COMPARING NOISY NEURAL POPULATION DYNAMICS USING OPTIMAL TRANSPORT DISTANCES

Amin Nejatbakhsh[1], Victor Geadah[1,2], Alex H. Williams[1,3] & David Lipshutz[1,4]



[1] Center for Computation Neuroscience, Flatiron Institute; [2] Applied and Computational Mathematics, Princeton University; [3] Center for Neural Science, New York University; [4] Department of Neuroscience, Baylor College of Medicine

Central question: how similar are these representations?

Central question: how similar are these representations?

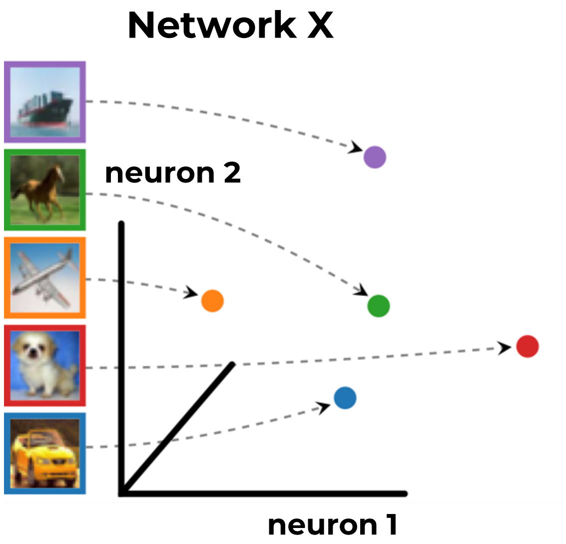

Central question: how similar are these representations?

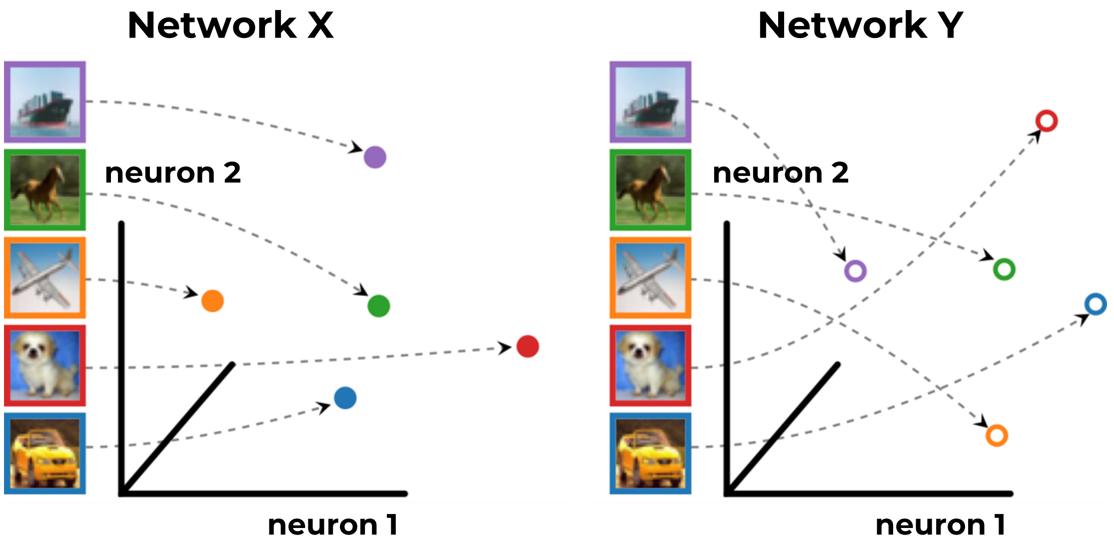

Central question: how similar are these representations?

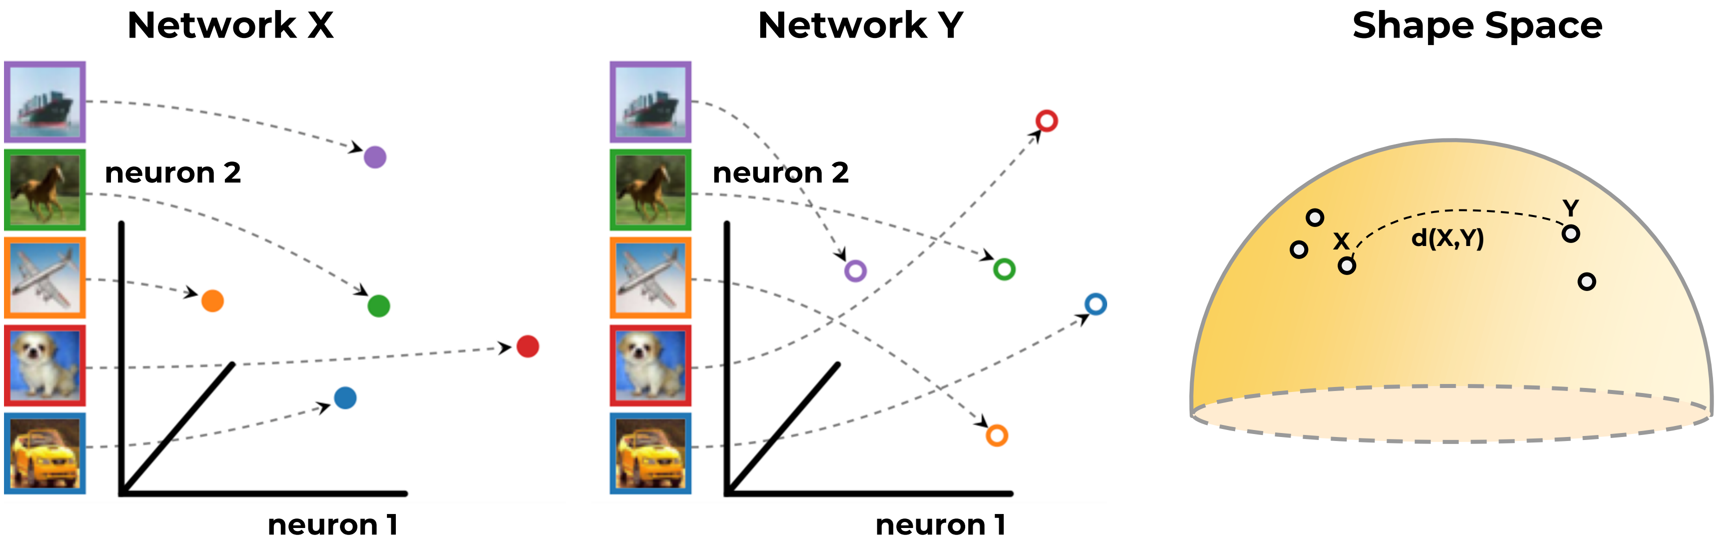

Central question: how similar are these representations?

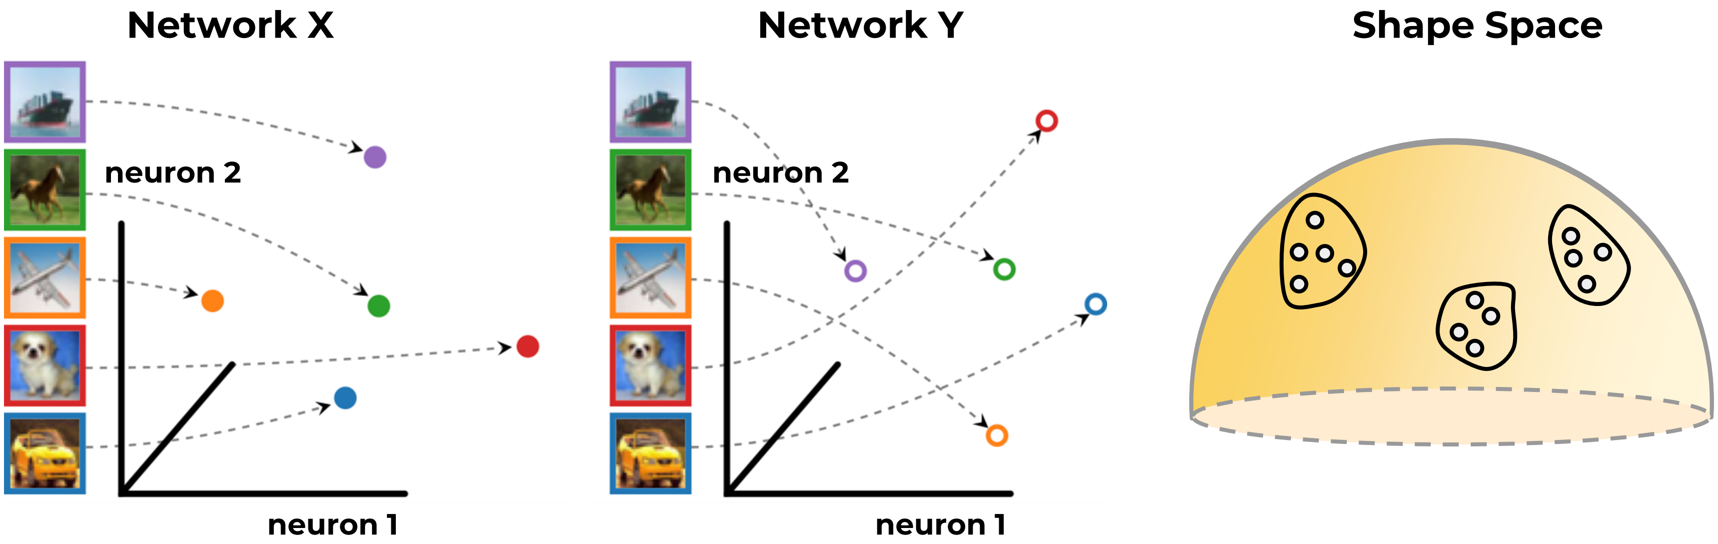

Central question: how similar are these representations?

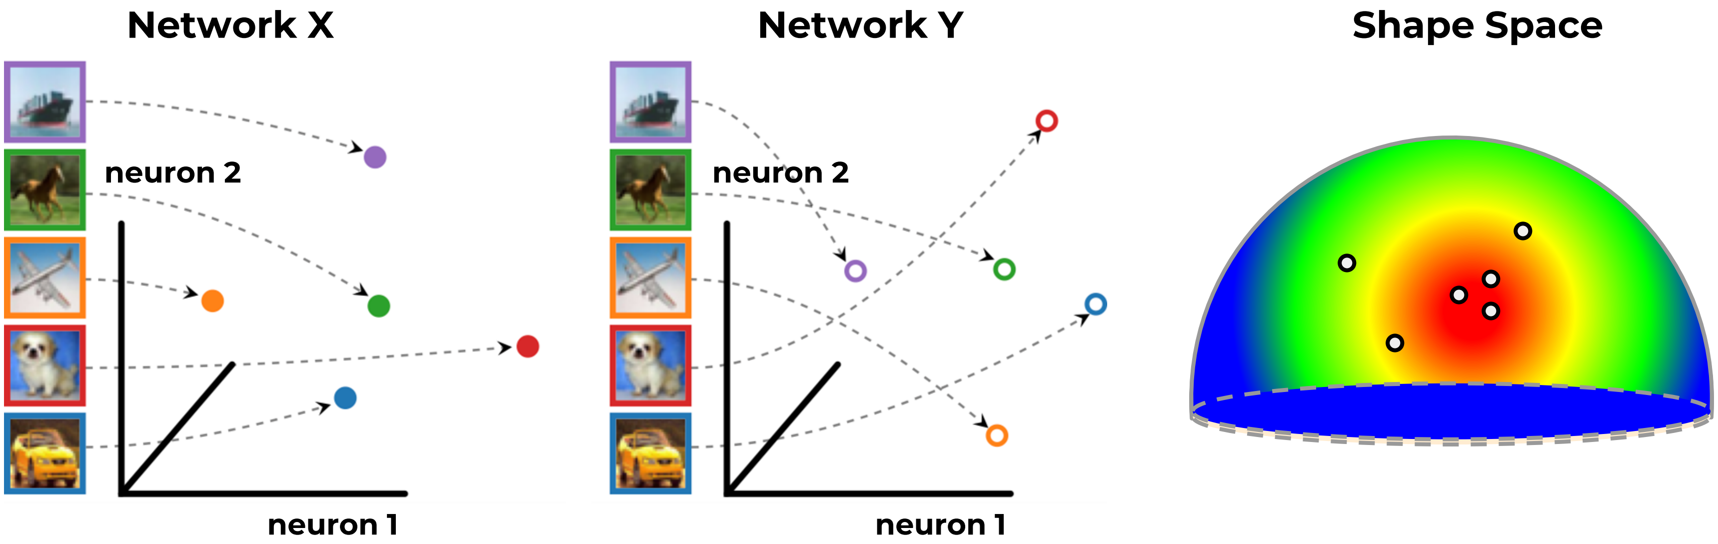

Representational similarity methods cover a range of assumptions

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

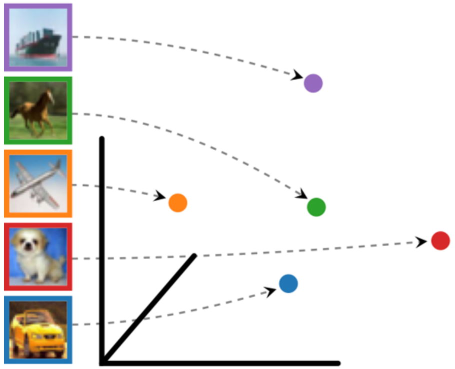

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

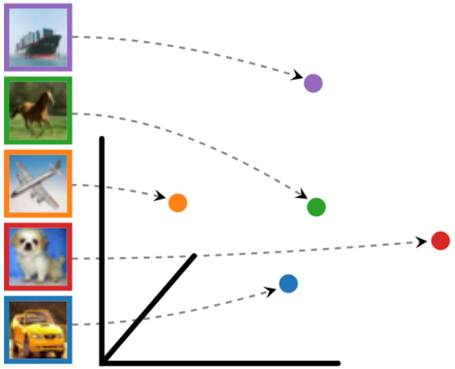

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

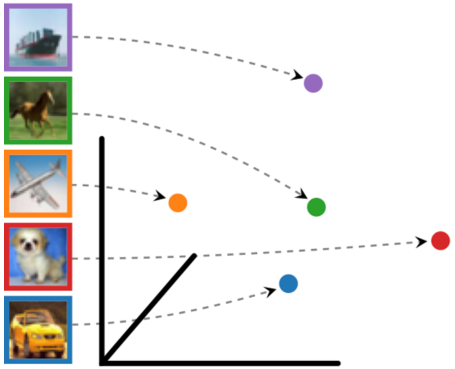

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Applications: ConvNets, MLPs, etc.

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

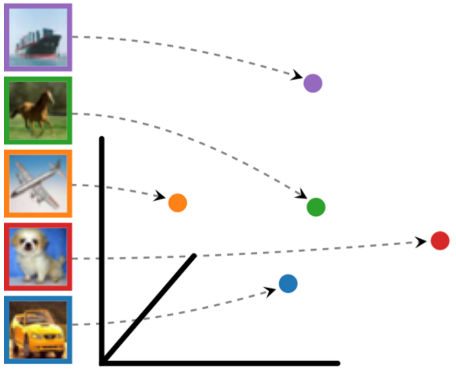

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Applications: ConvNets, MLPs, etc.

Stochastic responses

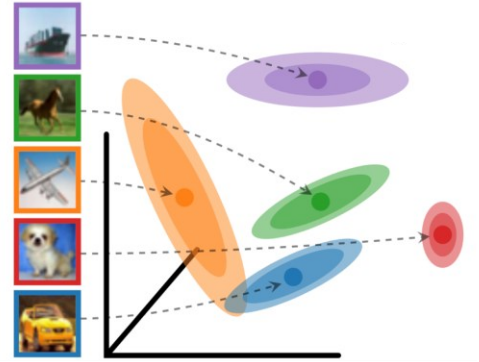

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

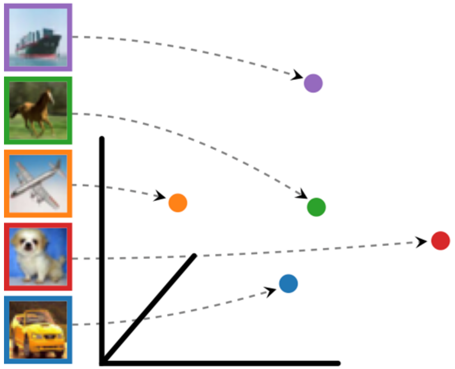

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Applications: ConvNets, MLPs, etc.

Stochastic responses

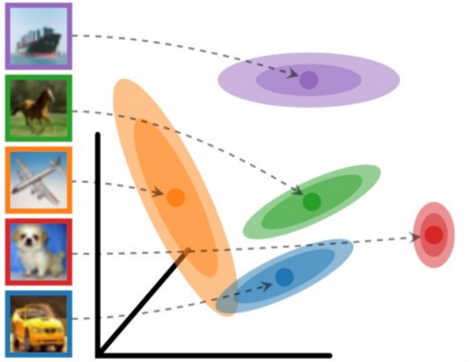

Stochastic Shape Distances (Duong et al. 2023)

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

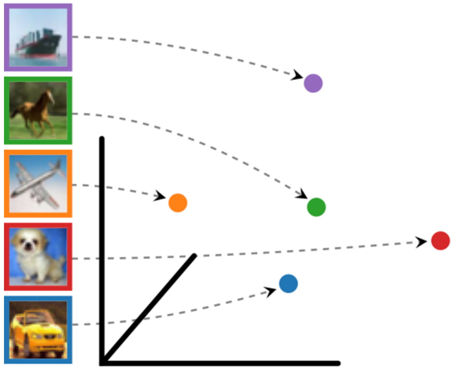

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Applications: ConvNets, MLPs, etc.

Stochastic responses

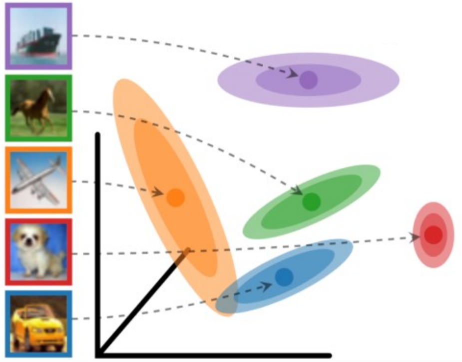

Stochastic Shape Distances (Duong et al. 2023)

Applications: VAEs, BNNs, Dropout, Noise Injection

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

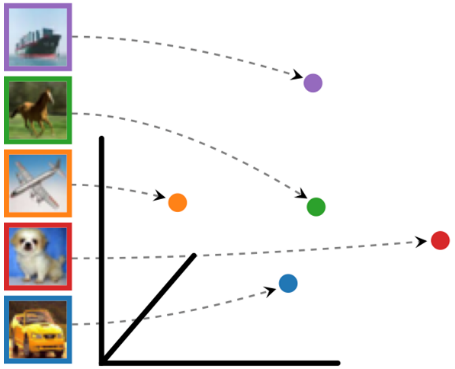

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Applications: ConvNets, MLPs, etc.

Stochastic responses

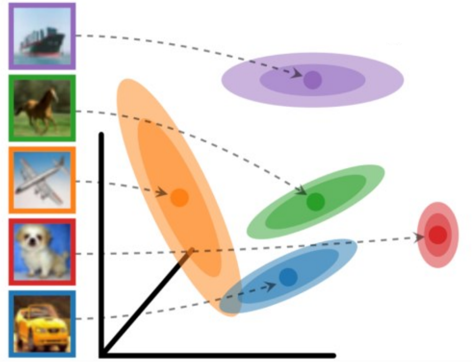

Stochastic Shape Distances (Duong et al. 2023)

Applications: VAEs, BNNs, Dropout, Noise Injection

Deterministic dynamic responses

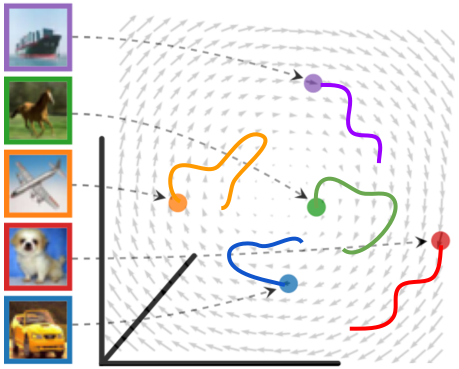

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

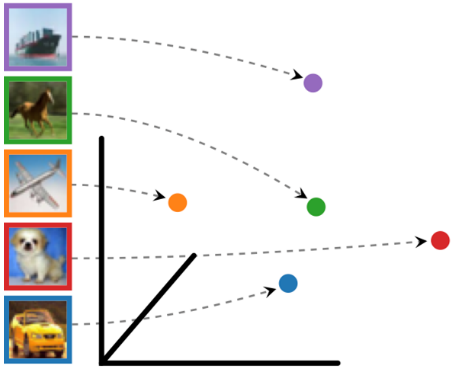

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Applications: ConvNets, MLPs, etc.

Stochastic responses

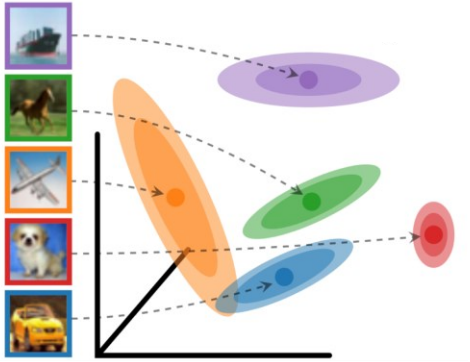

Stochastic Shape Distances (Duong et al. 2023)

Applications: VAEs, BNNs, Dropout, Noise Injection

Deterministic dynamic responses

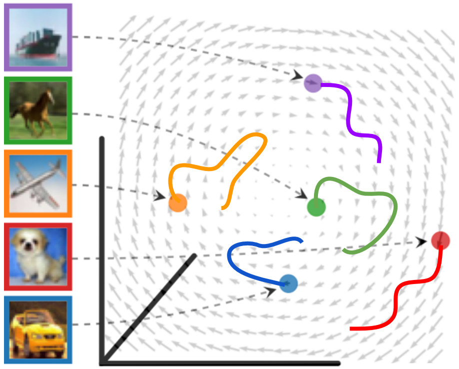

Dynamic Similarity Analysis (Ostrow et al. 2024)

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

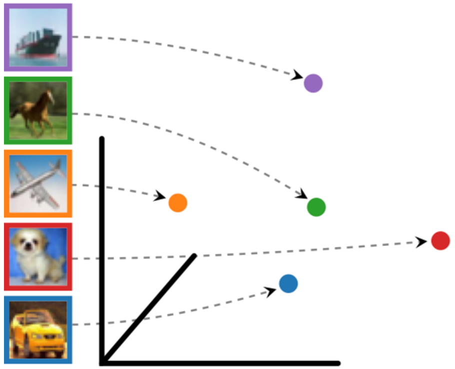

Linear regression (Yamins et al. 2014)
Representational Similarity Analysis (Kriegeskorte et al. 2008)
One-to-One Matching (Li et al. 2016)
Canonical Correlations Analysis (Raghu et al. 2017)
Centered Kernel Alignment (Kornblith et al. 2019)
Procrustes alignment (Degenhart et al. 2020)
Shape distances (Williams et al. 2021)

Applications: ConvNets, MLPs, etc.

Stochastic responses

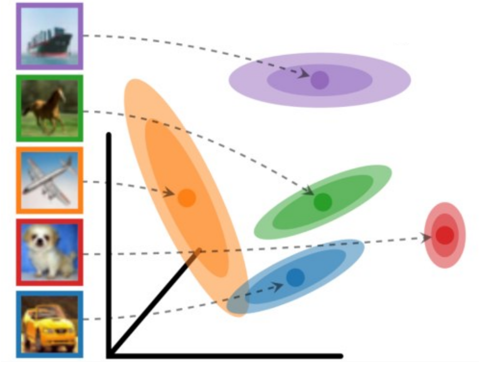

Stochastic Shape Distances (Duong et al. 2023)

Applications: VAEs, BNNs, Dropout, Noise Injection

Deterministic dynamic responses

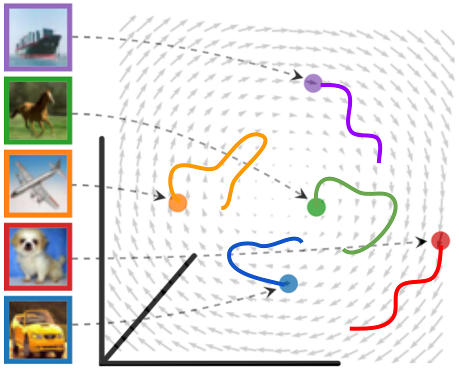

Dynamic Similarity Analysis (Ostrow et al. 2024)

Applications: RNNs, SSMs (e.g. MAMBA), Transformers, Diffusion Models, biological systems, etc.

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

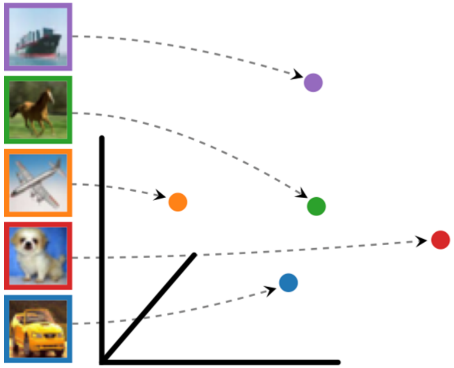

Stochastic responses

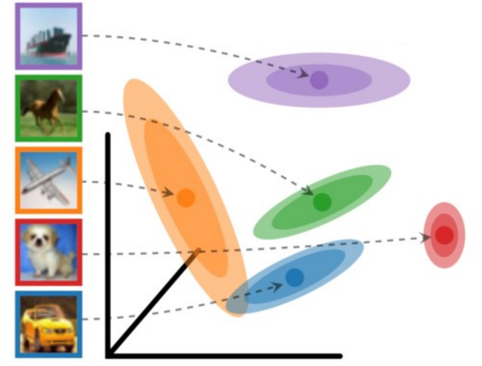

Deterministic dynamic responses

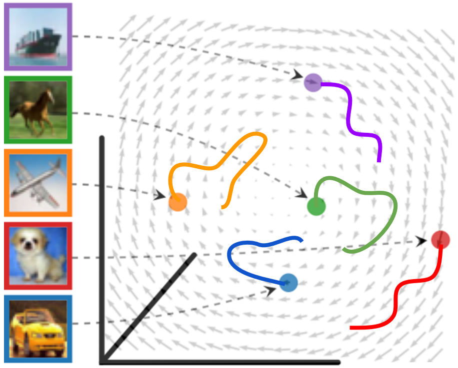

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

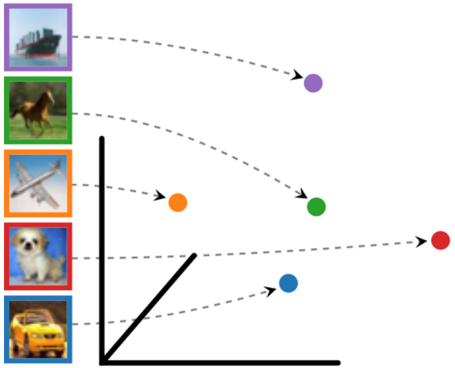

Stochastic responses

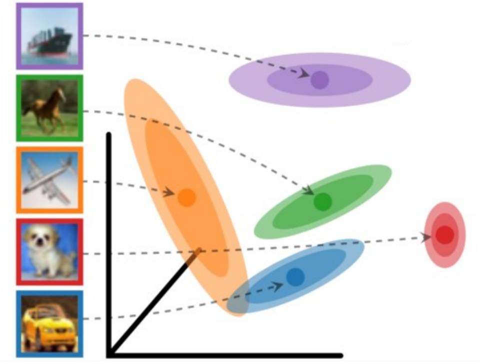

Deterministic dynamic responses

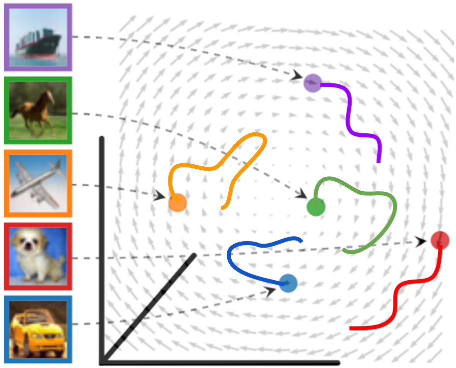

Procrustes Distance

$$\min_{\boldsymbol{Q}\in O(N)} \sum_{c=1}^{C} \|\underline{\boldsymbol{m}_x(c)} - \underline{\boldsymbol{Q}}\boldsymbol{m}_y(c)\|_2^2$$

Mean **Rotational alignment**

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

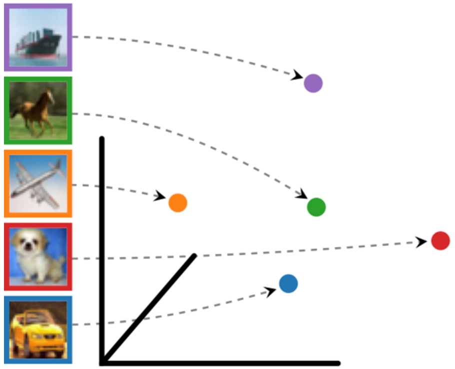

Stochastic responses

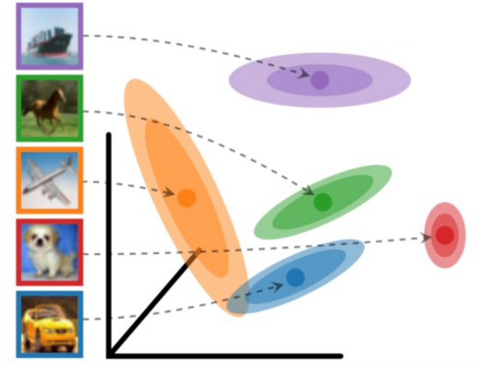

Deterministic dynamic responses

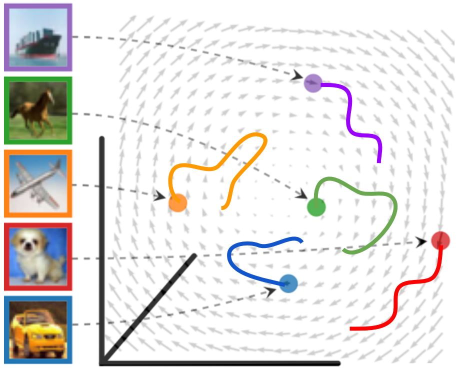

Procrustes Distance

$$\min_{Q\in O(N)} \sum_{c=1}^{C} \|\boldsymbol{m}_x(c) - \boldsymbol{Q}\boldsymbol{m}_y(c)\|_2^2$$

Mean

Rotational alignment

Stochastic Shape Distance (SSD)

$$\min_{Q\in O(N)} \sum_{c=1}^{C} \{(2-\alpha)\|\boldsymbol{m}_x(c) - \boldsymbol{Q}\boldsymbol{m}_y(c)\|^2 + \alpha\mathcal{B}^2\left(\boldsymbol{P}_x(c), \mathbf{Q}\boldsymbol{P}_y(c)\mathbf{Q}^\top\right)\},$$

Covariance

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

Stochastic responses

Deterministic dynamic responses

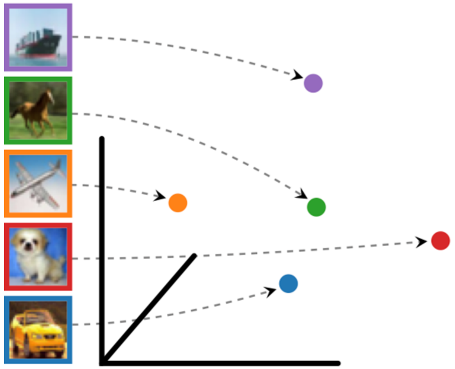

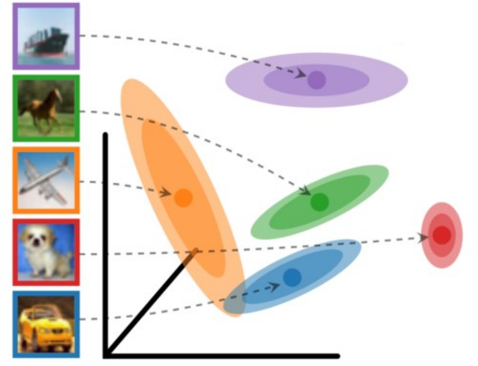

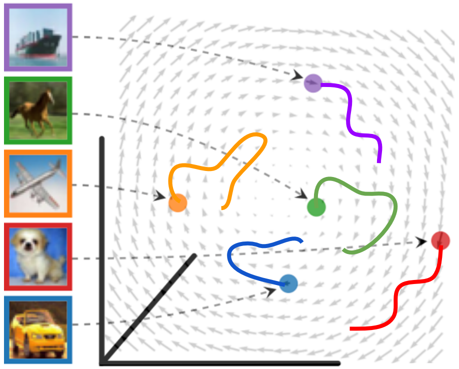

Procrustes Distance

$$\min_{Q\in O(N)} \sum_{c=1}^{C} \|\boldsymbol{m}_x(c) - \boldsymbol{Q}\boldsymbol{m}_y(c)\|_2^2$$

Mean

Rotational alignment

Stochastic Shape Distance (SSD)

$$\min_{Q\in O(N)} \sum_{c=1}^{C} \left\{(2-\alpha)\|\boldsymbol{m}_x(c) - \mathbf{Q}\boldsymbol{m}_y(c)\|^2 + \alpha\mathcal{B}^2\left(\boldsymbol{P}_x(c), \mathbf{Q}\boldsymbol{P}_y(c)\mathbf{Q}^\top\right)\right\},$$

Covariance

Dynamic Similarity Analysis (DSA)

$$\min_{Q\in O(N)} \|\boldsymbol{A}_x - \boldsymbol{Q}\boldsymbol{A}_y\boldsymbol{Q}^\top\|_F$$

Linearized flow field

Figures from [2,3]

Representational similarity methods cover a range of assumptions

Deterministic responses

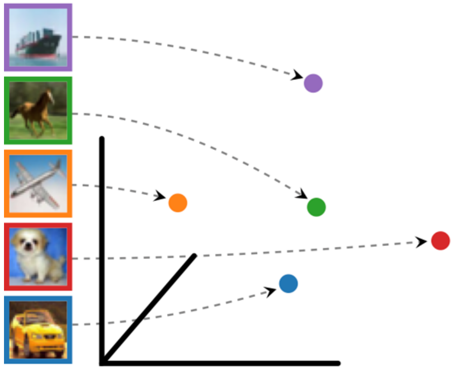

Procrustes Distance

$$\min_{Q \in O(N)} \sum_{c=1}^{C} \|\boldsymbol{m}_x(c) - \boldsymbol{Q}\boldsymbol{m}_y(c)\|_2^2$$

Mean

Rotational alignment

Stochastic responses

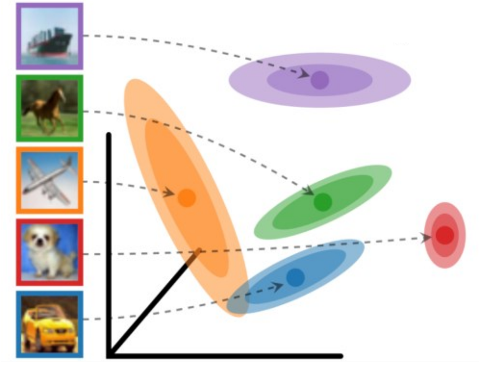

Stochastic Shape Distance (SSD)

$$\min_{Q \in O(N)} \sum_{c=1}^{C} \{(2-\alpha)\|\boldsymbol{m}_x(c) - \mathbf{Q}\boldsymbol{m}_y(c)\|^2 + \alpha\mathcal{B}^2\left(\boldsymbol{P}_x(c), \mathbf{Q}\boldsymbol{P}_y(c)\mathbf{Q}^\top\right)\},$$

Covariance

Deterministic dynamic responses

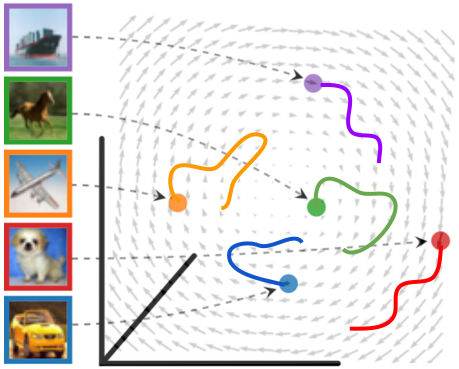

Dynamic Similarity Analysis (DSA)

$$\min_{Q \in O(N)} \|\boldsymbol{A}_x - \boldsymbol{Q}\boldsymbol{A}_y\boldsymbol{Q}^\top\|_F$$

Linearized flow field

Existing methods assume either dynamic or stochastic responses; not both.

Figures from [2,3]

Stochasticity alone is not sufficient to capture representational similarities in noisy dynamical systems

Stochasticity alone is not sufficient to capture representational similarities in noisy dynamical systems

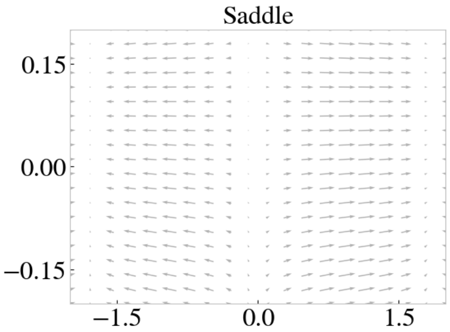

Stochasticity alone is not sufficient to capture representational similarities in noisy dynamical systems

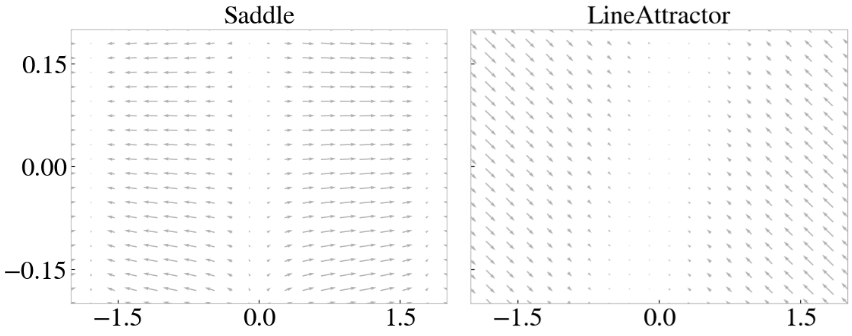

Stochasticity alone is not sufficient to capture representational similarities in noisy dynamical systems

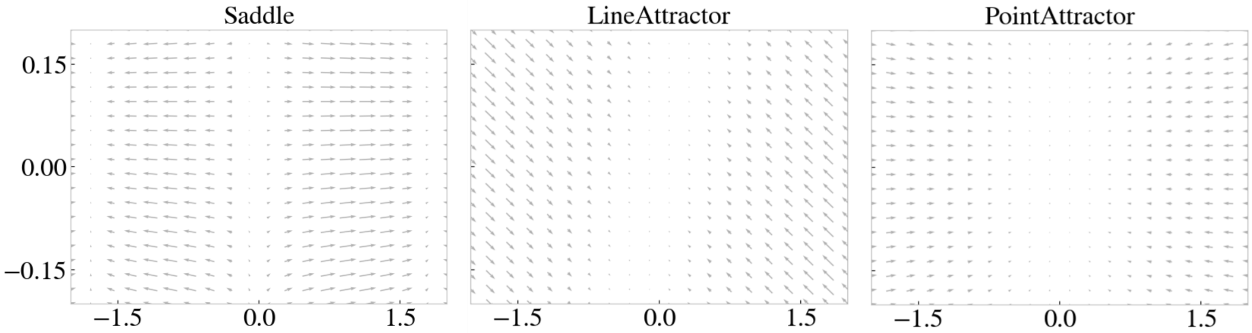

Stochasticity alone is not sufficient to capture representational similarities in noisy dynamical systems

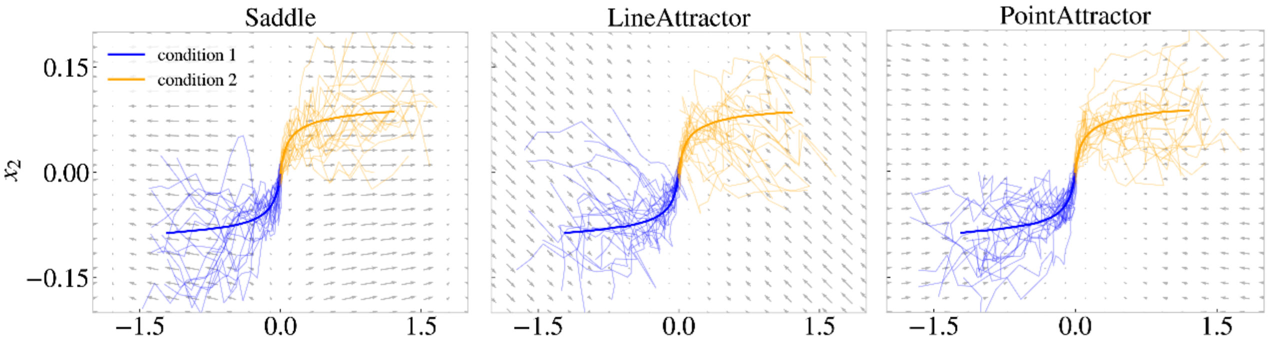

Stochasticity alone is not sufficient to capture representational similarities in noisy dynamical systems

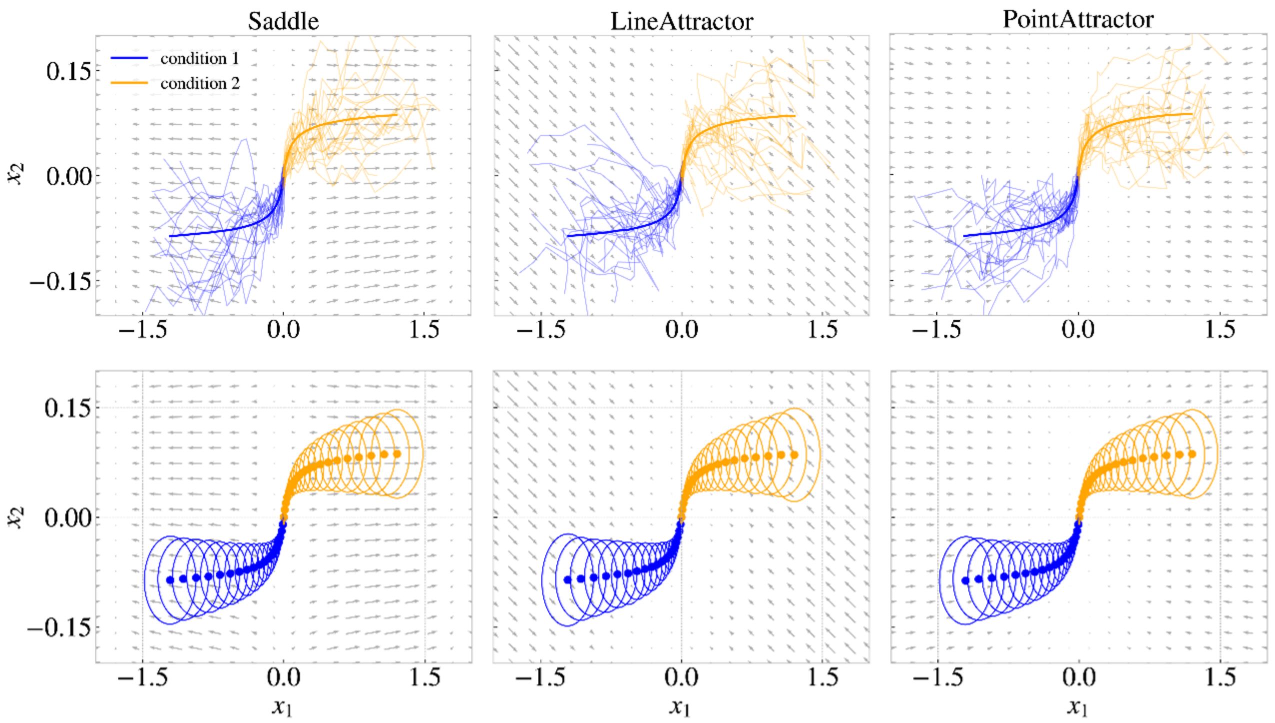

Key idea: compute distances on trajectories

Key idea: compute distances on trajectories

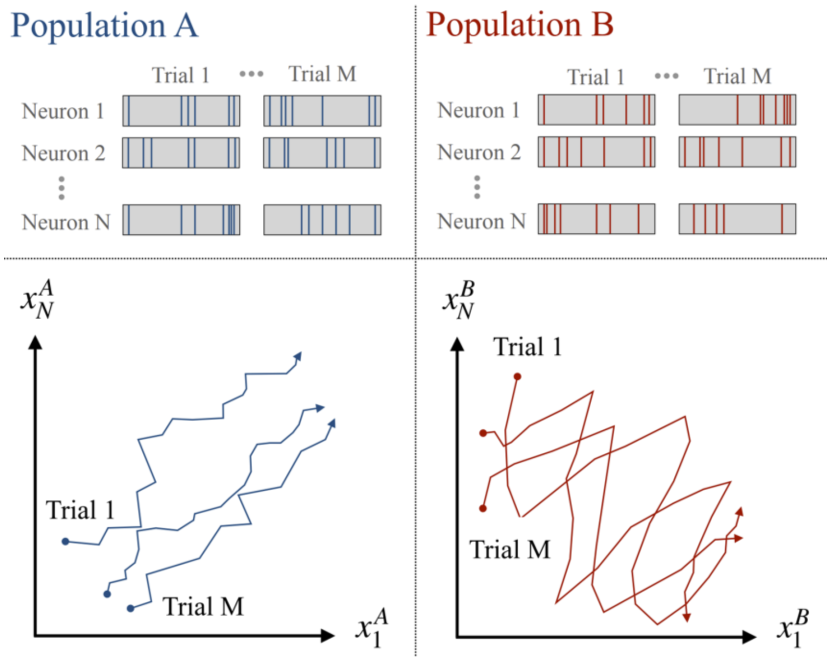

Key idea: compute distances on trajectories

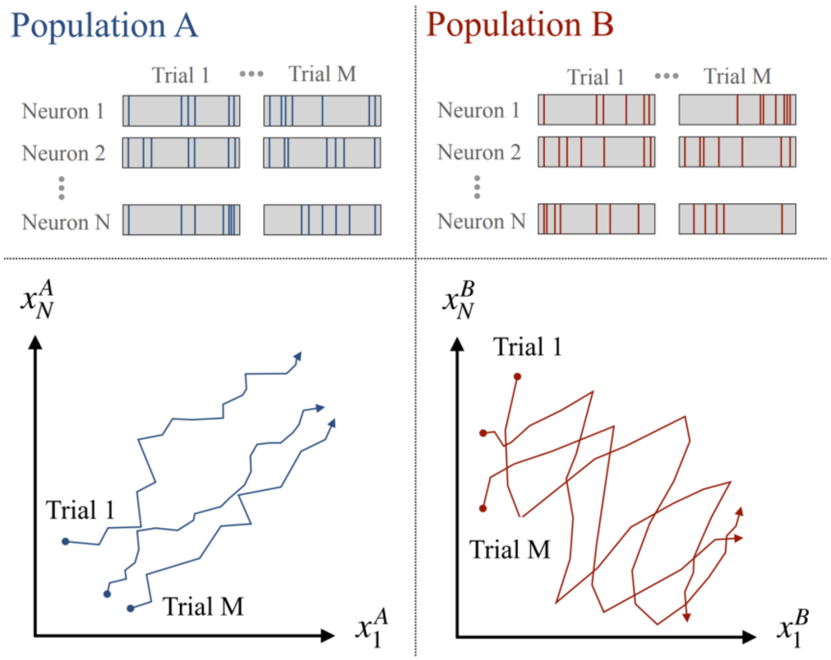

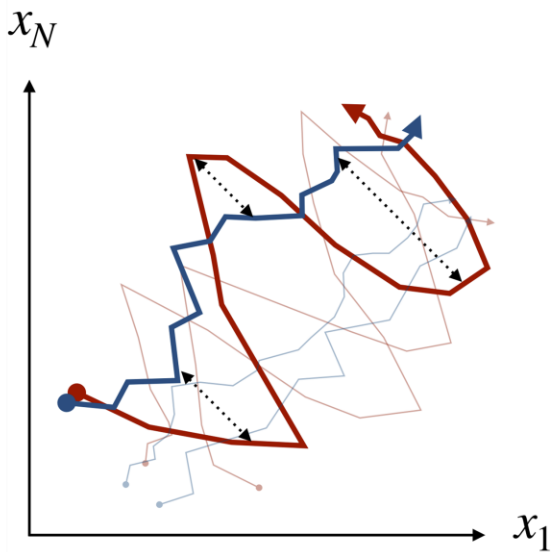

Causal Optimal Transport Distances

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

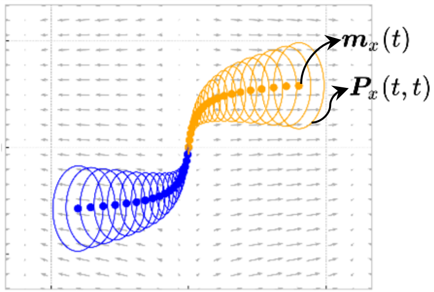

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

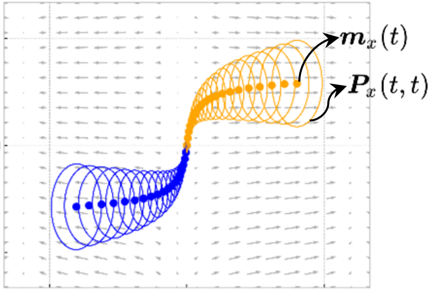

Causal OT

$$d_{\alpha\text{-causal}}(\mathbf{x},\mathbf{y}) := \min_{Q \in O(N)} \left\{ \sum_{t=1}^{T} (2-\alpha)\|\boldsymbol{m}_x(t) - \boldsymbol{Q}\boldsymbol{m}_y(t)\|^2 + \alpha \mathcal{AB}^2_{N,T}(\boldsymbol{C}_x, (\boldsymbol{I}_T \otimes \boldsymbol{Q})\boldsymbol{C}_y(\boldsymbol{I}_T \otimes \boldsymbol{Q}^\top)) \right\},$$

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

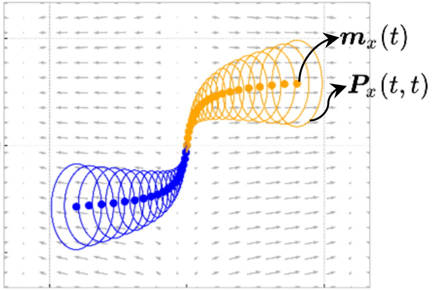

Causal OT

Aligned mean dist.

$$d_{\alpha\text{-causal}}(\mathbf{x},\mathbf{y}) := \min_{Q \in O(N)} \left\{ \sum_{t=1}^{T} (2-\alpha) \|\boldsymbol{m}_x(t) - \boldsymbol{Q}\boldsymbol{m}_y(t)\|^2 + \alpha \mathcal{AB}^2_{N,T}(\boldsymbol{C}_x, (\boldsymbol{I}_T \otimes \boldsymbol{Q})\boldsymbol{C}_y(\boldsymbol{I}_T \otimes \boldsymbol{Q}^\top)) \right\},$$

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

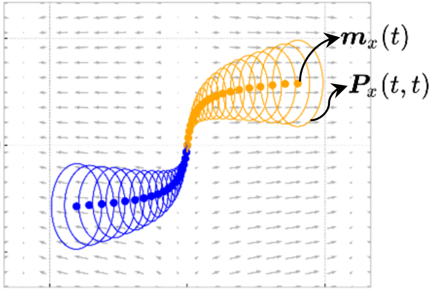

Causal OT

Aligned mean dist.

Aligned covariance dist.

$$d_{\alpha\text{-causal}}(\mathbf{x},\mathbf{y}) := \min_{Q \in O(N)} \left\{ \sum_{t=1}^{T} (2-\alpha) \|\boldsymbol{m}_x(t) - \boldsymbol{Q}\boldsymbol{m}_y(t)\|^2 + \alpha \mathcal{AB}^2_{N,T}(\boldsymbol{C}_x, (\boldsymbol{I}_T \otimes \boldsymbol{Q})\boldsymbol{C}_y(\boldsymbol{I}_T \otimes \boldsymbol{Q}^\top)) \right\},$$

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

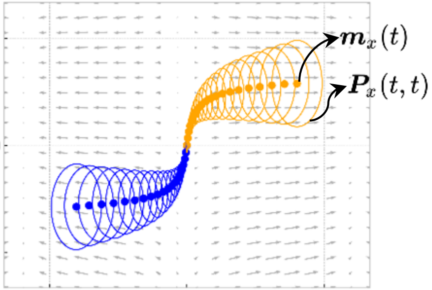

Causal OT

Aligned mean dist. **Aligned covariance dist.**

$$d_{\alpha\text{-causal}}(\mathbf{x},\mathbf{y}) := \min_{Q\in O(N)} \left\{ \sum_{t=1}^{T} (2-\alpha)\|\boldsymbol{m}_x(t) - \underline{\boldsymbol{Q}}\boldsymbol{m}_y(t)\|^2 + \alpha \mathcal{AB}^2_{N,T}(\boldsymbol{C}_x, (\boldsymbol{I}_T \otimes \boldsymbol{Q})\boldsymbol{C}_y(\boldsymbol{I}_T \otimes \boldsymbol{Q}^\top)) \right\},$$

Rotational alignment

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

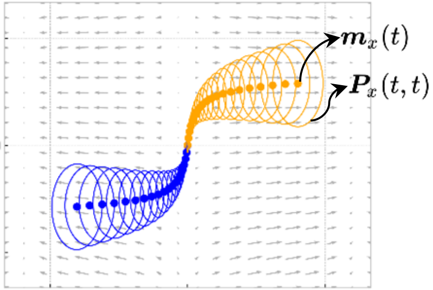

Causal OT

Aligned mean dist. **Aligned covariance dist.**

$$d_{\alpha\text{-causal}}(\mathbf{x},\mathbf{y}) := \min_{Q \in O(N)} \left\{ \sum_{t=1}^{T} (2-\alpha) \|\boldsymbol{m}_x(t) - \boldsymbol{Q}\boldsymbol{m}_y(t)\|^2 + \alpha \mathcal{AB}^2_{N,T}(\boldsymbol{C}_x, (\boldsymbol{I}_T \otimes \boldsymbol{Q})\boldsymbol{C}_y(\boldsymbol{I}_T \otimes \boldsymbol{Q}^\top)) \right\},$$

Rotational alignment

Adapted Bures dist.

Causal Optimal Transport Distances

Stochastic dynamics $\mathbf{x}(t)$

Mean trajectory $\boldsymbol{m}_x(t)$

Noise covariance $\boldsymbol{P}_x(s,t)$

Noise covariance

$$\boldsymbol{C}_x = \begin{bmatrix} \boldsymbol{P}_x(1,1) & \dots & \boldsymbol{P}_x(1,T) \\ \vdots & \ddots & \vdots \\ \boldsymbol{P}_x(T,1) & \dots & \boldsymbol{P}_x(T,T) \end{bmatrix}$$

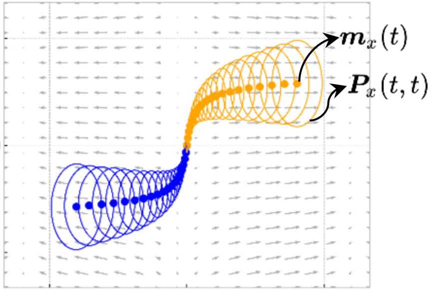

Causal OT

Aligned mean dist. Aligned covariance dist.

$$d_{\alpha\text{-causal}}(\mathbf{x},\mathbf{y}) := \min_{Q \in O(N)} \left\{ \sum_{t=1}^{T} (2-\alpha) \|\boldsymbol{m}_x(t) - \boldsymbol{Q}\boldsymbol{m}_y(t)\|^2 + \alpha \mathcal{AB}^2_{N,T}(\boldsymbol{C}_x, (\boldsymbol{I}_T \otimes \boldsymbol{Q})\boldsymbol{C}_y(\boldsymbol{I}_T \otimes \boldsymbol{Q}^\top)) \right\},$$

Rotational alignment

Adapted Bures dist.

Full NTxNT covariance

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

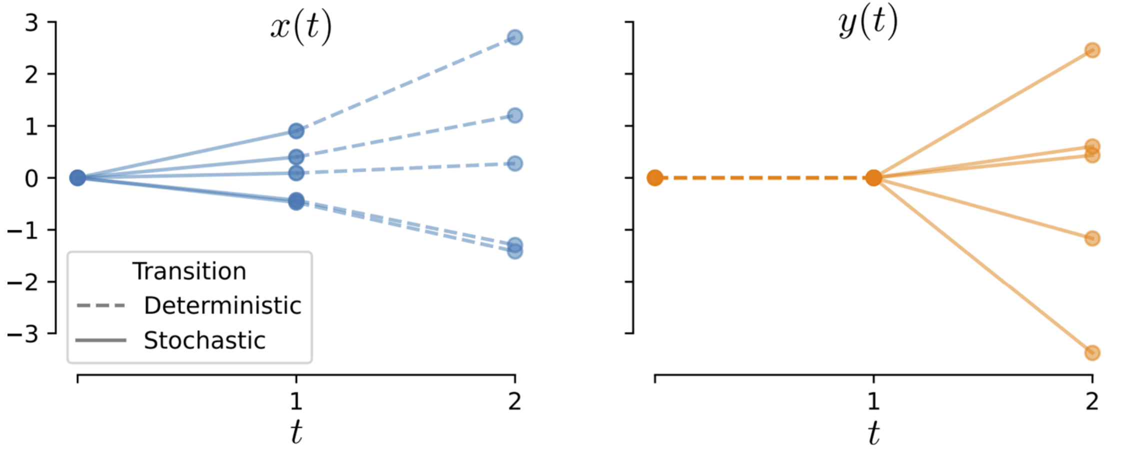

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

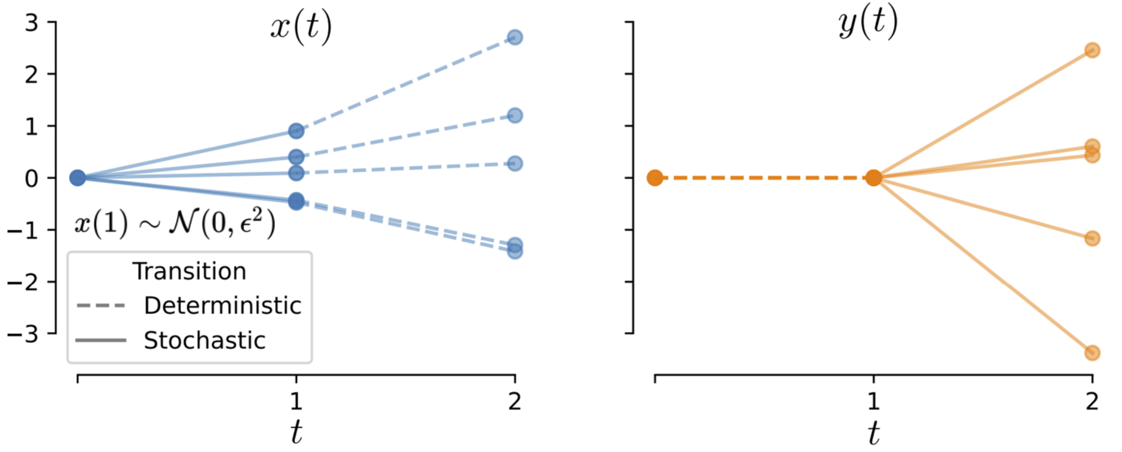

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

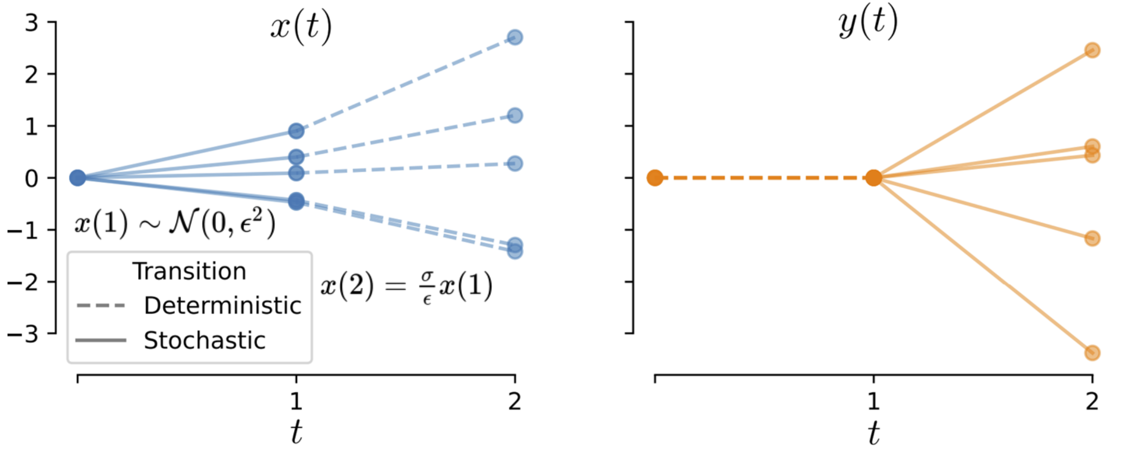

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

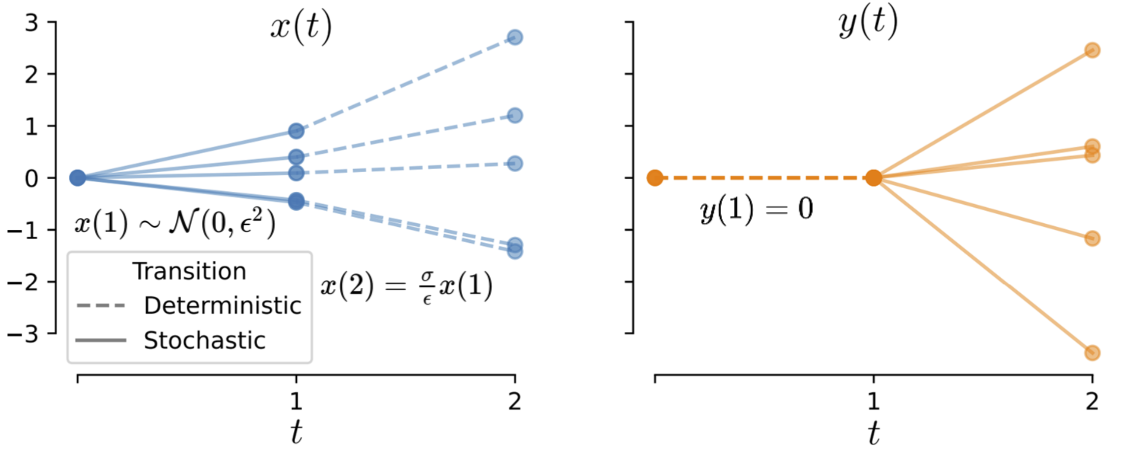

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

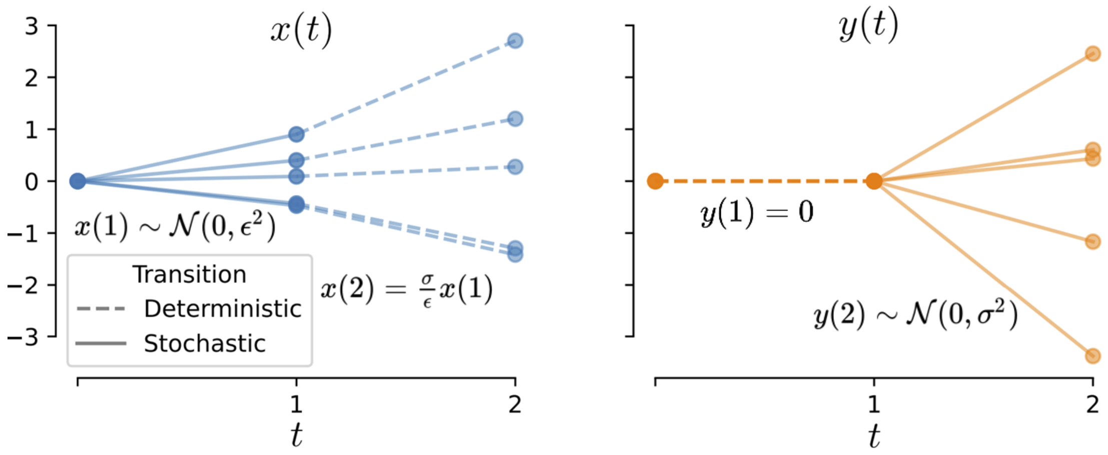

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

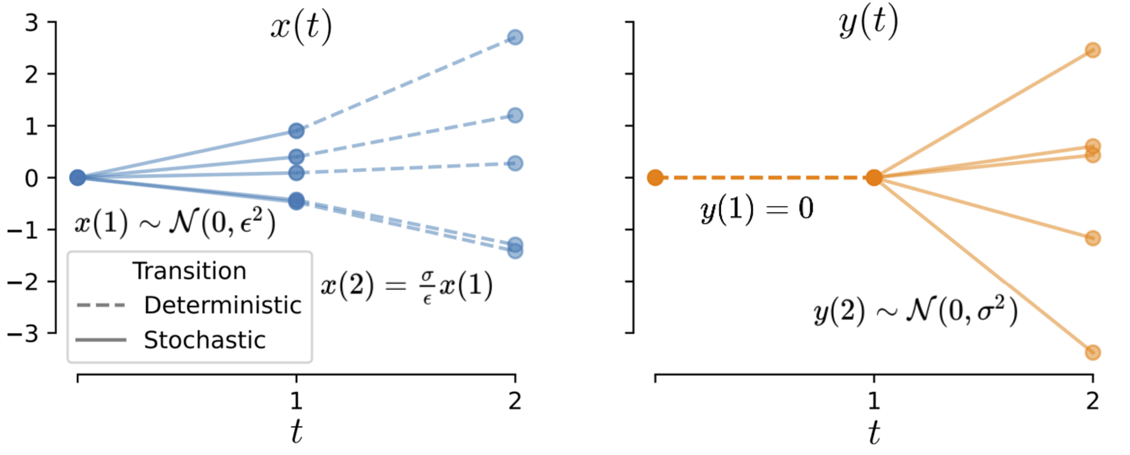

$d_{\text{Procrustes}}(x, y) = 0$ $\quad d_{\text{Wasserstein}}(x, y) = \epsilon$ $\quad d_{\text{1-SSD}}(x, y) = \epsilon$

Why is **Causal** Optimal Transport desirable? Scalar case for intuition building

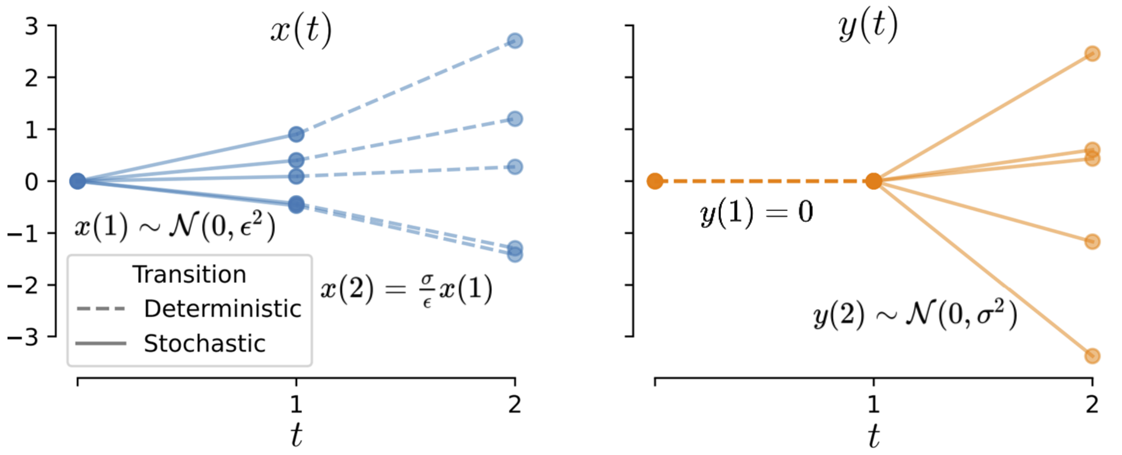

$d_{\text{Procrustes}}(x, y) = 0$ $\quad d_{\text{Wasserstein}}(x, y) = \epsilon$ $\quad d_{\text{1-SSD}}(x, y) = \epsilon$

$$d_{\text{1-causal}}(x, y) = \sqrt{\sigma^2 + \epsilon^2}$$

Motor preparatory dynamics in the null space of cortical activity

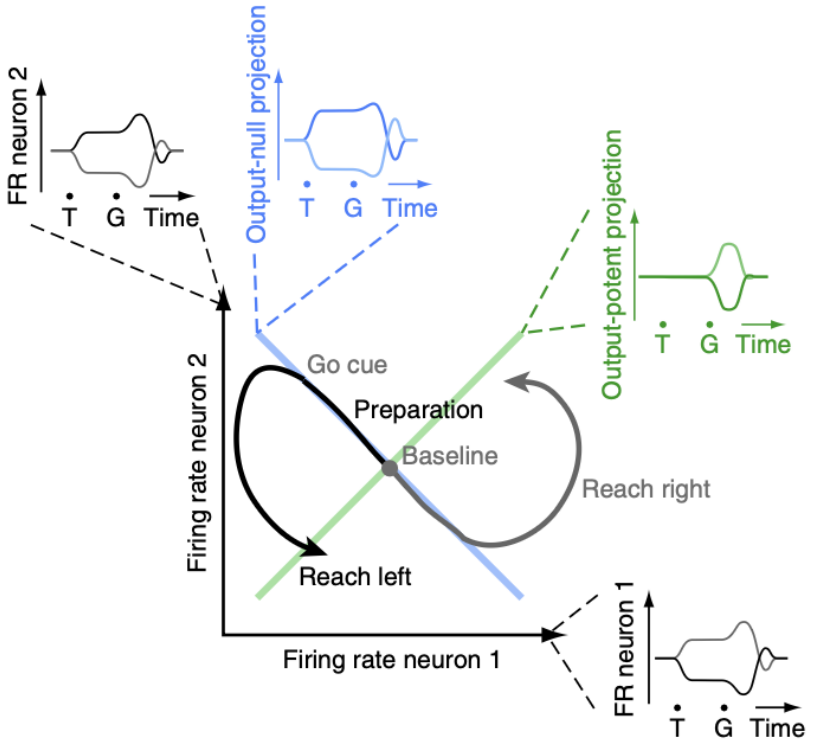

Figure from Kaufman et al, Nat. Neuro, 2014

Motor preparatory dynamics in the null space of cortical activity

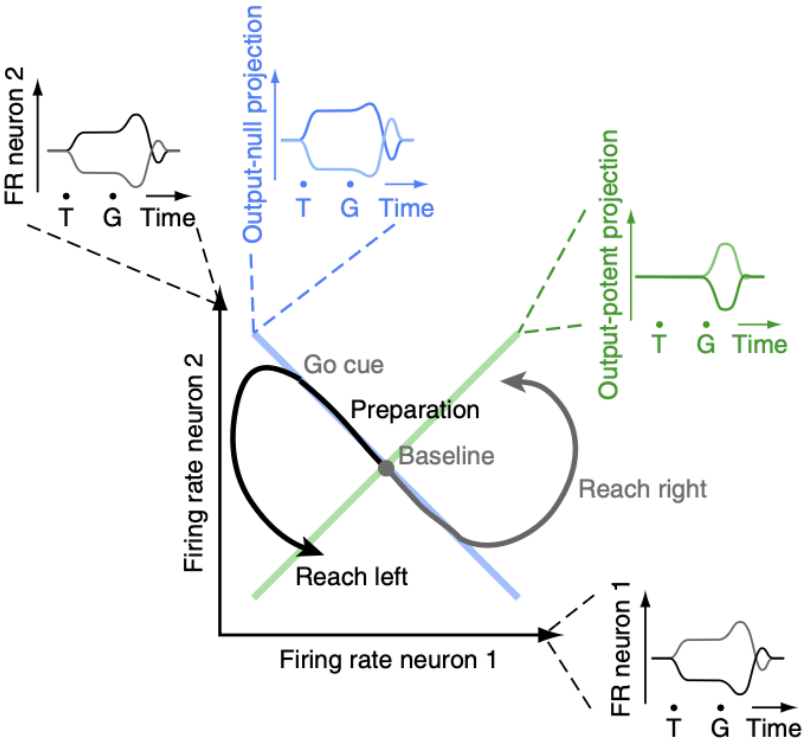

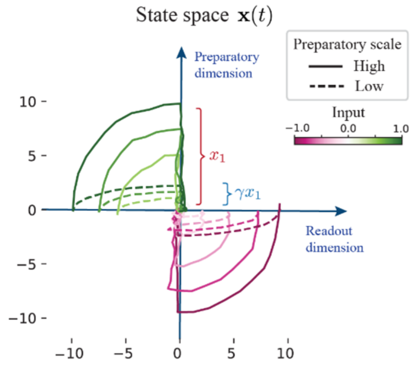

Figure from Kaufman et al, Nat. Neuro, 2014

Disambiguating recurrent and input driven dynamics

Disambiguating recurrent and input driven dynamics

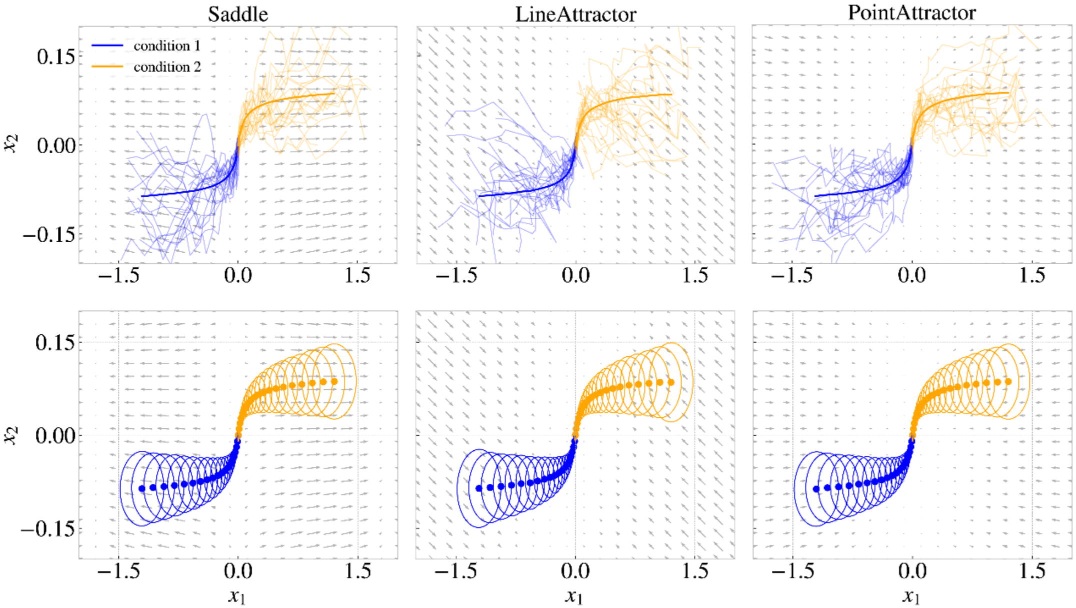

Disambiguating recurrent and input driven dynamics

Dynamics of mean and covariance

$$\begin{cases} \boldsymbol{m}_x(t) = \boldsymbol{A}(t)\boldsymbol{m}_x(t-1) + \boldsymbol{b}(t) \\ \boldsymbol{P}_x(t) = \boldsymbol{A}(t)\boldsymbol{P}_x(t-1)\boldsymbol{A}(t)^\top + \boldsymbol{\Sigma}(t)\boldsymbol{\Sigma}(t)^\top. \end{cases}$$

Adversarially tune inputs and latent noise $\begin{cases} \boldsymbol{b}(t) \\ \boldsymbol{\Sigma}(t) \end{cases}$

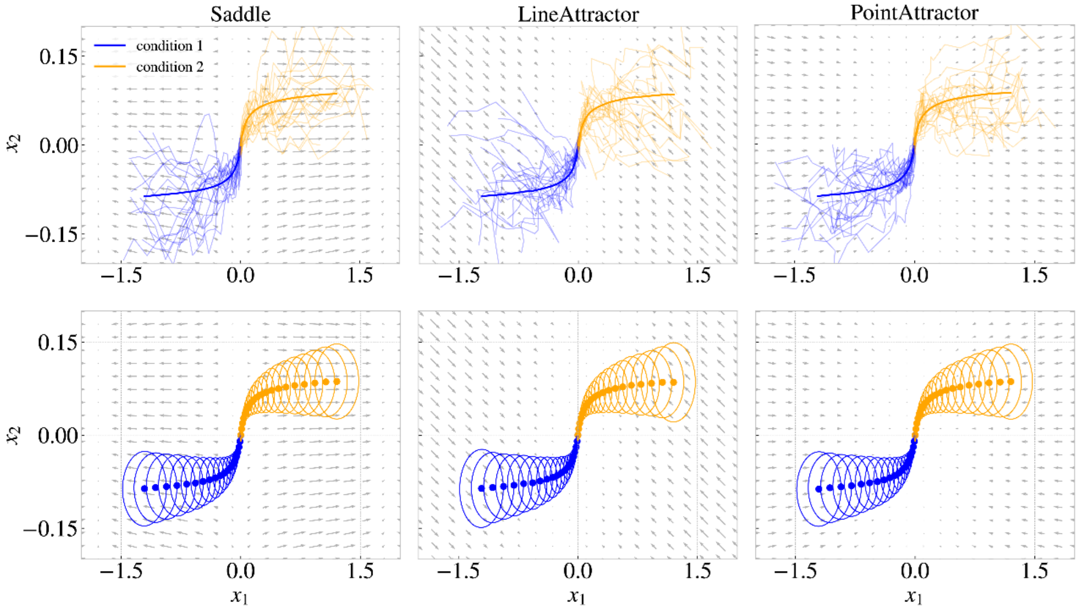

Disambiguating recurrent and input driven dynamics

Dynamics of mean and covariance

$$\begin{cases} \boldsymbol{m}_x(t) = \boldsymbol{A}(t)\boldsymbol{m}_x(t-1) + \boldsymbol{b}(t) \\ \boldsymbol{P}_x(t) = \boldsymbol{A}(t)\boldsymbol{P}_x(t-1)\boldsymbol{A}(t)^\top + \boldsymbol{\Sigma}(t)\boldsymbol{\Sigma}(t)^\top. \end{cases}$$

Adversarially tune inputs and latent noise $\begin{cases} \boldsymbol{b}(t) \\ \boldsymbol{\Sigma}(t) \end{cases}$

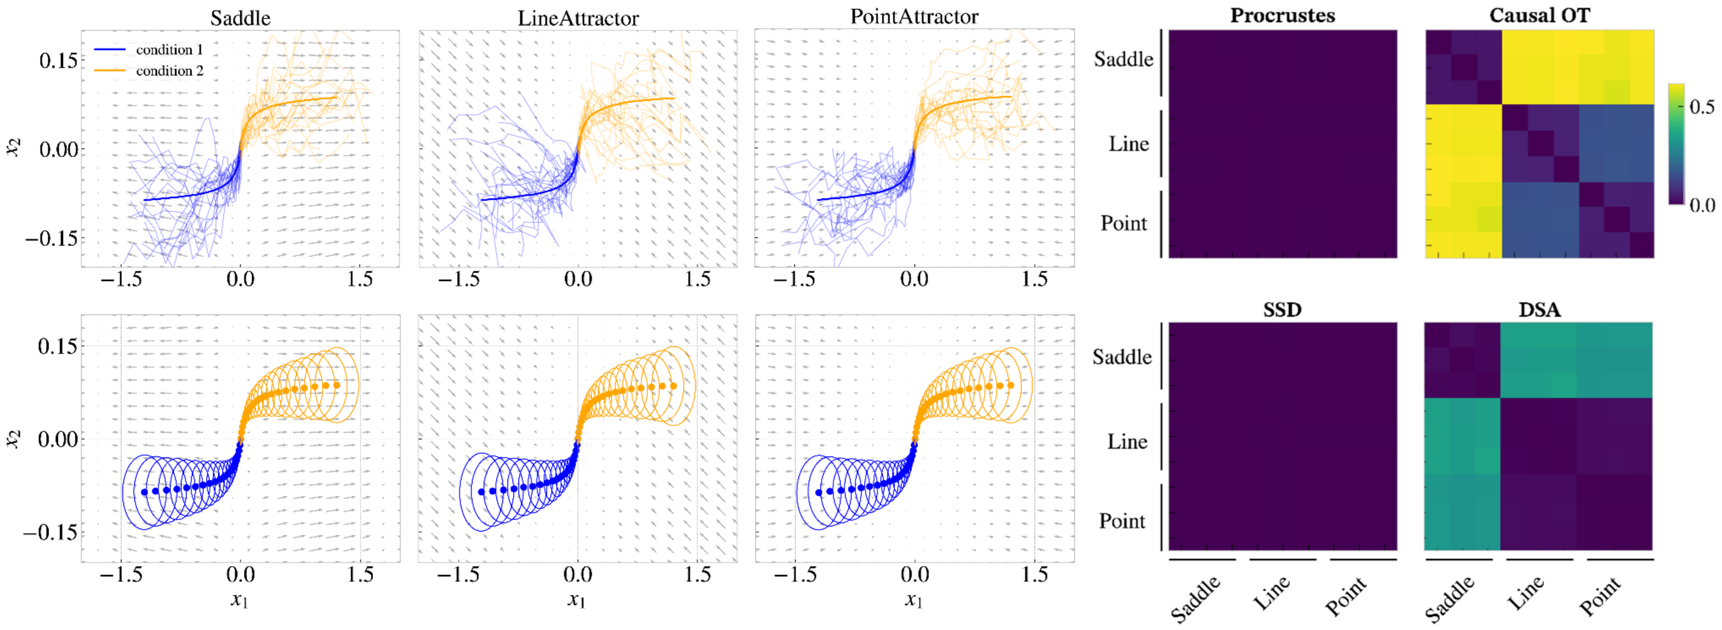

Disambiguating recurrent and input driven dynamics

Dynamics of mean and covariance
$$\begin{cases} \boldsymbol{m}_x(t) = \boldsymbol{A}(t)\boldsymbol{m}_x(t-1) + \boldsymbol{b}(t) \\ \boldsymbol{P}_x(t) = \boldsymbol{A}(t)\boldsymbol{P}_x(t-1)\boldsymbol{A}(t)^\top + \boldsymbol{\Sigma}(t)\boldsymbol{\Sigma}(t)^\top. \end{cases}$$

Adversarially tune inputs and latent noise
$$\begin{cases} \boldsymbol{b}(t) \\ \boldsymbol{\Sigma}(t) \end{cases}$$

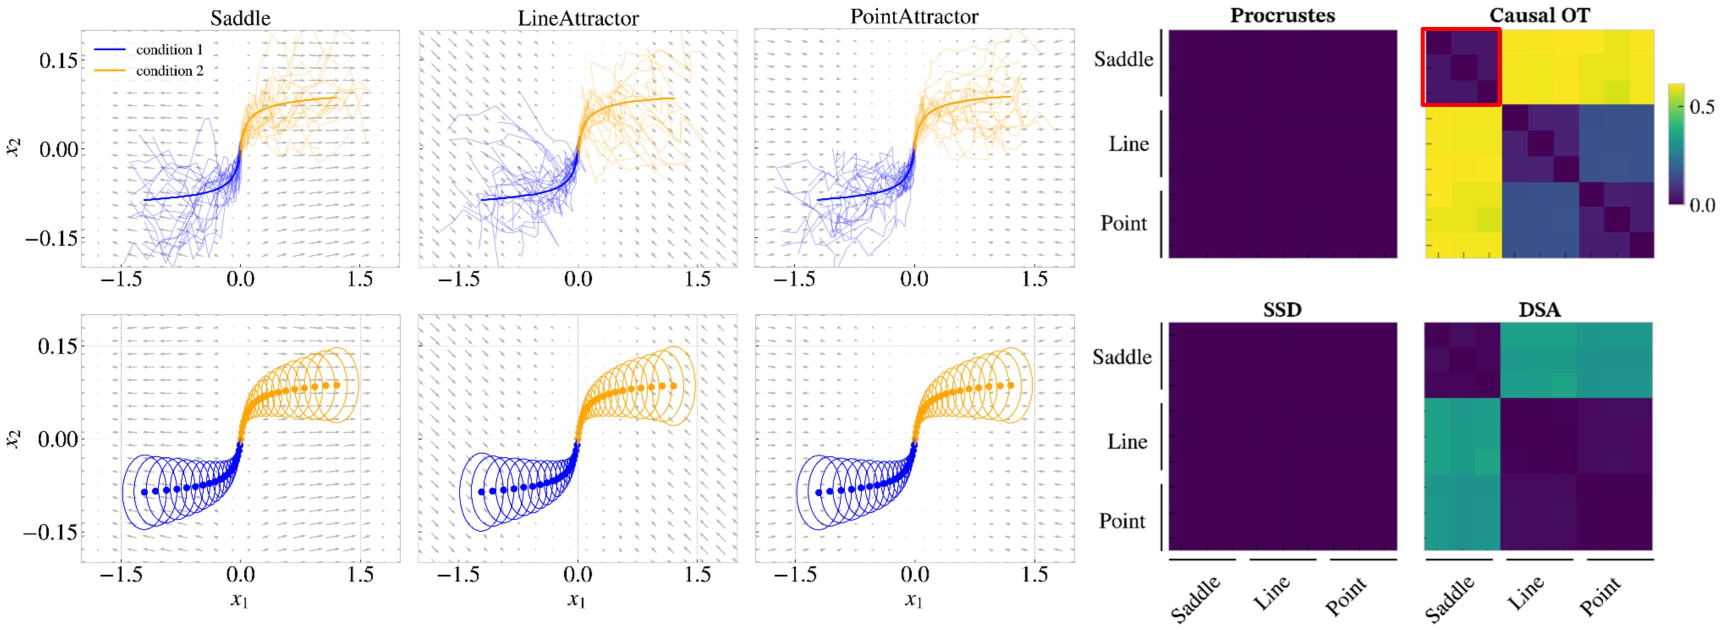

Disambiguating recurrent and input driven dynamics

Dynamics of mean and covariance
$$\begin{cases} \boldsymbol{m}_x(t) = \boldsymbol{A}(t)\boldsymbol{m}_x(t-1) + \boldsymbol{b}(t) \\ \boldsymbol{P}_x(t) = \boldsymbol{A}(t)\boldsymbol{P}_x(t-1)\boldsymbol{A}(t)^\top + \boldsymbol{\Sigma}(t)\boldsymbol{\Sigma}(t)^\top. \end{cases}$$

Adversarially tune inputs and latent noise
$$\begin{cases} \boldsymbol{b}(t) \\ \boldsymbol{\Sigma}(t) \end{cases}$$

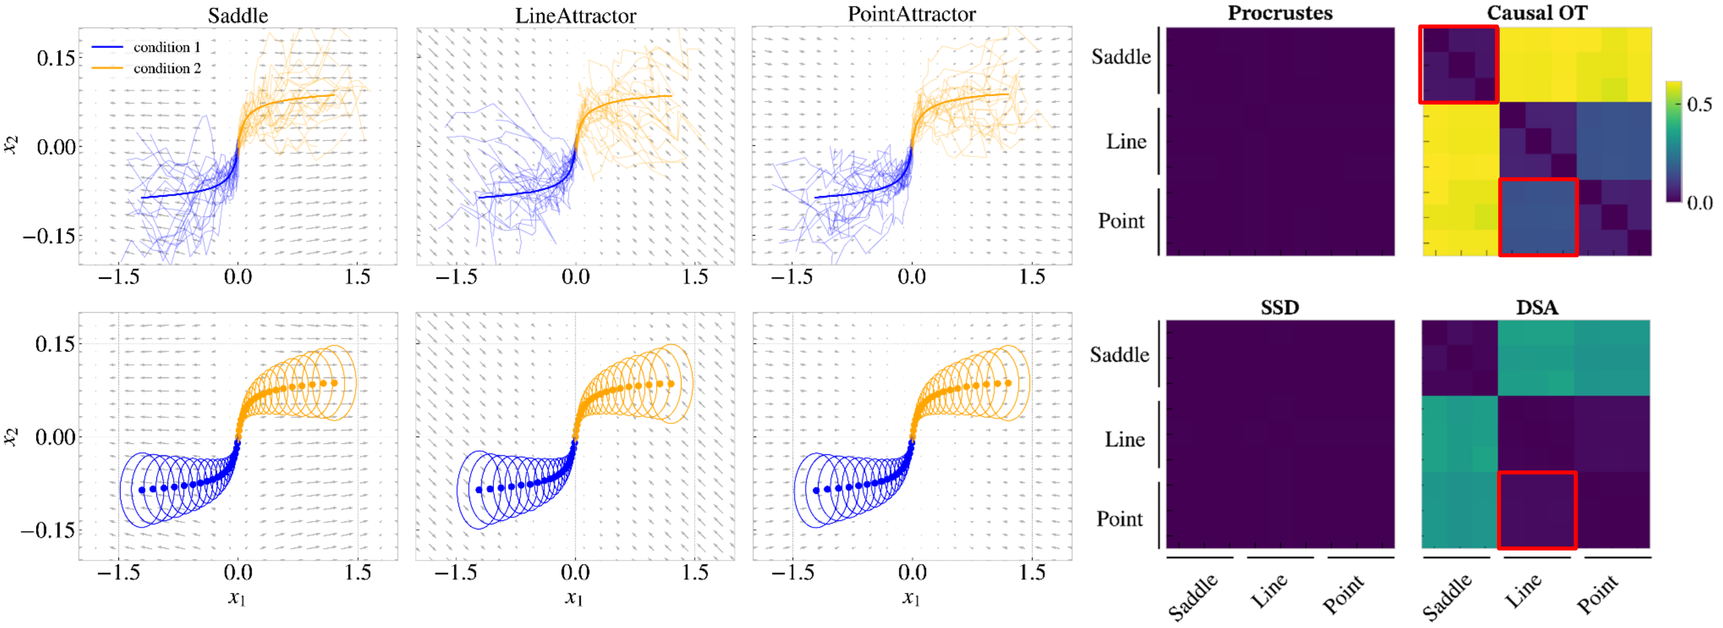

Latent text-to-image diffusion models

Latent text-to-image diffusion models

Trajectories form the stochastic process

$$\mathbf{x}^{\text{model M, prompt A}}$$

Latent text-to-image diffusion models

Trajectories form the stochastic process

$\mathbf{x}^{\text{model M, prompt A}}$

Primate

Sample 

Latent text-to-image diffusion models

Trajectories form the stochastic process

$\mathbf{x}^{\text{model M, prompt A}}$

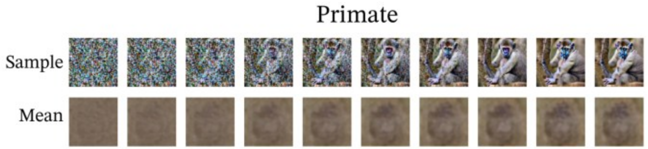

Latent text-to-image diffusion models

Trajectories form the stochastic process

$\mathbf{x}^{\text{model M, prompt A}}$

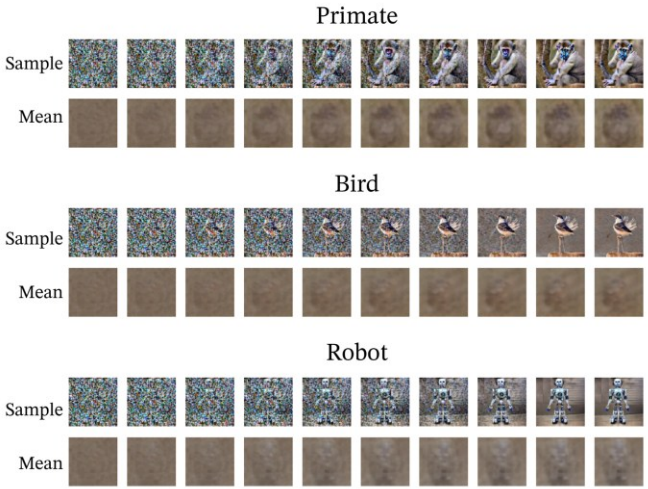

Latent text-to-image diffusion models

Trajectories form the stochastic process

$$\mathbf{x}^{\text{model M, prompt A}}$$

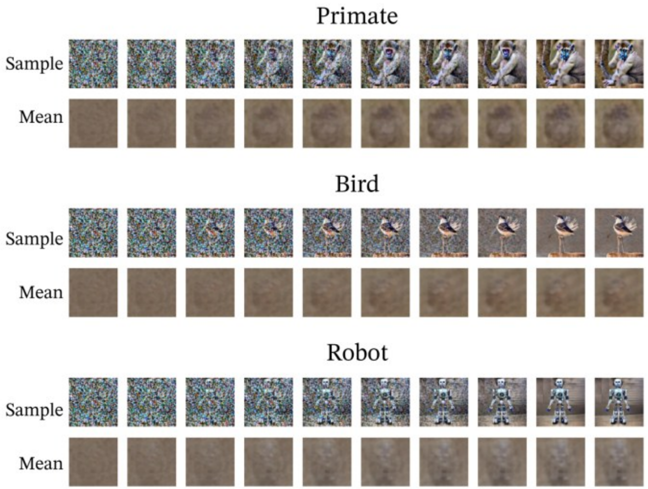

Computing distances between conditionals

$$d(\mathbf{x}^{\text{model M, prompt A}}, \mathbf{x}^{\text{model N, prompt B}})$$

Latent text-to-image diffusion models

Trajectories form the stochastic process

$$\mathbf{x}^{\text{model M, prompt A}}$$

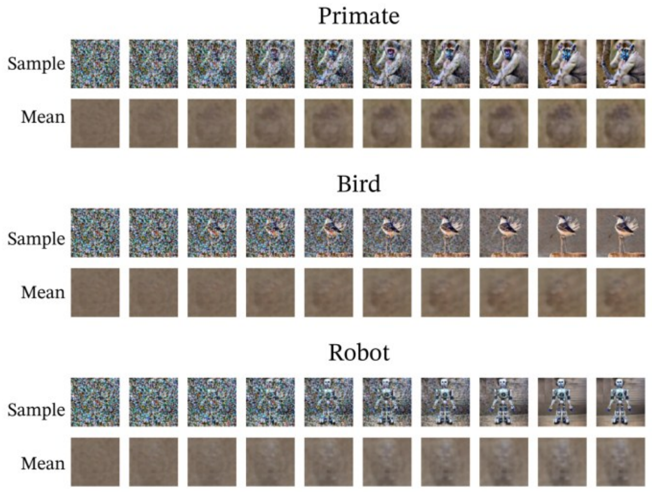

Computing distances between conditionals

$$d(\mathbf{x}^{\text{model M, prompt A}}, \mathbf{x}^{\text{model N, prompt B}})$$

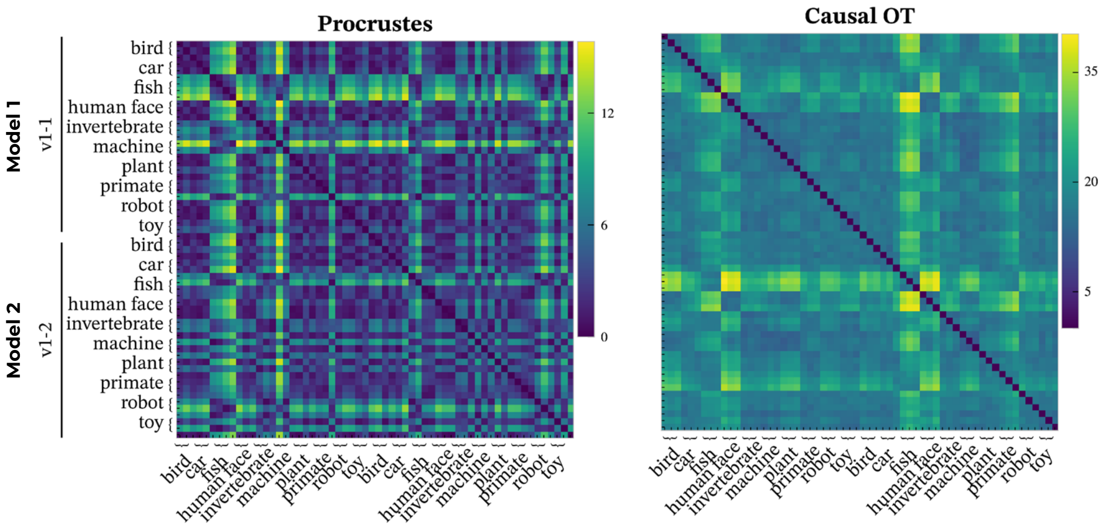

Latent text-to-image diffusion models

Trajectories form the stochastic process

$\mathbf{x}^{\text{model M, prompt A}}$

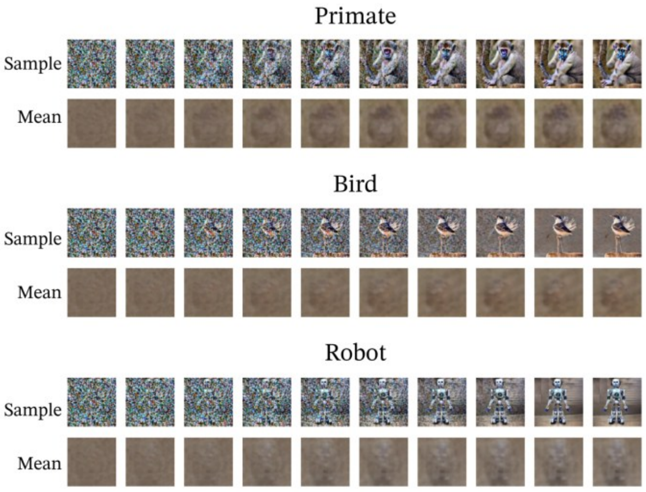

Computing distances between conditionals

$d(\mathbf{x}^{\text{model M, prompt A}}, \mathbf{x}^{\text{model N, prompt B}})$

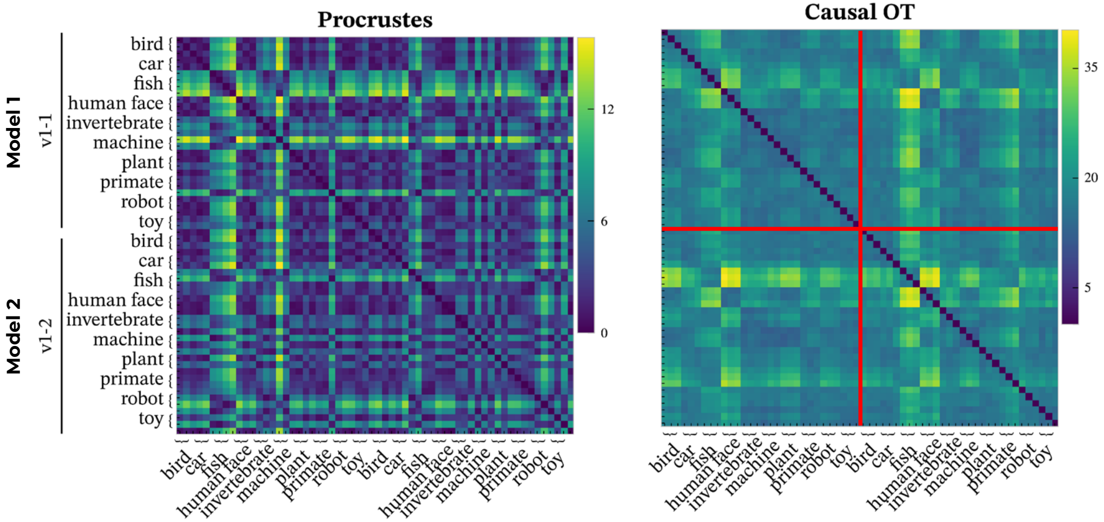

Latent text-to-image diffusion models

Trajectories form the stochastic process

$$\mathbf{x}^{\text{model M, prompt A}}$$

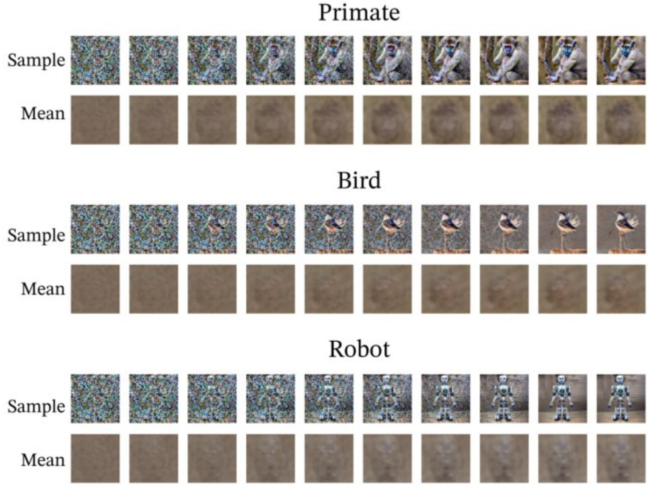

Computing distances between conditionals

$$d(\mathbf{x}^{\text{model M, prompt A}}, \mathbf{x}^{\text{model N, prompt B}})$$

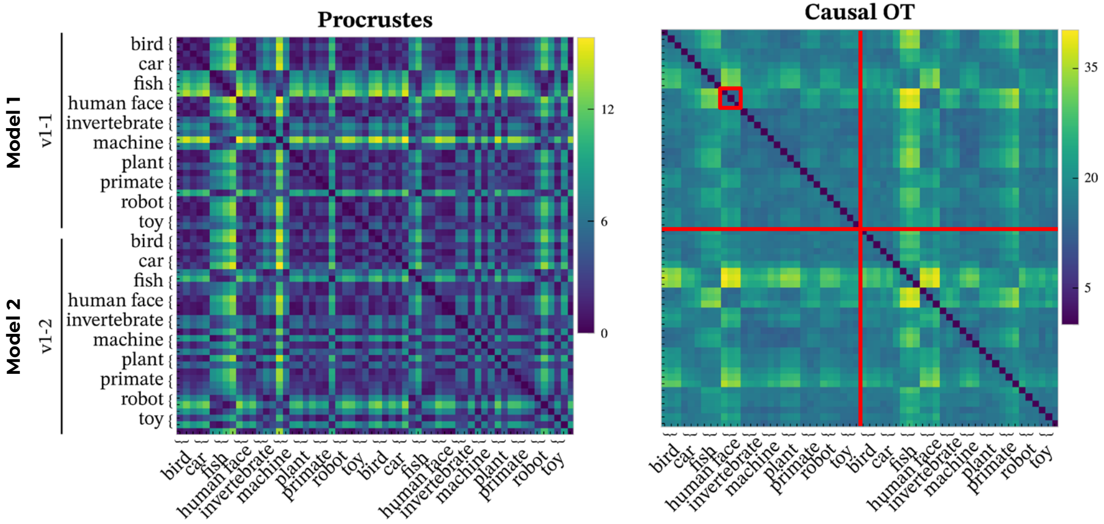

Latent text-to-image diffusion models

Trajectories form the stochastic process

$$\mathbf{x}^{\text{model M, prompt A}}$$

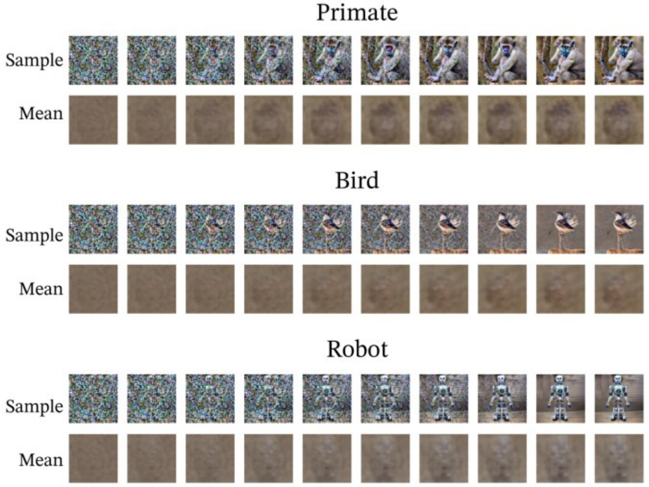

Computing distances between conditionals

$$d(\mathbf{x}^{\text{model M, prompt A}}, \mathbf{x}^{\text{model N, prompt B}})$$

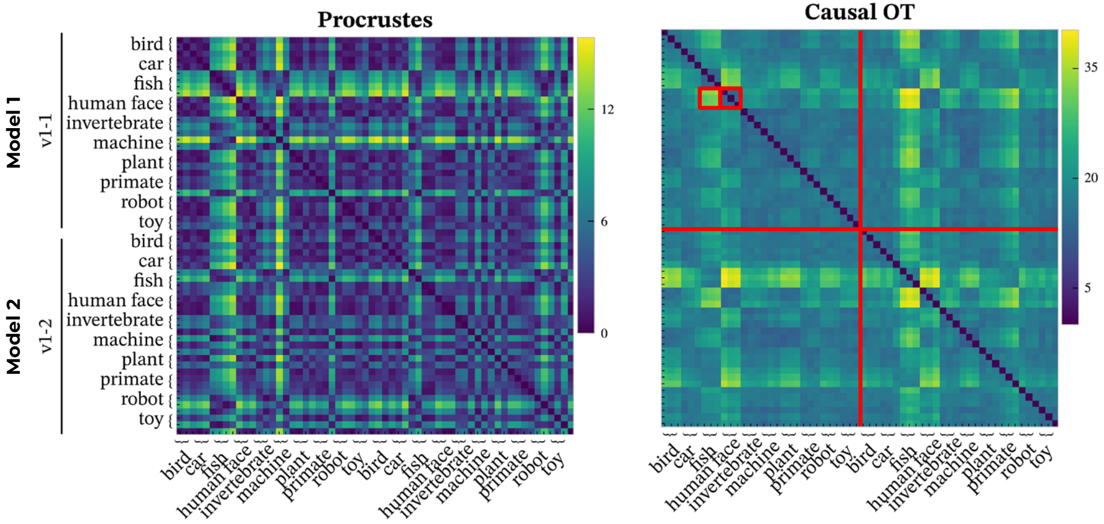

Summary

Summary

- **Representational similarity** enables studying shared computational principles across neural systems.

Summary

- **Representational similarity** enables studying shared computational principles across neural systems.

- Existing methods for comparing population activity either assume **dynamic** or **stochastic** responses; **not both.**

Summary

- **Representational similarity** enables studying shared computational principles across neural systems.
- Existing methods for comparing population activity either assume **dynamic** or **stochastic** responses; **not both.**
- We demonstrate that existing metrics can **fail to capture key differences between neural systems** with **noisy dynamic responses.**

Summary

- **Representational similarity** enables studying shared computational principles across neural systems.
- Existing methods for comparing population activity either assume **dynamic** or **stochastic** responses; **not both.**
- We demonstrate that existing metrics can **fail to capture key differences between neural systems** with **noisy dynamic responses.**
- We then propose a metric for comparing the geometry of noisy neural trajectories, which can be derived as an **optimal transport distance** between **Gaussian processes.**

Summary

- **Representational similarity** enables studying shared computational principles across neural systems.
- Existing methods for comparing population activity either assume **dynamic** or **stochastic** responses; **not both.**
- We demonstrate that existing metrics can **fail to capture key differences between neural systems** with **noisy dynamic responses.**
- We then propose a metric for comparing the geometry of noisy neural trajectories, which can be derived as an **optimal transport distance** between **Gaussian processes.**
- We use the metric to compare models of neural responses in different regions of the motor system and to compare the dynamics of latent diffusion models for text-to-image synthesis.

Summary

- **Representational similarity** enables studying shared computational principles across neural systems.

- Existing methods for comparing population activity either assume **dynamic** or **stochastic** responses; **not both.**

- We demonstrate that existing metrics can **fail to capture key differences between neural systems** with **noisy dynamic responses.**

- We then propose a metric for comparing the geometry of noisy neural trajectories, which can be derived as an **optimal transport distance** between **Gaussian processes.**
- We use the metric to compare models of neural responses in different regions of the motor system and to compare the dynamics of latent diffusion models for text-to-image synthesis.

Future direction: Computing Causal-OT requires many trials. The latent variable model in **[1]** can reduce the computational cost and data requirements.

References and Acknowledgments

References and Acknowledgments

arXiv

GitHub

References and Acknowledgments

arXiv

GitHub

Poster #64
(Hall 3)

References and Acknowledgments

References

[1] Geadah, Victor, et al. "Modeling Neural Activity with Conditionally Linear Dynamical Systems." arXiv preprint arXiv:2502.18347 (2025).

[2] Lipshutz, David, et al. "Disentangling recurrent neural dynamics with stochastic representational geometry." ICLR 2024 Workshop on Representational Alignment. 2024.

[3] Barbosa, Joao, et al. "Quantifying Differences in Neural Population Activity With Shape Metrics." bioRxiv (2025): 2025-01.

[4] Ostrow, Mitchell, et al. "Beyond geometry: Comparing the temporal structure of computation in neural circuits with dynamical similarity analysis." NeurIPS (2023): 33824-33837.

Poster #64
(Hall 3)

References and Acknowledgments

References

[1] Geadah, Victor, et al. "Modeling Neural Activity with Conditionally Linear Dynamical Systems." arXiv preprint arXiv:2502.18347 (2025).

[2] Lipshutz, David, et al. "Disentangling recurrent neural dynamics with stochastic representational geometry." ICLR 2024 Workshop on Representational Alignment. 2024.

[3] Barbosa, Joao, et al. "Quantifying Differences in Neural Population Activity With Shape Metrics." bioRxiv (2025): 2025-01.

[4] Ostrow, Mitchell, et al. "Beyond geometry: Comparing the temporal structure of computation in neural circuits with dynamical similarity analysis." NeurIPS (2023): 33824-33837.



Poster #64 (Hall 3)

References and Acknowledgments

References

[1] Geadah, Victor, et al. "Modeling Neural Activity with Conditionally Linear Dynamical Systems." arXiv preprint arXiv:2502.18347 (2025).

[2] Lipshutz, David, et al. "Disentangling recurrent neural dynamics with stochastic representational geometry." ICLR 2024 Workshop on Representational Alignment. 2024.

[3] Barbosa, Joao, et al. "Quantifying Differences in Neural Population Activity With Shape Metrics." bioRxiv (2025): 2025-01.

[4] Ostrow, Mitchell, et al. "Beyond geometry: Comparing the temporal structure of computation in neural circuits with dynamical similarity analysis." NeurIPS (2023): 33824-33837.

arXiv

GitHub

SIMONS FOUNDATION

Poster #64 (Hall 3)