Data-driven reduced order models: prediction and stabilisation

Taraneh Sayadi

M2N, CNAM Paris

Mathématiques à flots · Institut Henri Poincaré, Paris · 23 November 2026

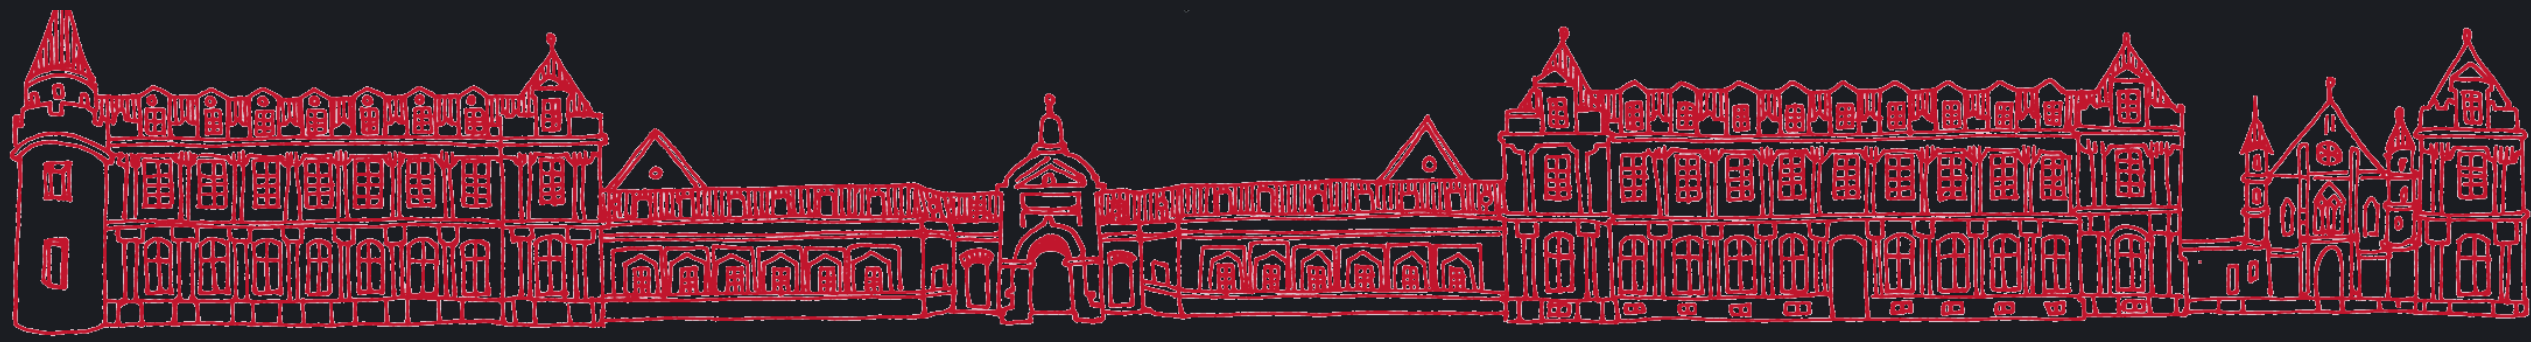

Numerical models (Digital Twin)

Fidelity

Governing equations (GE)

Approximated GE

Reduced-Order Models (ROM)

Efficiency

Ivey et al. 2011

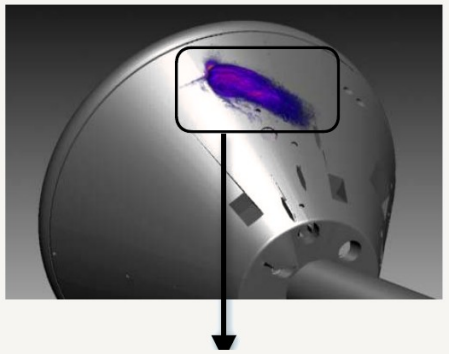

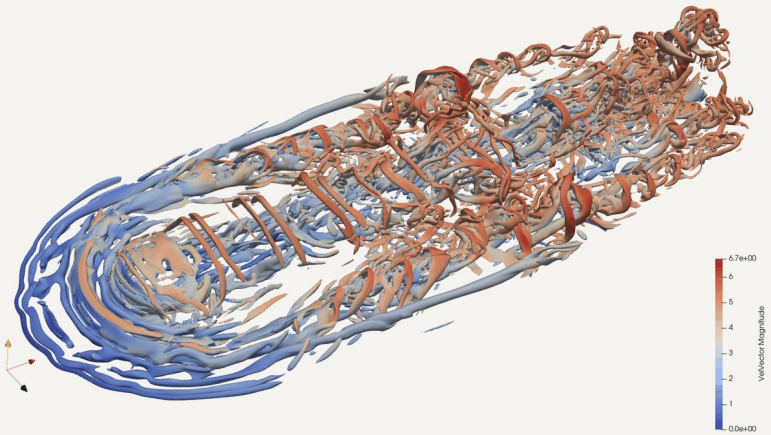

Requirements

Multi-query applications

⟶ O(100) function evaluations (**operating conditions of the system**)



ROM

Requirements:

- Cheap
- Predictive
- Robust
- Parametrised

Uncertainty-aware and parametrized dynamic reduced-order model: Application to unsteady flows

I. Zighed, N. Thome, P. Gallinari & T. Sayadi — Phys. Rev. Fluids 10, 114902 (2025)

High-dimensional Problems

Assumption:

Although the system is high-dimensional, the dynamics evolve on a lower dimensional manifold

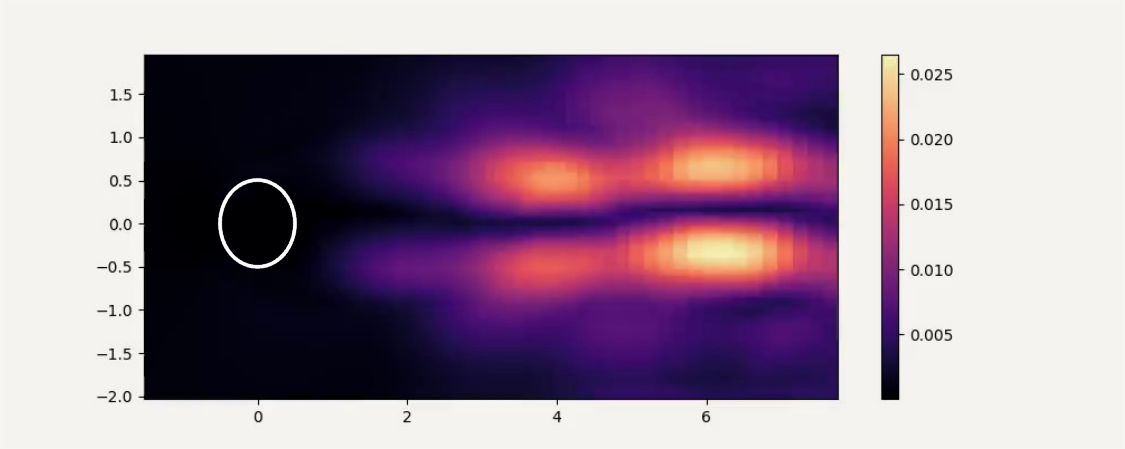

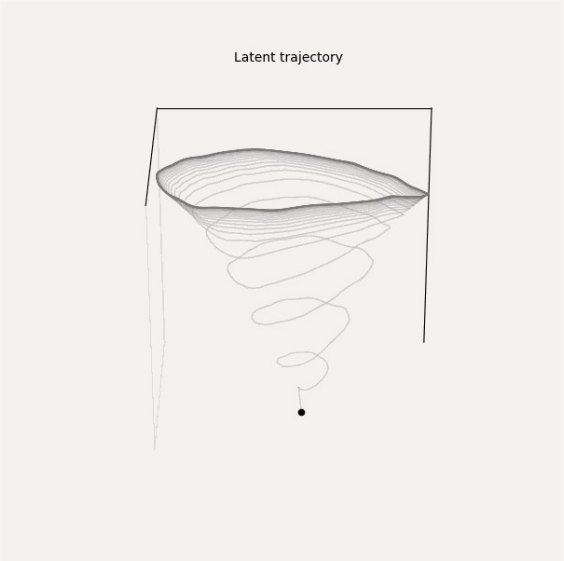

Approach

1. **Compression** ⟶ dimensionality reduction
2. **Dynamics** in the lower-dimensional manifold
3. **Decompression**

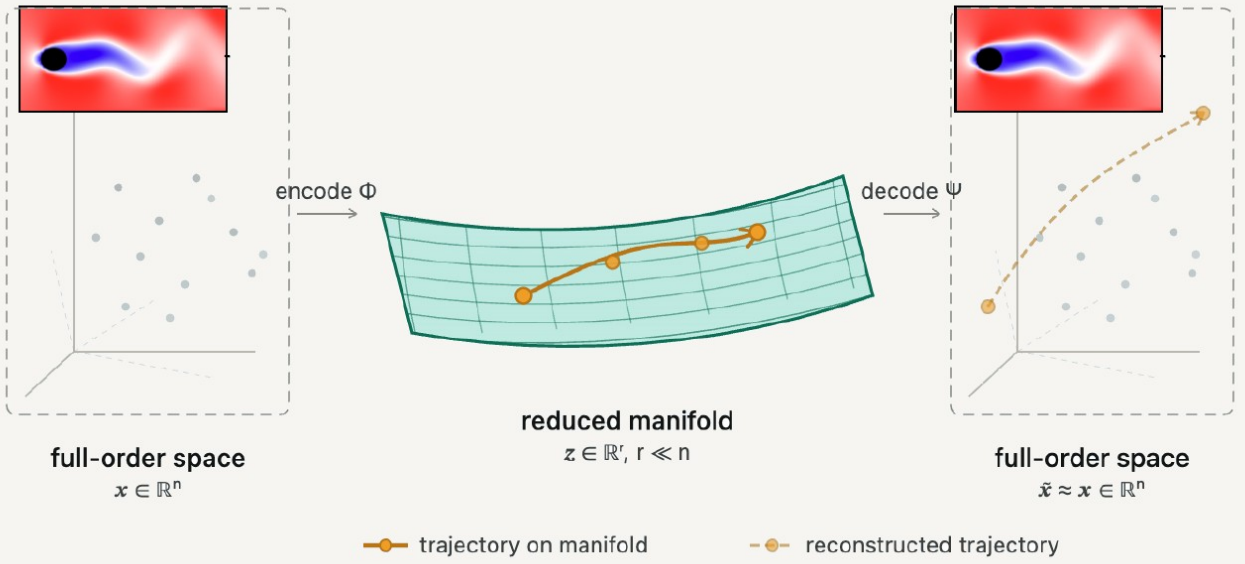

ML-based Compression

Compression/decompression

- PCA ⟶ Linear
- AE / VAE ⟶ nonlinear

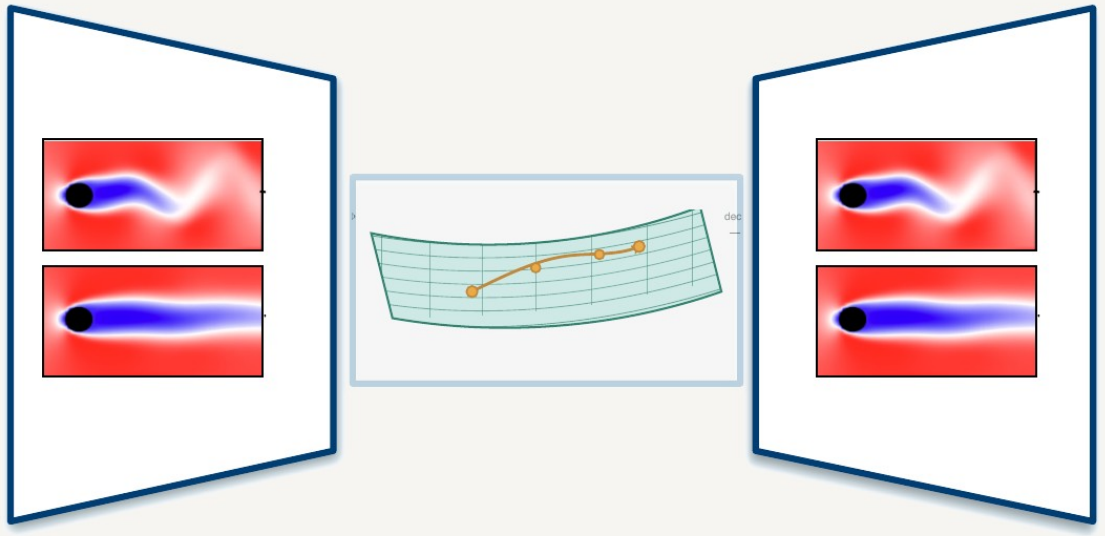

Keep the dimensions of the system (number of latent variables/observables) small!

MZ – formalism

Data-driven learning for the Mori–Zwanzig formalism: a generalization of the Koopman learning framework

Yen Ting Lin, Yifeng Tian, Marian Anghel, Daniel Livescu

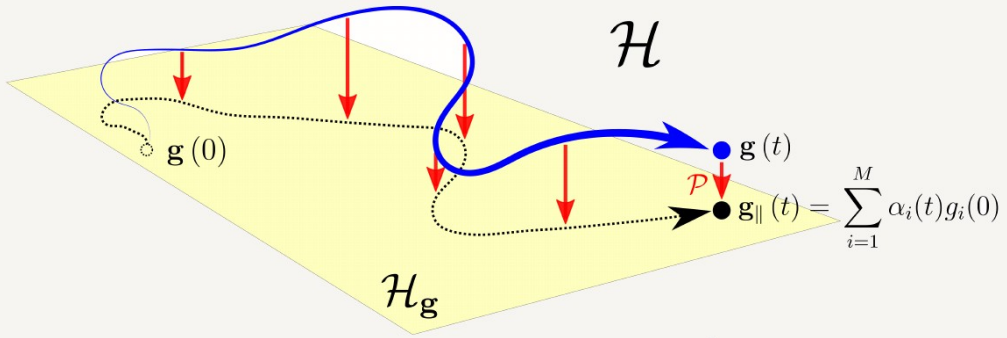

$$\frac{d\boldsymbol{\Phi}(t)}{dt} = S(\boldsymbol{\Phi}(t)), \quad \boldsymbol{\Phi}(0) = \boldsymbol{\Phi}_0,$$

$$\frac{\partial g(x,t)}{\partial t} = \mathcal{L}\{g(x,t)\} \quad , \quad g(0,t) = g_0$$

$$\frac{\partial}{\partial t}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0, t) = \underbrace{\mathcal{P}\mathcal{L}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0, t)}_{\text{Markov}} + \underbrace{\int_0^t \mathcal{P}\mathcal{L}e^{s\mathcal{Q}\mathcal{L}}\mathcal{Q}\mathcal{L}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0, t-s)ds}_{\text{Memory}} + \underbrace{e^{t\mathcal{Q}\mathcal{L}}\mathcal{Q}\mathcal{L}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0)}_{\text{Noise}}$$

Dynamic Reconstruction

Dynamics in the latent space

- Neural ODEs [1]
- ESN [2]
- LSTM [3]
- **Transformers**

Haller et al. 2022

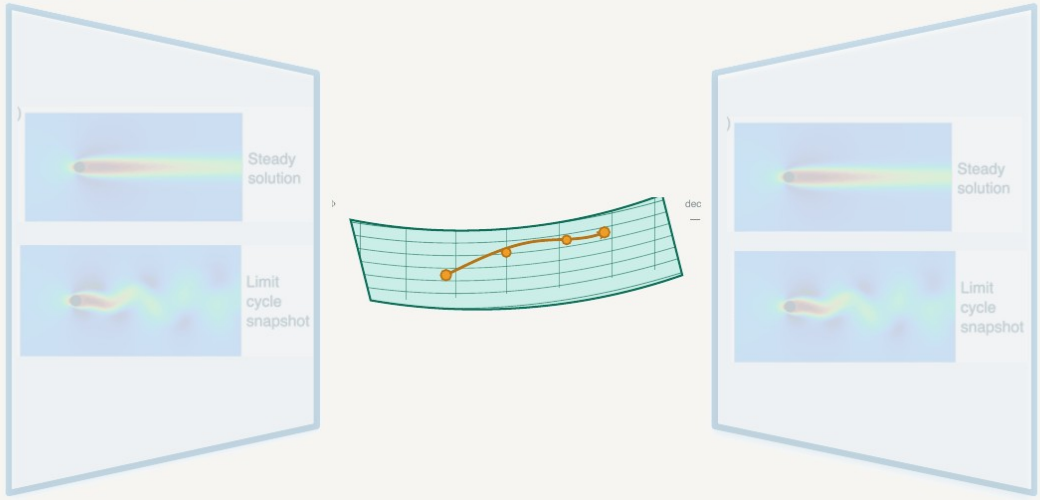

[1] E. Menier, M.A. Bucci, M. Yagoubi, L. Mathelin, M. Schoenauer, CD-ROM: Complemented Deep – Reduced order model, Computer Methods in Applied Mechanics and Engineering 410 (2023).
[2] A. Racca, N.A.K. Doan, L. Magri, Predicting turbulent dynamics with the convolutional autoencoder echo state network, Journal of Fluid Mechanics 975:A2, doi:10.1017/jfm.2023.716 (2023).
[3] P. Gupta, P.J. Schmid, D. Sipp, T. Sayadi & G. Rigas, Mori–Zwanzig latent space Koopman closure for nonlinear autoencoder, Proc. R. Soc. A 481, doi.org/10.1098/rspa.2024.0259 (2025).

MZ formalism – ML architecture

Observable space : $\hat{\mathbf{g}}(\boldsymbol{\Phi}_0, t)$

$$\frac{d\boldsymbol{\Phi}(t)}{dt} = S(\boldsymbol{\Phi}(t)), \quad \boldsymbol{\Phi}(0) = \boldsymbol{\Phi}_0,$$

$$\frac{\partial}{\partial t}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0, t) = \underbrace{\mathcal{P}\mathcal{L}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0, t)}_{\text{Markov}} + \underbrace{\int_0^t \mathcal{P}\mathcal{L}e^{s\mathcal{Q}\mathcal{L}}\mathcal{Q}\mathcal{L}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0, t-s)ds}_{} + \underbrace{e^{t\mathcal{Q}\mathcal{L}}\mathcal{Q}\mathcal{L}\hat{\mathbf{g}}(\boldsymbol{\Phi}_0)}_{\text{Noise}}$$

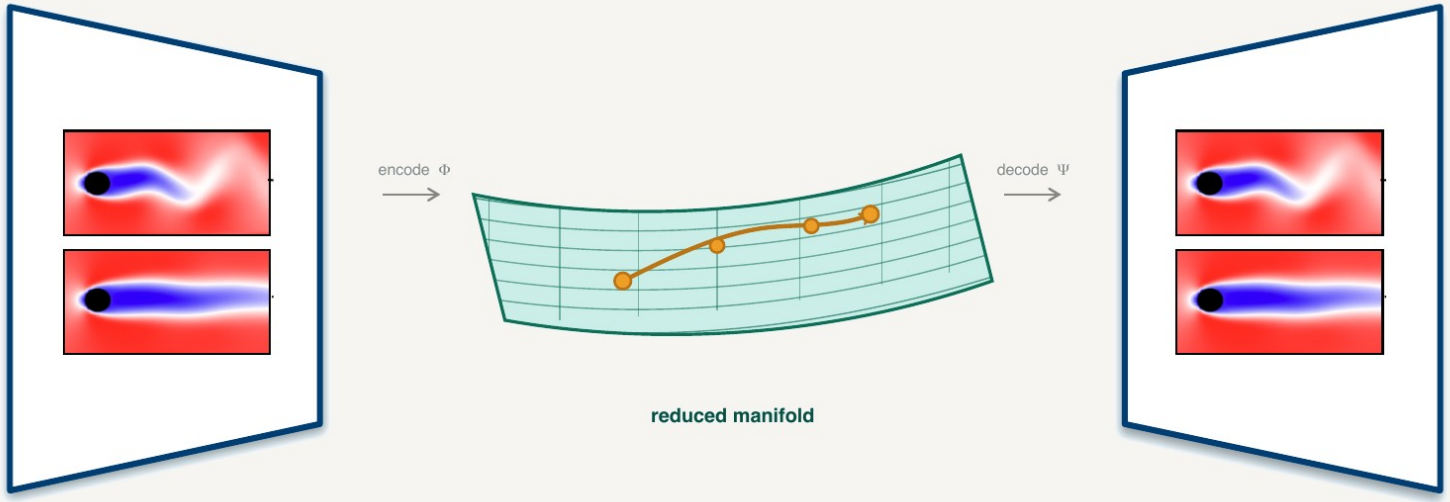

Dynamical system: flow around a cylinder



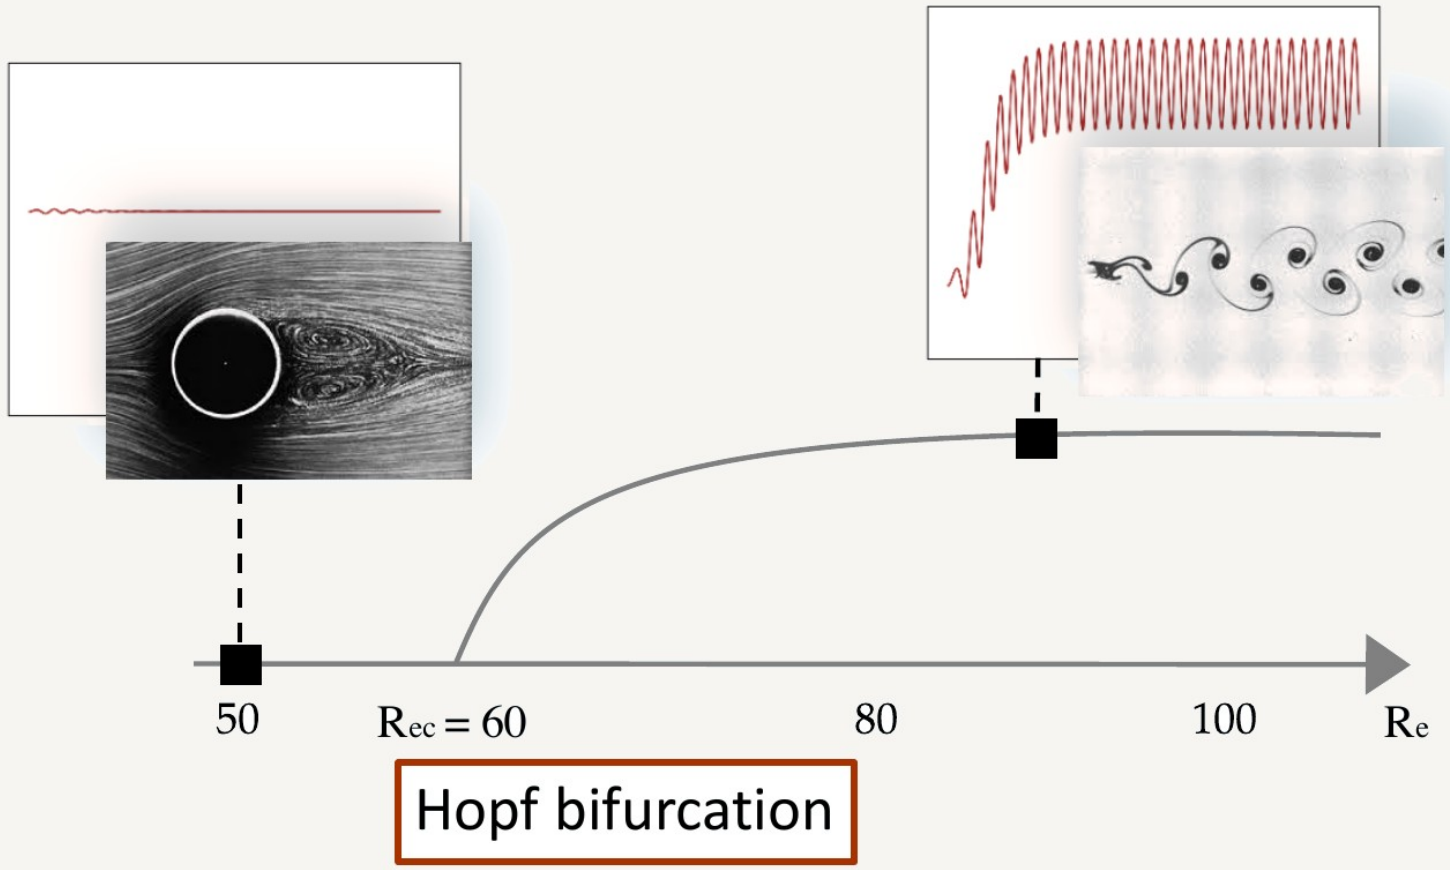
50
$R_{ec} = 60$
80
100
R_e
Hopf bifurcation

Model prediction – out of distribution

- In-distribution works well

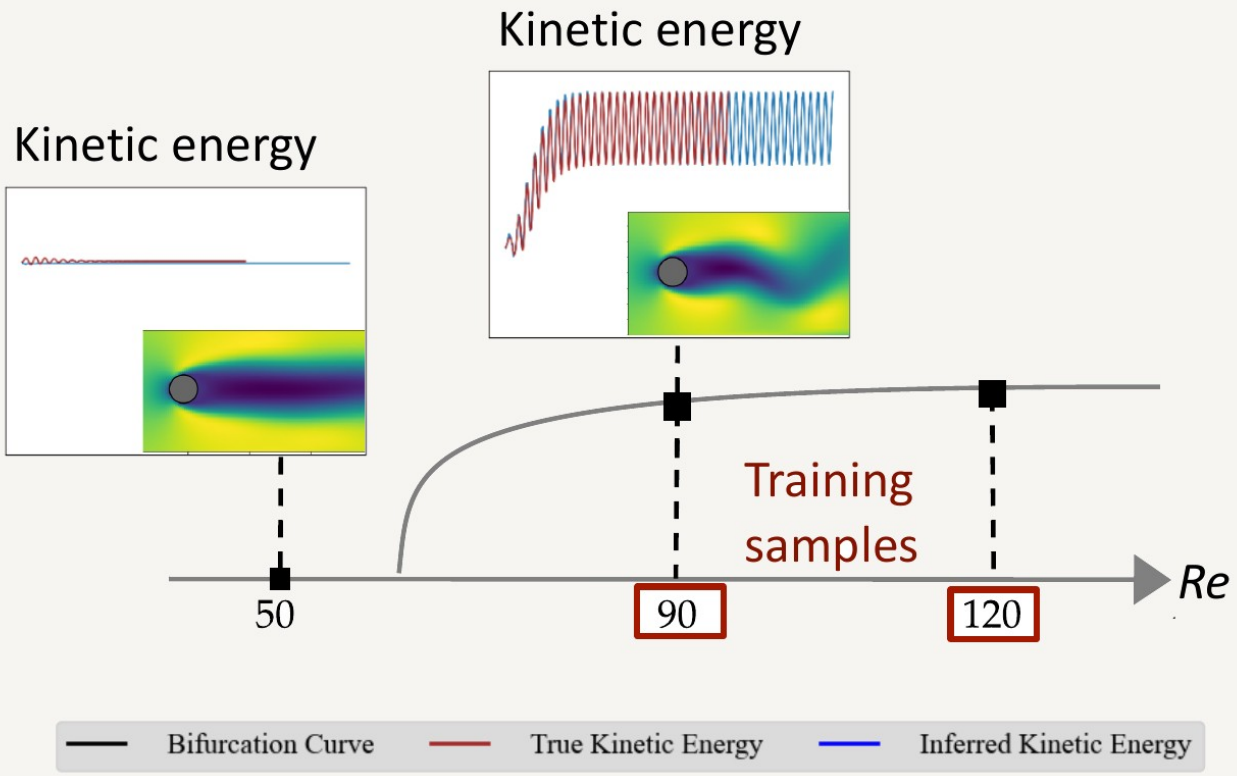

What about chaotic systems?

Kolmogorov flow

- NS equation
- Re = 90

$$\begin{cases} \dfrac{\partial \mathrm{u}}{\partial t} + (\mathrm{u}\cdot\nabla)\mathrm{u} = -\nabla p + \nu\nabla^2\mathrm{u} + \mathrm{f}, \quad \text{with } \mathrm{f} = \gamma \sin\left(n_f \dfrac{2\pi y}{L}\right) \mathrm{e}_x \\ \nabla\cdot\mathrm{u} = 0 \end{cases}$$

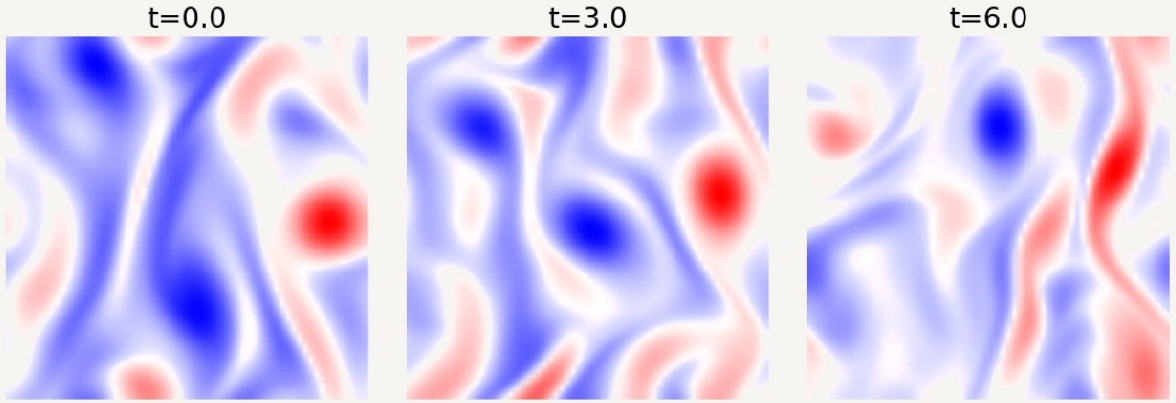

Figure 1: Vorticity snapshots

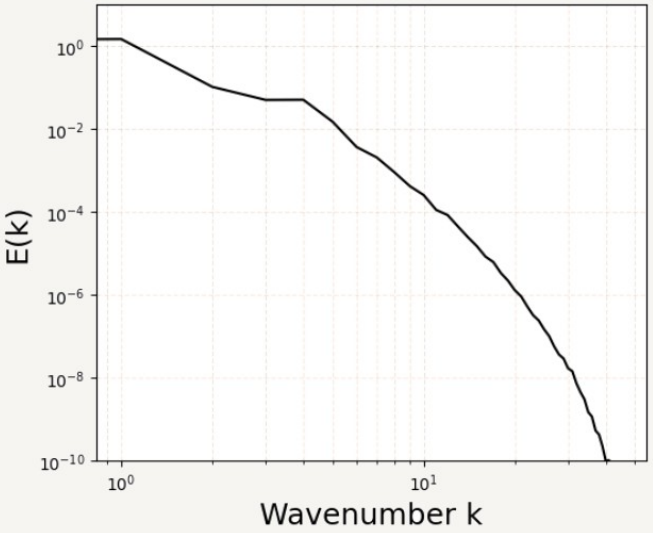

Figure 2: Energy spectrum

Modeling

- Nonlinear dynamical system

$$\frac{dx_t}{dt} = \mathcal{F}(x_t)$$

- Discretisation in time

$$x_{t+1} = x_t + \Delta t\, \mathcal{F}(x_t)$$

- Replace the operator with a model (emulator)

$$x_{t+1} = x_t + \Phi_\theta(x_t)$$

Fourier Neural Operator (FNO)

Long-time prediction

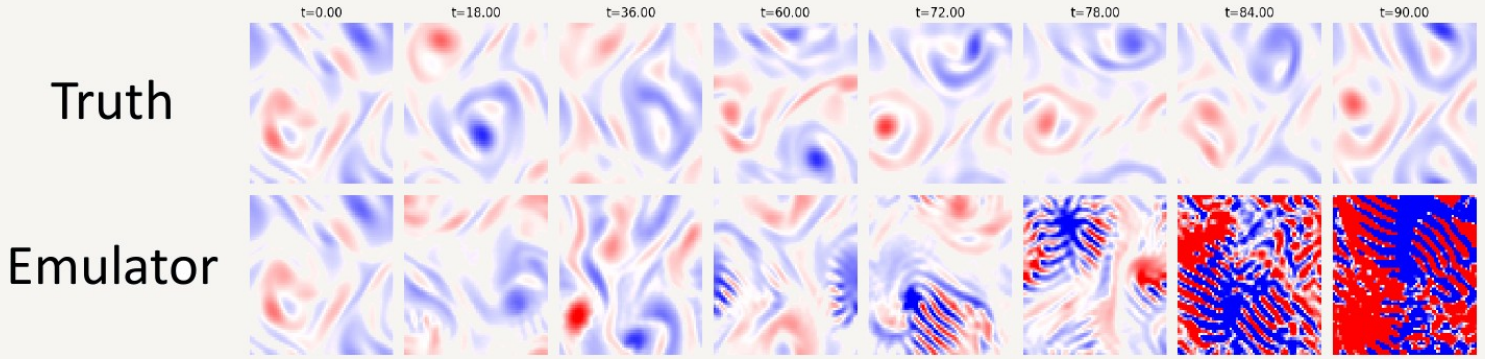

Blow up !

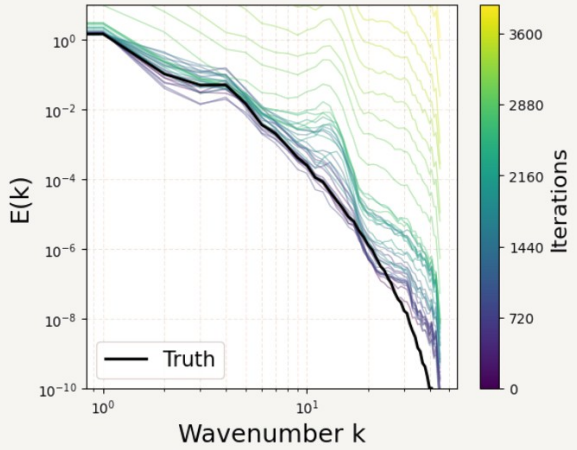

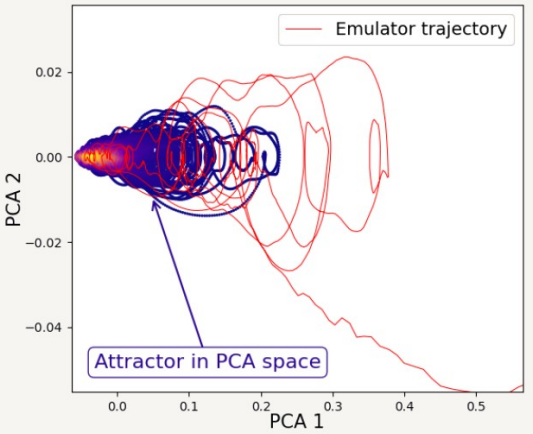

Two solutions I propose here

- Stabilisation of the trajectories
- Unroll probabilistically rather than deterministically

Stochastic representation

Give up on point by point prediction !

What is the ultimate goal of these models ?

- Image reconstruction ?
- Acceptable dynamical/statistical behaviour

Stabilisation

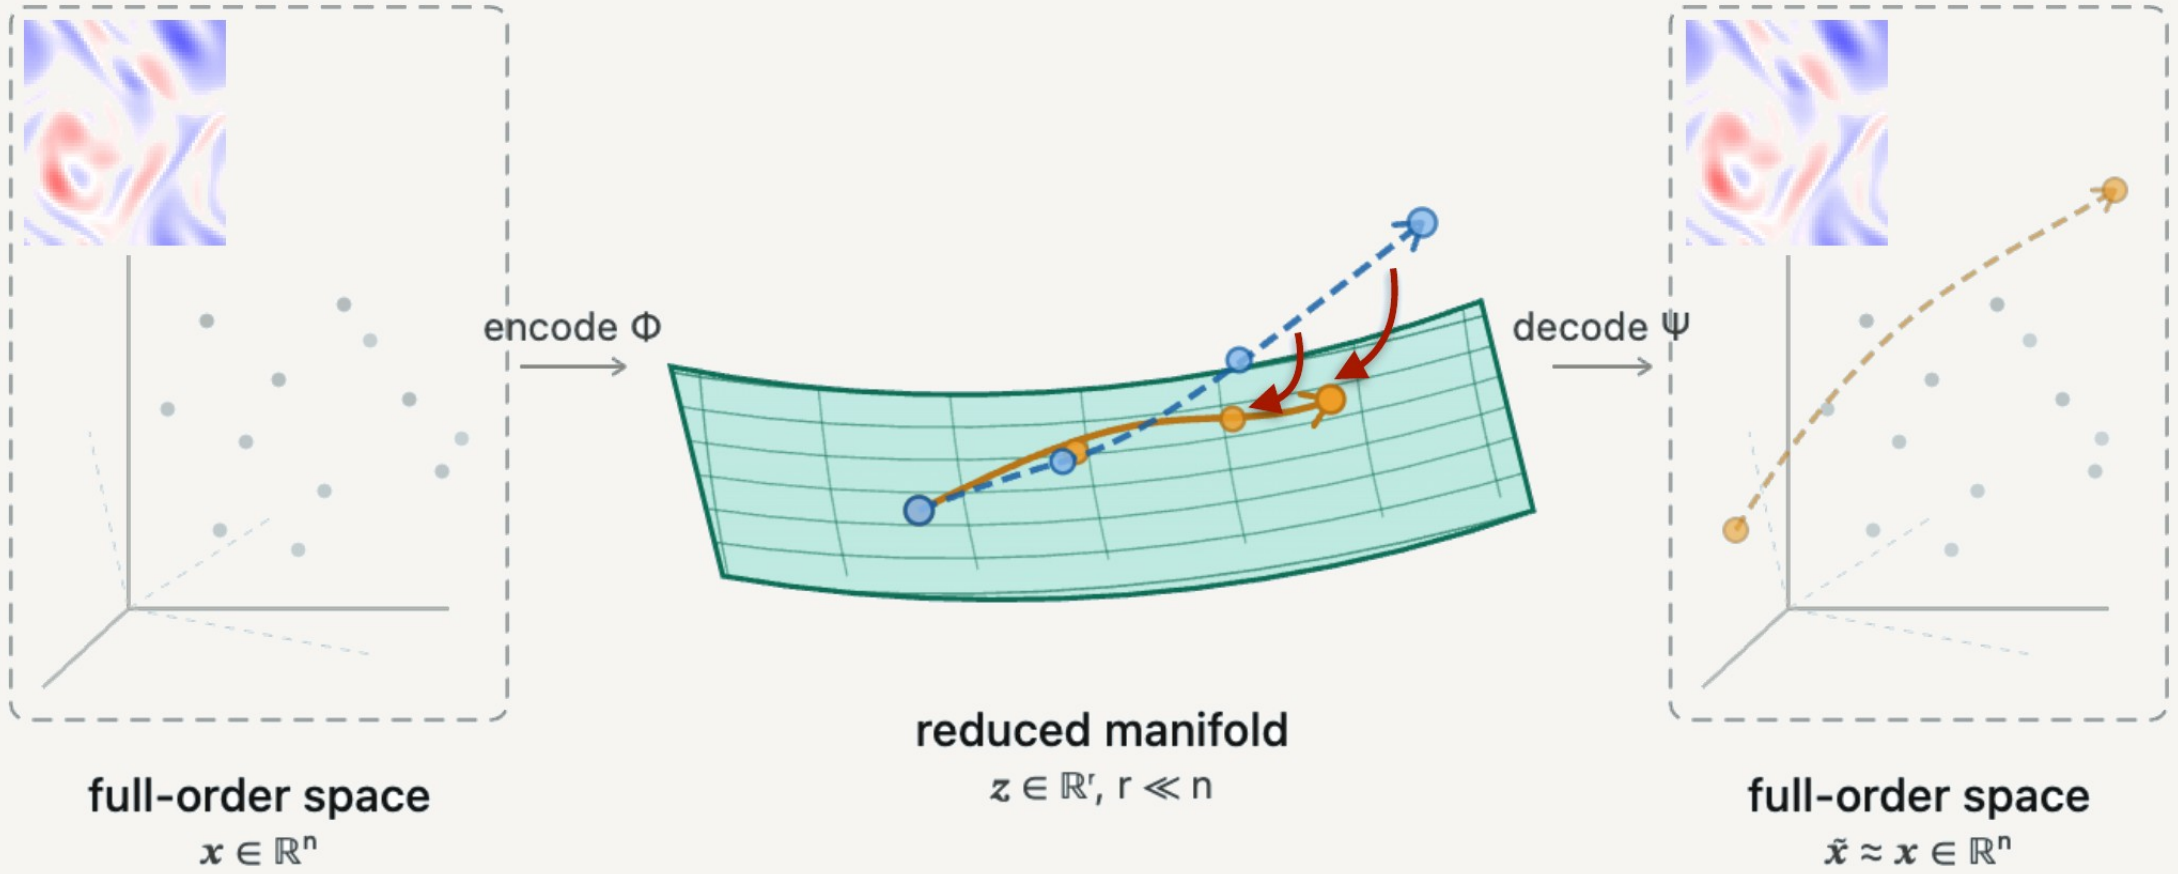

Probabilistic formulation

- Probabilistic Markov formulation

$$\pi(x_{t+1} \mid x_t)$$

- Gaussian transition kernel

$$\pi(x_{t+1} \mid x_t) = \mathcal{N}\left(x_t + \Phi_\theta(x_t),\ \tau^2 I\right)$$

- Stochastic dynamical system

$$x_{t+1} = x_t + \Phi_\theta(x_t) + \tau\epsilon$$

$$\epsilon \sim \mathcal{N}(0, I)$$

Thermalizer: Stable autoregressive neural emulation of spatiotemporal chaos
Christian Pedersen, Laure Zanna, Joan Bruna

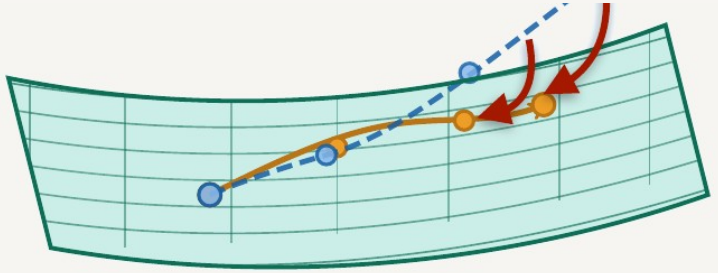

How to fix the distribution ?

- Denoising Diffusion Probabilistic Model (DDPM, Ho *et al.* 2020)

Forward ←

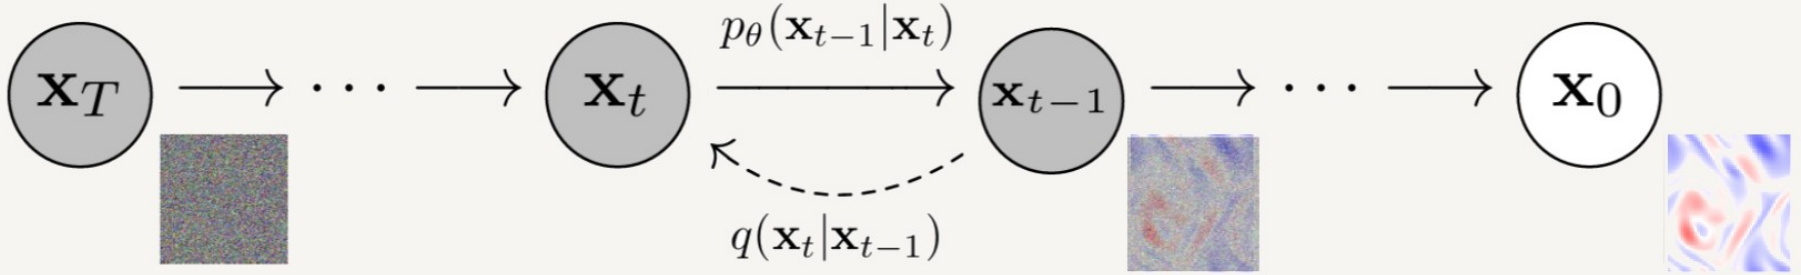

→ **Backward**

Diffusion strategy: destroy structure with noise (easy, fixed process), learn to invert it (hard, learned process)

Forward Process

Forward ←

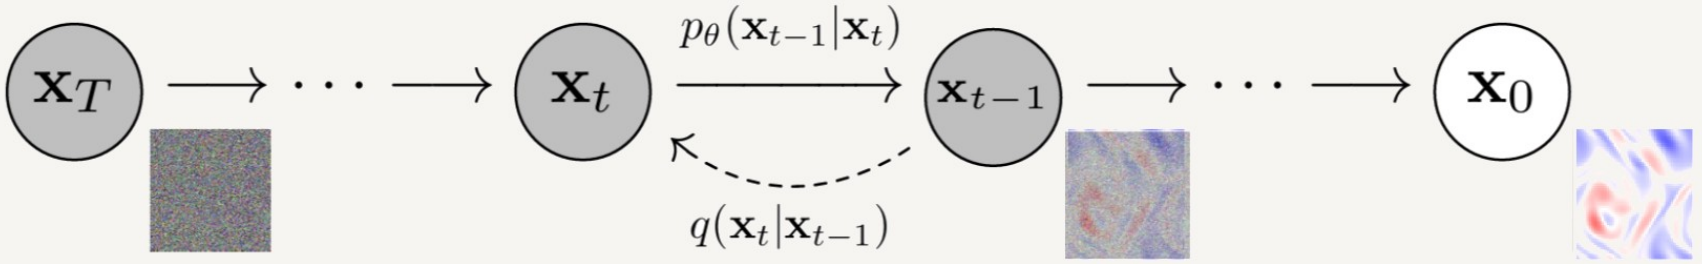

General form

$$dx = f(x,t)\,dt + g(t)\,dW$$

DDPM: variance-preserving SDE

$$dx = -\tfrac{1}{2}\beta(t)\,x\,dt + \sqrt{\beta(t)}\,dW$$

Linear (Ornstein–Uhlenbeck-type) SDE → the transition kernel stays Gaussian, in closed form:

$$x_t \mid x_0 \;\sim\; \mathcal{N}\big(\sqrt{\bar{\alpha}_t}\,x_0,\; (1-\bar{\alpha}_t)I\big)$$

Reverse Time Process

Backward

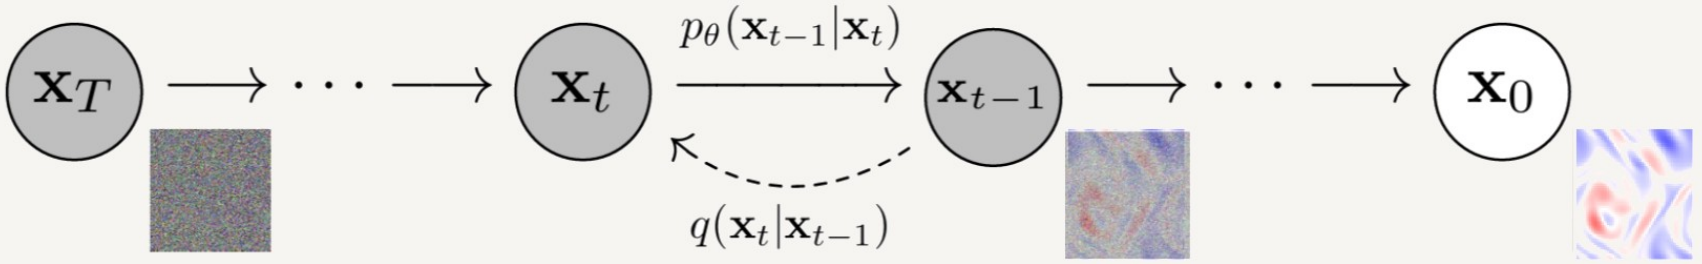

Any forward diffusion has a reverse-time diffusion that recovers the data distribution from noise:

$$dx = \left[f(x,t) - g(t)^2 \nabla_x \log p_t(x) \right] dt + g(t)\, d\bar{W}$$

Everything about generation reduces to knowing one quantity:

$$\nabla_x \log p_t(x)$$

the score of the noised data distribution

How to fix the distribution ?

- Denoising Diffusion Probabilistic Model (DDPM, Ho *et al.* 2020)

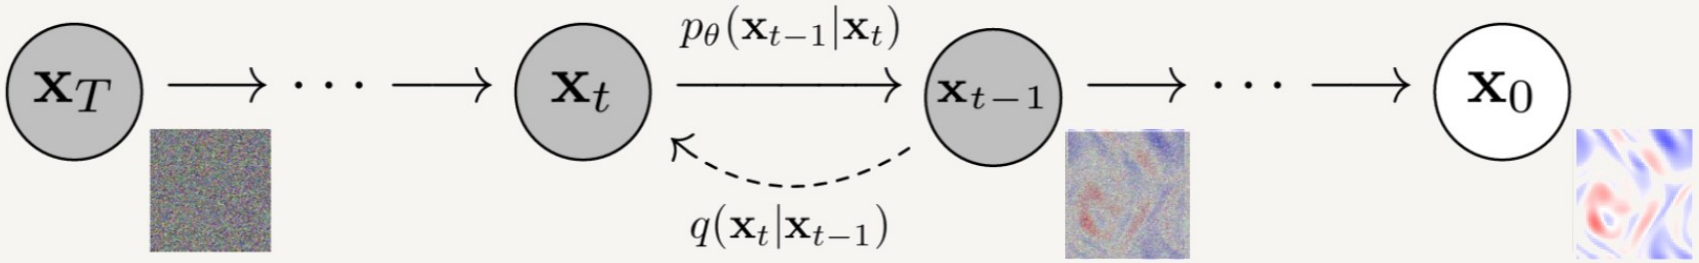

- Training procedure
- During inference we need to evaluate how “far” the trajectory is from the center
- Add a **classifier** to the DDPM to evaluate the noise level

Reconstruction

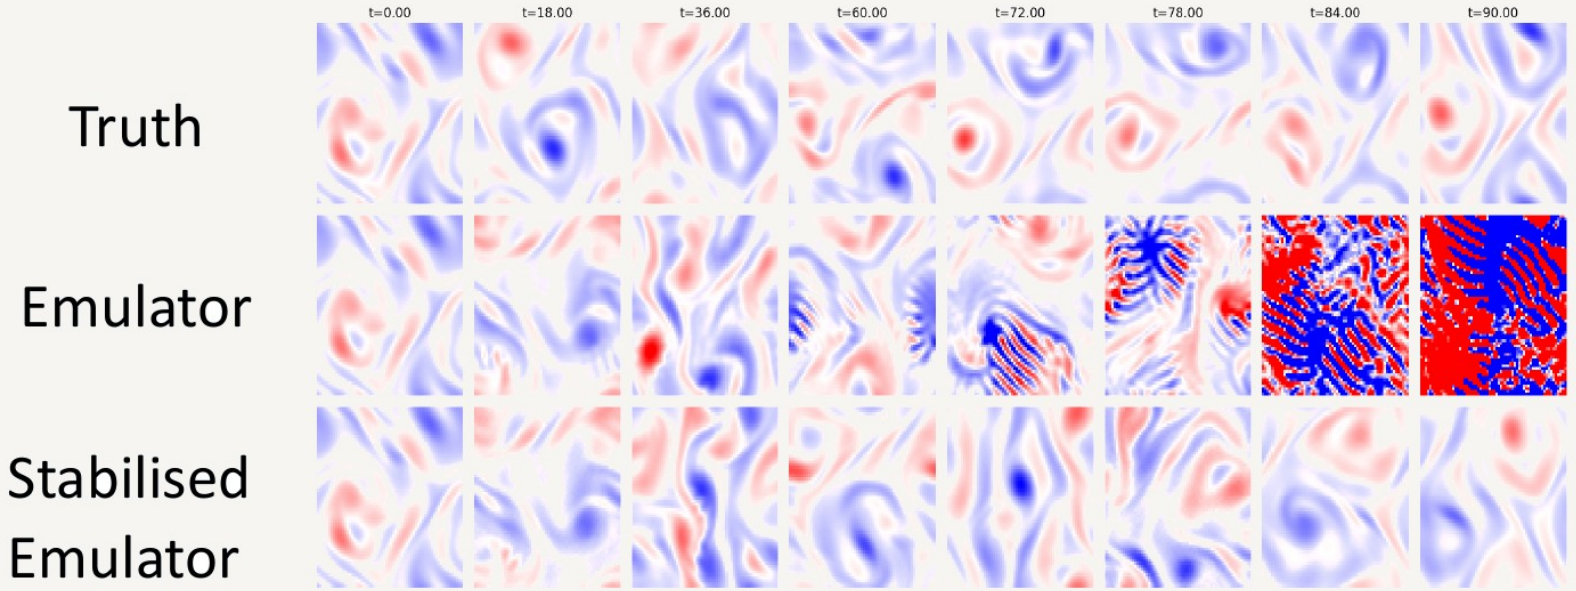

Stable for very long time roll-out

Point-wise prediction does not match, **statistical measures** do !

Spectrum

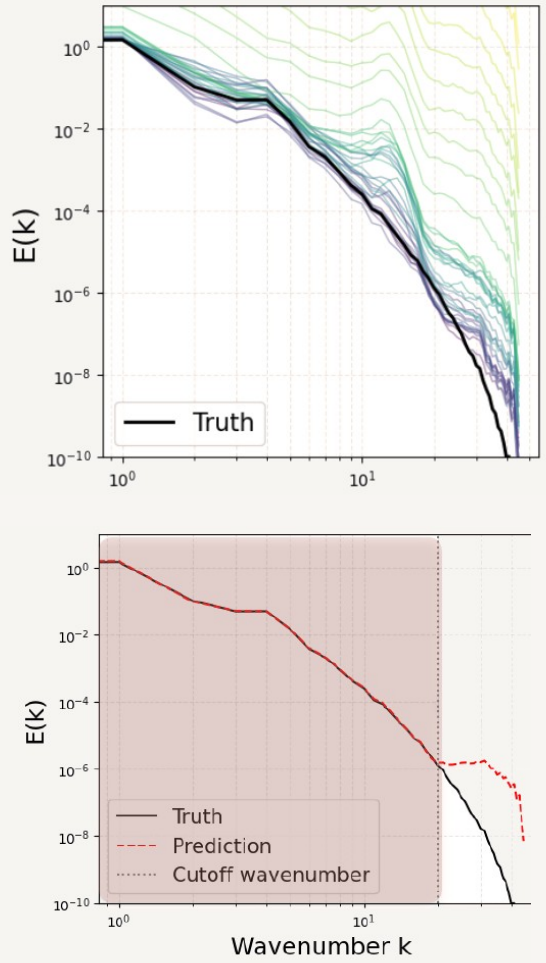

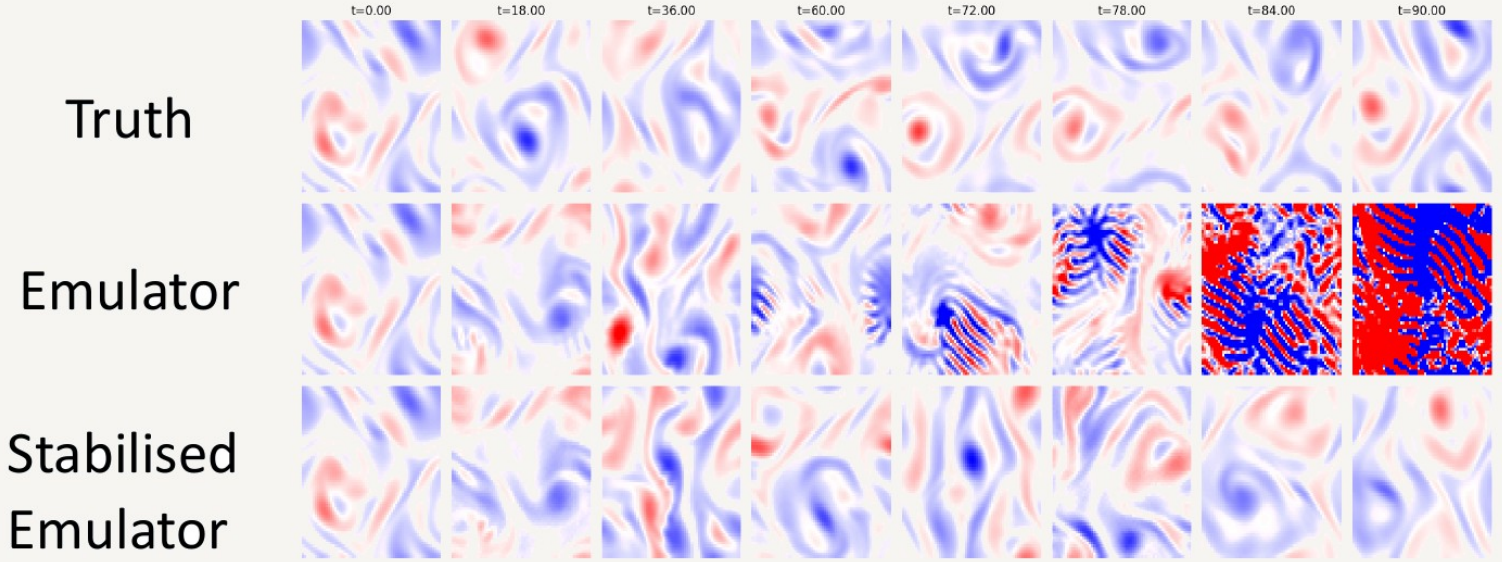

Limit of predictability of the original emulator

Attractor

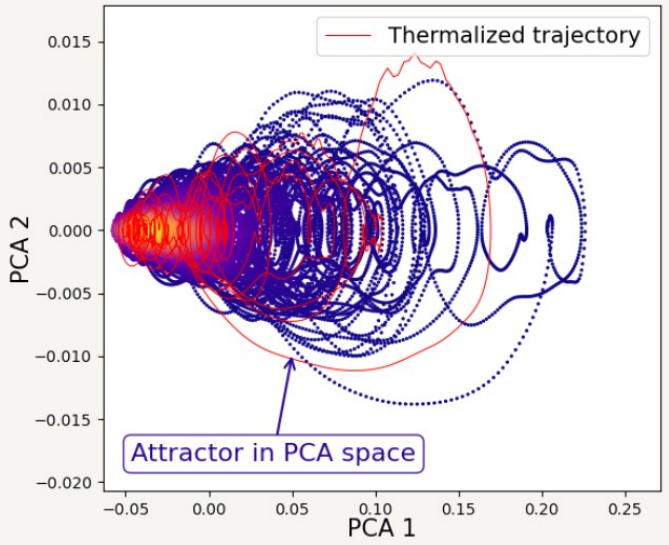
Thermalized trajectory
Attractor in PCA space
PCA 1
PCA 2
0.015
0.010
0.005
0.000
−0.005
−0.010
−0.015
−0.020
−0.05
0.00
0.05
0.10
0.15
0.20
0.25

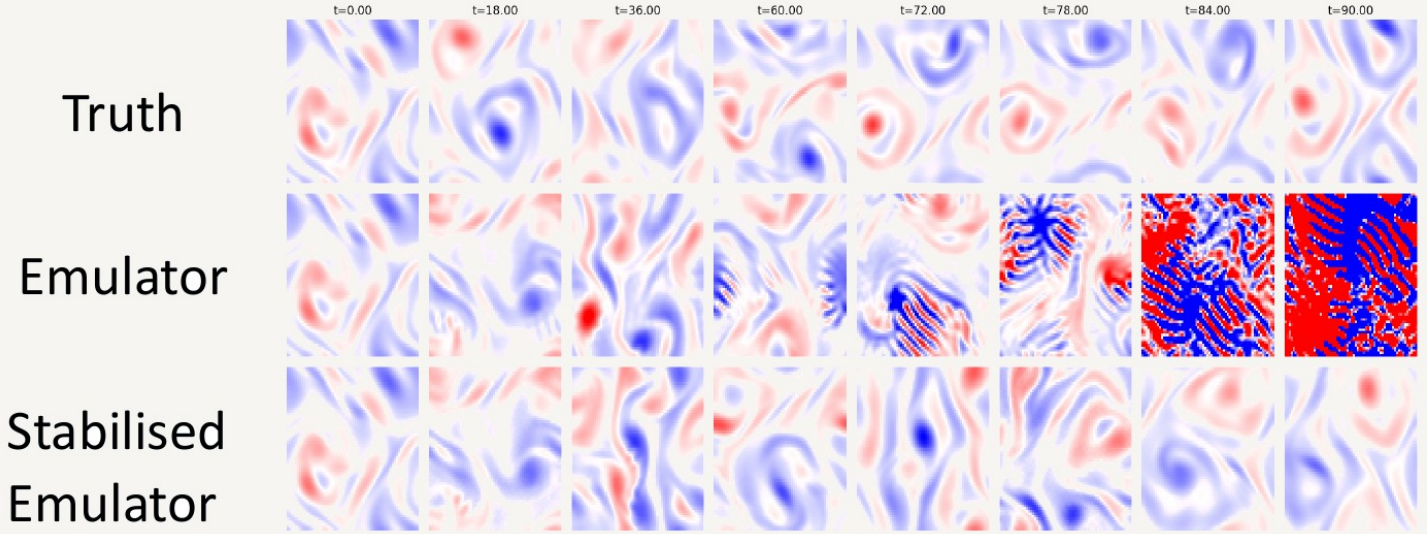
Truth
Emulator
Stabilised Emulator
t=0.00
t=18.00
t=36.00
t=60.00
t=72.00
t=78.00
t=84.00
t=90.00

Conclusions

- Moving to a **stochastic perspective** stabilises the model
- Trajectories cannot be predicted — the **statistics are correct**

Point-by-point prediction is given up; what the model preserves is the distribution the dynamics live on.

Stable for very long roll-out

Spectrum recovered

Attractor recovered

Parametrisation

- Parametrise **both** the model and the diffusion
- One emulator across a range of ***Re***, not one per operating point
- Conditioning mechanism decides how far it extrapolates

Trained at $Re \in \{30, 80\}$; error stays low between the two, and grows outside

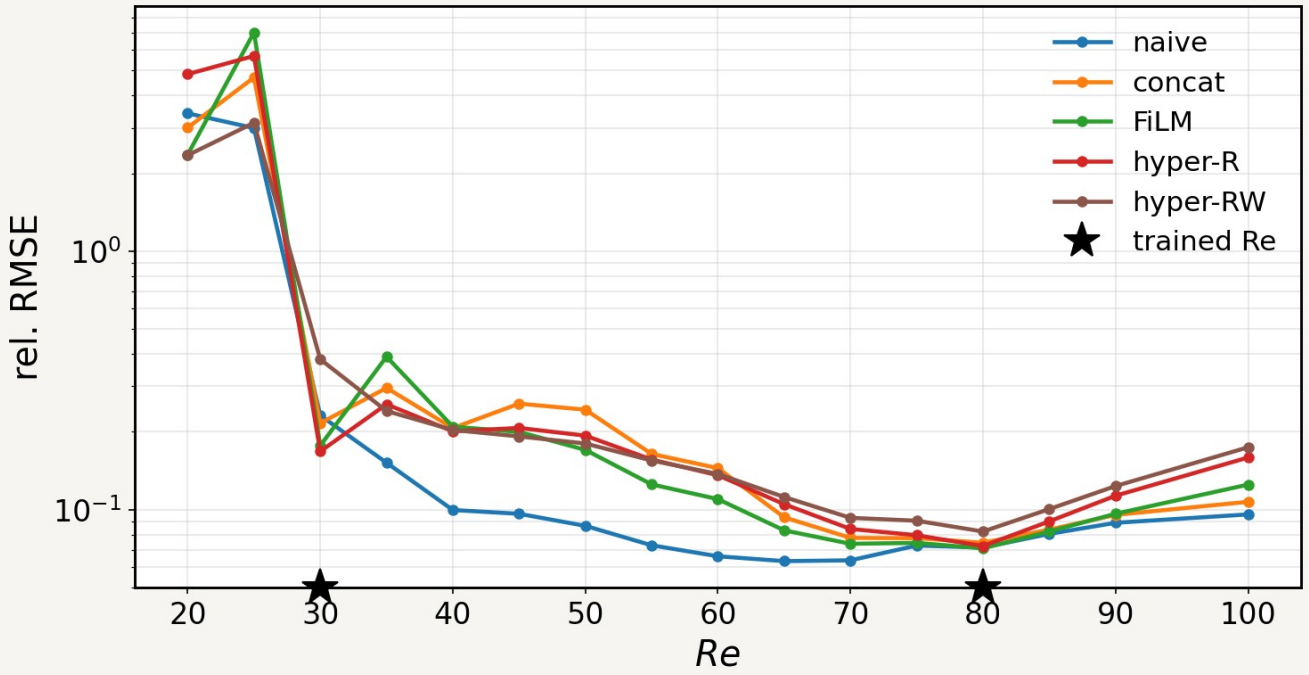

Predicting stochastically in the latent space

- A CNN **encoder** compresses each flow snapshot to a latent state
- A conditional **diffusion transformer** samples the next latent state from the last L
- A CNN **decoder** maps it back to physical space

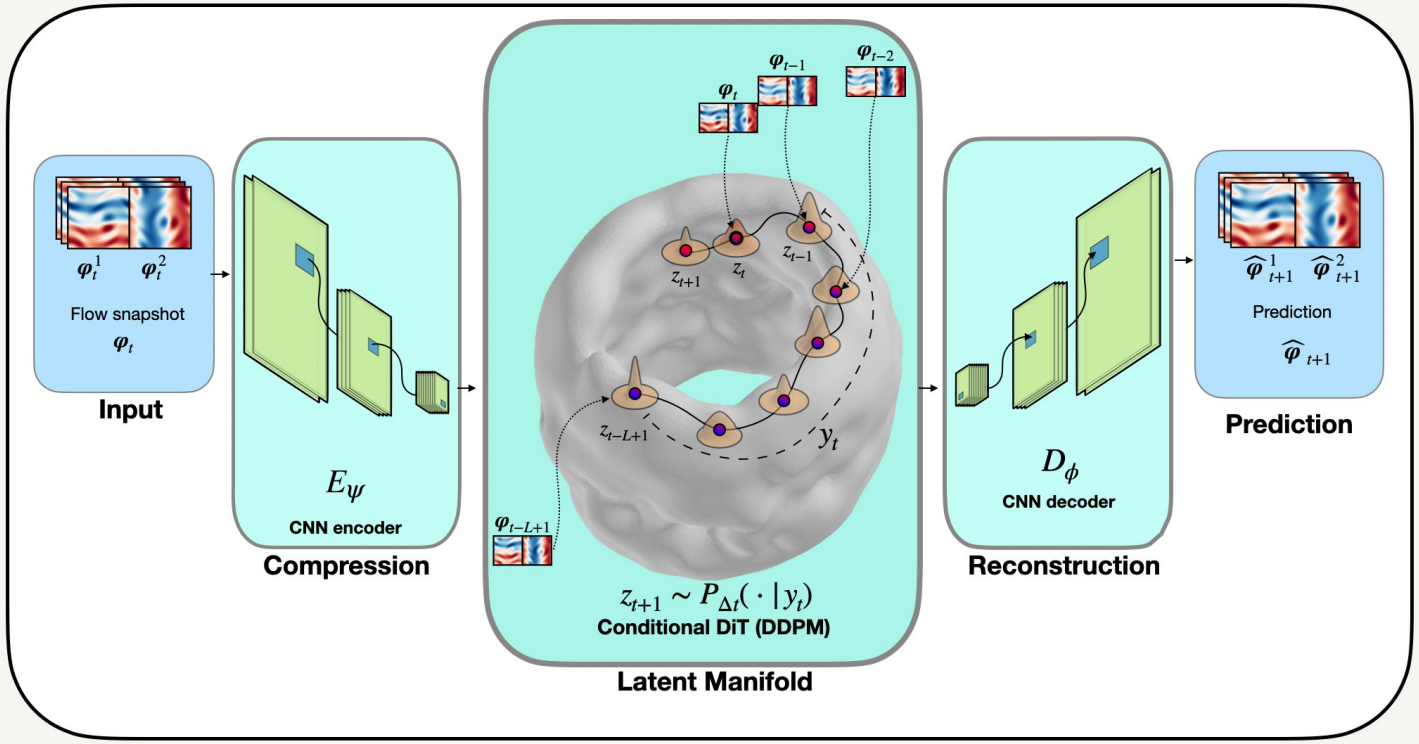

K. Taneja, F. Gonzalez, I. Mortazavi & T. Sayadi — *Latent diffusion transformers for stable, statistically faithful reduced-order modelling of chaotic flows*, in preparation (2026)