Training Neural Samplers with Reverse Diffusive KL Divergence

Jiajun He[1,*], Wenlin Chen[1,3,*], Mingtian Zhang[2,*], David Barber[2], José Miguel Hernández-Lobato[1]

[1]University of Cambridge, [2]University College London, [3]Max Planck Institute for Intelligent Systems

***Equal contribution**

AISTATS 2025

https://arxiv.org/abs/2410.12456

UNIVERSITY OF CAMBRIDGE

UCL

MAX PLANCK INSTITUTE FOR INTELLIGENT SYSTEMS

Background

Sampling from unnormalized density

$$p_d(x) = \exp(-E(x))/Z$$

- Bayesian posterior
- Boltzmann distributions
- Easy to evaluate, hard to sample

Simiulation with a generative model

$$p_\theta(x) \approx p_d(x)$$

Reverse KL Minimization

$$\mathrm{KL}(p_\theta || p_d) = \int (\log p_\theta(x) - \log p_d(x)) p_\theta(x) dx$$
$$= \int (\log p_\theta(x) + E(x)) p_\theta(x) dx + \log Z,$$

- Not defined for implicit model

$$p_\theta(x) = \int \delta(x - g_\theta(z)) p(z) dz.$$

- Mode collapse problem

Our Method

Diffusive KL Divergence (DiKL)

$$\mathrm{DiKL}_{\mathcal{K}}(p||q) \equiv \sum_{t=1}^{T} w(t) \mathrm{KL}(p * k_t || q * k_t),$$

$$k_t(x_t|x) = \mathcal{N}(x_t|\alpha_t x, \sigma_t I) \quad \mathrm{DiKL}[p||q] = 0 \Leftrightarrow p = q$$

DiKL Encourages Mode Coverage

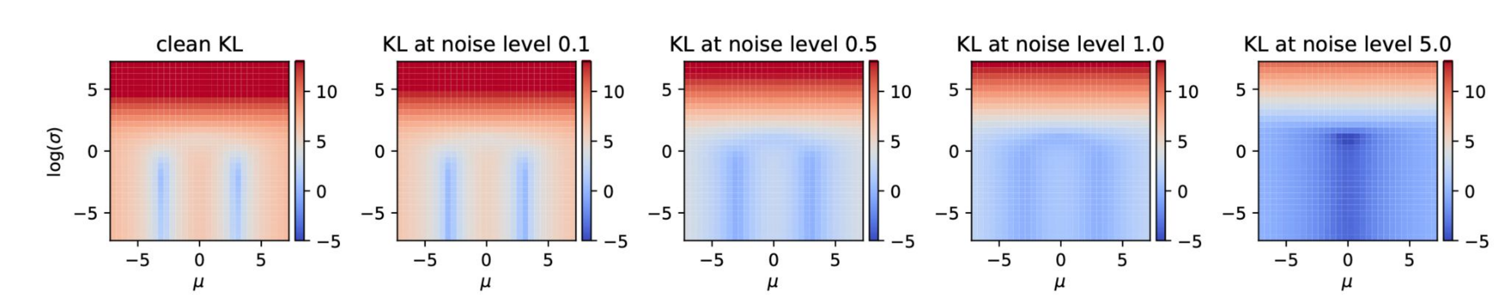

DiKL Gradient Estimator

$$\nabla_\theta \mathrm{DiKL}_{k_t}(p_\theta || p_d) = \nabla_\theta \mathrm{KL}(p_\theta * k_t || p_d * k_t)$$
$$= \int p_\theta(x_t) \left(\nabla_{x_t} \log p_\theta(x_t) - \nabla_{x_t} \log p_d(x_t) \right) \frac{\partial x_t}{\partial \theta} dx_t,$$

- Estimating the noisy model score with denoising score matching (DSM)

$$\min_\phi \iint \| s_\phi(x_t) - \nabla_{x_t} \log k(x_t|x) \|_2^2 k(x_t|x) p_\theta(x) dx dx_t$$

- Estimating the noisy target score with mixed score identity (MSI)

$$\nabla_{x_t} \log p_d(x_t) = \int (\alpha_t (x + \nabla_x \log p_d(x)) - x_t) p_d(x|x_t) dx.$$

Experiments

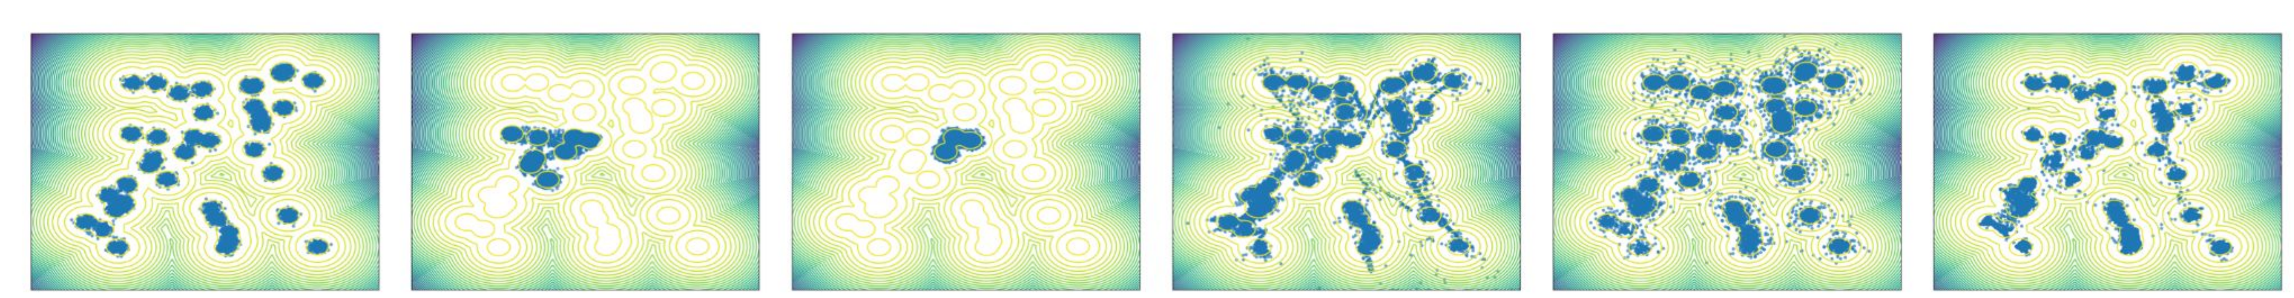

(a) Ground Truth (b) R-KL SM (c) R-KL Bound (d) FAB (e) iDEM (f) DiKL (ours)

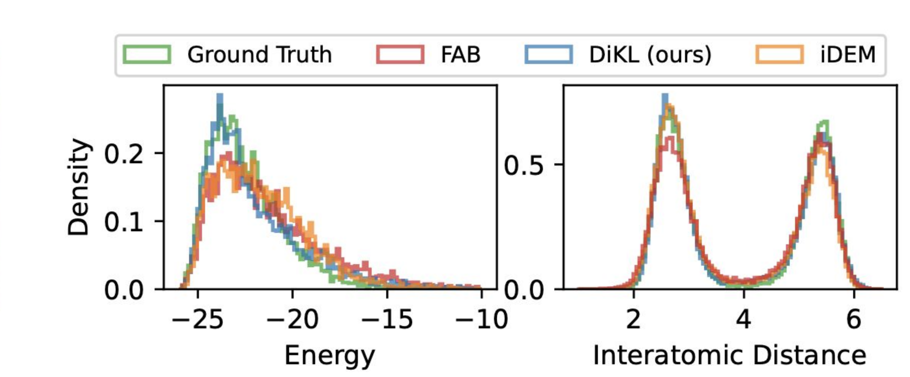

(a) DW-4

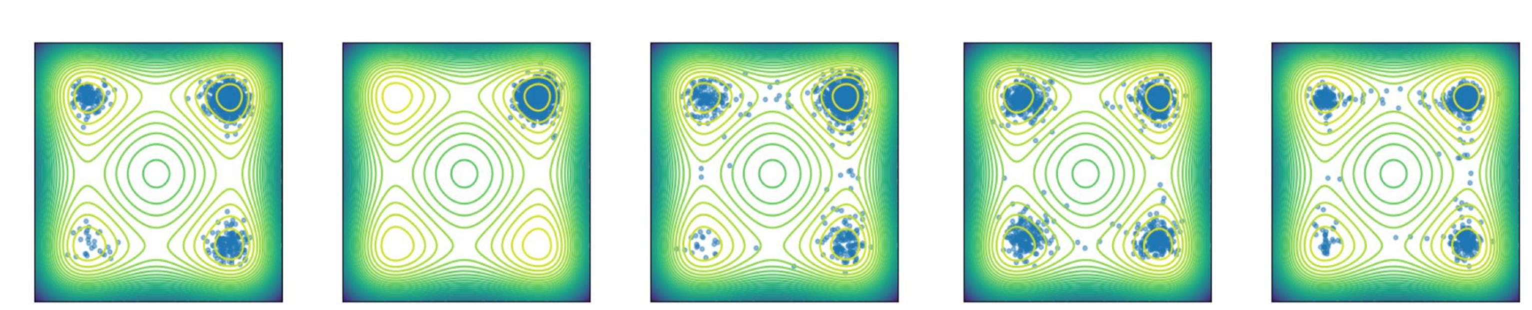

(a) Ground Truth (b) KL (c) FAB (d) iDEM (e) DiKL (ours)

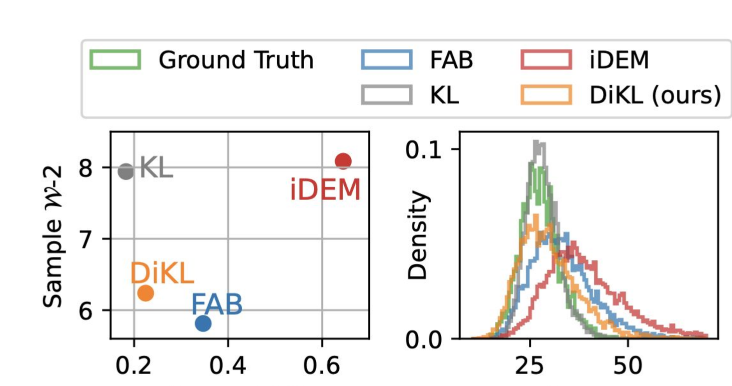

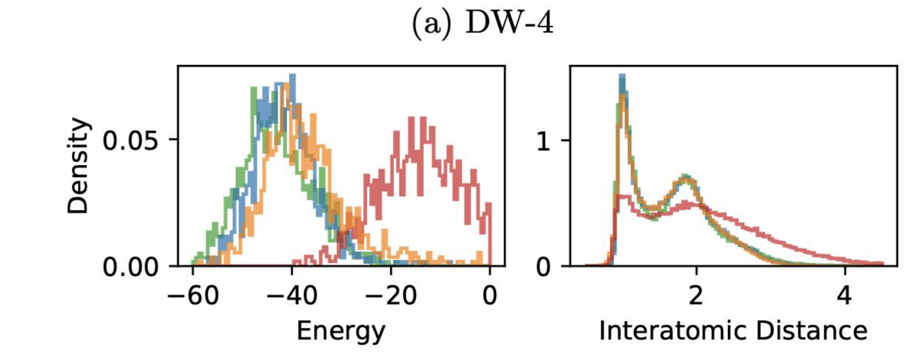

(b) LJ-13