Subspace Locally Competitive Algorithms

Dylan M Paiton
Steven Shepard
Kwan Ho Ryan Chan
Bruno Olshausen

Neurally Inspired Computational Elements, ~~2020~~ 2021

Useful representations of natural signals

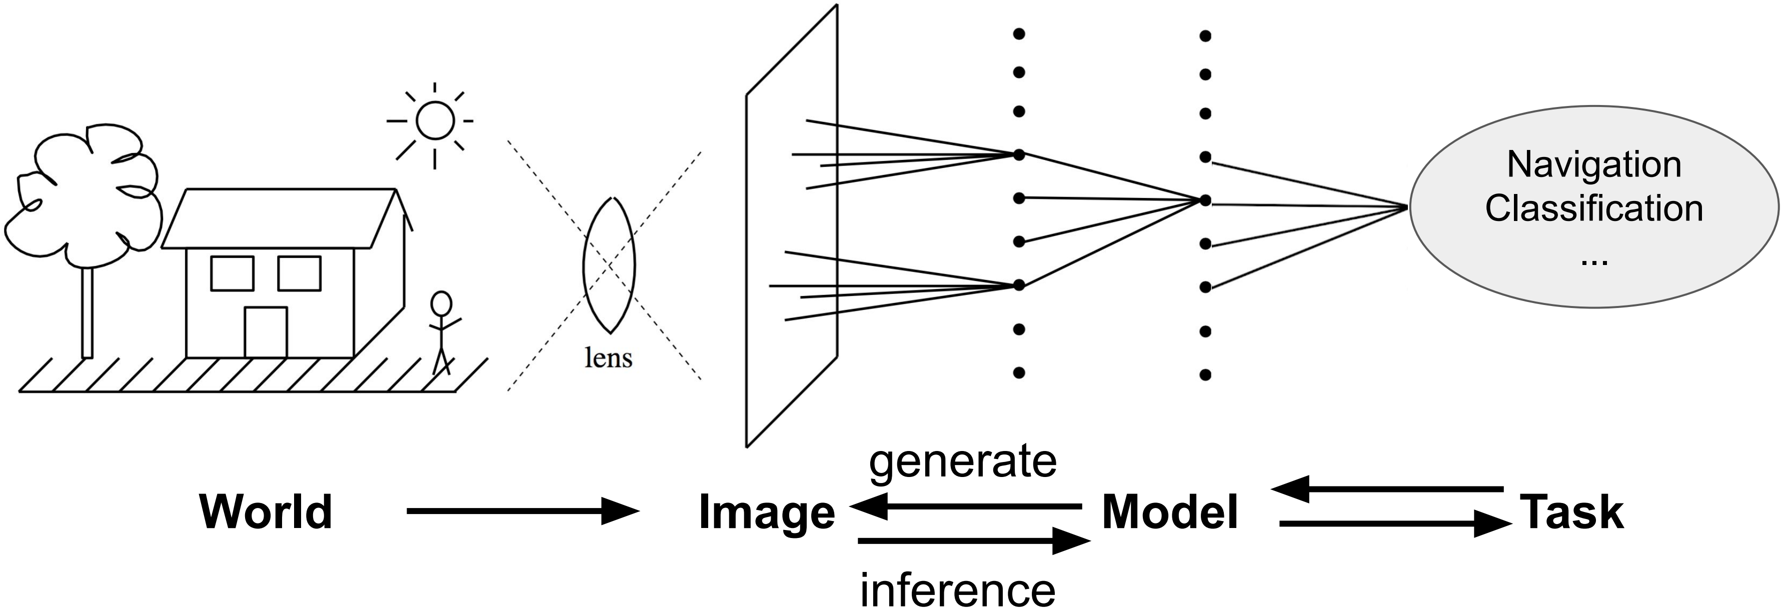

Slide adapted from BA Olshausen

Useful representations of natural signals

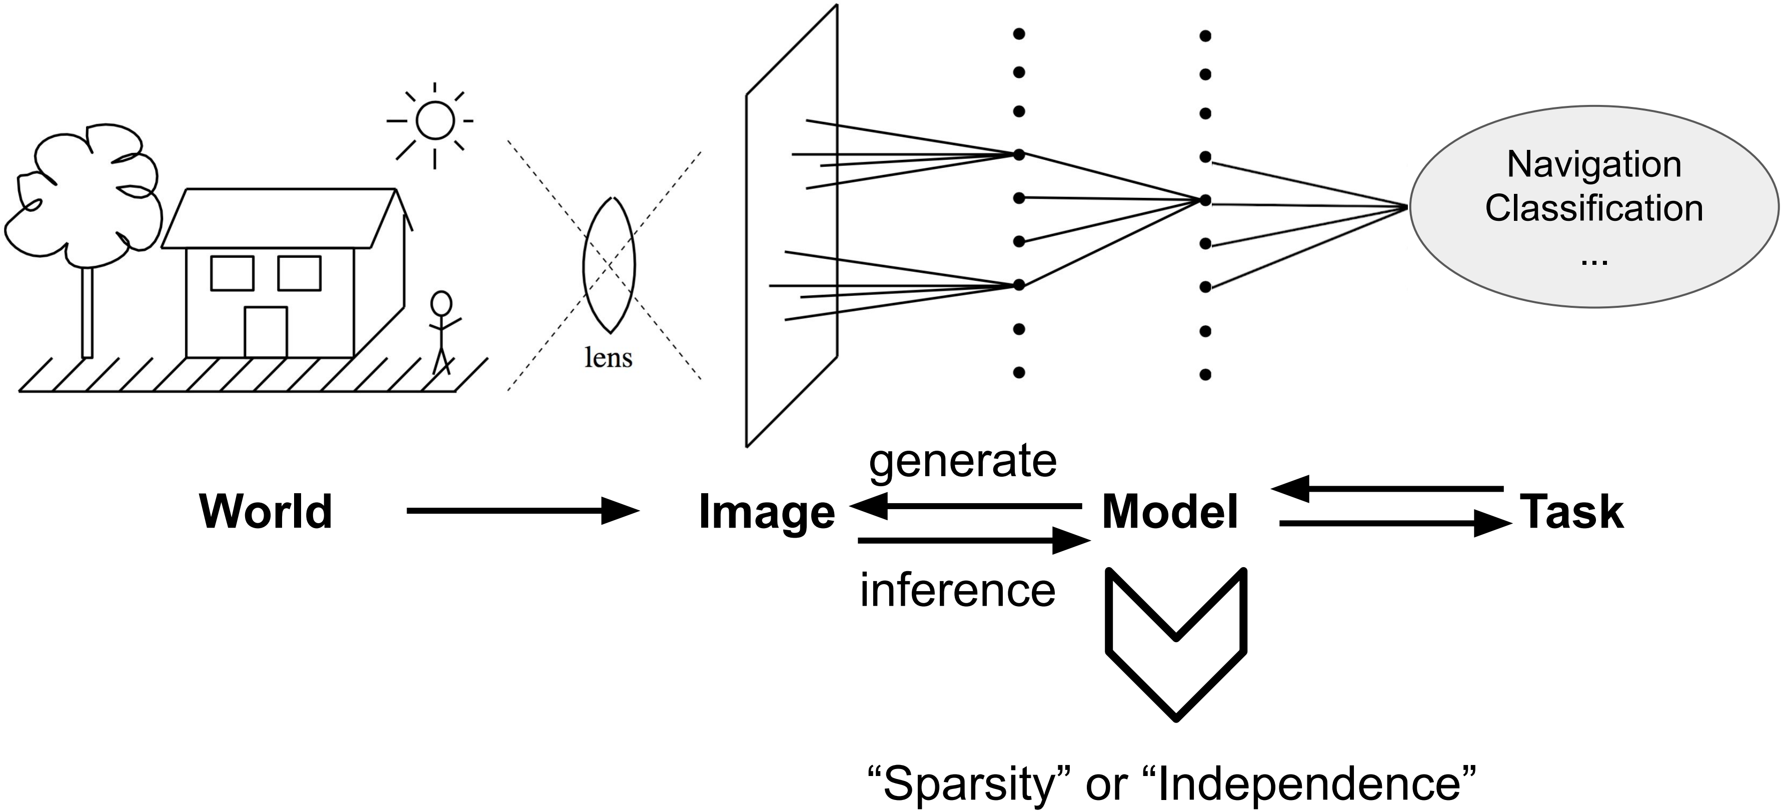

Slide adapted from BA Olshausen

“Independence” model fits biological simple cells

Parameters for Gabor Fitting of Receptive Fields (Cycles)

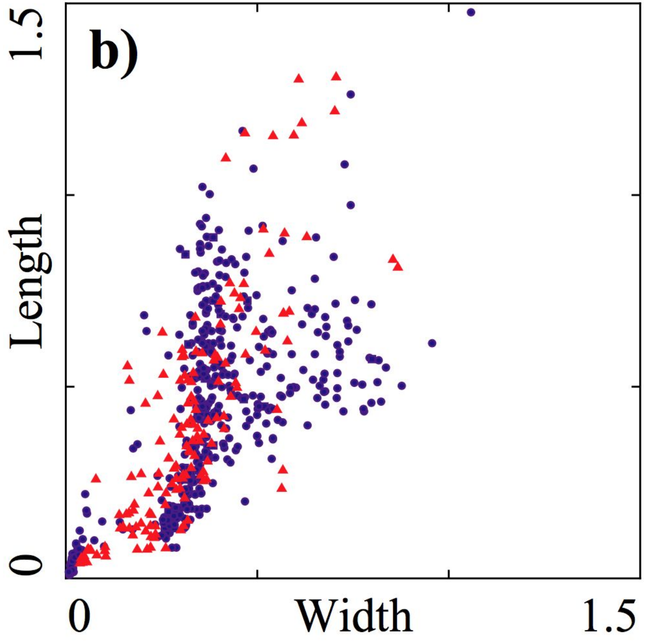

Red - Macaque; Blue - Model

Rehn and Sommer. *Journal of computational neuroscience*. 2007

Contrast Invariant Tuning in Mammals vs Model

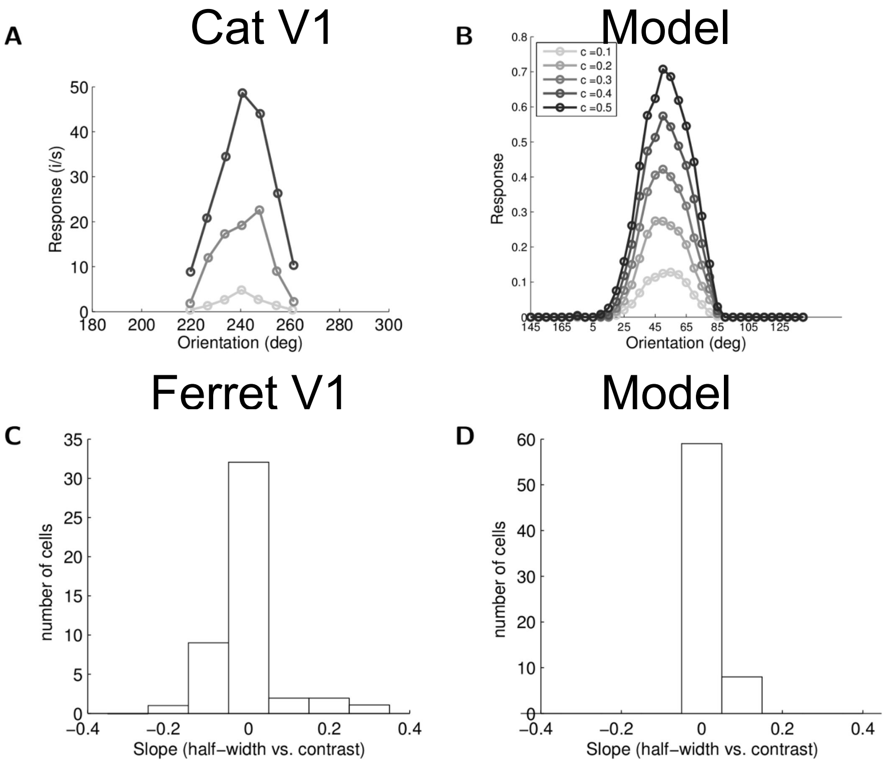

Zhu and Rozell. *PLoS Computational Biology*. 2013

Overcomplete sparse inference

Sparse Coding via Locally Competitive Algorithms (LCA)

$$E_{\mathrm{LCA}} = \frac{1}{2}||s - \Phi a||_2^2 + \lambda \sum_k C(a_k)$$

Preserve as much information as possible

Minimal neuron activity

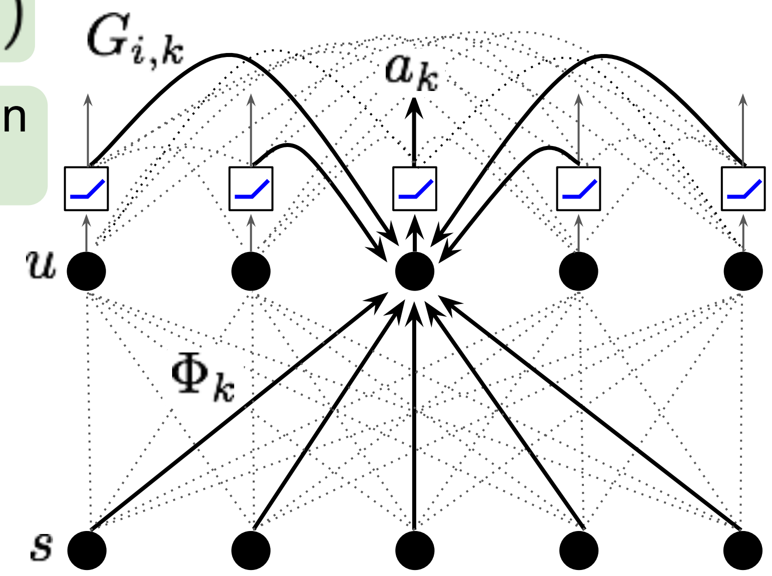

CJ Rozell et al. *Neural Computation*. 2008

Overcomplete sparse inference

Sparse Coding via Locally Competitive Algorithms

$$E_{\mathrm{LCA}} = \frac{1}{2}||s - \Phi a||_2^2 + \lambda \sum_k C(a_k)$$

Preserve as much information as possible

Minimal neuron activity

$$\dot{u} = \frac{\partial E}{\partial a}$$

$$\tau \dot{u}_k + u_k = \Phi_k^\top s - G_k^\top a$$

Membrane leak

Forward drive

Lateral inhibition

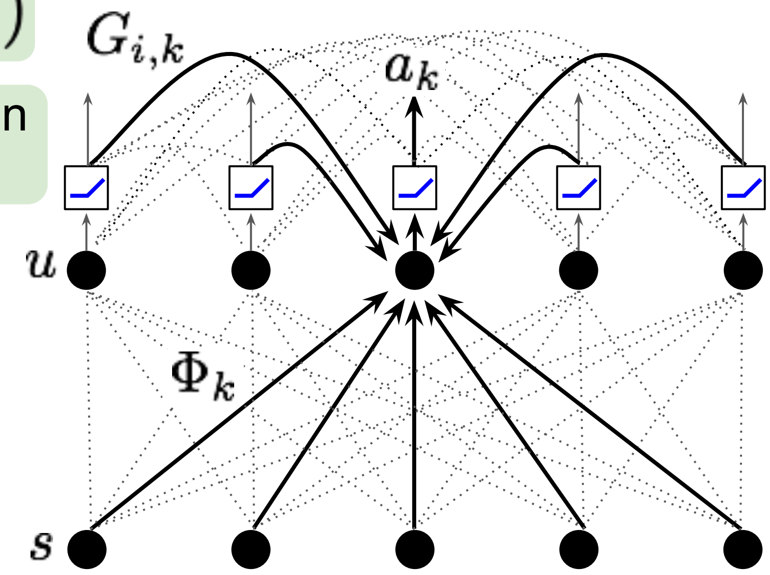

CJ Rozell et al. *Neural Computation*. 2008

Overcomplete sparse inference

Sparse Coding via Locally Competitive Algorithms

$$E_{\mathrm{LCA}} = \frac{1}{2}||s - \Phi a||_2^2 + \lambda \sum_k C(a_k)$$

Preserve as much information as possible

Minimal neuron activity

$$\dot{u} = \frac{\partial E}{\partial a}$$

$$\tau \dot{u}_k + u_k = \Phi_k^\top s - G_k^\top a$$

Membrane leak

Forward drive

Lateral inhibition

$$a_k = f_\lambda(u_k)$$

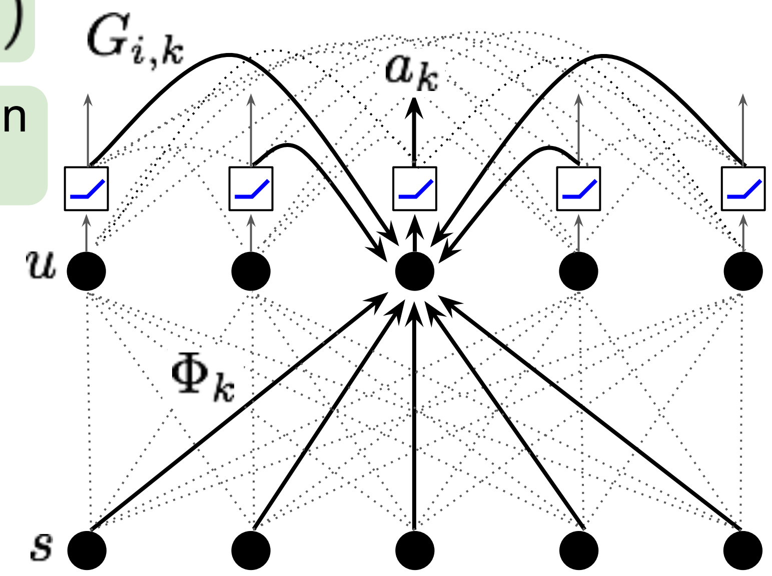

CJ Rozell et al. *Neural Computation*. 2008

Overcomplete sparse inference

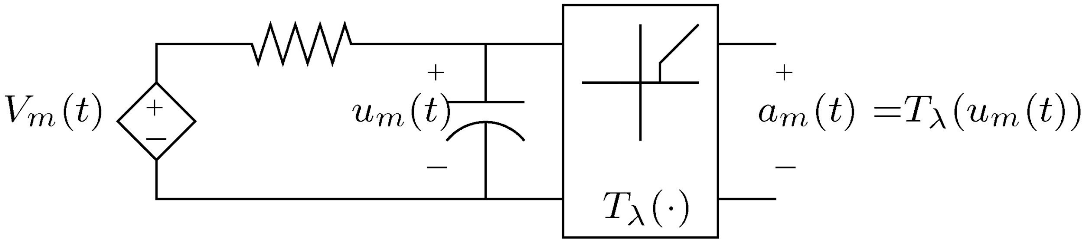

$$V_m(t) = \langle \varphi_m, s(t) \rangle - \sum_{n \neq m} T_\lambda(u_n(t)) \langle \varphi_m, \varphi_n \rangle$$

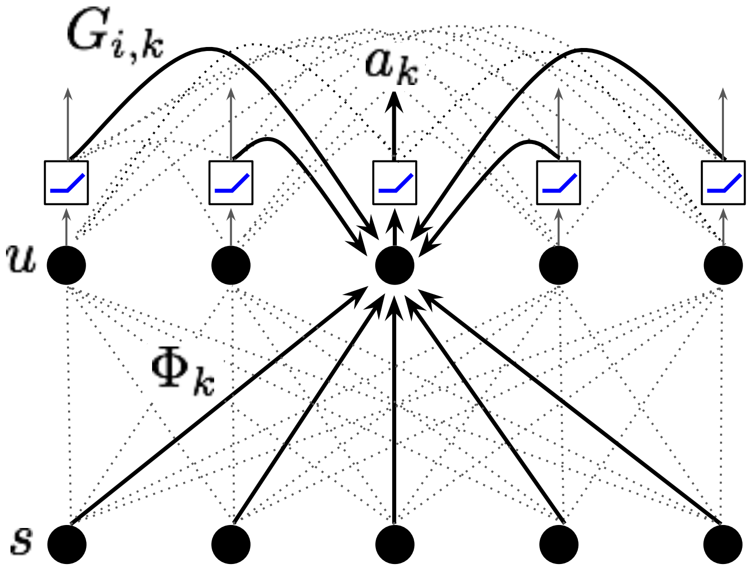

$$\tau \dot{u}_k + u_k = \Phi_k^\top s - G_k^\top a$$

$$a_k = f_\lambda(u_k)$$

CJ Rozell et al. *Neural Computation*. 2008

Comparing to an "energy model" complex cell

Subspace Independent Component Analysis (ISA)

$$u_{nm} = \Phi_{nm} s$$

$$c_m = \sqrt{\sum_n u_{nm}^2}$$

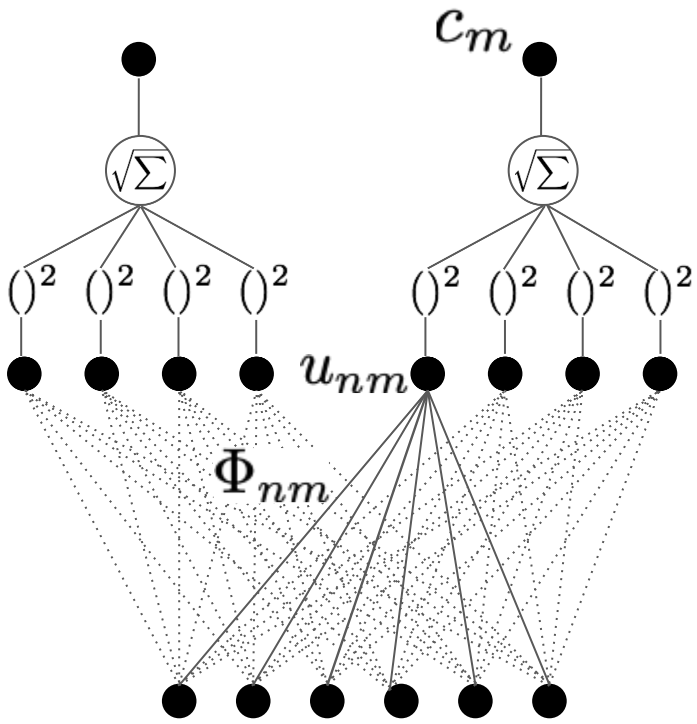

Subspace locally competitive algorithms

layer 1 neuron activation

$$T_\lambda(u_{nm}) := f_\lambda^{-1}(u_{nm}) = \begin{cases} 0, & ||u_m||_2 \leq \lambda \\ (||u_m||_2 - \lambda)\dfrac{u_{nm}}{||u_m||_2}, & ||u_m||_2 > \lambda. \end{cases}$$

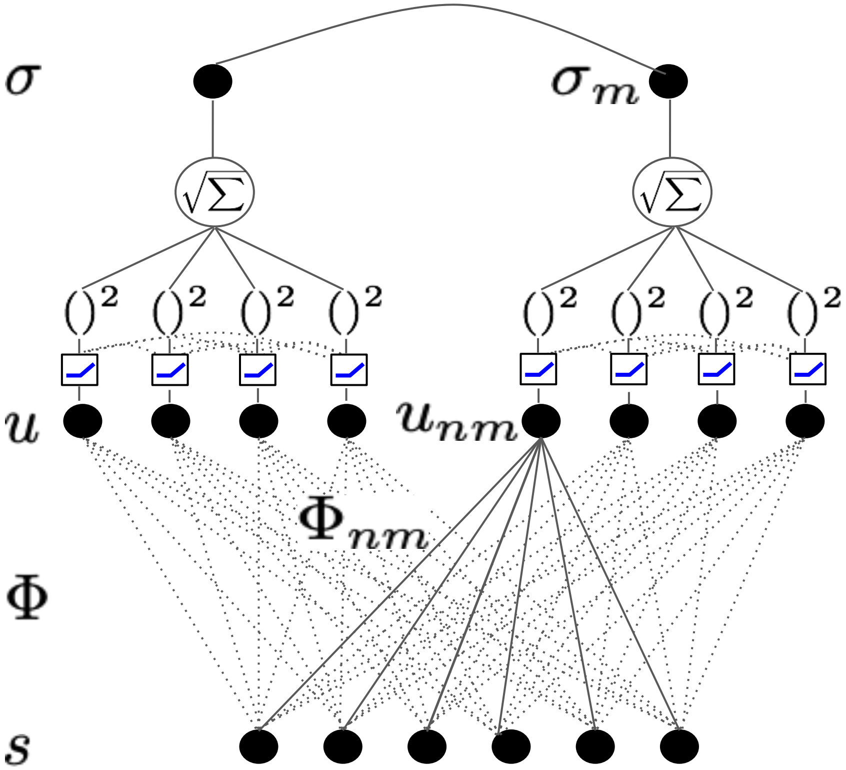

Subspace locally competitive algorithms

layer 1 neuron activation

$$T_\lambda(u_{nm}) := f_\lambda^{-1}(u_{nm}) = \begin{cases} 0, & ||u_m||_2 \leq \lambda \\ (||u_m||_2 - \lambda)\dfrac{u_{nm}}{||u_m||_2}, & ||u_m||_2 > \lambda. \end{cases}$$

layer 2 group amplitude

$$\sigma_m = ||a_m||_2 = \sqrt{\sum_{i=1}^{N} a_{im}^2}$$

Complex cell output

Simple cell output

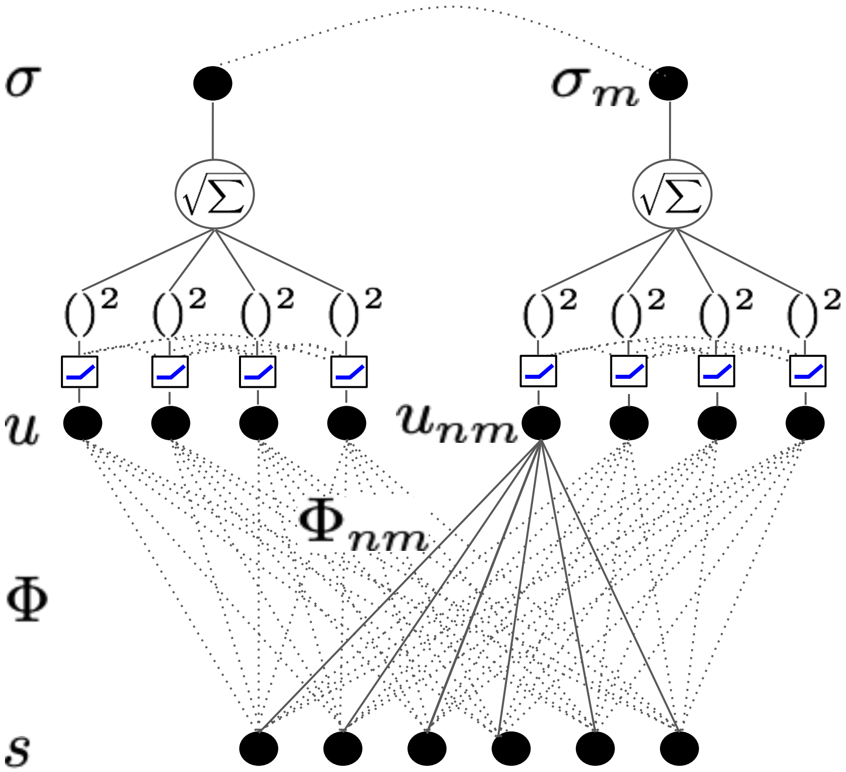

Subspace locally competitive algorithms

layer 1 neuron activation

$$T_\lambda(u_{nm}) := f_\lambda^{-1}(u_{nm}) = \begin{cases} 0, & ||u_m||_2 \leq \lambda \\ (||u_m||_2 - \lambda)\dfrac{u_{nm}}{||u_m||_2}, & ||u_m||_2 > \lambda. \end{cases}$$

layer 2 group amplitude

$$\sigma_m = ||a_m||_2 = \sqrt{\sum_{i=1}^{N} a_{im}^2}$$

steering vector (phase)

$$z_{nm} = \frac{a_{nm}}{\sigma_m}$$

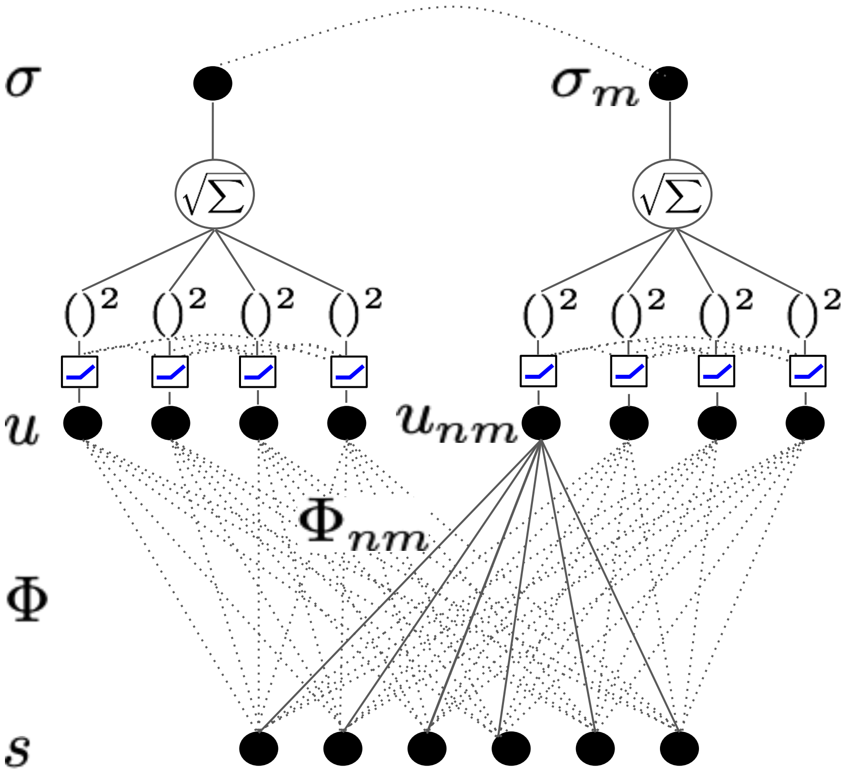

Parameter Sweep

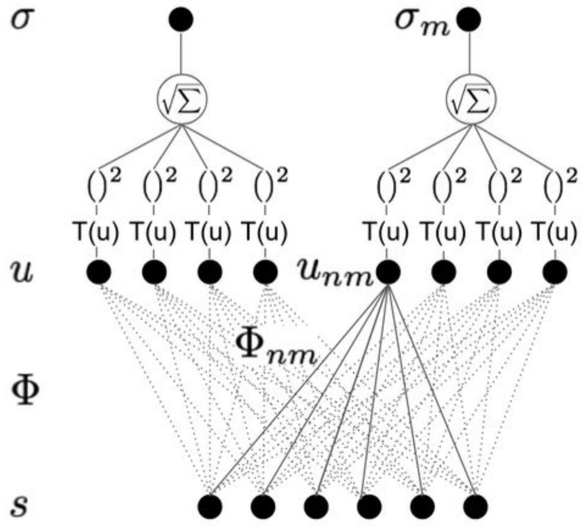

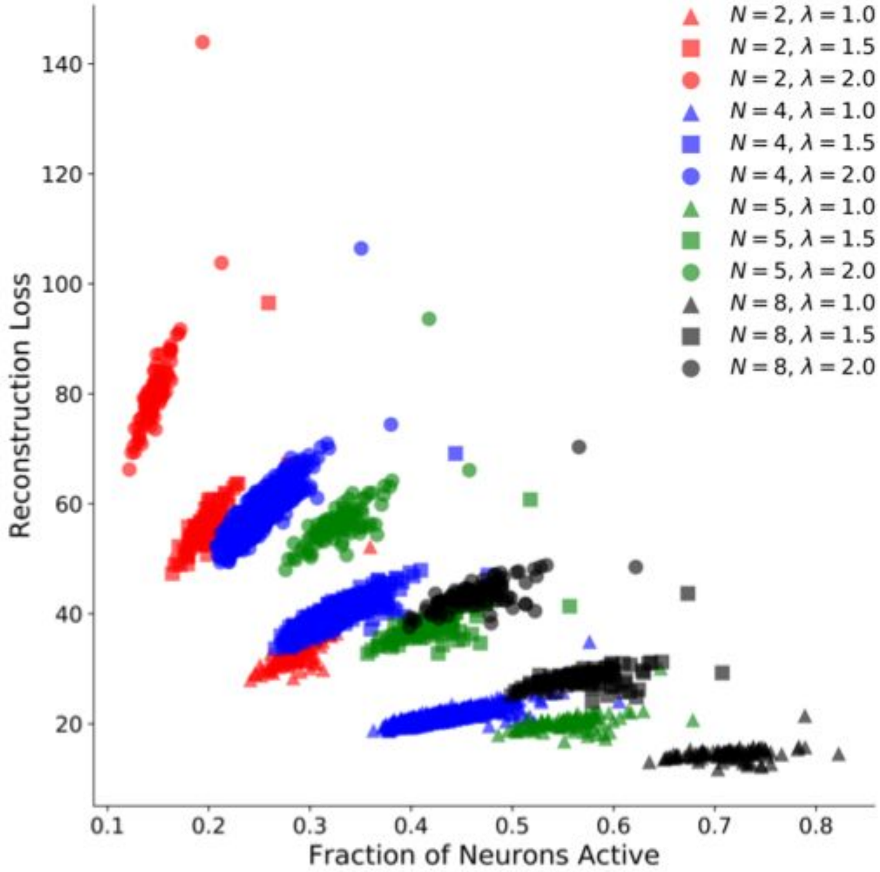

SLCA learns to sample the space of generators

Filters learned

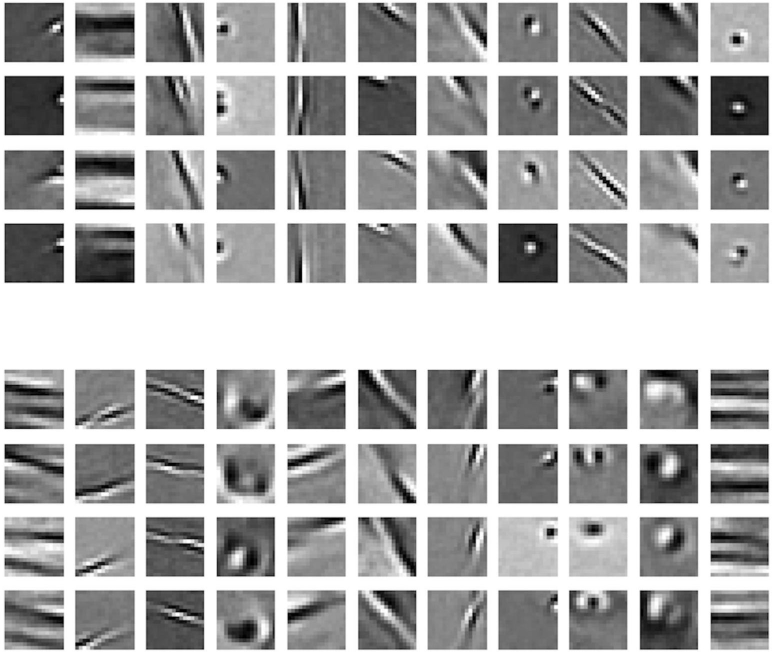

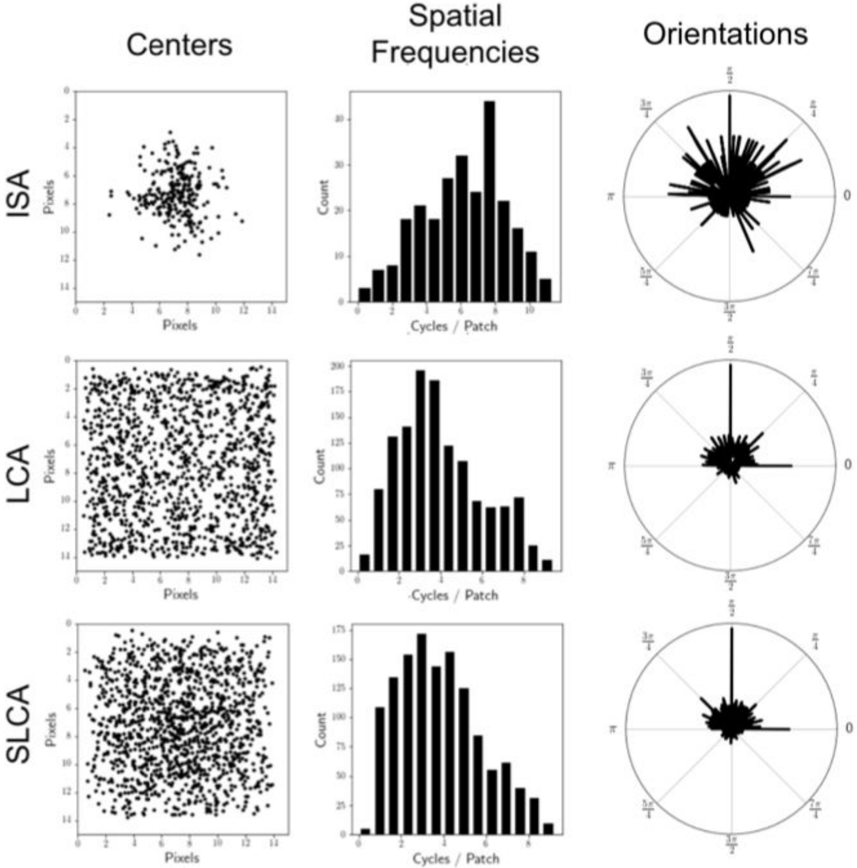

SLCA learned invariances

complex cell output

$$\sigma_m = ||a_m||_2$$

steering vector (angle)

$$z_{nm} = \frac{a_{nm}}{\sigma_m}$$

Bases

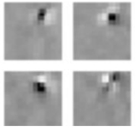

SLCA learned invariances

complex cell output

$$\sigma_m = ||a_m||_2$$

steering vector (angle)

$$z_{nm} = \frac{a_{nm}}{\sigma_m}$$

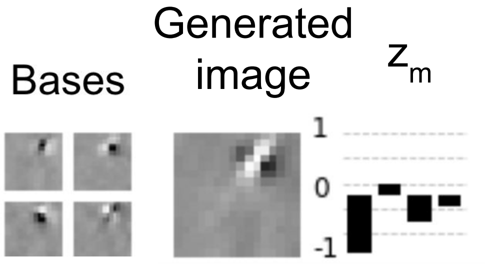

SLCA learned invariances

complex cell output

$$\sigma_m = ||a_m||_2$$

steering vector (angle)

$$z_{nm} = \frac{a_{nm}}{\sigma_m}$$

Bases

Generated images and corresponding activations

phase

SLCA learned invariances

complex cell output

$$\sigma_m = ||a_m||_2$$

steering vector (angle)

$$z_{nm} = \frac{a_{nm}}{\sigma_m}$$

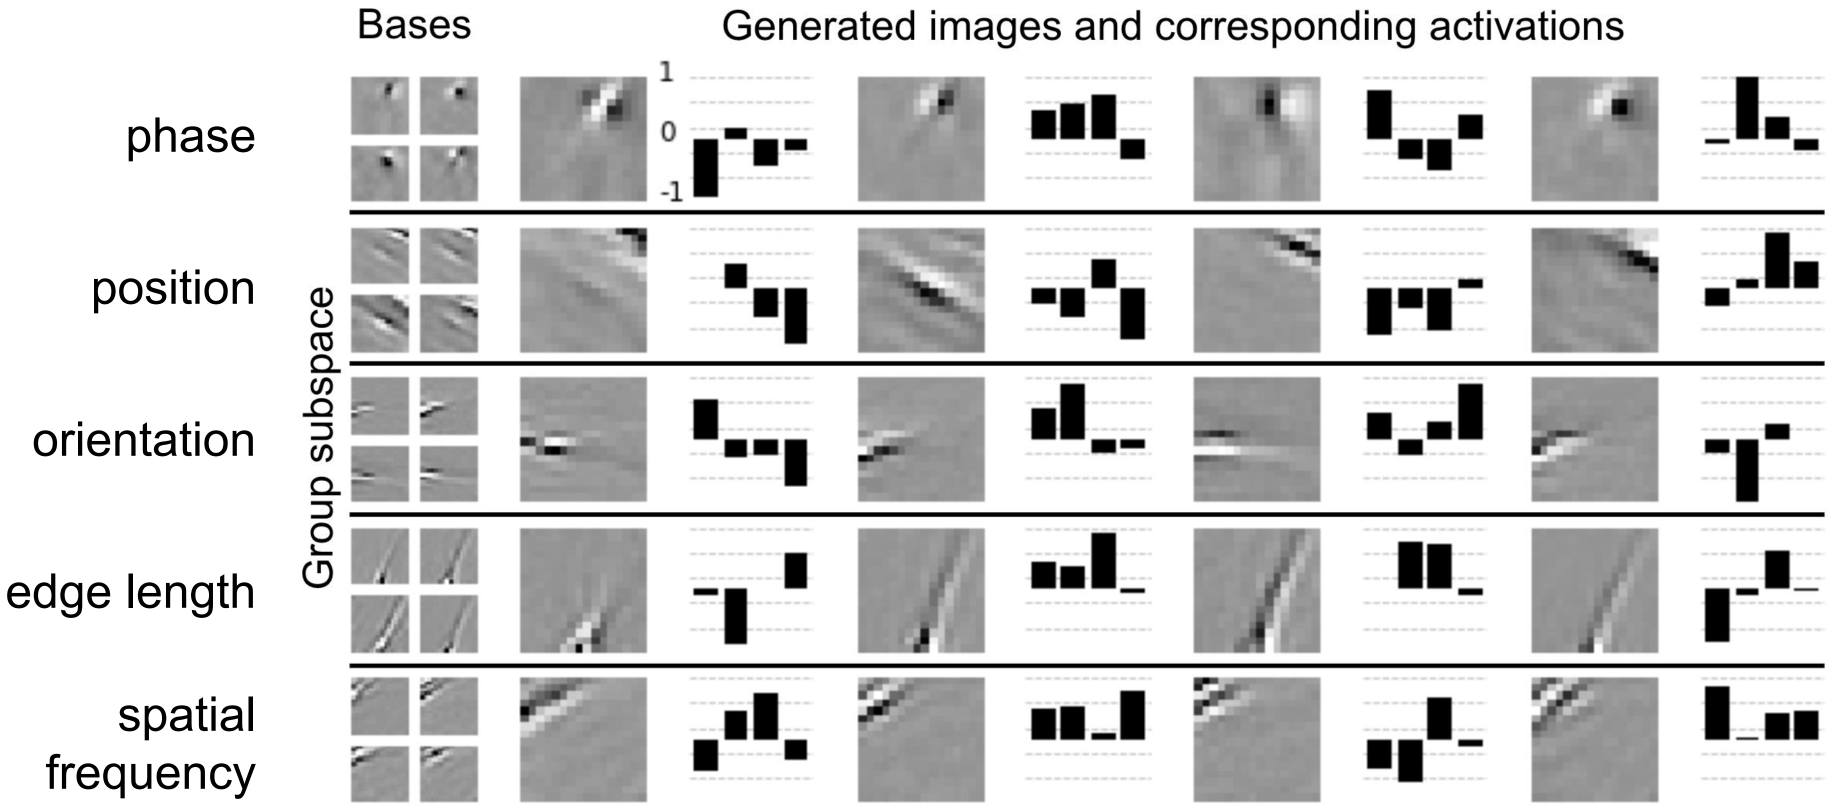

Response surface geometry

- PxP pixel image lives in $N=P^2$ dimensional space
- We select 2D cross-sections (gray in the image to the right)

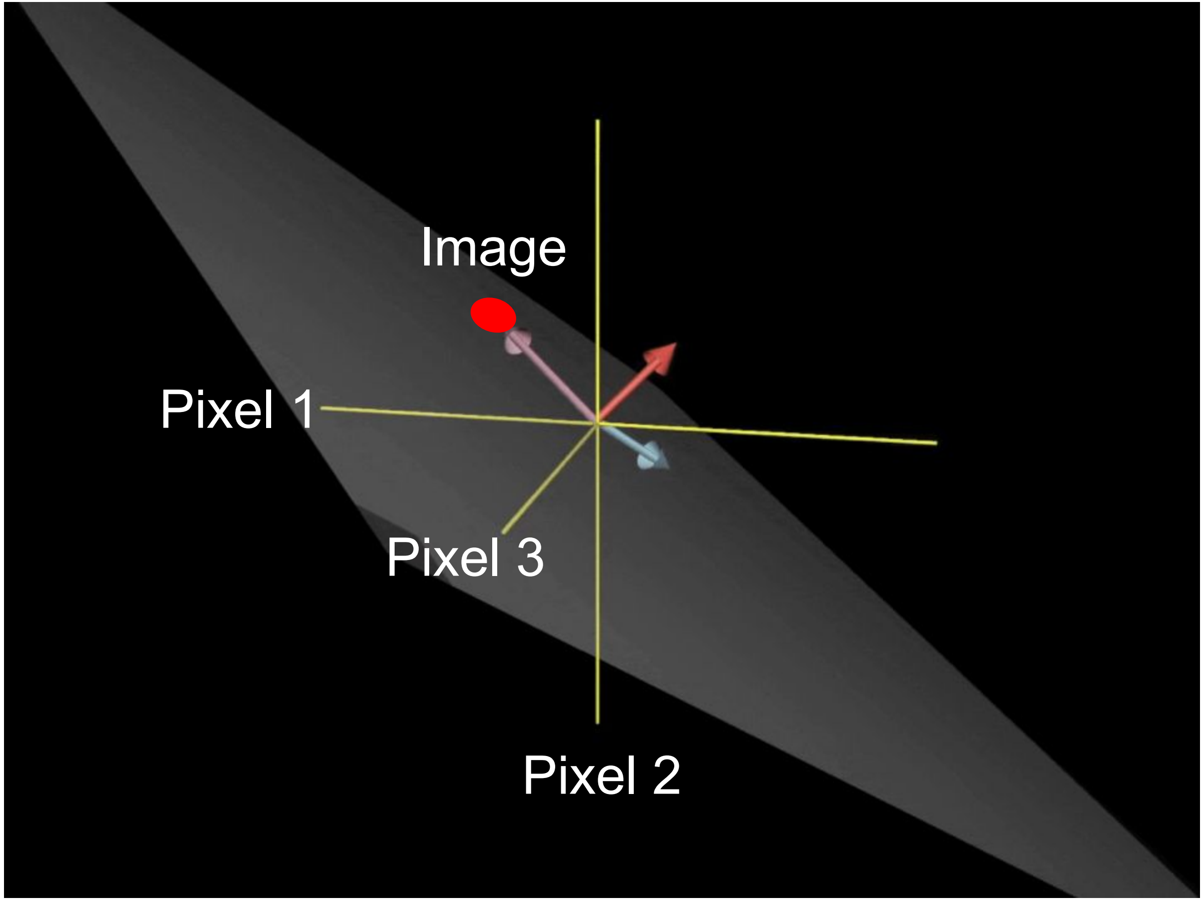

Response surface geometry

- PxP pixel image lives in $N=P^2$ dimensional space
- We select 2D cross-sections (gray in the image to the right)
- Next we measure the neuron response for a tiling of images on this cross-section

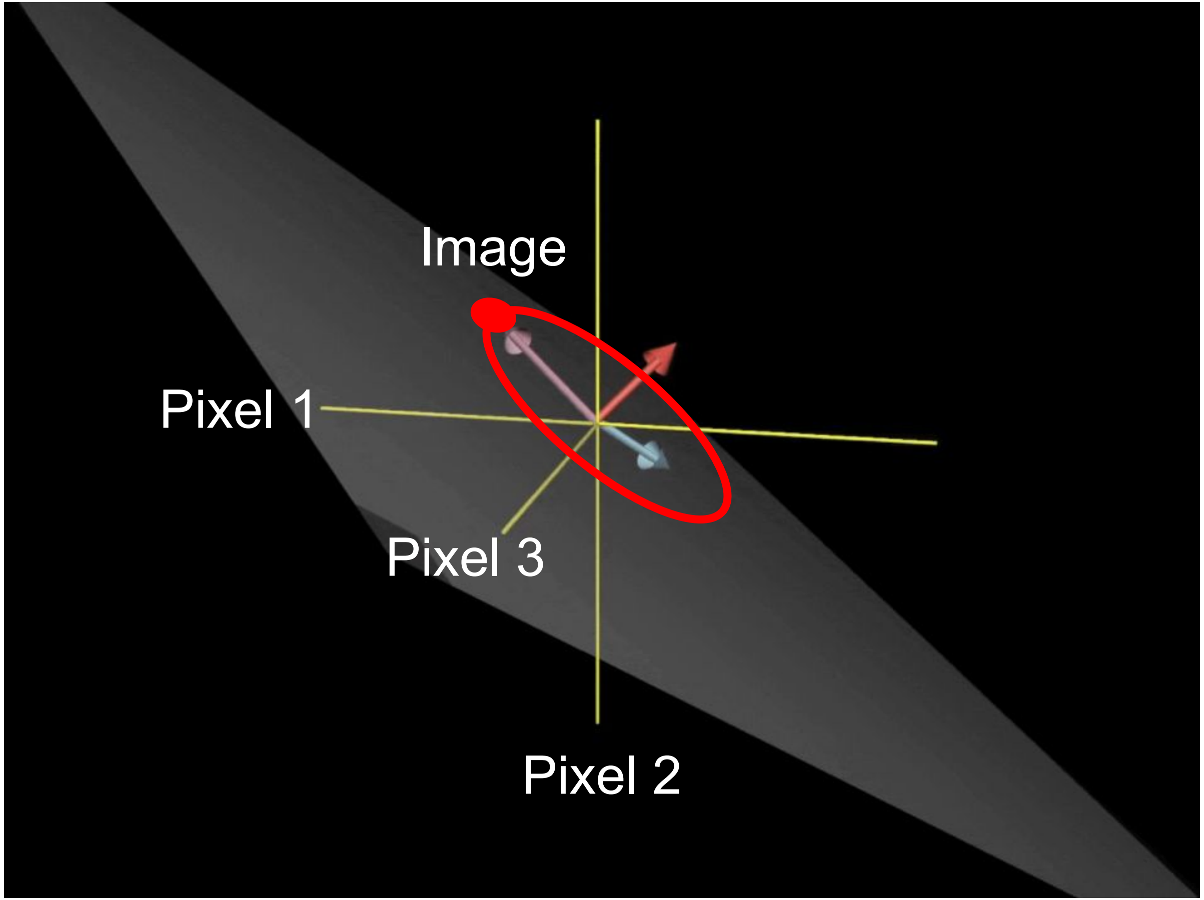

Response surface geometry - selectivity & invariance

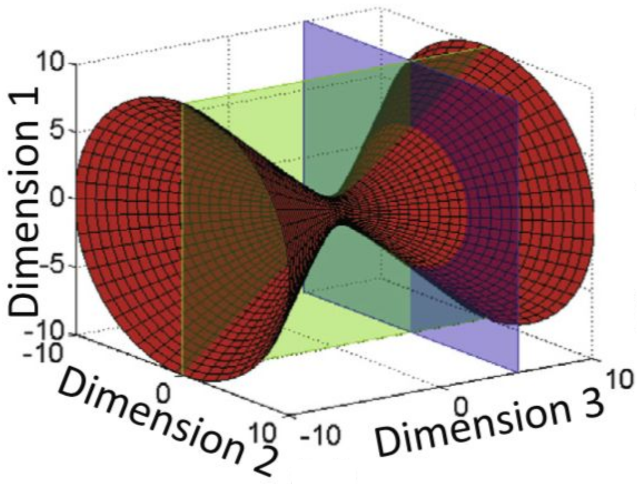

Selectivity

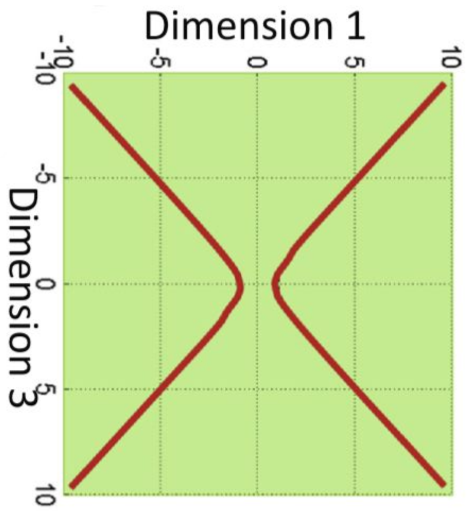

e.g. orientation selectivity

Invariance

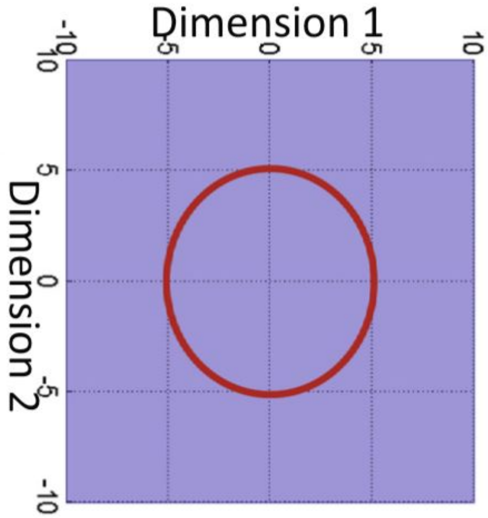

e.g. phase invariance

Golden et al. *Vision Research*. 2016

Complex cells produce “endo-origin” curvature

$$u_{nm} = \Phi_{nm} s$$

$$c_m = \sqrt{\sum_n u_{nm}^2}$$

Endo-Origin

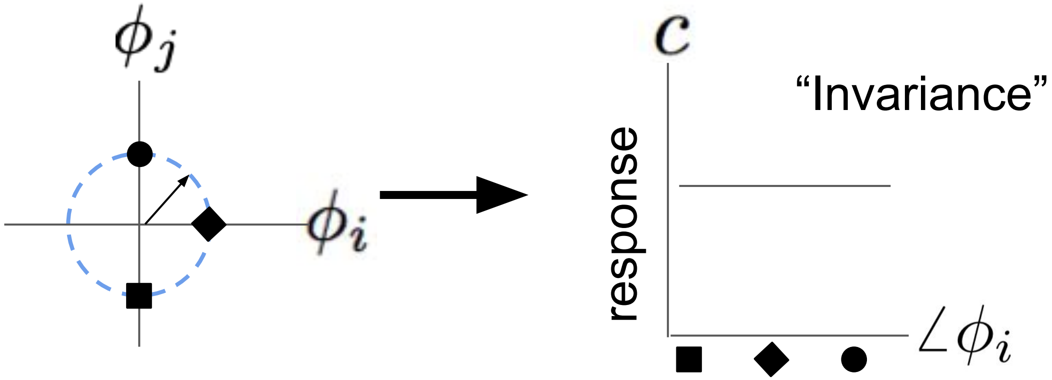

Subspace Independent Component Analysis (ISA)

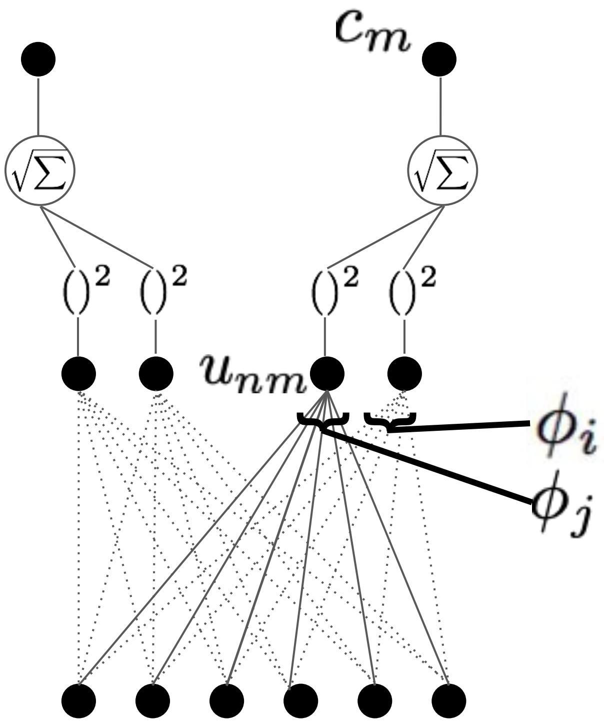

LCA neurons produce “exo-origin” curvature

L/NL neuron

ICA

$$a_k = f_\lambda(\Phi_k^\top s)$$

Flat

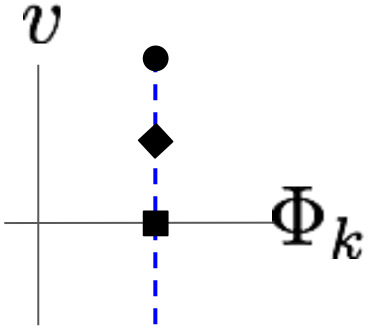

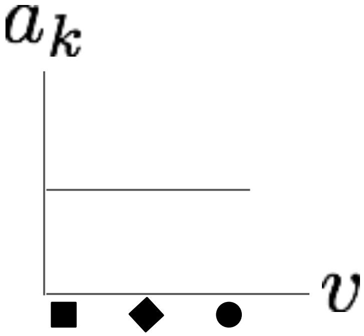

Exo-Origin

LCA

$$\tau \dot{u}_k + u_k = \Phi_k^\top s - G_k^\top a$$
$$a_k = f_\lambda(u_k)$$

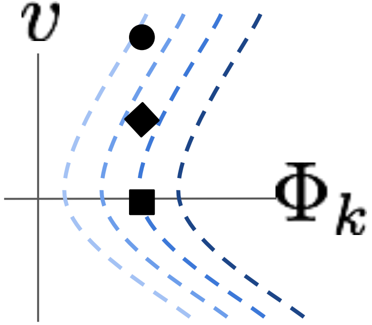

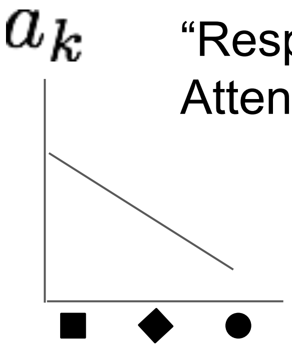

SLCA neurons exhibit both types of curvature

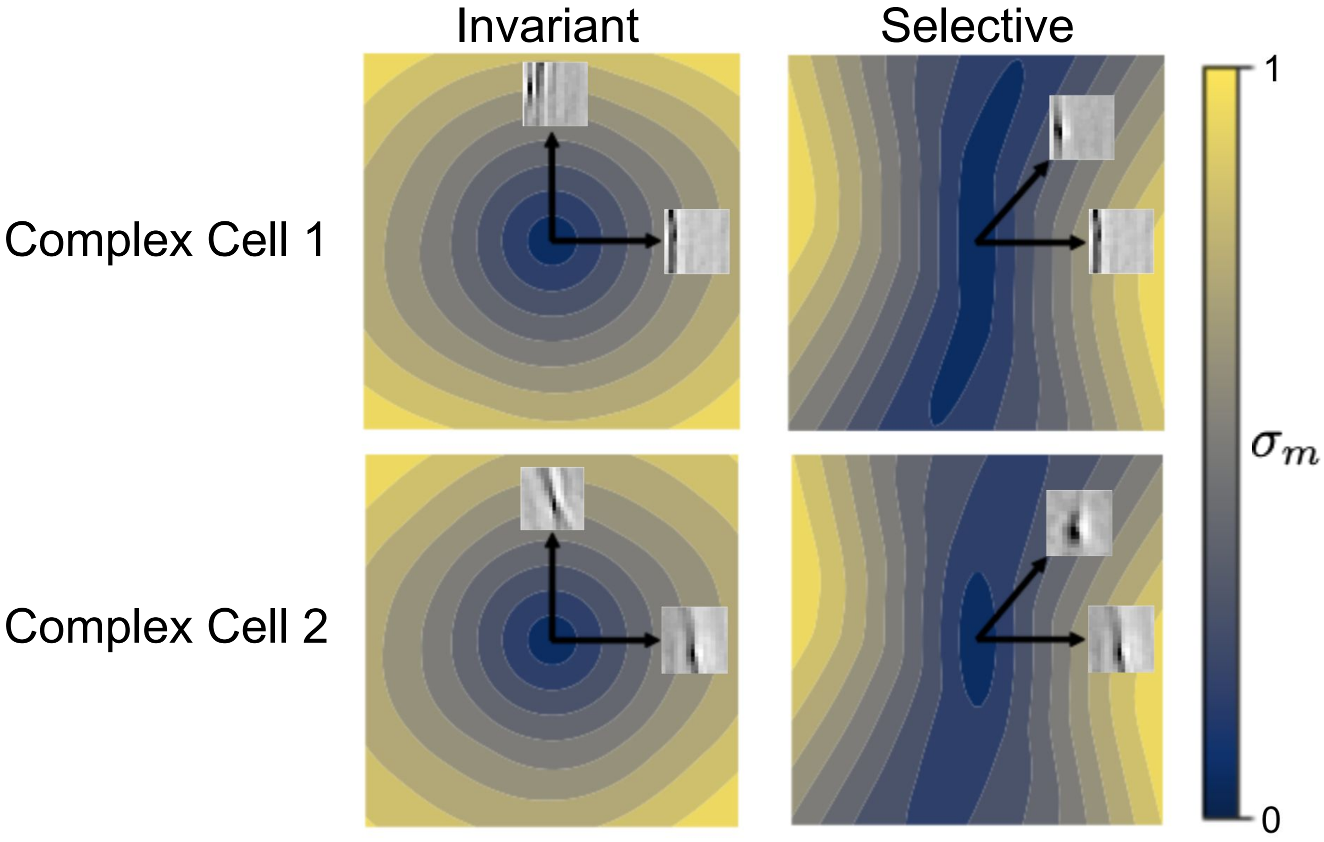

Network types

ICA

- Linear first layer

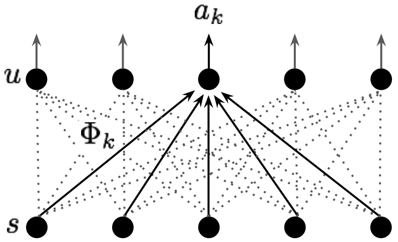

Subspace ICA (ISA)

- Linear first layer
- Energy second layer

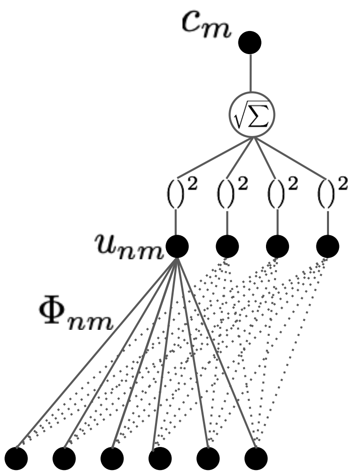

LCA

- Lateral interactions in first layer

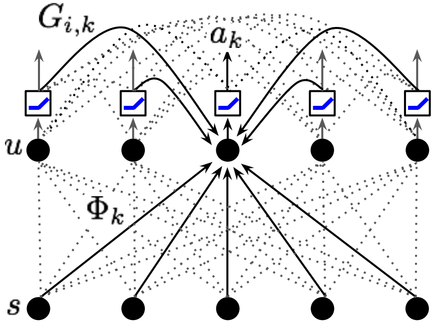

Subspace LCA

- Lateral interactions in first layer
- Energy second layer
- Lateral interactions second layer

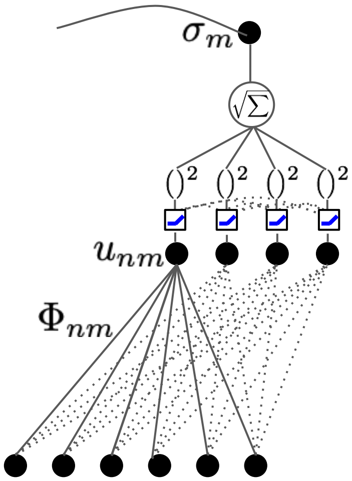

SLCA has selective neurons **and** invariant neurons

ICA

- minimal selectivity
- no invariance

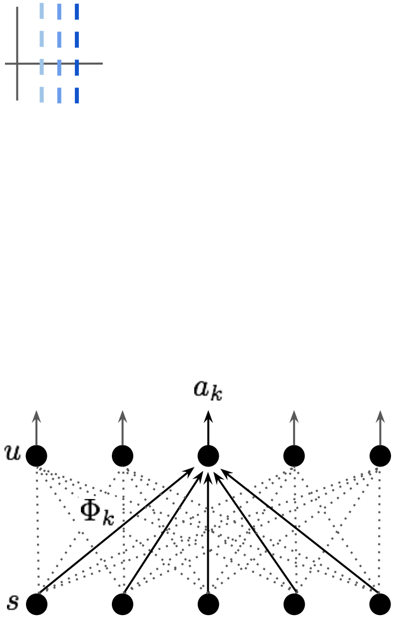

Subspace ICA (ISA)

- minimal selectivity
- invariance

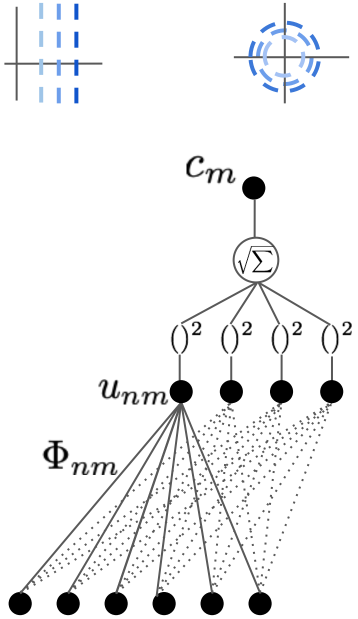

LCA

- high selectivity
- minimal invariance

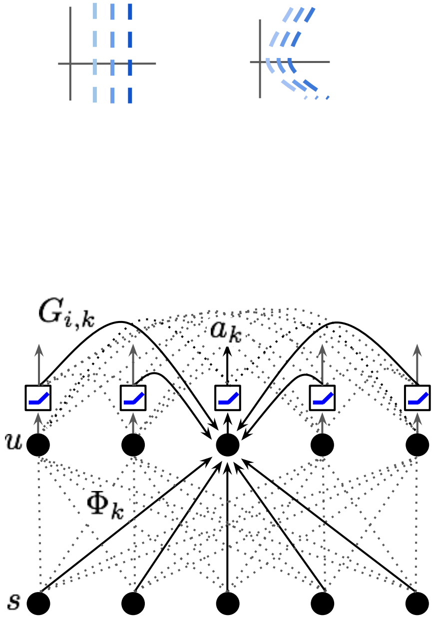

Subspace LCA

- high selectivity
- high invariance

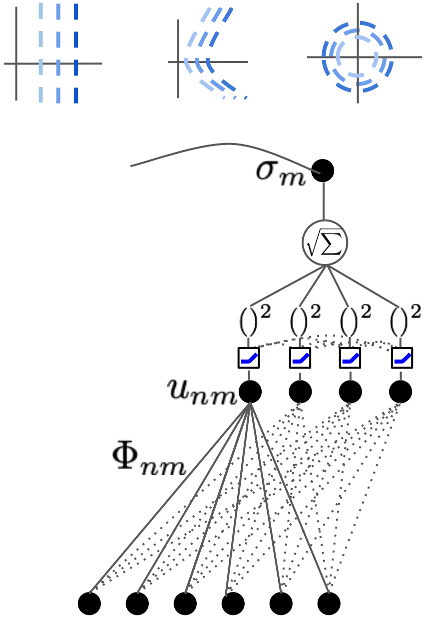

SLCA produces more stable video representations

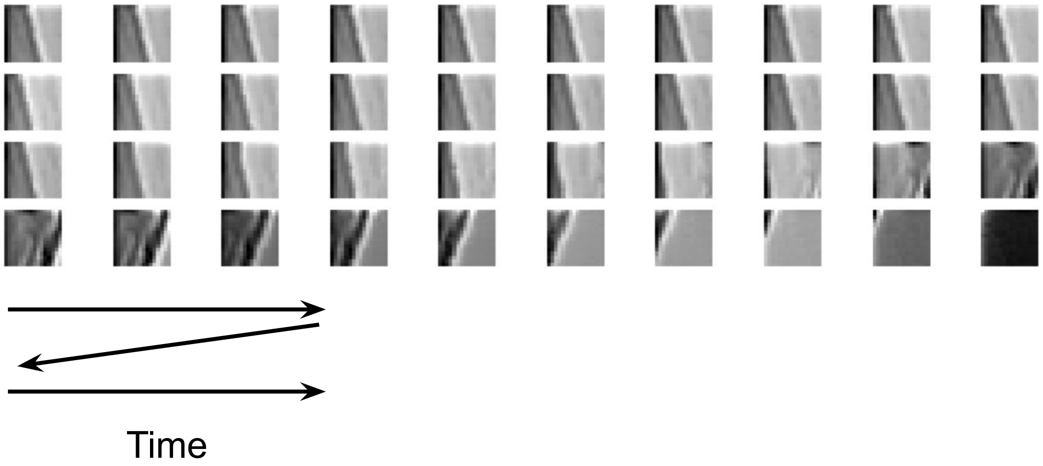

SLCA produces more stable video representations

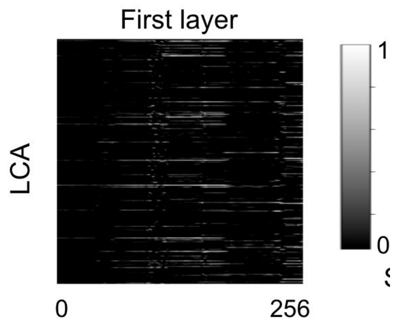

Time
(frame)

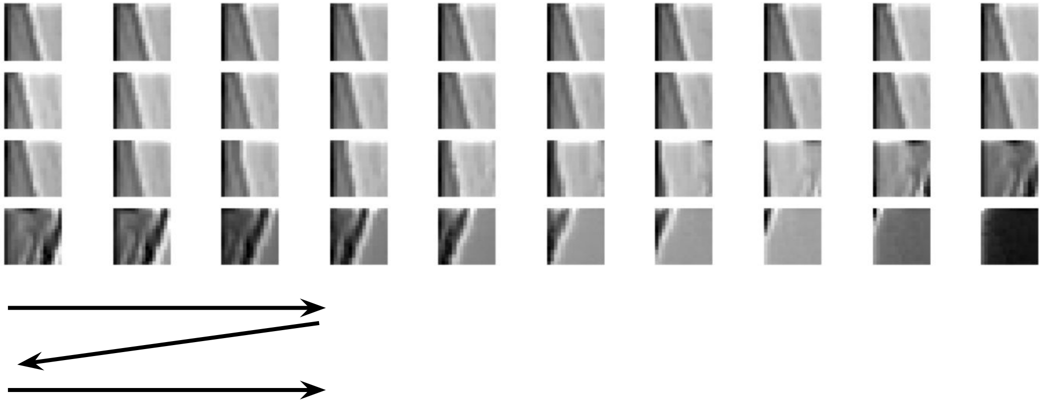

Time

SLCA produces more stable video representations

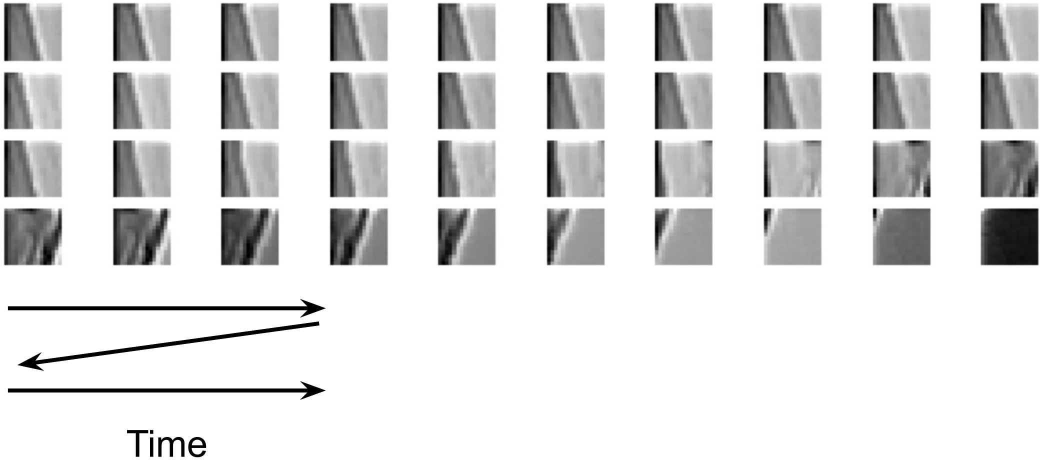

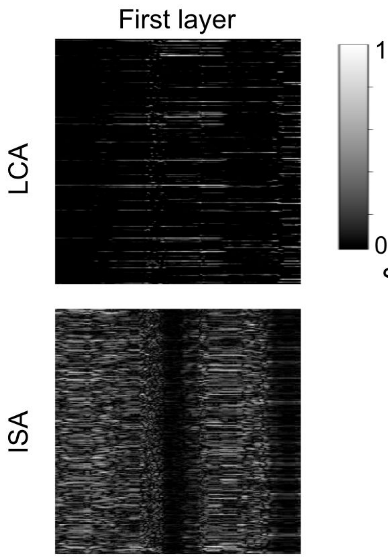

SLCA produces more stable video representations

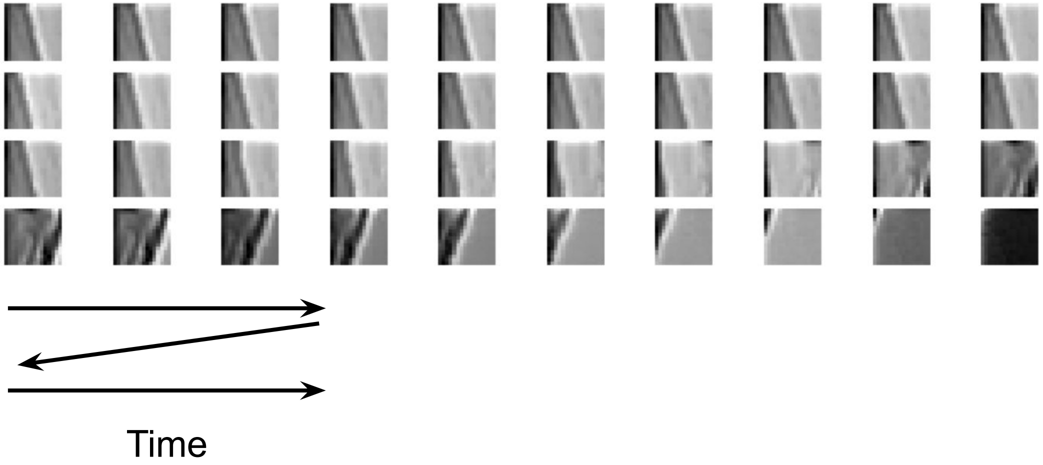

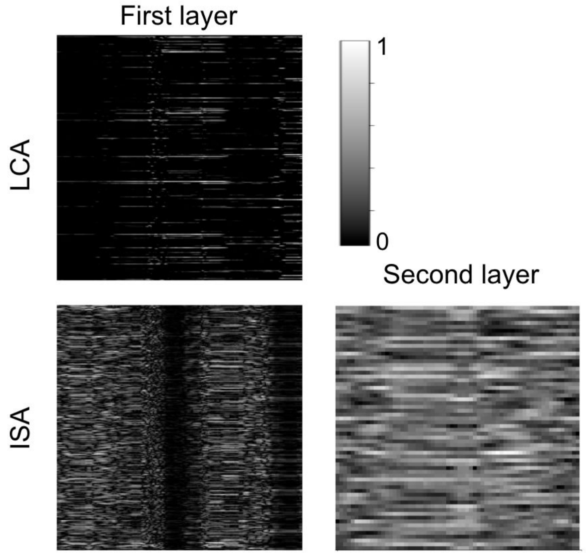

SLCA produces more stable video representations

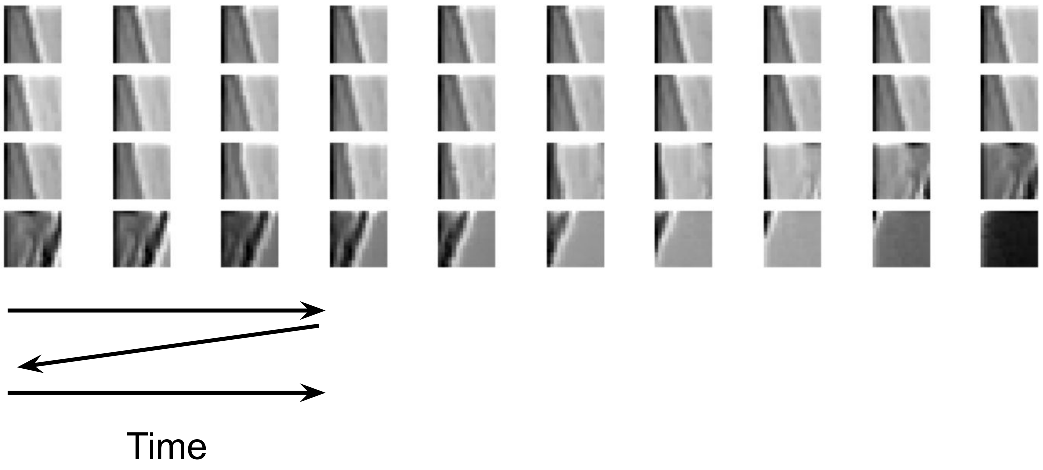

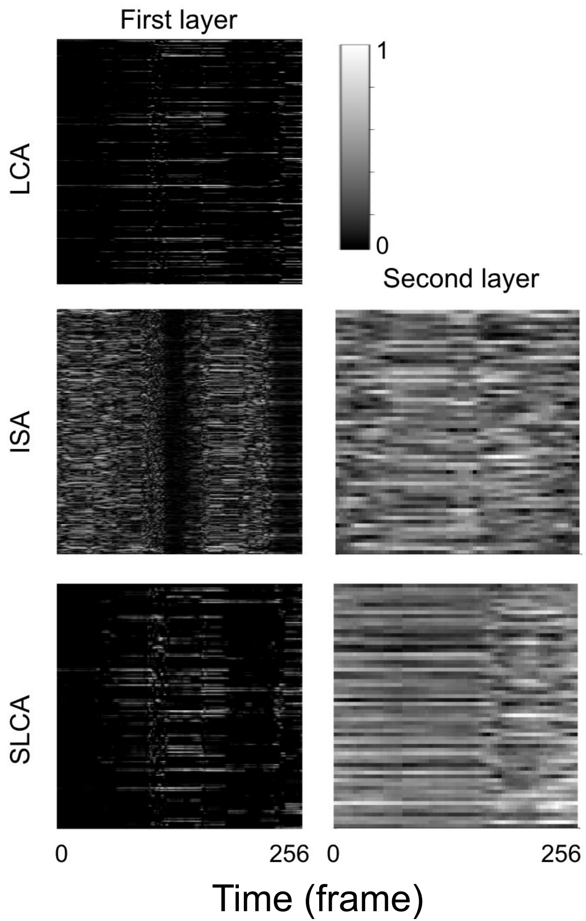

Please read our paper for more!

Authors:

Dylan Paiton | Steven Shepard | Kwan Ho “Ryan” Chan | Bruno Olshausen



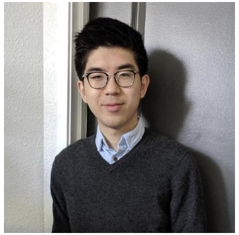


@dylanpaiton
dpaiton@gmail.com

Funding:
ENIGMA NSF research program; Berkeley AI Research group