Sampling with diffusion models and Schrödinger bridges

Julius Berner

Department of Computing and Mathematical Sciences
California Institute of Technology

Caltech

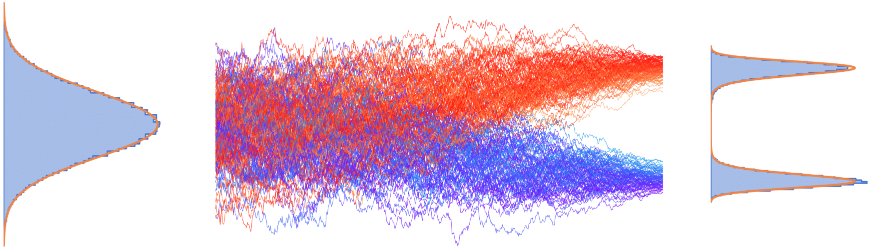

Collaborators

Lorenz Richter
Zuse Institute Berlin, dida Datenschmiede GmbH

Karen Ullrich
Meta AI

Sampling via learned diffusions

Goal: Sample from a distribution $p_{\mathrm{target}} = q_{\mathrm{target}}/\mathcal{Z}$ using auxiliary prior p_{prior}.

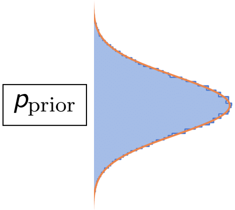

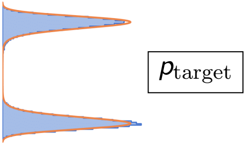

Sampling via learned diffusions

Goal: Sample from a distribution $p_{\mathrm{target}} = q_{\mathrm{target}}/\mathcal{Z}$ using auxiliary prior p_{prior}.

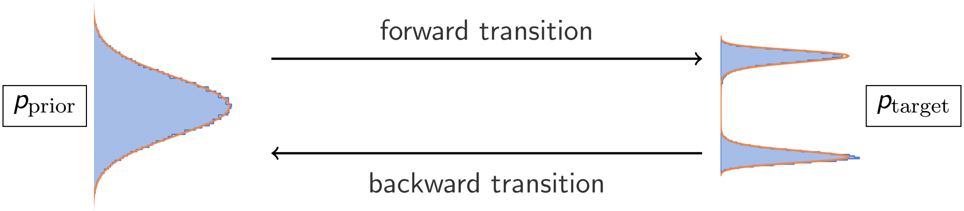

Reversals: Forward & backward transitions between prior p_{prior} and target p_{target}.

Sampling via learned diffusions

Goal: Sample from a distribution $p_{\text{target}} = q_{\text{target}}/\mathcal{Z}$ using auxiliary prior p_{prior}.

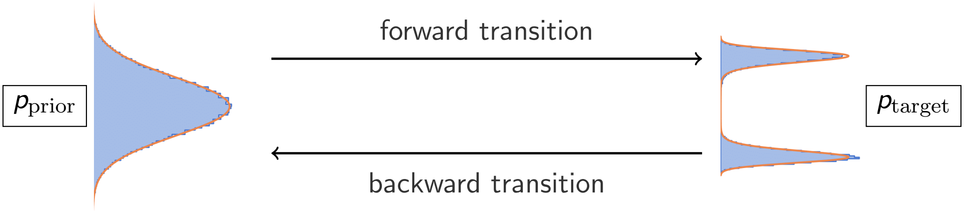

Reversals: Forward & backward transitions between prior p_{prior} and target p_{target}.
Examples: VI in latent variable models (VAEs), (optimal) transport, (Schrödinger) bridges.

Sampling via learned diffusions

Goal: Sample from a distribution $p_{\mathrm{target}} = q_{\mathrm{target}}/\mathcal{Z}$ using auxiliary prior p_{prior}.

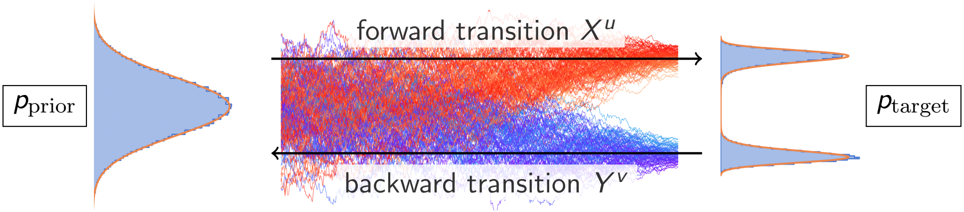

Reversals: Forward & backward transitions between prior p_{prior} and target p_{target}.

Examples: VI in latent variable models (VAEs), (optimal) transport, (Schrödinger) bridges.

This talk: Controlled SDEs (notation: $\overleftarrow{\sigma}(t) := \sigma(T-t)$)

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Sampling via learned diffusions

Goal: Sample from a distribution $p_{\mathrm{target}} = q_{\mathrm{target}}/\mathcal{Z}$ using auxiliary prior p_{prior}.

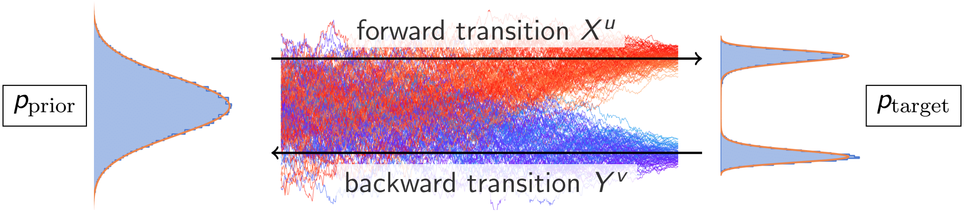

Reversals: Forward & backward transitions between prior p_{prior} and target p_{target}.

Examples: VI in latent variable models (VAEs), (optimal) transport, (Schrödinger) bridges.

This talk: Controlled SDEs (notation: $\overleftarrow{\sigma}(t) := \sigma(T-t)$)

$$
\begin{aligned}
\mathrm{d}X_s^u &= (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},\\
\mathrm{d}Y_s^v &= (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.
\end{aligned}
$$

New Goal: Learn u, v such that X^u is the time-reversal of Y^v.

Time-Reversals of SDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

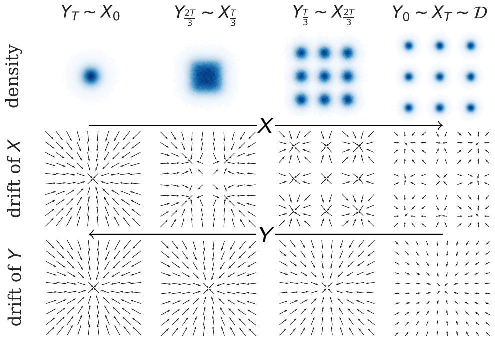

Time-Reversals of SDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Time-reversal (Proof: Fokker-Planck eq.):

$$\begin{cases} \mathrm{d}\overleftarrow{Y}_s^v = (f + \sigma\tilde{u})(\overleftarrow{Y}_s^v, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \\ \overleftarrow{Y}_0^v \sim p_{\mathrm{prior}} \end{cases}$$

with $\tilde{u} = \sigma\nabla \log \overleftarrow{p}_{Y^v} - v$, where $p_{Y^v}(\cdot, s)$ is the density of Y_s^v.

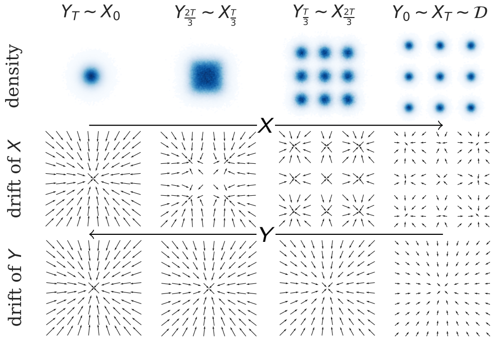

B. D. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Time-reversal (Proof: Fokker-Planck eq.):

$$\begin{cases} \mathrm{d}\overleftarrow{Y}_s^v = (f + \sigma\tilde{u})(\overleftarrow{Y}_s^v, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \\ \overleftarrow{Y}_0^v \sim p_{\mathrm{prior}} \end{cases}$$

with $\tilde{u} = \sigma\nabla \log \overleftarrow{p}_{Y^v} - v$, where $p_{Y^v}(\cdot, s)$ is the density of Y_s^v.

$\Rightarrow$ **Nelson's identity:**

$$u^* + v^* = \sigma\nabla \log p_{X^{u^*}} = \sigma\nabla \log \overleftarrow{p}_{Y^{v^*}}$$

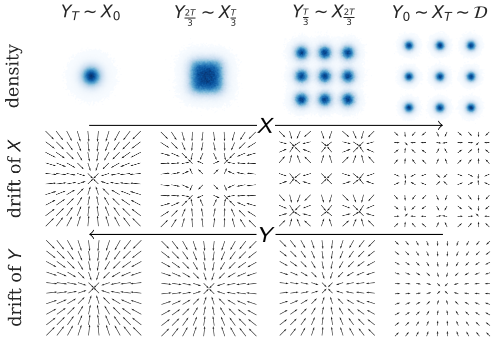

E. Nelson. Dynamical theories of Brownian motion. *Press, Princeton, NJ*, 1967

B. D. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982

Sampling via learned diffusions: Special cases

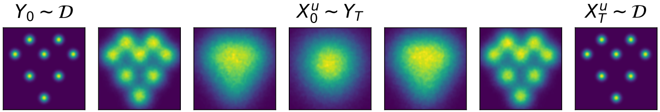

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Nelson's identity: $u^* + v^* = \sigma \nabla \log p_{X^{u^*}} = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$

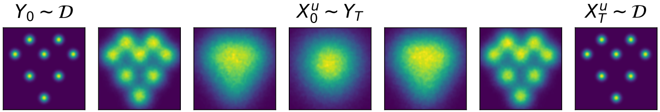

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Nelson's identity: $u^* + v^* = \sigma \nabla \log p_{X^{u^*}} = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$

Example (fix one control): Choose f, v^* such that $Y_T^{v^*} \approx p_{\mathrm{prior}}$ (independent of p_{target}).

Sampling via learned diffusions: Special cases

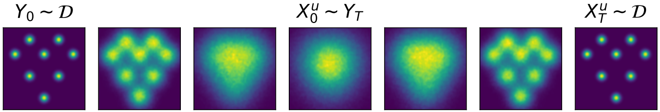

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Nelson's identity: $u^* + v^* = \sigma \nabla \log p_{X^{u^*}} = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$

Example (fix one control): Choose f, v^* such that $Y_T^{v^*} \approx p_{\mathrm{prior}}$ (independent of p_{target}).
$\Rightarrow$ Typical choices: $p_{\mathrm{prior}} = \mathcal{N}(0, \nu^2 \mathrm{I})$ or $p_{\mathrm{prior}} = \delta_{x_0}$.

Sampling via learned diffusions: Special cases

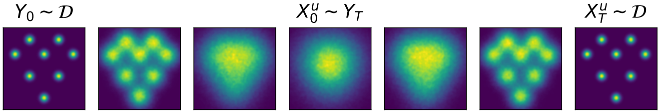

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Nelson's identity: $u^* + v^* = \sigma \nabla \log p_{X^{u^*}} = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$

Example (fix one control): Choose f, v^* such that $Y_T^{v^*} \approx p_{\mathrm{prior}}$ (independent of p_{target}).
$\Rightarrow$ Typical choices: $p_{\mathrm{prior}} = \mathcal{N}(0, \nu^2 \mathrm{I})$ or $p_{\mathrm{prior}} = \delta_{x_0}$.

- **Score-based models:** $v^* = 0 \Rightarrow u^* = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$

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 Representations*, 2020

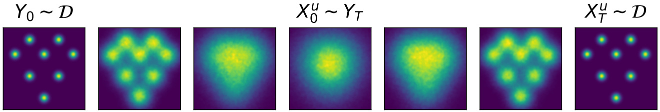

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \qquad Y_0^v \sim p_{\mathrm{target}}.$$

Nelson's identity: $u^* + v^* = \sigma\nabla \log p_{X^{u^*}} = \sigma\nabla \log \overleftarrow{p}_{Y^{v^*}}$

Example (fix one control): Choose f, v^* such that $Y_T^{v^*} \approx p_{\mathrm{prior}}$ (independent of p_{target}).
$\Rightarrow$ Typical choices: $p_{\mathrm{prior}} = \mathcal{N}(0, \nu^2 \mathrm{I})$ or $p_{\mathrm{prior}} = \delta_{x_0}$.

- **Score-based models:** $v^* = 0 \Rightarrow u^* = \sigma\nabla \log \overleftarrow{p}_{Y^{v^*}}$
- **Föllmer drift:** f such that p_{X^0} is tractable, $v^* = \sigma\nabla \log p_{X^0} \Rightarrow u^* = \sigma\nabla \log \frac{\overleftarrow{p}_{Y^{v^*}}}{p_{X^0}}$

P. Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied mathematics and Optimization*, 23(1):313–329, 1991
B. Tzen and M. Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pages 3084–3114. PMLR, 2019
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 Representations*, 2020

Sampling via learned diffusions: Special cases

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Examples: Ornstein-Uhlenbeck processes with tractable marginals conditioned on their initial value.
Let $\sigma(t) \coloneqq \sqrt{2\beta(t)}$ and $\alpha(t) \coloneqq \int_0^t \beta(s)\mathrm{d}s$ with $\alpha(T) \gg 1$.

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 Representations*, 2020

Sampling via learned diffusions: Special cases

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Examples: Ornstein-Uhlenbeck processes with tractable marginals conditioned on their initial value.

Let $\sigma(t) := \sqrt{2\beta(t)}$ and $\alpha(t) := \int_0^t \beta(s)\mathrm{d}s$ with $\alpha(T) \gg 1$.

- **Variance-Exploding (VE) SDE:** $f = 0, \quad v^* = 0 \Rightarrow X_T \approx \mathcal{N}(0, 2\alpha(T)\mathrm{I})$.

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 Representations*, 2020

Sampling via learned diffusions: Special cases

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Examples: Ornstein-Uhlenbeck processes with tractable marginals conditioned on their initial value.

Let $\sigma(t) := \sqrt{2\beta(t)}$ and $\alpha(t) := \int_0^t \beta(s)\mathrm{d}s$ with $\alpha(T) \gg 1$.

- **Variance-Exploding (VE) SDE:** $f = 0, \quad v^* = 0 \Rightarrow X_T \approx \mathcal{N}(0, 2\alpha(T)\mathrm{I})$.
- **Variance-Preserving (VP) SDE I:** $f(x,t) = \beta(t)x, \quad v^* = 0 \Rightarrow Y_T^0 \approx \mathcal{N}(0, \mathrm{I})$.

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 Representations*, 2020

Sampling via learned diffusions: Special cases

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Examples: Ornstein-Uhlenbeck processes with tractable marginals conditioned on their initial value.

Let $\sigma(t) \coloneqq \sqrt{2\beta(t)}$ and $\alpha(t) \coloneqq \int_0^t \beta(s)\mathrm{d}s$ with $\alpha(T) \gg 1$.

- **Variance-Exploding (VE) SDE:** $f = 0, \quad v^* = 0 \Rightarrow X_T \approx \mathcal{N}(0, 2\alpha(T)\mathrm{I})$.
- **Variance-Preserving (VP) SDE I:** $f(x,t) = \beta(t)x, \quad v^* = 0 \Rightarrow Y_T^0 \approx \mathcal{N}(0, \mathrm{I})$.
- **VP SDE II:** $f(x,t) = -\beta(t)x, \quad v^* = \sigma\nabla \log p_{X^0} = -\frac{2f}{\sigma} \Rightarrow Y_T^0 \approx \mathcal{N}(0, \mathrm{I})$.

F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023

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 Representations*, 2020

Sampling via learned diffusions: Special cases

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Examples: Ornstein-Uhlenbeck processes with tractable marginals conditioned on their initial value.

Let $\sigma(t) := \sqrt{2\beta(t)}$ and $\alpha(t) := \int_0^t \beta(s)\mathrm{d}s$ with $\alpha(T) \gg 1$.

- **Variance-Exploding (VE) SDE:** $f = 0, \quad v^* = 0$ $\Rightarrow X_T \approx \mathcal{N}(0, 2\alpha(T)\mathrm{I})$.
- **Variance-Preserving (VP) SDE I:** $f(x,t) = \beta(t)x, \quad v^* = 0$ $\Rightarrow Y_T^0 \approx \mathcal{N}(0, \mathrm{I})$.
- **VP SDE II:** $f(x,t) = -\beta(t)x, \quad v^* = \sigma \nabla \log p_{X^0} = -\frac{2f}{\sigma}$ $\Rightarrow Y_T^0 \approx \mathcal{N}(0, \mathrm{I})$.
- **Pinned Brownian motion:** $f = 0, \quad v^* = \sigma \nabla \log p_{X^0} = \frac{-\sigma x}{2\alpha(t)}$ $\Rightarrow Y_T^0 = \delta$.

Q. Zhang and Y. Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022

F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023

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 Representations*, 2020

Optimality PDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

- **Score-based models:** $u^* = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$
- **Föllmer drift:** $u^* = \sigma \nabla \log \frac{\overleftarrow{p}_{Y^{v^*}}}{p_{X^0}}$

Optimality PDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

- **Score-based models:** $u^* = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$
- **Föllmer drift:** $u^* = \sigma \nabla \log \frac{\overleftarrow{p}_{Y^{v^*}}}{p_{X^0}}$

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\tfrac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \tfrac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Optimality PDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

- **Score-based models:** $u^* = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$

$$g = -\log q_{\mathrm{target}}, \qquad h = -\operatorname{div}(f)$$

- **Föllmer drift:** $u^* = \sigma \nabla \log \frac{\overleftarrow{p}_{Y^{v^*}}}{p_{X^0}}$

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\tfrac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \tfrac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

- **Score-based models:** $u^* = \sigma\nabla \log \overleftarrow{p}_{Y^{v^*}}$ $\quad g = -\log q_{\mathrm{target}}, \quad h = -\operatorname{div}(f)$
- **Föllmer drift:** $u^* = \sigma\nabla \log \frac{\overleftarrow{p}_{Y^{v^*}}}{p_{X^0}}$ $\quad g = -\log \frac{q_{\mathrm{target}}}{p_{X_T^0}}, \quad h = 0$

Optimality PDE (HJB equation): $u^* = -\sigma\nabla V(x,t)$, where

$$\partial_t V = -\tfrac{1}{2}\sigma^2 \Delta V - f\cdot\nabla V - h + \tfrac{1}{2}\left\|\sigma\nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

P. Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied mathematics and Optimization*, 23(1):313–329, 1991

B. Tzen and M. Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pages 3084–3114. PMLR, 2019

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

Optimality PDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

- **Score-based models:** $u^* = \sigma \nabla \log \overleftarrow{p}_{Y^{v^*}}$ $\qquad g = -\log q_{\mathrm{target}}, \quad h = -\operatorname{div}(f)$
- **Föllmer drift:** $u^* = \sigma \nabla \log \frac{\overleftarrow{p}_{Y^{v^*}}}{p_{X^0}}$ $\qquad g = -\log \frac{q_{\mathrm{target}}}{p_{X_T^0}}, \quad h = 0$

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\tfrac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \tfrac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Proof: Schrödinger bridge optimality PDEs & Hopf-Cole transform of Fokker-Planck eq.

P. Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied mathematics and Optimization*, 23(1):313–329, 1991

B. Tzen and M. Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pages 3084–3114. PMLR, 2019

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\frac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \frac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Numerical solution of HJB equations

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\frac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \frac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Numerical solution:

- (Neural) PDE solver: PINNs, (B)SDE-based methods, Tensor methods

L. Richter, L. Sallandt, and N. Nüsken. Solving high-dimensional parabolic pdes using the tensor train format. In *International Conference on Machine Learning*, pages 8998–9009. PMLR, 2021

N. Nüsken and L. Richter. Interpolating between BSDEs and PINNs–deep learning for elliptic and parabolic boundary value problems. *arXiv preprint arXiv:2112.03749*, 2021

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

Numerical solution of HJB equations

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\frac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \frac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Numerical solution:

- (Neural) PDE solver: PINNs, (B)SDE-based methods, Tensor methods
- $\exp(V)$ satisfies Kolmogorov equation: Feynman-Kac formula provides MC estimator

F. Vargas, A. Ovsianas, D. Fernandes, M. Girolami, N. D. Lawrence, and N. Nüsken. Bayesian learning via neural Schrödinger–Föllmer flows. *Statistics and Computing*, 33(1):1–22, 2023

L. Richter and J. Berner. Robust SDE-based variational formulations for solving linear PDEs via deep learning. In *International Conference on Machine Learning*, pages 18649–18666. PMLR, 2022

L. Richter, L. Sallandt, and N. Nüsken. Solving high-dimensional parabolic pdes using the tensor train format. In *International Conference on Machine Learning*, pages 8998–9009. PMLR, 2021

N. Nüsken and L. Richter. Interpolating between BSDEs and PINNs–deep learning for elliptic and parabolic boundary value problems. *arXiv preprint arXiv:2112.03749*, 2021

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

Numerical solution of HJB equations

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\frac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \frac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Numerical solution:

- (Neural) PDE solver: PINNs, (B)SDE-based methods, Tensor methods
- $\exp(V)$ satisfies Kolmogorov equation: Feynman-Kac formula provides MC estimator
- Verification theorem from stochastic optimal control

N. Nüsken and L. Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021

W. H. Fleming and H. M. Soner. *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media, 2006

F. Vargas, A. Ovsianas, D. Fernandes, M. Girolami, N. D. Lawrence, and N. Nüsken. Bayesian learning via neural Schrödinger–Föllmer flows. *Statistics and Computing*, 33(1):1–22, 2023

L. Richter and J. Berner. Robust SDE-based variational formulations for solving linear PDEs via deep learning. In *International Conference on Machine Learning*, pages 18649–18666. PMLR, 2022

L. Richter, L. Sallandt, and N. Nüsken. Solving high-dimensional parabolic pdes using the tensor train format. In *International Conference on Machine Learning*, pages 8998–9009. PMLR, 2021

N. Nüsken and L. Richter. Interpolating between BSDEs and PINNs–deep learning for elliptic and parabolic boundary value problems. *arXiv preprint arXiv:2112.03749*, 2021

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

Verification theorem

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\frac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \frac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Verification theorem from stochastic optimal control (Proof: Ito's formula):

$$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \left(\frac{1}{2}\|u\|^2 + h\right)(X_s^u, s)\,\mathrm{d}s + g(X_T^u)\right]$$

Verification theorem

Optimality PDE (HJB equation): $u^* = -\sigma\nabla V(x,t)$, where

$$\partial_t V = -\tfrac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \tfrac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Verification theorem from stochastic optimal control (Proof: Ito's formula):

$$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \left(\tfrac{1}{2}\|u\|^2 + h\right)(X_s^u, s)\,\mathrm{d}s + g(X_T^u)\right]$$

- **Score-based models:** $g = -\log q_{\mathrm{target}}, \quad h = -\operatorname{div}(f)$

  $\Rightarrow$ Time-reversed diffusion sampler (DIS) with VPSDE I or VESDE

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

Verification theorem

Optimality PDE (HJB equation): $u^* = -\sigma \nabla V(x,t)$, where

$$\partial_t V = -\tfrac{1}{2}\sigma^2 \Delta V - f \cdot \nabla V - h + \tfrac{1}{2}\left\|\sigma \nabla V\right\|^2, \qquad V(\cdot, T) = g.$$

Verification theorem from stochastic optimal control (Proof: Ito's formula):

$$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \left(\tfrac{1}{2}\|u\|^2 + h\right)(X_s^u, s)\,\mathrm{d}s + g(X_T^u)\right]$$

- **Score-based models:** $g = -\log q_{\mathrm{target}}, \quad h = -\operatorname{div}(f)$

  $\Rightarrow$ Time-reversed diffusion sampler (DIS) with VPSDE I or VESDE

- **Föllmer drift:** $g = -\log \frac{q_{\mathrm{target}}}{p_{X_T^0}}, \quad h = 0$

  $\Rightarrow$ Path integral sampler (PIS) with pinned BM

  $\Rightarrow$ Denoising Diffusion Sampler (DDS) with VPSDE II

F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023

Q. Zhang and Y. Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022

L. Richter. *Solving high-dimensional PDEs, approximation of path space measures and importance sampling of diffusions*. PhD thesis, BTU Cottbus-Senftenberg, 2021

F. Vargas, A. Ovsianas, D. Fernandes, M. Girolami, N. D. Lawrence, and N. Nüsken. Bayesian learning via neural Schrödinger–Föllmer flows. *Statistics and Computing*, 33(1):1–22, 2023

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

Numerical Evaluation

Methods: PIS, DDS, DIS, and Continual Repeated Annealed Flow Transport (CRAFT).

Q. Zhang and Y. Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022

F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

A. Matthews, M. Arbel, D. J. Rezende, and A. Doucet. Continual repeated annealed flow transport monte carlo. In *International Conference on Machine Learning*, pages 15196–15219. PMLR, 2022

Methods: PIS, DDS, DIS, and Continual Repeated Annealed Flow Transport (CRAFT).
Metrics: Log-normalizing constant, Sinkhorn distance to ground truth, normalized effective sample size, and average stddev.

Q. Zhang and Y. Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022

F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023

J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024

A. Matthews, M. Arbel, D. J. Rezende, and A. Doucet. Continual repeated annealed flow transport monte carlo. In *International Conference on Machine Learning*, pages 15196–15219. PMLR, 2022

Numerical Evaluation

Methods: PIS, DDS, DIS, and Continual Repeated Annealed Flow Transport (CRAFT).
Metrics: Log-normalizing constant, Sinkhorn distance to ground truth, normalized effective sample size, and average stddev.

Problem	Method	$\Delta \log \mathcal{Z}$ (rw) $\downarrow$	$\mathcal{W}_\gamma^2 \downarrow$	ESS $\uparrow$	Δ std $\downarrow$
GMM	CRAFT	0.012	**0.020**	-	**0.019**
(d=2)	PIS	0.249	0.467	0.0051	1.937
	DDS	**0.005**	0.042	**0.0737**	2.161
	DIS	0.015	0.064	0.0226	2.522
Funnel	CRAFT	0.123	5.517	-	6.139
(d=10)	PIS	0.111	5.639	0.1333	6.921
	DDS	0.045	5.305	0.1446	6.133
	DIS	**0.032**	**5.120**	**0.1383**	**5.254**

Q. Zhang and Y. Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022
F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023
J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024
A. Matthews, M. Arbel, D. J. Rezende, and A. Doucet. Continual repeated annealed flow transport monte carlo. In *International Conference on Machine Learning*, pages 15196–15219. PMLR, 2022

Methods: PIS, DDS, DIS, and Continual Repeated Annealed Flow Transport (CRAFT).
Metrics: Log-normalizing constant, Sinkhorn distance to ground truth, normalized effective sample size, and average stddev.

Problem	Method	$\Delta \log \mathcal{Z}$ (rw) $\downarrow$	$\mathcal{W}_\gamma^2 \downarrow$	ESS $\uparrow$	Δ std $\downarrow$
GMM	CRAFT	0.012	**0.020**	-	**0.019**
(d=2)	PIS	0.249	0.467	0.0051	1.937
	DDS	**0.005**	0.042	**0.0737**	2.161
	DIS	0.015	0.064	0.0226	2.522
Funnel	CRAFT	0.123	5.517	-	6.139
(d=10)	PIS	0.111	5.639	0.1333	6.921
	DDS	0.045	5.305	0.1446	6.133
	DIS	**0.032**	**5.120**	**0.1383**	**5.254**

Gaussian Mixture:

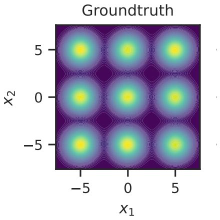

Q. Zhang and Y. Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022
F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023
J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024
A. Matthews, M. Arbel, D. J. Rezende, and A. Doucet. Continual repeated annealed flow transport monte carlo. In *International Conference on Machine Learning*, pages 15196–15219. PMLR, 2022

Methods: PIS, DDS, DIS, and Continual Repeated Annealed Flow Transport (CRAFT).
Metrics: Log-normalizing constant, Sinkhorn distance to ground truth, normalized effective sample size, and average stddev.

Problem	Method	$\Delta \log \mathcal{Z}$ (rw) $\downarrow$	$\mathcal{W}_\gamma^2 \downarrow$	ESS $\uparrow$	Δ std $\downarrow$
GMM	CRAFT	0.012	**0.020**	-	**0.019**
(d=2)	PIS	0.249	0.467	0.0051	1.937
	DDS	**0.005**	0.042	**0.0737**	2.161
	DIS	0.015	0.064	0.0226	2.522
Funnel	CRAFT	0.123	5.517	-	6.139
(d=10)	PIS	0.111	5.639	0.1333	6.921
	DDS	0.045	5.305	0.1446	6.133
	DIS	**0.032**	**5.120**	**0.1383**	**5.254**

Gaussian Mixture:

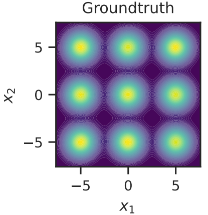

Funnel (test for MCMC methods):

$$p_{\mathrm{target}}(x) = \mathcal{N}(x_1; 0, 9) \prod_{i=2}^{10} \mathcal{N}(x_i; 0, e^{x_1})$$

Q. Zhang and Y. Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022
F. Vargas, W. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023
J. Berner, L. Richter, and K. Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024
A. Matthews, M. Arbel, D. J. Rezende, and A. Doucet. Continual repeated annealed flow transport monte carlo. In *International Conference on Machine Learning*, pages 15196–15219. PMLR, 2022

Application: Posterior sampling

Föllmer drift: $$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log \frac{q_{\mathrm{target}}}{p_{X_T^0}}(X_T^u)\right],$$
$$\text{where } \mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s,\ X_0^u \sim p_{\mathrm{prior}}.$$

Application: Posterior sampling

Föllmer drift: $$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log \frac{q_{\mathrm{target}}}{p_{X_T^0}}(X_T^u)\right],$$
$$\text{where } \mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s,\ \ X_0^u \sim p_{\mathrm{prior}}.$$

Posterior sampling: Assume we have data and measurement distributions $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ and want to sample from the posterior $p_{\mathrm{target}} = p_{\mathcal{X}|\mathcal{Y}}(\cdot|y)$ for a given measurement y.

Application: Posterior sampling

Föllmer drift: $$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log \frac{q_{\mathrm{target}}}{p_{X_T^0}}(X_T^u)\right],$$
$$\text{where } \mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \ \ X_0^u \sim p_{\mathrm{prior}}.$$

Posterior sampling: Assume we have data and measurement distributions $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ and want to sample from the posterior $p_{\mathrm{target}} = p_{\mathcal{X}|\mathcal{Y}}(\cdot|y)$ for a given measurement y.

Bayes' theorem: $p_{\mathcal{X}|\mathcal{Y}}(\cdot|y) \propto p_{\mathcal{Y}|\mathcal{X}}(y|\cdot)p_{\mathcal{X}}$.

Application: Posterior sampling

Föllmer drift: $$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log \frac{q_{\mathrm{target}}}{p_{X_T^0}}(X_T^u)\right],$$
$$\text{where } \mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \;\; X_0^u \sim p_{\mathrm{prior}}.$$

Posterior sampling: Assume we have data and measurement distributions $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ and want to sample from the posterior $p_{\mathrm{target}} = p_{\mathcal{X}|\mathcal{Y}}(\cdot|y)$ for a given measurement y.

Bayes' theorem: $p_{\mathcal{X}|\mathcal{Y}}(\cdot|y) \propto p_{\mathcal{Y}|\mathcal{X}}(y|\cdot)p_{\mathcal{X}}$.

Prior SDE: Assume the drift f is pretrained such that the uncontrolled SDE X^0 samples from the data distribution $X_T^0 \sim p_{\mathcal{X}}$.

Föllmer drift: $$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log \frac{q_{\mathrm{target}}}{p_{X_T^0}}(X_T^u)\right],$$
$$\text{where } \mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \;\; X_0^u \sim p_{\mathrm{prior}}.$$

Posterior sampling: Assume we have data and measurement distributions $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ and want to sample from the posterior $p_{\mathrm{target}} = p_{\mathcal{X}|\mathcal{Y}}(\cdot|y)$ for a given measurement y.

Bayes' theorem: $p_{\mathcal{X}|\mathcal{Y}}(\cdot|y) \propto p_{\mathcal{Y}|\mathcal{X}}(y|\cdot)p_{\mathcal{X}}$.

Prior SDE: Assume the drift f is pretrained such that the uncontrolled SDE X^0 samples from the data distribution $X_T^0 \sim p_{\mathcal{X}}$.

$\Rightarrow$ **Objective** (prior cancels in terminal costs):

$$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log p_{\mathcal{Y}|\mathcal{X}}(y|X_T^u)\right]$$

Application: Posterior sampling

Föllmer drift: $$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log \frac{q_{\mathrm{target}}}{p_{X_T^0}}(X_T^u)\right],$$
$$\text{where } \mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \;\; X_0^u \sim p_{\mathrm{prior}}.$$

Posterior sampling: Assume we have data and measurement distributions $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ and want to sample from the posterior $p_{\mathrm{target}} = p_{\mathcal{X}|\mathcal{Y}}(\cdot|y)$ for a given measurement y.

Bayes' theorem: $p_{\mathcal{X}|\mathcal{Y}}(\cdot|y) \propto p_{\mathcal{Y}|\mathcal{X}}(y|\cdot)p_{\mathcal{X}}$.

Prior SDE: Assume the drift f is pretrained such that the uncontrolled SDE X^0 samples from the data distribution $X_T^0 \sim p_{\mathcal{X}}$.

$\Rightarrow$ **Objective** (prior cancels in terminal costs):

$$u^* = \arg\min_u \mathbb{E}\left[\int_0^T \tfrac{1}{2}\|u(X_s^u, s)\|^2 \,\mathrm{d}s - \log p_{\mathcal{Y}|\mathcal{X}}(y|X_T^u)\right]$$

👍 Online method (only needs unnormalized likelihood and no samples).

Sampling via learned diffusions: General case

General case: Two controlled SDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

General case: Two controlled SDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Path space perspective: Consider path measures $\mathbb{P}_{X^u}$ and $\mathbb{P}_{\overleftarrow{Y}^v}$ on $C([0,T],\mathbb{R}^d)$.

Sampling via learned diffusions: General case

General case: Two controlled SDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Path space perspective: Consider path measures $\mathbb{P}_{X^u}$ and $\mathbb{P}_{\overleftarrow{Y}^v}$ on $C([0,T],\mathbb{R}^d)$.
Goal: Identify controls u^*, v^* via divergence D between these measures

$$u^*, v^* \in \operatorname*{arg\,min}_{u,v} D\left(\mathbb{P}_{X^u} \middle| \mathbb{P}_{\overleftarrow{Y}^v}\right).$$

Sampling via learned diffusions: General case

General case: Two controlled SDEs

$$\mathrm{d}X_s^u = (f + \sigma u)(X_s^u, s)\,\mathrm{d}s + \sigma(s)\,\mathrm{d}W_s, \qquad X_0^u \sim p_{\mathrm{prior}},$$
$$\mathrm{d}Y_s^v = (-\overleftarrow{f} + \overleftarrow{\sigma}\overleftarrow{v})(Y_s^v, s)\,\mathrm{d}s + \overleftarrow{\sigma}(s)\,\mathrm{d}W_s, \quad Y_0^v \sim p_{\mathrm{target}}.$$

Path space perspective: Consider path measures $\mathbb{P}_{X^u}$ and $\mathbb{P}_{\overleftarrow{Y}^v}$ on $C([0,T],\mathbb{R}^d)$.
Goal: Identify controls u^*, v^* via divergence D between these measures

$$u^*, v^* \in \operatorname*{arg\,min}_{u,v} D\left(\mathbb{P}_{X^u} \middle| \mathbb{P}_{\overleftarrow{Y}^v}\right).$$

Tool: Radon-Nikodym derivative

$$\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w) = \int_0^T \left((u+v)\cdot\left(w + \frac{v-u}{2}\right) + \nabla\cdot(\sigma v - f)\right)(X_s^w, s)\,\mathrm{d}s + \int_0^T (u+v)(X_s^w, s)\cdot\mathrm{d}W_s + \log\frac{p_{\mathrm{prior}}(X_0^w)}{p_{\mathrm{target}}(X_T^w)}$$

F. Vargas and N. Nüsken. Transport, variational inference and diffusions: with applications to annealed flows and Schrödinger bridges. *arXiv preprint arXiv:2307.01050*, 2023

L. Richter and J. Berner. Improved sampling via learned diffusions. *International Conference on Learning Representations*, 2024

$$\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w) = \int_0^T \Big((u+v) \cdot \Big(w + \frac{v-u}{2} \Big) + \nabla \cdot (\sigma v - f) \Big)(X_s^w, s)\,\mathrm{d}s + \int_0^T (u+v)(X_s^w, s) \cdot \mathrm{d}W_s + \log \frac{p_{\mathrm{prior}}(X_0^w)}{p_{\mathrm{target}}(X_T^w)}$$

Regularizer: General objective $\min_{v,u} D\big(\mathbb{P}_{X^u} \big| \mathbb{P}_{\overleftarrow{Y}^v}\big)$ is not unique.

Sampling via learned diffusions: General case

$$\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w) = \int_0^T \Big((u+v) \cdot \Big(w + \frac{v-u}{2} \Big) + \nabla \cdot (\sigma v - f) \Big)(X_s^w, s)\, \mathrm{d}s + \int_0^T (u+v)(X_s^w, s) \cdot \mathrm{d}W_s + \log \frac{p_{\mathrm{prior}}(X_0^w)}{p_{\mathrm{target}}(X_T^w)}$$

Regularizer: General objective $\min_{v,u} D\big(\mathbb{P}_{X^u} \big| \mathbb{P}_{\overleftarrow{Y}^v}\big)$ is not unique.
$\Rightarrow$ Constrain controls or add additional regularizers to obtain, e.g., Schrödinger bridges or prescribed probability flows (as in our special cases with fixed v^*).

F. Vargas and N. Nüsken. Transport, variational inference and diffusions: with applications to annealed flows and Schrödinger bridges. *arXiv preprint arXiv:2307.01050*, 2023

G.-H. Liu, T. Chen, O. So, and E. A. Theodorou. Deep generalized Schrödinger bridge. *arXiv preprint arXiv:2209.09893*, 2022

L. Richter and J. Berner. Improved sampling via learned diffusions. *International Conference on Learning Representations*, 2024

Sampling via learned diffusions: General case

$$\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w) = \int_0^T \left((u+v) \cdot \left(w + \frac{v-u}{2} \right) + \nabla \cdot (\sigma v - f) \right)(X_s^w, s)\,\mathrm{d}s + \int_0^T (u+v)(X_s^w, s) \cdot \mathrm{d}W_s + \log \frac{p_{\mathrm{prior}}(X_0^w)}{p_{\mathrm{target}}(X_T^w)}$$

Regularizer: General objective $\min_{v,u} D\left(\mathbb{P}_{X^u} \middle| \mathbb{P}_{\overleftarrow{Y}^v}\right)$ is not unique.
$\Rightarrow$ Constrain controls or add additional regularizers to obtain, e.g., Schrödinger bridges or prescribed probability flows (as in our special cases with fixed v^*).

Divergence: Standard choice: $D_{\mathrm{KL}}(\mathbb{P}_{X^u} | \mathbb{P}_{\overleftarrow{Y}^v}) = \mathbb{E}\left[\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^u)\right]$.

$\Rightarrow$ Recovers previously discussed methods when fixing v^*.

F. Vargas and N. Nüsken. Transport, variational inference and diffusions: with applications to annealed flows and Schrödinger bridges. *arXiv preprint arXiv:2307.01050*, 2023
G.-H. Liu, T. Chen, O. So, and E. A. Theodorou. Deep generalized Schrödinger bridge. *arXiv preprint arXiv:2209.09893*, 2022
L. Richter and J. Berner. Improved sampling via learned diffusions. *International Conference on Learning Representations*, 2024

Sampling via learned diffusions: General case

$$\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w) = \int_0^T \Big((u+v) \cdot \Big(w + \frac{v-u}{2} \Big) + \nabla \cdot (\sigma v - f) \Big)(X_s^w, s)\,\mathrm{d}s + \int_0^T (u+v)(X_s^w, s) \cdot \mathrm{d}W_s + \log \frac{p_{\mathrm{prior}}(X_0^w)}{p_{\mathrm{target}}(X_T^w)}$$

Regularizer: General objective $\min_{v,u} D\big(\mathbb{P}_{X^u} \big| \mathbb{P}_{\overleftarrow{Y}^v}\big)$ is not unique.
$\Rightarrow$ Constrain controls or add additional regularizers to obtain, e.g., Schrödinger bridges or prescribed probability flows (as in our special cases with fixed v^*).

Divergence: Standard choice: $D_{\mathrm{KL}}(\mathbb{P}_{X^u} | \mathbb{P}_{\overleftarrow{Y}^v}) = \mathbb{E}\left[\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^u)\right].$

$\Rightarrow$ Recovers previously discussed methods when fixing v^*.

Reverse KL divergence suffers from mode collapse.

F. Vargas and N. Nüsken. Transport, variational inference and diffusions: with applications to annealed flows and Schrödinger bridges. *arXiv preprint arXiv:2307.01050*, 2023
G.-H. Liu, T. Chen, O. So, and E. A. Theodorou. Deep generalized Schrödinger bridge. *arXiv preprint arXiv:2209.09893*, 2022
L. Richter and J. Berner. Improved sampling via learned diffusions. *International Conference on Learning Representations*, 2024

Path space perspective: Different divergences

Other divergence: We propose the log-variance divergence

$$D_{\mathrm{LV}}^{w}(\mathbb{P}_{X^u}, \mathbb{P}_{\overleftarrow{Y}^v}) := \mathrm{Var}\left(\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w)\right).$$

L. Richter, A. Boustati, N. Nüsken, F. Ruiz, and O. D. Akyildiz. VarGrad: low-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems*, 33:13481–13492, 2020

N. Nüsken and L. Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021

Other divergence: We propose the log-variance divergence

$$D^{w}_{\mathrm{LV}}(\mathbb{P}_{X^u}, \mathbb{P}_{\overleftarrow{Y}^v}) := \mathrm{Var}\left(\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w)\right).$$

👍 **off-policy:** (almost) arbitrary choice for w to balance exploration and exploitation.

L. Richter, A. Boustati, N. Nüsken, F. Ruiz, and O. D. Akyildiz. VarGrad: low-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems*, 33:13481–13492, 2020

N. Nüsken and L. Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021

Path space perspective: Different divergences

Other divergence: We propose the log-variance divergence

$$D^w_{\mathrm{LV}}(\mathbb{P}_{X^u}, \mathbb{P}_{\overleftarrow{Y}^v}) := \mathrm{Var}\left(\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w)\right).$$

👍 **off-policy:** (almost) arbitrary choice for w to balance exploration and exploitation.
👍 **sticking-the-landing:** zero variance at the optimum.

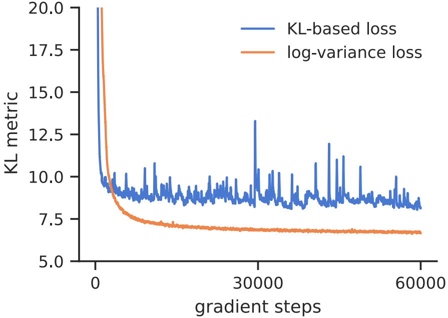

L. Richter, A. Boustati, N. Nüsken, F. Ruiz, and O. D. Akyildiz. VarGrad: low-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems*, 33:13481–13492, 2020

N. Nüsken and L. Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021

Other divergence: We propose the log-variance divergence

$$D_{\mathrm{LV}}^{w}(\mathbb{P}_{X^u}, \mathbb{P}_{\overleftarrow{Y}^v}) := \mathrm{Var}\left(\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w)\right).$$

- **off-policy:** (almost) arbitrary choice for w to balance exploration and exploitation.
- **sticking-the-landing:** zero variance at the optimum.
- **black-box SDE solver / more efficient:** no differentiation through the SDE solver.

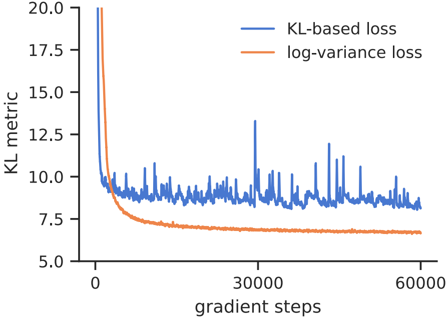

L. Richter, A. Boustati, N. Nüsken, F. Ruiz, and O. D. Akyildiz. VarGrad: low-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems*, 33:13481–13492, 2020
N. Nüsken and L. Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021

Other divergence: We propose the log-variance divergence

$$D^{w}_{\mathrm{LV}}(\mathbb{P}_{X^u}, \mathbb{P}_{\overleftarrow{Y}^v}) := \mathrm{Var}\left(\log \frac{\mathrm{d}\mathbb{P}_{X^u}}{\mathrm{d}\mathbb{P}_{\overleftarrow{Y}^v}}(X^w)\right).$$

- **off-policy:** (almost) arbitrary choice for w to balance exploration and exploitation.
- **sticking-the-landing:** zero variance at the optimum.
- **black-box SDE solver / more efficient:** no differentiation through the SDE solver.
- **zero-th order method:** No gradients of p_{target} required (recall that the RND contains $-\log p_{\mathrm{target}}(X^w_T)$).

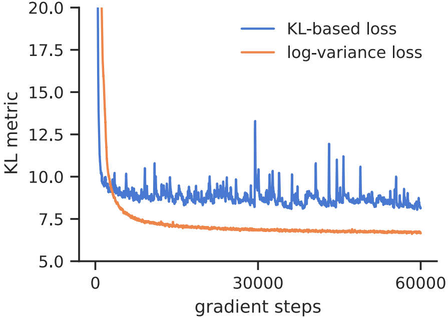

L. Richter, A. Boustati, N. Nüsken, F. Ruiz, and O. D. Akyildiz. VarGrad: low-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems*, 33:13481–13492, 2020

N. Nüsken and L. Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021

Gaussian mixture

Better mode coverage: Improved performance with D_{LV} (compared against D_{KL}) for PIS, DIS, DDS, and the general bridge sampler.

Better mode coverage: Improved performance with D_{LV} (compared against D_{KL}) for PIS, DIS, DDS, and the general bridge sampler.

Method	Divergence	$\Delta \log \mathcal{Z}$ (rw) ↓	$\mathcal{W}_\gamma^2$ ↓	ESS ↑	Δstd ↓
CRAFT		0.012	0.020	-	0.019
PIS	KL	0.249	0.467	0.0051	1.937
	LV	**0.001**	**0.020**	**0.9093**	**0.023**
DIS	KL	**0.015**	0.064	0.0226	2.522
	LV	0.017	**0.020**	**0.8660**	**0.004**
DDS	KL	0.005	0.042	0.0737	2.161
	LV	**0.001**	**0.020**	**0.8929**	**0.006**
Bridge	KL	0.560	0.393	0.0180	0.698
	LV	**0.100**	**0.020**	**0.9669**	**0.010**

Better mode coverage: Improved performance with D_{LV} (compared against D_{KL}) for PIS, DIS, DDS, and the general bridge sampler.

Method	Divergence	$\Delta \log \mathcal{Z}$ (rw) ↓	$\mathcal{W}_\gamma^2$ ↓	ESS ↑	Δstd ↓
CRAFT		0.012	0.020	-	0.019
PIS	KL	0.249	0.467	0.0051	1.937
	LV	**0.001**	**0.020**	**0.9093**	**0.023**
DIS	KL	**0.015**	0.064	0.0226	2.522
	LV	0.017	**0.020**	**0.8660**	**0.004**
DDS	KL	0.005	0.042	0.0737	2.161
	LV	**0.001**	**0.020**	**0.8929**	**0.006**
Bridge	KL	0.560	0.393	0.0180	0.698
	LV	**0.100**	**0.020**	**0.9669**	**0.010**

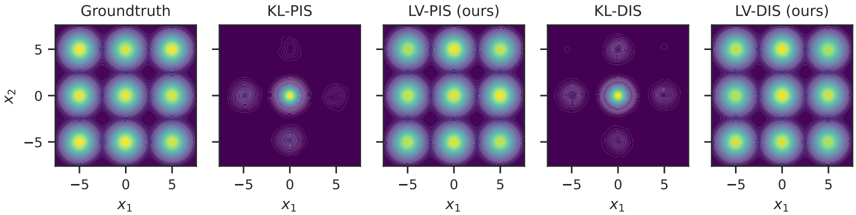

Two-dimensional image distribution

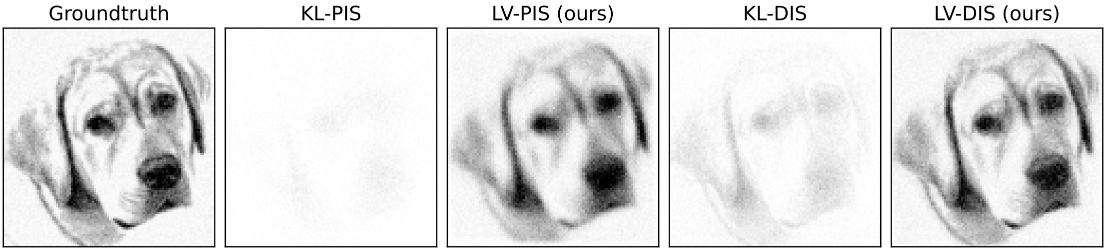

Many-Well: resembles typical problems in molecular dynamics with

$$\log q_{\mathrm{target}}(x) = -\sum_{i=1}^{w}(x_i^2 - \delta)^2 - \frac{1}{2}\sum_{i=w+1}^{d} x_i^2.$$

Many-Well: resembles typical problems in molecular dynamics with

$$\log q_{\mathrm{target}}(x) = -\sum_{i=1}^{w}(x_i^2-\delta)^2 - \frac{1}{2}\sum_{i=w+1}^{d} x_i^2.$$

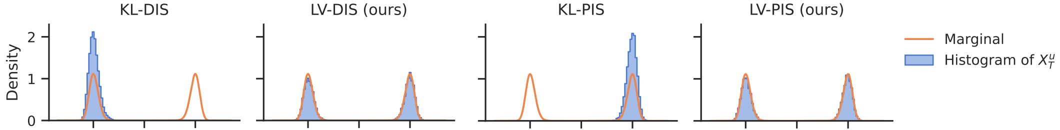

Many-Well: resembles typical problems in molecular dynamics with

$$\log q_{\mathrm{target}}(x) = -\sum_{i=1}^{w}(x_i^2-\delta)^2 - \frac{1}{2}\sum_{i=w+1}^{d} x_i^2 .$$

Problem	Method	$\Delta \log \mathcal{Z} \downarrow$	$\mathcal{W}_\gamma^2 \downarrow$	ESS $\uparrow$	Δ std $\downarrow$
Many-Well	PIS-KL	3.567	1.699	0.0004	1.409
($d=5, w=5, \delta=4$)	**PIS-LV**	**0.214**	**0.121**	**0.6744**	**0.001**
	DIS-KL	1.462	1.175	0.0012	0.431
	DIS-LV	**0.375**	**0.120**	**0.4519**	**0.001**
Many-Well	PIS-KL	0.101	**6.821**	0.8172	0.001
($d=50, w=5, \delta=2$)	**PIS-LV**	**0.087**	6.823	**0.8453**	**0.000**
	DIS-KL	1.785	**6.854**	0.0225	0.009
	DIS-LV	**1.783**	6.855	**0.0227**	0.009

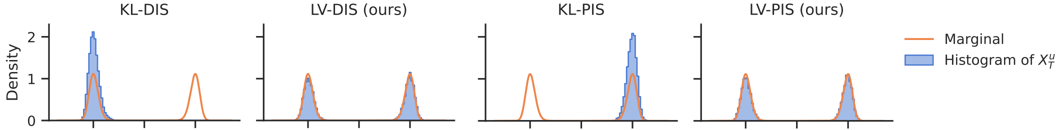

Conclusion and Outlook

Summary:

- Unifying framework for diffusion-based samplers.

Conclusion and Outlook

Summary:

- Unifying framework for diffusion-based samplers.
- Connections between sampling/generative modeling, optimal control, PDEs, and path space measures.

Conclusion and Outlook

Summary:

- Unifying framework for diffusion-based samplers.
- Connections between sampling/generative modeling, optimal control, PDEs, and path space measures.
- Path space perspective allows the usage of general divergences: Log-variance divergence improves upon previous methods.

Conclusion and Outlook

Summary:

- Unifying framework for diffusion-based samplers.
- Connections between sampling/generative modeling, optimal control, PDEs, and path space measures.
- Path space perspective allows the usage of general divergences: Log-variance divergence improves upon previous methods.
- Competitive to samplers based on SMC and normalizing flows.

Conclusion and Outlook

Summary:

- Unifying framework for diffusion-based samplers.
- Connections between sampling/generative modeling, optimal control, PDEs, and path space measures.
- Path space perspective allows the usage of general divergences: Log-variance divergence improves upon previous methods.
- Competitive to samplers based on SMC and normalizing flows.

Outlook:

- Leverage the underlying PDEs and its stochastic representations.

Conclusion and Outlook

Summary:

- Unifying framework for diffusion-based samplers.
- Connections between sampling/generative modeling, optimal control, PDEs, and path space measures.
- Path space perspective allows the usage of general divergences: Log-variance divergence improves upon previous methods.
- Competitive to samplers based on SMC and normalizing flows.

Outlook:

- Leverage the underlying PDEs and its stochastic representations.
- Further explore connections to posterior sampling.

Conclusion and Outlook

Summary:

- Unifying framework for diffusion-based samplers.
- Connections between sampling/generative modeling, optimal control, PDEs, and path space measures.
- Path space perspective allows the usage of general divergences: Log-variance divergence improves upon previous methods.
- Competitive to samplers based on SMC and normalizing flows.

Outlook:

- Leverage the underlying PDEs and its stochastic representations.
- Further explore connections to posterior sampling.
- Optimization without simulation of entire trajectories.

Thank you for your attention!

Github: `https://github.com/juliusberner/sde_sampler`
Mail: `mail@jberner.info`
Website: `https://jberner.info`

References

- J. Berner, L. Richter, K. Ullrich. *An optimal control perspective on diffusion-based generative modeling*. TMLR, 2024.
- L. Richter, J. Berner. *Improved sampling via learned diffusions*. ICLR, 2024.