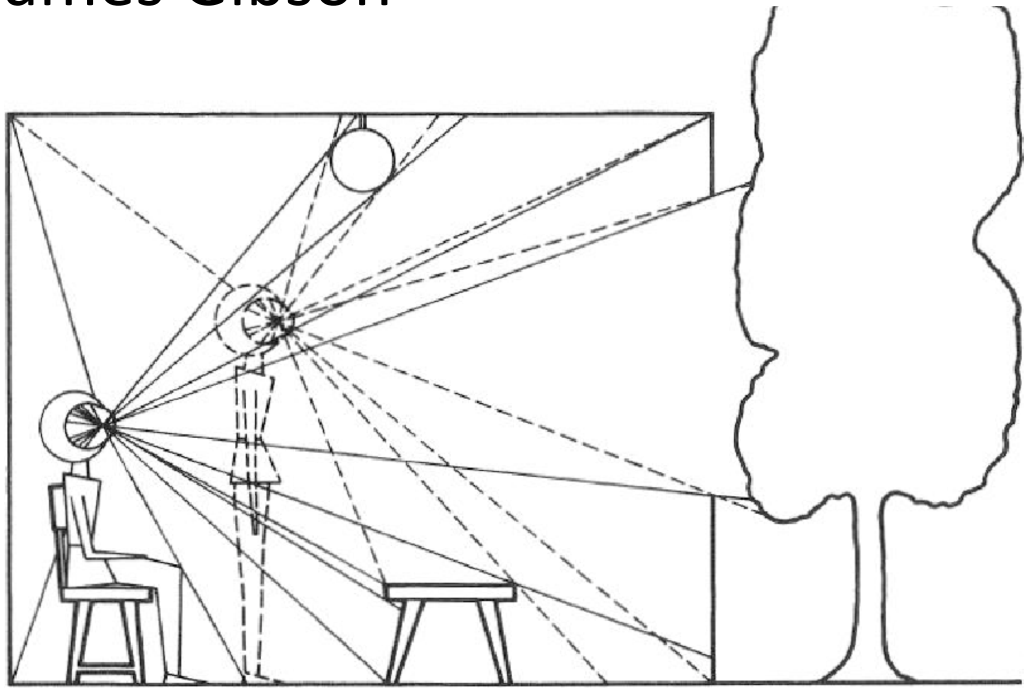

"A great many properties of the array are lawfully or regularly variant with change of observation point, and this means that in each case a property defined by the law is invariant."

SENSORY LAWS

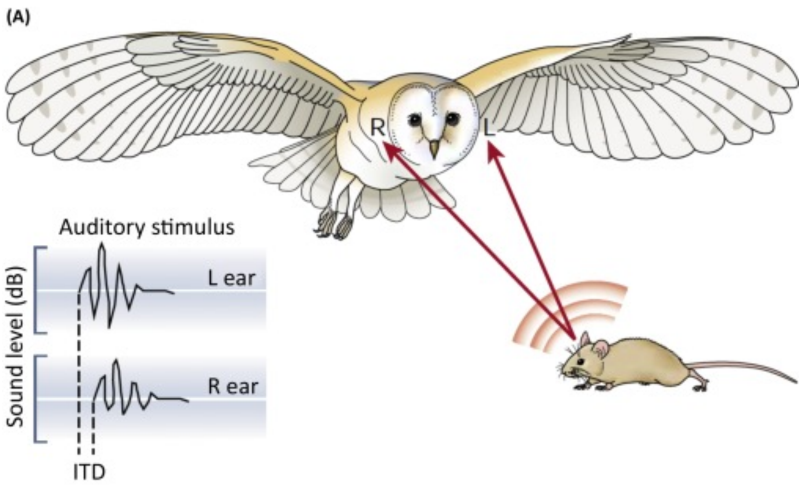

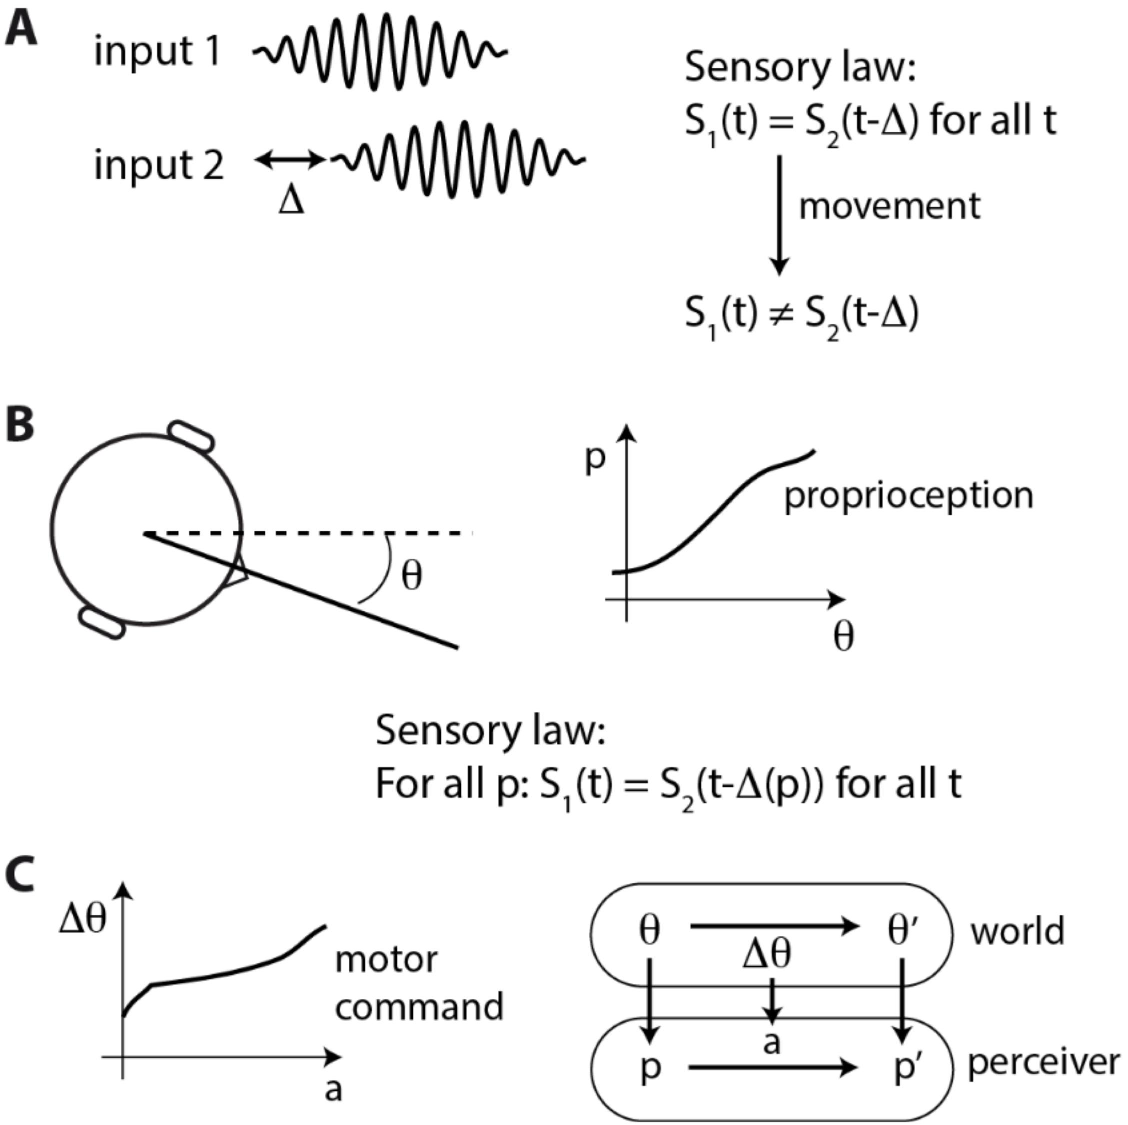

Romain Brette, "Subjective Physics", 2013. Arxiv:1311.3129

SENSORY LAWS

"Therefore, one basic challenge for neural models of perceptual systems is to identify and respond to laws that unfold in time." – Brette, 2013

"Synchrony is a temporal invariant, and if spike trains are caused by sensory signals, then a particular pattern of synchrony in neural population reflects the occurrence of a particular sensory law." (Brette 2012; D. F. M. Goodman and Brette 2010)

Can "sensory laws" be expressed in a temporal logic built from neural activity?

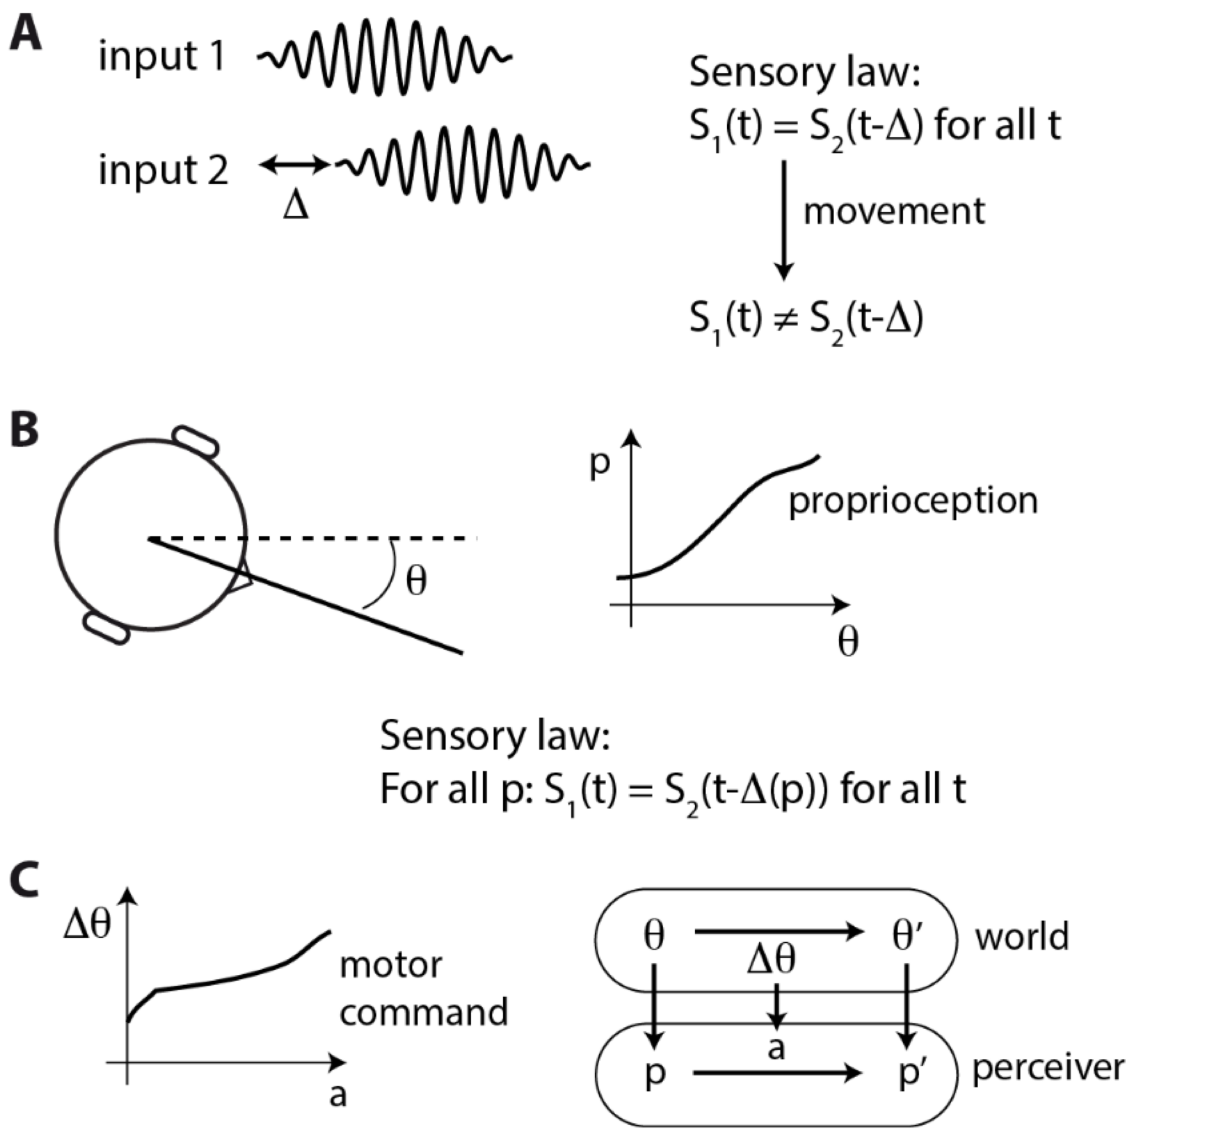

Romain Brette, "Subjective Physics", 2013. Arxiv:1311.3129

THANK YOU

Email: btheilma@ucsd.edu
Website: bradtheilman.com
Twitter: @bradtheilman