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

Presented at the IEEE Conference on Decision and Control, Rio de Janeiro, Brazil, 2025

Amirhossein Taghvaei
Joint work with Yuhang Mei and Ali Pakniyat

Department of Aeronautics & Astronautics
University of Washington, Seattle

Dec 10, 2025

Overview

Steering the state of a stochastic system

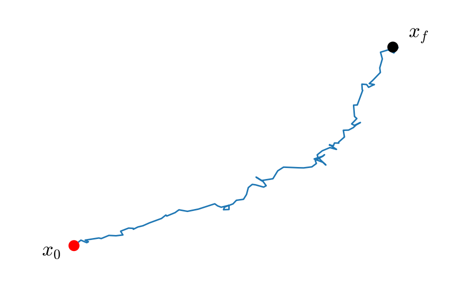

- Focus on computability rather than optimality
- **Methodology:** Time-reversal of diffusions
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview

Steering the state of a stochastic system

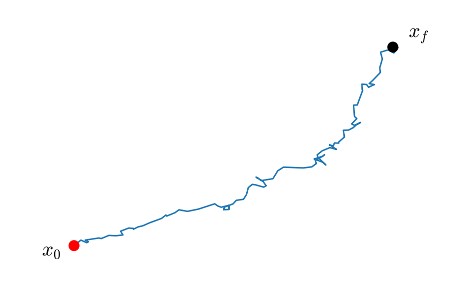

- Focus on computability rather than optimality
- **Methodology:** Time-reversal of diffusions
- **Algorithm:** Simulate and solve a nonlinear regression problem

Steering the state of a stochastic system

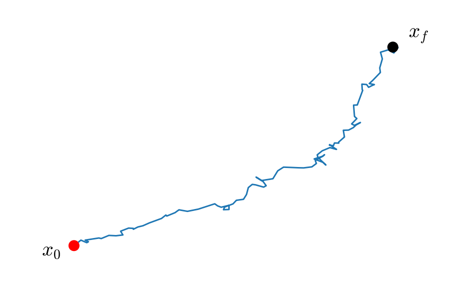

- Focus on computability rather than optimality
- **Methodology:** Time-reversal of diffusions
- **Algorithm:** Simulate and solve a nonlinear regression problem

Overview

Steering the state of a stochastic system

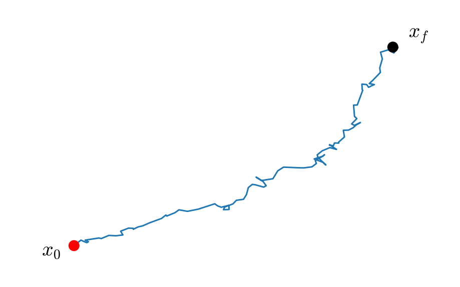

- Focus on computability rather than optimality
- **Methodology:** Time-reversal of diffusions
- **Algorithm:** Simulate and solve a nonlinear regression problem

Outline

- Time-reversal of diffusions
- Proposed methodology

Outline

- Time-reversal of diffusions
- Proposed methodology

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

Time-reversal of diffusions

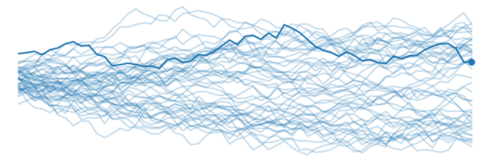

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

Time-reversal of diffusions

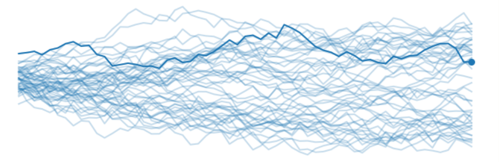

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

Time-reversal of diffusions

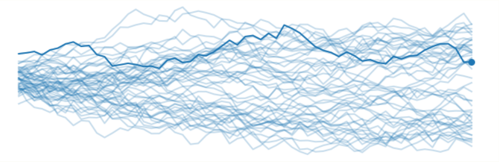

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

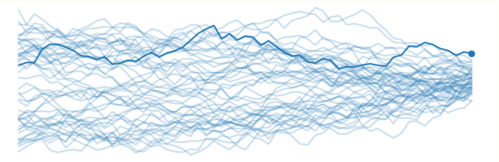

reversed process: $\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = ?$

Time-reversal of diffusions

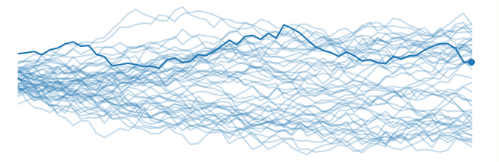

forward process: $\mathrm{d}Z_t = h(Z_t)\mathrm{d}t + g(Z_t)\mathrm{d}W_t$

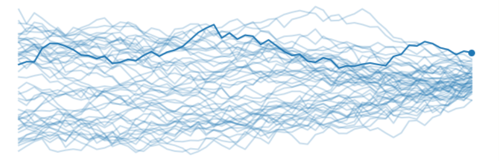

reversed process: $\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = ?$

Time-reversal of diffusions
Application in generative modeling

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

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?

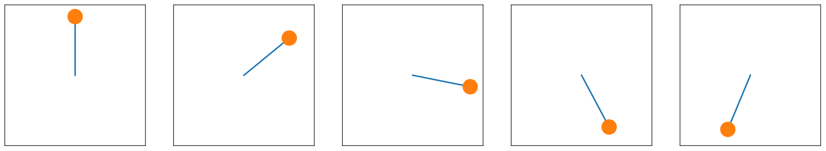

$\mathrm{d}Z_t =$ "pendulum dynamics"

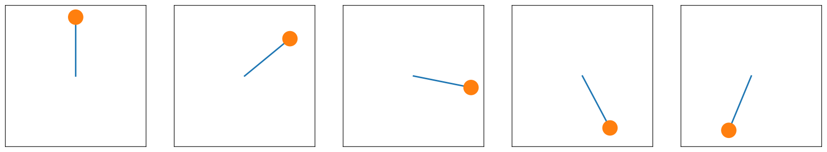

$$\mathrm{d}Z_t = \text{“pendulum dynamics”}$$

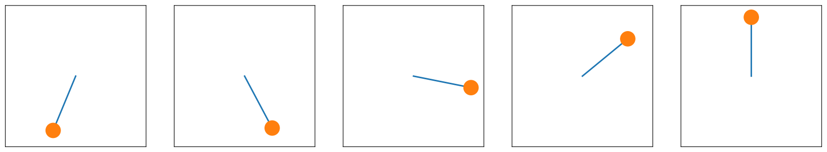

$$\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = \text{“steer to upward position”}$$

Time-reversal of diffusions

Application in control?

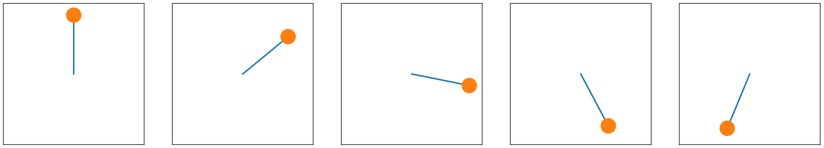

$$\mathrm{d}Z_t = \text{“pendulum dynamics”}$$

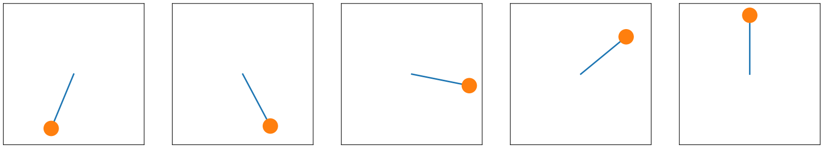

$$\widetilde{Z}_t := Z_{T-t}, \quad \mathrm{d}\widetilde{Z}_t = \text{“steer to upward position”}$$

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

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

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

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

Time-reversal formula [Haussmann & Pardoux, 1986]

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

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

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

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

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

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

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

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

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

Time-reversal formula [Haussmann & Pardoux, 1986]

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

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

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

- In order to numerically approximate

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

- define the objective function

$$J(s) = \mathbb{E}[\frac{1}{2}\|s(t,Z_t) - \nabla \log(p(t,Z_t))\|^2]$$

- expand and apply the integration by parts

$$J(s) = \mathbb{E}\left[\frac{1}{2}\|s(t,Z_t)\|^2 + \nabla \cdot s(t,Z_t)\right] + \text{const.}$$

- 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

$$J(s) = \mathbb{E}[\frac{1}{2}\|s(t,Z_t) - \nabla \log(p(t,Z_t))\|^2]$$

- expand and apply the integration by parts

$$J(s) = \mathbb{E}\left[\frac{1}{2}\|s(t,Z_t)\|^2 + \nabla \cdot s(t,Z_t)\right] + \text{const.}$$

- 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

$$J(s) = \mathbb{E}[\frac{1}{2}\|s(t,Z_t) - \nabla \log(p(t,Z_t))\|^2]$$

- expand and apply the integration by parts

$$J(s) = \mathbb{E}\left[\frac{1}{2}\|s(t,Z_t)\|^2 + \nabla \cdot s(t,Z_t)\right] + \text{const.}$$

- 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

$$J(s) = \mathbb{E}[\frac{1}{2}\|s(t,Z_t) - \nabla \log(p(t,Z_t))\|^2]$$

- expand and apply the integration by parts

$$J(s) = \mathbb{E}\left[\frac{1}{2}\|s(t,Z_t)\|^2 + \nabla \cdot s(t,Z_t)\right] + \text{const.}$$

- 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

Outline

- Time-reversal of diffusions
- Proposed methodology

Outline

Problem formulation

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

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?

Problem formulation

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

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?

Problem formulation

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

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 $\mathbb{E}[\|X_T - x_f\|^2] \leq \epsilon$.

$\boxed{\text{Can we use time-reversal method to solve the problem?}}$

Proposed methodology

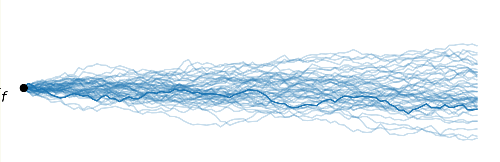

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

Proposed methodology

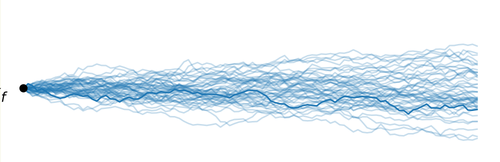

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

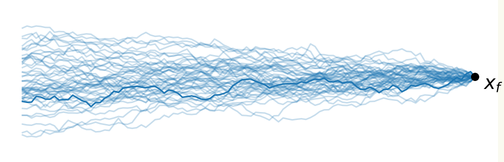

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

Proposed methodology

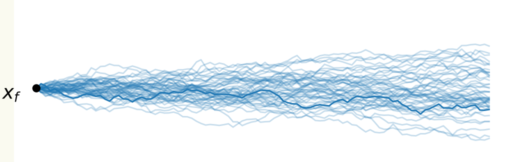

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

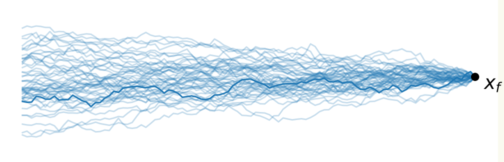

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

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

Proposed methodology

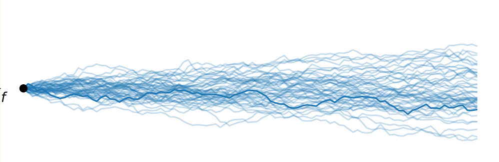

Auxiliary process: $\mathrm{d}Z_t = [-f(Z_t) + \nabla \cdot (gg^\top)(Z_t)]\mathrm{d}t + g(Z_t)\mathrm{d}W_t, \quad Z_0 = x_f$

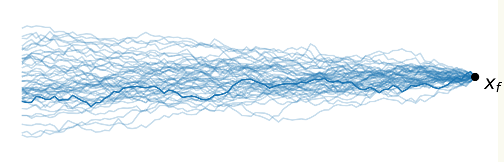

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_T = x_f$

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

Does it solve the exact steering problem?

Main result

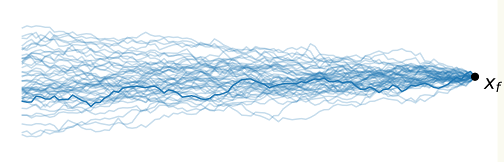

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

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

Main result

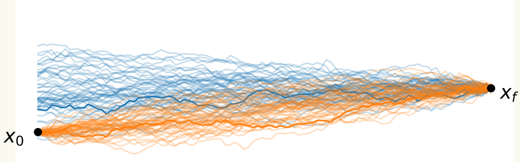

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

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

Main result

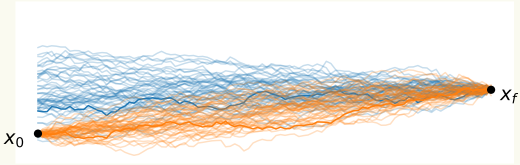

Reversed Auxiliary process: $\mathrm{d}\widetilde{Z}_t = [f(\widetilde{Z}_t) + g(\widetilde{Z}_t)k^*(t, \widetilde{Z}_t)]\mathrm{d}t + g(\widetilde{Z}_t)\mathrm{d}W_t, \quad \widetilde{Z}_0 \sim \widetilde{P}_0$

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

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

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

- Similar to the control law in [Chen & Georgiou, 2015] [Pakniyat & Tsiotras, 2021]

Y. Chen and T. Georgiou, "Stochastic bridges of linear systems," IEEE Transactions on Automatic Control, 2015.
A. Pakniyat and P. Tsiotras, "Steering the state of linear stochastic systems: a constrained minimum principle formulation," American Control Conference (ACC), 2021.

- 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

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

- Similar to the control law in [Chen & Georgiou, 2015] [Pakniyat & Tsiotras, 2021]

Y. Chen and T. Georgiou, "Stochastic bridges of linear systems," IEEE Transactions on Automatic Control, 2015.
A. Pakniyat and P. Tsiotras, "Steering the state of linear stochastic systems: a constrained minimum principle formulation," American Control Conference (ACC), 2021.

- 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

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

- Similar to the control law in [Chen & Georgiou, 2015] [Pakniyat & Tsiotras, 2021]

Y. Chen and T. Georgiou, "Stochastic bridges of linear systems," IEEE Transactions on Automatic Control, 2015.
A. Pakniyat and P. Tsiotras, "Steering the state of linear stochastic systems: a constrained minimum principle formulation," American Control Conference (ACC), 2021.

- 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

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

- Similar to the control law in [Chen & Georgiou, 2015] [Pakniyat & Tsiotras, 2021]

Y. Chen and T. Georgiou, "Stochastic bridges of linear systems," IEEE Transactions on Automatic Control, 2015.
A. Pakniyat and P. Tsiotras, "Steering the state of linear stochastic systems: a constrained minimum principle formulation," American Control Conference (ACC), 2021.

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

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

- Similar to the control law in [Chen & Georgiou, 2015] [Pakniyat & Tsiotras, 2021]

Y. Chen and T. Georgiou, "Stochastic bridges of linear systems," IEEE Transactions on Automatic Control, 2015.
A. Pakniyat and P. Tsiotras, "Steering the state of linear stochastic systems: a constrained minimum principle formulation," American Control Conference (ACC), 2021.

Auxiliary process:

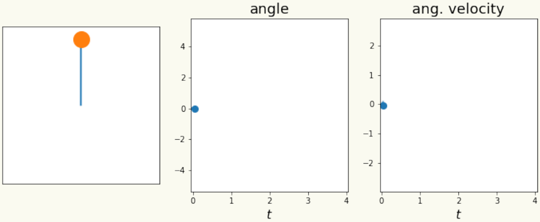

Actual controlled process:

Auxiliary process:

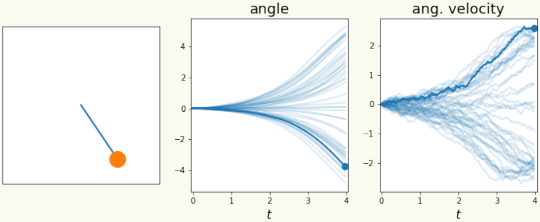

Actual controlled process:

Auxiliary process:

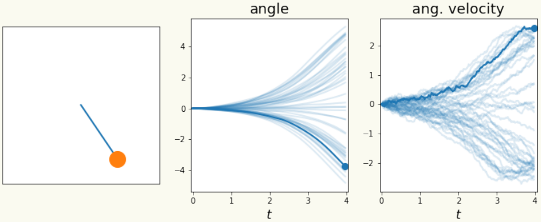

Actual controlled process:

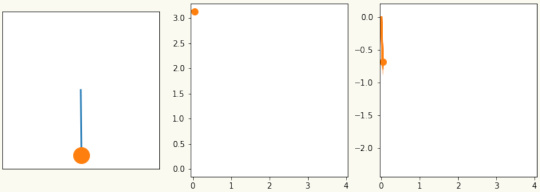

Auxiliary process:

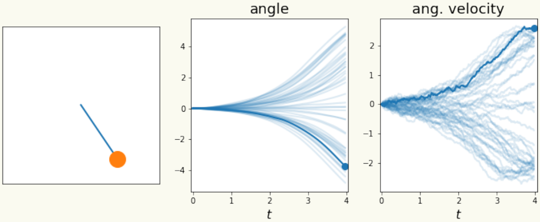

Actual controlled process:

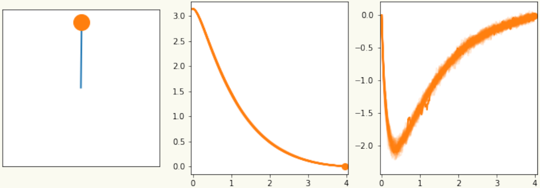

Joint work with:

Yuhang Mei

Ali Pakniyat

Reference:

- *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

Thank you for your attention!

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,

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

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

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,

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

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

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,

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

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

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

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

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

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

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

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