ICML International Conference On Machine Learning

Duke UNIVERSITY

Dynamic Compression Flows for Neuroscience Data

Ganchao Wei*
Daniela de Albuquerque*
Miles Martinez
Shiyang Pan
John Pearson

P[λ]ab Brains & Computation

Learning Low-D Dynamics from Neuroscience Data

- Neuroscience data often exhibit low-dimensional structure. We are interested in learning low-D dynamics from data.

- Existing approaches fall into 3 main categories
  - Compression focused
  - Dynamical Systems focused
  - Hybrid, often with strong modeling assumptions

- Use a set of *coupled-flows* to learn *identifiable* and flexible *low-D* representations that *preserve dynamics*

Building Our Coupled Flows

Data $\mathbf{x}_t \in \mathbb{R}^D$

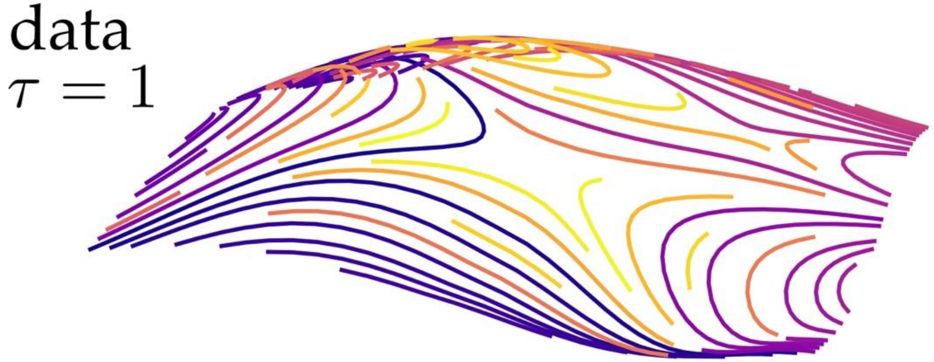

Building Our Coupled Flows

- Data $\mathbf{x}_t \in \mathbb{R}^D$
- Wish to learn low-D latent dynamics

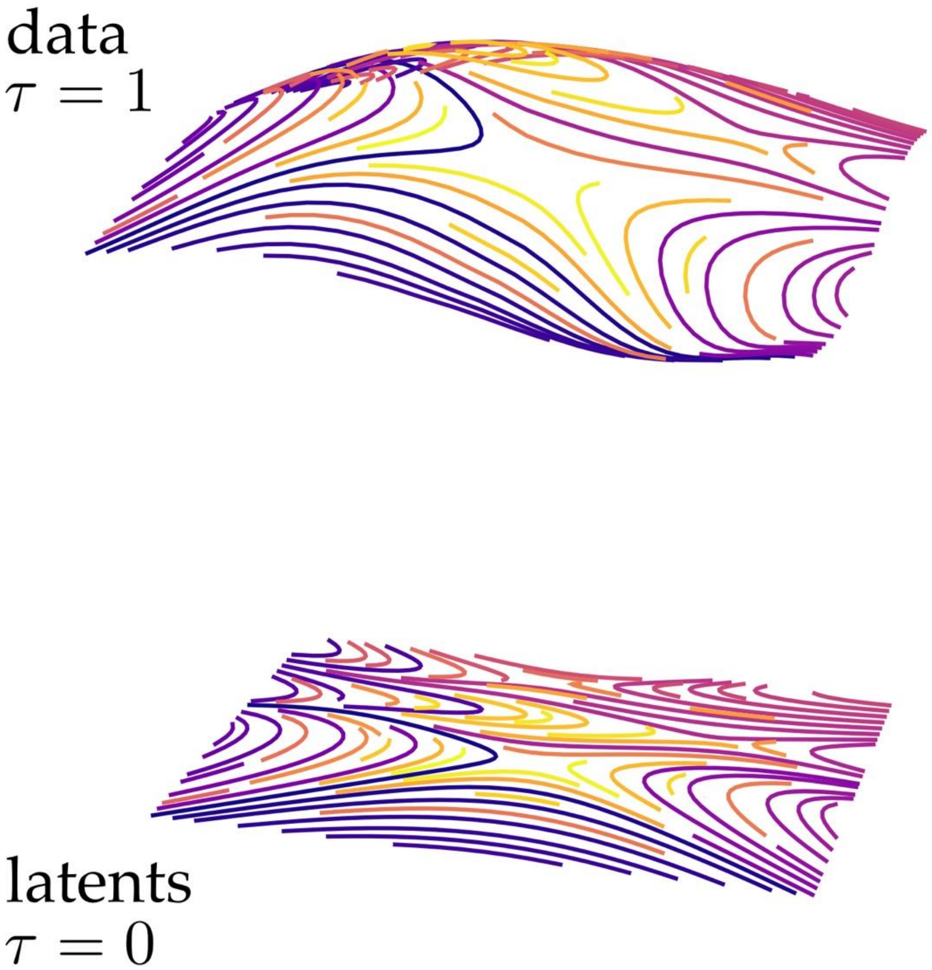

Building Our Coupled Flows

- Data $\mathbf{x}_t \in \mathbb{R}^D$
- Wish to learn low-D latent dynamics
- Compressive/Generative flow transports data between ambient and compressed spaces

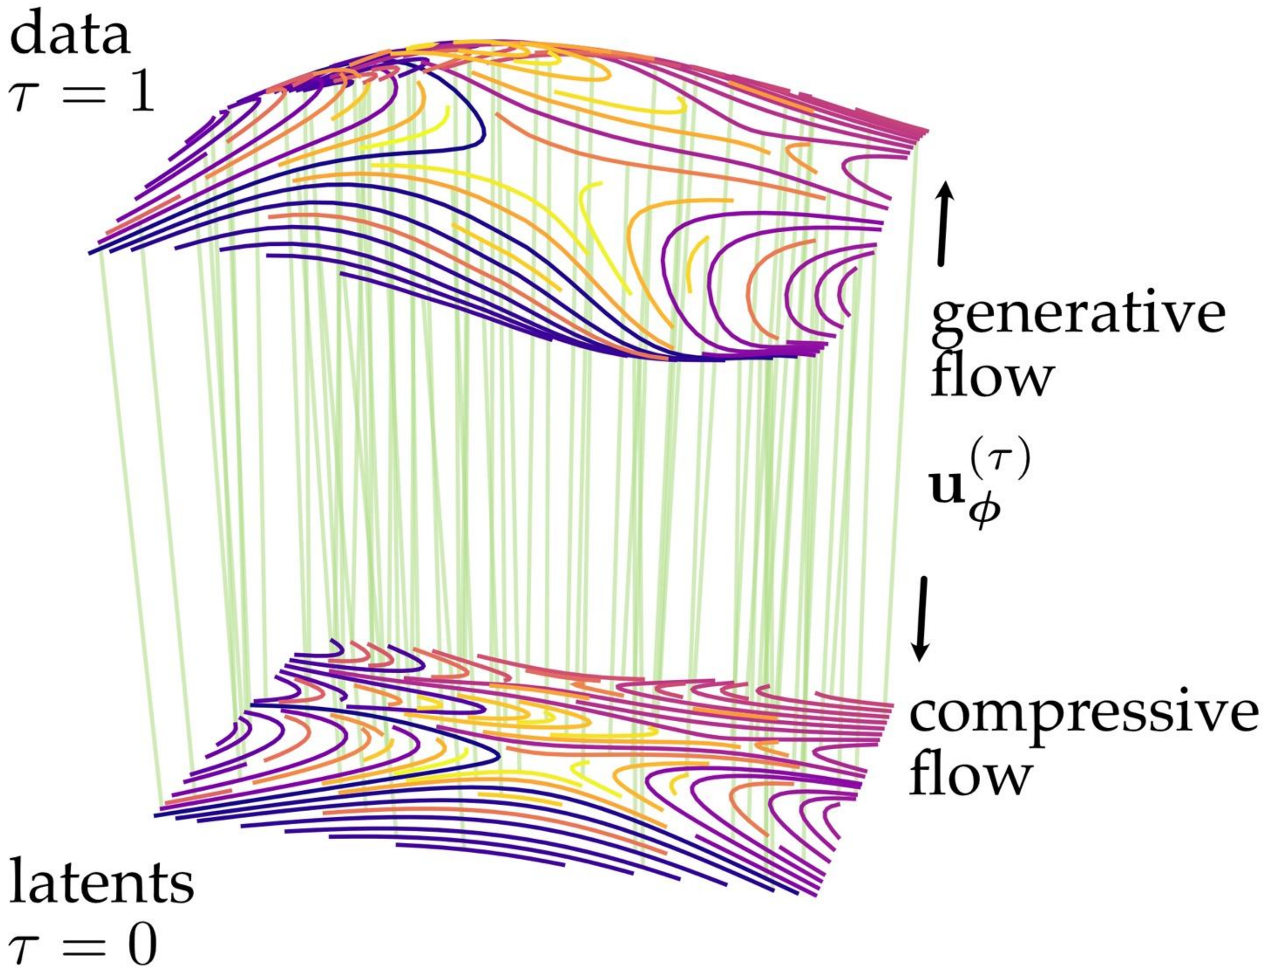

Building Our Coupled Flows

- Data $\mathbf{x}_t \in \mathbb{R}^D$
- Wish to learn low-D latent dynamics
- Compressive/Generative flow transports data between ambient and compressed spaces
- Dynamical flow captures temporal structure across all compression levels

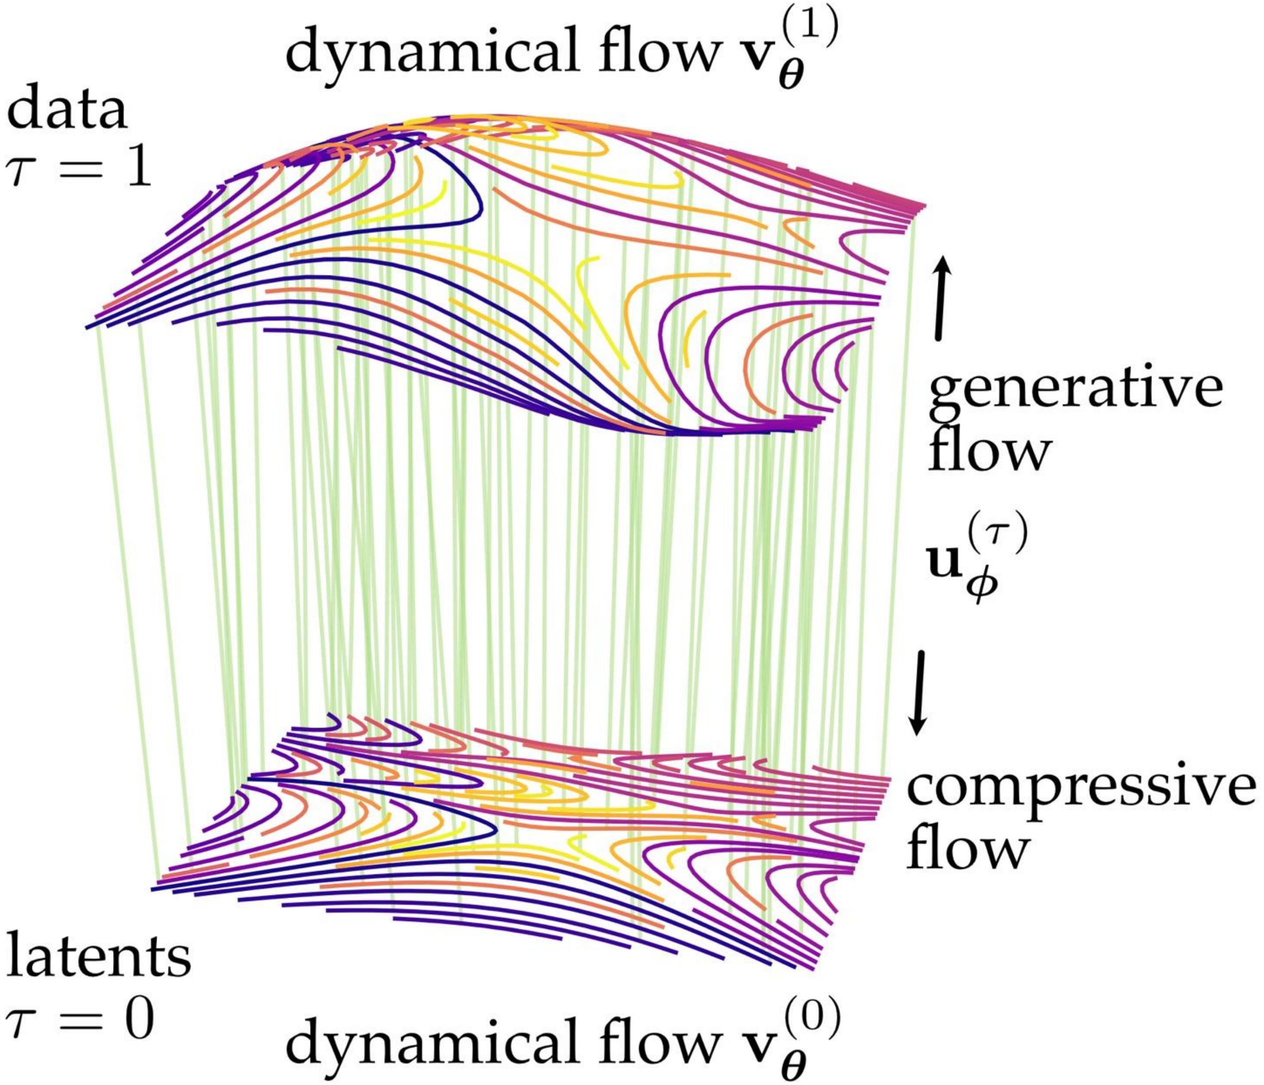

Compression Target Distribution

- Use an encoder to specify compressed distribution

$$\mathbf{x}_t^{(0)} = \mathbf{b} + \mathbf{L}\mathbf{D}^{1/2} \cdot \boldsymbol{\mu}_\psi \left(\mathbf{x}_t^{(1)} \right)$$

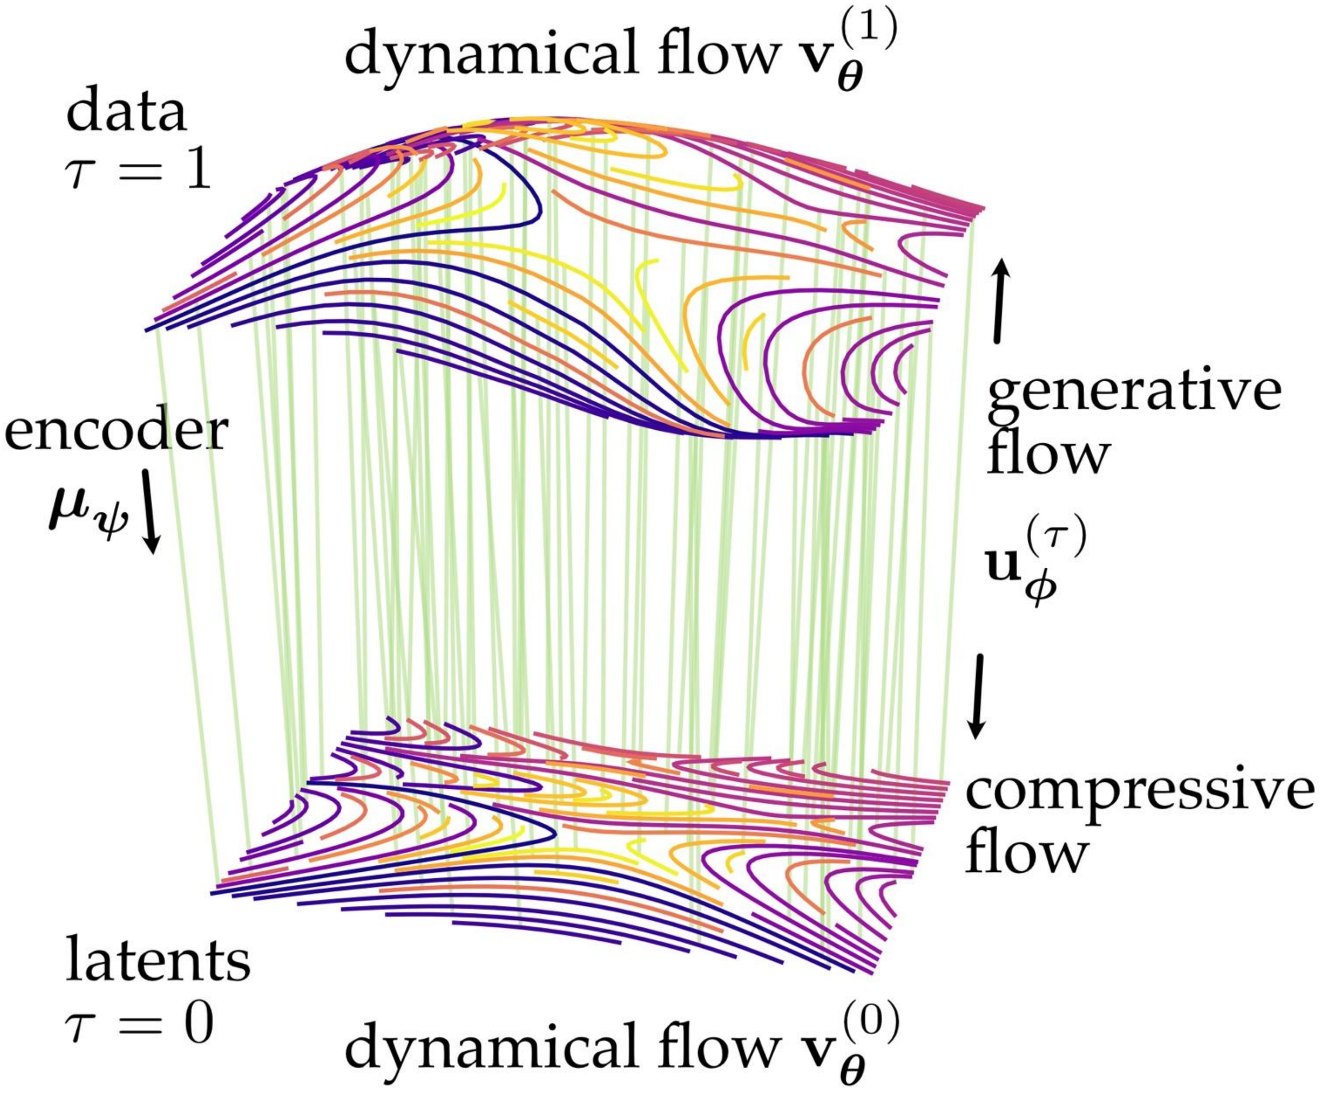

Compression Target Distribution

🍃Use an encoder to specify compressed distribution

$$\mathbf{x}_t^{(0)} = \mathbf{b} + \mathbf{L}\mathbf{D}^{1/2} \cdot \boldsymbol{\mu}_{\psi}\left(\mathbf{x}_t^{(1)}\right)$$

Orthonormal Non-pivoted QR decomposition

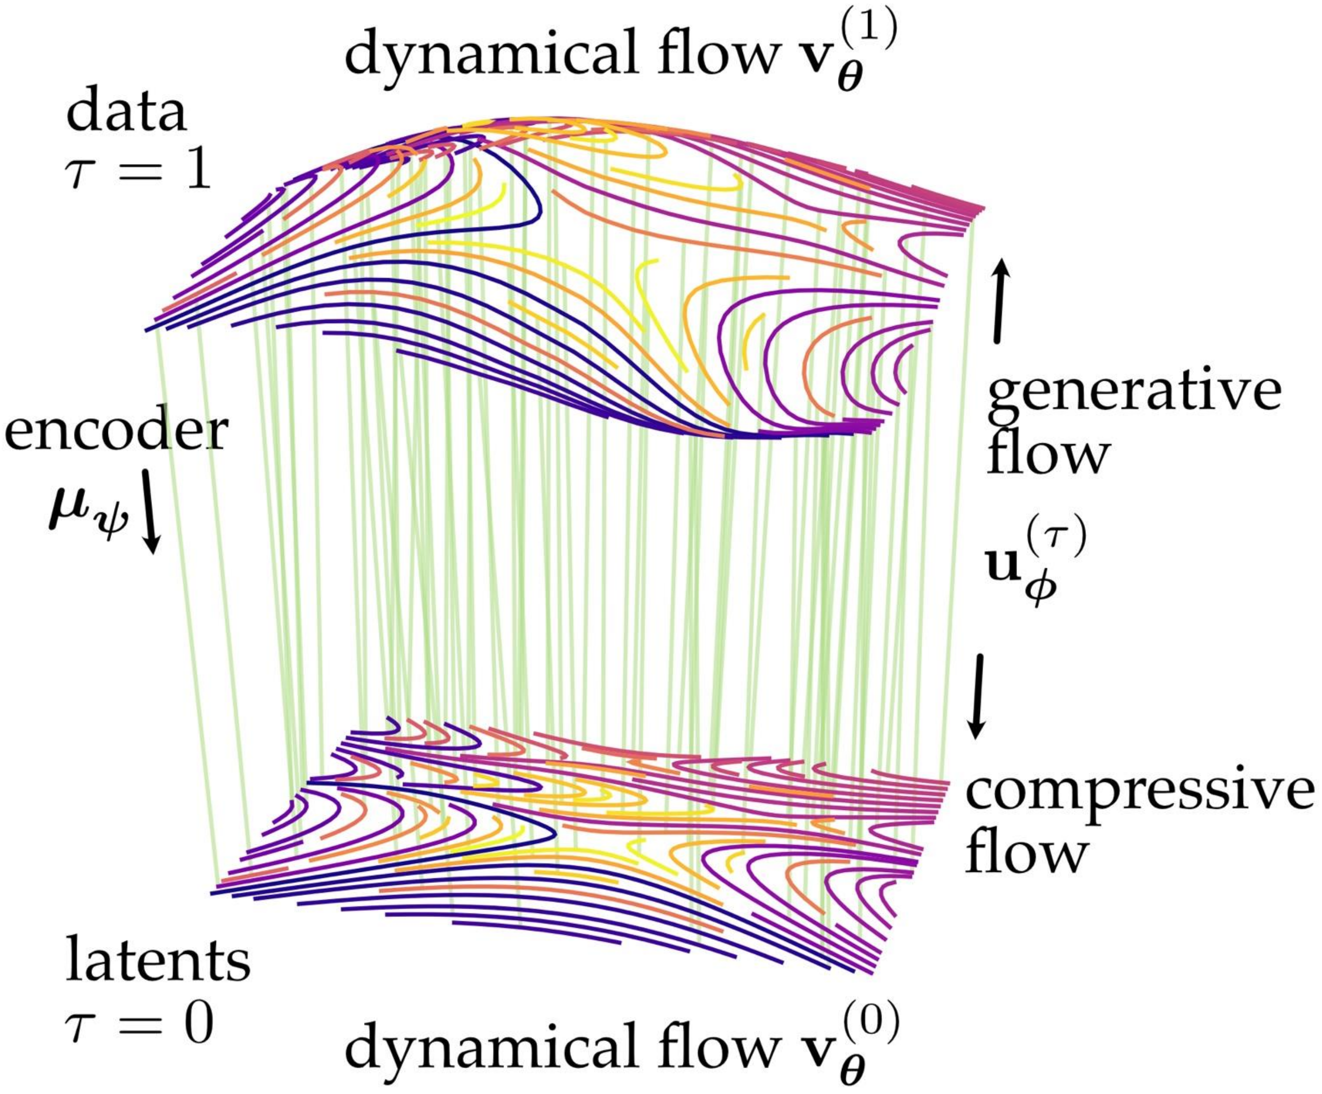

Compression Target Distribution

- Use an encoder to specify compressed distribution

$$\mathbf{x}_t^{(0)} = \mathbf{b} + \mathbf{L}\mathbf{D}^{1/2} \cdot \boldsymbol{\mu}_\psi \left(\mathbf{x}_t^{(1)} \right)$$

Orthonormal Non-pivoted QR decomposition

- Learn an encoding map that minimizes rank of D
  - Nested Dropout

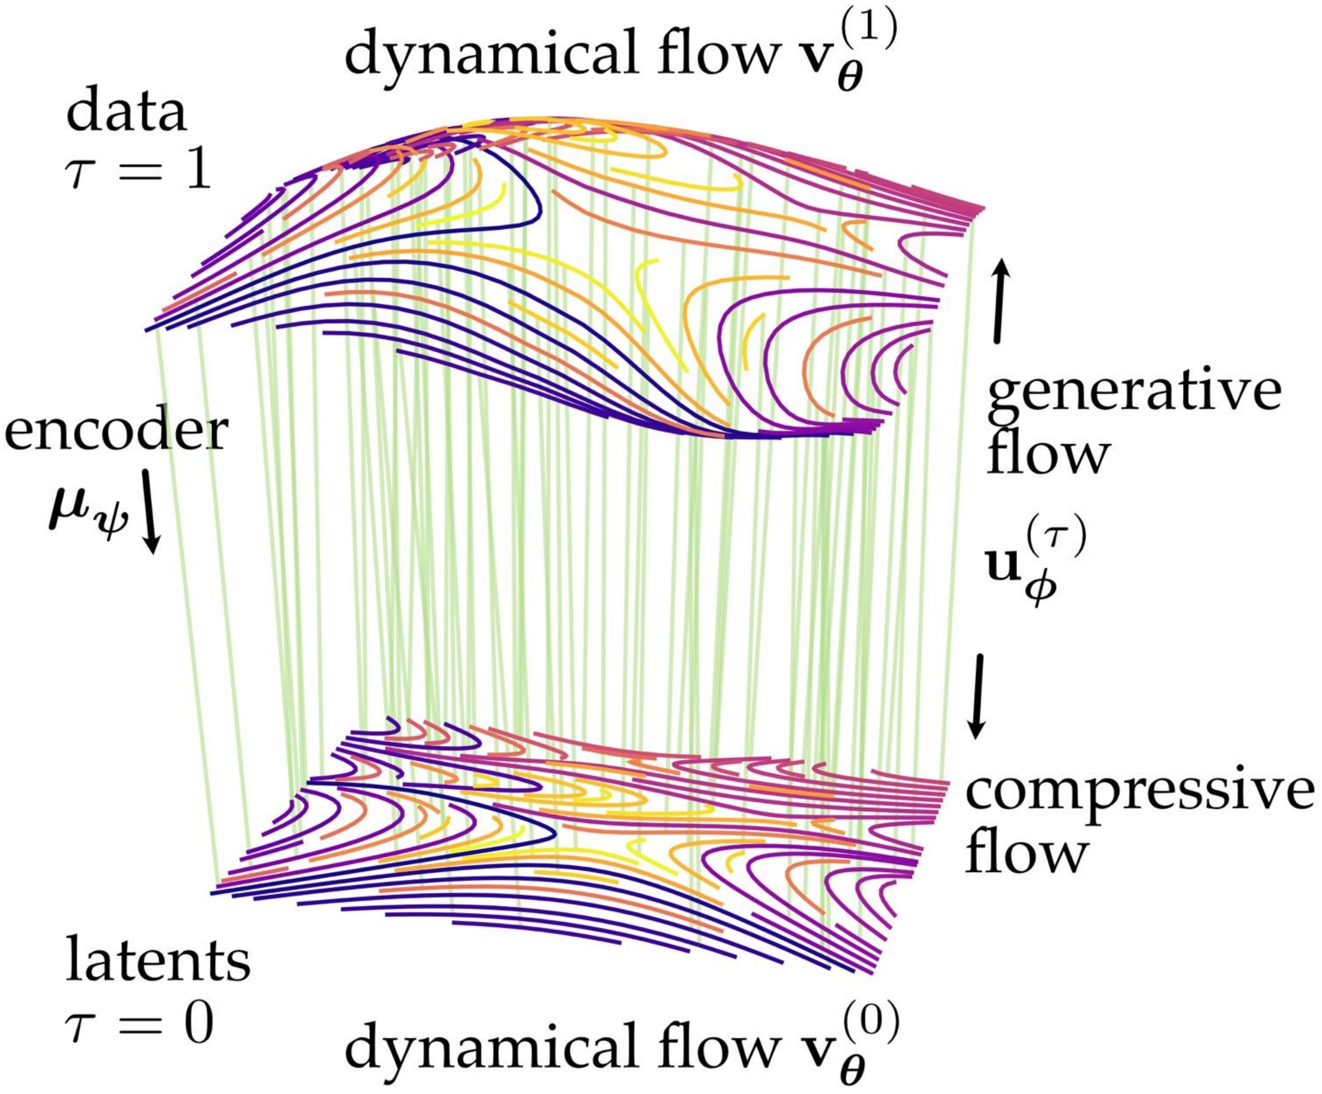

Compression Target Distribution

Nested Dropout

- Sample prefix length

$$K \sim \mathrm{Geom}(p),\ \mathbb{E}(K) = 1/p$$

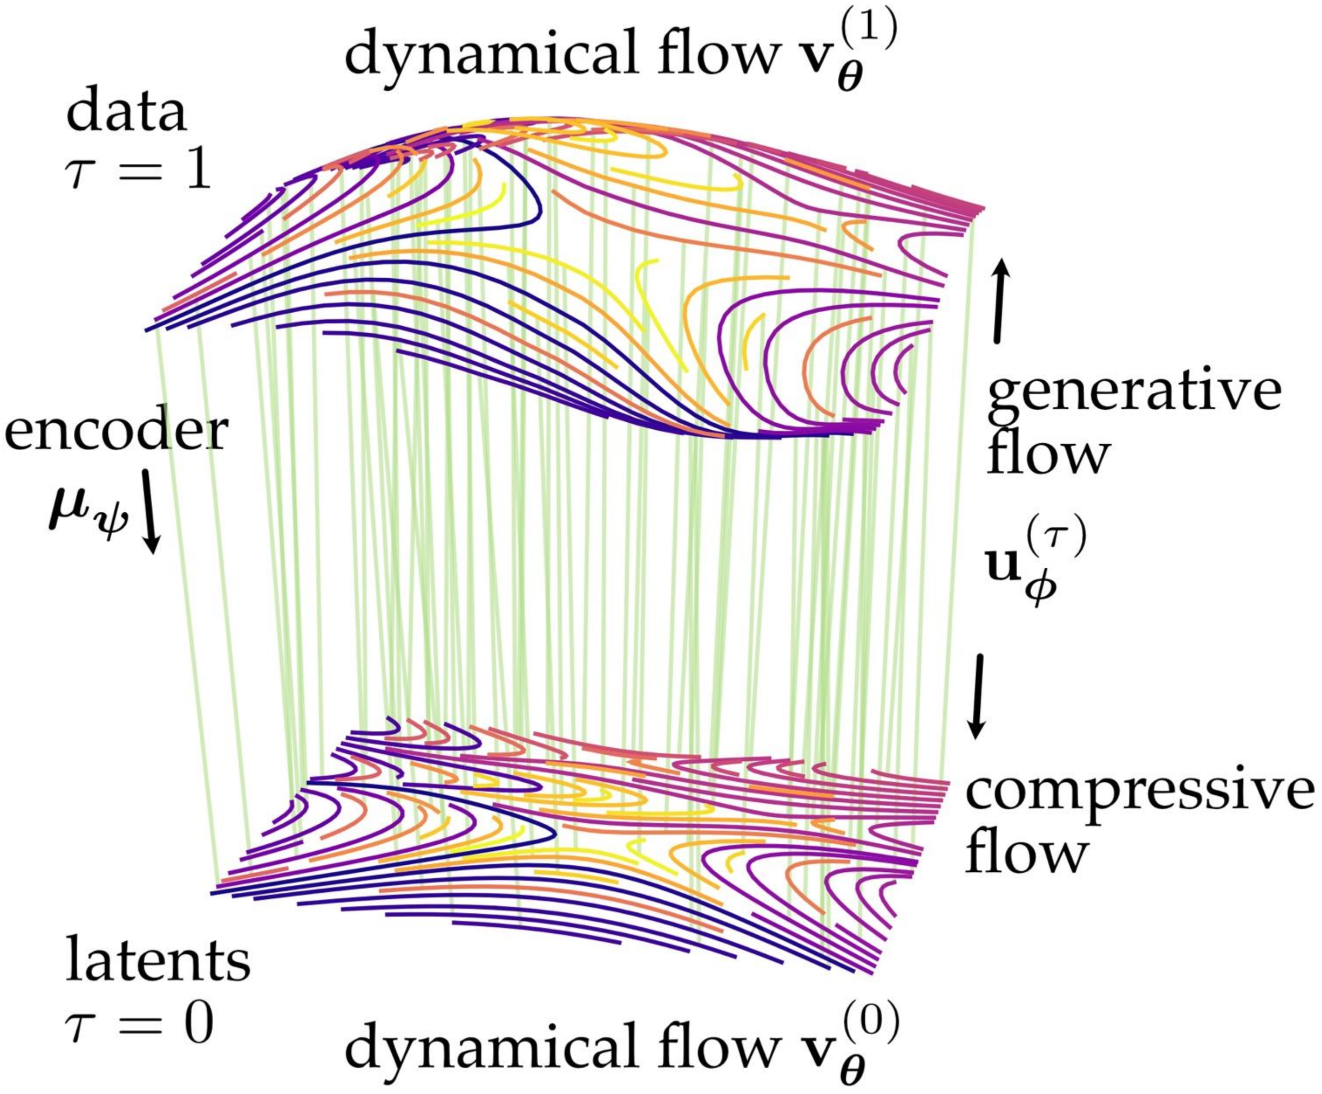

Compression Target Distribution

Nested Dropout

- Sample prefix length

$$K \sim \mathrm{Geom}(p),\ \mathbb{E}(K) = 1/p$$

- Apply prefix mask

$$\mathbf{x}_t^{(0,K)} = \mathbf{b} + \mathbf{L}\mathbf{D}^{1/2}\boldsymbol{\mu}_{\boldsymbol{\psi}}^{(K)}(\mathbf{x}_t^{(1)})$$

$$\mathbf{m}_K(i) = \mathbf{1}\{i \leq K\}, \qquad \boldsymbol{\mu}_{\boldsymbol{\psi}}^{(K)}(\mathbf{x}) = \mathbf{m}_K \odot \boldsymbol{\mu}_{\boldsymbol{\psi}}(\mathbf{x})$$

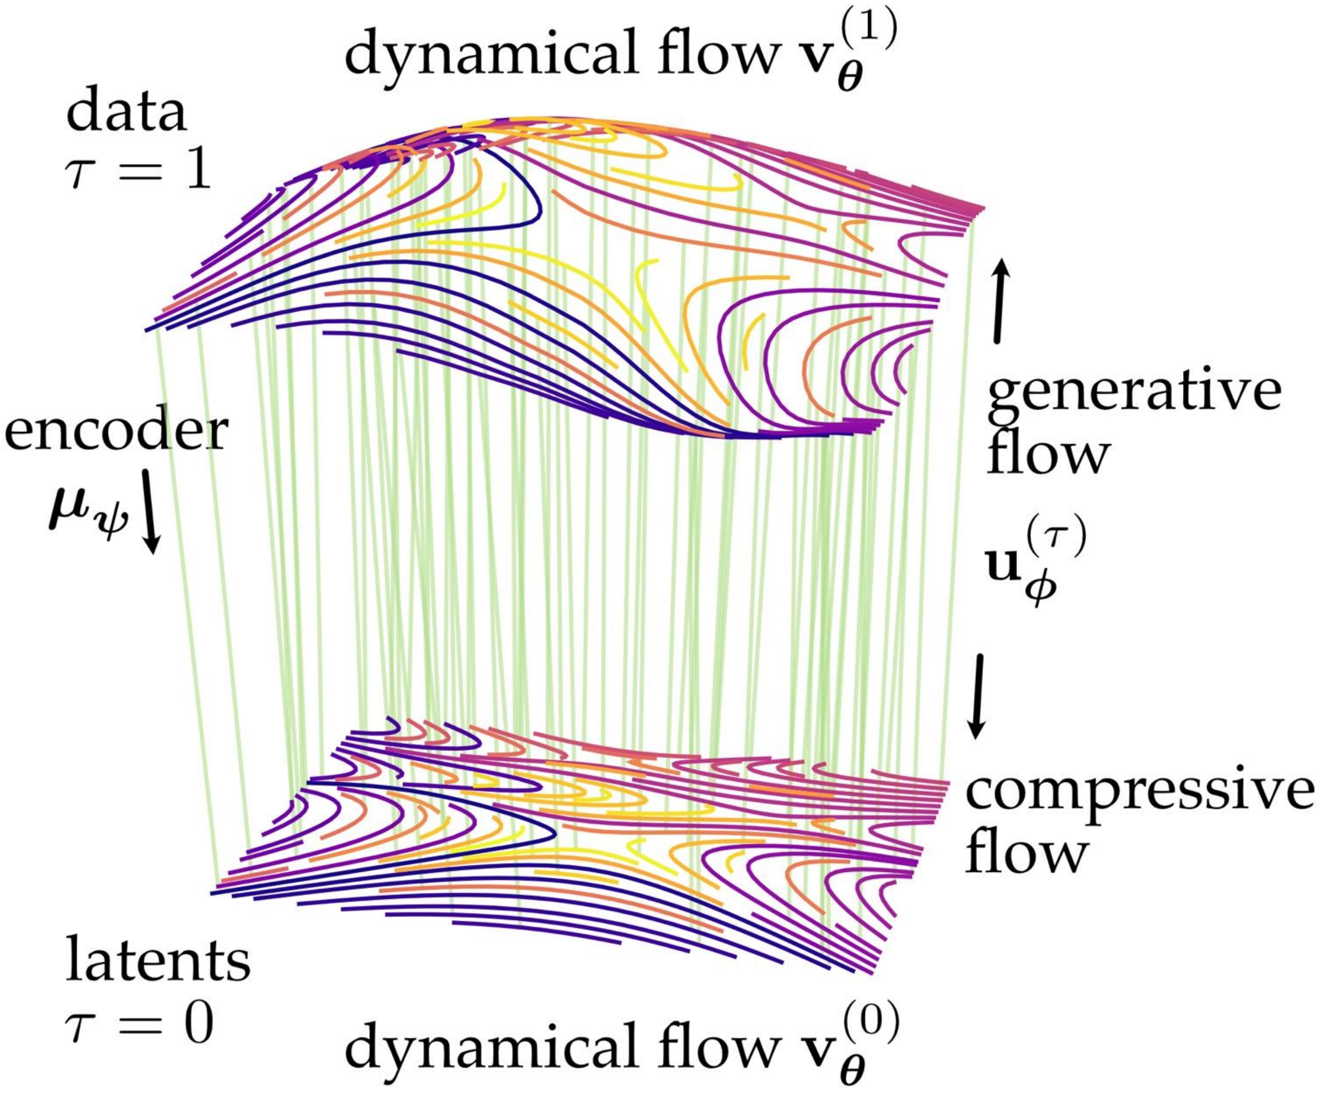

Compression Target Distribution

Nested Dropout

- Sample prefix length

$$K \sim \mathrm{Geom}(p),\ \mathbb{E}(K) = 1/p$$

- Apply prefix mask

$$\mathbf{x}_t^{(0,K)} = \mathbf{b} + \mathbf{L}\mathbf{D}^{1/2}\boldsymbol{\mu}_{\boldsymbol{\psi}}^{(K)}(\mathbf{x}_t^{(1)})$$

$$\mathbf{m}_K(i) = \mathbf{1}\{i \leq K\}, \qquad \boldsymbol{\mu}_{\boldsymbol{\psi}}^{(K)}(\mathbf{x}) = \mathbf{m}_K \odot \boldsymbol{\mu}_{\boldsymbol{\psi}}(\mathbf{x})$$

- This construction yields
  - *Controllable* dimensionality reduction
  - *Ordered* and *identifiable* (up to sign) latent representations

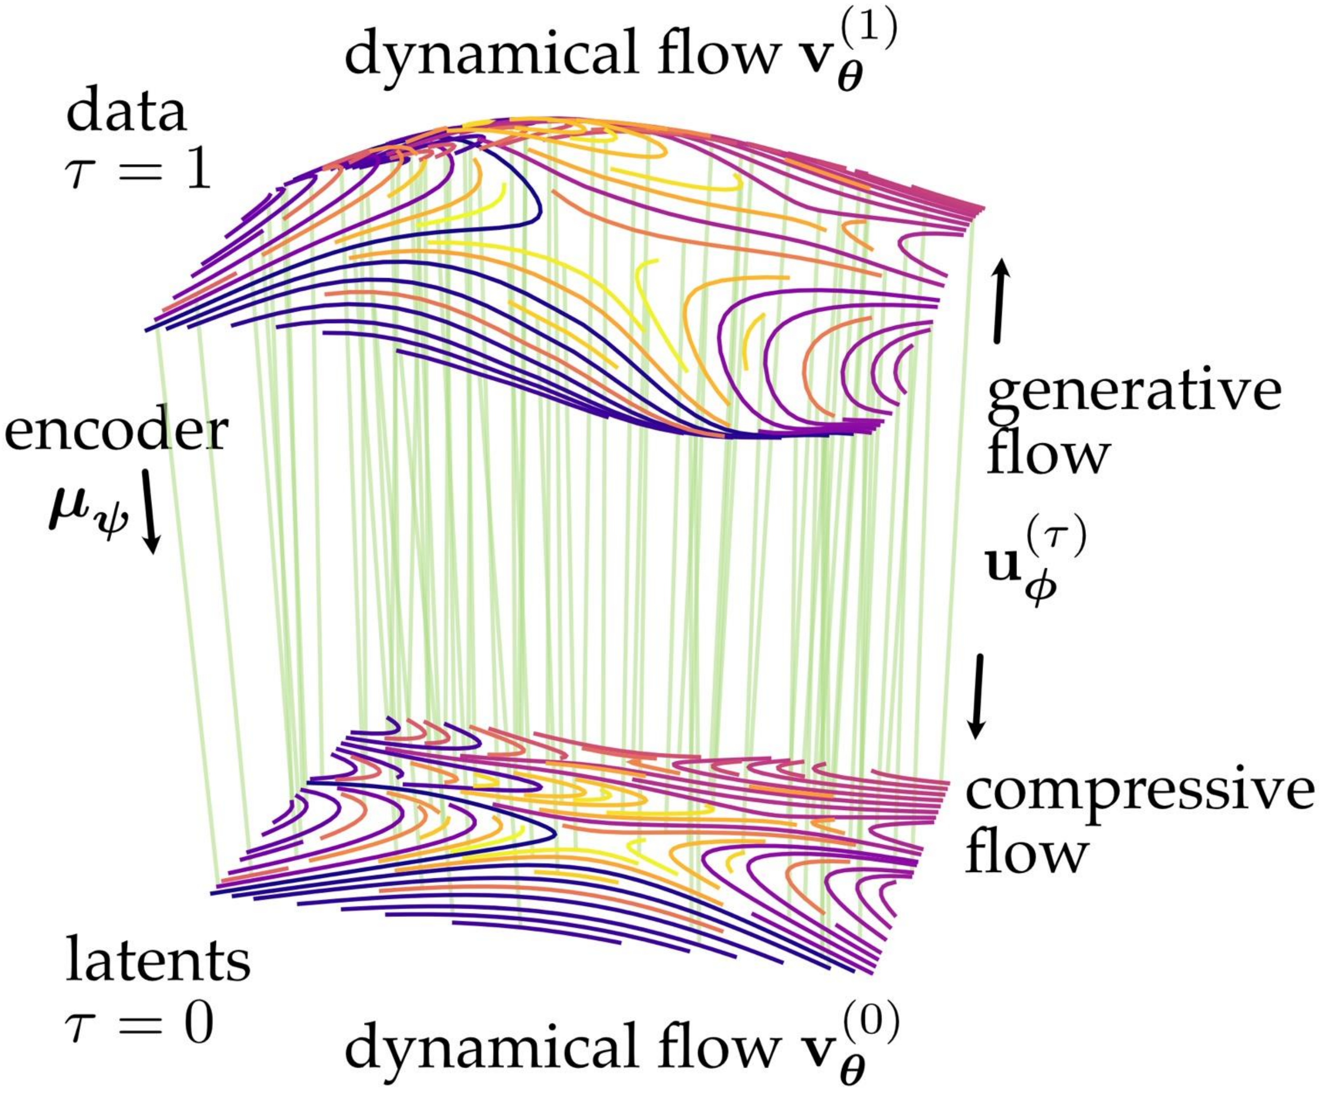

Coupled Flow Training

Objective

$$\mathcal{L} = \alpha\,\mathcal{L}_{\text{cf}} + \beta\,\mathcal{L}_{\text{df}} + \eta\,\mathcal{L}_{\text{align}}$$

$$\mathcal{L}_{\text{cf}} = \mathbb{E}_{\mathbf{x}}\mathbb{E}_{\tau}\left[\|\mathbf{u}_{\phi}(\mathbf{x}_t^{(\tau)}, \tau) - \mathbf{u}_t^{\star}\|_2^2\right]$$

Flow-Matching Loss Compressive Flow

$$\mathcal{L}_{\text{df}} = \mathbb{E}_{\mathbf{x}}\mathbb{E}_{\tau,s}\left[\|\mathbf{v}_{\boldsymbol{\theta}}(\mathbf{x}_t^{(\tau)}, \tau, s, \mathbf{x}_{\text{hist}}^{(\tau)}) - \mathbf{v}_{k,\tau}^{\star}\|_2^2\right]$$

Flow-Matching Loss Dynamical Flow

$$\mathcal{L}_{\text{align}} = \mathbb{E}_{\mathbf{x}}\mathbb{E}_{K}\left[\|\mathbf{x}_t^{(1)} - \mathbf{x}_t^{(0,K)}\|_2^2\right]$$

Latent Alignment Loss Encoder

Applications: Quick Overview

Mouse Behavioral Data

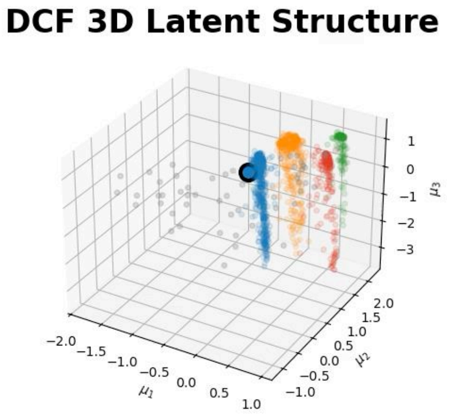

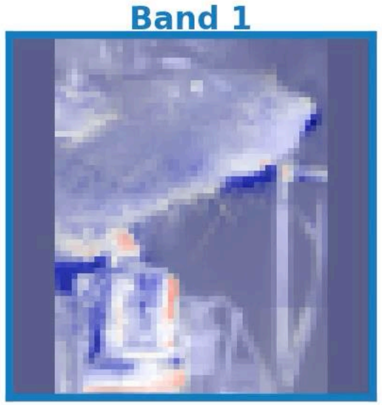

Bird Audio Data

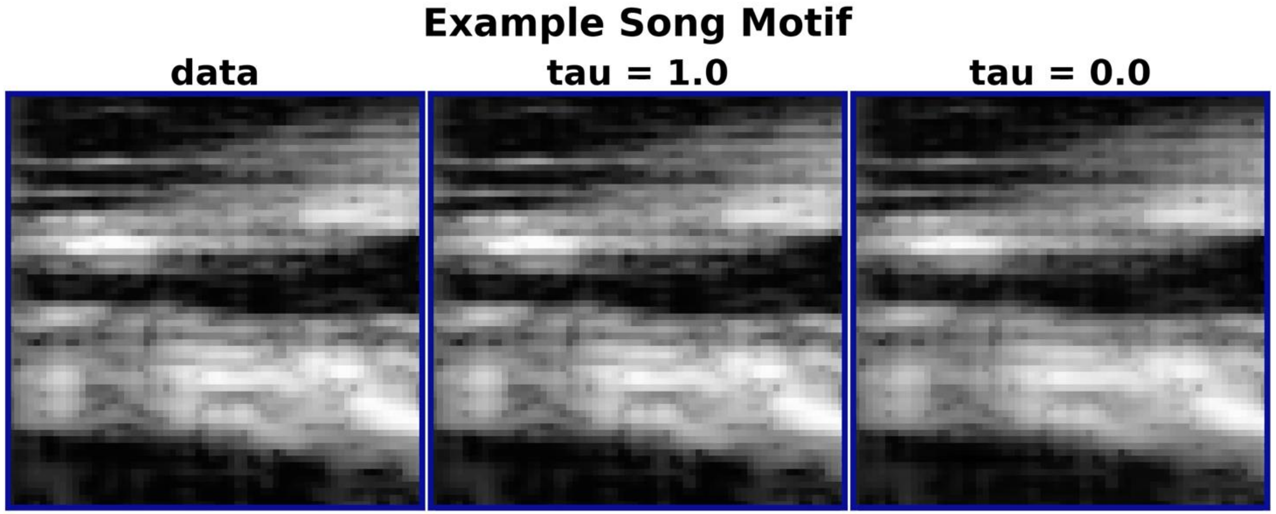

Neural Population Data

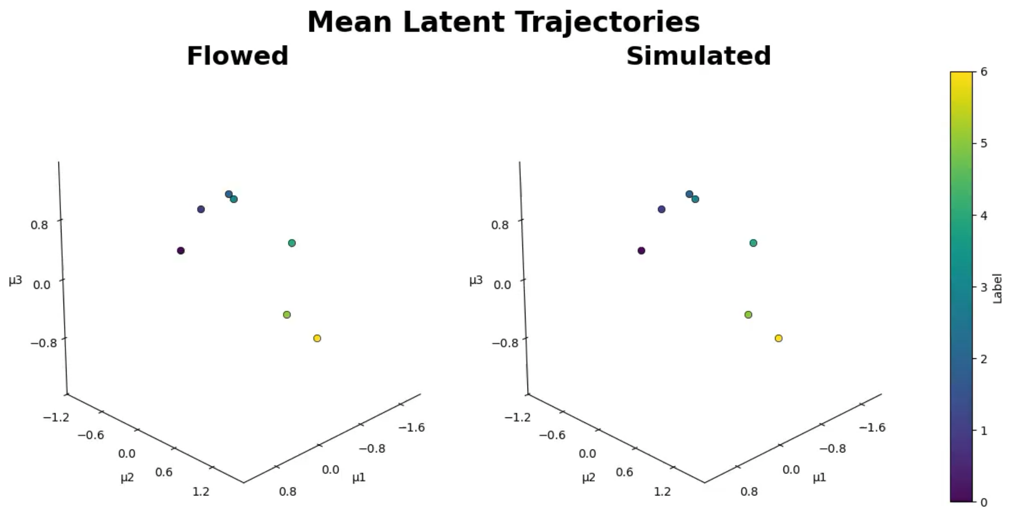

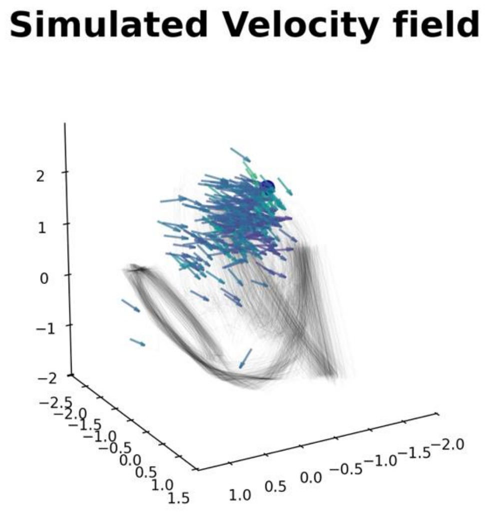

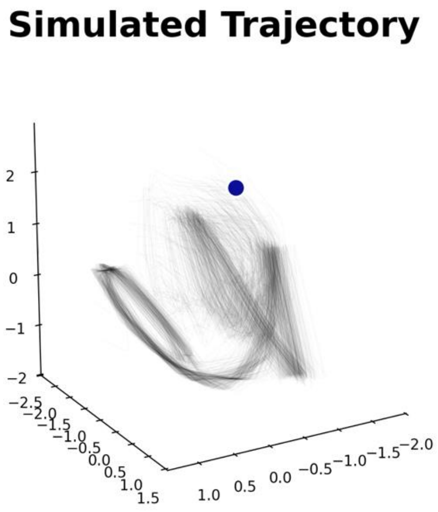

Thank You!

Pearson Lab Members:

John Pearson
Ganchao Wei*
Daniela de Albuquerque*
Trevor Alston
Ziyi Gong
Caitlin Lewis
Miles Martinez
Liz O’Gorman
Shiyang Pan
David St-Amand

NIH National Institute of Mental Health

NIH National Institute of Neurological Disorders and Stroke

NIH National Institute on Drug Abuse

RF1 DA056376
F30 MH129086
F31 NS132469

Duke University School of Medicine

Duke
PIERRE R. LAMOND DEPARTMENT *of*
ELECTRICAL & COMPUTER ENGINEERING