From diffusion models to stochastic control: a time-reversal methodology for feedback control design

Amirhossein Taghvaei

Seminars in Applied Mathematicss
University of Washington, Seattle

Oct 15, 2025

Collaboration with the Applied Math department
Optimal transportation methods in nonlinear filtering


M. Al-Jarrah


M. martino


N. Jin


A. Hsu


B. Hosseini

- *An optimal transport formulation of Bayes' law for nonlinear filtering algorithms*
  IEEE Conference on Decision and Control (CDC), Cancun, 2022
- *Nonlinear Filtering with Brenier Optimal Transport Maps*
  International Conference of Machine Learning (ICML), Vienna, 2024
- *Data-Driven Approximation of Stationary Nonlinear Filters with Optimal Transport Maps*
  IEEE Conference on Decision and Control (CDC), Milan, 2024
- Conditional Optimal Transport on Function Spaces
  SIAM/ASA Journal on Uncertainty Quantification 13.1 (2025): 304-338.

This talk is about stochastic control

Collaboration with the Applied Math department

Optimal transportation methods in nonlinear filtering

M. Al-Jarrah

M. martino

N. Jin

A. Hsu

B. Hosseini

- *An optimal transport formulation of Bayes' law for nonlinear filtering algorithms*
  IEEE Conference on Decision and Control (CDC), Cancun, 2022
- *Nonlinear Filtering with Brenier Optimal Transport Maps*
  International Conference of Machine Learning (ICML), Vienna, 2024
- *Data-Driven Approximation of Stationary Nonlinear Filters with Optimal Transport Maps*
  IEEE Conference on Decision and Control (CDC), Milan, 2024
- Conditional Optimal Transport on Function Spaces
  SIAM/ASA Journal on Uncertainty Quantification 13.1 (2025): 304-338.

This talk is about stochastic control

Research theme: “Controlling” probability distributions
Applications

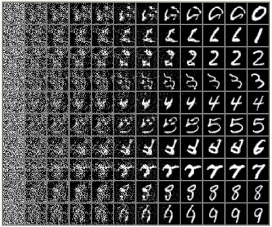

Generative models

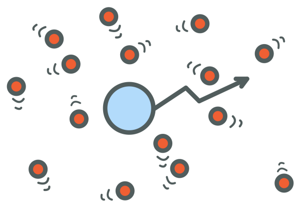

Stochastic thermodynamics

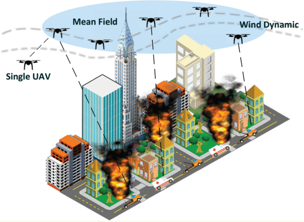

Robotic swarms

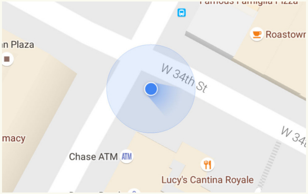

Navigation

Research theme: "Controlling" probability distributions
Literature

History:

- Controlling Liouville equation [Brocket 1997]
- Mean-field games [Lasry & Lions, 2006] [Huang et. al., 2006]
- Mean-field optimal control [Bensoussan 2013] [Carmona & Delarue, 2018]
- Optimal Transport and Schrodinger bridges [Chen et. al. 2016]
- ...

Main Challenge: numerical methods for non-Gaussian and nonlinear settings

Research objective: bridge the theory-computation gap

History:

- Controlling Liouville equation [Brocket 1997]
- Mean-field games [Lasry & Lions, 2006] [Huang et. al., 2006]
- Mean-field optimal control [Bensoussan 2013] [Carmona & Delarue, 2018]
- Optimal Transport and Schrodinger bridges [Chen et. al. 2016]
- ...

Main Challenge: numerical methods for non-Gaussian and nonlinear settings

Research objective: bridge the theory-computation gap

History:

- Controlling Liouville equation [Brocket 1997]
- Mean-field games [Lasry & Lions, 2006] [Huang et. al., 2006]
- Mean-field optimal control [Bensoussan 2013] [Carmona & Delarue, 2018]
- Optimal Transport and Schrodinger bridges [Chen et. al. 2016]
- . . .

Main Challenge: numerical methods for non-Gaussian and nonlinear settings

Research objective: bridge the theory-computation gap

Overview of the talk

Part 1: point to point steering

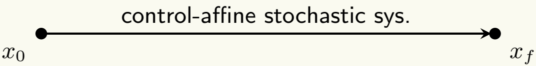

Part 2: distribution to distribution steering

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

Overview of the talk

Part 1: point to point steering

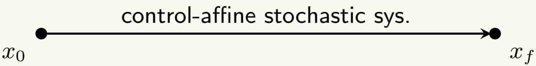

Part 2: distribution to distribution steering

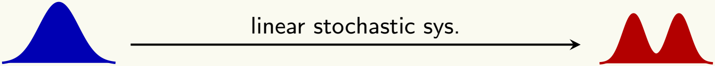

- Focus on computability rather than optimality
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview of the talk

Part 1: point to point steering

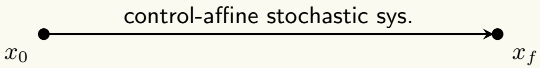

Part 2: distribution to distribution steering

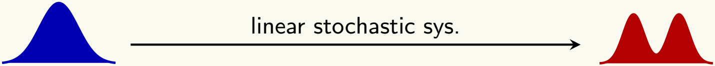

- Focus on computability rather than optimality
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview of the talk

Part 1: point to point steering

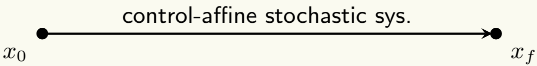

Part 2: distribution to distribution steering

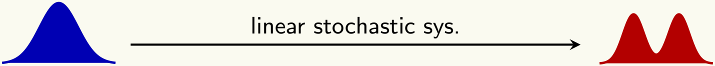

- Focus on computability rather than optimality
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview of the talk

Part 1: point to point steering

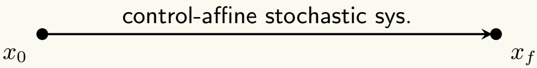

Part 2: distribution to distribution steering

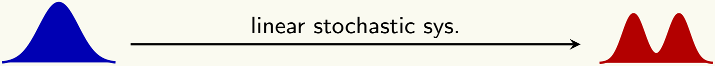

- Focus on computability rather than optimality
- **Methodology:** Time-reversal and flow matching
- **Algorithm:** Simulate and solve a nonlinear regression problem

Outline

- **Part 0:** Problem setup
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Outline

- **Part 0:** Problem setup
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Model:

$$\dot{x}(t) = u(t), \quad x(0) = x_0$$

- $x(t) \in \mathbb{R}^n$ is the state
- $u(t) \in \mathbb{R}^n$ is the control input

Objective: For any target x_f, find $\{u(t); 0 \leq t \leq 1\}$ so that $x(1) = x_f$.

Solution:

$$u(t) = x_f - x_0$$

Model:

$$\dot{x}(t) = u(t), \quad x(0) = x_0$$

- $x(t) \in \mathbb{R}^n$ is the state
- $u(t) \in \mathbb{R}^n$ is the control input

Objective: For any target x_f, find $\{u(t); 0 \leq t \leq 1\}$ so that $x(1) = x_f$.

Solution:

$$u(t) = x_f - x_0$$

Problem setup
Point to point steering (simple case)

Model:

$$\dot{x}(t) = u(t), \quad x(0) = x_0$$

- $x(t) \in \mathbb{R}^n$ is the state
- $u(t) \in \mathbb{R}^n$ is the control input

Objective: For any target x_f, find $\{u(t); 0 \leq t \leq 1\}$ so that $x(1) = x_f$.

Solution:

$$u(t) = x_f - x_0$$

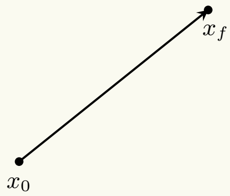

Model:

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0$$

- $x(t) \in \mathbb{R}^n$ is the state
- $u(t) \in \mathbb{R}^m$ is the control input (typically $m < n$)

Objective: For any target x_f, find $\{u(t); 0 \leq t \leq 1\}$ so that $x(1) = x_f$.

Solution:

$$u(t) = B^\top e^{(1-t)A^\top} \Phi_1^{-1}(x_f - e^A x_0)$$

$$\Phi_1 := \int_0^1 e^{tA} B B^\top e^{tA^\top} \, \mathrm{d}t \quad \text{(Cont. Gramian)}$$

(A, B) is controllable $\Leftrightarrow$ Φ_1 is invertible $\Leftrightarrow$ $[B; AB; \ldots; A^{n-1}B]$ is full-rank

Model:

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0$$

- $x(t) \in \mathbb{R}^n$ is the state
- $u(t) \in \mathbb{R}^m$ is the control input (typically $m < n$)

Objective: For any target x_f, find $\{u(t); 0 \le t \le 1\}$ so that $x(1) = x_f$.

Solution:

$$u(t) = B^\top e^{(1-t)A^\top} \Phi_1^{-1}(x_f - e^A x_0)$$

$$\Phi_1 := \int_0^1 e^{tA} B B^\top e^{tA^\top} \mathrm{d}t \quad \text{(Cont. Gramian)}$$

(A, B) is controllable $\Leftrightarrow$ Φ_1 is invertible $\Leftrightarrow$ $[B; AB; \dots; A^{n-1}B]$ is full-rank

Problem setup
Point to point steering (linear system)

Model:

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0$$

- $x(t) \in \mathbb{R}^n$ is the state
- $u(t) \in \mathbb{R}^m$ is the control input (typically $m < n$)

Objective: For any target x_f, find $\{u(t); 0 \leq t \leq 1\}$ so that $x(1) = x_f$.

Solution:

$$u(t) = B^\top e^{(1-t)A^\top} \Phi_1^{-1}(x_f - e^A x_0)$$

$$\Phi_1 := \int_0^1 e^{tA} BB^\top e^{tA^\top} \mathrm{d}t \quad \text{(Cont. Gramian)}$$

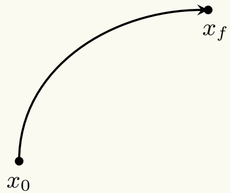

(A, B) is controllable $\Leftrightarrow$ Φ_1 is invertible $\Leftrightarrow$ $[B; AB; \ldots; A^{n-1}B]$ is full-rank

Model:

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0$$

- $x(t) \in \mathbb{R}^n$ is the state
- $u(t) \in \mathbb{R}^m$ is the control input (typically $m < n$)

Objective: For any target x_f, find $\{u(t); 0 \le t \le 1\}$ so that $x(1) = x_f$.

Solution:

$$u(t) = B^\top e^{(1-t)A^\top} \Phi_1^{-1}(x_f - e^A x_0)$$

$$\Phi_1 := \int_0^1 e^{tA} BB^\top e^{tA^\top} \mathrm{d}t \quad \text{(Cont. Gramian)}$$

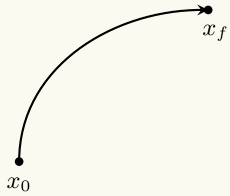

(A, B) is controllable $\Leftrightarrow$ Φ_1 is invertible $\Leftrightarrow$ $[B; AB; \ldots; A^{n-1}B]$ is full-rank

Problem setup

Point to point steering (stochastic linear system)

Model:

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

- $W_t \in \mathbb{R}^m$ is the Wiener process
- $\mathcal{F}_t := \sigma(W_s; 0 \le s \le t)$

Objective: For any target x_f, find $\mathcal{F}_t$-adapted U_t so that $X_1 = x_f$.

Solution:

$$U_t = B^\top e^{(1-t)A^\top}\Phi_{1-t}^{-1}(x_f - e^{(1-t)A}X_t) =: k(t, X_t)$$

$$\Phi_t := \int_0^t e^{sA}BB^\top e^{sA^\top}\mathrm{d}s \quad \text{(Cont. Gramian)}$$

$\rightarrow$ Feedback control structure

Problem setup

Point to point steering (stochastic linear system)

Model:

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

- $W_t \in \mathbb{R}^m$ is the Wiener process
- $\mathcal{F}_t := \sigma(W_s; 0 \le s \le t)$

Objective: For any target x_f, find $\mathcal{F}_t$-adapted U_t so that $X_1 = x_f$.

Solution:

$$U_t = B^\top e^{(1-t)A^\top} \Phi_{1-t}^{-1}(x_f - e^{(1-t)A} X_t) =: k(t, X_t)$$

$$\Phi_t := \int_0^t e^{sA} BB^\top e^{sA^\top} \mathrm{d}s \quad \text{(Cont. Gramian)}$$

$\rightarrow$ Feedback control structure

Problem setup
Point to point steering (stochastic linear system)

Model:

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

- $W_t \in \mathbb{R}^m$ is the Wiener process
- $\mathcal{F}_t := \sigma(W_s; 0 \le s \le t)$

Objective: For any target x_f, find $\mathcal{F}_t$-adapted U_t so that $X_1 = x_f$.

Solution:

$$U_t = B^\top e^{(1-t)A^\top}\Phi_{1-t}^{-1}(x_f - e^{(1-t)A}X_t) =: k(t, X_t)$$

$$\Phi_t := \int_0^t e^{sA}BB^\top e^{sA^\top}\mathrm{d}s \quad \text{(Cont. Gramian)}$$

→ Feedback control structure

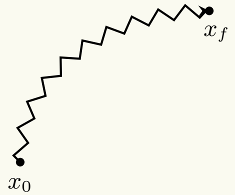

Problem setup

Point to point steering (stochastic control-affine)

Model:

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

- $W_t \in \mathbb{R}^m$ is the Wiener process
- $\mathcal{F}_t := \sigma(W_s; 0 \leq s \leq t)$

Objective: For any target x_f, find $\mathcal{F}_t$-adapted U_t so that $X_1 = x_f$.

Solution:

$U_t =$ Part (I) of the talk

Model:

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

- $W_t \in \mathbb{R}^m$ is the Wiener process
- $\mathcal{F}_t := \sigma(W_s; 0 \le s \le t)$

Objective: For any target x_f, find $\mathcal{F}_t$-adapted U_t so that $X_1 = x_f$.

Solution:

$U_t =$ Part (I) of the talk

Problem setup

Point to point steering (stochastic control-affine)

Model:

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

- $W_t \in \mathbb{R}^m$ is the Wiener process
- $\mathcal{F}_t := \sigma(W_s; 0 \leq s \leq t)$

Objective: For any target x_f, find $\mathcal{F}_t$-adapted U_t so that $X_1 = x_f$.

Solution:

$U_t =$ Part (I) of the talk

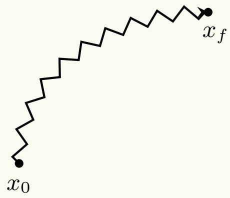

Problem setup
Distribution to distribution steering (stochastic control-affine)

Model:

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

- P_0 is the initial distribution
- P_f is the target distribution

Objective: Find $\mathcal{F}_t$-adapted U_t so that $X_1 \sim P_f$.

Solution:

$$U_t = \text{Part (II) of the talk}$$

Model:

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

- P_0 is the initial distribution
- P_f is the target distribution

Objective: Find $\mathcal{F}_t$-adapted U_t so that $X_1 \sim P_f$.

Solution:

$U_t =$ Part (II) of the talk

Problem setup

Distribution to distribution steering (stochastic control-affine)

Model:

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

- P_0 is the initial distribution
- P_f is the target distribution

Objective: Find $\mathcal{F}_t$-adapted U_t so that $X_1 \sim P_f$.

Solution:

$U_t =$ Part (II) of the talk

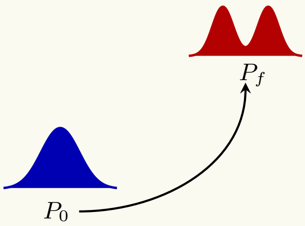

Outline

- **Part 0:** Problem setup
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Outline

- **Part 0:** Problem setup
- **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

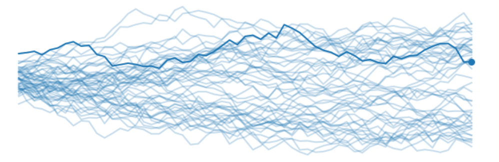

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

Time-reversal of diffusions

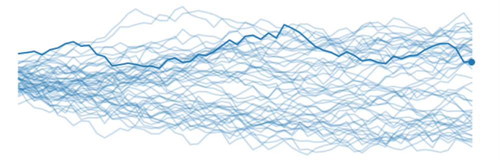

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

Time-reversal of diffusions

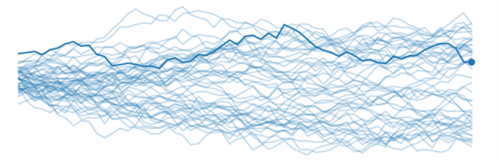

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

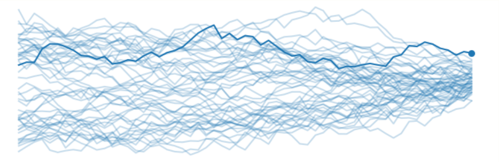

reversed process: $\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = ?$

Time-reversal of diffusions

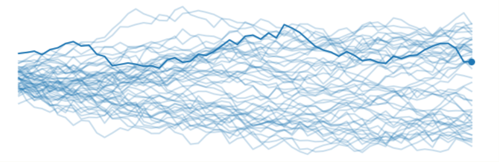

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

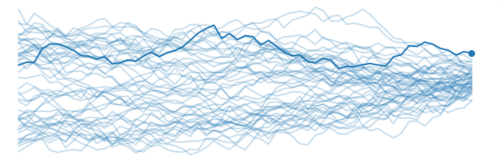

reversed process: $\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = ?$

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

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

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

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

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?

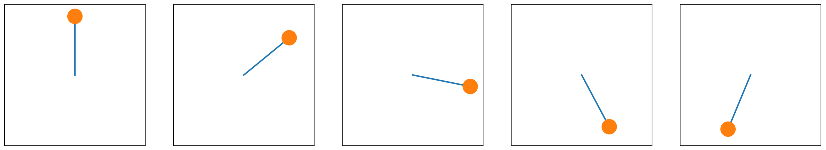

$$\xrightarrow{\mathrm{d}Z_t = \text{"pendulum dynamics"}}$$

Time-reversal of diffusions

Application in control?

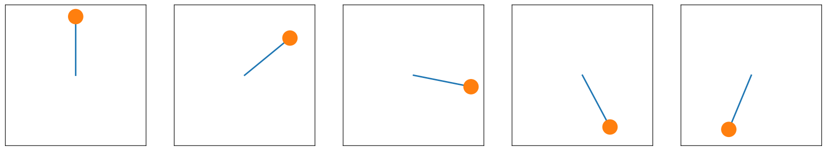

$\mathrm{d}Z_t =$ “pendulum dynamics”

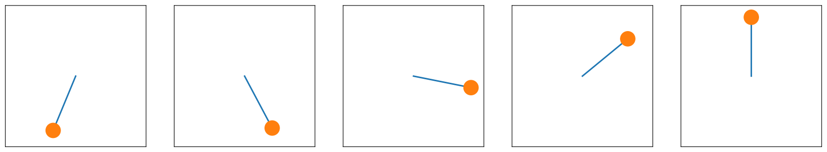

$\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t =$ “steer to upward position”

Time-reversal of diffusions
Application in control?

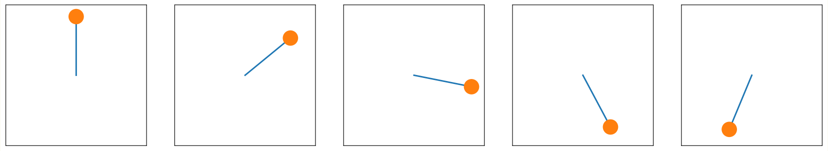

$\mathrm{d}Z_t =$ "pendulum dynamics"

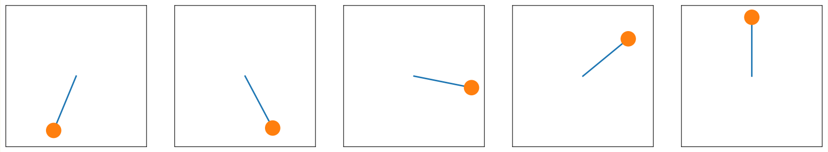

$\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t =$ "steer to upward position"

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

Time-reversal formula [Haussmann & Pardoux, 1986]

$$d\widetilde{Z}_t = [-h(\widetilde{Z}_t) + \nabla \cdot (gg^\top)(\widetilde{Z}_t)]dt + gg^\top(\widetilde{Z}_t)s(T-t, \widetilde{Z}_t)dt + g(\widetilde{Z}_t)dW_t$$

where $\quad s(t,x) = \nabla \log(p(t,x)) \quad p(t,\cdot) := \text{pdf}(Z_t)$

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

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

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

Time-reversal formula [Haussmann & Pardoux, 1986]

$$\mathrm{d}\widetilde{Z}_t = [-h(\widetilde{Z}_t) + \nabla \cdot (gg^\top)(\widetilde{Z}_t)]\mathrm{d}t + gg^\top(\widetilde{Z}_t)s(T-t, \widetilde{Z}_t)\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t$$

where $s(t,x) = \nabla \log(p(t,x)) \quad p(t,\cdot) := \text{pdf}(Z_t)$

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
Theory

- 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]

$$\mathrm{d}\widetilde{Z}_t = [-h(\widetilde{Z}_t) + \nabla \cdot (gg^\top)(\widetilde{Z}_t)]\mathrm{d}t + gg^\top(\widetilde{Z}_t)s(T-t, \widetilde{Z}_t)\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t$$

where $\quad s(t,x) = \nabla \log(p(t,x)) \quad p(t,\cdot) := \text{pdf}(Z_t)$

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

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

- 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]

$$d\widetilde{Z}_t = [-h(\widetilde{Z}_t) + \nabla \cdot (gg^\top)(\widetilde{Z}_t)]dt + gg^\top(\widetilde{Z}_t)s(T-t, \widetilde{Z}_t)dt + g(\widetilde{Z}_t)dW_t$$

where $s(t,x) = \nabla \log(p(t,x)) \quad p(t,\cdot) := \mathsf{pdf}(Z_t)$

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

Time-reversal of diffusions
Score function approximation

- In order to numerically approximate

$$s(t,x) = \nabla \log(p(t,x))$$

- define the objective function

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

- expand and apply the integration by parts

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

- Parameterize $s(t,x)$ 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

$$s(t, x) = \nabla \log(p(t, x))$$

- define the objective function

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

- expand and apply the integration by parts

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

- Parameterize $s(t, x)$ 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

$$s(t,x) = \nabla \log(p(t,x))$$

- define the objective function

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

- expand and apply the integration by parts

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

- Parameterize $s(t,x)$ 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

$$s(t,x) = \nabla \log(p(t,x))$$

- define the objective function

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

- expand and apply the integration by parts

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

- Parameterize $s(t,x)$ 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

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

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

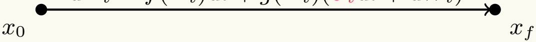

Exact steering: find a control law $U_t = k(t, X_t)$ so that $X_T = x_f$.

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?

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

$x_0 \longrightarrow x_f$

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
Proposed methodology

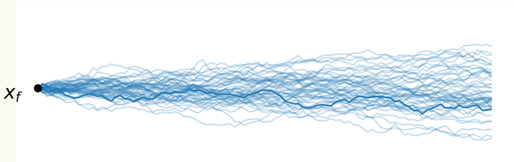

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

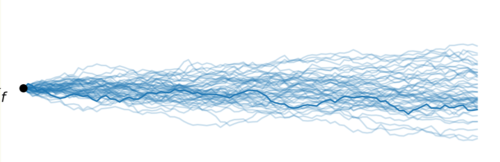

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

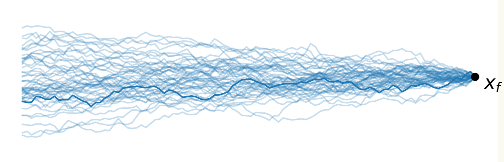

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

Point to point steering
Proposed methodology

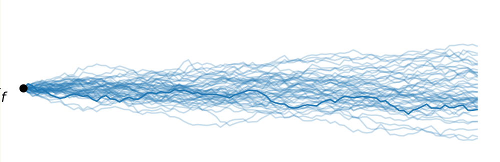

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

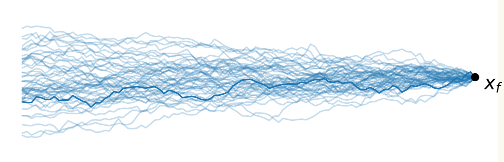

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

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

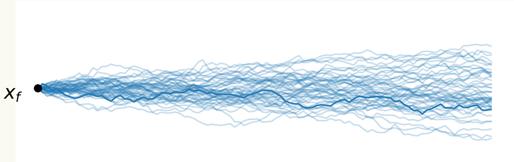

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

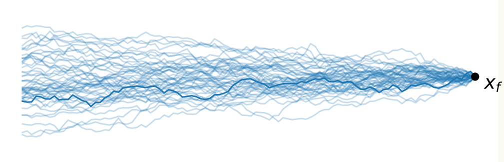

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

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

Does it solve the exact steering problem?

Point to point steering
Main result

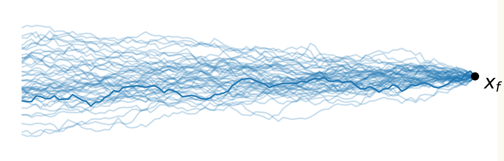

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t,\widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

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

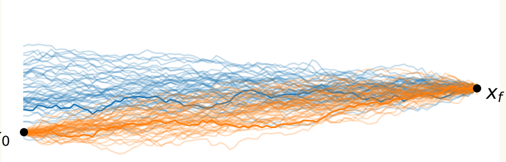

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

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

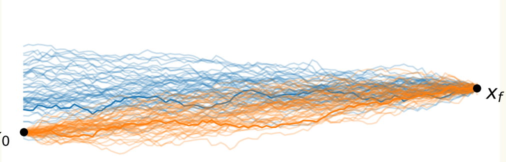

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

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

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

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

- 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

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

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

- 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

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

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

- 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

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

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

- 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

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

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

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

- 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,

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

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

Point to point steering
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,

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

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

Auxiliary process:

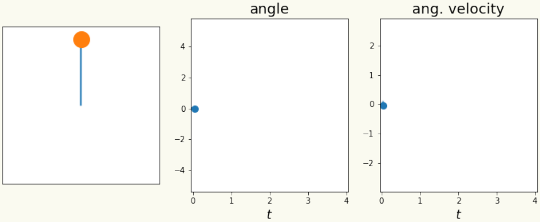

Actual controlled process:

Point to point steering
Numerical demonstration with inverted pendulum

Auxiliary process:

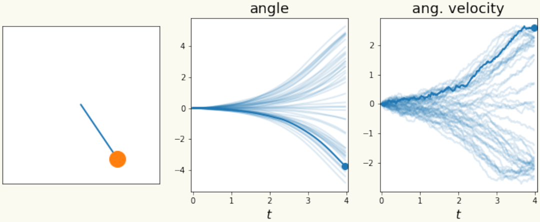

Actual controlled process:

Auxiliary process:

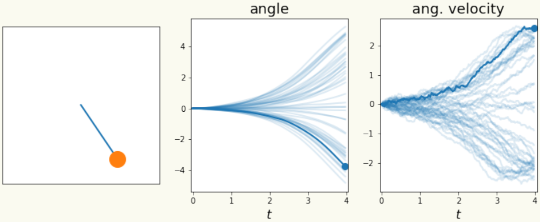

Actual controlled process:

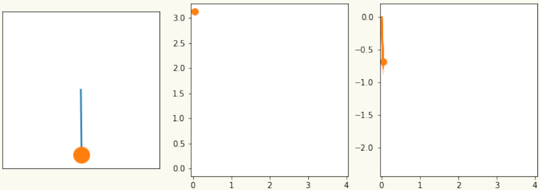

Point to point steering
Numerical demonstration with inverted pendulum

Auxiliary process:

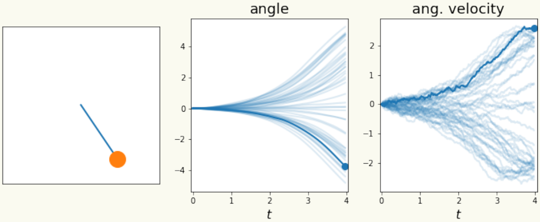

Actual controlled process:

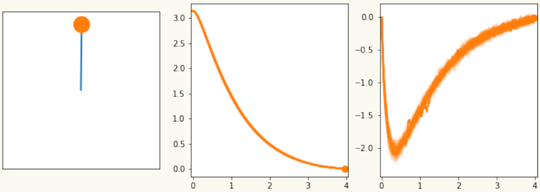

Outline

- **Part 0:** Preliminary
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Outline

- **Part 0:** Preliminary
- **Part 1:** Point to point steering
- **Part 2:** Distribution to distribution steering

Distribution to distribution steering

Problem formulation

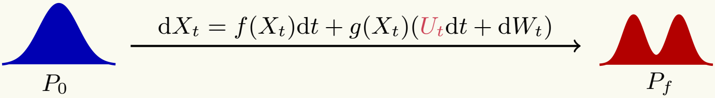

Problem: find a control law $U_t = k(t, X_t)$ so that $X_T \sim P_f$.

Can we use flow matching method to solve the problem?

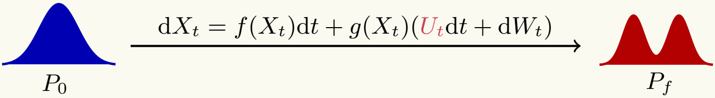

Problem: find a control law $U_t = k(t, X_t)$ so that $X_T \sim P_f$.

Can we use flow matching method to solve the problem?

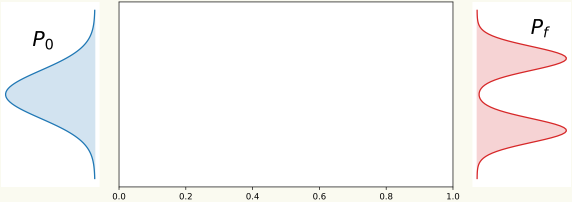

Actual process: $\dot{X}_t = k(t, X_t), \quad X_0 \sim P_0$

Bridge process: $X_t^\omega = (1-t)x_0 + t x_f \quad \omega = (x_0, x_f) \sim P_0 \otimes P_f$

Control law: $k(t,x) = \mathbb{E}[\dot{X}_t^\omega | X_t^\omega = x]$

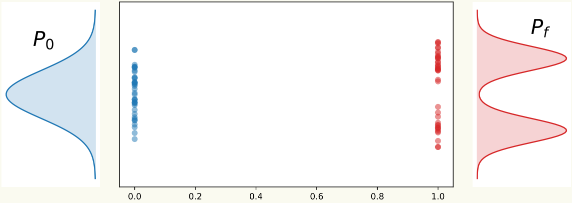

Actual process: $\dot{X}_t = k(t, X_t), \quad X_0 \sim P_0$

Bridge process: $X_t^\omega = (1-t)x_0 + t x_f \quad \omega = (x_0, x_f) \sim P_0 \otimes P_f$

Control law: $k(t,x) = \mathbb{E}[\dot{X}_t^\omega | X_t^\omega = x]$

Distribution to distribution steering

Flow matching (simple case)

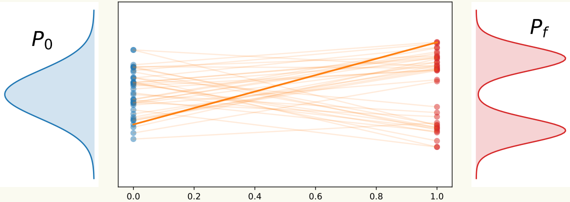

Actual process: $\dot{X}_t = k(t, X_t), \quad X_0 \sim P_0$

Bridge process: $X_t^\omega = (1-t)x_0 + t x_f \quad \omega = (x_0, x_f) \sim P_0 \otimes P_f$

Control law: $k(t,x) = \mathbb{E}[\dot{X}_t^\omega | X_t^\omega = x]$

Distribution to distribution steering

Flow matching (simple case)

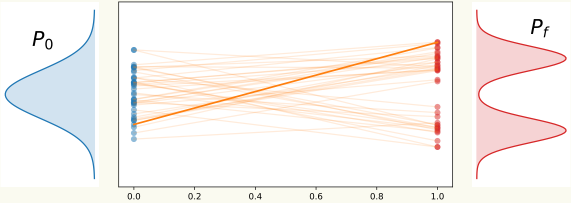

Actual process: $\dot{X}_t = k(t, X_t), \quad X_0 \sim P_0$

Bridge process: $X_t^\omega = (1-t)x_0 + t x_f \quad \omega = (x_0, x_f) \sim P_0 \otimes P_f$

Control law: $k(t,x) = \mathbb{E}[\dot{X}_t^\omega | X_t^\omega = x]$

$$= \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^\omega) - \dot{X}_t^\omega\|^2],$$

Distribution to distribution steering

Flow matching (simple case)

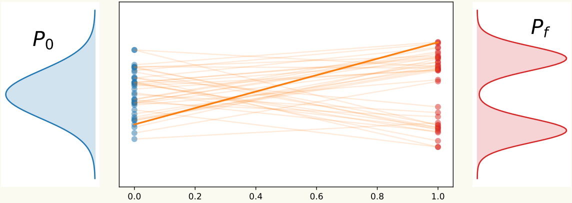

Actual process: $\dot{X}_t = k(t, X_t), \quad X_0 \sim P_0$

Bridge process: $X_t^\omega = (1-t)x_0 + t x_f \quad \omega = (x_0, x_f) \sim P_0 \otimes P_f$

Control law:

$$k(t,x) = \mathbb{E}[\dot{X}_t^\omega | X_t^\omega = x]$$

$$= \arg\min_\psi \mathbb{E}[\|\psi(t, X_t^\omega) - \dot{X}_t^\omega\|^2],$$

Distribution to distribution steering

Flow matching (simple case)

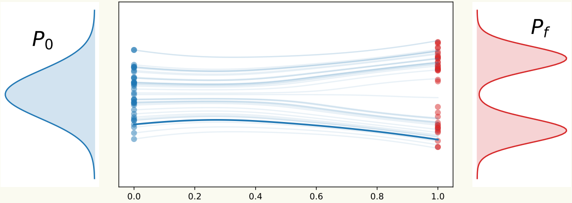

Actual process: $\dot{X}_t = k(t, X_t), \quad X_0 \sim P_0$

Bridge process: $X_t^\omega = (1-t)x_0 + t x_f \quad \omega = (x_0, x_f) \sim P_0 \otimes P_f$

Control law:

$$k(t,x) = \mathbb{E}[\dot{X}_t^\omega | X_t^\omega = x]$$

$$= \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^\omega) - \dot{X}_t^\omega\|^2],$$

Distribution to distribution steering

Flow matching (Stochastic control affine)

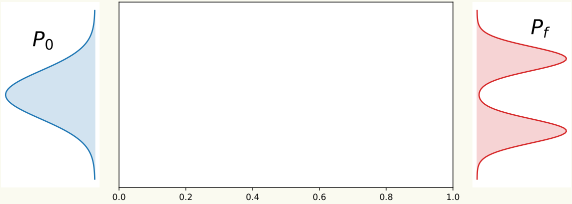

Actual process: $\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0$

Bridge process: $\mathrm{d}X_t^\omega = f(X_t^\omega)\mathrm{d}t + g(X_t^\omega)(U_t^\omega \mathrm{d}t + \mathrm{d}W_t), \quad (X_0^\omega, X_T^\omega) \sim P_0 \otimes P_f$

Control law: $k(t, x) = \mathbb{E}[U_t^\omega | X_t^\omega = x]$

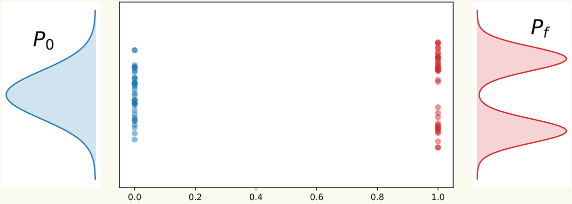

Actual process: $\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0$

Bridge process: $\mathrm{d}X_t^\omega = f(X_t^\omega)\mathrm{d}t + g(X_t^\omega)(U_t^\omega \mathrm{d}t + \mathrm{d}W_t), \quad (X_0^\omega, X_T^\omega) \sim P_0 \otimes P_f$

Control law: $k(t, x) = \mathbb{E}[U_t^\omega | X_t^\omega = x]$

Distribution to distribution steering

Flow matching (Stochastic control affine)

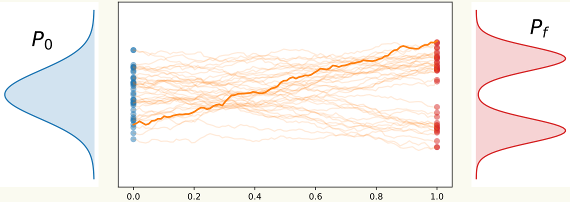

Actual process: $\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0$

Bridge process: $\mathrm{d}X_t^\omega = f(X_t^\omega)\mathrm{d}t + g(X_t^\omega)(U_t^\omega \mathrm{d}t + \mathrm{d}W_t), \quad (X_0^\omega, X_T^\omega) \sim P_0 \otimes P_f$

Control law: $k(t,x) = \mathbb{E}[U_t^\omega | X_t^\omega = x]$

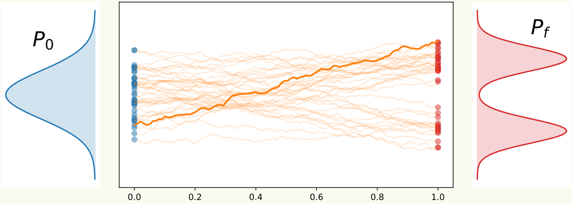

Actual process: $\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0$

Bridge process: $\mathrm{d}X_t^\omega = f(X_t^\omega)\mathrm{d}t + g(X_t^\omega)(U_t^\omega \mathrm{d}t + \mathrm{d}W_t), \quad (X_0^\omega, X_T^\omega) \sim P_0 \otimes P_f$

Control law: $k(t, x) = \mathbb{E}[U_t^\omega | X_t^\omega = x]$

$$= \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^\omega) - U_t^\omega\|^2],$$

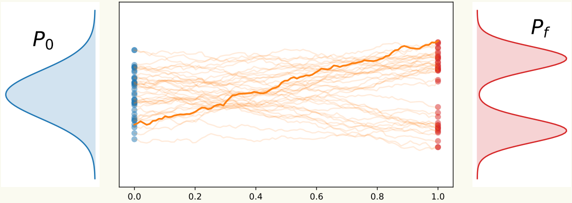

Actual process: $\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0$

Bridge process: $\mathrm{d}X_t^\omega = f(X_t^\omega)\mathrm{d}t + g(X_t^\omega)(U_t^\omega\mathrm{d}t + \mathrm{d}W_t), \quad (X_0^\omega, X_T^\omega) \sim P_0 \otimes P_f$

Control law:
$$k(t,x) = \mathbb{E}[U_t^\omega | X_t^\omega = x]$$
$$= \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^\omega) - U_t^\omega\|^2],$$

Distribution to distribution steering

Flow matching (Stochastic control affine)

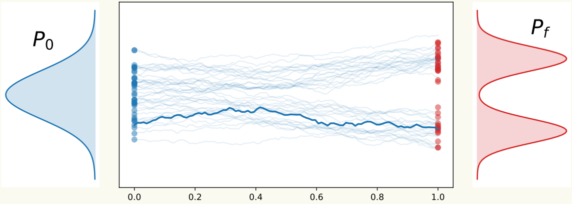

Actual process: $\mathrm{d}X_t = f(X_t)\mathrm{d}t + g(X_t)(k(t, X_t)\mathrm{d}t + \mathrm{d}W_t), \quad X_0 \sim P_0$

Bridge process: $\mathrm{d}X_t^\omega = f(X_t^\omega)\mathrm{d}t + g(X_t^\omega)(U_t^\omega \mathrm{d}t + \mathrm{d}W_t), \quad (X_0^\omega, X_T^\omega) \sim P_0 \otimes P_f$

Control law: $k(t, x) = \mathbb{E}[U_t^\omega | X_t^\omega = x]$

$$= \arg\min_\psi \mathbb{E}[\|\psi(t, X_t^\omega) - U_t^\omega\|^2],$$

Distribution to distribution steering

Theoretical justification

- A Markov process X is said to be the Markovian projection of a process η on $[0,T]$

$$\text{if} \quad X_t \stackrel{d}{=} \eta_t, \quad \forall t \in [0,T] \quad \text{(matching condition)}$$

- If η is an Itö process

$$\mathrm{d}\eta_t = \mathsf{f}_t \mathrm{d}t + \mathsf{g}_t \mathrm{d}W_t$$

- The Markovian projection of η is

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

$$\text{where} \quad f(t,x) = \mathbb{E}[\mathsf{f}_t | \eta_t = x], \quad g(t,x) = \mathbb{E}[\mathsf{g}_t \mathsf{g}_t^\top | \eta_t = x]]^{\frac{1}{2}}$$

- The (actual process) X is the Markovian projection of the (bridge process) X^ω over $[0,T]$. Therefore

$$X_t \stackrel{d}{=} X_t^\omega, \quad \forall t \in [0,T] \quad \Rightarrow \quad X_T \stackrel{d}{=} X_T^\omega \sim P_f$$

I. Gyongy "Mimicking the one-dimensional marginal distributions of processes having an Itô differential", Probab. Theory Relat. Fields, 71 (1986), pp. 501–516.

- A Markov process X is said to be the Markovian projection of a process η on $[0,T]$

$$\text{if} \quad X_t \overset{d}{=} \eta_t, \quad \forall t \in [0,T] \quad \text{(matching condition)}$$

- If η is an Itö process

$$\mathrm{d}\eta_t = \mathsf{f}_t \mathrm{d}t + \mathsf{g}_t \mathrm{d}W_t$$

- The Markovian projection of η is

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

$$\text{where} \quad f(t,x) = \mathbb{E}[\mathsf{f}_t|\eta_t = x], \quad g(t,x) = \mathbb{E}[\mathsf{g}_t\mathsf{g}_t^\top|\eta_t = x]^{\frac{1}{2}}$$

- The (actual process) X is the Markovian projection of the (bridge process) X^ω over $[0,T]$. Therefore

$$X_t \overset{d}{=} X_t^\omega, \quad \forall t \in [0,T] \quad \Rightarrow \quad X_T \overset{d}{=} X_T^\omega \sim P_f$$

I. Gyongy "Mimicking the one-dimensional marginal distributions of processes having an Itô differential", Probab. Theory Relat. Fields, 71 (1986), pp. 501–516.

- A Markov process X is said to be the Markovian projection of a process η on $[0,T]$

$$\text{if} \quad X_t \overset{d}{=} \eta_t, \quad \forall t \in [0,T] \quad \text{(matching condition)}$$

- If η is an Itö process

$$\mathrm{d}\eta_t = \mathsf{f}_t \mathrm{d}t + \mathsf{g}_t \mathrm{d}W_t$$

- The Markovian projection of η is

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

$$\text{where} \quad f(t,x) = \mathbb{E}[\mathsf{f}_t|\eta_t = x], \quad g(t,x) = \mathbb{E}[\mathsf{g}_t \mathsf{g}_t^\top|\eta_t = x]]^{\frac{1}{2}}$$

- The (actual process) X is the Markovian projection of the (bridge process) X^ω over $[0,T]$. Therefore

$$X_t \overset{d}{=} X_t^\omega, \quad \forall t \in [0,T] \quad \Rightarrow \quad X_T \overset{d}{=} X_T^\omega \sim P_f$$

I. Gyongy "Mimicking the one-dimensional marginal distributions of processes having an Itô differential", Probab. Theory Relat. Fields, 71 (1986), pp. 501–516.

Distribution to distribution steering

Theoretical justification

- A Markov process X is said to be the Markovian projection of a process η on $[0,T]$

$$\text{if} \quad X_t \overset{d}{=} \eta_t, \quad \forall t \in [0,T] \quad \text{(matching condition)}$$

- If η is an Itö process

$$\mathrm{d}\eta_t = \mathsf{f}_t \mathrm{d}t + \mathsf{g}_t \mathrm{d}W_t$$

- The Markovian projection of η is

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

$$\text{where} \quad f(t,x) = \mathbb{E}[\mathsf{f}_t|\eta_t = x], \quad g(t,x) = \mathbb{E}[\mathsf{g}_t \mathsf{g}_t^\top |\eta_t = x|]^{\frac{1}{2}}$$

- The (actual process) X is the Markovian projection of the (bridge process) X^ω over $[0,T]$. Therefore

$$X_t \overset{d}{=} X_t^\omega, \quad \forall t \in [0,T] \quad \Rightarrow \quad X_T \overset{d}{=} X_T^\omega \sim P_f$$

I. Gyongy "Mimicking the one-dimensional marginal distributions of processes having an Itô differential", Probab. Theory Relat. Fields, 71 (1986), pp. 501–516.

Distribution to distribution steering
Linear system

- Consider the linear system

$$dX_t = AX_t dt + B(U_t dt + dW_t)$$

- The bridge process X_t^ω admits an explicit distribution

$$X_t^\omega \sim \mathcal{N}(R_t x_0 + S_t x_f, \Sigma_t), \quad U_t^\omega = K_t(x_f - e^{(1-t)A} X_t^\omega)$$

where R_t, S_t, Σ_t and K_t only depend on (A, B)

- This allows us to find an explicit formula for the control law in certain settings (Gaussian $\to$ Mixture of Gaussians)

$$k(t,x) = \mathbb{E}[U_t^\omega | X_t^\omega = x] = \sum_{i=1}^{L} w_l(t,x)(Q_l(t)x - \mu_l(t))$$

- And easily sample from the bridge process for solving the nonlinear regression for any distributions

$$k(t,x) = \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^\omega) - U_t^\omega\|^2]$$

Distribution to distribution steering
Linear system

- Consider the linear system

$$\mathrm{d}X_t = AX_t\mathrm{d}t + B(U_t\mathrm{d}t + \mathrm{d}W_t)$$

- The bridge process X_t^ω admits an explicit distribution

$$X_t^\omega \sim \mathcal{N}(R_t x_0 + S_t x_f, \Sigma_t), \quad U_t^\omega = K_t(x_f - e^{(1-t)A} X_t^\omega)$$

where R_t, S_t, Σ_t and K_t only depend on (A, B)

- This allows us to find an explicit formula for the control law in certain settings (Gaussian $\to$ Mixture of Gaussians)

$$k(t,x) = \mathbb{E}[U_t^\omega | X_t^\omega = x] = \sum_{i=1}^{L} w_l(t,x)(Q_l(t)x - \mu_l(t))$$

- And easily sample from the bridge process for solving the nonlinear regression for any distributions

$$k(t,x) = \arg\min_{\psi} \mathbb{E}[\|\psi(t, X_t^\omega) - U_t^\omega\|^2]$$

- Consider the linear system

$$\mathrm{d}X_t = AX_t\mathrm{d}t + B(U_t\mathrm{d}t + \mathrm{d}W_t)$$

- The bridge process X_t^ω admits an explicit distribution

$$X_t^\omega \sim \mathcal{N}(R_t x_0 + S_t x_f, \Sigma_t), \quad U_t^\omega = K_t(x_f - e^{(1-t)A} X_t^\omega)$$

where R_t, S_t, Σ_t and K_t only depend on (A, B)

- This allows us to find an explicit formula for the control law in certain settings (Gaussian $\rightarrow$ Mixture of Gaussians)

$$k(t,x) = \mathbb{E}[U_t^\omega | X_t^\omega = x] = \sum_{i=1}^{L} w_l(t,x)(Q_l(t)x - \mu_l(t))$$

- And easily sample from the bridge process for solving the nonlinear regression for any distributions

$$k(t,x) = \arg\min_\psi \mathbb{E}[\|\psi(t, X_t^\omega) - U_t^\omega\|^2]$$

Distribution to distribution steering
Linear system

- Consider the linear system

$$\mathrm{d}X_t = AX_t\mathrm{d}t + B(U_t\mathrm{d}t + \mathrm{d}W_t)$$

- The bridge process X_t^ω admits an explicit distribution

$$X_t^\omega \sim \mathcal{N}(R_t x_0 + S_t x_f, \Sigma_t), \quad U_t^\omega = K_t(x_f - e^{(1-t)A}X_t^\omega)$$

where R_t, S_t, Σ_t and K_t only depend on (A, B)

- This allows us to find an explicit formula for the control law in certain settings (Gaussian $\rightarrow$ Mixture of Gaussians)

$$k(t,x) = \mathbb{E}[U_t^\omega | X_t^\omega = x] = \sum_{i=1}^{L} w_l(t,x)(Q_l(t)x - \mu_l(t))$$

- And easily sample from the bridge process for solving the nonlinear regression for any distributions

$$k(t,x) = \arg\min_\psi \mathbb{E}[\|\psi(t, X_t^\omega) - U_t^\omega\|^2]$$

Distribution to distribution steering
Numerical demonstration

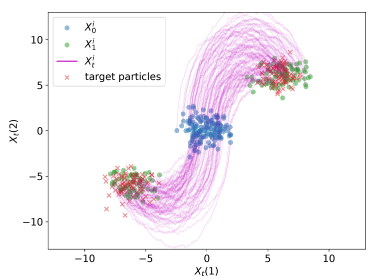

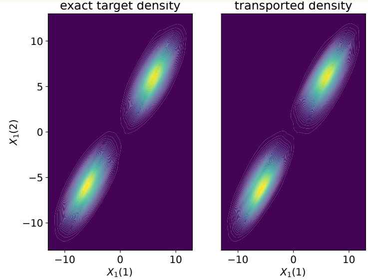

Linear system: (double integrator) $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$

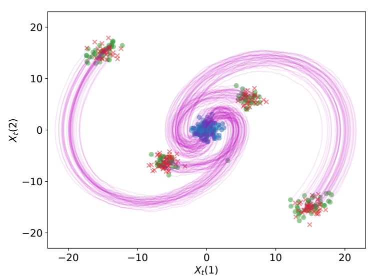

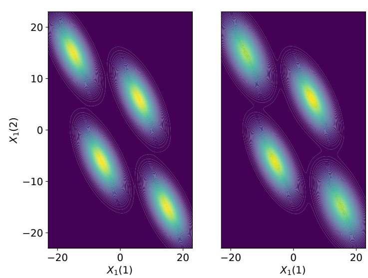

Linear system: (Oscillator) $A = \begin{bmatrix} 0 & 5 \\ -5 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$

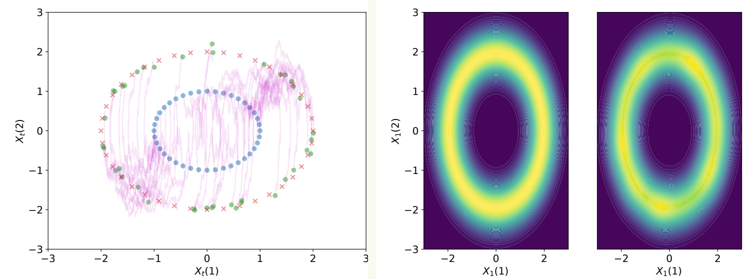

Linear system: (Nyquist Johnson resistor) $A = \begin{bmatrix} 0 & 1 \\ -1 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$

Thank you for your attention!

Joint work with:

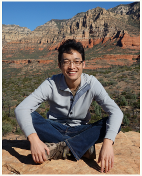

Yuhang Mei   Mohammad Al-Jarrah   Ali Pakniyat   Yongxin Chen

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

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