

Efficient and Unbiased Sampling from Boltzmann Distributions via Variance-Tuned Diffusion Models

Anonymous authors

Paper under double-blind review

Abstract

Score-based diffusion models (SBDMs) are powerful amortized samplers for Boltzmann distributions; however, imperfect score estimates bias downstream Monte Carlo estimates. Classical importance sampling (IS) can correct this bias, but computing exact likelihoods requires solving the probability-flow ordinary differential equation (PF-ODE), a procedure that is prohibitively costly and scales poorly with dimensionality. We introduce Variance-Tuned Diffusion Importance Sampling (VT-DIS), a lightweight post-training method that adapts the per-step noise covariance of a pretrained SBDM by minimizing the α -divergence ($\alpha = 2$) between its forward diffusion and reverse denoising trajectories. VT-DIS assigns a single trajectory-wise importance weight to the joint forward–reverse process, yielding unbiased expectation estimates at test time with negligible overhead compared to standard sampling. On the DW-4, LJ-13, and alanine-dipeptide benchmarks, VT-DIS achieves effective sample sizes of approximately 80%, 35%, and 3.5%, respectively, while using only a fraction of the computational budget required by vanilla diffusion + IS or PF-ODE-based IS.

1 Introduction

Sampling from Boltzmann distributions is central to many applications in statistical physics, chemistry, and materials science. Accurate equilibrium samples enable *in silico* prediction of thermodynamic observables, accelerating the discovery of catalysts, drugs, and novel materials and reducing reliance on costly wet-lab experiments. However, the high dimensionality and multimodal nature of realistic energy landscapes render traditional Markov chain Monte Carlo (MCMC) (Metropolis et al., 1953; Hastings, 1970) and molecular dynamics (MD) simulations (Frenkel & Smit, 2023) prohibitively time-consuming for many practical problems.

Normalizing-flow (NF) Boltzmann generators (Noé et al., 2019; Tan et al., 2025; Köhler et al., 2020) amortize sampling by learning an expressive, invertible mapping from a simple base distribution to the target, and exact density evaluation of generated samples permits unbiased importance sampling (IS). However, the requirement of tractable Jacobian determinants constrains architectural flexibility, becoming a bottleneck in high dimensions. Continuous normalizing flows (CNFs) (Chen et al., 2018; Grathwohl et al., 2019) mitigate this issue by parameterizing the transformation as an ordinary differential equation (ODE) and have been combined with E(3)-equivariant architectures for molecular systems (Garcia Satorras et al., 2021; Klein et al., 2023). Yet even CNFs demand expensive Jacobian computations for log-density evaluation.

Score-based diffusion models (SBDMs) (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2020) offer an expressive alternative that does not require invertibility. As with any learned sampler, imperfect score estimates bias downstream Monte Carlo estimates. Two IS-based corrections exist for a pretrained score network: (i) solve the probability flow ODE to obtain exact likelihoods—an approach that is computationally intensive, and whose approximations reintroduce bias; or (ii) solve the reverse stochastic differential equation (SDE) and perform IS over the joint forward–reverse trajectory, whose density factorizes into Gaussian terms; however, naively applying this approach typically yields a very low effective sample size (ESS) because the joint proposal distribution cannot fully cover the target distribution.

We adopt the second approach and show that a simple post-training covariance adjustment makes it practical. We introduce **variance-tuned diffusion importance sampling** (VT-DIS), which freezes the score network

and optimizes only the per-step noise covariance by minimizing the α -divergence ($\alpha = 2$) between the forward diffusion and reverse sampling processes. VT-DIS yields an unbiased trajectory-wise estimator and achieves high ESS at an orders-of-magnitude lower computational cost than vanilla diffusion + IS or NF-based samplers. Our contributions are:

- **Algorithm.** We propose VT-DIS, a post-training procedure that endows diffusion-based Boltzmann samplers with high-ESS, trajectory-wise importance sampling while incurring negligible overhead relative to sample generation. Furthermore, we extend VT-DIS to learn both diagonal and full covariance structures, moving beyond isotropic noise.
- **Equivariant extension.** We demonstrate how VT-DIS integrates with E(3)-equivariant diffusion models, enabling the unbiased generation of molecular samples that respect desired symmetry constraints.
- **Empirical validation.** On Gaussian mixtures, toy many-particle systems, and molecular benchmarks (DW-4, LJ-13, and alanine dipeptide), VT-DIS significantly increases ESS relative to naive diffusion + IS with the same number of sampling steps, while incurring only a fraction of the computational cost required for exact likelihood computation via the probability-flow ODE.

2 Related Work

Normalizing-flow (NF) Boltzmann generators learn an invertible mapping from a simple base distribution to the Boltzmann target, enabling exact evaluation of sample densities and unbiased importance sampling via the change-of-variables formula (Noé et al., 2019). Subsequent work improved expressiveness and scalability: Continuous Normalizing Flows (CNFs) based on neural ODEs (Chen et al., 2018; Grathwohl et al., 2018); symmetry-aware flows introduce E(3)-equivariance for molecular systems (Klein et al., 2023). All CNF variants, however, still require computationally expensive Jacobian determinants or ODE integration for exact likelihood evaluation, making unbiased sampling prohibitive for large biomolecules.

Score-based diffusion models learn the noise-perturbed score and generate samples either by reversing the stochastic differential equation (SDE) (Sohl-Dickstein et al., 2015; Ho et al., 2020) or by integrating the probability-flow ODE, which yields exact likelihoods (Song & Ermon, 2019). Although diffusion methods have proven effective for molecular data (Hooeboom et al., 2022), evaluating proposal log densities via the PF ODE still requires Jacobian-trace computation and becomes expensive in high dimensions.

Beyond the two popular approaches considered above, neural samplers (Vargas et al., 2023a; Richter & Berner, 2023; Vargas et al., 2023b; Zhang & Chen, 2021) can, in principle, generate unbiased samples directly from a target energy function without requiring training data. However, He et al. (2025) show that these methods typically incur substantial computational overhead and are often less efficient than generating training data via MCMC and subsequently fitting a diffusion model. Scalability is also a concern: to the best of our knowledge, no neural sampler currently performs well on molecules in 3D Cartesian coordinates—the focus of this work. We therefore omit neural-sampler baselines from our experiments.

VT-DIS: Bridging the Gap. VT-DIS retains the efficient stochastic reverse trajectory of standard diffusion models while enabling unbiased reweighting. Instead of integrating the ODE, VT-DIS learns per-step Gaussian covariances by minimizing the α -divergence ($\alpha = 2$) between the forward diffusion and reverse sampling processes, thereby maximizing ESS. Concurrent work has tuned covariances under the Kullback–Leibler divergence to improve sample quality (Ou et al., 2024; Bao et al., 2022); VT-DIS differs by explicitly optimizing the α -divergence objective for importance-sampling efficiency.

3 Background

Before presenting our method, we review three key ingredients: (i) score-based diffusion models (DMs), (ii) their E(3)-equivariant extensions, and (iii) the role of importance sampling (IS) in obtaining unbiased estimators at inference time.

3.1 Score-Based Diffusion Models

Score-based diffusion models (DMs) (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2020) generate data by first *diffusing* a sample from the data distribution $\pi_{\text{data}}(\mathbf{x})$ with progressively increasing Gaussian noise and then *denoising* it in reverse time. The forward process is described by the Itô stochastic differential equation (SDE)

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t) dt + g(t) d\mathbf{w}_t, \quad (1)$$

where $t \in [0, T]$ with $T > 0$, \mathbf{f} is the drift, g the diffusion coefficient, and \mathbf{w}_t a standard Wiener process. Throughout this work we follow the commonly used *variance-exploding* setting of Karras *et al.* (Karras et al., 2022) and set $\mathbf{f} \equiv \mathbf{0}$ and $g(t) = \sqrt{2t}$. The forward SDE admits two equivalent reverse-time formulations:

$$d\mathbf{x}_t = -2t \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t) dt + \sqrt{2t} d\bar{\mathbf{w}}_t, \quad (\text{reverse SDE}) \quad (2)$$

$$d\mathbf{x}_t = -t \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t) dt, \quad (\text{probability-flow ODE}), \quad (3)$$

where $p_t(\mathbf{x}_t)$ is the marginal density at time t and $\bar{\mathbf{w}}_t$ is a Wiener process running backwards from T to 0. In practice we approximate the (intractable) score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t)$ with a neural network $s_{\theta}(\mathbf{x}_t, t)$ trained by score matching (Hyvärinen & Dayan, 2005; Song & Ermon, 2019).

To draw a sample, we integrate either equation 2 or equation 3 backward from $t = T$ to a small cutoff $\epsilon \approx 0$, and take \mathbf{x}_{ϵ} as an approximate draw from π_{data} . Solving the reverse SDE with the discrete DDPM sampler (Ho et al., 2020) amounts to the update rule $p_{\theta}(\mathbf{x}_{n-1} | \mathbf{x}_n) = p(\mathbf{x}_{n-1} | \mathbf{x}_n, \hat{\mathbf{x}}_0(\mathbf{x}_n, t_n; \theta))$, where $\hat{\mathbf{x}}_0$ is the network’s prediction of the clean signal at step n . Given a time grid $\{t_n\}_{n=0}^N$, we obtain the full trajectory density $p_{\theta}(\mathbf{x}_{0:N}) = p(\mathbf{x}_N) \prod_{n=1}^N p_{\theta}(\mathbf{x}_{n-1} | \mathbf{x}_n)$. Both the learned reverse-time transition $p_{\theta}(\mathbf{x}_{n-1} | \mathbf{x}_n)$ and the reference forward kernel $q(\mathbf{x}_n | \mathbf{x}_{n-1})$ are Gaussian; their closed-form expressions are provided in Appendix A.3. Alternatively, one may solve the probability-flow ODE equation 3 with any off-the-shelf ODE solver to obtain samples at $t = \epsilon$.

3.2 E(3)–Equivariant Score-Based Diffusion Models

Standard diffusion models ignore the geometric symmetries that characterise many structured datasets. In molecular systems, for instance, the potential-energy surface is invariant under global *translations*, *rotations*, and *reflections*—the Euclidean group E(3)—as well as permutations of identical atoms. Encoding these symmetries explicitly is critical for producing physically valid and statistically efficient samplers.

Equivariant backbone. Hoogeboom *et al.* (Hoogeboom et al., 2022) introduced E(3)-*equivariant diffusion models* by replacing the standard score network with an Equivariant Graph Neural Network (EGNN) (Satorras et al., 2021). Following the analysis of Xu *et al.* (Xu et al., 2022), a generative process that is equivariant to E(3) transformations must satisfy two conditions:

1. **Rotation and reflection invariance.** For any orthogonal transformation $R \in \text{O}(3)$,

$$p(\mathbf{x}_N) = p(R\mathbf{x}_N), \quad p(\mathbf{x}_{n-1} | \mathbf{x}_n) = p(R\mathbf{x}_{n-1} | R\mathbf{x}_n),$$

which implies that the marginal $p(\mathbf{x}_n)$ is invariant at every time step.

2. **Translation invariance.** Perfect translation invariance cannot hold for non-degenerate continuous densities; instead, it is imposed by restricting all configurations to the *centre-of-mass (CoM) subspace* $\sum_i \mathbf{x}_i = \mathbf{0}$, i.e. every sample is translated so that its CoM is at the origin.

Practical implementation. These constraints are enforced by replacing ordinary Gaussians $\mathcal{N}(\cdot)$ with Gaussians projected onto the zero-CoM subspace, denoted $\mathcal{N}_x(\cdot)$ (sampling and density evaluation details are given in Appendix B). Sampling proceeds exactly as in Sec. 3.1: we integrate either the reverse SDE or the probability-flow ODE, but all intermediate states now live in the zero-CoM subspace and all conditional kernels are the projected Gaussians \mathcal{N}_x . The resulting pipeline produces molecular configurations that are E(3)-equivariant by construction.

3.3 Diffusion Models with Importance Sampling

Assume that the (unnormalized) target density $\pi_{\text{data}}(\cdot)$ can be evaluated point-wise, whereas drawing i.i.d. samples is infeasible. Directly using raw outputs from a pretrained diffusion model therefore yields biased estimates of expectations $\mathbb{E}_{\pi_{\text{data}}}[\phi(\mathbf{x})]$. IS removes this bias through the weight $\pi_{\text{data}}(\mathbf{x})/p_0(\mathbf{x})$, where p_0 denotes the diffusion sampler’s marginal at $t = 0$. Unfortunately, p_0 is typically intractable, and any approximation re-introduces bias. Because we may generate samples either by integrating the reverse SDE or by solving the probability-flow ODE, two corresponding IS strategies arise.

Reverse SDE pathway. When sampling with a DDPM-type discretization of the reverse SDE, define

$$q(\mathbf{x}_{0:N}) = \pi_{\text{data}}(\mathbf{x}_0) \prod_{n=0}^{N-1} q(\mathbf{x}_{n+1} | \mathbf{x}_n), \quad p_{\theta}(\mathbf{x}_{0:N}) = p(\mathbf{x}_N) \prod_{n=0}^{N-1} p_{\theta}(\mathbf{x}_n | \mathbf{x}_{n+1}), \quad (4)$$

where both kernels are Gaussian and therefore tractable. The trajectory weight $w = \frac{q(\mathbf{x}_{0:N})}{p_{\theta}(\mathbf{x}_{0:N})}$ yields an unbiased estimator (see Appendix A.3). In practice, however, the mismatch introduced by a finite time discretization causes the effective sample size (ESS) to collapse.

PF ODE pathway. If instead we solve the probability-flow ODE equation 3, the likelihood of the terminal sample can be written as (Chen et al., 2018; Song et al., 2020)

$$\log p_0(\mathbf{x}(0)) = \log p_T(\mathbf{x}(T)) + \int_0^T \nabla_{\mathbf{x}} \cdot (-t s_{\theta}(\mathbf{x}(t), t)) dt. \quad (5)$$

where s_{θ} is the score model. Computing the divergence term exactly is prohibitively expensive and is typically approximated with the Hutchinson–Skilling trace estimator (Skilling, 1989; Hutchinson, 1989). Such stochastic estimates, together with time-discretisation error, re-introduce bias and defeat the purpose of IS.

Limitations. Although both pathways are theoretically sound, each faces practical obstacles: the reverse-SDE approach yields vanishing ESS, whereas the ODE approach incurs excessive compute or bias. In the next section we address the first limitation by introducing a covariance-adjusted scheme that dramatically increases ESS in practice.

4 Method

In this section, we introduce the post-training method VT-DIS to improve ESS for diffusion models with importance sampling. We will first give an overview of the whole algorithm and explain each part of the algorithm in more details and the rationale behind them, then we will explain how our method can also be applied to E(3)-Equivariant diffusion models for molecular generations.

4.1 Overview

We build on a pretrained score-based diffusion model, freezing its score network s_{θ} , and introduce a lightweight post-training step that adjusts only the per-step noise covariance $\{\Sigma_{\phi}(n)\}_{n=1}^N$, with ϕ being learnable parameters. The objective is to make the learned reverse-time process better agree with the true forward (noising) process by minimizing the $\alpha = 2$ divergence, which directly controls the variance of importance weights and hence the ESS. In practice, we: (1) Simulate target trajectories $\{\mathbf{x}_{0:N}\}$ from the forward SDE; (2) Compute the log-density ratio between the forward and reverse joint distributions given current $\{\Sigma_{\phi}(n)\}_{n=1}^N$; (3) Update the covariance parameters by minimizing the Monte Carlo estimate of the $\alpha = 2$ divergence; (4) At inference, sample with the tuned covariances and assign each trajectory a single importance weight. We summarize the whole process in Algorithm 1.

4.2 Covariance Parameterization

We explore three complementary ways to parameterize the per-step covariance $\Sigma_{\phi}(n)$.

Algorithm 1 Post-training Covariance Tuning (VT-DIS)**Require:** Score network s_θ , initial covariances $\{\Sigma_\phi\}$, time grid $\{t_n\}_{n=0}^N$, batch size M **Ensure:** Tuned covariances $\{\Sigma_{\phi^*}\}$

```

1: repeat
2:   Sample  $\{\mathbf{x}_{0:N}^{(i)}\}_{i=1}^M \sim q(\mathbf{x}_{0:N})$ 
3:    $\ell^{(i)}(\phi) \leftarrow \log q(\mathbf{x}_{0:N}^{(i)}) - \log p_{\{\Sigma_\phi\}}(\mathbf{x}_{0:N}^{(i)})$ 
4:    $L(\phi) \leftarrow \frac{1}{M} \sum_{i=1}^M \exp(\ell^{(i)}(\phi))$ 
5:    $\{\Sigma_\phi\} \leftarrow \{\Sigma_\phi\} - \eta \nabla_\phi L(\phi)$ 
6: until convergence
7:  $\{\Sigma_{\phi^*}\} \leftarrow \{\Sigma_\phi\}$ 

```

Isotropic Covariance The standard DDPM proposal $p_\theta(\mathbf{x}_{n-1}|\mathbf{x}_n)$ uses an isotropic variance $\Sigma_n = \sigma_{\text{ddpm}}^2(n)\mathbf{I}$ (where $\sigma_{\text{ddpm}}^2(n)$ is defined in Appendix A.3, equation 21). To tune this, for each $n = 1, \dots, N-1$ we introduce a nonnegative scalar η_n so that

$$\Sigma_\phi(n) = \eta_n \sigma_{\text{ddpm}}^2(n) \mathbf{I} \quad (6)$$

with $\eta_n \geq 0$ enforced (e.g. via **SoftPlus**) and initialized at $\eta_n = 1$ to recover the standard DDPM baseline. We use ϕ to represent all tunable parameters η_n for each n .

Structured covariances To capture per-coordinate anisotropy or correlations, we replace the scalar multiplier with vector or matrix parameters. A diagonal form is

$$\Sigma_\phi(n) = \sigma_{\text{ddpm}}^2(n) \text{diag}(\eta_{n,1}, \dots, \eta_{n,d}), \quad (7)$$

with $\eta_{n,i} \geq 0$ and initialized with each $\eta_{n,i} = 1$ so that baseline is recovered. A fully correlated covariance uses a learnable matrix $L_n \in \mathbb{R}^{d \times d}$: $\Sigma_\phi(n) = \sigma_{\text{ddpm}}^2(n) L_n L_n^\top$, with L_n initialized to the identity to recover the baseline. To limit parameter growth which makes optimization hard especially when number of time steps is large, we further propose a practical low-rank plus isotropic decomposition:

$$\Sigma_\phi(n) = \sigma_{\text{ddpm}}^2(n) (A_n A_n^\top + \alpha_n \mathbf{I}), \quad (8)$$

where $A_n \in \mathbb{R}^{d \times k}$ where $k < d$ is chosen depending on the computing budget and $\alpha_n \geq 0$ are learned, and $A_n A_n^\top + \alpha_n \mathbf{I}$ is initialized to \mathbf{I} by initializing A_n to be zero matrix and $\alpha_n = 1$.

Sample-dependent models In the case that the number of time steps is large and even the low-rank version (equation 8) becomes infeasible, we can use a neural network with input t_n to predict the optimal covariance matrix at each step. Additionally, for maximal flexibility, $\Sigma_\phi(n)$ can be made a function of the current state \mathbf{x}_n (and t_n) via a neural network. That is:

$$\Sigma_\phi(n) = \text{NN}_\phi(t_n) \quad \text{or} \quad \Sigma_\phi(n) = \text{NN}_\phi(\mathbf{x}_n, t_n)$$

Although this state-dependent form remains valid, our experiments did not show significant ESS gains over the time-only or time-and-sample dependent parameterizations and the additional function evaluation to get covariance increases the computational cost during inference of diffusion models. Therefore, we defer its exploration to future work.

4.3 Extension to E(3)-Equivariant Diffusion

Our VT-DIS covariance tuning readily extends to E(3)-equivariant diffusion models: the only modification is that all trajectories lie in the zero-center-of-mass (CoM) subspace. Isotropic covariances require no change, since it commutes with rotations, reflections, and CoM projection.

For diagonal or fully correlated covariances, the following form of covariance matrix will satisfy E(3)-equivariance and permutation equivariance:

$$\Sigma_\phi(n) = B_\phi(n) \otimes I_3, \quad (9)$$

where $B_\phi(n) \in \mathbb{R}^{M \times M}$ acts on the M particles, I_3 on the three Cartesian coordinates and \otimes denotes the Kronecker product. The matrices $B_\phi(n)$ are learned per step. This form automatically satisfies E(3)-equivariance (proof in Appendix B). To enforce permutation symmetry, $B_\phi(n)$ must respect the specific particle-identities of the dataset. We illustrate two representative cases below.

Many-particle systems All particles are identical, so a diagonal $B_\phi(n)$ reduces to an isotropic covariance and adds no flexibility. Instead, we parameterize the full covariance

$$B_\phi(n) = (b_n - a_n)\mathbf{I} + a_n\mathbf{1}\mathbf{1}^\top$$

where a_n and b_n are learnable parameters with $b_n > 0$ and $-\frac{b_n}{M-1} < a_n < b_n$ to ensure $B_\phi(n) \succ 0$. Appendix B shows that this choice satisfies both E(3)-equivariance and full permutation invariance, as needed for systems such as n -body systems.

Molecular distributions. Real-world molecules introduce chemical symmetries that are more subtle than the “all-particles-indistinguishable” setting of the above section. For instance, in alanine dipeptide only the three hydrogens bound to the same carbon atom may be permuted freely; the remaining atoms are fixed. To encode such partial symmetries we assign every atom a symmetry label $L_i \in \{1, \dots, K\}$, $i = 1, \dots, M$, where atoms sharing the same label are interchangeable while atoms with different labels are not. The covariance $B_\phi(n)$ at sampling step n is then constrained to depend on the pair (i, j) *only* through the labels (L_i, L_j) . Proof in Appendix B.3.

1. **Diagonal (per-class variance).** Let $\eta_{n,\ell} \geq 0$ be a learnable variance for class ℓ .

$$B_\phi(n) = \text{diag}(\eta_{n,L_1}, \eta_{n,L_2}, \dots, \eta_{n,L_M}).$$

Positive definiteness is automatic when every $\eta_{n,\ell} > 0$.

2. **Block-structured (full class-class covariance).** For any learnable matrix $A_n \in \mathbb{R}^{K \times K}$ and a learnable scalar $\alpha_n > 0$. Define

$$[B_\phi(n)]_{ij} = (A_n A_n^\top)_{L_i L_j} + \alpha_n \delta_{ij},$$

where δ_{ij} is the Kronecker delta. The added diagonal term $\alpha_n I$ guarantees $B_\phi(n) \succ 0$.

4.4 Objective Function

Our aim is to align the reverse-process joint proposal $p_\theta(\mathbf{x}_{0:N})$ with the forward-diffusion joint $q(\mathbf{x}_{0:N})$ (see equation 4), in a manner that directly improves importance-sampling performance. Although one could minimize the Kullback–Leibler (KL) divergence $D_{\text{KL}}(q\|p_\phi)$ (and indeed closed-form optimal covariances under this criterion have been derived (Ou et al., 2024)), such minimization does not control tail mismatch. Poor alignment in the tails leads to high variance of importance weights and thus low ESS when sample size is large. Instead, we minimize the α -divergence with $\alpha = 2$, which is proportional to the variance of importance weights:

$$D_2(q\|p_\phi) = \int \frac{q(\mathbf{x}_{0:N})}{p_\phi(\mathbf{x}_{0:N})} q(\mathbf{x}_{0:N}) d\mathbf{x}_{0:N} \approx \frac{1}{M} \sum_{i=1}^M \frac{q(\mathbf{x}_{0:N}^{(i)})}{p_\phi(\mathbf{x}_{0:N}^{(i)})}, \quad \mathbf{x}_{0:N}^{(i)} \sim q \quad (10)$$

In practice, we draw M trajectories $\{\mathbf{x}_{0:N}^{(i)}\}$ from the forward process¹ q to form the unbiased Monte Carlo estimate as shown above. Optimizing this loss by gradient descent yields covariance parameters that minimize importance weight variance and thereby maximize ESS.

¹Empirically, we found that attempting to estimate D_2 with samples from the learned reverse process p_θ often induces mode collapse and degrades forward ESS.

4.5 Choice of Sampling Schedule

Karras et al. (2022) propose a family of continuous-time schedules controlled by a parameter ρ :

$$t_n = \left(t_N^{1/\rho} + \frac{n}{N-1} (t_0^{1/\rho} - t_N^{1/\rho}) \right)^\rho, \quad t_0 = \epsilon, \quad t_N = T.$$

In the limit $\rho \rightarrow \infty$, this reduces to a fixed geometric grid $t_n = t_0 \left(t_N / t_0 \right)^{n/(N-1)}$, which we found consistently yields lower $\alpha = 2$ divergence and higher ESS across all benchmarks. We also experimented with jointly tuning the schedule and per-step covariances during VT-DIS. While joint tuning improves performance on simple Gaussian mixtures, it proved unstable on molecular datasets. We defer a robust treatment of schedule learning to future work.

4.6 Complexity & Inference-Time Cost

Relative to a standard DDPM equipped with importance sampling, VT-DIS adds only a *single, offline* optimization step: tuning the per-step noise covariances $\{\Sigma_\phi(n)\}$. After this one-time cost, generate samples and importance weights only adds negligible overhead to standard sampling from diffusion models—no extra function evaluations. Unlike flow-based Boltzmann generators or diffusion generate samples by solving PF ODE, VT-DIS never has to integrate an ODE to evaluate marginal proposal log-densities, a procedure that becomes prohibitively expensive in high-dimensional settings.

5 Experiments

Experimental setup. We assume access to both the energy function and the training data used to fit the diffusion model (same assumption as in, e.g., Klein et al. (2023); Tan et al. (2025)). We first train a diffusion model with an EGNN backbone (Satorras et al., 2021) and EDM preconditioning (Karras et al., 2022) using a standard training procedure. We then apply VT-DIS to the same training set to optimize the proposal covariances. For evaluation, we compare VT-DIS with a DDPM+IS baseline under a fixed compute budget measured by the number of function evaluations (NFEs; i.e., sampling steps). We also compare against importance sampling with exact likelihoods computed via the probability-flow ODE (PF-ODE). Performance is reported as both forward and reverse effective sample size (ESS).

We consider three representative classes of Boltzmann systems: (i) **Synthetic**: a high-dimensional two-mode Gaussian mixture (GMM-2); (ii) **Many-particle**: double-well (DW-4) and Lennard-Jones-13 (LJ-13) clusters (Köhler et al., 2020); and (iii) **Small-molecule**: alanine dipeptide (ALDP). For all targets, the (unnormalized) density π_{data} is available, enabling exact importance weights. We use the atomic coordinate representation for each sample and train our models on these coordinates. Further implementation details are provided in Appendix C.

5.1 GMM-2

We begin with a two-component isotropic Gaussian mixture in \mathbb{R}^d :

$$\pi_{\text{data}}(\mathbf{x}) = w_1 \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_1, \sigma_1^2 \mathbf{I}) + w_2 \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_2, \sigma_2^2 \mathbf{I}),$$

where the means, variances and weights are defined as $\boldsymbol{\mu}_1 = 1$, $\boldsymbol{\mu}_2 = -2$, $\sigma_1^2 = \sigma_2^2 = 0.15$, and $w_1 = \frac{2}{3}$, $w_2 = \frac{1}{3}$. We evaluate two ambient dimensions, $d \in \{50, 100\}$. Figure 1 shows forward and reverse ESS as a function of the number of function evaluations (NFE) for VT-DIS with three noise-covariance parameterizations—*isotropic*, *diagonal*, and *full rank*—together with the DDPM + IS baseline.

The plots reveal three key trends: (i) VT-DIS consistently outperforms the baseline in forward ESS, reverse ESS, or both, for every covariance parameterization. (ii) The full-rank covariance yields the highest forward and reverse ESS for both values of d . For $d = 100$, however, the forward-ESS error bars are wide when $d = 100$, indicating residual mismatch between the forward and reverse processes at very high dimension. (iii) With isotropic and diagonal covariances, reverse ESS improves as NFE grows, whereas forward ESS stays low, underscoring the importance of modeling cross-dimensional correlations in this setting.

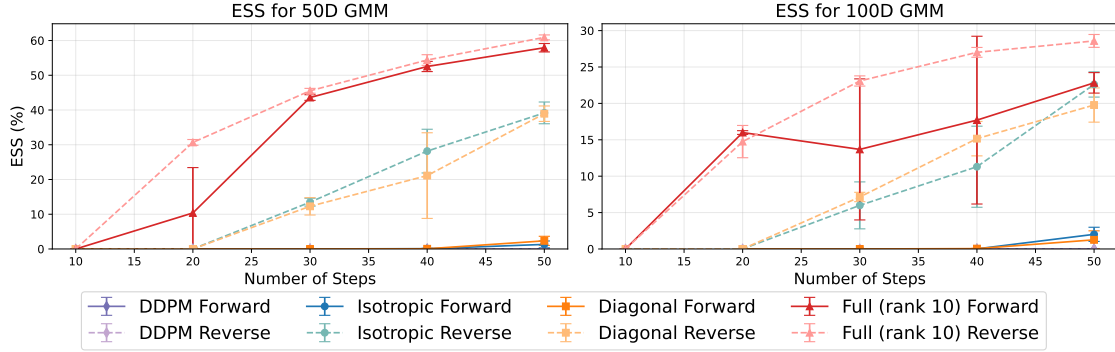


Figure 1: **GMM-2**: ESS versus NFEs for dimensions $d = 50$ and $d = 100$, comparing VT-DIS with the DDPM + IS baseline. ESS estimates are computed using 10^5 samples.

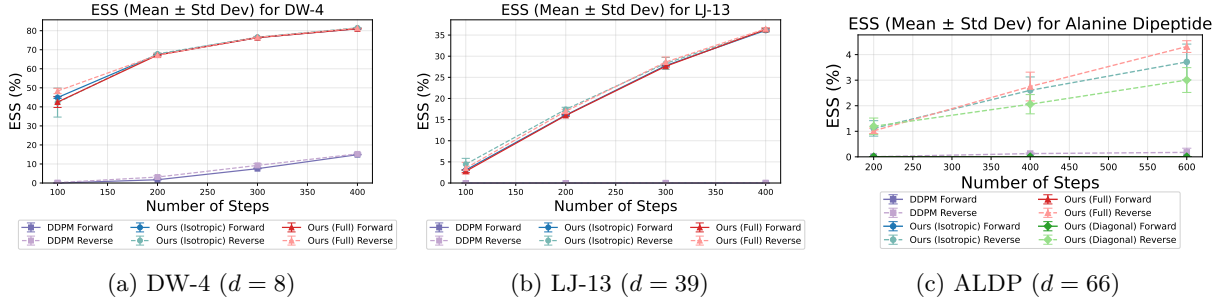


Figure 2: **ESS versus NFEs**. VT-DIS consistently outperforms the standard DDPM + IS baseline across all tasks under the same computational budget. ESS values are estimated using 10^5 samples.

5.2 Many-particle systems: DW-4 and LJ-13

We next evaluate VT-DIS on two prototypical many-body systems with pair-wise interactions: *DW-4* (double-well potential, four particles in \mathbb{R}^2 , $d = 8$) and *LJ-13* (Lennard-Jones cluster, thirteen particles in \mathbb{R}^3 , $d = 39$) (Köhler et al., 2020).

ESS versus NFEs. Figure 2 reports forward and reverse ESS as a function of the number of NFEs. On the low-dimensional DW-4 system, VT-DIS raises both ESS measures from $\leq 1\%$ to $\approx 40\%$ using only 100 NFEs. The improvement is even more pronounced on the higher-dimensional LJ-13 cluster: the standard DDPM + IS baseline yields virtually 0% ESS up to 400 NFEs, whereas VT-DIS achieves $\approx 35\%$ ESS on both forward and reverse measures with the same budget. Due to the permutation-equivariant constraint of the many-particle system, the full covariance matrix has only one additional degree of freedom compared to an isotropic covariance. In this setting, we observe no significant ESS gain from using the full covariance, although both forward and reverse ESS exhibit lower standard deviations than with isotropic noise. Across both tasks, ESS increases monotonically with the number of NFEs, facilitating a clear accuracy-cost trade-off.

Bias correction. To illustrate bias correction we compare three energy histograms for each task: (i) *ground-truth* samples from the test dataset, (ii) *unweighted* VT-DIS generated samples, and (iii) *reweighted* samples. Figure 3 shows that the raw proposals are biased, but importance reweighting realigns the distributions with the target, corroborating the ESS improvements.

5.3 Small Molecule: Alanine Dipeptide (ALDP)

We evaluate VT-DIS on the 66-dimensional Boltzmann distribution of alanine dipeptide (ALDP) in implicit solvent at temperature $T = 800$ K. The unnormalized Boltzmann density is defined by a classical force-field

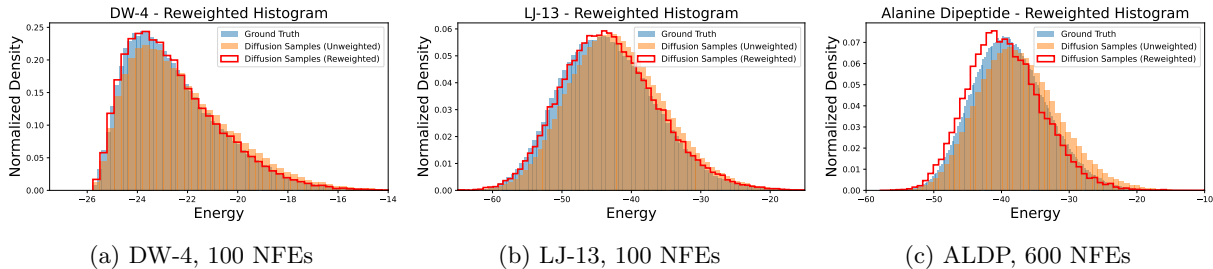


Figure 3: **Energy distributions.** Unweighted samples generated by VT-DIS with isotropic covariance (orange) are biased relative to ground-truth samples (blue), but importance-weighted samples (red) effectively correct the discrepancy. Histograms are plotted using 5×10^5 samples.

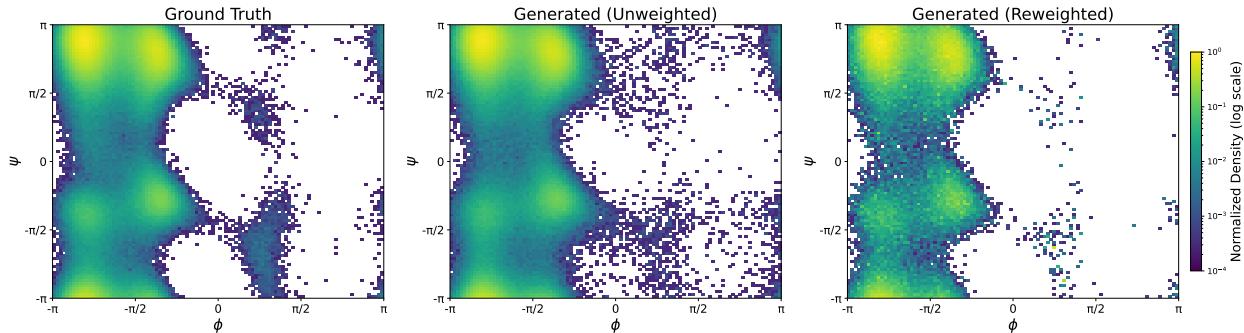


Figure 4: Ramachandran plots from left to right: ground-truth samples generated by MD, VT-DIS (isotropic) samples before and after reweighting. All plots are generated using 10^6 samples.

potential, and the availability of its exact log density enables unbiased importance-sampling evaluation. Training and implementation details are provided in Appendix C.

Results. Figure 2c shows forward and reverse ESS over 200–600 NFEs. VT-DIS consistently improves reverse ESS, reaching approximately 2.5% at 400 NFEs compared to under 0.5% for the baseline. Again, we do not observe a significant improvement in ESS when using diagonal or full covariance parameterizations. This may be because the atomic coordinates exhibit weak correlations, making an isotropic covariance sufficient.

Figure 3c displays the energy histogram, and Figure 4 presents the Ramachandran plot. Both demonstrate that importance sampling effectively corrects bias in the generated samples. Notably, the molecular dynamics (MD) “ground-truth” samples exhibit a slight bias, as the reweighted energy histogram shifts toward lower energies than the MD samples.

Forward ESS remains near zero for both VT-DIS and the baseline, indicating that neither proposal distribution fully covers the joint target. This was also observed with flow-based samplers (Midgley et al., 2023), which we leave for future investigation.

5.4 Computational Cost versus Probability–Flow ODE

While VT-DIS and PF-ODE + IS both yield unbiased estimators, they differ substantially in computational efficiency. Here, we compare the ESS and computational cost of each method.

Experimental setup. To isolate inference cost, we reuse the *same* pretrained score network and employ an identical number of integration steps for both methods. Computational effort is measured in *GFLOPs per sample*, and accuracy is quantified by forward and reverse ESS. Results for DW-4, LJ-13, and ALDP are summarized in Table 1.

Table 1: Performance on DW-4, LJ-13 and ALDP with 10^4 samples on a single A100. The number in parentheses after each dataset denotes the sampling steps. \uparrow = higher is better, \downarrow = lower is better.

	DW-4 (100)		LJ-13 (200)		ALDP (200)	
	Rev. SDE	PF ODE	Rev. SDE	PF ODE	Rev. SDE	PF ODE
Forward ESS (%) \uparrow	44.32	94.29	15.87	79.21	≈ 0	34.81
Reverse ESS (%) \uparrow	46.76	83.60	17.69	77.42	1.05	35.39
GFlops / sample \downarrow	0.81	7.70	26.2	2173	33.7	45967

Results. PF-ODE + IS attains higher ESS but incurs a steep computational penalty: its GFLOPs per sample are one to two orders of magnitude greater than VT-DIS’s, and this gap widens with dimensionality. In practice, PF-ODE requires a *minimum* number of ODE steps to control proposal-density bias; each additional step further increases cost. VT-DIS, in contrast, trades some ESS for dramatically lower compute. Its estimator is asymptotically unbiased as the sample number increases, regardless of the time-step budget, and each additional reverse-SDE step directly boosts ESS.

6 Conclusion and Limitations

Contributions. We have introduced VT-DIS, a post-training procedure that can be applied to any pretrained diffusion model, and demonstrated its effectiveness in molecular generation. By computing trajectory-wise importance weights, VT-DIS achieves high effective sample size (ESS) across multiple benchmarks while using only a fraction of the computational budget required to solve the PF ODE for exact likelihood estimation. Furthermore, VT-DIS provides asymptotically unbiased expectation estimates as the number of samples grows—independent of the number of sampling steps. In contrast, PF-ODE solvers guarantee asymptotically unbiased sampling only when both the number of samples and the number of discretization steps increase, which limits their practical utility. These findings highlight VT-DIS’s potential for diffusion-based Boltzmann generators.

Limitations and Future Work. Despite its advantages, VT-DIS has several limitations. First, the post-training procedure must be repeated for each sampling schedule, which can incur significant computational and memory costs. Figure 5 in Appendix D.1 shows a clear pattern in optimal hyperparameters across different step counts; future work could leverage this pattern to develop a unified post-training algorithm that supports arbitrary schedules in one stage. Second, optimizing the sampling schedule itself to minimize the α -divergence remains an open question; existing methods from (Sabour et al., 2024; Williams et al., 2024) may offer a promising starting point.

References

- Tara Akhoun-Sadegh, Jarrod Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, et al. Iterated denoising energy matching for sampling from boltzmann densities. *arXiv preprint arXiv:2402.06121*, 2024.
- Fan Bao, Chongxuan Li, Jiacheng Sun, Jun Zhu, and Bo Zhang. Estimating the optimal covariance with imperfect mean in diffusion probabilistic models. *arXiv preprint arXiv:2206.07309*, 2022.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Daan Frenkel and Berend Smit. *Understanding molecular simulation: from algorithms to applications*. Elsevier, 2023.
- Victor Garcia Satorras, Emiel Hoogeboom, Fabian Fuchs, Ingmar Posner, and Max Welling. E (n) equivariant normalizing flows. *Advances in Neural Information Processing Systems*, 34:4181–4192, 2021.

- Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.
- Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, and David Duvenaud. Scalable reversible generative models with free-form continuous dynamics. In *International Conference on Learning Representations*, volume 3, 2019.
- W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970.
- Jiajun He, Yuanqi Du, Francisco Vargas, Dinghuai Zhang, Shreyas Padhy, RuiKang OuYang, Carla Gomes, and José Miguel Hernández-Lobato. No trick, no treat: Pursuits and challenges towards simulation-free training of neural samplers. *arXiv preprint arXiv:2502.06685*, 2025.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Emiel Hoogetboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, pp. 8867–8887. PMLR, 2022.
- Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.
- Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. *Advances in Neural Information Processing Systems*, 36:59886–59910, 2023.
- Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for symmetric densities. In *International conference on machine learning*, pp. 5361–5370. PMLR, 2020.
- 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 Midgley, Vincent Stimper, Javier Antorán, Emile Mathieu, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Se (3) equivariant augmented coupling flows. *Advances in Neural Information Processing Systems*, 36:79200–79225, 2023.
- Laurence Illing Midgley, Vincent Stimper, Gregor NC Simm, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Flow annealed importance sampling bootstrap. *arXiv preprint arXiv:2208.01893*, 2022.
- 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.
- Zijing Ou, Mingtian Zhang, Andi Zhang, Tim Z Xiao, Yingzhen Li, and David Barber. Improving probabilistic diffusion models with optimal covariance matching. *arXiv preprint arXiv:2406.10808*, 2024.
- Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. *arXiv preprint arXiv:2307.01198*, 2023.
- Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your steps: Optimizing sampling schedules in diffusion models. *arXiv preprint arXiv:2404.14507*, 2024.
- Victor Garcia Satorras, Emiel Hoogetboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.

- John Skilling. The eigenvalues of mega-dimensional matrices. *Maximum Entropy and Bayesian Methods: Cambridge, England, 1988*, pp. 455–466, 1989.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. PMLR, 2015.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- Charlie B Tan, Avishek Joey Bose, Chen Lin, Leon Klein, Michael M Bronstein, and Alexander Tong. Scalable equilibrium sampling with sequential boltzmann generators. *arXiv preprint arXiv:2502.18462*, 2025.
- Francisco Vargas, Will Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. *arXiv preprint arXiv:2302.13834*, 2023a.
- Francisco Vargas, Shreyas Padhy, Denis Blessing, and Nikolas Nüsken. Transport meets variational inference: Controlled monte carlo diffusions. *arXiv preprint arXiv:2307.01050*, 2023b.
- Christopher Williams, Andrew Campbell, Arnaud Doucet, and Saifuddin Syed. Score-optimal diffusion schedules. *arXiv preprint arXiv:2412.07877*, 2024.
- Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022.
- Yilun Xu, Shangyuan Tong, and Tommi Jaakkola. Stable target field for reduced variance score estimation in diffusion models. *arXiv preprint arXiv:2302.00670*, 2023.
- Qinsheng Zhang and Yongxin Chen. Path integral sampler: a stochastic control approach for sampling. *arXiv preprint arXiv:2111.15141*, 2021.

A Score-based Diffusion Models with Importance Sampling

In this appendix, we present the mathematical derivations for combining score-based diffusion models with importance sampling. We first review the fundamentals of importance sampling and then show how it can be integrated into score-based diffusion to correct bias in generated samples.

A.1 Review of Importance Sampling

Importance sampling (IS) provides an estimator for expectations of the form

$$\mathbb{E}_{\mathbf{x} \sim \pi_{\text{data}}}[\phi(\mathbf{x})] = \int \phi(\mathbf{x}) \pi_{\text{data}}(\mathbf{x}) d\mathbf{x},$$

where π_{data} is the (possibly unnormalized) target density and ϕ is a test function. Directly sampling from π_{data} may be infeasible, but we assume that we can evaluate its density up to a normalization constant. Instead, we draw samples $\{\mathbf{x}^{(n)}\}_{n=1}^N$ from a proposal distribution $p(\mathbf{x})$ and form the self-normalized estimator

$$\mathbb{E}_{\mathbf{x} \sim \pi_{\text{data}}}[\phi(\mathbf{x})] \approx \sum_{n=1}^N \bar{w}_n \phi(\mathbf{x}^{(n)}),$$

where

$$w_n = \frac{\pi_{\text{data}}(\mathbf{x}^{(n)})}{p(\mathbf{x}^{(n)})}, \quad \bar{w}_n = \frac{w_n}{\sum_{m=1}^N w_m}.$$

This self-normalized importance sampling (SNIS) estimator is consistent and asymptotically unbiased, with its variance decreasing as N increases.

A.2 Effective Sample Size

The quality of an IS estimator is often measured by its effective sample size (ESS). Two common definitions are:

- *Forward ESS*, which evaluates the overlap with the target distribution:

$$\widehat{\text{ESS}}_{\text{fwd}} = \frac{N^2}{Z^{-1} \sum_{i=1}^N w(\mathbf{x}^{(i)})}, \quad \mathbf{x}^{(i)} \sim \pi_{\text{data}}, \quad (11)$$

where $Z^{-1} \approx \frac{1}{N} \sum_{n=1}^N \frac{1}{w_n}$ is the inverse normalizing constant.

- *Reverse ESS*, which measures the concentration of weights under the proposal:

$$\widehat{\text{ESS}}_{\text{rev}} = \frac{(\sum_{i=1}^N w(\mathbf{y}^{(i)}))^2}{\sum_{i=1}^N w(\mathbf{y}^{(i)})^2}, \quad \mathbf{y}^{(i)} \sim p. \quad (12)$$

While reverse ESS is widely used for evaluation, it can overestimate performance if p suffers from mode collapse. Forward ESS, by contrast, tends to provide a more reliable measure of overlap between the target and proposal.

A.3 Combining Score-based Diffusion Models with Importance Sampling

We now consider task of estimating the integral $\mathbb{E}_{\mathbf{x} \sim \pi_{\text{data}}}[\phi(\mathbf{x})]$ where normalized or unnormalized π_{data} is the target density which we can evaluate, and ϕ is the test function of interest. We aim to train a diffusion models to act as the proposal distribution which we can draw samples from. Suppose we have a time interval $[\epsilon, T]$ and we discretize the time interval into N sections with $N+1$ time steps such that $\epsilon = t_0 < t_1 < \dots < t_N = T$.

Assume the conditional proposal distribution of the diffusion model is $p_{\theta}(\mathbf{x}_{n-1}|\mathbf{x}_n)$, and the conditional noise distribution is $q(\mathbf{x}_n|\mathbf{x}_{n-1})$ for $n = 1, \dots, N$. Using SNIS, the integral can be estimated as

$$\mathbb{E}_{\mathbf{x} \sim \pi}[\phi(\mathbf{x})] \approx \mathbb{E}_{\mathbf{x}_0 \sim \pi}[\phi(\mathbf{x}_0)] \quad (13)$$

$$= \int \phi(\mathbf{x}_0) \pi(\mathbf{x}_0) d\mathbf{x}_0 \quad (14)$$

$$= \int \phi(\mathbf{x}_0) \pi(\mathbf{x}_0) \left(\prod_{n=1}^N q(\mathbf{x}_n|\mathbf{x}_{n-1}) \right) d\mathbf{x}_{0:N} \quad (15)$$

$$= \int \phi(\mathbf{x}_0) \left(\frac{\pi(\mathbf{x}_0) \prod_{n=1}^N q(\mathbf{x}_n|\mathbf{x}_{n-1})}{p_{\theta}(\mathbf{x}_N) \prod_{n=1}^N p_{\theta}(\mathbf{x}_{n-1}|\mathbf{x}_n)} \right) p_{\theta}(\mathbf{x}_{0:N}) d\mathbf{x}_{0:N} \quad (16)$$

$$\approx \frac{1}{K} \sum_{k=1}^K \underbrace{\left(\frac{\pi(\mathbf{x}_{t_0}^{(k)}) \prod_{n=1}^N q(\mathbf{x}_{t_n}^{(k)}|\mathbf{x}_{t_{n-1}}^{(k)})}{p_{\theta}(\mathbf{x}_N^{(k)}) \prod_{n=1}^N p_{\theta}(\mathbf{x}_{t_{n-1}}^{(k)}|\mathbf{x}_{t_n}^{(k)})} \right)}_{w_k} \phi(\mathbf{x}_0) \quad (17)$$

where $\mathbf{x}_{0:N}^{(k)}$ for $k = 1, \dots, K$ are samples from the joint proposal distribution p_{θ} and w_k for $k = 1, \dots, K$ are the importance weights. Note that ϵ is chosen such that the error for the first approximation is negligible.

We now need to specify the exact forms of $q(\mathbf{x}_n|\mathbf{x}_{n-1})$ and $p_{\theta}(\mathbf{x}_{n-1}|\mathbf{x}_n)$. Recall the forward SDE with zero drift ($\mathbf{f} \equiv \mathbf{0}$) and diffusion coefficient $g(t) = \sqrt{2t}$:

$$d\mathbf{x}_t = \sqrt{2t} d\mathbf{w}_t.$$

Given a time discretization $\epsilon = t_0 < t_1 < \dots < t_N = T$, the conditional *target* distribution of the forward process is Gaussian:

$$p(\mathbf{x}_n | \mathbf{x}_{n-1}) = \mathcal{N}(\mathbf{x}_n; \mathbf{x}_{n-1}, (t_n^2 - t_{n-1}^2) \mathbf{I}), \quad t_n > t_{n-1}. \quad (18)$$

By Bayes' rule, the DDPM-style *reverse* (proposal) distribution conditional on \mathbf{x}_n and \mathbf{x}_0 is

$$\begin{aligned} p(\mathbf{x}_{n-1} | \mathbf{x}_n, \mathbf{x}_0) &\propto p(\mathbf{x}_n | \mathbf{x}_{n-1}) p(\mathbf{x}_{n-1} | \mathbf{x}_0) \\ &= \mathcal{N}(\mathbf{x}_{n-1}; \boldsymbol{\mu}(\mathbf{x}_n, \mathbf{x}_0), \sigma_{\text{ddpm}}^2(n) \mathbf{I}), \end{aligned} \quad (19)$$

where

$$\boldsymbol{\mu}(\mathbf{x}_n, \mathbf{x}_0) = \frac{t_n^2}{t_{n-1}^2} \mathbf{x}_n + \left(1 - \frac{t_n^2}{t_{n-1}^2}\right) \mathbf{x}_0, \quad (20)$$

$$\sigma_{\text{ddpm}}^2(n) = \frac{t_{n-1}^2 (t_n^2 - t_{n-1}^2)}{t_n^2}. \quad (21)$$

In practice, \mathbf{x}_0 will be predicted via a neural network and thus the reverse conditional is defined as $p_{\theta}(\mathbf{x}_{n-1}|\mathbf{x}_n) = p(\mathbf{x}_{n-1}|\mathbf{x}_n, \hat{\mathbf{x}}_0(\mathbf{x}_n, t_n; \boldsymbol{\theta}))$. Since both forward and reverse conditionals are Gaussian with closed-form parameters, the joint density ratio over trajectories can be evaluated exactly with minimal overhead.

B VT-DIS for E(3)-Equivariant Diffusion Models

In this appendix we detail how the proposed velocity-tempered diffusion importance sampling (VT-DIS) post-training procedure can be integrated with E(3)-equivariant diffusion models.

Notation. Let $M \geq 2$ denote the number of particles and let n be the spatial dimension. Concatenate the coordinates of all particles into the column vector

$$\mathbf{x} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)}) \in \mathbb{R}^{Mn}, \quad \mathbf{x}^{(i)} \in \mathbb{R}^n.$$

Unless stated otherwise we restrict our attention to the zero-centre-of-mass (CoM) subspace

$$\mathcal{X}_0 = \left\{ \mathbf{x} \in \mathbb{R}^{Mn} : \frac{1}{M} \sum_{i=1}^M \mathbf{x}^{(i)} = \mathbf{0} \right\}.$$

Change-of-basis matrix. To remove the CoM degree of freedom, we construct a change-of-basis matrix P that maps any $\mathbf{x} \in \mathcal{X}_0$ from the Mn -dimensional ambient space to an $(M-1)n$ -dimensional Euclidean space. Choose an orthonormal matrix $V_M \in \mathbb{R}^{(M-1) \times M}$ whose rows span the subspace $\{\mathbf{z} \in \mathbb{R}^M : \mathbf{1}^\top \mathbf{z} = 0\}$; one explicit choice is obtained by applying a QR factorization to $I_M - \frac{1}{M} \mathbf{1} \mathbf{1}^\top$. Define

$$P := V_M \otimes I_n \in \mathbb{R}^{(M-1)n \times Mn},$$

where \otimes denotes the Kronecker product. For any $\mathbf{x} \in \mathcal{X}_0$ we set $\tilde{\mathbf{x}} = P\mathbf{x} \in \mathbb{R}^{(M-1)n}$. The identities

$$PP^\top = I_{(M-1)n}, \quad P^\top P = (I_M - \frac{1}{M} \mathbf{1} \mathbf{1}^\top) \otimes I_n, \quad (22)$$

follow immediately from $V_M V_M^\top = I_{M-1}$ and $V_M^\top V_M = I_M - \frac{1}{M} \mathbf{1} \mathbf{1}^\top$. The first identity shows that P is an isometry between \mathcal{X}_0 and $\mathbb{R}^{(M-1)n}$.

Evaluating and Sampling from \mathcal{N}_x with full covariance For $\boldsymbol{\mu} \in \mathcal{X}_0$ and a positive-definite $\Sigma \in \mathbb{R}^{Mn \times Mn}$ whose image lies in \mathcal{X}_0 , define the density

$$\mathcal{N}_x(\mathbf{x}; \boldsymbol{\mu}, \Sigma) = \mathcal{N}(\tilde{\mathbf{x}}; \tilde{\boldsymbol{\mu}}, P\Sigma P^\top) \quad (23)$$

$$= \frac{1}{(2\pi)^{\frac{(M-1)n}{2}} \sqrt{\det(P\Sigma P^\top)}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^\top P^\top P \Sigma^{-1} P^\top P (\mathbf{x} - \boldsymbol{\mu})\right]. \quad (24)$$

To get samples from this distribution, we have to first sample from the normal distribution with dimension $(M-1)n$ and then we map it back to the Mn dimensional ambient space so that its center of gravity equals zero using the change-of-basis matrix P . This is the general way to get samples for \mathcal{N}_x with full covariance. Things will simplify if the covariance becomes isotropic.

If $\Sigma = \sigma^2 I_{Mn}$ is isotropic, equation 24 will becomes:

$$\mathcal{N}_x(\mathbf{x}; \boldsymbol{\mu}, \sigma^2 \mathbf{I}) := \frac{1}{(2\pi\sigma^2)^{\frac{(M-1)n}{2}}} \exp\left[-\frac{1}{2\sigma^2}(\mathbf{x} - \boldsymbol{\mu})^\top (\mathbf{x} - \boldsymbol{\mu})\right].$$

because of the fact that $PP^\top = I_{M-1}$. This is the Gaussian distribution constrained on zero CoM subspace with isotropic covariance (Hoogetboom et al., 2022; Xu et al., 2022). To sample from this distribution, we could sample in the Mn dimensional ambient space directly, and subtract its center of mass (Xu et al., 2022; Satorras et al., 2021).

B.1 Proposed Covariance Structure Satisfies E(3)-Equivariance

A diffusion model is E(3)-equivariant if, for every orthogonal transformation $R \in \text{O}(Mn)$,² (i) the forward-process endpoint distribution $p(\mathbf{x}_T)$ is invariant, and (ii) the reverse-process conditionals $p(\mathbf{x}_{t-1} | \mathbf{x}_t)$ are equivariant:

$$p(R\mathbf{x}_{t-1} | R\mathbf{x}_t) = p(\mathbf{x}_{t-1} | \mathbf{x}_t) \quad \forall t \in \{1, \dots, T\}.$$

Invariance of the isotropic Gaussian prior $p(\mathbf{x}_T) = \mathcal{N}_x(\mathbf{0}, T^2 I_{Mn})$ is immediate. We therefore focus on the conditional law.

Lemma 1. *Let the reverse transition be*

$$p(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}_x(\mathbf{x}_{t-1}; \mathbf{x}_t, \Sigma), \quad \Sigma = B \otimes I_n,$$

where $B \in \mathbb{R}^{M \times M}$ is positive definite. Then the reverse kernel is E(3)-equivariant.

²Translations are excluded because we restrict to the zero-CoM subspace \mathcal{X}_0 ; cf. Sec. B.

Proof. Every spatial rotation acts identically on all M particles and can therefore be written as the block-orthogonal matrix

$$R = I_M \otimes R_n, \quad R_n \in O(n).$$

Recall that the reverse kernel is defined on the $(M-1)n$ -dimensional coordinate $\tilde{\mathbf{x}} := P\mathbf{x}$ via

$$p(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(P\mathbf{x}_{t-1}; P\mathbf{x}_t, P\Sigma P^\top), \quad \Sigma = B \otimes I_n, \quad B \succ 0.$$

For arbitrary $\mathbf{x}, \mathbf{y} \in \mathcal{X}_0$ we have

$$\begin{aligned} \log p(R\mathbf{y} | R\mathbf{x}) &= \log \mathcal{N}(PR\mathbf{y}; PR\mathbf{x}, P\Sigma P^\top) \\ &= -\frac{1}{2}(PR\mathbf{y} - PR\mathbf{x})^\top P\Sigma^{-1}P^\top (PR\mathbf{y} - PR\mathbf{x}) + \text{const} \\ &= -\frac{1}{2}(\mathbf{y} - \mathbf{x})^\top R^\top P^\top P\Sigma^{-1}P^\top PR(\mathbf{y} - \mathbf{x}) \\ &\stackrel{(a)}{=} -\frac{1}{2}(\mathbf{y} - \mathbf{x})^\top P^\top PR^\top \Sigma^{-1}RP^\top P(\mathbf{y} - \mathbf{x}) \\ &\stackrel{(b)}{=} -\frac{1}{2}(\mathbf{y} - \mathbf{x})^\top P^\top P\Sigma^{-1}P^\top P(\mathbf{y} - \mathbf{x}) \\ &= -\frac{1}{2}(P\mathbf{y} - P\mathbf{x})^\top P\Sigma^{-1}P^\top (P\mathbf{y} - P\mathbf{x}) \\ &= \log \mathcal{N}(P\mathbf{y}; P\mathbf{x}, P\Sigma P^\top) \\ &= \log p(\mathbf{y} | \mathbf{x}), \end{aligned}$$

where step (a) uses the fact that $P^\top P$ commutes with R :

$$RP^\top P = (I_n \otimes R_n)(V_M \otimes I_n) = V_M \otimes R_n = P^\top PR$$

and step (b) uses the fact that

$$\begin{aligned} R^\top \Sigma^{-1} R &= (R^\top \Sigma R)^{-1} \\ &= ((I_M \otimes R_n^\top)(B \otimes I_n)(I_M \otimes R_n))^{-1} \\ &= (B \otimes R_n^\top R_n)^{-1} \\ &= (B \otimes I_n)^{-1} \\ &= \Sigma^{-1} \end{aligned}$$

Hence we have proved that our proposed covariance form satisfies the E(3)-equivariance required. \square

B.2 Permutation Equivariance for Many-Particle Systems

Datasets such as DW-4 and LJ-13 treat all particles as indistinguishable; hence the reverse kernel must be equivariant under any permutation of the M particles.

Lemma 2 (Permutation equivariance). *Let the covariance in the reduced coordinate space be*

$$\Sigma = B \otimes I_n, \quad B = (b - a)I_{M-1} + a\mathbf{1}\mathbf{1}^\top,$$

with parameters $b > 0$ and $a \in \mathbb{R}$ chosen such that $b - a > 0$ and $b + (M-2)a > 0$, ensuring $B \succ 0$. For any permutation matrix $S_M \in \mathbb{R}^{M \times M}$ and the block permutation $S := S_M \otimes I_n$, the reverse transition satisfies $p(S\mathbf{y} | S\mathbf{x}) = p(\mathbf{y} | \mathbf{x})$ for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}_0$.

Proof. (i) **Commutation with the CoM projector.** Using $P^\top P = (V_M \otimes I_n)$ (Sec. B) and the fact that $S_M \mathbf{1} = \mathbf{1}$ and $\mathbf{1}^\top S_M = \mathbf{1}^\top$,

$$SP^\top P = (S_M(I_M - \frac{1}{M}\mathbf{1}\mathbf{1}^\top)) \otimes I_n = (S_M - \frac{1}{M}\mathbf{1}\mathbf{1}^\top) \otimes I_n = P^\top PS,$$

so S commutes with the projector onto the zero-CoM subspace.

(ii) Invariance of the covariance. Because S_M merely reorders rows and columns,

$$S_M B S_M^\top = (b - a) I_{M-1} + a S_M \mathbf{1} \mathbf{1}^\top S_M^\top = (b - a) I_{M-1} + a \mathbf{1} \mathbf{1}^\top = B,$$

whence $S \Sigma S^\top = \Sigma$.

(iii) Equality of densities. With the same quadratic-form argument used in the E(3)-equivariance proof,

$$(\mathbf{y} - \mathbf{x})^\top P^\top P \Sigma^{-1} P^\top P (\mathbf{y} - \mathbf{x}) = (S\mathbf{y} - S\mathbf{x})^\top P^\top P \Sigma^{-1} P^\top P (S\mathbf{y} - S\mathbf{x}),$$

establishing $p(S\mathbf{y} \mid S\mathbf{x}) = p(\mathbf{y} \mid \mathbf{x})$. \square

B.3 Permutation Equivariance for General Molecular Distributions

Again, we aim to show that our choice of B_ϕ satisfies the permutation-equivariance constraint under the parametrization proposed in Section 4.3.

Permutation equivariance. Let S_M denote any permutation matrix that reorders atoms *within* each label class (i.e., $L_{S_M(i)} = L_i$). Analogous to the proof in Section B.2, because any permutation matrix commutes with the CoM projector, this condition is automatically satisfied. All that remains is to prove the invariance of the covariance. Because the two parameterizations above depend on the indices only through the labels, we have

$$S_M B_\phi(n) S_M^\top = B_\phi(n).$$

Consequently, the lifted covariance

$$\Sigma_\phi(n) = B_\phi(n) \otimes I_n$$

satisfies

$$S \Sigma_\phi(n) S^\top = \Sigma_\phi(n)$$

for every block permutation $S = S_M \otimes I_n$, so the reverse kernel remains permutation-equivariant by the argument of the above section.

C Implementation Details

All experiments were run on a single NVIDIA A100 GPU (80 GB); hence every “GPU hour” refers to wall-clock hours on that device.

C.1 Datasets

- **GMM-2.** A two-component Gaussian mixture used as a toy benchmark. Training and evaluation samples are drawn directly from the target distribution.
- **DW-4 and LJ-13.** We use the 10^7 configurations released by Klein et al. (2023). A subset of 10^5 configurations trains the score model and VT-DIS; the remainder forms the test set.
- **Alanine dipeptide.** We adopt the molecular dataset of Midgley et al. (2023), which provides 10^5 training samples and 10^6 test samples.

C.2 Training the Score Model

Our training pipeline—including parameterization and preconditioning—follows Karras et al. (2022).

- **GMM-2.** A seven-layer MLP (hidden width 512), identical to the network used by Akhoun-Sadegh et al. (2024) for the GMM-40 task. Training runs for 100,000 iterations with a batch size of 1,024 for both $d = 50$ and $d = 100$.

Table 2: Score-model training time on a single A100.

Dataset	Architecture	Layers	Hidden dim	Batch Size	Iterations	GPU-h
GMM-2 ($d=50$)	MLP	7	512	5000	100 000	0.6
GMM-2 ($d=100$)	MLP	7	512	5000	100 000	1.0
DW-4	EGNN	4	128	512	100 000	0.5
LJ-13	EGNN	6	128	512	100 000	8.0
Alanine dipeptide	EGNN	7	256	512	100 000	8.0

Table 3: VT-DIS post-training time on a single A100.

Dataset	Batch Size	Iterations	Steps K	GPU-h
GMM-2 ($d = 50$)	512	5 000	100	1.8
GMM-2 ($d = 100$)	512	5 000	100	3.5
DW-4	512	5 000	100	0.9
LJ-13	512	5 000	100	2.7
Alanine dipeptide	128	5 000	100	4.3

- **DW-4.** EGNN backbone with 4 layers and a hidden width of 128.
- **LJ-13.** EGNN backbone with 6 layers and a hidden width of 128.
- **Alanine Dipeptide.** EGNN backbone with 7 layers and a hidden width of 256; node features include atom types and bond information. Scaling atomic coordinates by 10 during training improves sample quality; network outputs are divided by 10 at inference.

All models are trained for 100,000 iterations with a batch size of 512 using the Adam optimizer, a learning rate of 0.001, and a cosine annealing schedule with a minimum learning rate of 1×10^{-6} . Table 2 summarizes these details.

C.3 VT-DIS Post-training

We follow Algorithm 1 exactly. With the score network frozen, we optimize the per-step covariance matrices by minimizing the α -divergence ($\alpha = 2$) between the forward and reverse processes using the same training set as for score learning. Each run comprises 5,000 optimization steps with a batch size of 512 (256 for alanine dipeptide). All runs use the Adam optimizer with a learning rate of 0.01 and a cosine annealing schedule with a minimum learning rate of 1×10^{-6} . Table 3 summarizes these settings.

The reported training time corresponds to tuning the isotropic covariance; diagonal and full covariance forms incur similar runtimes. We fix the number of time steps to 100; empirically, with fixed epochs and batch sizes, training time increases linearly with the number of sampling steps.

C.4 Evaluation Metrics

We measure efficiency using forward and reverse effective sample sizes (ESS), defined in Eqs. equation 11 and equation 12. The number of Monte Carlo draws for each estimate is specified in the corresponding figure or table caption. Table 4 summarizes the sampling time for each target distribution. Empirically, sampling time increases linearly with the number of sampling steps.

Table 4: VT-DIS sampling time on a single A100.

Dataset	Batch Size	Steps K	GPU-min
GMM-2 ($d = 50$)	100,000	100	0.25
GMM-2 ($d = 100$)	100,000	100	0.5
DW-4	100,000	100	1.3
LJ-13	100,000	100	4.3
Alanine dipeptide	10,000	100	3

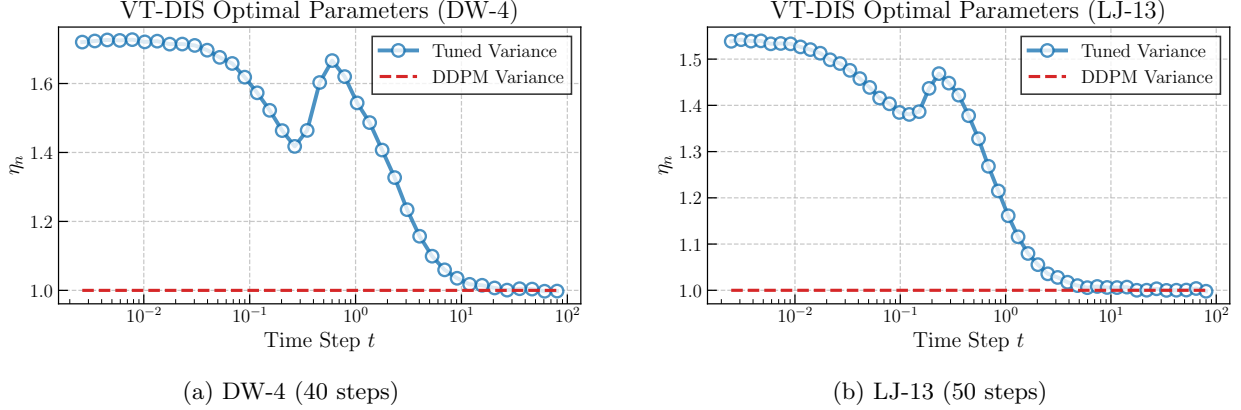


Figure 5: **Optimal variance scaling η_n .** VT-DIS inflates early-time variance to improve coverage, adapts non-monotonically across the transition regime, and converges to the DDPM value ($\eta_n = 1$, red dashed line) in the pure-noise limit.

D Additional Experimental Results

D.1 Optimized Covariance Matrices

Figure 5 contrasts the per-step variance-scaling factor of VT-DIS, η_n (see equation 6), with the unit baseline of the original DDPM. Although we display the isotropic case, the same qualitative pattern appears for diagonal and low-rank covariances.

Three regimes. Reading the t -axis from left ($t \approx 10^{-3}$) to right ($t \approx 10^2$) follows the reverse SDE from data space to pure noise. The learned profile is *not* monotone but exhibits three characteristic regimes:

1. **Data-dominated regime** ($t \lesssim 10^{-1}$). Here the signal-to-noise ratio (SNR) is high, the score network is the main source of model error, and under-dispersion relative to the true forward kernel results in heavy importance-weight tails. VT-DIS compensates by *inflating* the variance ($\eta_n \approx 1.6$ – 1.7), which widens the proposal and improves coverage.
2. **Transition regime** ($10^{-1} \lesssim t \lesssim 10^1$). As the SNR drops, the target conditional gradually morphs from highly non-Gaussian to approximately Gaussian. The mismatch between the analytic DDPM variance (optimal for the *Gaussian* limit) and the true kernel therefore decreases. VT-DIS first reduces the scaling (a local minimum at $t \approx 0.3$), then overshoots to a peak at $t \approx 1$, before descending. The intermediate bump arises where the learned score D_{θ} tends to *over-estimate* the gradient magnitude (a common artefact of score matching when the SNR is intermediate, e.g. $\mathcal{O}(1)$) (Xu et al., 2023). Increasing the proposal variance temporarily counter-balances this over-confidence, reducing the second-moment of the importance weights.

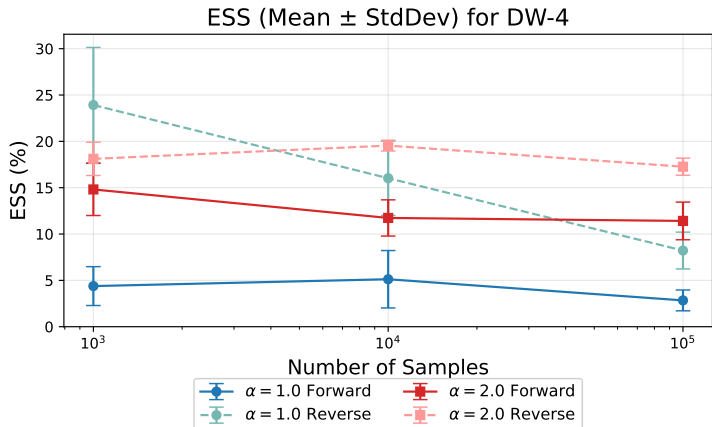


Figure 6: **KL ($\alpha = 1$) vs. $\alpha = 2$ divergence on DW-4.** $\alpha = 2$ tuning (red) attains higher, size-independent ESS, whereas KL tuning (blue) collapses for large M .

- Noise regime** ($t \gtrsim 10^1$). The state distribution is now close to $\mathcal{N}(0, t^2 I)$, for which the DDPM variance is *exactly* optimal. Accordingly, η_n relaxes to 1, confirming that VT-DIS does not perturb steps that are already statistically efficient.

Implications. The variance inflations at early steps are crucial: without them, trajectories started near the data manifold place negligible mass in regions where the forward kernel has significant probability, causing weight explosion and the near-zero forward ESS observed for the vanilla baseline. Conversely, keeping $\eta_n \rightarrow 1$ in the noise regime ensures that VT-DIS retains the favourable properties of the original DDPM sampler at large time steps, adding no unnecessary dispersion..

D.2 KL versus α -divergence

The VT-DIS objective in equation 10 minimises the α -divergence with $\alpha = 2$, which is proportional to the *variance* of importance weights and therefore maximizes ESS directly (Midgley et al., 2022). For completeness we retrain the covariance parameters on DW-4 using the forward KL divergence ($\alpha = 1$) and compare the resulting samplers.

Figure 6 plots forward (solid) and reverse (dashed) ESS as functions of the sample size M . Tuning with $\alpha = 2$ yields consistently higher and *stable* ESS in both directions. In contrast, KL-tuned parameters give (i) low forward ESS and (ii) reverse ESS that *decreases* with M , a classic symptom of heavy weight tails caused by poor proposal coverage. These observations confirm that controlling the second moment of weights ($\alpha = 2$) is crucial for robust importance sampling, whereas KL divergence alone fails to regularise the tails.

E Broader Impacts

Our method enhances the efficiency and accessibility of unbiased molecular sampling, potentially speeding up the discovery of catalysts, drugs, and new materials while reducing reliance on costly experiments. At the same time, stronger samplers could be misused to design toxic or environmentally harmful compounds more rapidly. There is also a risk that institutions with advanced computing resources will widen the gap with those lacking such infrastructure, reinforcing inequities in scientific research.