Understanding and learning generative models for physical systems

Nisha Chandramoorthy

The University of Chicago

Thanks: Adriaan de Clercq, Sebastian Deery (U Chicago)

Manifold hypothesis
Pope et al 2021
Benjio Courville
Vincent 2013

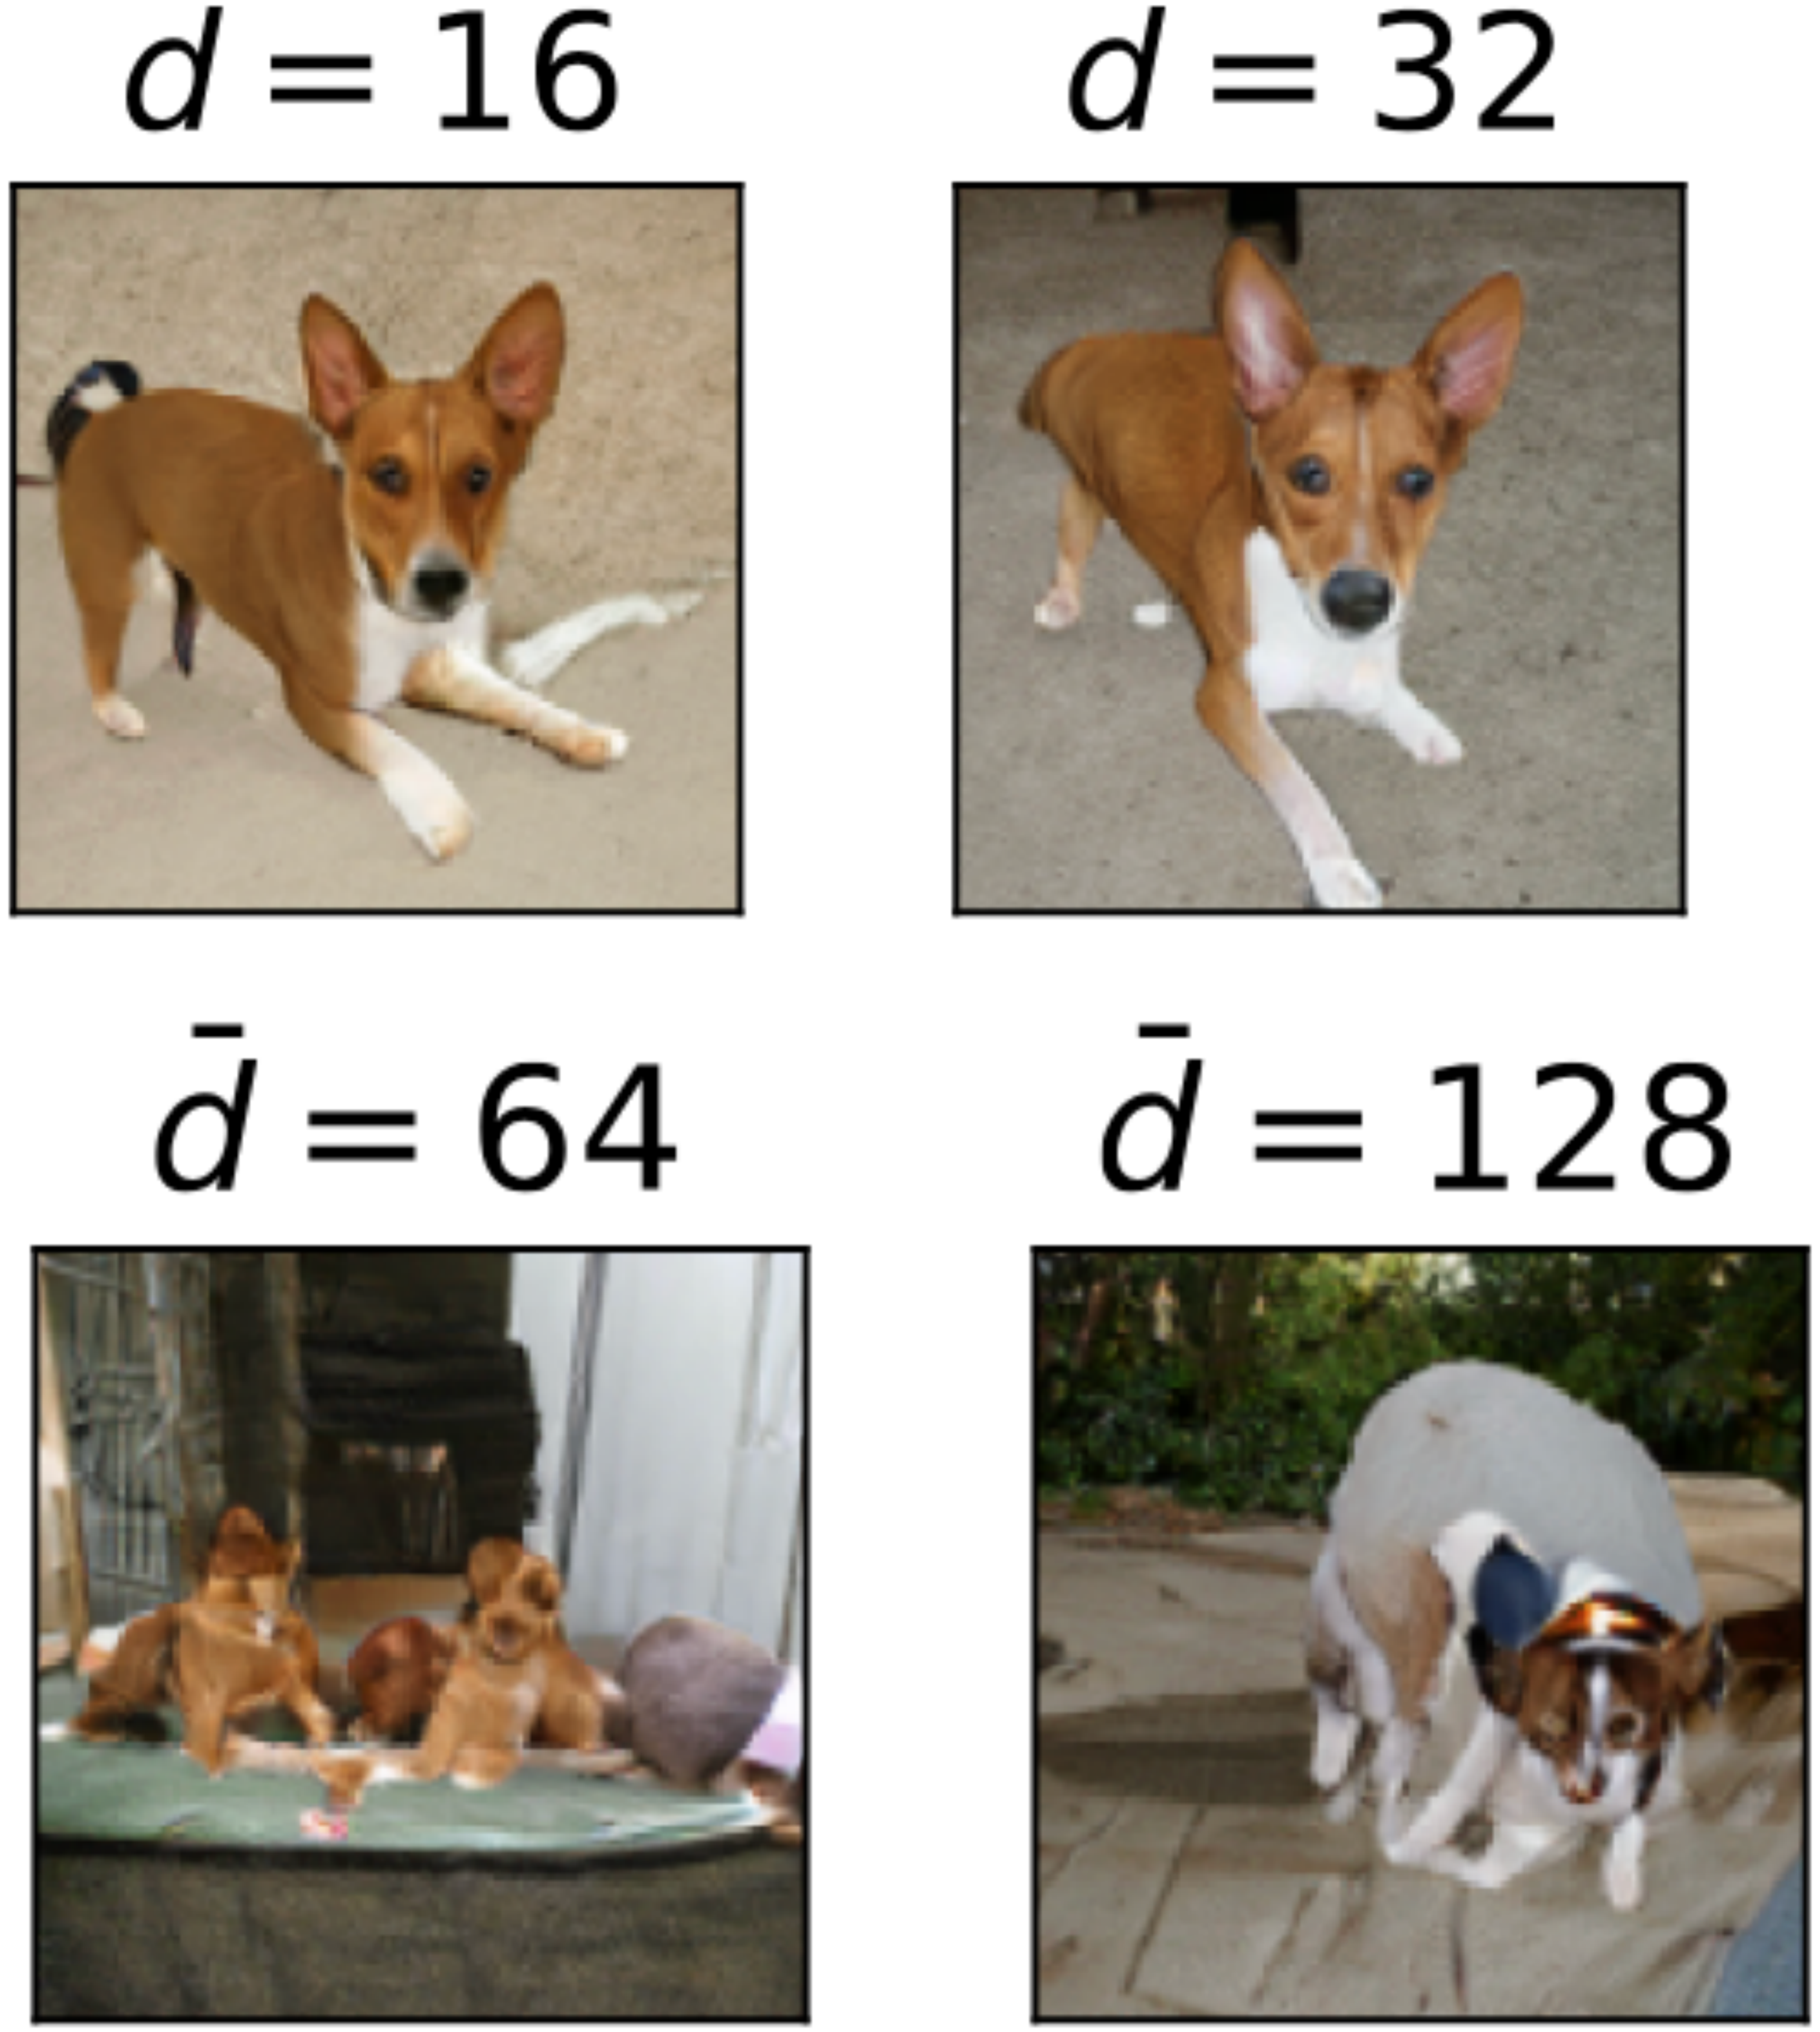

Low-dimensional absolutely continuous structure — in the physical sciences

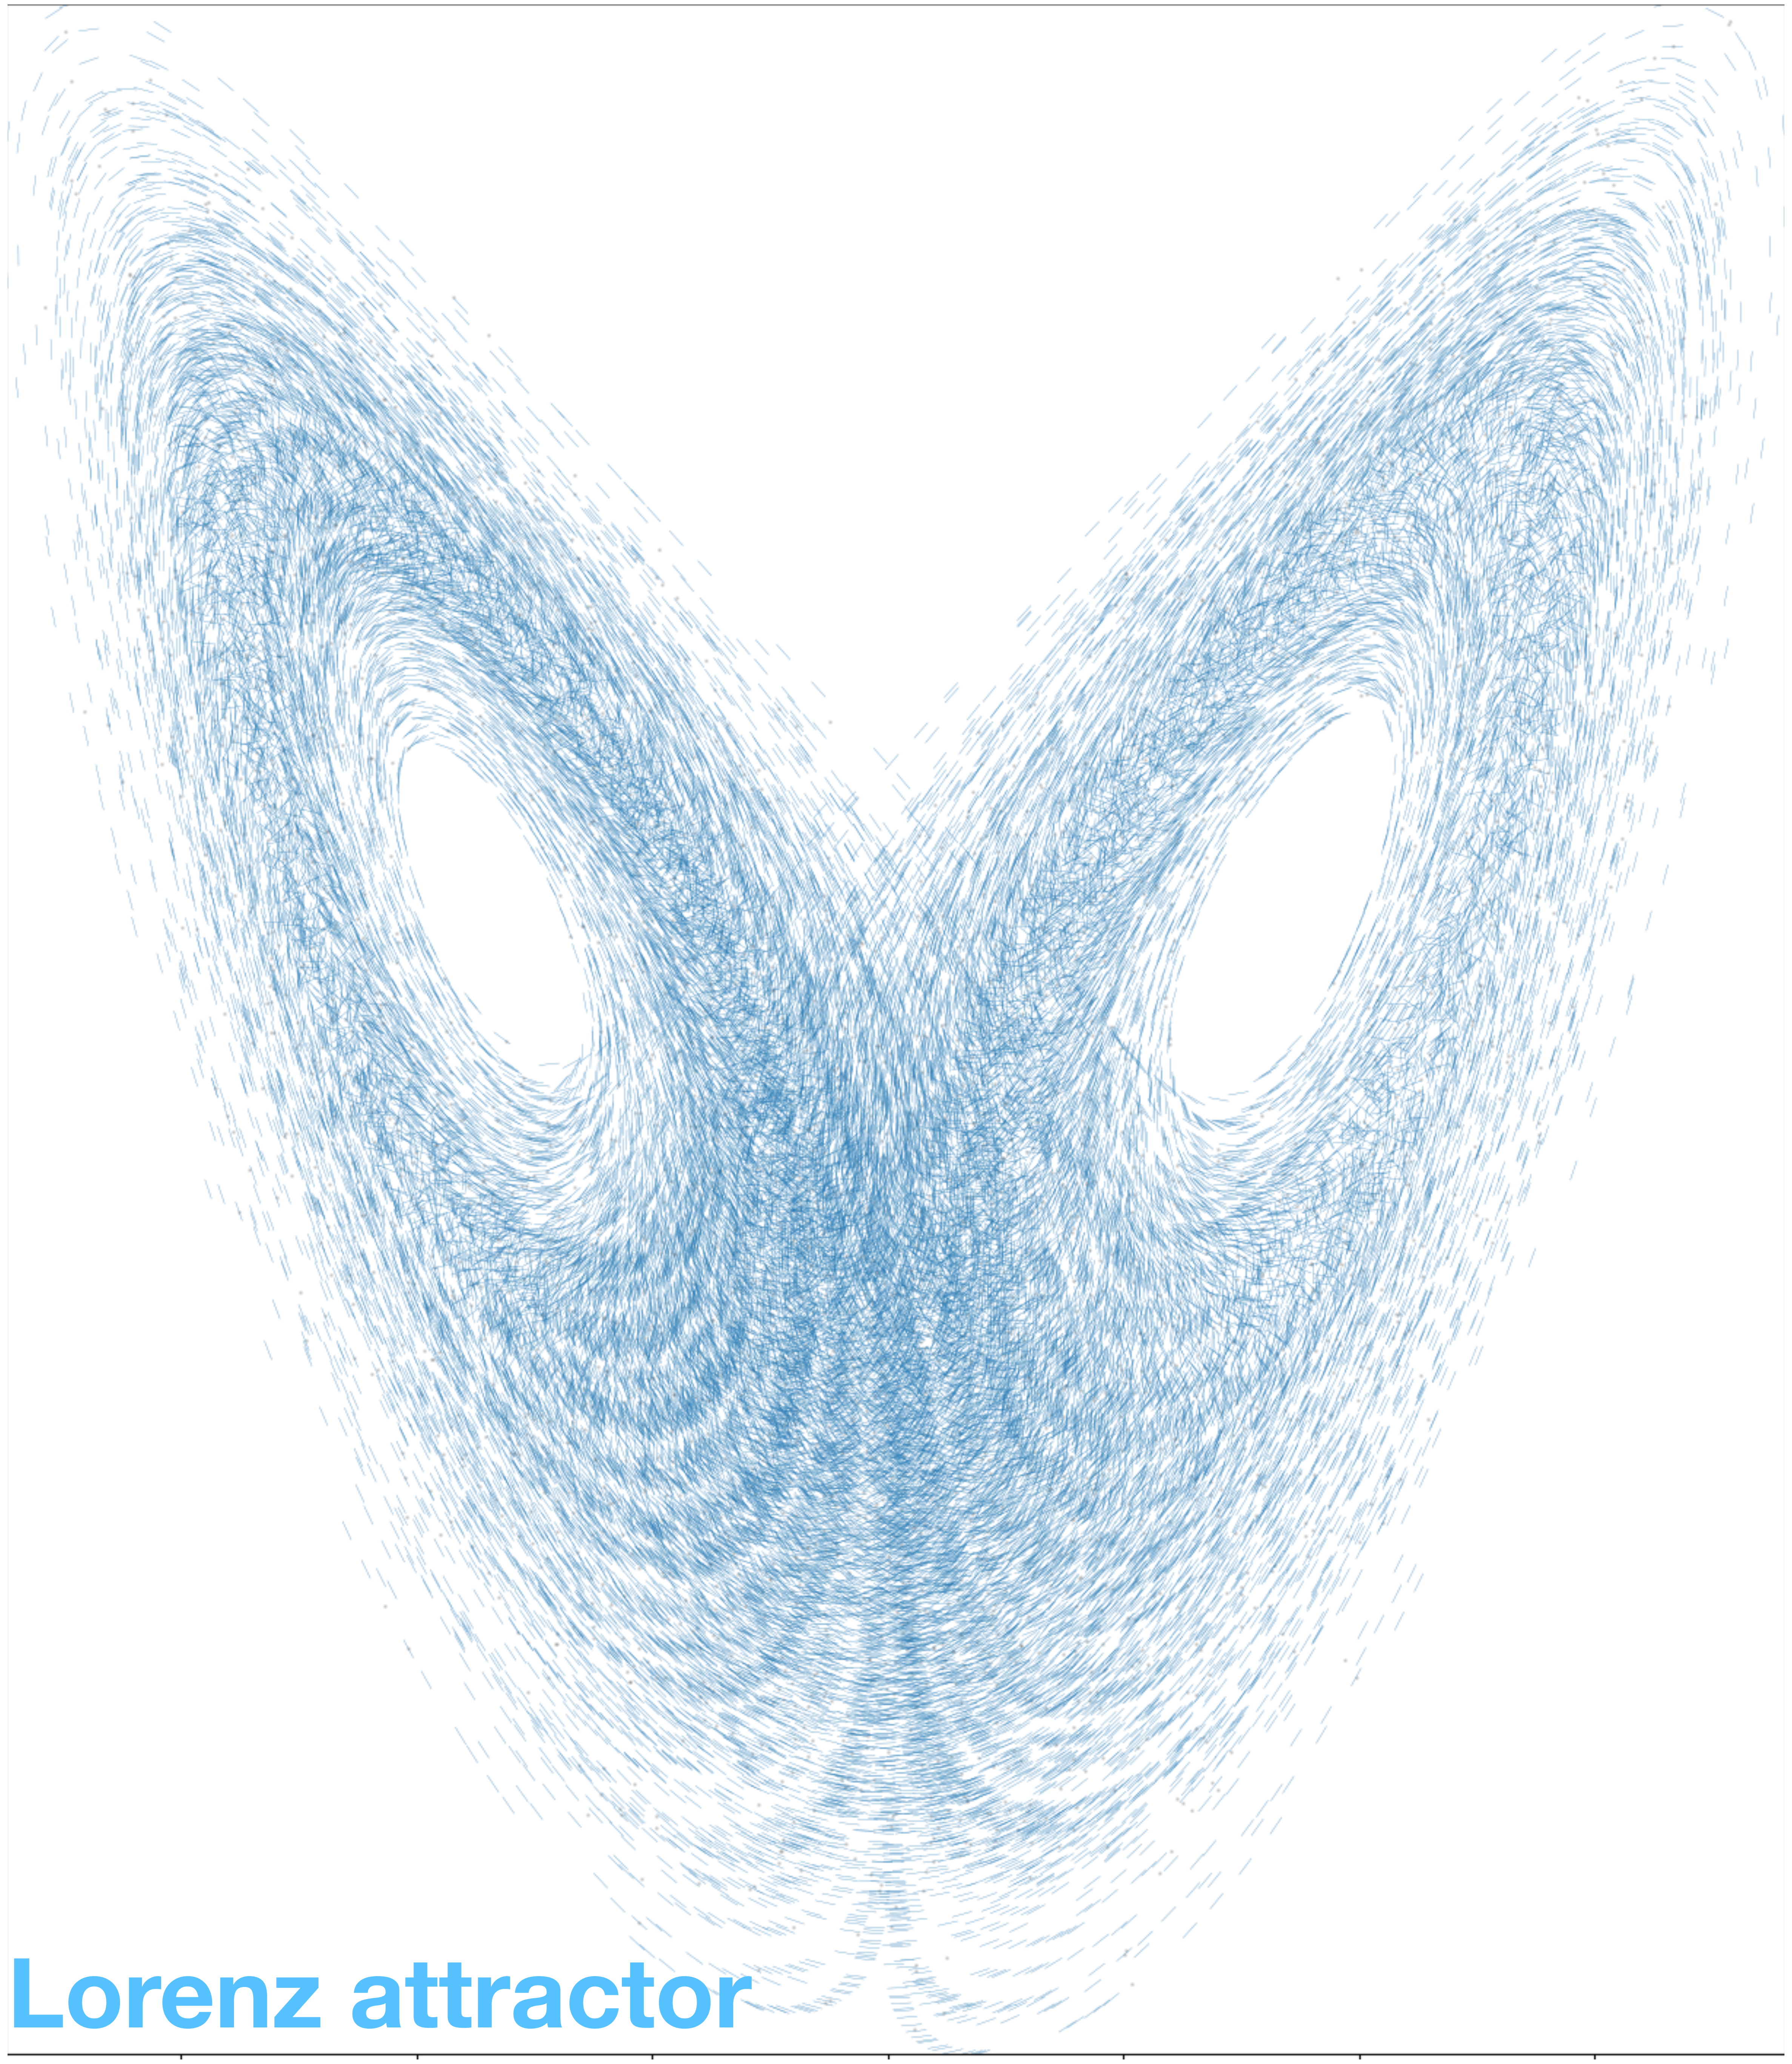

The unstable subspaces on the Lorenz attractor

Inherited low-dimensional structure in Bayesian computation

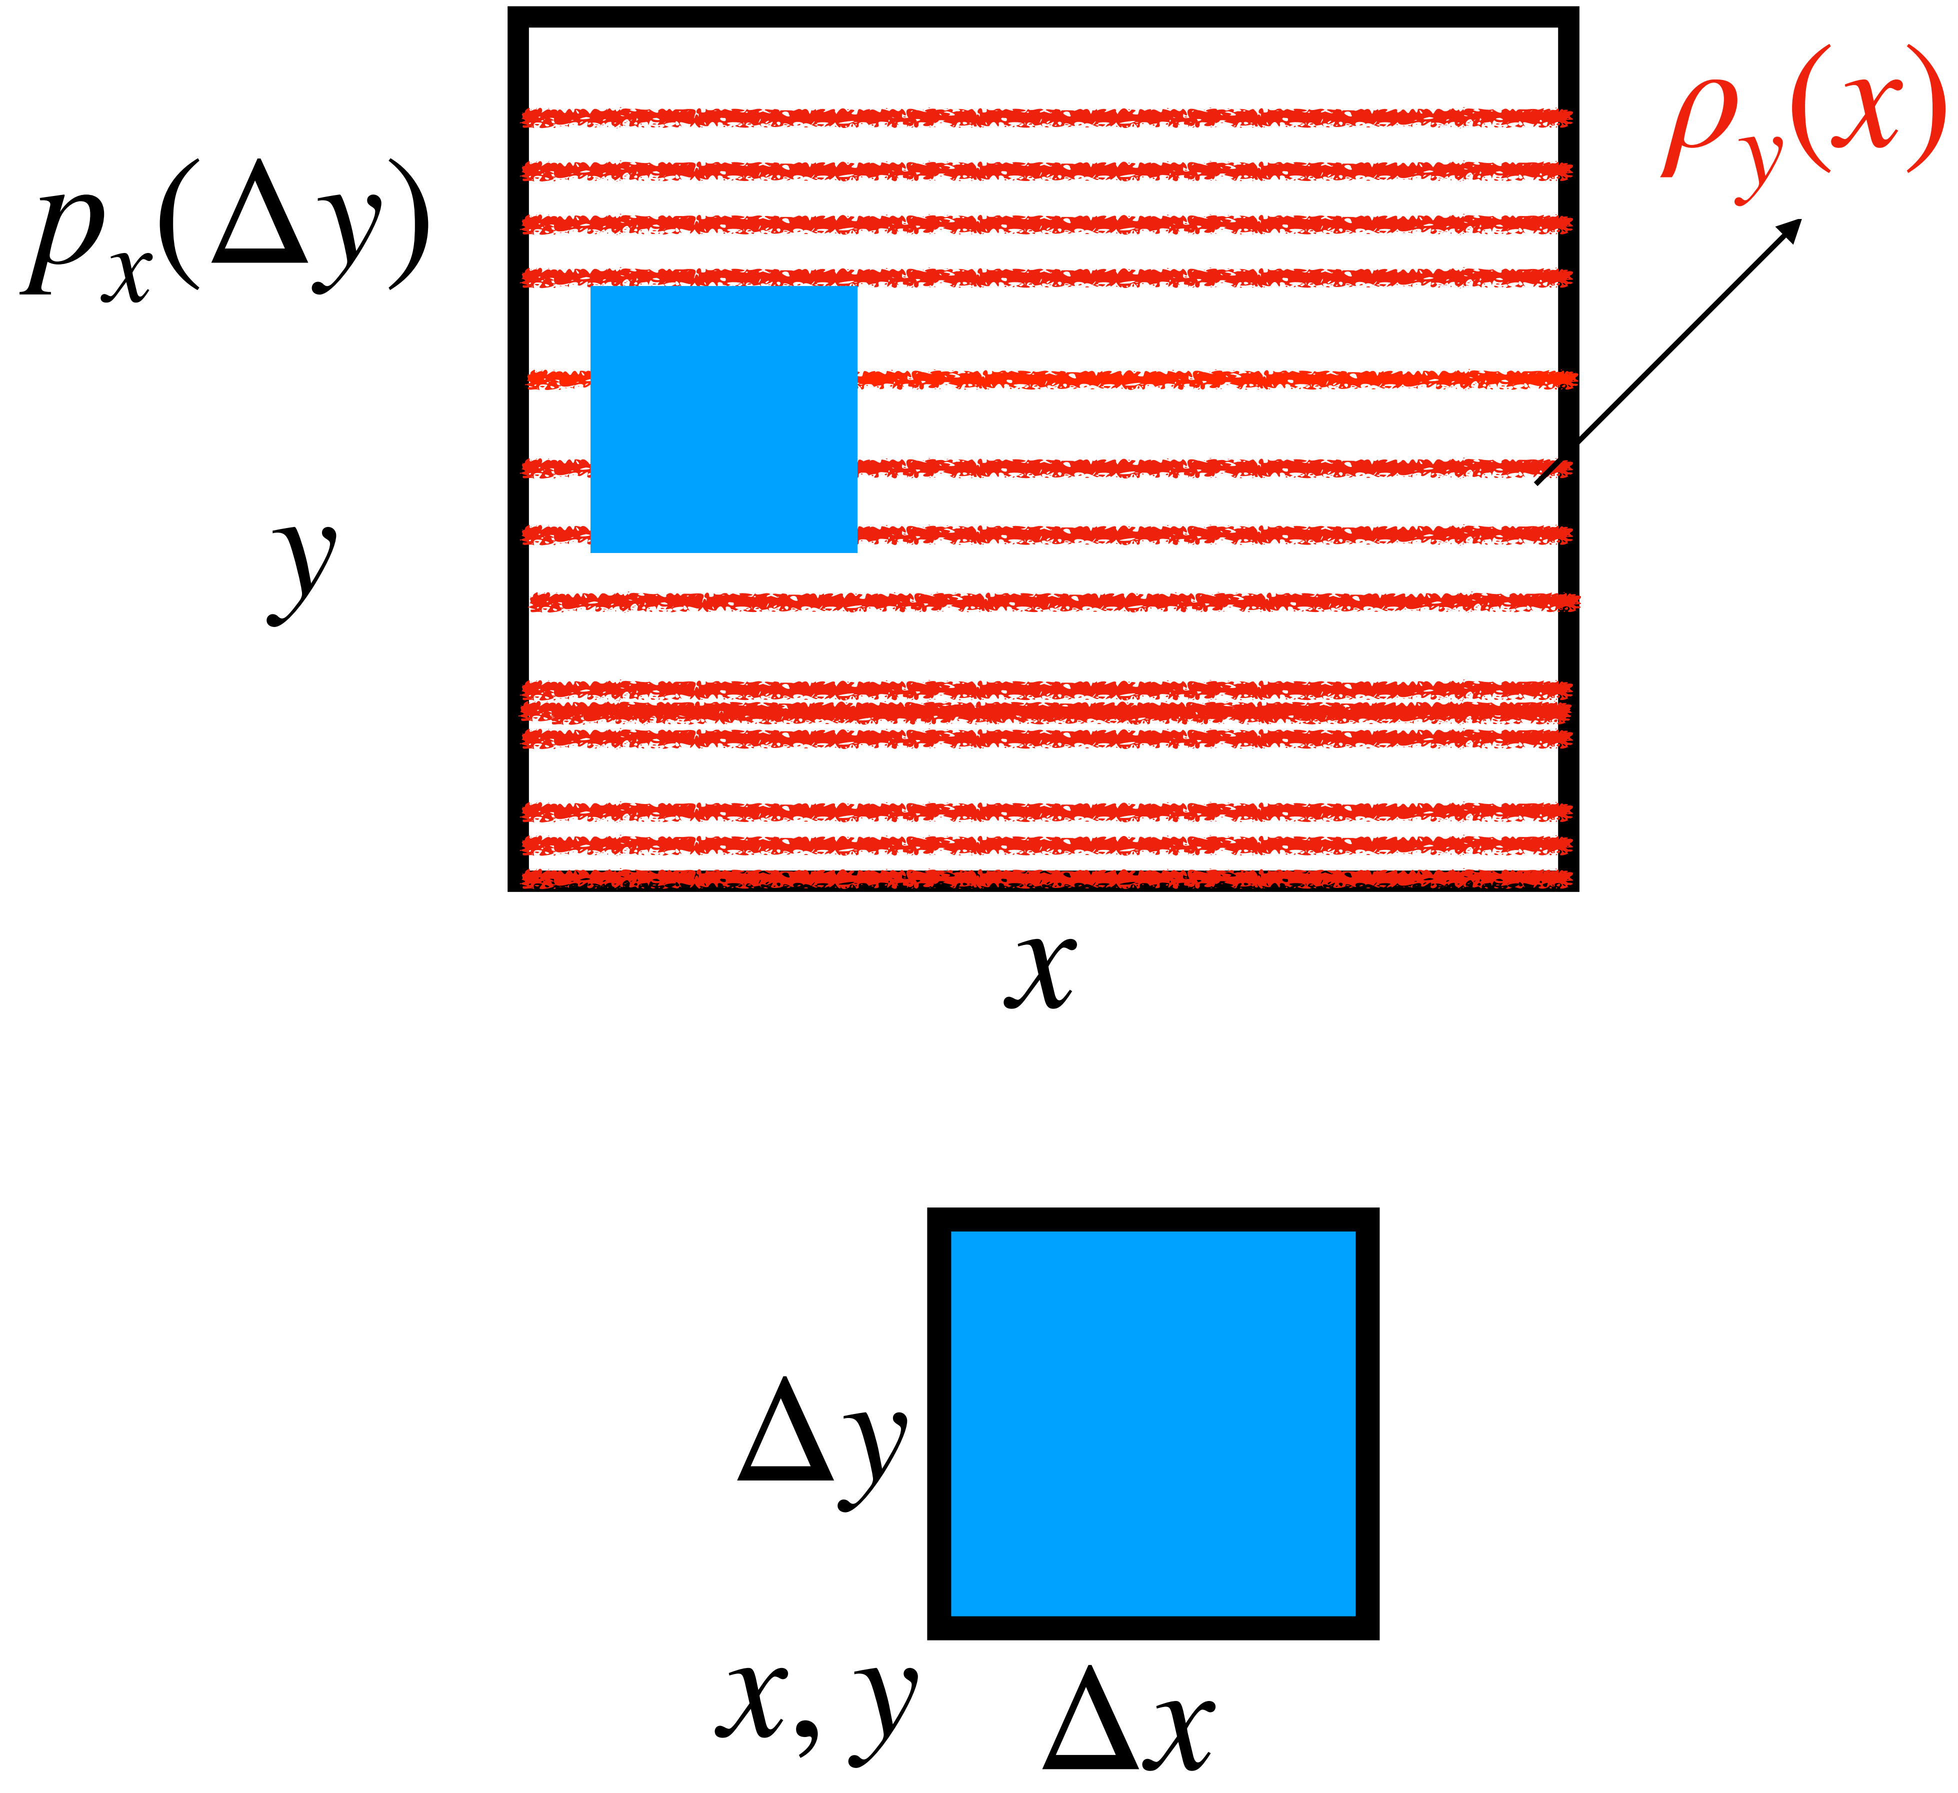

- Bayesian inference/data assimilation, denoising/conditional generation, flow matching — all involve some dynamics/transport
- Only want to learn dynamics on x
- Dimension reduction?

Data manifold (sometimes) unknown

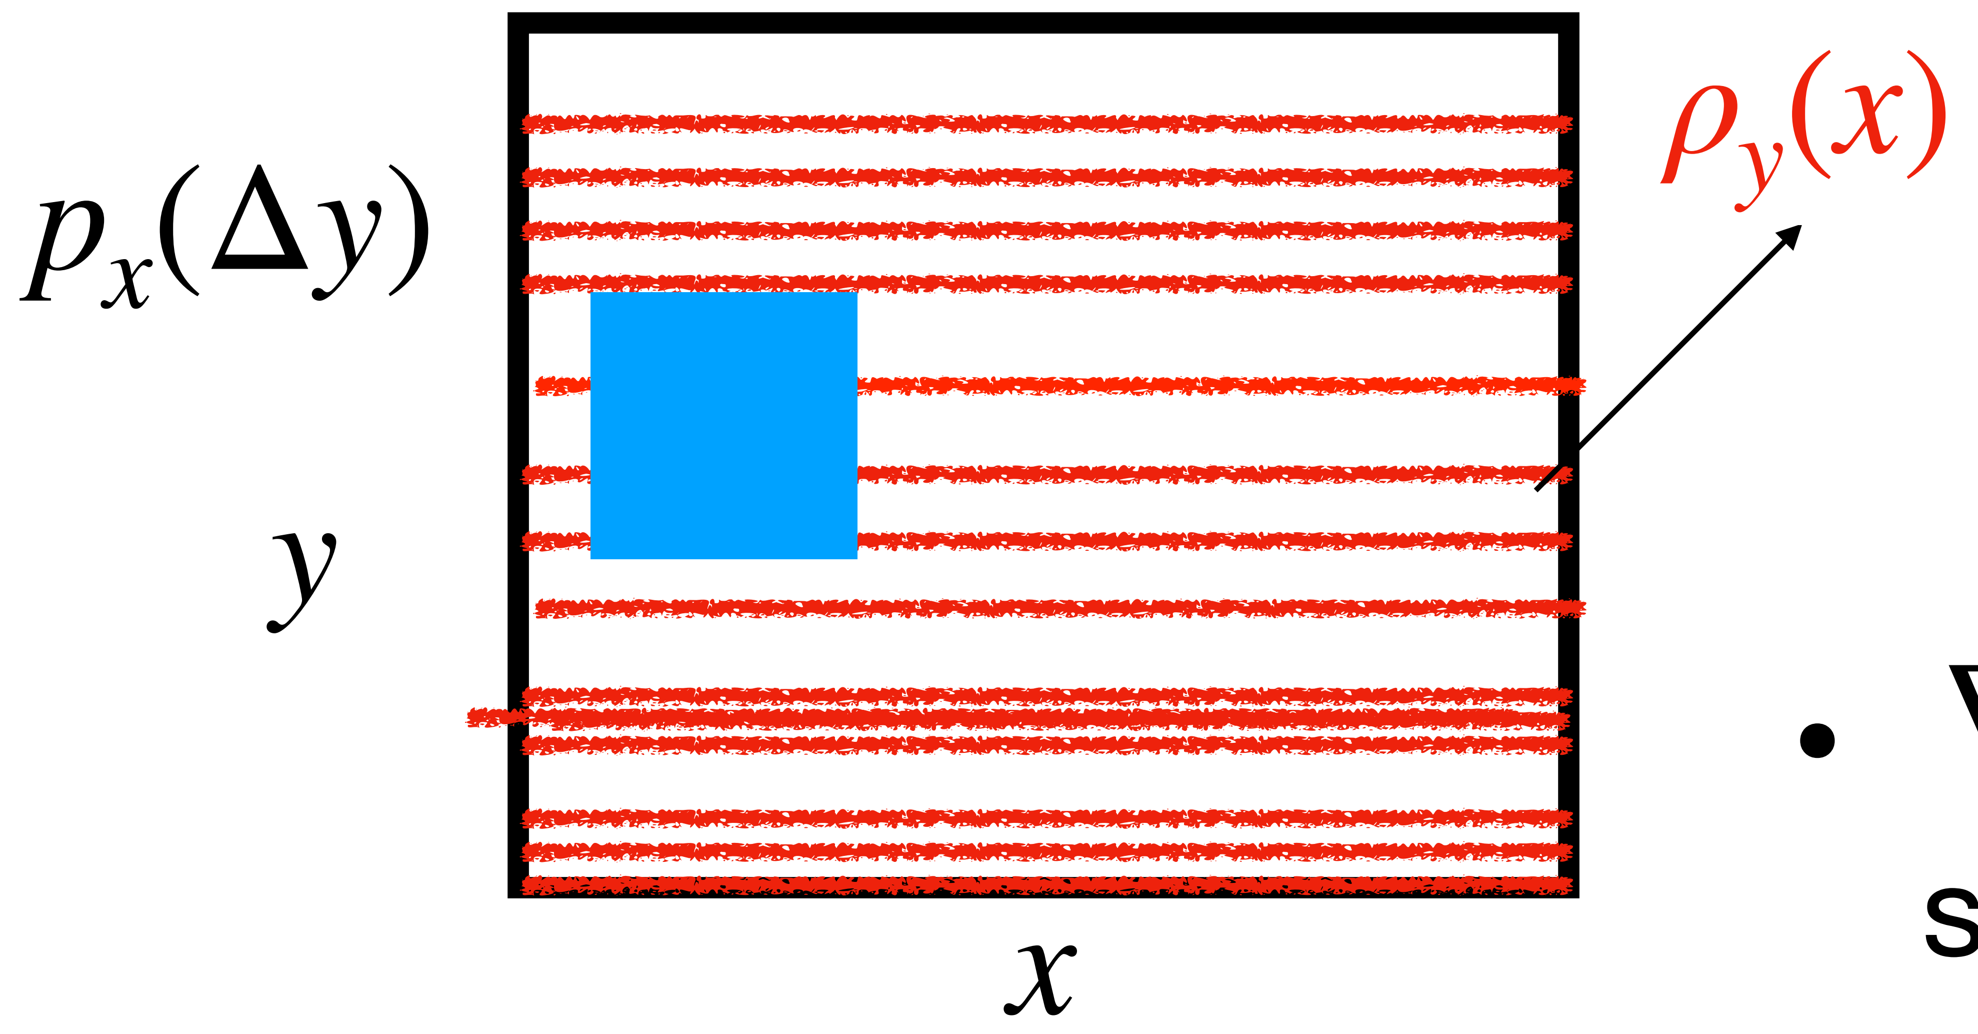

$$dp_1 = \frac{\rho_y \circ F^{-1} dx\, dp_x \circ F^{-1}}{(\|\nabla F e_x\| \|\nabla F e_y\|) \circ F^{-1}}$$

- $\nabla_x \log \rho_y$: may be known/easy to compute in some chaotic systems
- Factorization persists under transformation

Goal-oriented generative models

- Can GMs be used to learn the data manifold?
- Inevitable errors in GMs: how do they affect the predicted distribution?
- Do some errors not matter for conditional sampling?
- Generative models for priors used for posterior sampling?
- Errors in GMs for priors: can “important” directions still be preserved?

Dynamical systems approach to generative modeling

The support of diffusion models is robust to learning errors

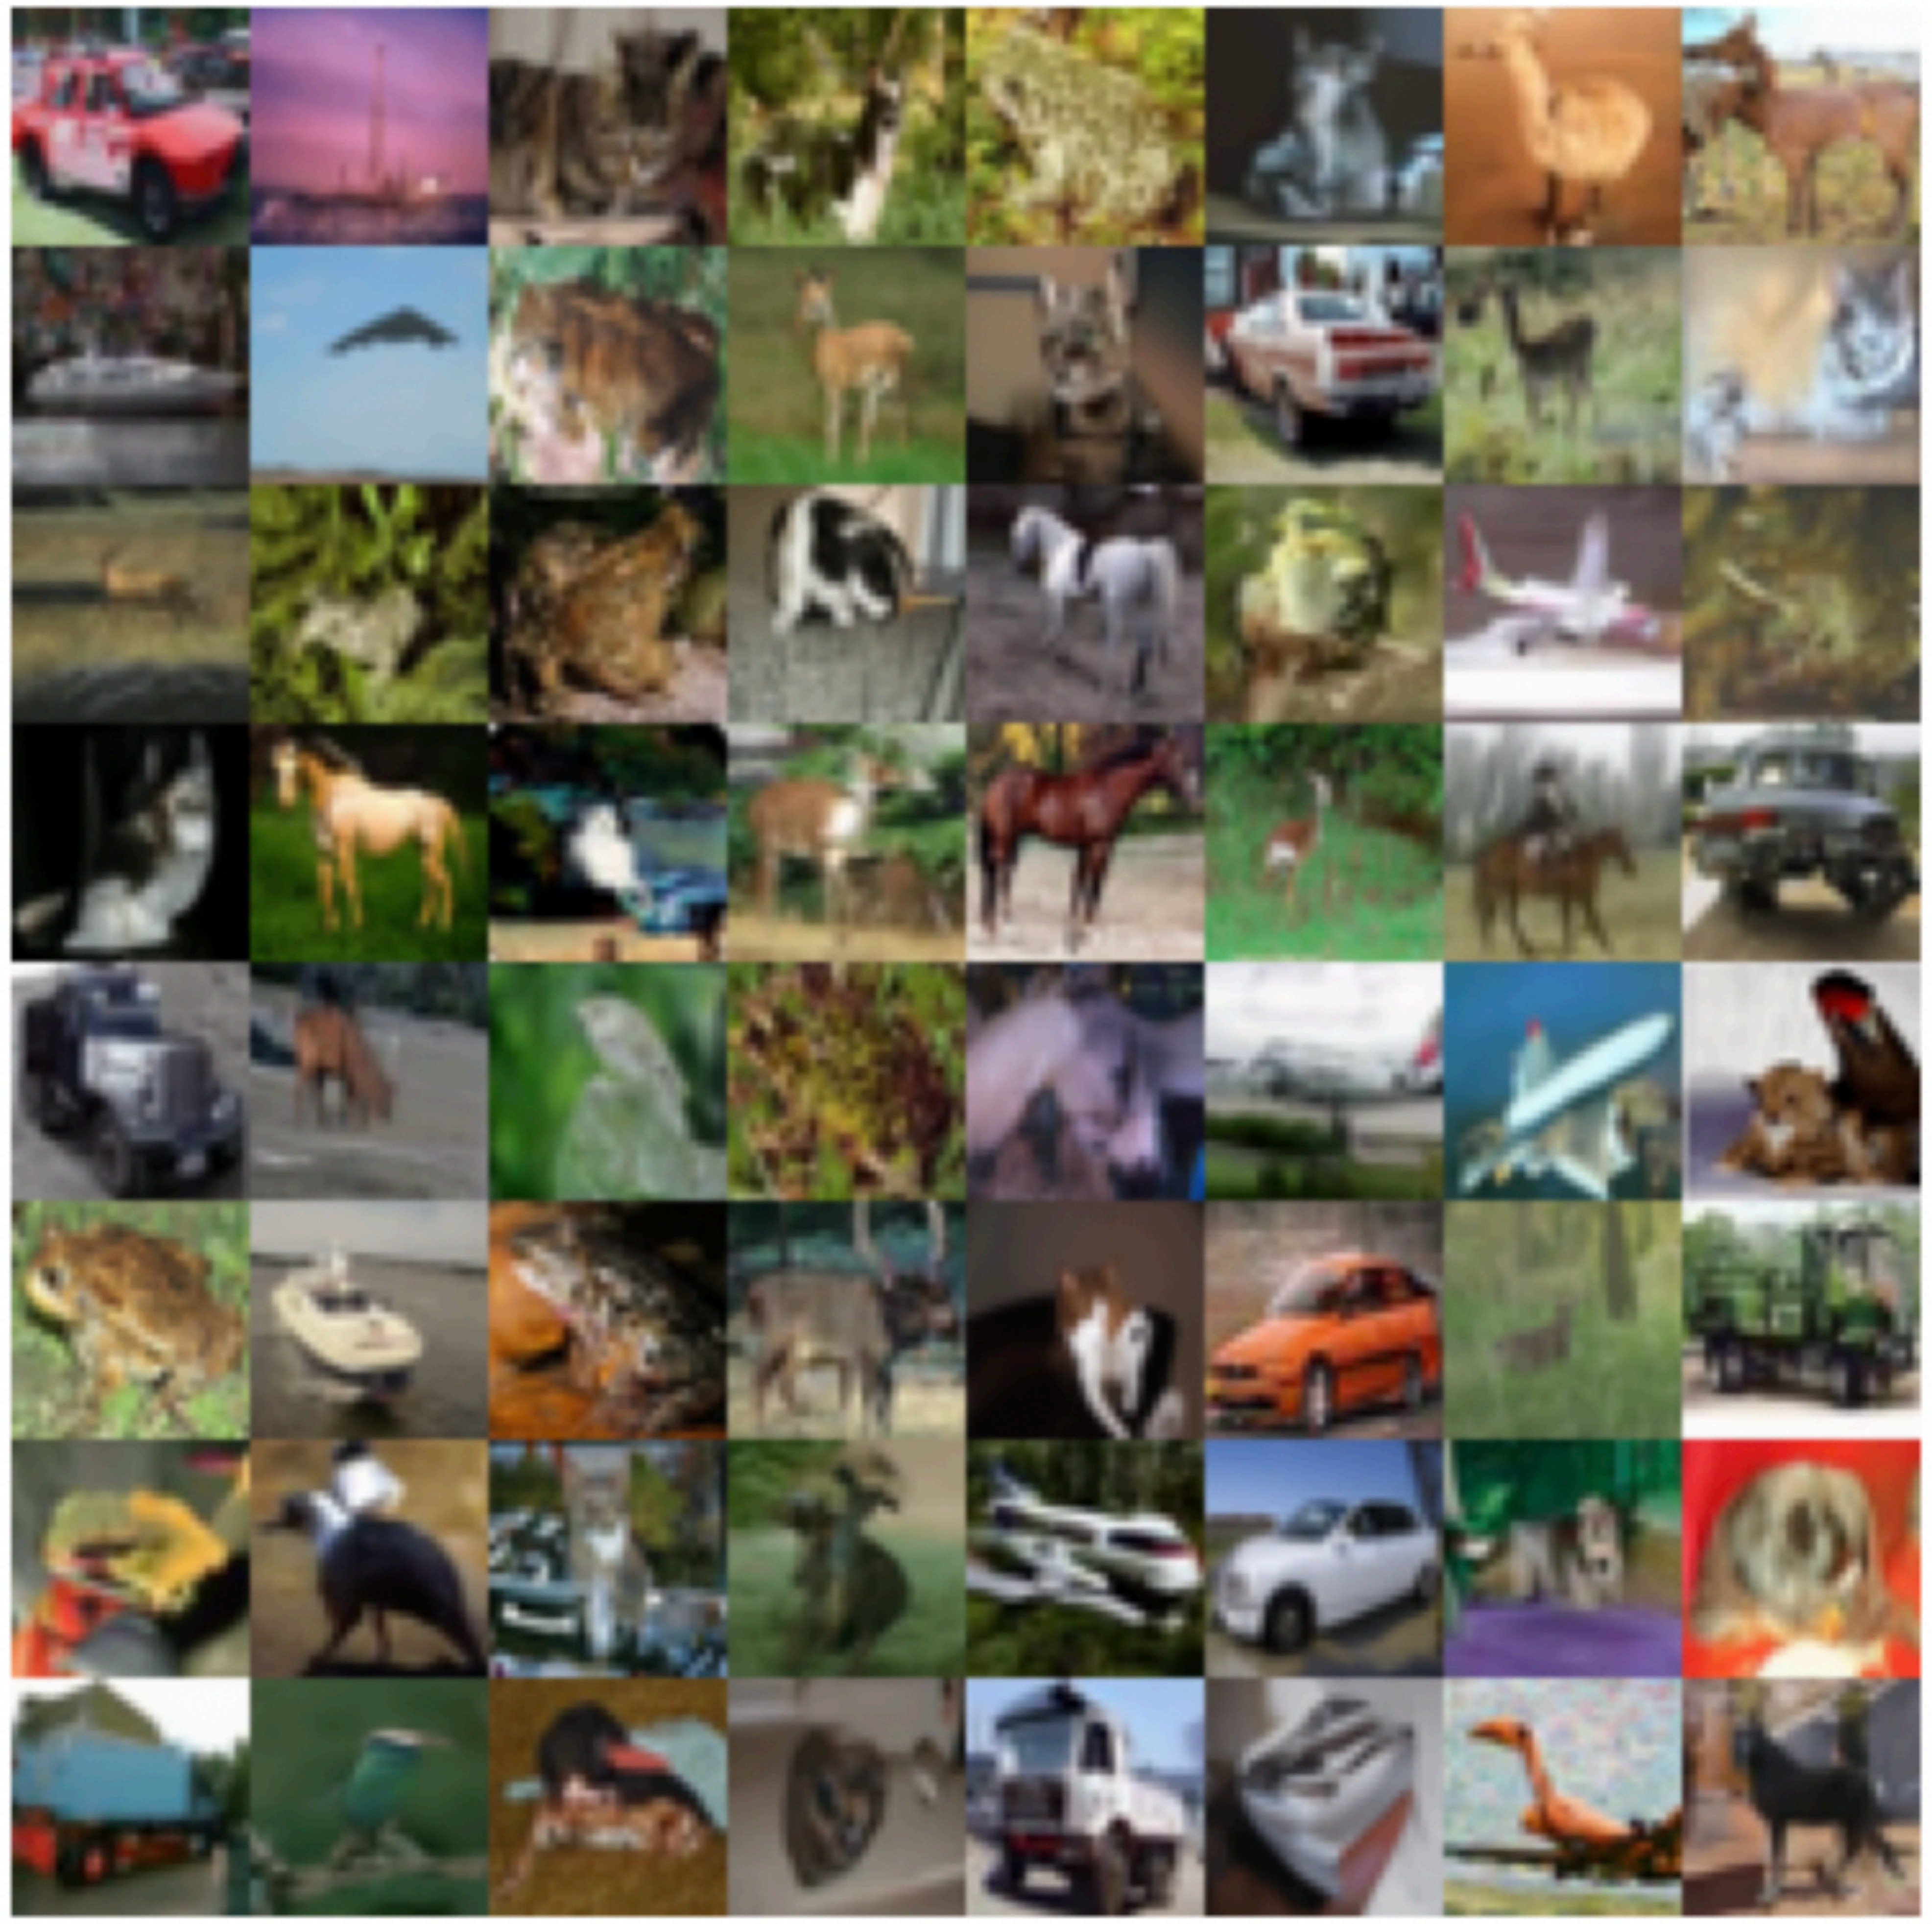

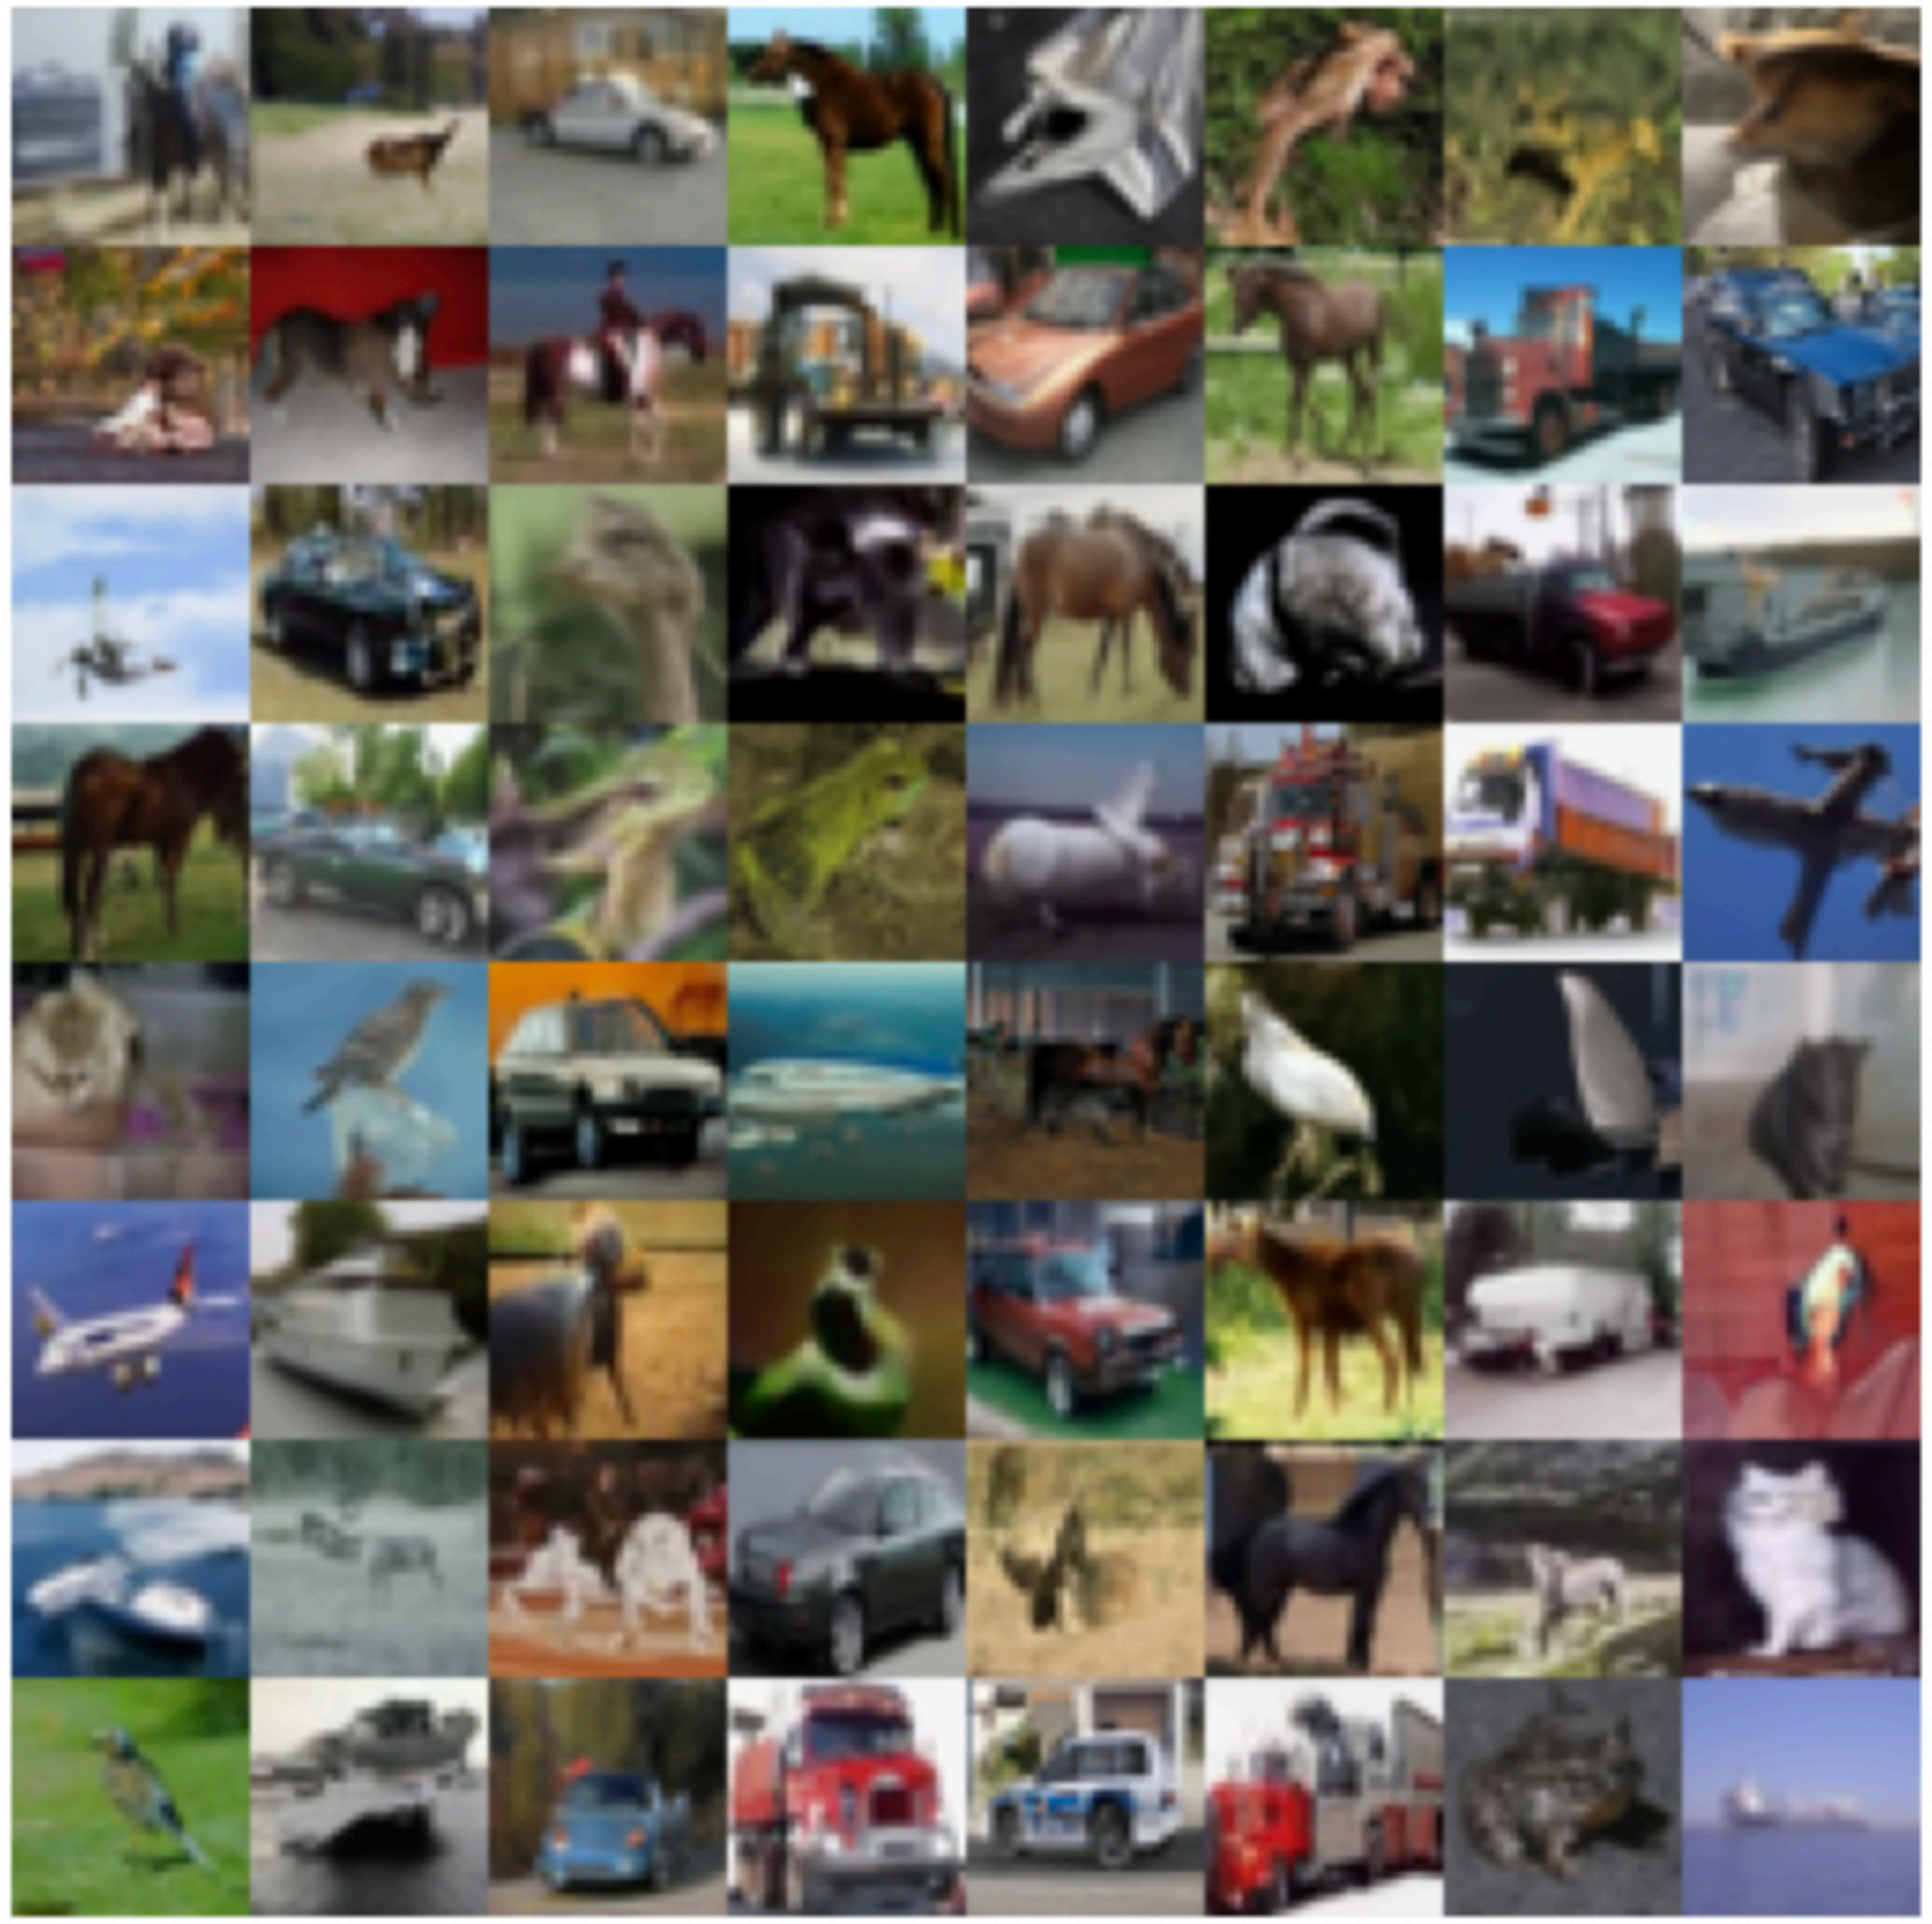

Robustness of the support

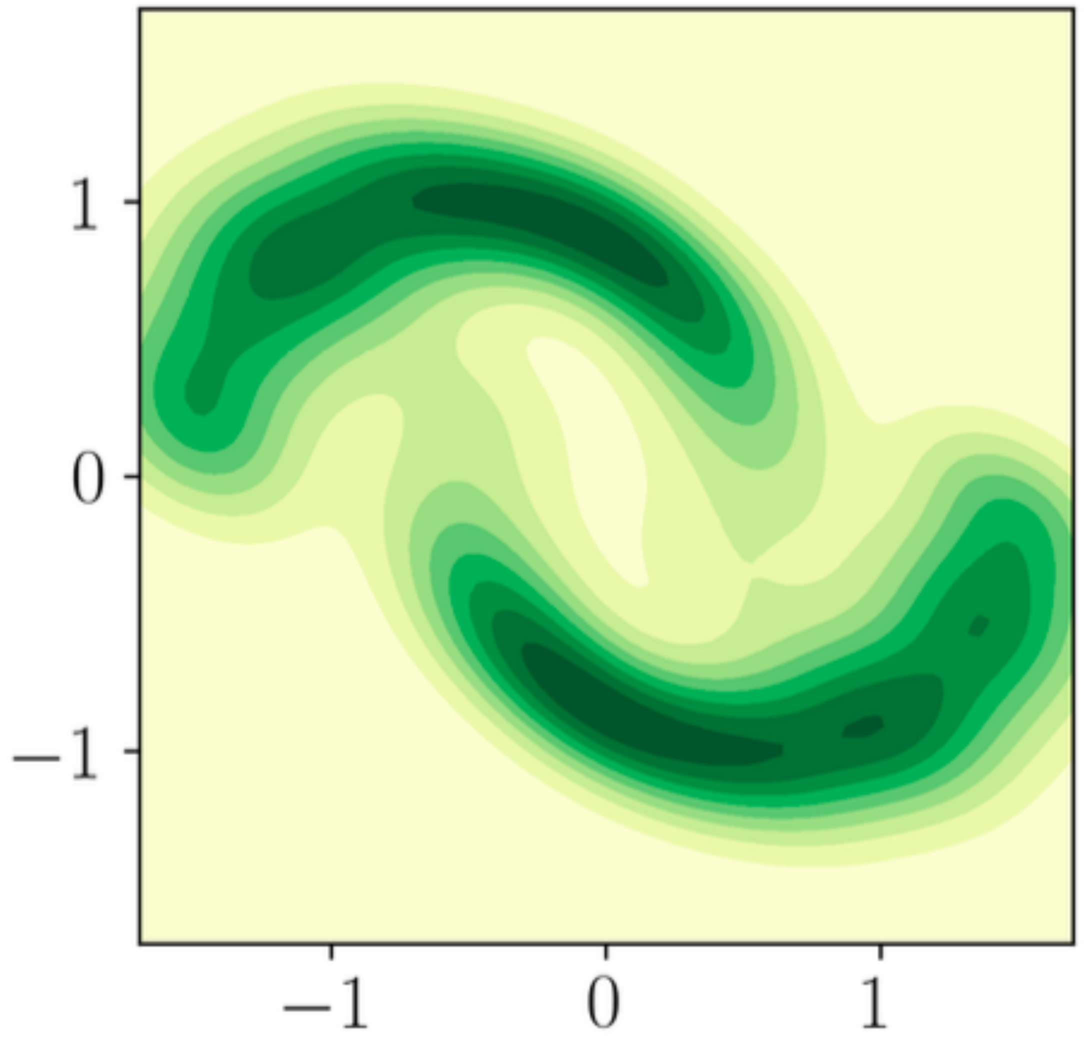

Diffusion model

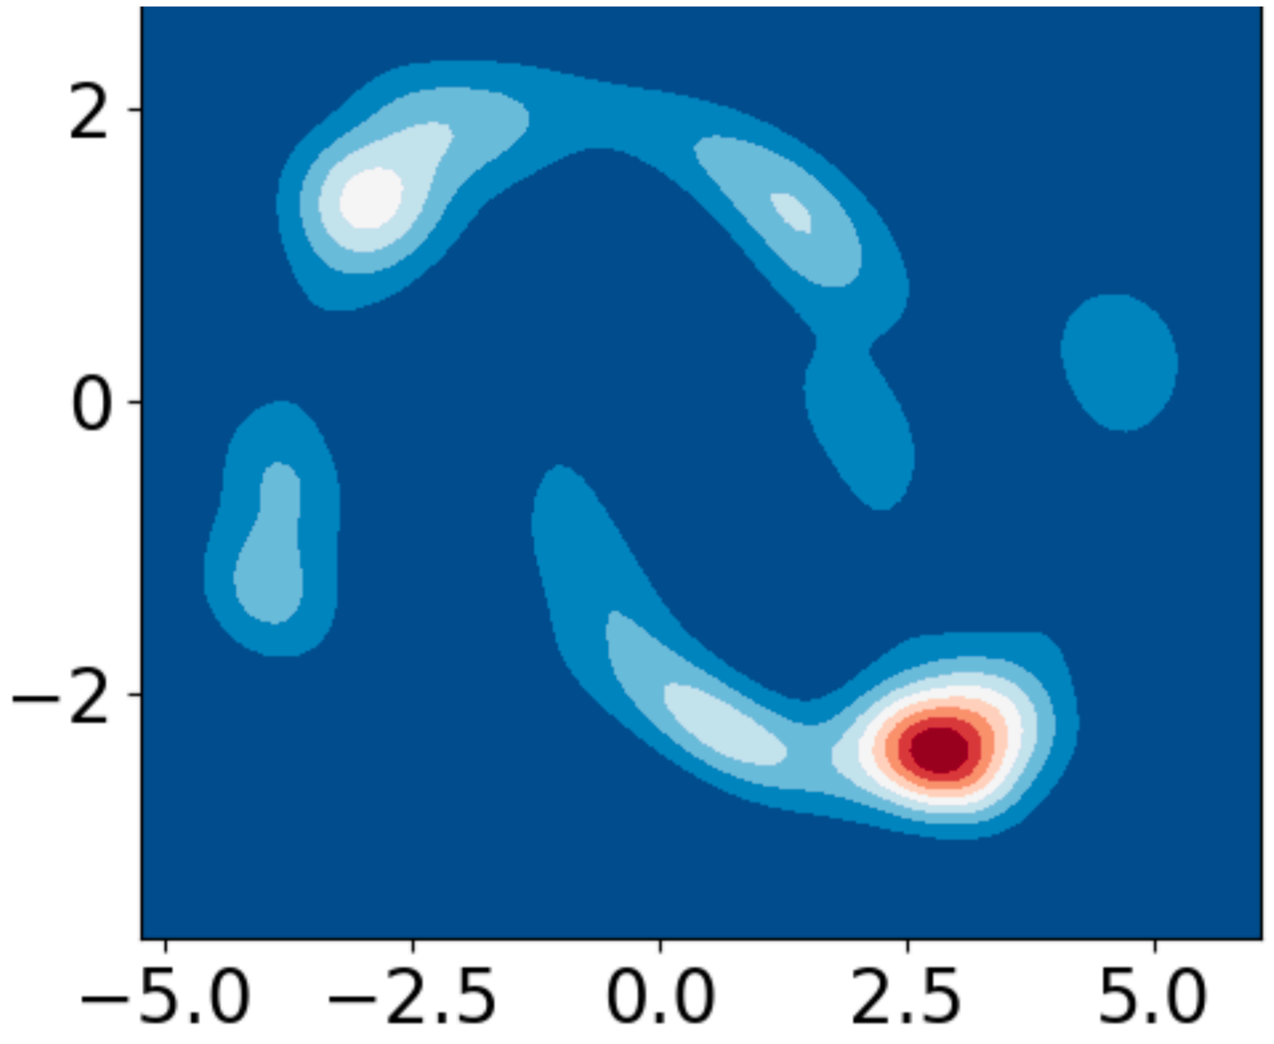

Conditional Flow matching

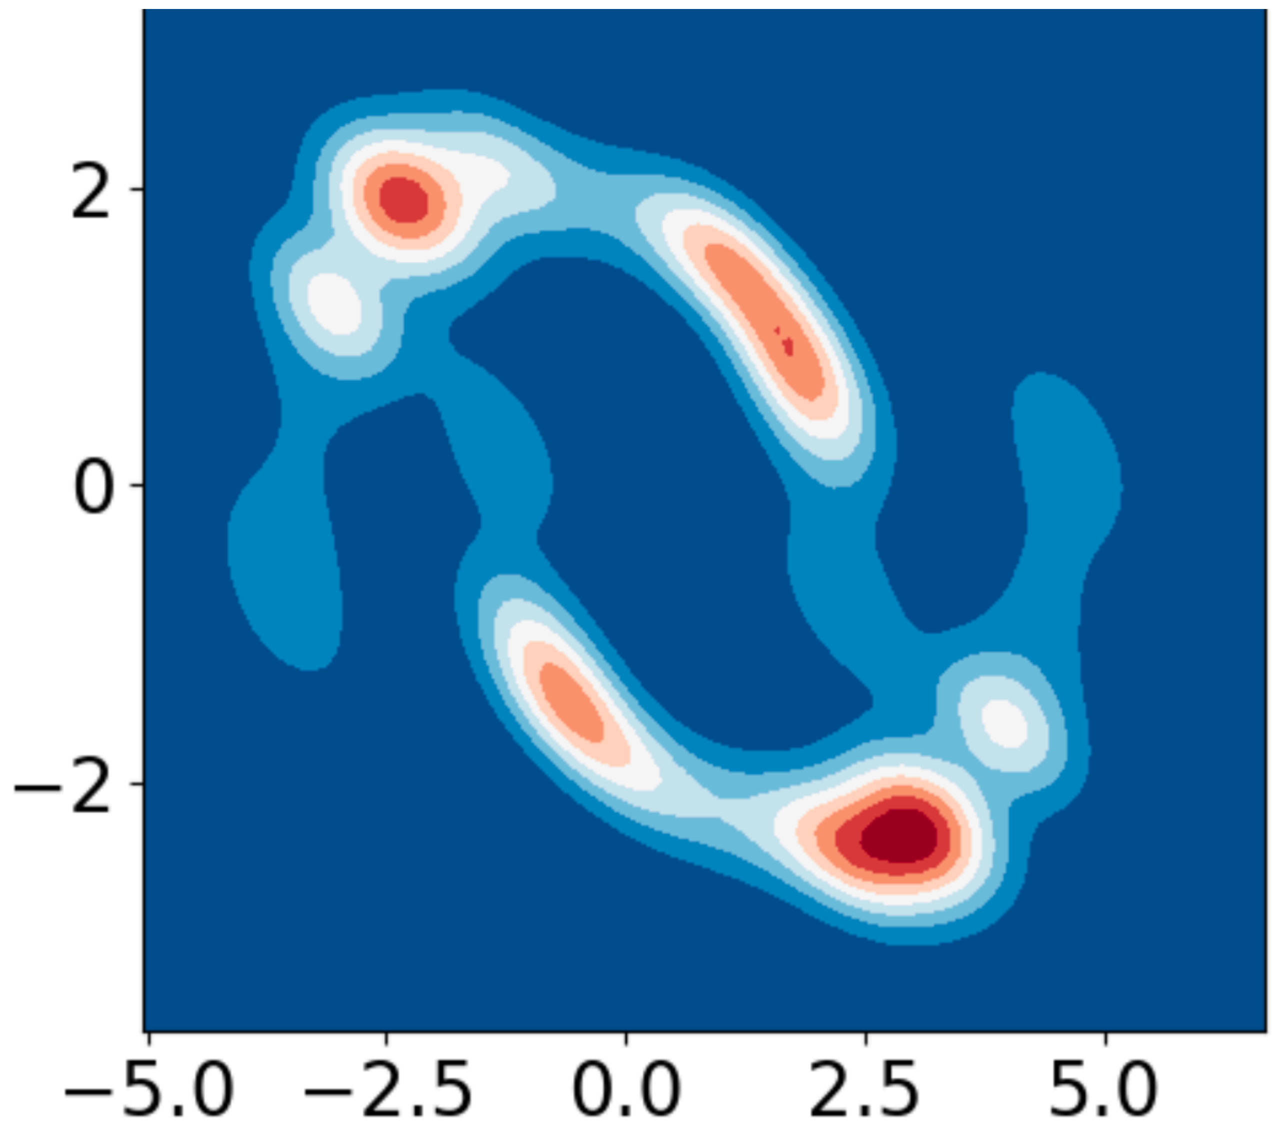

OT - Conditional Flow matching

All models are incorrect, but some are usefully incorrect

Can inexact generative models learn data manifold?

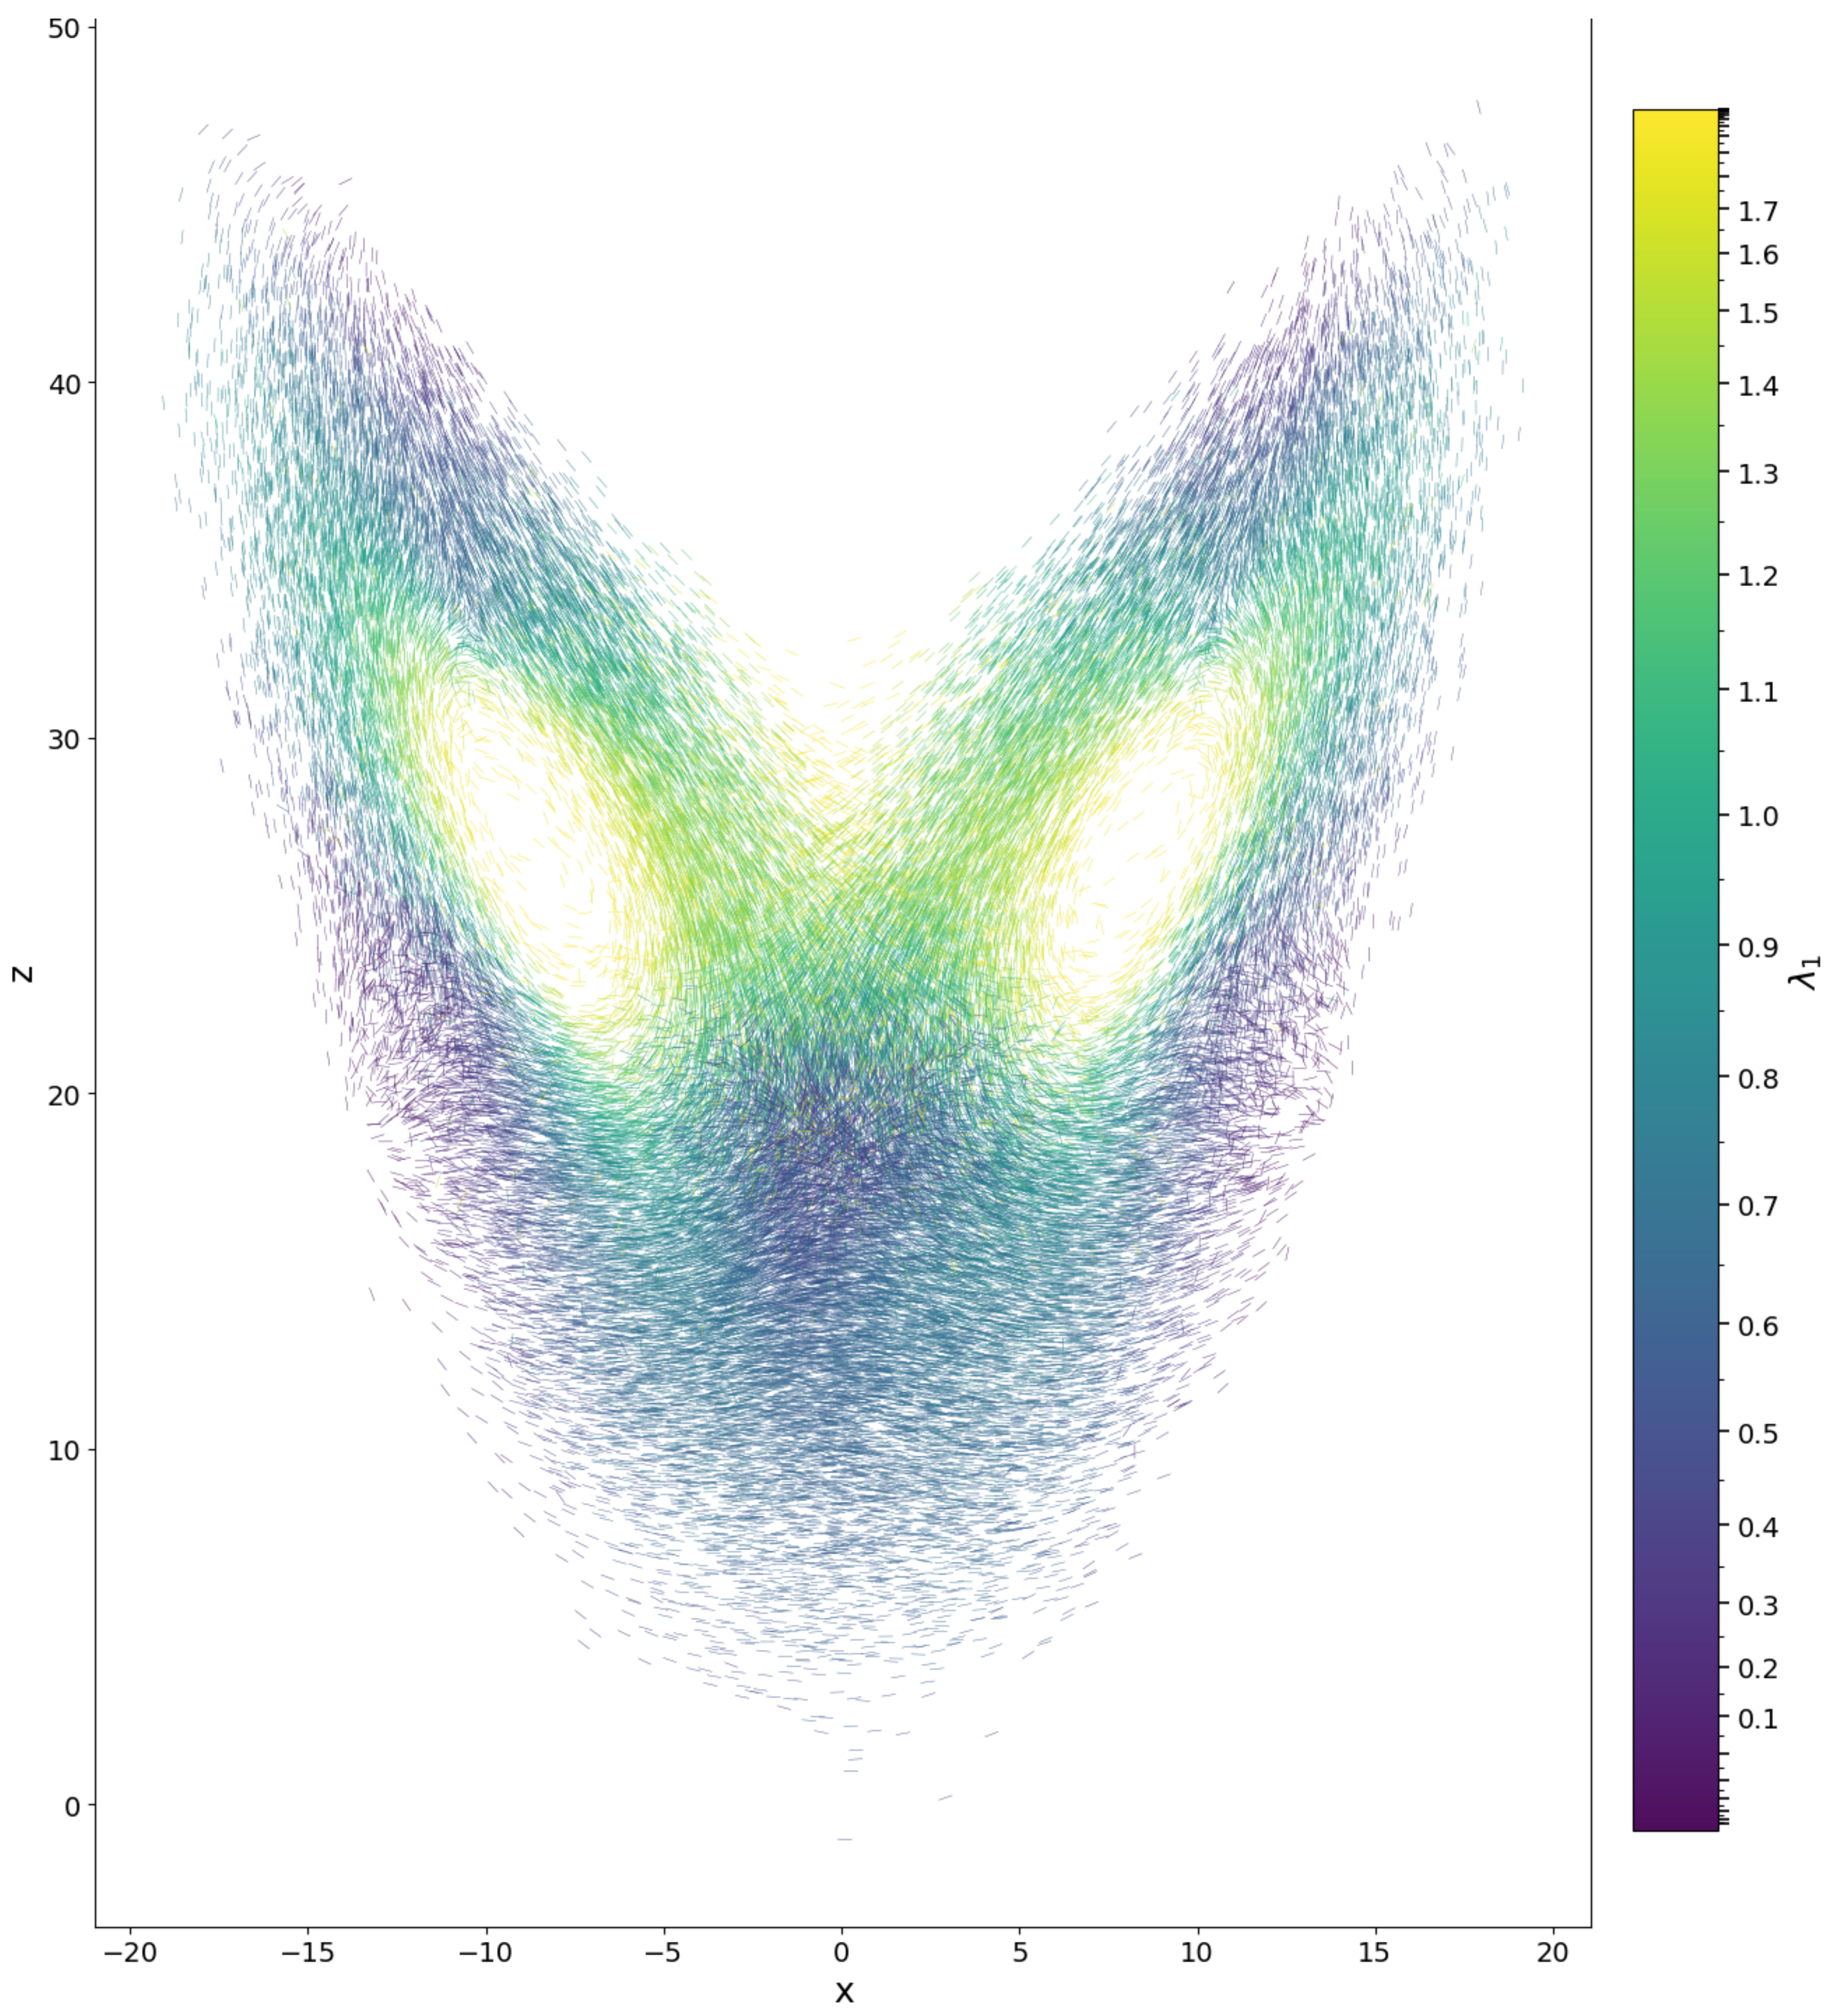

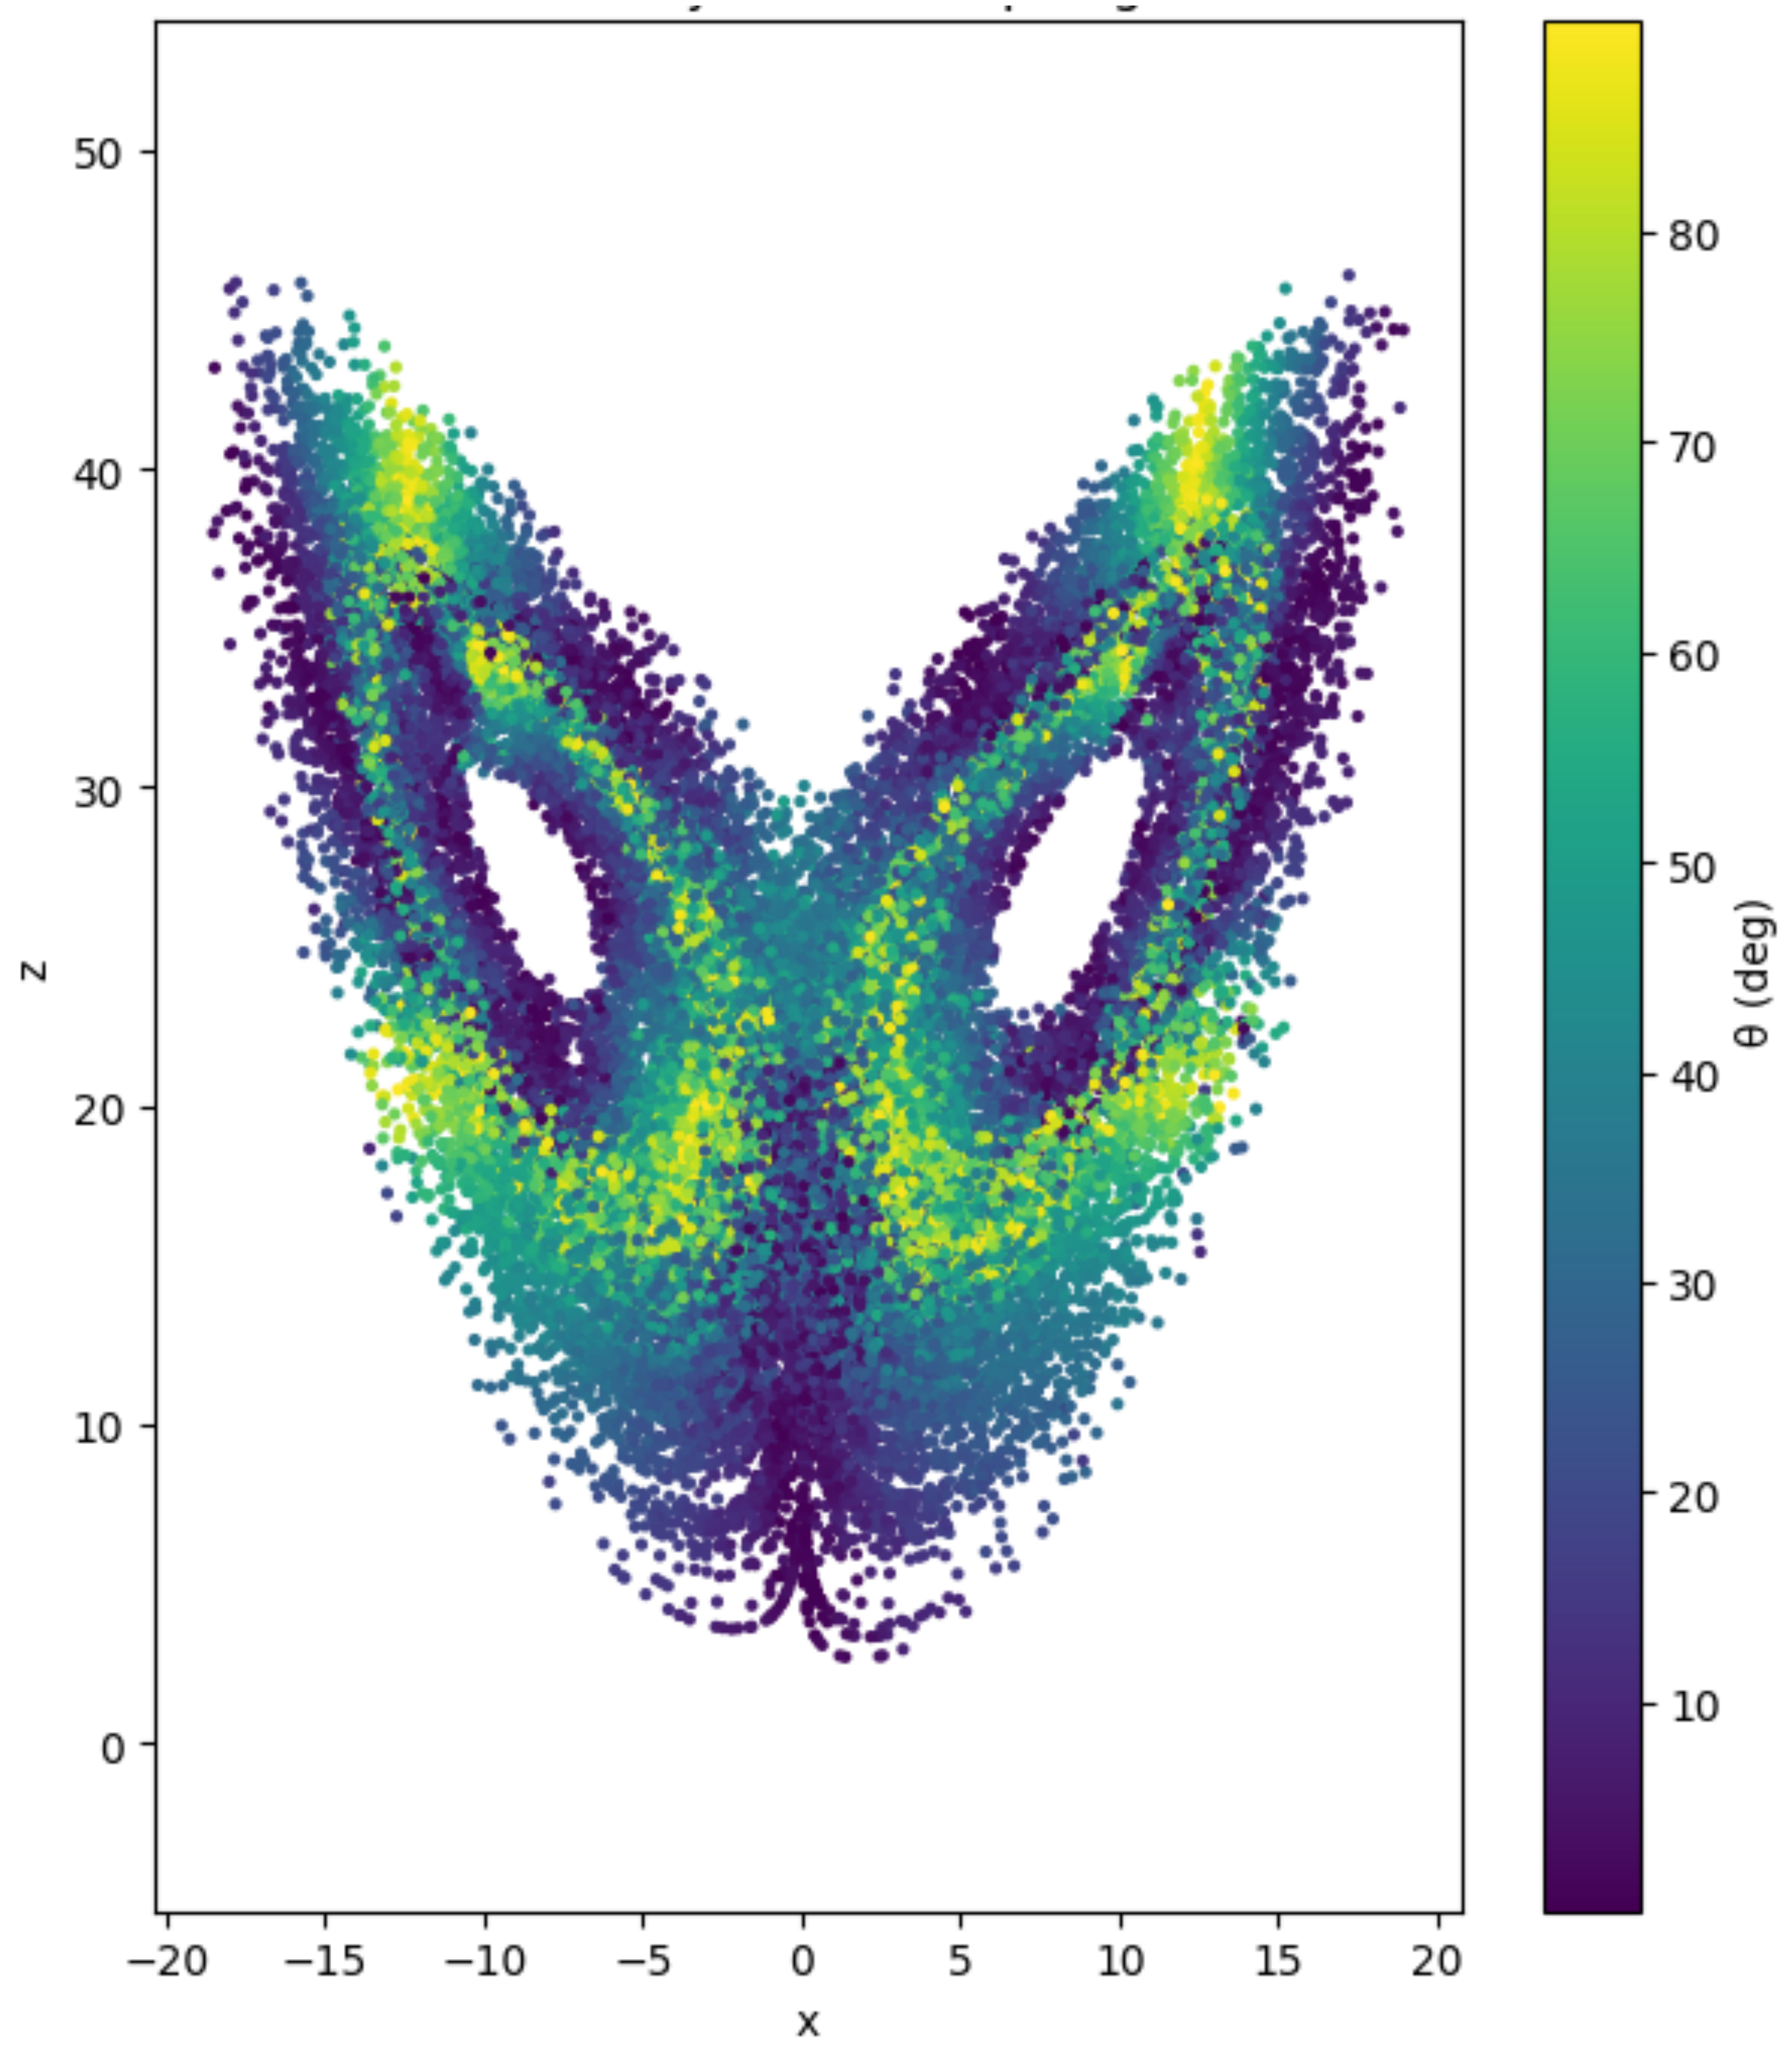

Generative models: setup

- Given samples $x_1, x_2, \cdots, x_m \sim p_{\text{data}}$, generate more samples
- $X \in \mathbb{R}^D$, but support of p_{data} is usually dimension $d < D$
- Typically achieved by learning a dynamical system so that $X_0 \sim p_0$ (usually with density) and $X_\tau \sim p_{\text{data}}$
- Dynamical system can be continuous, vector field $x \to v_t(x)$

Generative models: setup

- Learned vector field: $x \to v_t(x) \in T_x\mathbb{R}^D$
- In score-based diffusion [**Anderson 1982; De Bortoli et al, Song and Ermon 2019, Song et al 2020, Sohl-Dickstein et al 2015**], $v_t(x) \equiv s(x, t)$, scores of densities of a noising process initialized with $x_i \sim p_{\mathrm{data}}$.
- In conditional flow matching variants, stochastic interpolant variants [**Liu 2022, Lipman et al 2023, Tong et al 2023, Albergo et al 2023**], flow of $v_t(x)$ transports probability densities from p_0 (easy) to p_{data} (target)
- $\ell(\theta) = \mathbb{E}_{t, X_t \sim p_t} \| v_{\theta,t}(X_t) - v_t(X_t) \|^2$,
- $\partial p_t / \partial t = -\mathrm{div}(v_t\, p_t)$, with $p_\tau \equiv p_{\mathrm{data}}$.

Random Dynamical Systems

$F^{t,W} = F^W_{t-1} \circ \cdots \circ F^W_0$ is a random dynamical system, where at each time t, we choose $F^W_t \sim \pi_t$.

Discrete-time RDS: $W = \{W_t\}$ iid standard normal RVs.

For SGM: $F^W_t(x) = x + (\delta t)s(x, \tau - t) + \sqrt{\delta t}\, W_t$, $W_t \sim \mathcal{N}(0, \mathrm{Id})$.

One step flow map: can come from any integration scheme.

Classical field: Kifer, Young, Ledrappier, Pesin, Arnold …

GM with learning errors

$F_{t,\epsilon}(x) = F_t(x) + \epsilon\, \chi_t(x)$, map at time t with learning errors.

$$u_{t+1}(F_t(x)) = \nabla F_t(x)\, u_t(x) + \chi_t(x)$$

Tracking infinitesimal errors

Probability space

Perturbed evolution/-pushforward operator: $p_{t+1,\epsilon} := F_{t,\epsilon\sharp} p_{t,\epsilon}$

When $p_{t,\epsilon}$ has density ρ_t,

$$\begin{aligned} \rho_{t+1,\epsilon} &= \mathcal{L}_{t,\epsilon}\ \rho_{t,\epsilon} \\ &:= \rho_{t,\epsilon} \circ F_{t,\epsilon}^{-1} / |\mathrm{det}\, dF_{t,\epsilon}| \circ F_{t,\epsilon}^{-1} \end{aligned}$$

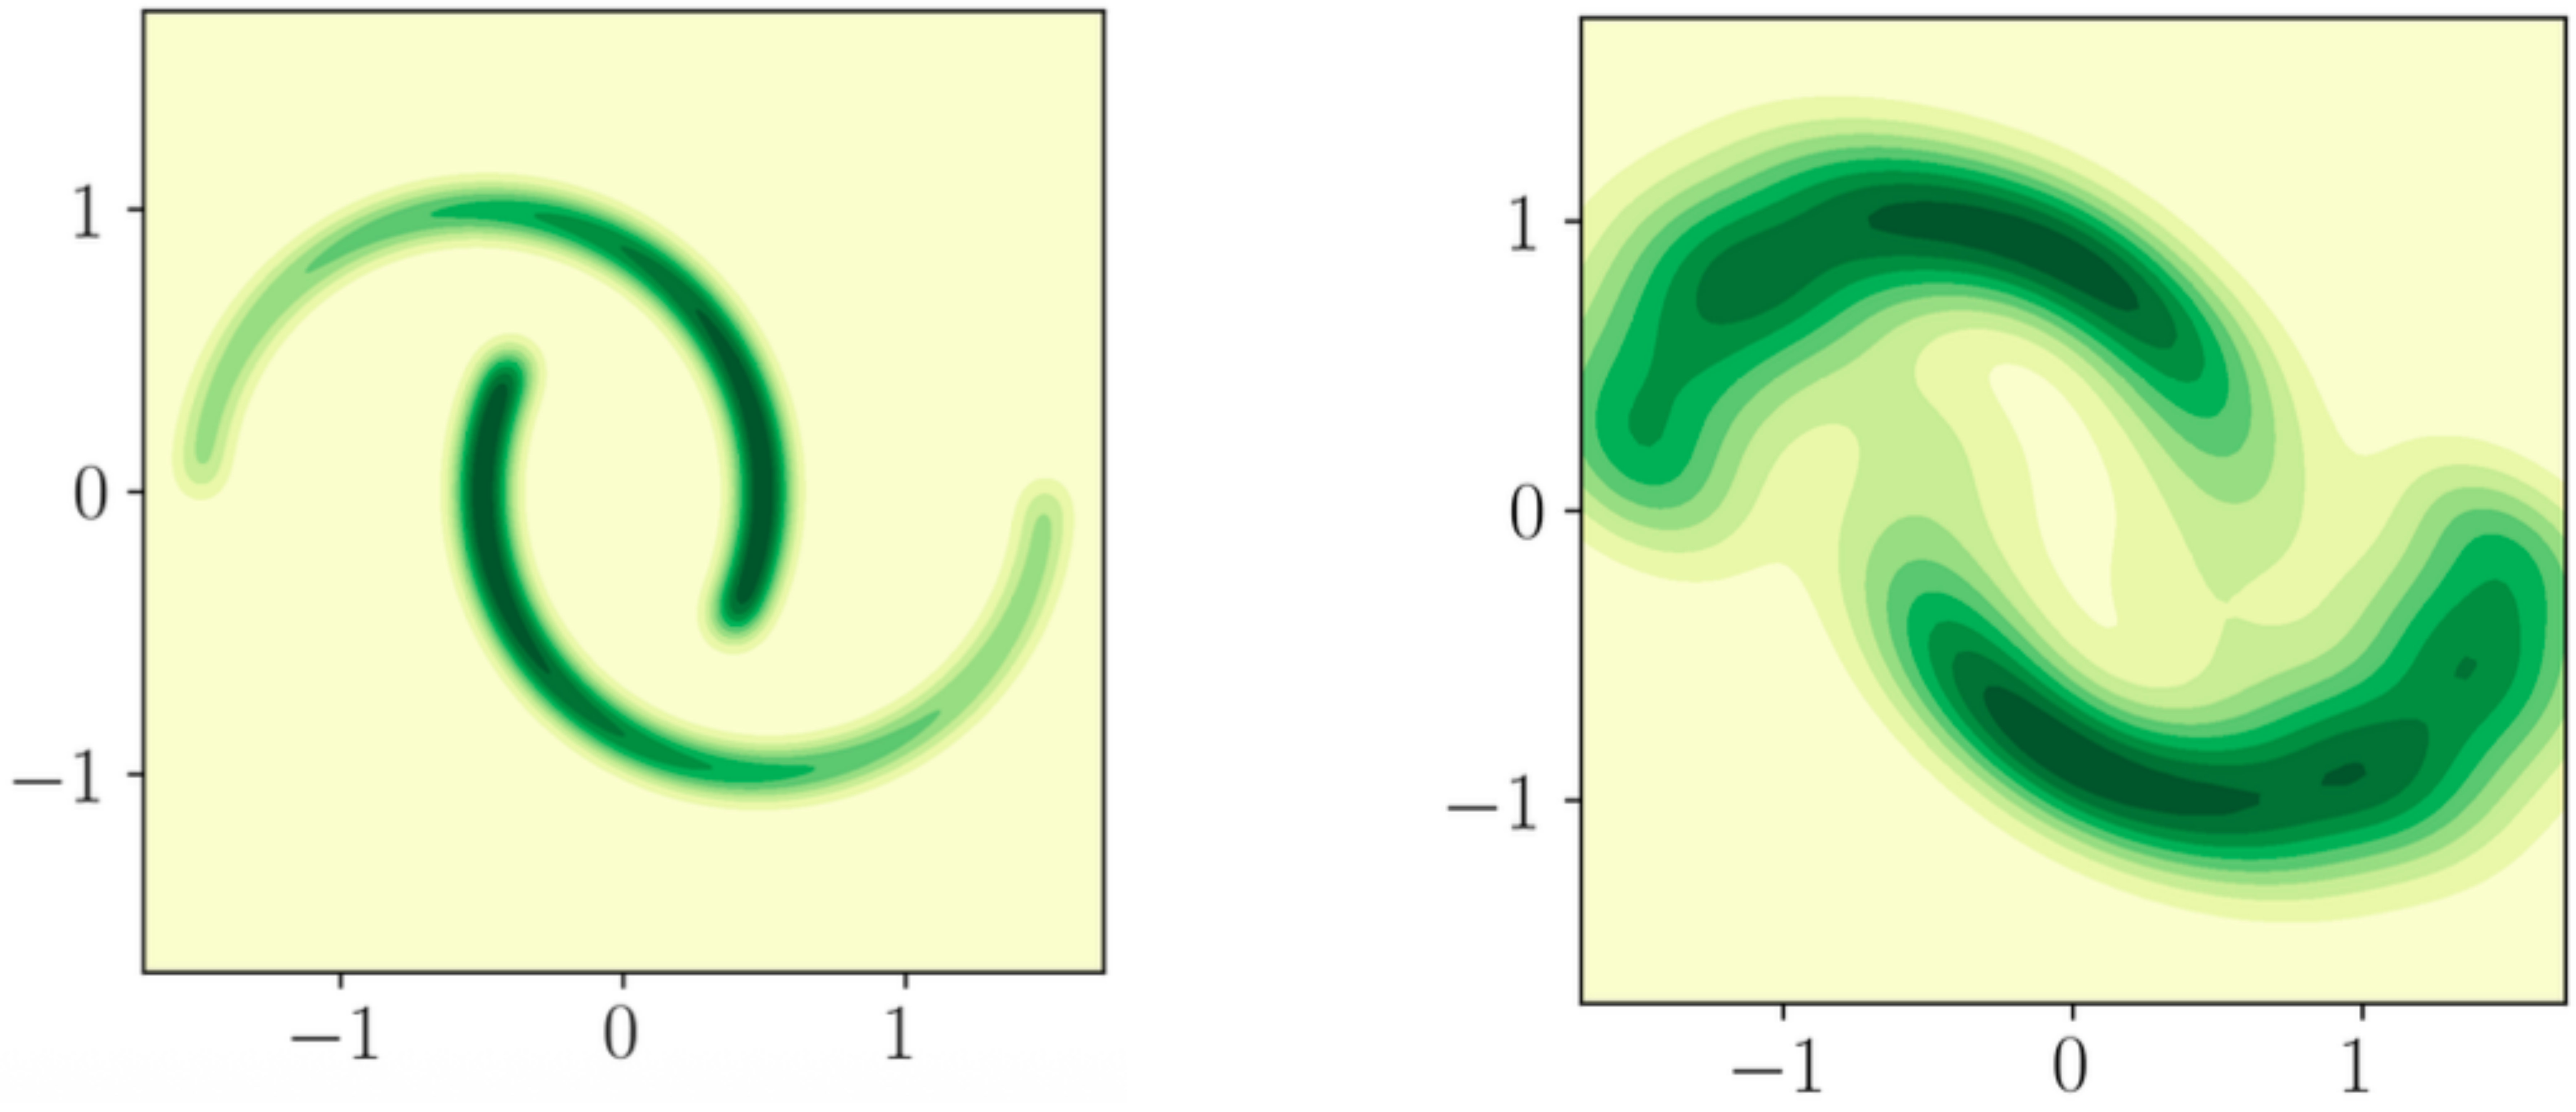

The support appears robust

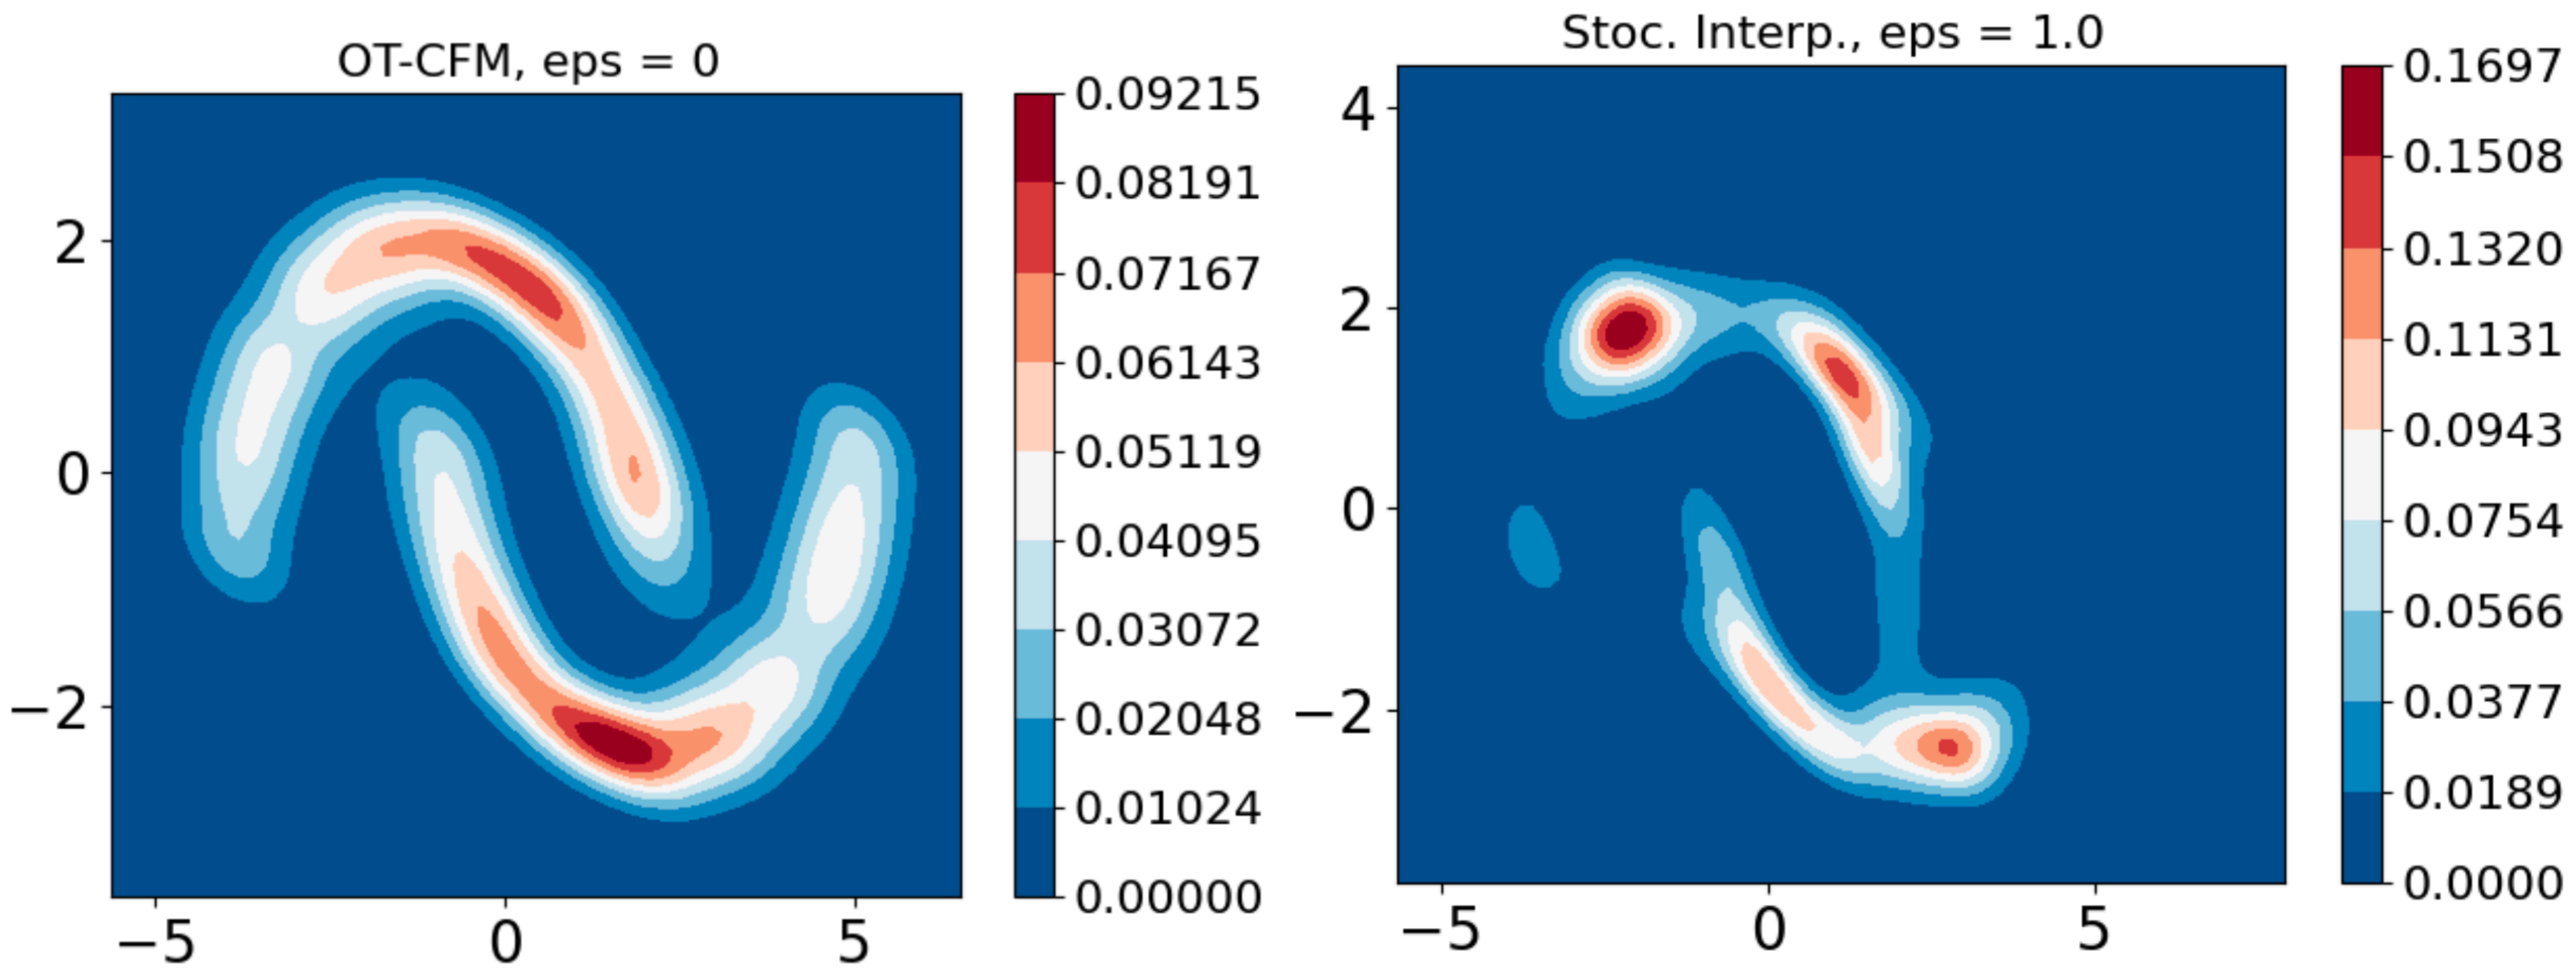

The support is not robust

How do densities change under errors in the dynamics?

$$\partial_\epsilon \rho_\epsilon(x_\tau) = \rho(x_\tau) \sum_{t<\tau} (\mathrm{div}(\chi_t) + \chi_t \cdot s_t)(x_{\tau-t})$$

A bounded continuous scalar field

- Assuming absolute continuity, and regularity of F_t
- Most generative models should have robust support to small perturbations

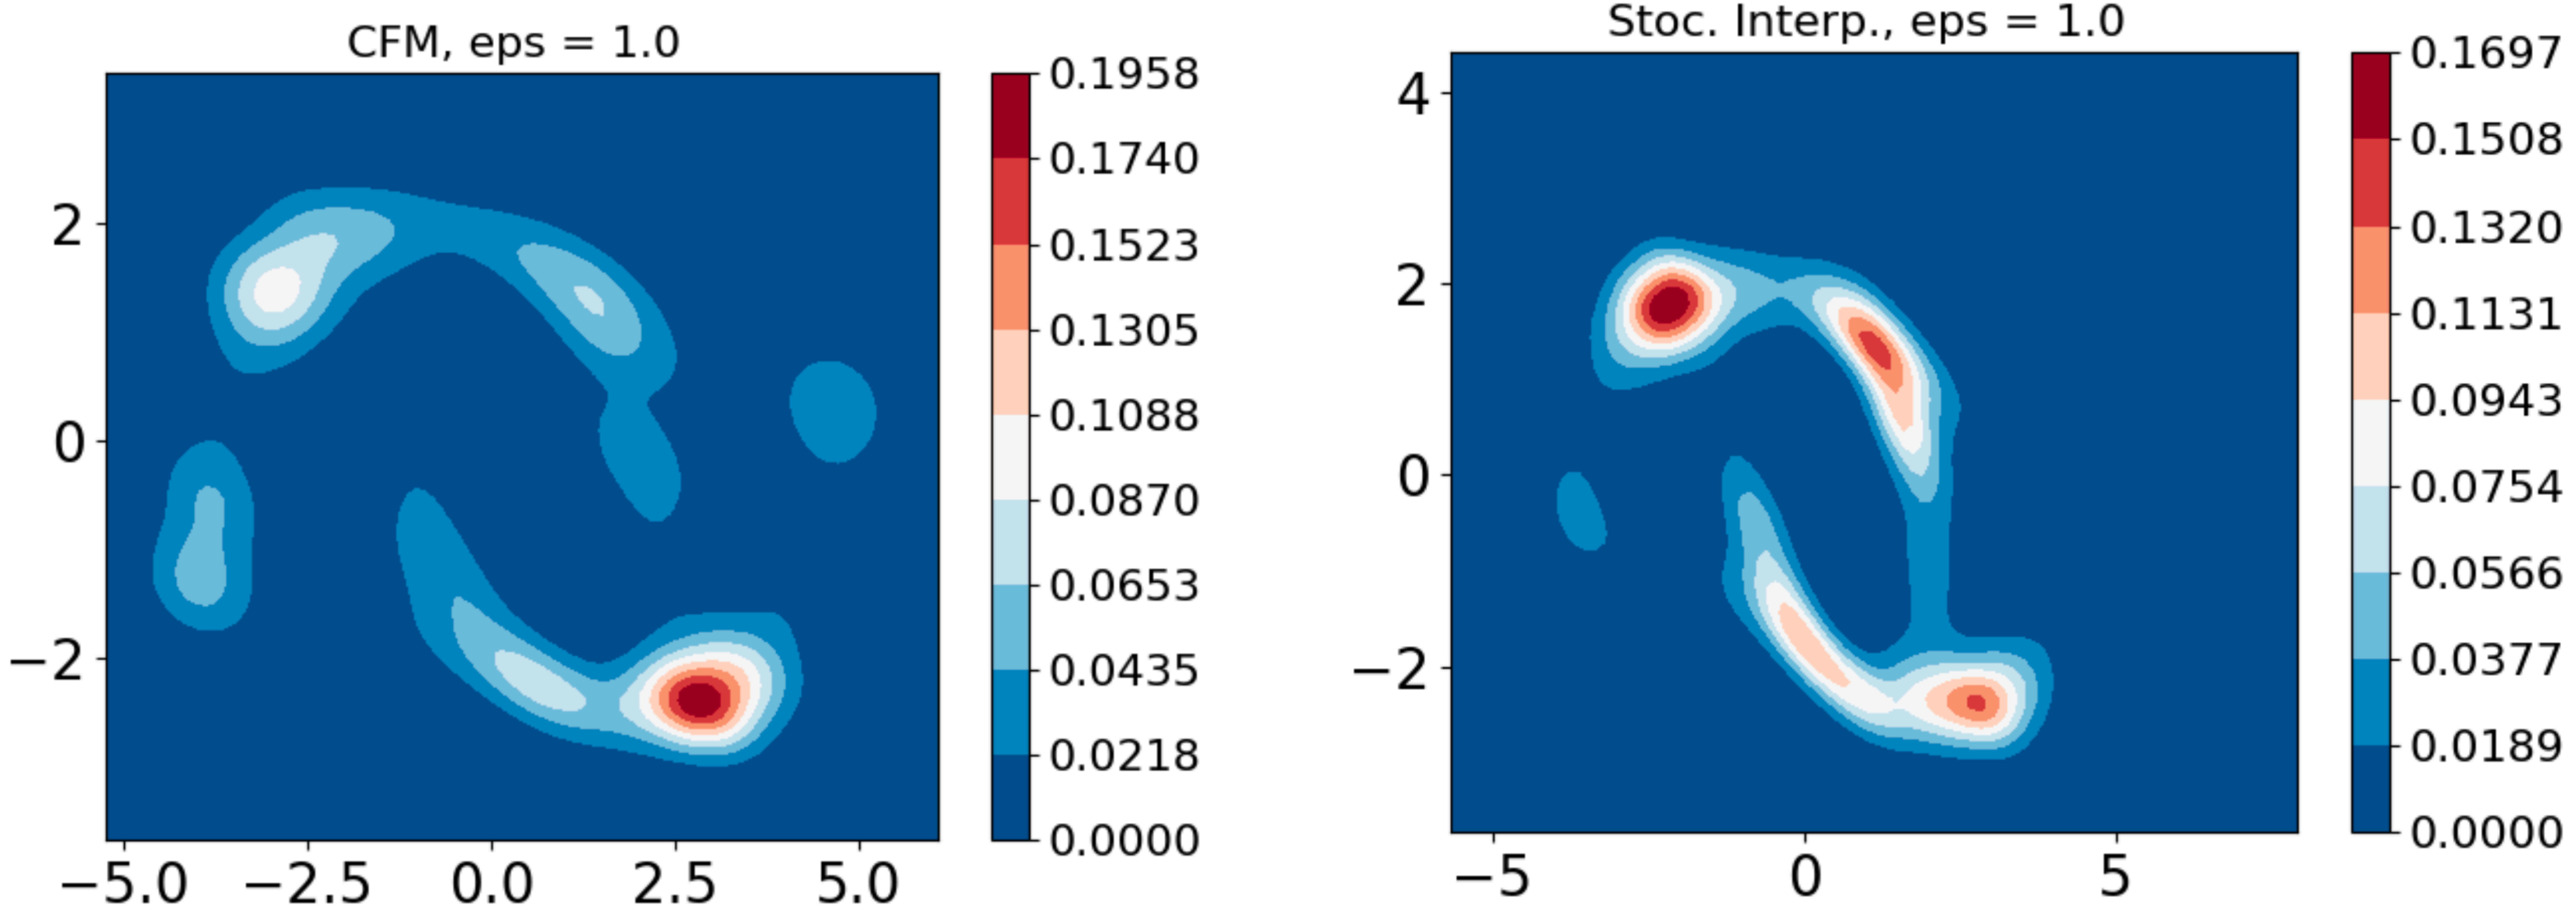

Does not explain lack of robustness

No directionality/control in analysis

Defining most sensitive subspaces

- Recall Jacobian map: $\nabla F_t : \mathbb{R}^D \to \mathbb{R}^D$.
- Let $E_0(x) \in \mathbb{R}^{D\times d}$ be a random subspace at each x
- $E_{t+1}R_{t+1} := \nabla F_t \, E_t$ from a QR decomposition
- E_τ : most sensitive subspace at time τ

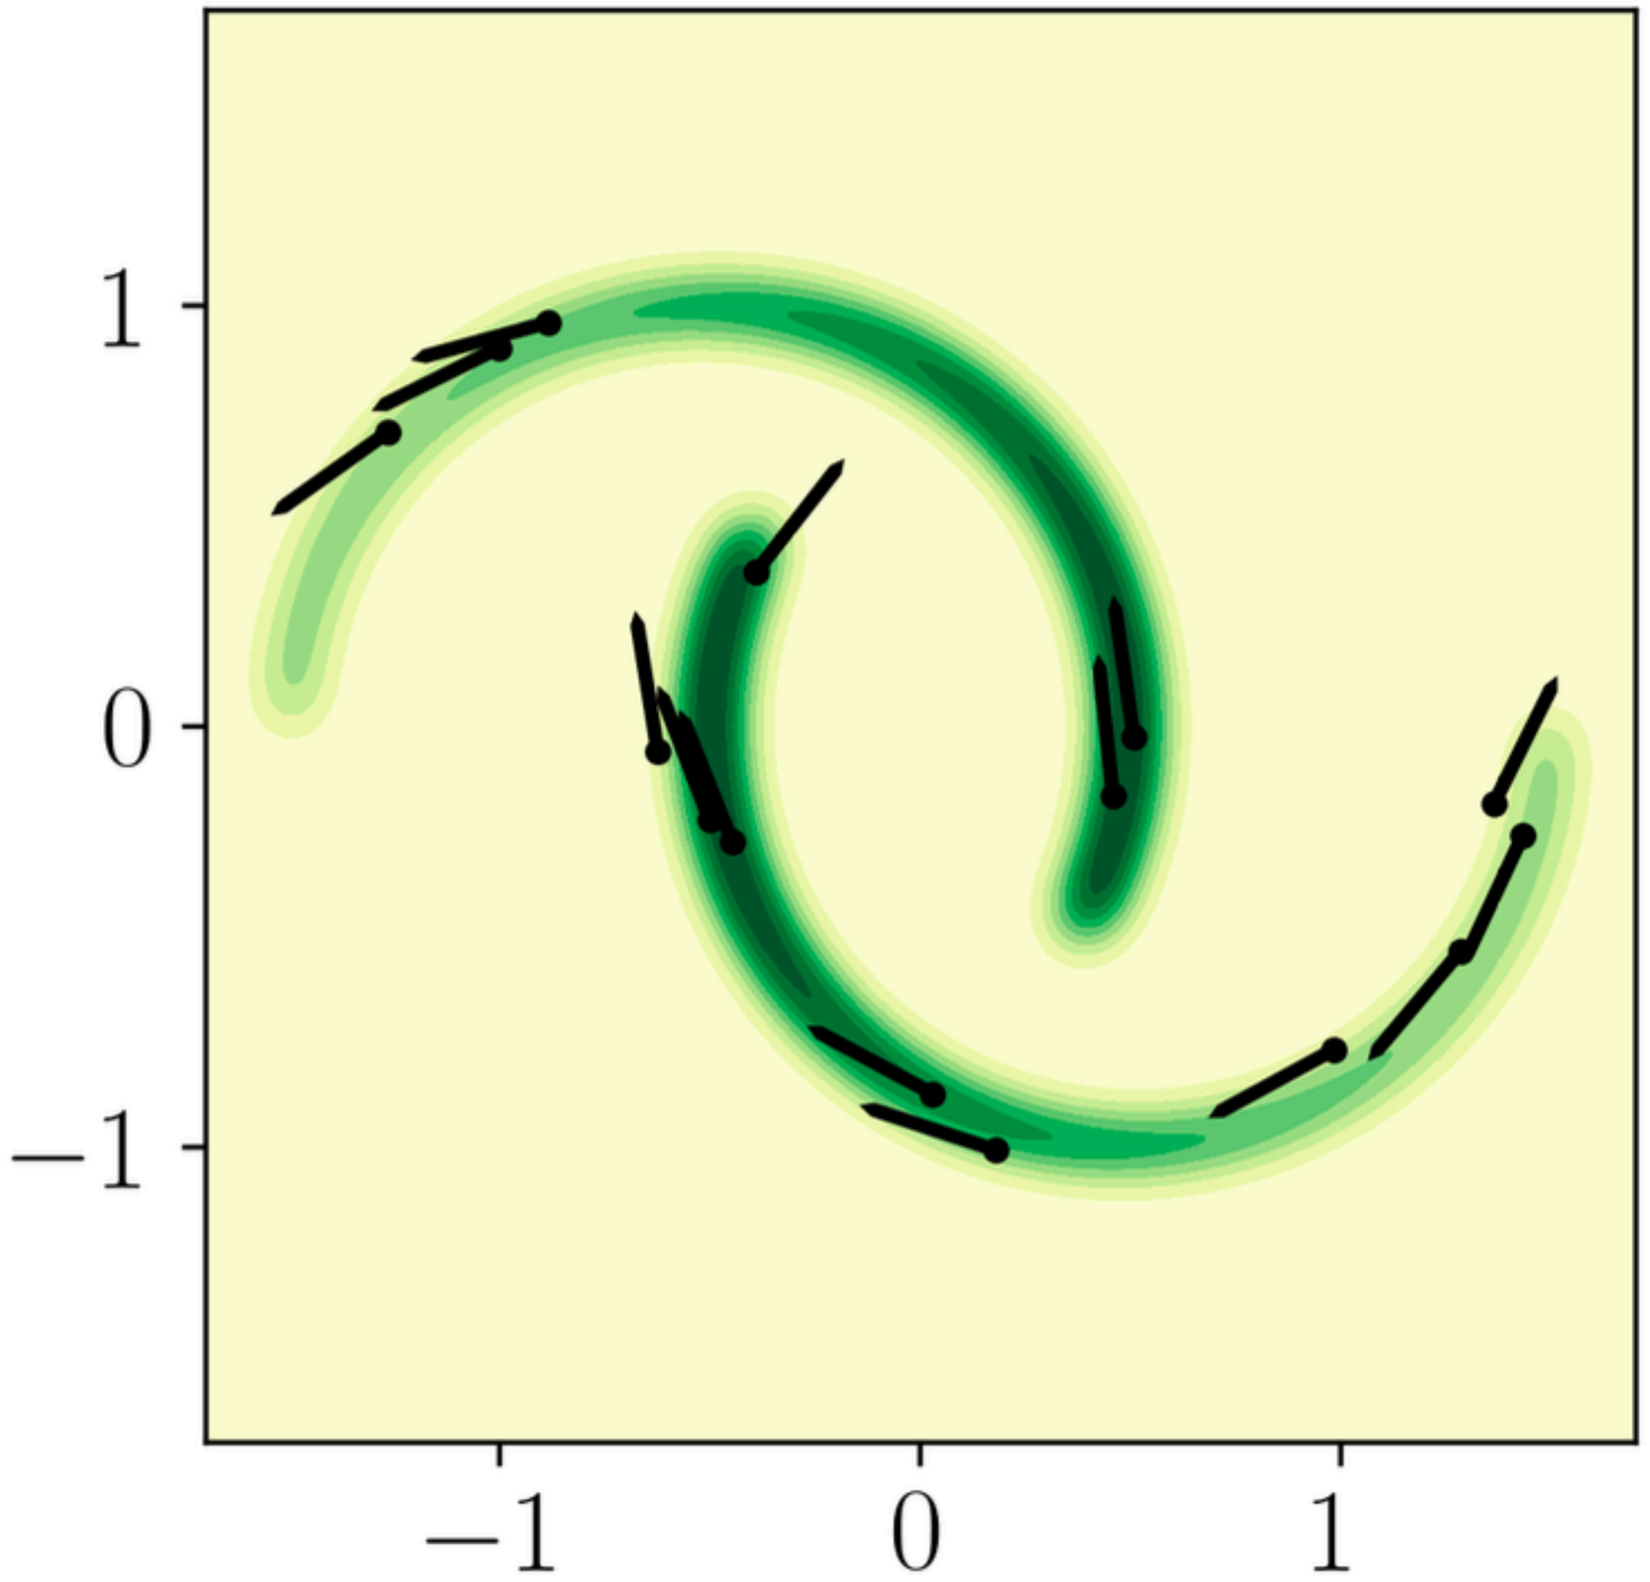

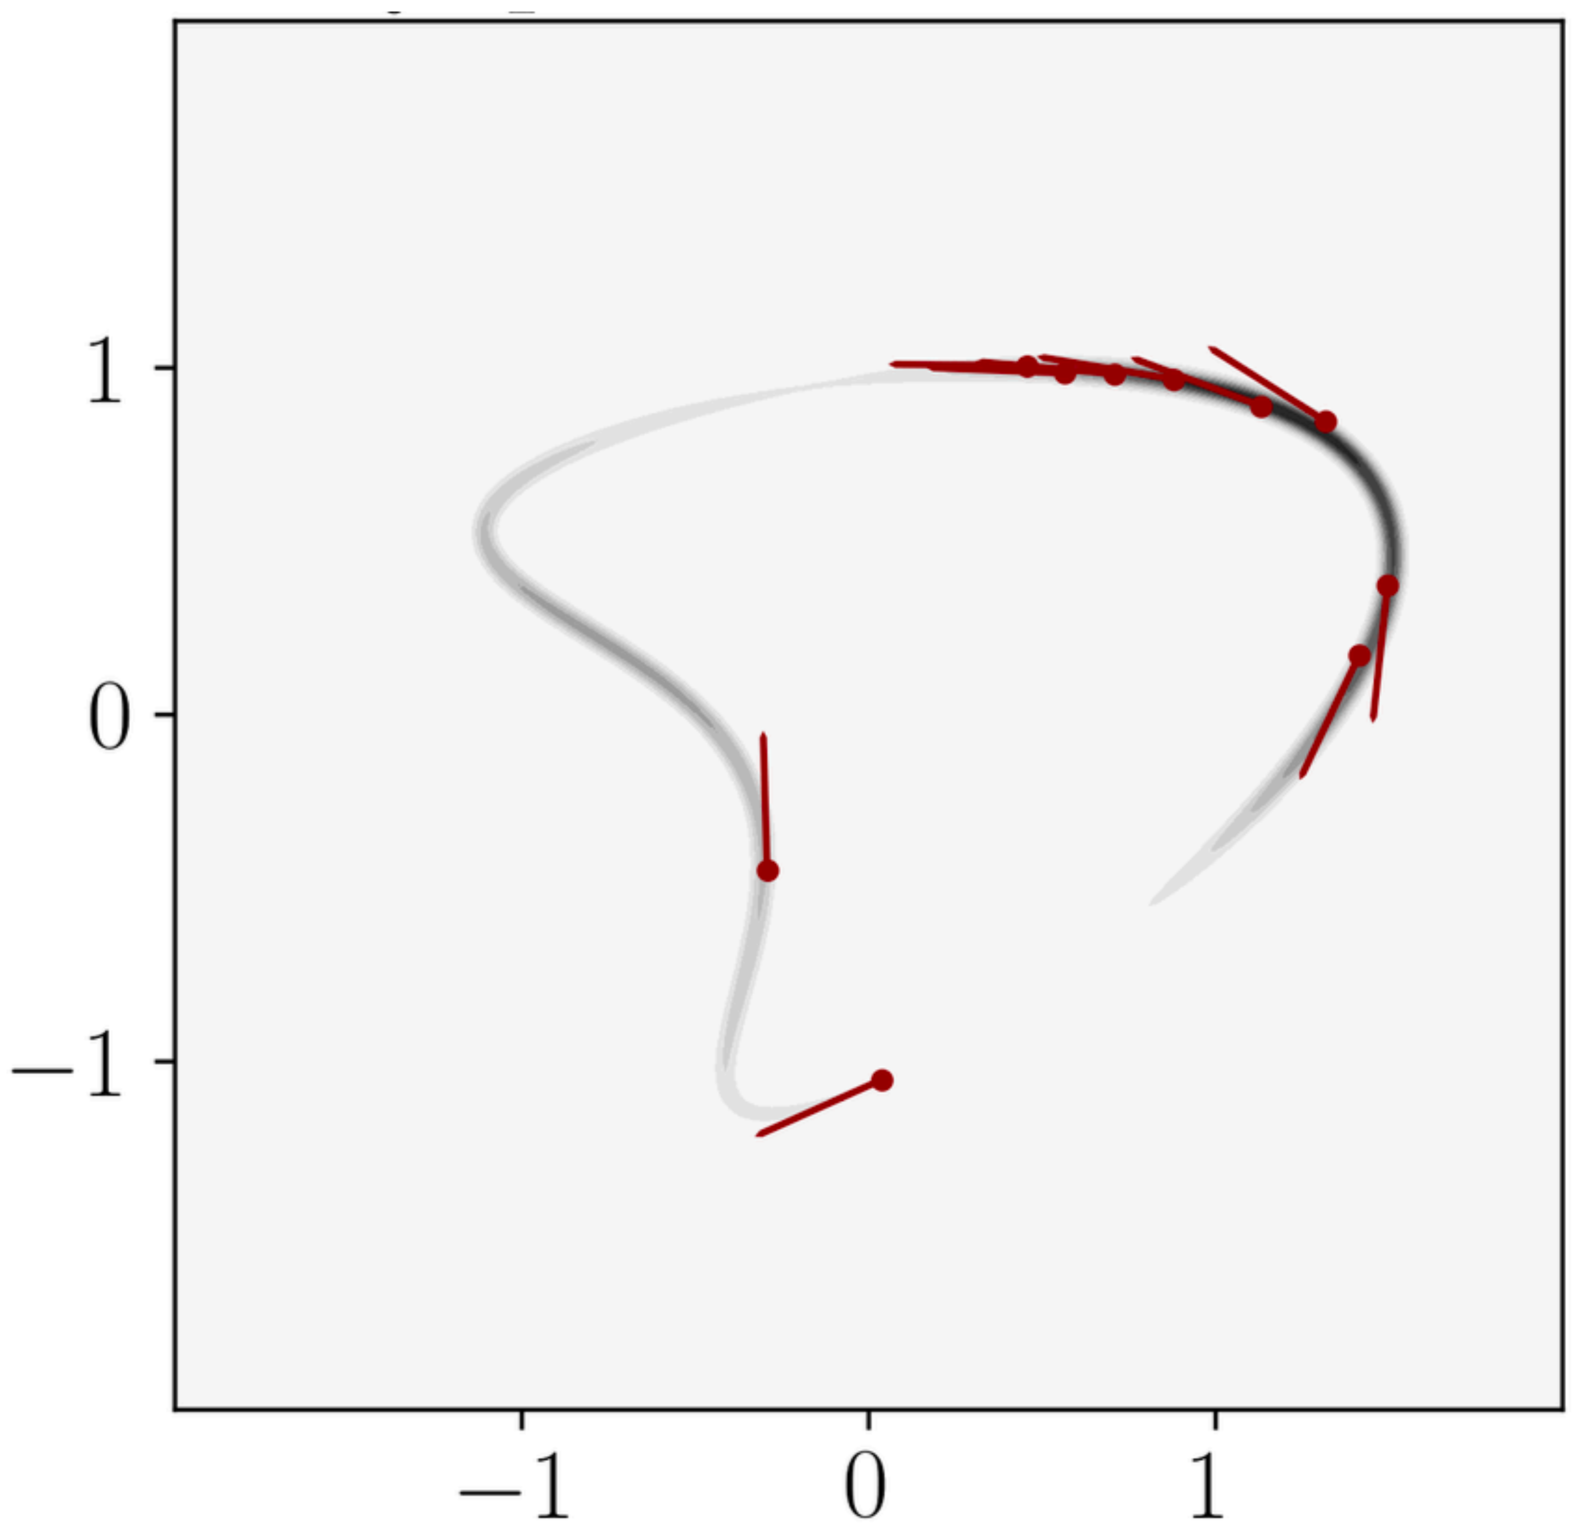

$E_\tau \approx$ top d Lyapunov Vectors

- d: intrinsic dimension.
- $\approx$ top d eigenvectors of $\nabla F^\tau (\nabla F^\tau)^\top$: Cauchy-Green deformation tensor in fluid mechanics

Alignment with support of p_{data}

- M: support of p_{data}
- A GM is ***aligned*** if TM is parallel to E_τ

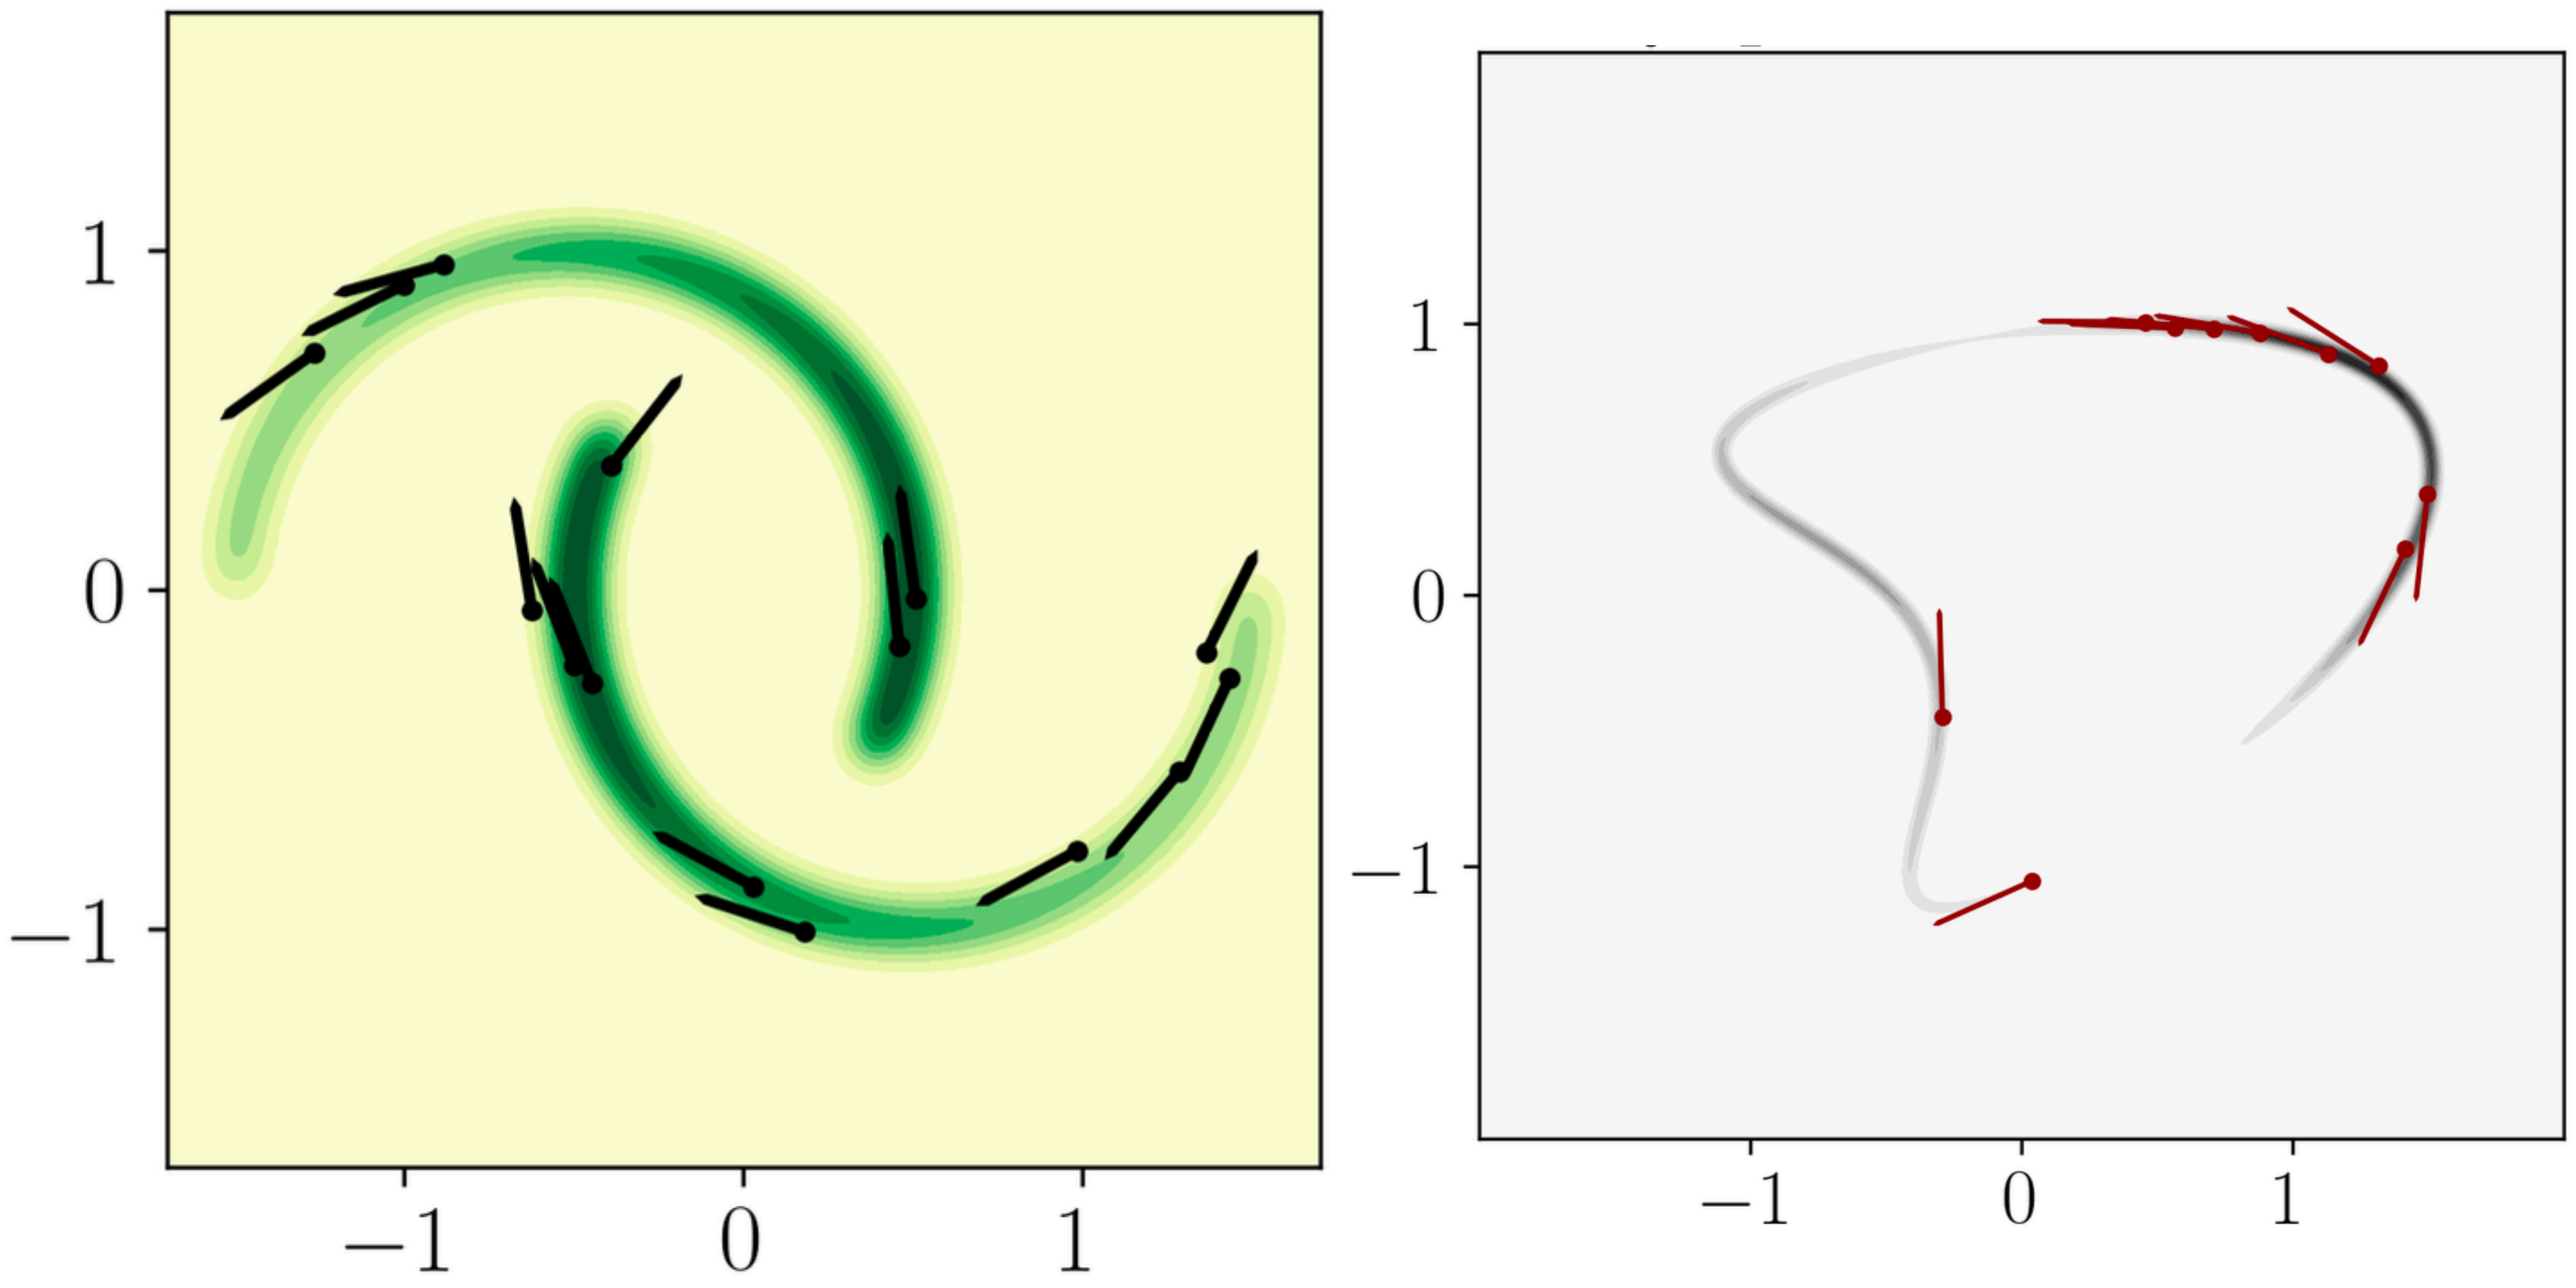

Learning the support

Informal: A convergent generative model learns the support of the target if **aligned**

- Convergent: optimal transport map of small norm exists between generated samples and the target samples; $||x_i - y_i|| \sim \mathcal{O}(\epsilon^c)$ (e.g., Lee, Lu, Tan 2023)
- Support estimation $\equiv$ Learning one-class classifier
- Margin does not change under alignment, hence has same generalization error

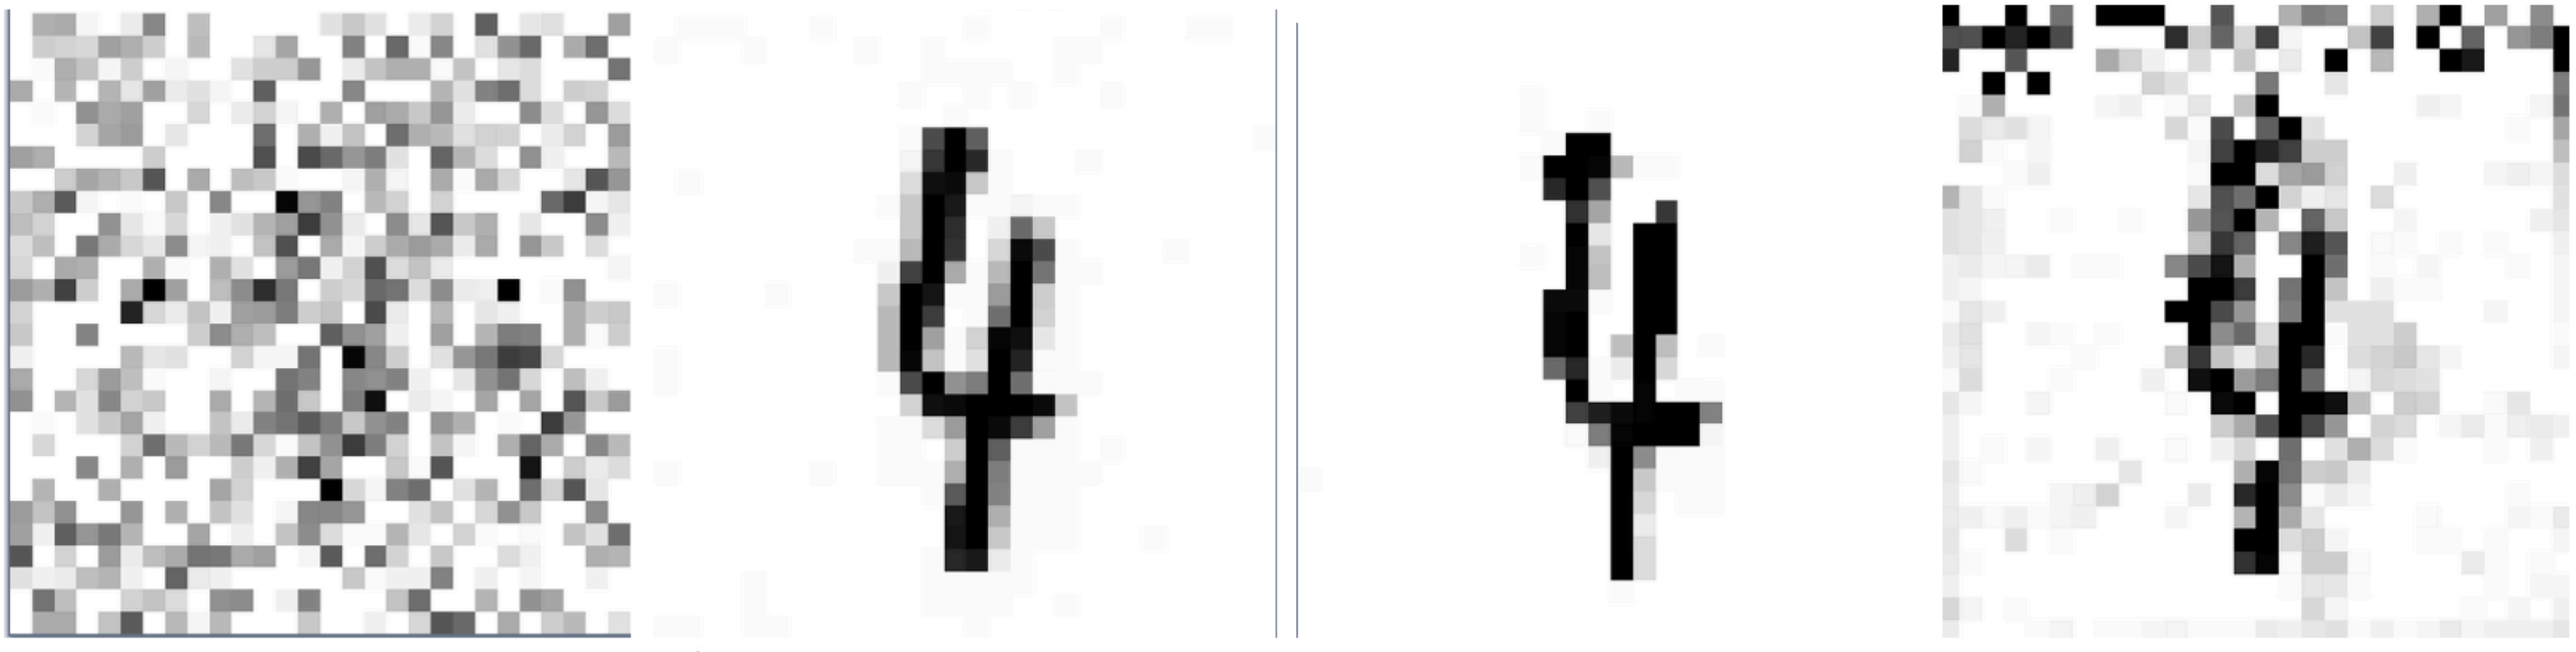

Noise image

Generated by SGM

+ LV 1 (1st column of E_τ)

+ LV 100 (100th column of E_τ)

Alignment in practice

Alignment leads to learning the dimension of the data manifold

Finite-time Lyapunov exponents (mean of log diagonal of R_t)

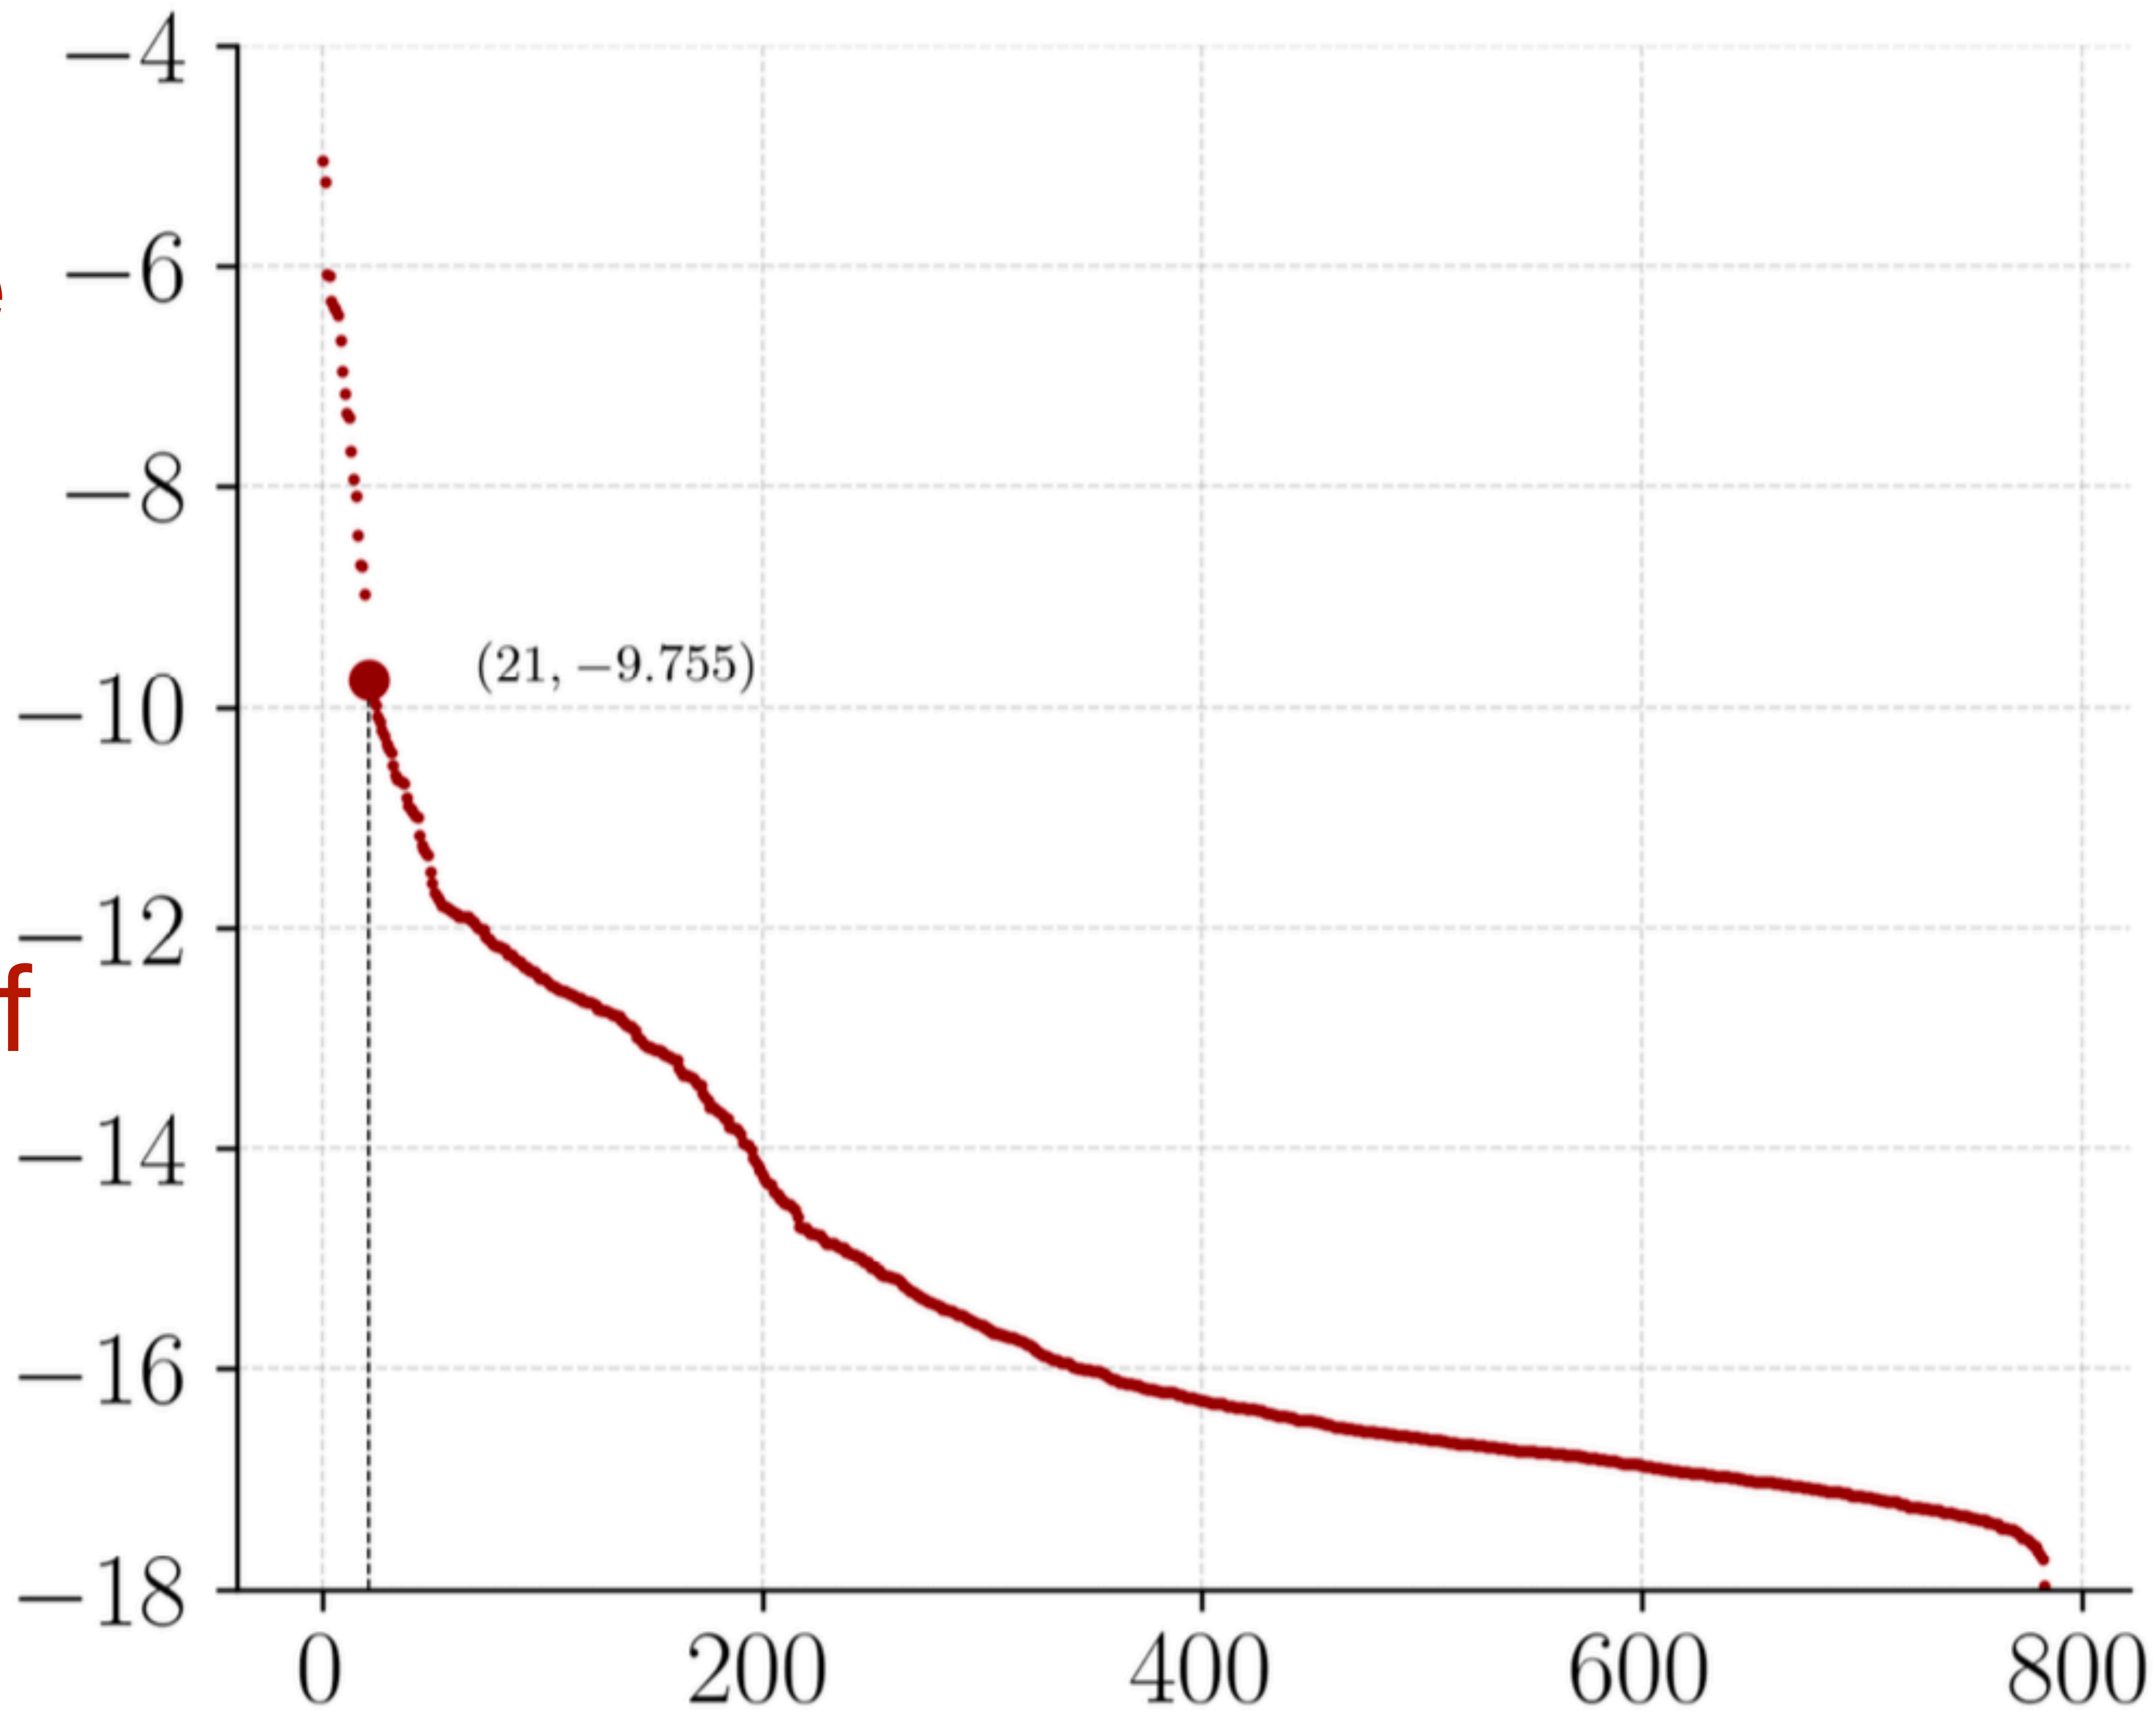

Index

Pidstrigach 2022; Stanczuk et al 2024; Kadkhodaie et al 2024; Chen, Huang, Zhao, and Wang 2023; Lee Lu Tan 2023; Mimikos-Stamatopoulos, Zhang, Katsoulakis 2024

Alignment leads to learning the data manifold

- The most sensitive subspaces E_τ are constructively defined; scalable computation
- Is **alignment** preserved under perturbations? Yes, under smooth perturbations.
- What kind of dynamics leads to **alignment**?
- How to ensure **alignment**?

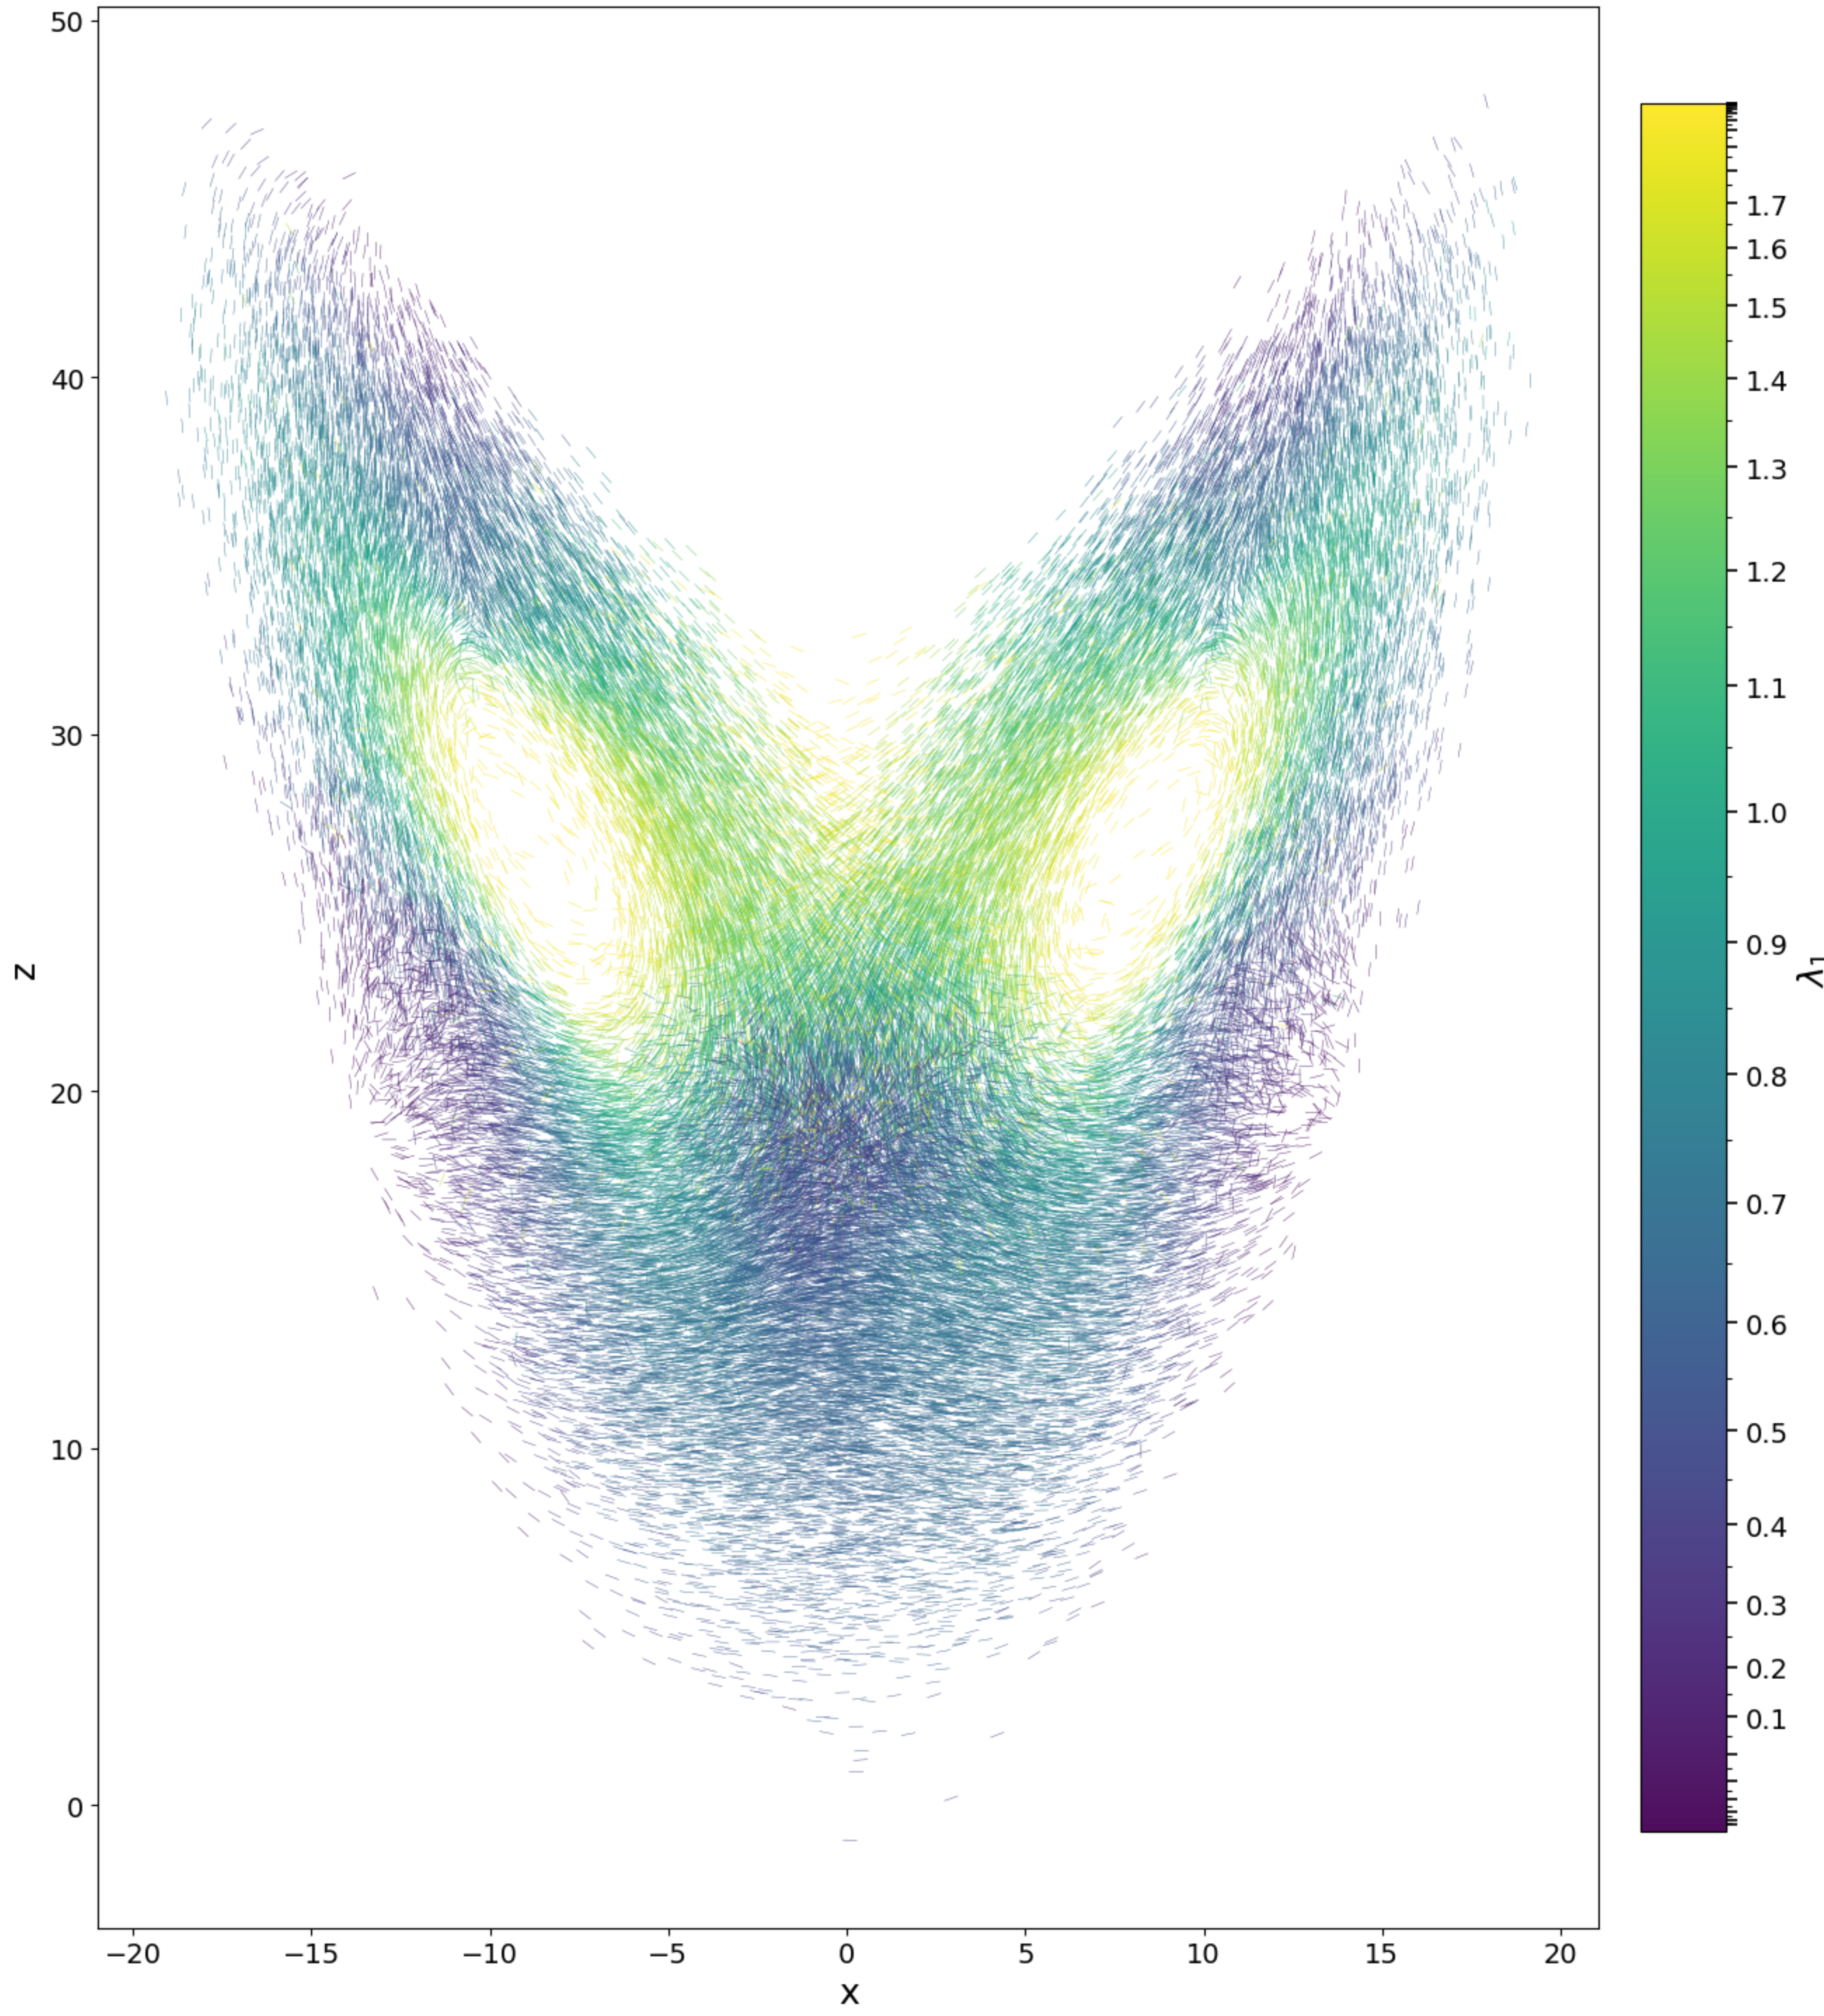

Vector fields (scores) in diffusion models

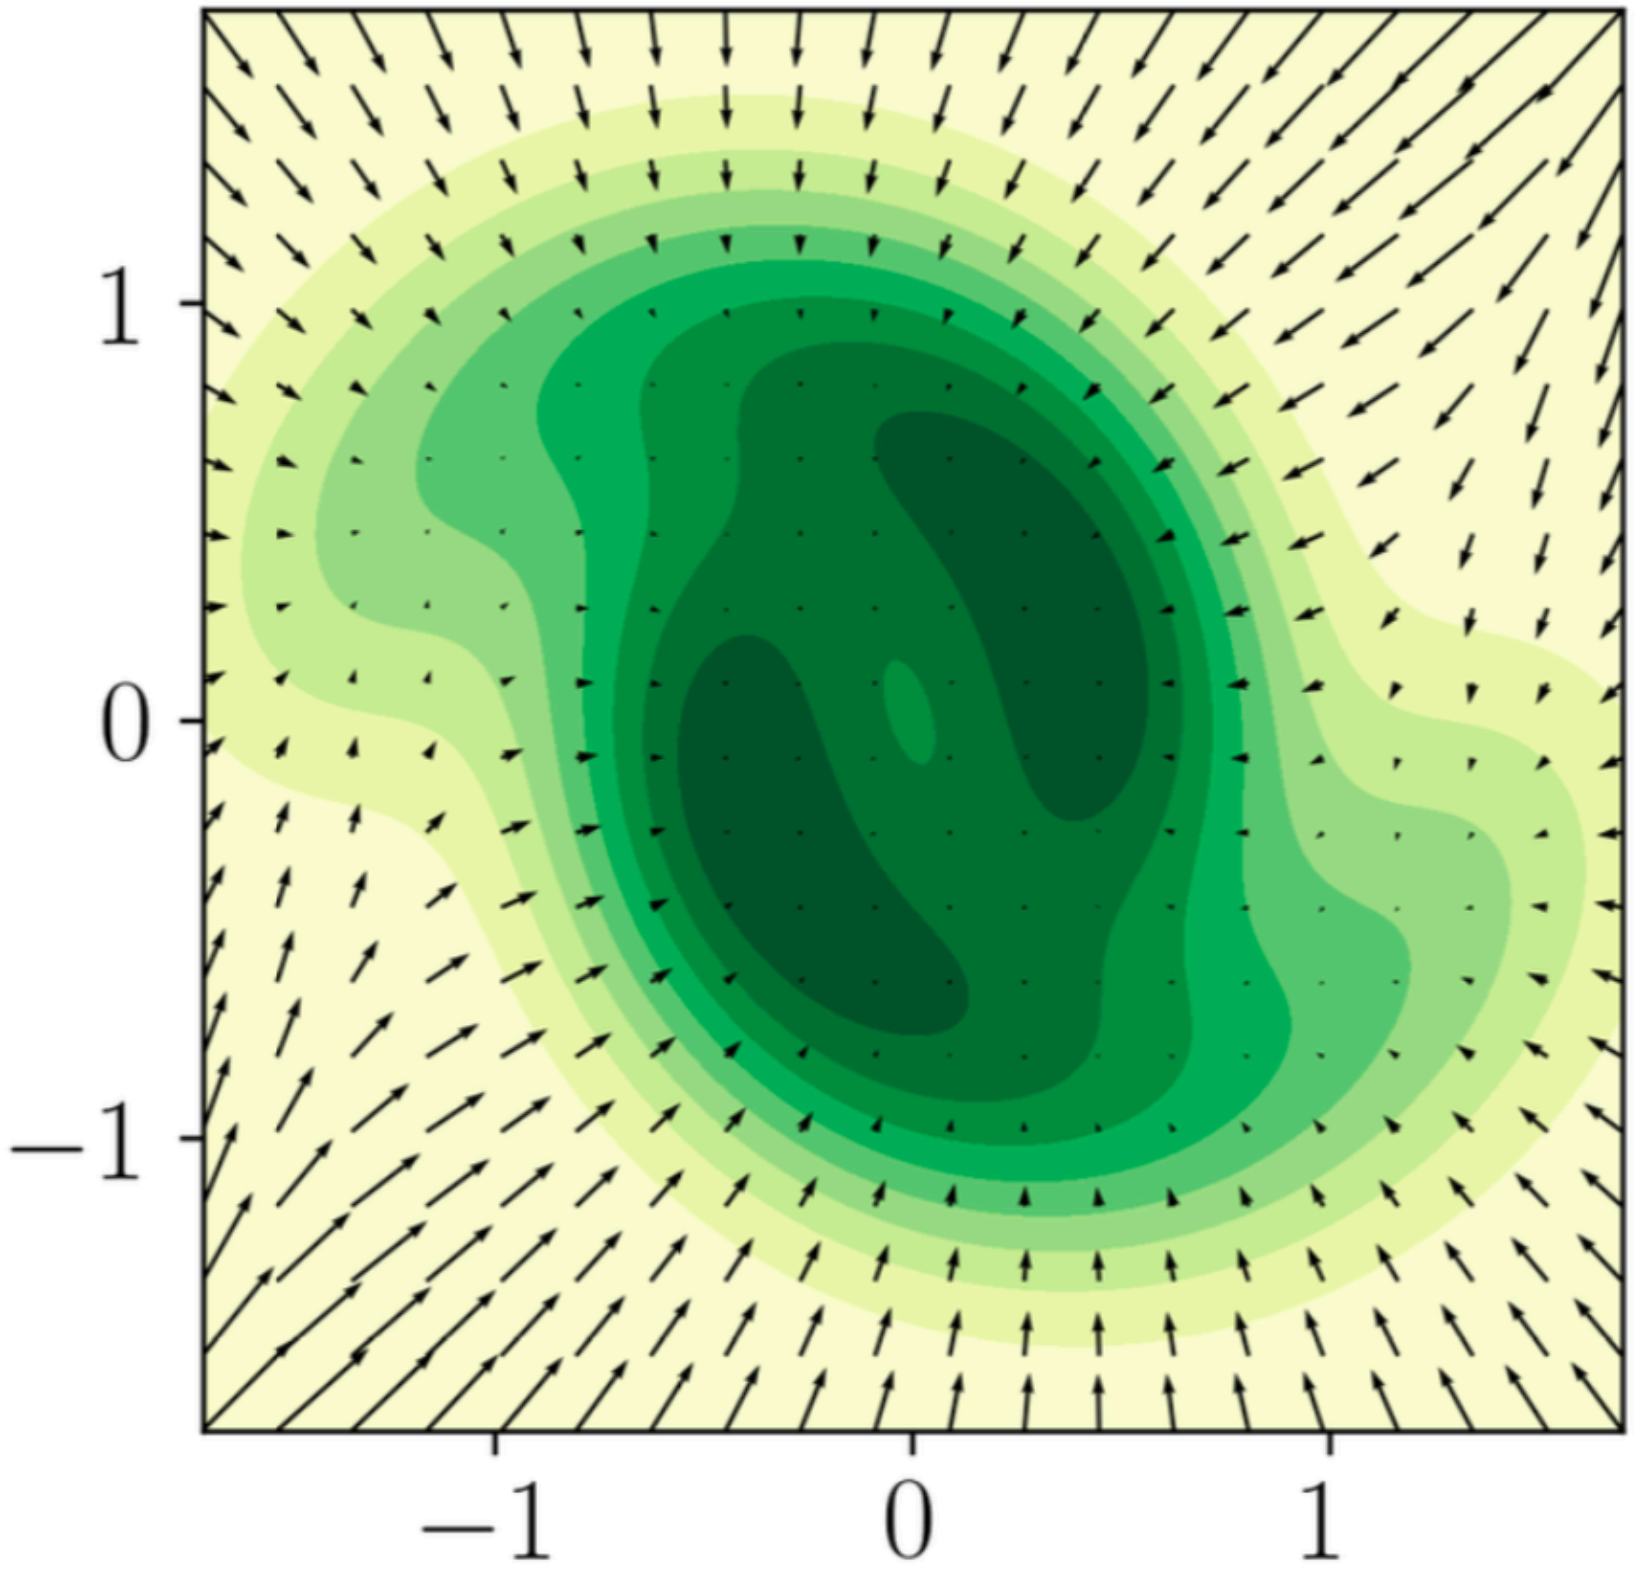

$t \approx 0$

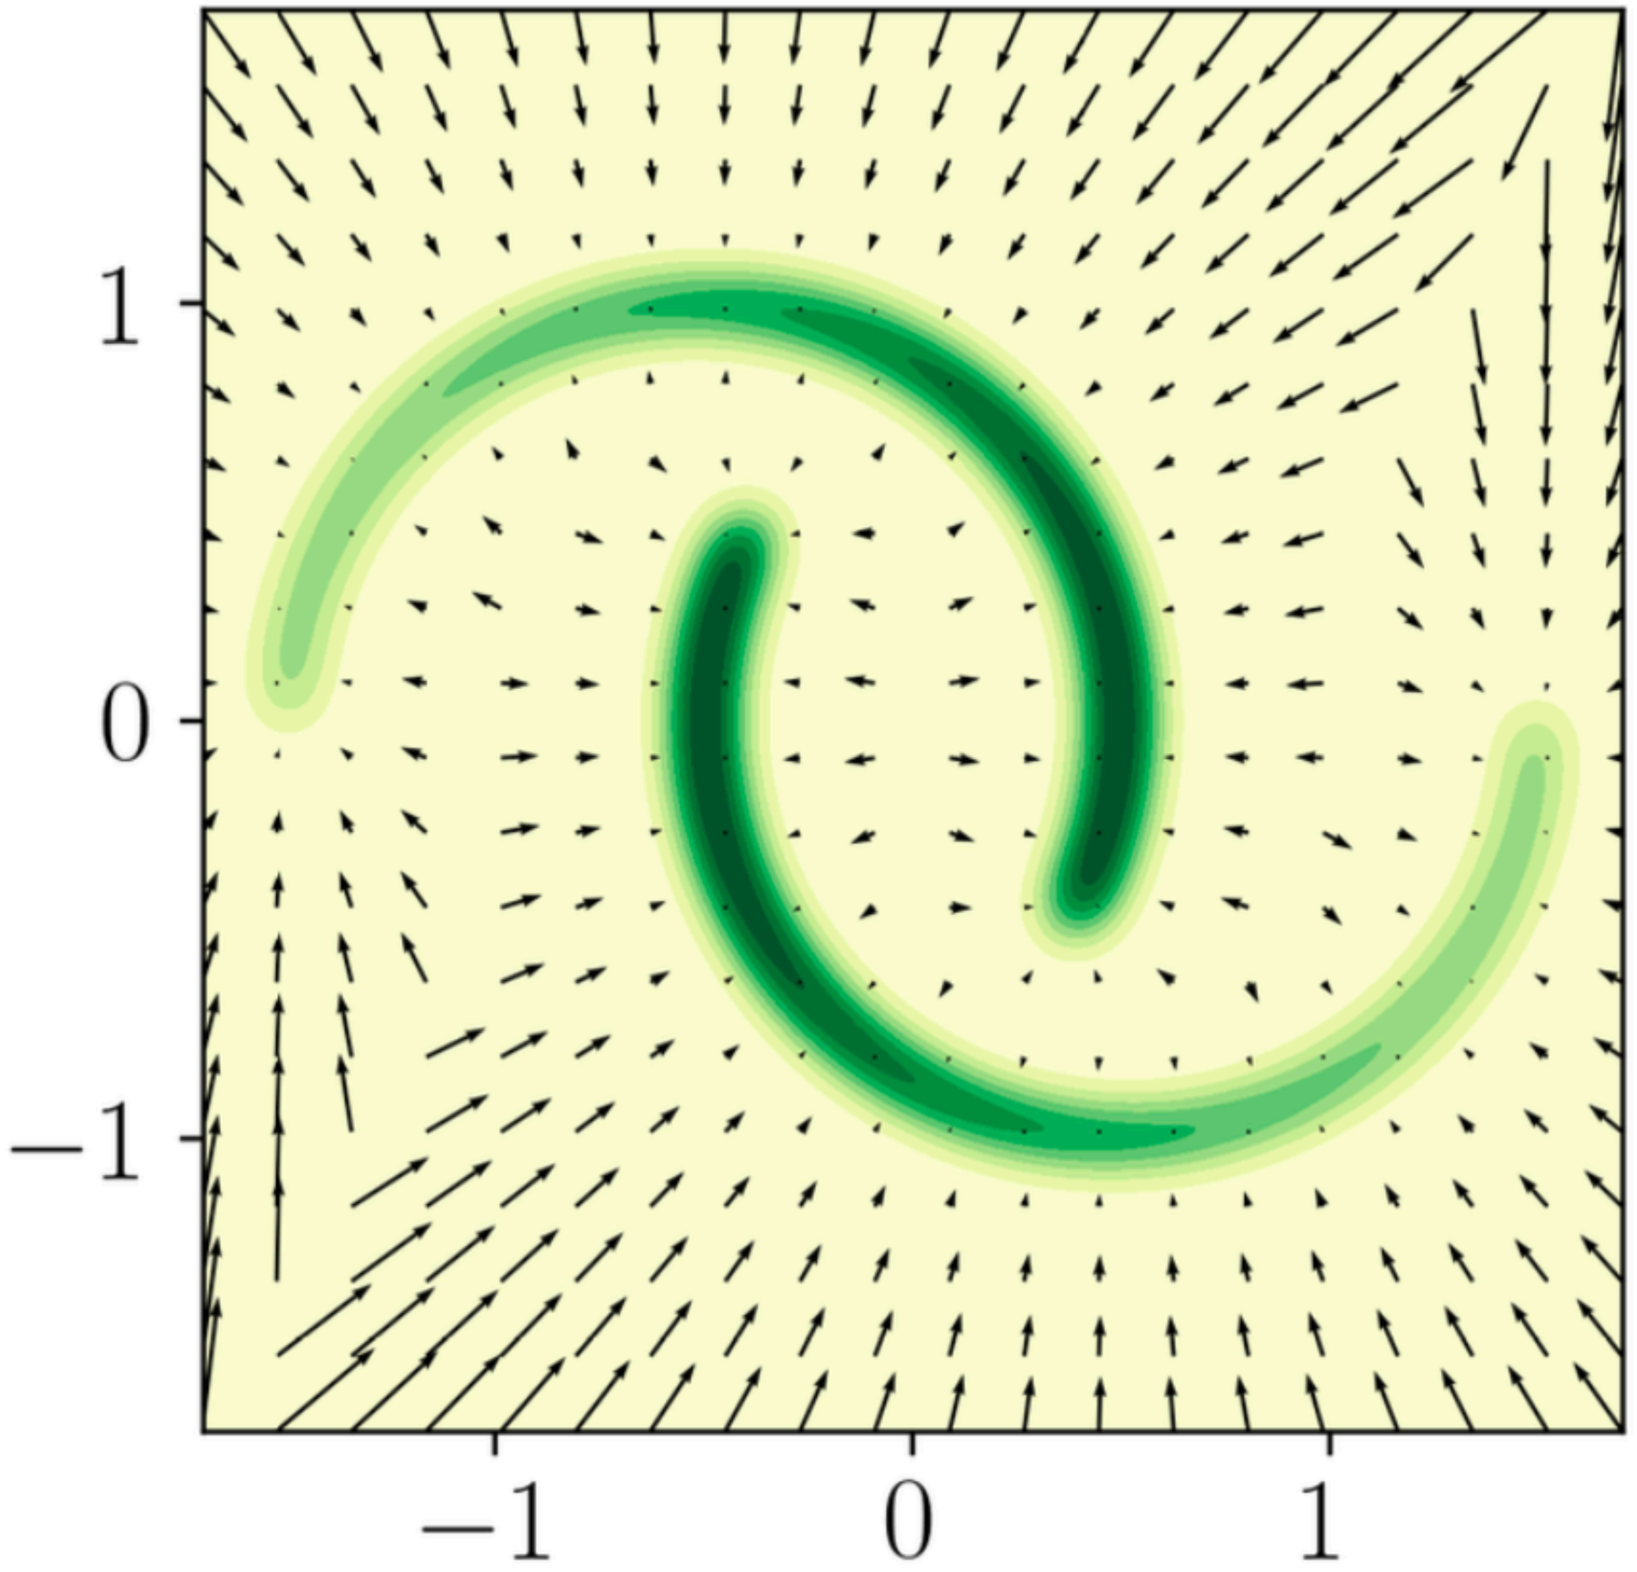

$t \approx \tau$

Sufficient conditions for alignment

Theorem (informal): If F^τ is compressive overall, v_t is uniformly compressive for t close to τ, and v_t has small cross-derivatives, alignment holds.

$$(s_{t+1}\, E_{t+1}) \circ F_t = s_t\, E_t\, R_t^{-1} - \mathrm{tr}((dF_t^{-1} d^2 F_t)\, E_t R_t^{-1}.$$

- $s_\tau \perp TM$ when target is singular
- Expansive dynamics leads to contraction above
- Compression can lead to alignment as well

Summary

- GMs can be viewed as random dynamical systems
- This perspective explains their behavior under learning errors
- **Alignment** of the d most sensitive subspaces with the tangent space of the d-dimensional data manifold leads to robustness of the support
- Aligned generative models can learn the data manifold

C and de Clercq, NeuRIPS 2025, https://arxiv.org/abs/2508.07581

Takeaways for digital twins

- Probability measures of interest to DA and DT often have a factorizable low-dimensional structure.
- We can use Oseledets theory for dynamical GMs to exploit this structure
- How to produce alignment?
- How to control GMs for rare events or regions of interest? Diffusion guidance [Ho Salimans 2022, Song Shen Xing and Ermon 2021]