Learning Controlled Stochastic Differential Equations

Luc Brogat-Motte[1], Riccardo Bonalli[2], Alessandro Rudi[3]

[1] Istituto Italiano di Tecnologia
[2] L2S, CentraleSupélec
[3] SDA Bocconi

WCCM-ECCOMAS, July 2026

UniGe | MaLGa

Problem setting

Controlled Stochastic Dynamical Systems

Dynamical systems:

Lorenz 63

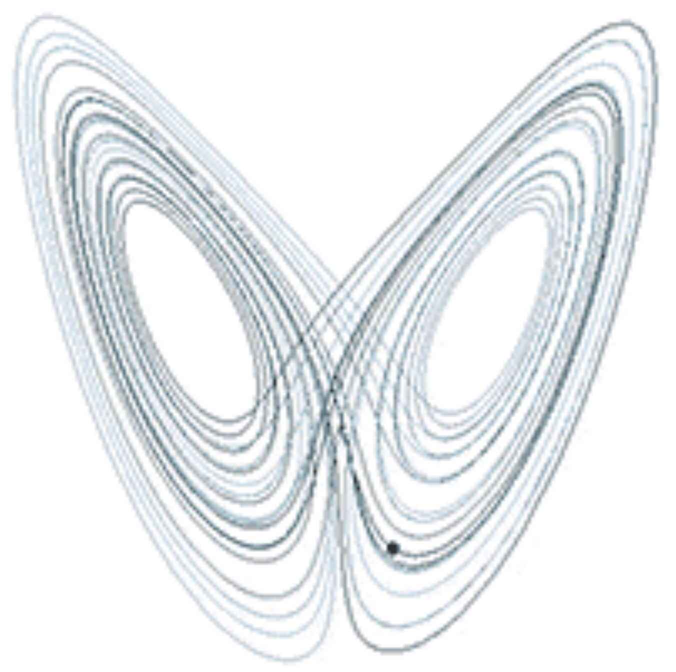

Molecular dynamics

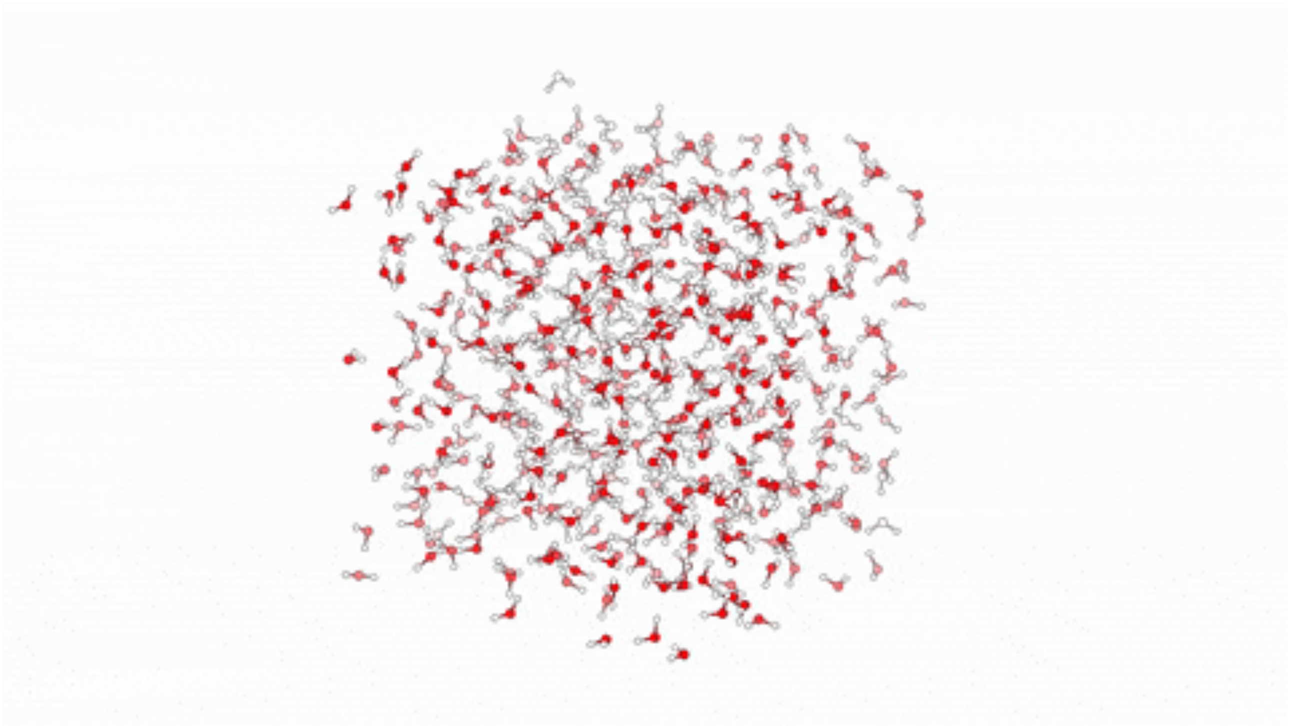

Robotic arm [1, 2]

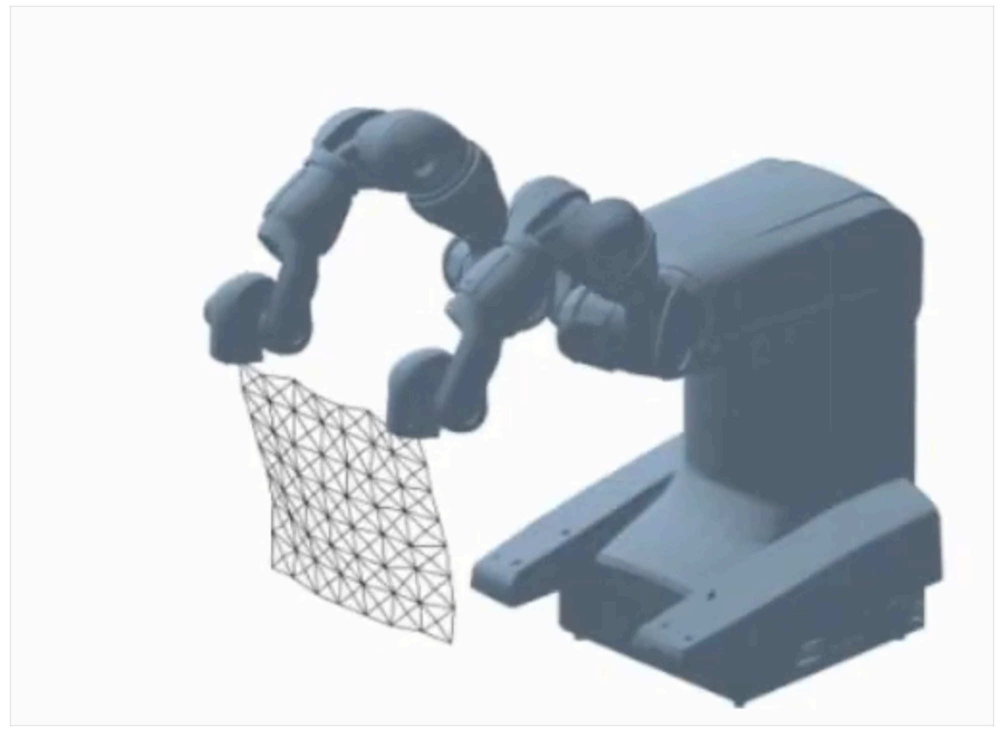

Models of dynamical systems:

Continuous-time

$x_t \in \mathbb{R}^n$: system's state at time t

$$dx_t = b(x_t)\,dt$$

Stochastic

W_t : Brownian motion

$$dX_t = \underbrace{b(X_t)}_{\text{drift}}\,dt + \underbrace{\sigma(X_t)}_{\text{diffusion}}dW_t$$

Controlled

$u_t \in \mathbb{R}^d$: control input at time t

$$dX_t = b(X_t, u_t)\,dt + \sigma(X_t, u_t)dW_t$$

Why models of controlled dynamical systems?

Predict system's evolution:

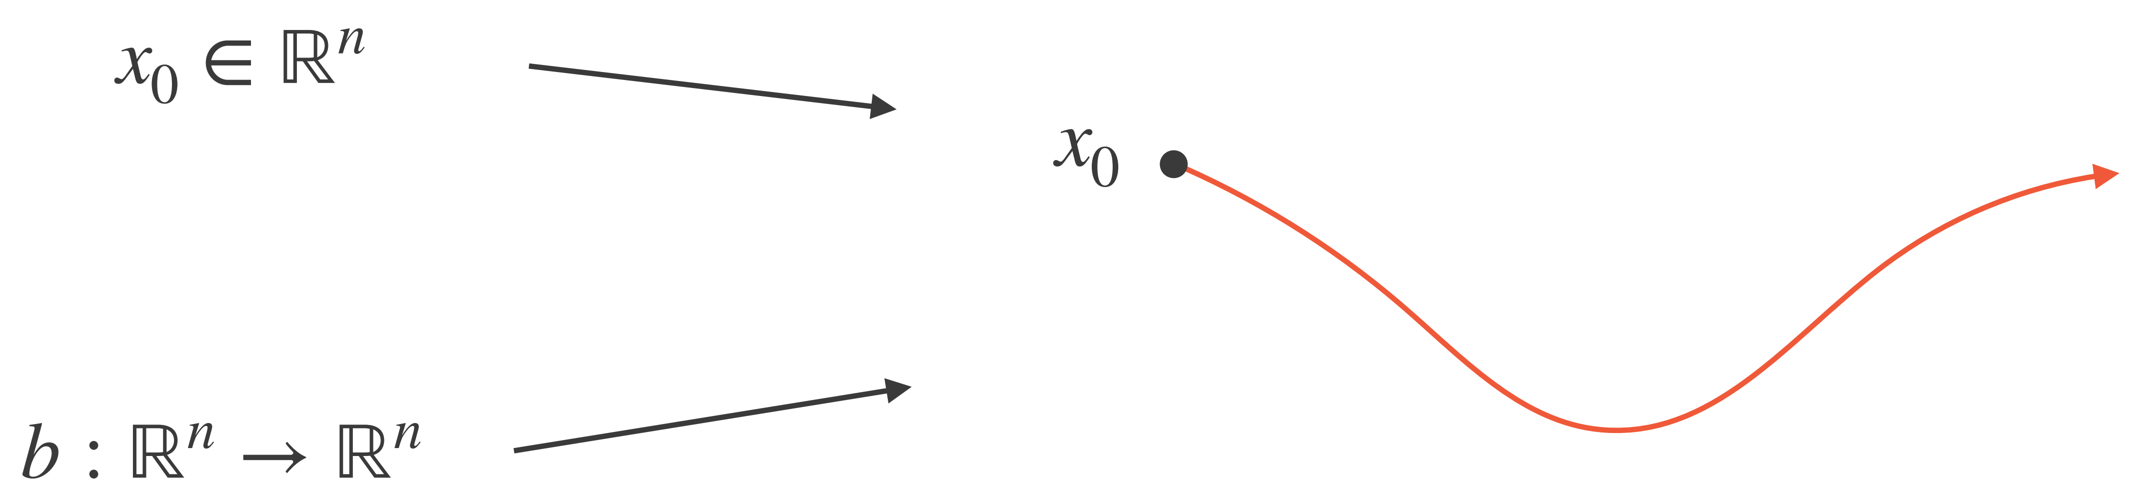

Why models of controlled dynamical systems?

Predict system's evolution:

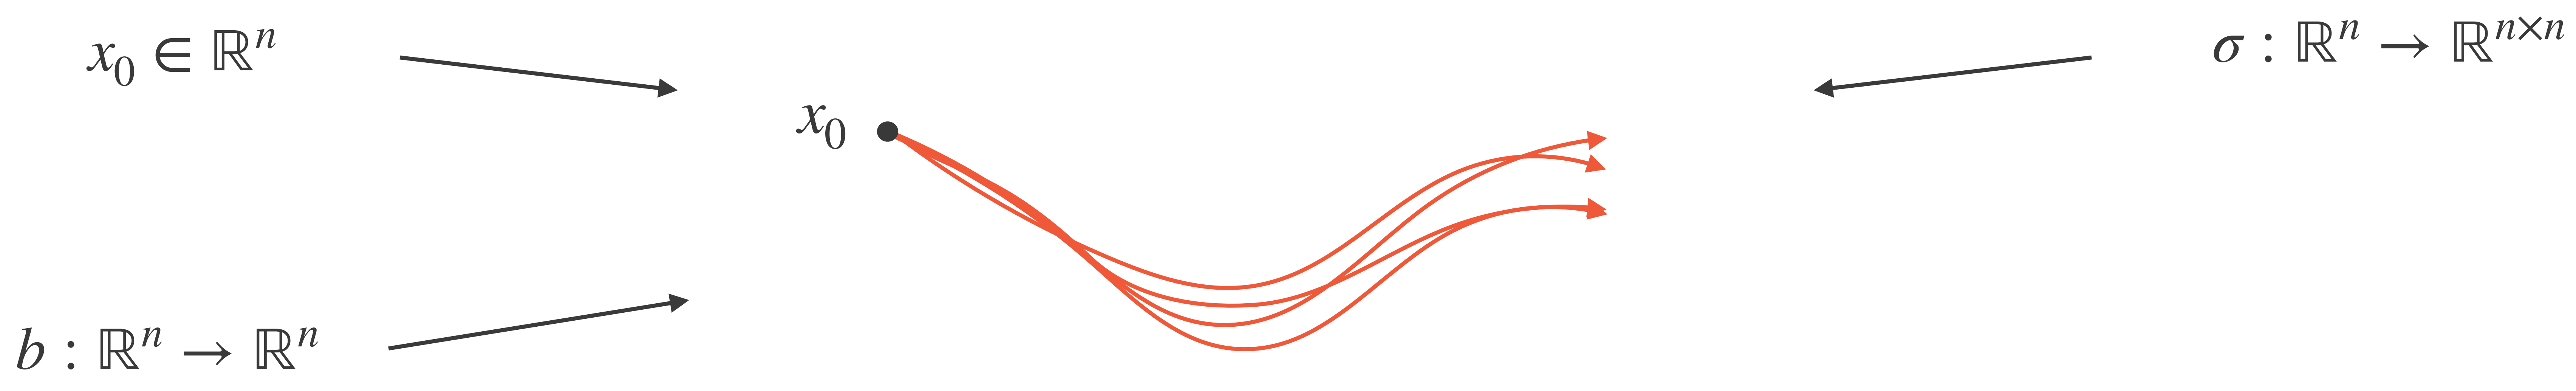

Why models of controlled dynamical systems?

Predict system's evolution:

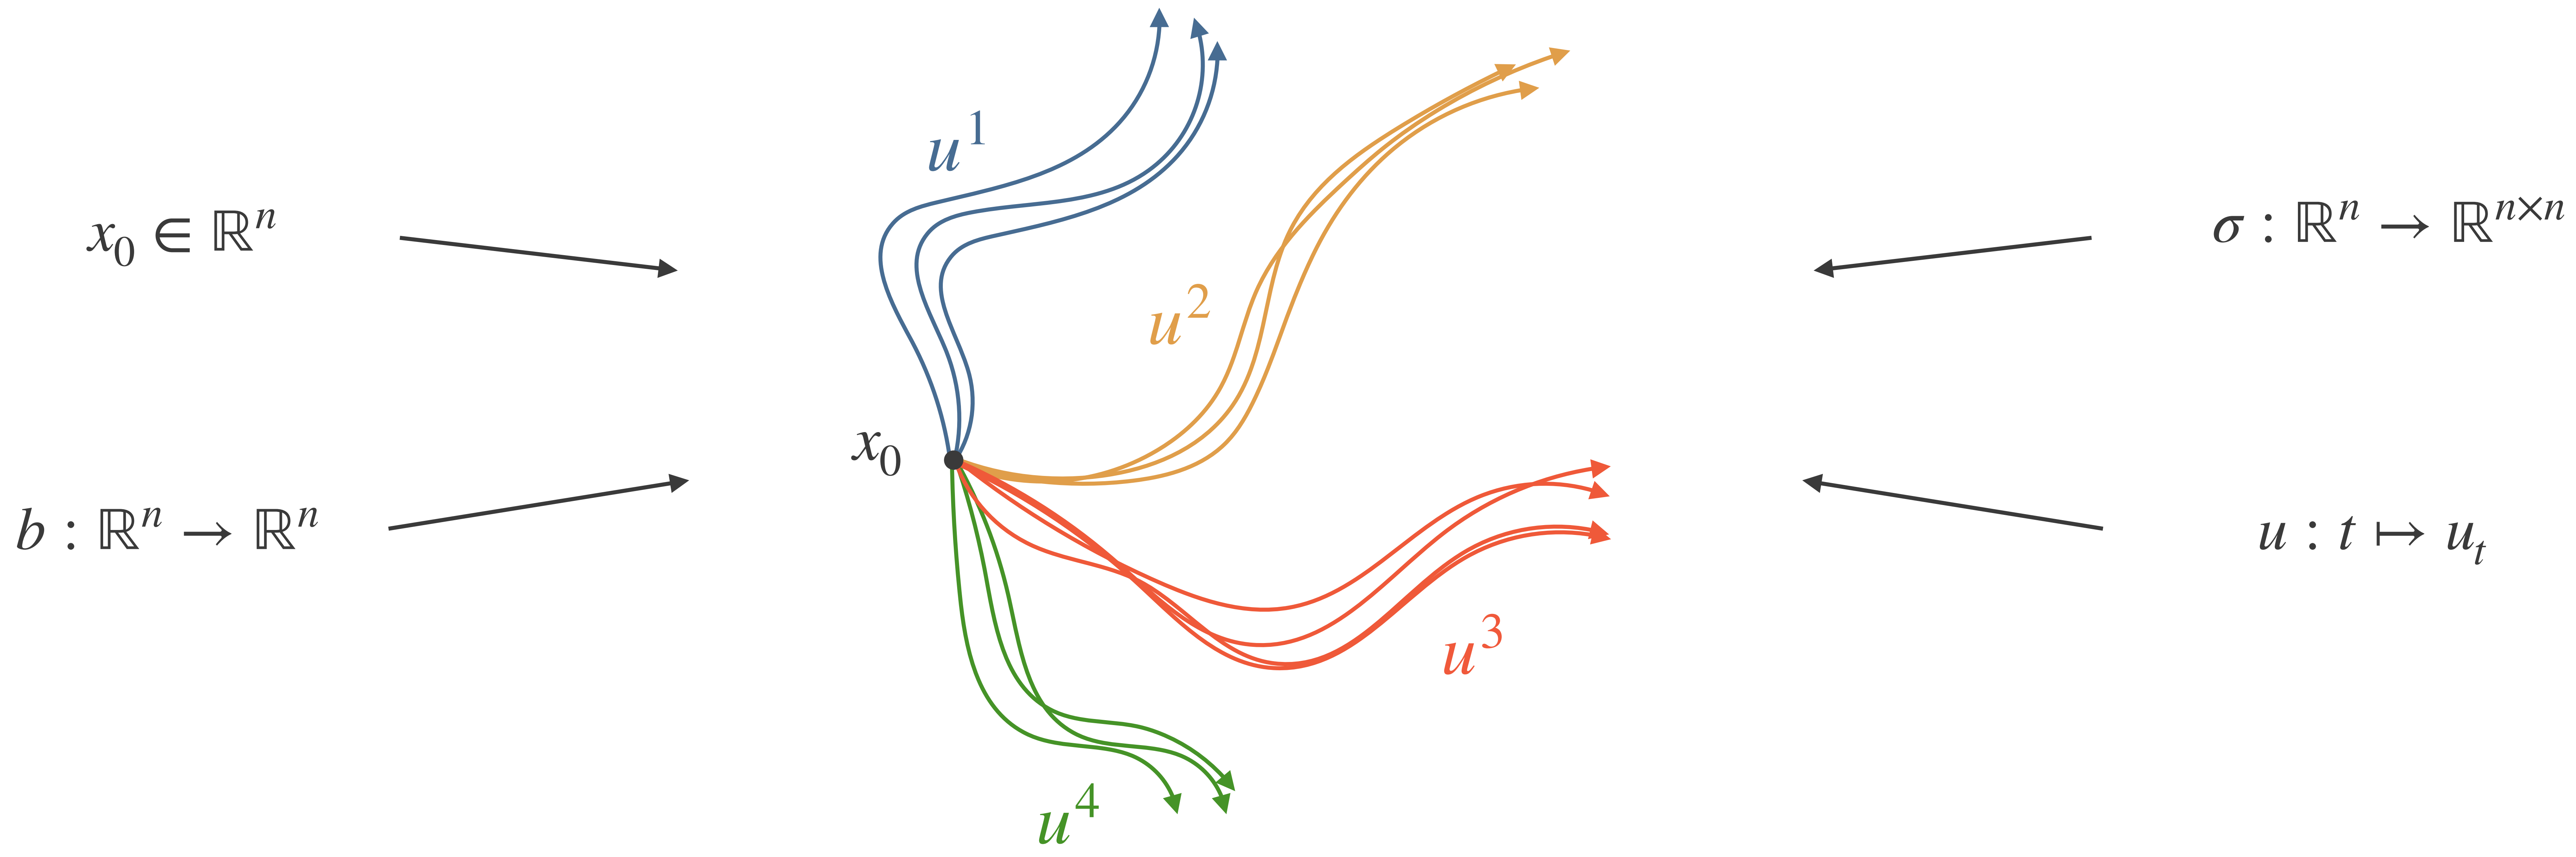

Why models of controlled dynamical systems?

Predict system's evolution:

$x_0 \in \mathbb{R}^n$

$b : \mathbb{R}^n \to \mathbb{R}^n$

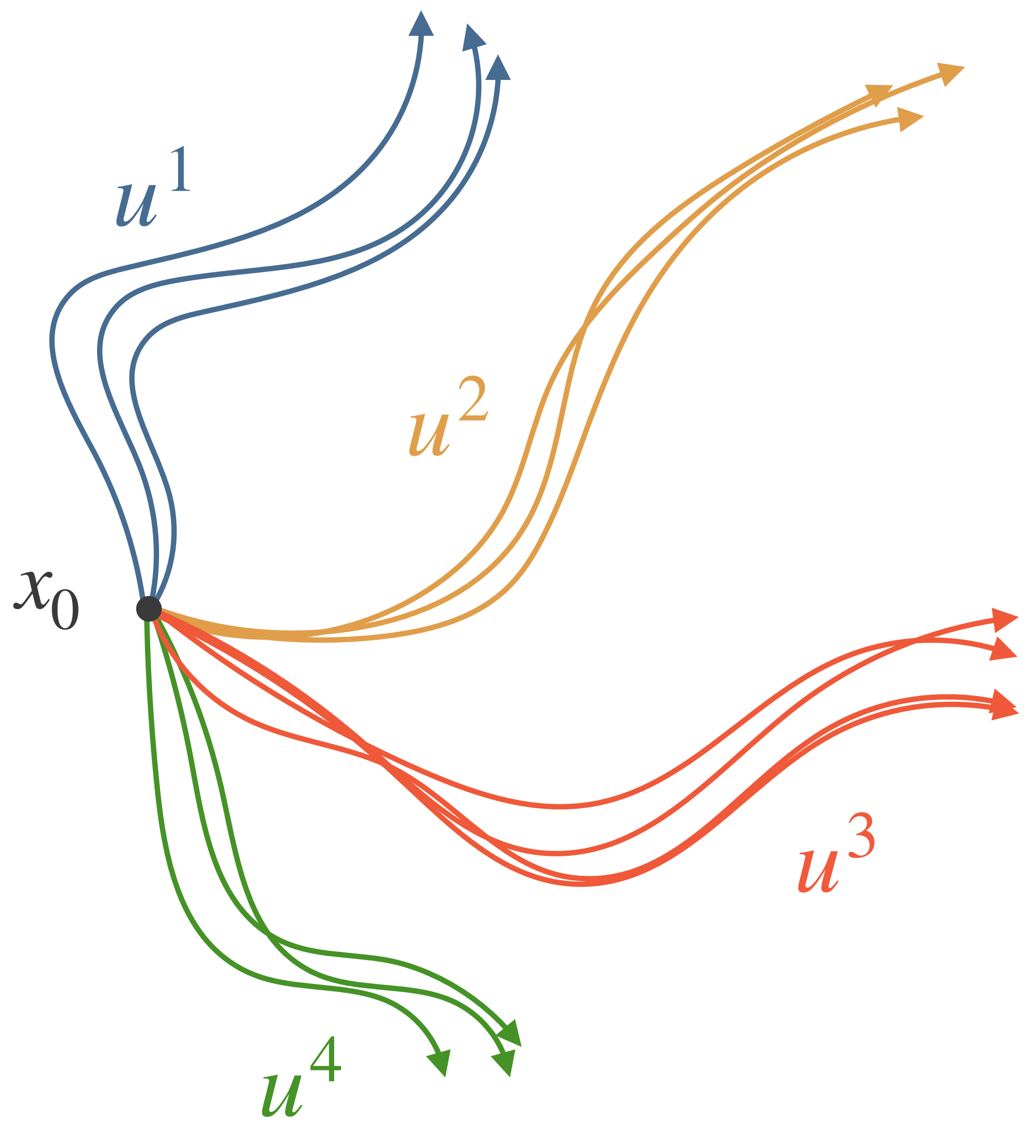

$\sigma : \mathbb{R}^n \to \mathbb{R}^{n \times n}$

$u : t \mapsto u_t$

Numerous applications: design controls to achieve desired behavior

Industrial processes

Robotic systems

Medical systems [3]

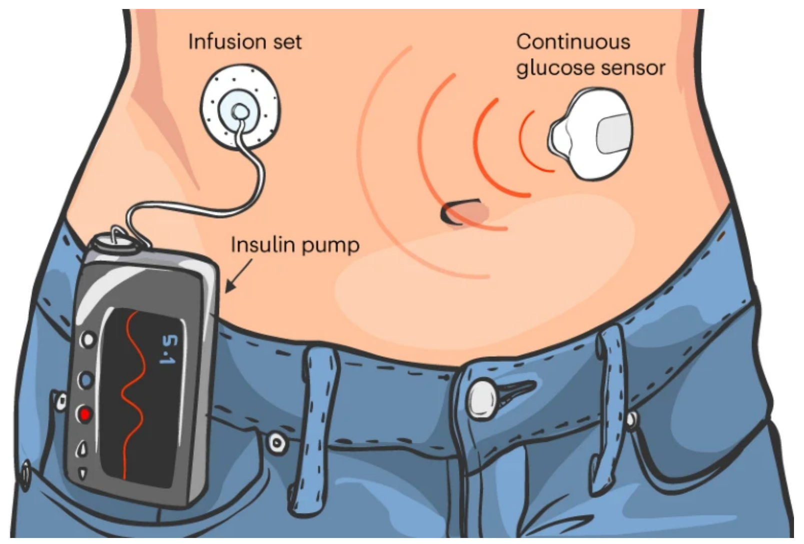

Learning problem

In many applications, the model is unknown:

$$(b, \sigma) = ?$$

Learning problem

In many applications, the model is unknown:

$$(b, \sigma) = (b_0, \sigma_0) + ?$$

Learning problem

In many applications, the model is unknown:

$$(b, \sigma) = ?$$

Learning problem

In many applications, the model is unknown:

$$(b, \sigma) = \ ?$$

Parametrized control space:

bounded set

$$u_\theta : [0,T] \to \mathbb{R}^d \qquad\qquad \theta \in \Theta \subset \mathbb{R}^m$$

Controlled SDE estimation: learning (b, σ) by collecting trajectories under various controls $u_{\theta_1}, u_{\theta_2}, \ldots u_{\theta_K}$

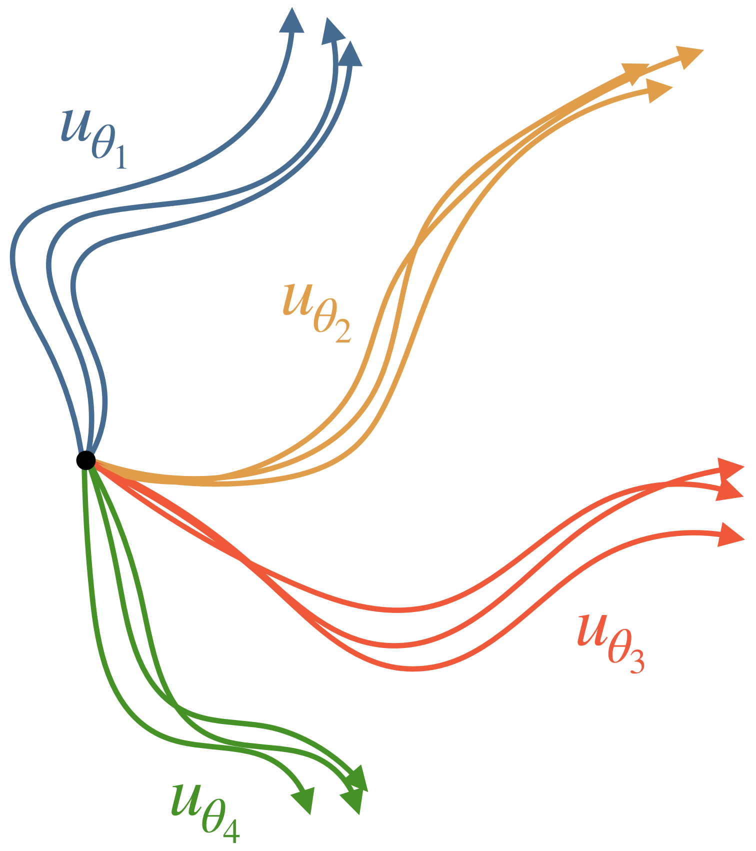

Learning problem

In many applications, the model is unknown:

$$(b, \sigma) = ?$$

Parametrized control space:

bounded set

$$u_\theta : [0,T] \to \mathbb{R}^d \qquad \theta \in \Theta \subset \mathbb{R}^m$$

Controlled SDE estimation: learning (b, σ) by collecting trajectories under various controls $u_{\theta_1}, u_{\theta_2}, \dots u_{\theta_K}$

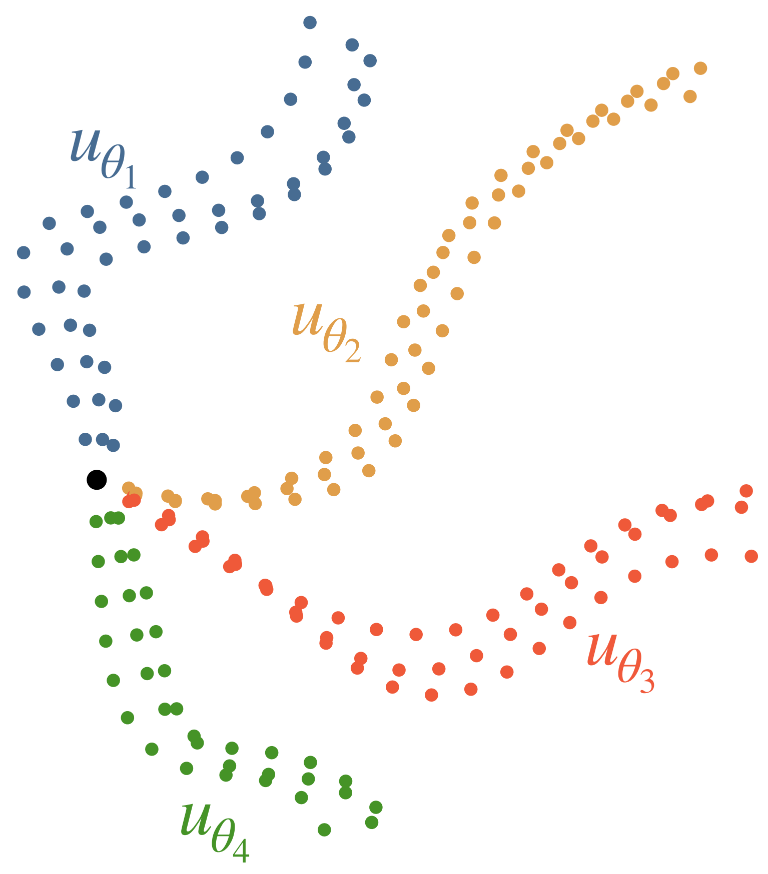

$$\longrightarrow$$

$$\{(\theta_k, X_{\theta_k}(w_i^k, t_l)\}_{k,i,l} \in \mathbb{R}^{K \times N \times M}$$

$$K = 4, N = 3, M \approx 25$$

Focus of this work

Density flows: probability density of $X_\theta(t)$ under control u_θ

$$\theta, t, x, \quad \mapsto \quad p_{b,\sigma}(\theta, t, x)$$

Focus: *reproduce the density flow across controls*

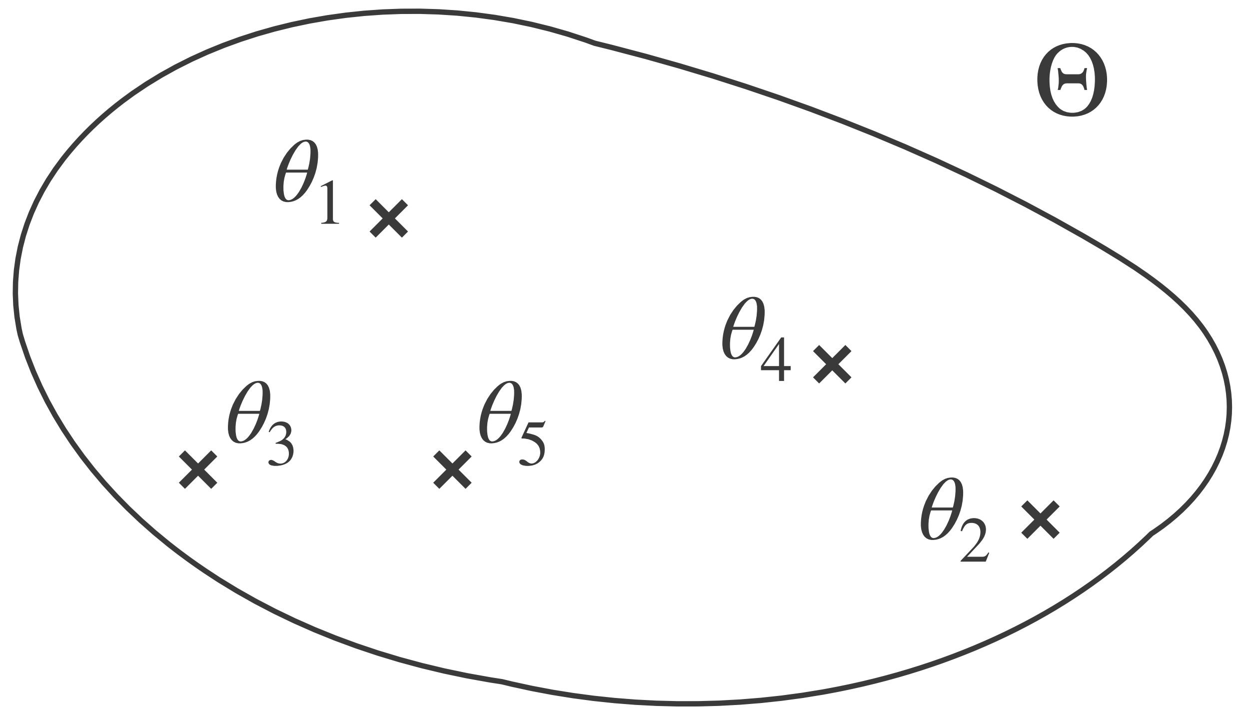

$$\boxed{\forall \theta \in \Theta, \qquad p_{\hat{b},\hat{\sigma}}(\theta) \approx p_{b,\sigma}(\theta)}$$

Remark: different drift–diffusion pairs may induce the same density flow

Trajectories under u_θ

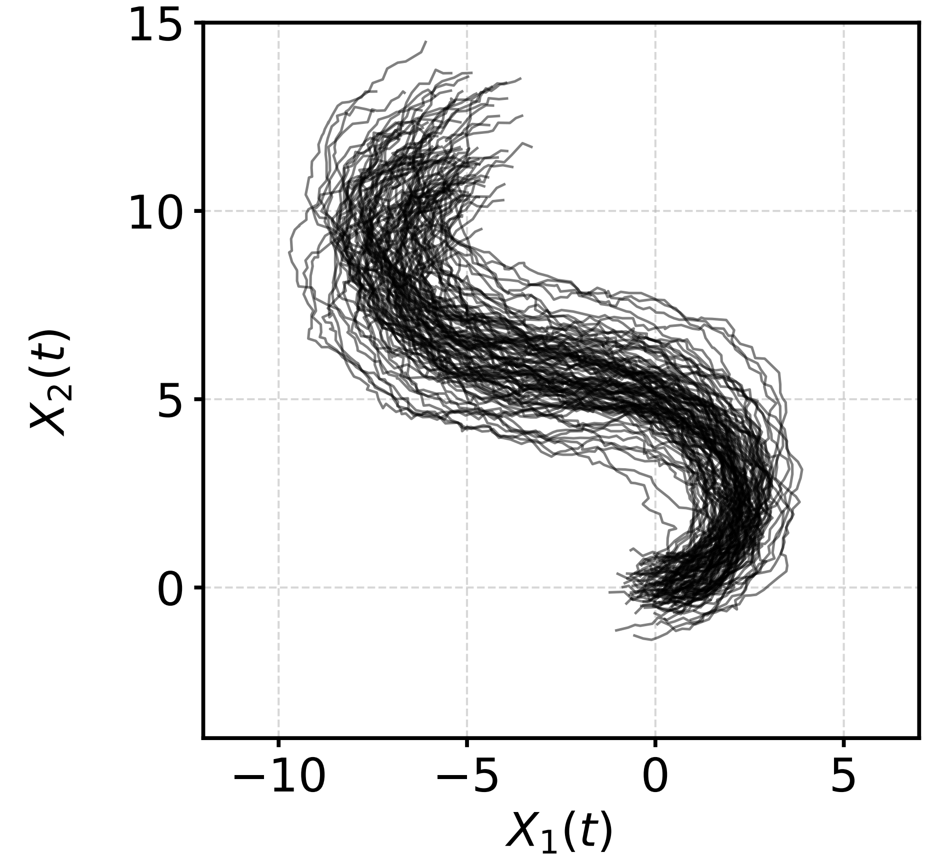

Density flow under u_θ

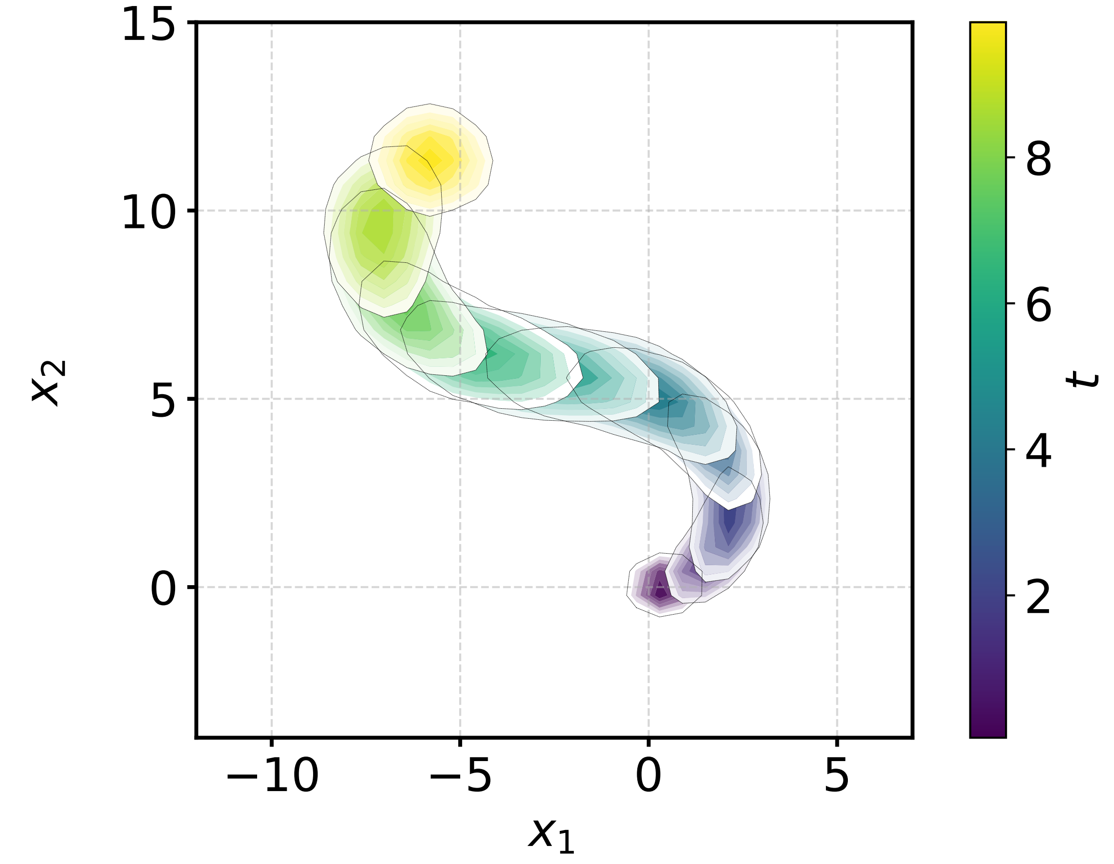

Proposed method

The Fokker-Planck equation

Fokker-Planck equation: let $X_\theta(t)$ governed by

$$dX_t = b(t, X_t, u_\theta(t))\,dt + \sigma(t, X_t, u_\theta(t))\,dW_t$$

Trajectories under u_θ

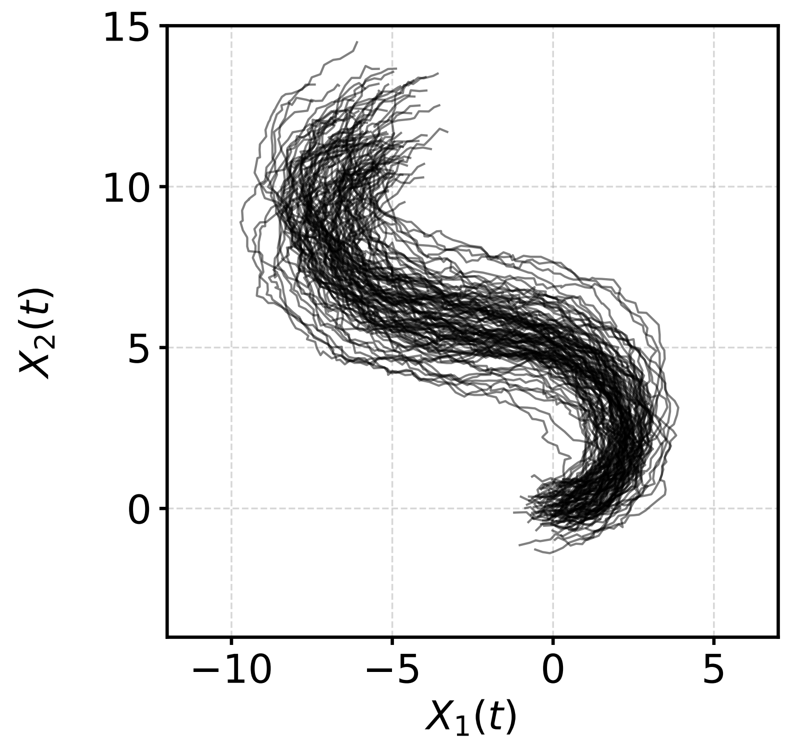

then under suitable regularity assumptions, its density flow p satisfies

$$\boxed{\partial_t p(\theta, t, x) = L^{(b,a),\theta} p(\theta, t, x),}$$

where

$$a = \sigma\sigma^T, \qquad L^{(b,a),\theta} q(t,x) = \frac{1}{2}\sum_{i,j=1}^{n} \partial_{x_i x_j}\Big(a_{ij}(t, x, u_\theta(t))\, q(t,x)\Big) - \sum_{i=1}^{n} \partial_{x_i}\big(b_i(t, x, u_\theta(t))\, q(t,x)\big)$$

Density flow under u_θ

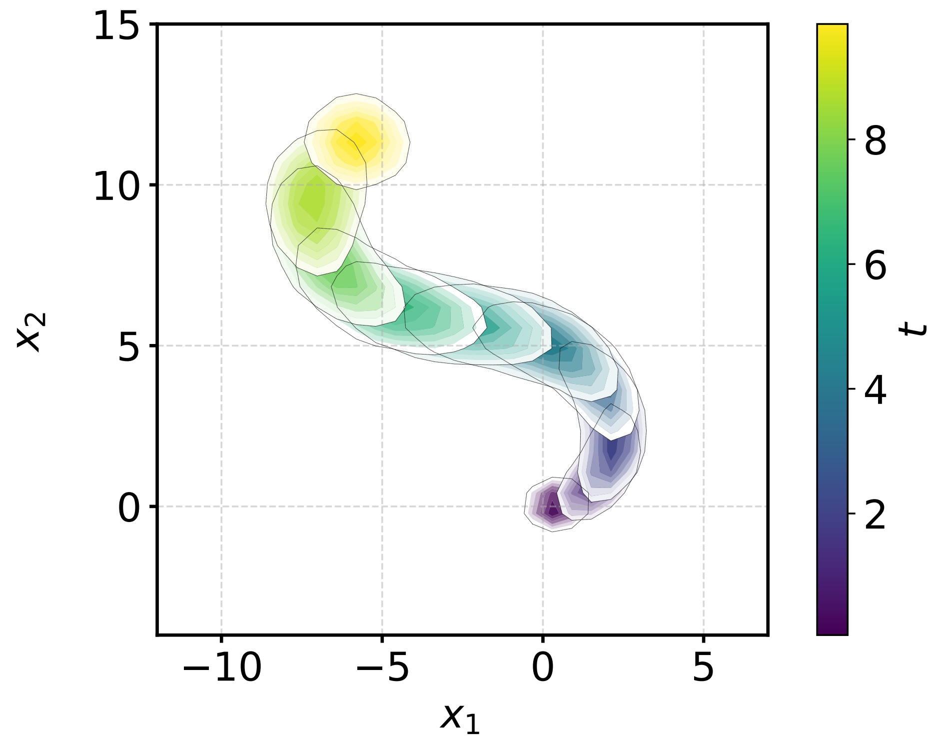

Proposed estimator

Two-step method:

1. Estimate the density flows $\hat{p}_k \approx p(\theta_k, \cdot\,, \cdot\,)$ for each of the controls in the dataset $u_{\theta_1}, \ldots, u_{\theta_K}$

2. Estimate coefficients thanks to the Fokker–Planck equation:

$$\arg\min_{(\hat{b},\hat{\sigma})\in\mathscr{F}} \left\| \frac{\partial p_{b,\sigma}}{\partial t} - L^{\hat{b},\hat{\sigma}} p_{b,\sigma} \right\|_{L^2}^2$$

Model $\mathscr{F}$ for the coefficients:

$$b(t,x,v) = \underbrace{\xi(x)}_{\text{localization}}\,\underbrace{b_0(t,x,v)}_{\text{RKHS model}} \qquad a = \sigma\sigma^T \qquad a(t,x,v) = \underbrace{\kappa I_n}_{\text{ellipticity floor}} + \overbrace{\xi(x)}^{\text{localization}}\,\underbrace{a_0(t,x,v)}_{\text{positive-semidefinite RKHS model [4]}}$$

Statistical guarantees

Main theoretical guarantees

Statistical guarantee: under sufficiently accurate density flow estimations, with probability at least $1-\delta$

$$\left\| p_{\hat{b},\hat{a}} - p_{b,a} \right\|_{L^2(\Theta\times[0,T]\times\mathbb{R}^n)} \lesssim \left(\frac{\log(K/\delta)}{\sqrt{K}} \right)^{\frac{1}{1-\min(r,s)}}$$

with r, s measuring the regularity-dimension of the true controlled system in the state–time and control directions.

Example: for a ν-Sobolev-regular controlled system *

$$r < 1 - \frac{\dim(X)+1}{2\nu}, \qquad s < 1 - \frac{\dim(\Theta)}{2\nu}$$

Take-home message:

Higher regularity and lower dimension $\implies$ faster generalization

(★) $(\theta,t) \mapsto u_\theta(t)$, $(\theta,t,x) \mapsto p(\theta,t,x)$, $(t,x,v) \mapsto (b,a)(t,x,v)$ are sufficiently Sobolev-regular so that $(\theta,t,x) \mapsto L^{u_\theta}_{b,a} p(\theta,t,x)$ has H^ν-regular sections in θ and (t,x)

Beyond average performance?

Beyond average accuracy?

$L^2(\Theta)$ bounds give accuracy on average over controls

Uniform guarantees: under same assumptions, with probability at least $1-\delta$

$$\left\| p_{\hat{b},\hat{a}} - p_{b,a} \right\|_{L^\infty(\Theta; L^2([0,T]\times\mathbb{R}^n))} \lesssim \left(\frac{\log(K/\delta)}{\sqrt{K}} \right)^{\frac{r\wedge s}{1-r\wedge s}}$$

slower than L^2

Tail-risk guarantees: if $\mathbb{E}\left[\|X_0\|^\beta\right] < \infty, \beta > 2$, with probability at least $1-\delta$

slower than L^∞

$$\sup_{\theta\in\Theta} \left\| \mathrm{CVaR}_\rho\big(f(X^\theta_{\hat{b},\hat{a}}(\,\cdot\,))\big) - \mathrm{CVaR}_\rho\big(f(X^\theta_{b,a}(\,\cdot\,))\big) \right\|_{L^2(0,T)} \lesssim \frac{\|f\|_\infty}{\rho} \left(\frac{\log(K/\delta)}{\sqrt{K}} \right)^{\frac{r\wedge s}{1-r\wedge s}\frac{\beta}{\beta+n/2}}$$

Numerical experiments

Numerical illustration

Ornstein–Uhlenbeck process: training on $K = 10$ controls, $N = 1000$ trajectories.

$$dX(t) = \frac{1}{2}(u(t) - X(t))\,dt + \frac{1}{\sqrt{8}}\,dW(t), \quad u(t) = \begin{cases} \theta_0 & t < \theta_3 \\ \theta_1 & t \geq \theta_3 \end{cases}$$

u_1 u_2 u_3

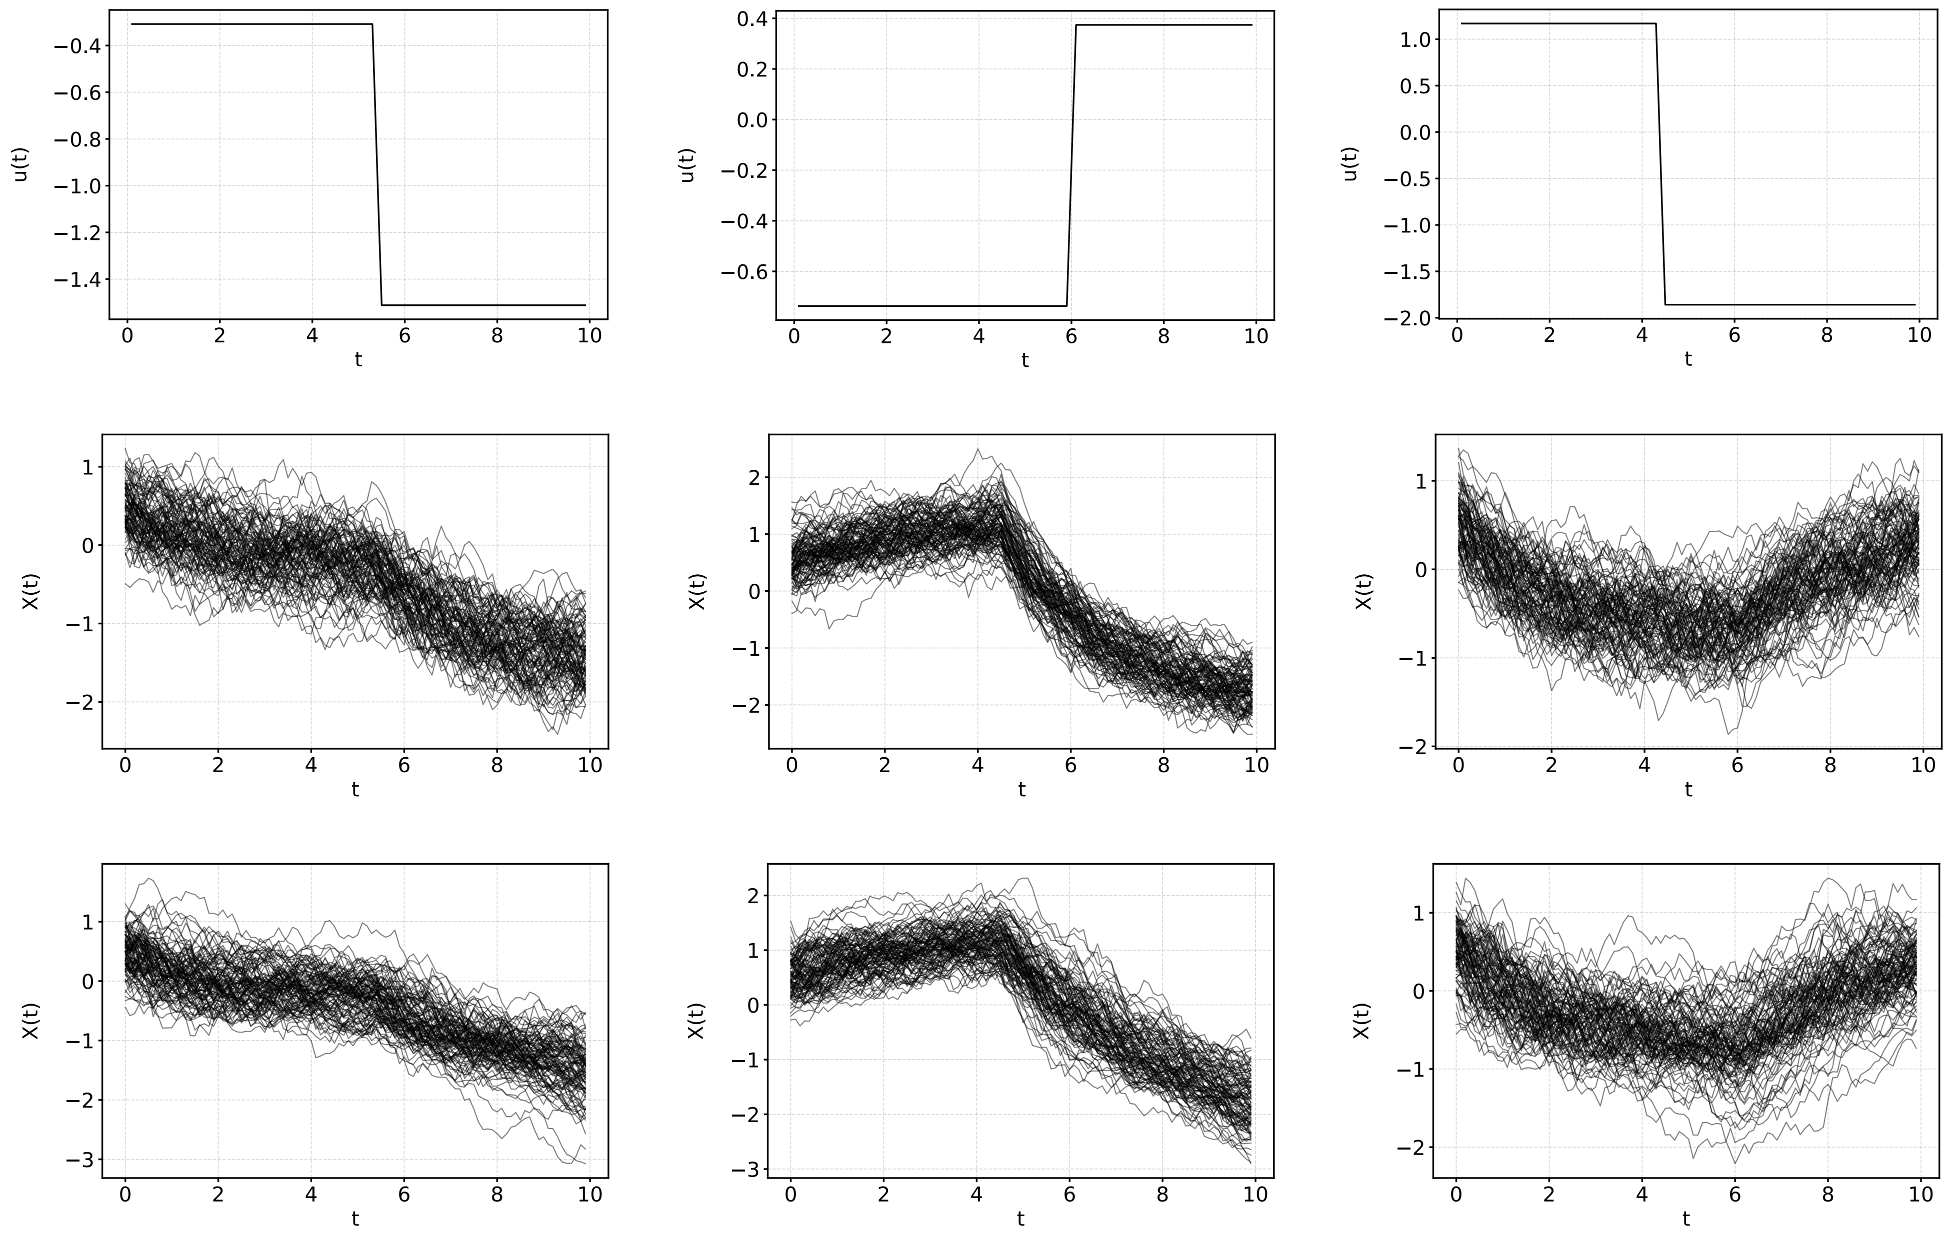

Test controls

True trajectories from (b, a)

Simulated trajectories from $(\hat{b}, \hat{a})$

Conclusion and perspectives

Takeaways and perspectives

Takeaways:

- Learn controlled dynamics from finitely many controls
- Reproduce density flow through the *Fokker-Planck equation*
- Obtain statistical convergence rates governed by regularity and dimension

Code: GitHub lmotte/sde-learn lmotte/controlled-sde-learn

Perspectives:

- From open-loop to feedback controls
- From exact kernels to scalable approximations
- From simulations to experiments on real systems

References

(1) Amadio, Fabio, et al. "Controlled Gaussian process dynamical models with application to robotic cloth manipulation: F. Amadio et al." *International Journal of Dynamics and Control* 11.6 (2023): 3209-3219.

(2) Caldarelli, Edoardo, et al. "Dynamic robotic cloth folding with efficient Koopman operator-based model predictive control." *arXiv preprint arXiv:2605.18373* (2026).

(3) Khan, M. (2024). "Navigating Insulin Pumps: Technology, Benefits, and What's Next." *News-Medical.Net.*

(4) Muzellec, Boris, Francis Bach, and Alessandro Rudi. "Learning PSD-valued functions using kernel sums-of-squares." *arXiv preprint arXiv:2111.11306* (2021).