A time-reversal methodology for steering the state of control-affine stochastic systems

Presented at the 61th Annual Allerton Conference on Communication, Control, and Computing, Urbana, Illinois

Amirhossein Taghvaei

Department of Aeronautics & Astronautics
University of Washington, Seattle

Sep 17, 20245

Overview

Part 1: point to point steering

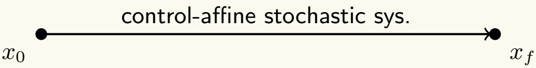

Part 2: distribution to distribution steering

- Focus on computability rather than optimality
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview

Part 1: point to point steering

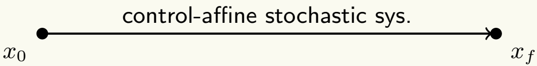

Part 2: distribution to distribution steering

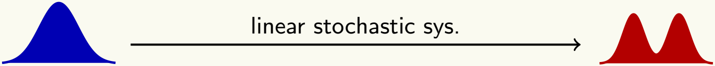

- Focus on computability rather than optimality
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview

Part 1: point to point steering

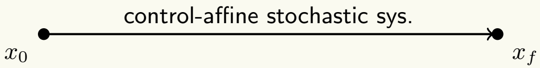

Part 2: distribution to distribution steering

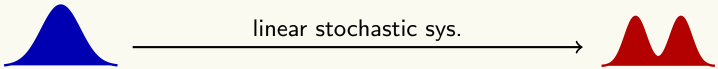

- Focus on computability rather than optimality
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview

Part 1: point to point steering

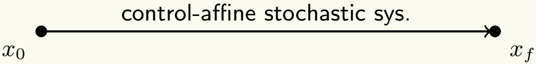

Part 2: distribution to distribution steering

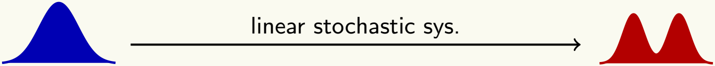

- Focus on <u>computability</u> rather than <u>optimality</u>
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview

Part 1: point to point steering

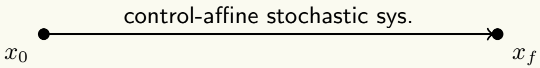

Part 2: distribution to distribution steering

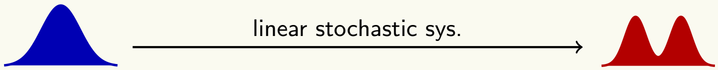

- Focus on <u>computability</u> rather than <u>optimality</u>
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Outline

- **Part 0:** Background on time-reversal of diffusions
- **Part 1:** point to point steering
- **Part 2:** distribution to distribution steering

Outline

- **Part 0:** Background on time-reversal of diffusions
- **Part 1:** point to point steering
- **Part 2:** distribution to distribution steering

Time-reversal of diffusions

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

Time-reversal of diffusions

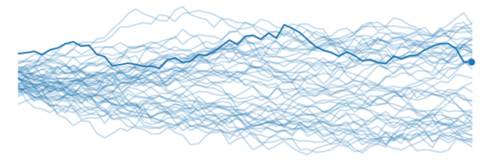

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

Time-reversal of diffusions

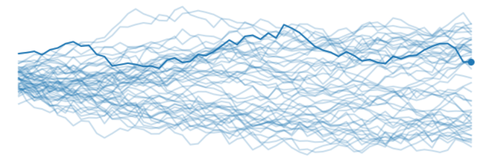

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

Time-reversal of diffusions

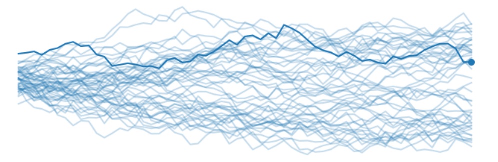

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

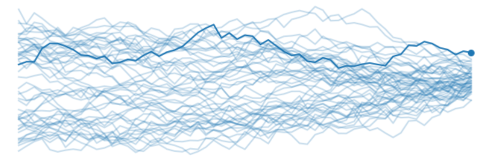

reversed process: $\widetilde{Z}_t := Z_{T-t},\quad \mathrm{d}\widetilde{Z}_t = ?$

Time-reversal of diffusions

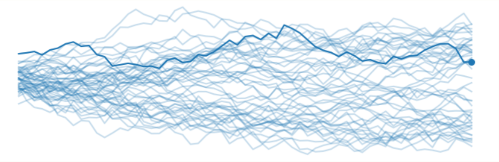

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

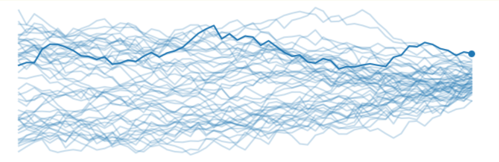

reversed process: $\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = ?$

Time-reversal of diffusions
Application in generative modeling

$$\mathrm{d}Z_t = -Z_t\mathrm{d}t + \sqrt{2}\mathrm{d}W_t$$

Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in International Conference on Learning Rep- resentations, 2021

Time-reversal of diffusions
Application in generative modeling

$$\mathrm{d}Z_t = -Z_t\mathrm{d}t + \sqrt{2}\mathrm{d}W_t$$

$$\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = \text{“generative model”}$$

Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” in International Conference on Learning Rep- resentations, 2021

Time-reversal of diffusions

Application in generative modeling

$$\mathrm{d}Z_t = -Z_t \mathrm{d}t + \sqrt{2}\mathrm{d}W_t$$

$$\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = \text{"generative model"}$$

Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in International Conference on Learning Rep- resentations, 2021

Time-reversal of diffusions
Application in control?

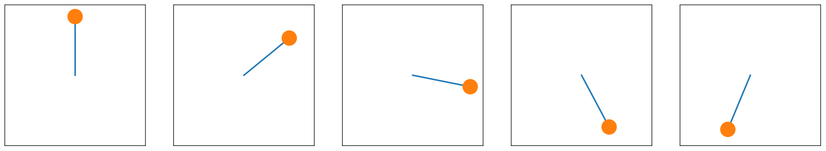

$\mathrm{d}Z_t =$ "pendulum dynamics" $\longrightarrow$

Time-reversal of diffusions

Application in control?

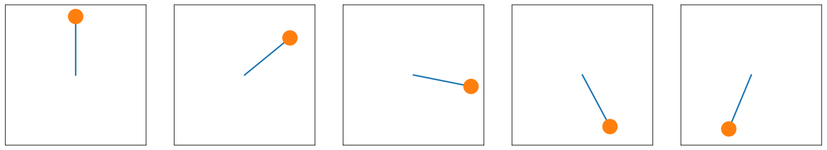

$\mathrm{d}Z_t =$ "pendulum dynamics"

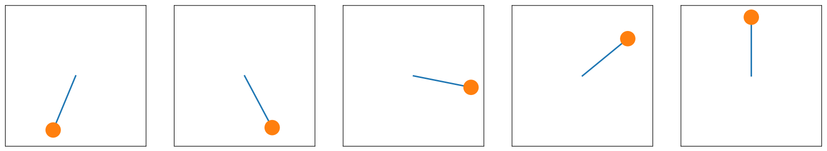

$\widetilde{Z}_t := Z_{T-t}$, $\mathrm{d}\widetilde{Z}_t =$ "steer to upward position"

Time-reversal of diffusions
Application in control?

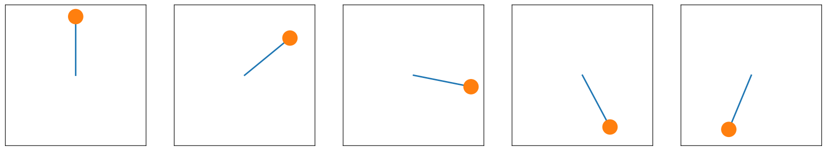

$\mathrm{d}Z_t =$ "pendulum dynamics"

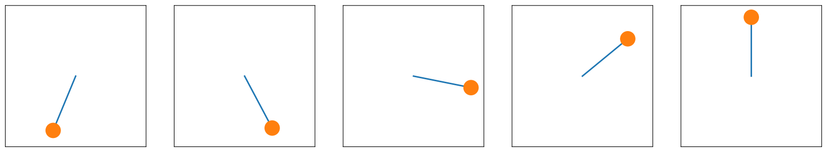

$\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t =$ "steer to upward position"

- Forward process

$$dZ_t = h(Z_t)dt + g(Z_t)dW_t$$

- Hörmander condition $\rightarrow$ allows for degenerate diffusions

$$\text{Span}(\text{Lie}\{h(x), g_1(x), \ldots, g_m(x)\}) = \mathbb{R}^n, \quad \forall x \in \mathbb{R}^n$$

- Time-reversed process $\tilde{Z} := \{\tilde{Z}_t = Z_{T-t}; 0 \leq t \leq T\}$

Time-reversal formula [Haussmann & Pardoux, 1986]

$$d\widetilde{Z}_t = -h(\widetilde{Z}_t)dt + s(T-t, \widetilde{Z}_t)dt + g(\widetilde{Z}_t)dW_t$$

$$\text{where} \quad s(t,x) = \frac{1}{p(t,x)} \nabla \cdot (p(t,x) g g^\top(x)), \quad p(t,\cdot) := \text{pdf}(Z_t)$$

B. D. Anderson, "Reverse-time diffusion equation models," Stochastic Processes and their Applications, vol. 12, no. 3, pp. 313–326, 1982
U. G. Haussmann and E. Pardoux, "Time reversal of diffusions," The Annals of Probability, pp. 1188–1205, 1986
P. Cattiaux, G. Conforti, I. Gentil, and C. Leonard, "Time reversal of diffusion processes under a finite entropy condition", 2021

- Forward process

$$\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$$

- Hörmander condition $\rightarrow$ allows for degenerate diffusions

$$\text{Span}(\text{Lie}\{h(x), g_1(x), \ldots, g_m(x)\}) = \mathbb{R}^n, \quad \forall x \in \mathbb{R}^n$$

- Time-reversed process $\tilde{Z} := \{\tilde{Z}_t = Z_{T-t};\, 0 \leq t \leq T\}$

Time-reversal formula [Haussmann & Pardoux, 1986]

$$\mathrm{d}\widetilde{Z}_t = -h(\widetilde{Z}_t)\mathrm{d}t + s(T-t, \widetilde{Z}_t)\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t$$

where $\quad s(t,x) = \dfrac{1}{p(t,x)} \nabla \cdot (p(t,x) g g^\top(x)), \quad p(t,\cdot) := \text{pdf}(Z_t)$

B. D. Anderson, "Reverse-time diffusion equation models," Stochastic Processes and their Applications, vol. 12, no. 3, pp. 313–326, 1982
U. G. Haussmann and E. Pardoux, "Time reversal of diffusions," The Annals of Probability, pp. 1188–1205, 1986
P. Cattiaux, G. Conforti, I. Gentil, and C. Leonard, "Time reversal of diffusion processes under a finite entropy condition", 2021

- Forward process

$$\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$$

- Hörmander condition $\rightarrow$ allows for degenerate diffusions

$$\mathsf{Span}(\mathsf{Lie}\{h(x), g_1(x), \ldots, g_m(x)\}) = \mathbb{R}^n, \quad \forall x \in \mathbb{R}^n$$

- Time-reversed process $\tilde{Z} := \{\tilde{Z}_t = Z_{T-t};\, 0 \le t \le T\}$

Time-reversal formula [Haussmann & Pardoux, 1986]

$$\mathrm{d}\widetilde{Z}_t = -h(\widetilde{Z}_t)\mathrm{d}t + s(T-t, \widetilde{Z}_t)\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t$$

$$\text{where} \quad s(t,x) = \frac{1}{p(t,x)} \nabla \cdot (p(t,x) g g^\top(x)), \quad p(t,\cdot) := \mathsf{pdf}(Z_t)$$

B. D. Anderson, "Reverse-time diffusion equation models," Stochastic Processes and their Applications, vol. 12, no. 3, pp. 313–326, 1982
U. G. Haussmann and E. Pardoux, "Time reversal of diffusions," The Annals of Probability, pp. 1188–1205, 1986
P. Cattiaux, G. Conforti, I. Gentil, and C. Leonard, "Time reversal of diffusion processes under a finite entropy condition", 2021

- Forward process

$$\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$$

- Hörmander condition $\rightarrow$ allows for degenerate diffusions

$$\text{Span}(\text{Lie}\{h(x), g_1(x), \ldots, g_m(x)\}) = \mathbb{R}^n, \quad \forall x \in \mathbb{R}^n$$

- Time-reversed process $\tilde{Z} := \{\tilde{Z}_t = Z_{T-t};\, 0 \le t \le T\}$

Time-reversal formula [Haussmann & Pardoux, 1986]

$$\mathrm{d}\widetilde{Z}_t = -h(\widetilde{Z}_t)\mathrm{d}t + s(T-t, \widetilde{Z}_t)\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t$$

where $s(t,x) = \frac{1}{p(t,x)} \nabla \cdot (p(t,x) g g^\top(x)), \quad p(t,\cdot) := \text{pdf}(Z_t)$

B. D. Anderson, "Reverse-time diffusion equation models," Stochastic Processes and their Applications, vol. 12, no. 3, pp. 313–326, 1982
U. G. Haussmann and E. Pardoux, "Time reversal of diffusions," The Annals of Probability, pp. 1188–1205, 1986
P. Cattiaux, G. Conforti, I. Gentil, and C. Leonard, "Time reversal of diffusion processes under a finite entropy condition", 2021

Time-reversal of diffusions
Score function approximation

- In order to numerically approximate

$$s(t,x) = \frac{1}{p(t,x)} \nabla \cdot (p(t,x) gg^\top(x))$$

- define the objective function

$$J(\psi) = \mathbb{E}[\|\psi(t,Z_t) - s(t,Z_t)\|^2]$$

- expand and apply the integration by parts

$$J(\psi) = \mathbb{E}\left[\|\psi(t,Z_t)\|^2 + 2\text{Tr}(gg^\top(Z_t)\nabla\psi(t,Z_t))\right] + \text{const.}$$

- Parameterize ψ and apply a stochastic optimization algorithm

A. Hyvärinen and P. Dayan, "Estimation of non-normalized statistical models by score matching." Journal of Machine Learning Research, vol. 6, no. 4, 2005
W. Tang and H. Zhao, "Score-based diffusion models via stochastic differential equations–a technical tutorial," arXiv:2402.07487,2024

Time-reversal of diffusions
Score function approximation

- In order to numerically approximate

$$s(t,x) = \frac{1}{p(t,x)} \nabla \cdot (p(t,x) g g^\top (x))$$

- define the objective function

$$J(\psi) = \mathbb{E}[\|\psi(t, Z_t) - s(t, Z_t)\|^2]$$

- expand and apply the integration by parts

$$J(\psi) = \mathbb{E}\left[\|\psi(t, Z_t)\|^2 + 2\mathrm{Tr}(g g^\top (Z_t) \nabla \psi(t, Z_t))\right] + \text{const.}$$

- Parameterize ψ and apply a stochastic optimization algorithm

A. Hyvärinen and P. Dayan, "Estimation of non-normalized statistical models by score matching." Journal of Machine Learning Research, vol. 6, no. 4, 2005
W. Tang and H. Zhao, "Score-based diffusion models via stochastic differential equations–a technical tutorial," arXiv:2402.07487,2024

- In order to numerically approximate

$$s(t,x) = \frac{1}{p(t,x)} \nabla \cdot (p(t,x) g g^\top(x))$$

- define the objective function

$$J(\psi) = \mathbb{E}[\|\psi(t, Z_t) - s(t, Z_t)\|^2]$$

- expand and apply the integration by parts

$$J(\psi) = \mathbb{E}\left[\|\psi(t, Z_t)\|^2 + 2\mathsf{Tr}(g g^\top(Z_t) \nabla \psi(t, Z_t))\right] + \text{const.}$$

- Parameterize ψ and apply a stochastic optimization algorithm

A. Hyvärinen and P. Dayan, "Estimation of non-normalized statistical models by score matching." Journal of Machine Learning Research, vol. 6, no. 4, 2005
W. Tang and H. Zhao, "Score-based diffusion models via stochastic differential equations–a technical tutorial," arXiv:2402.07487,2024

- In order to numerically approximate

$$s(t,x) = \frac{1}{p(t,x)} \nabla \cdot (p(t,x) g g^\top(x))$$

- define the objective function

$$J(\psi) = \mathbb{E}[\|\psi(t, Z_t) - s(t, Z_t)\|^2]$$

- expand and apply the integration by parts

$$J(\psi) = \mathbb{E}\left[\|\psi(t, Z_t)\|^2 + 2\mathsf{Tr}(g g^\top(Z_t) \nabla \psi(t, Z_t))\right] + \text{const.}$$

- Parameterize ψ and apply a stochastic optimization algorithm

---

A. Hyvärinen and P. Dayan, "Estimation of non-normalized statistical models by score matching." Journal of Machine Learning Research, vol. 6, no. 4, 2005
W. Tang and H. Zhao, "Score-based diffusion models via stochastic differential equations–a technical tutorial," arXiv:2402.07487,2024

Outline

- **Part 0:** Time-reversal of diffusions
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Outline

- **Part 0:** Time-reversal of diffusions
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

$$x_0 \quad \xrightarrow{\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(U_t\mathrm{d}t + \mathrm{d}W_t)} \quad x_f$$

Exact steering: find a control law $U_t = k(t, X_t)$ so that $X_T = x_f$.

Approximate steering: find a control law $U_t = k(t, X_t)$ so that $\mathbb{E}[\|X_T - x_f\|^2] \leq \epsilon$.

Can we use time-reversal method to solve the problem?

Point to point steering
Problem formulation

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(U_t\mathrm{d}t + \mathrm{d}W_t)$$

x_0 ⟶ x_f

Exact steering: find a control law $U_t = k(t, X_t)$ so that $X_T = x_f$.

Approximate steering: find a control law $U_t = k(t, X_t)$ so that $\mathbb{E}[\|X_T - x_f\|^2] \leq \epsilon$.

Can we use time-reversal method to solve the problem?

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(U_t\mathrm{d}t + \mathrm{d}W_t)$$

x_0 → x_f

Exact steering: find a control law $U_t = k(t, X_t)$ so that $X_T = x_f$.

Approximate steering: find a control law $U_t = k(t, X_t)$ so that $\mathbb{E}[\|X_T - x_f\|^2] \leq \epsilon$.

Can we use time-reversal method to solve the problem?

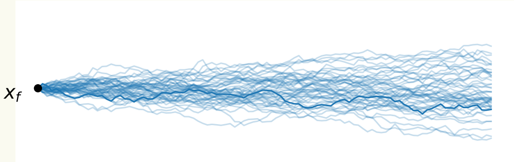

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

Point to point steering
Proposed methodology

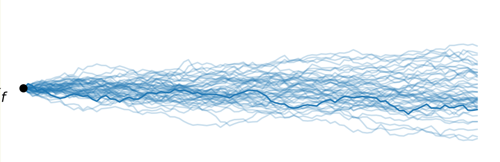

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

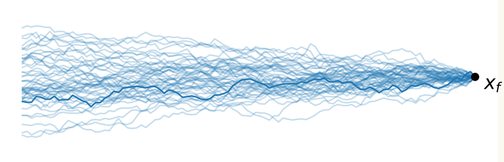

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

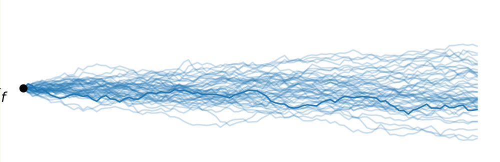

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

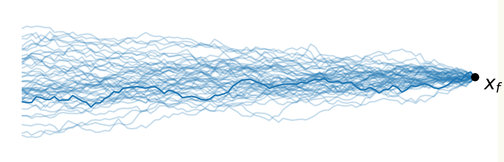

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

$$\boxed{k^*(t,x) = g(x)^\top \nabla \log(p(T-t,x))}$$

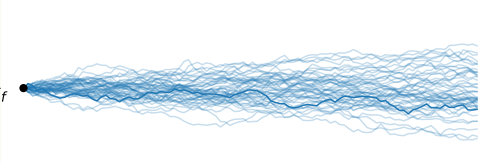

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

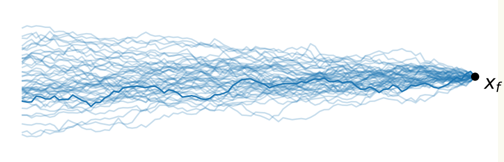

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

$$\boxed{k^*(t,x) = g(x)^\top \nabla \log(p(T-t,x))}$$

Does it solve the exact steering problem?

Point to point steering
Main result

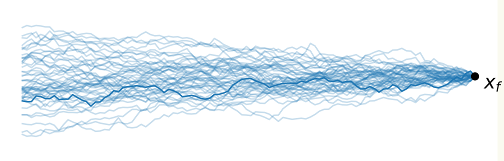

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

Actual process: $\mathrm{d}X_t = [f(X_t) + g(X_t)k^*(t, X_t)]\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$

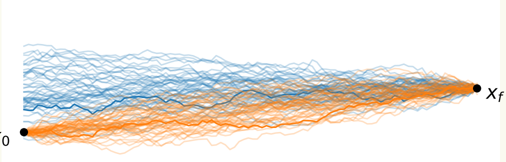

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t) k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

Actual process: $\mathrm{d}X_t = [f(X_t) + g(X_t) k^*(t, X_t)]\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$

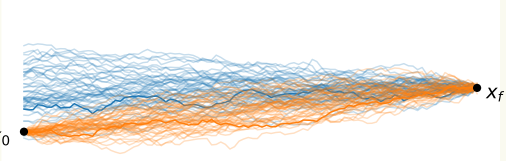

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

Actual process: $\mathrm{d}X_t = [f(X_t) + g(X_t)k^*(t, X_t)]\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$

Solution to the exact steering problem

If $\widetilde{P}_0(x_0) > 0$, then $k^*(t,x)$ solves the exact steering problem, i.e.

$$X_T = x_f \quad \text{a.s.}$$

- Model

$$\mathrm{d}X_t = AX_t + B(U_t + \mathrm{d}W_t), \quad X_0 = x_0$$

- Hörmander condition $\Rightarrow$ (A, B) is a controllable pair
- Auxiliary process

$$\mathrm{d}Z_t = -AZ_t + B\mathrm{d}W_t, \quad Z_0 = x_f$$

$$\Rightarrow \quad Z_t \sim \mathcal{N}(m_t, \Sigma_t)$$

- Resulting control law

$$k(t,x) = -B^\top \Sigma_{T-t}^{-1}(x - m_{T-t})$$

- Special case $A = 0$ and $B = 1$:

$$\mathrm{d}X_t = \frac{1}{T-t}(x_f - X_t) + \mathrm{d}W_t \quad \rightarrow \quad \text{Brownian bridge}$$

Point to point steering
Linear setting

- Model

$$\mathrm{d}X_t = AX_t + B(U_t + \mathrm{d}W_t), \quad X_0 = x_0$$

- Hörmander condition $\Rightarrow$ (A, B) is a controllable pair
- Auxiliary process

$$\mathrm{d}Z_t = -AZ_t + B\mathrm{d}W_t, \quad Z_0 = x_f$$

$$\Rightarrow \quad Z_t \sim \mathcal{N}(m_t, \Sigma_t)$$

- Resulting control law

$$k(t, x) = -B^\top \Sigma_{T-t}^{-1}(x - m_{T-t})$$

- Special case $A = 0$ and $B = 1$:

$$\mathrm{d}X_t = \frac{1}{T-t}(x_f - X_t) + \mathrm{d}W_t \quad \rightarrow \quad \text{Brownian bridge}$$

Point to point steering
Linear setting

- Model

$$\mathrm{d}X_t = AX_t + B(U_t + \mathrm{d}W_t), \quad X_0 = x_0$$

- Hörmander condition $\quad\Rightarrow\quad$ (A, B) is a controllable pair
- Auxiliary process

$$\mathrm{d}Z_t = -AZ_t + B\mathrm{d}W_t, \quad Z_0 = x_f$$

$$\Rightarrow \quad Z_t \sim \mathcal{N}(m_t, \Sigma_t)$$

- Resulting control law

$$k(t,x) = -B^\top \Sigma_{T-t}^{-1}(x - m_{T-t})$$

- Special case $A = 0$ and $B = 1$:

$$\mathrm{d}X_t = \frac{1}{T-t}(x_f - X_t) + \mathrm{d}W_t \quad \rightarrow \quad \text{Brownian bridge}$$

- Model

$$\mathrm{d}X_t = AX_t + B(U_t + \mathrm{d}W_t), \quad X_0 = x_0$$

- Hörmander condition $\Rightarrow$ (A, B) is a controllable pair
- Auxiliary process

$$\mathrm{d}Z_t = -AZ_t + B\mathrm{d}W_t, \quad Z_0 = x_f$$

$$\Rightarrow \quad Z_t \sim \mathcal{N}(m_t, \Sigma_t)$$

- Resulting control law

$$k(t, x) = -B^\top \Sigma_{T-t}^{-1}(x - m_{T-t})$$

- Special case $A = 0$ and $B = 1$:

$$\mathrm{d}X_t = \frac{1}{T-t}(x_f - X_t) + \mathrm{d}W_t \quad \rightarrow \quad \text{Brownian bridge}$$

Point to point steering
Linear setting

- Model

$$\mathrm{d}X_t = AX_t + B(U_t + \mathrm{d}W_t), \quad X_0 = x_0$$

- Hörmander condition $\Rightarrow$ (A, B) is a controllable pair
- Auxiliary process

$$\mathrm{d}Z_t = -AZ_t + B\mathrm{d}W_t, \quad Z_0 = x_f$$

$$\Rightarrow \quad Z_t \sim \mathcal{N}(m_t, \Sigma_t)$$

- Resulting control law

$$k(t, x) = -B^\top \Sigma_{T-t}^{-1}(x - m_{T-t})$$

- Special case $A = 0$ and $B = 1$:

$$\mathrm{d}X_t = \frac{1}{T-t}(x_f - X_t) + \mathrm{d}W_t \quad \rightarrow \quad \text{Brownian bridge}$$

- Singularity of the control law: if $x \neq x_f$

$$k(t,x) = g(x)^\top \nabla \log(p(T-t,x)) \to \infty \quad \text{as} \quad t \to T$$

- Regularize the initial distribution of the auxiliary process:

$$Z_0 \sim \mathcal{N}(x_f, \delta I) \quad \text{instead of} \quad Z_0 = x_f$$

- Denote the resulting control law and trajectory by $k^\delta(t,x)$ and X_t^δ.

Accuracy of the regularized control in the linear Gaussian setting

$k^\delta(t,x)$ solves the approximate steering problem. In particular,

$$\mathbb{E}[\|X_T^\delta - x_f\|^2] \leq \delta^2 \|e^{TA}x_0 - x_f\|_{M^2}^2 + \delta(n - \mathrm{Tr}(M)) \to 0 \quad \text{as} \quad \delta \to 0$$

where $M = (\delta I + \int_0^T e^{tA} B B^\top e^{tA^\top} \mathrm{d}t)^{-1}$.

- Singularity of the control law: if $x \neq x_f$

$$k(t,x) = g(x)^\top \nabla \log(p(T-t,x)) \to \infty \quad \text{as} \quad t \to T$$

- Regularize the initial distribution of the auxiliary process:

$$Z_0 \sim \mathcal{N}(x_f, \delta I) \quad \text{instead of} \quad Z_0 = x_f$$

- Denote the resulting control law and trajectory by $k^\delta(t,x)$ and X_t^δ.

Accuracy of the regularized control in the linear Gaussian setting

$k^\delta(t,x)$ solves the approximate steering problem. In particular,

$$\mathbb{E}[\|X_T^\delta - x_f\|^2] \le \delta^2 \|e^{TA}x_0 - x_f\|_{M^2}^2 + \delta(n - \mathsf{Tr}(M)) \to 0 \quad \text{as} \quad \delta \to 0$$

where $M = (\delta I + \int_0^T e^{tA} BB^\top e^{tA^\top} \mathrm{d}t)^{-1}$.

- Singularity of the control law: if $x \neq x_f$

$$k(t,x) = g(x)^\top \nabla \log(p(T-t,x)) \to \infty \quad \text{as} \quad t \to T$$

- Regularize the initial distribution of the auxiliary process:

$$Z_0 \sim \mathcal{N}(x_f, \delta I) \quad \text{instead of} \quad Z_0 = x_f$$

- Denote the resulting control law and trajectory by $k^\delta(t,x)$ and X_t^δ.

Accuracy of the regularized control in the linear Gaussian setting

$k^\delta(t,x)$ solves the approximate steering problem. In particular,

$$\mathbb{E}[\|X_T^\delta - x_f\|^2] \leq \delta^2 \|e^{TA}x_0 - x_f\|_{M^2}^2 + \delta(n - \mathrm{Tr}(M)) \to 0 \quad \text{as} \quad \delta \to 0$$

where $M = (\delta I + \int_0^T e^{tA} B B^\top e^{tA^\top} \mathrm{d}t)^{-1}$.

Point to point steering
Avoiding singularity

- Singularity of the control law: if $x \neq x_f$

$$k(t,x) = g(x)^\top \nabla \log(p(T-t,x)) \to \infty \quad \text{as} \quad t \to T$$

- Regularize the initial distribution of the auxiliary process:

$$Z_0 \sim \mathcal{N}(x_f, \delta I) \quad \text{instead of} \quad Z_0 = x_f$$

- Denote the resulting control law and trajectory by $k^\delta(t,x)$ and X_t^δ.

Accuracy of the regularized control in the linear Gaussian setting

$k^\delta(t,x)$ solves the approximate steering problem. In particular,

$$\mathbb{E}[\|X_T^\delta - x_f\|^2] \leq \delta^2 \|e^{TA}x_0 - x_f\|_{M^2}^2 + \delta(n - \mathsf{Tr}(M)) \to 0 \quad \text{as} \quad \delta \to 0$$

where $M = (\delta I + \int_0^T e^{tA} BB^\top e^{tA^\top} \mathrm{d}t)^{-1}$.

Point to point steering
Optimality and relationship to diffusion bridges

- Diffusion process (with no control)

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$$

- Condition on the event that $\{X_T = x_f\}$.
- The conditioned process satisfies (Doob's h-transform)

$$\mathrm{d}\widetilde{X}_t = f(\widetilde{X}_t)\mathrm{d}t + g(\widetilde{X}_t)\nabla \log P(X_t = x | X_T = x_f)\mathrm{d}t + g(\widetilde{X}_t)\mathrm{d}W_t, \quad \widetilde{X}_0 = x_0$$

- The additional term also serves as a control that ensures $X_T = x_f$
- Our proposed control law is different in general, but identical in the linear setting

J. Heng, V. De Bortoli, A. Doucet, and J. Thornton, "Simulating diffusion bridges with score matching," Biometrika, p. asaf048, 07 2025

Point to point steering

Optimality and relationship to diffusion bridges

- Diffusion process (with no control)

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$$

- Condition on the event that $\{X_T = x_f\}$.
- The conditioned process satisfies (Doob's h-transform)

$$\mathrm{d}\widetilde{X}_t = f(\widetilde{X}_t)\mathrm{d}t + g(\widetilde{X}_t)\nabla \log P(X_t = x | X_T = x_f)\mathrm{d}t + g(\widetilde{X}_t)\mathrm{d}W_t, \quad \widetilde{X}_0 = x_0$$

- The additional term also serves as a control that ensures $X_T = x_f$
- Our proposed control law is different in general, but identical in the linear setting

J. Heng, V. De Bortoli, A. Doucet, and J. Thornton, "Simulating diffusion bridges with score matching," Biometrika, p. asaf048, 07 2025

Point to point steering

Optimality and relationship to diffusion bridges

- Diffusion process (with no control)

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$$

- Condition on the event that $\{X_T = x_f\}$.
- The conditioned process satisfies (Doob's h-transform)

$$\mathrm{d}\widetilde{X}_t = f(\widetilde{X}_t)\mathrm{d}t + g(\widetilde{X}_t)\nabla \log P(X_t = x | X_T = x_f)\mathrm{d}t + g(\widetilde{X}_t)\mathrm{d}W_t, \quad \widetilde{X}_0 = x_0$$

- The additional term also serves as a control that ensures $X_T = x_f$
- Our proposed control law is different in general, but identical in the linear setting

J. Heng, V. De Bortoli, A. Doucet, and J. Thornton, "Simulating diffusion bridges with score matching," Biometrika, p. asaf048, 07 2025

Point to point steering
Optimality and relationship to diffusion bridges

- Diffusion process (with no control)

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$$

- Condition on the event that $\{X_T = x_f\}$.
- The conditioned process satisfies (Doob's h-transform)

$$\mathrm{d}\widetilde{X}_t = f(\widetilde{X}_t)\mathrm{d}t + g(\widetilde{X}_t)\nabla \log P(X_t = x | X_T = x_f)\mathrm{d}t + g(\widetilde{X}_t)\mathrm{d}W_t, \quad \widetilde{X}_0 = x_0$$

- The additional term also serves as a control that ensures $X_T = x_f$
- Our proposed control law is different in general, but identical in the linear setting

J. Heng, V. De Bortoli, A. Doucet, and J. Thornton, "Simulating diffusion bridges with score matching," Biometrika, p. asaf048, 07 2025

Point to point steering

Optimality and relationship to diffusion bridges

- Diffusion process (with no control)

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)\mathrm{d}W_t, \quad X_0 = x_0$$

- Condition on the event that $\{X_T = x_f\}$.
- The conditioned process satisfies (Doob's h-transform)

$$\mathrm{d}\widetilde{X}_t = f(\widetilde{X}_t)\mathrm{d}t + g(\widetilde{X}_t)\nabla \log P(X_t = x | X_T = x_f)\mathrm{d}t + g(\widetilde{X}_t)\mathrm{d}W_t, \quad \widetilde{X}_0 = x_0$$

- The additional term also serves as a control that ensures $X_T = x_f$
- Our proposed control law is different in general, but identical in the linear setting

J. Heng, V. De Bortoli, A. Doucet, and J. Thornton, "Simulating diffusion bridges with score matching," Biometrika, p. asaf048, 07 2025

Auxiliary process:

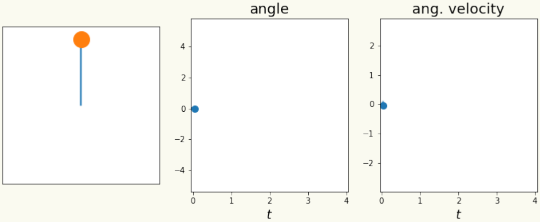

Actual controlled process:

Auxiliary process:

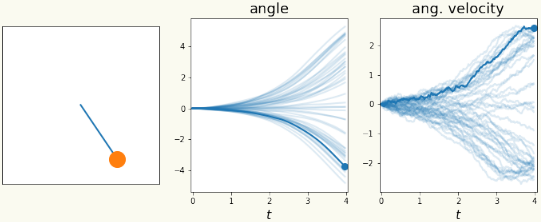

Actual controlled process:

Auxiliary process:

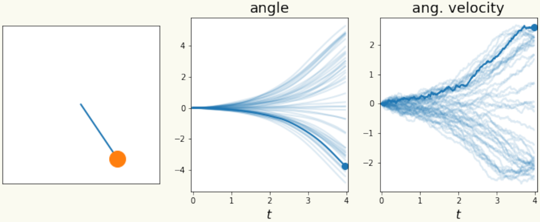

Actual controlled process:

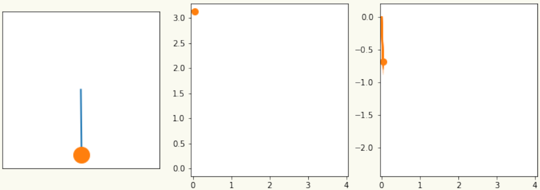

Point to point steering
Numerical demonstration with inverted pendulum

Auxiliary process:

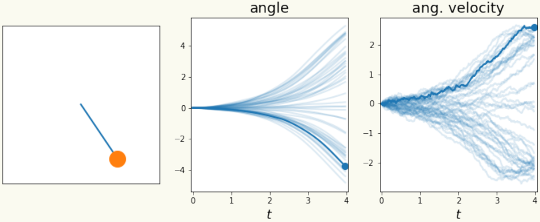

Actual controlled process:

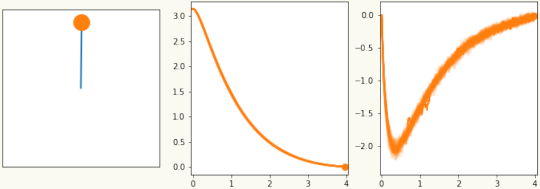

Outline

- **Part 0:** Time-reversal of diffusions
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Outline

- **Part 0:** Time-reversal of diffusions
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Distribution to distribution steering

Flow matching methodology

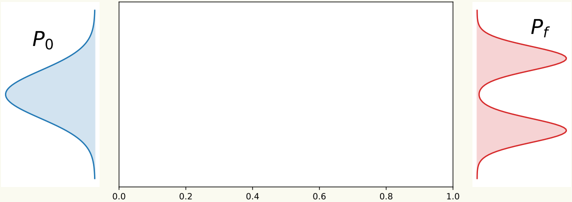

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t,X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0, \quad X_T \sim X_f$$

$$\mathrm{d}X_t^i = f(X_t^i)\mathrm{d}t + g(X_t)(U_t^i\mathrm{d}t + \mathrm{d}W_t), \quad X_0^i = x_0^i, \quad X_T^i = x_f^i$$

$$k(t,x) = \arg\min_{\psi} \mathbb{E}[\|\psi(t,X_t^i) - U_t^i\|^2], \quad (x_0^i, x_f^i) \sim P_0 \otimes P_f$$

Distribution to distribution steering
Flow matching methodology

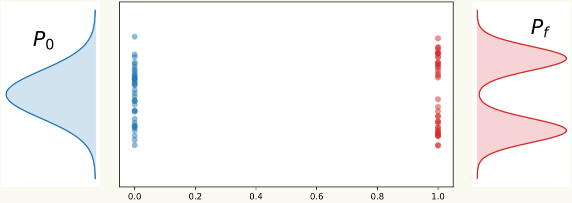

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t,X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0, \quad X_T \sim X_f$$

$$\mathrm{d}X_t^i = f(X_t^i)\mathrm{d}t + g(X_t)(U_t^i\mathrm{d}t + \mathrm{d}W_t), \quad X_0^i = x_0^i, \quad X_T^i = x_f^i$$

$$k(t,x) = \arg\min_{\psi} \mathbb{E}[\|\psi(t,X_t^i) - U_t^i\|^2], \quad (x_0^i, x_f^i) \sim P_0 \otimes P_f$$

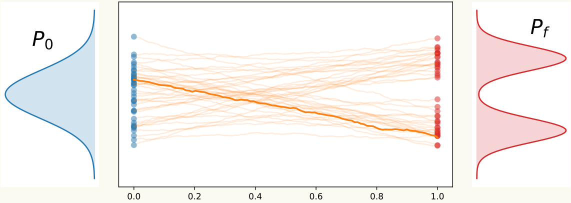

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0, \quad X_T \sim X_f$$

$$\mathrm{d}X_t^i = f(X_t^i)\mathrm{d}t + g(X_t)(U_t^i\mathrm{d}t + \mathrm{d}W_t), \quad X_0^i = x_0^i, \quad X_T^i = x_f^i$$

$$k(t,x) = \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^i) - U_t^i\|^2], \quad (x_0^i, x_f^i) \sim P_0 \otimes P_f$$

Distribution to distribution steering

Flow matching methodology

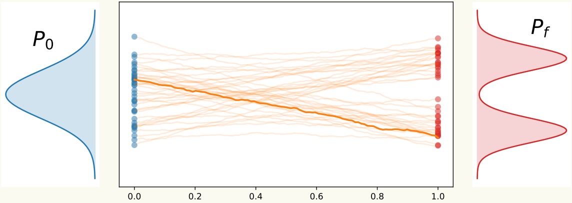

$$dX_t = f(X_t)dt + g(X_t)(k(t, X_t)dt + dW_t), \quad X_0 \sim P_0, \quad X_T \sim X_f$$

$$dX_t^i = f(X_t^i)dt + g(X_t)(U_t^i dt + dW_t), \quad X_0^i = x_0^i, \quad X_T^i = x_f^i$$

$$k(t, x) = \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^i) - U_t^i\|^2], \quad (x_0^i, x_f^i) \sim P_0 \otimes P_f$$

Distribution to distribution steering

Flow matching methodology

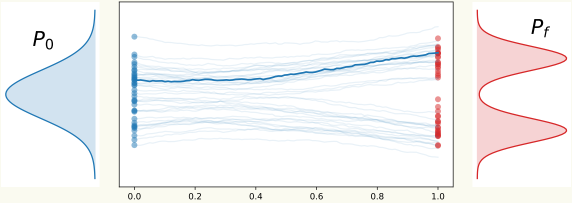

$$\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0, \quad X_T \sim X_f$$

$$\mathrm{d}X_t^i = f(X_t^i)\mathrm{d}t + g(X_t)(U_t^i\mathrm{d}t + \mathrm{d}W_t), \quad X_0^i = x_0^i, \quad X_T^i = x_f^i$$

$$k(t,x) = \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^i) - U_t^i\|^2], \quad (x_0^i, x_f^i) \sim P_0 \otimes P_f$$

Distribution to distribution steering

Numerical demonstration

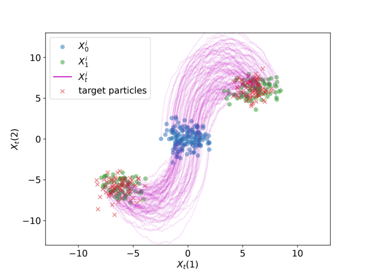

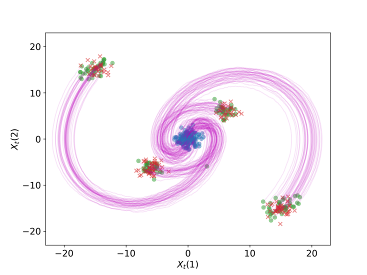

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Thank you for your attention!

Joint work with:

   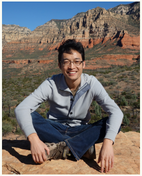

Yuhang Mei | Mohammad Al-Jarrah | Ali Pakniyat | Yongxin Chen

References:

- *A Time-Reversal Control Synthesis for Steering the State of Stochastic Systems*
  Yuhang Mei, Amirhossein Taghvaei, Ali Pakniyat
  IEEE Conference on Decision and Control (CDC), 2025
- *Flow matching for stochastic linear control systems*
  Yuhang Mei, Mohammad Al-Jarrah, Amirhossein Taghvaei, Yongxin Chen
  7th Annual Learning for Dynamics & Control Conference (L4DC), 2025