Neural computation as a transformation of similarity

Dmitri "Mitya" Chklovskii

FLATIRON INSTITUTE
Division of Simons Foundation

NYU NEUROSCIENCE

Imaging neural activity *in vivo*

50μm

Raw data: Yuste lab

CaImAn: Giovannucci, Friedrich et al, Chklovskii, Pnevmatikakis

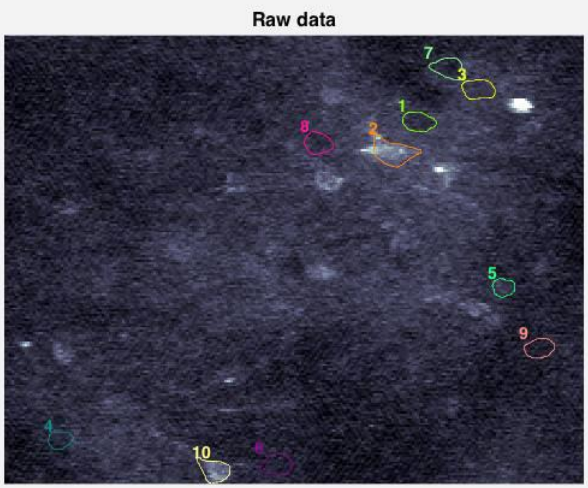

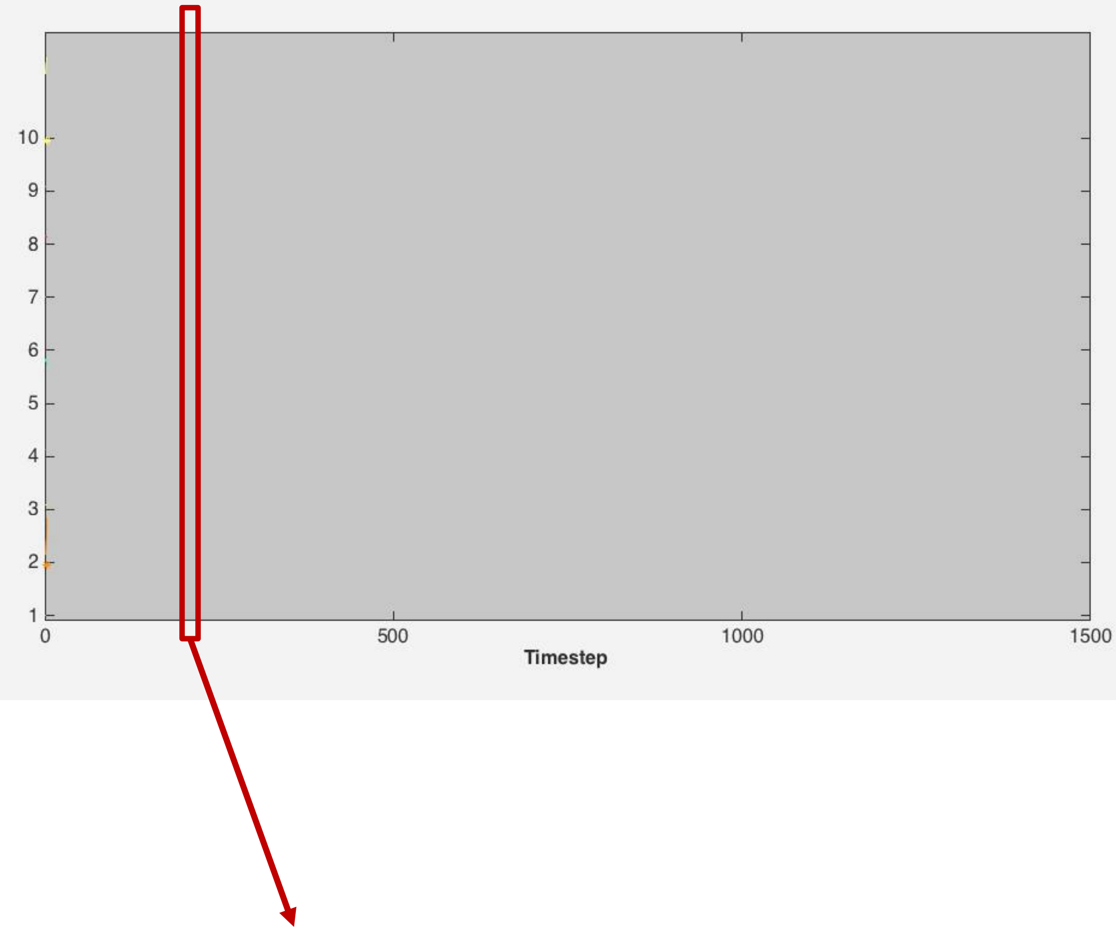

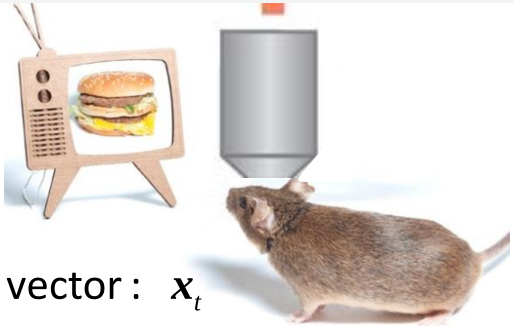

Stimulus vector : x_t

Neural activity vector : y_t

What does neural activity represent?

How is neural representation computed?

Reconstruction approach: linear decoding

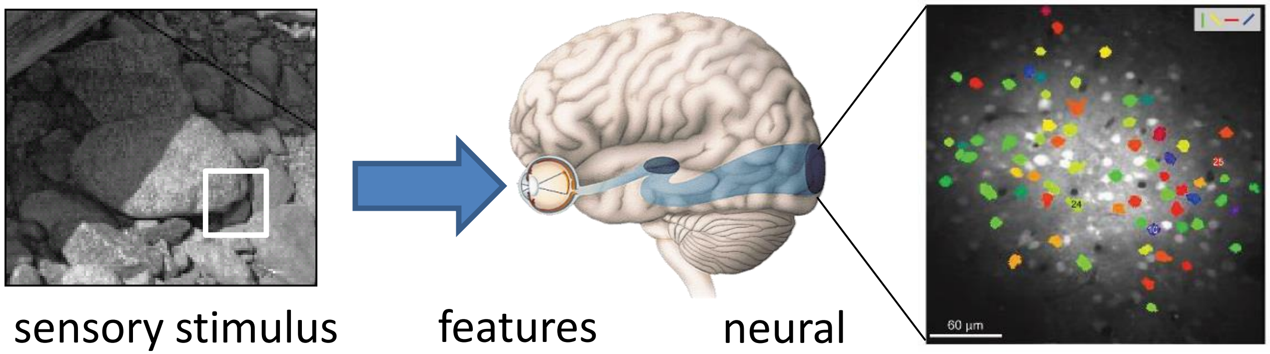

sensory stimulus — features — neural activities

$$\boldsymbol{x} \approx \sum_k \boldsymbol{w}_k y_k$$

$\boldsymbol{w}_{1:n}$:

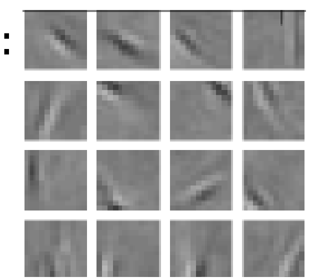

Abbott, Atick, Bialek, Chklovskii, Olshausen & Field, …

Challenges

- Activity patterns vary across individuals
- Nonlocal synaptic learning rules

Similarity of neural activity patterns in IT

Human Cortex

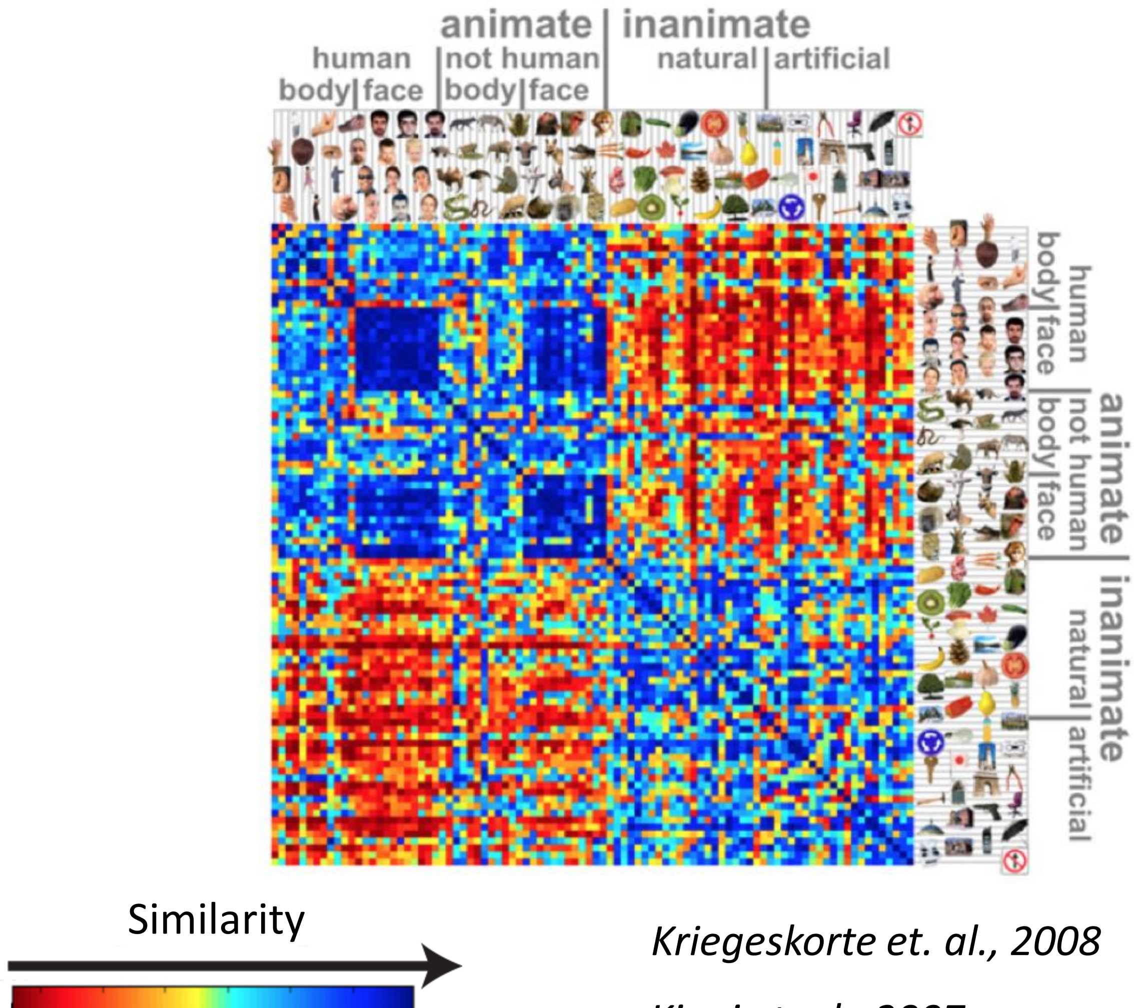

Kriegeskorte et. al., 2008

Kiani et. al., 2007

Similarity of neural activity patterns in IT

Monkey Cortex

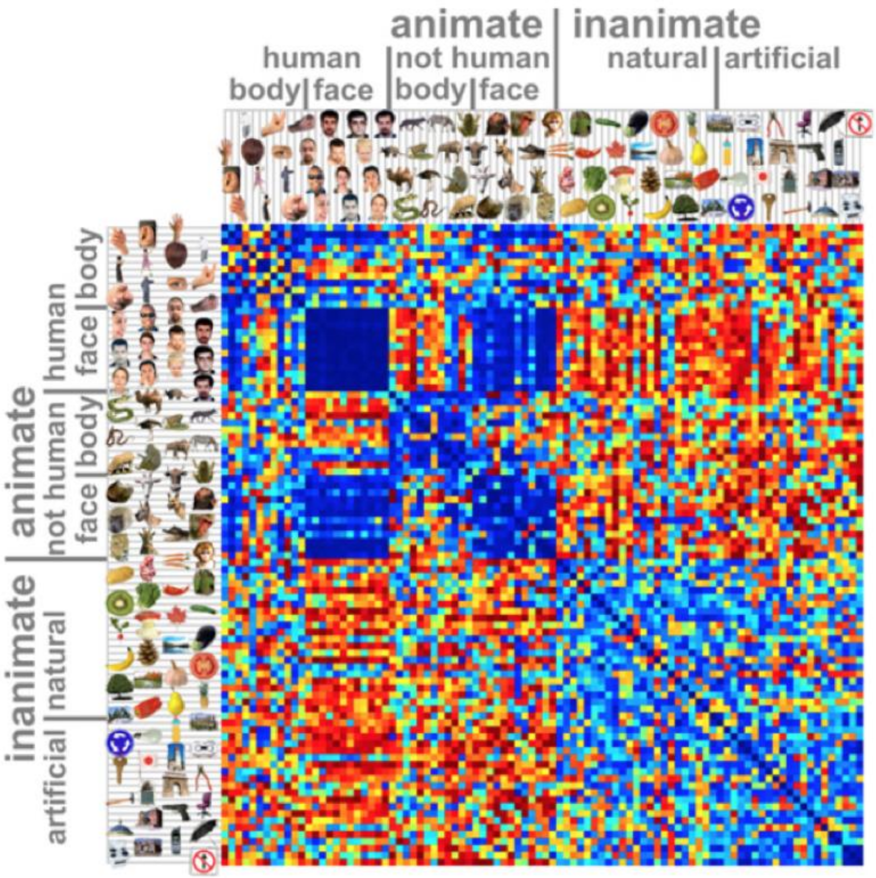

Human Cortex

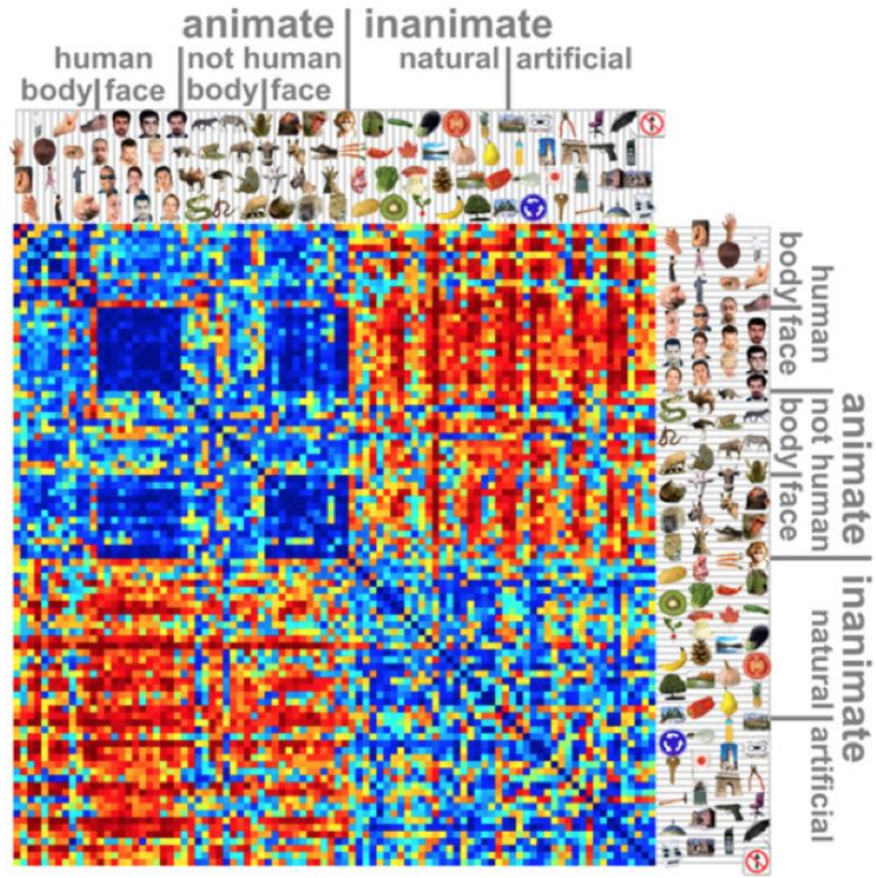

Similarity

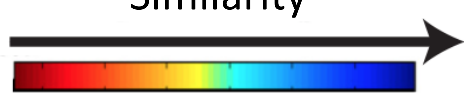

Kriegeskorte et. al., 2008

Kiani et. al., 2007

Invariance of similarity across individuals may account for the invariance of concepts

Neural representation is representation of similarities
Shimon Edelman (1998)

Similarity along the visual pathway

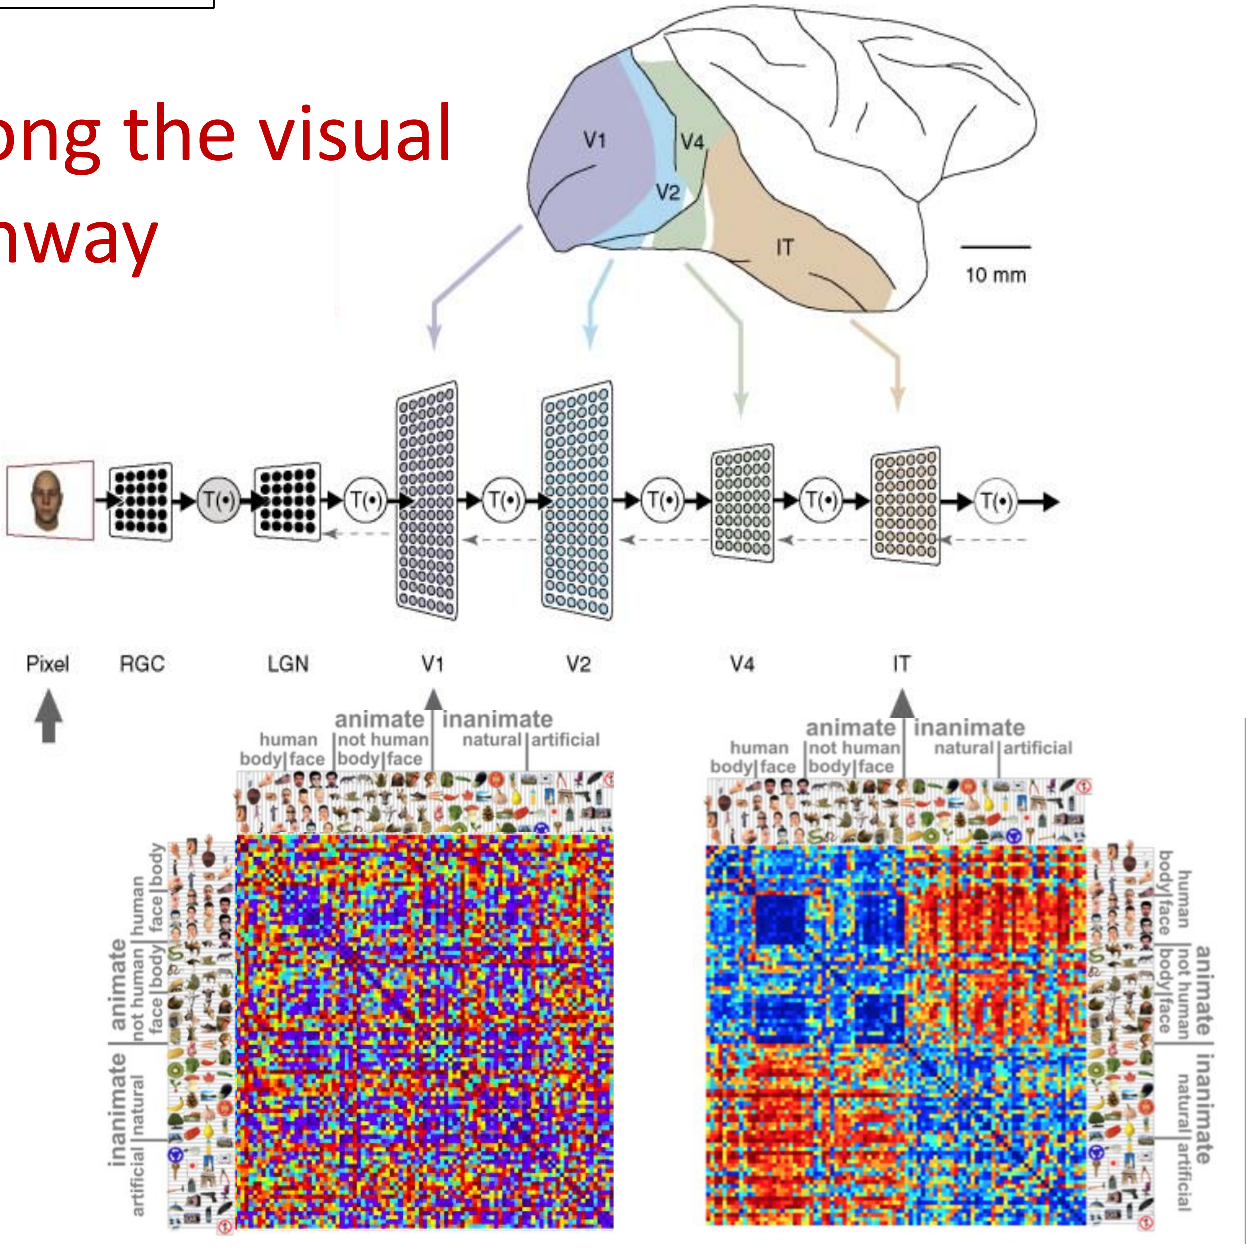

Kiani et. al., 2007
Kriegeskorte et. al., 2008

What does neural activity represent?

Similarity along the visual pathway

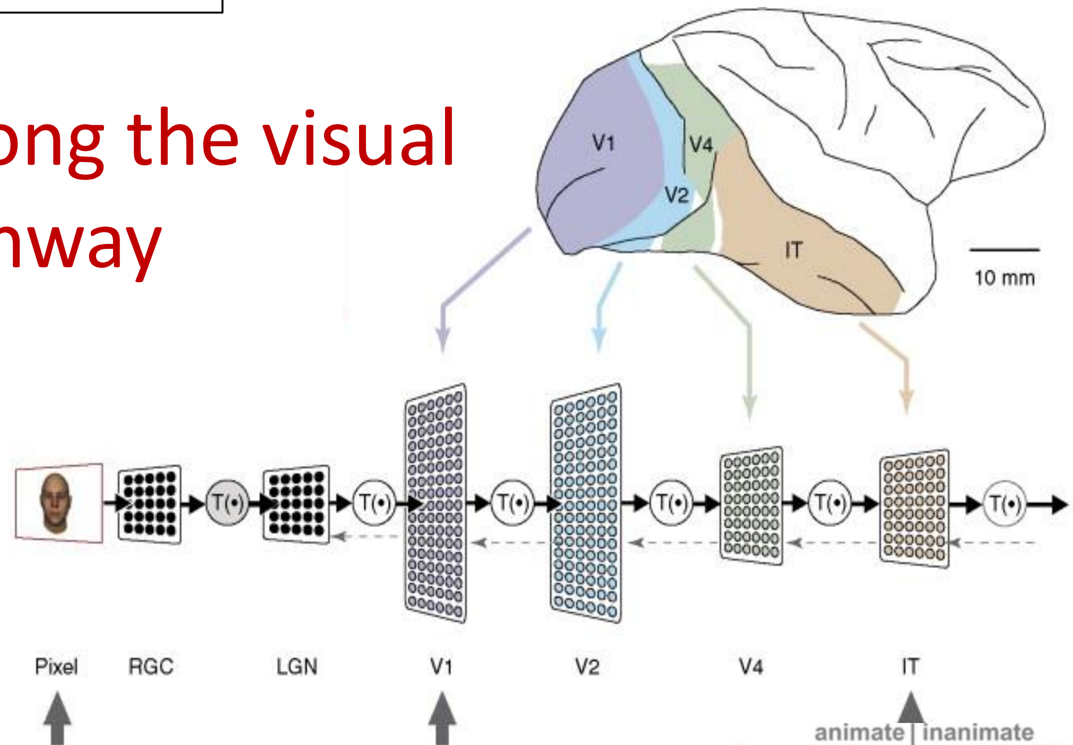

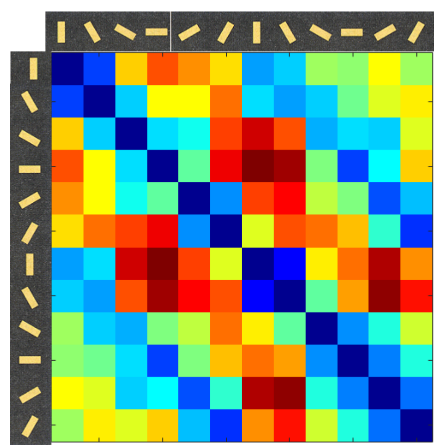

Kiani lab

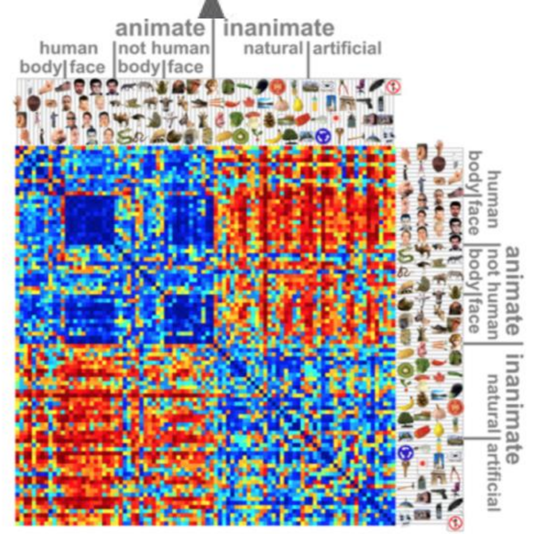

Kriegeskorte et. al., 2008

Similarity alignment: Similar input activity patterns evoke similar output activity patterns

Hu, Pehlevan, Chklovskii (2014)
Pehlevan, Chklovskii (2014)
Pehlevan, Hu, Chklovskii (2015)
Pehlevan, Chklovskii (2015)
Bahroun, Hunsicker, Soltoggio (2017)
Seung & Zung (2017)

Categorization

pixel intensity space

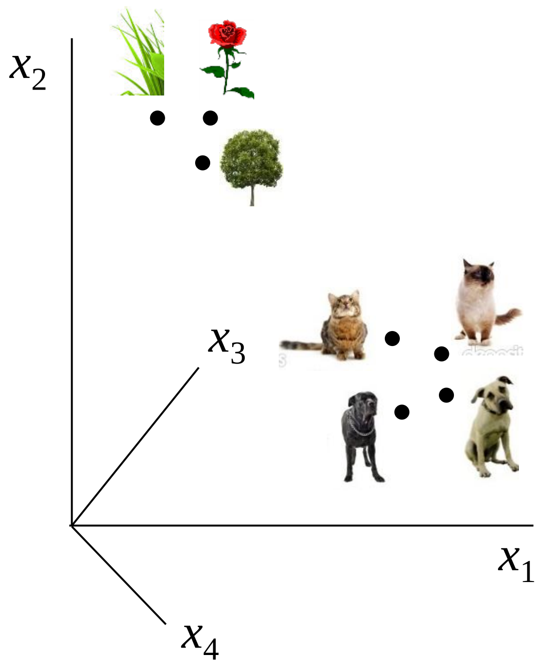

Part I: Unsupervised clustering

pixel intensity space

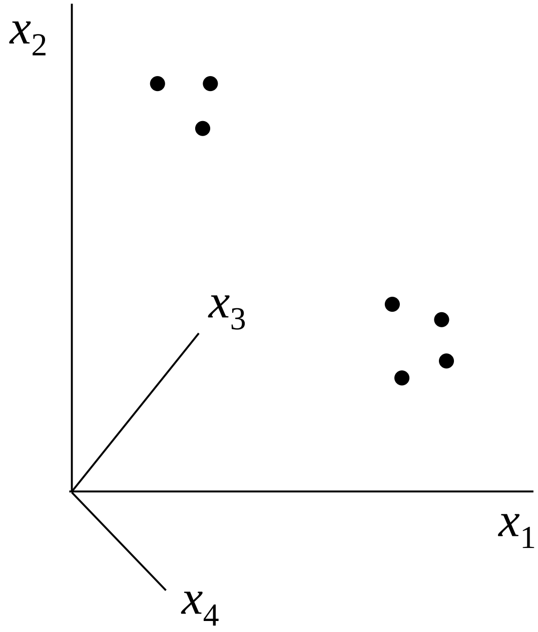

Normative algorithmic approach

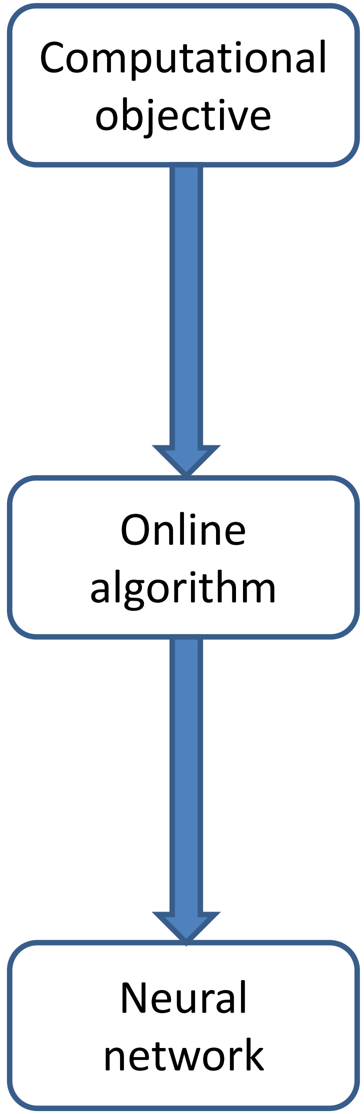

Clustering

Real-time processing, limited memory storage

Dynamics of neural activity and synaptic plasticity, local learning rules

Clustering by similarity alignment

Input

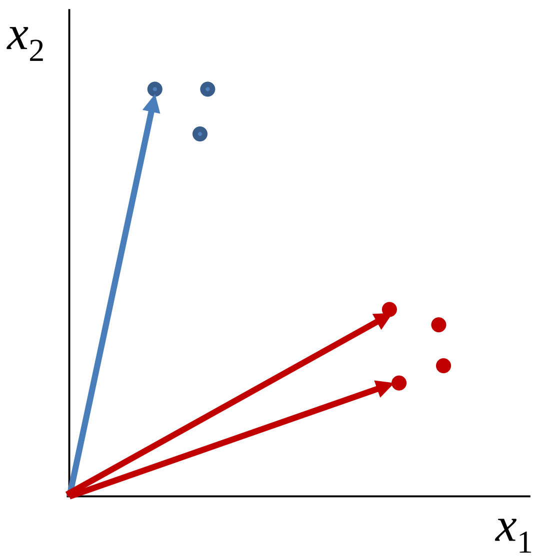

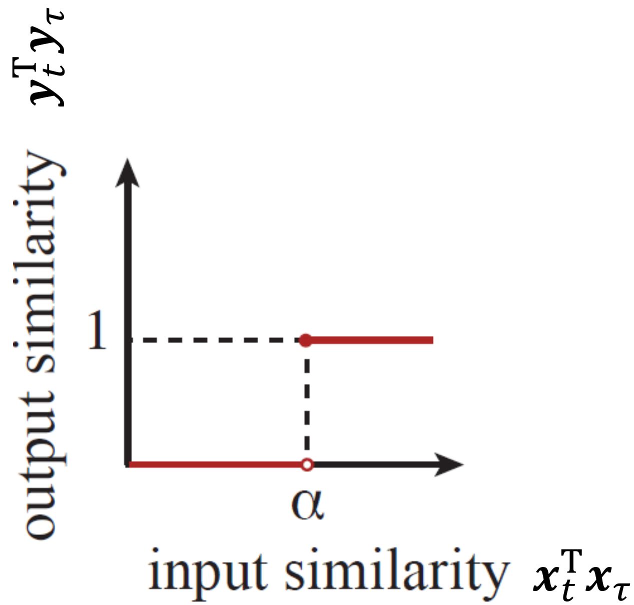

Output

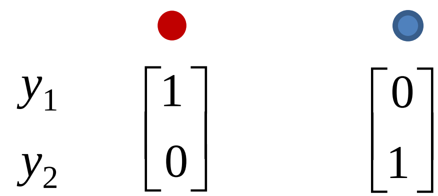

$$\max_{\boldsymbol{y}_t \geq 0, \boldsymbol{y}_\tau \geq 0} \left(\boldsymbol{x}_t^{\mathrm{T}}\boldsymbol{x}_\tau - \alpha\right)\boldsymbol{y}_t^{\mathrm{T}}\boldsymbol{y}_\tau \quad \text{s.t.} \quad \|\boldsymbol{y}_t\| \leq 1, \|\boldsymbol{y}_\tau\| \leq 1$$

Similarity alignment objective

$$\max_{y_t \geq 0, y_\tau \geq 0} \frac{1}{T} \sum_{t=1}^{T} \sum_{\tau=1}^{T} \left(\mathbf{x}_t^\top \mathbf{x}_\tau - \alpha \right) \mathbf{y}_t^\top \mathbf{y}_\tau \quad \text{s.t.} \quad \| \mathbf{y}_t \| \leq 1, \| \mathbf{y}_\tau \| \leq 1$$

similarity of pixel intensity

$$\begin{pmatrix} \mathbf{x}_1^\top \mathbf{x}_1 - \alpha & \mathbf{x}_1^\top \mathbf{x}_2 - \alpha & \cdots & \mathbf{x}_1^\top \mathbf{x}_T - \alpha \\ \mathbf{x}_2^\top \mathbf{x}_1 - \alpha & \mathbf{x}_2^\top \mathbf{x}_2 - \alpha & \cdots & \mathbf{x}_2^\top \mathbf{x}_T - \alpha \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{x}_T^\top \mathbf{x}_1 - \alpha & \mathbf{x}_T^\top \mathbf{x}_2 - \alpha & \cdots & \mathbf{x}_T^\top \mathbf{x}_T - \alpha \end{pmatrix}$$

similarity of neural activity

$$\begin{pmatrix} \mathbf{y}_1^\top \mathbf{y}_1 & \mathbf{y}_1^\top \mathbf{y}_2 & \cdots & \mathbf{y}_1^\top \mathbf{y}_T \\ \mathbf{y}_2^\top \mathbf{y}_1 & \mathbf{y}_2^\top \mathbf{y}_2 & \cdots & \mathbf{y}_2^\top \mathbf{y}_T \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}_T^\top \mathbf{y}_1 & \mathbf{y}_T^\top \mathbf{y}_2 & \cdots & \mathbf{y}_T^\top \mathbf{y}_T \end{pmatrix}$$

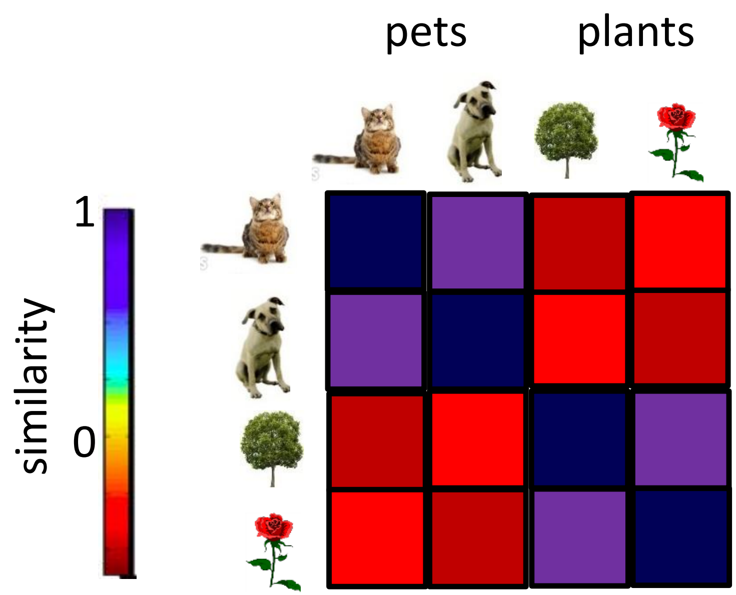

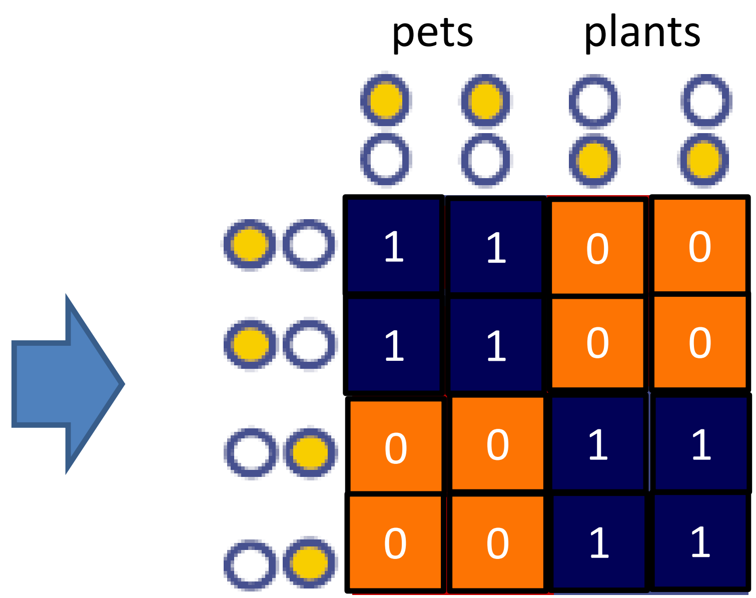

Deriving a neural network

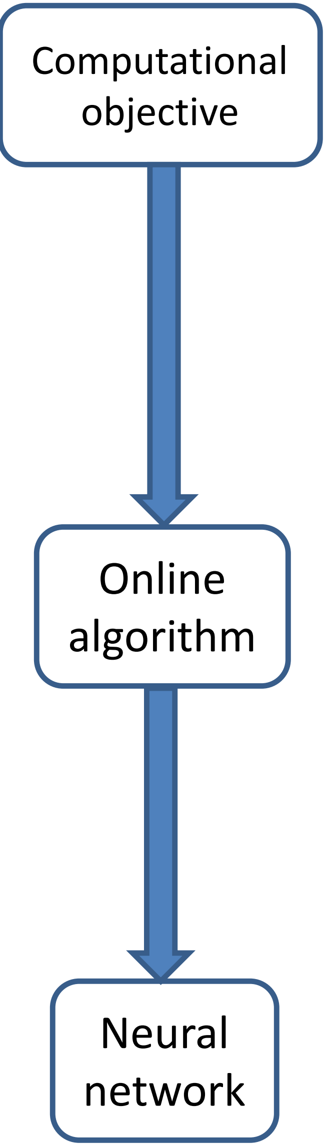

$$\max_{\boldsymbol{y}_t \geq 0} \min_{\boldsymbol{z}_t \geq 0} \frac{1}{T} \sum_{t=1}^{T} \sum_{\tau=1}^{T} \left(\boldsymbol{x}_t^{\top} \boldsymbol{x}_\tau - \alpha \right) \boldsymbol{y}_t^{\top} \boldsymbol{y}_\tau - \left(\boldsymbol{y}_t^{\top} \boldsymbol{y}_\tau - 1 \right) \boldsymbol{z}_t^{\top} \boldsymbol{z}_\tau$$

$$\frac{1}{T} \sum_{t=1}^{T} \sum_{\tau=1}^{T} \boldsymbol{x}_t^{\top} \boldsymbol{x}_\tau \boldsymbol{y}_t^{\top} \boldsymbol{y}_\tau = \sum_{t=1}^{T} \boldsymbol{y}_t^{\top} \left(\frac{1}{T} \sum_{\tau=1}^{T} \boldsymbol{y}_\tau \boldsymbol{x}_\tau^{\top} \right) \boldsymbol{x}_t = \sum_{t=1}^{T} \boldsymbol{y}_t^{\top} \boldsymbol{W}^{YX} \boldsymbol{x}_t$$

same for $\boldsymbol{z}_t, \boldsymbol{W}^{YZ}$

neural activity: $\boldsymbol{y}_t \leftarrow \left[\boldsymbol{y}_t + \gamma \left(\boldsymbol{W}^{YX} \boldsymbol{x}_t - \boldsymbol{W}^{YZ} \boldsymbol{z}_t - \alpha \boldsymbol{b}_y \right) \right]_+$

synaptic plasticity: $\boldsymbol{W}_{i,j}^{YX} \leftarrow \boldsymbol{W}_{i,j}^{YX} + \eta \left(y_{t,i} x_{t,j} - \boldsymbol{W}_{i,j}^{YX} \right)$

Local learning rule!

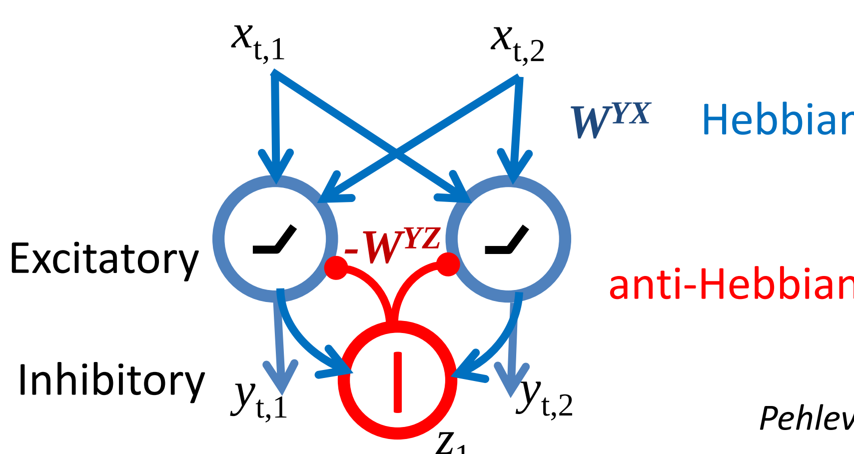

Pehlevan, Genkin & Chklovskii (2017)

Similarity alignment can (softly) cluster

input, ***x***

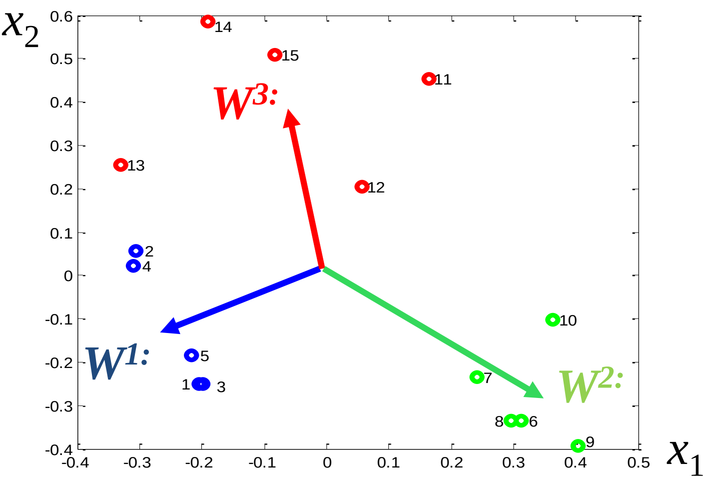

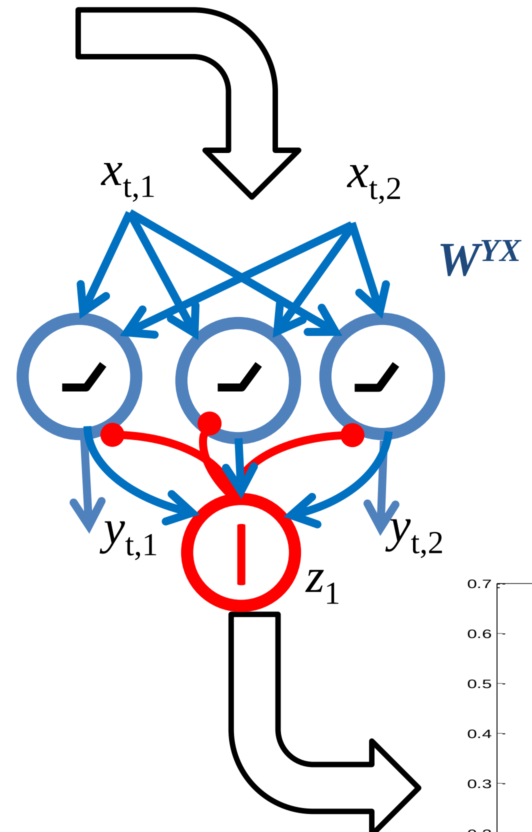

output, ***y***

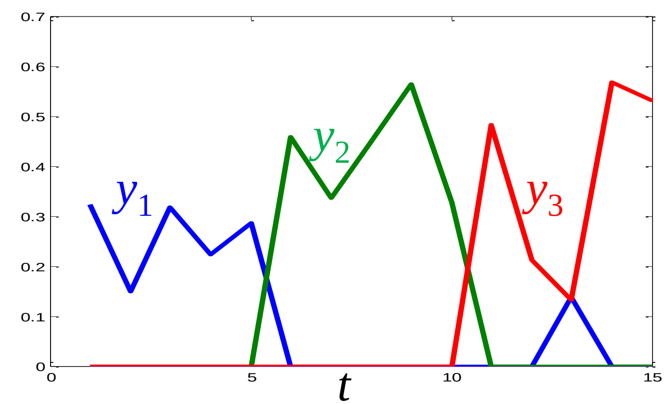

Experimental test in fly larva antennal lobe

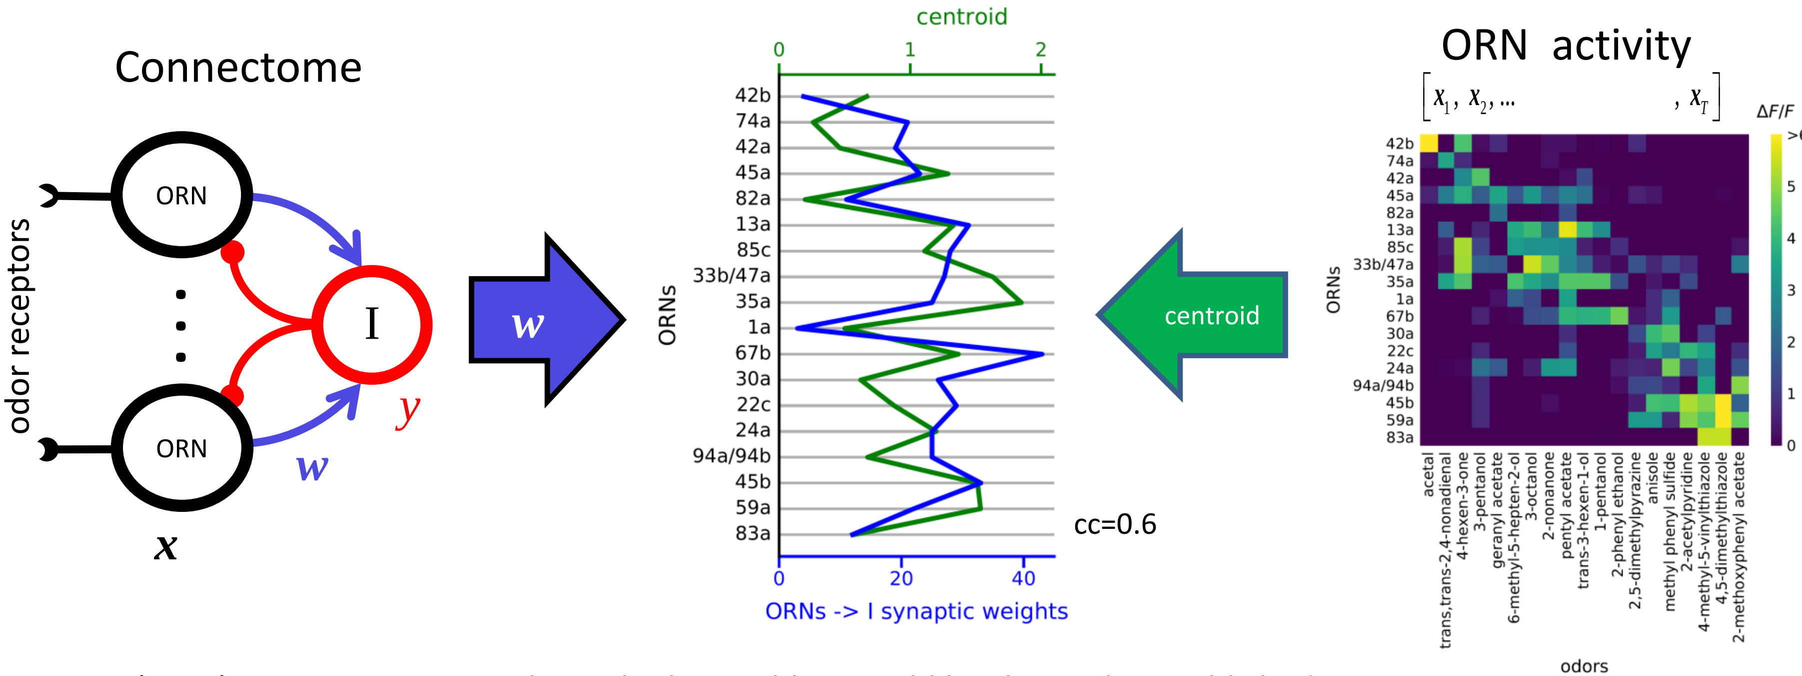

Berck et al, 2016   *Chapochnikov, Pehlevan, Chklovskii et al, unpublished*   *Samuel lab*

Clustering model of insect olfaction

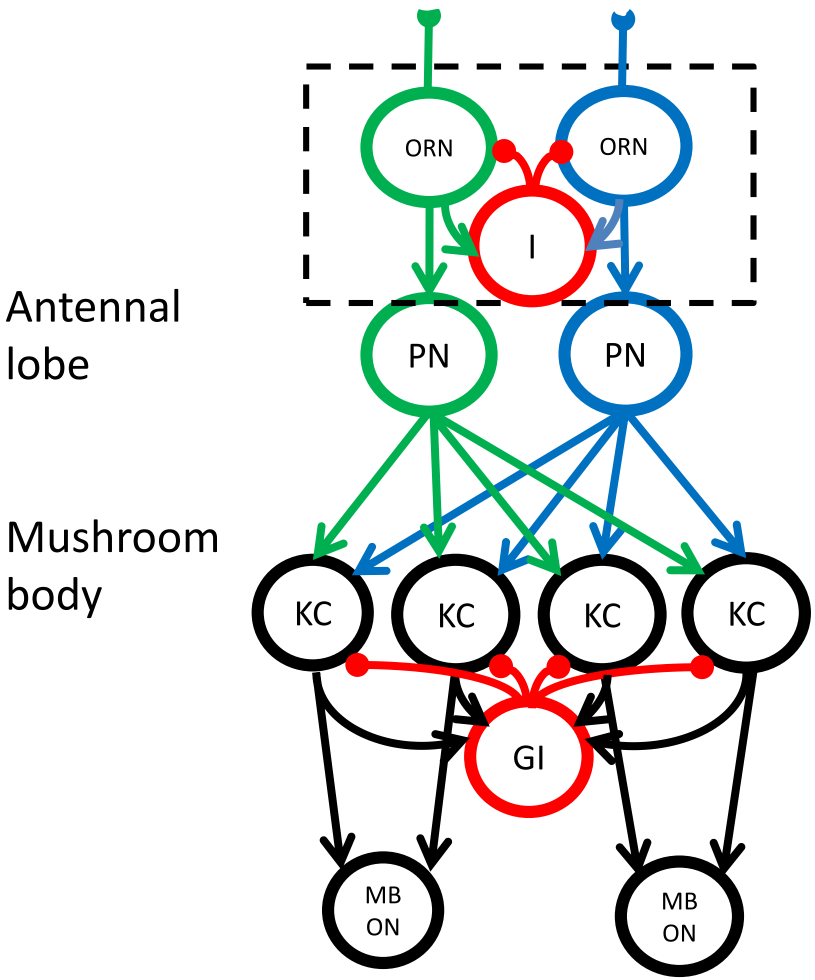

Decorrelation by interneuron:
ORN->I synaptic weight vector
~ centroid of neural activity

Clustering model of insect olfaction

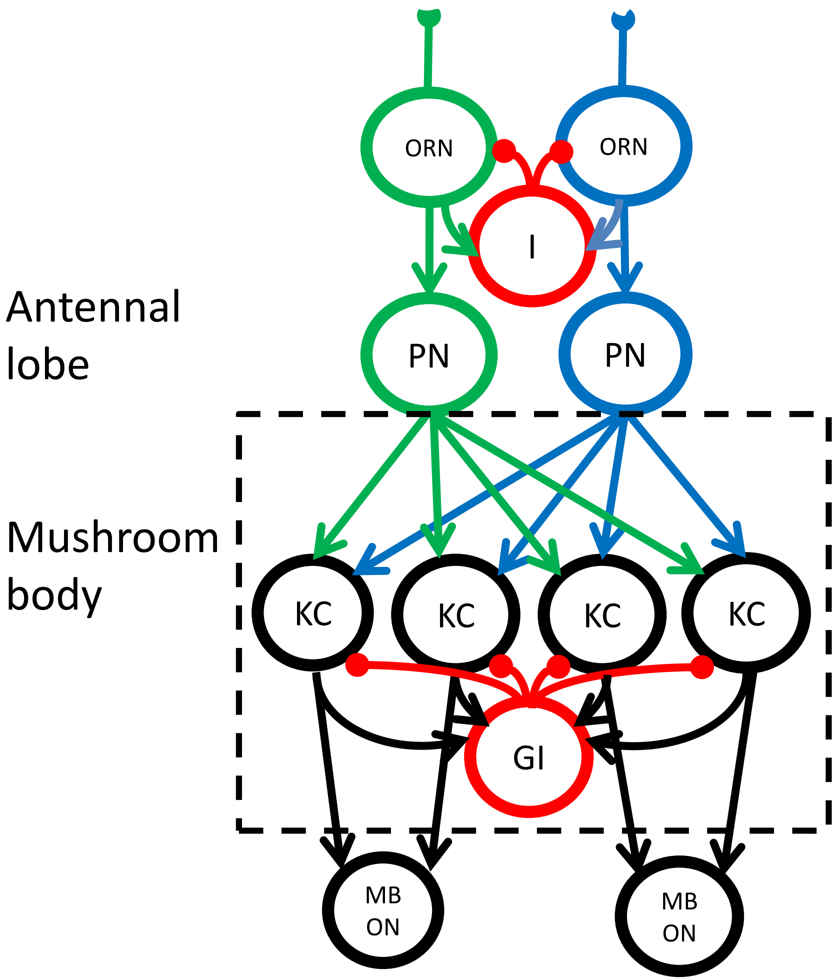

- Rectification by KCs
- Single giant interneuron (GI)
- Non-random connectivity (Eichler et al, 2017)
- Sparse over-complete representation = soft clustering

Pehlevan, Genkin & Chklovskii (2017)

Part II: Manifold learning by similarity alignment

Categorization as manifold disentanglement

DiCarlo & Cox, 2007

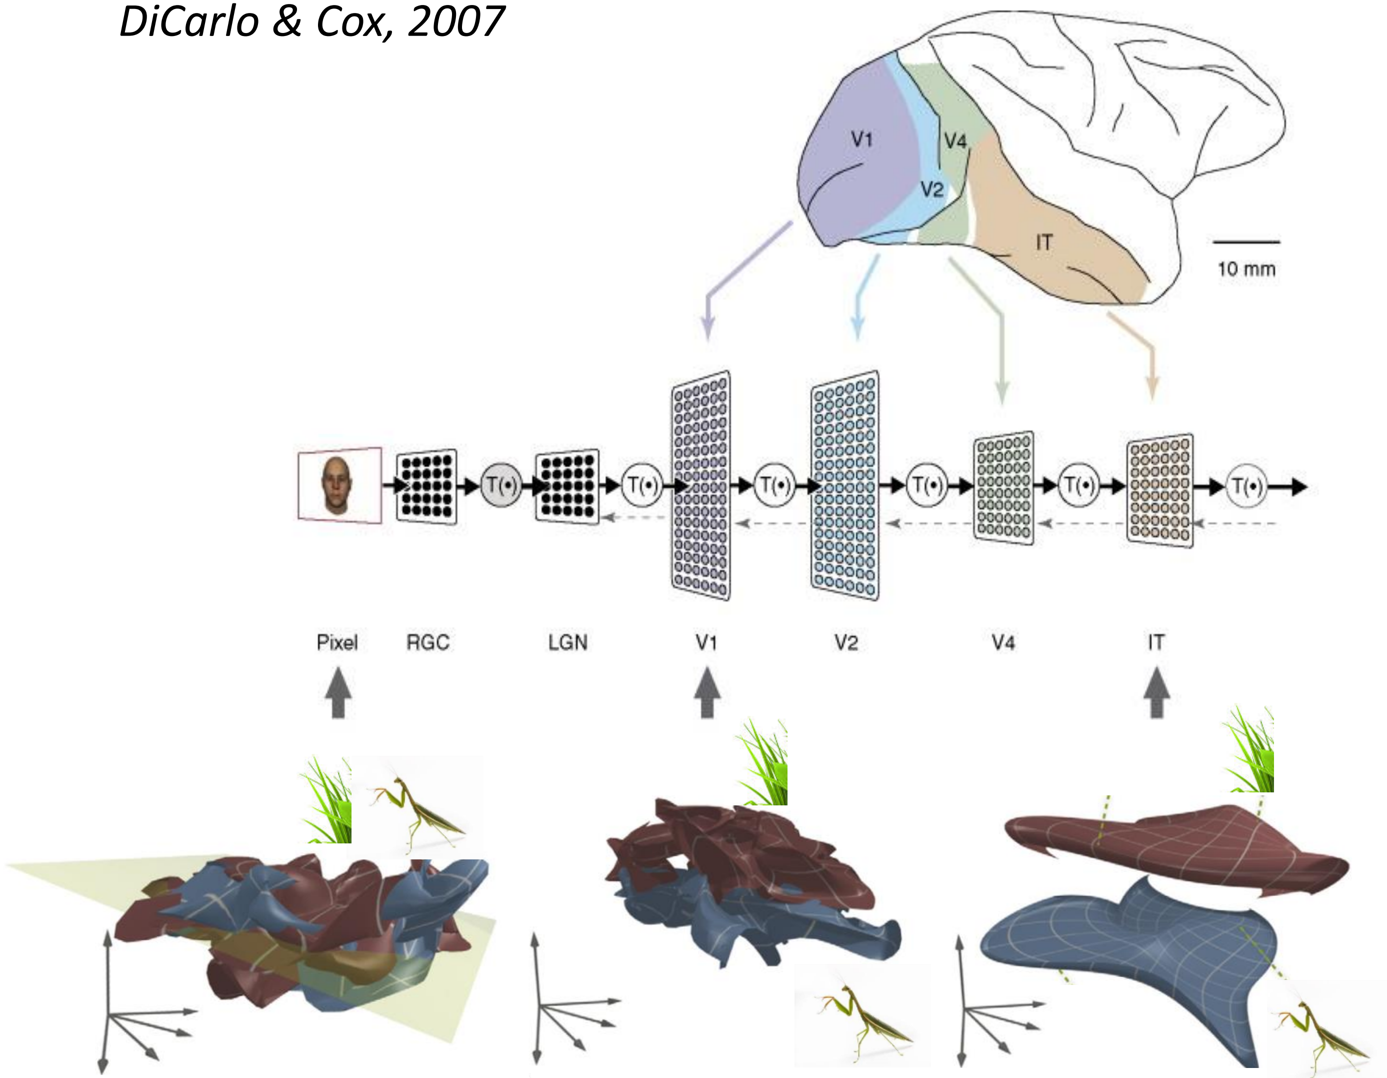

pixel intensity space

V1 neural activity space

IT neural activity space

Manifold learning

$$\max_{\boldsymbol{y}_t\geq 0,\boldsymbol{y}_\tau\geq 0} \frac{1}{T}\sum_{t=1}^{T}\sum_{\tau=1}^{T}\left(\boldsymbol{x}_t^{\top}\boldsymbol{x}_\tau-\alpha\right)\boldsymbol{y}_t^{\top}\boldsymbol{y}_\tau \quad \text{s.t.} \quad \|\boldsymbol{y}_t\|\leq 1, \|\boldsymbol{y}_\tau\|\leq 1$$

$$\max_{\boldsymbol{y}_t\geq 0,\boldsymbol{y}_\tau\geq 0} \frac{1}{T}\sum_{t=1}^{T}\sum_{\tau=1}^{T}\left(\alpha-\|\boldsymbol{x}_t-\boldsymbol{x}_\tau\|_2^2\right)\boldsymbol{y}_t^{\top}\boldsymbol{y}_\tau \quad \text{s.t.} \quad \|\boldsymbol{y}_t\|\leq 1, \|\boldsymbol{y}_\tau\|\leq 1$$

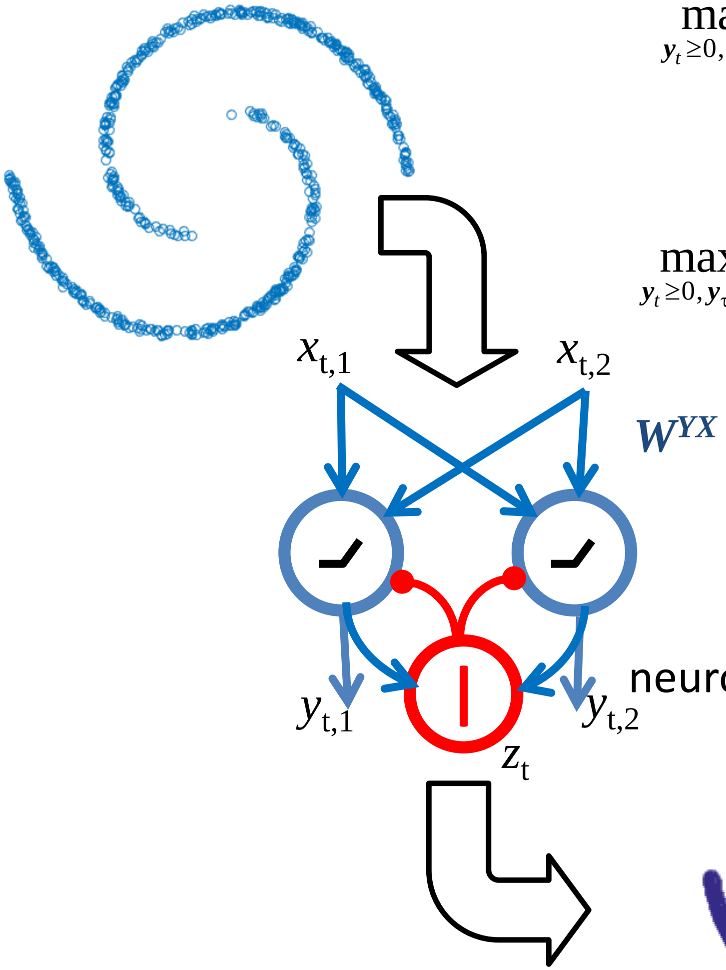

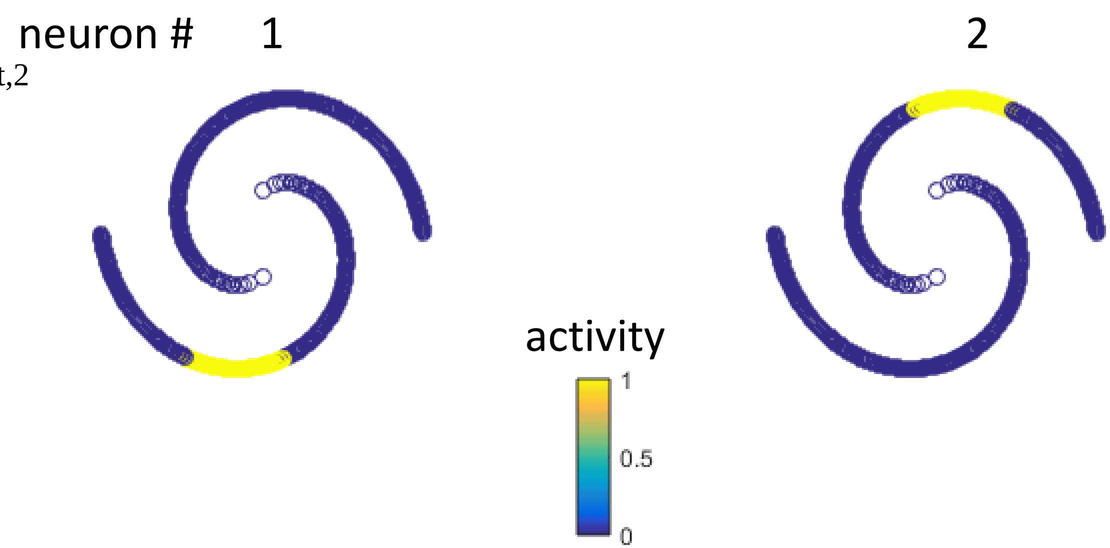

Similarity alignment learns manifolds

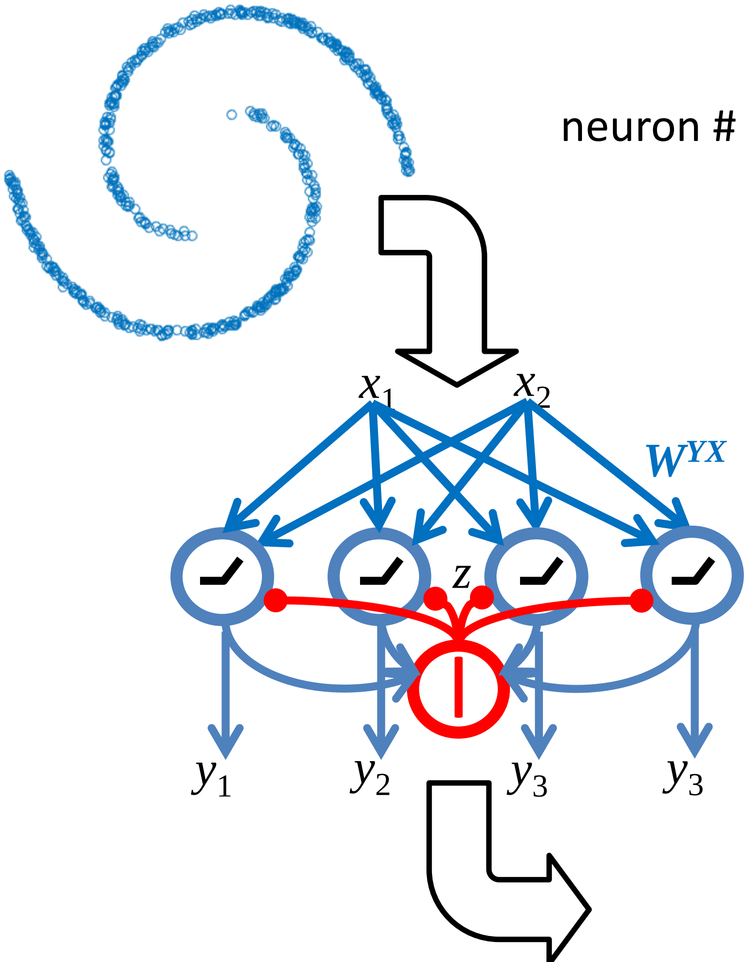

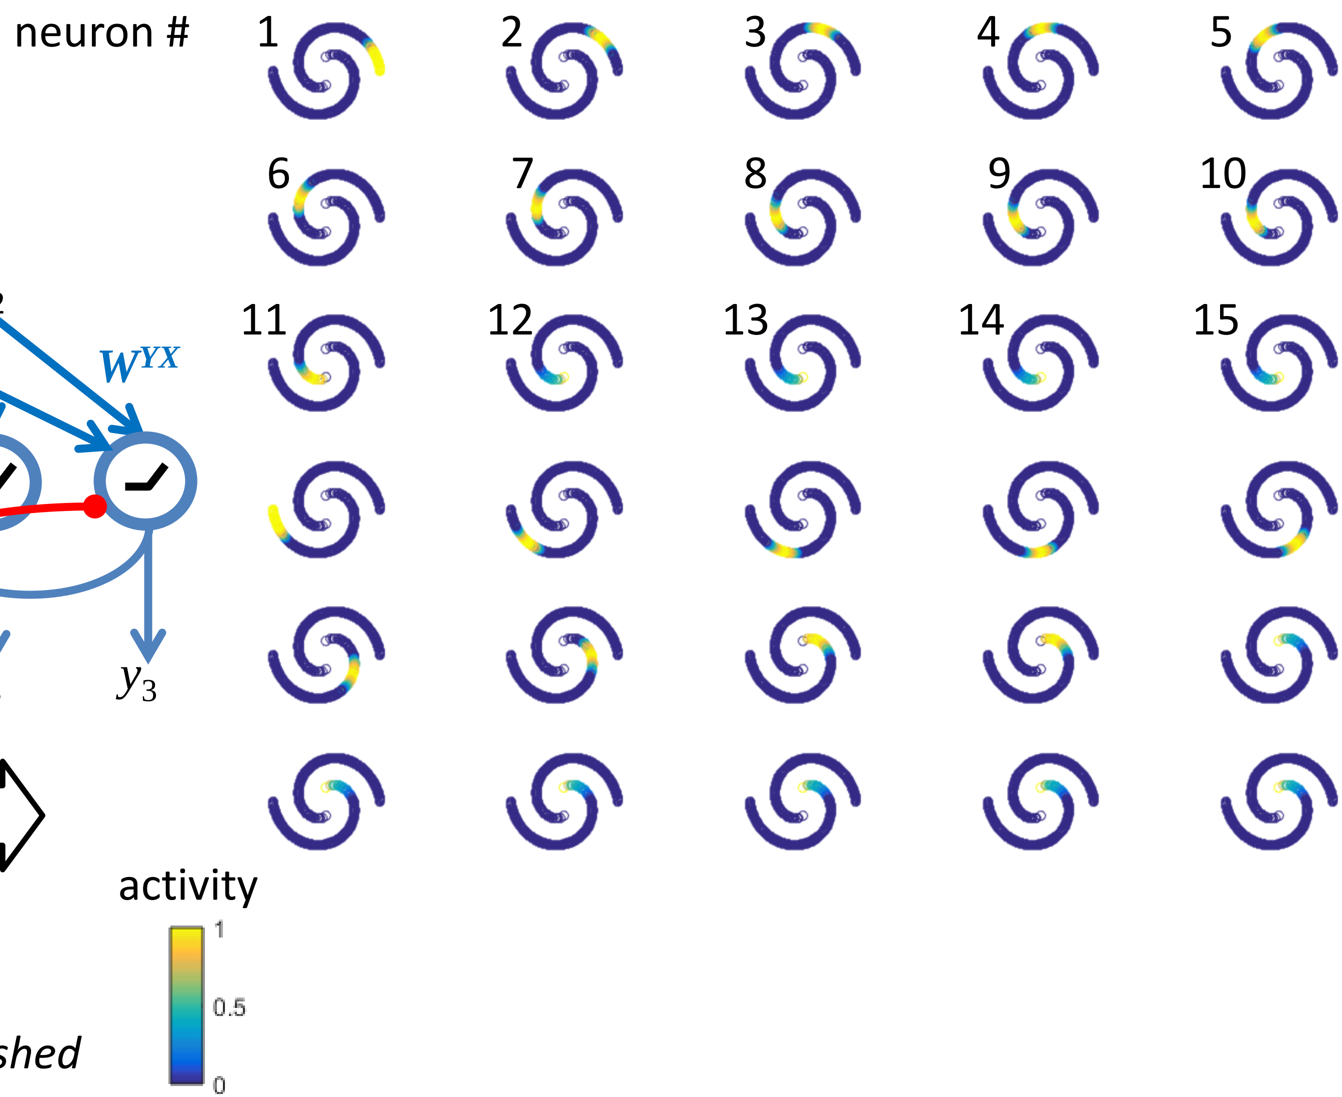

Pehlevan et al, unpublished

Similarity alignment network learns V1 features from natural images

natural images

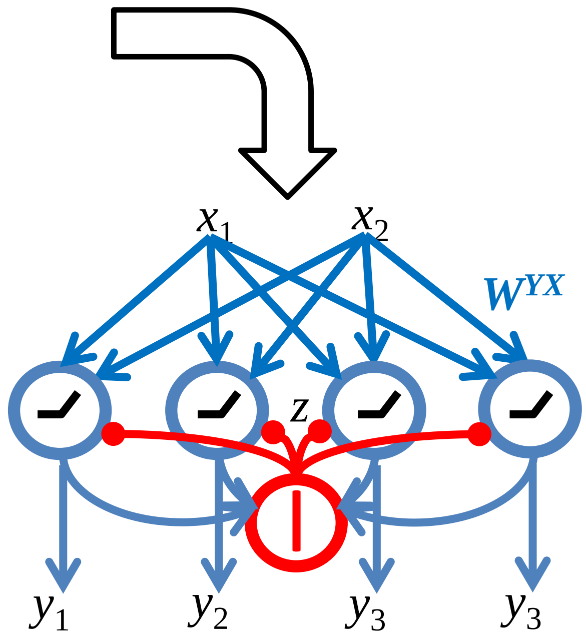

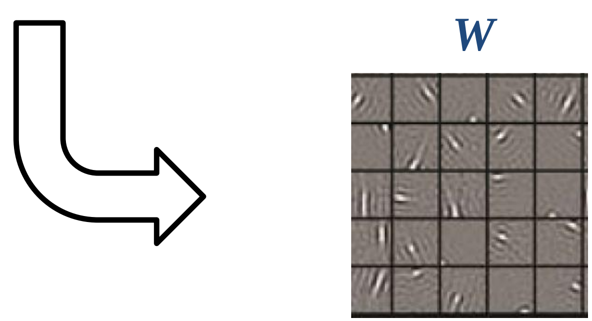

Pehlevan & Chklovskii (2014)

Manifold learning

Similarity along the visual pathway

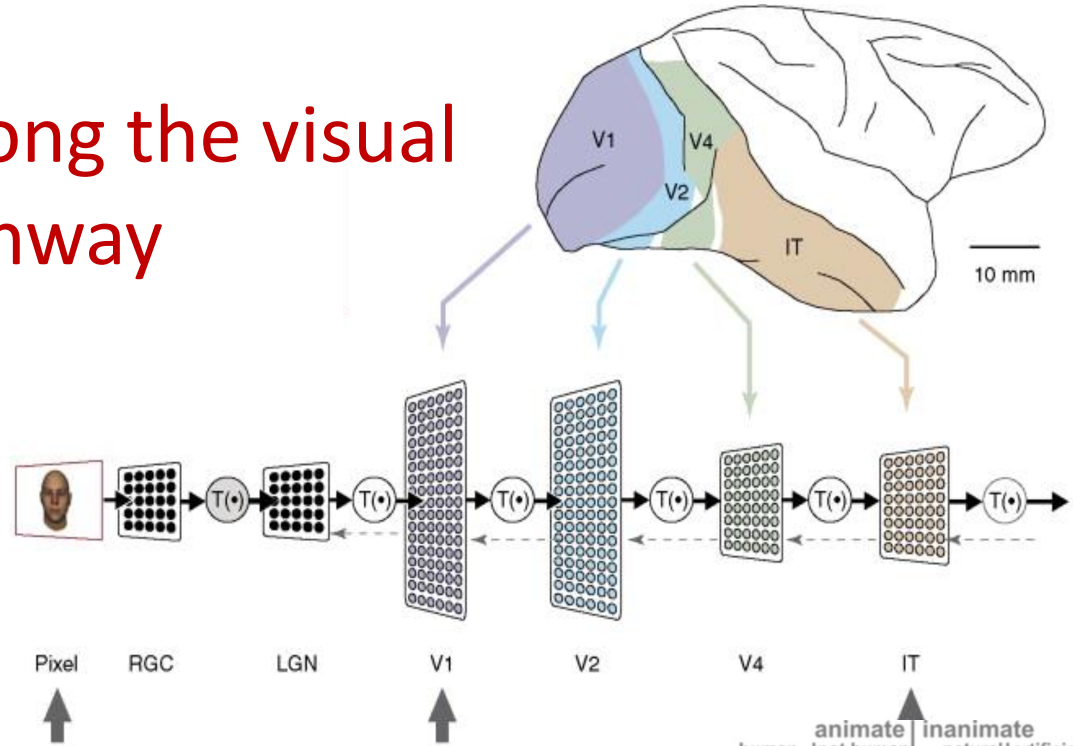

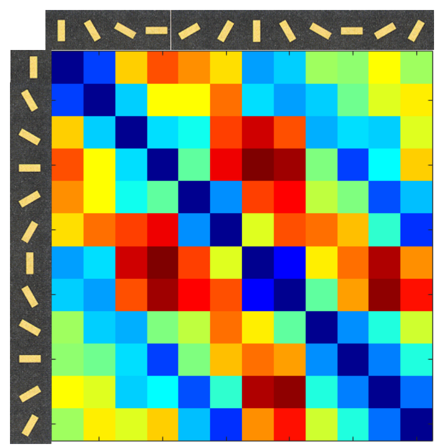

Kiani lab

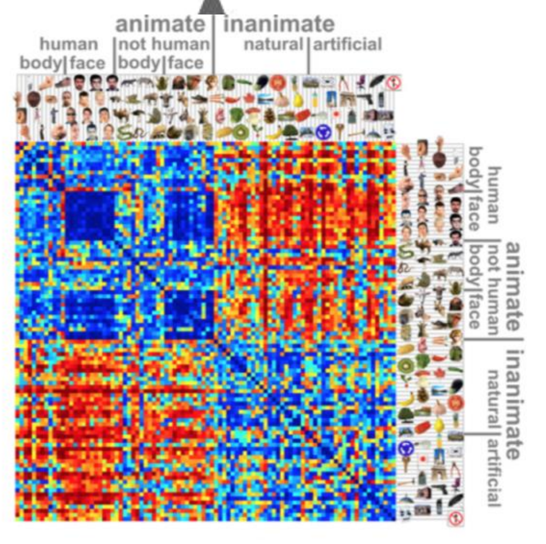

Kriegeskorte et. al., 2008

Unsupervised manifold disentangling and clustering

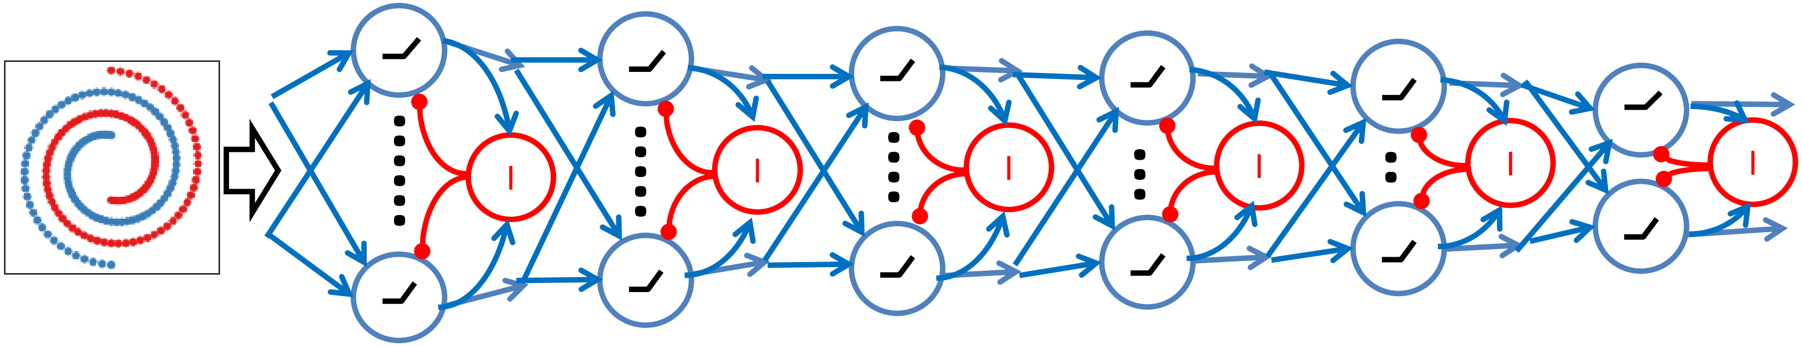

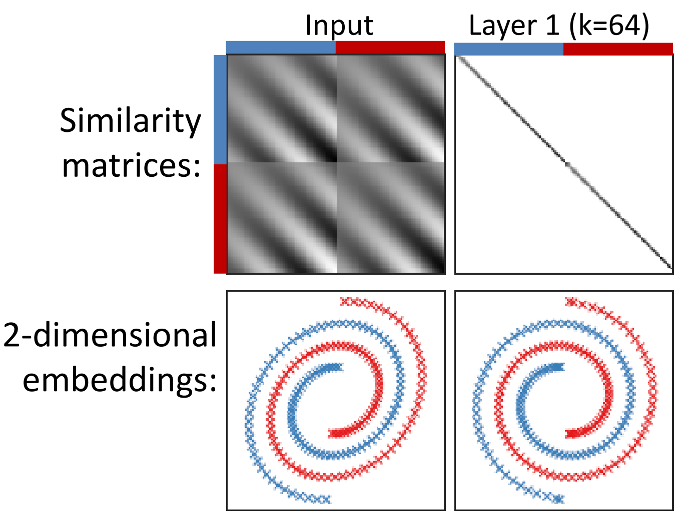

Tepper, Sengupta & Chklovskii (2017)
Pehlevan, Genkin, Chklovskii (2017) and unpublished

Family of similarity alignment networks

BIOLOGICAL FEATURE	**MATHEMATICAL CONSTRUCT**
Hebbian plasticity	Similarity alignment
Neural rectification, local RFs	Nonnegativity constraint
Adaptive neural thresholds	Rank and sparsity regularizers
Anti-Hebbian inhibitory interneurons	Constrained output similarity matrix
2-compartment neuron	Canonical correlation analysis (CCA)
…	…

Hu, Pehlevan & Chklovskii (2014), Pehlevan & Chklovskii (2014), Pehlevan, Hu, Chklovskii (2015), Pehlevan & Chklovskii (2015), Pehlevan & Chklovskii (2016), Pehlevan, Mohan & Chklovskii (2017), Pehlevan, Sengupta & Chklovskii (2017), Tepper, Sengupta & Chklovskii (2017), Pehlevan, Genkin & Chklovskii (2017), Bahroun, Hunsicker, Soltoggio (2017), Seung & Zung (2017)

The similarity alignment approach yields biologically plausible networks for nontrivial computations

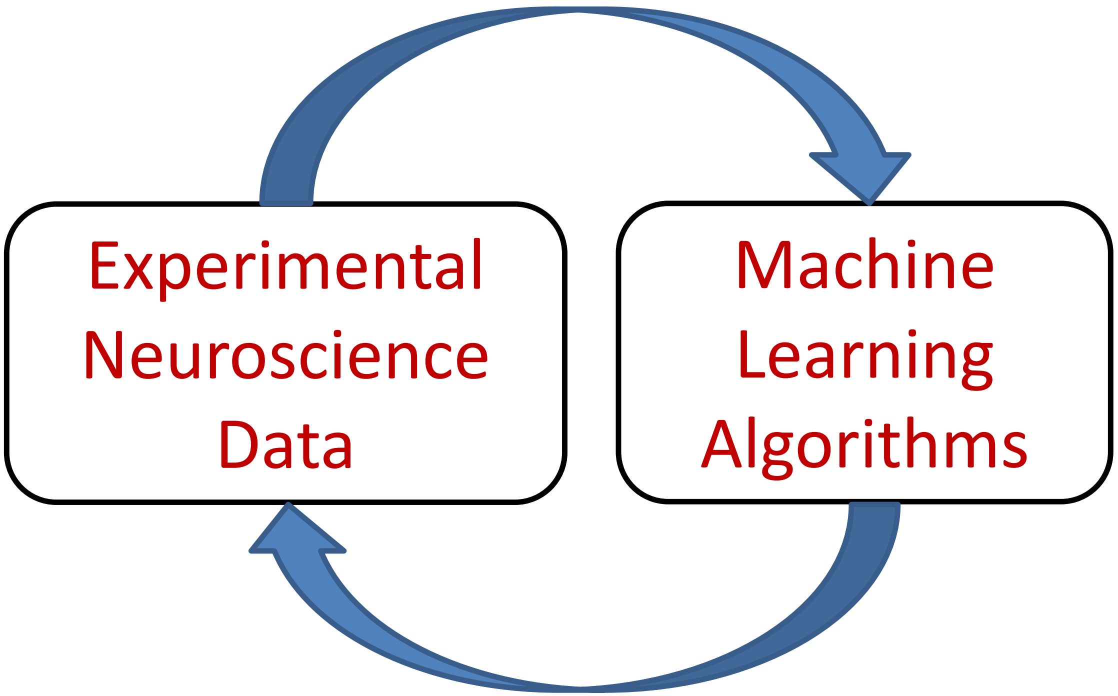

Acknowledgements

Cengiz Pehlevan

Mariano Tepper

Alex Genkin

Anirvan Sengupta

Victor Minden

Nikolai Chapochnikov

Tao Hu

Sreyas Mohan