EXACT GLOBAL MCMC WITH DENOISING DIFFUSION

Anonymous authors
Paper under double-blind review

ABSTRACT

This work shows that diffusion models learned with the standard denoising loss provide exact and effective global MCMC proposals for complex high-dimensional target densities. A forward noising process followed by its exact reversal is a Markov chain that leaves the target invariant, and a denoising diffusion model is a learned approximation of that reversal. A Metropolis–Hastings (MH) step whose acceptance ratio includes the forward and reverse path densities of a discrete-time SDE makes the chain exact for any denoiser, with an acceptance rate that measures how closely the denoiser approximates the ideal reversal. We train denoising diffusion models on locally convergent samples from the Metropolis-adjusted Langevin algorithm (MALA), whose weighting across modes need not match the target since the MH step absorbs the mismatch, and call the composition of the resulting global path sampler with a local MALA sampler Denoising Diffusion Monte Carlo (DDMC). Across Gaussian mixtures, particle systems, and a 550-dimensional Bayesian neural network posterior, DDMC matches ground-truth energy distributions at the measured estimator floor and recovers correct mode weights. On the LJ-38 double-funnel potential below its melting temperature, DDMC obtains the equilibrium funnel occupancy from a mis-weighted training dataset, agrees with parallel tempering, and attains an order of magnitude lower variance at equal total compute, dataset generation and denoiser training included.

1 INTRODUCTION

An ideal diffusion model is already a Markov chain Monte Carlo (MCMC) sampler. If X is a draw from a target density q, noising it to $Y = X + \sigma Z$ and denoising Y with an ideal reverse process gives $\hat{X} \sim q$, so the noise-to-denoise map is a Markov transition that leaves the target invariant. This ideal kernel is the proximal sampler (Lee et al., 2021; Chen et al., 2022), whose second step samples the intractable denoising posterior. At large σ it moves freely between regions of the state space that no local sampler can connect. Such a model would be a global sampler for whatever density it represents. We build upon this observation and propose a practical method for sampling from unnormalized multimodal densities, a core problem in Bayesian inference and the physical sciences.

Denoiser training is an ordinary regression whose optimum is the exact denoising posterior mean and which provides an approximate model of the reverse diffusion (Vincent, 2011; Ho et al., 2020; Karras et al., 2022). A learned diffusion cannot be ideal, so a practical sampler must account for model and discretization error in the reverse path. A Metropolis–Hastings (MH) step on the joint density of a discrete diffusion path does this and makes the chain exact for any reverse model. An inaccurate model costs acceptance, never correctness. Since the ideal reversal is always accepted, acceptance measures denoiser quality, and better denoising training means more accepted global jumps.

Training data is an obstacle: exact samples are unavailable, since producing them is the problem being solved, but locally converged samples are cheap. Metropolis-adjusted Langevin algorithm (MALA) chains equilibrate within basins, and from a suitable initialization they can cover the basins even if the resulting dataset weights them incorrectly. The path MH step that absorbs reverse-model error also absorbs this dataset bias. At sampling time we compose the path move with MALA steps for local refinement. The two are natural partners: the diffusion path proposes global moves between basins that MALA cannot make, and the MALA steps return the chain to the equilibrated configurations the denoiser was trained on. We call the composed sampler *Denoising Diffusion Monte Carlo* (DDMC).

A normalizing flow supplies a tractable density which has been widely used for exact sampling by importance reweighting or by MH proposal (Noé et al., 2019; Köhler et al., 2020; Klein et al., 2023; Gabrié et al., 2022). A diffusion model has no closed-form marginal density, so diffusion-based

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

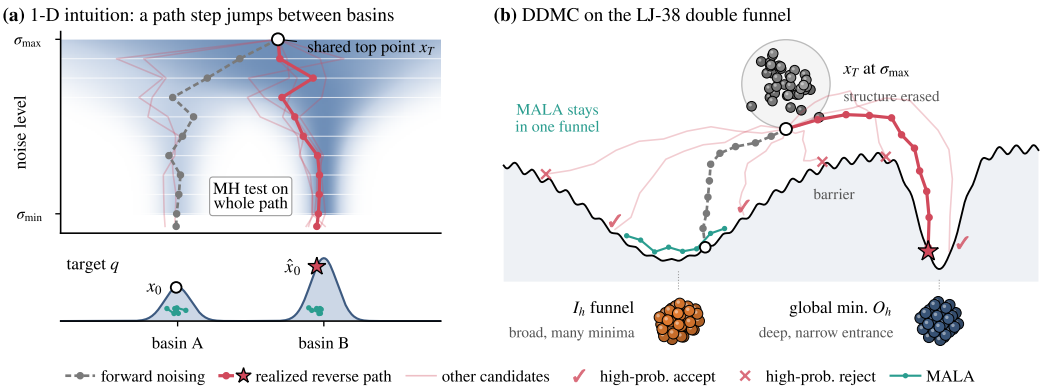


Figure 1: **(a)** One DDMC step on a two-basin 1-D target, with shading q_σ per level. The state x_0 is noised to x_T, and the realized reverse path from the same x_T lands in the other basin (star). An MH test on the whole path keeps the move exact, and MALA refines locally. Paths use the exact denoiser of this mixture. **(b)** On LJ-38 below melting, MALA is confined to the broad I_h funnel, while path proposals pass through a noised configuration and can enter the narrow O_h funnel. Acceptance is determined by the quality of the reverse path and the denoised sample it proposes.

samplers have instead been trained variationally from the unnormalized density (Zhang & Chen, 2022; Vargas et al., 2023; 2024; Guo et al., 2026). MH corrections of such variationally trained paths were proposed concurrently and independently by MAD-Path (Chen et al., 2026a) and SPS (Chen et al., 2026b). Variational objectives carry structural limits. The training signal comes from the model's own trajectories, and the model must simultaneously solve the problems of target exploration and sample generation. Many variational objectives differentiate through the sampling trajectory, in which case memory grows with the number of noise levels. Denoising regression has neither limit: local MCMC training samples supply a stable external source of information about the target, and the denoiser is fit by ordinary regression on clean and noisy pairs. Exactness comes from the path density, so nothing constrains the architecture. Data quantity is not a constraint either since the unnormalized density is a data generator that local MCMC equilibration turns into training samples in parallel.

Our experiments span Gaussian mixtures, a 550-dimensional Bayesian neural network posterior, and particle potentials. DDMC matches the energy distribution of ground-truth samples at the measured estimator floor wherever a trustworthy reference exists and recovers correct mode weights from mis-weighted training datasets. Our main result is on the 38-particle Lennard-Jones (LJ) cluster, a double-funnel landscape whose narrow entrance to the global-minimum funnel greatly reduces the efficiency of parallel tempering (PT) with local MALA samplers. There DDMC agrees with PT on the funnel occupancy at two temperatures below melting and attains roughly an order of magnitude lower variance at equal total compute, dataset generation and training included. Our main contributions are:

- **DDMC**, an exact MCMC sampler whose global proposals come from a standard denoising diffusion model trained on locally converged but possibly mis-weighted samples, composed with local MALA steps, and requiring no change to the diffusion modeling or learning formulation.
- **A practical recipe** for training data generation, moment calibration of the reverse kernel variances, an acceptance diagnostic that tracks denoiser viability for MCMC during training, and an i-SIR extension that amplifies acceptance by exploiting GPU throughput.
- **Empirical validation** on targets up to $d = 550$, including exactness at measured ground-truth floors, repair of mis-weighted datasets, and funnel occupancies on LJ-38 below melting, a long-standing hard case across MCMC methods (Doye et al., 1999; Neirotti et al., 2000; Mandelshtam & Frantsuzov, 2006; Pártay et al., 2010), in agreement with PT and at an order of magnitude lower variance at equal compute.

2 Background

Metropolis–Hastings and local samplers. We wish to draw samples from a density q on $\mathbb{R}^d$ known only up to a constant: we can evaluate $\tilde{q}$ with $q = \tilde{q}/\mathcal{Z}$ for an intractable $\mathcal{Z}$, and we have the score $\nabla \log \tilde{q} = \nabla \log q$, which covers Bayesian posteriors and Boltzmann distributions $\tilde{q} = e^{-U}$. Given

any proposal kernel $R(x' \mid x)$, the MH correction accepts $x \to x'$ with probability

$$\alpha(x, x') = \min\left\{1,\ \frac{\tilde{q}(x')\,R(x \mid x')}{\tilde{q}(x)\,R(x' \mid x)}\right\}, \tag{1}$$

and otherwise leaves the chain at x (Metropolis et al., 1953; Hastings, 1970). The ratio needs only $\tilde{q}$, and q is stationary for any R. The proposal determines only the mixing speed. MALA (Roberts & Tweedie, 1996) takes one Euler–Maruyama step of the overdamped Langevin diffusion, $R_h(x' \mid x) = \mathcal{N}(x'; x + \frac{h^2}{2}\nabla \log q(x), h^2 I_d)$, followed by (1). Well-tuned h sits near the smallest local standard deviation of the target, so MALA mixes within a mode but crosses barriers only through many small steps, with crossing times exponential in the barrier height. Hamiltonian Monte Carlo (Duane et al., 1987; Neal, 2011) and NUTS (Hoffman & Gelman, 2014) share this limitation.

Denoising diffusion. We use the EDM formulation of Karras et al. (2022). Diffusion models act on the smoothed densities $q_\sigma = q * \mathcal{N}(0, \sigma^2 I_d)$ for $\sigma \in [\sigma_{\min}, \sigma_{\max}]$, where $q_{\sigma_{\min}} \approx q$ and $q_{\sigma_{\max}}$ is close to $\mathcal{N}(0, \sigma_{\max}^2 I_d)$. A denoiser $D_\theta(y, \sigma)$ is fit by regression on samples $X \sim q$,

$$\mathcal{L}(\theta) = \mathbb{E}_\sigma\left[\lambda(\sigma)\,\mathbb{E}_{X,\,Z}\,\|D_\theta(X + \sigma Z, \sigma) - X\|^2\right], \tag{2}$$

whose minimizer is the posterior mean $\mathbb{E}[X \mid X + \sigma Z = y]$, related to the score by Tweedie's identity $\nabla_y \log q_\sigma(y) = (D^*(y, \sigma) - y)/\sigma^2$ (Efron, 2011). Training is simulation-free, each update noising one sample at one level, and is continuous in σ, so any ladder may be used at sampling time.

Path densities. Samples are generated by a discrete-time SDE in the style of DDPM ancestral sampling (Ho et al., 2020; Song et al., 2021), whose path densities are available in closed form. Let x_0 be a clean state and let $x_1, \ldots, x_T$ sit at levels $\sigma_{\min} = \sigma_1 < \cdots < \sigma_T = \sigma_{\max}$, and write $\Delta_k^2 = \sigma_k^2 - \sigma_{k-1}^2$ with $\Delta_1^2 = \sigma_1^2$. The forward process adds independent Gaussian increments,

$$p(x_k \mid x_{k-1}) = \mathcal{N}\big(x_k;\, x_{k-1},\, \Delta_k^2 I_d\big), \qquad p(x_{0:T}) = q(x_0)\prod_{k=1}^{T} p(x_k \mid x_{k-1}), \tag{3}$$

so that $x_k \mid x_0 \sim \mathcal{N}(x_0, \sigma_k^2 I_d)$ and $x_k \sim q_{\sigma_k}$. Generation starts at $x_T \sim \mathcal{N}(0, \sigma_{\max}^2 I_d)$ and runs the ladder in reverse with one denoiser evaluation per level,

$$p_\theta(x_{k-1} \mid x_k) = \mathcal{N}\big(x_{k-1};\, \gamma_k x_k + (1 - \gamma_k) D_\theta(x_k, \sigma_k),\, \tau_k^2 I_d\big), \qquad \gamma_k = \sigma_{k-1}^2/\sigma_k^2, \quad \gamma_1 = 0, \tag{4}$$

giving the backward path density $p_\theta(x_{0:T}) = \mathcal{N}(x_T; 0, \sigma_{\max}^2 I_d)\prod_k p_\theta(x_{k-1} \mid x_k)$. The mean $\mu_\theta(x_k, \sigma_k) = \gamma_k x_k + (1 - \gamma_k) D_\theta(x_k, \sigma_k)$ in (4) is the exact posterior mean of x_{k-1} given x_k with the unknown clean sample replaced by D_θ. The variance τ_k^2 is a modeling choice and any positive value gives a valid sampler. Both path densities are products of Gaussian kernels with closed-form means and variances, so any trajectory's density is evaluated exactly in either direction with T denoiser evaluations and no marginal q_{σ_k}.

3 Method

An ideal diffusion model is already an MCMC sampler for the density it models. Let q be a target density on $\mathbb{R}^d$ and $q_\sigma = q * \mathcal{N}(0, \sigma^2 I_d)$. A diffusion model defines a conditional distribution $\rho(x \mid y, \sigma)$ over endpoints of the reverse path started from y at noise level σ, and is *ideal* if it transports q_σ to q exactly. For $X \sim q$ and $Y = X + \sigma Z$, an ideal model returns $\hat{X} \sim q$, so $X \to \hat{X}$ leaves q invariant for every $\sigma > 0$. It is ergodic (Appendix A), and large σ makes it global. This transition is the proximal sampler of Lee et al. (2021), a Gibbs sampler on the pair $(X, X + \sigma Z)$ whose second step draws from the exact denoising posterior. Chen et al. (2022) read that step as the time reversal of the noising process. That reversal is exactly what a diffusion model approximates: the denoiser supplies the score of every q_σ through Tweedie's identity, the reverse SDE driven by that score is ρ, and the denoising objective (2) is the standard way to learn it from samples. We show that the standard formulation, unchanged, realizes the kernel as an exact and effective global proposal.

In practice the model is not ideal, and an exact sampler must account for the imperfection. We define the sampler as a two-block Gibbs sweep on the augmented target (3), whose x_0 marginal is q, with blocks $V = (X_0, \ldots, X_{T-1})$ and the top point X_T. Given the current state X_0, drawing the forward increments $X_1, \ldots, X_{T-1}$ completes V. The update $X_T \mid V$ is the Gaussian refresh $p(\cdot \mid X_{T-1})$,

and the update $V \mid X_T$ is Metropolis-within-Gibbs and moves the chain. By Bayes' rule and (3) the conditional is $p(v \mid x_T) = q(x_0) \prod_k p(x_k \mid x_{k-1}) / q_{\sigma_T}(x_T)$, whose denominator is the intractable marginal that cancels because X_T is held fixed, as does the normalizing constant of q. The proposal runs the learned reverse path from that same top point,

$$R(\hat{v} \mid v, x_T) \;=\; p_\theta(\hat{v} \mid x_T) \;=\; \prod_{k=1}^{T} p_\theta(\hat{x}_{k-1} \mid \hat{x}_k), \qquad \hat{x}_T = x_T, \tag{5}$$

and does not depend on the current path. Substituting into (1) gives

$$\alpha(v, \hat{v}) = \min\left\{1,\; \frac{\tilde{q}(\hat{x}_0) \prod_{k=1}^{T} p(\hat{x}_k \mid \hat{x}_{k-1}) \prod_{k=1}^{T} p_\theta(x_{k-1} \mid x_k)}{\tilde{q}(x_0) \prod_{k=1}^{T} p(x_k \mid x_{k-1}) \prod_{k=1}^{T} p_\theta(\hat{x}_{k-1} \mid \hat{x}_k)}\right\}, \qquad \hat{x}_T = x_T. \tag{6}$$

Both block updates leave p invariant, so the sweep does, and since the x_0 marginal of p is q, the induced transition $X_0 \to \hat{X}_0$ leaves q invariant. The chain is ergodic since the Gaussian kernels have full support and the acceptance probability is positive wherever $\tilde{q} > 0$. The construction is exact for any p_θ.

For an ideal reverse process the proposal equals the true conditional $p(\hat{v} \mid x_T)$, every term in (6) cancels, and acceptance is one (Appendix A). For a learned reverse process, the ratio accounts for the combined model and discretization gap and acceptance is a direct measure of denoiser quality: the ratio approaches one as the trained model approaches its optimum and the ladder is refined.

A correction of this form was proposed concurrently and independently by MAD-Path (Chen et al., 2026a) and SPS (Chen et al., 2026b), both of which train the score variationally from the density rather than by denoising regression on samples. Appendix F.2 compares the two approaches on LJ-55.

3.1 Denoising Diffusion Monte Carlo

We propose *Denoising Diffusion Monte Carlo* (DDMC), whose global moves are the path-space MH step (6) driven by a denoiser trained with the standard objective (2) on locally converged MALA samples, and whose local moves are MALA steps:

$$\mathcal{K}_{\mathrm{DDMC}}(x' \mid x) \;=\; \int_{\mathbb{R}^d} \mathcal{K}_{\mathrm{path}}(u \mid x)\, \mathcal{K}_h^N(x' \mid u)\, \mathrm{d}u. \tag{7}$$

Here $\mathcal{K}_{\mathrm{path}}$ is the kernel of the bottom path endpoint under one path move, and $\mathcal{K}_h^N$ is the N-fold composition of the MALA kernel of Section 2. On rejection each places its mass at the current state. One application of (7), a path step followed by N MALA steps, is one DDMC step. DDMC chain lengths throughout the paper are counted in DDMC steps. Both kernels leave q invariant, so the composition does, and alternating a global proposal with local refinement in this way is an established pattern (Samsonov et al., 2022; Gabrié et al., 2022; Chen et al., 2024). The MALA steps return the chain toward the locally equilibrated states the denoiser was trained on, so repeated path rejections move it toward configurations the model handles well. In the limit of zero path acceptance DDMC degrades to plain MALA.

3.2 Creating a MALA Dataset

Dataset generation has two stages. The first is a target-dependent seeding stage that determines which regions the dataset will represent, starting from an initialization distribution and optionally including gradient ascent on $\log \tilde{q}$ and MALA legs with adaptive step sizes, and the second is a final MALA stage at a fixed calibrated step size that thermalizes each chain wherever seeding left it (Appendix C.1). Training a generative model on local MCMC or molecular dynamics samples is established practice in the Boltzmann generator literature (Noé et al., 2019; Köhler et al., 2020; Klein et al., 2023; Wirnsberger et al., 2022).

The dataset need not match the target's weighting across metastable regions. Since (6) is exact for any p_θ, a moderate mismatch does not preclude attaining useful acceptance. The acceptance cost grows with the mismatch, including the extreme case of region omission: a region carrying appreciable target probability that appears in the dataset only very rarely or never is generated by the denoiser with negligible probability, and a chain that nevertheless reaches it is slow to leave, since the model explains the current state poorly and most moves out are rejected. What the dataset must supply is coverage of the regions that carry mass, in reasonable proportion rather than exact proportion.

Retraining on the sampler's own output, as explored for normalizing flows (Gabrié et al., 2022), could improve the match between proportions and is left for future work.

For most targets seeding starts from a normal or uniform distribution spanning the high-probability region. This suffices when every metastable region captures enough of the initial draws, that is, when the initialization places appreciable mass in each region's basin of attraction under the seeding dynamics. LJ-38 below its melting temperature is an exception. Its global-minimum funnel captured only a tiny fraction of random starts under every seeding schedule we tried, roughly 0.018% for a direct quench, so our seeding instead draws from a database of local minima found by basin hopping (Wales & Doye, 1997; Doye et al., 1999). Appendix C.2 gives the details.

3.3 Path Tuning and Variance Calibration

Two objects specify the discrete path: the noise ladder $\{\sigma_k\}_{k=1}^{T}$ and the reverse variances $\{\tau_k^2\}_{k=1}^{T}$. We use the EDM spacing of Karras et al. (2022) throughout, so the ladder is determined by its endpoints $\sigma_{\min}$ and $\sigma_{\max}$ together with the number of levels T. Here $\sigma_{\max}$ must be large enough to bridge the distance between metastable regions, and $\sigma_{\min}$ small enough that the final denoising step lands within the thermal shell of the target. In most cases acceptance increases or remains stable with larger T, provided the reverse variances are calibrated as described below.

Exactness holds for any τ_k^2, so their calibration affects only acceptance. Given n_{cal} held-out states we simulate a forward path from each and set the variance at each level to the mean squared residual of the reverse mean,

$$\hat{\tau}_k^2 \;=\; \frac{1}{n_{\text{cal}}\, d} \sum_{i=1}^{n_{\text{cal}}} \left\| x_{k-1}^{(i)} - \mu_\theta\big(x_k^{(i)}, \sigma_k\big) \right\|^2, \tag{8}$$

which is the moment-matched and KL-optimal choice for a Gaussian kernel with mean μ_θ. Uncalibrated or poorly calibrated reverse variances are costly because per-level errors accumulate along the ladder. For the targets in this work (8) suffices given a large enough n_{cal}, and calibration is a one-time cached cost per configuration. With the spacing rule and the calibration procedure fixed, the DDMC path step has four hyperparameters: T, $\sigma_{\min}$, $\sigma_{\max}$, and n_{cal}. Appendices D and H provide tuning details and the diagnostics we monitor during training: the log MH ratio $\log r$ of (6) on held-out states, decomposed into a target term δ_q and a path term δ_{path}.

i-SIR parallel proposals. At low acceptance most reverse passes are spent on rejected paths. The proposal (5) depends only on the shared top point and not on the current path, so any number of proposals can be drawn from the same distribution. This is the setting of iterated sampling importance resampling (Andrieu et al., 2018; Samsonov et al., 2022), closely related to multiple-try Metropolis (Liu et al., 2000). Given the current path v and top point x_T, set $\hat{v}^{(1)} = v$, draw $K-1$ independent reverse paths, and give each candidate the target-to-proposal ratio

$$w^{(j)} \;=\; \frac{\tilde{q}\big(\hat{x}_0^{(j)}\big) \prod_{k=1}^{T} p\big(\hat{x}_k^{(j)} \mid \hat{x}_{k-1}^{(j)}\big)}{\prod_{k=1}^{T} p_\theta\big(\hat{x}_{k-1}^{(j)} \mid \hat{x}_k^{(j)}\big)}, \qquad \hat{x}_T^{(j)} = x_T, \tag{9}$$

selecting the next state with probability proportional to $w^{(j)}$. Including the current state in the pool leaves the conditional target invariant, again for any p_θ, and the probability of leaving the current state is nondecreasing in K, since each added candidate only enlarges the denominator of the selection probabilities. The motivation is hardware: one reverse pass leaves a modern GPU far from saturation for the models in this work, and the $K-1$ independent proposals batch into one pass.

4 Related Work

Path-space corrections. Accepting or rejecting an entire trajectory has been explored in several prior works. Tempered transitions (Neal, 1996) apply one test to a sequence of annealed moves, and the forward/reverse path density at the heart of (6) appears in PDDS (Phillips et al., 2024) as an SMC weight. The single-level ideal kernel is the proximal sampler (Lee et al., 2021; Chen et al., 2022), whose denoising posterior DiGS (Chen et al., 2024) samples with an inner Metropolis-corrected chain. Concurrently with this work, MAD-Path (Chen et al., 2026a) and SPS (Chen et al., 2026b) propose closely related MH corrections of learned reverse paths. Proposal acceptance requires a precise model at the target level and at each step along the path.

Exact MCMC with learned global moves. Gabrié et al. (2022) use flows as MH proposals, adapt them on the chain's own samples, and demonstrate repair of mode weights. Ex^{2}MCMC (Samsonov et al., 2022) composes global i-SIR proposals with local MALA refinement, which is the pattern (7) follows. NeuTra (Hoffman et al., 2019) instead reparameterizes HMC with a flow. This approach requires a learned model that is sufficiently accurate for global proposals.

Generative proposals trained from samples. Boltzmann generators train a normalizing flow on molecular dynamics or MCMC samples and correct it exactly through its tractable density, by importance reweighting or as an MH proposal (Noé et al., 2019; Köhler et al., 2020; Wirnsberger et al., 2022; Klein et al., 2023; Midgley et al., 2023). A flow must bridge noise and target in one invertible map with a tractable Jacobian, which constrains the architecture.

Diffusion samplers from the density alone. A large family trains a diffusion or controlled SDE by minimizing a path-space divergence to the target (Zhang & Chen, 2022; Vargas et al., 2023; Richter & Berner, 2024; Berner et al., 2024; Vargas et al., 2024; Chen et al., 2025; Havens et al., 2025; Liu et al., 2025; Guo et al., 2026), and a second replaces the variational objective with Monte Carlo estimates of the noised score, either learned (Akhound-Sadegh et al., 2024) or formed at sampling time (Huang et al., 2024; Grenioux et al., 2024). Across these methods, what is known about the target comes from states the sampler itself produces, and none applies an exact correction.

LJ-38. The double-funnel landscape and its solid-solid transition were established by basin hopping and disconnectivity analysis (Wales & Doye, 1997; Doye et al., 1999), and its thermodynamics by parallel tempering (Neirotti et al., 2000), exchange Monte Carlo (Mandelshtam & Frantsuzov, 2006), and nested sampling (Pártay et al., 2010). Exact inter-funnel moves from a Gaussian mixture fitted around known minima have been applied to LJ-38 (Finkler & Goedecker, 2020), but the mixture cannot capture non-Gaussian structures of the funnels and the method falls back on parallel tempering when the Gaussian approximation breaks down.

5 Experiments

We examine targets which have different degrees of multimodality and geometric complexity. Gaussian mixtures cannot be traversed by MALA but are otherwise simple. High-temperature particle potentials and a BNN posterior can be traversed by MALA on its own, but prior learned global samplers fall short on them. On LJ-38 below melting MALA cannot cross between funnels and PT is expensive, and there is the potential for global MCMC to make sampling much more viable. For each target we train a standard denoising diffusion model on a dataset of locally converged MALA samples and sample with the DDMC kernel (7). Producer chains are initialized from a normal or uniform distribution spanning the high-probability region of the target, except for LJ-38 below melting, where they are seeded from local minima found by basin hopping (Appendix C.2). All sampling chains are initialized from a normal or uniform distribution on every target, LJ-38 included. Unless stated otherwise, chains run for $10{,}000$ DDMC steps across 4 seeds, discard the first 400, and are thinned by four. We report acceptance, integrated autocorrelation time (IACT), and the energy Wasserstein-2 distance $E(\cdot)\,\mathcal{W}_2$ computed with the PDNS evaluation code (Guo et al., 2026), plus the sample distance $\mathcal{W}_2$ for particle targets. Alongside each distance we report the *ground-truth floor*, the same metric between two independent reference sets, at which a sampler is statistically indistinguishable from exact at this sample size. Results for the MH version (6) are labeled DDMC and those for the K-replicate i-SIR version (9) DDMC$_K$. Metrics, references and per-target configurations are in Appendix E, and Appendix H shows $\log r$ improving with denoiser training on every target.

5.1 Gaussian Mixtures

We start with Gaussian mixture models (GMMs) to test whether the MH correction repairs a misweighted dataset. MoG-40 is the 2-dimensional, 40-component mixture introduced by FAB (Midgley et al., 2023) and used by iDEM (Akhound-Sadegh et al., 2024) and DiGS (Chen et al., 2024). GMM-256 is the 256-dimensional, 2-component mixture of Grenioux et al. (2025) at mode separation $a = 10$, whose components have opposed diagonal covariances and uneven $1/3$–$2/3$ weights. Despite their simplicity, MALA and related local samplers cannot traverse these targets. In both cases the dataset is built from producer chains initialized at random and is far from the target's weighting (Table 1): roughly $50/50$ between the two modes of GMM-256 against the true $1/3$–$2/3$, and

Table 1: Gaussian mixtures against exact draws. Upper: TV is the total-variation distance of the mode weights from the truth, for the dataset and the DDMC chain. Lower: occupancy of the least-occupied mode for DDMC, for exact draws at matched sample size, and its true weight. For MoG-40 all 40 weights are equal, so the minimum over modes falls below $1/40$ even for exact draws. Every chain visits every mode, and all DDMC entries are statistically consistent with exact-draw controls at matched effective sample size.

Target	d	accept ↑	IACT ↓	$E(\cdot)\,\mathcal{W}_2$ ↓	*GT floor*	TV: dataset → DDMC
MoG-40	2	0.724 ± 0.004	1.14 ± 0.04	0.11 ± 0.03	0.09 ± 0.03	$0.251 \to 0.074 \pm 0.017$
GMM-256	256	0.390 ± 0.004	3.89 ± 0.73	0.71 ± 0.25	0.65 ± 0.13	$0.165 \to 0.014 \pm 0.010$

Least-occupied mode	DDMC	exact	true weight
MoG-40	0.017 ± 0.001	0.018 ± 0.001	0.025
GMM-256	0.345 ± 0.014	0.332 ± 0.009	0.333

Table 2: Particle targets, following the protocol of Guo et al. (2026): Wasserstein-2 distances on samples and on energies at $n = 2000$. Baselines as reported by Guo et al. (2026). Ours are mean ± sd across independent chains under their evaluation code and ground-truth samples. Bold marks the best mean per column and any entry whose uncertainty reaches it.

	DW-4 (d=8)		LJ-13 (d=39)		LJ-55 (d=165)	
Method	$\mathcal{W}_2$ ↓	$E(\cdot)\,\mathcal{W}_2$ ↓	$\mathcal{W}_2$ ↓	$E(\cdot)\,\mathcal{W}_2$ ↓	$\mathcal{W}_2$ ↓	$E(\cdot)\,\mathcal{W}_2$ ↓
PDDS (Phillips et al., 2024)	0.92 ± 0.08	0.58 ± 0.25	4.66 ± 0.87	56.01 ± 10.80	—	—
SCLD (Chen et al., 2025)	1.30 ± 0.64	0.40 ± 0.19	2.93 ± 0.19	27.98 ± 1.26	—	—
PIS (Zhang & Chen, 2022)	0.68 ± 0.28	0.65 ± 0.25	1.93 ± 0.07	18.02 ± 1.12	4.79 ± 0.45	228.70 ± 131.27
DDS (Vargas et al., 2023)	0.92 ± 0.11	0.90 ± 0.37	1.99 ± 0.13	24.61 ± 8.99	4.60 ± 0.09	173.09 ± 18.01
LV-PIS (Richter & Berner, 2024)	1.04 ± 0.29	1.89 ± 0.89	—	—	—	—
iDEM (Akhound-Sadegh et al., 2024)	0.70 ± 0.06	0.55 ± 0.14	1.61 ± 0.01	30.78 ± 24.46	4.69 ± 1.52	93.53 ± 16.31
AS (Havens et al., 2025)	0.62 ± 0.06	0.55 ± 0.12	1.67 ± 0.01	2.40 ± 1.25	4.50 ± 0.05	58.04 ± 20.98
ASBS (Liu et al., 2025)	$\mathbf{0.38 \pm 0.05}$	0.19 ± 0.03	1.59 ± 0.00	1.28 ± 0.22	4.00 ± 0.03	27.69 ± 3.86
PDNS (Guo et al., 2026)	0.51 ± 0.04	0.21 ± 0.03	$\mathbf{1.57 \pm 0.01}$	1.01 ± 0.18	3.95 ± 0.01	21.97 ± 3.14
DDMC (ours)	0.50 ± 0.05	$\mathbf{0.14 \pm 0.04}$	1.59 ± 0.01	$\mathbf{0.42 \pm 0.09}$	3.91 ± 0.00	$\mathbf{1.22 \pm 0.43}$
DDMC$_K$ (ours)	0.50 ± 0.05	$\mathbf{0.14 \pm 0.04}$	1.58 ± 0.00	$\mathbf{0.42 \pm 0.10}$	$\mathbf{3.88 \pm 0.00}$	$\mathbf{0.83 \pm 0.31}$
GT floor	0.31 ± 0.02	0.12 ± 0.03	1.44 ± 0.01	0.40 ± 0.13	2.57 ± 0.04	0.87 ± 0.23

unevenly across the equally weighted modes of MoG-40. DDMC nonetheless sits at the ground-truth floor on the energy distance and recovers the true mode weights within statistical uncertainty.

5.2 Particle Potentials

High-temperature particle potentials have more complex geometry than the Gaussian mixtures and allow direct comparison with diffusion-based samplers from prior work. Well-tuned MALA samples these targets on its own, so our results primarily demonstrate that the DDMC path proposal attains useful acceptance on a complex geometry (Table 3). As in prior work on these benchmarks (Köhler et al., 2020; Klein et al., 2023; Akhound-Sadegh et al., 2024; Guo et al., 2026), the denoiser is equivariant to permutation of particle labels and to rotation, since the potentials are invariant to both. Energy functions and temperatures are in Appendix B.1 and architectures in Appendix E.1. Table 2 shows that DDMC attains energy distances at or within uncertainty of the ground-truth floor on all three systems, with the largest margin over the published baselines on LJ-55, the highest-dimensional system.

MAD-Path (Chen et al., 2026a) is concurrent work that applies the same path-space Metropolis–Hastings step (6) to a variationally trained network. We ran their released code on LJ-55 and found that the sampled values of $-\log \tilde{q}$ settle roughly 750 units above the target's band. In a matched setup that does not include MALA, DDMC's chains instead sit in the band, with a median within one unit of the reference. Appendix F.2 gives the setup and the scope of this experiment.

5.3 Bayesian Neural Network

We now study the 550-dimensional BNN posterior presented by DiGS (Chen et al., 2024). Local MALA covers this density, so as in the previous section we primarily test trainability and exactness at scale, here with a DiT (Peebles & Xie, 2023) that is equivariant to permutations of the hidden

Table 3: Chain statistics on the particle targets. Paired entries are single-proposal / i-SIR.

Target	Sampler	accept ↑	IACT ↓
DW-4	DDMC / DDMC_{33}	0.545/0.941	1.30/1.03
LJ-13	DDMC / DDMC_{33}	0.277/0.831	5.37/1.37
LJ-55	DDMC / DDMC_{64}	0.028/0.342	29.8/4.67

Table 4: BNN posterior ($d = 550$, $K = 32$) against a pooled PT reference. Upper: mean $\pm$ sd over six runs at $T = 640$ and one long run at $T = 2560$. Lower: burn-in from random initialization to the reference energy band, with wall-clock for 1 and 512 parallel chains on one A6000. The DiGS comparison is in Appendix F.1.

Sampler	accept ↑	IACT ↓	$E(\cdot)$ $\mathcal{W}_2$ ↓	*GT floor*
DDMC_{32} (T=640)	0.132 ± 0.004	4.87 ± 0.27	1.03 ± 0.26	1.17 ± 0.27
DDMC_{32} (T=2560)	0.275	3.61	1.37	1.17 ± 0.27

Method	to converge	1 chain	512 chains
DDMC (T=640)	1–2 DDMC steps	18–36 s	403–806 s
MALA (adapted step)	71,200 MALA steps	99 s	130 s

units. This was our most difficult learning scenario: a fixed dataset of 6.4M samples eventually stopped improving $\log r$, with the path term drifting negative and erasing gains from the target term, so we trained the denoiser concurrently with producer workers generating fresh MALA samples. The resulting model with i-SIR proposals accepts global proposals at 13–28% in 550 dimensions and matches an independent parallel tempering reference at its energy floor (Table 4). From a random initialization the first proposal is accepted with probability near one and carries the chain into the equilibrium band, so convergence takes one to two DDMC steps, 2.8–5.5$\times$ faster than tuned MALA for a single chain, though a single device batching hundreds of chains favors MALA. The global jump of DiGS, run with its released code, was essentially never accepted on this target (Appendix F.1).

5.4 LJ-38 Below the Melting Temperature

Every target considered so far can be sampled with reasonable efficiency by either MALA or PT with MALA chains. Our final target cannot be traversed by MALA and presents a significant challenge for PT. The LJ-38 cluster is a canonical double-funnel landscape (Doye et al., 1999): its global minimum is a face-centered cubic truncated octahedron (O_h), competing with a far broader funnel of icosahedral (I_h) structures whose lowest member lies at only a slightly higher energy. The competition is between energy and entropy: the O_h funnel is deeper but narrow, with few low-lying minima, while the I_h funnel is shallower and contains thousands. The entrance to the global minimum is extremely narrow, and quenching random configurations reaches it roughly 0.018% of the time. We attribute PT's difficulty to this entrance: hot replicas mix freely and carry the structures they find down the ladder, which becomes inefficient when a region that carries high probability at low temperature remains rarely entered at high temperature.

Above melting LJ-38 behaves like the other Lennard-Jones targets: plain MALA is already effective, and DDMC attains non-trivial acceptance of its path steps (Appendix G.4). We focus instead on estimating the equilibrium occupancy $\pi(I_h)$ of the icosahedral funnel at two temperatures below the melting point, by gathering many samples and assigning each to O_h or I_h according to the local minimum found by quenching it. Such occupancies matter wherever a free energy difference decides which structure is observed, as in polymorph and conformer prediction. The funnel occupancy of LJ-38 is very challenging to obtain across MCMC methods.

Both the global proposals and the exactness of DDMC are necessary here. MALA cannot cross between funnels at either solid temperature, so the path step is needed for the occupancy estimate to do more than reproduce its initialization. At the lower temperature $\pi(I_h) \approx 0.7\%$, while the training dataset is built with an uninformative 50/50 split between the two funnels (Appendix C.2). Recovering the occupancy is therefore a reweighting by roughly two orders of magnitude, which the MH step performs and an uncorrected proposal could not. Table 5 reports $\pi(I_h)$ at the two temperatures for DDMC and PT. The two estimates are statistically consistent at both temperatures, to 0.50σ and 0.76σ. Agreement with PT is evidence against bias from incomplete mixing, the one way

Table 5: LJ-38 funnel occupancy and cost to precision below melting. $\pi(I_h)$ with standard error over 128 independent replicas per method and regime. C is total compute, including dataset generation and training for DDMC. $\mathrm{SEM}^2 C$ charges each method for the burn-in it discards. The asymptotic column excludes burn-in and DDMC's training. Lower is better in both.

T_{LJ}	Method	$\pi(I_h)$	burn-in	C (GPU-h)	$\mathrm{SEM}^2 C \downarrow$	asympt. $\downarrow$
0.085	DDMC	0.00706 ± 0.00009	0	734	6.38×10^{-6}	5.84×10^{-6}
	PT	0.00726 ± 0.00039	0.50	734	1.14×10^{-4}	5.72×10^{-5}
	agree, 0.50σ				17.9×	9.8×
0.135	DDMC	0.35212 ± 0.00146	0	737	1.57×10^{-3}	1.43×10^{-3}
	PT	0.35716 ± 0.00646	0.10	737	3.08×10^{-2}	2.77×10^{-2}
	agree, 0.76σ				19.6×	19.4×

that training imperfections such as funnel mis-weighting could still affect a finite run. The remaining question is efficiency.

Both methods were run to the same total compute, 734 and 737 GPU-hours split across 128 independent chains for each method and temperature. PT starts from random quenched configurations at $T_{\mathrm{LJ}} = 0.135$, and at $T_{\mathrm{LJ}} = 0.085$ half its ensembles start in each funnel from the same basin-hopping minima used to seed the DDMC dataset. The compute time for DDMC includes dataset generation and denoiser training in addition to sampling, while PT's includes only sampling. We measure precision per unit compute as $\mathrm{SEM}^2 \times C$, the squared standard error over independent chains times GPU-hours spent, a decorrelation time expressed in compute. At equal compute DDMC attains the same occupancy with $17.9\times$ and $19.6\times$ less variance. Burn-in fractions follow our pre-registered rule (Appendix G.2). DDMC's burn-in is zero because its first accepted proposal already lands in the equilibrium band. The asymptotic column excludes training and burn-in from both cost and sample set to measure the efficiency of an equilibrated chain. It is insensitive to the burn-in cut at $T_{\mathrm{LJ}} = 0.135$, where PT has converged. At $T_{\mathrm{LJ}} = 0.085$ PT's asymptotic figure was still improving slowly at the end of the run, so its ratio is an upper bound on DDMC's asymptotic advantage, though the slowing improvement suggests the bound is not far from the converged value.

Both methods have a compute requirement that must be met before producing usable data at all. PT required 324–389 GPU-hours at $T_{\mathrm{LJ}} = 0.085$ to (mostly) reach equilibrium from its initialization, against the 62 GPU-hours DDMC spends once on its dataset and denoiser, after which a fresh chain equilibrates within a short time. DDMC's initial cost is amortized: training is paid once per target, whereas PT requires fresh burn-in on every run.

A limitation of DDMC is that, unlike PT, a learned proposal might bypass or get stuck in regions that received insufficient weight in the training data. At $T_{\mathrm{LJ}} = 0.135$ one chain of 128 entered a basin of small but nonzero target mass and remained there for the last 7,443 of its 67,500 DDMC steps. Chains restarted inside the basin all escaped, with a mean dwell time of about 7,000 DDMC steps (Appendix G.3).

6 Conclusion

We proposed Denoising Diffusion Monte Carlo, which composes an exact path-space Metropolis–Hastings step driven by a standard denoising diffusion model with local MALA moves, and trains the denoiser by ordinary regression on locally converged MALA samples. It matches reference samples at the measured estimator floor, and on LJ-38 below its melting temperature it is consistent with parallel tempering on the funnel occupancy while attaining roughly an order of magnitude lower variance at equal total compute. Several limitations stand out. The dataset must cover the regions that carry mass, and a failure to do so can be difficult to detect without domain knowledge of the target. Each proposal costs two (MH) or more (i-SIR) full reverse paths. Harder targets are likely to need finer ladders and require more expensive reverse paths. A variety of directions remain open for future work. The unnormalized density plays no direct part in training since the denoiser is fit to samples alone, and making the energy available to the model is a natural extension. Our denoisers are trained on only a single target, and a unified model which supports different temperatures and related potentials would enable sampling across a family of targets. Since acceptance is driven by the optimality of a standard denoiser, whether sampling improves with the scaling behavior that denoising diffusion shows in the image domain is a question this work leaves open.

Reproducibility statement

The method is specified completely in Section 3: the path-space Metropolis–Hastings ratio (6), the i-SIR weights (9), the reverse-variance calibration (8), and the composed DDMC kernel (7). The dataset-generation procedure is described in Section 3.2 and Appendix C, and the reverse-path hyperparameters and their tuning in Section 3.3 and Appendix D. Per-target architectures, dataset sizes, training steps, and reverse-path configurations are listed in Appendix E.1. Metric definitions and the measured ground-truth floors are in Appendix E.2, the construction of every reference sample set, including the parallel tempering reference we constructed for the BNN, is in Appendix E.3, and compute requirements are in Appendix E.4. Training code, dataset-generation schedules, evaluation scripts, and model checkpoints will be released.

AI use statement

Generative AI tools (ChatGPT 5.5 from OpenAI, and Claude Opus 4.8, Claude Opus 5, Claude Fable 5, and Claude Fable 5.1 from Anthropic) were used during this project for implementing coding methods, creating scientific figures, suggesting experimental parameters, providing feedback on research methodology or experiments, interpreting results, assistance with writing and editing, literature search, and checking of derivations and experimental bookkeeping. Generative AI tools were not used to generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, propose or refine hypotheses, design research methodology or experiments, assist with translation, clean and reformat dataset, or support qualitative and thematic data analysis. We take the final responsibility for all content of this paper.

References

Tara Akhound-Sadegh, Jarrid Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, and Alexander Tong. Iterated denoising energy matching for sampling from Boltzmann densities. In *International Conference on Machine Learning*, 2024.

Christophe Andrieu, Anthony Lee, and Matti Vihola. Uniform ergodicity of the iterated conditional SMC and geometric ergodicity of particle Gibbs samplers. *Bernoulli*, 24(2):842–872, 2018.

Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024.

Han Chen, Sifan Liu, and Jun Yang. Markov chain Monte Carlo with diffusion paths. *arXiv preprint arXiv:2607.11631*, 2026a.

Junhua Chen, Lorenz Richter, Julius Berner, Denis Blessing, Gerhard Neumann, and Anima Anandkumar. Sequential controlled Langevin diffusions. In *International Conference on Learning Representations*, 2025.

Shiyang Chen, Moxian Qian, Gert Aarts, Biagio Lucini, and Kai Zhou. Stochastic path sampler for lattice field theory. *arXiv preprint arXiv:2606.13790*, 2026b.

Wenlin Chen, Mingtian Zhang, Brooks Paige, José Miguel Hernández-Lobato, and David Barber. Diffusive Gibbs sampling. In *International Conference on Machine Learning*, 2024.

Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. Improved analysis for a proximal algorithm for sampling. In *Conference on Learning Theory*, 2022.

Jonathan P. K. Doye, Mark A. Miller, and David J. Wales. The double-funnel energy landscape of the 38-atom Lennard-Jones cluster. *The Journal of Chemical Physics*, 110(14):6896–6906, 1999.

Simon Duane, Anthony D. Kennedy, Brian J. Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics Letters B*, 195(2):216–222, 1987.

Bradley Efron. Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011.

Jonas A. Finkler and Stefan Goedecker. Funnel hopping Monte Carlo: An efficient method to overcome broken ergodicity. *The Journal of Chemical Physics*, 152(16):164106, 2020.

Marylou Gabrié, Grant M. Rotskoff, and Eric Vanden-Eijnden. Adaptive Monte Carlo augmented with normalizing flows. *Proceedings of the National Academy of Sciences*, 119(10):e2109420119, 2022.

Charles J. Geyer. Practical Markov chain Monte Carlo. *Statistical Science*, 7(4):473–483, 1992.

Louis Grenioux, Maxence Noble, Marylou Gabrié, and Alain Oliviero Durmus. Stochastic localization via iterative posterior sampling. In *International Conference on Machine Learning*, 2024.

Louis Grenioux, Maxence Noble, and Marylou Gabrié. Improving the evaluation of samplers on multi-modal targets. In *ICLR Workshop on Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025.

Wei Guo, Jaemoo Choi, Yuchen Zhu, Molei Tao, and Yongxin Chen. Proximal diffusion neural sampler. In *International Conference on Learning Representations*, 2026.

W. Keith Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

Aaron J. Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Daniel S. Levine, Brandon M. Wood, Bin Hu, Brandon Amos, Brian Karrer, Xiang Fu, Guan-Horng Liu, and Ricky T. Q. Chen. Adjoint sampling: Highly scalable diffusion samplers via adjoint matching. In *International Conference on Machine Learning*, 2025.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.

Matthew Hoffman, Pavel Sountsov, Joshua V. Dillon, Ian Langmore, Dustin Tran, and Srinivas Vasudevan. NeuTra-lizing bad geometry in Hamiltonian Monte Carlo using neural transport. *arXiv preprint arXiv:1903.03704*, 2019.

Matthew D. Hoffman and Andrew Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.

Emiel Hoogeboom, Víctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3D. In *International Conference on Machine Learning*, 2022.

Xunpeng Huang, Hanze Dong, Yifan Hao, Yi-An Ma, and Tong Zhang. Reverse diffusion Monte Carlo. In *International Conference on Learning Representations*, 2024.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 26565–26577, 2022.

Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. In *Advances in Neural Information Processing Systems*, 2023.

Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for symmetric densities. In *International Conference on Machine Learning*, 2020.

Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured logconcave sampling with a restricted Gaussian oracle. In *Conference on Learning Theory*, pp. 2993–3050, 2021.

Guan-Horng Liu, Jaemoo Choi, Yongxin Chen, Benjamin Kurt Miller, and Ricky T. Q. Chen. Adjoint Schrödinger bridge sampler. In *Advances in Neural Information Processing Systems*, 2025.

Jun S. Liu, Faming Liang, and Wing Hung Wong. The multiple-try method and local optimization in Metropolis sampling. *Journal of the American Statistical Association*, 95(449):121–134, 2000.

Vladimir A. Mandelshtam and Pavel A. Frantsuzov. Multiple structural transformations in Lennard-Jones clusters: Generic versus size-specific behavior. *The Journal of Chemical Physics*, 124(20): 204511, 2006.

Kerrie L. Mengersen and Richard L. Tweedie. Rates of convergence of the Hastings and Metropolis algorithms. *The Annals of Statistics*, 24(1):101–121, 1996.

Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.

Laurence Illing Midgley, Vincent Stimper, Gregor N. C. Simm, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Flow annealed importance sampling bootstrap. In *International Conference on Learning Representations*, 2023.

Radford M. Neal. Sampling from multimodal distributions using tempered transitions. *Statistics and Computing*, 6(4):353–366, 1996.

Radford M. Neal. Mcmc using hamiltonian dynamics. In Steve Brooks, Andrew Gelman, Galin L. Jones, and Xiao-Li Meng (eds.), *Handbook of Markov Chain Monte Carlo*, pp. 113–162. Chapman & Hall/CRC, 2011.

J. P. Neirotti, F. Calvo, David L. Freeman, and J. D. Doll. Phase changes in 38-atom lennard-jones clusters. i. a parallel tempering study in the canonical ensemble. *The Journal of Chemical Physics*, 112(23):10340–10349, 2000.

Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.

Lívia B. Pártay, Albert P. Bartók, and Gábor Csányi. Efficient sampling of atomic configurational spaces. *The Journal of Physical Chemistry B*, 114(32):10502–10512, 2010.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.

Angus Phillips, Hai-Dang Dau, Michael John Hutchinson, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. Particle denoising diffusion sampler. In *International Conference on Machine Learning*, 2024.

Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. In *International Conference on Learning Representations*, 2024.

Gareth O. Roberts and Richard L. Tweedie. Exponential convergence of langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.

Sergey Samsonov, Evgeny Lagutin, Marylou Gabrié, Alain Durmus, Alexey Naumov, and Eric Moulines. Local-global MCMC kernels: the best of both worlds. In *Advances in Neural Information Processing Systems*, volume 35, pp. 5178–5193, 2022.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

Luke Tierney. Markov chains for exploring posterior distributions. *The Annals of Statistics*, 22(4): 1701–1728, 1994.

Francisco Vargas, Will Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. In *International Conference on Learning Representations*, 2023.

Francisco Vargas, Shreyas Padhy, Denis Blessing, and Nikolas Nüsken. Transport meets variational inference: Controlled Monte Carlo diffusions. In *International Conference on Learning Representations*, 2024.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.

David J. Wales and Jonathan P. K. Doye. Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A*, 101(28):5111–5116, 1997.

Peter Wirnsberger, George Papamakarios, Borja Ibarz, Sébastien Racanière, Andrew J. Ballard, Alexander Pritzel, and Charles Blundell. Normalizing flows for atomic solids. *Machine Learning: Science and Technology*, 3(2):025009, 2022.

Qinsheng Zhang and Yongxin Chen. Path integral sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022.

A Ergodicity of the Noise-to-Denoise Transition

Throughout this section $\mathbb{R}^d$ denotes the space on which the target is a proper density. For the particle targets that space is the mean-free subspace rather than the ambient coordinates (Appendix B.2).

The general principle: a Markov kernel that leaves q invariant and whose transition distribution has a component with a density positive at every pair (x, x') with $\tilde{q}(x') > 0$ is q-irreducible and aperiodic, and therefore ergodic for q (Tierney, 1994; Mengersen & Tweedie, 1996). It remains to check positivity in each case.

Reversibility at a single noise level. For an ideal reverse process the noise-to-denoise transition K_σ satisfies

$$q(x)\, K_\sigma(\hat{x} \mid x) \;=\; q(x)\, q(\hat{x}) \int \mathcal{N}(y; x, \sigma^2 I)\, \mathcal{N}(y; \hat{x}, \sigma^2 I) \,/\, q_\sigma(y)\, \mathrm{d}y, \tag{10}$$

which is symmetric in x and $\hat{x}$, so the transition is q-reversible and not merely q-invariant. This is the reversibility of the proximal sampler (Lee et al., 2021), of which the path construction of Section 3 is the multi-level form. In this limit the intermediate path carries no information beyond x_T and could be integrated out. Paths are needed only because the ideal conditional can be neither sampled nor evaluated in practice.

Positivity. The forward step $Y = X + \sigma Z$ has the everywhere-positive density $\mathcal{N}(y; x, \sigma^2 I_d)$, and the reverse pass composes Gaussian kernels with strictly positive variances, so the composite transition density is strictly positive on $\mathbb{R}^d \times \mathbb{R}^d$. The learned means can move the kernels anywhere but cannot remove support.

DDMC transition. Each component of a DDMC step leaves q invariant, so the step does. The forward increments (3) and reverse kernels (4) are positive-variance Gaussians and the acceptance probability (6) is strictly positive wherever $\tilde{q} > 0$, so the accepted-move component of the step satisfies the positivity condition of the general principle above, and the DDMC chain is ergodic for q.

B Particle Targets

B.1 Definitions

DW-4 is the four-particle double-well system of Köhler et al. (2020) in two dimensions, with energy

$$E(x) = \sum_{i<j} \left[0.9\,(d_{ij} - 4)^4 - 4\,(d_{ij} - 4)^2\right] \tag{11}$$

in the pair distances d_{ij} and $\log q = -E$, following the released iDEM code (Akhound-Sadegh et al., 2024). The Lennard-Jones targets combine a Lennard-Jones pair potential with a harmonic tether to the center of mass, which keeps the cluster bound and the density normalizable on the mean-free subspace (Köhler et al., 2020). Classical cluster simulations usually confine the atoms in a hard spherical shell instead (Neirotti et al., 2000; Mandelshtam & Frantsuzov, 2006), but we follow the harmonic tether used throughout the machine learning benchmarks. As in the released iDEM code, whose convention the subsequent machine learning literature has largely adopted, the pair term is summed over *ordered* pairs and therefore carries an effective factor of two. With the usual unordered-pair Lennard-Jones energy in reduced units, well depth and pair-minimum distance both equal to one,

$$E_{\mathrm{LJ}}(x) \;=\; \sum_{i<j} (d_{ij}^{-12} - 2\, d_{ij}^{-6}) \tag{12}$$

and $\bar{x}$ the center of mass, the energy is

$$E(x) \;=\; 2\,E_{\mathrm{LJ}}(x) \;+\; \frac{w}{2}\sum_i \|x_i - \bar{x}\|^2, \qquad \log q(x) = -\beta E(x), \tag{13}$$

with $w = 1$ for LJ-13 and LJ-55 and $w = 0.25$ for LJ-38. Because the pair term is doubled, the reduced temperature is $T_{\mathrm{LJ}} = 1/(2\beta)$. All particle energies are translation-invariant, so each target is proper only on the mean-free subspace. Appendix B.2 gives the consequences for the sampler.

DW-4, LJ-13 and LJ-55. These follow the benchmark convention established by Köhler et al. (2020) and Klein et al. (2023) and used by all of the baselines in Table 2, with $\beta = 1$, which for LJ-13 and LJ-55 is $T_{\mathrm{LJ}} = 0.5$ in the units above. LJ-13 and LJ-55 are magic-number clusters, so their landscapes have one dominant funnel.

LJ-38. The LJ-38 landscape has two funnels (Doye et al., 1999), which we label O_h and I_h after the octahedral and icosahedral motifs of their structures. The labels name the motifs rather than the exact symmetry of every minimum. The O_h funnel ends in the global minimum, a face-centered cubic truncated octahedron with $E_{\mathrm{LJ}} = -173.928$. The I_h funnel is built from icosahedral structures and contains many low-lying minima, the lowest of which lies at $E_{\mathrm{LJ}} = -173.252$. The funnels are separated by a barrier of about 4.2 in E_{LJ}. The O_h funnel is lower in energy but narrow, while the I_h funnel is wider and carries more entropy, so their relative weight changes with temperature. We assign each sample to a funnel by quenching it to its local minimum. The rare minima that belong to neither funnel are counted separately and are negligible below melting. LJ-38 uses $w = 0.25$ at $T_{\mathrm{LJ}} = 0.5$, 0.135 and 0.085. Appendix G.1 explains the choice of tether and temperatures.

B.2 Translation Invariance and the Mean-Free Subspace

The particle targets require one extension of the construction in Section 3. Their energies (11) and (13) depend only on pair displacements and on displacements from the center of mass, so they are exactly translation-invariant: $E(x + c \otimes \mathbf{1}) = E(x)$ for every $c \in \mathbb{R}^k$. Consequently $e^{-\beta E}$ is not integrable over the ambient space $\mathbb{R}^d$ with $d = nk$, and there is no invariant probability measure there. The target is proper only on the mean-free subspace

$$\mathcal{M} \;=\; \Big\{\, x \in \mathbb{R}^d \;:\; \textstyle\sum_{i=1}^n x_i = 0 \,\Big\}, \qquad \dim \mathcal{M} = D = (n-1)\,k, \tag{14}$$

with $D = 162$ for LJ-55 and $D = 111$ for LJ-38 against ambient $d = 165$ and 114. Tables quote the ambient d, following the convention of the benchmark literature.

Following prior equivariant diffusion models (Hoogeboom et al., 2022), we project every noise draw and every denoiser output onto $\mathcal{M}$, so all states stay mean-free and every Gaussian in (3) and (4) is a D-dimensional isotropic density on $\mathcal{M}$. The Gaussian normalizers cancel in (6), since forward and reverse kernels are paired at equal variances, so the acceptance ratio is unchanged. The calibration (8) divides by $n_{\mathrm{cal}}D$ rather than $n_{\mathrm{cal}}d$, and the ergodicity argument of Appendix A applies on $\mathcal{M}$.

C Dataset Production

C.1 Dataset Calibration

Beyond the choice of initialization discussed in Section 3.2, dataset generation has two phases: calibrating MALA and measuring the energy spectrum of locally converged states, then converting that calibration into the cheapest fixed schedule that produces new samples.

The first phase runs gradient ascent on $\log \tilde{q}$ with Adam from the chosen initialization until convergence, which locates a set of modes. We then run MALA from the modes with a per-chain adaptive step size until the mean of $\log \tilde{q}$ plateaus. The quantiles of $\log \tilde{q}$ at the plateau define a reference energy band against which all later convergence checks are measured, and the adapted step sizes at the plateau give the sampling step size for all later MALA legs. To tune the MALA step size across different regions, we take the largest h which does not freeze or leave the target band across equilibrated samples so that the step size is compatible with the tightest energy constraints across the target.

The second phase fixes a schedule. The producer runs gradient ascent followed by MALA, and its one free parameter is how far the ascent may overshoot the reference band before MALA takes over:

the ascent stops at $L(f) = L_{\text{band}} + f\,(L_{\text{mode}} - L_{\text{band}})$ for $f \in [0, 1]$, where $f = 0$ stops as soon as the band's energy is reached and $f = 1$ descends fully to the mode. Both extremes are defensible a priori: small f is cheap but a point on a gradient trajectory at the right energy is not yet a thermal sample and may not have settled into a basin, while large f guarantees a genuine mode but pays for the thermal climb twice. We sweep f, run each candidate to the plateau, and accept it only if no chain freezes and the final energy quantiles match the reference band within tolerance. The cheapest valid schedule by gradient evaluations per sample is selected, re-validated on a fresh batch, and then executed to produce a dataset.

C.2 Structure-Seeded Initialization for LJ-38

Below its melting temperature, the high-probability regions of LJ-38 cannot be effectively covered by the method above. Random configurations quenched to their inherent structure reach the O_h funnel roughly 0.018% of the time, and a large proportion of them end in defect minima that belong to neither funnel and have negligible mass below the melting point. In the liquid regime at $T_{\text{LJ}} = 0.5$, random initialization is sufficient as for other targets we study because no funnel structure exists.

To handle low-temperature initialization we draw the producer's starting configurations from a database of local minima obtained by basin hopping (Wales & Doye, 1997; Doye et al., 1999). Only the seeding stage differs from the procedure of Appendix C.1. Each seed is perturbed, quenched, and thermalized with MALA at the production temperature exactly as for every other target.

The two funnels are seeded differently. The O_h funnel is narrow and, at these temperatures, effectively a single basin, so its seeds are its global minimum. The I_h funnel contains many low-lying minima, and its seeds are drawn from the basin-hopping database with weight proportional to $\exp(-\beta E_i)$ at the production temperature, where E_i is the energy of minimum i. This weighting is deliberately naive: it uses minimum energies alone and ignores the geometry of the basin. Like the funnel split below, the mismatch is left for the MH step to correct.

Each seed is perturbed by isotropic Gaussian noise of scale 0.15, quenched to the nearest local minimum, and thermalized with $4{,}000$ MALA steps at the production temperature. The quench returns most seeds to the basin they came from. About 3% land in higher-energy minima that belong to neither funnel, and about 1% of O_h seeds cross into I_h. Within-basin equilibration takes at most $1{,}000$ MALA steps at both temperatures so the $4{,}000$-step budget is sufficient. Each temperature regime's dataset has 3M states.

The proportion between the funnels is set to an even $50/50$. The acceptance of a global proposal between well-separated regions is at most one minus the total variation distance between the dataset's and the target's weights on those regions. An even split bounds that distance by $1/2$, and hence keeps the bound on acceptance at or above $1/2$, so learning a useful sampler should be viable whatever the true occupancy turns out to be.

D Tuning the Reverse Path and the Training Diagnostic

D.1 Path Hyperparameters

Noise levels and ladder length. Of the four quantities in Section 3.3, $\sigma_{\max}$ controls whether global moves are possible at all and must be large enough to bridge the state-space distance between distinct modes. Above the data scale the true reverse transitions are near-Gaussian regardless of model quality, so those levels contribute little to the path term and an over-large $\sigma_{\max}$ is not a major risk provided T is large enough. Too small a $\sigma_{\min}$ pays a penalty because reverse-path errors at low noise are amplified by $1/\sigma^2$. Setting it a few times below the tuned MALA step size and sweeping nearby while watching acceptance is effective, and the value found at moderate T usually transfers to larger T. Larger T can in many cases improve acceptance. Simpler targets might only require $T \leq 80$ while more complex ones benefit from $T \geq 320$. We use the T that best balances acceptance against compute, which is easy to find with a grid sweep. Figure 2 shows a $\sigma_{\min}$ sweep on LJ-13.

The spacing rule itself is not tuned: we use the EDM ladder of Karras et al. (2022) for every target. We tried alternative spacings and did not observe an improvement in acceptance, so the ladder enters only through $\sigma_{\min}$, $\sigma_{\max}$ and T.

Calibration set size. Too few calibration states in (8) depress acceptance in a way that can be mistaken for a deficient model, because any departure from the optimal variance raises the per-level

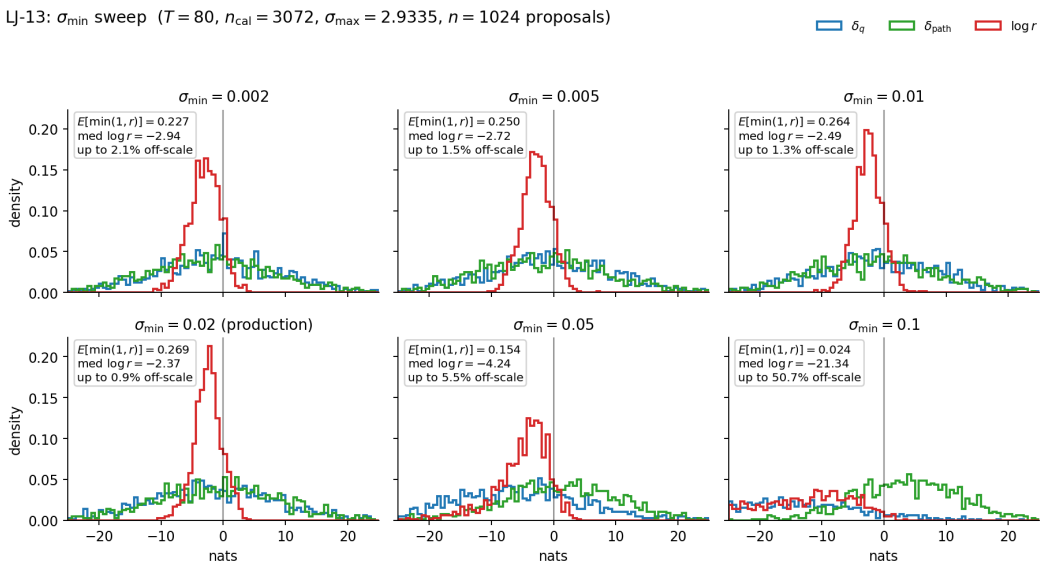


Figure 2: Single-proposal diagnostic on LJ-13 across $\sigma_{\min}$ at $T = 80$ and $n_{\text{cal}} = 3072$. Histograms of δ_q, δ_{path} and $\log r$ over 1024 proposals from held-out states. Expected acceptance rises gently as $\sigma_{\min}$ increases to the production value 0.02, then falls sharply once the final denoising step no longer lands within the thermal shell of the target and δ_q shifts strongly negative.

cost and estimation noise accumulates across levels as a downward bias rather than averaging out. Beyond a sufficient size the returns of more calibration samples vanish. In practice we raise n_{cal} at moderate T until acceptance stops improving and reuse that value for all T (Figure 3). The calibration states are disjoint from both the training set and the held-out states used for evaluation.

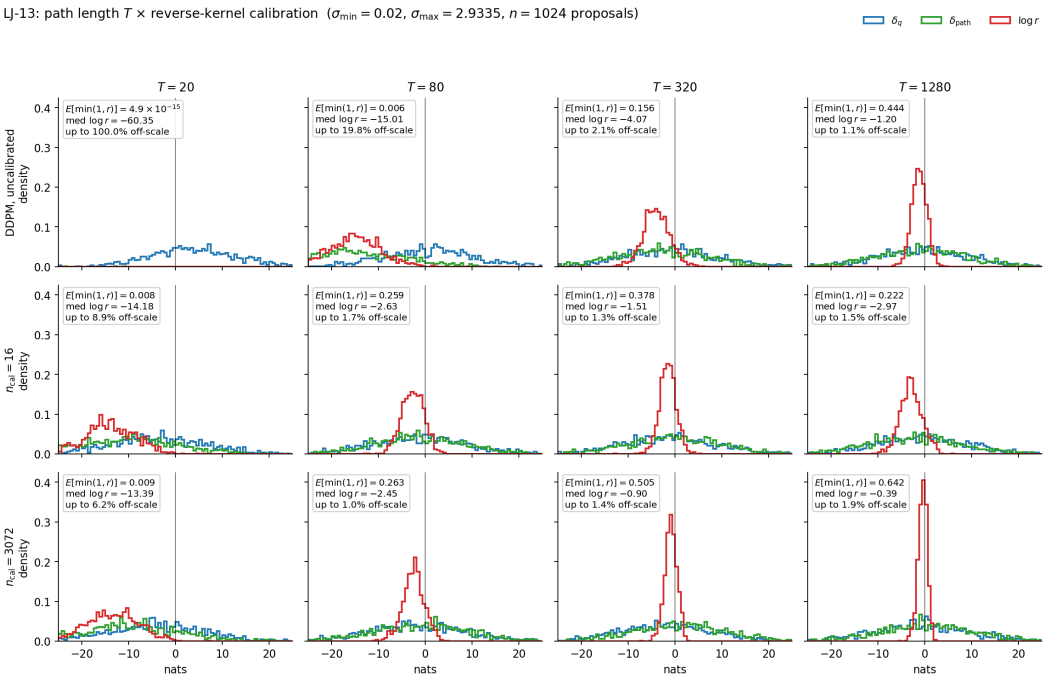


Figure 3: Single-proposal diagnostic on LJ-13 across path length T and reverse-variance calibration, at $\sigma_{\min} = 0.02$. Rows: uncalibrated DDPM variances, and calibration (8) with $n_{\text{cal}} = 16$ and 3072. Without calibration acceptance is negligible at small T and recovers only at large T. With too few calibration states acceptance stops improving and eventually degrades as T grows, because estimation noise accumulates across levels, while $n_{\text{cal}} = 3072$ improves monotonically with T.

D.2 The Single-Proposal Diagnostic

Monitoring whether a denoiser is a viable sampler drives decisions about data size, architecture and training length, but rolling out long chains to measure acceptance is costly. Held-out training samples are approximate equilibrium states, so a single-step proposal from one estimates path-kernel performance cheaply. We track the expected acceptance of one path proposal from a held-out state,

$$\bar{\alpha} \;=\; \mathbb{E}\big[\,\min\{1, e^{\log r}\}\big], \tag{15}$$

where $\log r$ is the log MH ratio of (6) and the expectation runs over the held-out state, its forward path, and one reverse proposal. It decomposes into a target term and a path term,

$$\log r = \delta_q + \delta_{\text{path}}, \qquad \delta_q = \log \tilde{q}(\hat{x}_0) - \log \tilde{q}(x_0), \tag{16}$$

$$\delta_{\text{path}} = \sum_{k=1}^{T} \Big[\log p(\hat{x}_k \mid \hat{x}_{k-1}) - \log p(x_k \mid x_{k-1}) \Big] + \sum_{k=1}^{T} \Big[\log p_\theta(x_{k-1} \mid x_k) - \log p_\theta(\hat{x}_{k-1} \mid \hat{x}_k) \Big], \tag{17}$$

with $\hat{x}_T = x_T$. The term δ_q measures whether the proposal improves on the current state under the target, and δ_{path} compares how well the forward and reverse processes are explained by the learned model. Both are needed: a chain driven by δ_q alone freezes at high $\log q$ instead of moving across the equilibrium energy spectrum. Since $\bar{\alpha}$ is governed by the right tail of $\log r$, we track its mean and standard deviation and inspect its histogram. A workable acceptance needs appreciable mass above zero or within a few nats below it.

For any MH ratio evaluated at a state drawn from the target, the untruncated ratio has unit mean, $\mathbb{E}[r] = 1$, because the density the pair is sampled from cancels the denominator of r and leaves the reversed joint, which integrates to one. This holds for (6) with any p_θ, and Jensen's inequality gives $\mathbb{E}[\log r] \le 0$ where the expectation is taken over the target. In our training runs we always observed $\mathbb{E}[\log r] \le 0$, even with the expectation taken over mis-weighted training data. For the ideal diffusion model we have $r = 1$ because the terms in the MH ratio cancel. Then $\delta_q = -\delta_{\text{path}}$ for every realization and the two terms are perfectly negatively correlated. A learned model should show a strong negative correlation that gets closer to -1 as denoiser quality improves, which is what we observe on every target. Figure 4 shows this cancellation for the final checkpoints on LJ-13, LJ-38 and LJ-55.

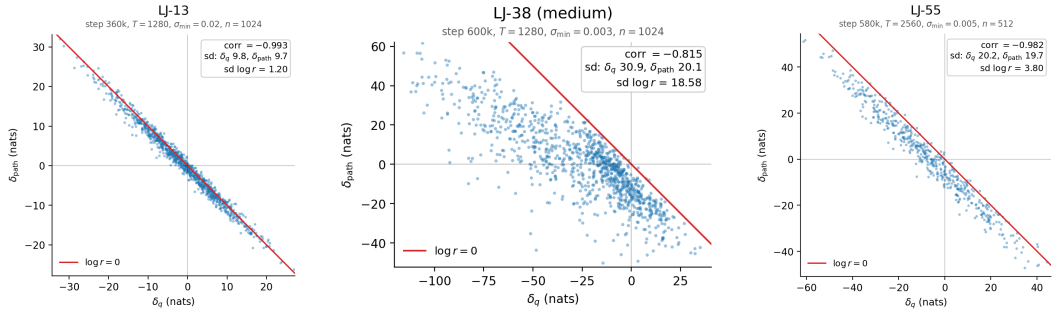


Figure 4: Per-proposal δ_{path} against δ_q from held-out states for the final checkpoints on LJ-13, LJ-38 ($T_{\text{LJ}} = 0.135$) and LJ-55. For the ideal model $\delta_q = -\delta_{\text{path}}$ and every point lies on the line $\log r = 0$. The learned models show strong negative correlation, and the two terms cancel, leaving $\log r$ with a smaller spread than either term.

E Experimental Settings

E.1 Per-Target Configurations

Table 6 summarizes the configuration used for each target. Full training configurations, dataset-generation schedules, and evaluation scripts will be released with the code.

Table 6: Per-target configuration. "Samples" is the number of MALA training states, and for the BNN the training set is refreshed continuously during training from a 6.4M-state seed. "Steps" is the training step of the shipped checkpoint. K is the i-SIR pool size where pooling is used, N is the number of MALA steps per DDMC step, and n_{cal} is the number of held-out states used to calibrate the reverse variances. LJ-38 uses one configuration for all three temperatures, with a separate model trained per temperature. Where two values are given, the first applies below melting ($T_{\text{LJ}} = 0.085$, 0.135) and the second above ($T_{\text{LJ}} = 0.5$). Producer chains are randomly initialized on every target except LJ-38 below melting (Appendix C.2).

Target	d	Samples	Architecture	Params	Steps	T	$\sigma_{\max}$	$\sigma_{\min}$	K	N	n_{cal}
MoG-40	2	10M	MLP (512×6)	3M	60K	320	0.38	0.005	—	20	3072
GMM-256	256	2M	token MLP (128×4)	0.9M	50K	320	18.92	0.003	—	20	3072
DW-4	8	4M	equivariant DiT (256×6)	7.5M	400K	40	1.62	0.05	33	20	3072
LJ-13	39	4M	equivariant DiT (256×6)	7.5M	360K	80	2.93	0.02	33	20	3072
LJ-55	165	6M	equivariant DiT (256×6)	7.5M	580K	320	4.0	0.005	64	20	2048
LJ-38	114	3M / 2M	equivariant DiT (256×6)	7.5M	600K	320	2.0	0.003	17 / 33	20	2048
BNN	550	6.4M	DiT (1024×16)	305M	1.4M	640/2560	20.0	0.003	32	1000	3072

E.2 Metrics

The energy Wasserstein-2 distance $E(\cdot)\,\mathcal{W}_2$ and the sample Wasserstein-2 distance $\mathcal{W}_2$ are computed with the evaluation code released by Guo et al. (2026), unmodified, at $n = 2000$ samples per chain. All metrics are computed per chain and never pooled across chains, and the reported uncertainty is the standard deviation across independent chains. Ground-truth floors are measured by evaluating the same metric between two independent reference sample sets at the same n: disjoint splits of a stored reference bank, or fresh exact draws when the target is analytically samplable, averaged over 20 to 50 replicate pairs. The floor is the value of the metric when both sample sets are exact, so results at the floor are statistically indistinguishable from exact sampling at this sample size, and rankings between values inside the floor are not meaningful. The integrated autocorrelation time (IACT) is estimated on the raw energy trace of each chain, in DDMC steps, before thinning, using the initial-positive-sequence estimator of Geyer (1992). IACT measures the number of chain steps needed to obtain one effectively independent sample, so a trace of length n carries roughly n/IACT independent samples and a value near one indicates nearly independent consecutive states.

E.3 Reference Samples

For MoG-40 and GMM-256 the reference is exact draws from the analytically samplable target. For the particle targets the reference is the published iDEM/PDNS ground-truth test split for each system, and we verified that the DW-4 split shipped with the PDNS code is byte-identical to iDEM's. For the BNN we additionally constructed a reference by parallel tempering: seven independent runs pooled, each with a 40-rung likelihood-only ladder, totaling $2{,}748{,}992$ states with $9{,}632$ hot-to-cold traversals and minimum swap acceptance 0.33. For LJ-38 above melting we constructed a reference of $16{,}384$ independent states, each equilibrated with $200{,}000$ exact MALA steps (Appendix G.4).

E.4 Compute

All experiments ran on NVIDIA A6000 GPUs with PyTorch. Path log-densities are accumulated in double precision to obtain a numerically accurate MH acceptance probability. The largest model (BNN, 305M parameters) was trained on two A6000 GPUs for approximately 280 GPU-hours, with sample generation running concurrently on the remaining GPUs of the node. The particle and mixture models (0.9M–7.5M parameters) are one to two orders of magnitude cheaper to train. Sampling for all reported chains used one GPU per chain.

F Details of Comparisons with Prior Work

F.1 Comparison with DiGS on the BNN Target

Burn-in. Running the released DiGS implementation at its own hyperparameters on the BNN target, its global jump was never accepted in over three million attempts, so its burn-in proceeds through its MALA component alone, and it differs from tuned MALA only in step size and in the cost of proposing jumps that are never accepted. From a random initialization DiGS reaches the reference

Table 7: Test NLL on the BNN target under the protocol of Chen et al. (2024) (mean $\pm$ sd across 4 seeds, expected-NLL estimator). No pair of methods is statistically separable after correcting for multiple comparisons.

Method	test NLL $\downarrow$	wall-clock
DiGS	0.229 ± 0.020	1032 s
MALA (DiGS step size)	0.237 ± 0.038	852 s
MALA (tuned step size)	0.208 ± 0.037	740 s
HMC	0.216 ± 0.025	526 s
DDMC (one proposal + MALA)	0.194 ± 0.011	787 s

energy band after 129,800 steps, 330 s for one chain and 356 s for 512 chains on one A6000, against 18–36 s for the one to two DDMC steps of Table 4, a 9–18$\times$ difference for a single chain.

Test NLL. The original DiGS evaluation on the BNN target reports test negative log-likelihood under a fixed compute protocol: two chains, 75 recorded samples with 5,000 gradient steps between recordings. We repeat this protocol for five methods across four seeds, running DiGS with its exact released implementation. For DDMC, a single path proposal moves the chain to the equilibrium region (its wall-clock cost is converted to MALA steps and charged against the budget) and MALA is run for the remainder. Results are in Table 7. All methods have statistically indistinguishable performance on this metric after accounting for multiple comparisons.

F.2 MAD-Path on LJ-55

MAD-Path (Chen et al., 2026a) applies a path-space MH correction closely related to ours to a reverse process trained variationally from the density, and its authors release code. We ran their released pipeline on LJ-55, which lies outside the problems their paper evaluates. We tried several configurations, including longer training, a wider network, and a finer discretization, and report the best: their mode pre-pass with merge radius 60 and default training. MAD-Path uses no local moves, so for a matched comparison we disabled MALA in DDMC and ran single-proposal MH chains with the same checkpoint as Section 5.2. Both methods start from their own initializations, run 4 chains of 10,000 steps, and are scored with the same evaluation code against the same reference (Table 8).

Table 8: MAD-Path and DDMC on LJ-55 without local moves, 4 chains of 10,000 steps each, mean $\pm$ sd over chains. The median is of $-\log \tilde{q}$. The reference is the iDEM/PDNS ground-truth split, and its floor is at 2,000 independent states. Without MALA each DDMC chain retains only about 280 distinct states and more steps are needed to more closely approach the reference energy band.

Sampler	accept $\uparrow$	$E(\cdot)\,\mathcal{W}_2 \downarrow$	$-\log \tilde{q}$ median
MAD-Path (best configuration)	0.0005	783 ± 21	$+427 \pm 15$
DDMC, single proposal, no MALA	0.0290	2.50 ± 1.19	-325.71 ± 1.06
reference	—	0.83 ± 0.18	-325.80 ± 0.38

MAD-Path's chains remain far above the target's energy band, while DDMC's reach it with the first accepted proposal and match the reference quantiles. MAD-Path's few accepted moves come at the beginning of the trajectory, after which the chains appear frozen in a high-energy configuration.

G LJ-38 Experimental Details

G.1 Tether and Temperatures

The relative weighting of the two LJ-38 funnels is affected by the tether strength w. At the strong tether $w = 1$ the funnel probability masses never reach a $50/50$ balance below the melting onset. We therefore use $w = 0.25$, for which the balance point sits at $T_{\mathrm{LJ}} \approx 0.141$, comfortably below the measured melting onset of 0.168–0.177.

We report three temperatures, chosen to span the qualitatively distinct regimes of this landscape:

- $T_{\mathrm{LJ}} = 0.5$ (*high*), above melting. The cluster is liquid, the funnels are not distinct states, and local sampling mixes structural identity on its own.

- $T_{\mathrm{LJ}} = 0.135$ (*medium*), near the balance point. Both funnels carry substantial weight, so the occupancy is an ordinary two-state question and neither funnel is rare. O_h has roughly $2/3$ of the probability mass and I_h has roughly $1/3$.
- $T_{\mathrm{LJ}} = 0.085$ (*low*). The icosahedral funnel retains a small but non-negligible share of the mass, below 1%, which is the regime where an estimator must be both exact and able to reach a rare state.

G.2 Pre-Registered Comparison Protocol

The LJ-38 comparison rules were fixed before the runs. A replica is one DDMC chain or the coldest rung of one PT ensemble, and we report the mean of $\pi(I_h)$ over the 128 replicas per method and regime, with the standard error over replicas as the error bar rather than a pooled binomial error, since PT's replica scatter is several times the binomial expectation. Each method discards its own burn-in, which is charged to its cost but excluded from its data, and the cutoff is the smallest burn-in fraction, on a grid of 0, 0.1, 0.2, and so on, after which the estimate stops moving within the combined standard error. DDMC's estimate is flat from the first DDMC step in both regimes, so its cutoff is zero. Agreement is tested as $|\pi_{\mathrm{PT}} - \pi_{\mathrm{DDMC}}|$ in units of $\sqrt{\mathrm{SEM}_{\mathrm{PT}}^2 + \mathrm{SEM}_{\mathrm{DDMC}}^2}$, and the two estimates are never merged into a consensus reference. A variance estimated from 128 replicas is known only to about $\pm 13\%$ and the per-replica distributions are heavy-tailed, so we treat the efficiency ratios as order-of-magnitude statements.

G.3 Escape From a Long-Lived Defect Basin

One replica of the $T_{\mathrm{LJ}} = 0.135$ comparison entered a basin of defect structures at DDMC step $60{,}058$ and stayed there for the remaining $7{,}443$ steps of its $67{,}500$-step run. To measure the escape time directly, we restarted 32 independent replicas from snapshots of that chain inside the basin, using the same kernel as the comparison, and ran each until it left. All 32 escaped, each into one of the two funnels (16 to O_h and 16 to I_h). The maximum likelihood mean dwell time is 7.1×10^3 DDMC steps with individually observed escape times from 93 to 4.0×10^4 steps. Of the 32 replicas, 9 stayed at least as long as the original chain. The chain encountered a long but observable dwell time rather than becoming indefinitely stuck in an absorbing state.

G.4 LJ-38 Above Melting

At $T_{\mathrm{LJ}} = 0.5$ the cluster is liquid, the funnels are not distinct states, and the target is similar to the LJ systems treated in Section 5.2. We evaluate it the same way, against the energy spectrum of $16{,}384$ independently equilibrated long exact-MALA states, from chains started at a random initialization (Table 9). The i-SIR arm sits at the ground-truth floor. The single-proposal arm has not fully equilibrated at this budget, so its energy distance is an upper bound rather than a converged value.

Table 9: LJ-38 above its melting temperature, evaluated as the other particle targets. Paired entries are single-proposal / i-SIR.

Sampler	accept ↑	IACT ↓	$E(\cdot)\,\mathcal{W}_2$ ↓	*GT floor*
DDMC / DDMC_{33}	$0.137/0.631$	$11.4/2.17$	$1.06/0.71$	0.64 ± 0.17

H Diagnostic Terms During Training

The figures below show the diagnostic quantities of Appendix D.2 over training for each target: histograms of δ_q, δ_{path}, and $\log r$, together with the evolution of their means and of the acceptance $\bar{\alpha}$ (15). The sampling configurations used for these diagnostics were chosen for monitoring during training and may differ from the inference settings of Section 5, but the trends are similar across configurations even when the $\log r$ shown in a diagram is not the optimal one. In healthy training δ_{path} tends to remain clustered around zero. The forward and reverse path densities balance early, and the improvement in $\log r$ comes from δ_q, that is, from the diffusion placing its proposal mass closer and closer to the true target. This behavior deserves a closer investigation, which we leave to future work.

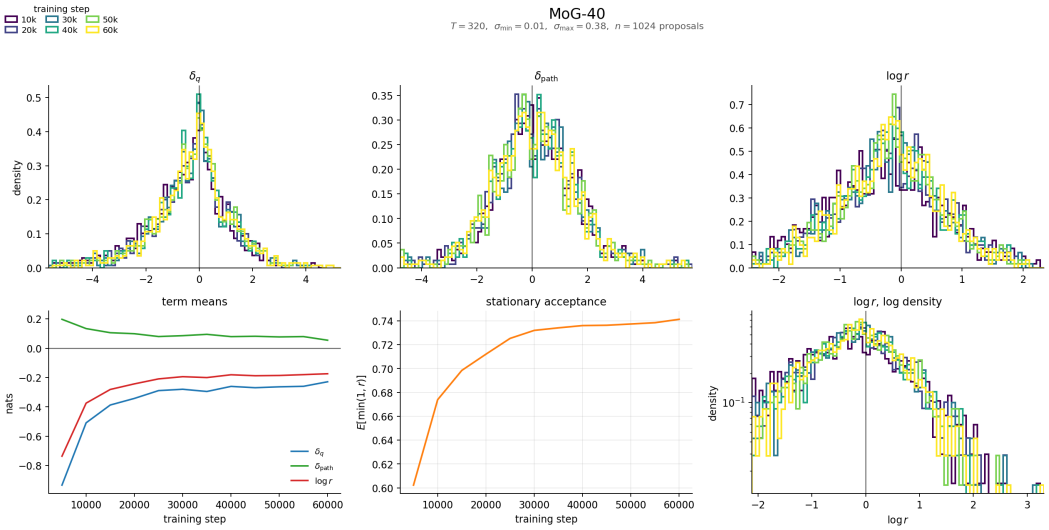


Figure 5: Training diagnostics for MoG-40.

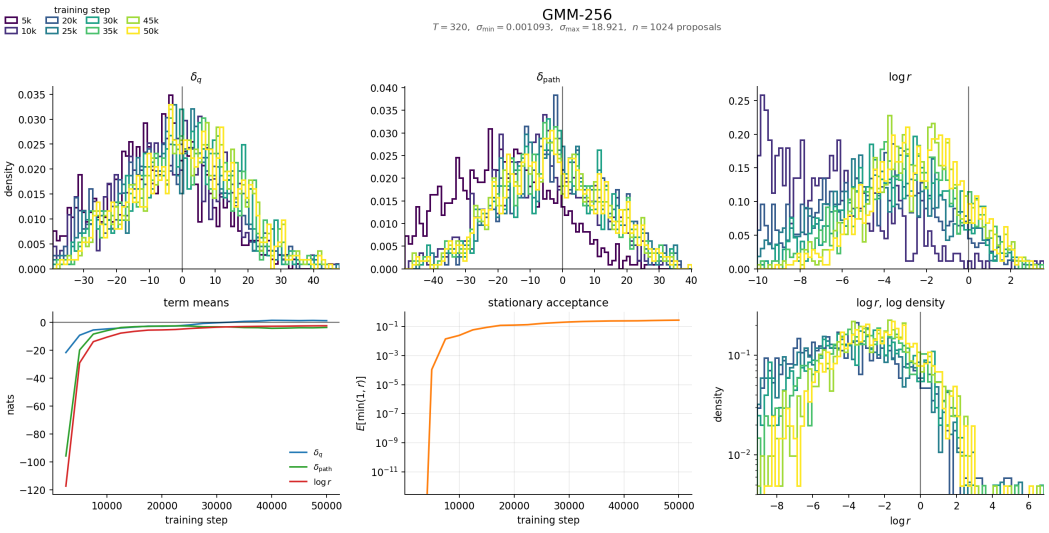


Figure 6: Training diagnostics for GMM-256.

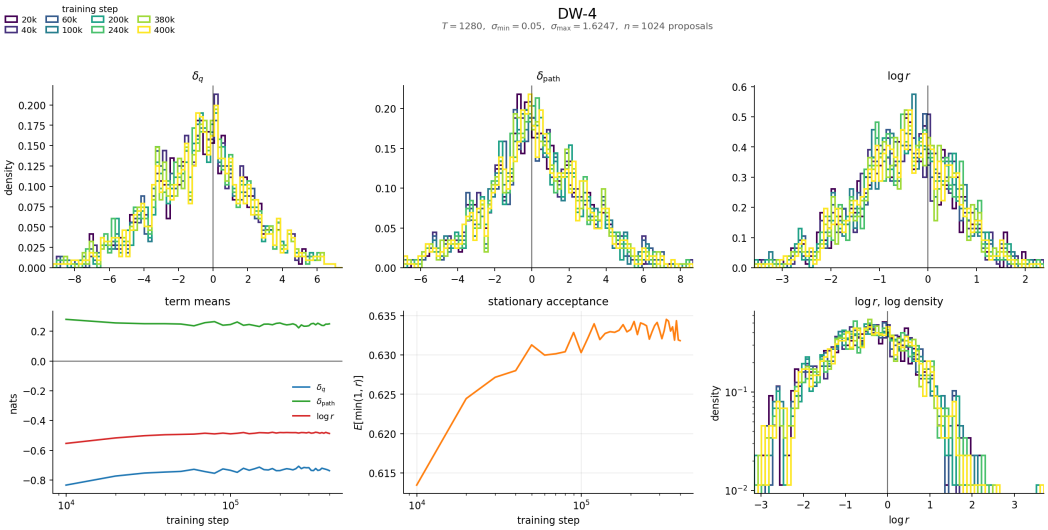


Figure 7: Training diagnostics for DW-4.

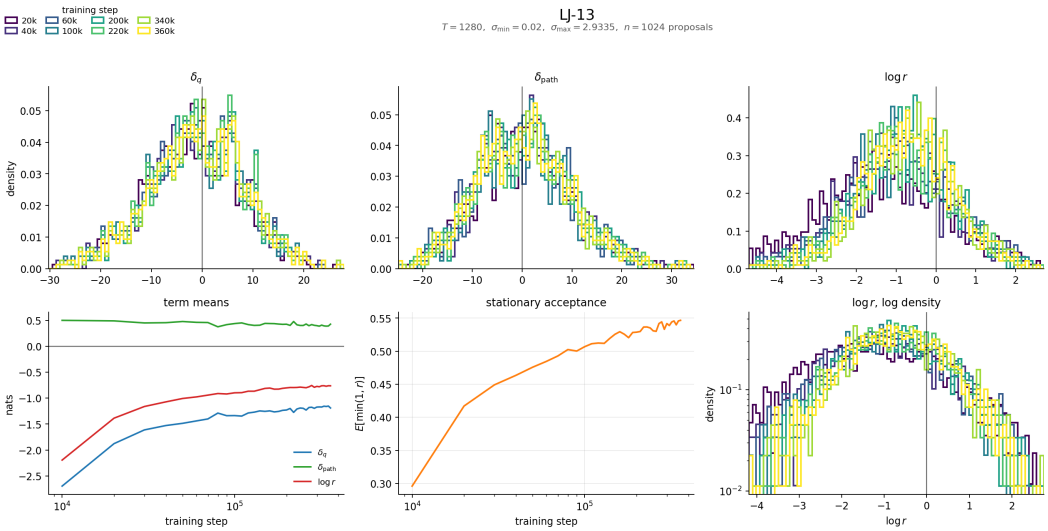


Figure 8: Training diagnostics for LJ-13.

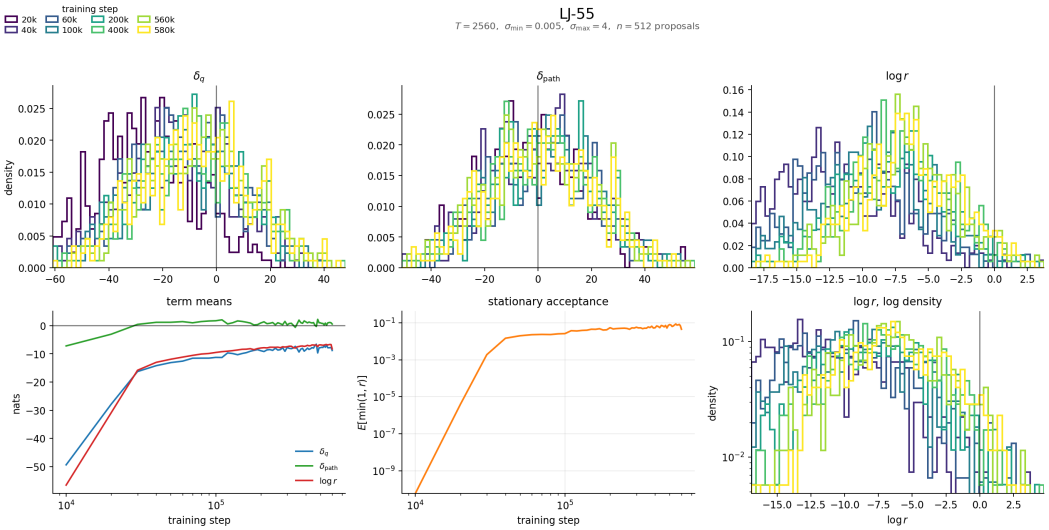


Figure 9: Training diagnostics for LJ-55.

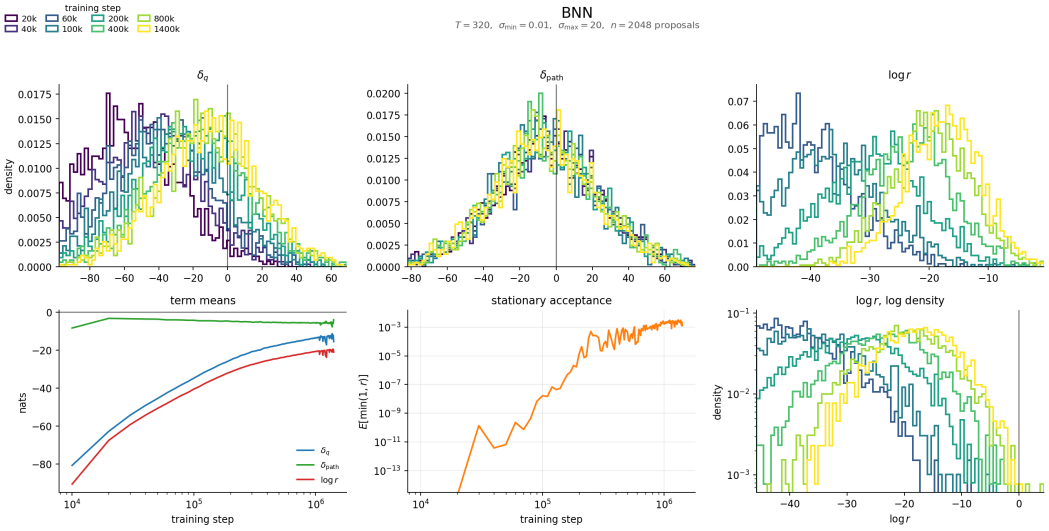


Figure 10: Training diagnostics for the BNN posterior.

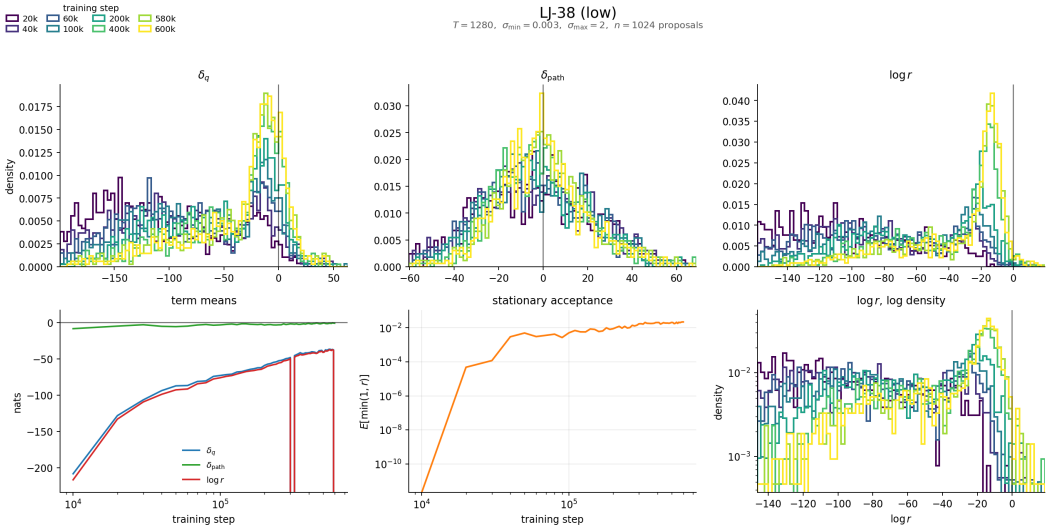


Figure 11: Training diagnostics for LJ-38 at $T_{\text{LJ}} = 0.085$. At a few evaluation snapshots the term means are dominated by a small number of strongly negative outliers, while most samples and the acceptance remain healthy.

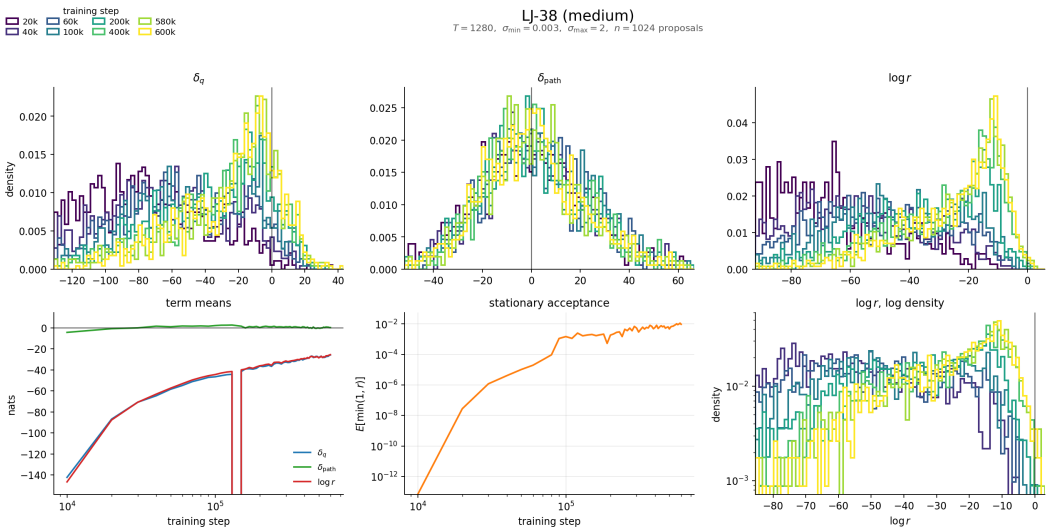


Figure 12: Training diagnostics for LJ-38 at $T_{\text{LJ}} = 0.135$. At a few evaluation snapshots the term means are dominated by a small number of strongly negative outliers, while most samples and the acceptance remain healthy.

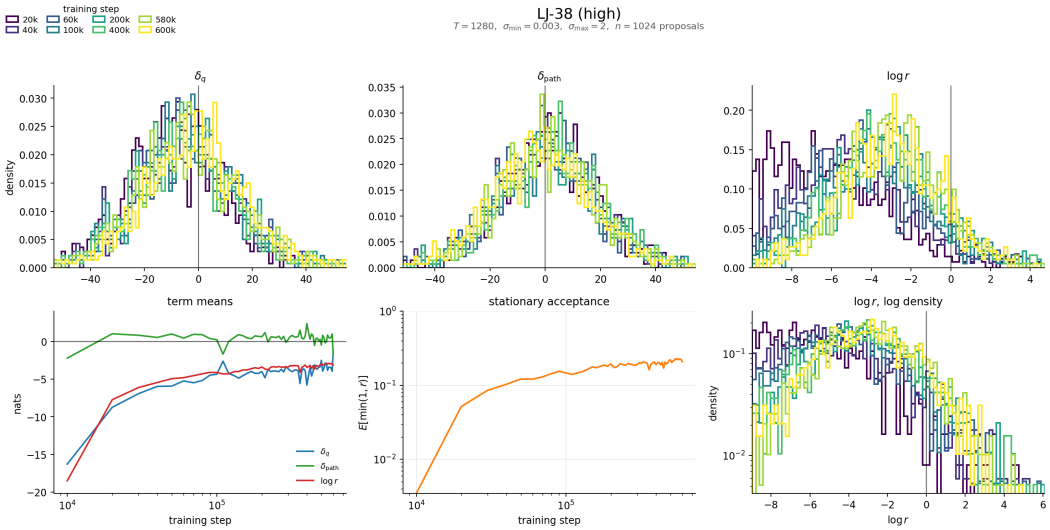


Figure 13: Training diagnostics for LJ-38 at $T_{\mathrm{LJ}} = 0.5$.