Organization

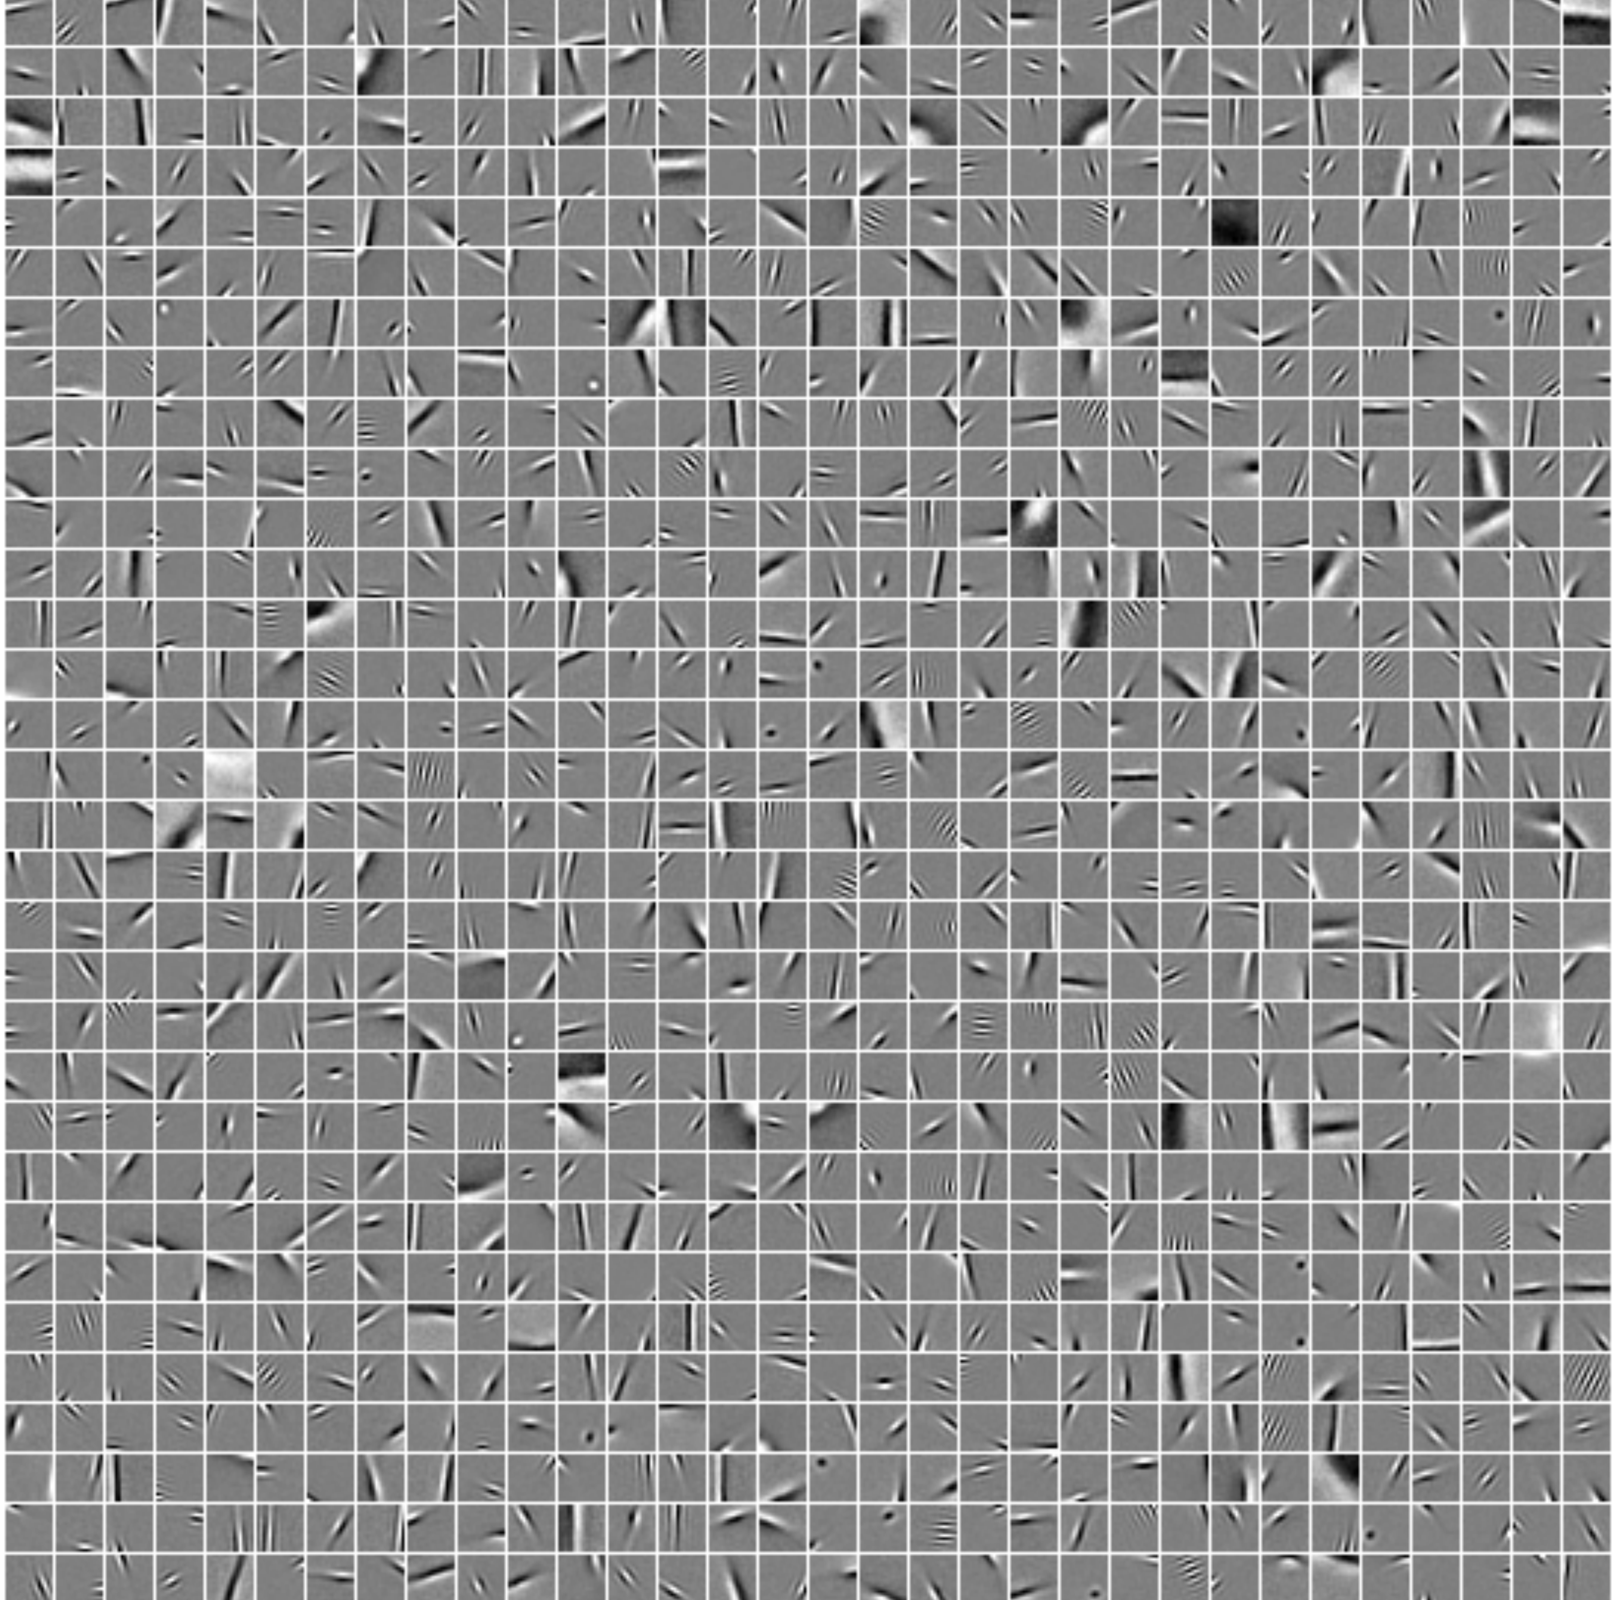

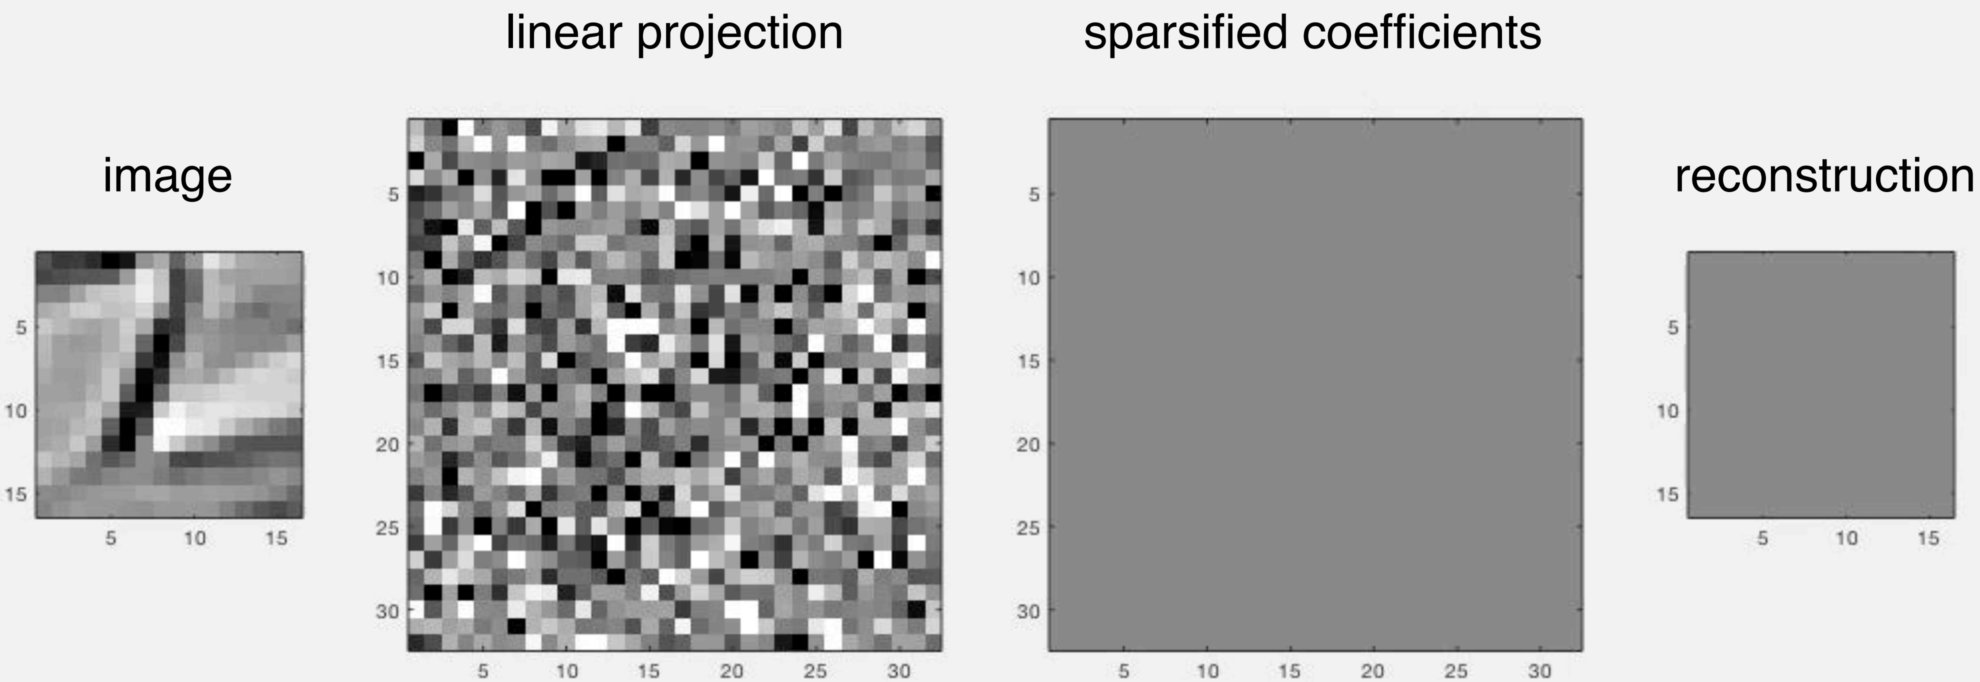

linear projection
sparsified coefficients
image
reconstruction

Visual cortex

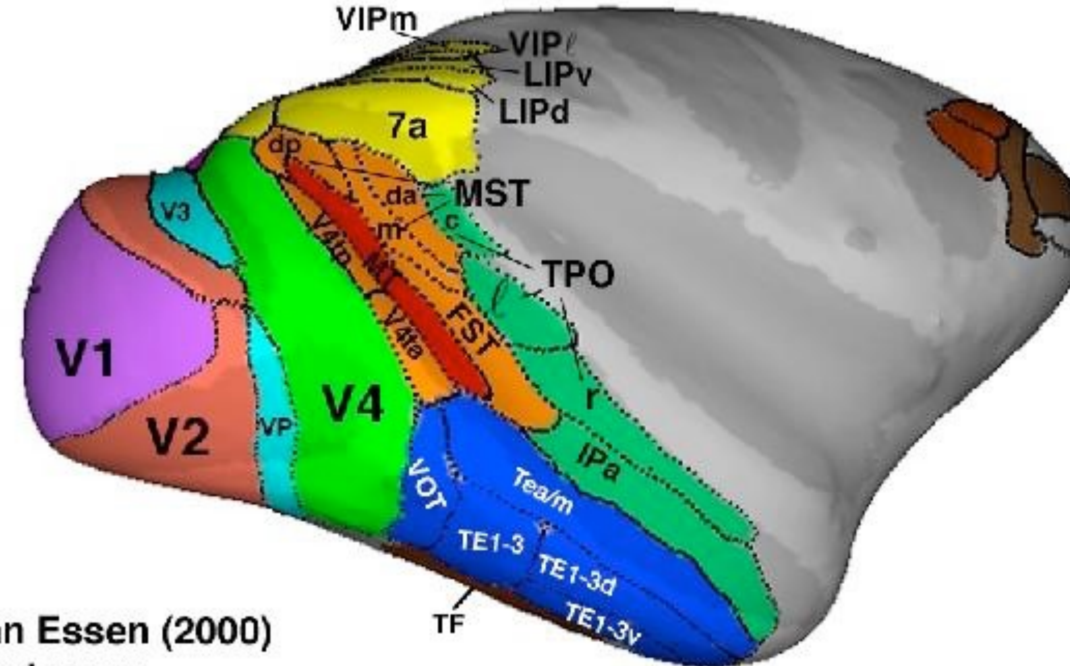

Lewis & Van Essen (2000)
Visual areas

V1
topographic
map

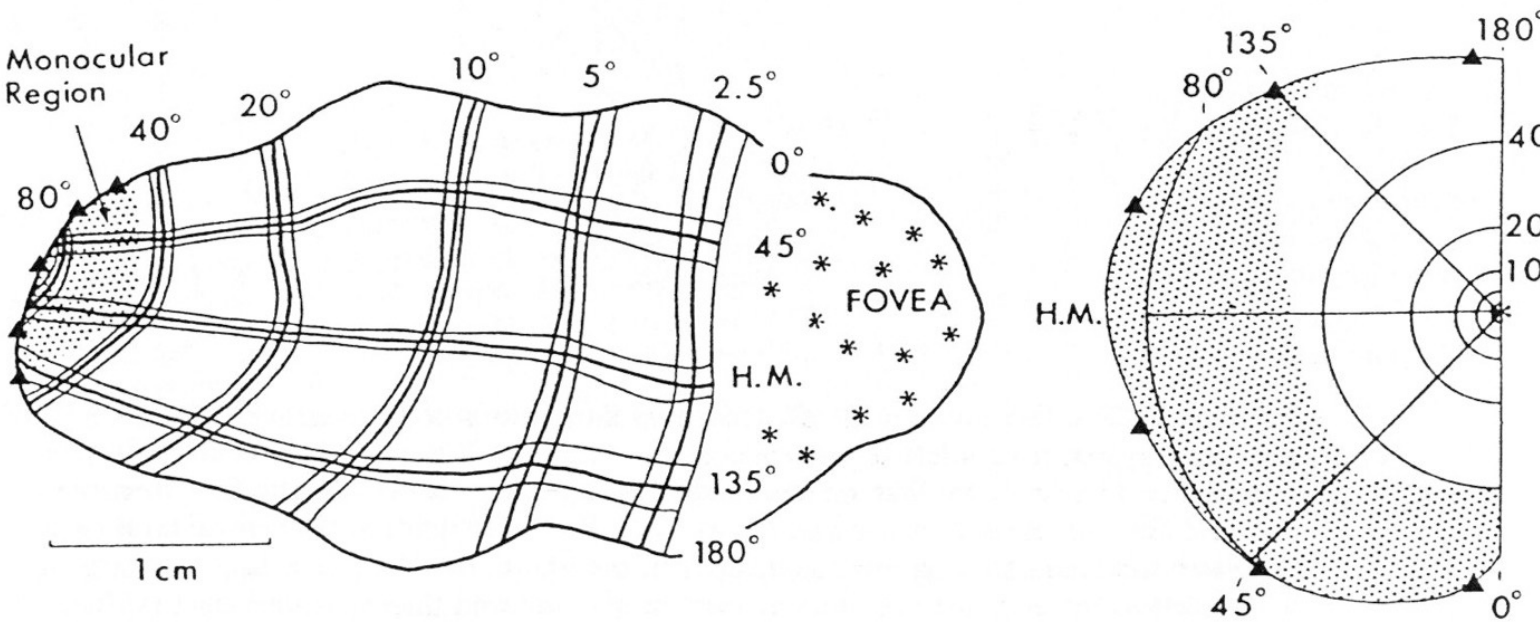

Somatosensory cortex

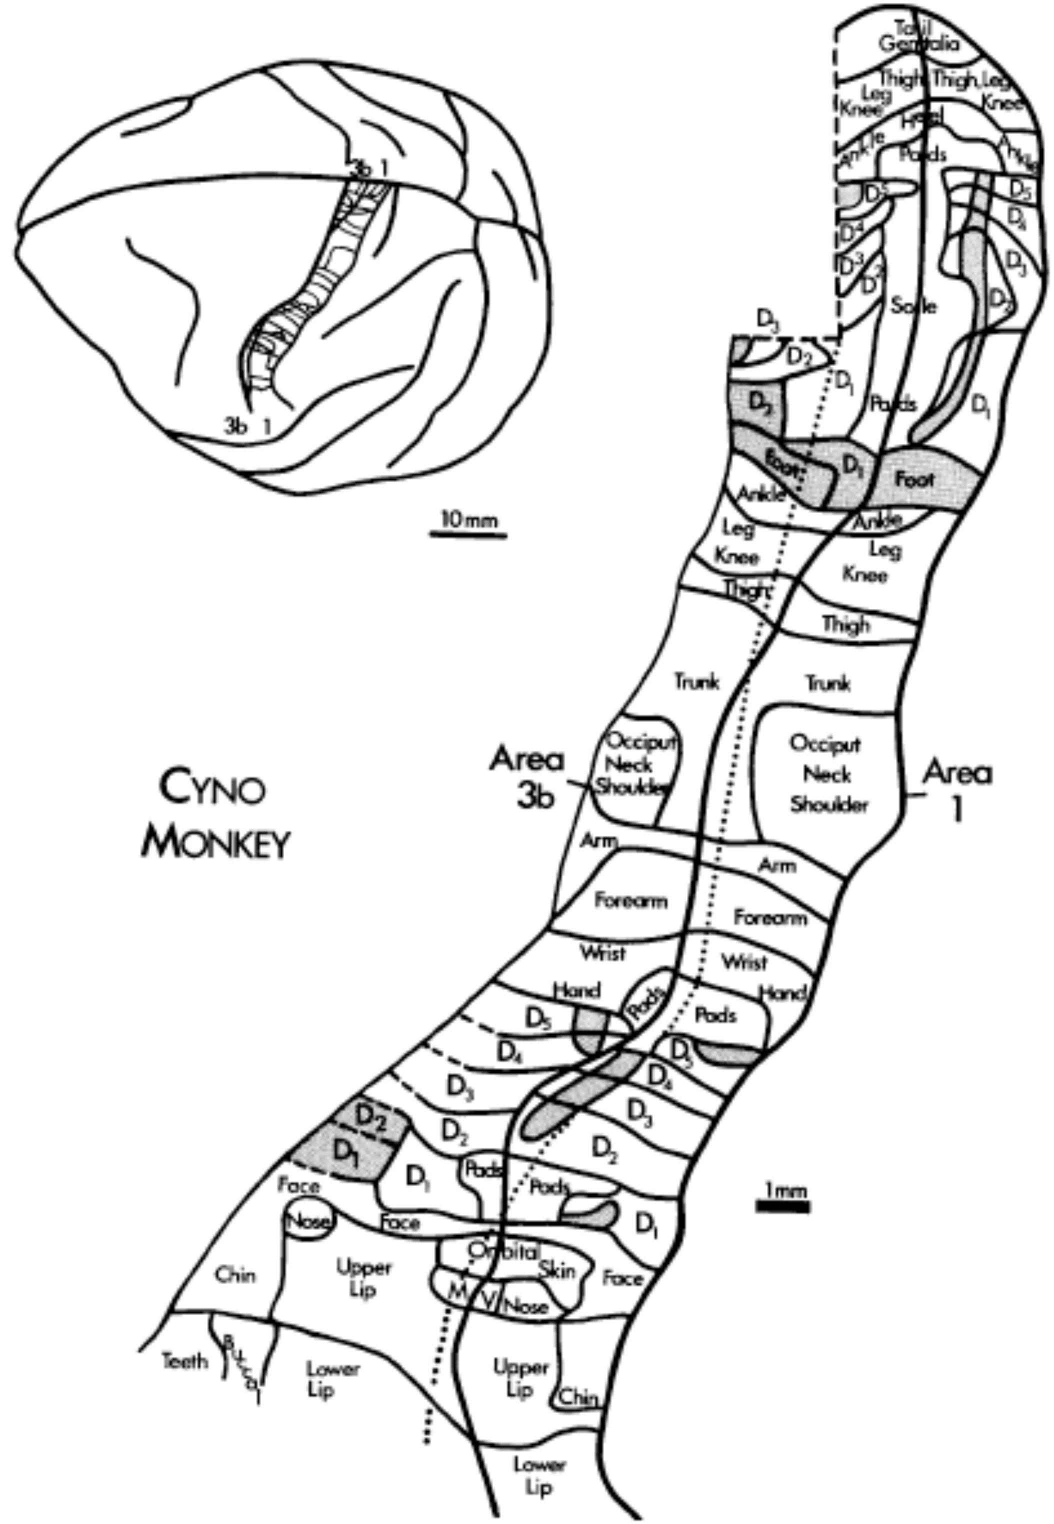

Orientation selectivity changes in a regular manner across the surface of V1

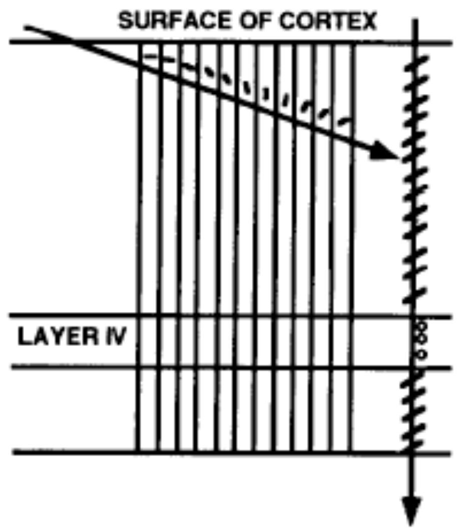

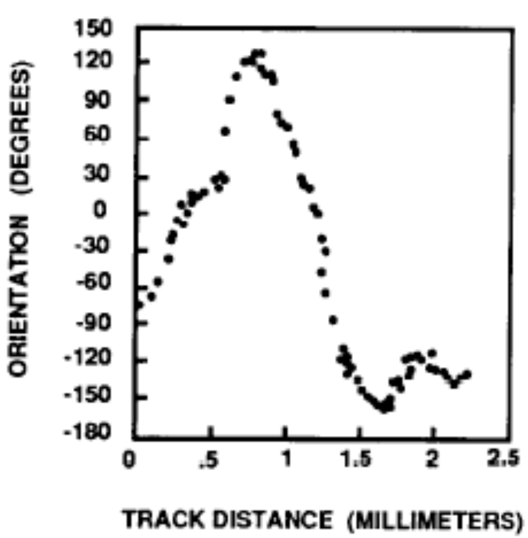

Hubel & Wiesel's "ice cube tray model"

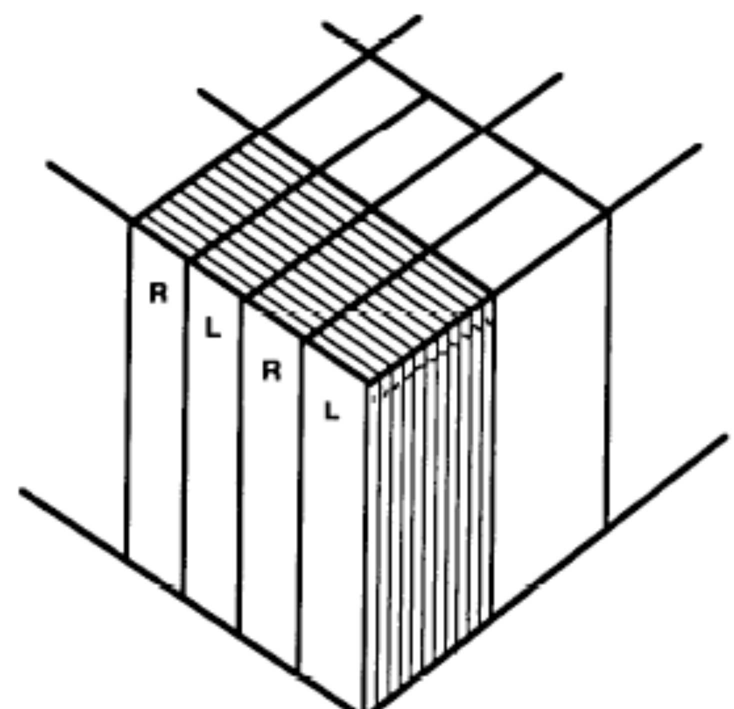

Orientation maps in V1 (Blasdel 1992)

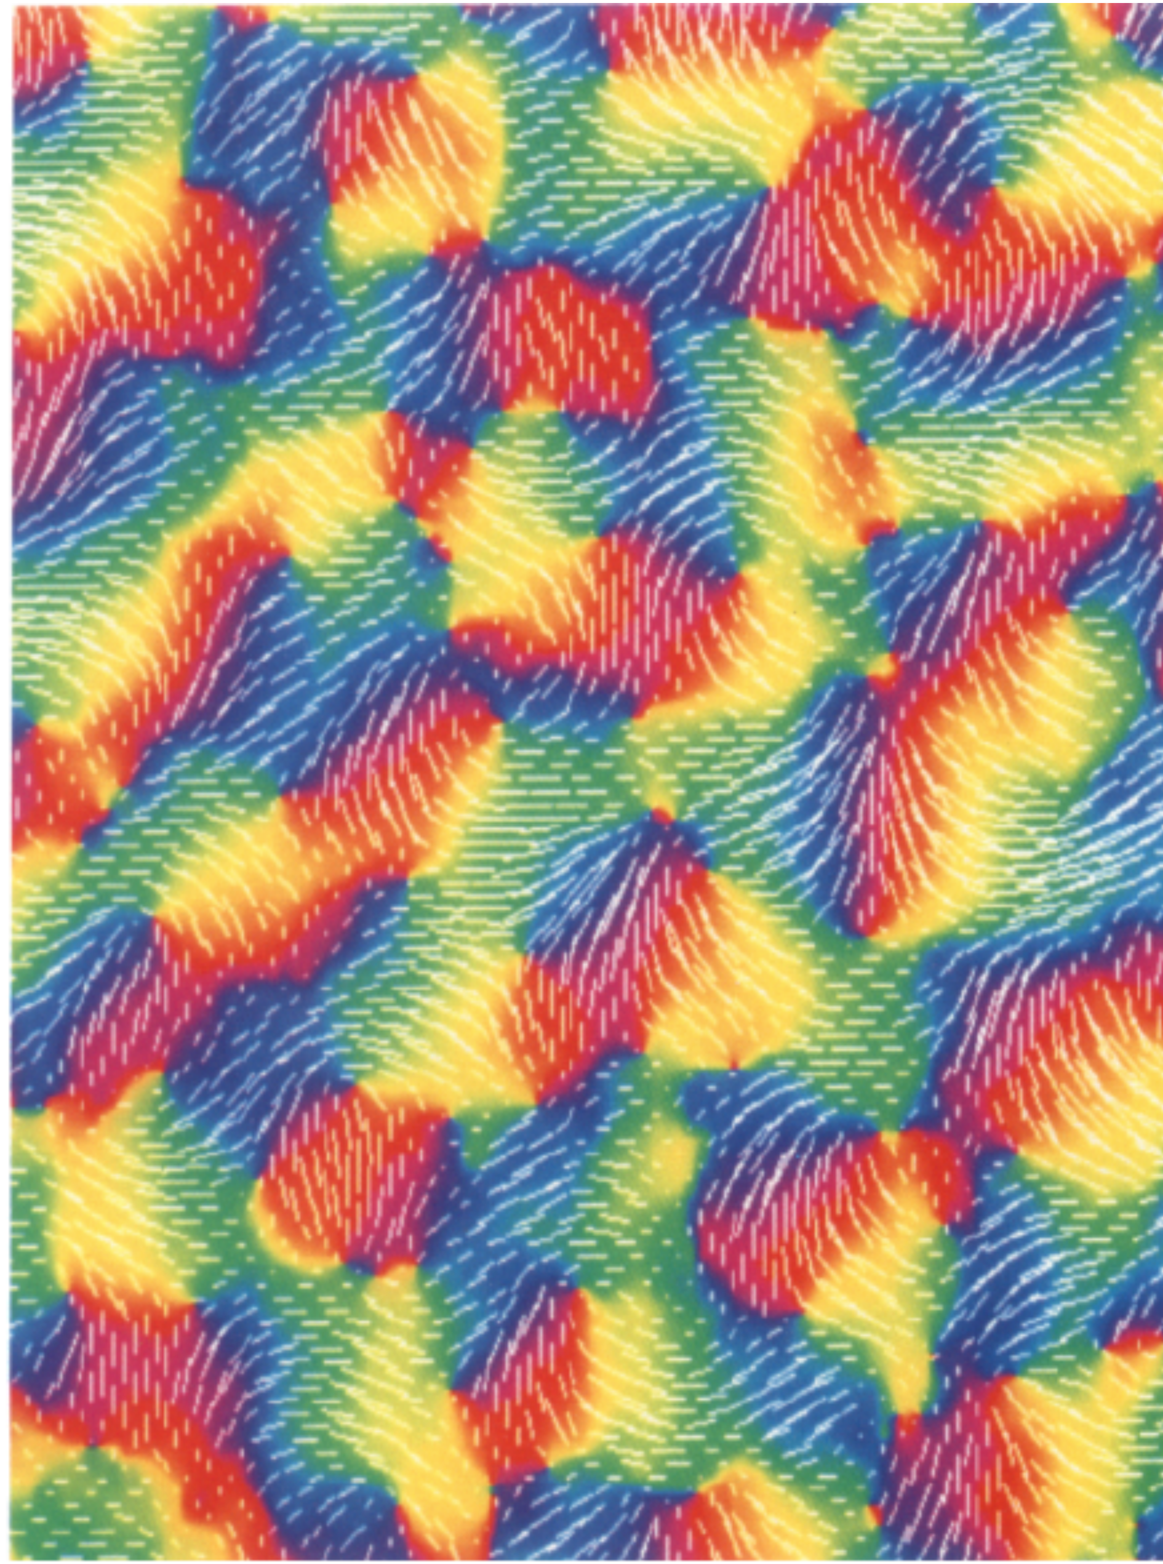

Joint map of orientation and ocular dominance demonstrates tradeoff between feature diversity vs. smoothness within each feature dimension (from Blasdel 1992)

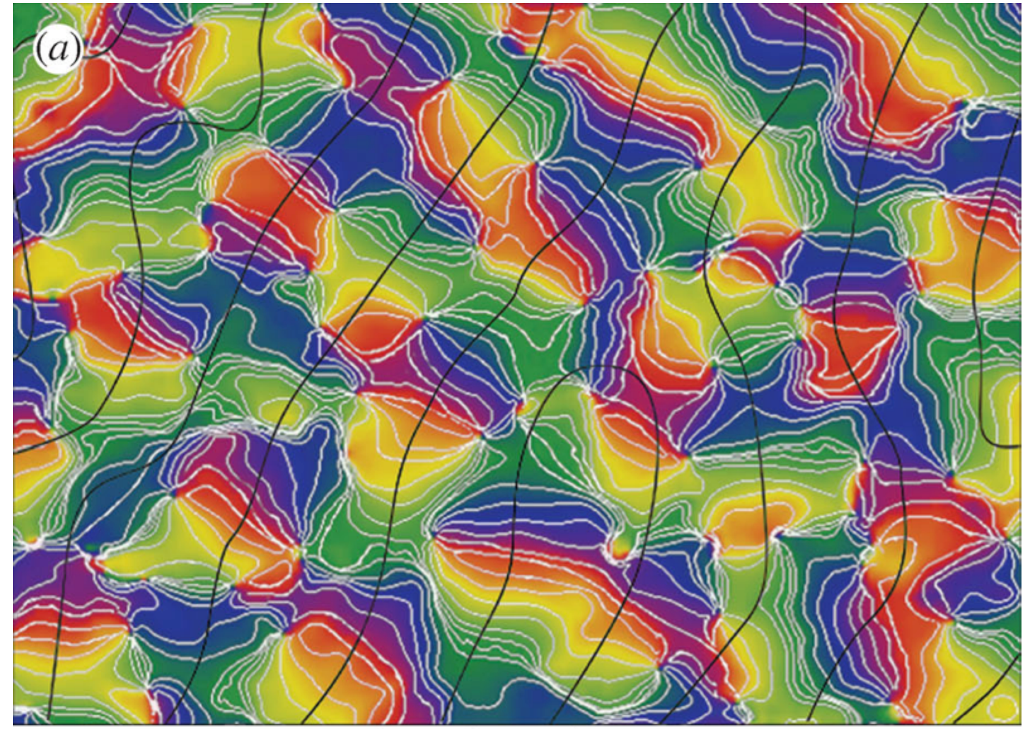

orientation columns

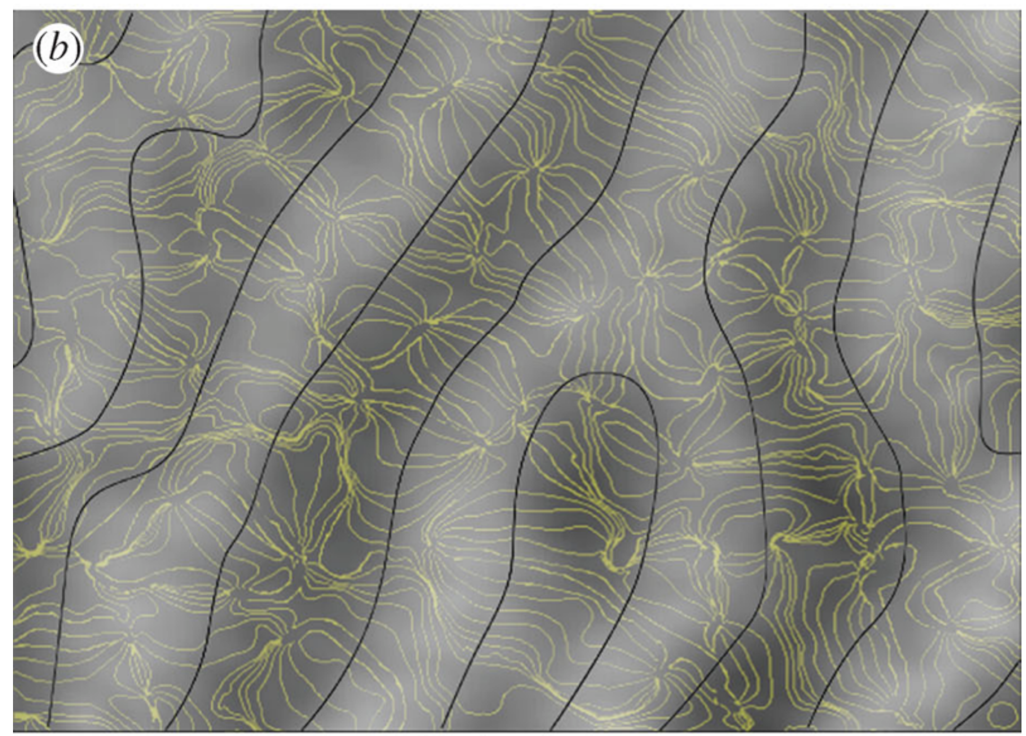

ocular dominance

Direction selective cells in cat area 18

(Ca++ imaging - Ohki, Chung, Ch'ng, Kara, Reid, 2005)

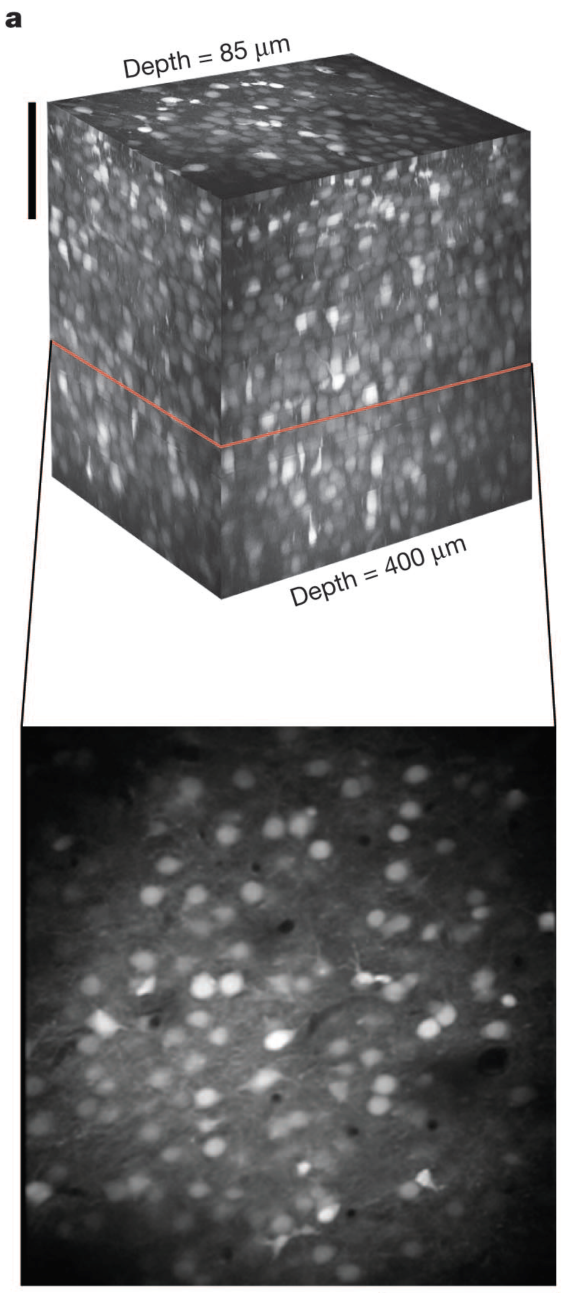

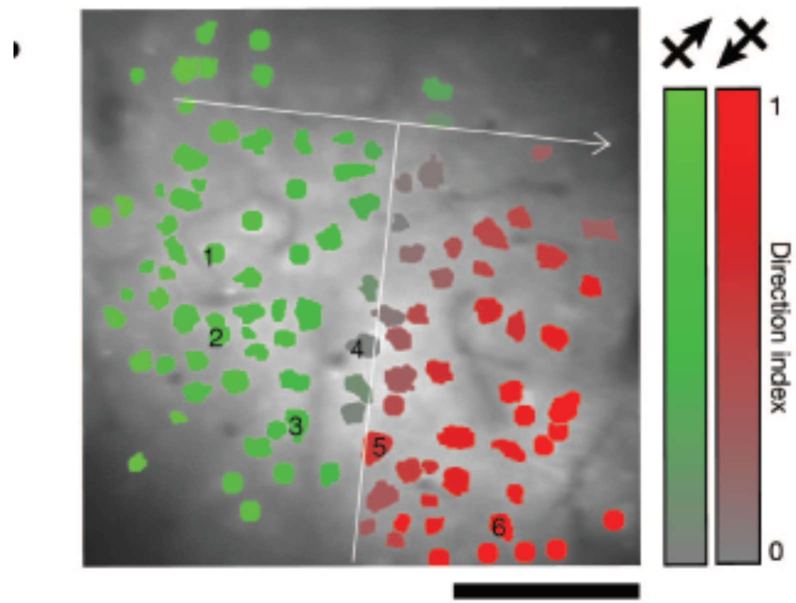

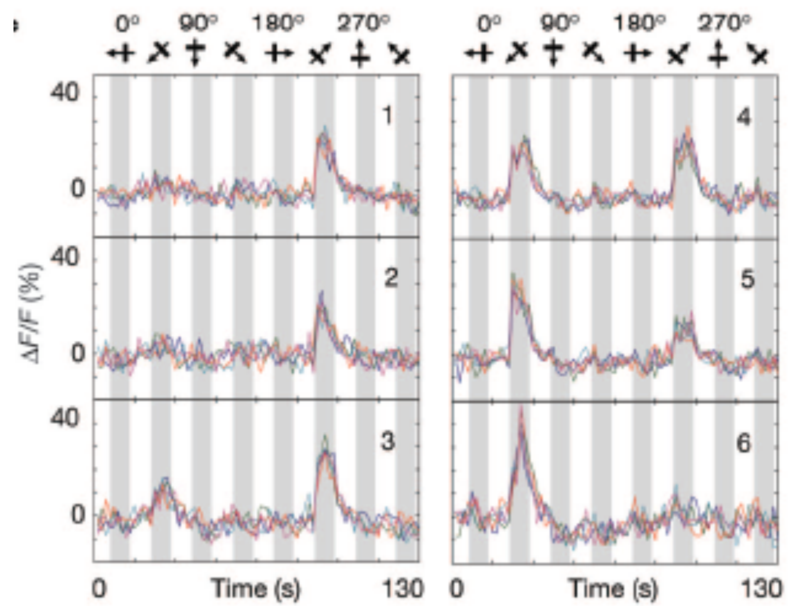

Self-organizing maps

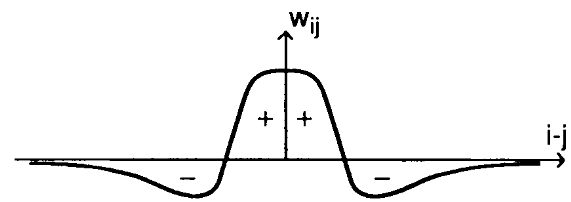

horizontal connections

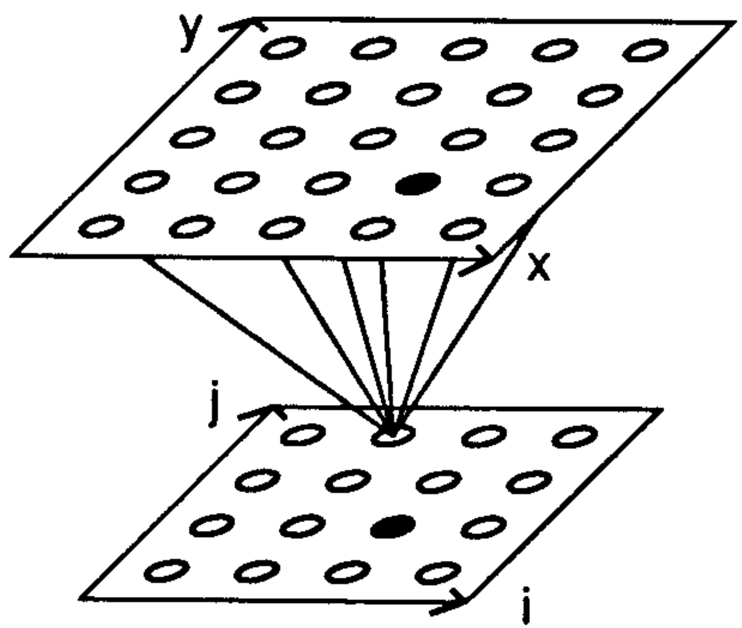

output

input

Kohonen's algorithm

$$\Delta \mathbf{w}_i = \eta \, \Lambda(i, i^*) \, (\mathbf{x} - \mathbf{w}_i)$$

$$\Lambda(i, i^*) = e^{-\frac{|\mathbf{r}_i - \mathbf{r}_{i^*}|^2}{2\sigma^2}}$$

Kohonen’s algorithm applied to a 2D input array

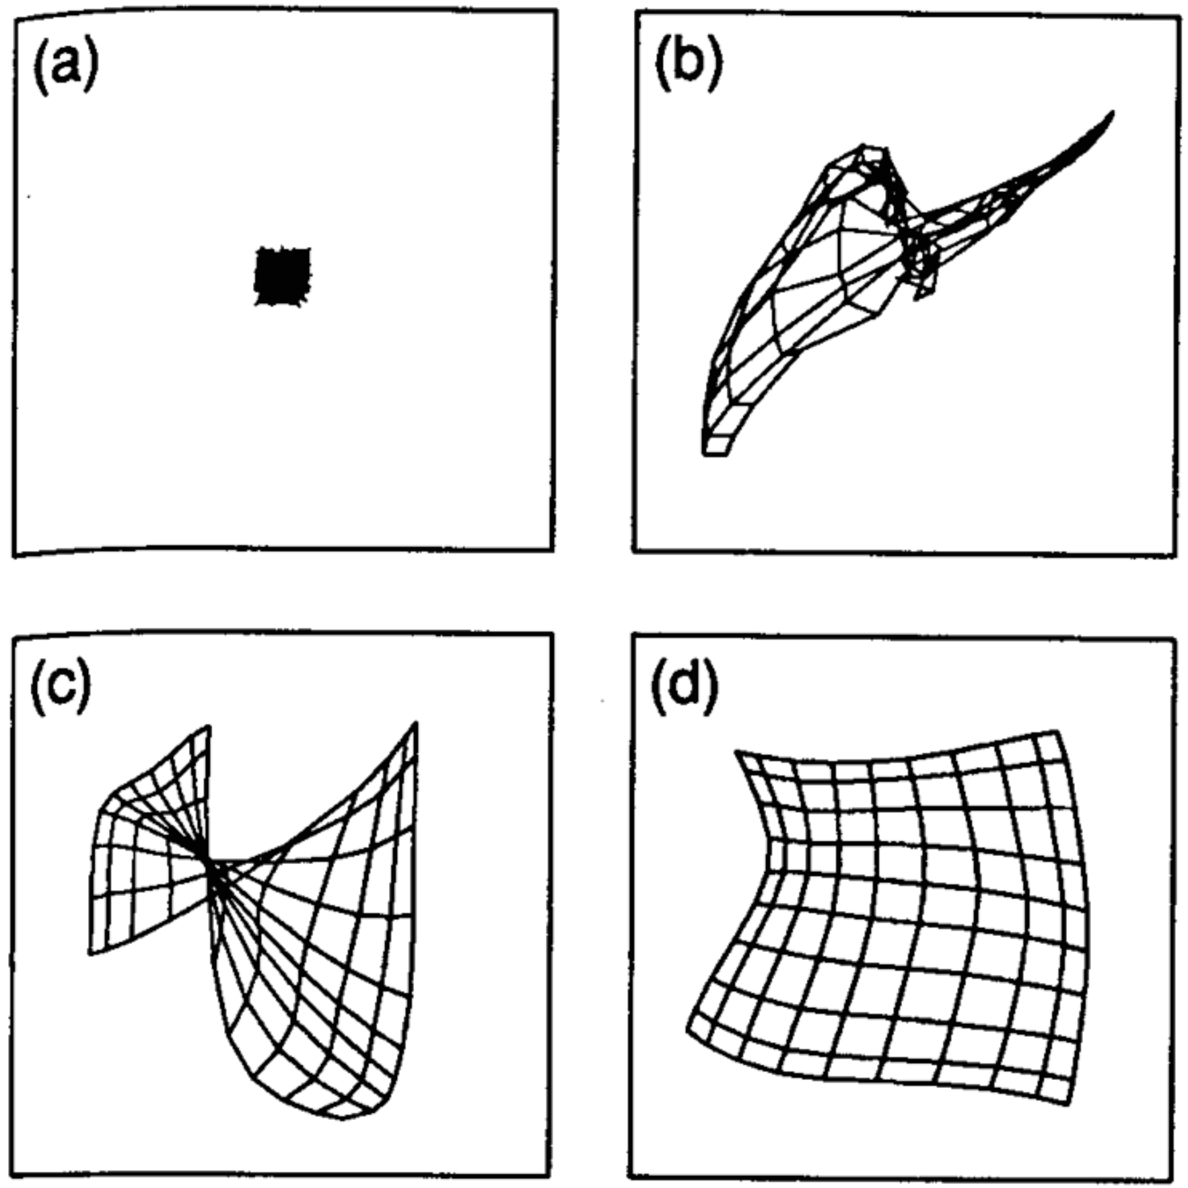

What is this organization good for?

Horizontal connections may enforce continuity among oriented elements

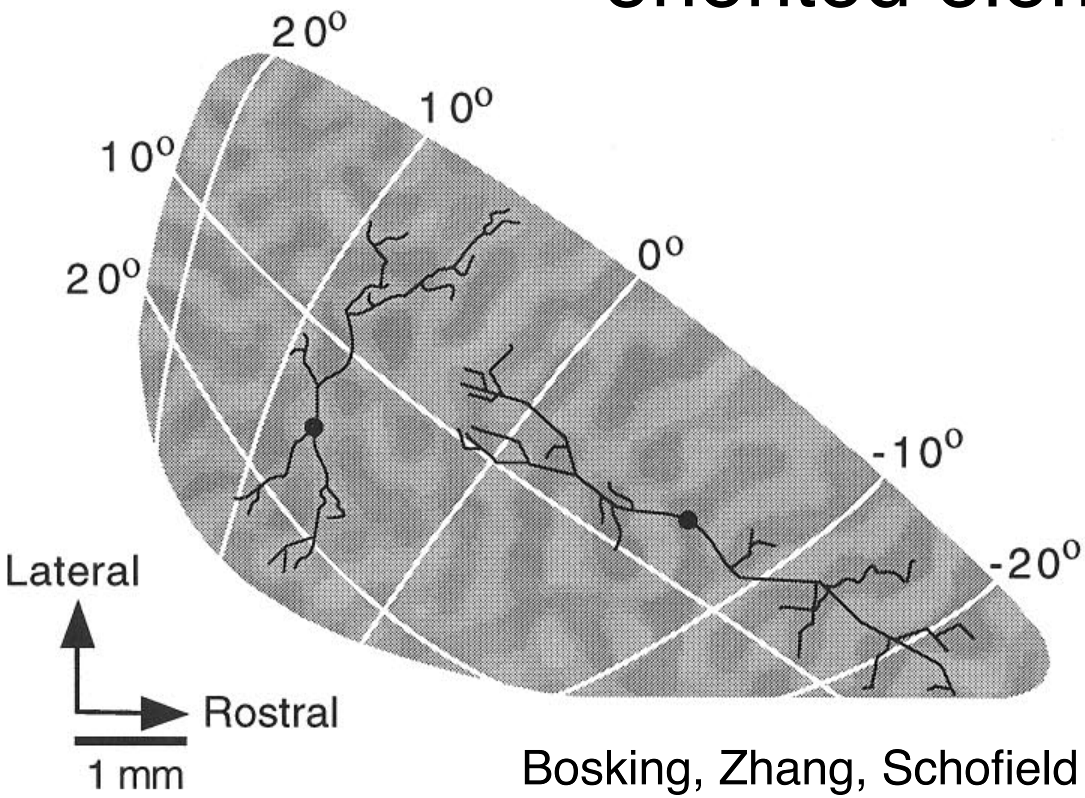

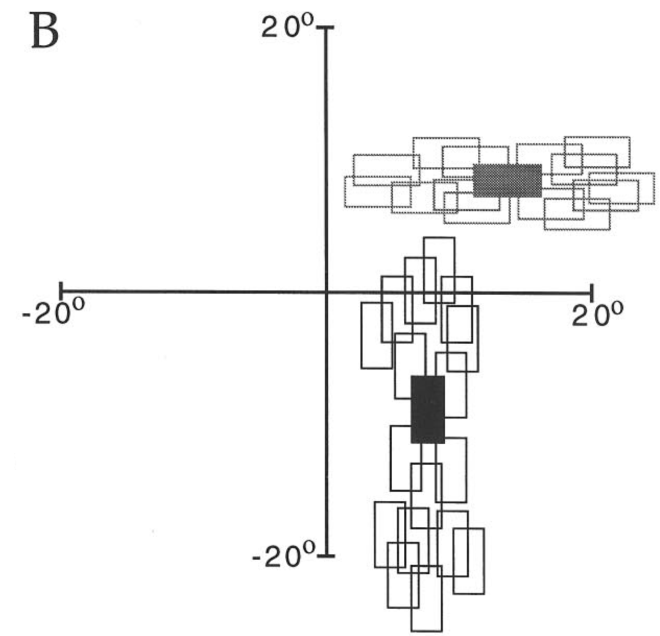

Bosking, Zhang, Schofield & Fitzpatrick (1993)

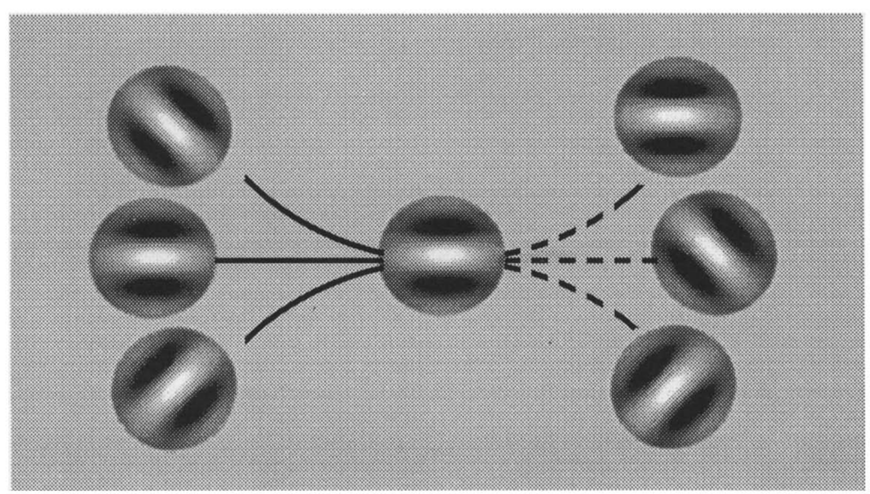

Field, Hayes & Hess (1993)

Statistical dependencies

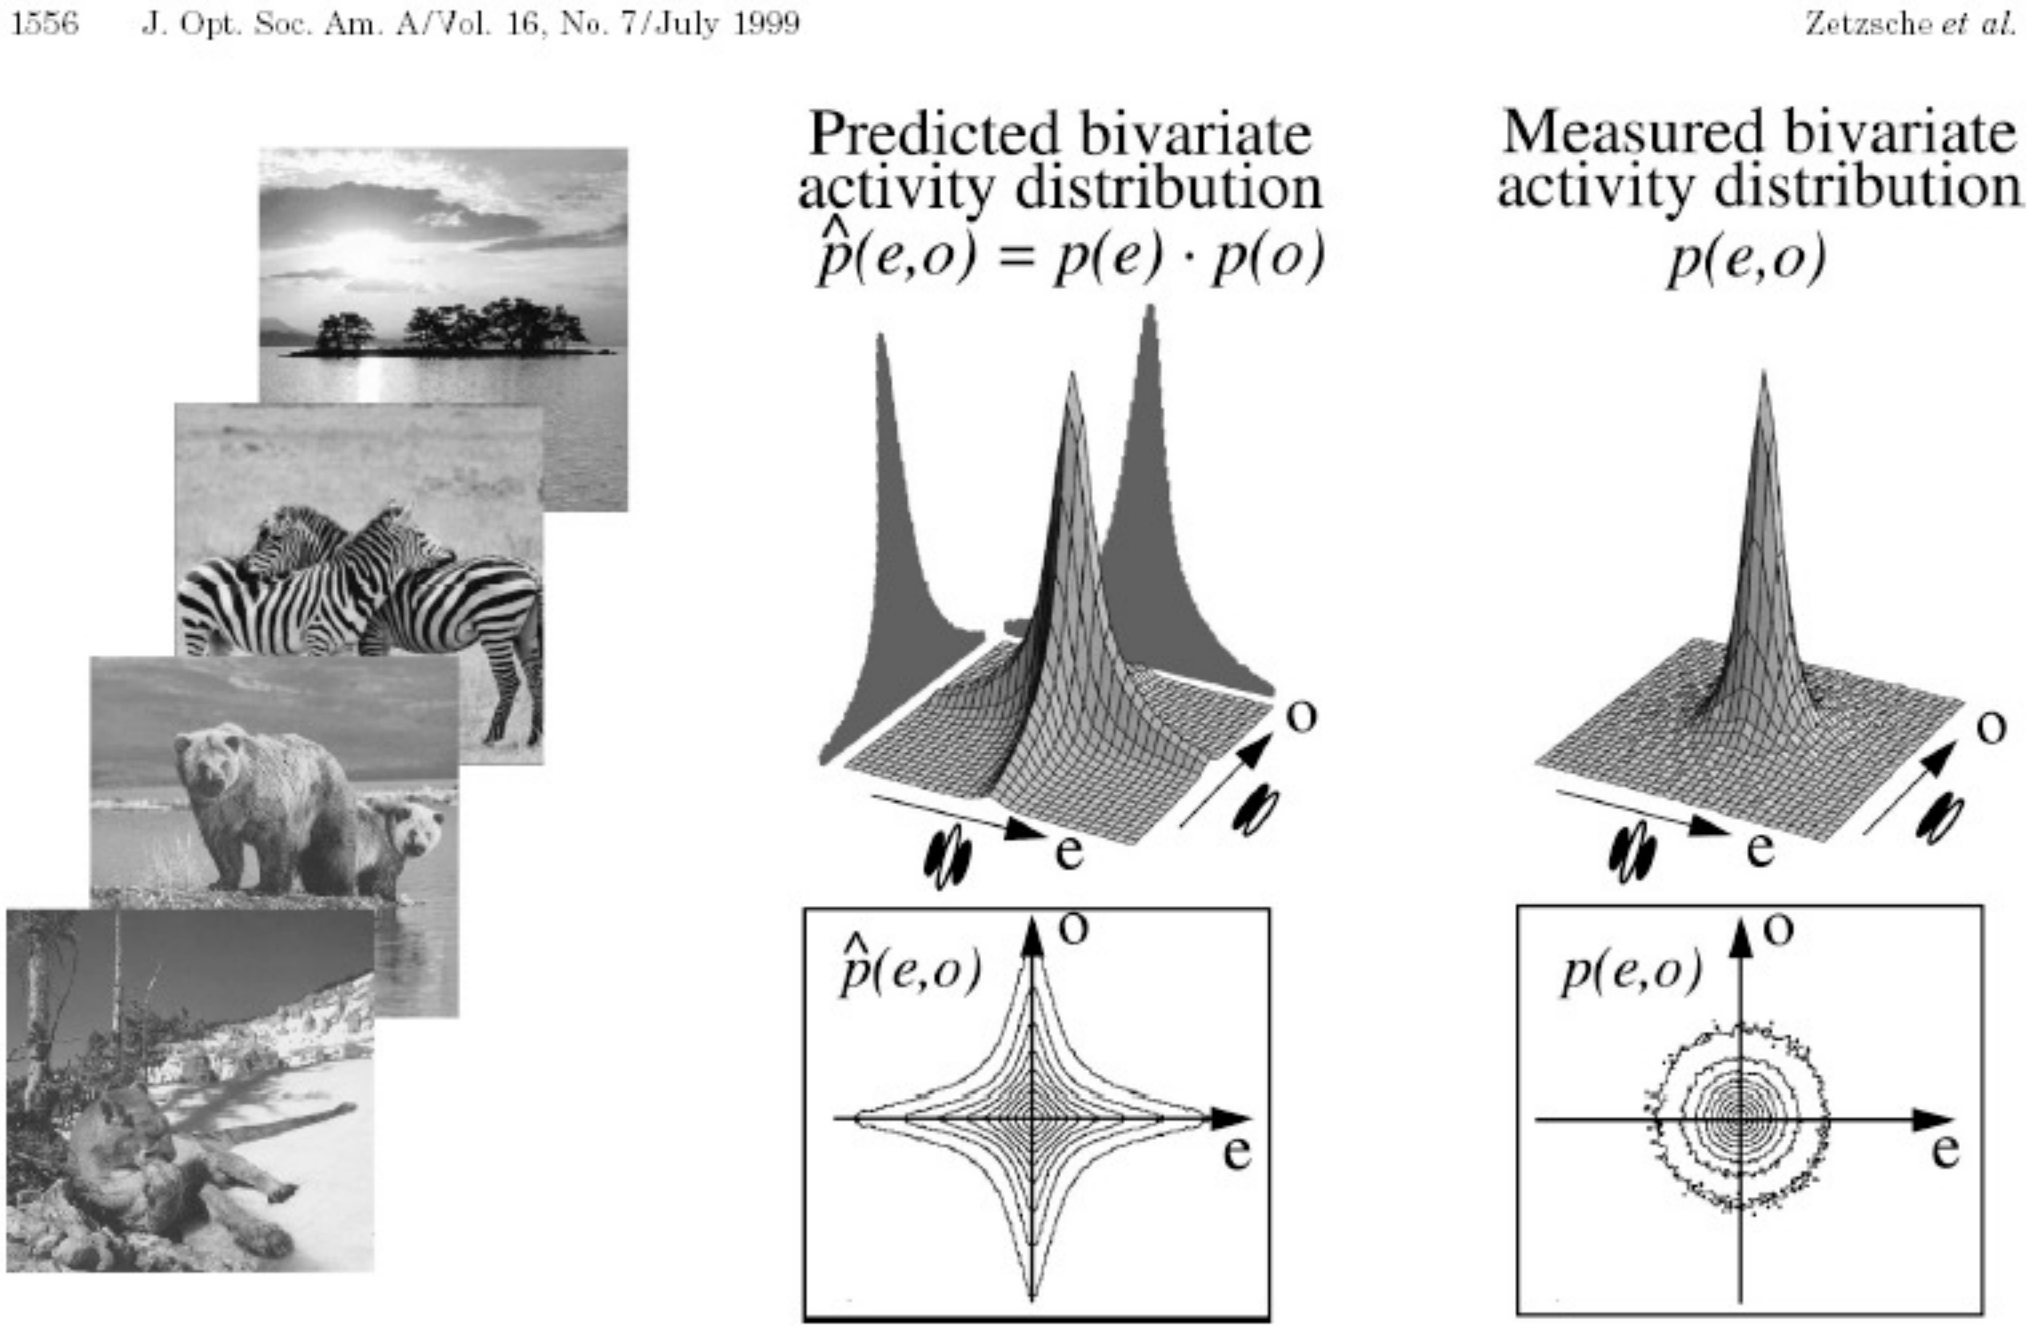

Joint coefficient activations in response to a moving edge

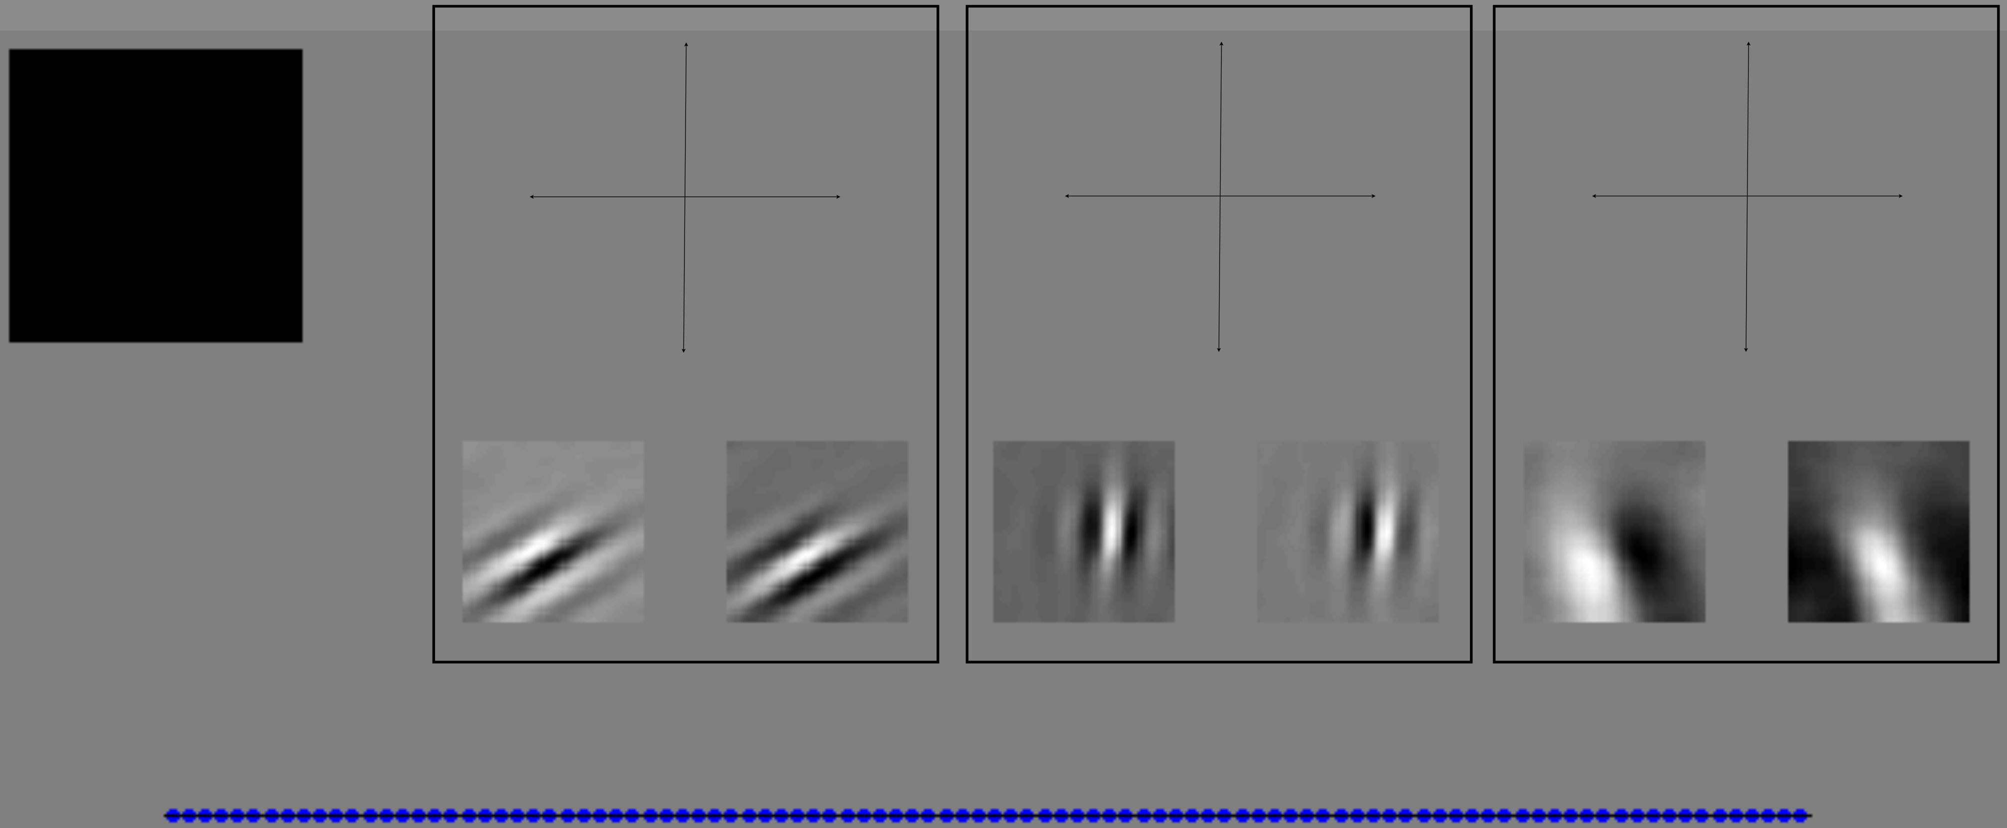

'Topographic ICA'
(Hyvarinen & Hoyer)

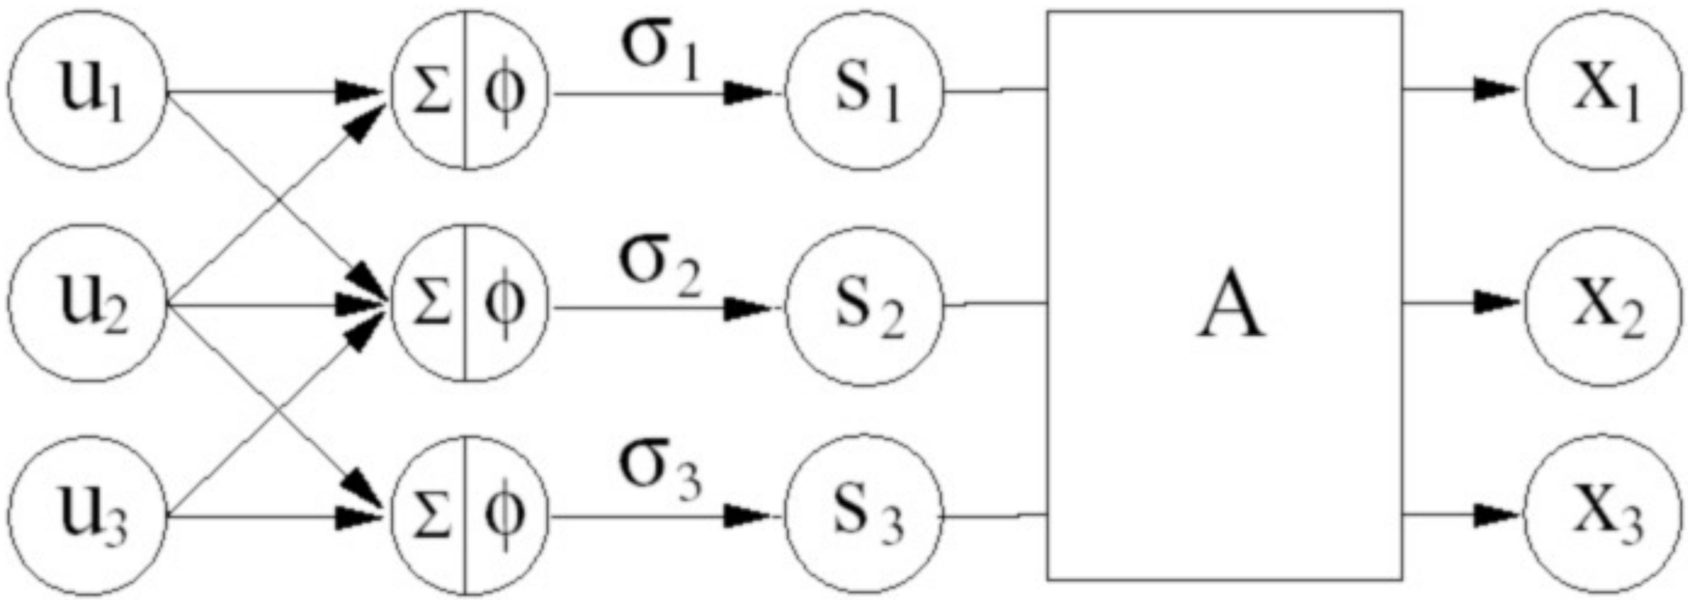

‘Topographic ICA’
(Hyvarinen & Hoyer)

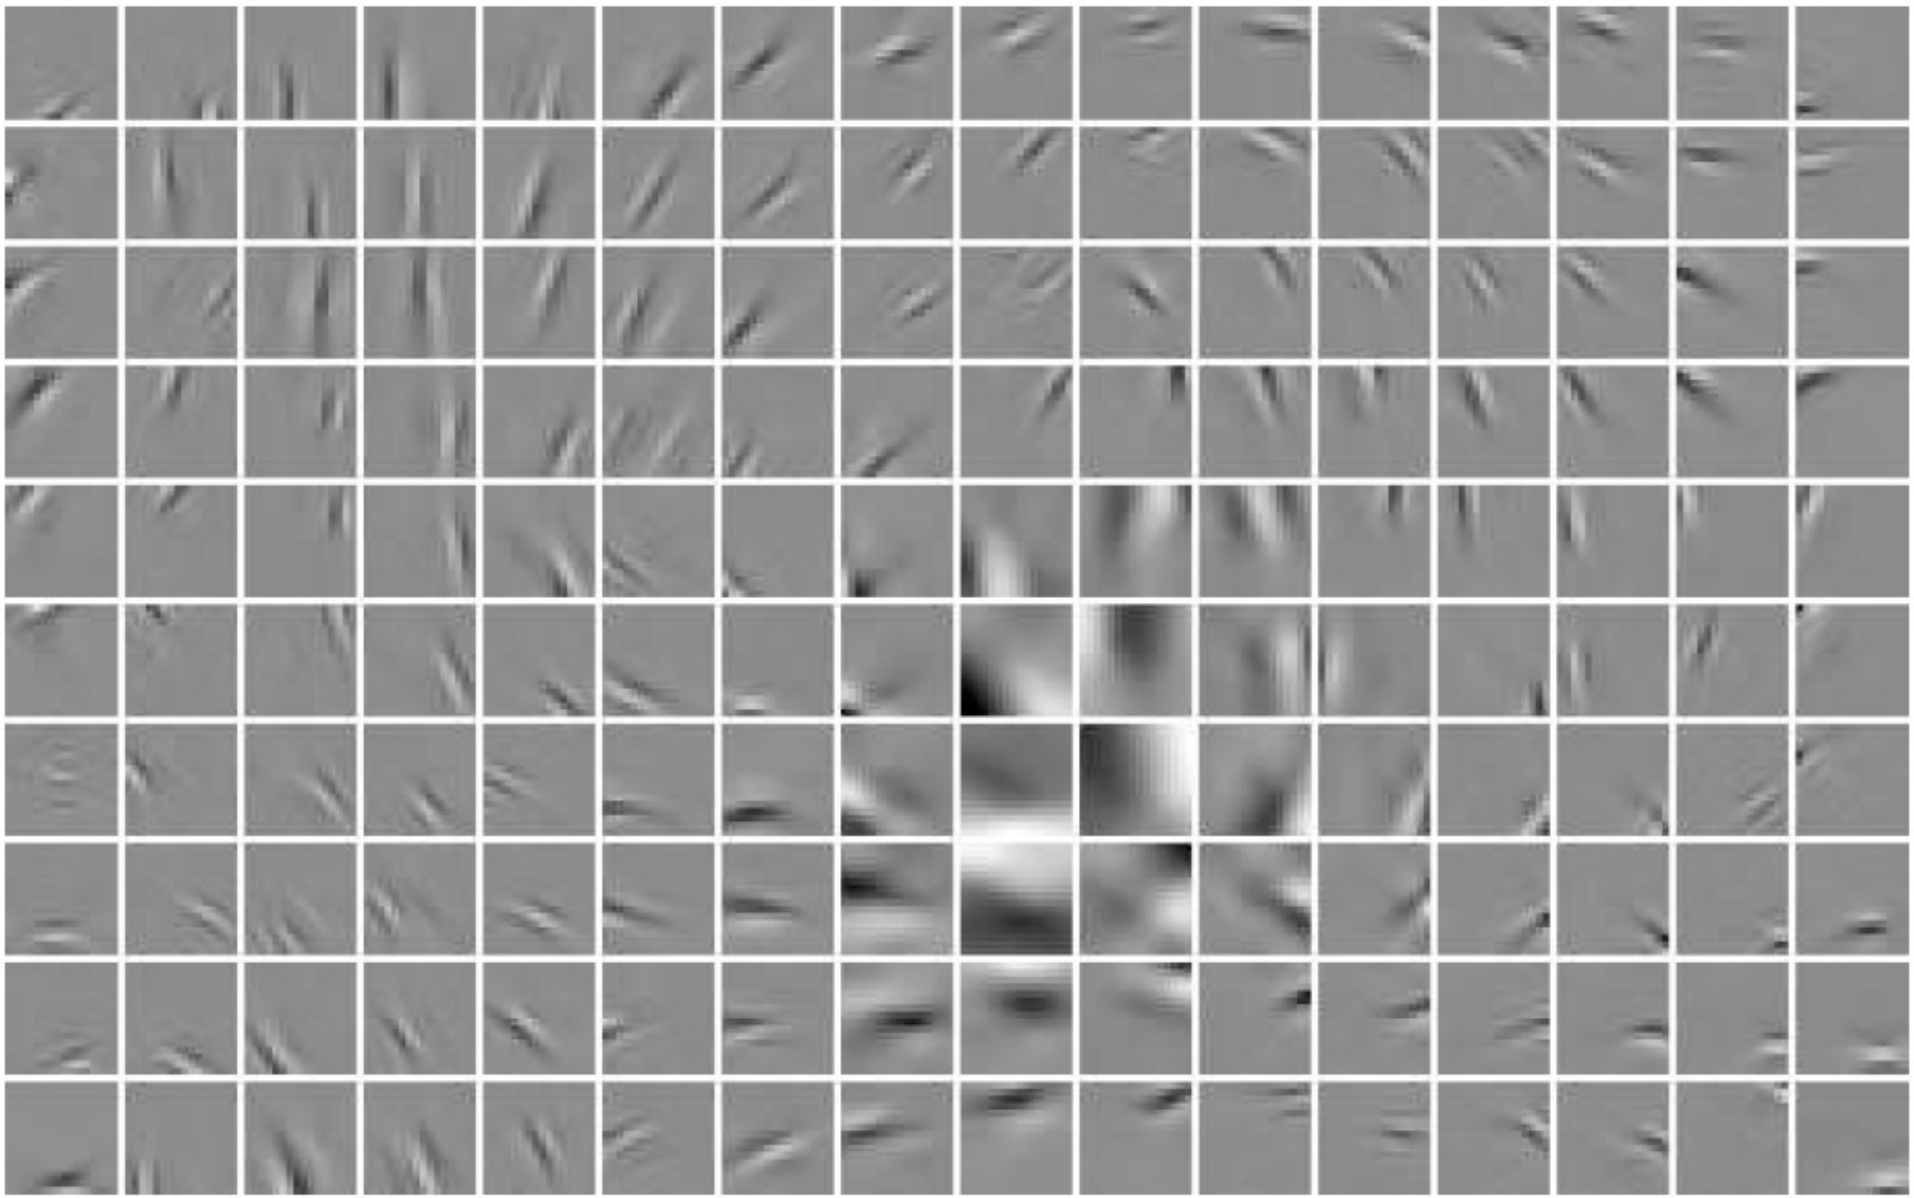

'Google Brain'

(Quoc Le et al. 2012)

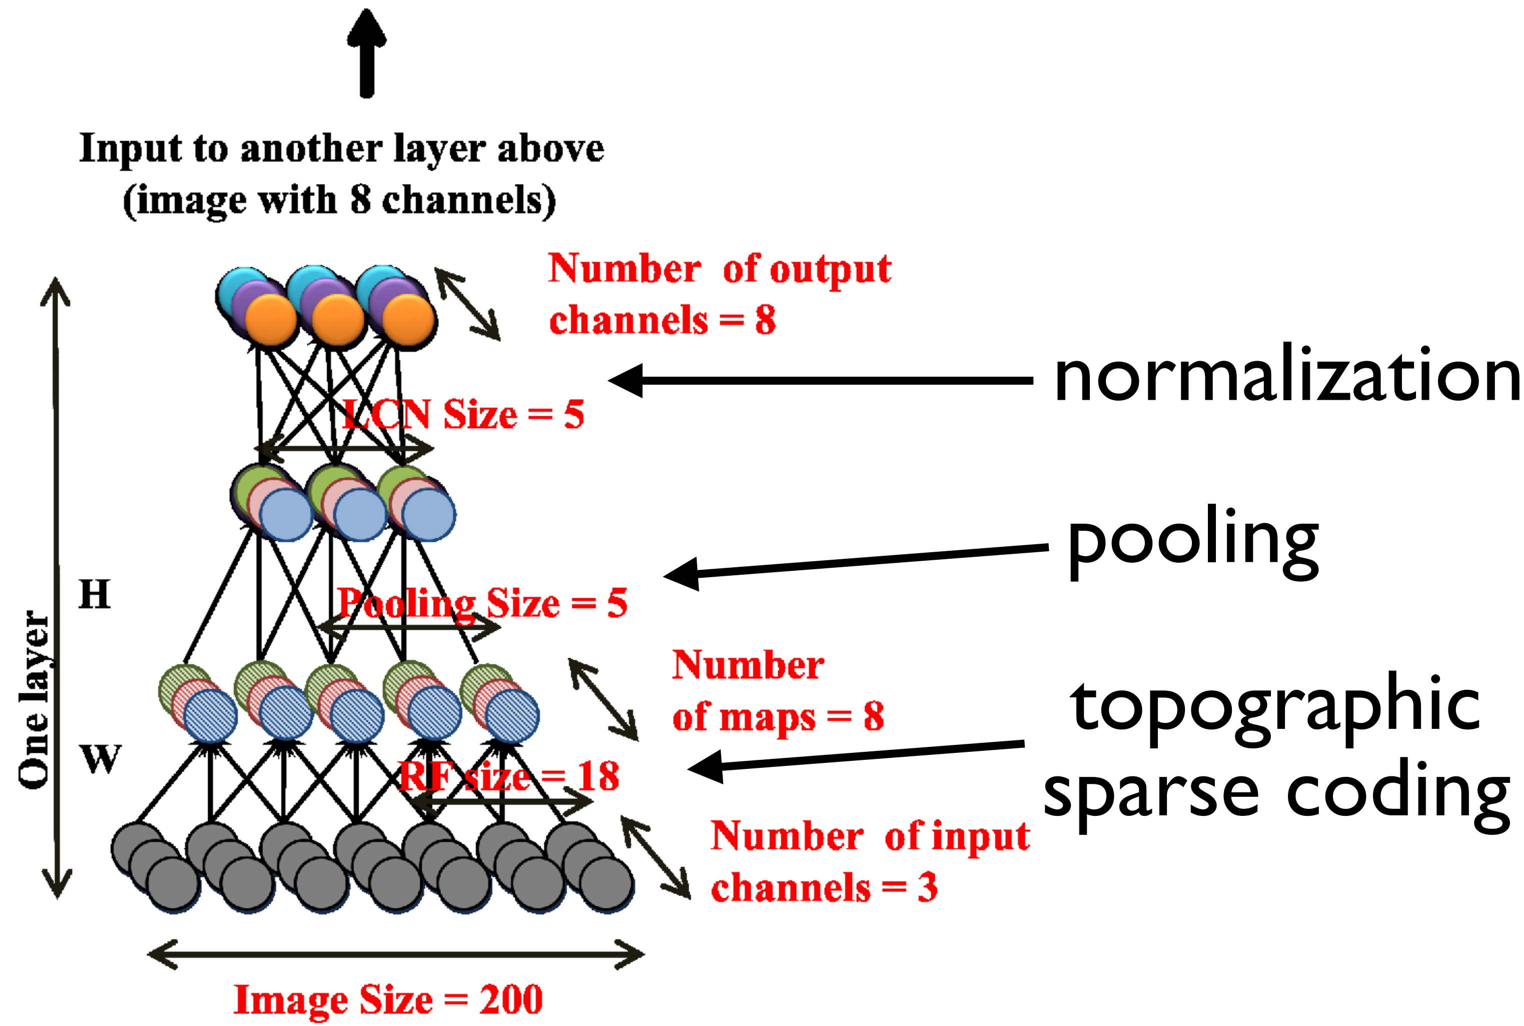

‘Google Brain’

(Quoc Le et al. 2012)

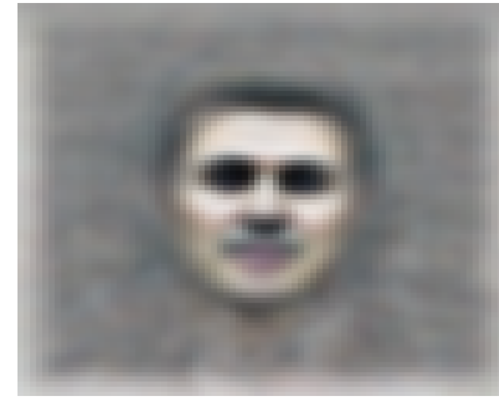

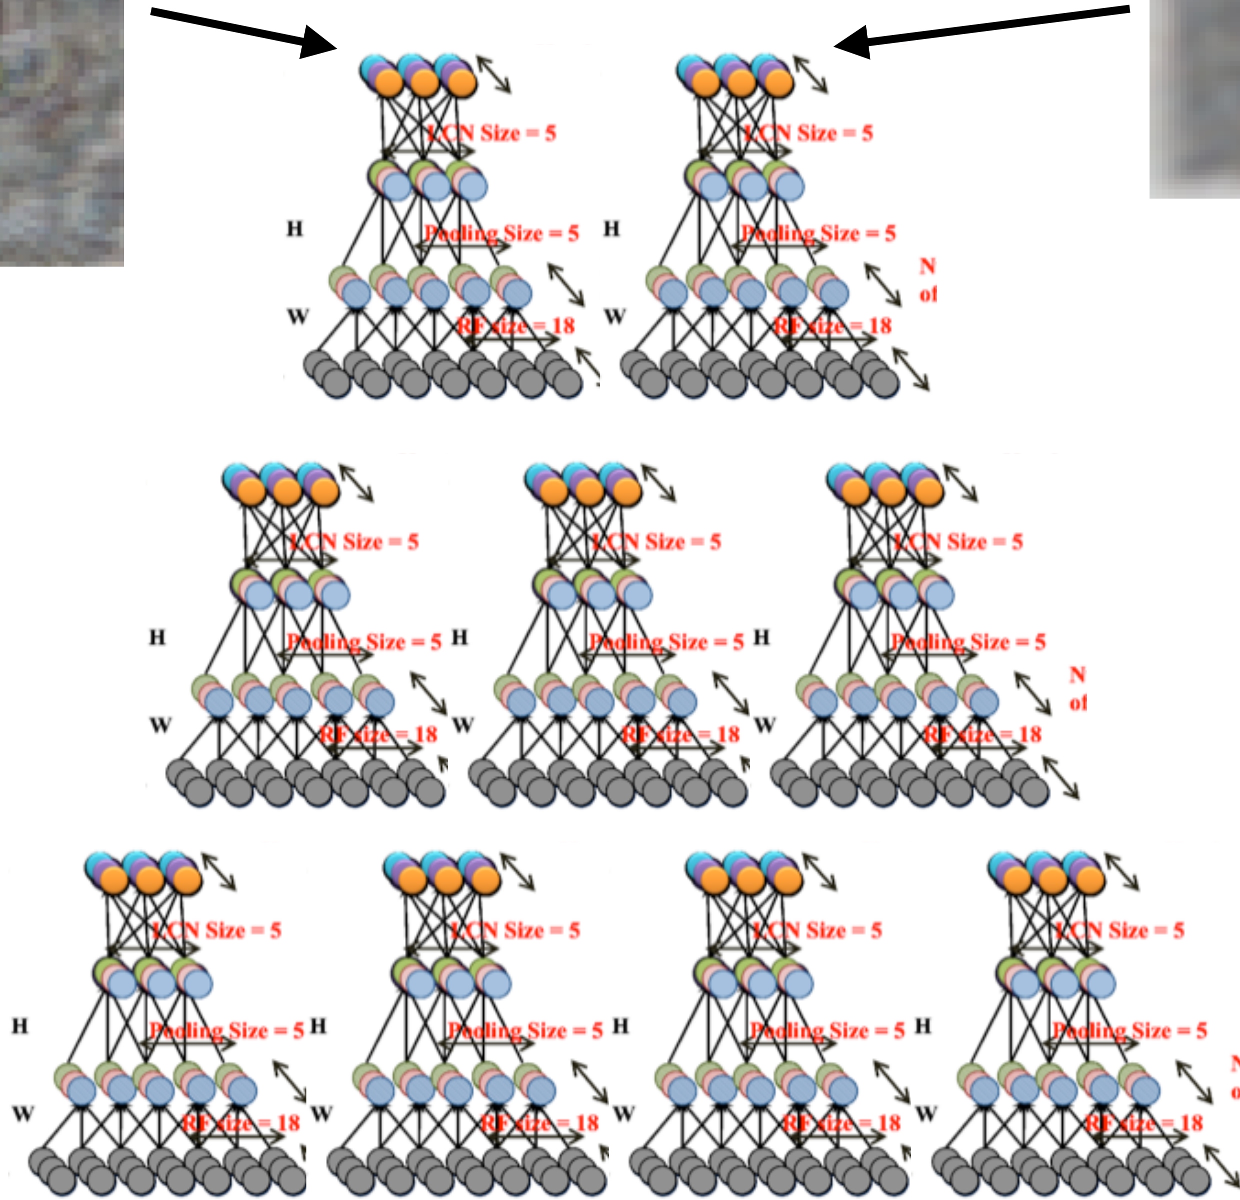

Manifolds

Many types of natural data lie along low dimensional manifolds embedded in a higher dimensional space.

Sparse representations tile the manifold of natural images in such a way that data points along the manifold are spanned by a small number of basis vectors.

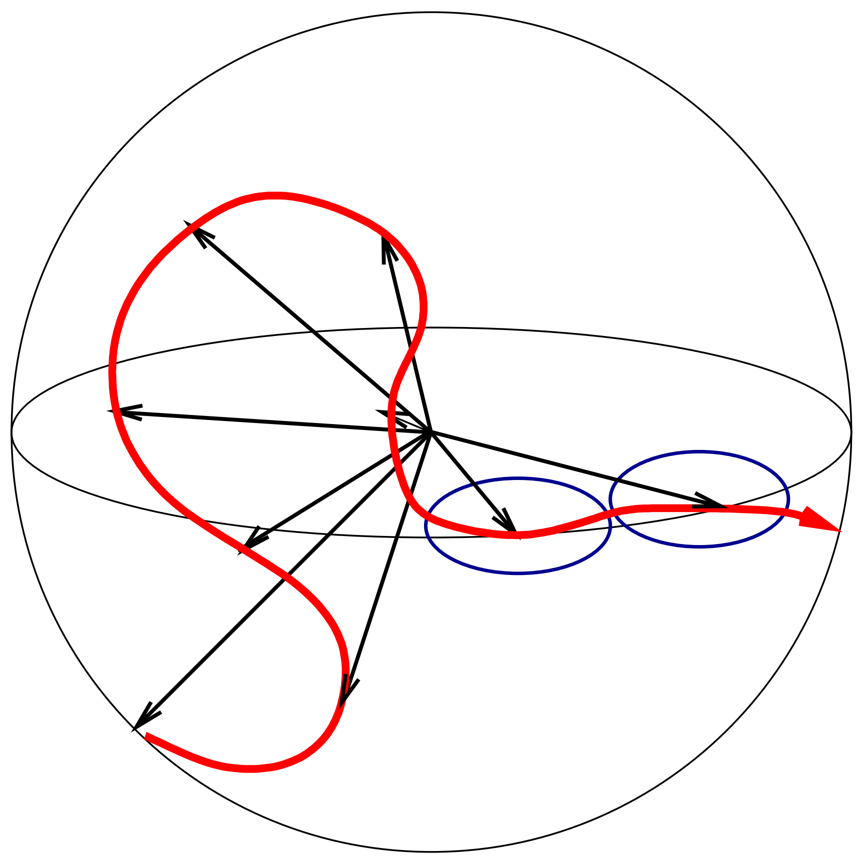

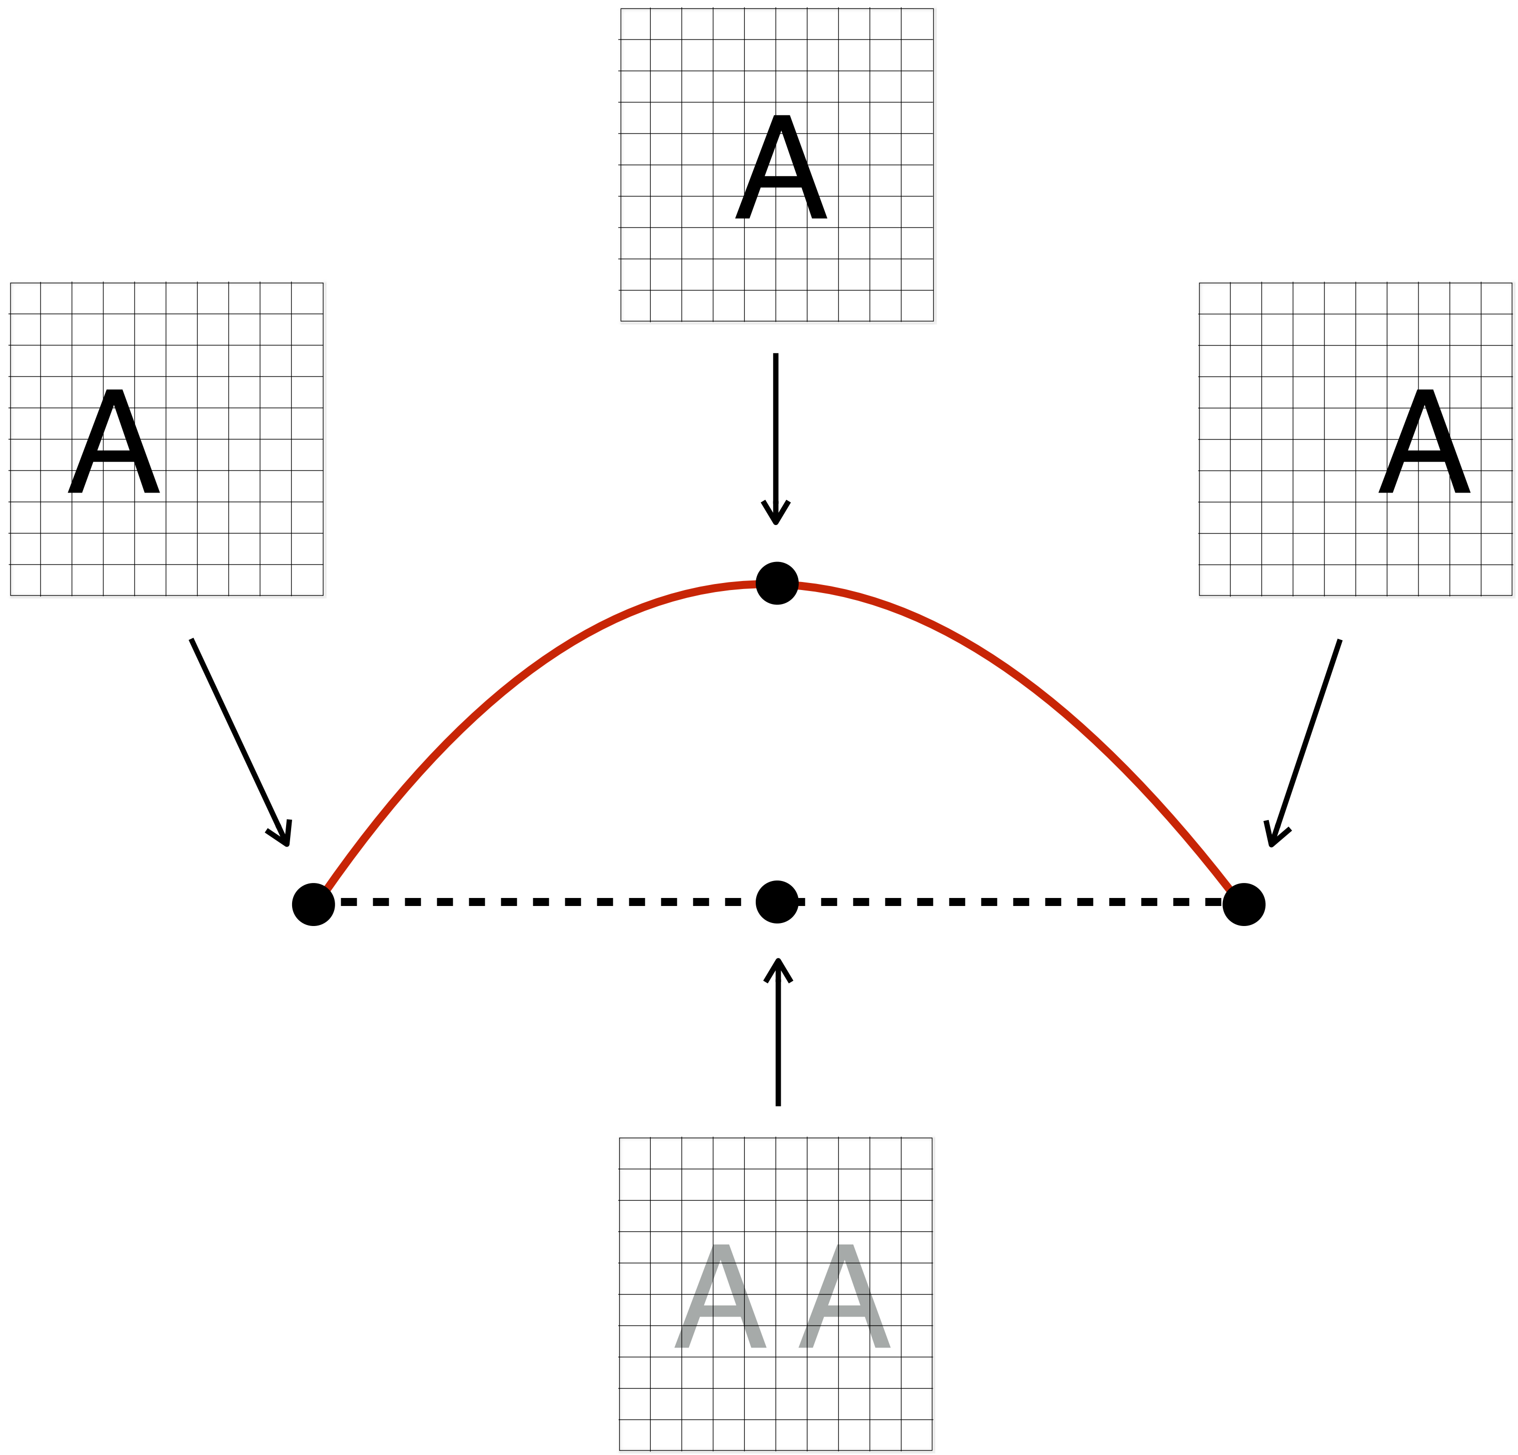
A
A
A
A A

Distribution of 3x3 pixel image patches drawn from natural scenes forms a Klein bottle (Carlsson et al., 2008)

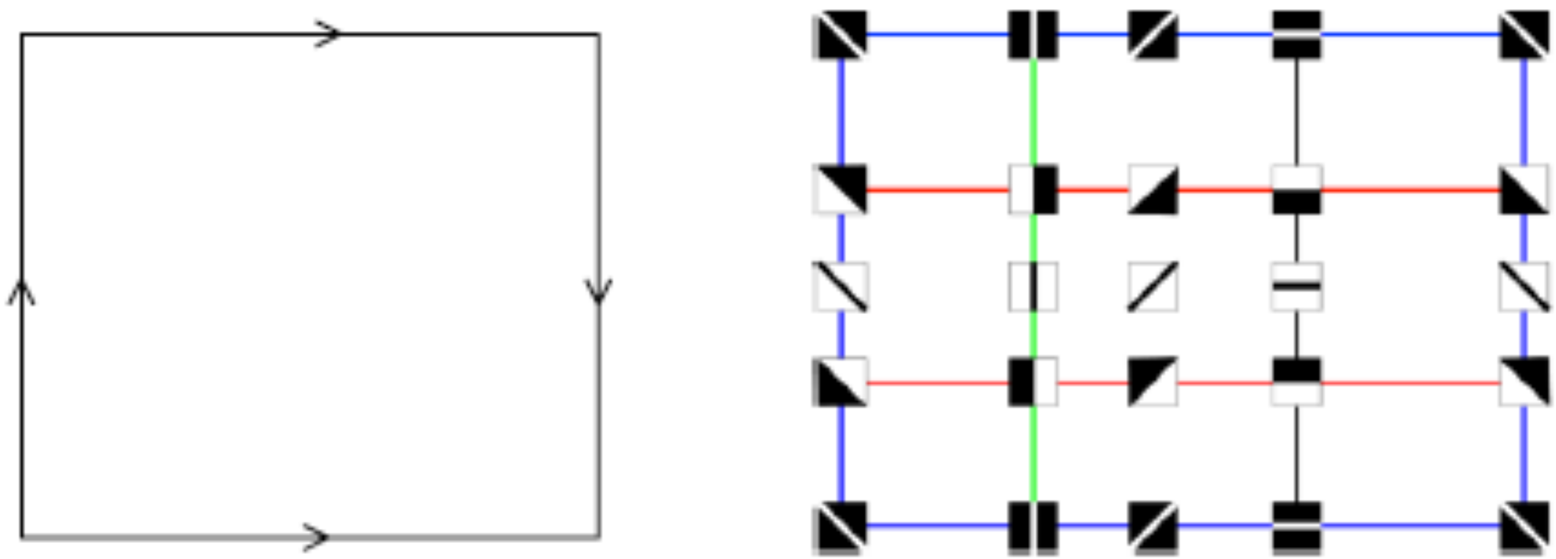

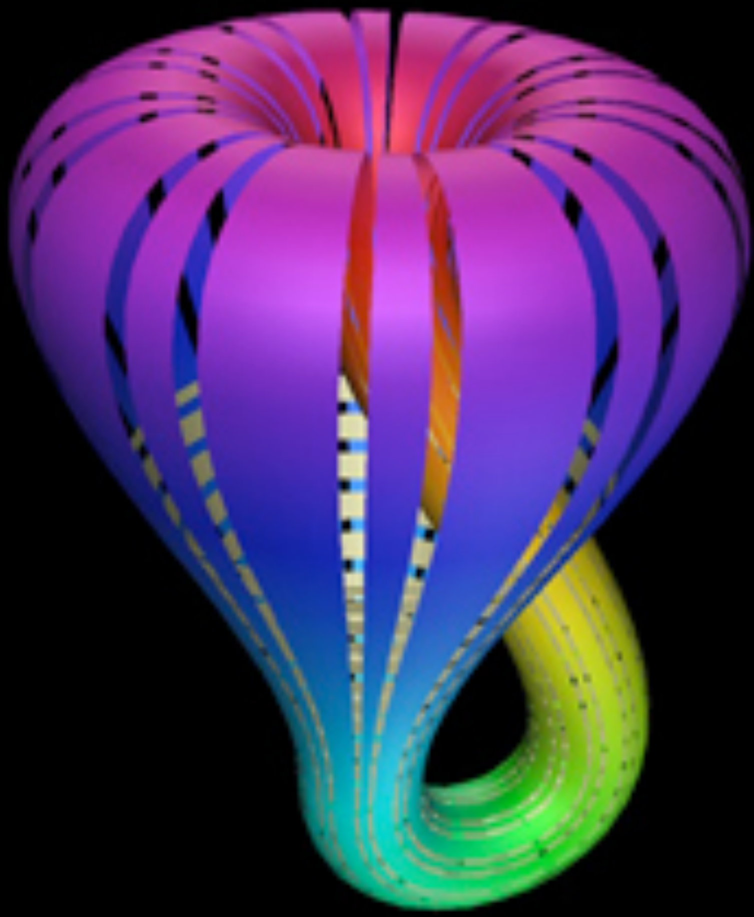

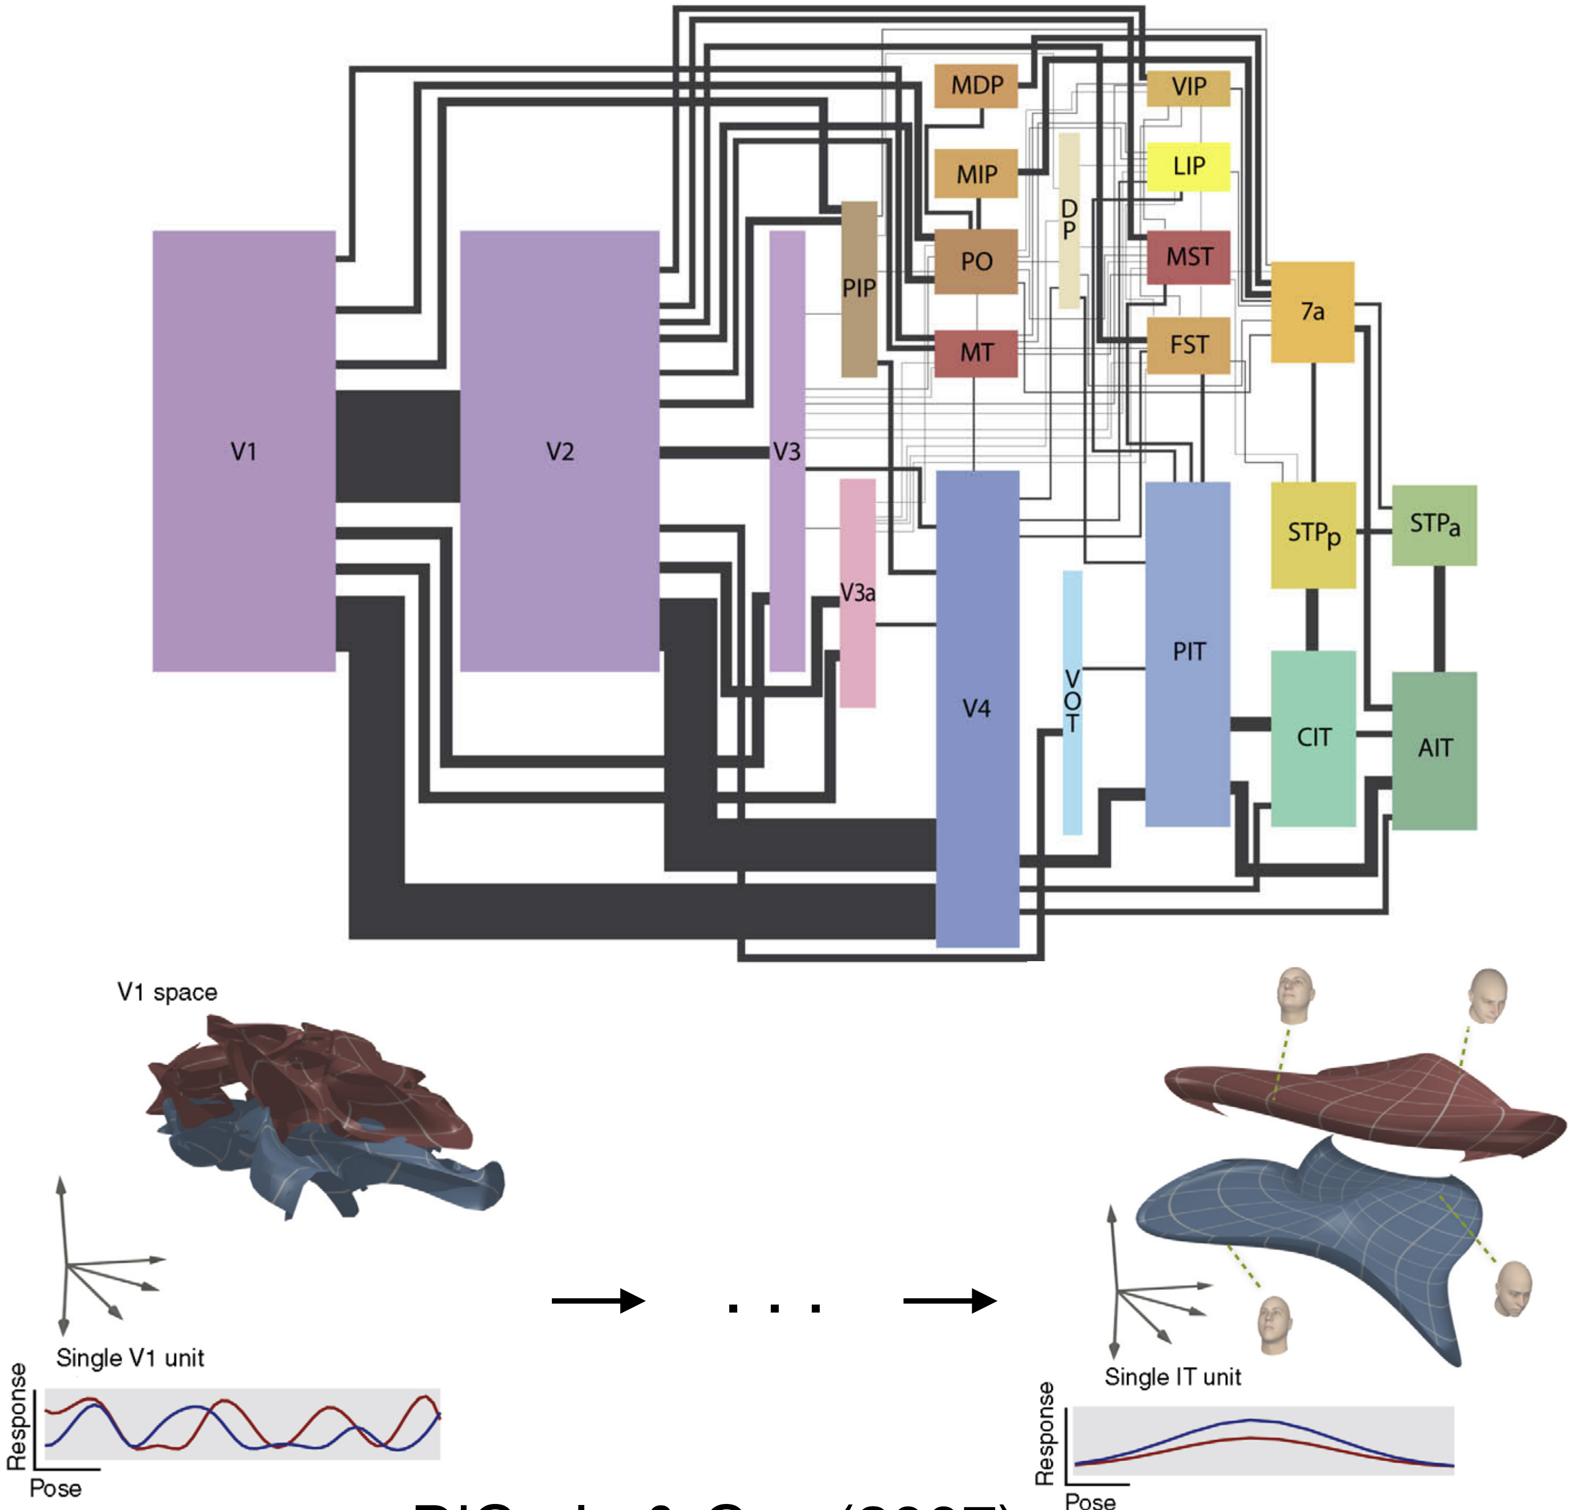

DiCarlo & Cox (2007)

Nonlinear Dimensionality Reduction by Locally Linear Embedding

Sam T. Roweis[1] and Lawrence K. Saul[2]

A Global Geometric Framework for Nonlinear Dimensionality Reduction

Joshua B. Tenenbaum,[1*] Vin de Silva,[2] John C. Langford[3]

Science, 22 Dec. 2000

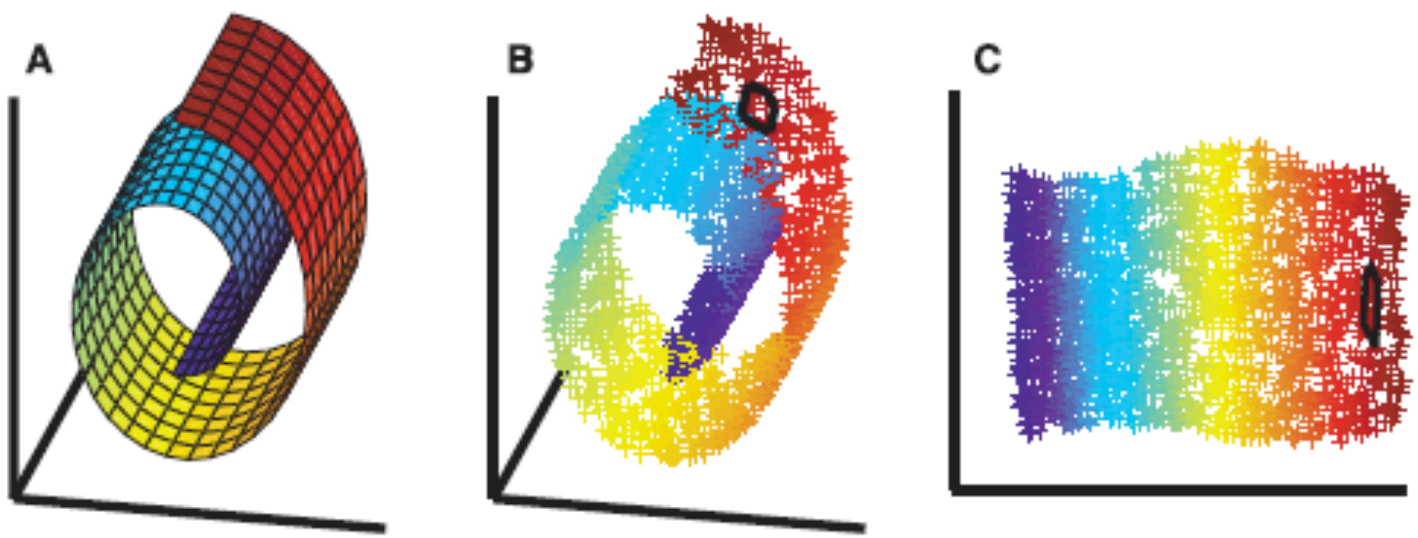

Local Linear Embedding (LLE)

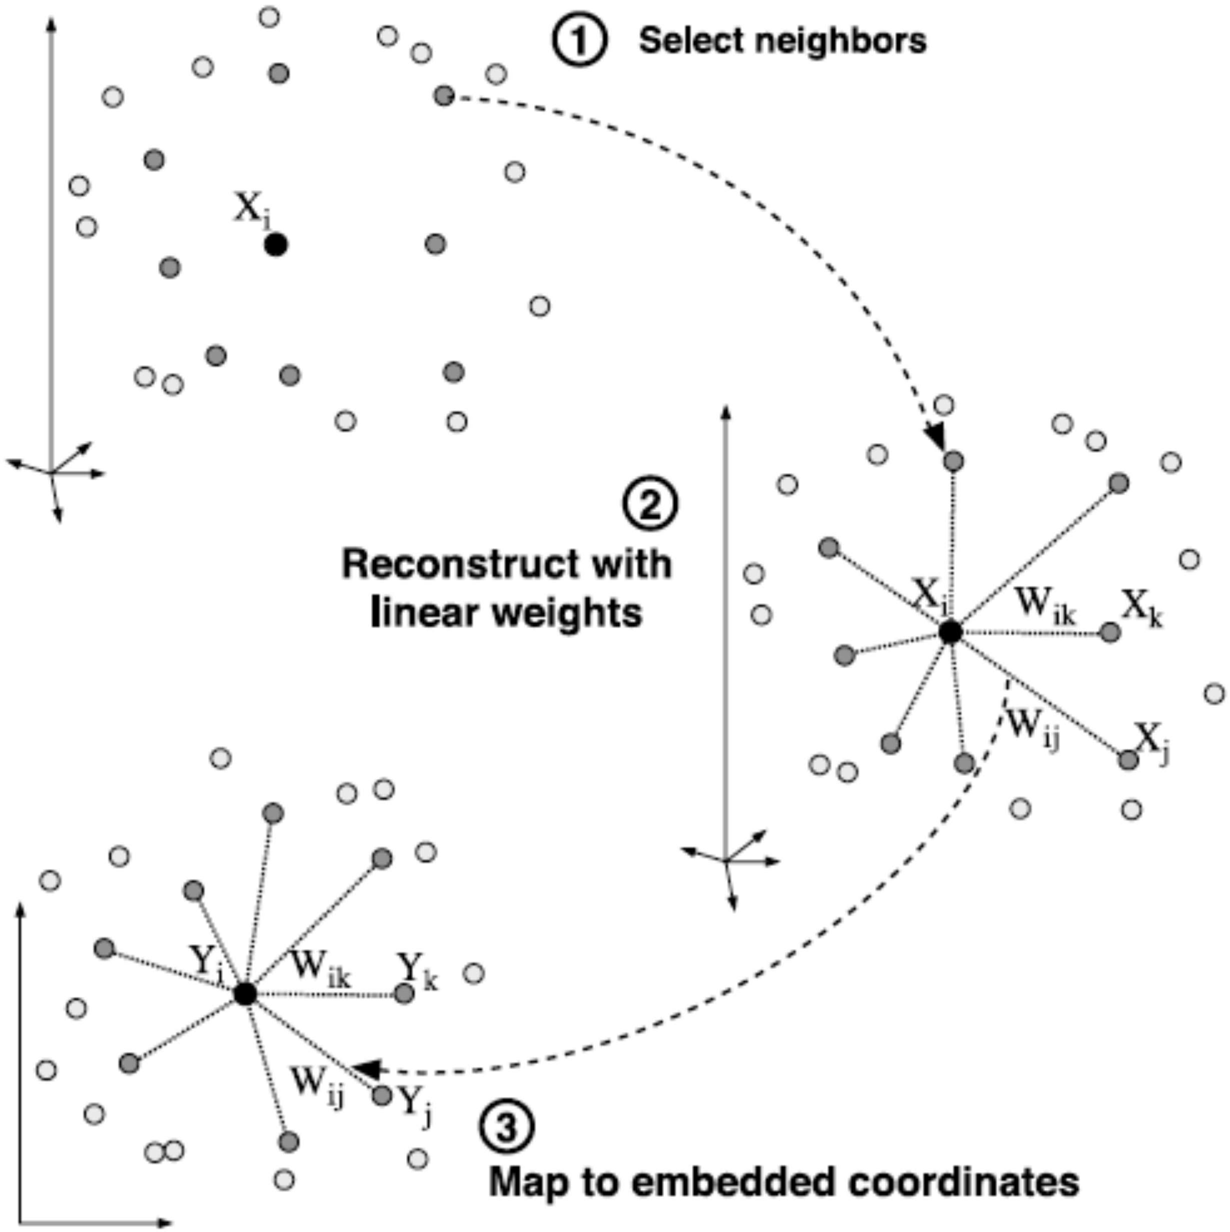

$$\varepsilon(W) = \sum_i \left| \vec{X}_i - \Sigma_j W_{ij} \vec{X}_j \right|^2$$

$$\Phi(Y) = \sum_i \left| \vec{Y}_i - \Sigma_j W_{ij} \vec{Y}_j \right|^2$$

Manifold of facial pose and lighting

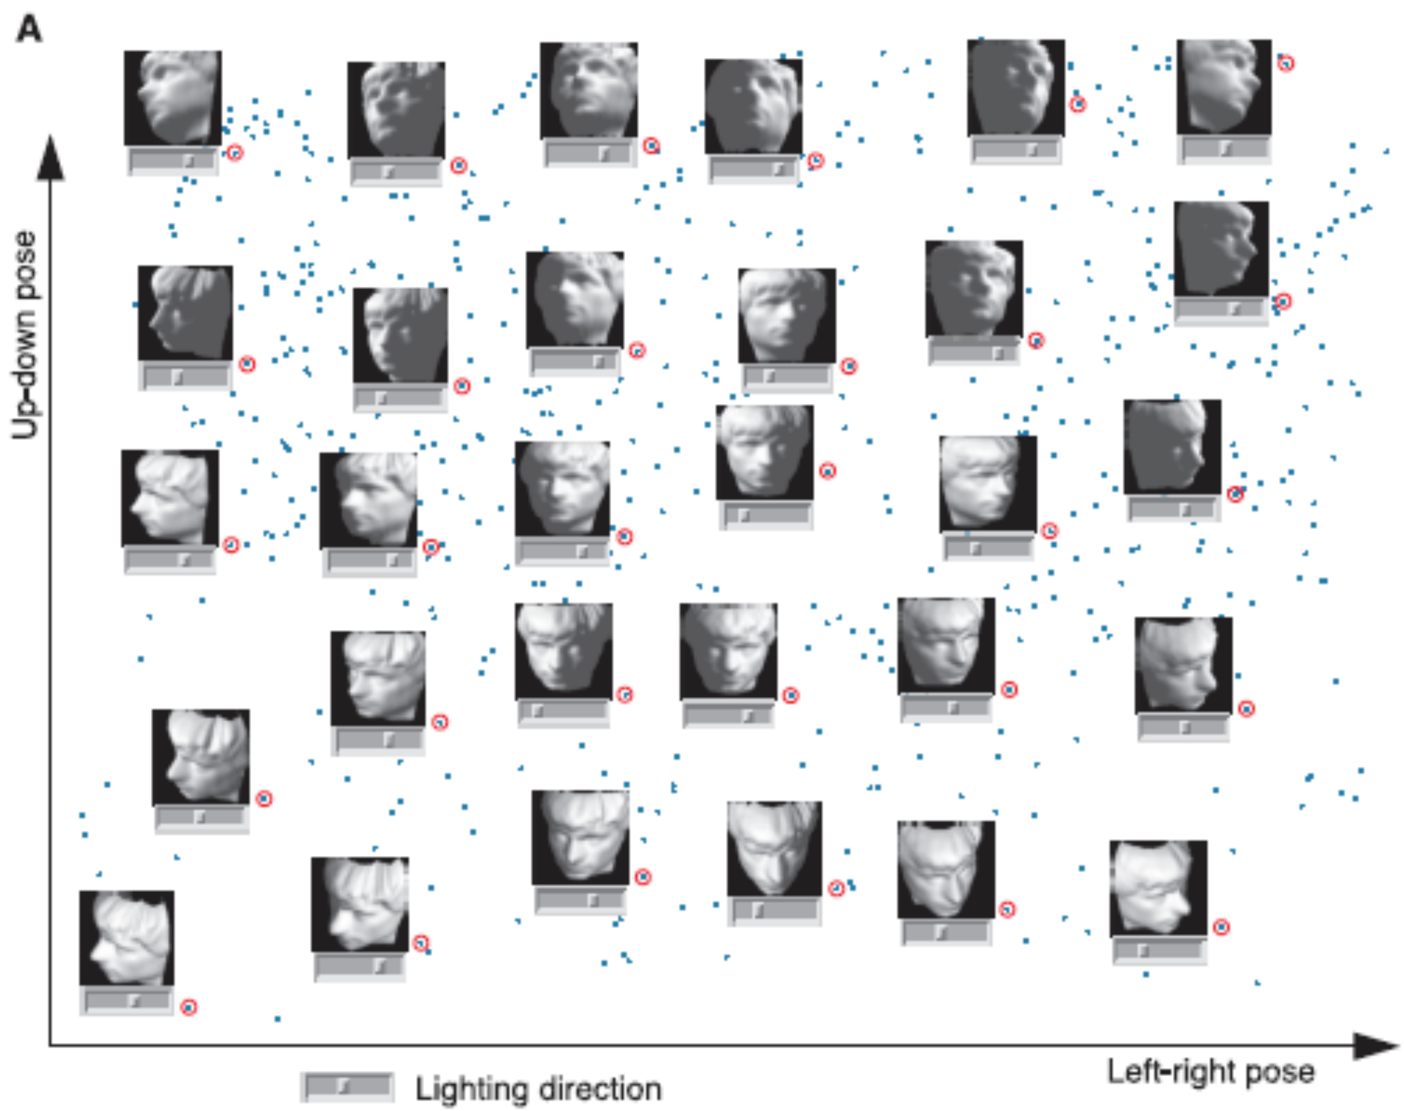

Hand-written digits

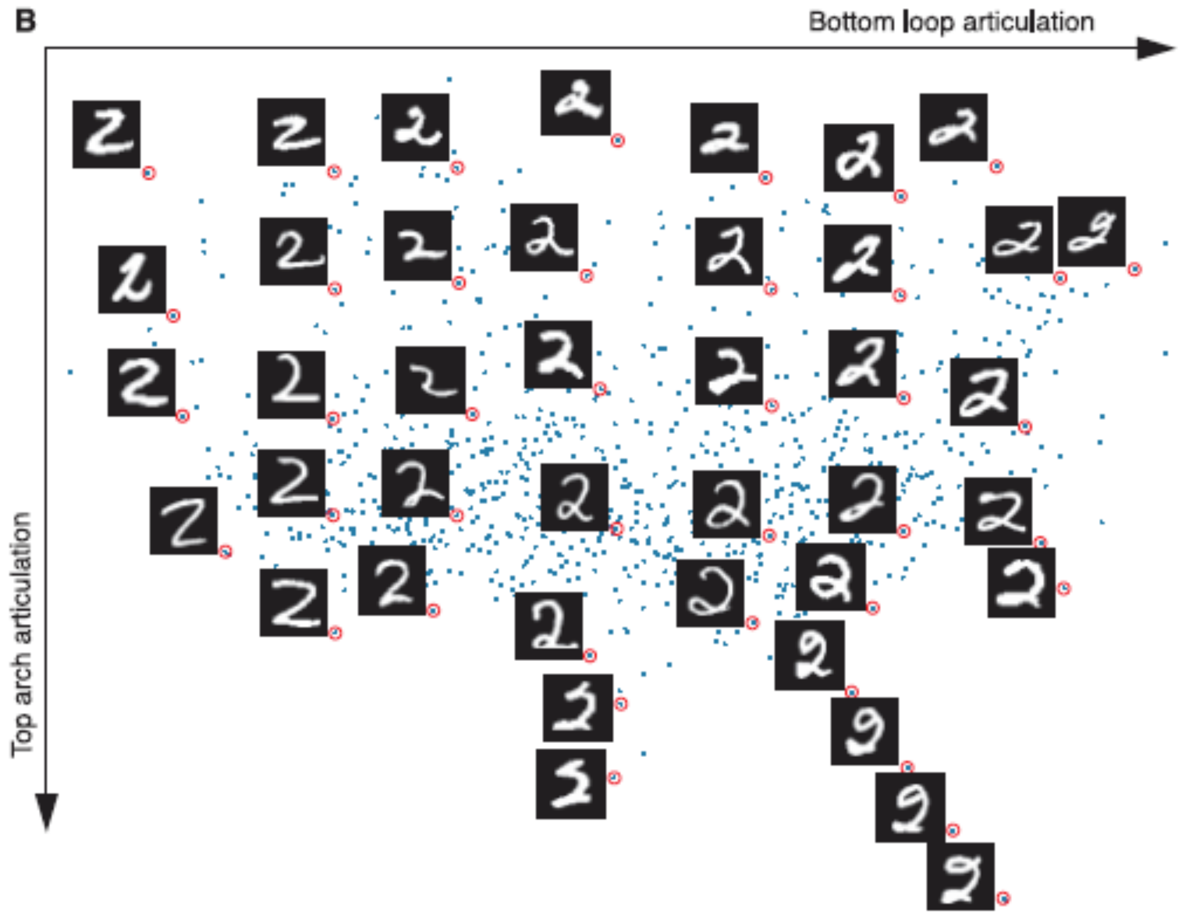

Local Linear Landmarks (LLL)
(Vladymyrov & Carreira-Perpinán, 2013)

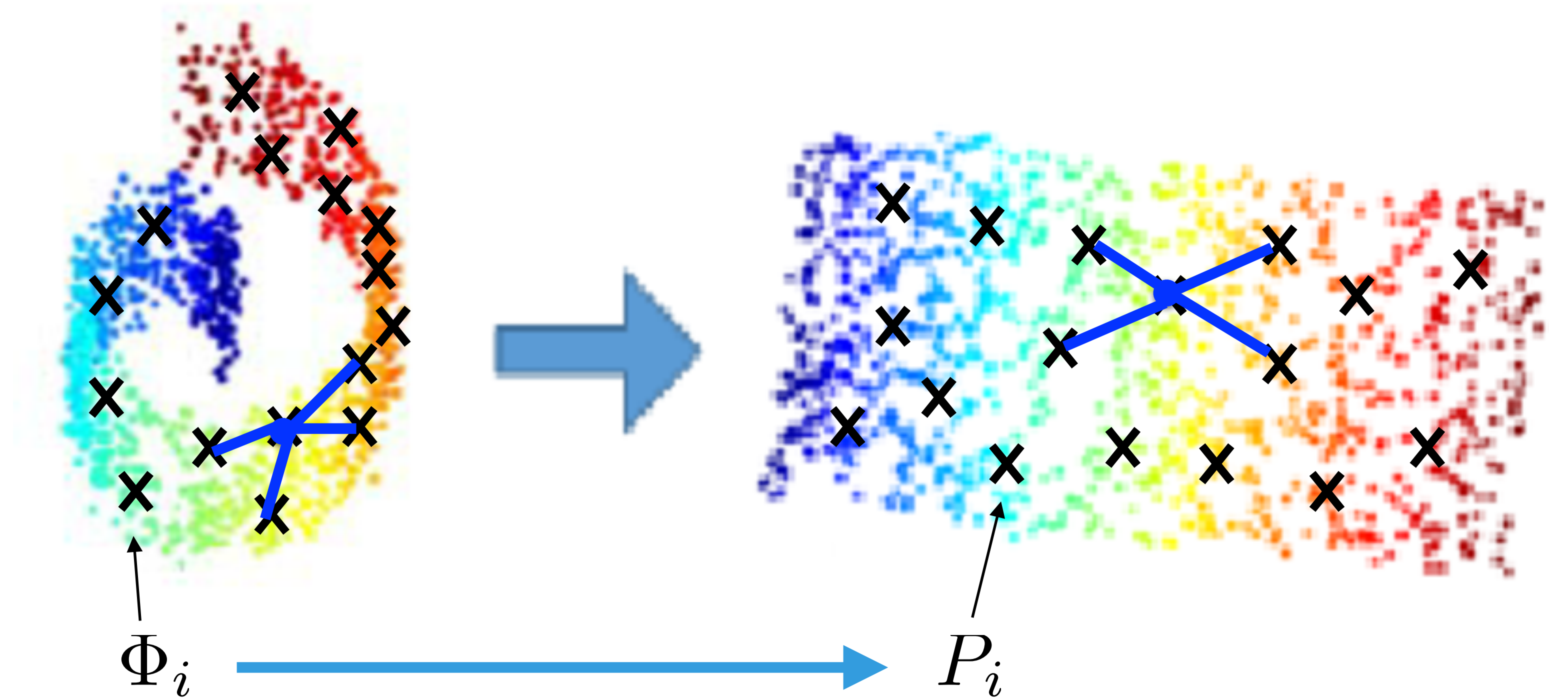

$$x = \Phi\,\alpha + n \longrightarrow y = P\,\alpha$$

local linear approximation → embedding

Basis functions learned by sparse coding form a locally linear approximation to the manifold of natural images

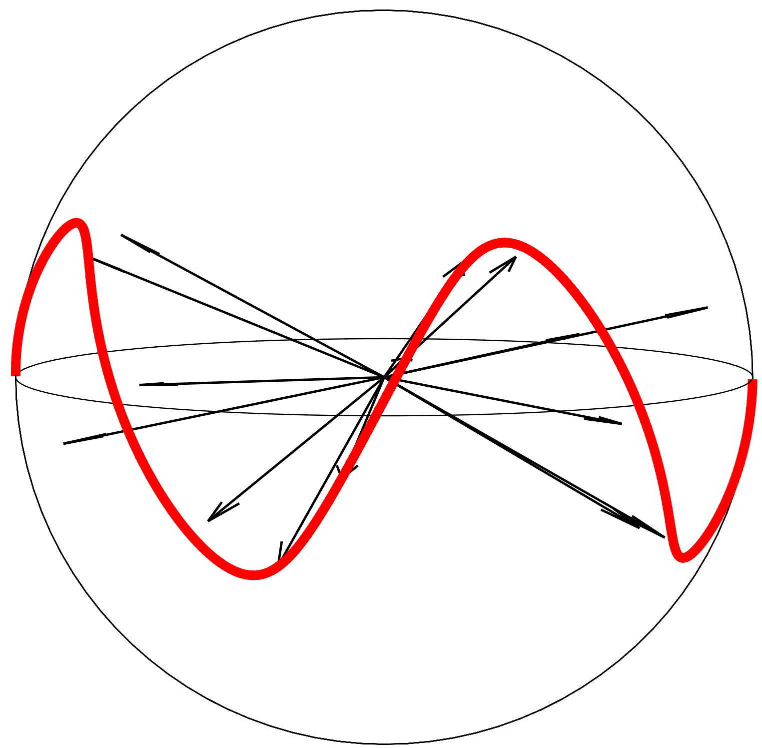

$$\mathbf{x}(t) = \Phi\,\alpha(t) + \mathbf{n}(t)$$

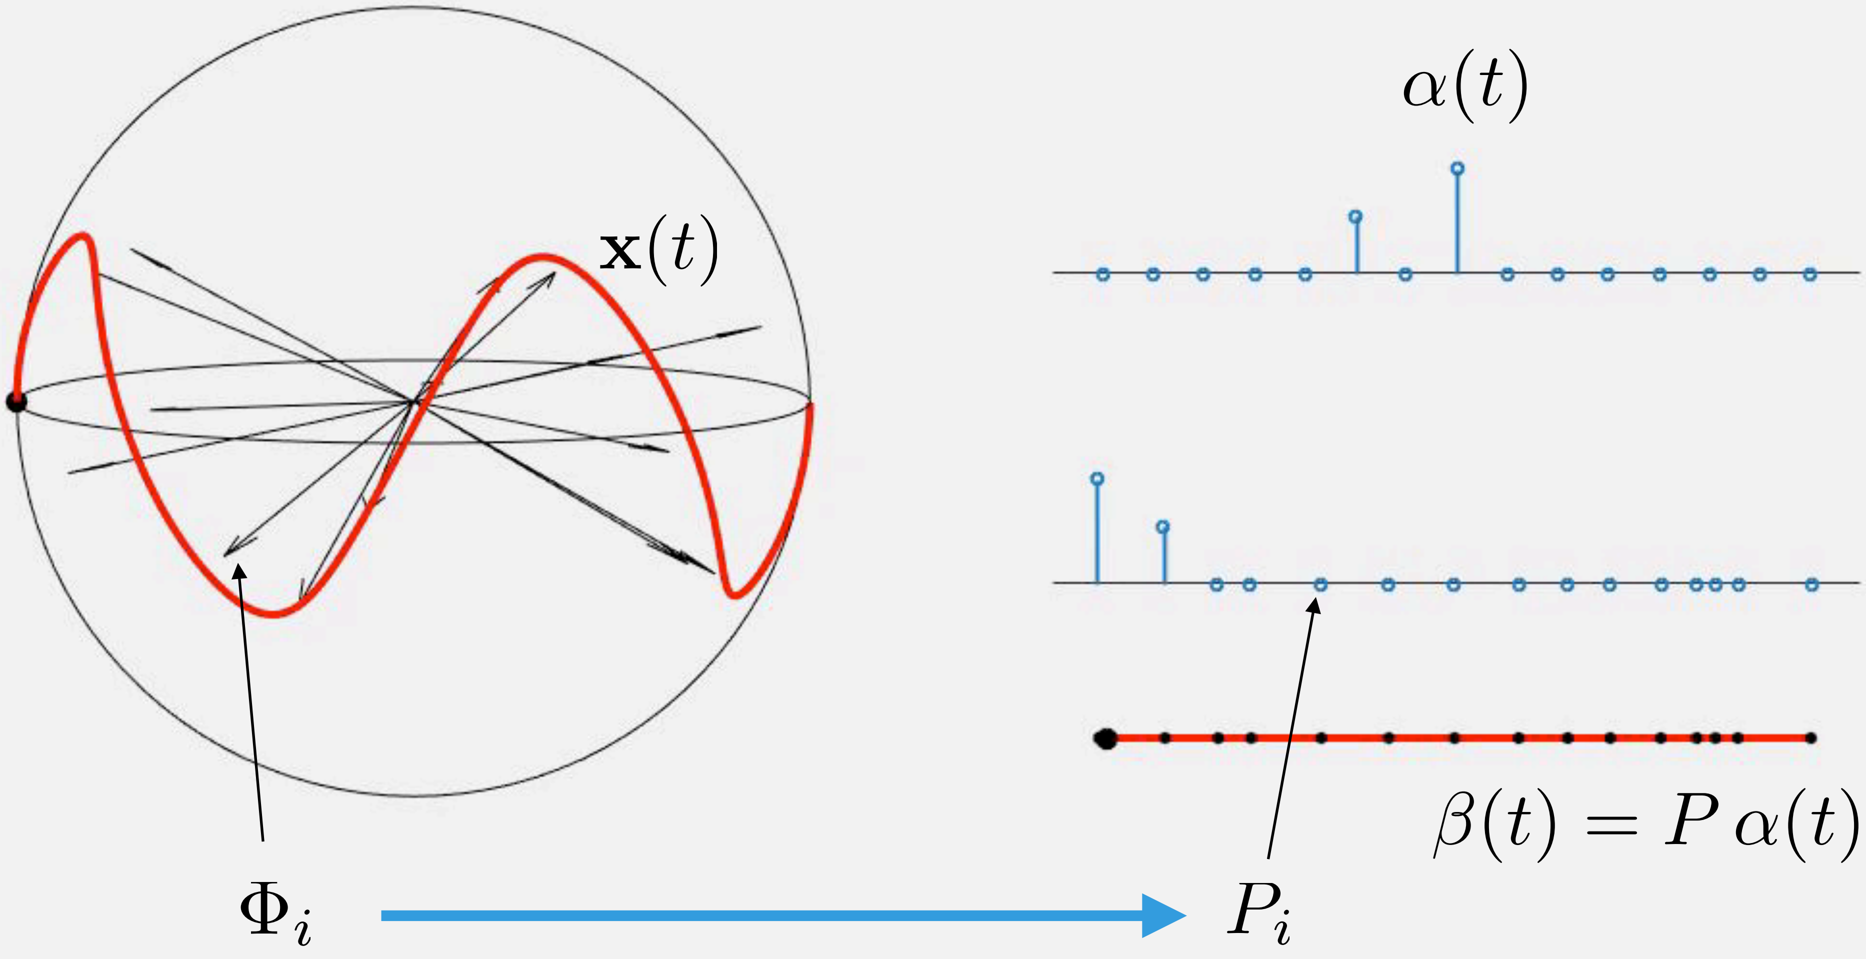

Sparse Manifold Transform
(Yubei Chen, Ph.D. thesis)

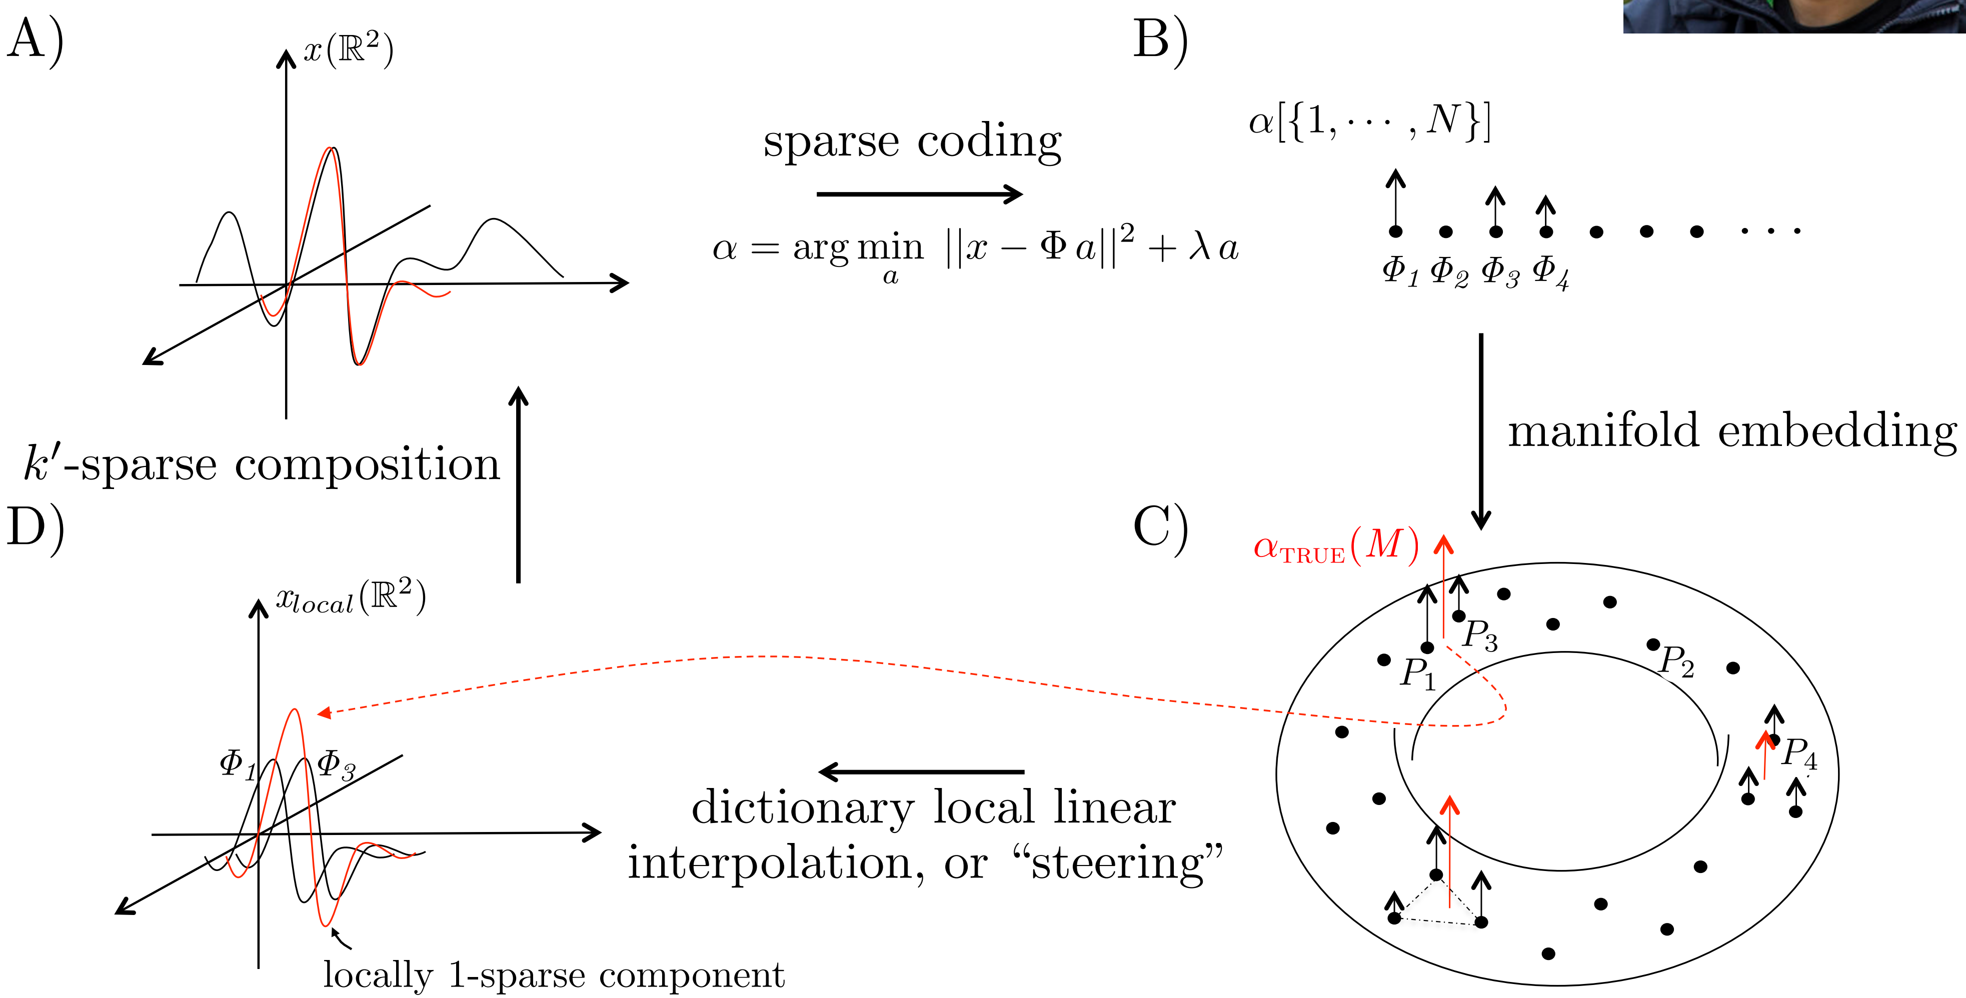

We seek a geometric mapping $f : \Phi \rightarrow P$, s.t. each of the dictionary elements is mapped to a new vector, $P_j = f(\Phi_j)$. Continuous temporal transformations in the input should have a linear flow on M and also in the geometrical embedding space.

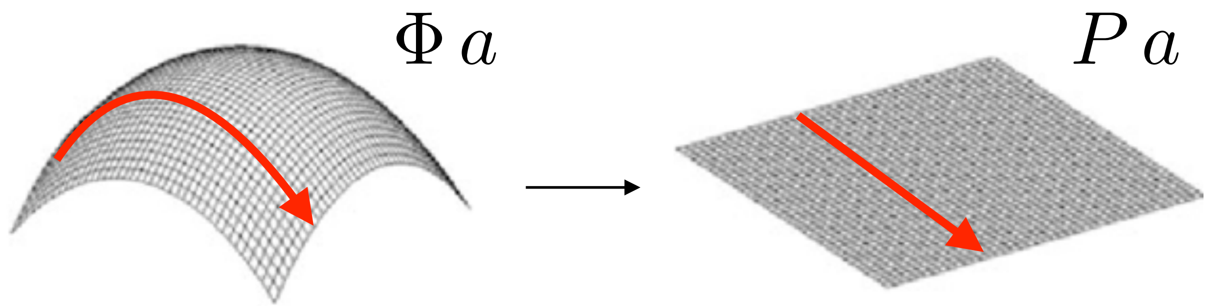

We desire: $$P a_t \approx \frac{1}{2} P a_{t-1} + \frac{1}{2} P a_{t+1}$$

Objective function: $$\min_P ||P\,A\,D||_F^2 + \gamma\,||P\,V||_1$$

s.t. $$P\,V\,P^T = I$$

$$V = \mathrm{Cov}(a)$$

$$D = \begin{bmatrix} 1 & -\frac{1}{2} & 0 & 0 & \dots & 0 \\ -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & \dots & 0 \\ 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & \dots & 0 \\ 0 & 0 & -\frac{1}{2} & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 0 & -\frac{1}{2} & 1 \end{bmatrix}$$

The sparse manifold transform

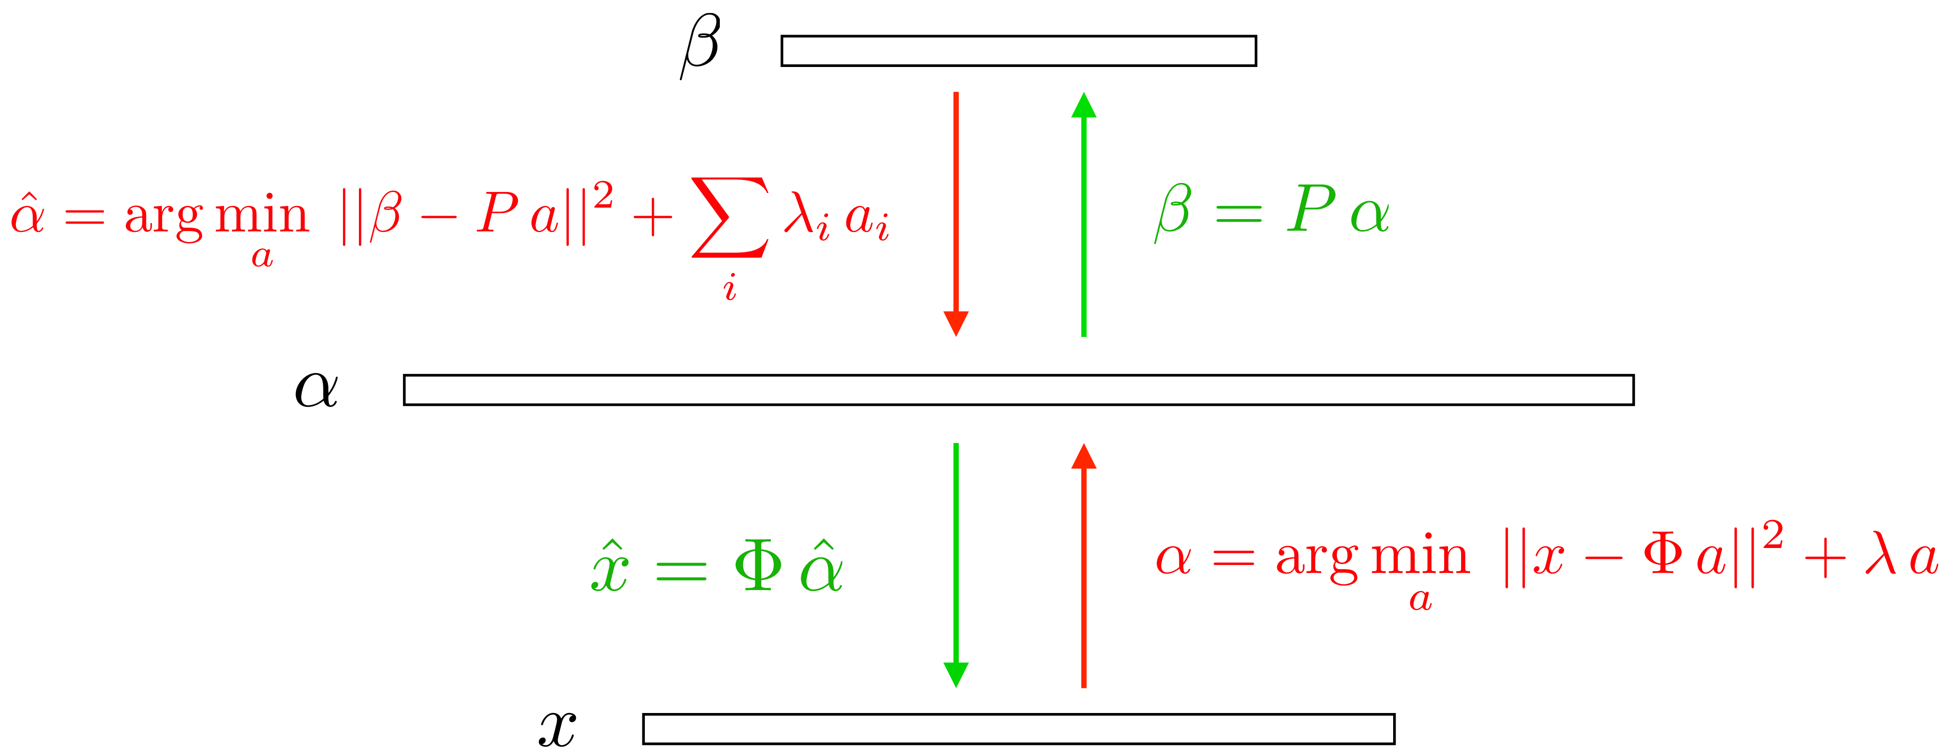

Encoding of a natural video sequence

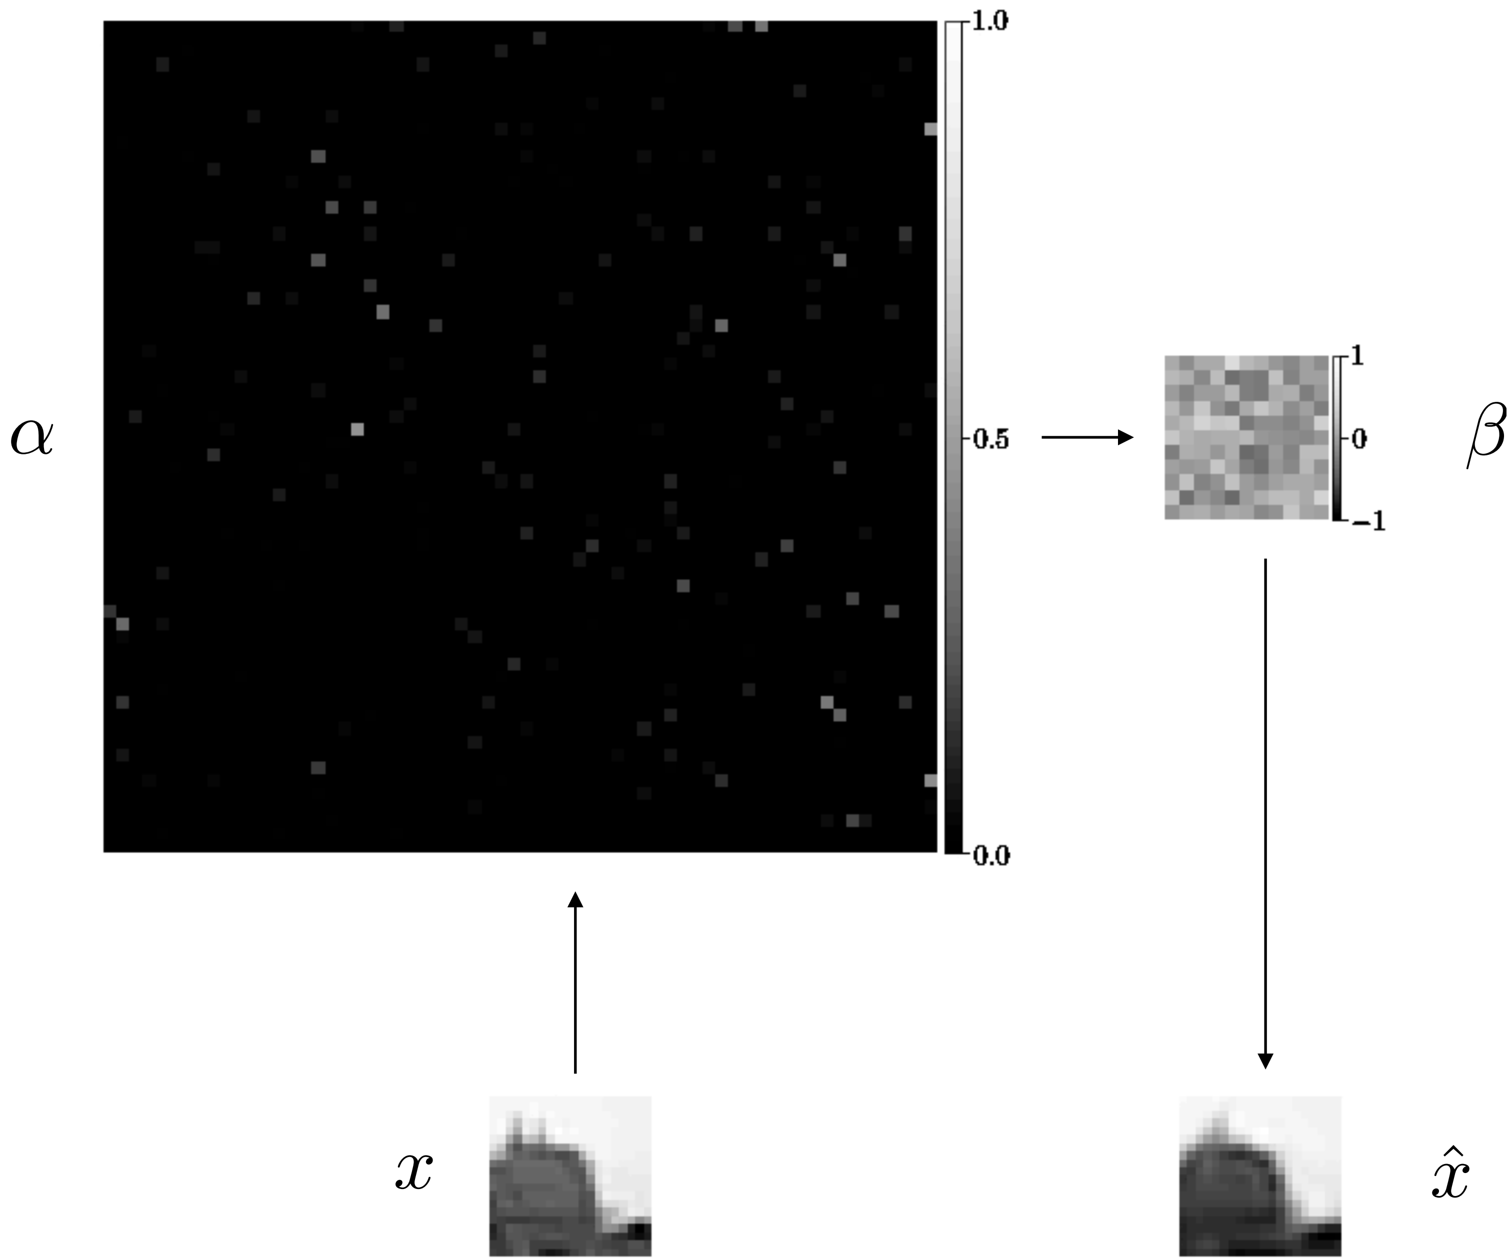

β units behave like complex-cells

Learned pooling of α subunits

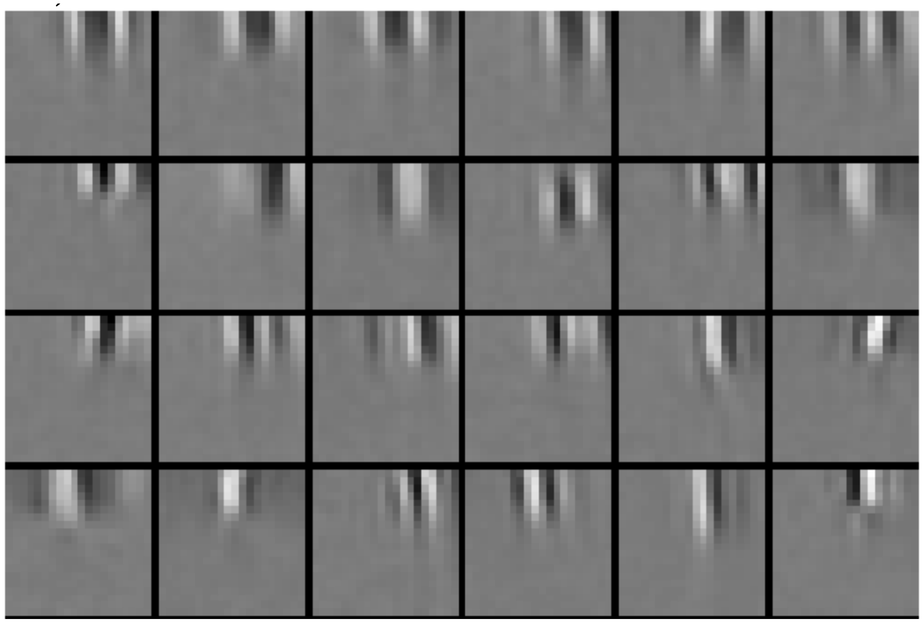

Random sample of 21 learned β units

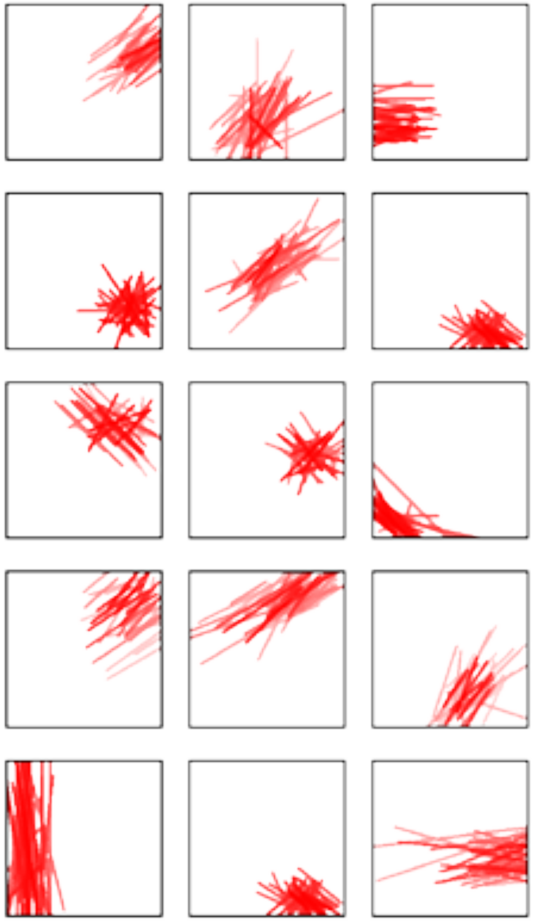

Phase invariance

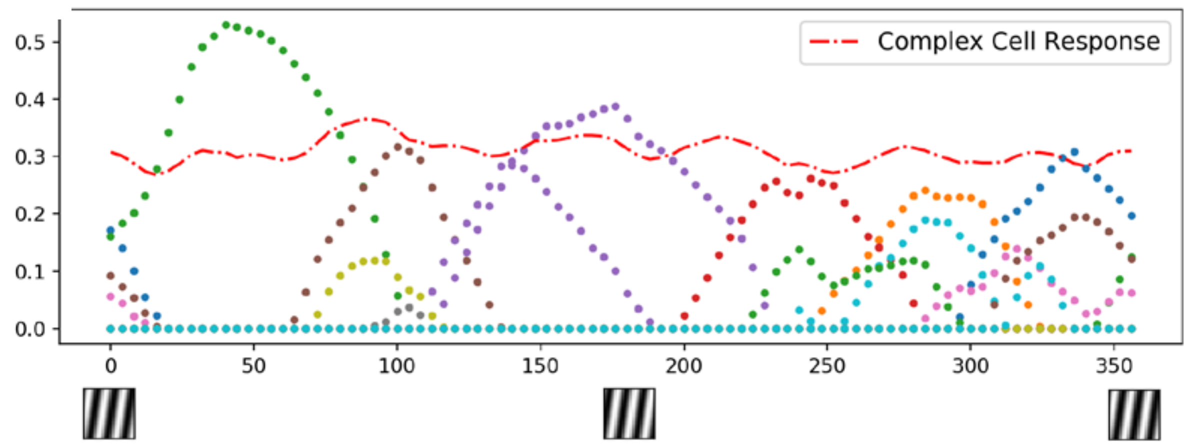

Orientation selectivity

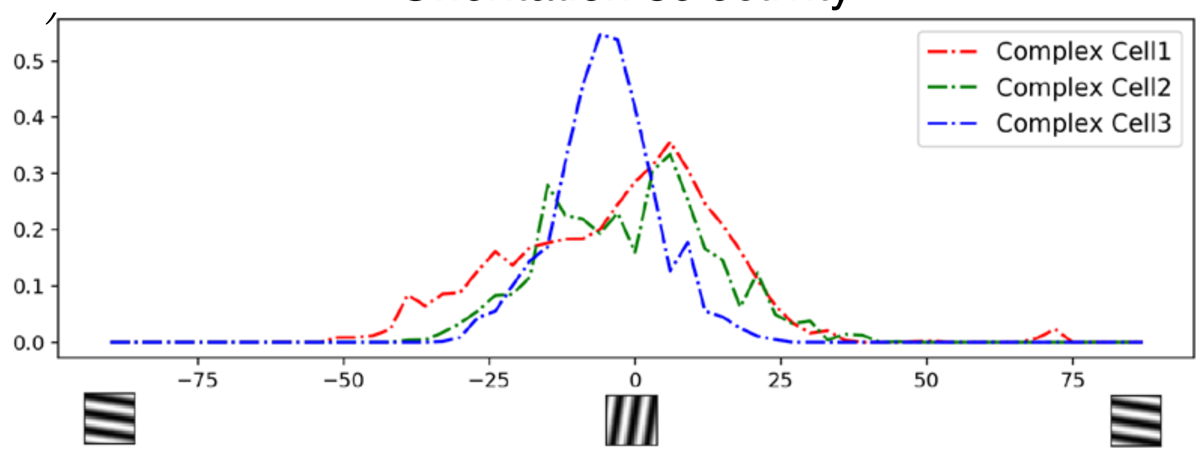

F1/F0 ratio

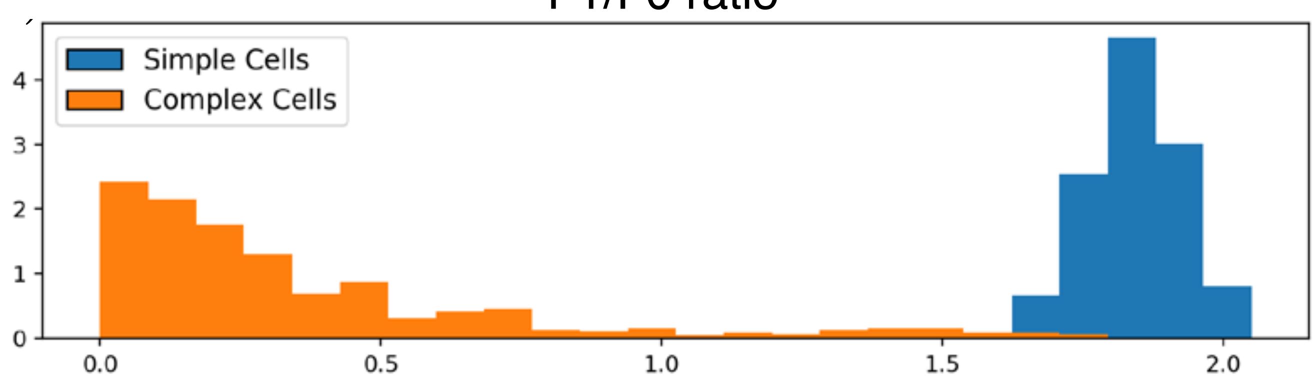

The *stacked* sparse manifold transform

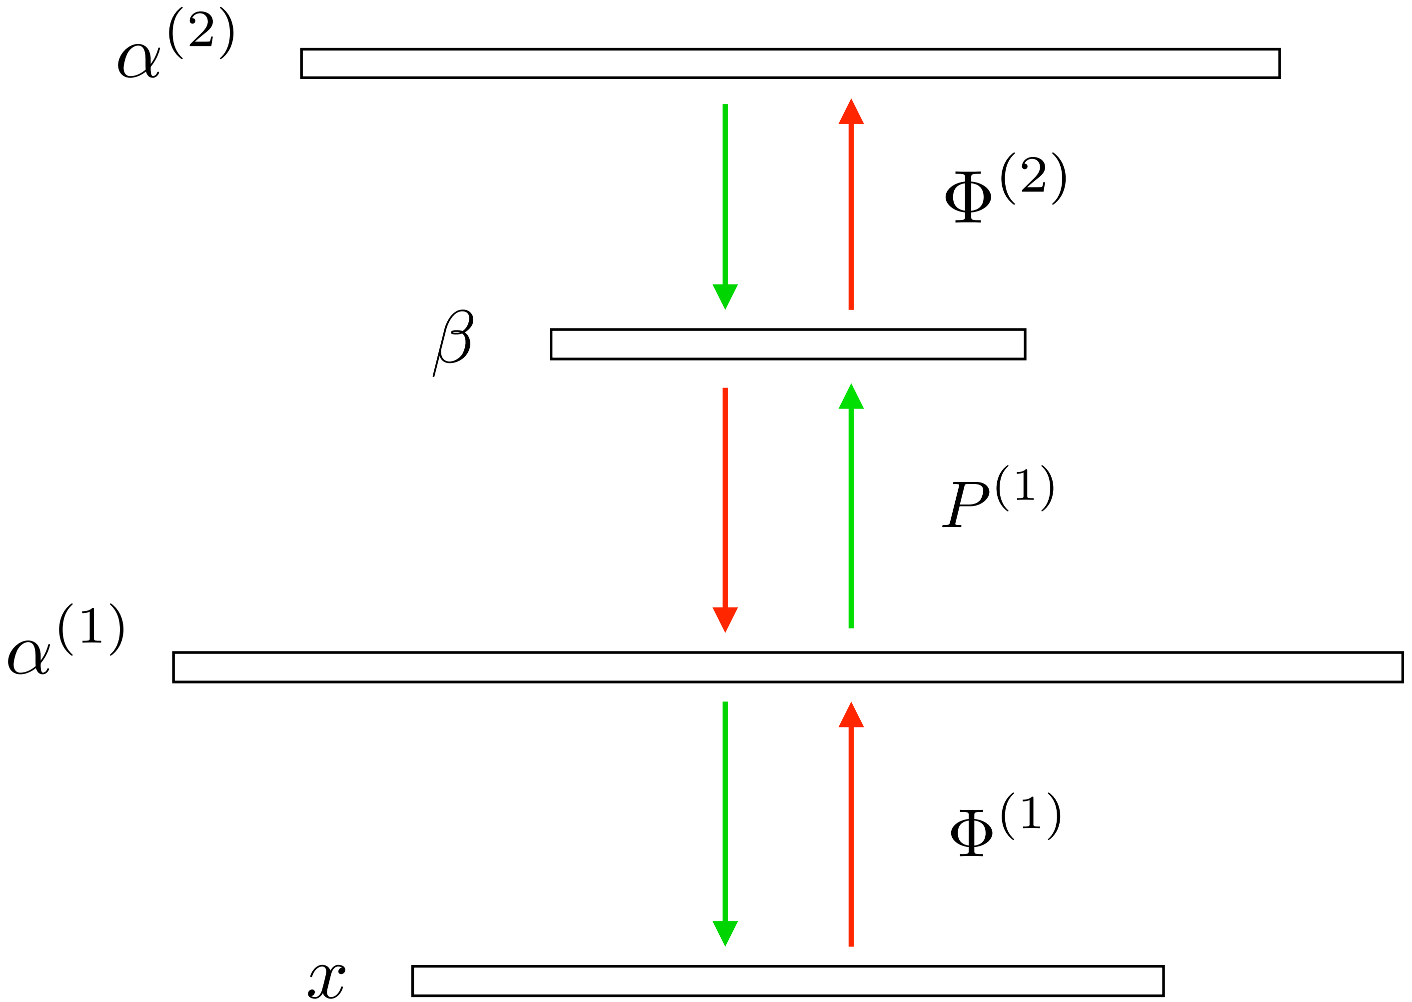

$$\min_{\alpha^{(1)},\alpha^{(2)}} \quad ||\Phi^{(2)}\alpha^{(2)} - P^{(1)}\alpha^{(1)}||^2 + ||x - \Phi^{(1)}\alpha^{(1)}||^2 + S(\alpha^{(1)}) + S(\alpha^{(2)})$$

Stacked Sparse Manifold Transform
(trained on MNIST)

Learned $\Phi^{(1)}$

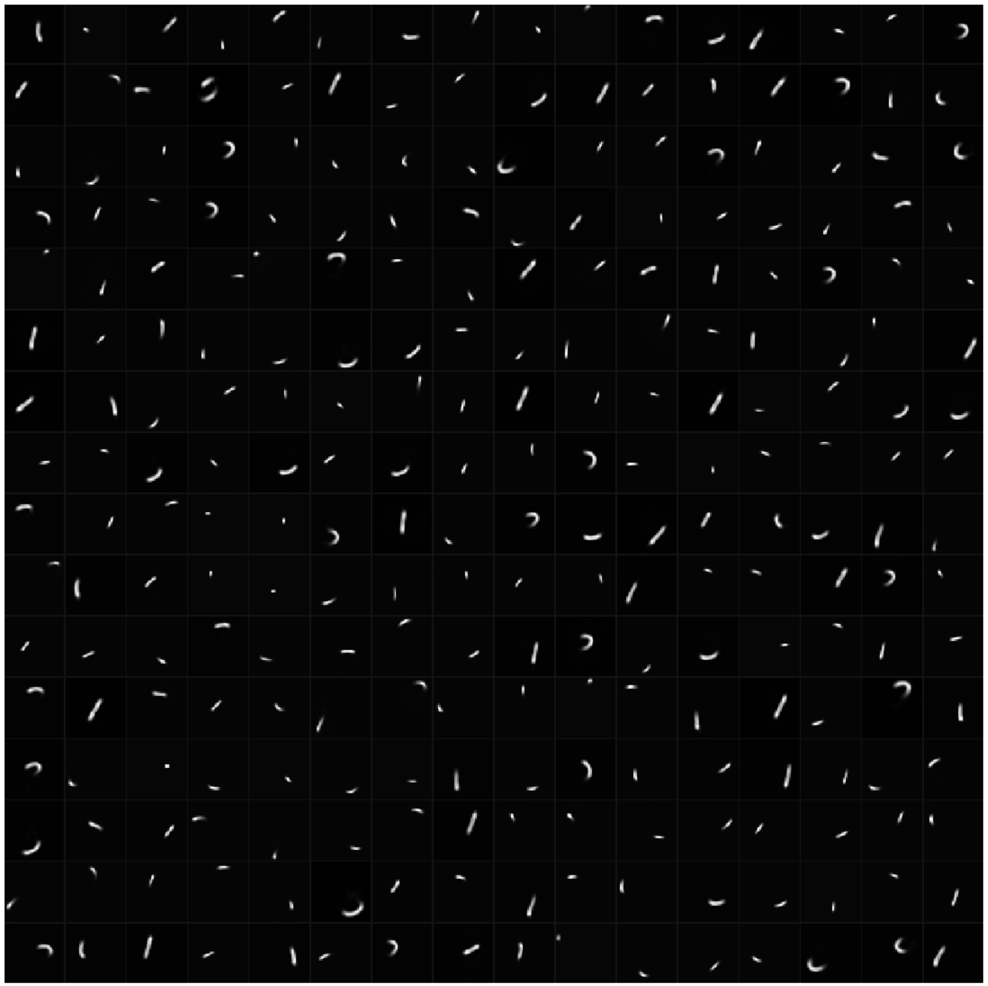

Learned $\Phi^{(2)}$