Statistical Analysis of Neural Data

Scott W. Linderman
Columbia University

April 29th, 2017

Chen et al., 2013

Ahrens et al., 2014

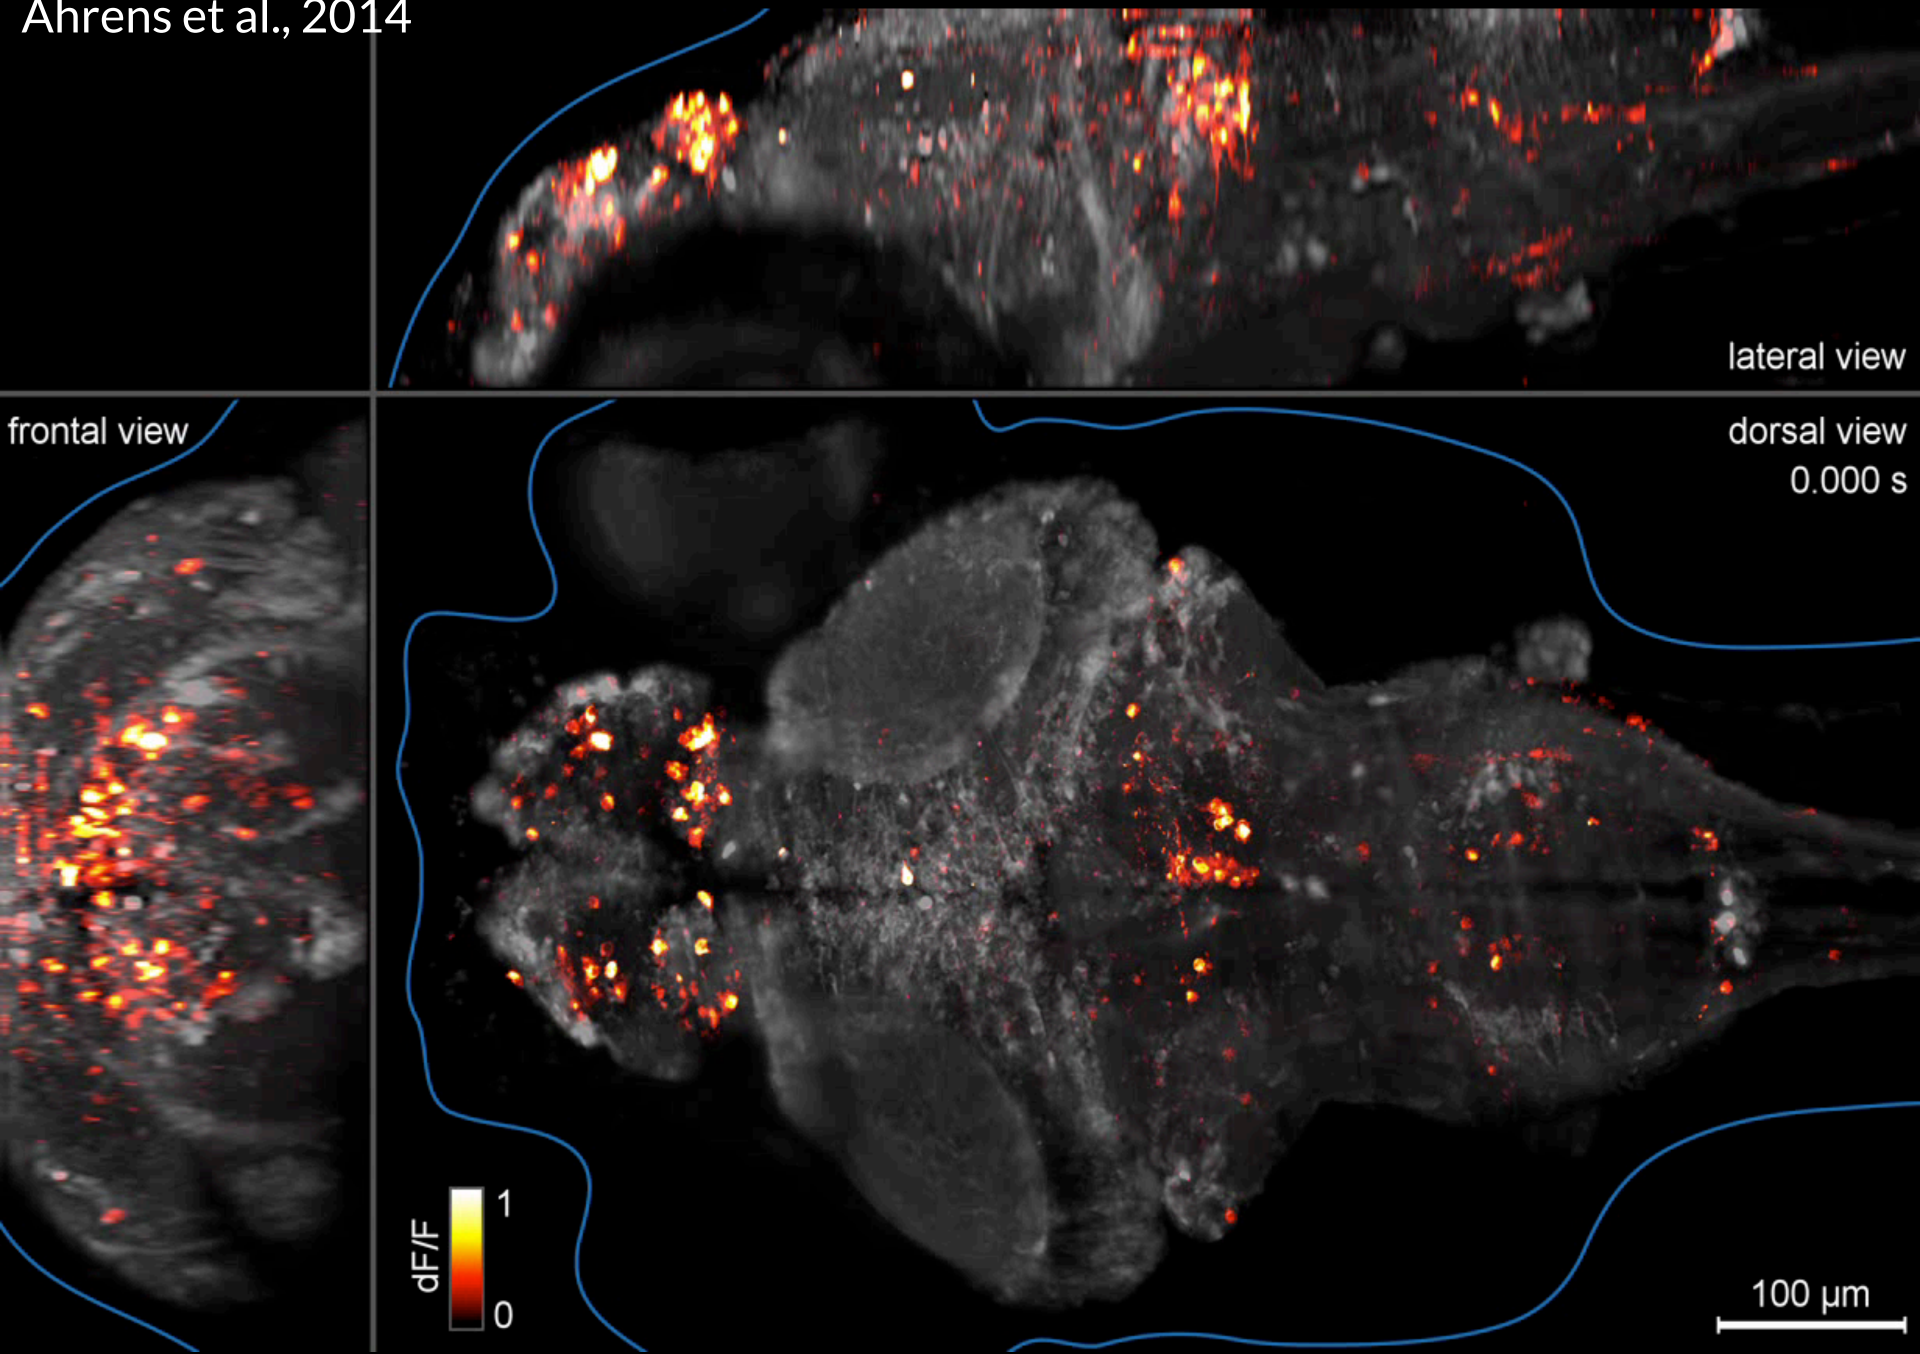

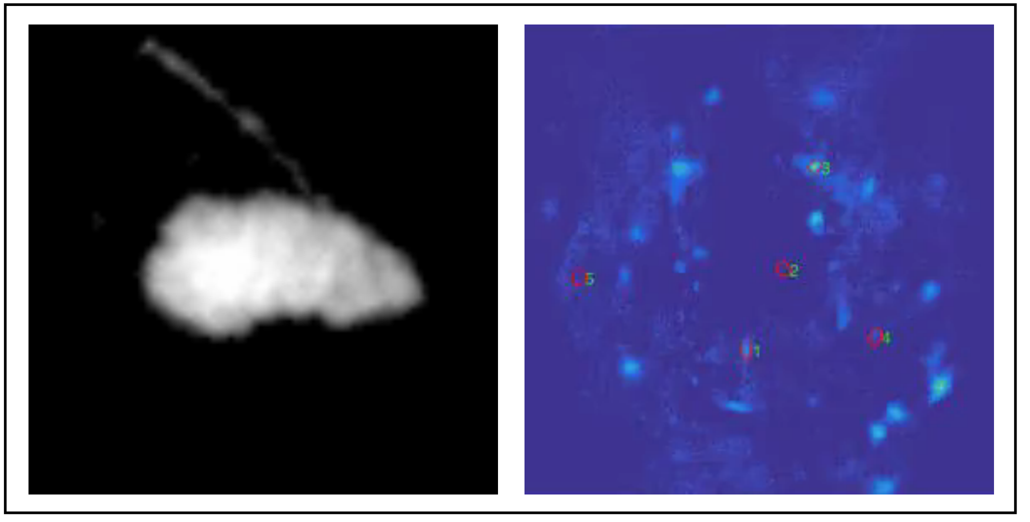

Datta Lab, Harvard Medical School

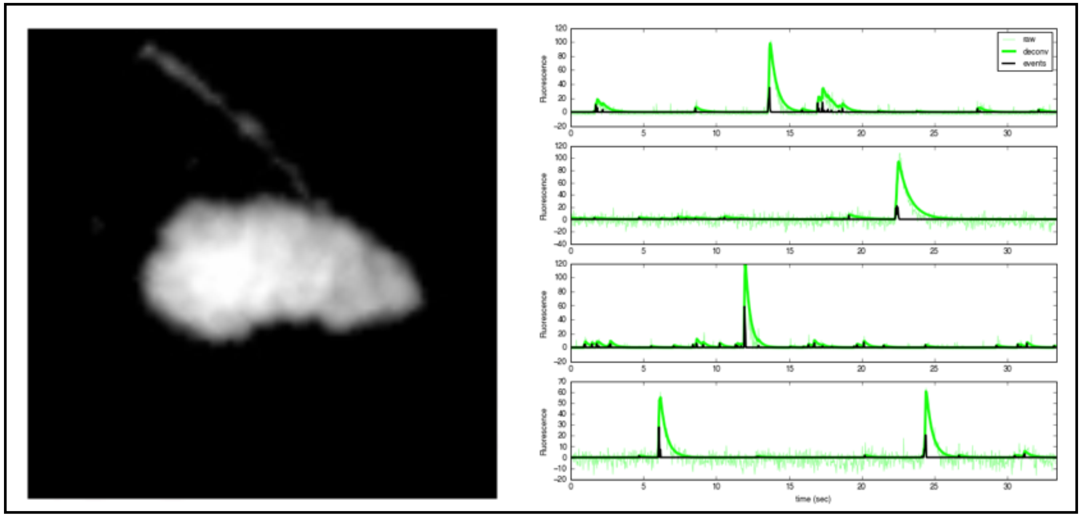

Datta Lab, Harvard Medical School

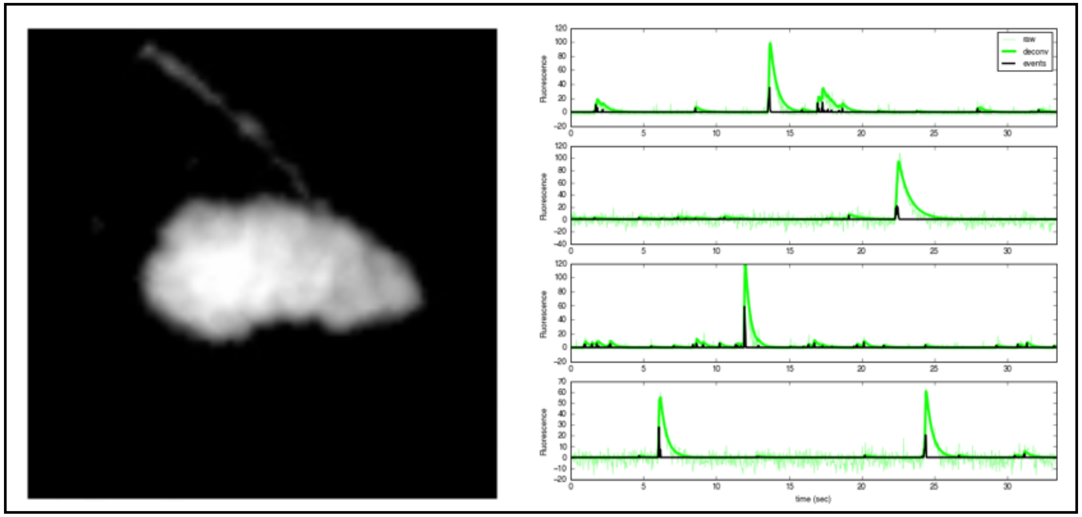

Datta Lab, Harvard Medical School

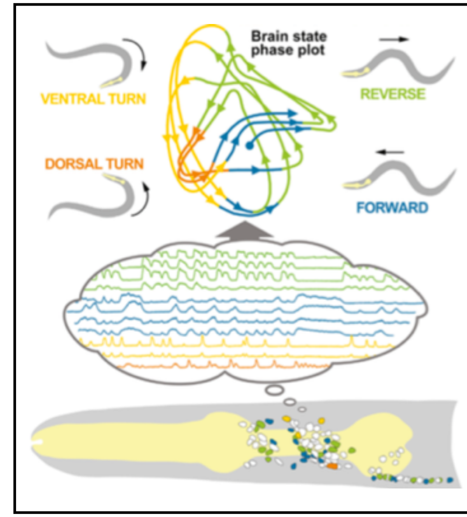

Zimmer Lab, IMP, Austria

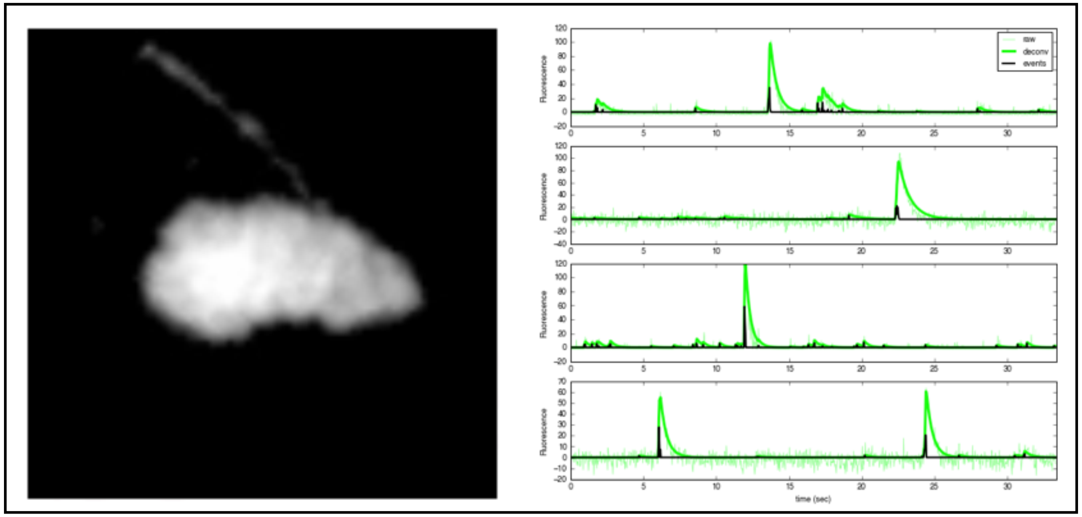

Datta Lab, Harvard Medical School

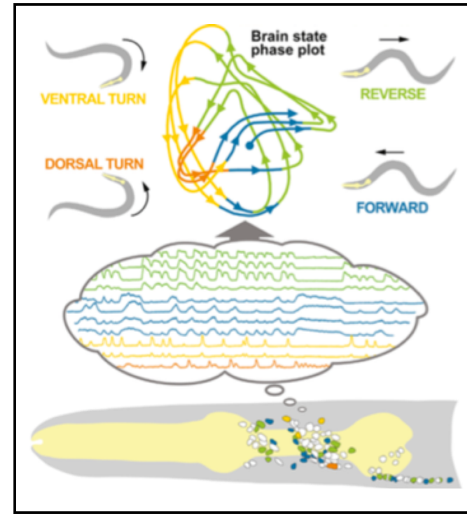

Zimmer Lab, IMP, Austria

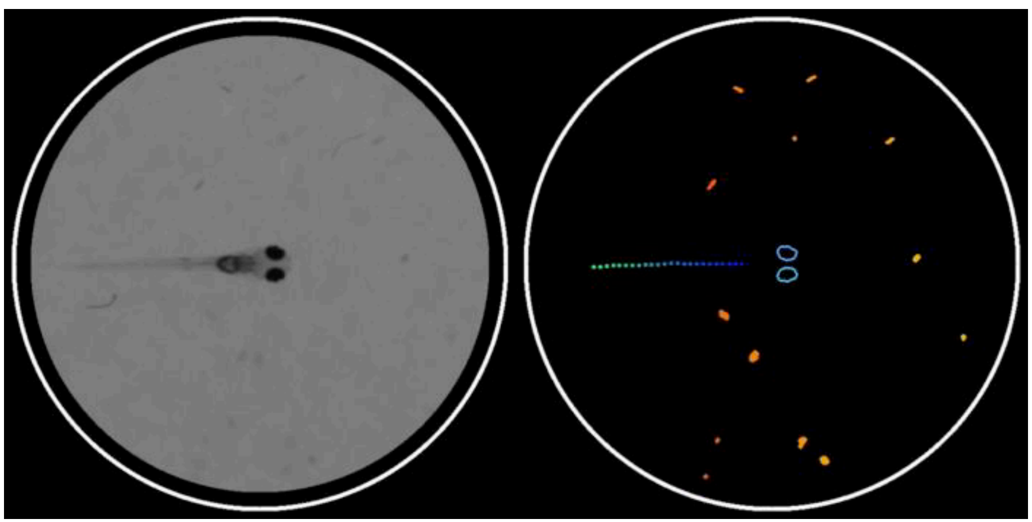

Engert Lab, Harvard

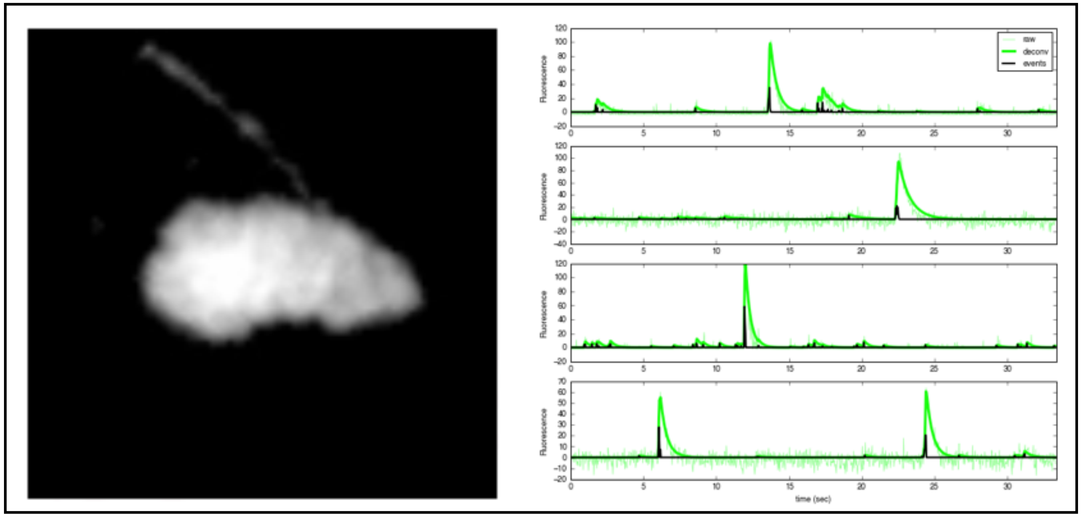

Datta Lab, Harvard Medical School

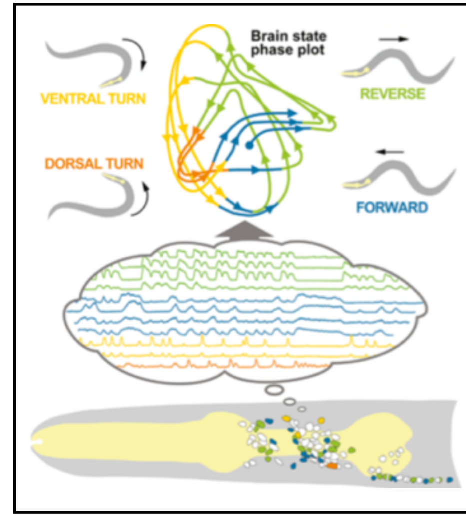

Zimmer Lab, IMP, Austria

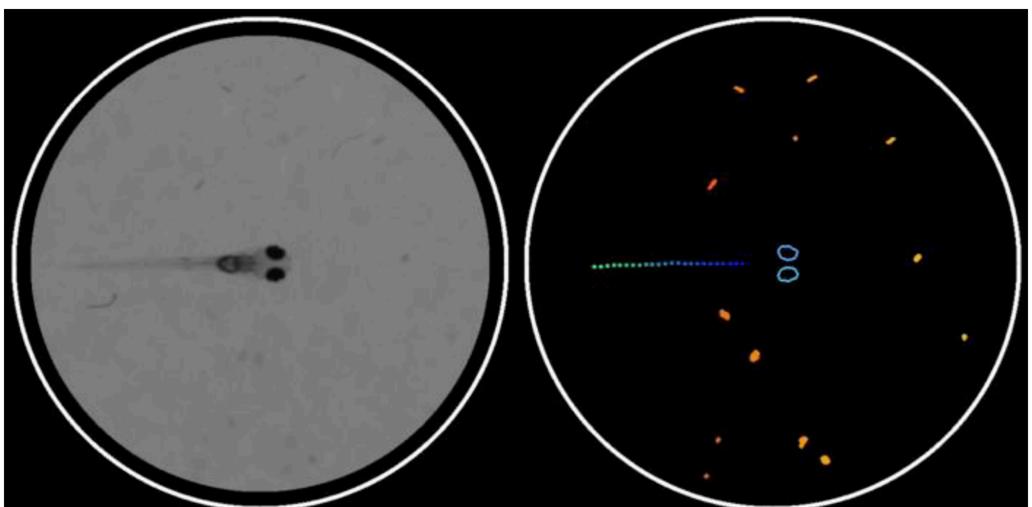

Engert Lab, Harvard

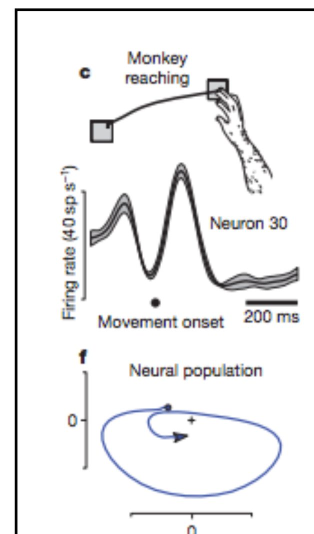

Churchland Lab, Columbia

High dimensional neural and behavioral data are the new norm in systems neuroscience.

These data often exhibit one or more of the following features:

1. *Low dimensional latent state;*
2. *Globally nonlinear dynamics;*
3. *Structure across dimensions and trials;*
4. *Missing or obscured data; and,*
5. *Count observations.*

A Probabilistic Approach

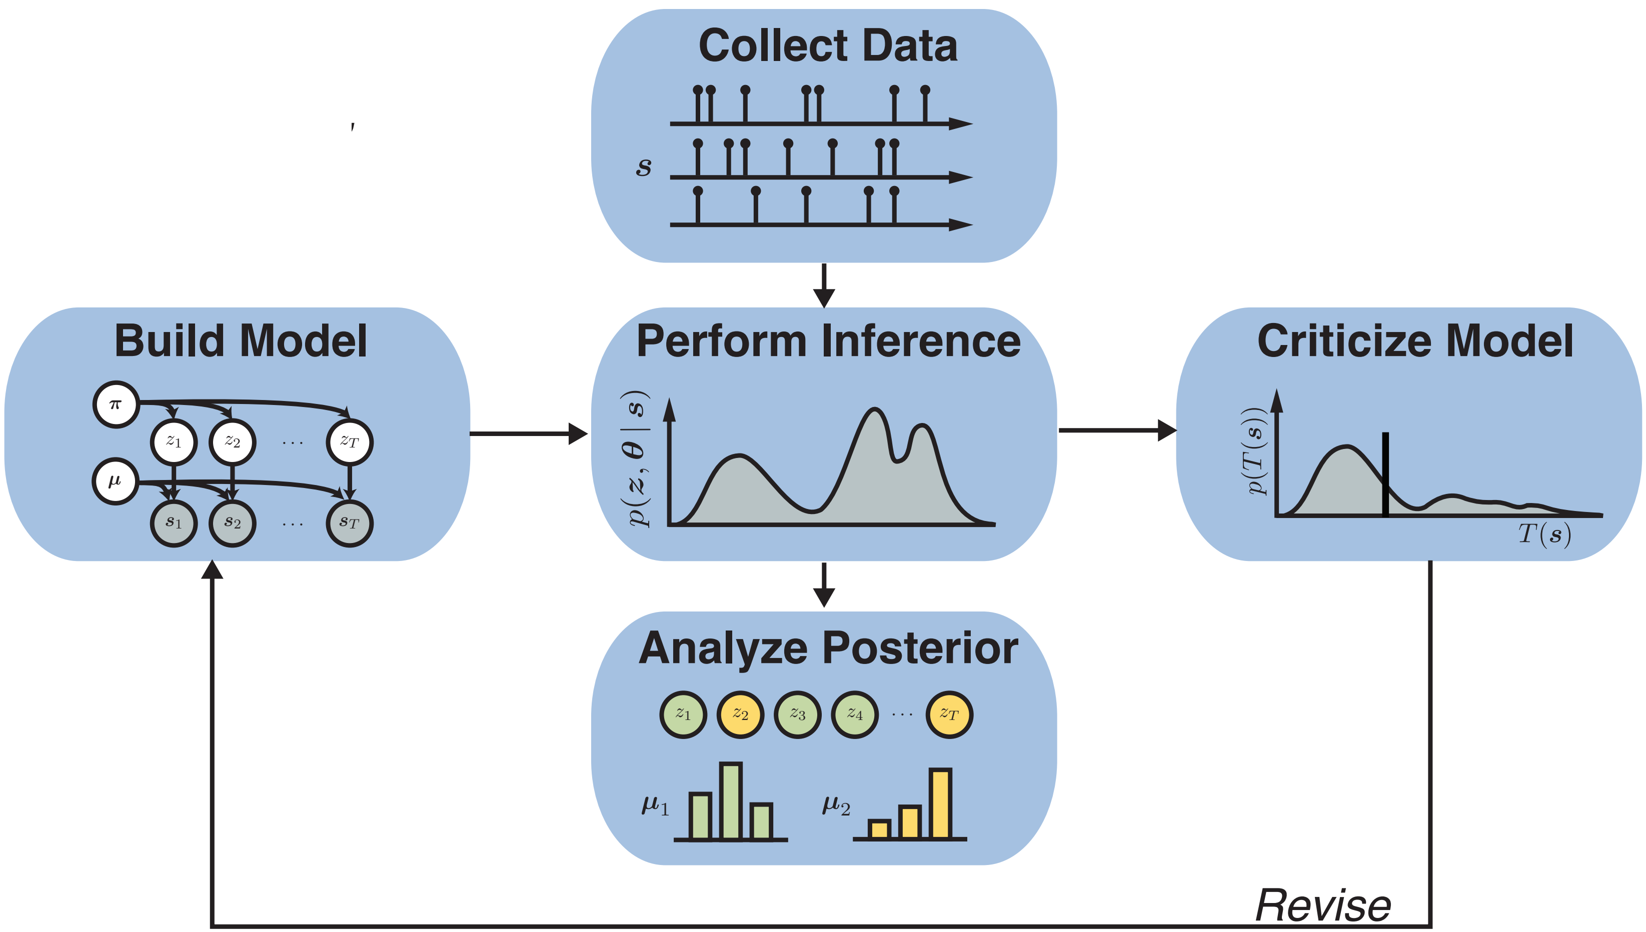

Adapted from Blei, 2014.

A Spectrum of Models

Limited capacity,
Specialized inference,
Data efficient,
Easy to fit and understand.

Highly flexible,
Generic inference,
Data intensive,
Harder to interpret.

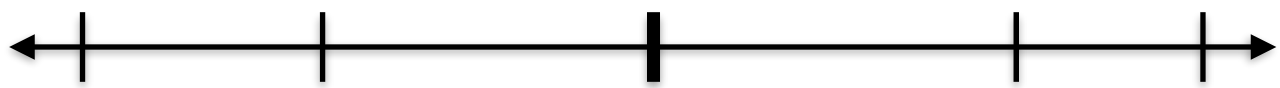

Principal Components Analysis,
Gaussian Mixture Models
Factor Analysis

State Space
Models

Recurrent
Neural Networks

Outline

- Background: State-space models
- Recurrent Switching Linear Dynamical Systems
- Whole-Brain Recordings of *C. Elegans*

State Space Models

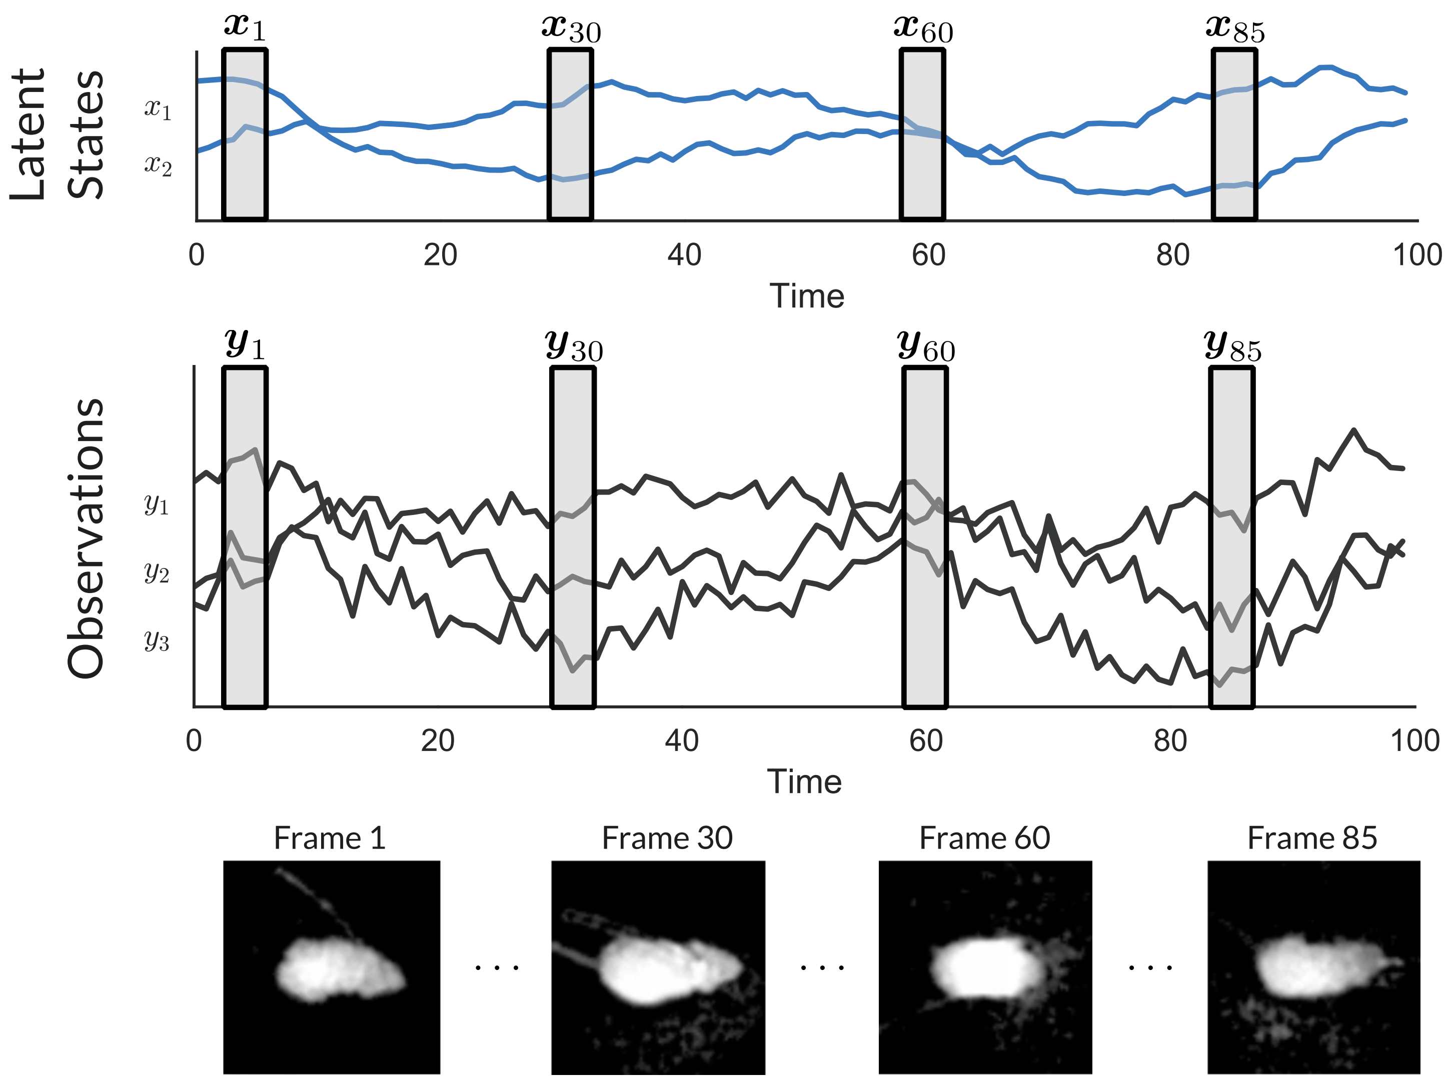

A Graphical Model of SSM's

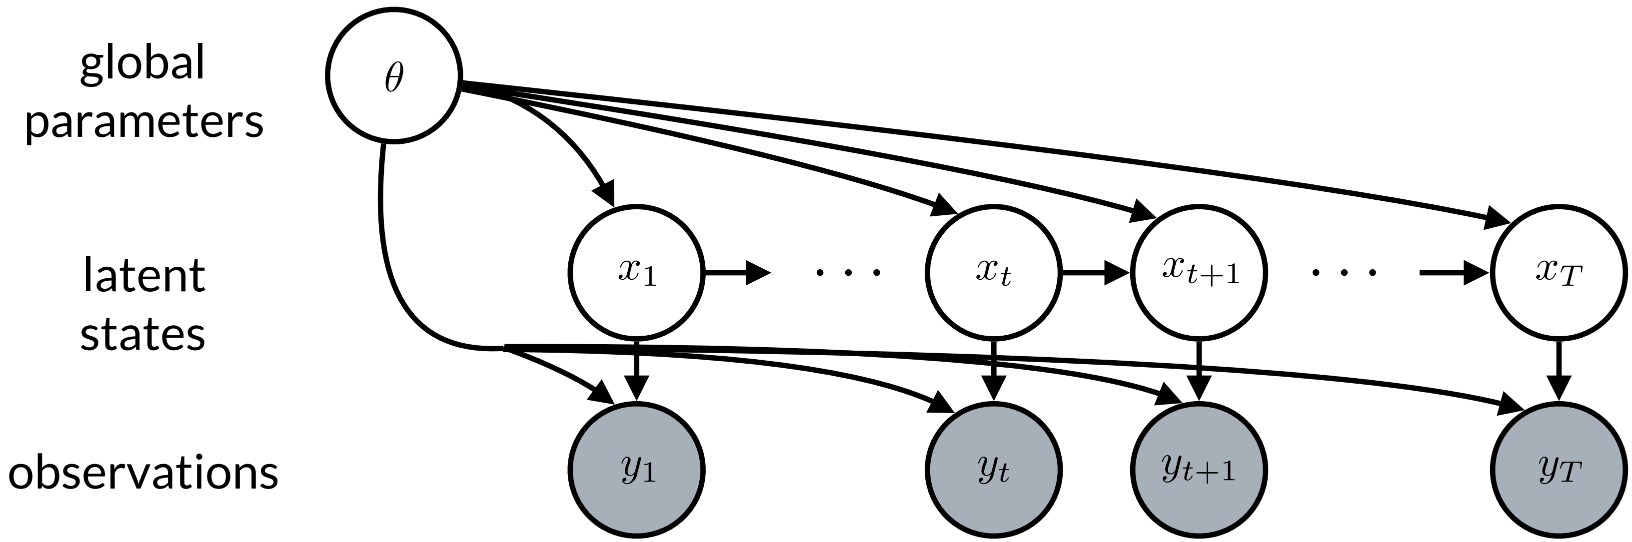

Linear Dynamical Systems

Latent
State Dynamics

Linear
Observation
Model

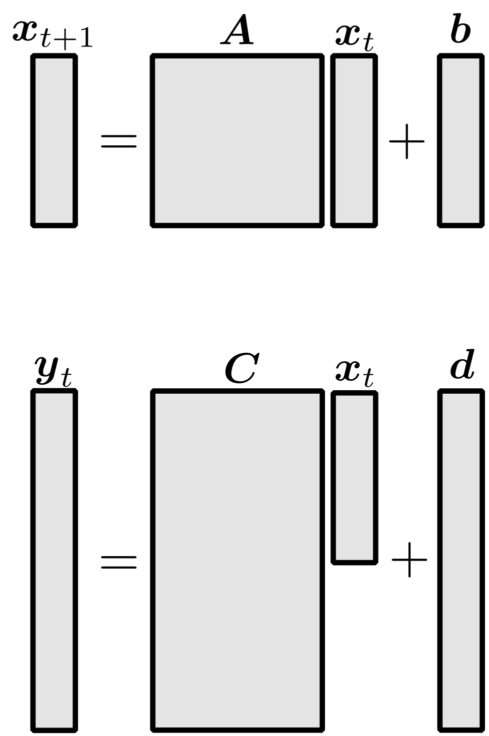

Visualizing Linear Observations

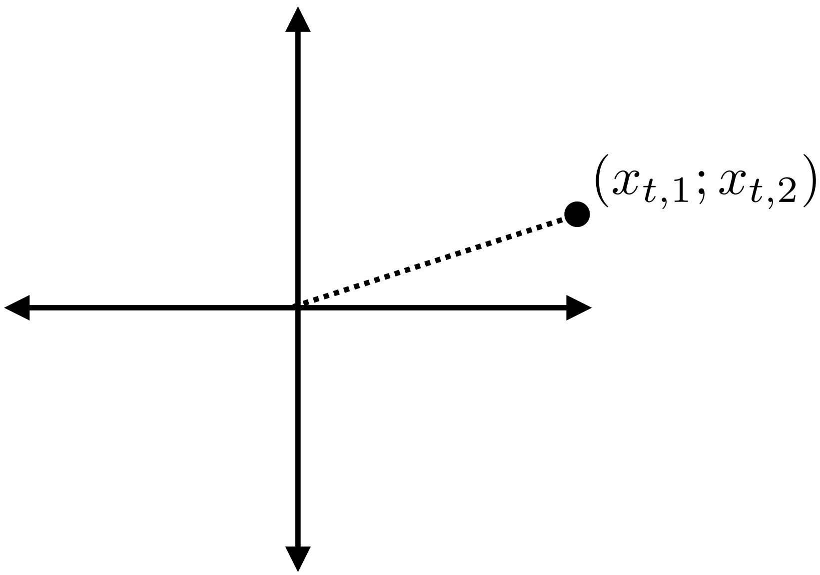

$$\boldsymbol{y}_t = \boldsymbol{C}\boldsymbol{x}_t + \boldsymbol{d}$$

$$\begin{bmatrix} y_{t,1} \\ y_{t,2} \\ y_{t,3} \\ \vdots \\ y_{t,N} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \\ \cdots & \\ c_{N1} & c_{N2} \end{bmatrix} \begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$= \begin{bmatrix} c_{11}x_{t,1} + c_{12}x_{t,2} \\ c_{21}x_{t,1} + c_{21}x_{t,2} \\ c_{31}x_{t,1} + c_{31}x_{t,2} \\ \vdots \\ c_{N1}x_{t,1} + c_{N2}x_{t,2} \end{bmatrix}$$

Visualizing Linear Observations

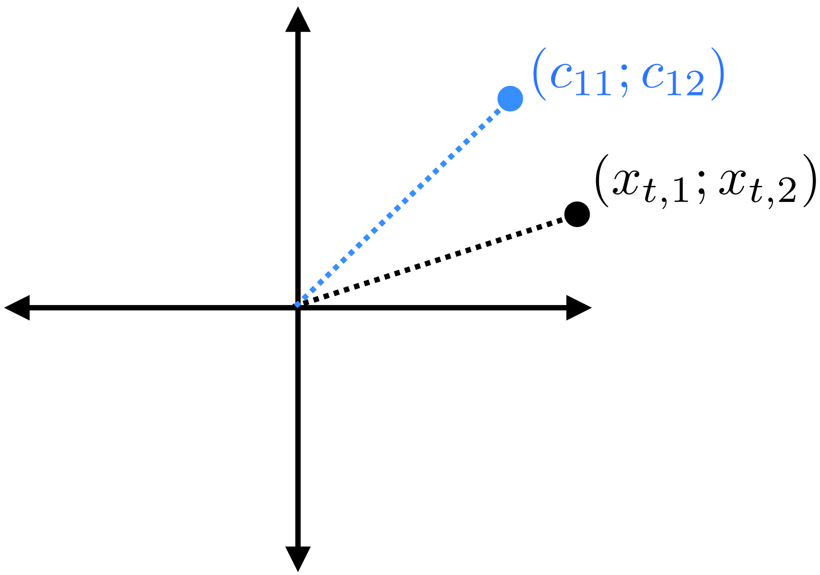

$$\boldsymbol{y}_t = \boldsymbol{C}\boldsymbol{x}_t + \boldsymbol{d}$$

$$\begin{bmatrix} y_{t,1} \\ y_{t,2} \\ y_{t,3} \\ \vdots \\ y_{t,N} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \\ \cdots & \\ c_{N1} & c_{N2} \end{bmatrix} \begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$= \begin{bmatrix} c_{11}x_{t,1} + c_{12}x_{t,2} \\ c_{21}x_{t,1} + c_{21}x_{t,2} \\ c_{31}x_{t,1} + c_{31}x_{t,2} \\ \vdots \\ c_{N1}x_{t,1} + c_{N2}x_{t,2} \end{bmatrix}$$

Visualizing Linear Observations

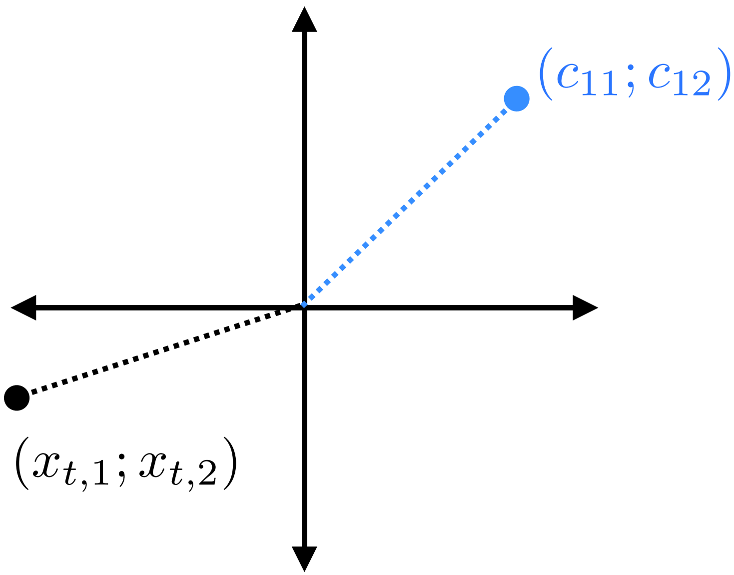

$$\boldsymbol{y}_t = \boldsymbol{C}\boldsymbol{x}_t + \boldsymbol{d}$$

$$\begin{bmatrix} y_{t,1} \\ y_{t,2} \\ y_{t,3} \\ \vdots \\ y_{t,N} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \\ \cdots & \\ c_{N1} & c_{N2} \end{bmatrix} \begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$= \begin{bmatrix} c_{11}x_{t,1} + c_{12}x_{t,2} \\ c_{21}x_{t,1} + c_{21}x_{t,2} \\ c_{31}x_{t,1} + c_{31}x_{t,2} \\ \vdots \\ c_{N1}x_{t,1} + c_{N2}x_{t,2} \end{bmatrix}$$

Visualizing Linear Observations

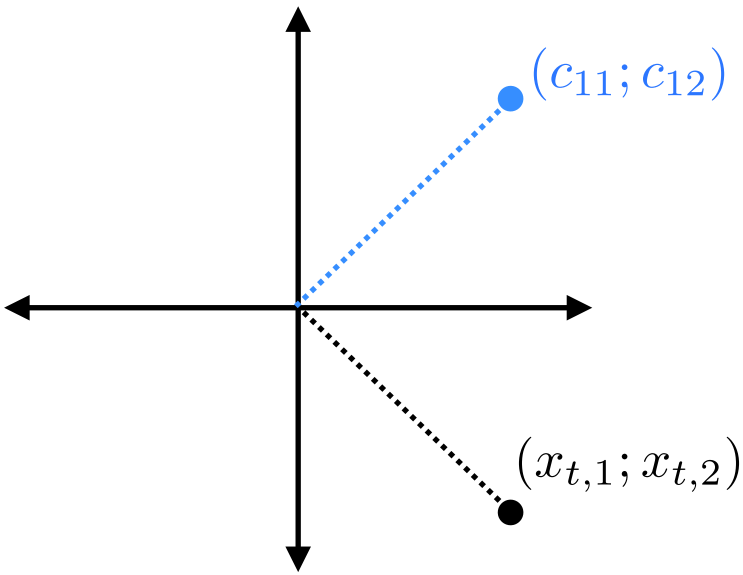

$$\boldsymbol{y}_t = \boldsymbol{C}\boldsymbol{x}_t + \boldsymbol{d}$$

$$\begin{bmatrix} y_{t,1} \\ y_{t,2} \\ y_{t,3} \\ \vdots \\ y_{t,N} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \\ \cdots & \\ c_{N1} & c_{N2} \end{bmatrix} \begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$= \begin{bmatrix} c_{11}x_{t,1} + c_{12}x_{t,2} \\ c_{21}x_{t,1} + c_{21}x_{t,2} \\ c_{31}x_{t,1} + c_{31}x_{t,2} \\ \vdots \\ c_{N1}x_{t,1} + c_{N2}x_{t,2} \end{bmatrix}$$

Visualizing Linear Observations

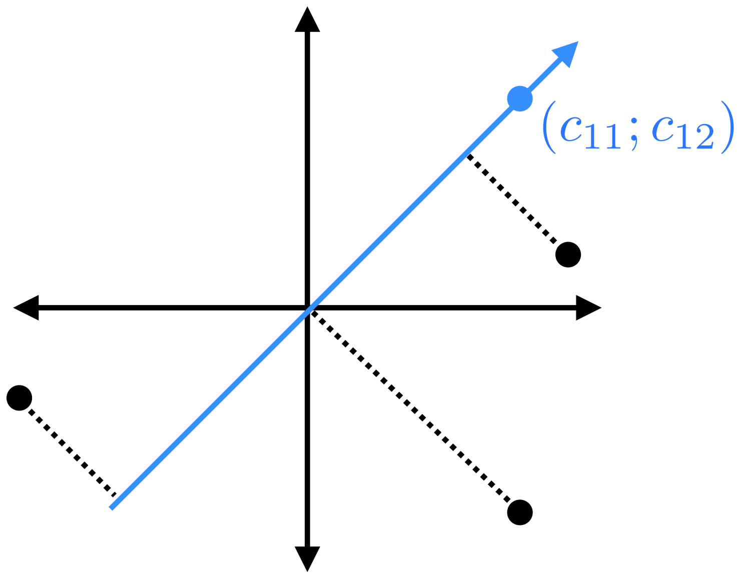

$$\boldsymbol{y}_t = \boldsymbol{C}\boldsymbol{x}_t + \boldsymbol{d}$$

$$\begin{bmatrix} y_{t,1} \\ y_{t,2} \\ y_{t,3} \\ \vdots \\ y_{t,N} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \\ \cdots & \\ c_{N1} & c_{N2} \end{bmatrix} \begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$= \begin{bmatrix} c_{11}x_{t,1} + c_{12}x_{t,2} \\ c_{21}x_{t,1} + c_{21}x_{t,2} \\ c_{31}x_{t,1} + c_{31}x_{t,2} \\ \vdots \\ c_{N1}x_{t,1} + c_{N2}x_{t,2} \end{bmatrix}$$

Visualizing Linear Observations

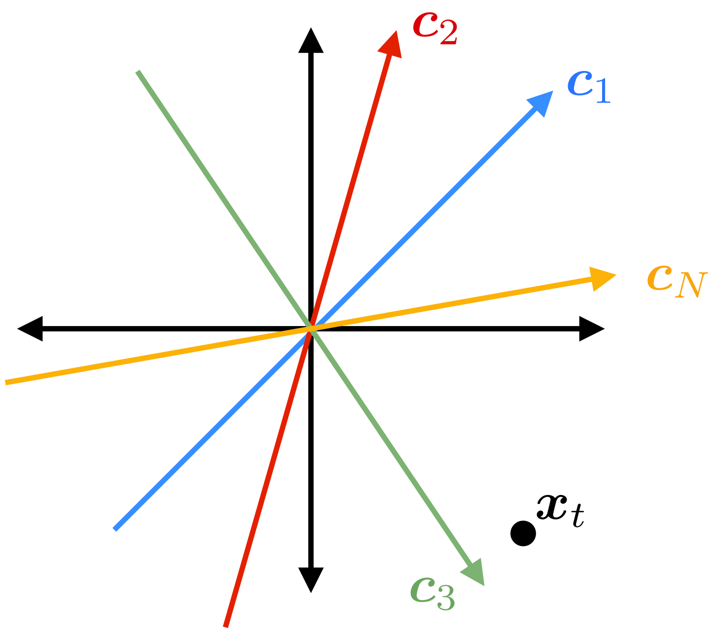

$$\boldsymbol{y}_t = \boldsymbol{C}\boldsymbol{x}_t + \boldsymbol{d}$$

$$\begin{bmatrix} y_{t,1} \\ y_{t,2} \\ y_{t,3} \\ \vdots \\ y_{t,N} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \\ c_{31} & c_{32} \\ \cdots & \\ c_{N1} & c_{N2} \end{bmatrix} \begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$= \begin{bmatrix} c_{11}x_{t,1} + c_{12}x_{t,2} \\ c_{21}x_{t,1} + c_{21}x_{t,2} \\ c_{31}x_{t,1} + c_{31}x_{t,2} \\ \vdots \\ c_{N1}x_{t,1} + c_{N2}x_{t,2} \end{bmatrix}$$

Linear Dynamical Systems

Latent
State Dynamics

Linear
Observation
Model

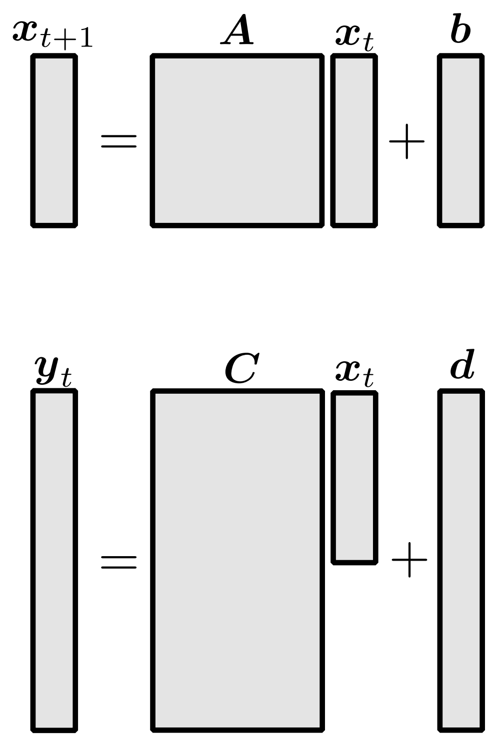

Visualizing Linear Dynamics

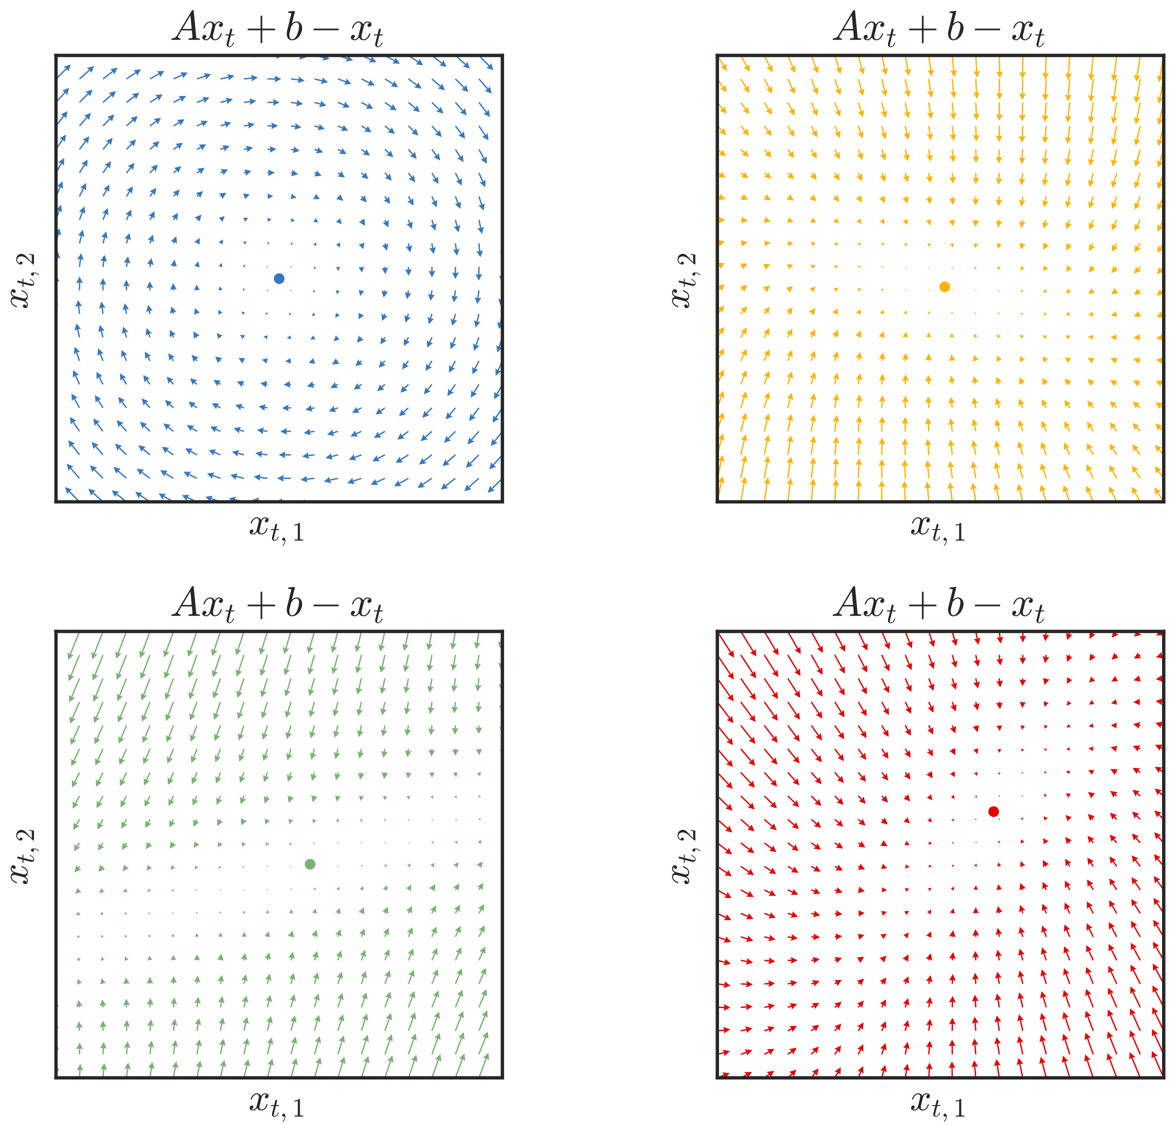

Rotational Dynamics

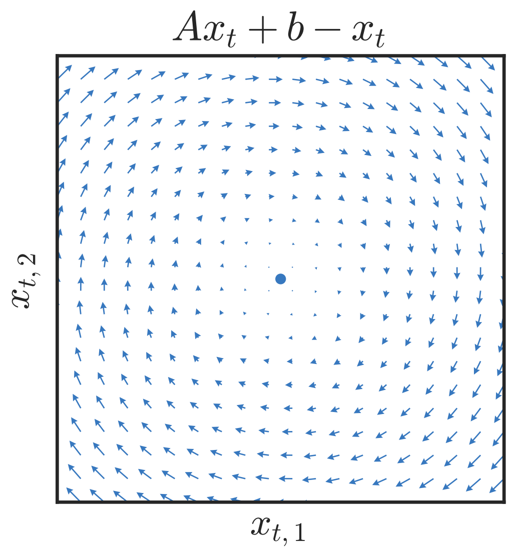

$$A = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \quad x_t = \begin{bmatrix} r\cos\phi \\ r\sin\phi \end{bmatrix} \quad b = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$\begin{aligned} Ax_t + b &= \begin{bmatrix} r\cos\theta\cos\phi - r\sin\theta\sin\phi \\ r\sin\theta\cos\phi + r\cos\theta\sin\phi \end{bmatrix} \\ &= \begin{bmatrix} r\cos(\theta+\phi) \\ r\sin(\theta+\phi) \end{bmatrix} \end{aligned}$$

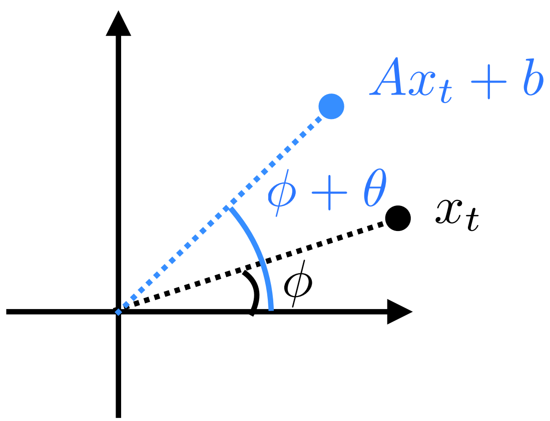

Gaussian Linear Dynamical System

Latent
State Dynamics

Observation
Model

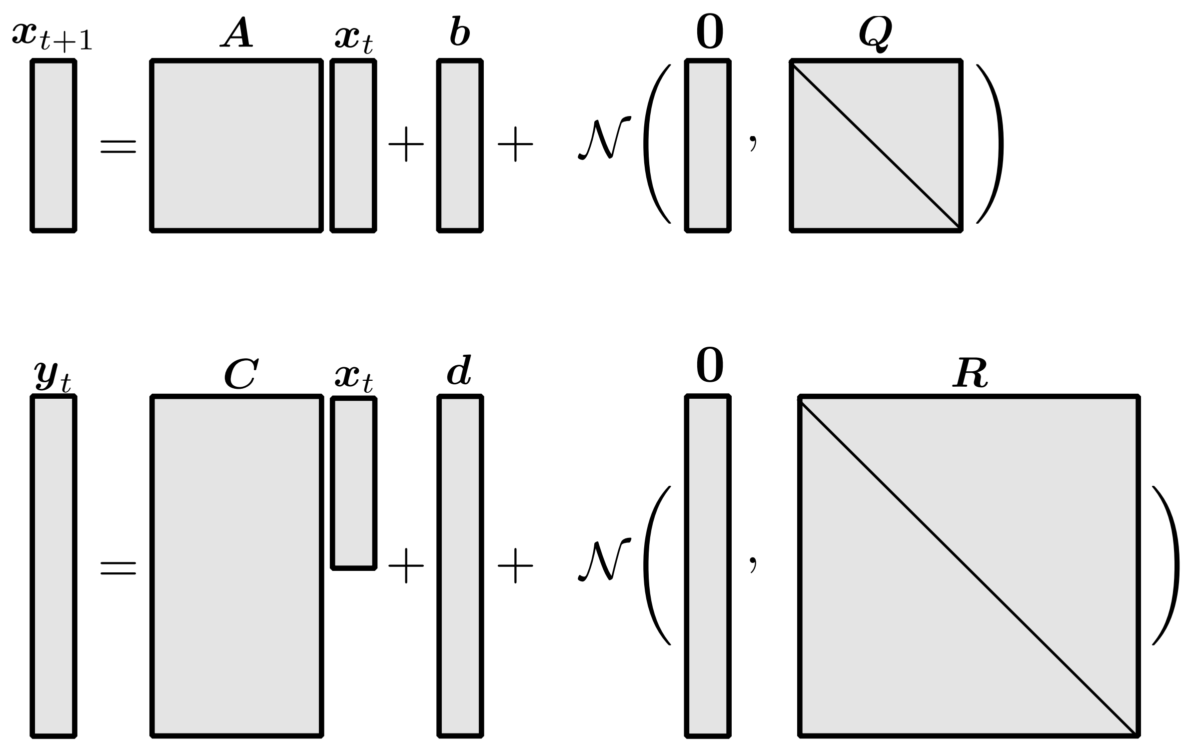

Multivariate Gaussian Distribution

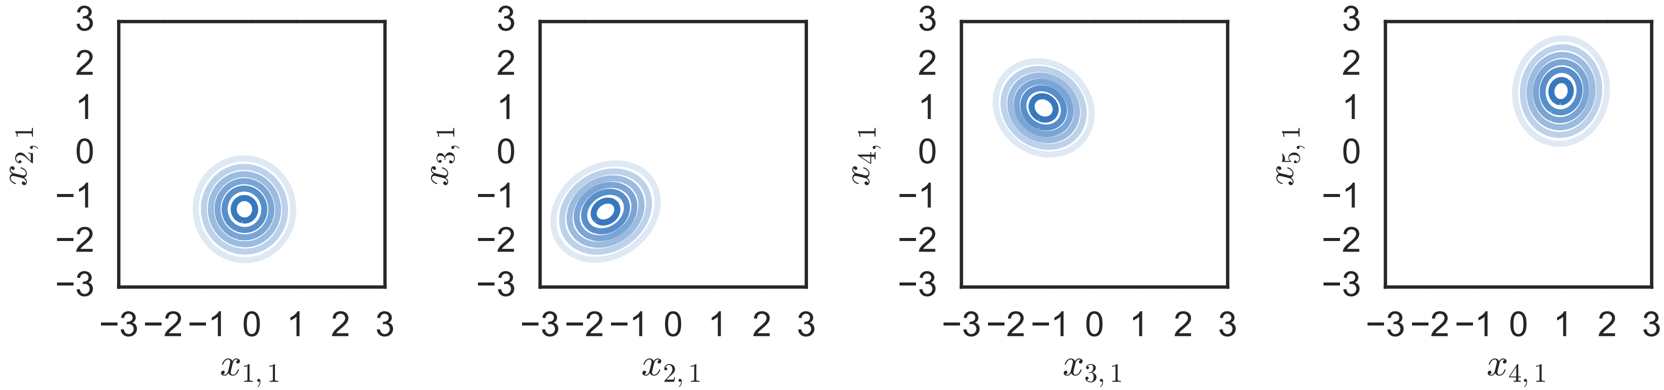

Switching Linear Dynamical Systems

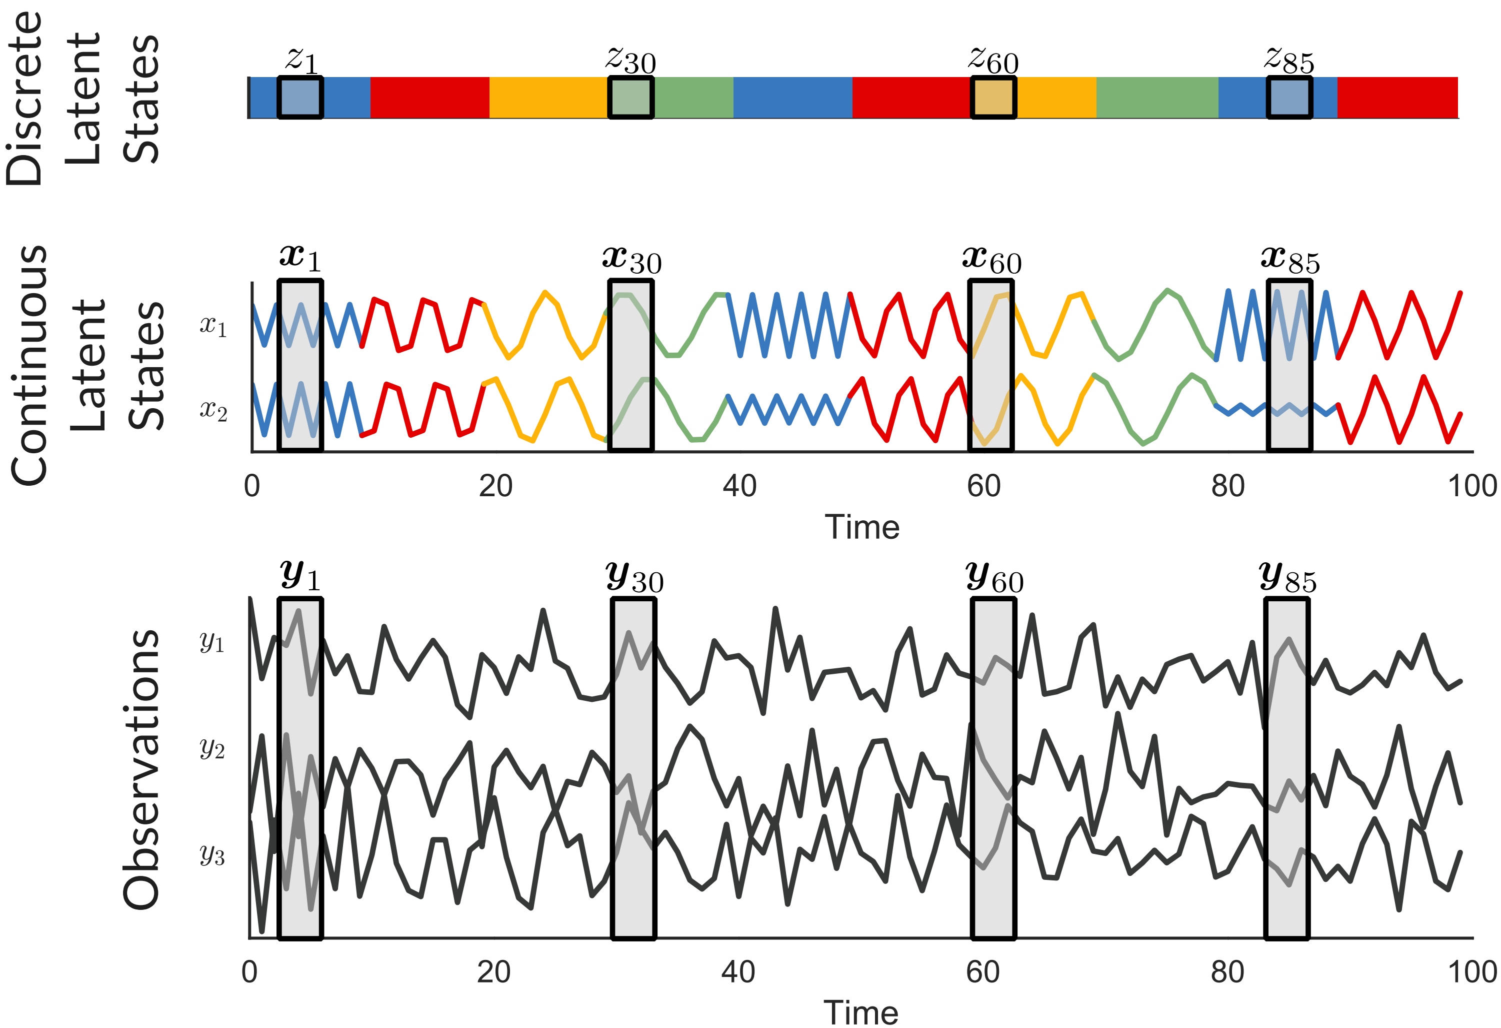

Ghahramani and Hinton (1996); Murphy (1998); Murphy (2013)

A Graphical Model of SLDS

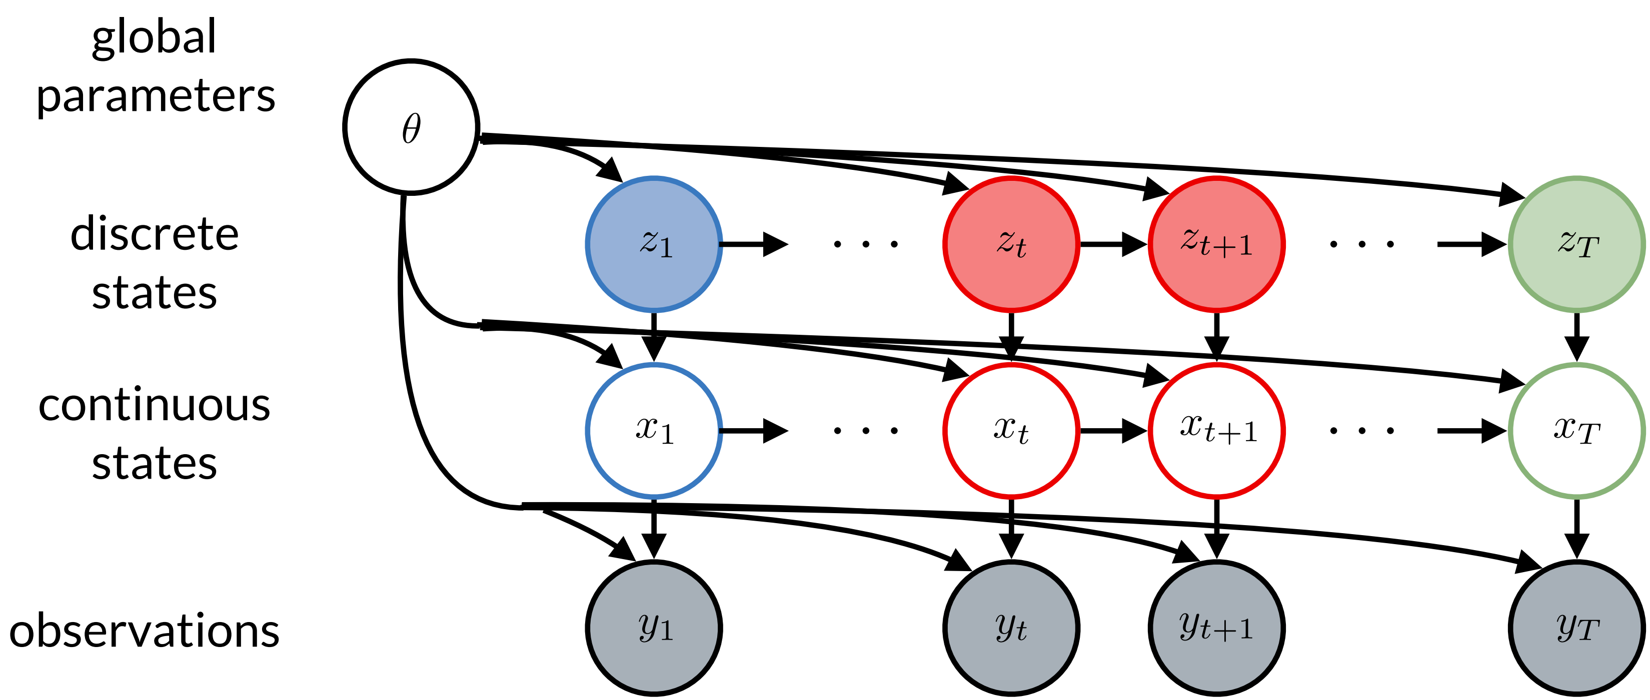

Switching Linear Dynamical Systems

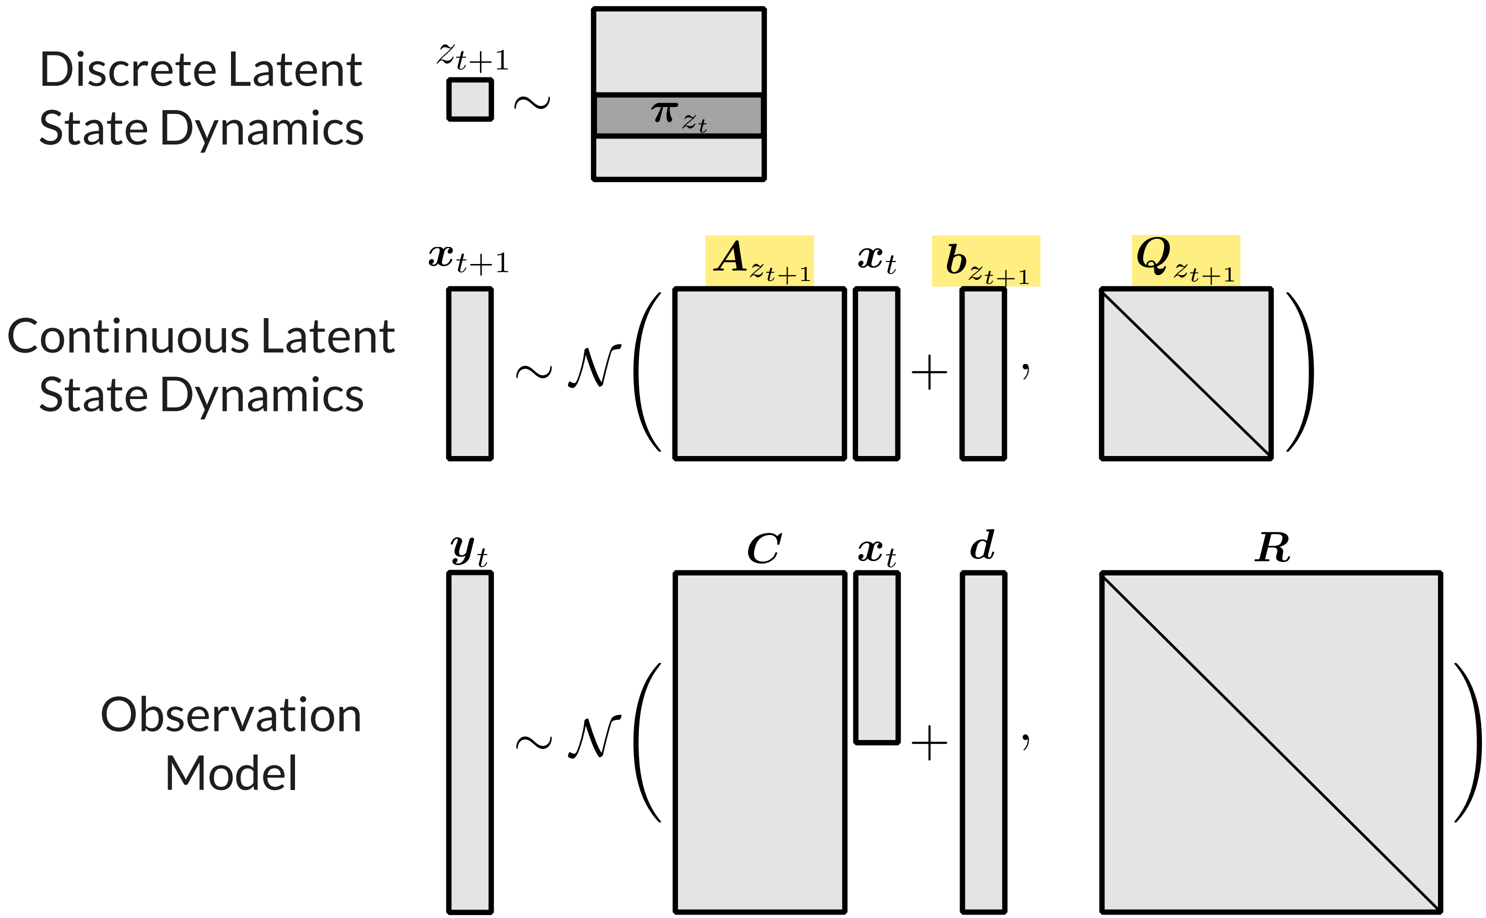

Discrete Markovian Dynamics

Discrete Latent State Dynamics

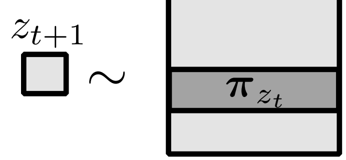

$$\boldsymbol{P} = \begin{bmatrix} - & \boldsymbol{\pi}_1 & - \\ - & \boldsymbol{\pi}_2 & - \\ & \cdots & \\ - & \boldsymbol{\pi}_K & - \end{bmatrix}$$

$$\text{e.g. } \boldsymbol{P} = \begin{bmatrix} 0.8 & 0.2 & 0.0 & 0.0 \\ 0.0 & 0.8 & 0.2 & 0.0 \\ 0.0 & 0.0 & 0.8 & 0.2 \\ 0.2 & 0.0 & 0.2 & 0.8 \end{bmatrix}$$

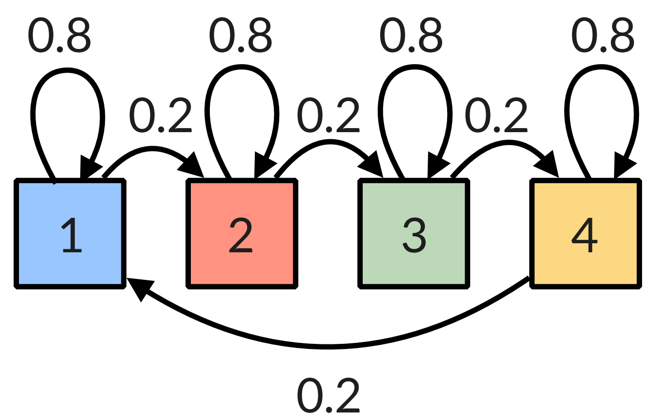

SLDS can Approximate Nonlinear Dynamics

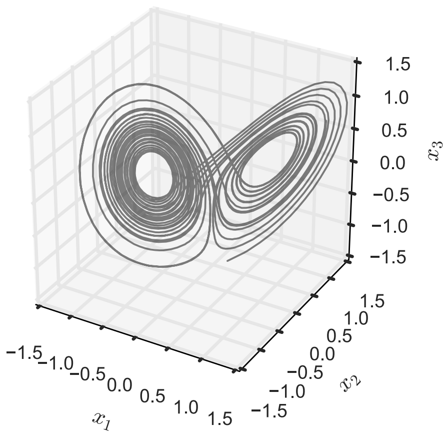

Lorenz Attractor: classic example of a nonlinear dynamical system.

$$\boldsymbol{x}_{t+1} = f(\boldsymbol{x}_t)$$

$$\begin{bmatrix} x_{t+1,1} \\ x_{t+1,2} \\ x_{t+1,3} \end{bmatrix} = \begin{bmatrix} (1-\alpha)x_{t,1} + \alpha x_{t,2} \\ x_{t,1}\beta - x_{t,1}x_{t,3} \\ (1-\gamma)x_{t,3} + x_{t,1}x_{t,2} \end{bmatrix}$$

SLDS can Approximate Nonlinear Dynamics

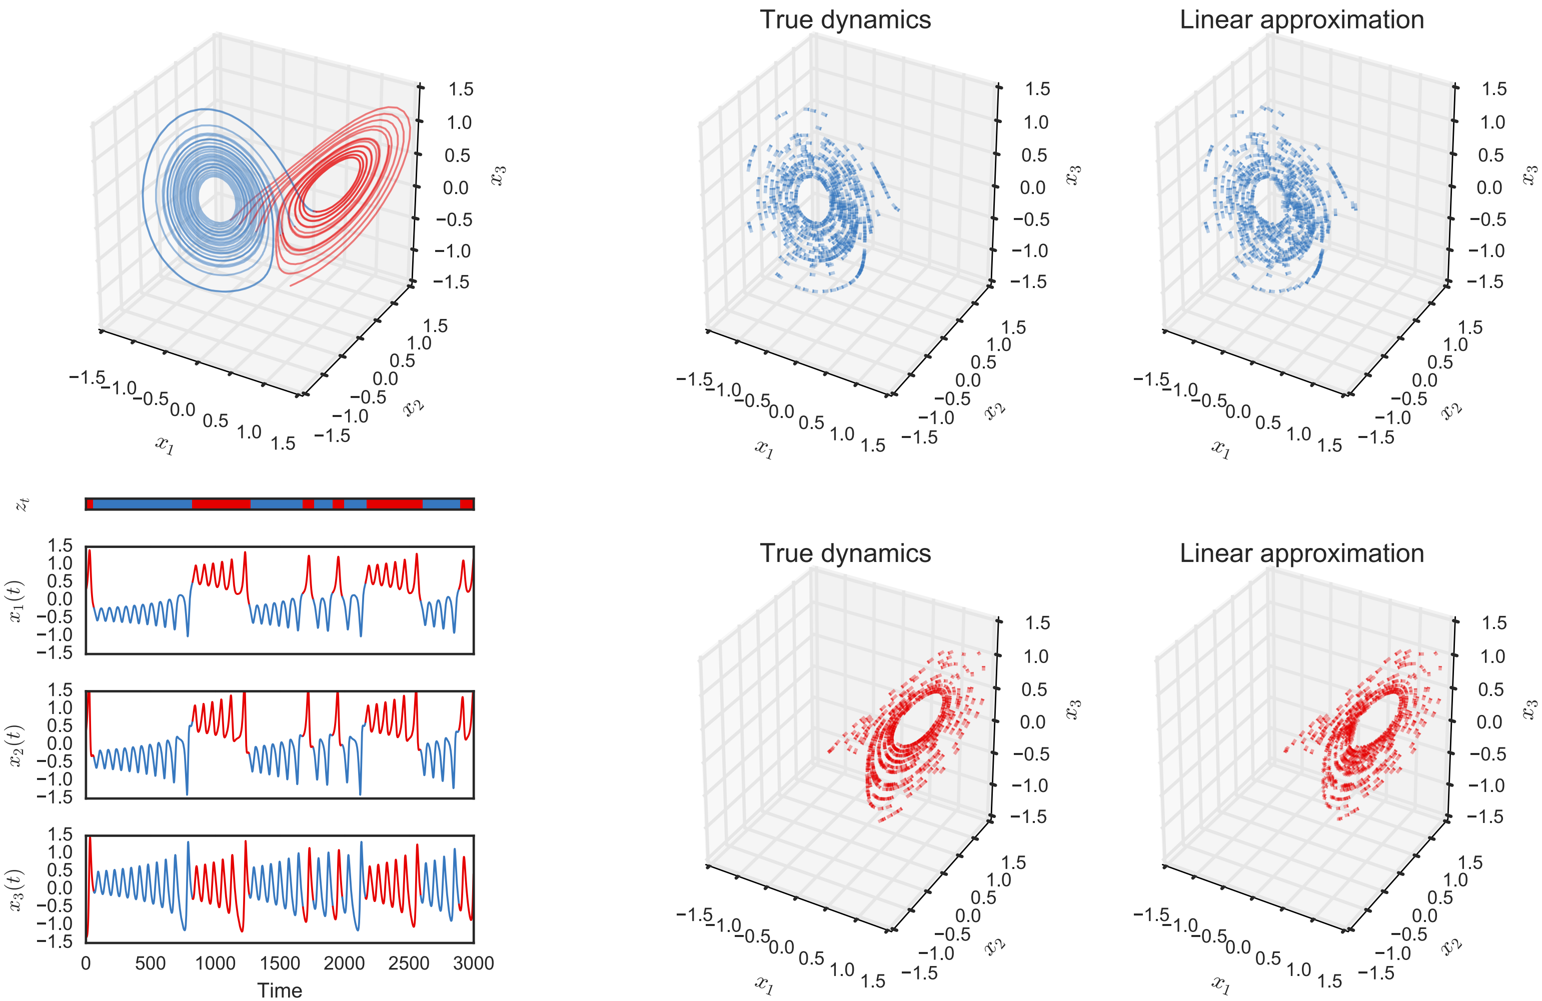

SLDS are Natural Models for Behavior

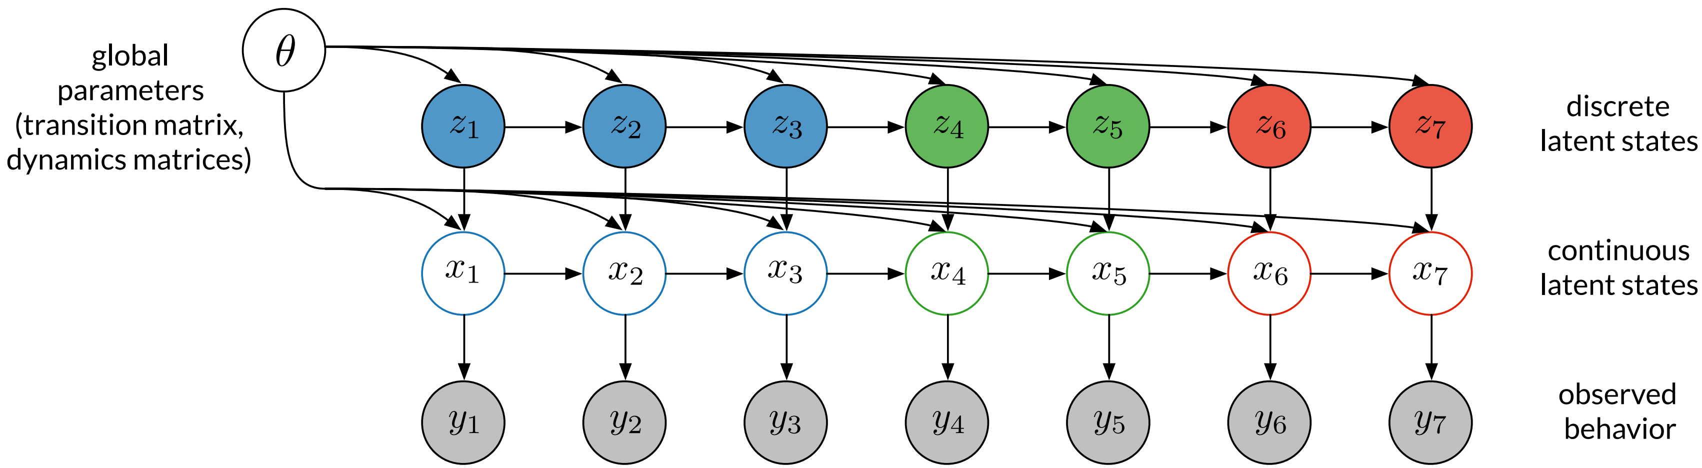

Kato et al., Cell 2015

Outline

- Background: Bayesian state-space models
- Recurrent Switching Linear Dynamical Systems
- Whole-Brain Recordings of *C. Elegans*

“Closed Loop” Models of Discrete Dynamics

The continuous and discrete states should be recurrently connected so that location in continuous space influences the distribution over next discrete states.

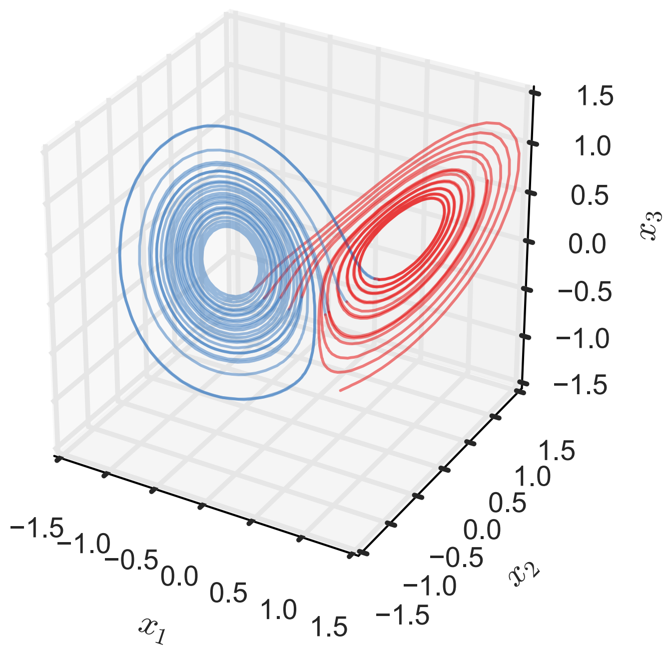

Recurrent Switching Linear Dynamical Systems (rSLDS)

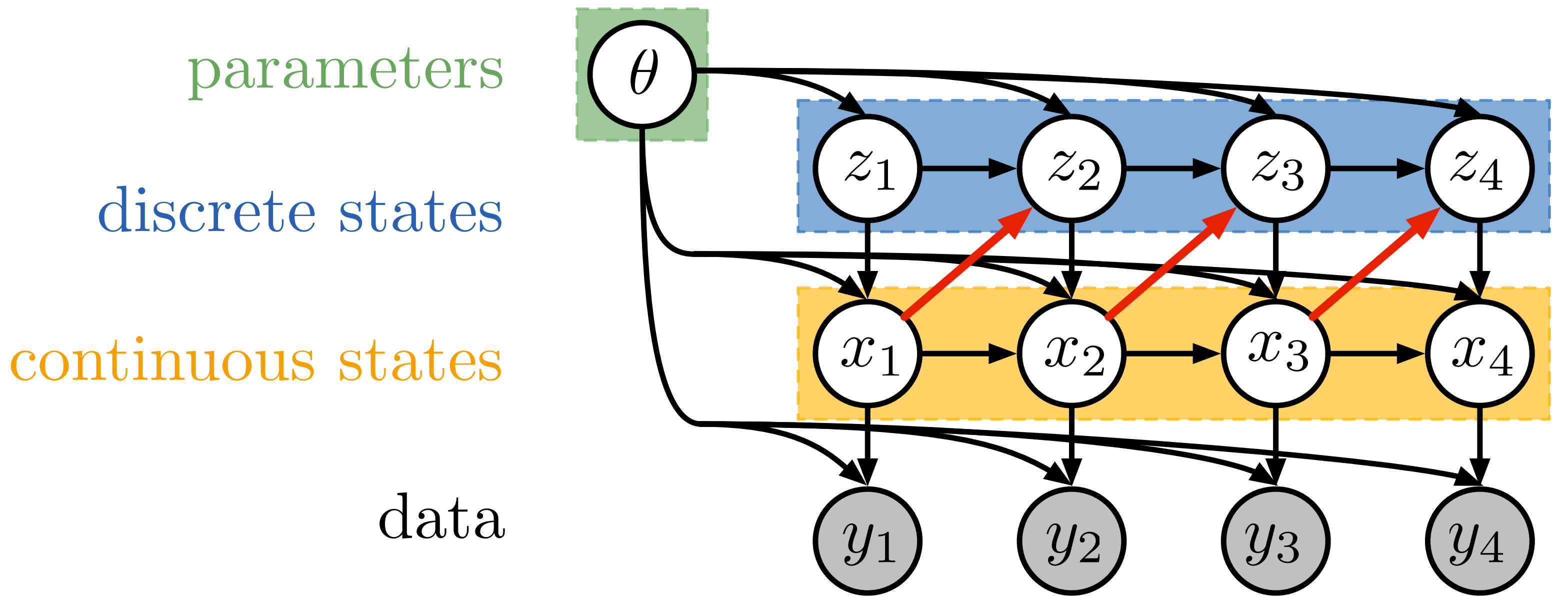

Also known as:

- *Hybrid systems*
- *Augmented SLDS (Barber, 2006)*

Linderman et al. AISTATS (2017)

Recurrent Switching Linear Dynamical Systems (rSLDS)

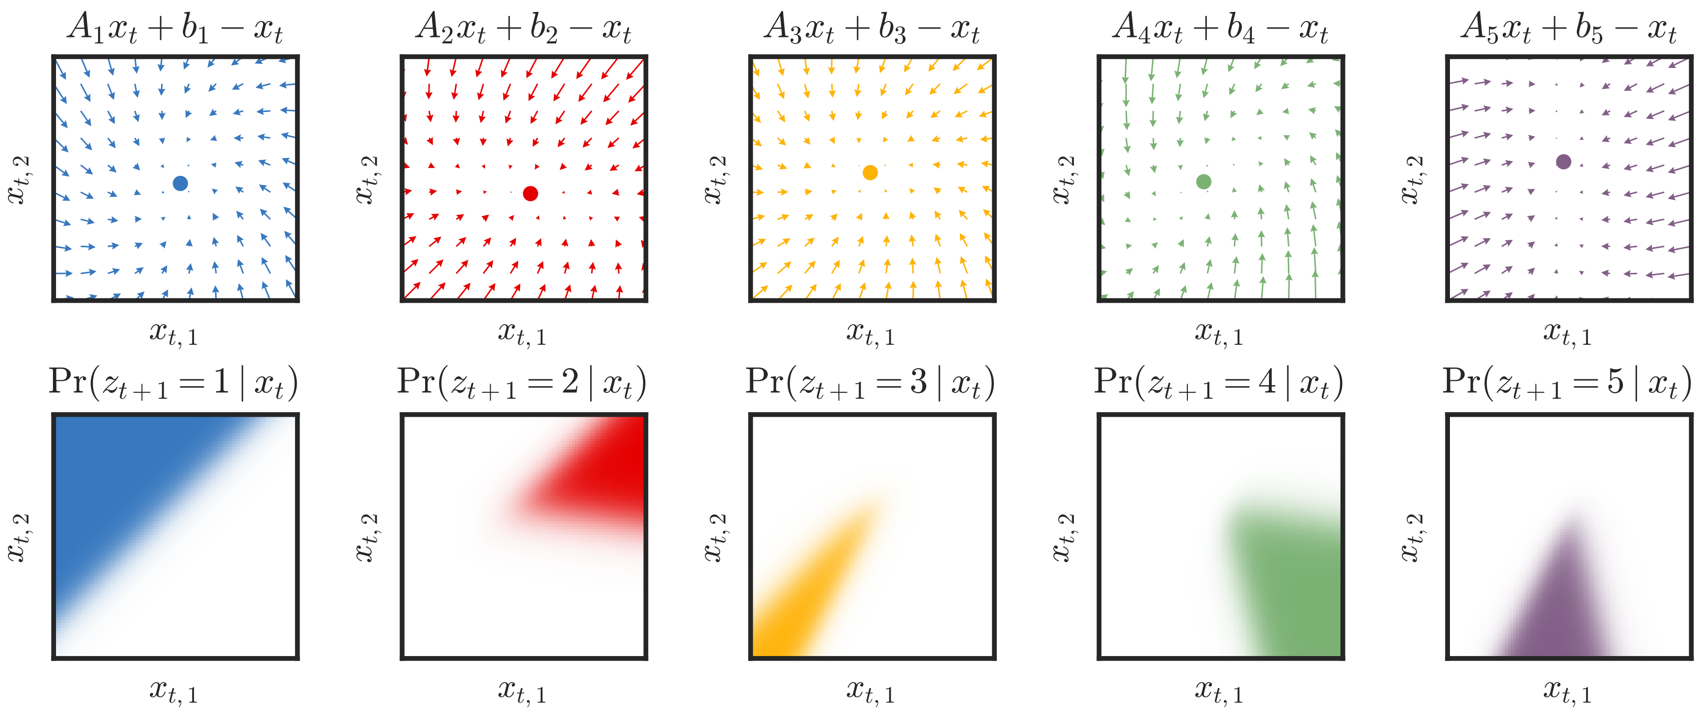

Linderman et al. AISTATS (2017)

Recurrent SLDS

Discrete Latent State Dynamics

$$z_{t+1} \sim \boldsymbol{\pi}_{z_t}, \qquad \boldsymbol{\pi}_{z_t} = f\left(\boldsymbol{R}\,\boldsymbol{x}_t\right)$$

Continuous Latent State Dynamics

$$\boldsymbol{x}_{t+1} \sim \mathcal{N}\left(\boldsymbol{A}_{z_{t+1}}\boldsymbol{x}_t + \boldsymbol{b}_{z_{t+1}},\ \boldsymbol{Q}_{z_{t+1}}\right)$$

Linear Activation Model

$$\boldsymbol{y}_t = \boldsymbol{C}\,\boldsymbol{x}_t + \boldsymbol{d}$$

Non-Gaussian Conditionals

The conditional distribution of $\boldsymbol{x}_{1:T}$ is no longer Gaussian.

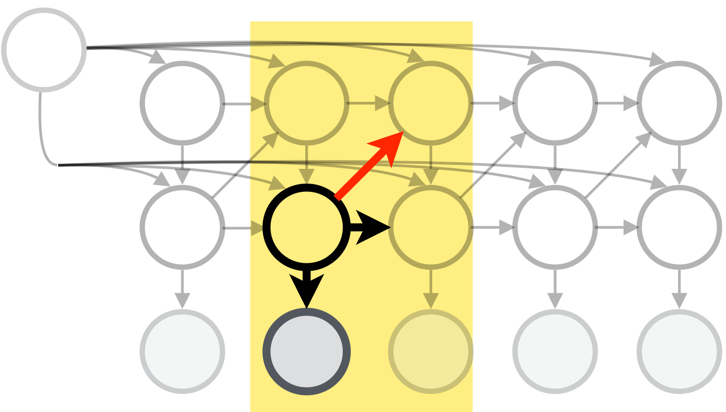

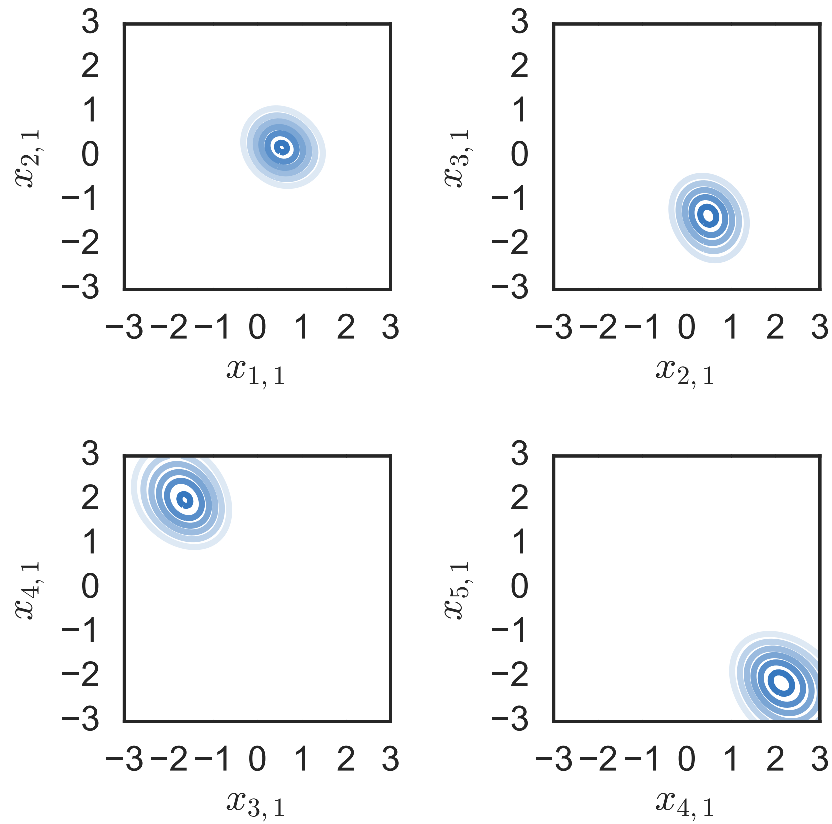

Pólya-gamma Augmentation

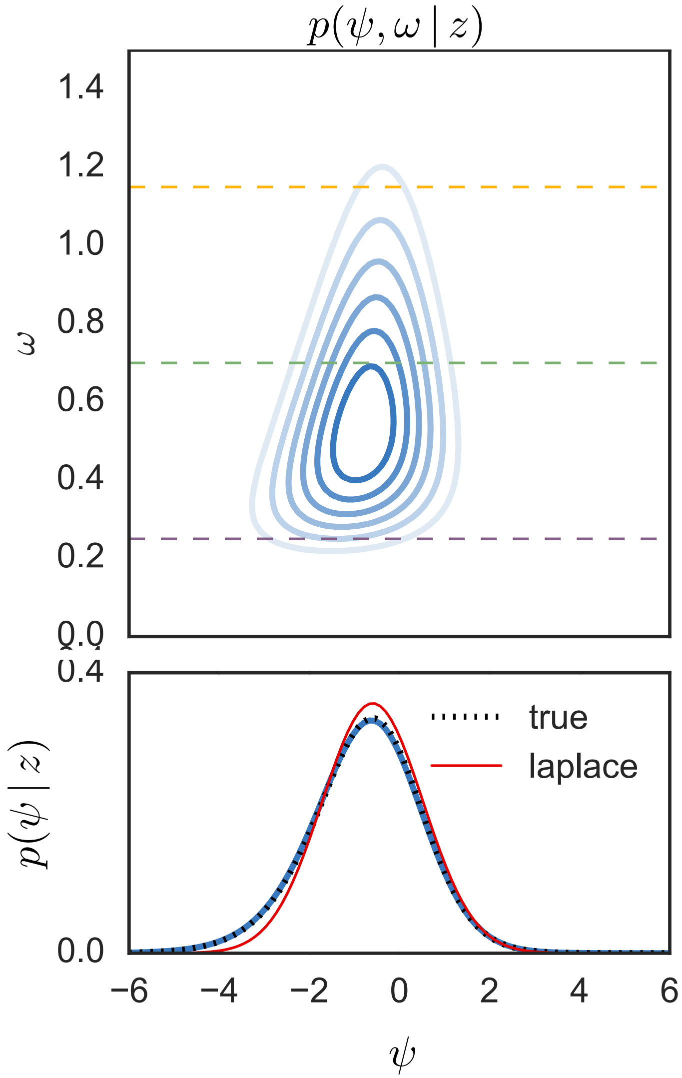

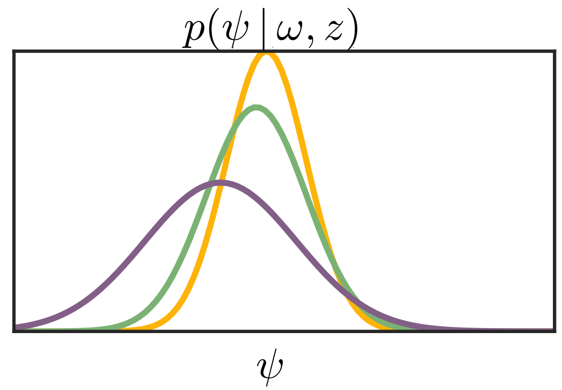

$$\psi \mid \omega, z \sim \mathcal{N}(\mu(\omega, z), \omega^{-1})$$

$$\omega \mid \psi, z \sim \mathrm{PG}(b(z), \psi)$$

Polson, Scott, and Windle. JASA (2013).

Works with Count Observations Too!

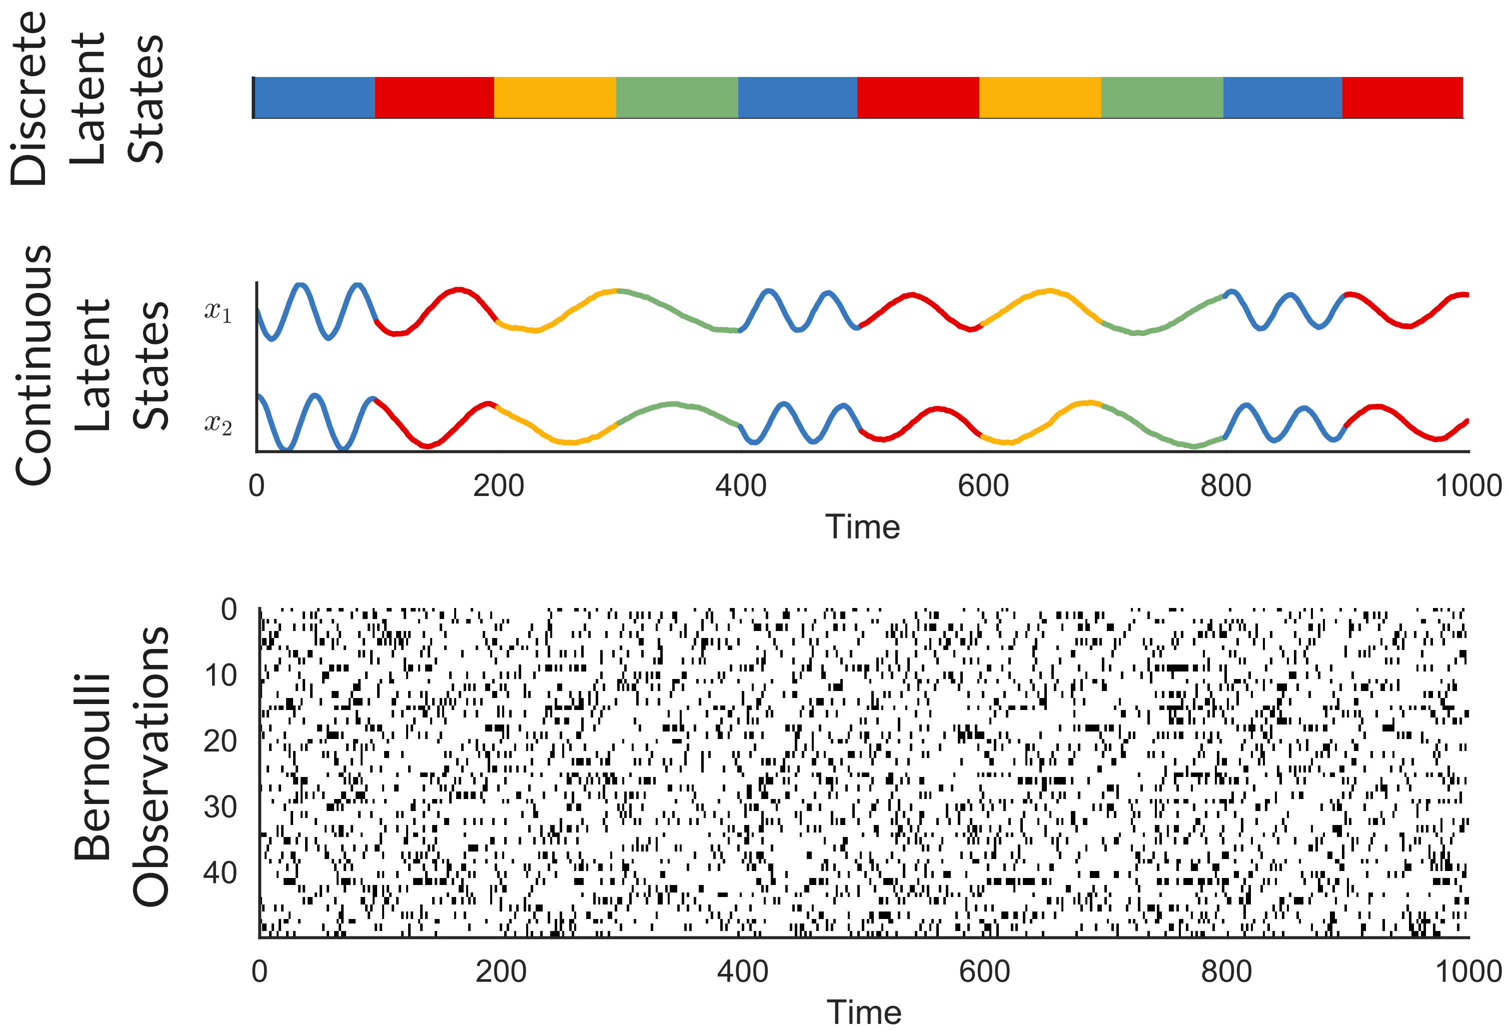

Synthetic NASCAR®

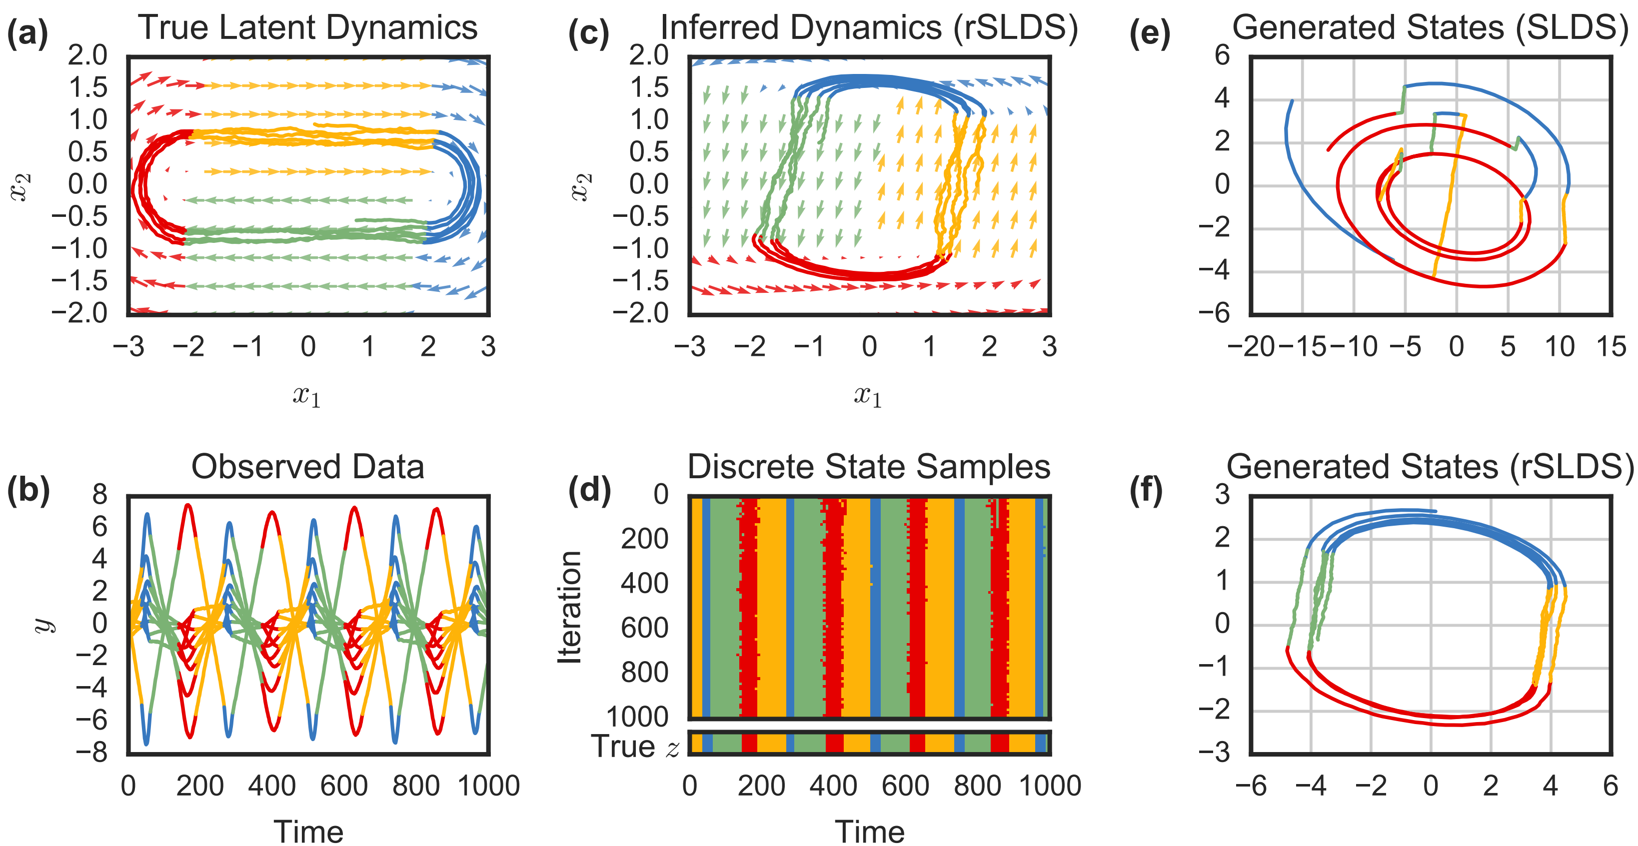

Bernoulli Lorenz

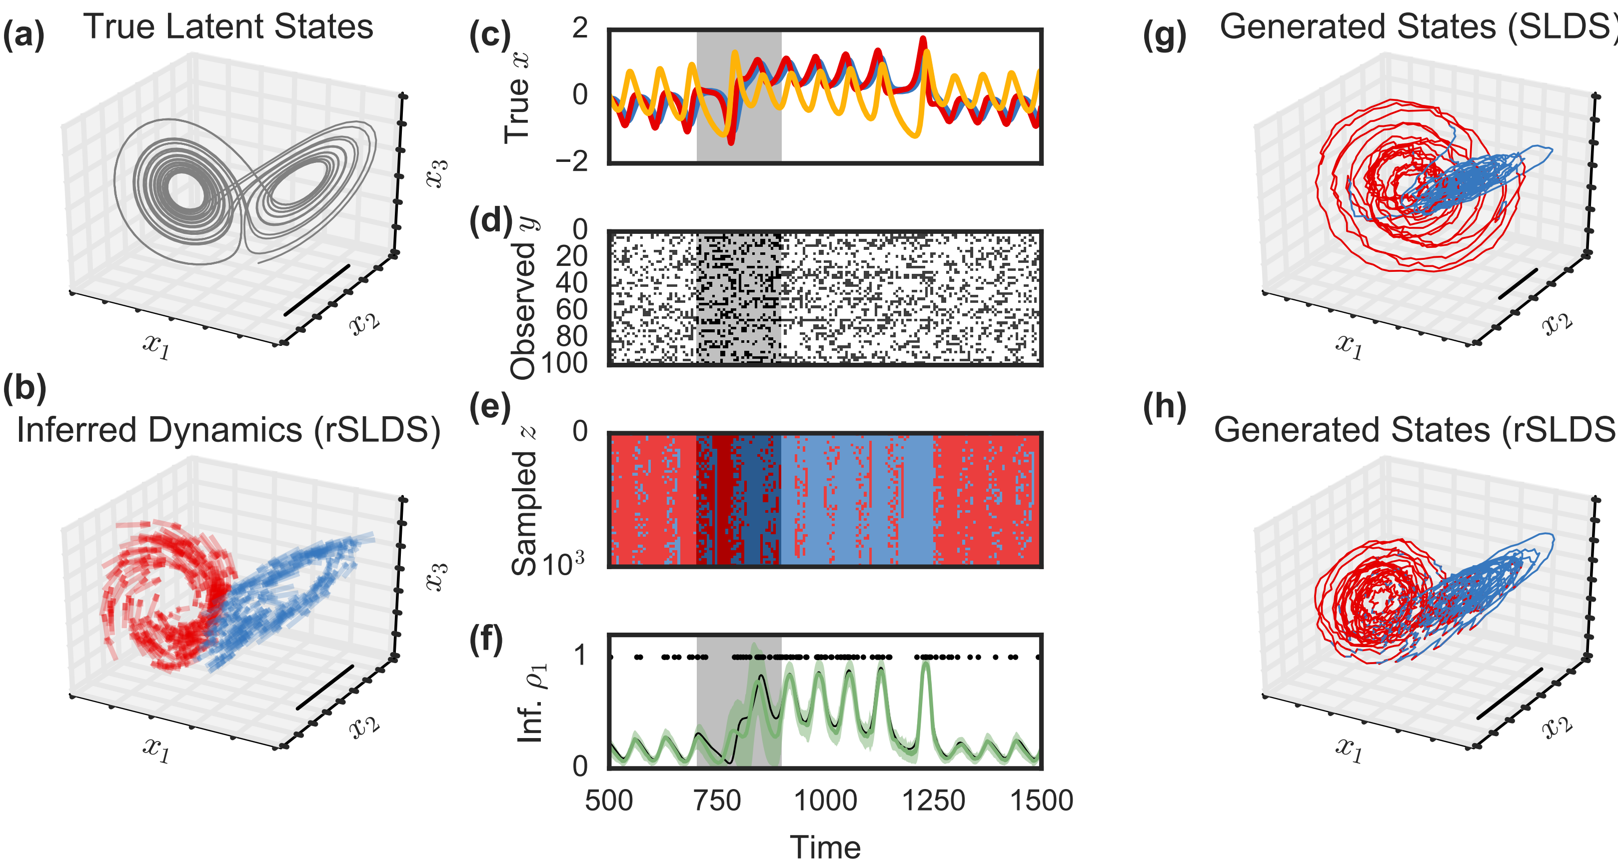

National Basketball Association

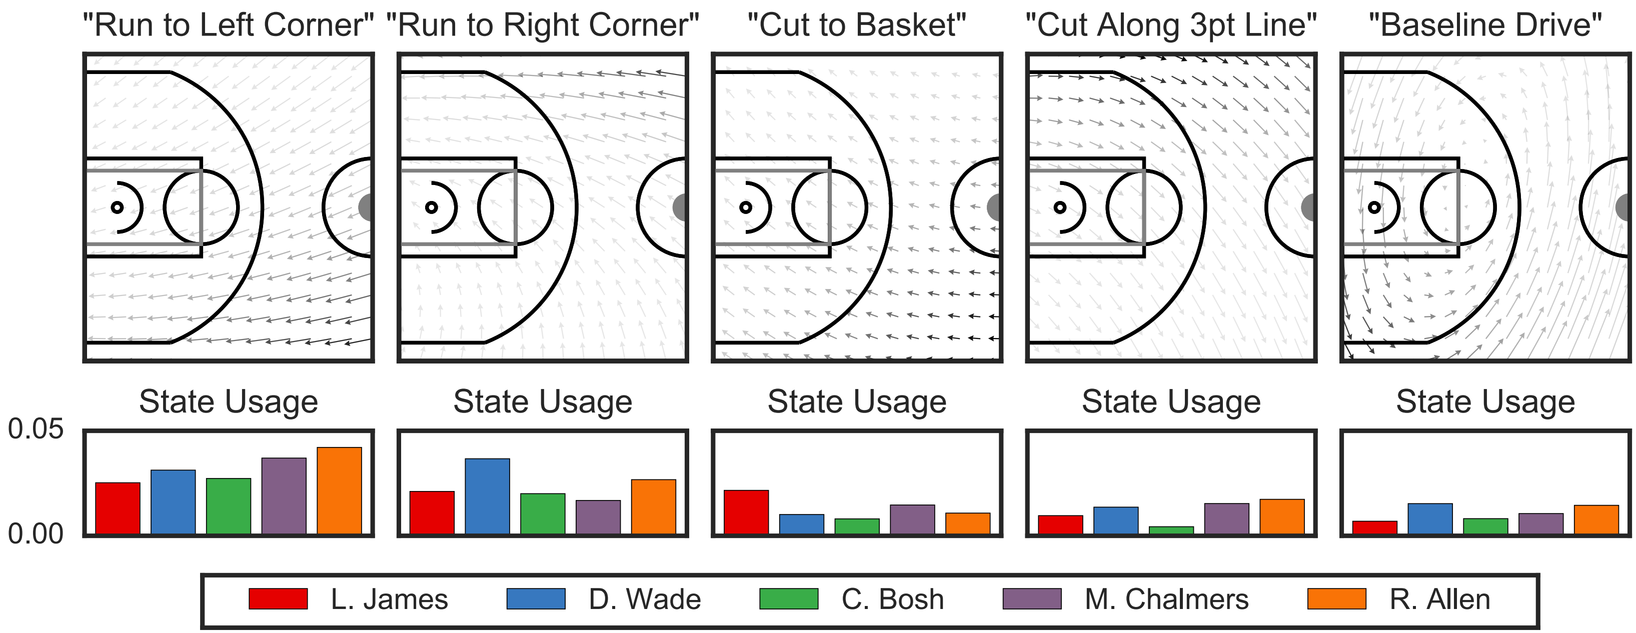

Outline

- Background: Bayesian state-space models
- Recurrent Switching Linear Dynamical Systems
- Whole-Brain Recordings of *C. Elegans*

“Whole-Brain” Recordings in C. *Elegans*

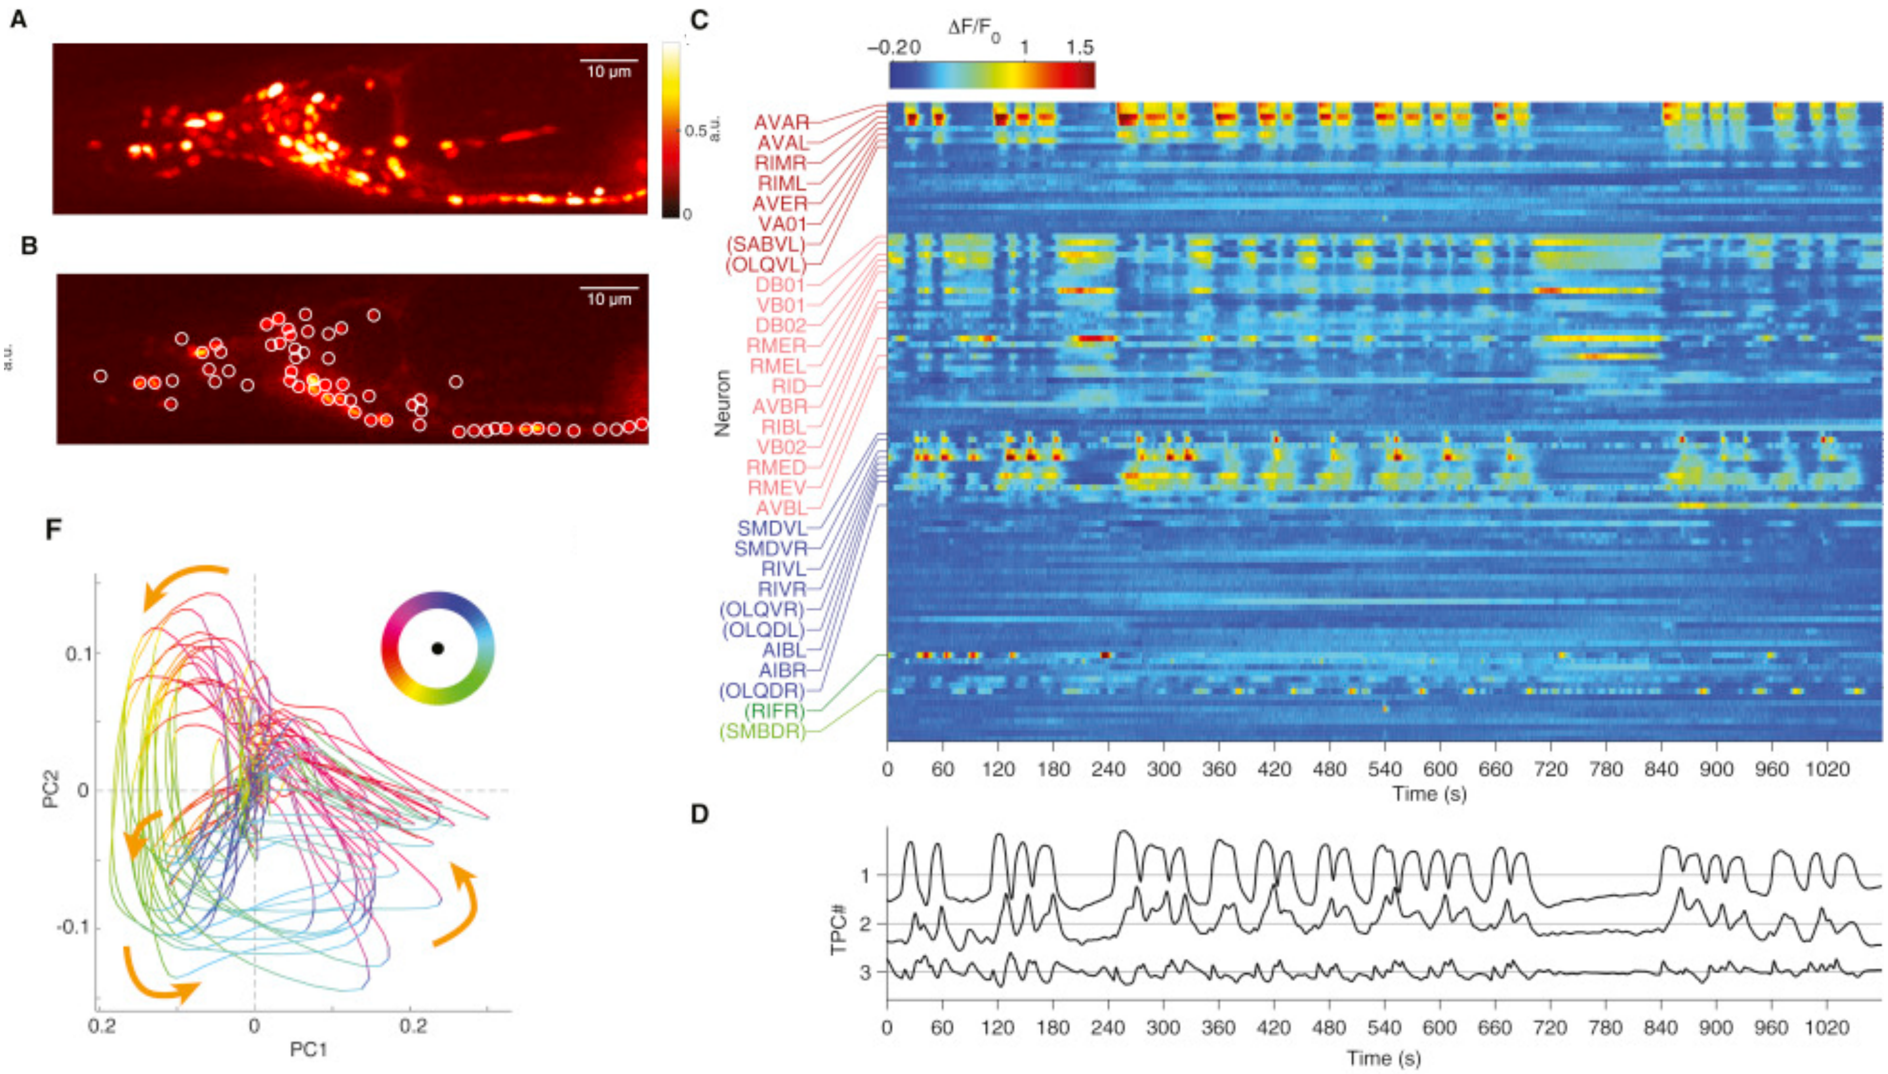

Kato et al., 2015

Principal Component Analysis

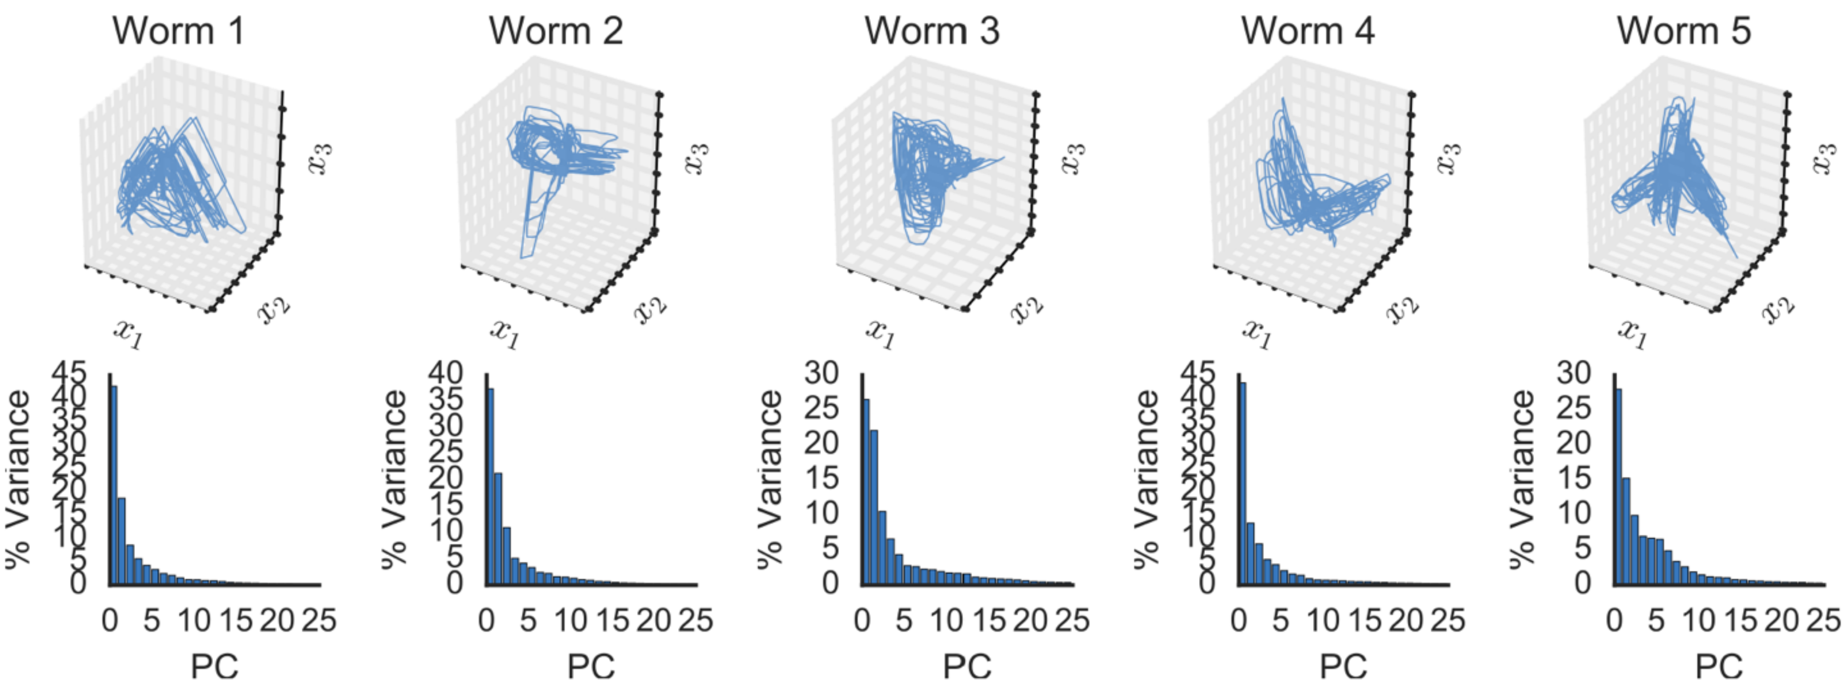

Sharing Partial Observations between Organisms

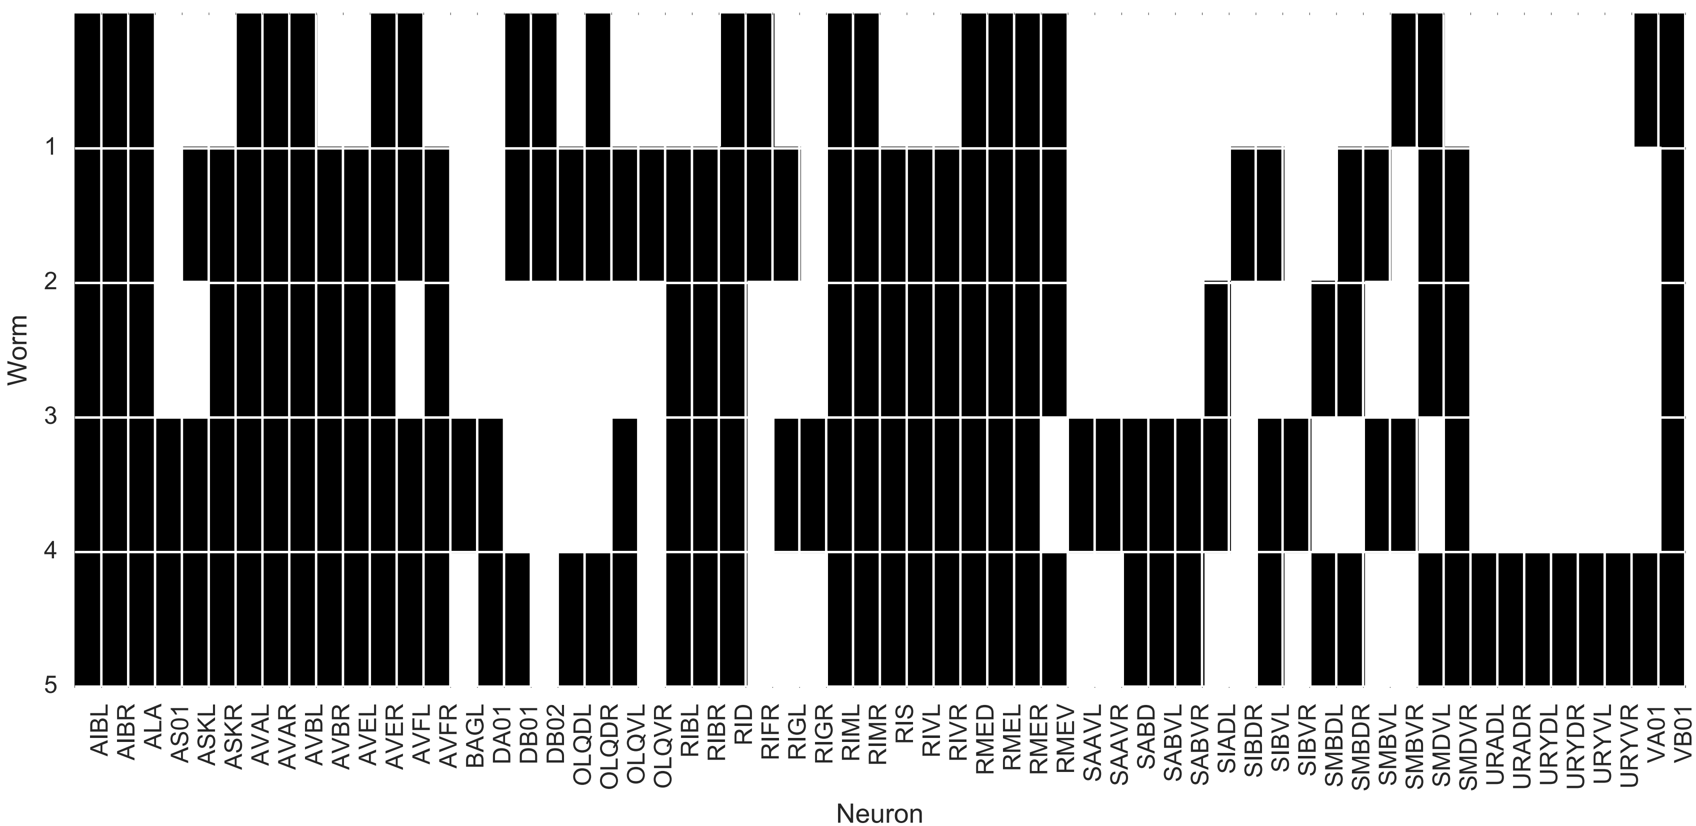

Sharing Partial Observations between Organisms

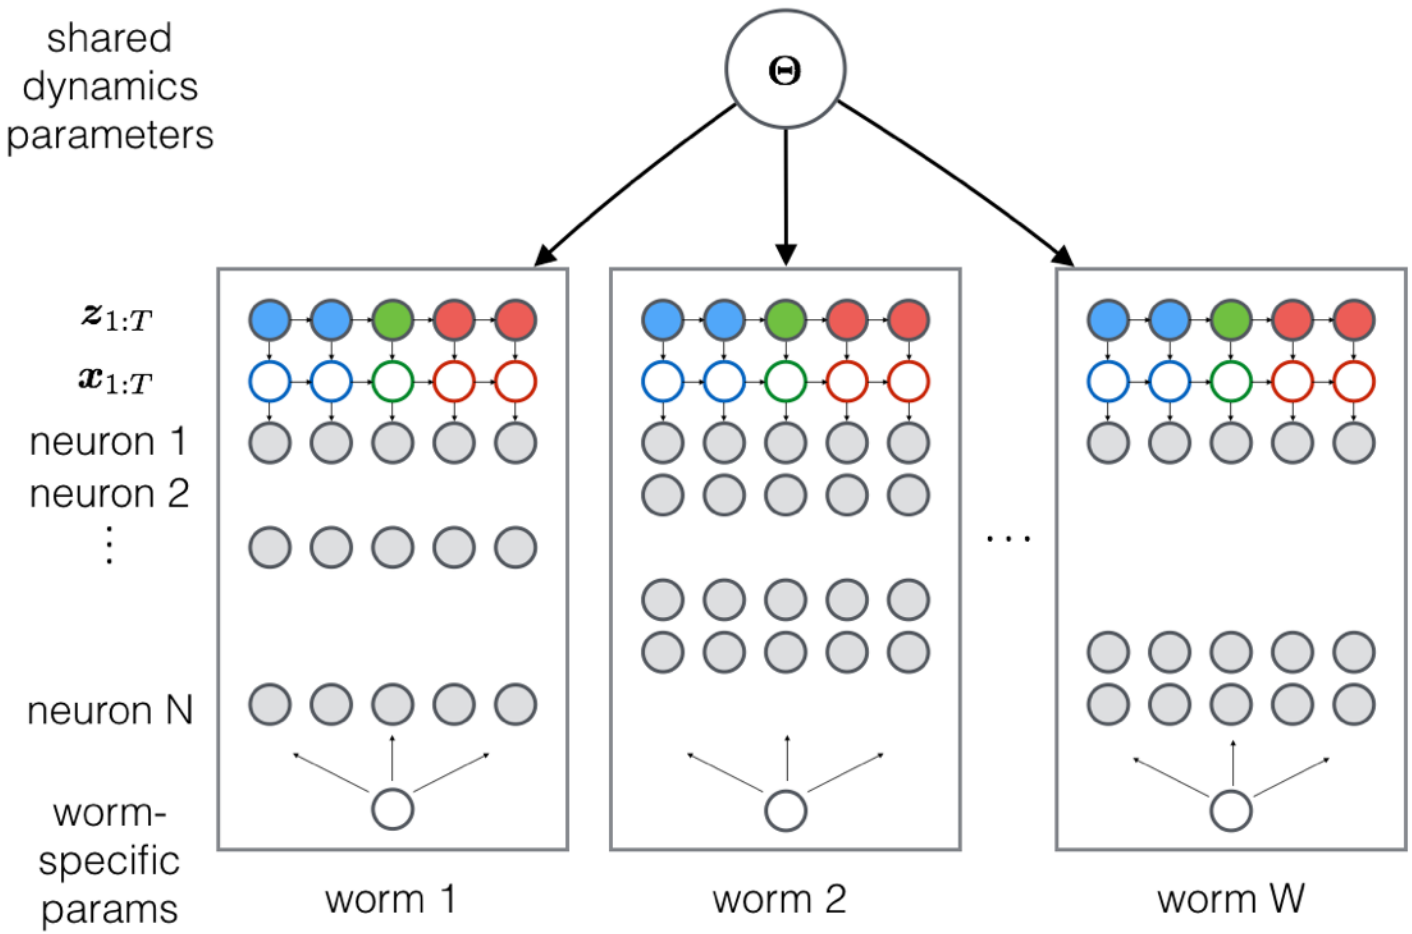

Composite Dynamics

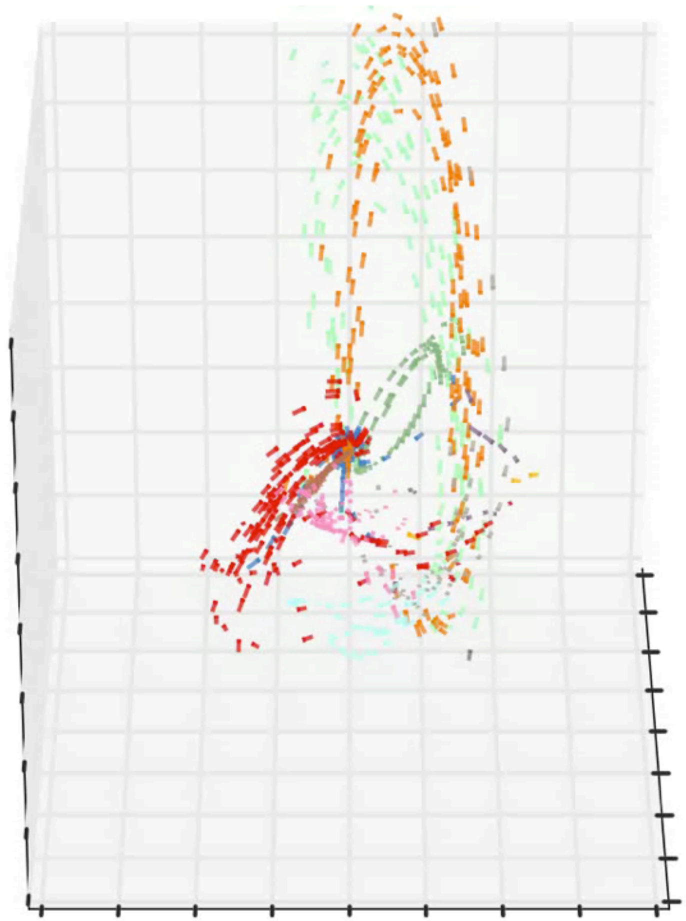

Inferred Change-points Match Manual Segmentation

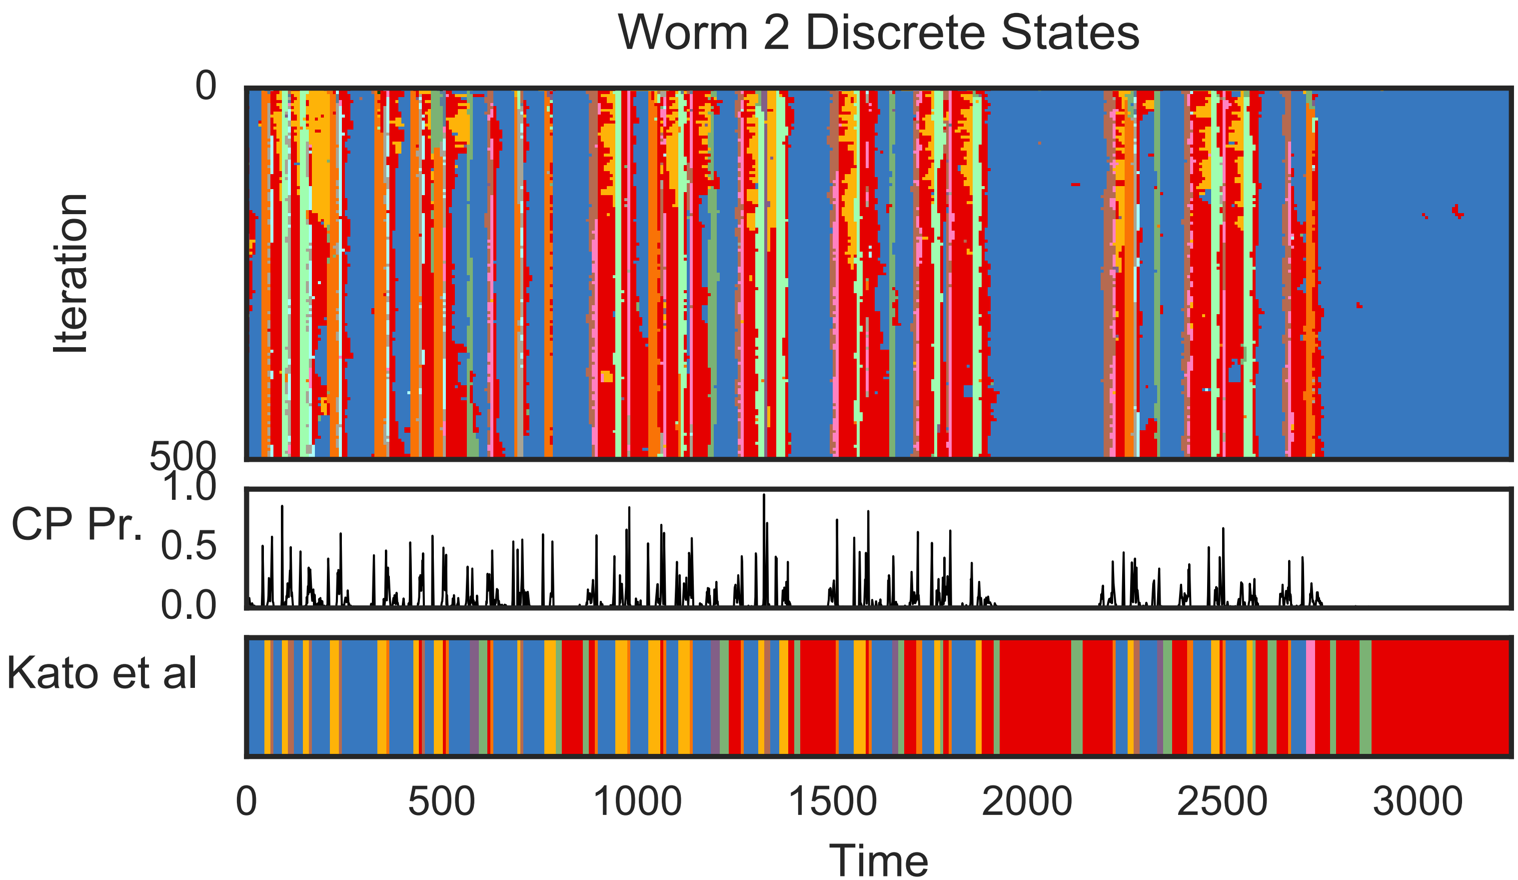

Inferred Dynamical States

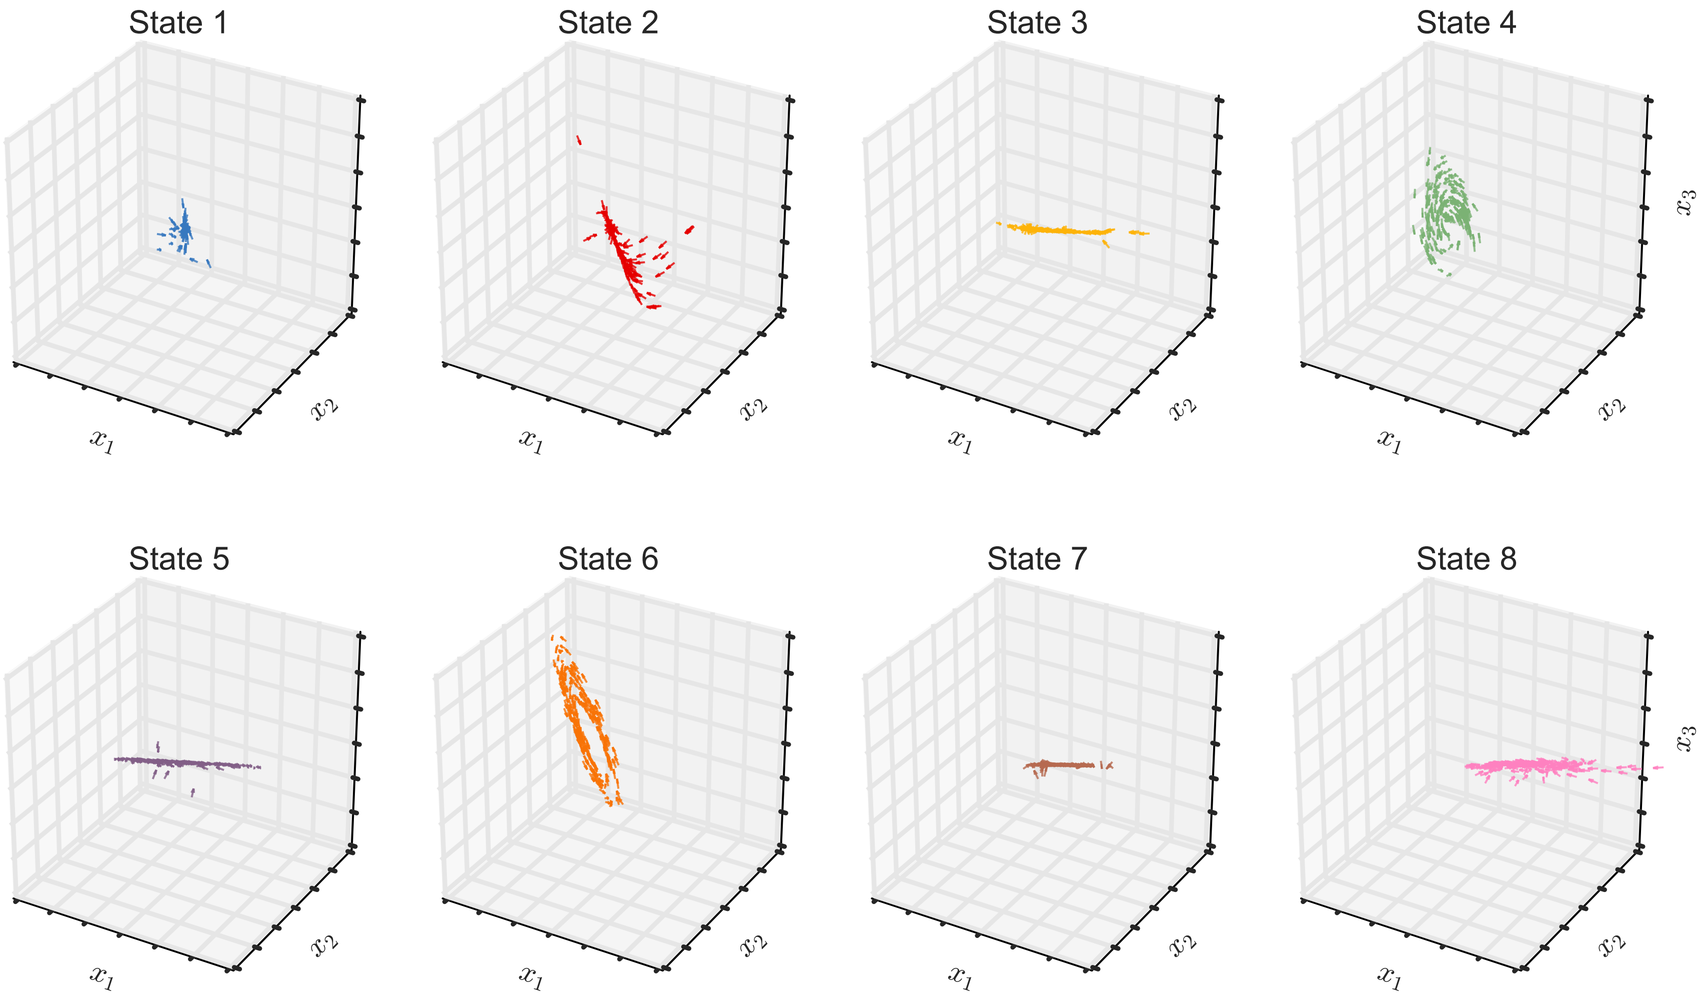

Reconstructions of Neural Activity

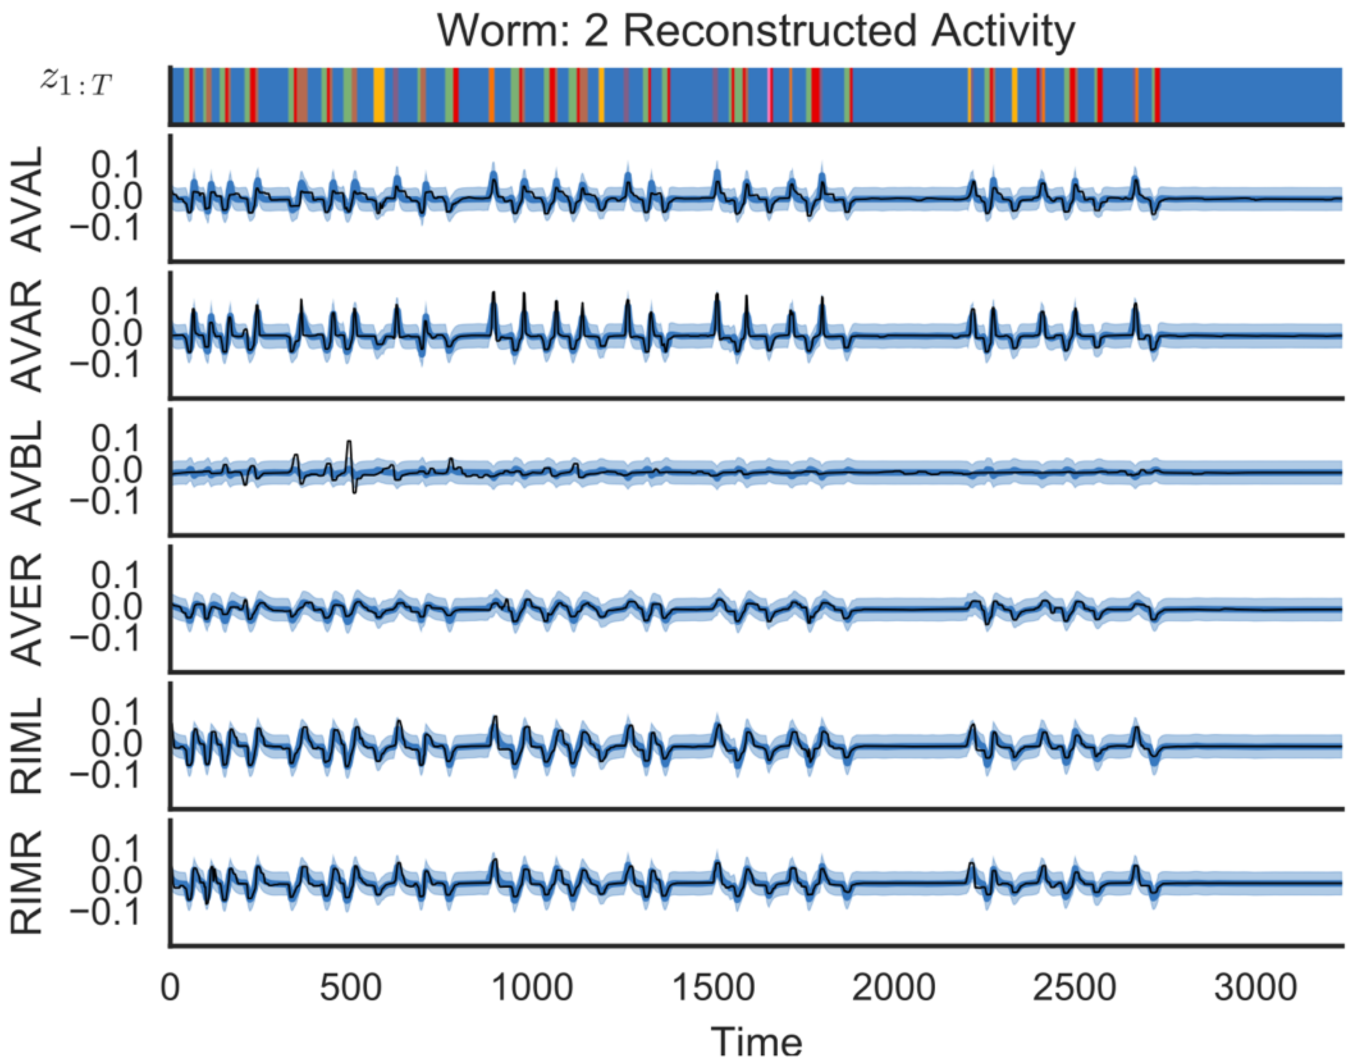

Generating from the Model

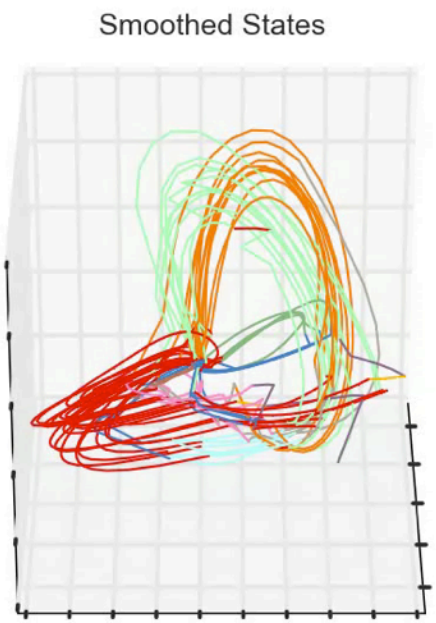

Generating from the Model

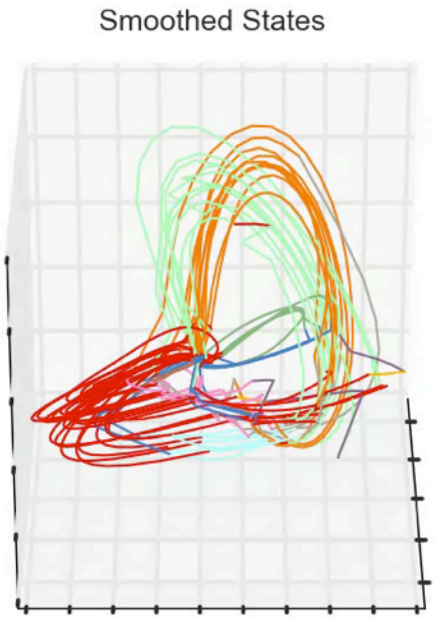

Generating from the Model

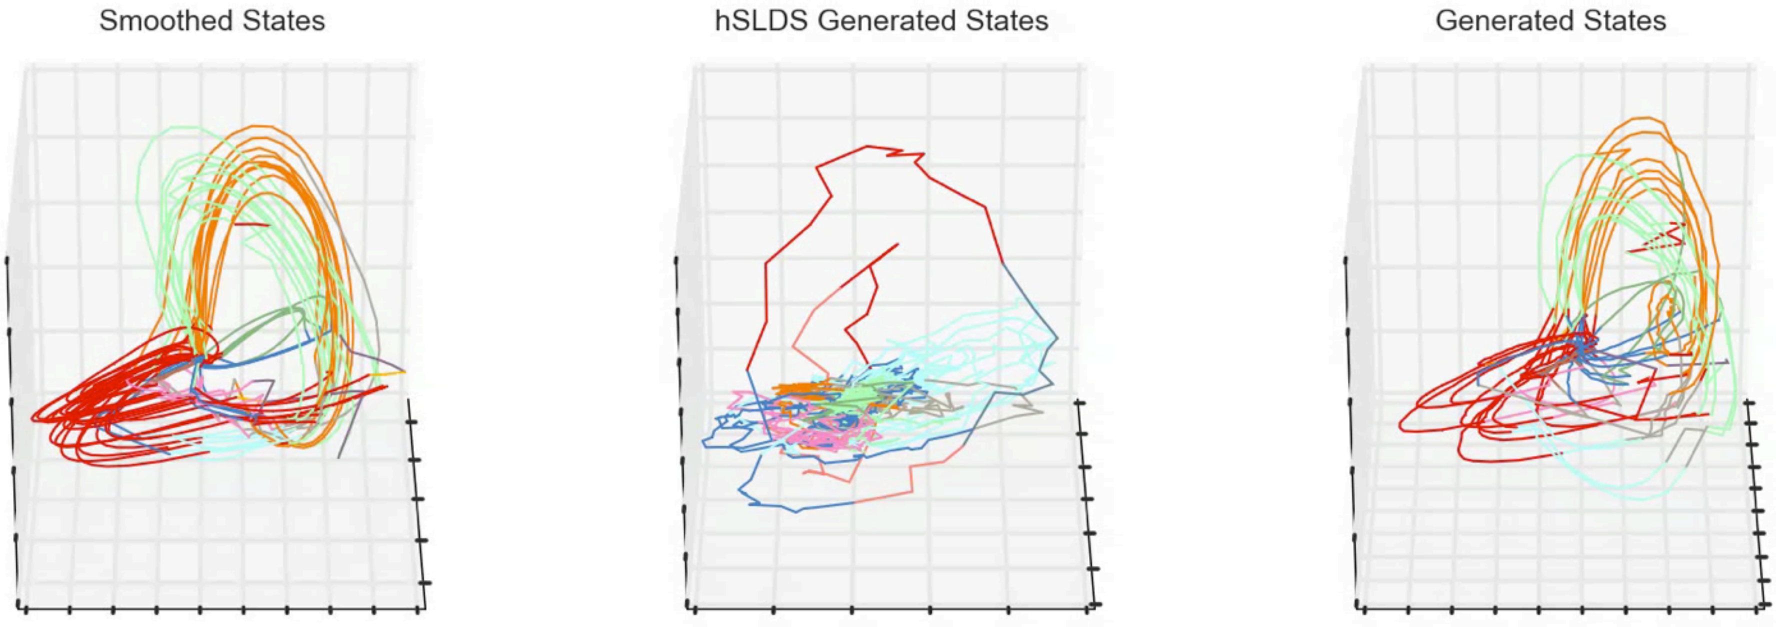

What's Next?

- *Joint observations of neural activity and behavior.*
- *Hierarchical structure in discrete state dynamics.*

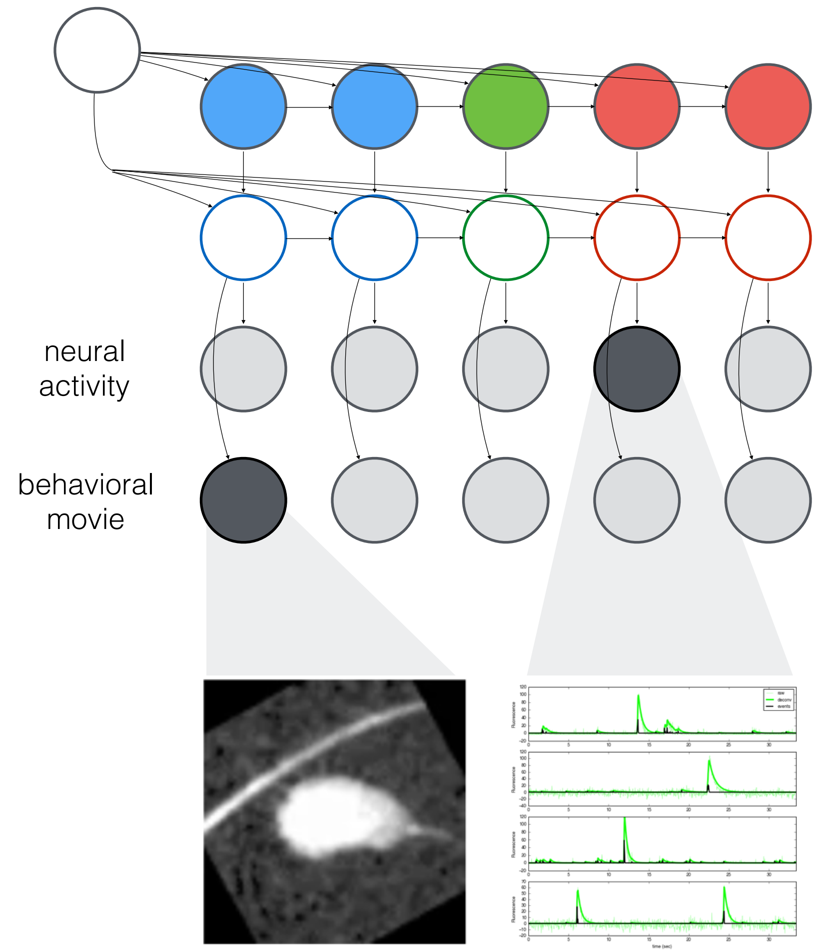

Conclusions

- Switching linear dynamical systems offer a parsimonious balance of flexibility, tractability, and interpretability.
- Recurrent SLDS add an important missing piece — “closed loop” dynamics.
- Augmentation schemes enable us to capitalize on conditionally linear Gaussian structure and efficient message passing algorithms.
- Bayesian formulations naturally handle missing data and hierarchical extensions.

Acknowledgments

Thank you!







Simons Collaboration on the Global Brain
Siebel Foundation