

The background features a complex network of thin, light gray lines connecting various sized gray circles of different shades. A large, semi-transparent light gray circle is centered on the slide, serving as a backdrop for the text.

Neural Samplers and PDE Solvers Based on Diffusion Processes

Julius Berner

Caltech

May 2, 2024

Overview


Overview

- SDE-based losses for neural PDE solver




Overview

- SDE-based losses for neural PDE solver
- The role of PDEs in generative modeling and sampling



Overview

- SDE-based losses for neural PDE solver
- The role of PDEs in generative modeling and sampling
- Diffusion models, Schrödinger bridges, and optimal control



Neural PDE solver

Partial differential equation (PDEs)

Differential equations [...] represent the most powerful tool humanity has ever created for making sense of the material world. { Steven Strogatz (2009)

Partial differential equation (PDEs)

Differential equations [...] represent the most powerful tool humanity has ever created for making sense of the material world. { Steven Strogatz (2009)

Fluid dynamics

Atmospheric and oceanic modeling

Pressure

Horizontal velocity

Initial geopotential height

After 1h

Partial differential equation (PDEs)

Differential equations [...] represent the most powerful tool humanity has ever created for making sense of the material world. { Steven Strogatz (2009)

Fluid dynamics

Atmospheric and oceanic modeling

Pressure

Horizontal velocity

Initial geopotential height

After 1h

Goal: Approximate the solution $u \in V$ of a PDE

$$Lu = 0 \quad \text{on } D;$$

given by a differential operator L on a domain D with suitable boundary conditions on ∂D .

Partial differential equation (PDEs)

Differential equations [...] represent the most powerful tool humanity has ever created for making sense of the material world. { Steven Strogatz (2009)

Fluid dynamics

Atmospheric and oceanic modeling

Pressure

Horizontal velocity

Initial geopotential height

After 1h

Goal: Approximate the solution $u \in V$ of a PDE

$$Lu = 0 \quad \text{on } D;$$

given by a differential operator L on a domain D with suitable boundary conditions on ∂D .

Diffusion models

- $D = \mathbb{R}^d \quad [0; T]$
- $Lu = \partial_t u + \text{div}(u \cdot \nabla) u - \frac{1}{2} \Delta u$
- $u(\cdot; 0) = p_{\text{target}}$

Finite difference method (FDM)

Approximate the differential operator L by difference quotients on a grid.

Finite difference method (FDM)

Approximate the differential operator L by difference quotients on a grid.

For instance:

$$\partial_{x_i x_i} u(x) \approx \frac{u(x_i - h) - 2u(x_i) + u(x_i + h))}{h^2}$$

Finite difference method (FDM)

Approximate the differential operator L by difference quotients on a grid.

For instance:

$$\partial_{x_i x_i} u(x) \approx \frac{u(x_i - h) - 2u(x_i) + u(x_i + h))}{h^2}$$

) linear system for

$$(u(x + jhe_i))_{i,j}$$

Selected classical methods

Finite difference method (FDM)

Approximate the differential operator L by difference quotients on a grid.

For instance:

$$\partial_{x_i x_i} u(x) \approx \frac{u(x_i - h) - 2u(x_i) + u(x_i + h))}{h^2}$$

) linear system for

$$(u(x + jhe_i))_{i,j}$$

Finite element method (FEM)

Approximate the function space V by a finite-dimensional space.

Selected classical methods

Finite difference method (FDM)

Approximate the differential operator L by difference quotients on a grid.

For instance:

$$\partial_{x_i}^2 u(x) \approx \frac{u(x_i - h) - 2u(x_i) + u(x_i + h))}{h^2}$$

) linear system for

$$(u(x + jh))_{i,j}$$

Finite element method (FEM)

Approximate the function space V by a finite-dimensional space.

For instance (Galerkin approx. for linear u):

$$u(x) \approx \sum_{i=1}^n c_i v_i(x);$$

where $(v_i)_{i=1}^n \subset V$ are suitable functions.

Selected classical methods

Finite difference method (FDM)

Approximate the differential operator L by difference quotients on a grid.

For instance:

$$\partial_{x_i}^2 u(x) \approx \frac{u(x_i - h) - 2u(x_i) + u(x_i + h))}{h^2}$$

) linear system for

$$(u(x + jh))_{i,j}$$

Finite element method (FEM)

Approximate the function space V by a finite-dimensional space.

For instance (Galerkin approx. for linear u):

$$u(x) \approx \sum_{i=1}^n c_i v_i(x);$$

where $(v_i)_{i=1}^n \subset V$ are suitable functions.

) linear system:

$$\sum_{i=1}^n c_i \mathbf{L}(v_i); v_j = \mathbf{L}f; v_j$$

Selected classical methods

Finite difference method (FDM)

Approximate the differential operator L by difference quotients on a grid.

For instance:

$$\partial_{x_i}^2 u(x) \approx \frac{u(x_i - h) - 2u(x_i) + u(x_i + h))}{h^2}$$

) linear system for

$$(u(x + jh))_{i,j}$$

Finite element method (FEM)

Approximate the function space V by a finite-dimensional space.

For instance (Galerkin approx. for linear L):

$$u(x) \approx \sum_{i=1}^n c_i v_i(x);$$

where $(v_i)_{i=1}^n \subset V$ are suitable functions.

) linear system:

$$\sum_{i=1}^n c_i L(v_i); v_j = f; v_j$$

high accuracies, theoretical guarantees

needs to be tailored to the PDE, slow/infeasible for fine grids or high dims.

Fig. 2: <https://de.wikipedia.org/wiki/Datei:Elmer-pump-heatequation.png>

Surrogate models / supervised:

Given data (from classical solvers, simulations, or real-world observations) learn the mapping from parameters/inputs to the PDE solution.

Initial Condition Prediction (5 steps) Ground Truth

Neural PDE Solvers

Surrogate models / supervised:

Given data (from classical solvers, simulations, or real-world observations) learn the mapping from parameters/inputs to the PDE solution.

Initial Condition Prediction (5 steps) Ground Truth

simple objective, PDE can be unknown
needs (high-resolution) training data

Neural PDE Solvers

Surrogate models / supervised:

Given data (from classical solvers, simulations, or real-world observations) learn the mapping from parameters/inputs to the PDE solution.

Variational approaches / unsupervised:

Minimize a variational formulation, where the minimizer corresponds to the PDE solution.

Initial Condition

Prediction (5 steps)

Ground Truth

simple objective, PDE can be unknown
needs (high-resolution) training data

Neural PDE Solvers

Surrogate models / supervised:

Given data (from classical solvers, simulations, or real-world observations) learn the mapping from parameters/inputs to the PDE solution.

Variational approaches / unsupervised:

Minimize a variational formulation, where the minimizer corresponds to the PDE solution.

- PDE-residual (PINNs, Deep Galerkin)

$$\min \|Lu - k\| + \text{BoundaryLoss}(u)$$

Initial Condition Prediction (5 steps) Ground Truth

simple objective, PDE can be unknown
needs (high-resolution) training data

Neural PDE Solvers

Surrogate models / supervised:

Given data (from classical solvers, simulations, or real-world observations) learn the mapping from parameters/inputs to the PDE solution.

Variational approaches / unsupervised:

Minimize a variational formulation, where the minimizer corresponds to the PDE solution.

- PDE-residual (PINNs, Deep Galerkin)
 $\min \|Lu - k\| + \text{BoundaryLoss}(u)$
- Weak formulations / energy-based (Deep-Ritz, Friedrichs learning, ...)

Initial Condition Prediction (5 steps) Ground Truth

PINN Deep Ritz Ground Truth

simple objective, PDE can be unknown
needs (high-resolution) training data

Neural PDE Solvers

Surrogate models / supervised:

Given data (from classical solvers, simulations, or real-world observations) learn the mapping from parameters/inputs to the PDE solution.

Variational approaches / unsupervised:

Minimize a variational formulation, where the minimizer corresponds to the PDE solution.

- PDE-residual (PINNs, Deep Galerkin)
 $\min \|Lu - k\| + \text{BoundaryLoss}(u)$
- Weak formulations / energy-based (Deep-Ritz, Friedrichs learning, ...)

Initial Condition Prediction (5 steps) Ground Truth

simple objective, PDE can be unknown
needs (high-resolution) training data

PINN Deep Ritz Ground Truth

applicable to most PDEs, unlimited data
sensitive to hyperparameters, (higher-order) derivatives needed

SDE-based neural PDE solver



Stochastic representations / (un)supervised:

For certain PDEs, solutions can be written as expectations over suitable random variables and we can regress noisy but cheap simulations.

Neural PDE solvers based on stochastic representations

Stochastic representations / (un)supervised:

For certain PDEs, solutions can be written as expectations over suitable random variables and we can regress noisy but cheap simulations.

without derivatives, include boundary conditions without penalty terms

Neural PDE solvers based on stochastic representations

Stochastic representations / (un)supervised:

For certain PDEs, solutions can be written as expectations over suitable random variables and we can regress noisy but cheap simulations.

without derivatives, include boundary conditions without penalty terms

Kolmogorov (backward) equations

$$\partial_t u = \frac{1}{2} \text{trace} \left(\sigma^T \sigma \right) \Delta u + r^T \nabla u$$

diffusion coefficient

Kolmogorov (backward) equations

$$\partial_t u = \frac{1}{2} \text{trace} \left(\sigma^T \sigma \right) u + r^T u$$

diffusion coefficient / drift coefficient

Kolmogorov (backward) equations

$$\partial_t u = \frac{1}{2} \text{trace} \left(\sigma^T \sigma \right) \Delta u + \mu^T \nabla u - r u; \quad u(x; T) = f(x):$$

diffusion coefficient / drift coefficient

terminal condition

Kolmogorov (backward) equations

$$\partial_t u + \frac{1}{2} \text{trace}(\sigma \sigma^T \nabla^2 u) + \mu \nabla u - r u = 0; \quad u(x; T) = g(x):$$

diffusion coefficient / drift coefficient

terminal condition

Method can be extended to models for drift-diffusion processes (e.g., diffusion models), optimal control problems (HJB equation), fluid dynamics, and option pricing.

Kolmogorov (backward) equations

$$\partial_t u + \frac{1}{2} \text{trace}(\sigma \sigma^T \partial_x^2 u) + \mu \partial_x u - r u = 0; \quad u(x; T) = V(x):$$

diffusion coefficient / drift coefficient

terminal condition

Method can be extended to models for drift-diffusion processes (e.g., diffusion models), optimal control problems (HJB equation), fluid dynamics, and option pricing.

Stochastic representation (Itô's lemma):

$$u(x; 0) = V(X_T^x) + \int_0^T \{Z_t\} dt;$$

noise with zero expectation

where X^x solves the stochastic differential equation (SDE)

$$dX_t^x = \mu(X_t^x) dt + \sigma(X_t^x) dB_t; \quad X_0^x = x:$$

Kolmogorov (backward) equations

$$\partial_t u + \frac{1}{2} \text{trace}(\sigma \sigma^T \partial_x^2 u) + \mu \partial_x u - r u = 0; \quad u(x; T) = V(x):$$

diffusion coefficient / drift coefficient

terminal condition

Method can be extended to models for drift-diffusion processes (e.g., diffusion models), optimal control problems (HJB equation), fluid dynamics, and option pricing.

Stochastic representation (Itô's lemma / Feynman-Kac formula):

$$u(x; 0) = V(X_T^x) \quad \mathbb{E}[\cdot | \mathcal{Z}_t] = \mathbb{E}[\cdot | X_t^x]$$

noise with zero expectation

where X^x solves the stochastic differential equation (SDE)

$$dX_t^x = \mu(X_t^x) dt + \sigma(X_t^x) dB_t; \quad X_0^x = x:$$

Euler-Maruyama scheme and Monte Carlo sampling

Discretize the SDE

$$dX_t^x = \mu(X_t^x) dt + \sigma(X_t^x) dB_t; \quad X_0^x = x;$$

using the Euler-Maruyama scheme with step size Δt :

$$\hat{X}_{t+\Delta t}^x = \hat{X}_t^x + \mu(\hat{X}_t^x) \Delta t + \sigma(\hat{X}_t^x) \left\{ \frac{B_{t+\Delta t} - B_t}{\sqrt{\Delta t}} \right\};$$

Euler-Maruyama scheme and Monte Carlo sampling

Discretize the SDE

$$dX_t^x = \mu(X_t^x) dt + \sigma(X_t^x) dB_t; \quad X_0^x = x;$$

using the Euler-Maruyama scheme with step size Δt :

$$\hat{X}_{t+\Delta t}^x = \hat{X}_t^x + \mu(\hat{X}_t^x) \Delta t + \sigma(\hat{X}_t^x) \left\{ \frac{B_{t+\Delta t} - B_t}{\sqrt{\Delta t}} \right\};$$

- Use Monte Carlo (MC) sampling to approximate the solution $u(x; 0) = E[\hat{X}_T^x]$ with MSE $O(\Delta t)$:

SDE-based neural PDE solver

Monte Carlo sampling gives us pointwise estimates of $u(x; 0)$ for fixed $x \in \mathbb{R}^d$.

Monte Carlo sampling gives us pointwise estimates of $u(x; 0)$ for fixed $x \in \mathbb{R}^d$.

- We can extend this framework to learn $u(x; 0)$ on a domain $K \subset \mathbb{R}^d$.

SDE-based neural PDE solver

Monte Carlo sampling gives us pointwise estimates of $u(x; 0)$ for fixed $x \in \mathbb{R}^d$.

- We can extend this framework to learn $u(\cdot; 0)$ on a domain $K \subset \mathbb{R}^d$.

Variational formulation: $u(\cdot; 0) = \operatorname{argmin}_v \mathbb{E} \left[\int_0^T (v(X_t) - v(\cdot))^2 dt \right]$:

- We extend the framework to parametric PDEs and space-time domains $K \subset [0; T] \times \mathbb{R}^d$.

- We extend the framework to parametric PDEs and space-time domains $K \subset [0; T] \times \mathbb{R}^d$.
- We prove bounds on the approx. and generalization error (can overcome the curse of dimensionality).

- We extend the framework to parametric PDEs and space-time domains $K \subset [0; T] \times \mathbb{R}^d$.
- We prove bounds on the approx. and generalization error (can overcome the curse of dimensionality).
- We provide variance-reduced losses (based on approximating g_u).

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

J. Berner, P. Grohs, and F. Voigtlaender. Learning ReLU networks to high uniform accuracy is intractable. *International Conference on Learning Representations*, 2023

J. Berner, P. Grohs, and A. Jentzen. Analysis of the generalization error [...]. *SIAM Journal on Mathematics of Data Science*, 2(3):631–657, 2020

J. Berner, M. Dablander, and P. Grohs. Numerically solving parametric families of high-dimensional Kolmogorov PDEs via deep learning. *Advances in Neural Information Processing Systems*, pages 16615–16627, 2020

SDE-based neural PDE solver

- We extend the framework to parametric PDEs and space-time domains $K \subset [0; T] \times \mathbb{R}^d$.
- We prove bounds on the approx. and generalization error (can overcome the curse of dimensionality).
- We provide variance-reduced losses (based on approximating $\mathbb{E}u$).
- We extend it to boundary value problems using walk-on-spheres.

H. C. Nam, J. Berner, and A. Anandkumar. Solving Poisson equations using neural walk-on-spheres. [ICLR 2024 Workshop on AI4Differential Equations](#). In *Science*, 2024

L. Richter and J. Berner. Robust SDE-based variational formulations for solving linear PDEs via deep learning. [International Conference on Machine Learning](#), pages 18649–18666, 2022

J. Berner, P. Grohs, and F. Voigtlaender. Learning ReLU networks to high uniform accuracy is intractable. [International Conference on Learning Representations](#), 2023

J. Berner, P. Grohs, and A. Jentzen. Analysis of the generalization error [...]. [SIAM Journal on Mathematics of Data Science](#), 2(3):631–657, 2020

J. Berner, M. Dablander, and P. Grohs. Numerically solving parametric families of high-dimensional Kolmogorov PDEs via deep learning. [Advances in Neural Information Processing Systems](#), pages 16615–16627, 2020

SDE-based neural PDE solver - Results

- We solve (parametric) PDEs from optimal control, chemistry, and finance with \triangleright 50 dims. (out-of-scope for classical methods).

SDE-based neural PDE solver - Results

- We solve (parametric) PDEs from optimal control, chemistry, and finance with $\gg 50$ dims. (out-of-scope for classical methods).
- Numerical evidence that our method can overcome the curse of dimensionality.

SDE-based neural PDE solver - Results

- We solve (parametric) PDEs from optimal control, chemistry, and finance with $\gg 50$ dims. (out-of-scope for classical methods).
- Numerical evidence that our method can overcome the curse of dimensionality.
- More efficient and accurate than competing methods (e.g., PINNs, Deep Ritz).

SDE-based neural PDE solver - Results

- We solve (parametric) PDEs from optimal control, chemistry, and finance with $\gg 50$ dims. (out-of-scope for classical methods).
- Numerical evidence that our method can overcome the curse of dimensionality.
- More efficient and accurate than competing methods (e.g., PINNs, Deep Ritz).
- Extensions: Combining SDE-based losses with neural operators for the Navier-Stokes eq.

R. Zhang, Q. Meng, R. Zhu, Y. Wang, W. Shi, S. Zhang, Z.-M. Ma, and T.-Y. Liu. Monte Carlo neural operator for learning PDEs via probabilistic representation. [arXiv:2302.05104](https://arxiv.org/abs/2302.05104), 2023

J. Berner, M. Dablander, and P. Grohs. Numerically solving parametric families of high-dimensional Kolmogorov PDEs via deep learning. [Advances in Neural Information Processing Systems](https://arxiv.org/abs/2006.16627), pages 16615-16627, 2020

L. Richter and J. Berner. Robust SDE-based variational formulations for solving linear PDEs via deep learning. [International Conference on Machine Learning](https://arxiv.org/abs/2206.18666), pages 18649-18666, 2022

H. C. Nam, J. Berner, and A. Anandkumar. Solving Poisson equations using neural walk-on-spheres. [ICLR 2024 Workshop on AI4Differential Equations in Science](https://arxiv.org/abs/2402.11441), 2024

SDE-based generative modeling and sampling




Image & Video Generation

Motivation: Selected applications of diffusion models

Image & Video Generation

Biology and Chemistry

Motivation: Selected applications of diffusion models

Image & Video Generation

Biology and Chemistry

Language modeling

Fig. 3: A. Lou, C. Meng, and S. Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution. [arXiv preprint arXiv:2310.16834](#), 2023

Fig. 2: B. L. Trippe, J. Yim, D. Tischer, T. Broderick, D. Baker, R. Barzilay, and T. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scoring problem. [arXiv preprint arXiv:2206.04119](#), 2022

Fig. 1: [https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_\(SDXL\).jpg](https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_(SDXL).jpg)

Motivation: Selected applications of diffusion models

Image & Video Generation

Biology and Chemistry

Language modeling

Inverse Problems

Fig. 4: H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye. Diffusion posterior sampling for general noisy inverse problems. [arXiv preprint arXiv:2209.14687](#), 2022

Fig. 3: A. Lou, C. Meng, and S. Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution. [arXiv preprint arXiv:2310.16834](#), 2023

Fig. 2: B. L. Trippe, J. Yim, D. Tischer, T. Broderick, D. Baker, R. Barzilay, and T. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scolding problem. [arXiv preprint arXiv:2206.04119](#), 2022

Fig. 1: [https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_\(SDXL\).jpg](https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_(SDXL).jpg)

Motivation: Selected applications of diffusion models

Image & Video Generation

Biology and Chemistry

Language modeling

Inverse Problems

Neural Compression

Fig. 5: <https://github.com/juliusberner/NeuralCompression>

Fig. 4: H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye. Diffusion posterior sampling for general noisy inverse problems. [arXiv preprint arXiv:2209.14687](https://arxiv.org/abs/2209.14687), 2022

Fig. 3: A. Lou, C. Meng, and S. Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution. [arXiv preprint arXiv:2310.16834](https://arxiv.org/abs/2310.16834), 2023

Fig. 2: B. L. Trippe, J. Yim, D. Tischer, T. Broderick, D. Baker, R. Barzilay, and T. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scoring problem. [arXiv preprint arXiv:2206.04119](https://arxiv.org/abs/2206.04119), 2022

Fig. 1: [https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_\(SDXL\).jpg](https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_(SDXL).jpg)

Motivation: Selected applications of diffusion models

Image & Video Generation

Biology and Chemistry

Language modeling

Inverse Problems

Neural Compression

Sampling Problems

Fig. 6: L. Richter and J. Berner. Improved sampling via learned diffusions. In *International Conference on Learning Representations*, 2024

Fig. 5: <https://github.com/juliusberner/NeuralCompression>

Fig. 4: H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022

Fig. 3: A. Lou, C. Meng, and S. Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution. *arXiv preprint arXiv:2310.16834*, 2023

Fig. 2: B. L. Trippe, J. Yim, D. Tischer, T. Broderick, D. Baker, R. Barzilay, and T. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scoring problem. *arXiv preprint arXiv:2206.04119*, 2022

Fig. 1: [https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_\(SDXL\).jpg](https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_(SDXL).jpg)

Motivation: Selected applications of diffusion models

Image & Video Generation

Biology and Chemistry

Language modeling

Inverse Problems

Neural Compression

Sampling Problems



Fig. 6: L. Richter and J. Berner. Improved sampling via learned diffusions. In *International Conference on Learning Representations*, 2024

Fig. 5: <https://github.com/juliusberner/NeuralCompression>

Fig. 4: H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye. Diffusion posterior sampling for general noisy inverse problems. [arXiv preprint arXiv:2209.14687](https://arxiv.org/abs/2209.14687), 2022

Fig. 3: A. Lou, C. Meng, and S. Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution. [arXiv preprint arXiv:2310.16834](https://arxiv.org/abs/2310.16834), 2023

Fig. 2: B. L. Trippe, J. Yim, D. Tischer, T. Broderick, D. Baker, R. Barzilay, and T. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scoring problem. [arXiv preprint arXiv:2206.04119](https://arxiv.org/abs/2206.04119), 2022

Fig. 1: [https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_\(SDXL\).jpg](https://en.wikipedia.org/wiki/File:Astronaut_Riding_a_Horse_(SDXL).jpg)

Sampling from high-dimensional distributions

Task: Sample from a high-dimensional distribution p_{target} .

Task: Sample from a high-dimensional distribution p_{target} .

p_{target} can be given in the form of:

1. samples $x^{(i)} \sim p_{\text{target}}$ (images, text, sound, ...).

Task: Sample from a high-dimensional distribution p_{target} .

p_{target} can be given in the form of:

1. samples $x^{(i)} \sim p_{\text{target}}$ (images, text, sound, ...).
2. an (unnormalized) density q_{target} with $p_{\text{target}} = q_{\text{target}} / Z$ (e.g., in Bayesian statistics, comp. physics, and chemistry).

Task: Sample from a high-dimensional distribution p_{target} .

p_{target} can be given in the form of:

1. samples $x^{(i)} \sim p_{\text{target}}$ (images, text, sound, ...).
2. an (unnormalized) density q_{target} with $p_{\text{target}} = \frac{q_{\text{target}}}{Z}$ (e.g., in Bayesian statistics, comp. physics, and chemistry).

Generative modeling

Task: Sample from a high-dimensional distribution p_{target} .

p_{target} can be given in the form of:

1. samples $x^{(i)} \sim p_{\text{target}}$ (images, text, sound, ...).
2. an (unnormalized) density q_{target} with $p_{\text{target}} = \frac{q_{\text{target}}}{Z}$ (e.g., in Bayesian statistics, comp. physics, and chemistry).

Generative modeling

Sampling problems

Overview of generative models

Connections: Diffusion models give rise to continuous-time normalizing flows and can be viewed as hierarchical VAEs with fixed encoder.

Connections: Diffusion models give rise to continuous-time normalizing flows and can be viewed as hierarchical VAEs with fixed encoder.

Tasks: Flow-based models have been used for Tasks 1 & 2. We show that this is also possible for diffusion models.

- We cannot use denoising score matching (or other adaptations of likelihood training / forward-KL divergence):

$$E_{x \sim p_{\text{target}}; n \sim N(0; 2I)} \sum_{i=1}^h u(x + n; \theta) x^{2^i}$$

Problems

- We cannot use denoising score matching (or other adaptations of likelihood training / forward-KL divergence):

$$E_{x \sim p_{\text{target}}; n \sim N(0; 2I)} \sum_{i=1}^h u(x + n; \theta) x^{2^i}$$

- Requires samples $x \sim p_{\text{target}}$ from the target distribution.

- We cannot use denoising score matching (or other adaptations of likelihood training / forward-KL divergence):

$$E_{x \sim p_{\text{target}}; n \sim N(0; 2I)} \sum_{i=1}^h u(x + n; \cdot) x^{2^i}$$

- Requires samples $x \sim p_{\text{target}}$ from the target distribution.
- Importance sampling does not scale to high dims.

Problems

- We cannot use denoising score matching (or other adaptations of likelihood training / forward-KL divergence):

$$E_{x \sim p_{\text{target}}; n \sim N(0; 2I)} \sum_{i=1}^h u(x + n; \theta) x^{2^i}$$

- Requires samples $x \sim p_{\text{target}}$ from the target distribution.
 - Importance sampling does not scale to high dims.
- We cannot use reverse-KL divergence of normalizing flows:

$$\min D_{\text{KL}}(p_{X_T} \parallel p_{\text{target}})$$

Problems

- We cannot use denoising score matching (or other adaptations of likelihood training / forward-KL divergence):

$$E_{x \sim p_{\text{target}}; n \sim N(0; \sigma^2 I)} \sum_{i=1}^h u(x + n; \theta) \cdot x^{2^i}$$

- Requires samples $x \sim p_{\text{target}}$ from the target distribution.
 - Importance sampling does not scale to high dims.
- We cannot use reverse-KL divergence of normalizing flows:

$$\min D_{\text{KL}}(p_{X_T} \parallel p_{\text{target}})$$

- The density p_{X_T} can only be evaluated for ODEs, i.e. $dX_s = f(X_s) ds$.

Problems

- We cannot use denoising score matching (or other adaptations of likelihood training / forward-KL divergence):

$$\mathbb{E}_{x \sim p_{\text{target}}; n \sim N(0; \sigma^2 I)} \sum_{i=1}^h u(x + n; \theta) x^{2^i}$$

- Requires samples $x \sim p_{\text{target}}$ from the target distribution.
 - Importance sampling does not scale to high dims.
- We cannot use reverse-KL divergence of normalizing flows:
$$\min D_{\text{KL}}(p_{X_T} \parallel p_{\text{target}})$$
 - The density p_{X_T} can only be evaluated for ODEs, i.e. $dX_s = \sigma(X_s) ds$.
 - Cannot use the probability flow ODE (score $\log p_X$ not available during training).

Problems

- We cannot use denoising score matching (or other adaptations of likelihood training / forward-KL divergence):

$$\mathbb{E}_{x \sim p_{\text{target}}; n \sim \mathcal{N}(0; 2I)} \sum_{i=1}^h u(x + n; \theta) \cdot x^2$$

- Requires samples $x \sim p_{\text{target}}$ from the target distribution.
 - Importance sampling does not scale to high dims.
- We cannot use reverse-KL divergence of normalizing flows:

$$\min D_{\text{KL}}(p_{X_T} \parallel p_{\text{target}})$$

- The density p_{X_T} can only be evaluated for ODEs, i.e. $\dot{X}_s = f(X_s)$.
- Cannot use the probability flow ODE (score $\log p_X$ not available during training).

Spoiler: Take the divergence on the path space $C([0; T]; \mathbb{R}^d)$ rather than the target space \mathbb{R}^d , i.e., $D_{\text{KL}}(P_{X^u} \parallel P_{\text{target}})$.

SDE-based generative modeling and sampling

Goal: Learn the drift of a SDE such that X_T approximates a given distribution.

Goal: $X_T \sim p_{\text{target}}$ with generative SDE

$$dX_s = \mu(X_s; s) ds + \sigma(X_s; s) dB_s; \quad X_0 \sim p_{\text{prior}};$$

$\sigma = 0$: continuous-time normalizing flow.

Connections to PDEs

Goal: $X_T \sim p_{\text{target}}$ with generative SDE

$$dX_s = \mu(X_s; s) ds + \sigma(X_s; s) dB_s; \quad X_0 \sim p_{\text{prior}} : \\ = 0: \text{continuous-time normalizing flow .}$$

Fokker-Planck equation for density p_X :

$$\partial_t p_X = -\text{div}(p_X \mu) + \frac{1}{2} \text{tr}(\sigma^T \sigma p_X) : \\ = 0: \text{continuity equation .}$$

Goal: $X_T \sim p_{\text{target}}$ with generative SDE

$$dX_s = \mu(X_s; s) ds + \sigma(X_s; s) dB_s; \quad X_0 \sim p_{\text{prior}}$$

$\sigma = 0$: continuous-time normalizing flow .

Fokker-Planck equation for density p_X :

$$\partial_t p_X = -\text{div}(p_X \mu) + \frac{1}{2} \text{tr}(\sigma^T \sigma p_X)$$

$\sigma = 0$: continuity equation .

Regularization: minimize energy

$$\inf_0^T E \int_0^T \int k(X_s; s)^2 ds$$

$\epsilon > 0$: Dynamic Schrödinger bridge .

$\epsilon = 0$: Optimal W_2 -transport .

Goal: $X_T \sim p_{\text{target}}$ with generative SDE

$$dX_s = \mu(X_s; s) ds + \sigma(X_s; s) dB_s; \quad X_0 \sim p_{\text{prior}}$$

$\sigma = 0$: continuous-time normalizing flow.

Fokker-Planck equation for density p_X :

$$\partial_t p_X = -\text{div}(p_X \mu) + \frac{1}{2} \text{tr}(\sigma^T \sigma p_X)$$

$\sigma = 0$: continuity equation.

Regularization: minimize energy

$$\inf_{\mu, \sigma} \int_0^T \mathbb{E} \| \mu(X_s; s) \|^2 ds$$

$\sigma > 0$: Dynamic Schrödinger bridge.
 $\sigma = 0$: Optimal W_2 -transport.

HJB equation for $\phi = r$:

$$\partial_t \phi = -\mu \cdot \nabla \phi + \frac{1}{2} \text{tr}(\sigma^T \sigma \nabla^2 \phi)$$

Connections to PDEs

Goal: $X_T \sim p_{\text{target}}$ with generative SDE

$$dX_s = \mu(X_s; s) ds + \sigma(X_s; s) dB_s; \quad X_0 \sim p_{\text{prior}}$$

$\sigma = 0$: continuous-time normalizing flow.

Fokker-Planck equation for density p_X :

$$\partial_t p_X = -\text{div}(p_X \mu) + \frac{1}{2} \text{tr}(\sigma^T \sigma p_X)$$

$\sigma = 0$: continuity equation.

Regularization: minimize energy

$$\inf_0^T \mathbb{E} \int_0^T k(X_s; s)^2 ds$$

> 0 : Dynamic Schrödinger bridge.
 $= 0$: Optimal W_2 -transport.

HJB equation for $v = r$:

$$\partial_t v = \frac{1}{2} \text{tr}(\sigma^T \sigma \nabla^2 v) + \mu^T \nabla v$$

- Losses of neural PDE solvers have been used as regularizers.

T. Koshizuka and I. Sato. Neural Lagrangian Schrödinger bridge. In *International Conference on Learning Representations*, 2022

L. Ruthotto, S. J. Osher, et al. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proceedings of the National Academy of Sciences*, 117(17):9183-9193, 2020

L. Yang and G. E. Karniadakis. Potential flow generator with L^2 optimal transport regularity for generative models. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):528-538, 2020

D. Onken, S. W. Fung, X. Li, and L. Ruthotto. OT-Flow: Fast and accurate continuous normalizing flows via optimal transport. In *Proceedings of the AAAI Conference on Artificial Intelligence*, number 10, pages 9223-9232, 2021

J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375-393, 2000

Connections to PDEs

Goal: $X_T \sim p_{\text{target}}$ with generative SDE

$$dX_s = \mu(X_s; s) ds + \sigma(X_s; s) dB_s; \quad X_0 \sim p_{\text{prior}}$$

$\sigma = 0$: continuous-time normalizing flow.

Fokker-Planck equation for density p_X :

$$\partial_t p_X = -\text{div}(p_X \mu) + \frac{1}{2} \text{tr}(\sigma^T \sigma p_X)$$

$\sigma = 0$: continuity equation.

Regularization: minimize energy

$$\inf_0^T \int E \|k(X_s; s)\|^2 ds$$

$\sigma > 0$: Dynamic Schrödinger bridge.
 $\sigma = 0$: Optimal W_2 -transport.

HJB equation for $v = r$:

$$\partial_t v = \frac{1}{2} \text{tr}(\sigma^T \sigma \nabla^2 v) + \mu^T \nabla v$$

- Losses of neural PDE solvers have been used as regularizers.
- Use neural PDE solver for the (log-transformed) Fokker-Planck eq. with $(X; T) = q_{\text{target}}$ to tackle sampling problems

J. Sun, J. Berner, K. Azizzadenesheli, and A. Anandkumar. Physics-informed neural networks for sampling. *ICLR 2024 Workshop on AI4Differential Equations*. In *Science*, 2024

T. Koshizuka and I. Sato. Neural Lagrangian Schrödinger bridge. In *International Conference on Learning Representations*, 2022

L. Ruthotto, S. J. Osher, et al. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proceedings of the National Academy of Sciences*, 117(17):9183-9193, 2020

L. Yang and G. E. Karniadakis. Potential flow generator with L^2 optimal transport regularity for generative models. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):528-538, 2020

D. Onken, S. W. Fung, X. Li, and L. Ruthotto. OT-Flow: Fast and accurate continuous normalizing flows via optimal transport. In *Proceedings of the AAAI Conference on Artificial Intelligence*, number 10, pages 9223-9232, 2021

J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375-393, 2000

Connections to PDEs

Goal: $X_T \sim p_{\text{target}}$ with generative SDE

$$dX_s = \mu(X_s; s) ds + \sigma(X_s; s) dB_s; \quad X_0 \sim p_{\text{prior}}$$

$\sigma = 0$: continuous-time normalizing flow.

Fokker-Planck equation for density p_X :

$$\partial_t p_X = -\text{div}(p_X \mu) + \frac{1}{2} \text{tr}(\sigma^T \sigma p_X)$$

$\sigma = 0$: continuity equation.

Regularization: minimize energy

$$\inf_0^T \int E_k(X_s; s) k^2 ds$$

$\sigma > 0$: Dynamic Schrödinger bridge.
 $\sigma = 0$: Optimal W_2 -transport.

HJB equation for $\phi = r$:

$$\partial_t \phi = \frac{1}{2} k^T r k - \frac{1}{2} \text{tr}(\sigma^T \sigma r)$$

- Losses of neural PDE solvers have been used as regularizers.
- Use neural PDE solver for the (log-transformed) Fokker-Planck eq. with $(X; T) = q_{\text{target}}$ to tackle sampling problems | but no clear benefit of diffusion-based models (i.e., $\sigma > 0$).

J. Sun, J. Berner, K. Azizzadenesheli, and A. Anandkumar. Physics-informed neural networks for sampling. *ICLR 2024 Workshop on AI4Differential Equations*. In *Science*, 2024

T. Koshizuka and I. Sato. Neural Lagrangian Schrödinger bridge. In *International Conference on Learning Representations*, 2022


L. Ruthotto, S. J. Osher, et al. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proceedings of the National Academy of Sciences*, 117(17):9183-9193, 2020

L. Yang and G. E. Karniadakis. Potential flow generator with L^2 optimal transport regularity for generative models. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):528-538, 2020

D. Onken, S. W. Fung, X. Li, and L. Ruthotto. OT-Flow: Fast and accurate continuous normalizing flows via optimal transport. In *Proceedings of the AAAI Conference on Artificial Intelligence*, number 10, pages 9223-9232, 2021

J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375-393, 2000

Diffusion models, Schrödinger bridges, and optimal control



Framework: Time-reversals


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

p_{prior}

p_{target}


Framework: Time-reversals

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



Framework: Time-reversals


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



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

Framework: Time-reversals

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




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

Our setting: Controlled SDEs (notation: $\gamma(t) := (T - t)$)

$$\begin{aligned}dX_s^u &= \gamma(s)u(X_s^u; s) ds + \gamma(s) dW_s; & X_0 & \sim p_{\text{prior}}; \\dY_s^v &= \gamma(s)v(Y_s^v; s) ds + \gamma(s) dW_s; & Y_0 & \sim p_{\text{target}};\end{aligned}$$

Framework: Time-reversals

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



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

Our setting: Controlled SDEs (notation: $\cdot(t) := (T - t)$)

$$\begin{aligned}dX_s^u &= \cdot(s)u(X_s^u; s) ds + \cdot(s) dW_s; & X_0 & \sim p_{\text{prior}}; \\dY_s^v &= \cdot(s)v(Y_s^v; s) ds + \cdot(s) dW_s; & Y_0 & \sim p_{\text{target}};\end{aligned}$$

New objective: Find $u; v$, such that X^u is the time-reversal \check{Y}^v of Y^v , i.e., consider the path measures \mathbb{P}_{X^u} and \mathbb{P}_{Y^v} , a divergence \mathcal{D} , and optimize $\min_{u;v} \mathcal{D}(\mathbb{P}_{X^u}, \mathbb{P}_{Y^v})$:

Unifying objective

How to calculate the divergence $\nabla \cdot (P_X^u P_Y^v)$?

Unifying objective

How to calculate the divergence $D_{KL} P_{X^u} P_{Y^v}$?

- We combine the PDE perspective with tools from optimal control to derive the Log-likelihood for path measures:

$$\log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) = \int_0^T (u + v) \left(w + \frac{v}{2} \frac{u}{v} \right) + \text{div}(v)(X_s^w; s) ds$$
$$+ \int_0^T (u + v)(X_s^w; s) dW_s + \log \frac{p_{\text{prior}}(X_0^w)}{p_{\text{target}}(X_T^w)}$$

| $\underbrace{\hspace{10em}}_{\text{zero expectation}}$ |

Unifying objective

How to calculate the divergence $D(P_{X^u} \| P_{Y^v})$?

- We combine the PDE perspective with tools from optimal control to derive the Log-likelihood for path measures:

$$\log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) = \int_0^T (u + v) \left(w + \frac{v}{2} \frac{u}{2} + \text{div}(v)(X_s^w; s) \right) ds + \int_0^T (u + v)(X_s^w; s) dW_s + \log \frac{p_{\text{prior}}(X_0^w)}{p_{\text{target}}(X_T^w)}$$

| $\underbrace{\hspace{10em}}_{\text{zero expectation}}$ |

Typical choice $D = D_{\text{KL}}$:

Unifying objective

How to calculate the divergence $D(P_{X^u} \| P_{Y^v})$?

- We combine the PDE perspective with tools from optimal control to derive the Log-likelihood for path measures:

$$\log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) = \int_0^{Z_T} (u + v) \left(w + \frac{v}{2} \frac{u}{2} + \text{div}(v) \right) (X_s^w; s) ds + \int_0^{Z_T} (u + v)(X_s^w; s) dW_s + \log \frac{p_{\text{prior}}(X_0^w)}{p_{\text{target}}(X_T^w)}$$

| $\underbrace{\hspace{10em}}_{\text{zero expectation}}$ |

Typical choice $D = D_{\text{KL}}$: Forward-KL recovers objectives from generative modeling:

- "Likelihood Schrödinger bridge" (non-unique)

Unifying objective

How to calculate the divergence $D(P_{X^u} \| P_{Y^v})$?

- We combine the PDE perspective with tools from optimal control to derive the Log-likelihood for path measures:

$$\log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) = \int_0^T (u + v) \left(w + \frac{v}{2} \frac{u}{v} \right) + \text{div}(v)(X_s^w; s) ds + \int_0^T (u + v)(X_s^w; s) dW_s + \log \frac{p_{\text{prior}}(X_0^w)}{p_{\text{target}}(X_T^w)}$$

| $\underbrace{\hspace{10em}}_{\text{zero expectation}}$ |

Typical choice $D = D_{\text{KL}}$: Forward-KL recovers objectives from generative modeling:

- "Likelihood Schrödinger bridge" (non-unique)
- Fixed v : ELBO for diffusion models

Unifying objective

Sampling problems: Consider the reverse-KL

$$D_{\text{KL}}(P_{X^u} \parallel P_{Y^v}) = E_{X^u} \left[\log \frac{dP_{X^u}}{dP_{Y^v}}(X^u) \right]$$

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL}(P_{X^u} \parallel P_{Y^v}) = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u)$:

Bridge sampler:

$$D_{KL}(P_{X^u} \parallel P_{Y^v}) = \int_0^1 \underbrace{\left[\frac{1}{2}ku + vk^2 + \text{div}(v)(X_s^u; s) \right]}_{\text{running cost}} ds + \underbrace{\log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)}}_{\text{initial / terminal cost}} + \underbrace{\log Z}_{\text{const.}}$$

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL}(P_{X^u} \parallel P_{Y^v}) = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u)$:

Bridge sampler:

$$D_{KL}(P_{X^u} \parallel P_{Y^v}) = \int_0^1 \underbrace{\frac{1}{2}ku + vk^2 + \text{div}(v)(X_s^u; s)}_{\text{running cost}} ds + \log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)} + \log \frac{Z}{Z} \Big|_{\text{const.}}$$

initial / terminal cost

We can also derive unique objectives:

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL}(P_{X^u} \parallel P_{Y^v}) = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u) : i$

Bridge sampler:

$$D_{KL}(P_{X^u} \parallel P_{Y^v}) = \int_0^1 \underbrace{\frac{1}{2}ku + vk^2 + \text{div}(v)(X_s^u; s)}_{\text{running cost}} ds + \underbrace{\log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)}}_{\text{initial / terminal cost}} + \underbrace{\log Z}_{\text{const.}}$$

We can also derive unique objectives:

- Fixed v : Time-Reversed Diffusion Sampler (DIS).

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL} P_{X^u} \parallel P_{Y^v} = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u) : i$

Bridge sampler:

$$D_{KL} P_{X^u} \parallel P_{Y^v} = \int_0^1 \frac{1}{2} k u + v k^2 + \text{div}(v) (X_s^u; s) ds + \log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)} + \log \frac{Z}{Z} \text{const.}$$

running cost
initial / terminal cost
const.

We can also derive unique objectives:

- Fixed v : Time-Reversed Diffusion Sampler (DIS).
- Reference process: Path Integral Sampler (PIS) & Denoising Diffusion Sampler (DDS)

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL}(P_{X^u} \parallel P_{Y^v}) = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u) : i$

Bridge sampler:

$$D_{KL}(P_{X^u} \parallel P_{Y^v}) = \int_0^{Z_T} \underbrace{\frac{1}{2}ku + vk^2 + \text{div}(v)}_{\text{running cost}}(X_s^u; s) ds + \log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)} + \log \frac{Z}{\text{const.}}$$

initial / terminal cost

We can also derive unique objectives:

- Fixed v : Time-Reversed Diffusion Sampler (DIS).
- Reference process: Path Integral Sampler (PIS) & Denoising Diffusion Sampler (DDS)

Short-comings:

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL}(P_{X^u} \parallel P_{Y^v}) = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u)$:

Bridge sampler:

$$D_{KL}(P_{X^u} \parallel P_{Y^v}) = \int_0^T \underbrace{\frac{1}{2}ku + vk^2 + \text{div}(v)(X_s^u; s)}_{\text{running cost}} ds + \underbrace{\log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)}}_{\text{initial / terminal cost}} + \underbrace{\log Z}_{\text{const.}}$$

We can also derive unique objectives:

- Fixed v : Time-Reversed Diffusion Sampler (DIS).
- Reference process: Path Integral Sampler (PIS) & Denoising Diffusion Sampler (DDS)

Short-comings:

No exploration: The simulation is "on-policy" (with the learned model).

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL}(P_{X^u} \parallel P_{Y^v}) = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u)$:

Bridge sampler:

$$D_{KL}(P_{X^u} \parallel P_{Y^v}) = \int_0^T \underbrace{\frac{1}{2}ku + vk^2 + \text{div}(v)}_{\text{running cost}}(X_s^u; s) ds + \underbrace{\log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)}}_{\text{initial / terminal cost}} + \underbrace{\log Z}_{\text{const.}}$$

We can also derive unique objectives:

- Fixed v : Time-Reversed Diffusion Sampler (DIS).
- Reference process: Path Integral Sampler (PIS) & Denoising Diffusion Sampler (DDS)

Short-comings:

No exploration: The simulation is "on-policy" (with the learned model).

Inefficient: Requires to backpropagate through the SDE solver.

Unifying objective

Sampling problems: Consider the reverse-KL $D_{KL}(P_{X^u} \parallel P_{Y^v}) = E^h \log \frac{dP_{X^u}}{dP_{Y^v}}(X^u) : i$

Bridge sampler:

$$D_{KL}(P_{X^u} \parallel P_{Y^v}) = \int_0^1 \underbrace{\frac{1}{2}ku + vk^2 + \text{div}(v)}_{\text{running cost}}(X_s^u; s) ds + \underbrace{\log \frac{p_{\text{prior}}(X_0^u)}{q_{\text{target}}(X_T^u)}}_{\text{initial / terminal cost}} + \underbrace{\log Z}_{\text{const.}}$$

We can also derive unique objectives:

- Fixed v : Time-Reversed Diffusion Sampler (DIS).
- Reference process: Path Integral Sampler (PIS) & Denoising Diffusion Sampler (DDS)

Short-comings:

No exploration: The simulation is "on-policy" (with the learned model).

Inefficient: Requires to backpropagate through the SDE solver.

Mode collapse: Reverse-KL divergence suffers from mode collapse.

- Our framework allows for other divergences: We propose the log-variance divergence

$$D_{LV}^w(P_{X^u}; P_{Y^v}) = V \log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) :$$

Log-variance divergence

- Our framework allows for other divergences: We propose the log-variance divergence

$$D_{LV}^w(P_{X^u}; P_{Y^v}) = V \log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) :$$

o -policy: choose w to balance exploration and exploitation.

Log-variance divergence

- Our framework allows for other divergences: We propose the log-variance divergence

$$D_{LV}^w(P_{X^u}; P_{Y^v}) = \mathbb{V} \log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) :$$

o-policy: choose w to balance exploration and exploitation.

more efficient: no differentiation through SDE solver.

Log-variance divergence

- Our framework allows for other divergences: We propose the log-variance divergence

$$D_{LV}^w(P_{X^u}; P_{Y^v}) = V \log \frac{dP_{X^u}}{dP_{Y^v}}(X^w) :$$

o-policy: choose w to balance exploration and exploitation.

more efficient: no differentiation through SDE solver.

sticking-the-landing: zero variance at the optimum.

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.

Metrics: Log-norm. const., Sinkhorn distance to ground truth, normalized effective sample size, and average stddev.

Method	Divergence	$\log Z$ (rw) #	W^2 #	ESS #	std #
CRAFT		0.012	<u>0.020</u>	-	0.019
PIS	KL	0.249	0.467	0.0051	1.937
	LV	<u>0.001</u>	<u>0.020</u>	0.9093	0.023
DIS	KL	0.015	0.064	0.0226	2.522
	LV	0.017	<u>0.020</u>	0.8660	<u>0.004</u>
DDS	KL	0.005	0.042	0.0737	2.161
	LV	<u>0.001</u>	<u>0.020</u>	0.8929	0.006
Bridge	KL	0.560	0.393	0.0180	0.698
	LV	0.100	<u>0.020</u>	<u>0.9669</u>	0.010

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

Metrics: Log-norm. const., Sinkhorn distance to ground truth, normalized effective sample size, and average stddev.

Method	Divergence	$\log Z(\text{rw}) \#$	$W^2 \#$	ESS ⁿ	std #
CRAFT		0.012	<u>0.020</u>	-	0.019
PIS	KL	0.249	0.467	0.0051	1.937
	LV	<u>0.001</u>	<u>0.020</u>	0.9093	0.023
DIS	KL	0.015	0.064	0.0226	2.522
	LV	0.017	<u>0.020</u>	0.8660	<u>0.004</u>
DDS	KL	0.005	0.042	0.0737	2.161
	LV	<u>0.001</u>	<u>0.020</u>	0.8929	0.006
Bridge	KL	0.560	0.393	0.0180	0.698
	LV	0.100	<u>0.020</u>	<u>0.9669</u>	0.010

Image density and funnel distribution

Funnel distribution (challenging benchmark for sampling methods):

$$p_{\text{target}}(\mathbf{x}) = \mathcal{N}(\mathbf{x}_1; \mathbf{0}; \mathbf{9}) \prod_{i=2}^D \mathcal{N}(\mathbf{x}_i; \mathbf{0}; e^{x_1})$$

Funnel distribution (challenging benchmark for sampling methods):

$$p_{\text{target}}(x) = \prod_{i=2}^{\gamma_0} \mathcal{N}(x_i; 0; e^{x_1})$$

Method	Divergence	logZ (rw) #	W ² #	ESS ⁿ	std #
CRAFT		0.123	5.517	-	6.139
PIS	KL	0.111	5.639	0.1333	6.921
	LV	0.097	5.593	0.0746	6.852
DIS	KL	0.032	5.120	0.1383	5.254
	LV	<u>0.028</u>	<u>5.075</u>	<u>0.2313</u>	<u>5.224</u>
DDS	KL	0.045	5.305	0.1446	6.133
	LV	0.043	5.305	0.1999	6.123

Funnel distribution (challenging benchmark for sampling methods):

$$p_{\text{target}}(x) = \prod_{i=1}^2 \mathcal{N}(x_i; 0; 9) \prod_{i=2}^{\gamma_0} \mathcal{N}(x_i; 0; e^{x_1})$$

Method	Divergence	logZ (rw) #	W ² #	ESS ⁿ	std #
CRAFT		0.123	5.517	-	6.139
PIS	KL	0.111	5.639	0.1333	6.921
	LV	0.097	5.593	0.0746	6.852
DIS	KL	0.032	5.120	0.1383	5.254
	LV	<u>0.028</u>	<u>5.075</u>	<u>0.2313</u>	<u>5.224</u>
DDS	KL	0.045	5.305	0.1446	6.133
	LV	0.043	5.305	0.1999	6.123

Our improved samplers based on diffusion processes are **competitive with state-of-the-art methods** based on SMC & normalizing flows (CRAFT).

Many-Well problem

Many-Well: typical problem in molecular dynamics with

$$\log p(x) = -\sum_{i=1}^w (x_i^2)^2 - \frac{1}{2} \sum_{i=w+1}^d x_i^2$$

Many-Well problem

Many-Well: typical problem in molecular dynamics with

$$\log p(x) = -\sum_{i=1}^w (x_i^2)^2 - \frac{1}{2} \sum_{i=w+1}^d x_i^2$$

Many-Well problem

Many-Well: typical problem in molecular dynamics with

$$\log (x) = \prod_{i=1}^w (x_i^2)^2 \frac{1}{2} \prod_{i=w+1}^d x_i^2 :$$

Problem	Method	$\log Z$ #	W^2 #	ESS "	std #
Many-Well ($d = 5; w = 5; = 4$)	PIS-KL	3.567	1.699	0.0004	1.409
	PIS-LV	<u>0.214</u>	<u>0.121</u>	<u>0.6744</u>	<u>0.001</u>
	DIS-KL	1.462	1.175	0.0012	0.431
	DIS-LV	<u>0.375</u>	<u>0.120</u>	<u>0.4519</u>	<u>0.001</u>
Many-Well ($d = 50; w = 5; = 2$)	PIS-KL	0.101	<u>6.821</u>	0.8172	0.001
	PIS-LV	<u>0.087</u>	6.823	<u>0.8453</u>	<u>0.000</u>
	DIS-KL	1.785	<u>6.854</u>	0.0225	0.009
	DIS-LV	<u>1.783</u>	6.855	<u>0.0227</u>	0.009



Thank you for your attention!

Website: `https://jberner.info`

Github: `https://github.com/juliusberner`

Mail: `mail@jberner.info`