Learning Spatial and Transformational Invariants for Visual Representation

Charles Cadieu
work with Bruno Olshausen

CIfAR Summer School
Toronto

REDWOOD CENTER
for Theoretical Neuroscience

Visual Representation

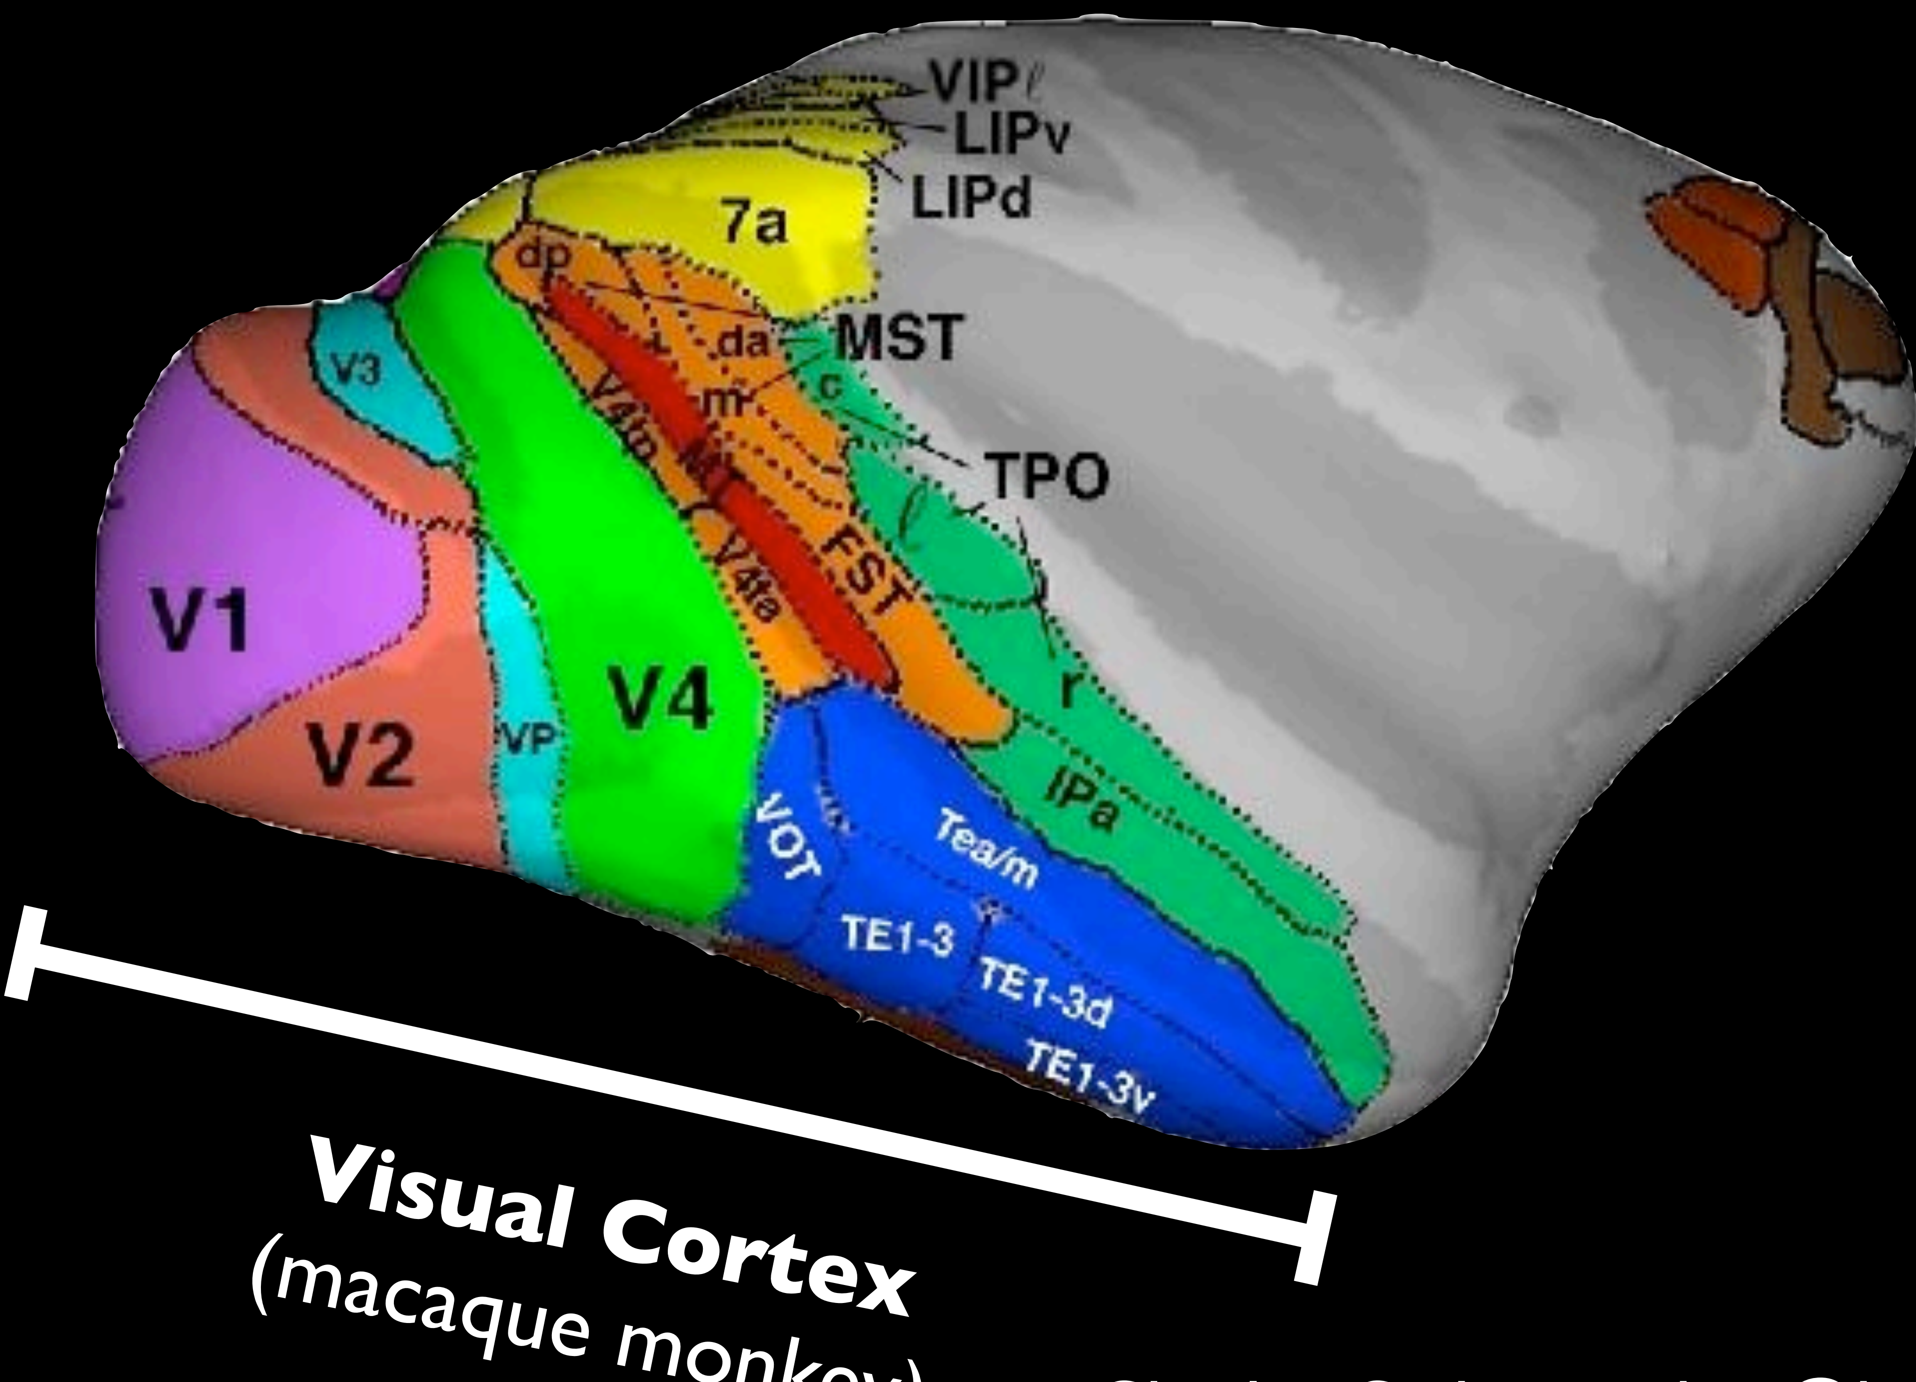

Charles Cadieu, cadieu@berkeley.edu

The Visual System Infers the Causes of Images

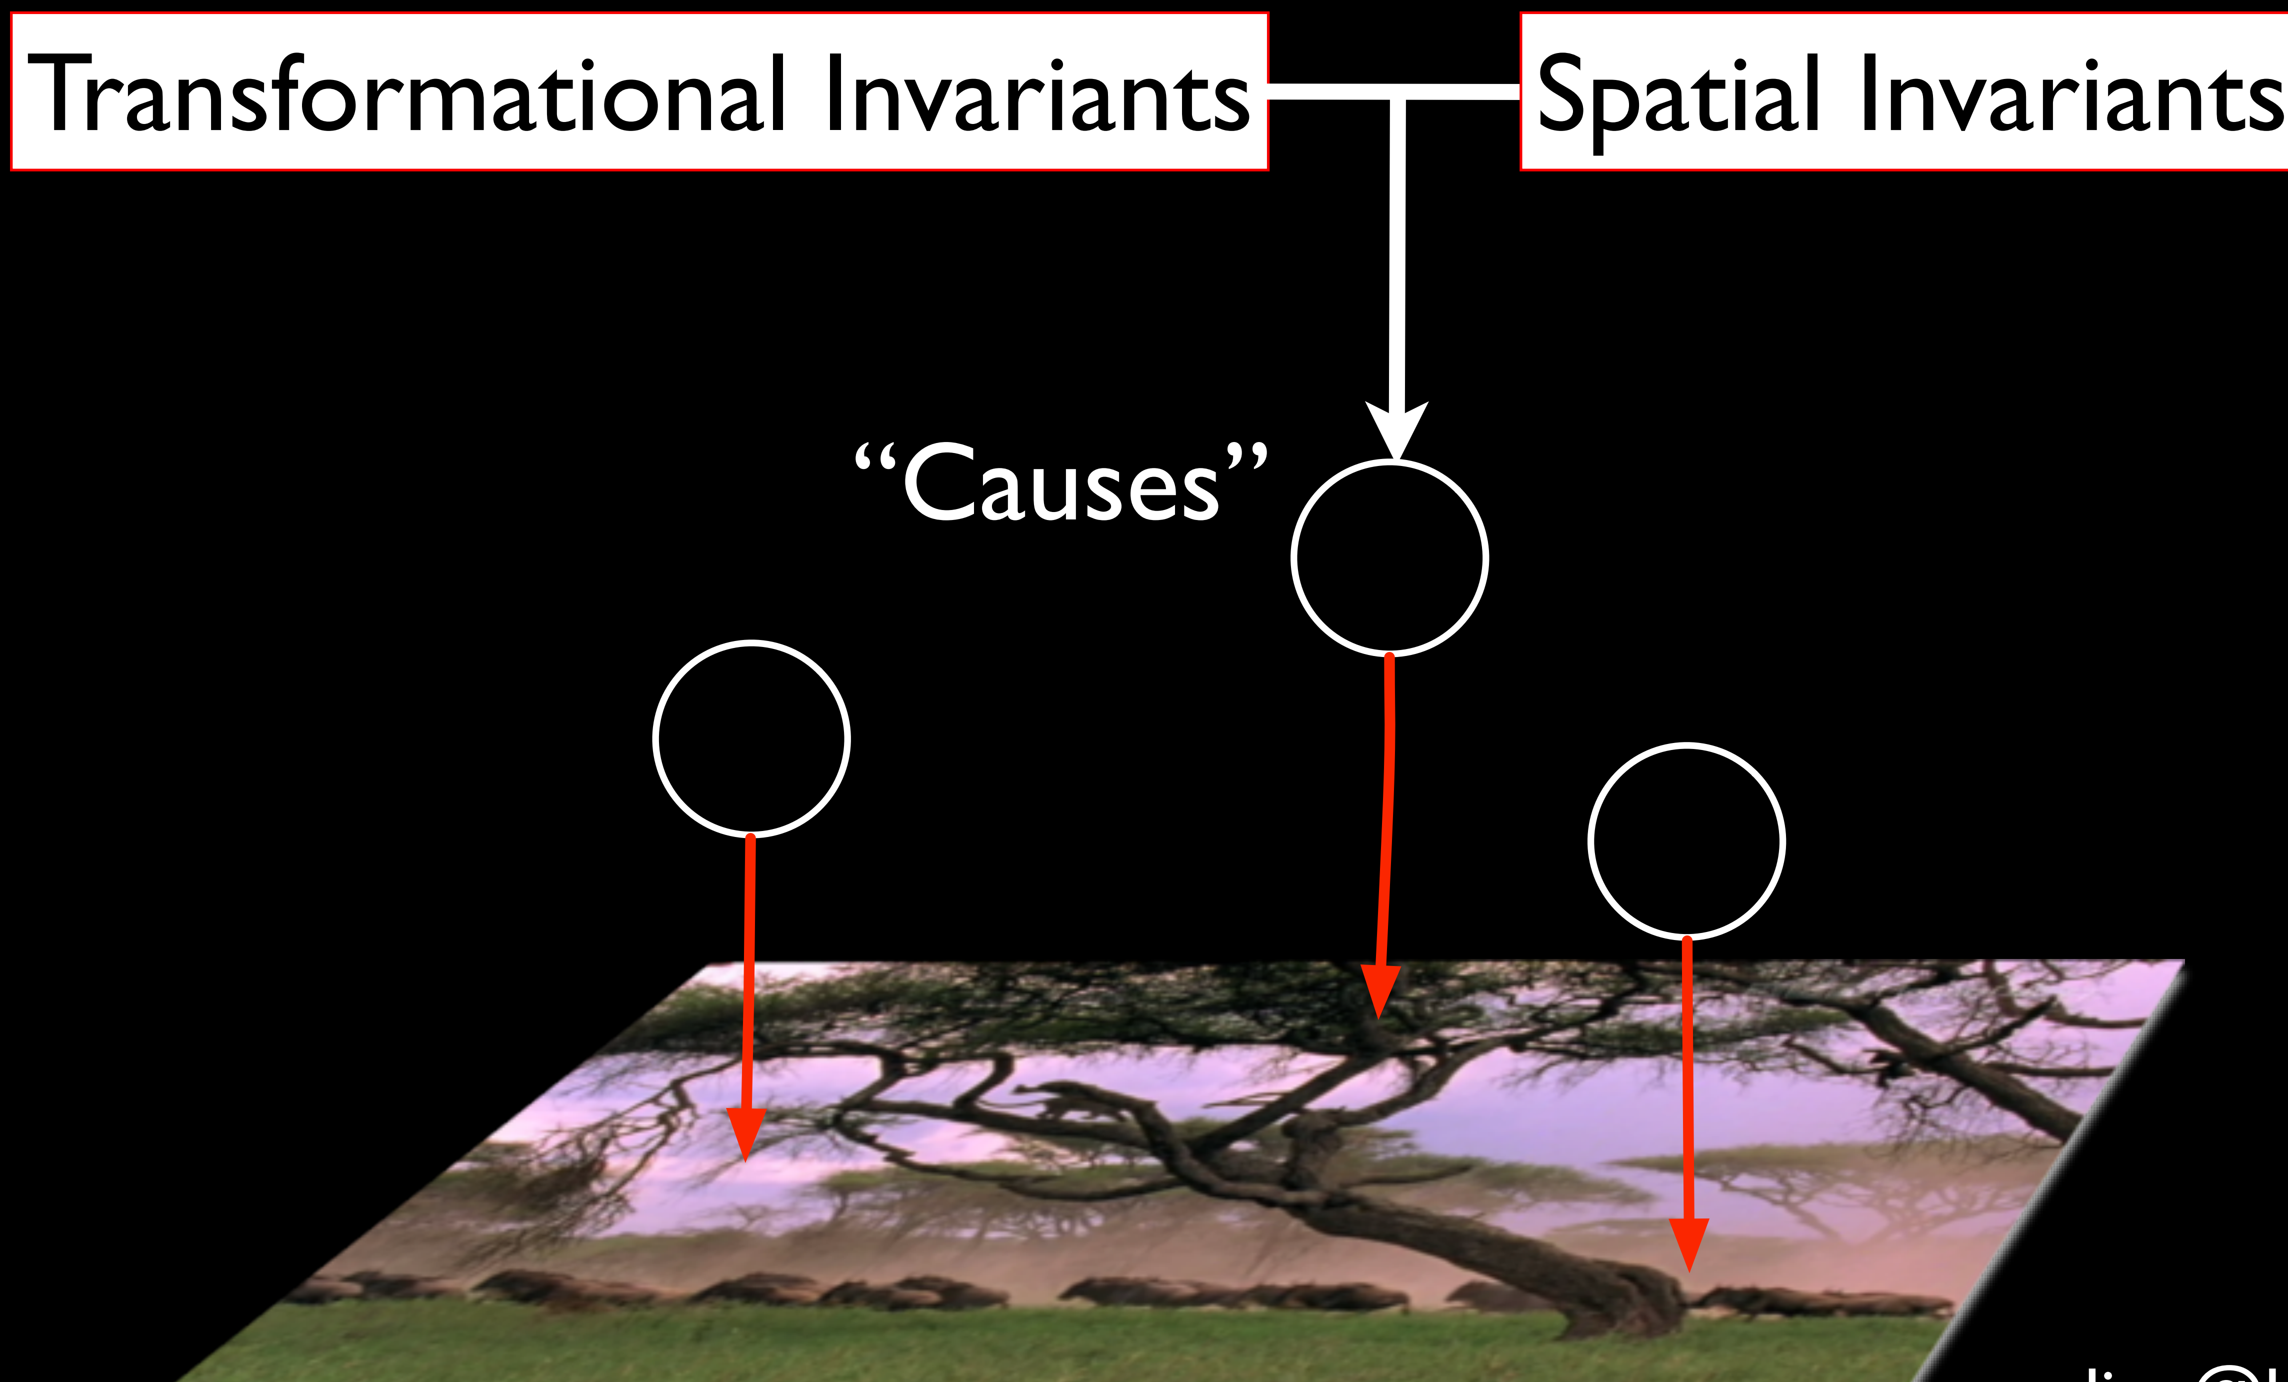

Charles Cadieu, cadieu@berkeley.edu

“Leopard” -
Spatial Invariant

“Galloping” -
Transformational Invariant

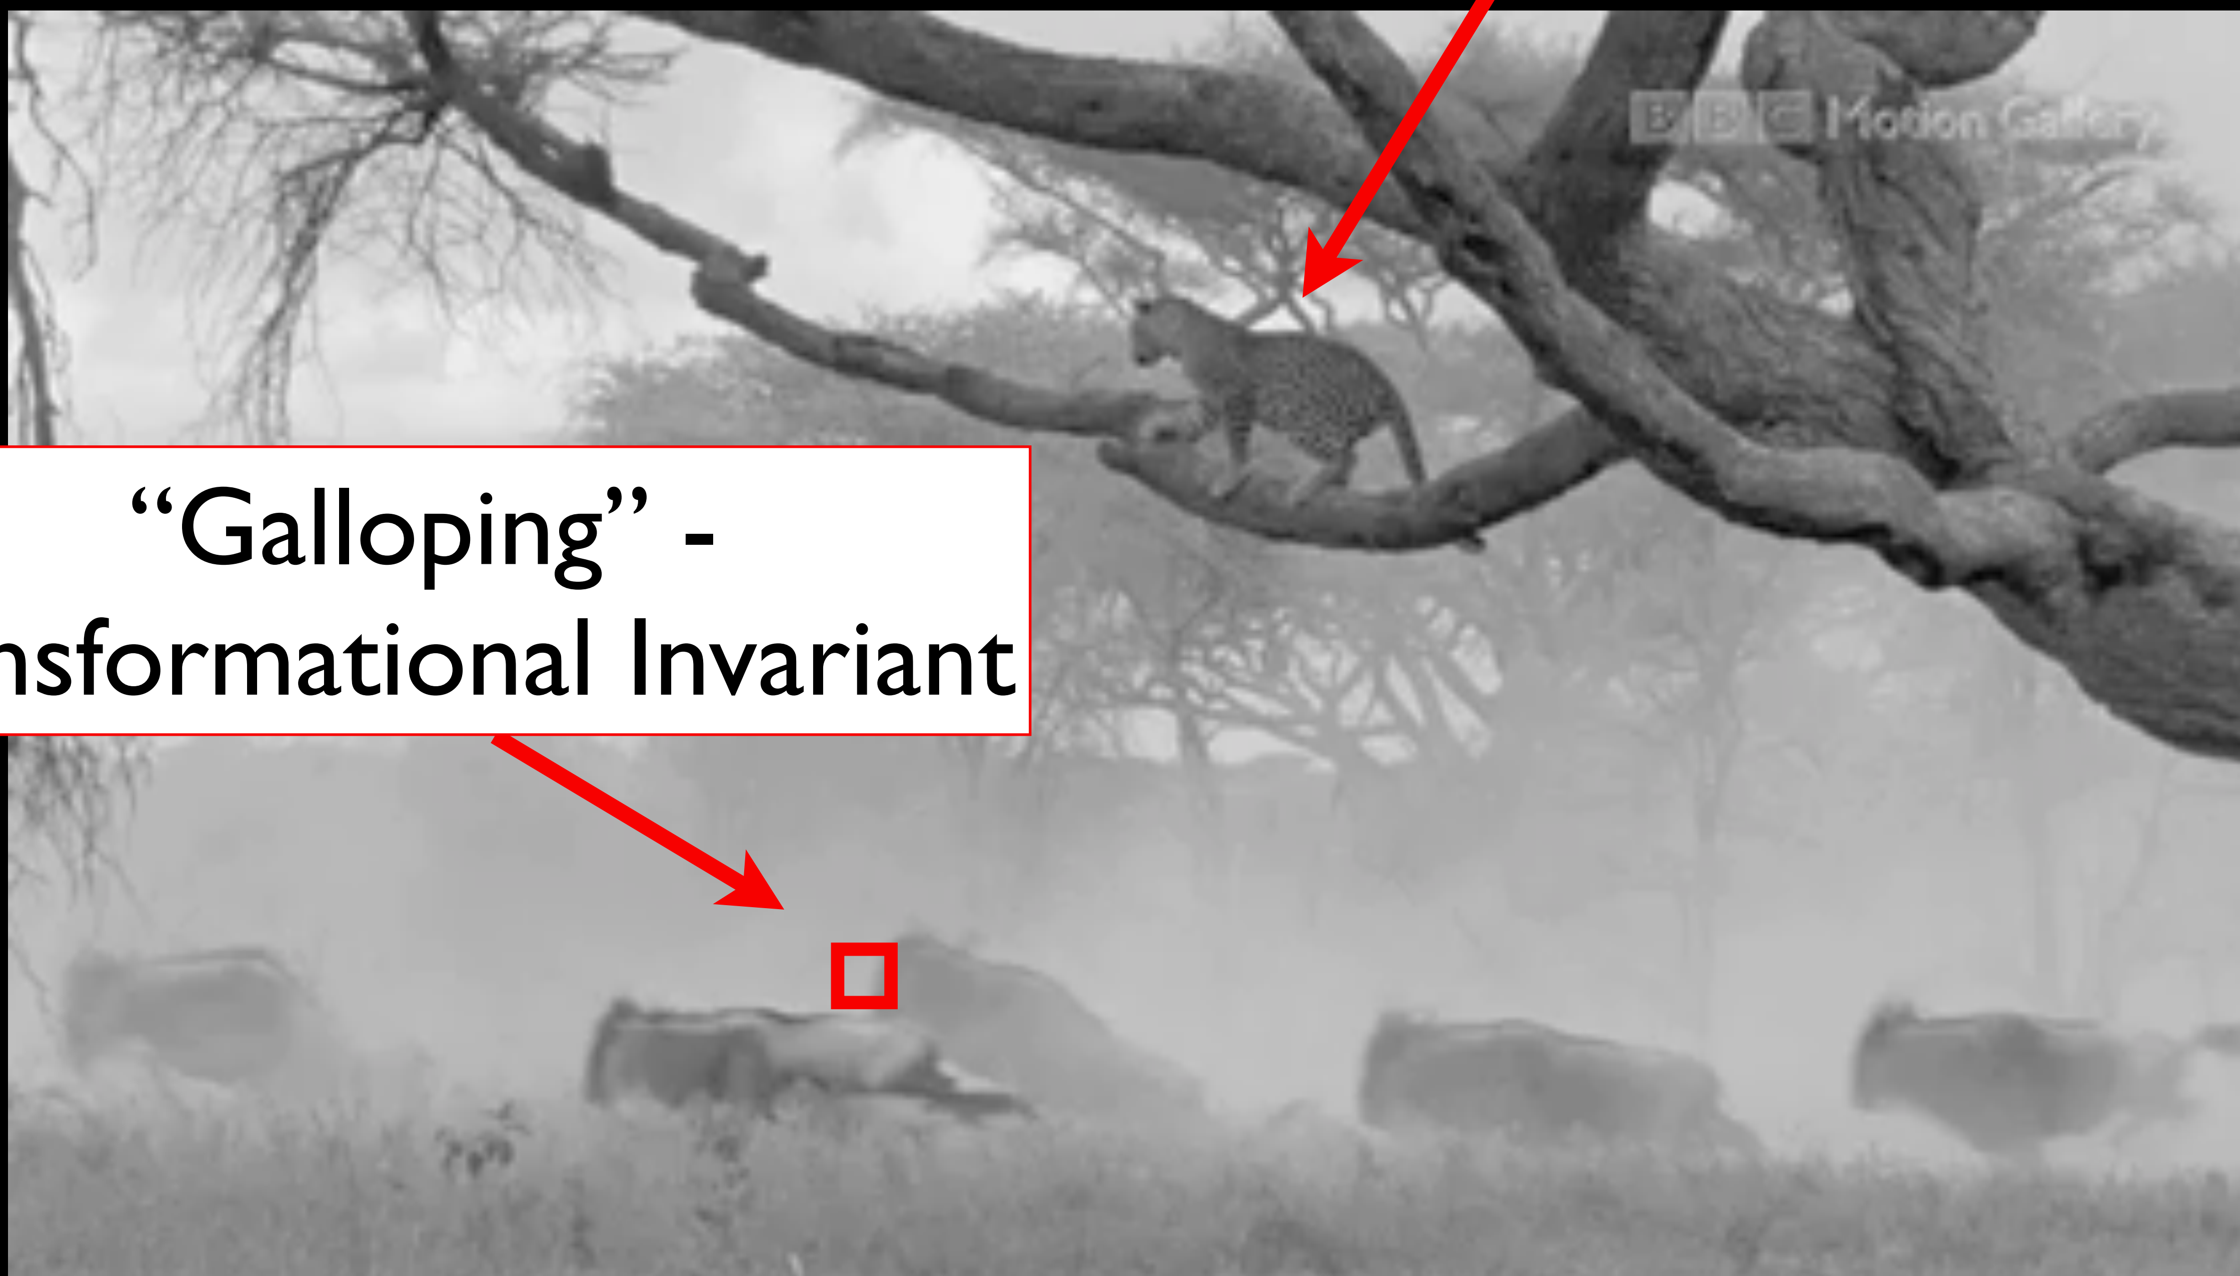

Charles Cadieu, cadieu@berkeley.edu

Structure within an Image Patch

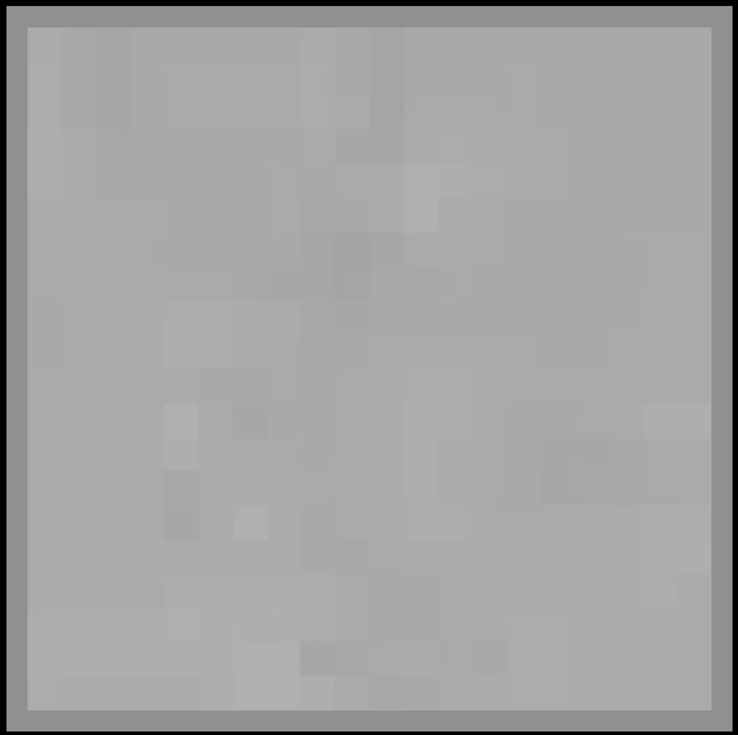

- How do we uncover the causes of this complicated data?

Charles Cadieu, cadieu@berkeley.edu

$$I(x,t) = \sum_i A_i(x) u_i(t) + \nu(x,t)$$

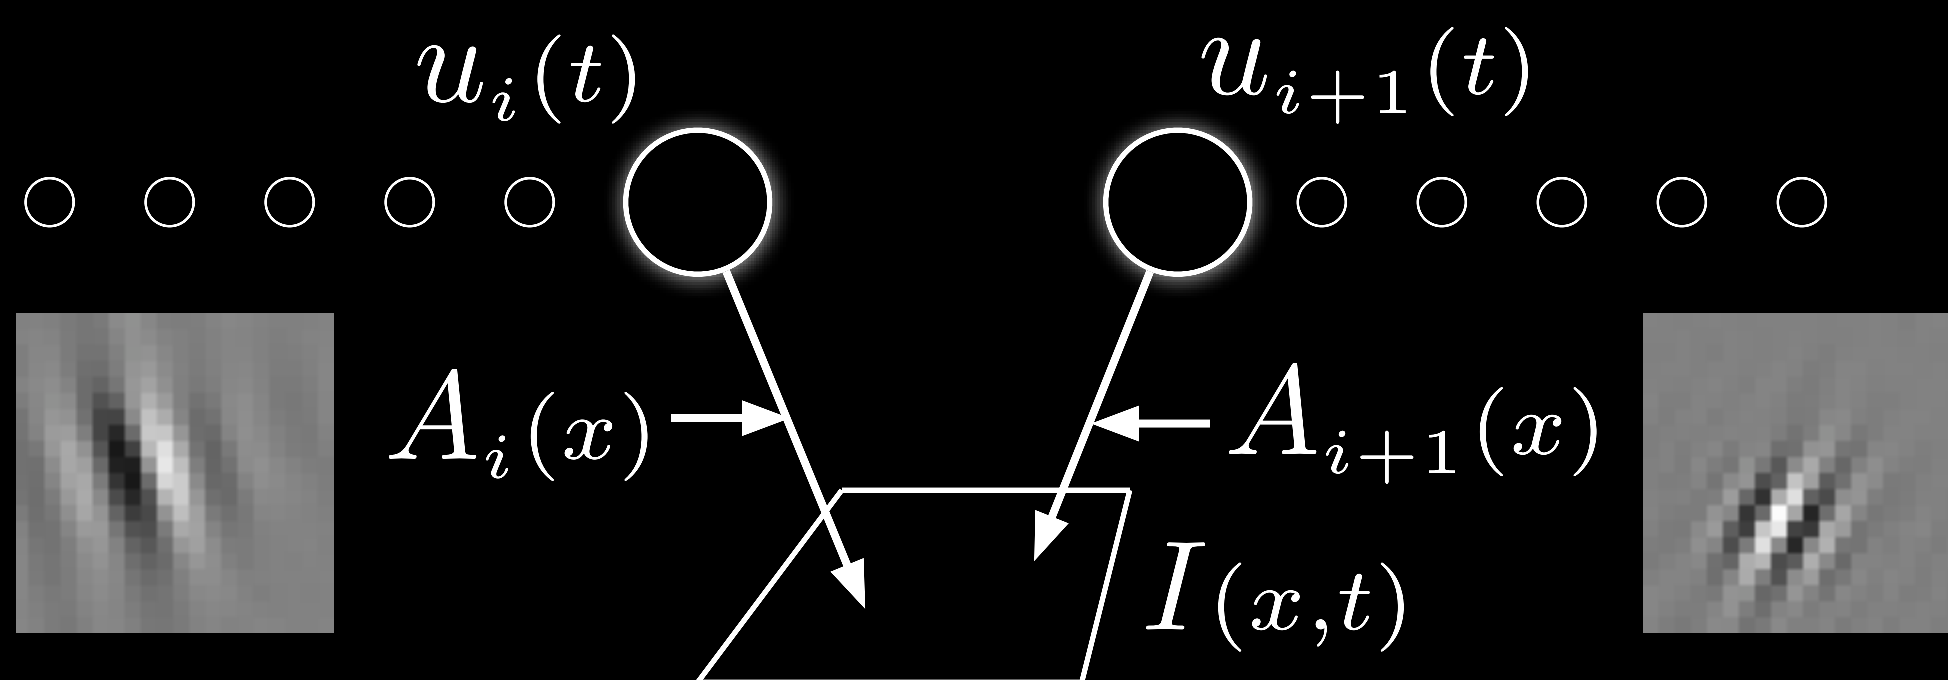

Charles Cadieu, cadieu@berkeley.edu

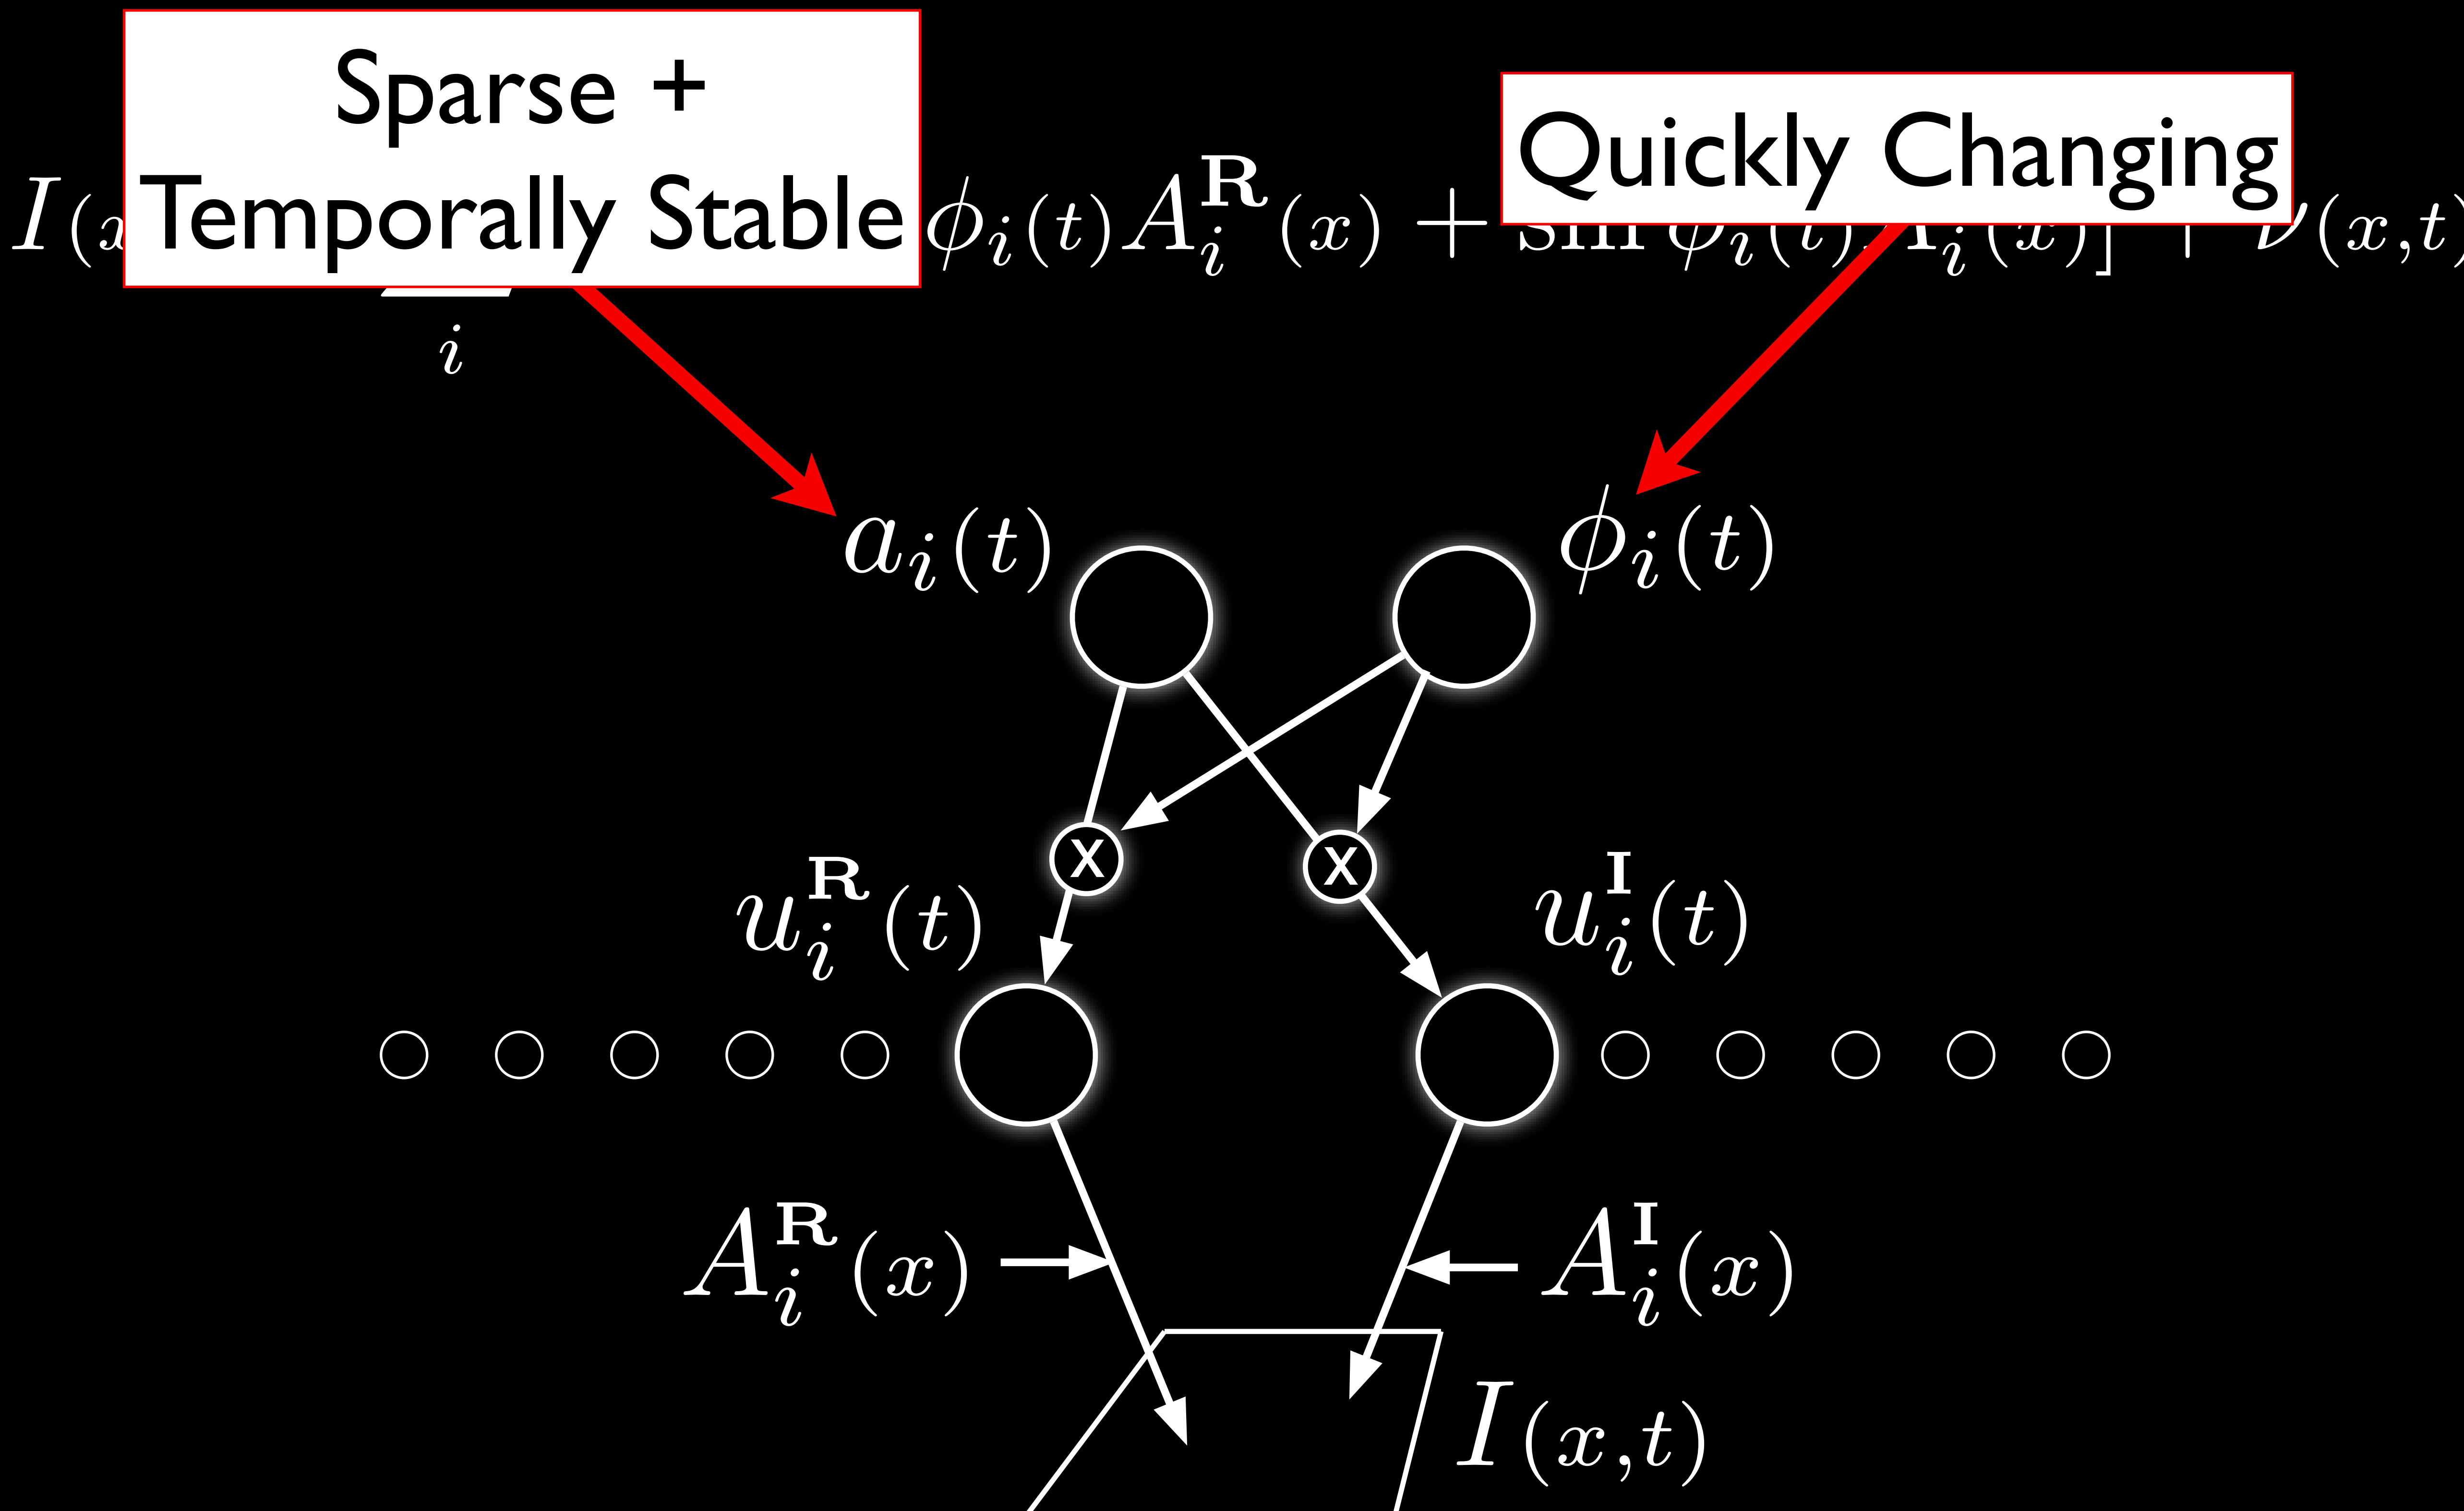

Charles Cadieu, cadieu@berkeley.edu

First Layer Probabilistic Model

$$P(I, a, \phi) \propto e^{-E_1}$$

Reconstruction Error | Sparse | Slow

$$E_1 = \sum_t \sum_x \frac{1}{\sigma_N^2} \left[I(x,t) - \sum_i \Re\{ z_i^*(t)\, A_i(x) \} \right]^2 + \lambda_{Sp} \sum_{i,t} a_i(t) + \lambda_{Sl} \sum_{i,t} (a_i(t) - a_i(t-1))^2$$

Adapt to Natural Movies

Charles Cadieu, cadieu@berkeley.edu

Learned Basis Functions

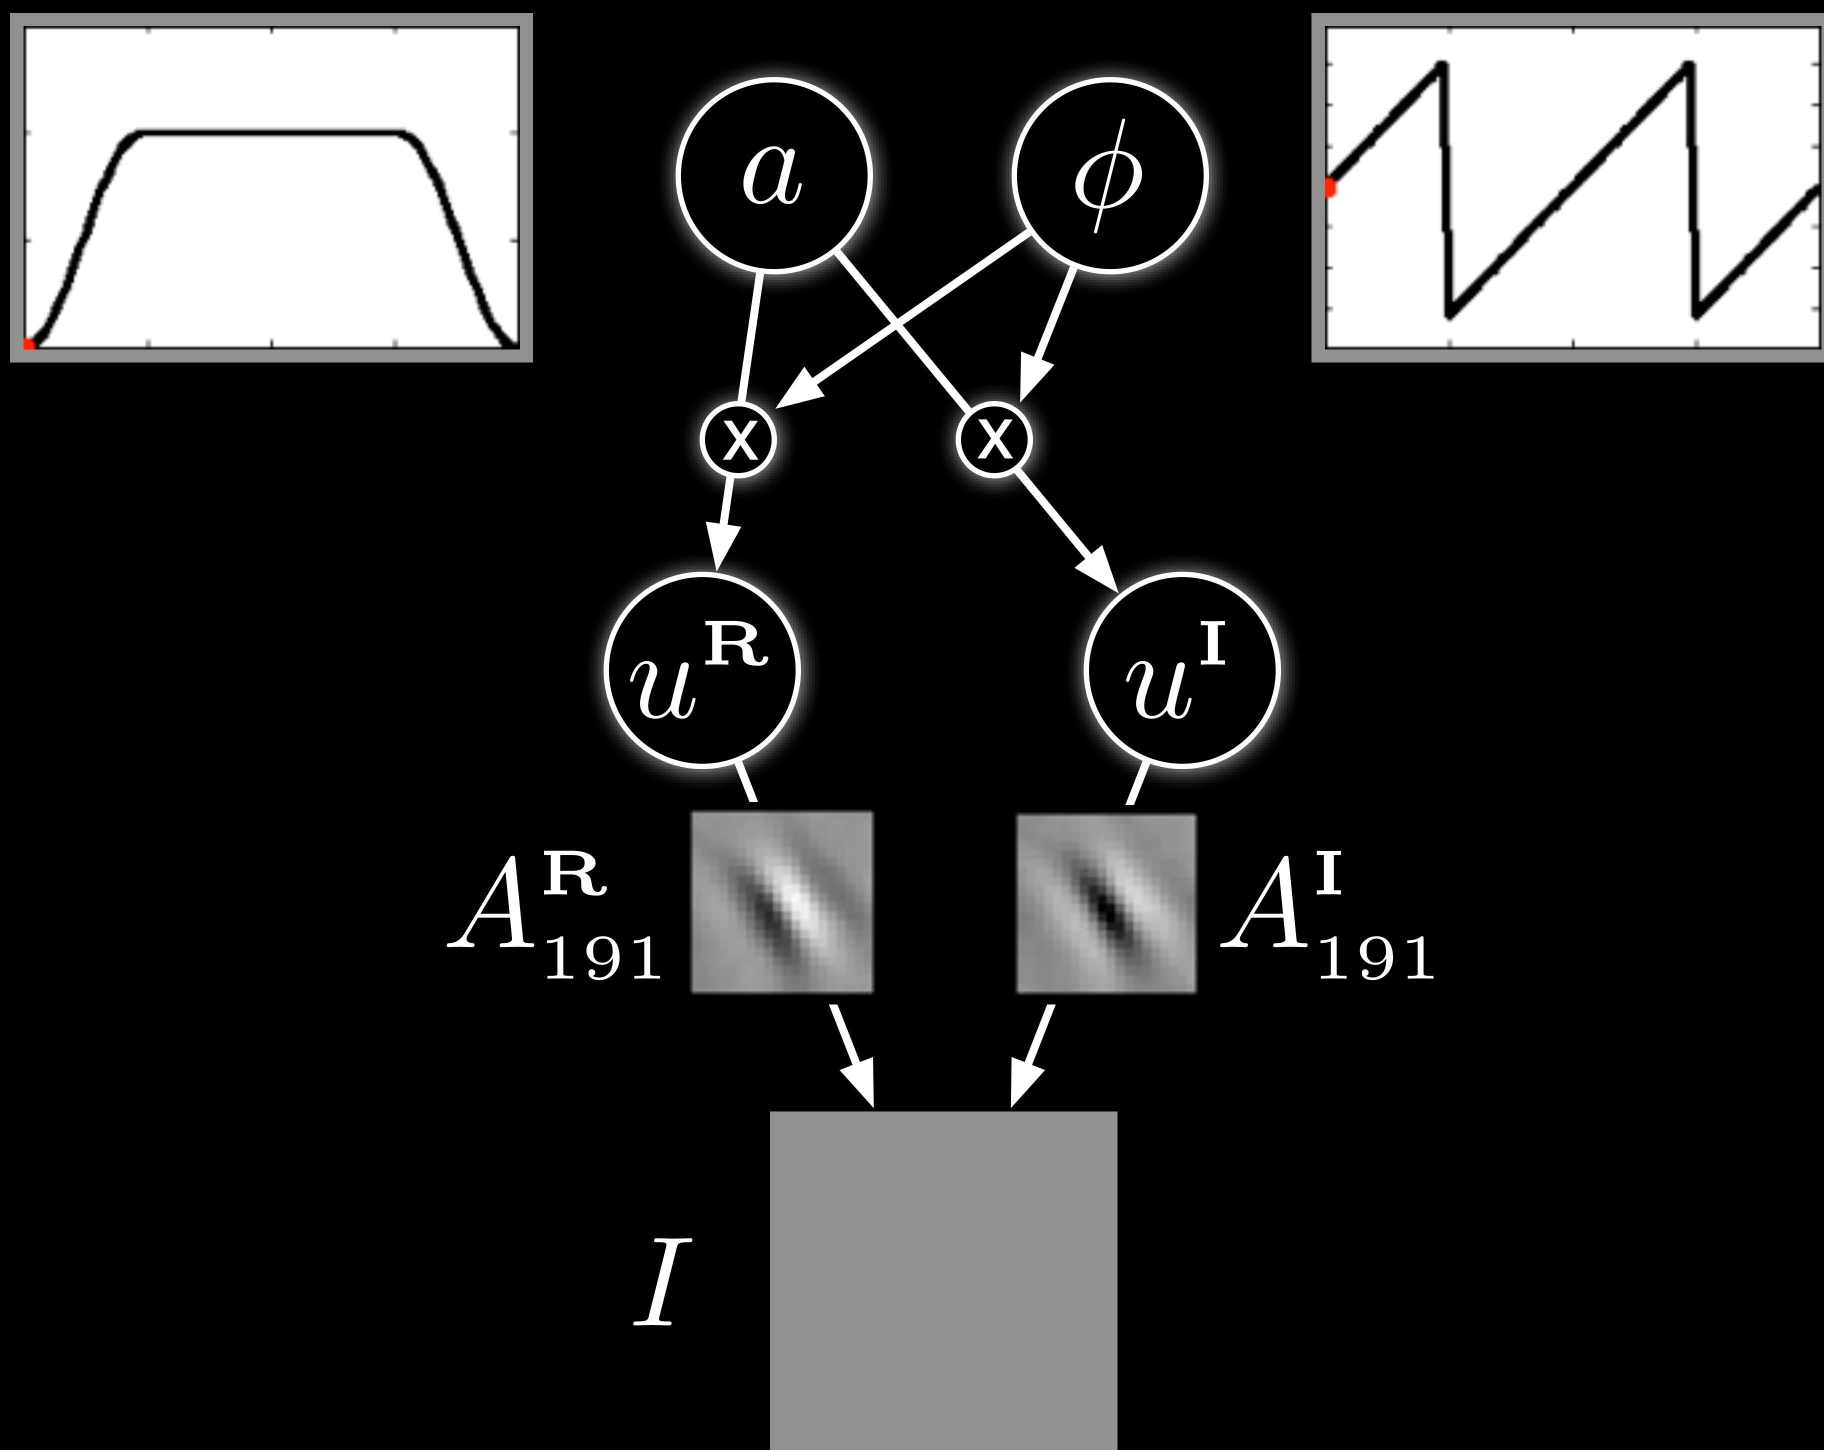

Charles Cadieu, cadieu@berkeley.edu

First Layer Basis Functions

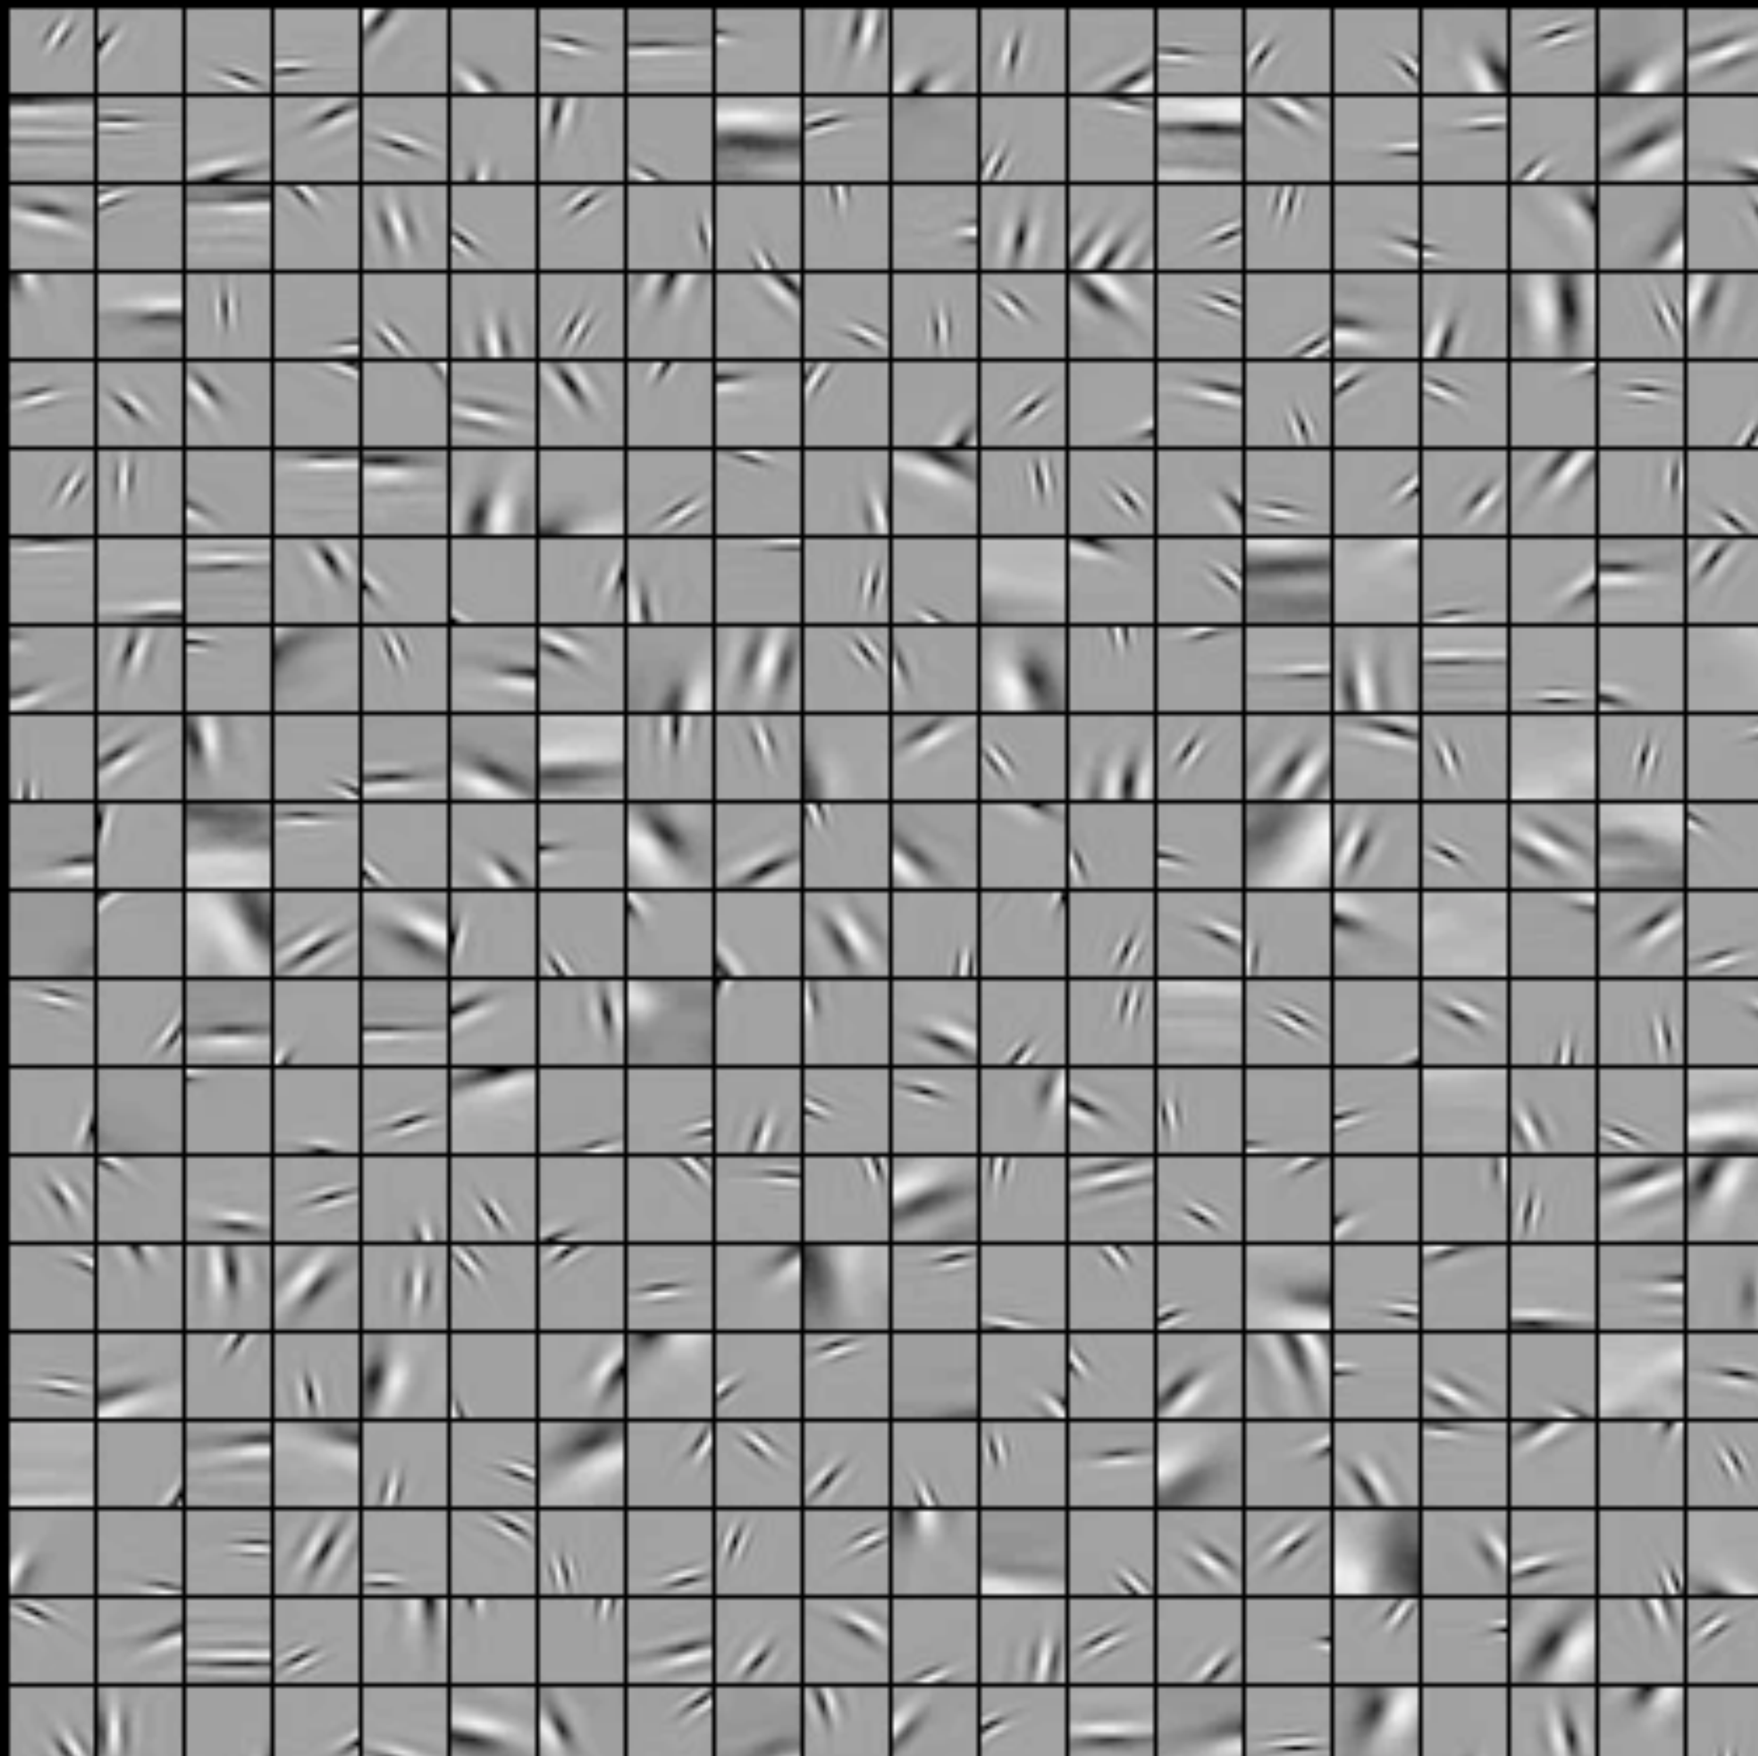

Charles Cadieu, cadieu@berkeley.edu

Motions Produce Patterns in Phase

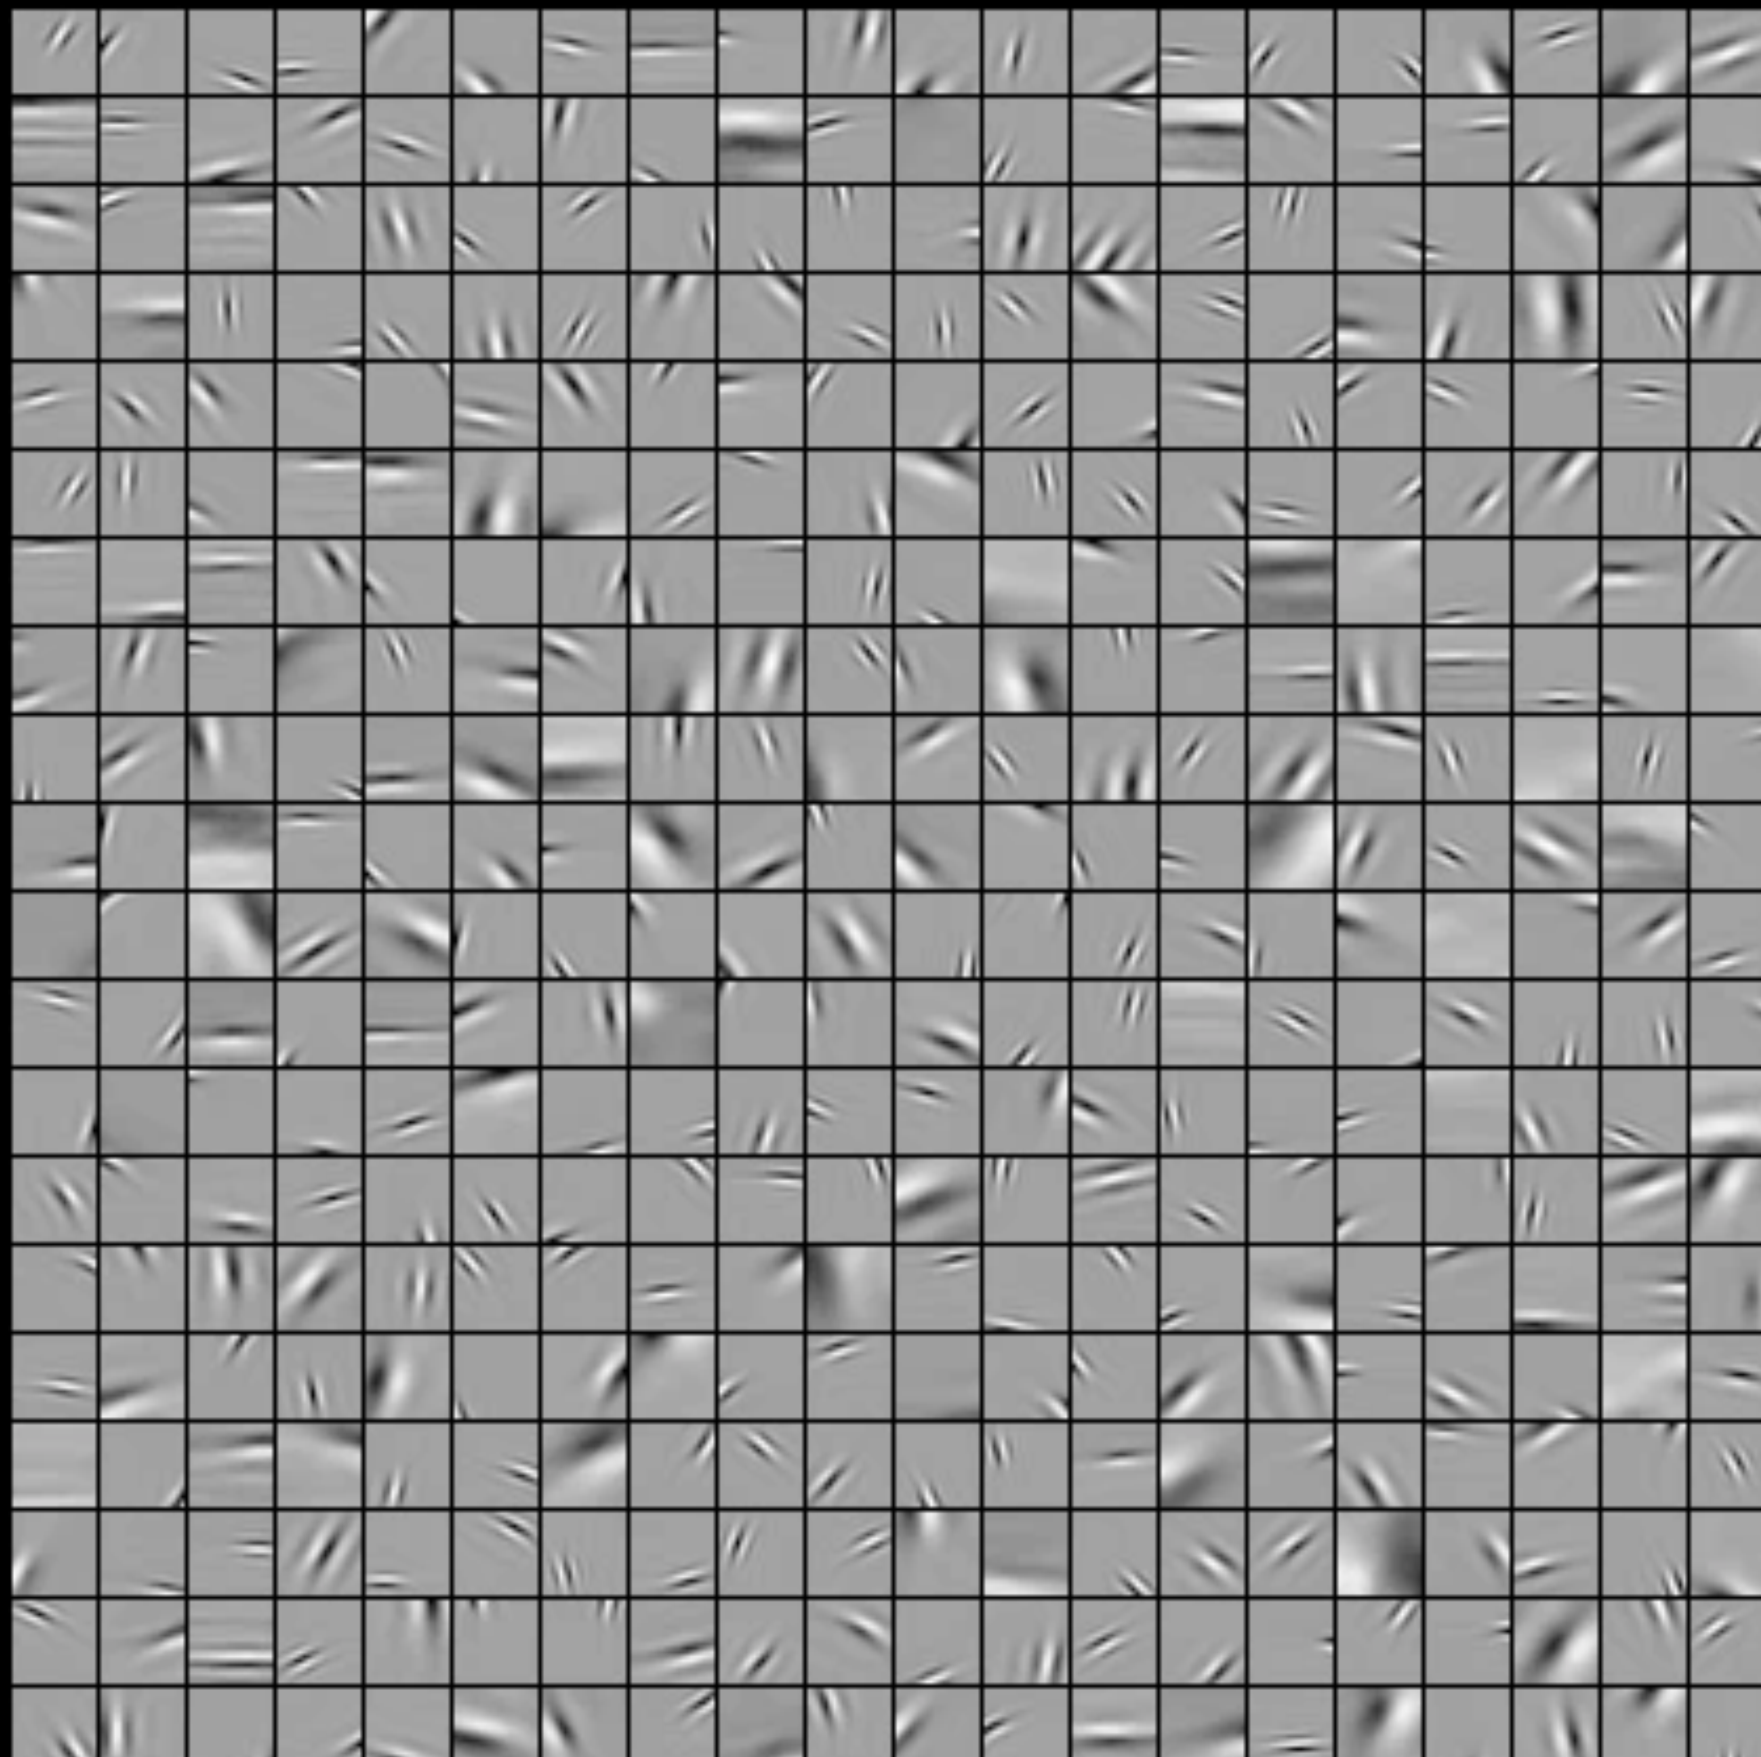

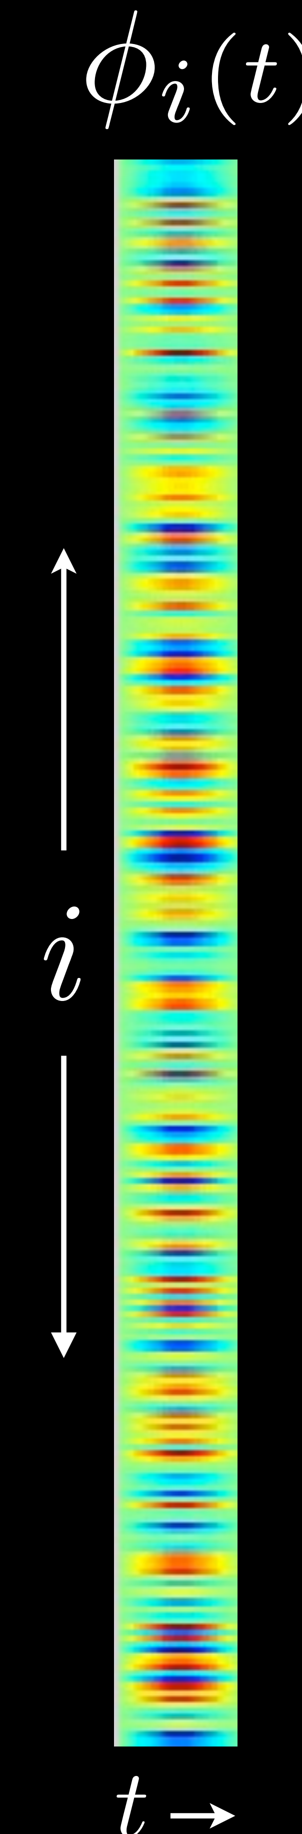

Charles Cadieu, cadieu@berkeley.edu

Model the changes in phase with a sparse, latent variable model

$$\phi_i(t) - \phi_i(t-1) = \sum_k D_{ik} w_k(t) + \delta_i(t)$$

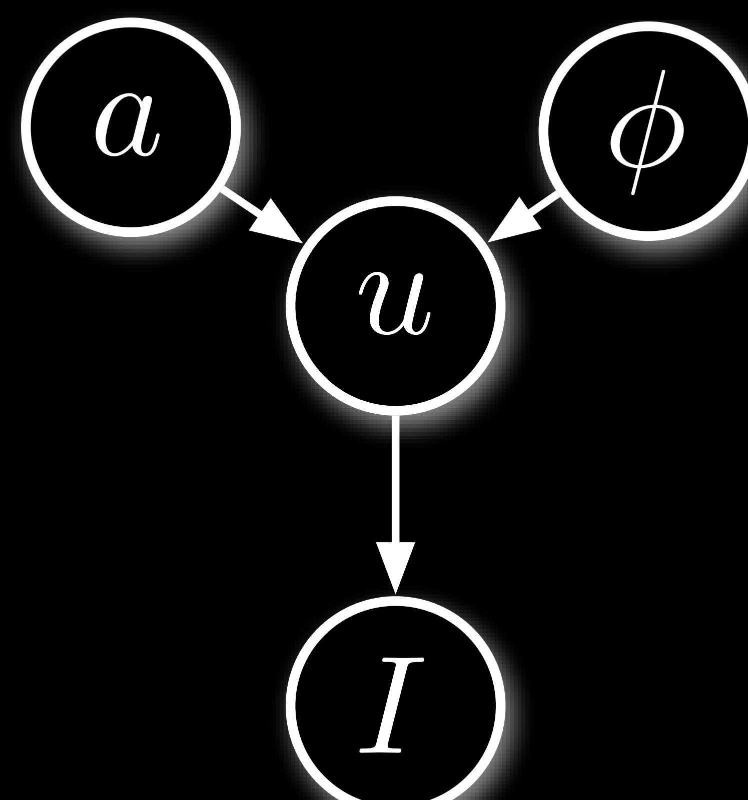

ı, cadieu@berkeley.edu

Second Layer Probabilistic Model: Transformational Invariants

$$P(I, a, \phi, w) \propto e^{-E_1 - E_2}$$

Error in Phase Dynamics — Sparse — Slow

$$E_2 = -\sum_t \sum_{i \in \{a_i(t) > 0\}} \kappa \cos(\dot{\phi}_i - [Dw(t)]_i) + \beta_{Sp} \sum_{k,t} |w_k(t)| + \beta_{Sl} \sum_{k,t} (w_k(t) - w_k(t-1))^2$$

Adapt to Natural Movies

Charles Cadieu, cadieu@berkeley.edu

Visualizing Learned Weights

D_{10}

Spatial Position

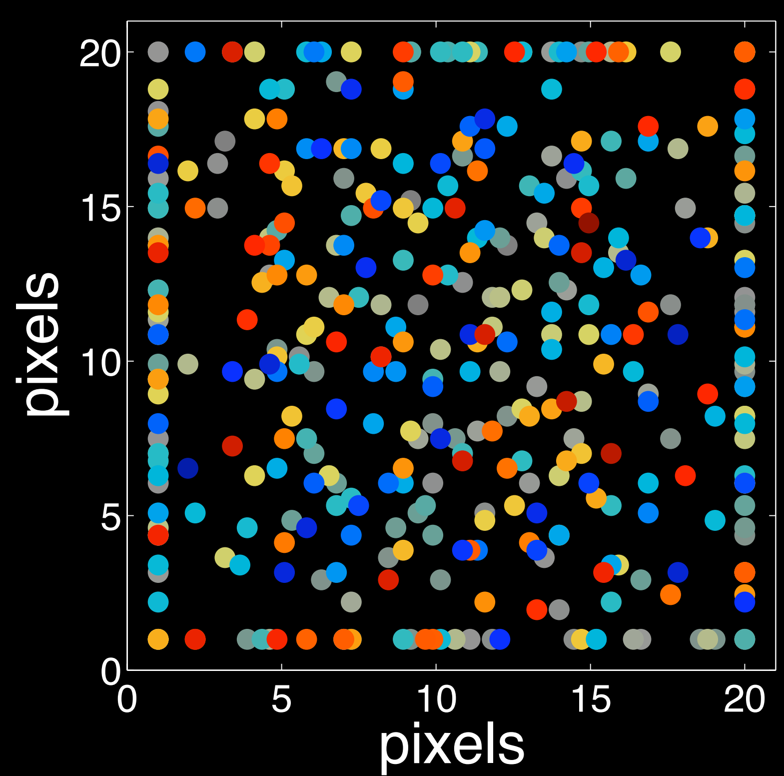

Spatial Frequency

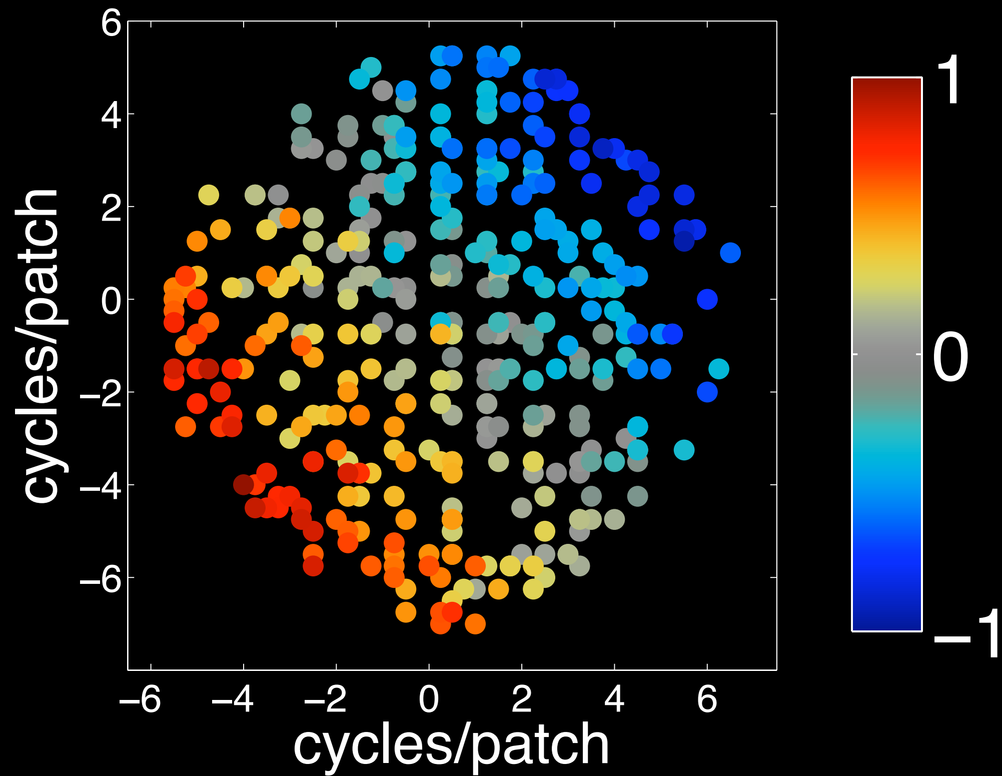

Charles Cadieu, cadieu@berkeley.edu

Transformation Components in Image Space

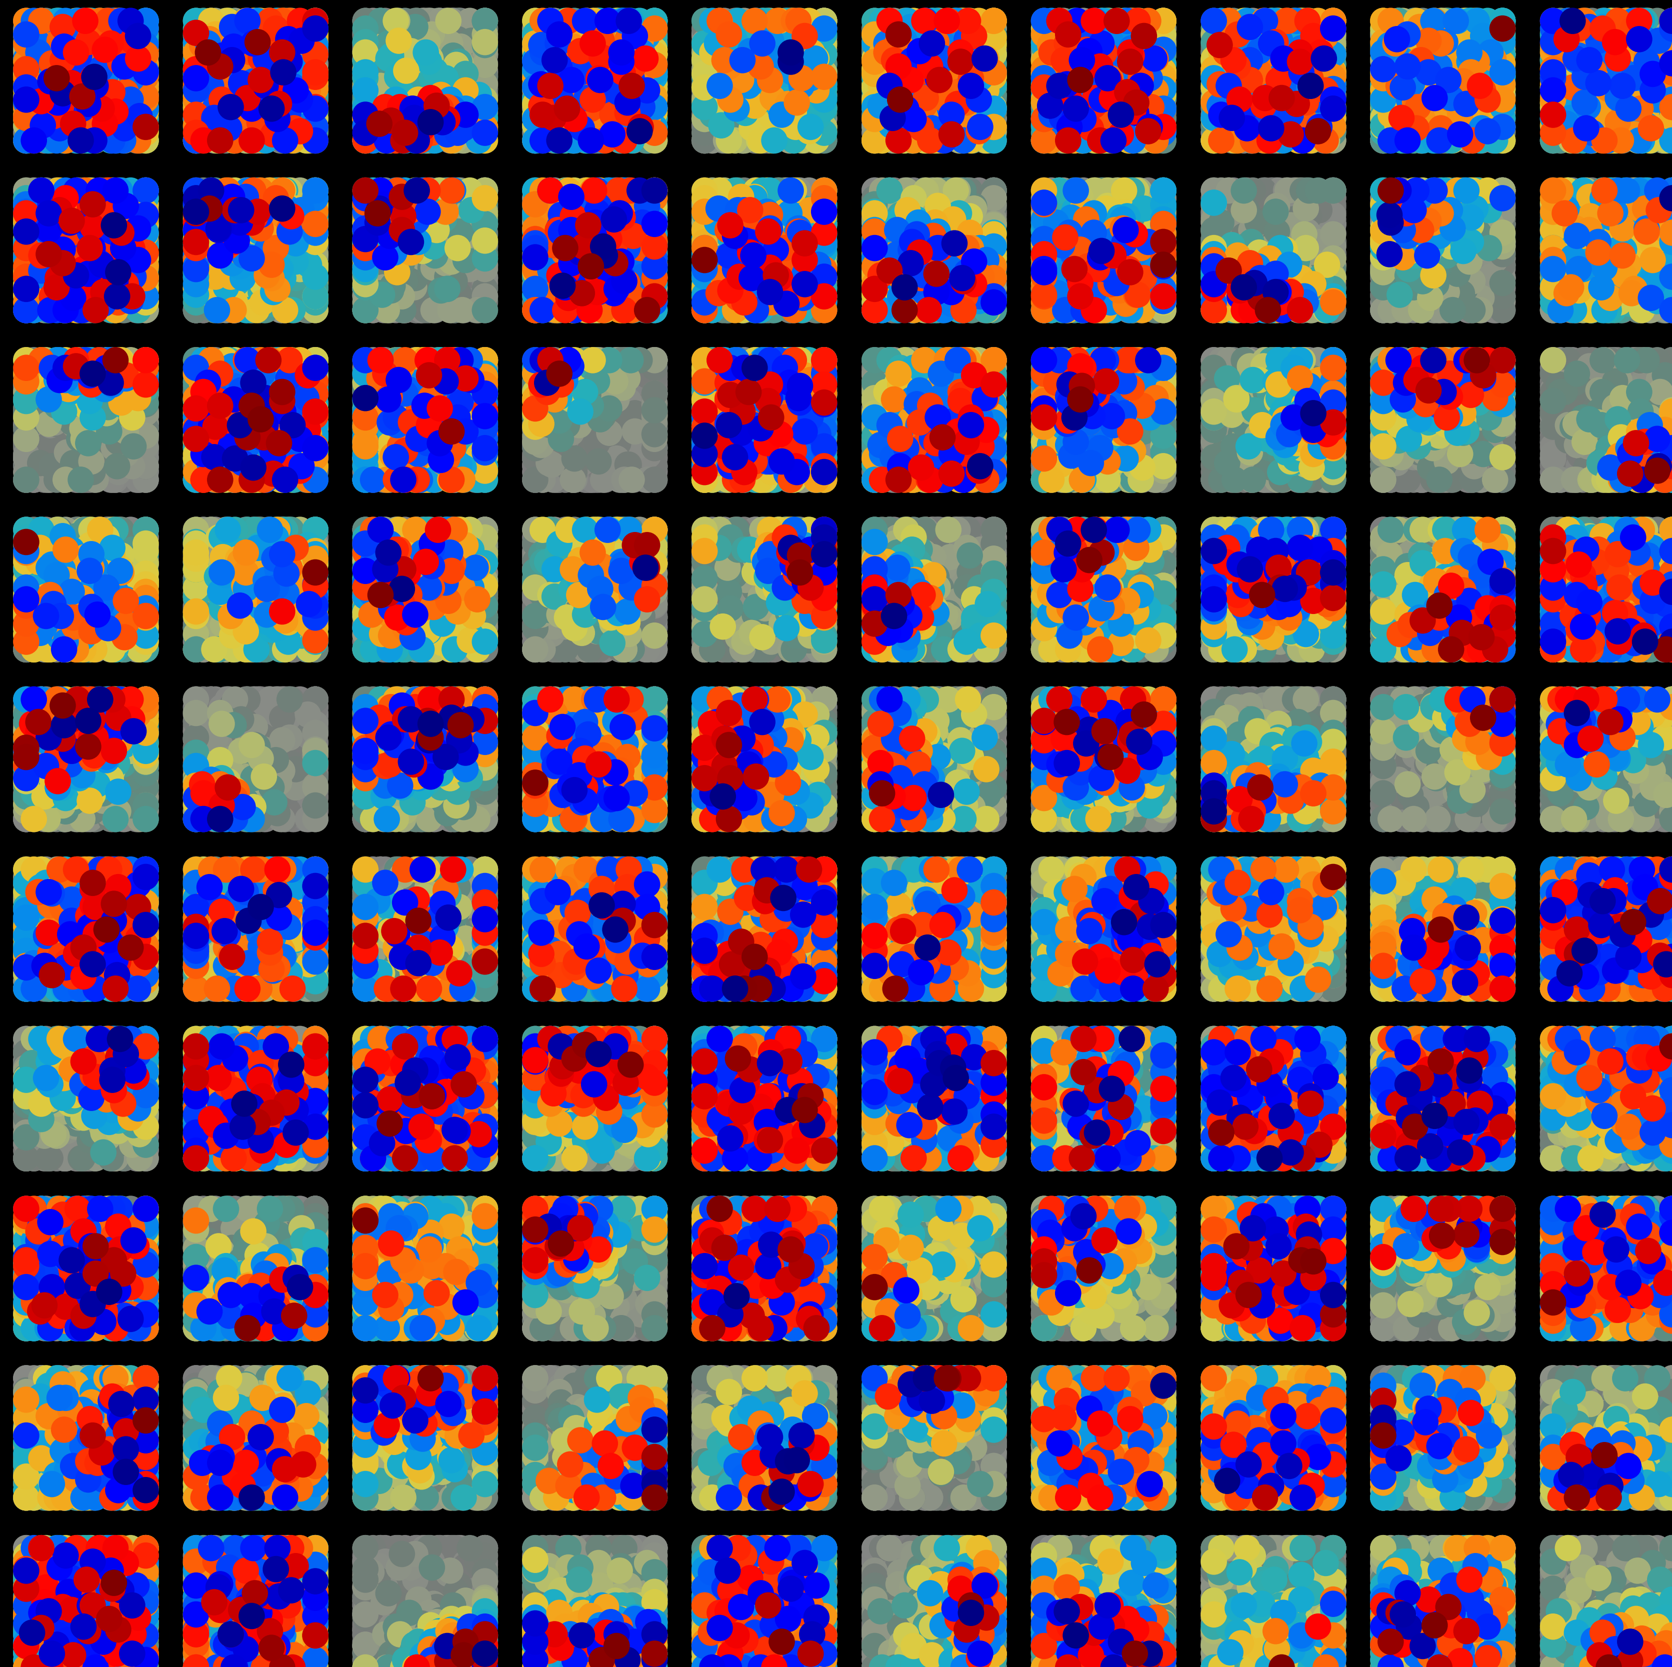

ı@berkeley.edu

Transformation Components in Frequency Space

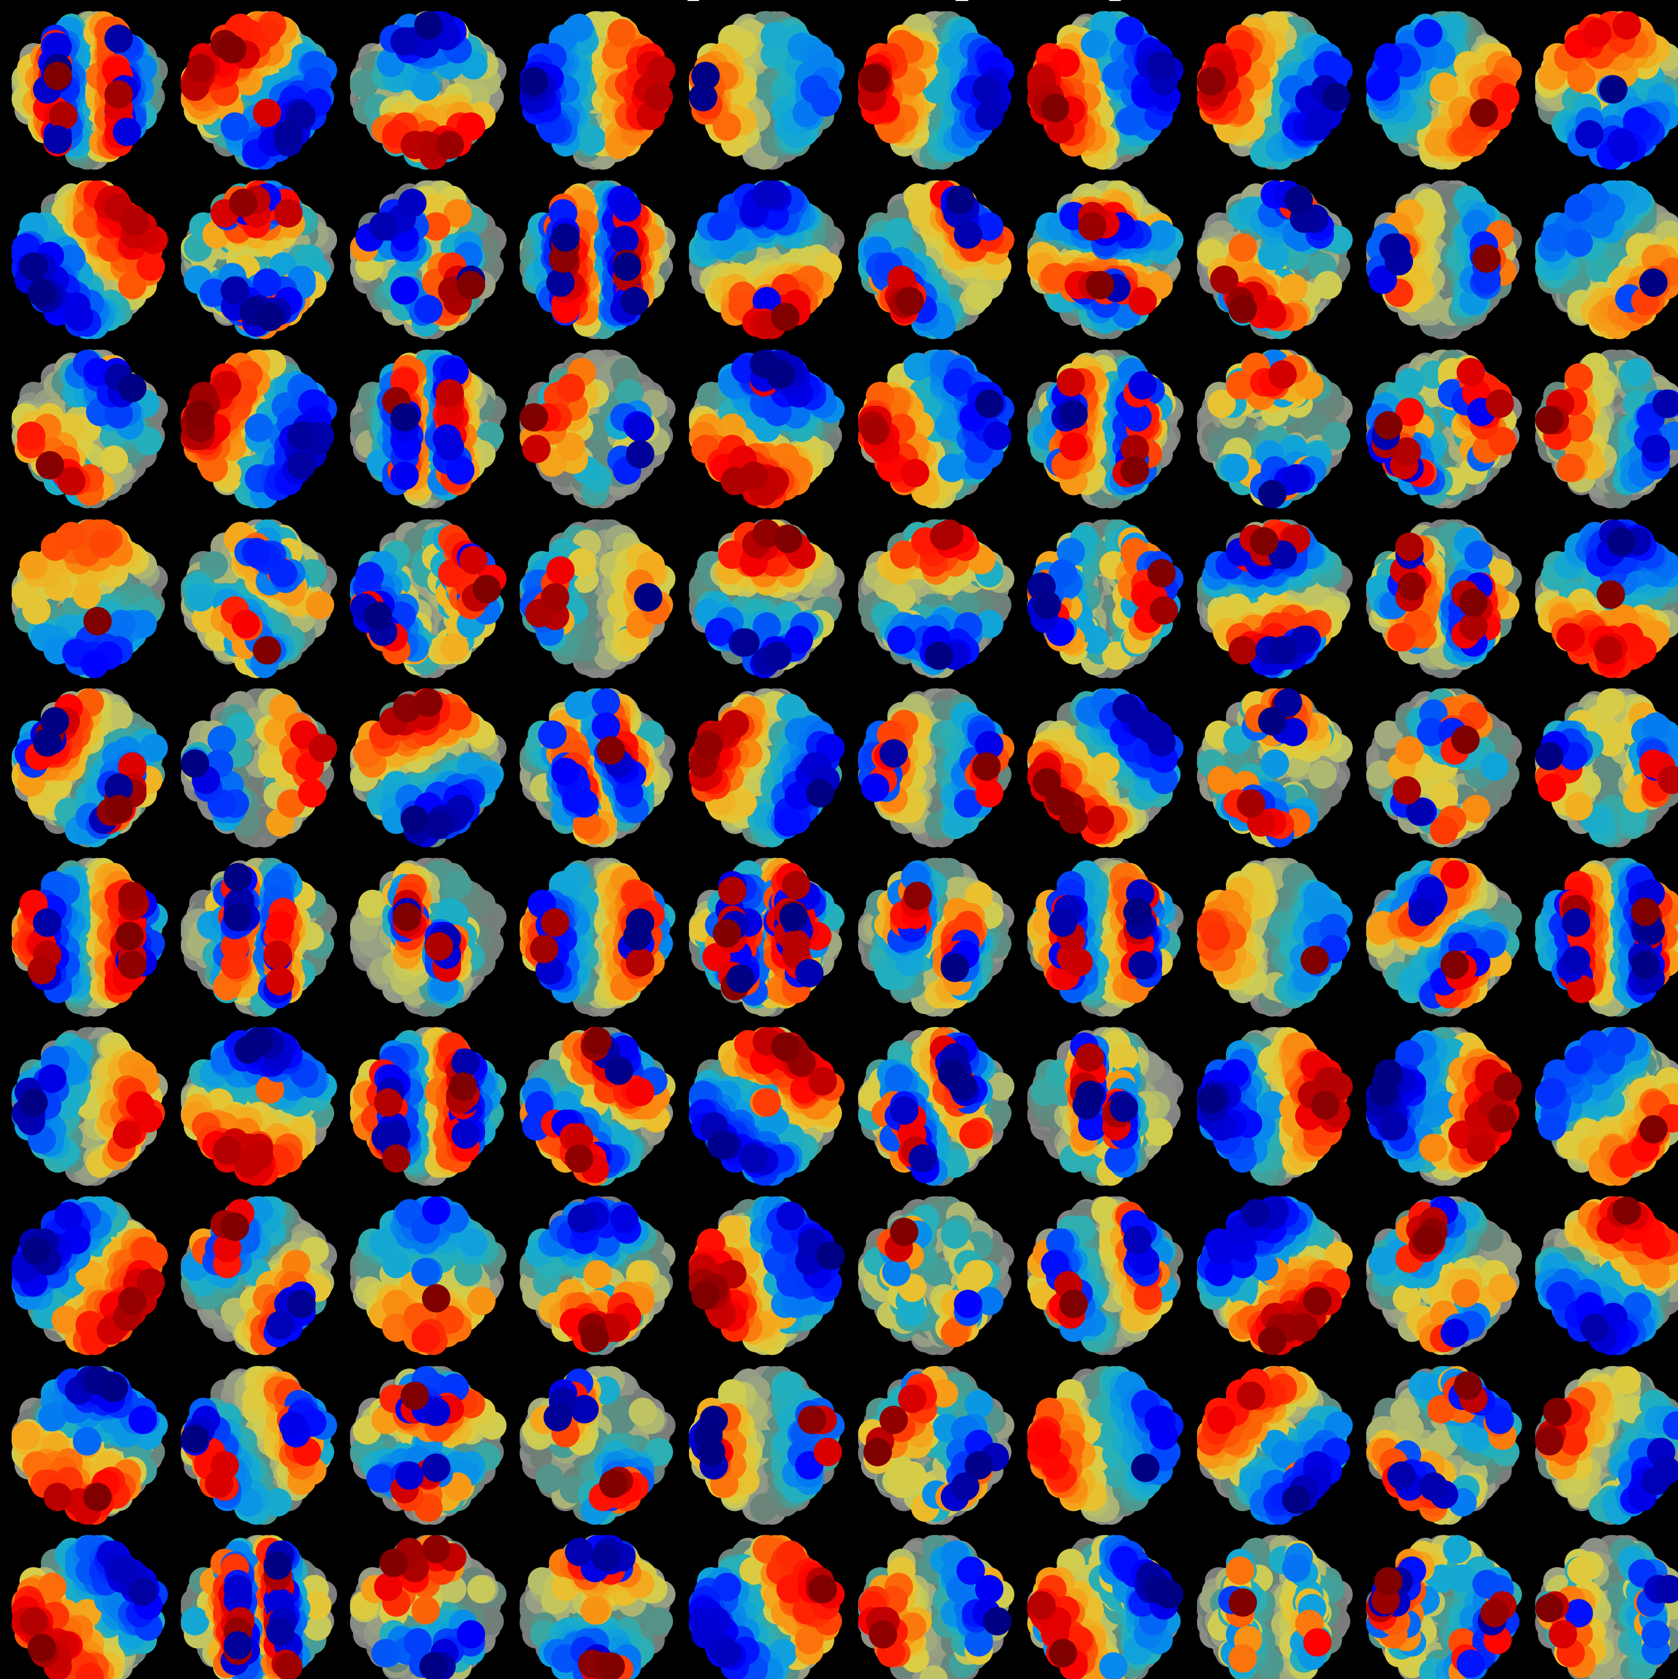

u@berkeley.edu

Learned Transformation Component

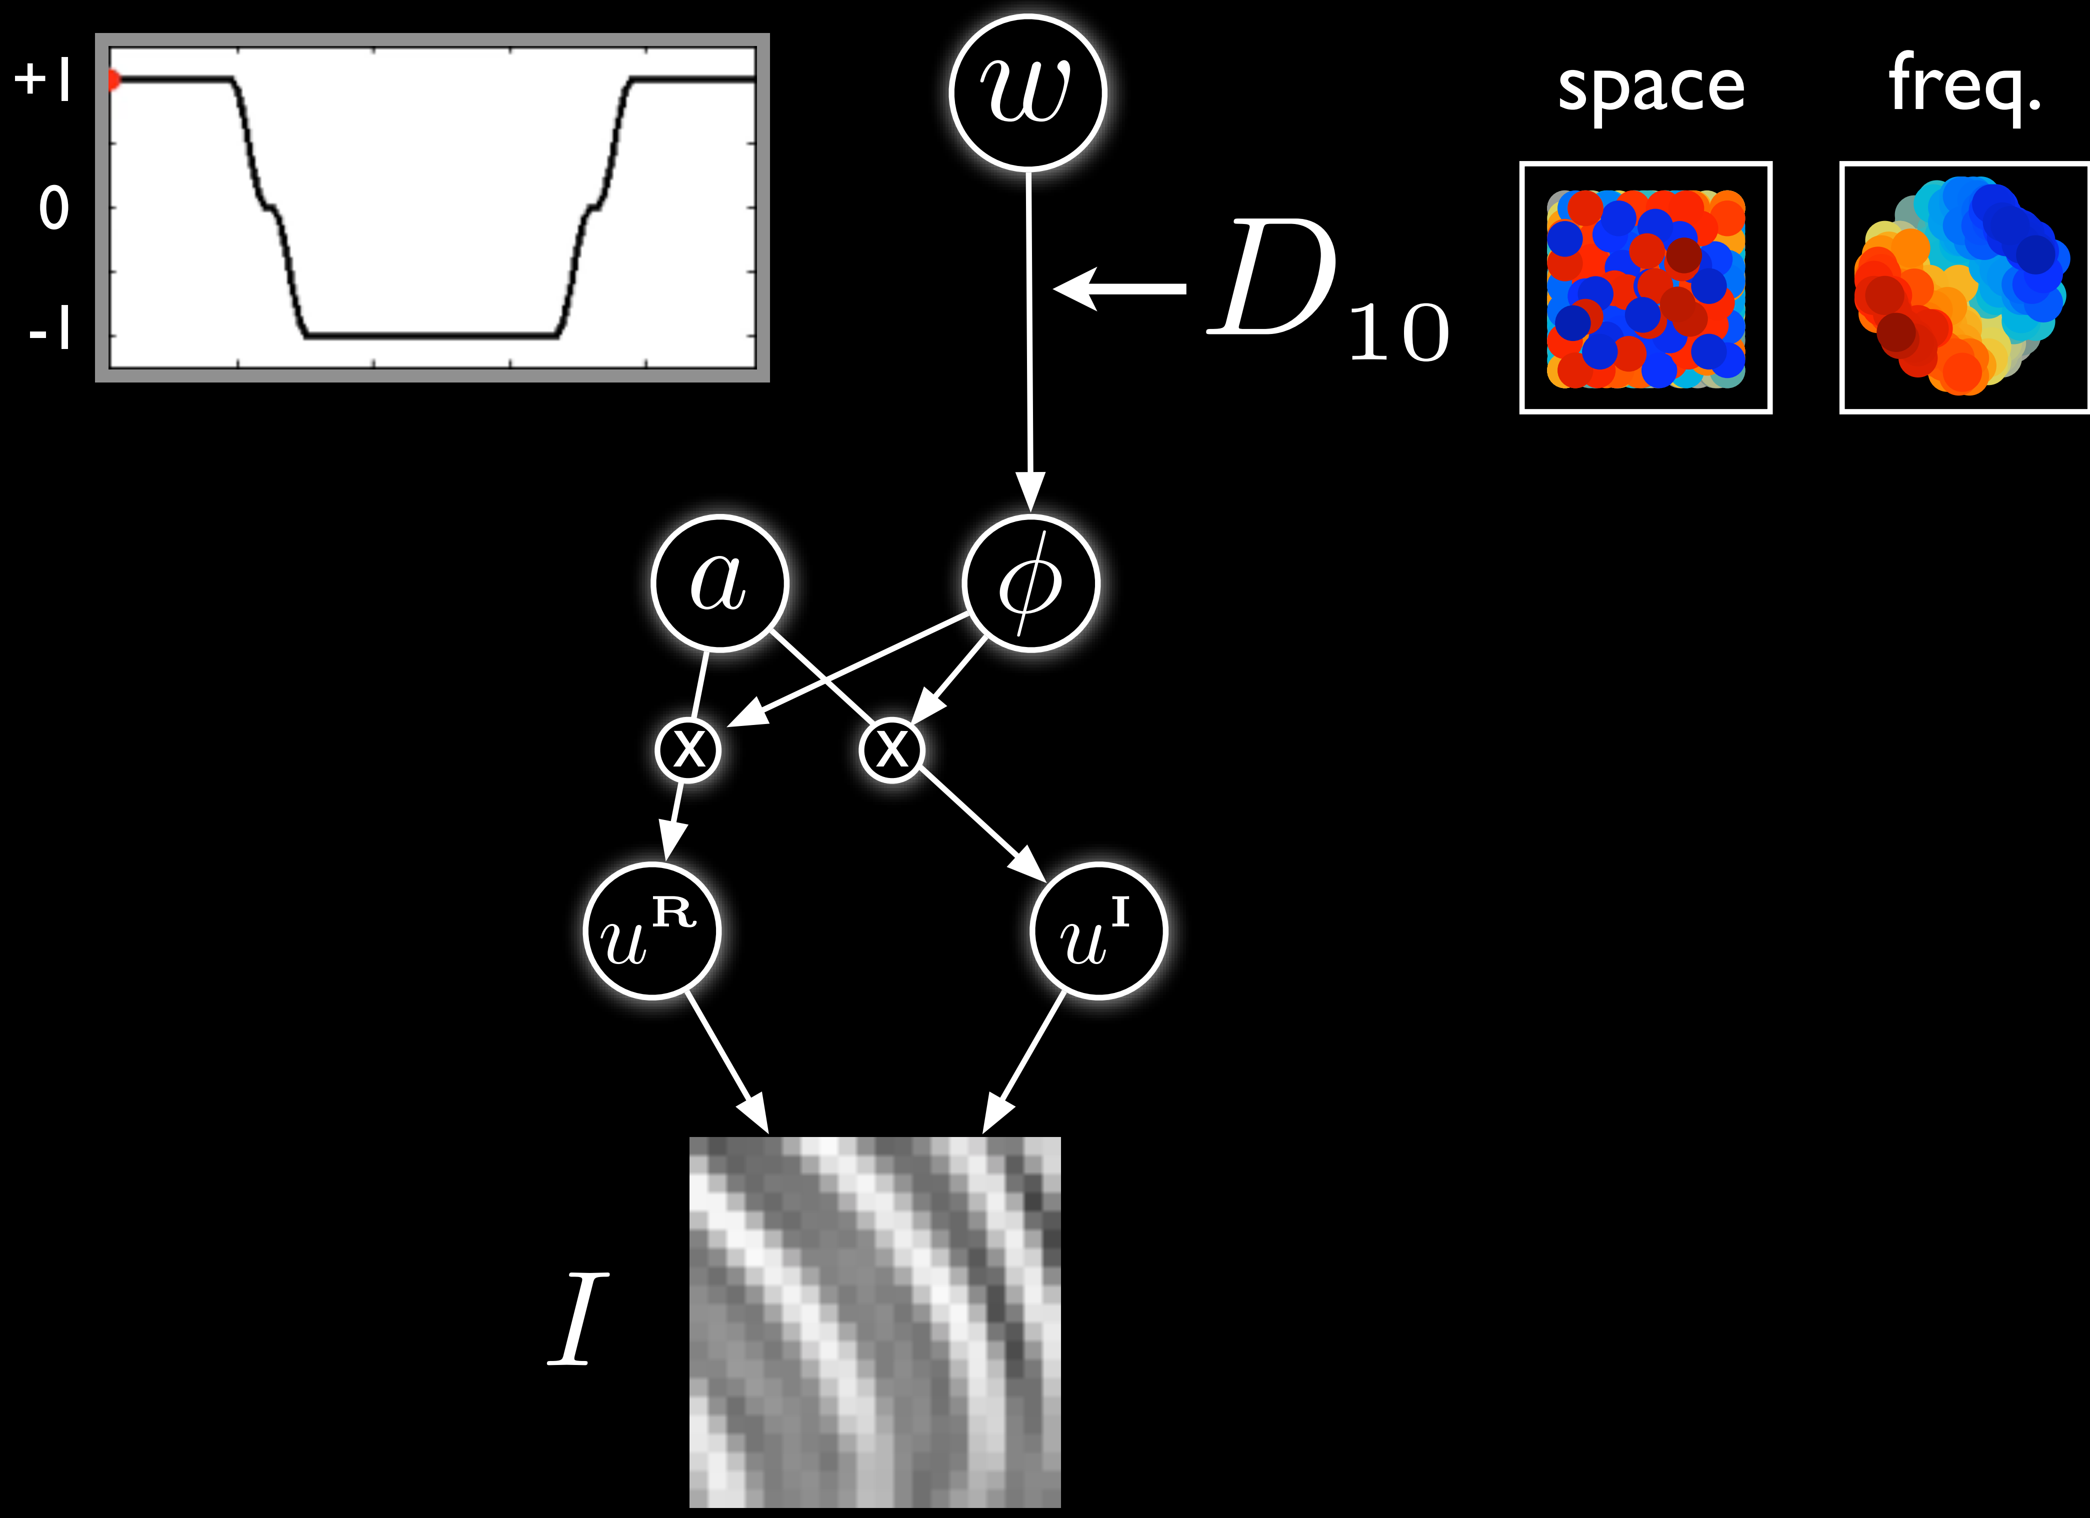

Charles Cadieu, cadieu@berkeley.edu

Learned Transformational Invariant

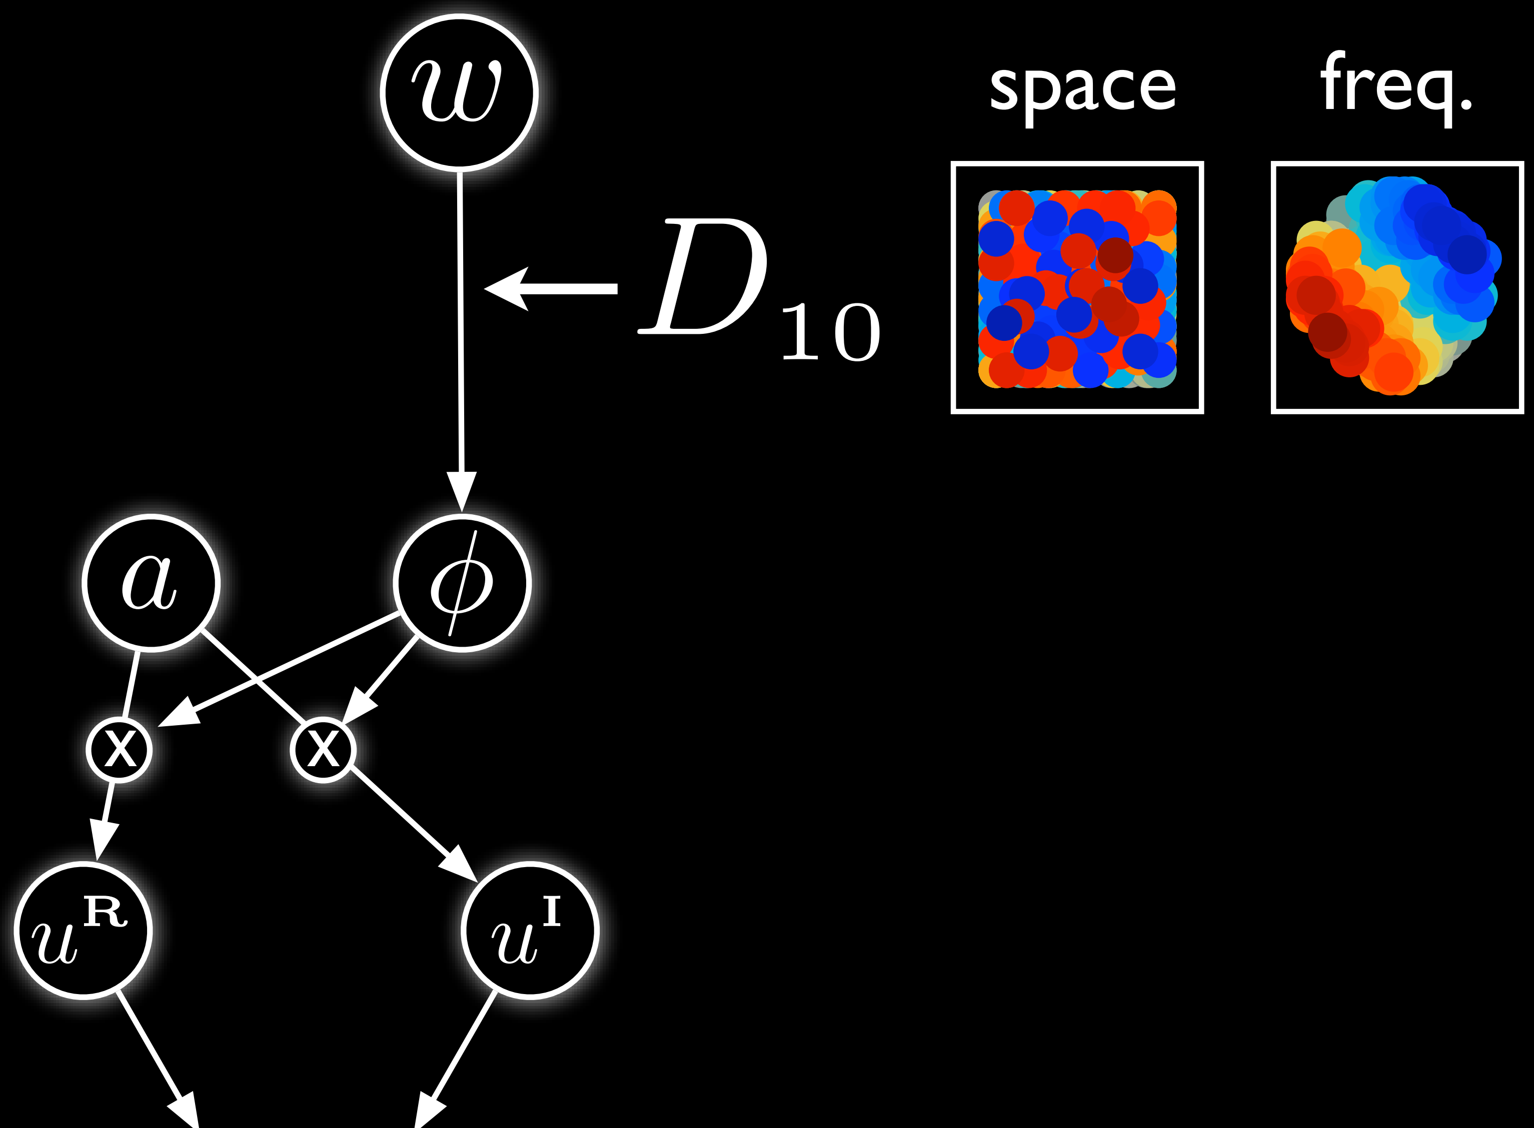

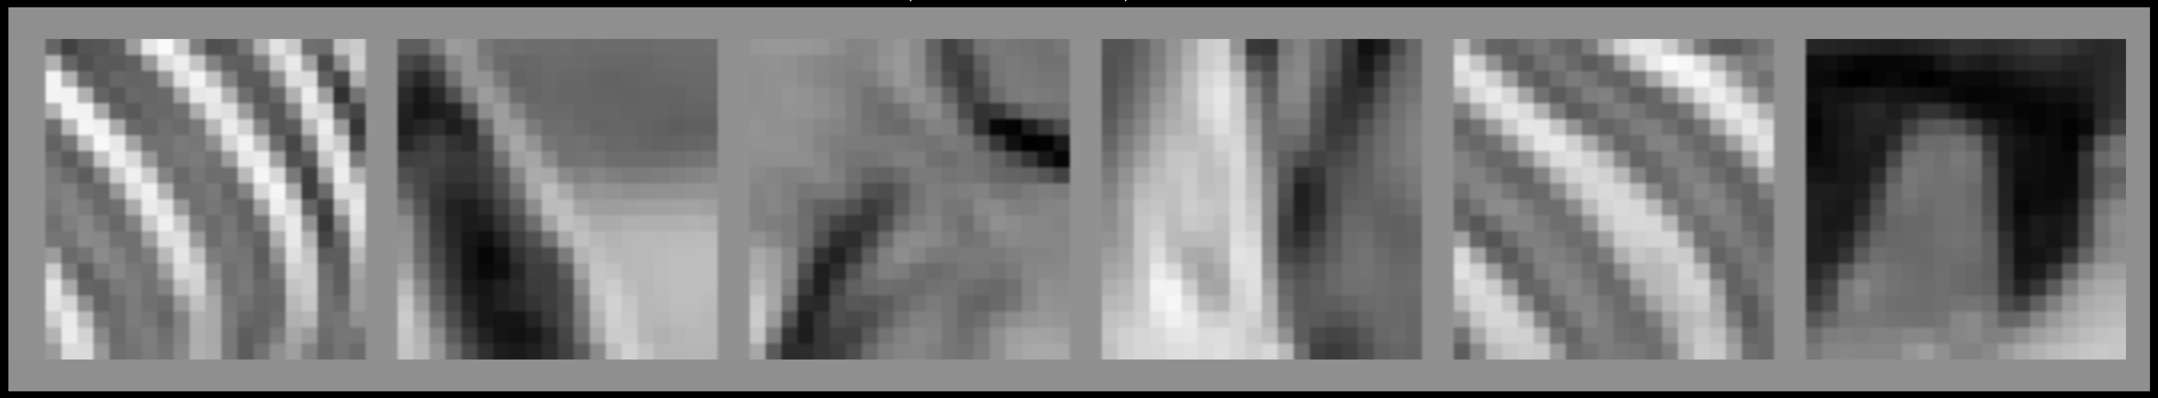

Charles Cadieu, cadieu@berkeley.edu

Learned Transformational Invariant

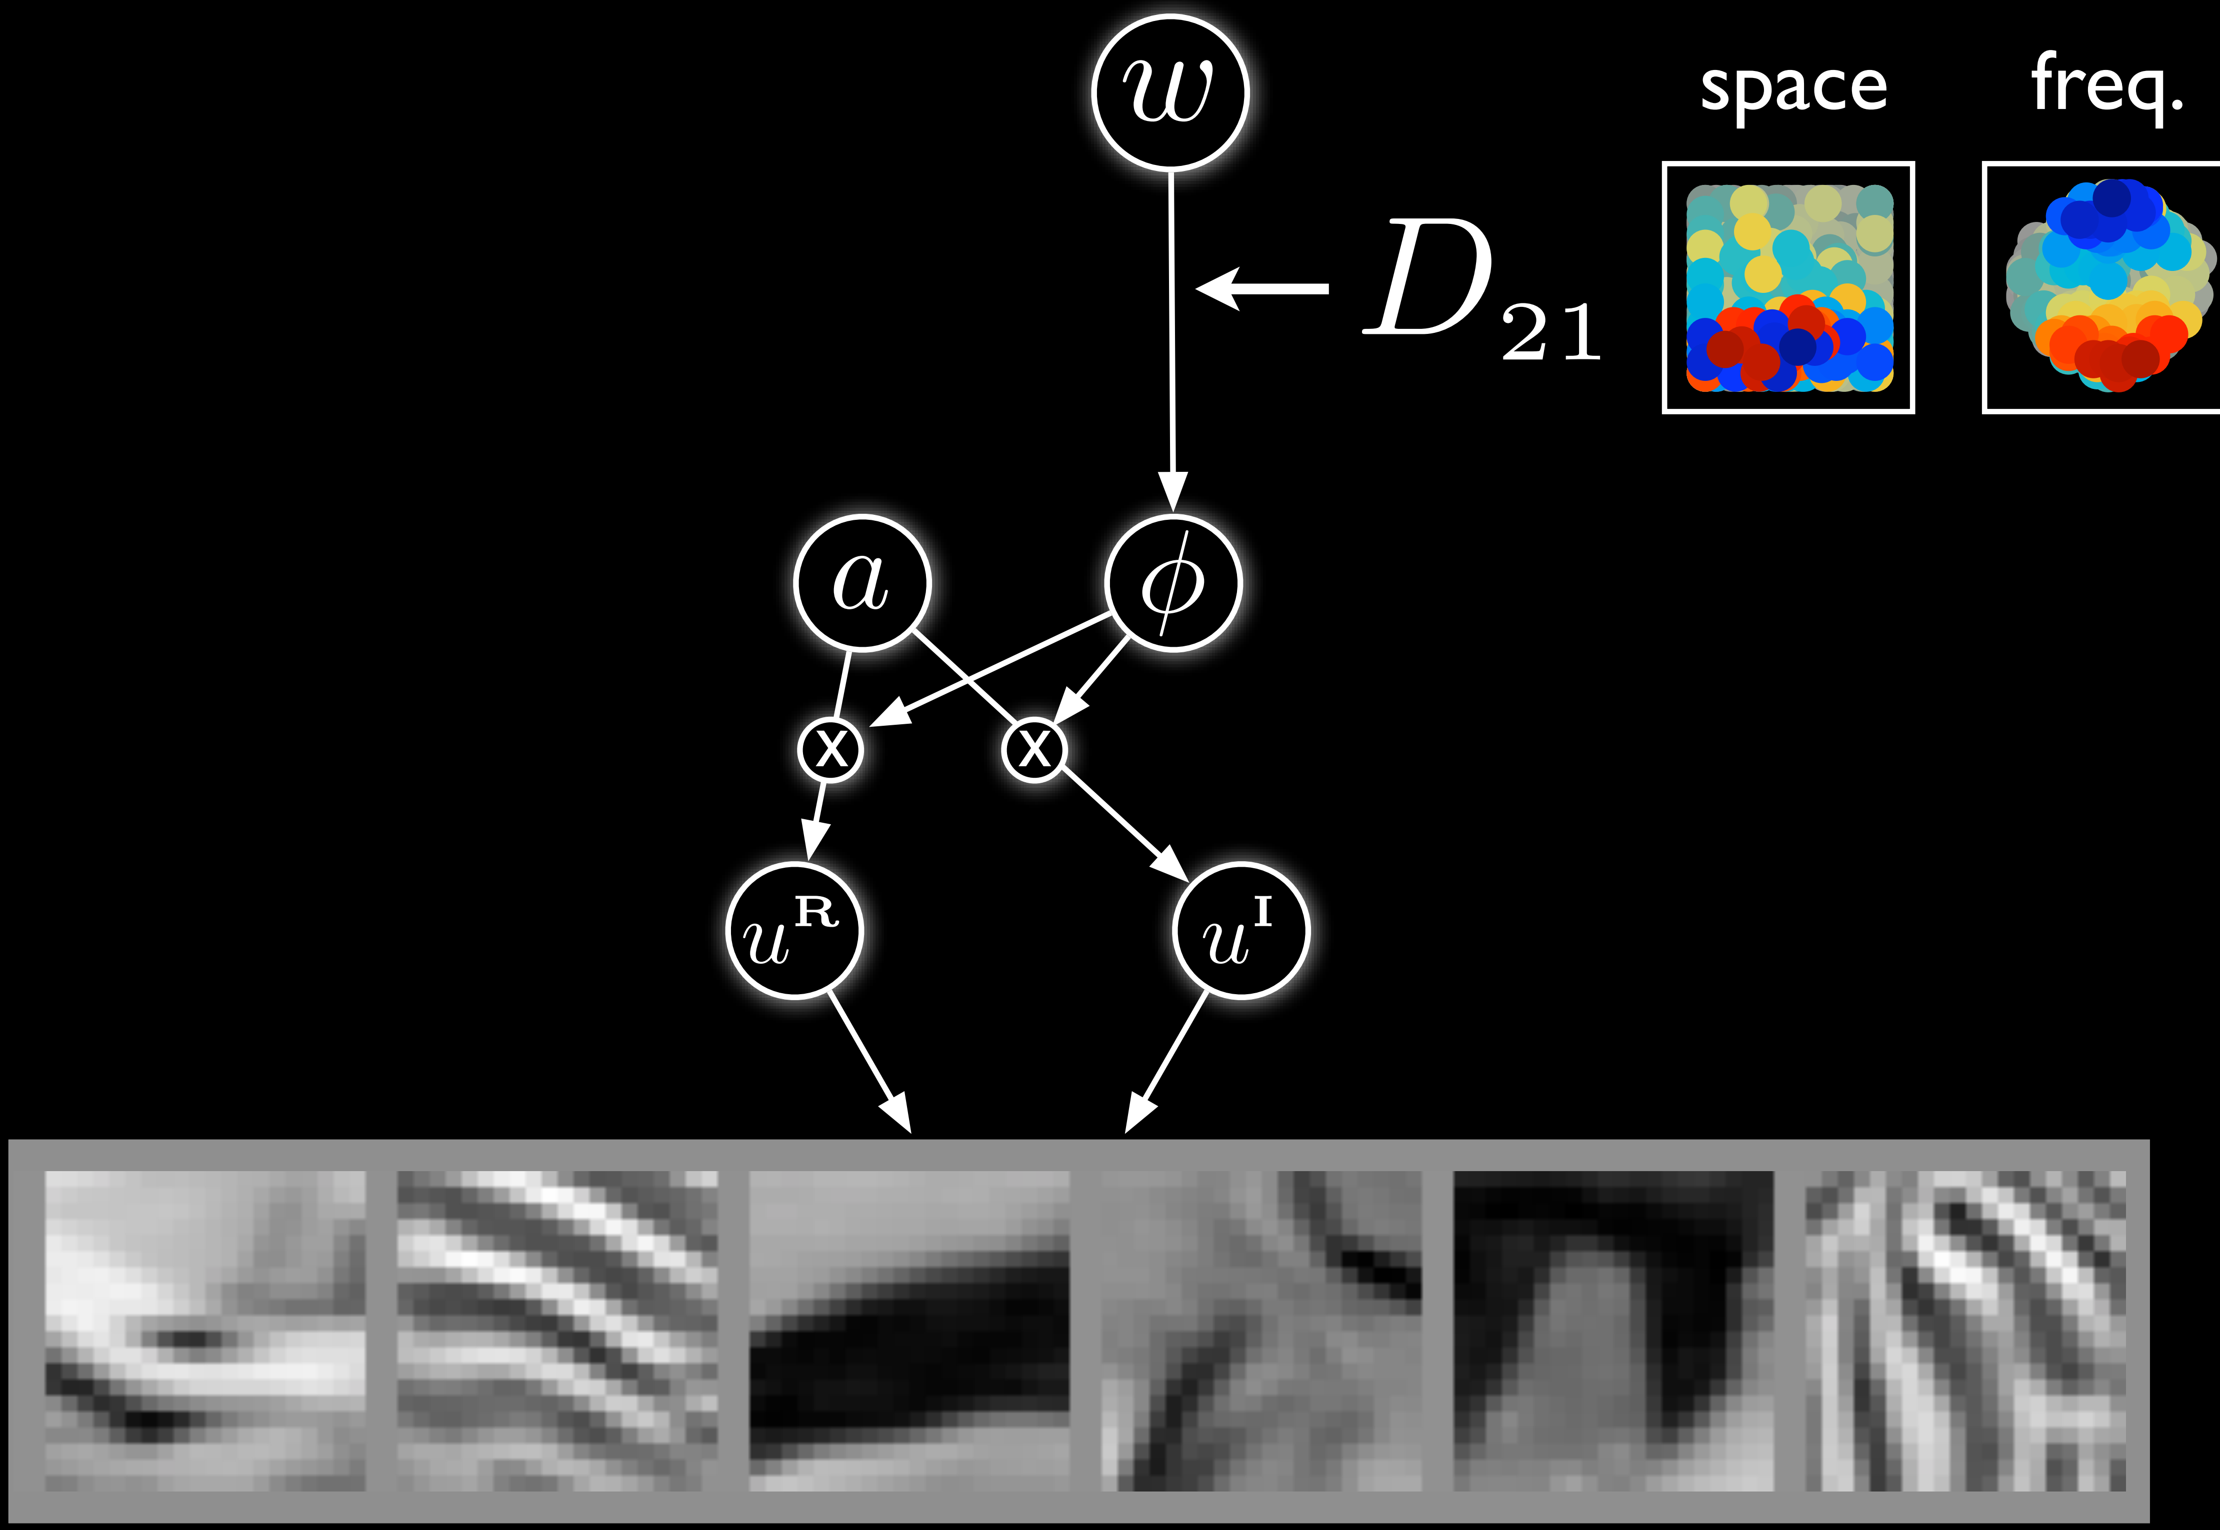

Charles Cadieu, cadieu@berkeley.edu

Learned Transformational Invariant

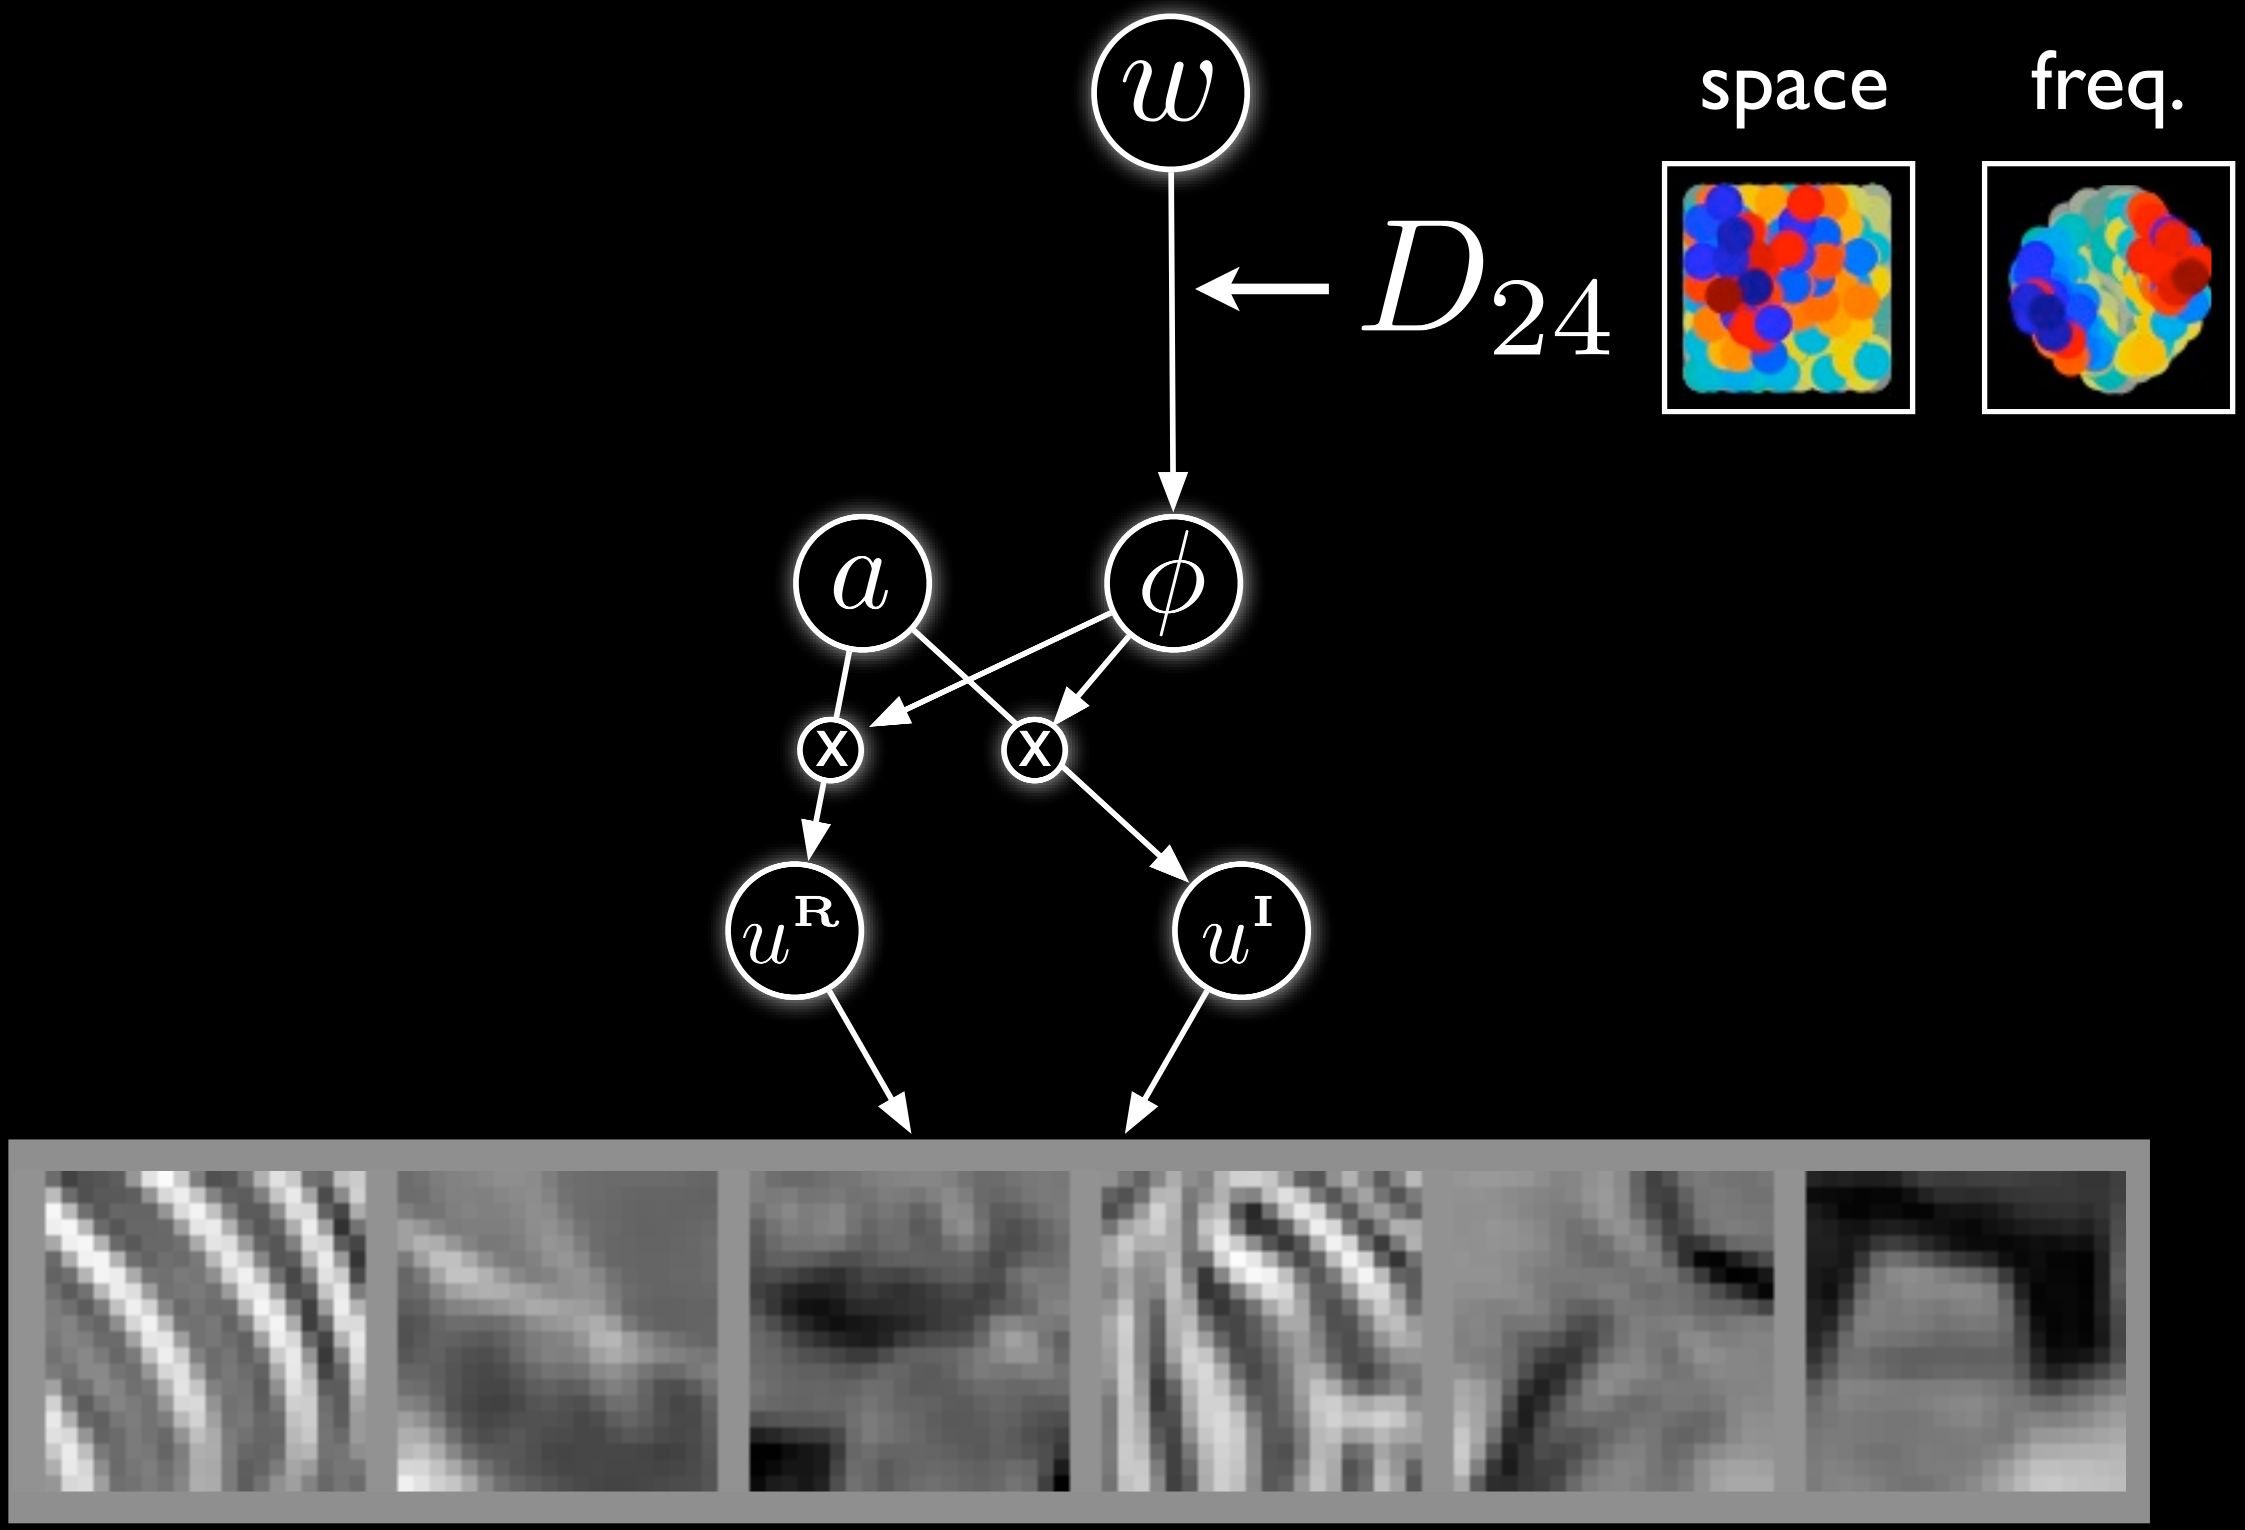

Charles Cadieu, cadieu@berkeley.edu

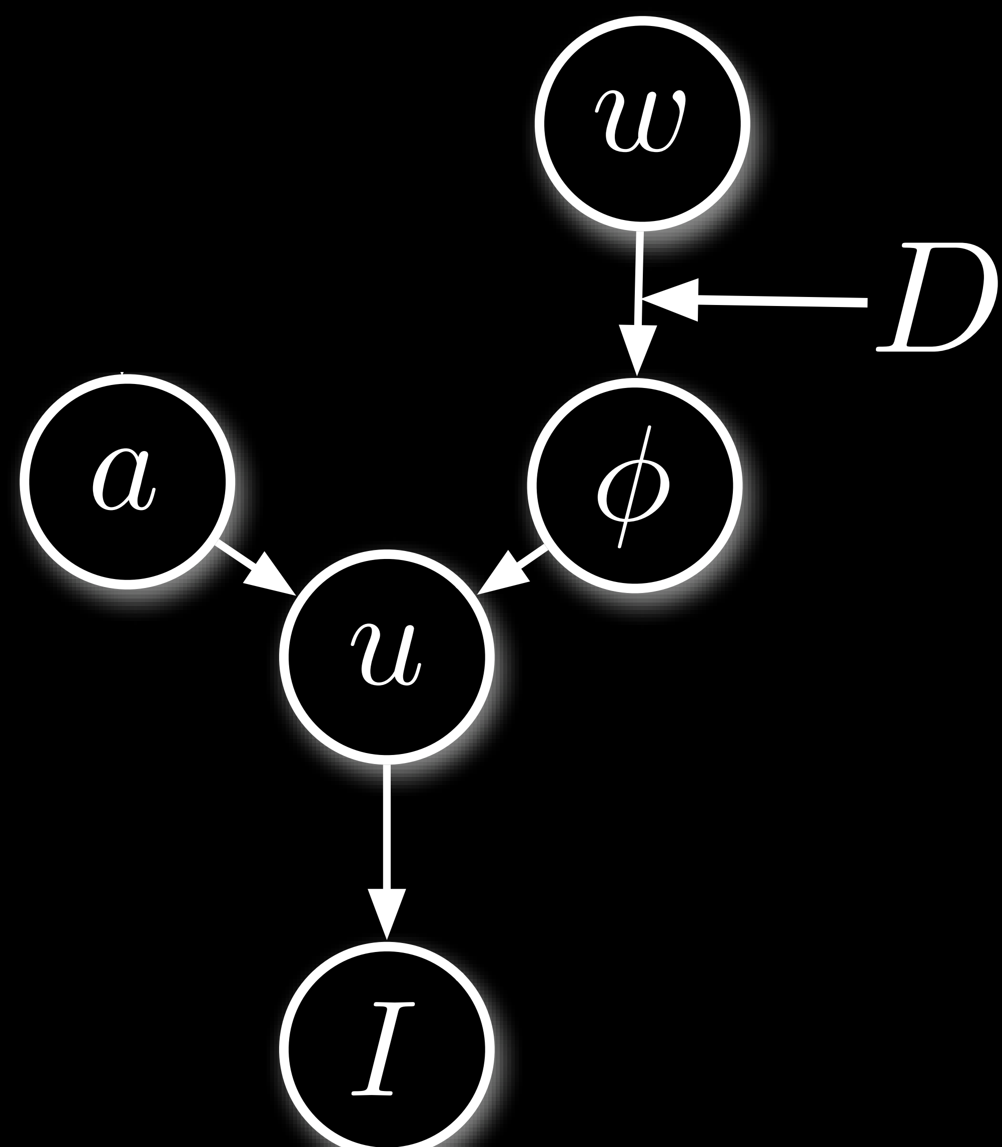

Charles Cadieu, cadieu@berkeley.edu

Second Layer Probabilistic Model: Spatial Invariants

$$P(I, a, \phi, v) \propto e^{-E_1 - E_2}$$

Error in log(amplitude)

Sparse

$$E_2 = -\sum_i \frac{1}{2\sigma_M^2}(\log(a_i) - [Bv]_i)^2 + \beta_{Sp} \sum_k |v_k|$$

Adapt to Natural Images

Charles Cadieu, cadieu@berkeley.edu

Learned Spatial Invariant

spatial domain

frequency domain

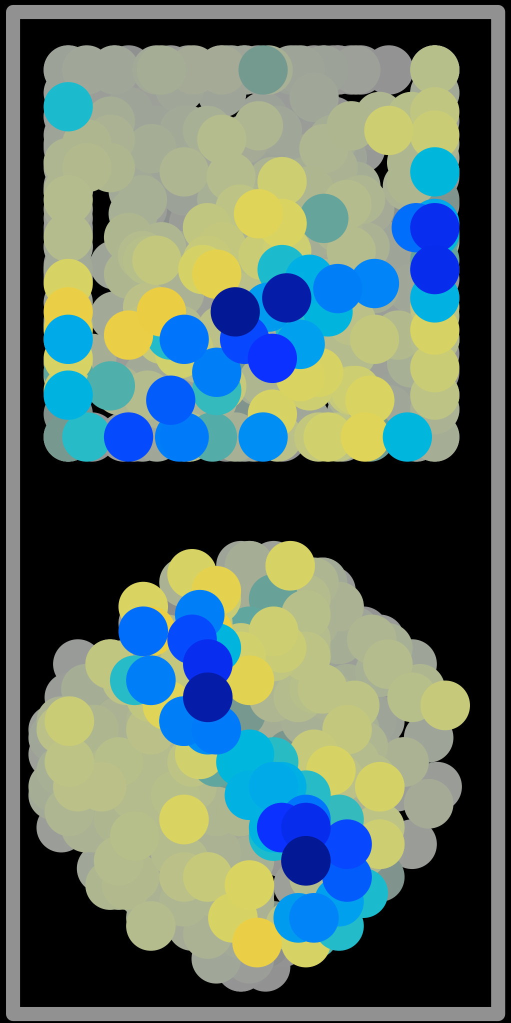

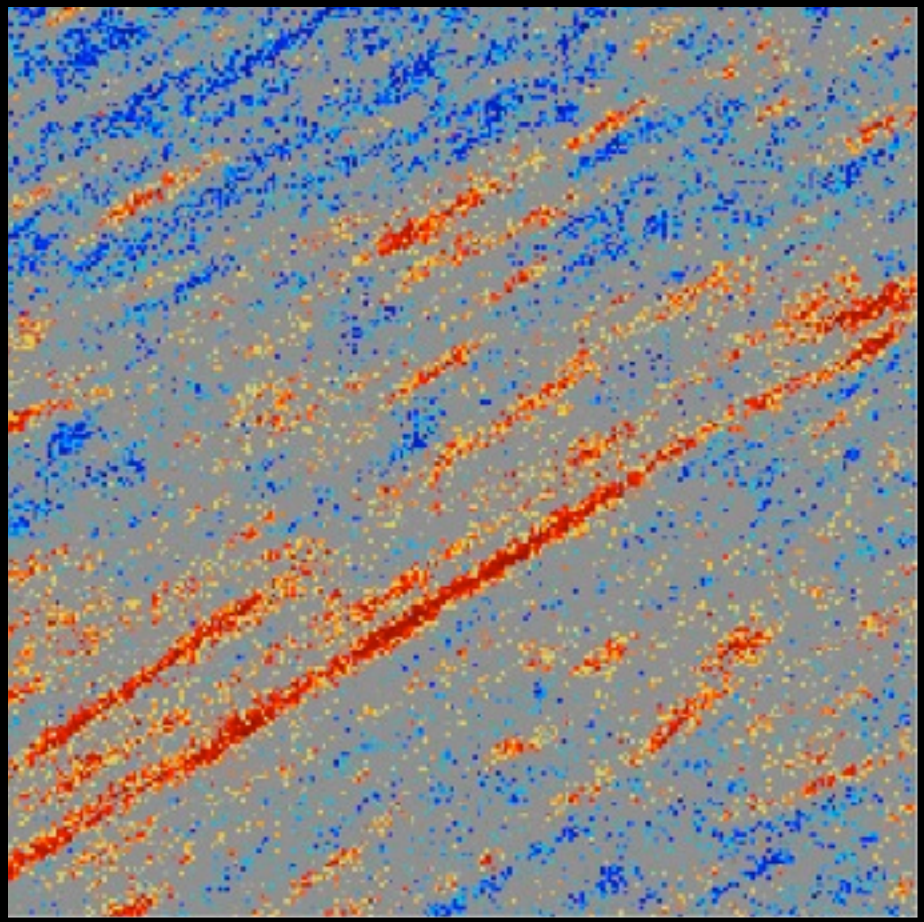

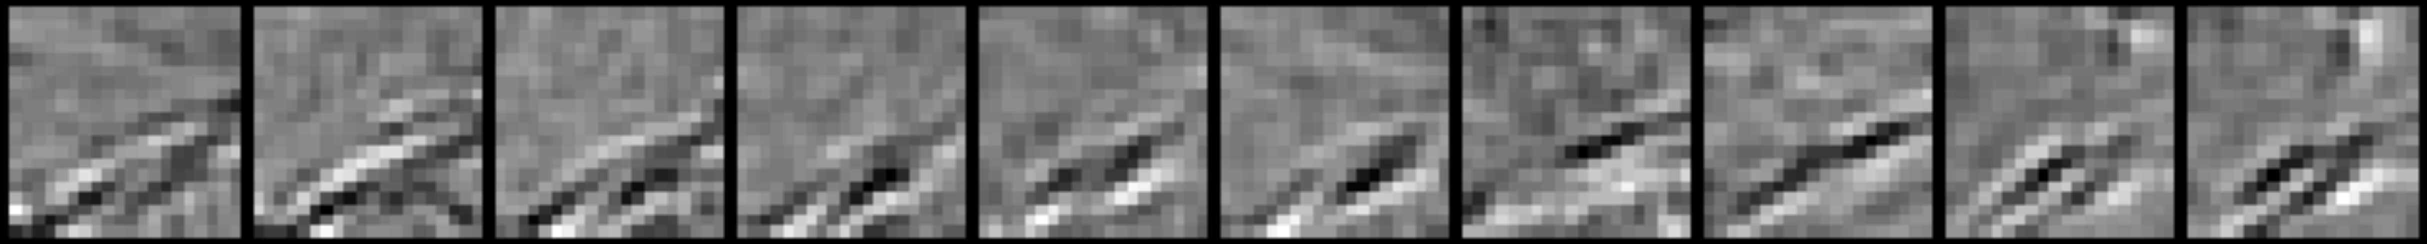

Charles Cadieu, cadieu@berkeley.edu

Learned Spatial Invariant

spatial domain

frequency domain

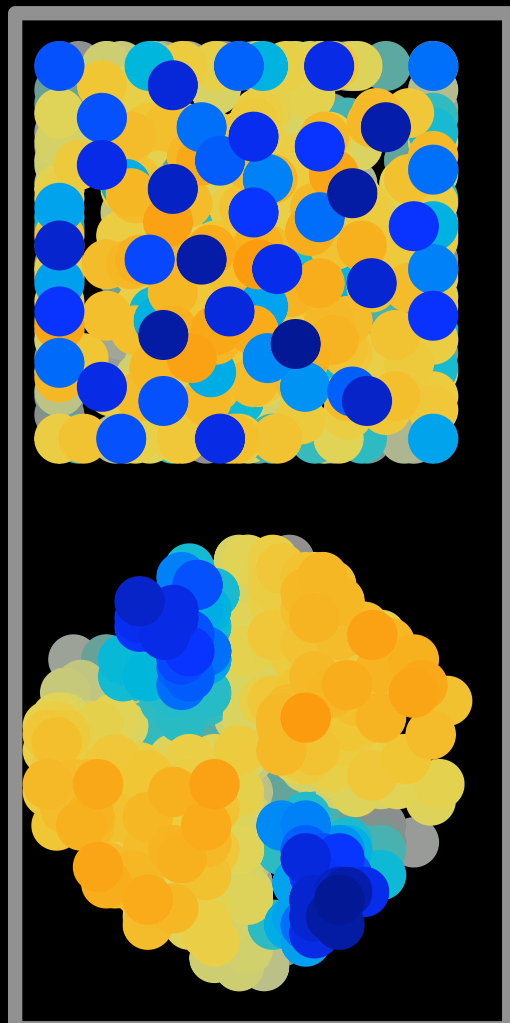

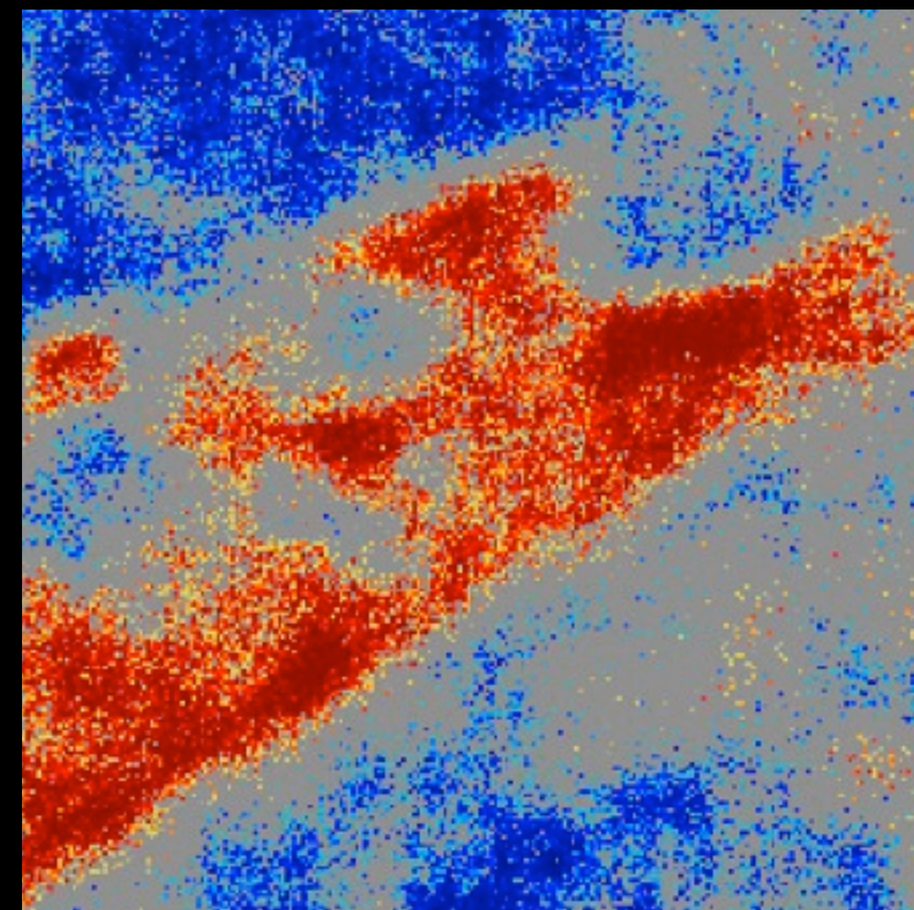

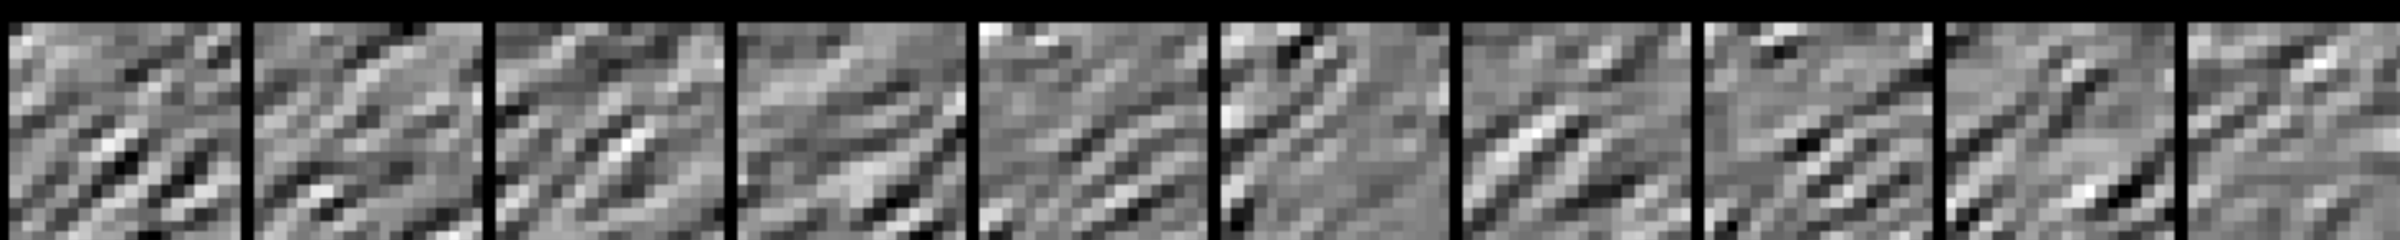

Charles Cadieu, cadieu@berkeley.edu

Learned Spatial Invariants

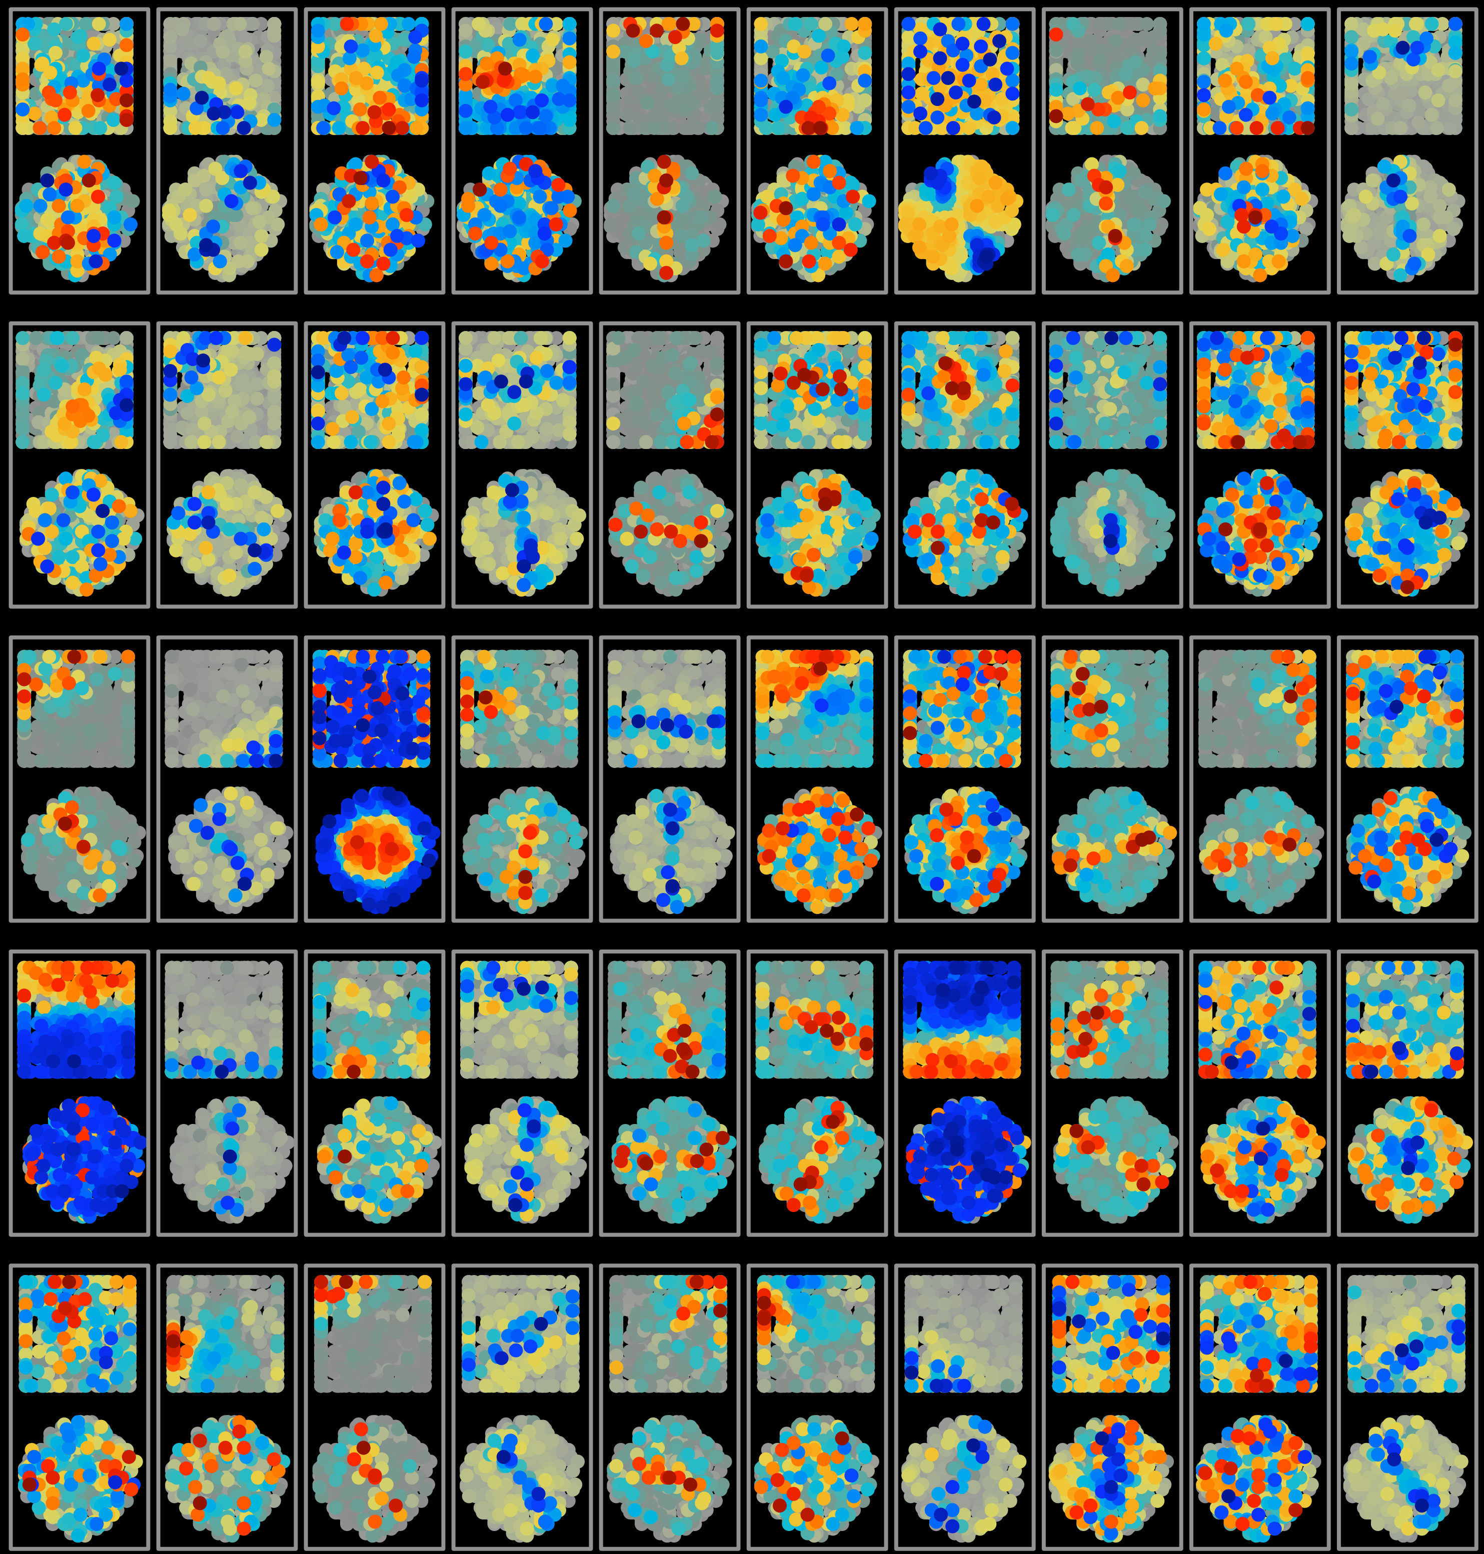

Charles Cadieu, cadieu@berkeley.edu

Feedback

$$P\big(\phi_i(t)\big|\phi_i(t{-}1), w(t)\big) \propto e^{k\cos(\phi_i(t)-\phi_i(t-1)-[D\,w]_i)}$$

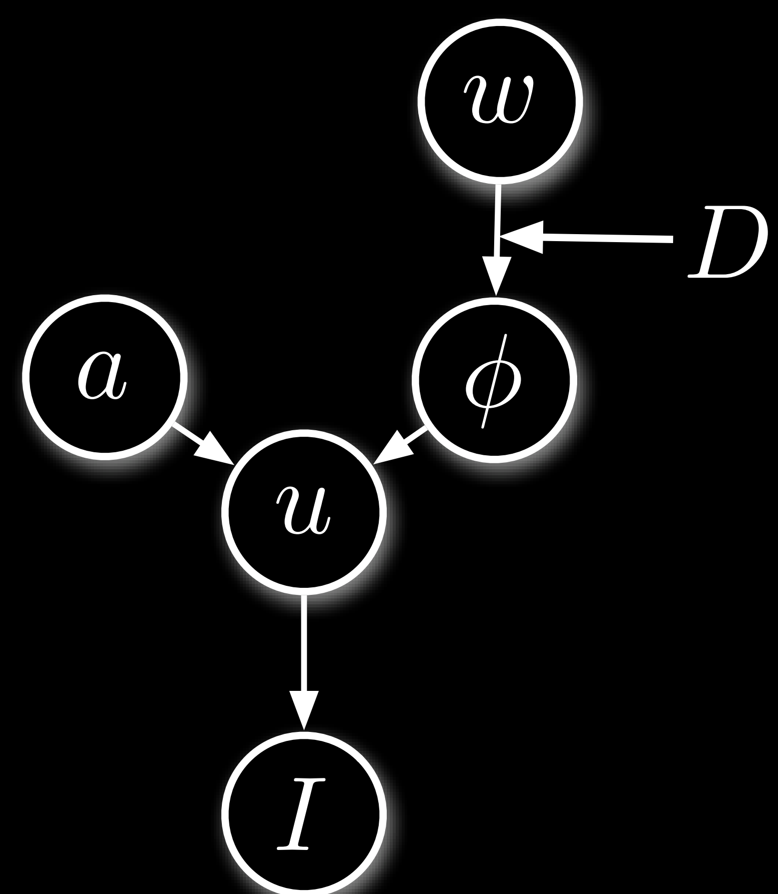

Charles Cadieu, cadieu@berkeley.edu

Image Denoising: testing an image model

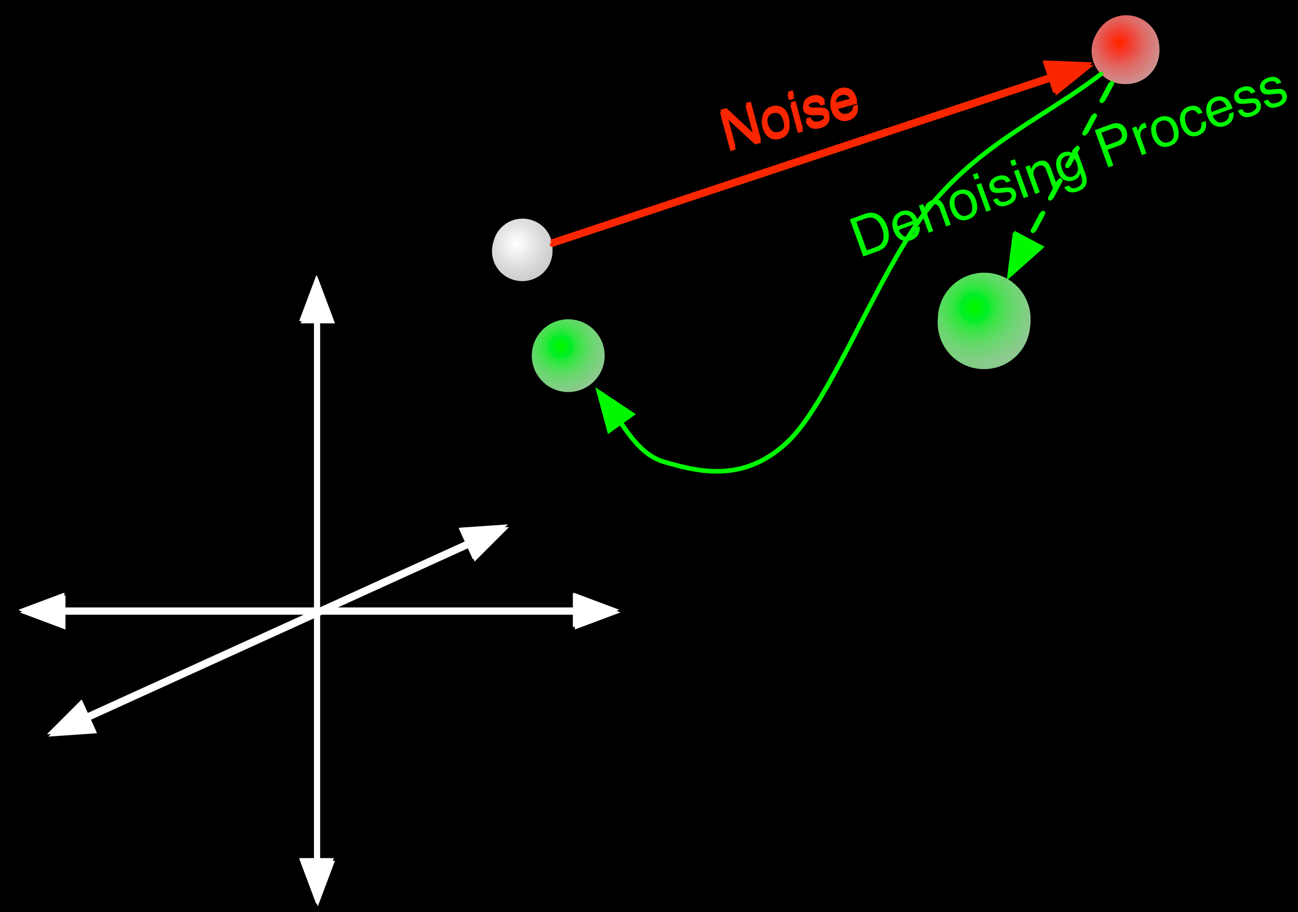

Charles Cadieu, cadieu@berkeley.edu

Original Noisy Denoised

SNR = -2.9 SNR = 6.7

Compare to Wiener, SNR = 3.9

Charles Cadieu, cadieu@berkeley.edu

Denoising Movies

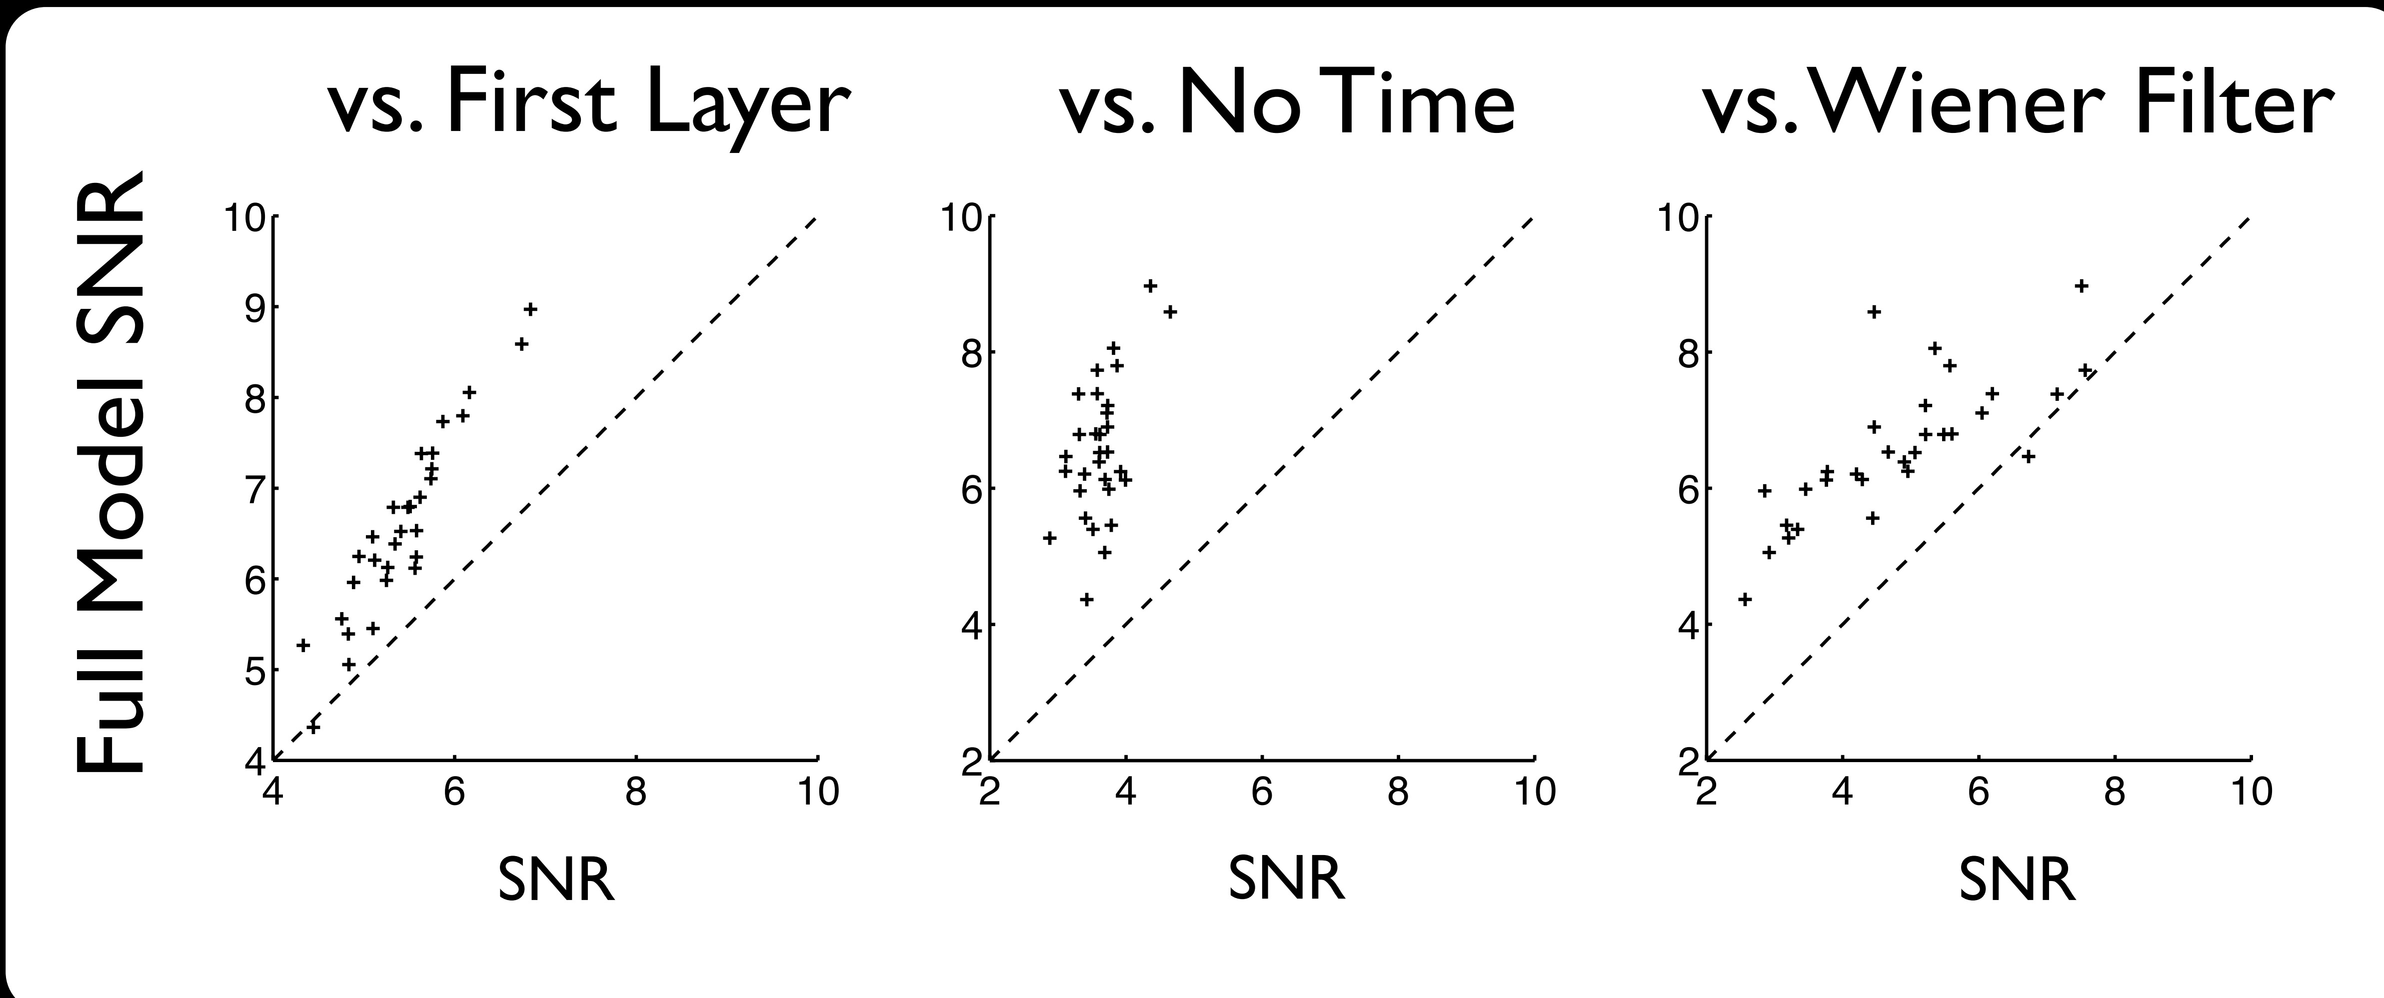

Mean SNR of Noisy Images = -2.0

Charles Cadieu, cadieu@berkeley.edu

Conclusions

We have,

- Motivated models the produce interpretations of the visual world,
- Learned transformational and spatial invariants from the natural world, and
- Shown how the model improves the interpretation of ambiguous inputs.

Charles Cadieu, cadieu@berkeley.edu

Thanks to

Bruno Olshausen

Redwood Center for Theoretical Neuroscience

NSF support

Charles Cadieu, cadieu@berkeley.edu

Extra Slides...

Charles Cadieu, cadieu@berkeley.edu

Feedback Changes First Layer Response

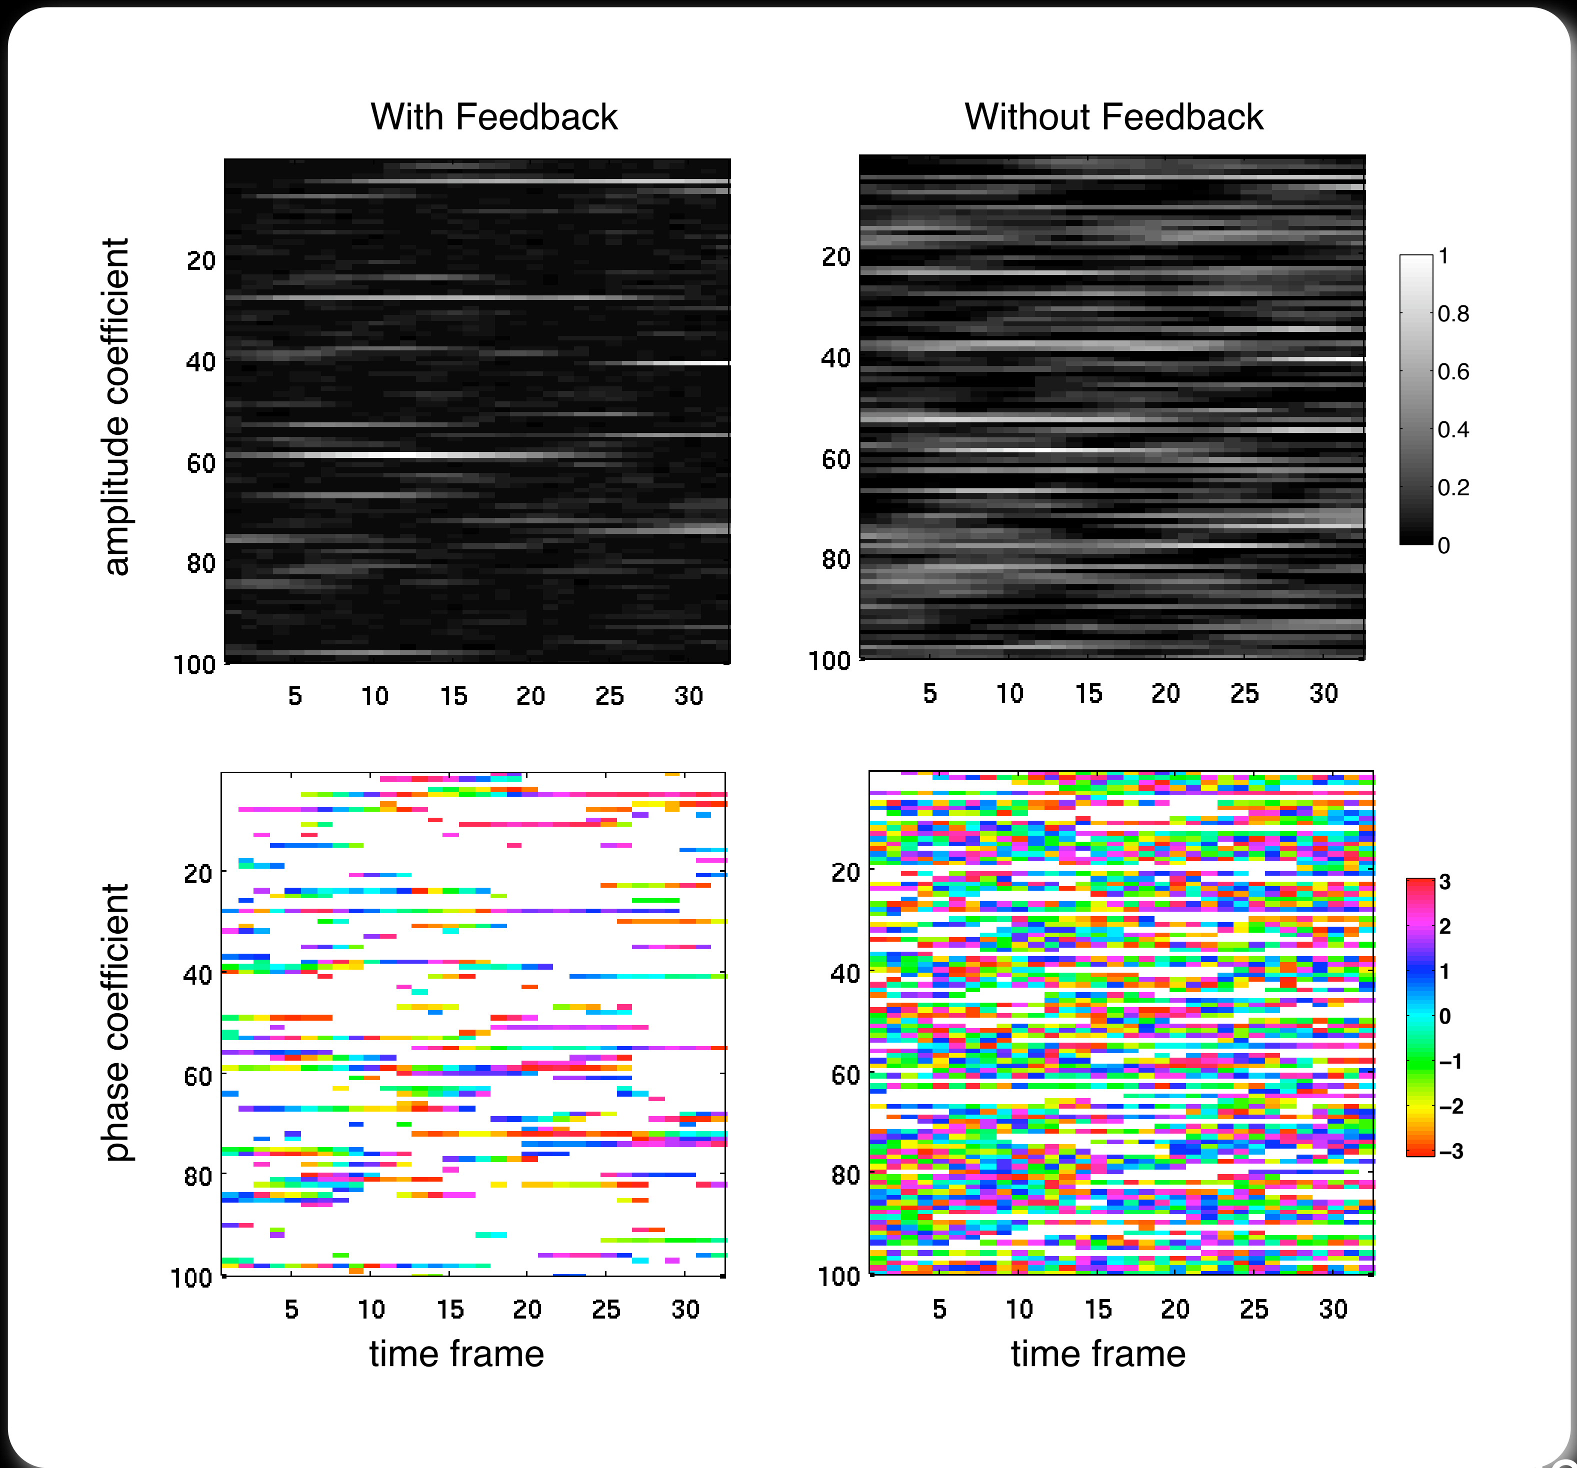

Charles Cadieu, cadieu@berkeley.edu

Feedback Changes First Layer Response

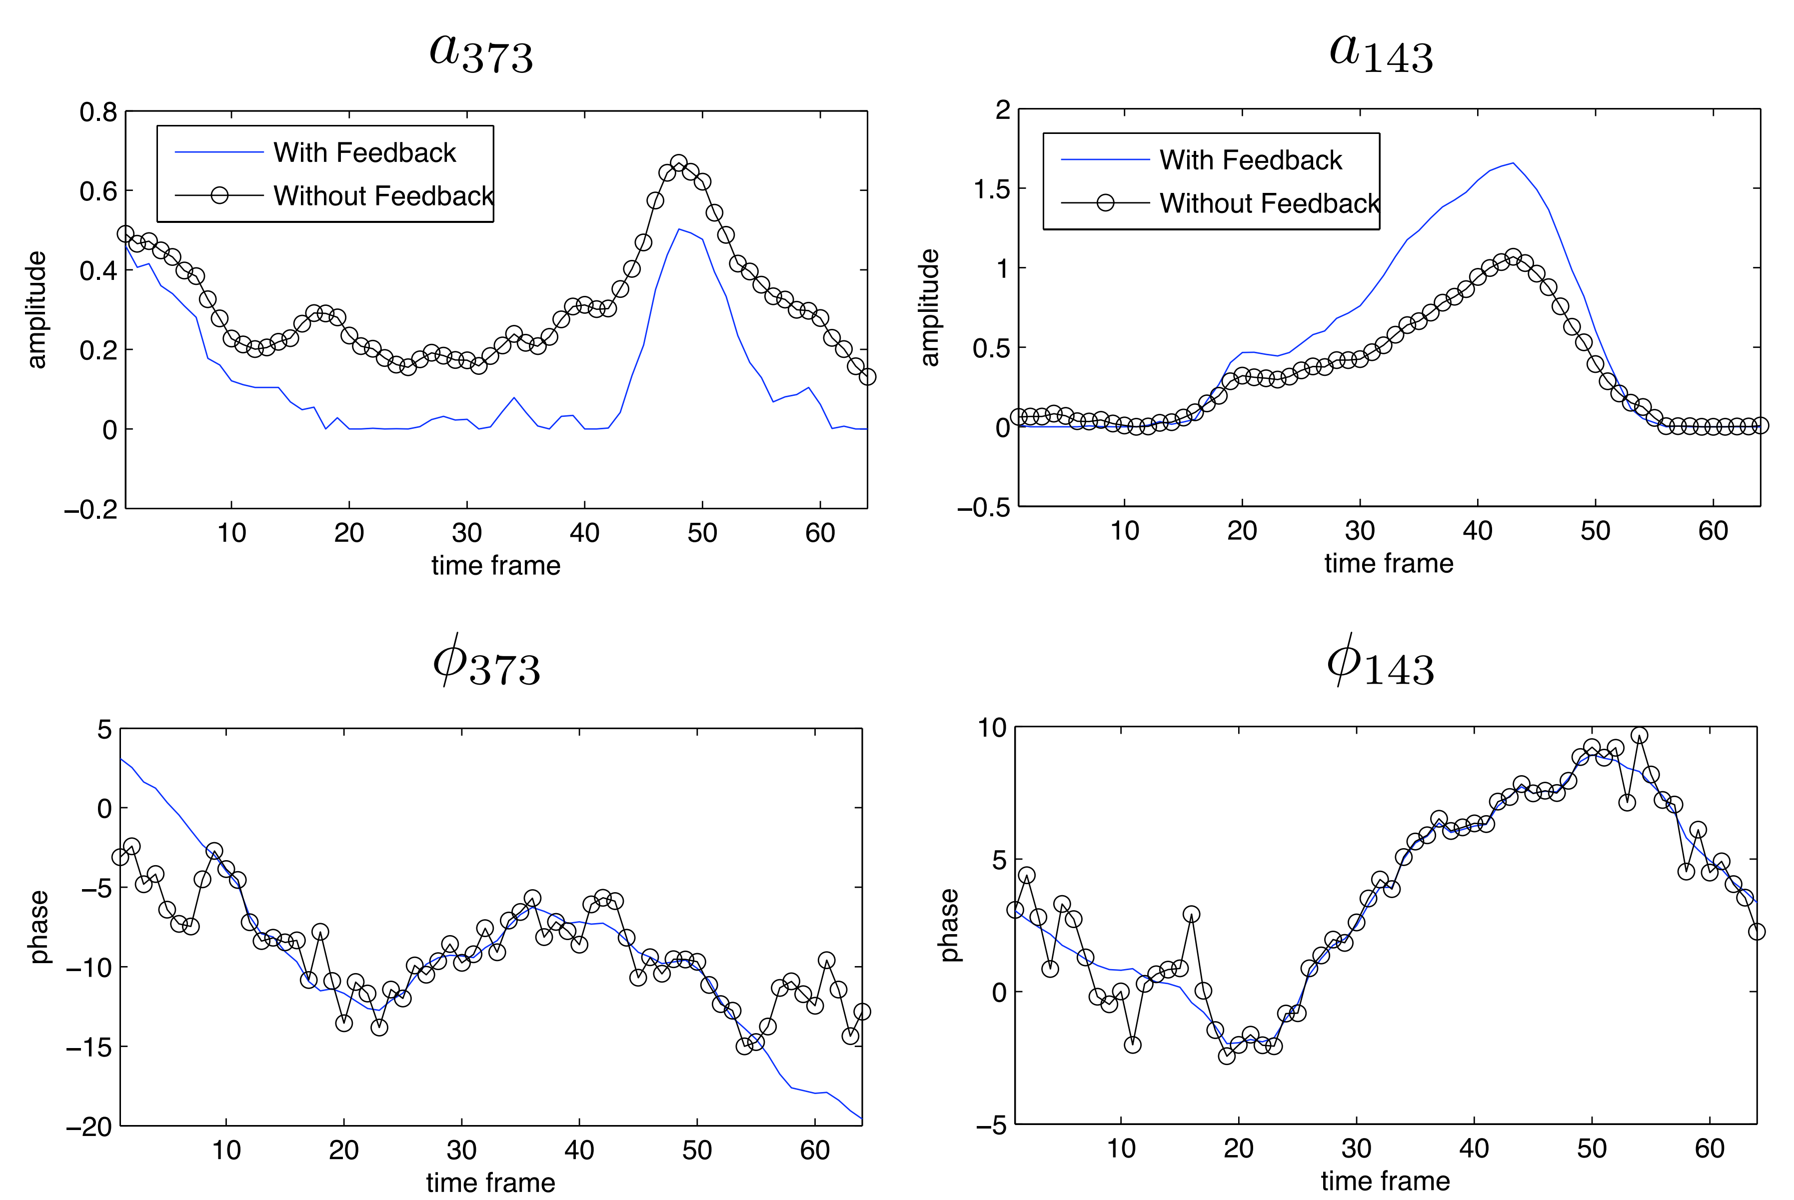

Charles Cadieu, cadieu@berkeley.edu