Learned Reference-based Diffusion Sampler for multi-modal distributions

Maxence Noble*, Louis Grenioux*, Marylou Gabrié, Alain Durmus

CMAP, École Polytechnique, France

Sampling from multi-modal distributions

Setting. We want **samples** from a target probability distribution π supported on $\mathbb{R}^d$, while only having access to its log-density $\log \pi$.

The difficulty. This task is especially difficult when π is *multi-modal* as the samples needs to reflect both the local properties (mean, covariance, ...) and the global properties (proportions between distinct non-zero probability areas).

Our (realistic) assumption. In this work, we simplify the problem by assuming access the the locations of the modes.

Standard sampling approaches

Classic MCMC. *Markov Chain Monte Carlo (MCMC)* builds chains using local transitions that preserve π, but such moves often prevent the chain from escaping modes, limiting global exploration.

Annealed MCMC. *Annealing methods* run MCMC over a sequence of distributions that bridge an easy-to-sample distribution ρ to the target π, typically via a *geometric interpolation*. While effective in low dimensions, these schemes struggle with high dimension.

Diffusion Models : from generative to sampling approach

Diffusion Models (DMs) are a powerful class of generative models based on a *noising diffusion process* $(X_t)_{t\in[0,T]}$, defined by

$$\mathrm{d}X_t = f(t)X_t\mathrm{d}t + g(t)\mathrm{d}B_t,\ X_0 \sim \pi$$

which gradually transforms the data into noise $X_T \sim \rho$, with $\mathbb{P}^\star$ denoting the associated path measure and p_t the density of $\mathbb{P}^\star{}_t$. To generate samples, DMs simulate the *time-reversed* process $(Y_t)_{t\in[0,T]} \sim (\mathbb{P}^\star)^R$

$$\mathrm{d}Y_t = -\left[f(T-t)Y_t - g^2(T-t)\nabla \log p_{T-t}(Y_t)\right]\mathrm{d}t + g(T-t)\mathrm{d}W_t$$

with $Y_0 \sim \rho$, which satisfies $Y_t \overset{\mathcal{L}}{=} X_{T-t}$, i.e., $Y_T \sim \pi$. Still, the score function $\nabla \log p_t$ is intractable and needs to be approximated. In practice, one aims to learn it with a *neural network* s_t^θ. We denote by $\mathbb{P}^\theta$ the path measure of the obtained *denoising diffusion process*.

Generative setting. Using samples from π, s_t^θ can be efficiently optimized via a *Denoising Score Matching* regression loss.

Sampling setting. This work tackles learning s_t^θ from model samples via a variational approach, without requiring access to π samples.

Reference-based variational approach for diffusion sampling

We aim to minimize the log-variance variational loss on path measures

$$\mathcal{L}(\theta) = \mathrm{Var}\left[\log\left(\mathrm{d}\mathbb{P}^\theta/\mathrm{d}(\mathbb{P}^\star)^R\right)\left(Y^{\hat\theta}_{[0,T]}\right)\right],\ Y^{\hat\theta}_{[0,T]} \sim \mathbb{P}^{\hat\theta}$$

where $\hat\theta = \mathrm{StopGrad}(\theta)$. As such, this loss is however not tractable...

Following [1], we introduce a *reference process* $\mathbb{P}^{\mathrm{ref}}$, defined as the exact noising process for a known distribution π^{ref}, which yields to simplifying the log-density ratio

$$\log \frac{\mathrm{d}\mathbb{P}^\theta}{\mathrm{d}(\mathbb{P}^\star)^R}\left(Y_{[0,T]}\right) = \log \frac{\mathrm{d}\mathbb{P}^\theta}{\mathrm{d}(\mathbb{P}^{\mathrm{ref}})^R}\left(Y_{[0,T]}\right) + \log \frac{\pi^{\mathrm{ref}}}{\pi}(Y_T)\ .$$

By further parameterizing s_t^θ as an additive *control* term

$$s_t^\theta = \nabla \log p_t^{\mathrm{ref}} + g(t)^{-1}\phi_t^\theta$$

we obtain the tractable (but not simulation free) objective

$$\mathcal{L}(\theta) = \mathrm{Var}\left[\int_0^T \left\{\tfrac{1}{2}\|\phi^\theta_{T-t}(Y_t^{\hat\theta})\|^2\mathrm{d}t + \phi^\theta_{T-t}(Y_t^{\hat\theta})^\top \mathrm{d}B_t\right\} + \log \tfrac{\pi^{\mathrm{ref}}}{\pi}(Y_T^{\hat\theta})\right]$$

Our novelty : learning the reference process from MCMC data

Previous designs of $\mathbb{P}^{\mathrm{ref}}$. Prior works [3, 2] use a Gaussian π^{ref} with fixed variance, which requires heavy tuning and strong constraints on the neural network ϕ_t^θ.

Our intuition on the role of the reference process

If well chosen, the reference process may drive Y_t^θ to high-density regions, thus simplifying the numerical optimization procedure.
$\Longrightarrow$ We propose to learn $\mathbb{P}^{\mathrm{ref}}$ based on approximate data from π !

LRDS methodology. We combine the three following steps:

1. Generate approximate samples from π via MCMC;
2. Fit a Diffusion Model on this data to obtain $\mathbb{P}^{\mathrm{ref}}$;
3. Minimize the log-variance loss $\mathcal{L}$ w.r.t. θ to obtain $\mathbb{P}^\theta$.

(1) MCMC — ✓ Local, × Global → (2) Learn $\mathbb{P}^{\mathrm{ref}}$ from MCMC data → (3) Learn $\mathbb{P}^\theta$ using $\mathbb{P}^{\mathrm{ref}}$ with $\mathcal{L}$ → (4) Run $\mathbb{P}^\theta$ — ✓ Local, ✓ Global

How to learn the scores and the densities of the reference process ?

The loss $\mathcal{L}$ requires both the scores $\nabla \log p_t^{\mathrm{ref}}$ and the density π^{ref}. We propose to learn them based on MCMC data in two ways.

- **Gaussian Mixture Model (GMM) – cheap but limited**
  Fit a GMM on MCMC data to model π^{ref} (see below); scores are computed analytically at no extra cost;

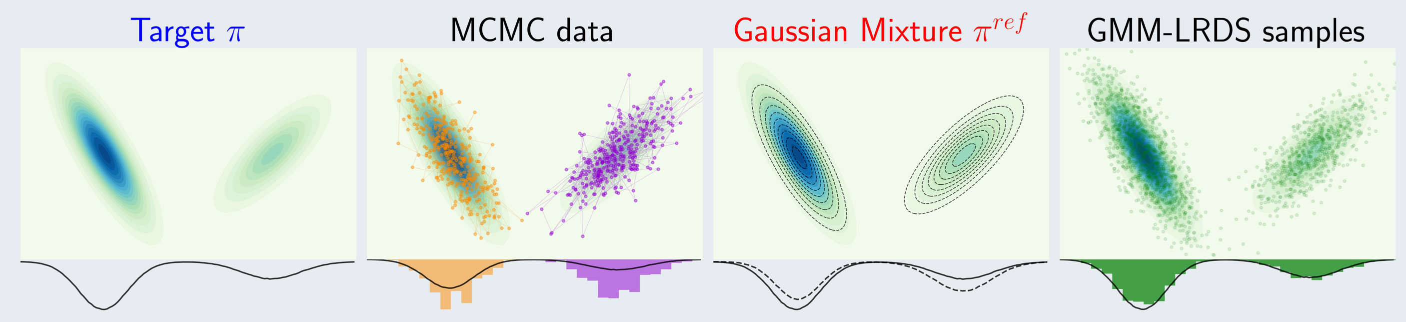

Figure 1. Learned π^{ref} and samples from GMM-LRDS on a Gaussian mixture.

- **Energy-Based Model (EBM) – costly but expressive**
  Learn p_t^{ref} via a new Maximum Likelihood approach using a neural net. Recover $\pi^{\mathrm{ref}} = p_0^{\mathrm{ref}}$ and get scores by differentiating the net.

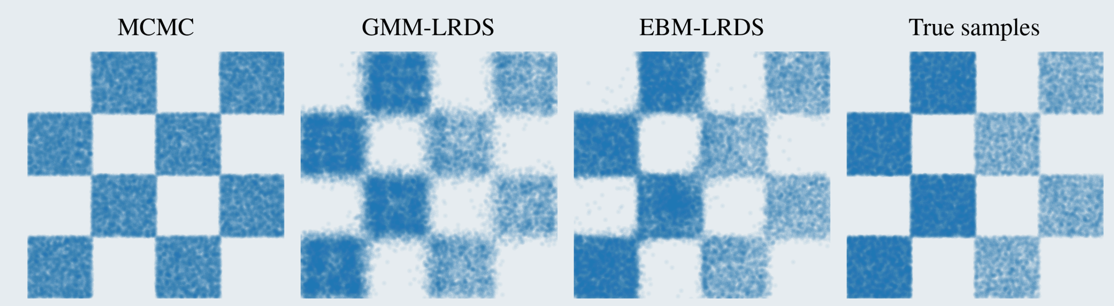

Figure 2. Samples from GMM-LRDS and EBM-LRDS on the checkerboard

Numerical results

We apply our methodology on a wide range of known multi-modal distributions and observe a significant performance gap against concurrent methods (annealed MCMC & diffusion-based sampling methods).

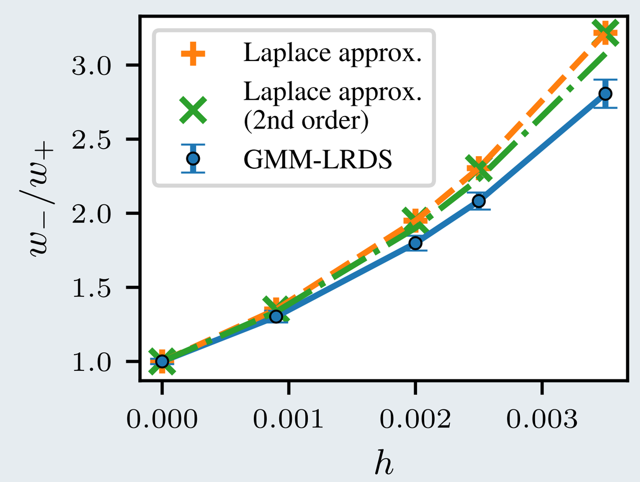

Figure 3. Estimation of the mode weight ratio in the challenging ϕ^4 system using GMM-LRDS. The external field h breaks the system's symmetry, creating an imbalance between modes. Laplace approximation serves as the reference ground truth.

Limitations. Learning both $\mathbb{P}^{\mathrm{ref}}$ and $\mathbb{P}^\theta$ may be expensive in practice. Moreover, similarly to previous approaches, LRDS doesn't scale well with high dimension or with high number of modes.

References

[1] L. Richter and J. Berner. Improved sampling via learned diffusions. In *ICLR*, 2024.

[2] F. Vargas, W. S. Grathwohl, and A. Doucet. Denoising diffusion samplers. In *ICLR*, 2023.

[3] Q. Zhang and Y. Chen. Path Integral Sampler. In *ICLR*, 2022.