Flow matching for stochastic linear control systems

Yuhang Mei, Mohammad Al-Jarrah, Amirhossein Taghvaei, Yongxin Chen
University of Washington, Seattle & Georgia Institute of Technology

The 7th Annual Learning for Dynamics & Control Conference (L4DC), Ann Arbor, United States, Jun, 2025

UNIVERSITY of WASHINGTON

Control of Probability Distributions

Fundamental problem in control:

$$x_0 \xrightarrow[\text{s.t.}\quad \dot{X}_t = AX_t + Bu_t]{\text{find}\quad u_t = k(t,x)} x_1$$

What if x_0, x_1 are characterized by probability distributions?

$$P_0 \xrightarrow[\text{s.t.}\quad \dot{X}_t = AX_t + Bu_t]{\text{find}\quad u_t = k(t,x)} P_1$$

$$\text{or}\quad \mathrm{d}X_t = AX_t\mathrm{d}t + B(u_t\mathrm{d}t + \epsilon\mathrm{d}W_t)$$

Motivation: Swarm control, stochastic thermodynamics,...

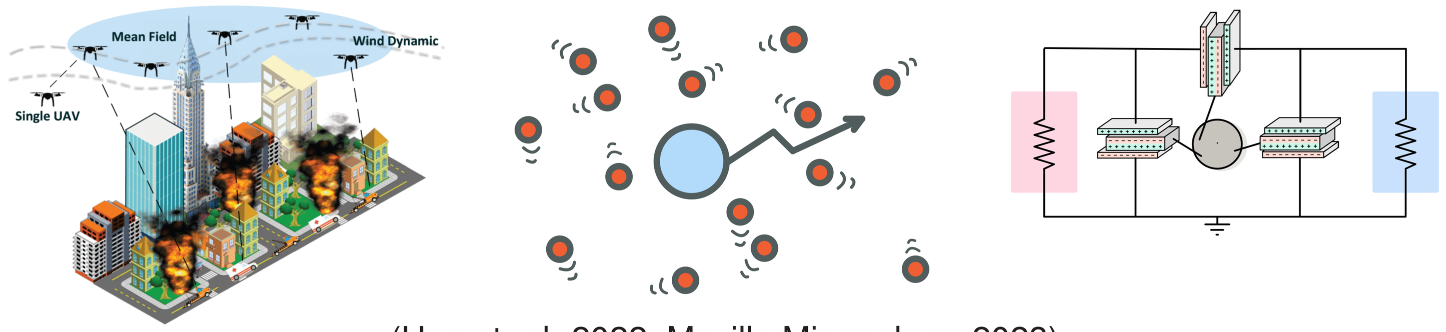

(Hao et. al. 2022; Movilla Miangolarra 2023)

Problem formulation

Model:

$$\mathrm{d}X_t = AX_t\,\mathrm{d}t + B(u_t\,\mathrm{d}t + \epsilon\,\mathrm{d}W_t) \qquad (\star)$$

- X_t is the state
- u_t is the control input
- W_t is the Brownian motion

Objective:

$$\begin{aligned} \text{Given a pair of distribution:}\quad & (P_0, P_1) \\ \text{Find a feedback control law:}\quad & u_t = k(t, X_t) \\ \text{s.t. If } X_0 \sim P_0, & \text{ then } X_1 \sim P_1, \end{aligned}$$

Flow matching methodology for generative models

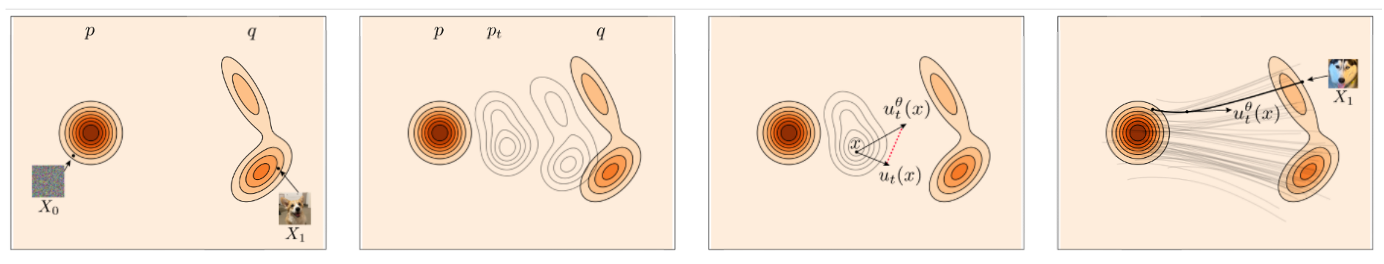

(a) Data. (b) Path design. (c) Training. (d) Sampling.

- A special case of $(\star)$, where $A = 0$, $B = I$, and $\epsilon = 0$.

$$\frac{\mathrm{d}X_t}{\mathrm{d}t} = u_t, \quad X_0 \sim P_{\text{initial}}$$

- Construct a probability flow $\{P_t; 0 \le t \le 1\}$ s.t. $(P_0 = P_{\text{initial}}, P_1 = P_{\text{target}})$
- Standard choice: linear interpolation

$$X_t^z = (1-t)x + ty, \quad z = (x,y) \sim \Pi := P_{\text{initial}} \otimes P_{\text{target}}$$

- Feedback control law $u_t = k(t, X_t)$:

$$k(t,\xi) = \mathbb{E}\left[\frac{\mathrm{d}X_t^z}{\mathrm{d}t} \Big| X_t^z = \xi\right], \quad \forall (t,\xi) \in [0,1] \times \mathbb{R}^n$$

- Approximate $k(t,\cdot)$ through least-squares regression:

$$\min_k \mathbb{E}_{z\sim\Pi}[\|k(t,X_t^z) - \frac{\mathrm{d}X_t^z}{\mathrm{d}t}\|^2]$$

Proposed methodology

Step 1: Construct stochastic interpolants for $(\star)$

$$\begin{aligned} \text{Given:}\quad & z = (x,y) \in \mathbb{R}^n \times \mathbb{R}^n, \text{ and } (\star) \\ \text{Find a feedback control law:}\quad & u_t = k^z(t, X_t) \\ \text{s.t.}\quad & X_0 = x, X_1 = y \end{aligned}$$

Solution:

$$k^z(t,x) = B^\top e^{(1-t)A^\top}\Phi_{1-t}^{-1}(y - e^{(1-t)A}x), \quad \text{where}$$

$$\Phi_t := \int_0^t e^{(t-s)A}BB^\top e^{(t-s)A^\top}\,\mathrm{d}s$$

Resulting trajectory:

$$X_t^z \overset{\mathrm{d}}{=} e^{tA}x + \Phi_t e^{(1-t)A^\top}\Phi_1^{-1}(y - e^{A}x) + \epsilon\Sigma_t^{\frac{1}{2}}G, \qquad (\star\star)$$

Illustration with three examples:

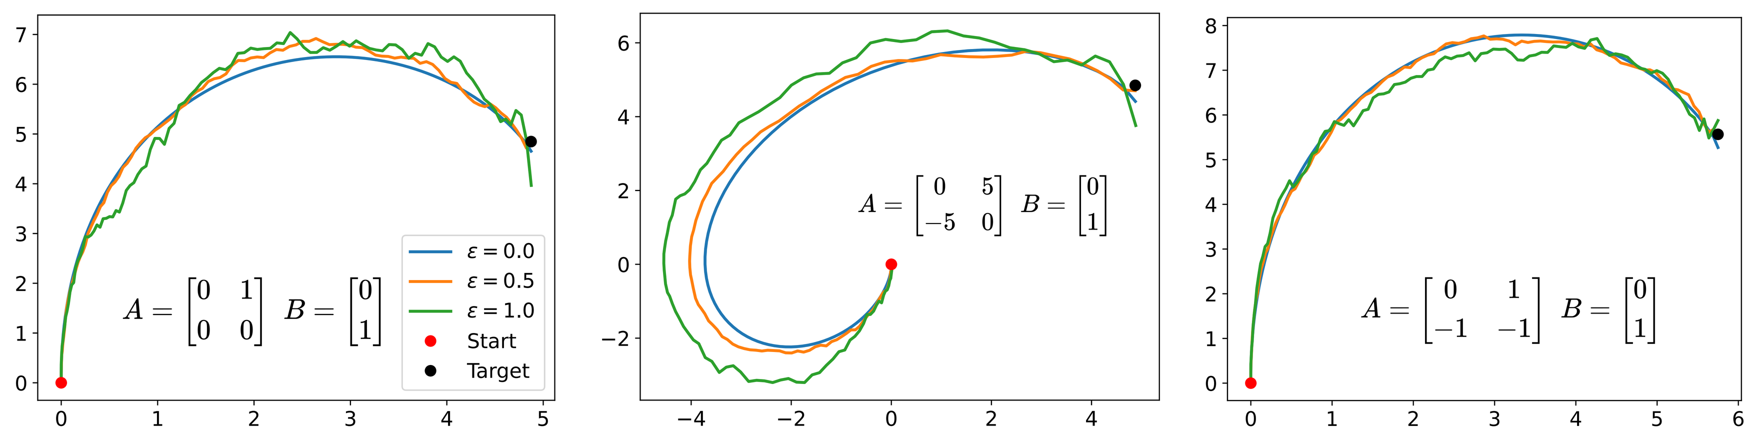

Step 2: Apply the flow matching methodology to the probability flow X_t^z defined in $(\star\star)$.

Feedback control law:

$$\bar{k}(t,\xi) = \mathbb{E}[k^z(t,X_t^z)|X_t^z = \xi]$$

Numerical approximation:

$$\min_{\bar{k}} \mathbb{E}[\|\bar{k}(t,X_t^z) - k^z(X_t^z,)\|^2] \approx \arg\min_{\bar{k}\in\mathcal{F}} \frac{1}{N}\sum_{i=1}^{N} \|\bar{k}(t,X_t^{z^i}) - k^{z^i}(t,X_t^{z^i})\|^2$$

with N independent samples of the pair $z^i = (x^i, y^i) \sim \Pi$

Theoretical justifications

Justification for step 1:

Consider the uncontrolled process $\tilde{X}_t$ satisfying

$$\mathrm{d}\tilde{X}_t = -A\tilde{X}_t\,\mathrm{d}t + \epsilon B\,\mathrm{d}\tilde{W}_t, \quad \tilde{X}_0 = y.$$

Then, $\tilde{X}_t \sim \mathcal{N}(m_t, Q_t)$, with $m_t = e^{-At}y$, and

$$Q_t = \epsilon^2\int_0^t e^{-sA}BB^\top e^{-sA^\top} = \epsilon^2 e^{-tA}\Phi_t e^{-tA^\top}.$$

The time-reversal $X_t := \tilde{X}_{1-t}$ satisfies the SDE

$$\mathrm{d}X_t = AX_t\,\mathrm{d}t + \epsilon B\,\mathrm{d}W_t - \epsilon^2 BB^\top Q_{1-t}^{-1}(X_t - m_{1-t})\,\mathrm{d}t.$$

Note that, by construction, $X_1 = \tilde{X}_0 = y$. This is true starting from any initial point $X_0 = x$.

Justification for step 2:

We show X_t defined by $\mathrm{d}X_t = AX_t + B\bar{k}(t,X_t)\,\mathrm{d}t + B\epsilon\,\mathrm{d}W_t$, and X_t^z have the same distribution. For any a smooth and bounded test function $f:\mathbb{R}^n \to \mathbb{R}$:

$$\textbf{Itô rule} \longrightarrow \quad \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[f(X_t)] = \mathbb{E}[\nabla f(X_t)(AX_t + B\bar{k}(t,X_t)) + \epsilon^2 Tr(BB^\top\nabla^2 f(X_t)]$$

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[f(X_t^z)] &= \mathbb{E}[\nabla f(X_t^z)(AX_t^z + Bk^z(t,X_t^z))) + \epsilon^2 Tr(BB^\top\nabla^2 f(X_t^z)] \\ \textbf{Tower property} \longrightarrow &= \mathbb{E}[\nabla f(X_t^z)(AX_t^z + B\mathbb{E}[k^z(t,X_t^z)|X_t^z]) + \epsilon^2 Tr(BB^\top\nabla^2 f(X_t^z)] \\ &= \mathbb{E}[\nabla f(X_t^z)(AX_t^z + B\bar{k}(t,X_t^z)) + \epsilon^2 Tr(BB^\top\nabla^2 f(X_t^z)] \end{aligned}$$

Numerical demonstrations

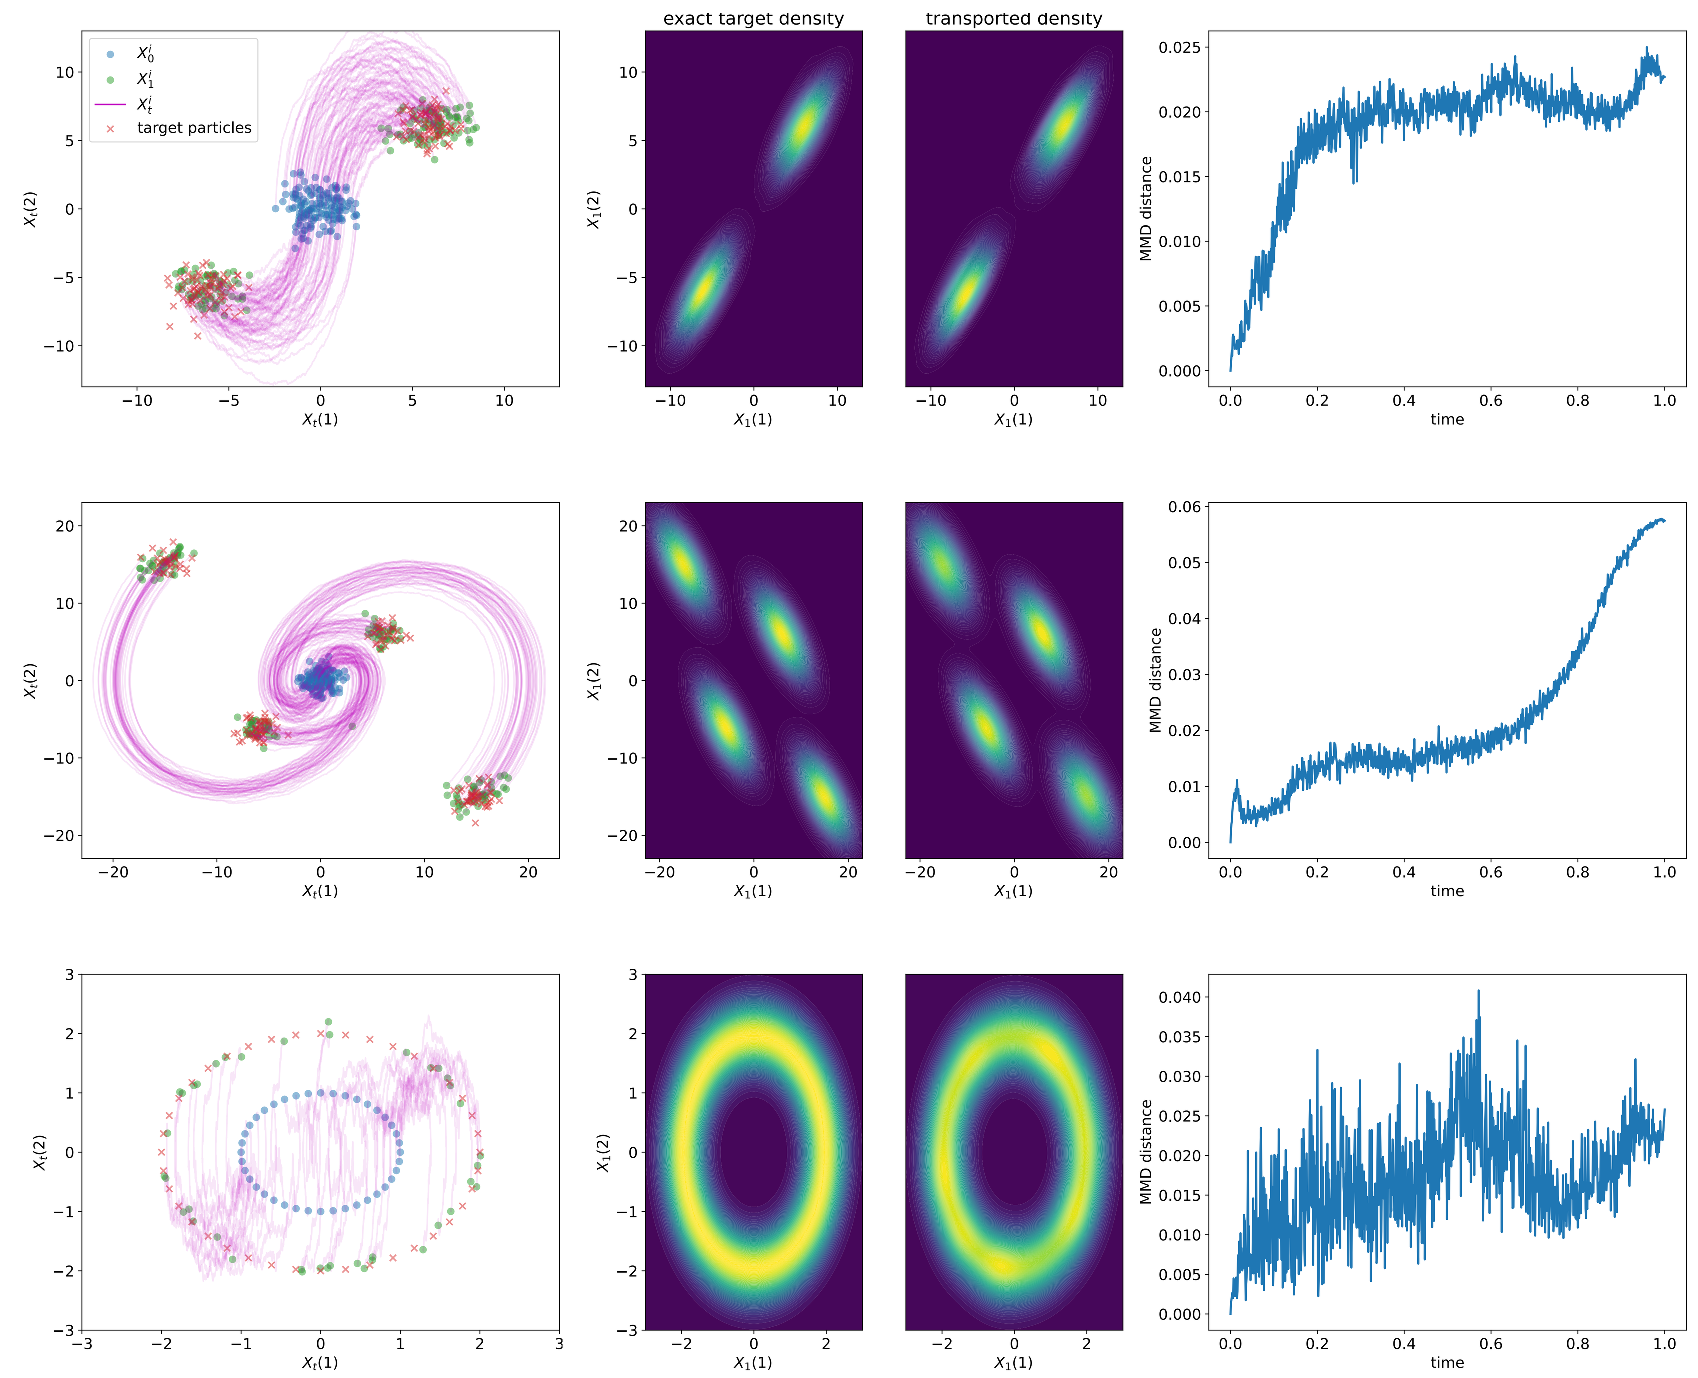

Numerical result: High-dimensional systems

Mass-spring model

$$A = \begin{bmatrix} 0, & I_{p\times p} \\ -\mathbb{T}, & -I_{p\times p} \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ I_{p\times p} \end{bmatrix}$$

$\mathbb{T} \in \mathbb{R}^{p\times p}$ is a Toeplitz matrix with 2 on main diagonal and -1 on the first super and sub-diagnal. System dimension $n = 2p$

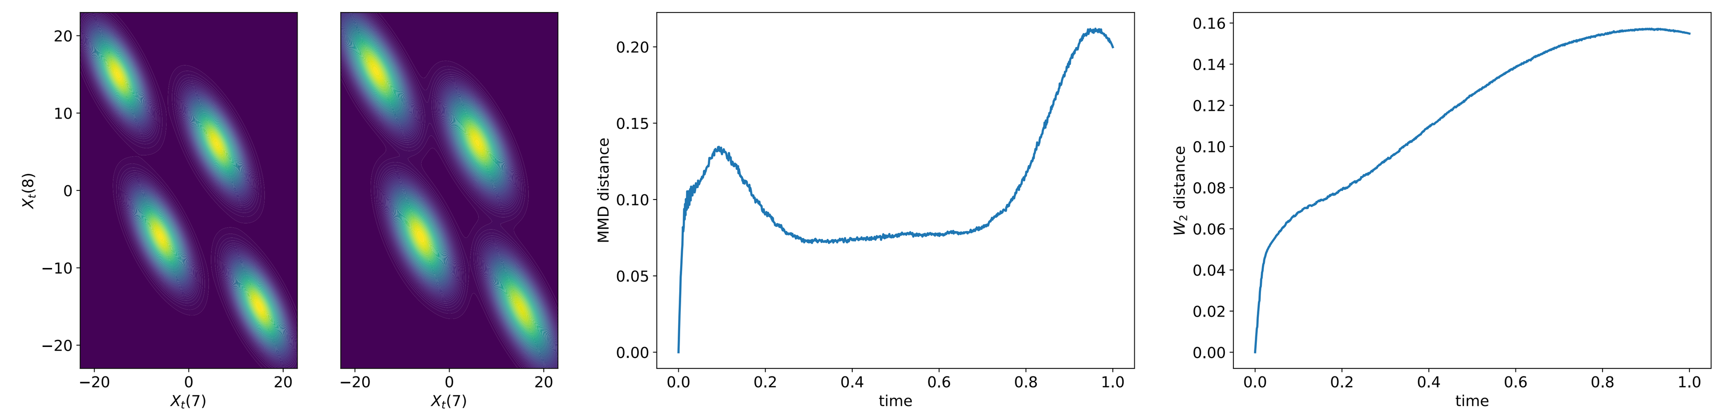

Conclusion

Main contribution: A novel flow matching framework within the context of deterministic and stochastic linear control systems.

Feature:

- High effectiveness in achieving distribution alignment
- Robustness across varying system parameters
- Scalability to diverse classes of distributions

Future directions:

- Extending the approach to control-affine systems
- Incorporating state constraints