
Sampling by averaging: A multiscale approach to score estimation

Paula Cordero-Encinar

Department of Mathematics
Imperial College London

paula.cordero-encinar22@imperial.ac.uk

Andrew B. Duncan

Department of Mathematics
Imperial College London
a.duncan@imperial.ac.uk

Sebastian Reich

Institut für Mathematik
Universität Potsdam
sereich@uni-potsdam.de

O. Deniz Akyildiz

Department of Mathematics
Imperial College London
deniz.akyildiz@imperial.ac.uk

Abstract

We introduce a novel framework for efficient sampling from complex, unnormalised target distributions by exploiting multiscale dynamics. Traditional score-based sampling methods either rely on learned approximations of the score function or involve computationally expensive nested Markov chain Monte Carlo (MCMC) loops. In contrast, the proposed approach leverages stochastic averaging within a slow-fast system of stochastic differential equations (SDEs) to estimate intermediate scores along a diffusion path without training or inner-loop MCMC. Two algorithms are developed under this framework: MULTALMC, which uses multiscale annealed Langevin dynamics, and MULTCDIFF, based on multiscale controlled diffusions for the reverse-time Ornstein-Uhlenbeck process. Both overdamped and underdamped variants are considered, with theoretical guarantees of convergence to the desired diffusion path. The framework is extended to handle heavy-tailed target distributions using Student’s t-based noise models and tailored fast-process dynamics. Empirical results across synthetic and real-world benchmarks, including multimodal and high-dimensional distributions, demonstrate that the proposed methods are competitive with existing samplers in terms of accuracy and efficiency, without the need for learned models.

1 Introduction

Efficiently sampling from complex unnormalised probability distributions is a fundamental challenge across many scientific domains, including statistics, chemistry, computational physics, and biology, see, e.g. Gelman et al. [2013], Lelièvre et al. [2010], Liu [2008]. Classical Markov Chain Monte Carlo (MCMC) methods provide asymptotically unbiased samples under mild assumptions on the target [Durmus and Moulines, 2017, Mousavi-Hosseini et al., 2023]. However, they often become computationally inefficient in practice, especially for high-dimensional, multimodal targets, due to the need for long Markov chains to ensure adequate mixing and convergence. To alleviate these issues, annealing-based strategies construct a sequence of smoother intermediate distributions bridging a simple base distribution and the target. This principle underlies algorithms such as Parallel Tempering [Geyer and Thompson, 1995, Swendsen and Wang, 1986], Annealed Importance Sampling [Neal, 2001], and Sequential Monte Carlo [Del Moral et al., 2006, Doucet et al., 2001].

Recently, score-based diffusion methods [Hyvärinen, 2005, Song and Ermon, 2020] have emerged as powerful models capable of generating high-quality samples. These approaches simulate a reverse

diffusion process guided by time-dependent score functions. While successful in generative modelling, they rely on access to training samples—a setting that differs from sampling problems, where only the unnormalised target density is available. Therefore, traditional score-matching techniques are not applicable. To address this, a growing body of work has explored ways to estimate the score function using only the target density, either through neural networks using variational objectives [Chen et al., 2025, Richter and Berner, 2024, Vargas et al., 2024] or training-free alternatives based on MCMC [Chen et al., 2024, Grenioux et al., 2024, Huang et al., 2024, Phillips et al., 2024, Saremi et al., 2024]. In this work, we present a novel training-free sampling framework that eliminates the need for the inner MCMC loop required in previous works to approximate the score function. Our main contributions are as follows:

- We demonstrate how a multiscale system of SDEs can be leveraged to enable efficient sampling along a diffusion path, obviating the need to estimate the score function.
- In particular, we propose two algorithms: **MULTALMC**: Multiscale Annealed Langevin Monte Carlo (Section 3.1) and **MULTCDIFF**: Multiscale Controlled Diffusions (Section 3.2) based on discretisations of two different multiscale SDEs: annealed Langevin dynamics for general noise schedules and the reverse SDE associated with an Ornstein-Uhlenbeck (OU) noising process, respectively. We also explore accelerated versions using underdamped dynamics.
- We provide a theoretical analysis of the proposed methods (Section 4) and demonstrate their effectiveness across different synthetic and real-world benchmarks (Section 5).

2 Background and related work

2.1 Diffusion paths

The reverse process in diffusion models consists in sampling along a path of probability distributions $(\mu_t)_{t \in [0,1]}$, which starts at a simple distribution and ends at the target distribution $\pi \propto e^{-V_\pi}$ on \mathbb{R}^d . Following Chehab and Korba [2024], the intermediate distributions can be expressed as follows

$$\mu_t(x) = \frac{\pi(x/\sqrt{\lambda_t})}{\lambda_t^{d/2}} * \frac{\nu(x/\sqrt{1-\lambda_t})}{(1-\lambda_t)^{d/2}}, \quad (1)$$

where $*$ denotes the convolution operation, ν describes the base or *noising* distribution, and λ_t is an increasing function called schedule, such that, $\lambda_t \in [0, 1]$ and $\lambda_1 = 1$. We refer to the path $(\mu_t)_{t \in [0,1]}$ in (1) as the *diffusion path*. By choosing $\lambda_t = e^{-2T(1-t)}$ for some fixed T , we recover the path corresponding to a forward OU noising process, a widely used approach in diffusion models [Benton et al., 2024, Chen et al., 2023, Song et al., 2021]. In this case, the reverse-time dynamics can be characterised [Anderson, 1982], requiring only to estimate the score functions $\nabla \log \mu_t$ along the path. However, for general diffusion paths, the reverse process cannot be described by a closed-form SDE due to intractable control terms in the drift, which can be estimated using neural networks as in Albergo et al. [2023]. Alternatively, sampling from the diffusion path can be achieved through *annealed Langevin dynamics* [Cordero-Encinar et al., 2025, Song and Ermon, 2019], which also relies on score estimation but avoids dealing with intractable drift terms.

The diffusion path has demonstrated very good performance in the generative modelling literature where the score of the intermediate distributions can be estimated empirically using score matching techniques [Song and Ermon, 2019]. In the context of sampling, however, estimating the score is more challenging, as samples from the data distribution are not available and instead we only have access to an unnormalised target probability distribution. In particular, under the assumption that ν, π are bounded and have finite second-order moments, and that $|\nabla \log \nu|$ and $|\nabla V_\pi|$ are bounded, the expression of the score of the intermediate distributions is given by

$$\nabla \log \mu_t(x) = \frac{1}{\sqrt{1-\lambda_t}} \mathbb{E}_{\rho_{t,x}(y)} \left[\nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_t}} \right) \right] = -\frac{1}{\sqrt{\lambda_t}} \mathbb{E}_{\rho_{t,x}(y)} \left[\nabla V_\pi \left(\frac{y}{\sqrt{\lambda_t}} \right) \right], \quad (2)$$

where $\rho_{t,x}(y) \propto \nu \left(\frac{x-y}{\sqrt{1-\lambda_t}} \right) e^{-V_\pi(y/\sqrt{\lambda_t})}$ and we have used that $\nabla(f * g) = (\nabla f) * g = f * (\nabla g)$ when f and g are differentiable functions. Notably, in the limits $\lambda_t = 0$ and $\lambda_t = 1$, the conditional distribution $\rho_{t,x}$ converges to a Dirac delta function. The score in (2) involves an expectation over $\rho_{t,x}$, which typically requires MCMC sampling and results in a computationally expensive nested-loop structure. We avoid this by exploiting multiscale dynamics, as described in the following section.

2.2 Multiscale dynamics and stochastic averaging

Consider the slow-fast system of SDEs

$$\begin{cases} dX_t = b(t, X_t, Y_t)dt + \sqrt{2}dB_t \\ dY_t = \frac{1}{\varepsilon}f(t, X_t, Y_t)dt + \sqrt{\frac{2}{\varepsilon}}d\tilde{B}_t \end{cases}$$

where $0 < \varepsilon \ll 1$ controls the scale separation between the slow X_t and the fast component Y_t . We assume that, for $\varepsilon = 1$ and fixed t and X_t , Y_t is ergodic. As $\varepsilon \rightarrow 0$, the fast component Y_t evolves on a much shorter timescale than X_t due to the time-rescaling property of Brownian motion. Intuitively, this means that Y_t rapidly reaches its stationary distribution while X_t remains substantially unchanged. In the limit $\varepsilon \rightarrow 0$, the slow-process X_t converges to the averaged dynamics \bar{X}_t which is given by

$$d\bar{X}_t = \bar{b}(t, \bar{X}_t)dt + \sqrt{2}dB_t, \quad \bar{b}(t, x) = \mathbb{E}_{Y \sim \rho_{t,x}}[b(t, x, Y)],$$

where $\rho_{t,x}$ is the invariant measure of the frozen process $dY_s = f(t, x, Y_s)ds + \sqrt{2}d\tilde{B}_s$. In computational statistics, stochastic slow-fast systems have been leveraged to simulate the averaged dynamics efficiently [Akyildiz et al., 2024, Harvey et al., 2011, Pavliotis and Stuart, 2008, Weinan et al., 2005]. In our context of sampling from diffusion paths, the averaged drift $\bar{b}(t, x)$ corresponds to the score $\nabla \log \mu_t(x)$, which is defined as an expectation (see Eq. (2)). By leveraging multiscale dynamics, we can approximate this expectation without relying on a nested-loop sampling structure, a key insight that underpins our method, as summarised in Section 3.

2.3 Related work

Diffusion-based samplers Several recent works have proposed sampling algorithms based on the diffusion path. We divide them into two categories: neural samplers and learning-free samplers. Neural samplers such as those in Chen et al. [2025], Noble et al. [2025], Richter and Berner [2024], Vargas et al. [2024], Zhang and Chen [2022] estimate the time-dependent drift function of the diffusion process using a parametrised model, typically a neural network, by solving a variational inference problem defined over the space of path measures. As in diffusion generative models, their performance is limited by the expressiveness of the chosen model class. In contrast, learning-free samplers [Chen et al., 2024, Grenioux et al., 2024, He et al., 2024b, Huang et al., 2024, Phillips et al., 2024, Saremi et al., 2024] do not require any training similar to our work. Instead, they estimate the score function $\nabla \log \mu_t$ using MCMC samples from the conditional distribution $\rho_{t,x}$. That is, these approaches involve a bi-level structure, with an inner MCMC loop used at each step of the outer diffusion-based sampler. In this sense, they implement a *Langevin-within-Langevin* strategy. Our approach avoids nested loops by leveraging multiscale dynamics along the diffusion path. This results in a *Langevin-by-Langevin* sampler that is both simpler and more computationally efficient.

Besides, while acceleration techniques based on underdamped dynamics [Duncan et al., 2017, Monmarché, 2020] have shown promise in the generative modelling literature [Dockhorn et al., 2022a,b, Holderrieth et al., 2024], and recent work by Blessing et al. [2025] introduces neural underdamped diffusion samplers, no training-free accelerated samplers based on diffusion processes have been proposed. Addressing this gap is one of the key contributions of our work.

Proximal samplers Our approach is also closely related to the class of proximal samplers [Chen et al., 2022, Lee et al., 2021b]. These methods define an auxiliary joint distribution of the form $\tilde{\pi}(x, z) \propto \pi(x) \exp(-\|x - z\|^2/(2\eta))$ and perform Gibbs sampling by alternating between the convolutional distribution $\tilde{\pi}(z|x)$ and the denoising posterior $\tilde{\pi}(x|z)$. In practice, sampling from the denoising posterior $\tilde{\pi}(x|z)$ involves finding a mode using gradient-based optimisation techniques, followed by rejection sampling around that mode. This step can be computationally expensive and inefficient, especially in high-dimensional settings. Our method improves upon these approaches by replacing the proximal-point-style posterior sampling step with a fast-timescale dynamics, enabling more efficient exploration of the denoising posterior. Furthermore, unlike proximal samplers that operate at a fixed noise level, our approach targets the entire diffusion path, smoothly interpolating between the base distribution and the target.

3 Sampling using multiscale dynamics

A key challenge in sampling from the diffusion path in (1) is accurately estimating the score function when we do not have direct access to samples. Since the score of the intermediate distributions along the diffusion path is computed as an expectation over the conditional distribution $\rho_{t,x}$ (see Eq. (2)), we propose to sample from the time-dependent conditional distribution $\rho_{t,x}$ using an alternative diffusion process with a shorter timescale than the original dynamics following the path in (1). This enables both processes to run in parallel, unlike existing methods that execute them sequentially. In those approaches, each iteration of the original diffusion process (outer loop) requires multiple time steps of the process targeting the conditional distribution $\rho_{t,x}$ (inner loop) to learn the score function, leading to increased computational cost and a higher number of evaluations of the target potential. By employing different timescales for each diffusion process, our method improves sampling efficiency while ensuring convergence to the correct target distribution, as justified by stochastic averaging theory [Liu et al., 2020]. We explore this approach in two settings: (i) when the original diffusion is a Langevin dynamics driven by the scores of the intermediate probability distribution $\nabla \log \mu_t$ (Section 3.1), and (ii) when we use the reverse SDE corresponding to an OU noising process (Section 3.2).

3.1 MULTALMC: Multiscale Annealed Langevin Monte Carlo

Following Cordero-Encinar et al. [2025], Song and Ermon [2019], we use a time inhomogeneous Langevin SDE to sample from the diffusion path (1) given by

$$d\tilde{X}_t = \nabla \log \hat{\mu}_t(\tilde{X}_t)dt + \sqrt{2}dB_t, \quad t \in [0, 1/\kappa], \quad (3)$$

where $\hat{\mu}_t$ denotes the reparametrised probability distributions from the diffusion path ($\hat{\mu}_t = \mu_{\kappa t}_{t \in [0, 1/\kappa]}$, for some $0 < \kappa < 1$, $\tilde{X}_0 \sim \mu_0 = \nu$ and $(B_t)_{t \geq 0}$ is a Brownian motion). We consider both Gaussian diffusion paths where the base distribution ν is a Normal distribution and heavy-tailed diffusions corresponding to a Student's t base distribution. Since the scores $\nabla \log \hat{\mu}_t$ are intractable and computed as expectations over $\rho_{t,x}$ (see Eq. (2)), we adopt a multiscale system to approximate the averaged dynamics in (3). The fast process targetting the conditional distribution $\rho_{t,x}$ can follow any fast-mixing dynamics. In particular, we explore the use of overdamped Langevin dynamics or a modified Itô diffusion [He et al., 2024a], but it is important to remark that our framework is more general. When using overdamped Langevin dynamics, the fast process is given by

$$\begin{aligned} dY_t &= \frac{1}{\varepsilon} \nabla \log \hat{\rho}_{t,X_t}(Y_t)dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t, \quad t \in [0, 1/\kappa], \\ \nabla \log \hat{\rho}_{t,X_t}(y) &= -\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) + \frac{1}{\sqrt{1 - \lambda_{\kappa t}}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1 - \lambda_{\kappa t}}} \right), \end{aligned}$$

where $\hat{\rho}_{t,X_t}$ is a reparametrised version of ρ_{t,X_t} and $(\tilde{B}_t)_{t \geq 0}$ is a Brownian motion on \mathbb{R}^d . This suggests sampling from the following stochastic slow-fast system

$$\begin{cases} dX_t &= \begin{cases} \frac{1}{\sqrt{1 - \lambda_{\kappa t}}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1 - \lambda_{\kappa t}}} \right) dt + \sqrt{2}dB_t & \text{if } \lambda_{\kappa t} < \tilde{\lambda} \\ -\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) dt + \sqrt{2}dB_t & \text{if } \lambda_{\kappa t} \geq \tilde{\lambda} \end{cases} \\ dY_t &= \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) + \frac{1}{\sqrt{1 - \lambda_{\kappa t}}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1 - \lambda_{\kappa t}}} \right) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t. \end{cases} \quad (4)$$

This system converges to the averaged dynamics described in (3) as $\varepsilon \rightarrow 0$, since taking the expectation of the drift term in the X_t dynamics with respect to $Y_t \sim \hat{\rho}_{t,X_t}$ recovers the score function $\nabla \log \hat{\mu}_t$ (see Equation (2) and Section 4). For numerical implementation, we use an adaptive scheme that alternates between the two expressions for X_t depending on the value of $\lambda_{\kappa t}$ in order to avoid numerical instabilities as $\lambda_{\kappa t}$ approaches 0 or 1. The resulting system of SDEs can be discretised using different numerical schemes, however, standard integrators such as Euler-Maruyama (EM) will be unstable when ε is small, due to the large-scale separation between the slow and fast processes. To mitigate this, we leverage discretisations which remain stable, despite the stiffness of the SDE, such as the exponential integrator [Hochbruck and Ostermann, 2010] and the SROCK method [Abdulle et al., 2018]. Furthermore, since $\rho_{t,x}$ converges to a Dirac delta distribution centred at 0 when $\lambda_{\kappa t} = 0$ and centred at x when $\lambda_{\kappa t} = 1$, we apply a further slight modification to the slow-fast system to improve numerical stability (see App. B.1.1 for details).

Accelerated MULTALMC When implemented in practice, the overdamped system in (4) requires a very small value of κ (which corresponds to a slowly varying dynamics driven by $\nabla \log \hat{\mu}_t$) to accurately follow the marginals of the path $(\hat{\mu}_t)_{t \in [0, 1/\kappa]}$ and ultimately sample from the target distribution. This results in a large number of discretisation steps, thus making the method computationally expensive, see App. E.1 for more details. In contrast, underdamped dynamics exhibit faster mixing times [Bou-Rabee and Eberle, 2023, Eberle and Lörler, 2024], which intuitively helps the system explore the space more efficiently and better track the intermediate distributions along the path $(\hat{\mu}_t)_{t \in [0, 1/\kappa]}$. These dynamics have also shown strong empirical performance in the neural sampling literature [Blessing et al., 2025], requiring fewer discretisation steps. Motivated by these advantages, we propose augmenting the state space with auxiliary velocity variables \bar{V}_t , and adopting the following underdamped Langevin diffusion to sample from the path $(\hat{\mu}_t)_{t \in [0, 1/\kappa]}$

$$\begin{cases} d\bar{X}_t &= M^{-1}\bar{V}_t dt \\ d\bar{V}_t &= \nabla \log \hat{\mu}_t(\bar{X}_t) dt - \Gamma M^{-1}\bar{V}_t dt + \sqrt{2\Gamma} dB_t. \end{cases} \quad \text{for } t \in [0, 1/\kappa]. \quad (5)$$

The mass parameter M determines the coupling between position X_t and velocity V_t , while the friction coefficient Γ controls the strength of noise injected into the velocity component. Both M and Γ are symmetric positive definite matrices, in practice we consider $M = mId$ and $\Gamma = \gamma Id$. The behaviour of the system is governed by the interplay between M and Γ [McCall, 2010]. Based on this, we propose to use the following underdamped slow-fast system to sample from the target distribution

$$\begin{cases} dX_t &= M^{-1}V_t dt \\ dV_t &= \begin{cases} \left(\frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1-\lambda_{\kappa t}}} \right) - \Gamma M^{-1}V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } \lambda_{\kappa t} < \tilde{\lambda} \\ \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_{\pi} \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) - \Gamma M^{-1}V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } \lambda_{\kappa t} \geq \tilde{\lambda} \end{cases} \\ dY_t &= \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_{\pi} \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) + \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1-\lambda_{\kappa t}}} \right) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t. \end{cases} \quad (6)$$

Note that in the underdamped setting, both X_t, V_t are treated as slow variables, and as $\varepsilon \rightarrow 0$, this system converges to the averaged dynamics in (5). To implement the system in practice, we combine a symmetric splitting scheme (OBABO) [Monmarché, 2020] for the slow-dynamics, and discretise the fast-dynamics using SROCK method [Abdulle et al., 2018]. We will evaluate the performance of the proposed sampler based on this underdamped slow-fast system (Algorithm 3), referred to as MULTALMC, under two choices of the reference distribution ν : a standard Gaussian and a Student's t distribution. A simplified version of the implementation is presented in Alg. 1, while a more detailed description is provided in App. B.1.2.

Heavy-tailed diffusions When the target distribution has heavy tails, standard Gaussian diffusions often fail to capture the correct tail behaviour [Pandey et al., 2025, Shariatian et al., 2025]. This motivates the use of heavy-tailed noising processes, which have been shown to offer theoretical guarantees in such settings [Cordero-Encinar et al., 2025].

We propose sampling along a heavy-tailed diffusion path $(\hat{\mu}_t)_{t \in [0, 1/\kappa]}$, where the reference distribution ν is chosen as a Student's t-distribution with tail index α , i.e., $\nu \propto (1 + \|x\|^2/\alpha)^{-(\alpha+d)/2}$. This can be implemented using the underdamped slow-fast system defined in (6). However, when the target distribution is heavy-tailed, the use of overdamped Langevin dynamics for the fast component results in slow convergence [Mousavi-Hosseini et al., 2023, Wang, 2006]. This motivates the consideration of alternative diffusion processes for the fast dynamics that mix more efficiently under such conditions. Specifically, we explore a modified Itô diffusion proposed in He et al. [2024a], defined as

$$\begin{aligned} dY_t &= -\frac{1}{\varepsilon} (\alpha + d - 1) \nabla U_{\hat{\mu}_t, X_t}(y) dt + \sqrt{\frac{2U_{\hat{\mu}_t, X_t}}{\varepsilon}} d\tilde{B}_t, \\ \hat{\mu}_{t, X_t}(y) &\propto \left(\sqrt{1 + \frac{\|x - y\|^2}{\alpha(1 - \lambda_{\kappa t})}} \pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right)^{-\frac{1}{\alpha+d}} \right)^{-(\alpha+d)} = U_{\hat{\mu}_t, X_t}(y)^{-(\alpha+d)}. \end{aligned} \quad (7)$$

This modified Itô diffusion for heavy-tailed distributions is shown to have faster convergence guarantees when the target has finite variance [He et al., 2024a]. By replacing the fast process in (6) with the diffusion defined in (7), we obtain samplers better suited for heavy-tailed targets, with improved convergence properties.

3.2 MULTCDIFF: Multiscale Controlled Diffusions

Annealed Langevin dynamics (3) offers a convenient approach to sampling from the target distribution for general schedules λ_t that interpolate between the base and target distributions in finite time.

Algorithm 1 MULTALMC sampler: accelerated version

Require: Schedule function λ_t , value for λ_δ and $\tilde{\lambda}$, friction coefficient Γ , mass parameter M , number of sampling steps L , time discretisation $0 = T_0 < \dots < T_L = 1/\kappa$, step size $h_l = T_{l+1} - T_l$. Constants μ, ν, κ for SROCK step from (14), (15).

Initial samples $X_0 \sim \mathcal{N}(0, I)$, $V_0 \sim \mathcal{N}(0, MI)$, $Y_0 \sim \mathcal{N}(0, I)$, $\lambda'_{\kappa t} = \begin{cases} \lambda_\delta & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \lambda_{\kappa t} & \text{if } \lambda_\delta \leq \lambda_{\kappa t} \leq 1. \end{cases}$

for $l = 0$ **to** L **do**

$\xi_l^{(1)}, \xi_l^{(2)}, \xi_l^{(3)} \sim \mathcal{N}(0, I)$

if $0 \leq \lambda_{\kappa T_l} < 1 - \lambda_\delta$ **then**

Half-step for velocity component

$V_l' = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right)V_l + \sqrt{h_l\Gamma}\xi_l^{(1)}, \quad V_l'' = \begin{cases} V_l' + \frac{h_l}{2\sqrt{1-\lambda'_{\kappa T_l}}}\nabla \log \nu\left(\frac{X_l - Y_l}{\sqrt{1-\lambda'_{\kappa T_l}}}\right) & \text{if } \lambda'_{\kappa T_l} < \tilde{\lambda} \\ V_l' - \frac{h_l}{2\sqrt{\lambda'_{\kappa T_l}}}\nabla V_\pi\left(\frac{Y_l}{\sqrt{\lambda'_{\kappa T_l}}}\right) & \text{if } \lambda'_{\kappa T_l} \geq \tilde{\lambda} \end{cases}$

Full EM-step for position component $X_{l+1} = X_l + h_l M^{-1} V_l''$

Full SROCK-step for fast component $Y_{l+1} = f(Y_l, X_{l+1}, \lambda'_{\kappa T_l}, h_l, \varepsilon, \mu, \nu, \kappa) + \sqrt{\frac{2h}{\varepsilon}}\xi_l^{(2)}$

Half-step for velocity component

$V_l''' = \begin{cases} V_l'' + \frac{h_l}{2\sqrt{1-\lambda'_{\kappa T_l}}}\nabla \log \nu\left(\frac{X_{l+1} - Y_{l+1}}{\sqrt{1-\lambda'_{\kappa T_l}}}\right) & \text{if } \lambda'_{\kappa T_l} < \tilde{\lambda}, \\ V_l'' - \frac{h_l}{2\sqrt{\lambda'_{\kappa T_l}}}\nabla V_\pi\left(\frac{Y_{l+1}}{\sqrt{\lambda'_{\kappa T_l}}}\right) & \text{if } \lambda'_{\kappa T_l} \geq \tilde{\lambda}, \end{cases}, \quad V_{l+1} = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right)V_l''' + \sqrt{h_l\Gamma}\xi_l^{(3)}$

if $1 - \lambda_\delta \leq \lambda_{\kappa T_l} \leq 1$ **then**

Half-step for velocity component

$V_l' = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right)V_l + \sqrt{h_l\Gamma}\xi_l^{(1)}, \quad V_l'' = V_l' - \frac{h_l}{2}\nabla V_\pi(X_l)$

Full EM-step for position component $X_{l+1} = X_l + h_l M^{-1} V_l''$

Half-step for velocity component

$V_l''' = V_l'' - \frac{h_l}{2}\nabla V_\pi(X_{l+1}), \quad V_{l+1} = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right)V_l''' + \sqrt{h_l\Gamma}\xi_l^{(3)}$

end for

return (X_L, V_L)

However, a key drawback of annealed Langevin dynamics is that it introduces bias, even if perfectly simulated (see Section 4.2). This bias can be corrected by incorporating control terms. Although such control terms are generally intractable for arbitrary schedules, they can be computed explicitly in the case of an OU noising process. The multiscale approach introduced in the previous section naturally extends to this controlled setting, in both overdamped and underdamped regimes. Given the improved efficiency of underdamped dynamics, we focus our discussion on the underdamped formulation (see App. B.2.1 for details on the overdamped case). When using an OU process as the forward noising model, the underdamped diffusion, initialised at $\vec{X}_0 \sim \pi$ and $\vec{V}_0 \sim \mathcal{N}(0, I)$, is defined as

$$\begin{cases} d\vec{X}_t = M^{-1}\vec{V}_t dt \\ d\vec{V}_t = \left(-\vec{X}_t - \Gamma M^{-1}\vec{V}_t\right) dt + \sqrt{2\Gamma} dB_t. \end{cases}$$

In this case, to ensure efficient and stable dynamics, we adopt the critical damping regime, where $\Gamma^2 = 4M$. This choice leads to fast convergence without oscillations [McCall, 2010]. The corresponding time-reversed SDE is

$$\begin{cases} d\vec{X}_t = -4\Gamma^{-2}\vec{V}_t dt \\ d\vec{V}_t = \left(\vec{X}_t + 4\Gamma^{-1}\vec{V}_t + 2\Gamma\nabla_v \log q_{T-t}(\vec{X}_t, \vec{V}_t)\right) dt + \sqrt{2\Gamma} dB_t, \end{cases} \quad (8)$$

where q_t is the marginal distribution of the forward process which has the following expression

$$q_t(x, v) = \int \pi(y)\varphi(v_0)\rho_t(x, v|y, v_0) dy dv_0,$$

with $\varphi = \mathcal{N}(0, I)$ and conditional distribution $\rho_t(x, v|y, v_0)$ normally distributed with mean $m_t(y, v_0)$ and covariance Σ_t given by

$$m_t(y, v_0) = e^{At} \otimes I \begin{pmatrix} y \\ v_0 \end{pmatrix}, \quad \Sigma_t = \int_0^t e^{A(t-s)} \begin{pmatrix} 0 & 0 \\ 0 & 2\Gamma \end{pmatrix} [e^{A(t-s)}]^\top ds \otimes I, \quad A = \begin{pmatrix} 0 & 4\Gamma^{-2} \\ -1 & -4\Gamma^{-1} \end{pmatrix}.$$

We observe that v_0 can be analytically integrated out in the expression for $\rho_t(x, v|y, v_0)$. Moreover, the conditional distribution $\rho_t(v|x, y)$ remains Gaussian, with mean $\tilde{m}_t(x, y)$ and covariance $\sigma_t^2 I$, where $\tilde{m}_t(x, y)$ is linear in y . This structure allows us to express the score $\nabla_v \log q_t(x, v)$ as

$$\nabla_v \log q_t(x, v) = \nabla_v \log q_t(v|x) = -\frac{v - \mathbb{E}_{Y|x, v}[\tilde{m}_t(x, Y)]}{\sigma_t^2} = -\frac{v - f_t(x, \mathbb{E}_{Y|x, v}[Y])}{\sigma_t^2}, \quad (9)$$

where f_t is a linear function of its arguments, see App. B.2 for a detailed derivation. Using the expressions above, we define a multiscale SDE system that enables sampling from the target distribution via the reverse SDE (8), without requiring prior estimation of the denoiser $\mathbb{E}_{Y|x, v}[Y]$ that appears in the expression for $\nabla_v \log q_{T-t}(x, v)$ (9). In this formulation, the fast process is modelled by an overdamped Langevin SDE that targets the conditional distribution $Y | (\tilde{X}_t, \tilde{V}_t)$, given by

$$dY_t = \frac{1}{\varepsilon} \left(\nabla \log \pi(Y_t) + \nabla_{Y_t} \log \rho_{T-t}(\tilde{X}_t, \tilde{V}_t | Y_t) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t.$$

Combining this fast process with the reverse SDE, we obtain the following stochastic slow-fast system, which we use for sampling from the target distribution

$$\begin{cases} d\tilde{X}_t = -4\Gamma^{-2} \tilde{V}_t dt \\ d\tilde{V}_t = \left(\tilde{X}_t + 4\Gamma^{-1} \tilde{V}_t - \frac{2\Gamma}{\sigma_{T-t}^2} \left(\tilde{V}_t - f_{T-t}(\tilde{X}_t, Y_t) \right) \right) dt + \sqrt{2\Gamma} dB_t \\ dY_t = \frac{1}{\varepsilon} \left(\nabla \log \pi(Y_t) + \nabla_{Y_t} \log \rho_{T-t}(\tilde{X}_t, \tilde{V}_t | Y_t) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t. \end{cases} \quad (10)$$

To implement this sampler, referred to as MULTCDIFF, we construct a novel integrator inspired by Dockhorn et al. [2022b]. Specifically, we leverage the symmetric Trotter splitting and the Baker–Campbell–Hausdorff formula [Strang, 1968, Trotter, 1959, Tuckerman, 2010] to design a stable and efficient symmetric splitting scheme, see App. B.2 for details. We also outline in the appendix the difficulties of extending controlled diffusions to the heavy-tailed scenario.

4 Theoretical guarantees

In this section, we identify the different sources of error in the proposed sampling algorithms. For MULTALMC, the total error consists of three components: the discretisation error from numerically solving the slow-fast system, the convergence error of the slow-fast system to its averaged dynamics, and the bias arising from the time-inhomogeneous averaged system, whose marginals differ from those of the diffusion path in (1)—this last component is discussed in 4.2. For MULTCDIFF, the error bound includes an initialisation error present in the error bounds of diffusion models [Benton et al., 2024, Chen et al., 2023], the discretisation error of the slow-fast system, and the convergence of the multiscale system to the averaged dynamics. The discretisation error depends on the choice of numerical integrator, underscoring the importance of an accurate and stable integrator scheme. We do not analyse this error in detail as it follows established methods [Leimkuhler et al., 2024, Monmarché, 2020]. We now study the convergence properties of the slow component X_t to the solution of the corresponding averaged dynamics that evolves along the time-inhomogeneous diffusion path in (1).

4.1 Convergence of the slow-fast system to the averaged dynamics

Building on results from stochastic averaging theory [Liu et al., 2020], our goal is to derive sufficient conditions on the coefficients of the multiscale systems—namely, (4), (6), and (10)—that guarantee strong convergence of the slow component to the averaged process as $\varepsilon \rightarrow 0$. We focus on the case where the base distribution is Gaussian $\nu \sim \mathcal{N}(0, I)$. Extending the convergence analysis to the case where ν is a heavy-tailed distribution, such as a Student’s t distribution, presents significant technical challenges, which fall outside of the scope of this work. We outline these challenges in App. C.2. To establish convergence guarantees, we assume the following regularity conditions.

A1. *The target distribution π has density with respect to the Lebesgue measure, which we write $\pi \propto e^{-V_\pi}$. The potential V_π has Lipschitz continuous gradients, with Lipschitz constant L_π . In addition, V_π is strongly convex with convexity parameter $M_\pi > 0$.*

A2. *The schedule $\lambda_t : \mathbb{R}^+ \rightarrow [0, 1]$ is a non-decreasing function of t , weakly differentiable and Hölder continuous with exponent γ_1 in $(0, 1]$.*

Under these assumptions, we obtain the following result.

Theorem 4.1. *Let the base distribution $\nu \sim \mathcal{N}(0, I)$. Suppose the target distribution π and the schedule function λ_t satisfy assumptions A1 and A2, respectively. Then, for any $\varepsilon \in (0, \varepsilon_0)$, where ε_0 is specified in App. C.1, and any given initial conditions, there exists unique solutions $\{(X_t^\varepsilon, Y_t^\varepsilon), t \geq 0\}$, $\{(X_t^\varepsilon, V_t^\varepsilon, Y_t^\varepsilon), t \geq 0\}$ and $\{\langle \bar{X}_t^\varepsilon, \bar{V}_t^\varepsilon, Y_t^\varepsilon \rangle, t \geq 0\}$ to the slow-fast stochastic systems (4), (6) and (10), respectively. Furthermore, for any $p > 0$, it holds that*

$$\lim_{\varepsilon \rightarrow 0} \mathbb{E} \left(\sup_{t \in [0, T]} \|X_t^\varepsilon - \bar{X}_t\|^p \right) = 0,$$

where \bar{X}_t denotes the solution of the averaged system and X_t^ε is the corresponding slow component of the multiscale systems (4), (6) and (10).

The proof is provided in App. C.1. The theorem applies to both the annealed Langevin dynamics and the controlled diffusion. We note that the strong convexity condition in A1—which guarantees exponential convergence of the fast process to its unique invariant measure when $\varepsilon = 1$ is fixed—is a restrictive assumption. However, as shown in Section 5, our proposed algorithms show strong empirical performance on benchmark distributions that do not satisfy this condition. Extending the theoretical analysis to cover non-log-concave targets remains an important future direction.

4.2 Bias of the annealed Langevin dynamics

The bias of the overdamped annealed Langevin dynamics (3) has been studied in prior work [Cordero-Encinar et al., 2025, Guo et al., 2025]. Here, we focus on quantifying the bias introduced by the underdamped averaged system (5) relative to the true diffusion path in the augmented state space with the velocity. At fixed time t , the invariant distribution of the underdamped averaged system takes the form

$$\hat{p}_t(x, v) \propto \exp \left(-\frac{1}{2} v^\top M^{-1} v + \log \hat{\mu}_t(x) \right). \quad (11)$$

Note that the Hamiltonian is separable, meaning that x and v are independent. It is important to emphasise that, even when simulated exactly, diffusion annealed Langevin dynamics introduces a bias, as the marginal distributions of the solution of (5) do not exactly match the prescribed distributions $(\hat{p}_t)_t$. A key quantity for characterising this bias will be the action of the curve of probability measures $\mu = (\mu_t)_{t \in [0, 1]}$ defined in (1) interpolating between the base distribution ν and the target distribution π , denoted by $\mathcal{A}(\mu)$. As noted by Guo et al. [2025], the action serves as a measure of the cost of transporting ν to π along the specified path. Formally, for an absolutely continuous curve of probability measures [Lisini, 2007] with finite second-order moment, the action is given by

$$\mathcal{A}(\mu) := \int_0^T \left(\lim_{\delta \rightarrow 0} |\delta|^{-1} W_2(\mu_{t+\delta}, \mu_t) \right)^2 dt.$$

Based on Theorem 1 from Guo et al. [2025], we have the following result.

Theorem 4.2. *Let $\mathbb{Q}_{\text{U-ALD}} = (q_{t, \text{U-ALD}})_{t \in [0, 1/\kappa]}$ be the path measure of the diffusion annealed Langevin dynamics (5), and $\mathbb{Q} = (\hat{p}_t)_{t \in [0, 1/\kappa]}$ that of a reference SDE such that the marginals at each time have distribution \hat{p}_t (11). If $q_{0, \text{U-ALD}} = \hat{p}_0$, the KL divergence between the path measures is given by*

$$\text{KL}(\mathbb{Q} \parallel \mathbb{Q}_{\text{U-ALD}}) = \kappa \mathcal{A}(\mu)/4.$$

The action $\mathcal{A}(\mu)$ is bounded when the target π has bounded second order moment and under mild conditions on the schedule function, see Cordero-Encinar et al. [2025, Lemmas 3.3 and 4.2]. Full details and a complete proof are provided in App. C.3.

5 Numerical experiments

We now evaluate the performance of both proposed underdamped samplers, based on the annealed Langevin dynamics (Section 3.1) and the controlled diffusions formulations (Section 3.2) across a range of benchmark distributions. Full details of each benchmark are provided in App. D.

- **Mixture of Gaussians (MoG) and mixture of Student’s t (MoS) distributions.**
- **Rings:** Two-dimensional distribution supported on a circular manifold.
- **Funnel:** 10-dimensional distribution defined by $\pi(x) \propto \mathcal{N}(x_1; 0, \eta^2) \prod_{i=2}^d \mathcal{N}(x_i; 0, e^{x_1})$ for $x = (x_i)_{i=1}^{10} \in \mathbb{R}^{10}$ with $\eta = 3$ [Neal, 2003].
- **Double well potential (DW):** d -dimensional distribution given by $\pi \propto \exp(-\sum_{i=1}^m (x_i^2 - \delta)^2 - \sum_{i=m+1}^d x_i^2)$ with $m \in \mathbb{N}$ and $\delta \in (0, \infty)$. Ground truth samples are obtained using rejection sampling with a Gaussian mixture proposal distribution [Midgley et al., 2023].
- **Examples from Bayesian statistics:** Posterior distributions arising from Bayesian logistic regression tasks on the Ionosphere (dimension 35) and Sonar (dimension 61) datasets.
- **Statistical physics model:** Sampling metastable states of the stochastic Allen-Cahn equation ϕ^4 model (dimension 100) [Albergo et al., 2019, Gabrié et al., 2022]. At the chosen temperature, the distribution has two well distinct modes with relative weight controlled by a parameter h . Following Grenioux et al. [2024], we estimate the relative weight between the two modes.

While we focus here on the underdamped versions of our algorithms, we have also evaluated their overdamped counterparts. These approximately require an order of magnitude more steps to achieve comparable performance. Detailed results and comparisons are presented in App. E.1. Besides, we have set the base distribution, denoted as ν in our algorithms, to be a standard Gaussian. In the case of the mixture of Student’s t benchmark distributions, we compare the performance of using either a standard Gaussian or a Student’s t base distribution. For the latter, we further examine the impact of using a modified Itô diffusion for the fast dynamics (7), as opposed to a standard Langevin diffusion.

We compare our approach against a representative selection of related sampling methods: Sequential Monte Carlo (SMC) [Del Moral et al., 2006, Doucet et al., 2001], Annealed Importance Sampling (AIS) [Neal, 2001], Underdamped Langevin Monte Carlo (ULMC) [Cheng et al., 2018, Neal et al., 2011], Parallel Tempering (PT) [Geyer and Thompson, 1995, Swendsen and Wang, 1986], Diffusive Gibbs Sampling (DiGS) [Chen et al., 2024], Reverse Diffusion Monte Carlo (RDMC) [Huang et al., 2024] and Stochastic Localisation via Iterative Posterior Sampling (SLIPS) [Grenioux et al., 2024]. We ensure a fair comparison by using the same number of energy evaluations and the same initialisation based on a standard Gaussian distribution for all baselines. For completeness, we also report the performance of SLIPS with optimal initialisation (App. D)—which assumes knowledge of the distribution’s scalar variance, unavailable in practice. For the ϕ^4 model, we only compare our results against a Laplace approximation and SLIPS algorithm, as the other methods provide degenerate samples. Since our approach is learning-free, we exclude comparisons with neural-based samplers as it would be difficult to make comparisons at equal computational budget.

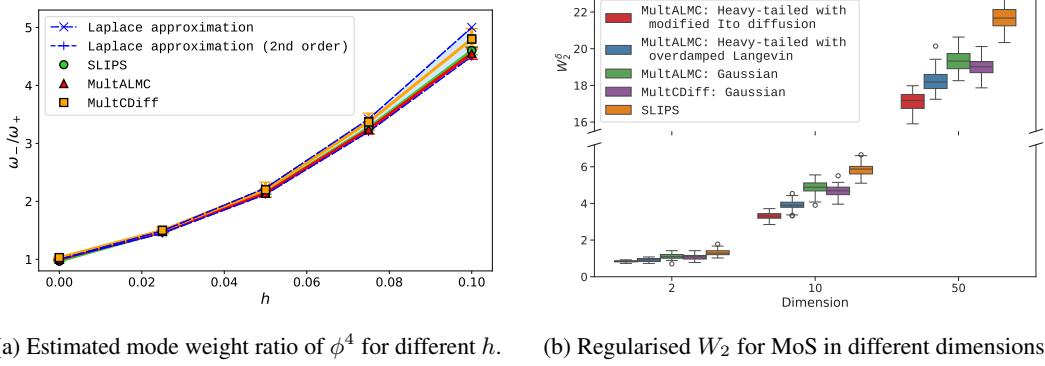
Our methods consistently achieve strong performance across all benchmarks, outperforming all baselines. They also exhibit low variance in results (see Tables 1 and 2). Among our samplers, the one based on controlled diffusions, MULTCDIFF, generally achieves slightly better results than the annealed Langevin-based sampler, MULTALMC. Additionally, Figure 1a illustrates the estimation of the relative weight between the two modes for the ϕ^4 model, where our algorithms demonstrate

Table 1: Metrics for different benchmarks averaged across 30 seeds. The metric for the mixture of Gaussian (MoG) and Rings is the entropy regularised Wasserstein-2 distance (with regularisation parameter 0.05) and the metric for the Funnel is the sliced Kolmogorov-Smirnov distance.

Algorithm	8-MoG (\downarrow) ($d = 2$)	40-MoG (\downarrow) ($d = 2$)	40-MoG (\downarrow) ($d = 50$)	Rings (\downarrow) ($d = 2$)	Funnel (\downarrow) ($d = 10$)
SMC	5.26 ± 0.19	4.79 ± 0.16	52.17 ± 1.32	0.29 ± 0.05	0.034 ± 0.005
AIS	5.53 ± 0.23	5.01 ± 0.20	65.83 ± 1.66	0.20 ± 0.03	0.040 ± 0.006
ULMC	6.90 ± 0.28	8.17 ± 0.72	109.26 ± 2.90	0.35 ± 0.06	0.123 ± 0.011
PT	3.91 ± 0.10	0.92 ± 0.08	21.04 ± 0.55	0.20 ± 0.04	0.033 ± 0.004
DiGS	1.22 ± 0.09	0.95 ± 0.03	21.87 ± 0.97	0.20 ± 0.02	0.037 ± 0.005
RDMC	1.01 ± 0.21	0.98 ± 0.07	36.10 ± 0.62	0.33 ± 0.01	0.080 ± 0.007
SLIPS	1.15 ± 0.12	1.04 ± 0.06	23.71 ± 0.65	0.26 ± 0.01	0.039 ± 0.006
MULTALMC	0.65 ± 0.07	0.91 ± 0.04	20.58 ± 0.71	0.18 ± 0.02	0.032 ± 0.005
MULTCDIFF	0.62 ± 0.10	0.93 ± 0.06	20.13 ± 0.59	0.19 ± 0.03	0.031 ± 0.005

Table 2: Performance metrics for different sampling benchmarks averaged across 30 seeds. The metric for the double well potential (DW) is the entropy regularised Wasserstein-2 distance (with regularisation parameter 0.05) and the metric for the Bayesian logistic regression on Ionosphere and Sonar datasets is the average predictive posterior log-likelihood on a test dataset.

Algorithm	5-dim DW (\downarrow) ($m = 5, \delta = 4$)	10-dim DW (\downarrow) ($m = 5, \delta = 3$)	50-dim DW (\downarrow) ($m = 5, \delta = 2$)	Ionosphere (\uparrow) ($d = 35$)	Sonar (\uparrow) ($d = 61$)
SMC	4.06 ± 0.13	11.86 ± 0.70	28.92 ± 0.99	-87.74 ± 0.10	-111.00 ± 0.11
AIS	4.11 ± 0.17	12.30 ± 0.61	39.20 ± 1.07	-88.11 ± 0.13	-111.15 ± 0.17
ULMC	7.87 ± 0.51	25.21 ± 1.00	62.74 ± 1.52	-116.37 ± 1.85	-173.61 ± 2.48
PT	2.39 ± 0.20	4.97 ± 0.43	16.13 ± 0.89	-87.91 ± 0.09	-112.99 ± 0.10
DiGS	2.51 ± 0.18	5.02 ± 0.42	17.04 ± 0.95	-88.02 ± 0.12	-110.53 ± 0.25
RDMC	3.68 ± 0.32	9.57 ± 0.75	25.21 ± 0.73	-108.44 ± 1.13	-130.20 ± 1.36
SLIPS	2.46 ± 0.21	5.09 ± 0.39	18.15 ± 0.50	-87.34 ± 0.11	-110.14 ± 0.10
MULTALMC	2.08 ± 0.16	4.48 ± 0.40	14.03 ± 0.63	-86.85 ± 0.09	-109.05 ± 0.13
MULTCDIFF	1.95 ± 0.24	4.23 ± 0.37	13.98 ± 0.56	-86.33 ± 0.10	-109.60 ± 0.21



(a) Estimated mode weight ratio of ϕ^4 for different h . (b) Regularised W_2 for MoS in different dimensions.

Figure 1: Results for sampling benchmarks.

competitive performance. For the mixture of Student’s t-distributions, Figure 1b shows that using the modified Itô diffusion for the fast dynamics, combined with a Student’s t base distribution ν (red boxplots) outperforms both our other proposed algorithms and the baselines. Further experimental details and additional results are provided in Appendices D and E.

6 Discussion

In this work, we introduce a framework based on multiscale diffusions for sampling from an unnormalised target density. In particular, we propose two samplers, MULTALMC and MULTCDIFF depending on the dynamics used for the slow process: annealed Langevin dynamics or controlled diffusions, respectively. We establish theoretical guarantees for the convergence of these sampling algorithms and illustrate their performance on a range of high-dimensional benchmark distributions.

Our approach has certain limitations. Notably, the theoretical guarantees rely on stringent assumptions, which we aim to relax in future work. Additionally, the current method requires manual tuning of hyperparameters such as step size δ , scale separation parameters ε and friction coefficient Γ . Automating this tuning process remains an important direction for future research. Further research could explore extending the controlled diffusion framework to heavy-tailed target distributions, which pose additional challenges or developing more efficient numerical schemes for implementing the proposed multiscale samplers.

Acknowledgements

PCE gratefully acknowledges support from the EPSRC through the Centre for Doctoral Training in Modern Statistics and Statistical Machine Learning (StatML), grant no. EP/S023151/1. SR's work was partially funded by the Deutsche Forschungsgemeinschaft (DFG) under Project-ID 318763901 – SFB1294.

References

Assyr Abdulle, Ibrahim Almuslimani, and Gilles Vilmart. Optimal Explicit Stabilized Integrator of Weak Order 1 for Stiff and Ergodic Stochastic Differential Equations. *SIAM/ASA Journal on Uncertainty Quantification*, 6(2):937–964, 2018.

O Deniz Akyildiz, Michela Ottobre, and Iain Souttar. A multiscale perspective on maximum marginal likelihood estimation. *arXiv preprint arXiv:2406.04187*, 2024.

M. S. Albergo, G. Kanwar, and P. E. Shanahan. Flow-based generative models for Markov chain Monte Carlo in lattice field theory. *Phys. Rev. D*, 100:034515, Aug 2019.

Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

Luigi Ambrosio and Bernd Kirchheim. Rectifiable sets in metric and Banach spaces. *Mathematische Annalen*, 318(3):527–555, 2000.

Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics. ETH Zürich. Birkhäuser Basel, 2008.

Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.

Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Nearly d -linear convergence bounds for diffusion models via stochastic localization. In *The Twelfth International Conference on Learning Representations*, 2024.

Richard Bertram and Jonathan E. Rubin. Multi-timescale systems and fast-slow analysis. *Mathematical Biosciences*, 287:105–121, 2017.

Denis Blessing, Xiaogang Jia, Johannes Esslinger, Francisco Vargas, and Gerhard Neumann. Beyond ELBOs: A large-scale evaluation of variational methods for sampling. In *Forty-first International Conference on Machine Learning*, 2024.

Denis Blessing, Julius Berner, Lorenz Richter, and Gerhard Neumann. Underdamped Diffusion Bridges with Applications to Sampling. In *The Thirteenth International Conference on Learning Representations*, 2025.

Nawaf Bou-Rabee and Andreas Eberle. Mixing time guarantees for unadjusted Hamiltonian Monte Carlo. *Bernoulli*, 29(1):75 – 104, 2023.

Omar Chehab and Anna Korba. A practical diffusion path for sampling. In *ICML 2024 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2024.

Junhua Chen, Lorenz Richter, Julius Berner, Denis Blessing, Gerhard Neumann, and Anima Anand-kumar. Sequential Controlled Langevin Diffusions. In *The Thirteenth International Conference on Learning Representations*, 2025.

Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *The Eleventh International Conference on Learning Representations*, 2023.

Wenlin Chen, Mingtian Zhang, Brooks Paige, José Miguel Hernández-Lobato, and David Barber. Diffusive Gibbs sampling. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 7731–7747. PMLR, 21–27 Jul 2024.

Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. Improved analysis for a proximal algorithm for sampling. In Po-Ling Loh and Maxim Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 2984–3014. PMLR, 02–05 Jul 2022.

Xiang Cheng, Niladri S. Chatterji, Peter L. Bartlett, and Michael I. Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 300–323. PMLR, 06–09 Jul 2018.

Sinho Chewi, Chen Lu, Kwangjun Ahn, Xiang Cheng, Thibaut Le Gouic, and Philippe Rigollet. Optimal dimension dependence of the Metropolis-Adjusted Langevin Algorithm. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 1260–1300. PMLR, 15–19 Aug 2021.

Paula Cordero-Encinar, O. Deniz Akyildiz, and Andrew B. Duncan. Non-asymptotic Analysis of Diffusion Annealed Langevin Monte Carlo for Generative Modelling. *arXiv preprint arXiv:2502.09306*, 2025.

Marco Cuturi, Laetitia Meng-Papaxanthos, Yingtao Tian, Charlotte Bunne, Geoff Davis, and Olivier Teboul. Optimal Transport Tools (OTT): A JAX Toolbox for all things Wasserstein. *arXiv preprint arXiv:2201.12324*, 2022.

Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.

Tim Dockhorn, Arash Vahdat, and Karsten Kreis. GENIE: Higher-Order Denoising Diffusion Solvers. In *Advances in Neural Information Processing Systems*, 2022a.

Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Score-Based Generative Modeling with Critically-Damped Langevin Diffusion. In *International Conference on Learning Representations (ICLR)*, 2022b.

Arnaud Doucet, Nando de Freitas, and Neil Gordon. *Sequential Monte Carlo Methods in Practice*. Springer New York, 2001.

A. B. Duncan, N. Nüsken, and G. A. Pavliotis. Using Perturbed Underdamped Langevin Dynamics to Efficiently Sample from Probability Distributions. *Journal of Statistical Physics*, 169(6): 1098–1131, 2017.

Alain Durmus and Éric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin Algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.

Andreas Eberle and Francis Lörler. Non-reversible lifts of reversible diffusion processes and relaxation times. *Probability Theory and Related Fields*, 2024.

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.

Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC., 3rd edition, 2013.

Charles J. Geyer and Elizabeth A. Thompson. Annealing Markov Chain Monte Carlo with Applications to Ancestral Inference. *Journal of the American Statistical Association*, 90(431):909–920, 1995.

Louis Grenioux, Alain Oliviero Durmus, Eric Moulines, and Marylou Gabrié. On sampling with approximate transport maps. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 11698–11733. PMLR, 23–29 Jul 2023.

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

Wei Guo, Molei Tao, and Yongxin Chen. Provable Benefit of Annealed Langevin Monte Carlo for Non-log-concave Sampling. In *The Thirteenth International Conference on Learning Representations*, 2025.

Emily Harvey, Vivien Kirk, Martin Wechselberger, and James Sneyd. Multiple Timescales, Mixed Mode Oscillations and Canards in Models of Intracellular Calcium Dynamics. *Journal of Nonlinear Science*, 21(5):639–683, 2011.

Ye He, Tyler Farghly, Krishnakumar Balasubramanian, and Murat A. Erdogdu. Mean-Square Analysis of Discretized Itô Diffusions for Heavy-tailed Sampling. *Journal of Machine Learning Research*, 25(43):1–44, 2024a.

Ye He, Kevin Rojas, and Molei Tao. Zeroth-Order Sampling Methods for Non-Log-Concave Distributions: Alleviating Metastability by Denoising Diffusion. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b.

Marlis Hochbruck and Alexander Ostermann. Exponential integrators. *Acta Numerica*, 19:209–286, 2010.

Peter Holderrieth, Yilun Xu, and Tommi Jaakkola. Hamiltonian Score Matching and Generative Flows. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Xunpeng Huang, Hanze Dong, Yifan Hao, Yian Ma, and Tong Zhang. Reverse Diffusion Monte Carlo. In *The Twelfth International Conference on Learning Representations*, 2024.

Aapo Hyvärinen. Estimation of Non-Normalized Statistical Models by Score Matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005.

Ioannis Karatzas and Steven E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer New York, NY, 1991.

Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Lower Bounds on Metropolized Sampling Methods for Well-Conditioned Distributions. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021a.

Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured Logconcave Sampling with a Restricted Gaussian Oracle. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 2993–3050. PMLR, 15–19 Aug 2021b.

Benedict Leimkuhler and Charles Matthews. Rational construction of stochastic numerical methods for molecular sampling. *Applied Mathematics Research eXpress*, 2013(1):34–56, 2013.

Benedict Leimkuhler, Daniel Paulin, and Peter A. Whalley. Contraction rate estimates of stochastic gradient kinetic Langevin integrators. *ESAIM: M2AN*, 58(6):2255–2286, 2024.

Tony Lelièvre, Mathias Rousset, and Gabriel Stoltz. *Free Energy Computations*. Imperial College Press, 2010.

Stefano Lisini. Characterization of absolutely continuous curves in Wasserstein spaces. *Calculus of Variations and Partial Differential Equations*, 28(1):85–120, 2007.

Jun S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer New York, NY, 2008.

Wei Liu, Michael Röckner, Xiaobin Sun, and Yingchao Xie. Averaging principle for slow-fast stochastic differential equations with time dependent locally Lipschitz coefficients. *Journal of Differential Equations*, 268(6):2910–2948, 2020.

Martin W. McCall. *Classical mechanics : from Newton to Einstein : a modern introduction*. Wiley, Hoboken, N.J., 2010.

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

Pierre Monmarché. High-dimensional MCMC with a standard splitting scheme for the underdamped Langevin diffusion. *Electronic Journal of Statistics*, 2020.

Alireza Mousavi-Hosseini, Tyler K. Farghly, Ye He, Krishna Balasubramanian, and Murat A. Erdogdu. Towards a Complete Analysis of Langevin Monte Carlo: Beyond Poincaré Inequality. In Gergely Neu and Lorenzo Rosasco, editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 1–35. PMLR, 12–15 Jul 2023.

Radford M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.

Radford M. Neal. Slice sampling. *The Annals of Statistics*, 31(3):705 – 767, 2003.

Radford M. Neal et al. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11):2, 2011.

Maxence Noble, Louis Grenioux, Marylou Gabrié, and Alain Oliviero Durmus. Learned reference-based diffusion sampler for multi-modal distributions. In *The Thirteenth International Conference on Learning Representations*, 2025.

Kushagra Pandey, Jaideep Pathak, Yilun Xu, Stephan Mandt, Michael Pritchard, Arash Vahdat, and Morteza Mardani. Heavy-Tailed Diffusion Models. In *The Thirteenth International Conference on Learning Representations*, 2025.

Grigoris A Pavliotis and Andrew Stuart. *Multiscale methods: averaging and homogenization*, volume 53. Springer Science & Business Media, 2008.

Gabriel Peyré and Marco Cuturi. Computational Optimal Transport: With Applications to Data Science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

Angus Phillips, Hai-Dang Dau, Michael John Hutchinson, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. Particle Denoising Diffusion Sampler. In *Forty-first International Conference on Machine Learning*, 2024.

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

Saeed Saremi, Ji Won Park, and Francis Bach. Chain of Log-Concave Markov Chains. In *The Twelfth International Conference on Learning Representations*, 2024.

Katharina Schuh and Iain Souttar. Conditions for uniform in time convergence: applications to averaging, numerical discretisations and mean-field systems. *arXiv preprint arXiv:2412.05239*, 2024.

Dario Shariatian, Umut Simsekli, and Alain Oliviero Durmus. Denoising Lévy Probabilistic Models. In *The Thirteenth International Conference on Learning Representations*, 2025.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.

Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12438–12448. Curran Associates, Inc., 2020.

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.

Gilbert Strang. On the Construction and Comparison of Difference Schemes. *SIAM Journal on Numerical Analysis*, 5(3):506–517, 1968.

Robert H. Swendsen and Jian-Sheng Wang. Replica Monte Carlo Simulation of Spin-Glasses. *Phys. Rev. Lett.*, 57:2607–2609, Nov 1986.

H. F. Trotter. On the Product of Semi-Groups of Operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.

Mark E Tuckerman. *Statistical mechanics: Theory and Molecular Simulation*. Oxford university press, 2010.

Francisco Vargas, Shreyas Padhy, Denis Blessing, and Nikolas Nüsken. Transport meets Variational Inference: Controlled Monte Carlo Diffusions. In *The Twelfth International Conference on Learning Representations*, 2024.

Fengyu Wang. *Functional inequalities Markov semigroups and spectral theory*. Elsevier, 2006.

E Weinan, Di Liu, and Eric Vanden-Eijnden. Analysis of multiscale methods for stochastic differential equations. *Communications on Pure and Applied Mathematics*, 58(11):1544–1585, 2005.

Keru Wu, Scott Schmidler, and Yuansi Chen. Minimax Mixing Time of the Metropolis-Adjusted Langevin Algorithm for Log-Concave Sampling. *Journal of Machine Learning Research*, 23(270):1–63, 2022.

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

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract summarises the main contributions. In addition, we describe these contributions in more detail at the end of Section 1. Each of our claims is clearly linked to the relevant sections of the paper for further detail.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We outline the limitations of our methodology in Section 6 and provide a more detailed discussion in Appendix F.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide full assumptions for all the theoretical results. Furthermore, the complete proofs of our results are included in Appendix C.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Detailed information necessary to reproduce the main experimental results—including the methodology and hyperparameter choices for both our algorithms and the baselines—is provided in Appendices B and D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code is available at the following anonymised link https://anonymous.4open.science/r/sampling_by_averaging-41F2.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The experimental setup for our algorithms is detailed in the main paper, specifically in Sections 3 and 5. Additional implementation details, including those for baseline methods, are provided in Section D of the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All numerical experiments report error estimates in the form of standard deviations computed over multiple runs with different random seeds.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: The compute resources and runtimes for the experiments are provided in Appendix D.4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We have read the code of conduct and ensure that the paper complies with it.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: We address broader impacts in Appendix G.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper proposes a general purpose algorithm for sampling from a probability distribution, which does not have potential for direct misuse or dual-use. We do not use any scraped datasets in the numerical experiments which could lead to potential misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We properly acknowledge and cite all assets and resources used in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper introduces a sampling method for unnormalised probability distributions and does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Preliminaries

Multiscale methods and stochastic averaging. Before presenting a formal definition, we first provide the intuition behind multiscale methods, following Schuh and Souttar [2024]. In many areas of science and engineering [Bertram and Rubin, 2017, Weinan et al., 2005], one often encounters systems that are difficult to analyse directly. However, there may exist a related system that is more tractable, though it does not produce exactly the same behaviour. In such cases, the simpler system can be viewed as an approximation of the more complex one. In our setting, the simpler system is a slow-fast system, while the complex system corresponds to the averaged dynamics. We develop these concepts in the following.

To conduct a theoretical analysis of the convergence of the proposed sampling algorithms, we examine the general stochastic slow-fast system studied in Liu et al. [2020]

$$\begin{cases} dX_t^\varepsilon = b(t, X_t^\varepsilon, Y_t^\varepsilon)dt + \sigma(t, X_t^\varepsilon)dB_t & X_0^\varepsilon = x \in \mathbb{R}^n \\ dY_t^\varepsilon = \frac{1}{\varepsilon}f(t, X_t^\varepsilon, Y_t^\varepsilon)dt + \frac{1}{\sqrt{\varepsilon}}g(t, X_t^\varepsilon, Y_t^\varepsilon)d\tilde{B}_t & Y_0^\varepsilon = y \in \mathbb{R}^m, \end{cases} \quad (12)$$

where ε is a small positive parameter describing the ratio of time scales between the slow component X_t^ε and the fast component Y_t^ε , and B_t and \tilde{B}_t are mutually independent standard Brownian motions on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and $\{\mathcal{F}_t, t \geq 0\}$ is the natural filtration generated by B_t and \tilde{B}_t . Let \bar{X}_t denote the solution of the following averaged equation

$$\begin{cases} d\bar{X}_t = \bar{b}(t, \bar{X}_t)dt + \sigma(t, \bar{X}_t)dB_t, \\ \bar{X}_0 = x, \end{cases}$$

where $\bar{b}(t, x) = \int_{\mathbb{R}^m} b(t, x, y) \rho_{t,x}(dy)$ and $\rho_{t,x}$ denotes the unique invariant measure for the transition semigroup of the corresponding frozen process, which can be informally defined as the fast process with $\varepsilon = 1$ and fixed slow dynamics $X_t^\varepsilon = x$,

$$\begin{cases} dY_s = f(t, x, Y_s)ds + g(t, x, Y_s)d\tilde{B}'_s \\ Y_0 = y, \end{cases}$$

where \tilde{B}'_s is a Brownian motion on another complete probability space.

We begin by presenting the assumptions underlying the analysis in Liu et al. [2020] and stating their main results, which will form the basis of our study in Appendix C.1.

The assumptions on the coefficients of the stochastic slow-fast system (12) are the following.

A3 (Coefficients of the slow process). (i) *There exists $\theta_1 \geq 0$ such that for any $t, R \geq 0$, $x_i \in \mathbb{R}^n$, $y \in \mathbb{R}^m$ with $\|x_i\| \leq R$,*

$$2\|b(t, x_1, y) - b(t, x_2, y)\| \|x_1 - x_2\| + \|\sigma(t, x_1, y) - \sigma(t, x_2, y)\|^2 \leq K_t(R)(1 + \|y\|^{\theta_1})\|x_1 - x_2\|^2,$$

where $K_t(R)$ is an \mathbb{R}_+ -valued \mathcal{F}_t -adapted process satisfying for all $R, T, p \in [0, \infty)$,

$$\alpha_T(R) := \int_0^T K_t(R)dt < \infty, \quad \text{on } \Omega, \quad \mathbb{E} \left[e^{p\alpha_T(1)} \right] < \infty, \quad \sup_{t \in [0, T]} \mathbb{E} \|K_t(1)\|^4 < \infty.$$

Furthermore, there exists $R_0 > 0$, such that for any $R \geq R_0$, $T \geq 0$,

$$\mathbb{E} \int_0^T [K_t(R)]^4 dt < \infty.$$

(ii) *There exist constants $\theta_2, \theta_3 \geq 1$ and $\gamma_1 \in (0, 1]$ such that for any $x \in \mathbb{R}^n$, $y, y_1, y_2 \in \mathbb{R}^m$ and $T > 0$ with $t, s \in [0, T]$,*

$$\|b(t, x, y_1) - b(t, x, y_2)\| \leq C_T \|y_1 - y_2\| \left[\|y_1\|^{\theta_2} + \|y_2\|^{\theta_2} + K_t(1) + \|x\|^{\theta_3} \right]$$

and

$$\|b(t, x, y) - b(s, x, y)\| \leq C_T \|t - s\|^{\gamma_1} \left[\|y\|^{\theta_2} + \|x\|^{\theta_3} + Z_T \right], \quad \text{on } \Omega,$$

where $C_T > 0$ and Z_t is some random variable satisfying $\mathbb{E} Z_T^2 < \infty$.

(iii) There exist $\lambda_1 \geq 0, C > 0, \theta_4 \geq 2$ and $\theta_5, \theta_6 \geq 1$ such that for any $t > 0, x \in \mathbb{R}^n, y \in \mathbb{R}^m$,

$$2\langle x, b(t, x, y) \rangle \leq K_t(1)(1 + \|x\|^2) + \lambda_1 \|y\|^{\theta_4},$$

and

$$\|b(t, x, y)\| \leq K_t(1) + C(\|x\|^{\theta_5} + \|y\|^{\theta_6}), \quad \|\sigma(t, x)\|^2 \leq K_t(1) + C\|x\|^2.$$

A4 (Coefficients of the fast process). (i) There exists $\beta \geq 0$ such that for any $t \geq 0, x \in \mathbb{R}^n, y_1, y_2 \in \mathbb{R}^m$,

$$2\langle f(t, x, y_1) - f(t, x, y_2), y_1 - y_2 \rangle + \|g(t, x, y_1) - g(t, x, y_2)\|^2 \leq -\beta \|y_1 - y_2\|^2.$$

(ii) For any $T > 0$, there exists $\gamma_2 \in (0, 1]$, $C_T > 0$, $\alpha_i \geq 1, i = 1, 2, 3, 4$ such that for any $t, s \in [0, T]$ and $x_i \in \mathbb{R}^n, y_i \in \mathbb{R}^m, i = 1, 2$,

$$\|f(t, x_1, y_1) - f(s, x_2, y_1)\| \leq C_T(|t - s|^{\gamma_2} + \|x_1 - x_2\|)(1 + \|x_1\|^{\alpha_1} + \|x_2\|^{\alpha_1} + \|y_1\|^{\alpha_2}),$$

$$\|g(t, x_1, y_1) - g(s, x_2, y_2)\| \leq C_T(|t - s|^{\gamma_2} + \|x_1 - x_2\| + \|y_1 - y_2\|),$$

$$\|f(t, x_1, y_1)\| \leq C_T(1 + \|x_1\|^{\alpha_3} + \|y_1\|^{\alpha_4}),$$

$$\|g(t, x_1, y_1)\| \leq C_T(1 + \|x_1\| + \|y_1\|).$$

(iii) For some fixed $k \geq 2$ and any $T > 0$, there exist $C_{T,k}, \beta_k > 0$ such that for any $t \in [0, T], x \in \mathbb{R}^n, y \in \mathbb{R}^m$,

$$2\langle y, f(t, x, y) \rangle + (k-1)\|g(t, x, y)\|^2 \leq -\beta_k \|y\|^2 - \lambda_2 \|y\|^{\theta_4} + C_{T,k} \left(\|x\|^{\frac{4}{\theta_4}} + 1 \right),$$

where $\lambda_2 = 0$ if $\lambda_1 = 0$, and $\lambda_2 > 0$ otherwise.

Note that assumption A4 on the coefficients of the fast process guarantees that the frozen process has a unique stationary measure and that the fast process converges to it exponentially fast. The following theorems establish the existence and uniqueness of solutions to the system in (12), as well as its convergence to the corresponding averaged dynamics.

Theorem A.1 (Theorem 2.2 [Liu et al., 2020]). *If assumptions A3 and A4 hold with $\lambda_1 > 0$, let $\varepsilon_0 = \frac{\lambda_2}{\lambda_1}$. Then for any $\varepsilon \in (0, \varepsilon_0)$ and any given initial values $x \in \mathbb{R}^n, y \in \mathbb{R}^m$, there exists a unique solution $\{(X_t^\varepsilon, Y_t^\varepsilon), t \geq 0\}$ to the system (12) and for all $T > 0$, we have $(X_t^\varepsilon, Y_t^\varepsilon) \in C([0, T]; \mathbb{R}^n) \times C([0, T]; \mathbb{R}^m)$, \mathbb{P} -almost surely. Furthermore, for all $t \in [0, T]$, the solution $(X_t^\varepsilon, Y_t^\varepsilon)$ is given by*

$$\begin{cases} X_t^\varepsilon = x + \int_0^t b(s, X_s^\varepsilon, Y_s^\varepsilon) ds + \int_0^t \sigma(s, X_s^\varepsilon) dB_s \\ Y_t^\varepsilon = y + \frac{1}{\varepsilon} \int_0^t f(s, X_s^\varepsilon, Y_s^\varepsilon) ds + \frac{1}{\sqrt{\varepsilon}} \int_0^t g(s, X_s^\varepsilon, Y_s^\varepsilon) d\tilde{B}_s. \end{cases}$$

Theorem A.2 (Theorem 2.3 [Liu et al., 2020]). *If assumptions A3 and A4 hold with $\lambda_1 > 0$ and $k > \tilde{\theta}_2$ where $\tilde{\theta}_2 = \max\{4\theta_1, 2\theta_2 + 2, 2\theta_6, 4\alpha_2, \theta_4\theta_5, 2\alpha_1\theta_4\}$. Then, for any $0 < p < \frac{2k}{\theta_4}$ we have*

$$\lim_{\varepsilon \rightarrow 0} \mathbb{E} \left(\sup_{t \in [0, T]} \|X_t^\varepsilon - \bar{X}_t\|^2 \right) = 0.$$

Optimal transport. Let $v = (v_t : \mathbb{R}^d \rightarrow \mathbb{R}^d)$ be a vector field and $\mu = (\mu_t)_{t \in [a, b]}$ be a curve of probability measures on \mathbb{R}^d with finite second-order moments. μ is generated by the vector field v if the continuity equation

$$\partial_t \mu_t + \nabla \cdot (\mu_t v_t) = 0,$$

holds for all $t \in [a, b]$. The metric derivative of μ at $t \in [a, b]$ is then defined as

$$|\dot{\mu}|_t := \lim_{\delta \rightarrow 0} \frac{W_2(\mu_{t+\delta}, \mu_t)}{|\delta|}.$$

If $|\dot{\mu}|_t$ exists and is finite for all $t \in [a, b]$, we say that μ is an absolutely continuous curve of probability measures. Ambrosio and Kirchheim [2000] establish weak conditions under which a curve of probability measures with finite second-order moments is absolutely continuous.

By Ambrosio et al. [2008, Theorem 8.3.1] we have that among all velocity fields v_t which produce the same flow μ , there is a unique optimal one with smallest $L^p(\mu_t; X)$ -norm. This is summarised in the following lemma.

Lemma A.3 (Lemma 2 from Guo et al. [2025]). *For an absolutely continuous curve of probability measures $\mu = (\mu_t)_{t \in [a,b]}$, any vector field $(v_t)_{t \in [a,b]}$ that generates μ satisfies $|\dot{\mu}|_t \leq \|v_t\|_{L^2(\mu_t)}$ for almost every $t \in [a,b]$. Moreover, there exists a unique vector field v_t^* generating μ such that $|\dot{\mu}|_t = \|v_t^*\|_{L^2(\mu_t)}$ almost everywhere.*

We also introduce the action of the absolutely continuous curve $(\mu_t)_{t \in [a,b]}$ since it will play a key role in our convergence results. In particular, we define the action $\mathcal{A}(\mu)$ as

$$\mathcal{A}(\mu) := \int_a^b |\dot{\mu}|_t^2 dt.$$

Girsanov's theorem. Consider the SDE

$$dX_t = b(X_t, t)dt + \sigma(X_t, t)dB_t,$$

for $t \in [0, T]$, where $(B_t)_{t \in [0, T]}$ is a standard Brownian motion in \mathbb{R}^d . Denote by \mathbb{P}^X the *path measure* of the solution $X = (X_t)_{t \in [0, T]}$ of the SDE, which characterises the distribution of X over the sample space Ω .

The KL divergence between two path measures can be characterised as a consequence of Girsanov's theorem [Karatzas and Shreve, 1991]. In particular, the following result will be central in our analysis.

Lemma A.4. *Consider the following two SDEs defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$*

$$dX_t = a_t(X)dt + \sqrt{2}dB_t, \quad dY_t = b_t(Y)dt + \sqrt{2}dB_t, \quad t \in [0, T]$$

with the same initial conditions $X_0, Y_0 \sim \mu_0$. Denote by \mathbb{P}^X and \mathbb{P}^Y the path measures of the processes X and Y , respectively. It follows that

$$\text{KL}(\mathbb{P}^X \parallel \mathbb{P}^Y) = \frac{1}{4} \mathbb{E}_{X \sim \mathbb{P}^X} \left[\int_0^T \|a_t(X) - b_t(X)\|^2 dt \right].$$

B Sampling using multiscale dynamics

In this section, we provide a detailed description of the implementation of our proposed algorithms including the numerical discretisation schemes used.

B.1 MULTALMC: Multiscale Annealed Langevin Monte Carlo

Here, we present further details of the overdamped and underdamped versions of the sampler.

B.1.1 Overdamped system

In the overdamped setting, we propose a new sampling strategy based on the following stochastic slow-fast system

$$\begin{cases} dX_t &= \begin{cases} \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1-\lambda_{\kappa t}}} \right) dt + \sqrt{2}dB_t & \text{if } \lambda_{\kappa t} < \tilde{\lambda} \\ -\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) dt + \sqrt{2}dB_t & \text{if } \lambda_{\kappa t} \geq \tilde{\lambda} \end{cases} \\ dY_t &= \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) + \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1-\lambda_{\kappa t}}} \right) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t. \end{cases}$$

As discussed in the main text, the conditional distribution $\rho_{t,x}$ converges to a Dirac delta centred at 0 when $\lambda_{\kappa t} = 0$, and to a Dirac at x when $\lambda_{\kappa t} = 1$. These degenerate limits can cause numerical instabilities at the endpoints of the diffusion schedule. To mitigate this, we define $0 < \lambda_\delta \ll 1$ and consider the modified dynamics

$$\begin{cases} dX_t &= \begin{cases} \frac{1}{\sqrt{1-\lambda_\delta}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1-\lambda_\delta}} \right) dt + \sqrt{2}dB_t & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1-\lambda_{\kappa t}}} \right) dt + \sqrt{2}dB_t & \text{if } \lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda} \\ -\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) dt + \sqrt{2}dB_t & \text{if } \tilde{\lambda} \leq \lambda_{\kappa t} < 1 - \lambda_\delta \\ -\nabla V_\pi(X_t) dt + \sqrt{2}dB_t & \text{if } 1 - \lambda_\delta \leq \lambda_{\kappa t} \leq 1 \end{cases} \\ dY_t &= \begin{cases} \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_\delta}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_\delta}} \right) + \frac{1}{\sqrt{1-\lambda_\delta}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1-\lambda_\delta}} \right) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) + \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1-\lambda_{\kappa t}}} \right) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t & \text{if } \lambda_\delta \leq \lambda_{\kappa t} \leq 1. \end{cases} \end{cases} \quad (13)$$

Note that if the slow-fast system converges to its corresponding averaged dynamics, then running the modified multiscale system up to time t_δ where $\lambda_{\kappa t_\delta} = 1 - \lambda_\delta$, we can approximate the distribution $\hat{\mu}_{t_\delta}$ that is close to the target π , provided λ_δ is sufficiently small. For $t \geq t_\delta$, the dynamics in (13) reduce to standard overdamped Langevin dynamics targetting π , initialised at $\hat{\mu}_{t_\delta}$. This warm start is known to significantly improve convergence rates to the target distribution [Chewi et al., 2021, Lee et al., 2021a, Wu et al., 2022].

We now analyse different numerical schemes for implementing the sampler in practice. As we have just mentioned, when $\lambda_{\kappa t} \geq 1 - \lambda_\delta$, the dynamics in (13) reduce to standard overdamped Langevin dynamics, which can be discretised using any preferred numerical method. Therefore, we focus on the regime $0 \leq \lambda_{\kappa t} < 1 - \lambda_\delta$.

For general base distributions ν , we propose to use the following numerical scheme that combines Euler-Maruyama for the slow dynamics with the SROCK method [Abdulle et al., 2018] for the fast dynamics. Before presenting the complete algorithm in Algorithm 2, we introduce the necessary notation and coefficients for the SROCK update.

Denote by T_s the Chebychev polynomials of first kind and define the coefficients in the SROCK update as follows

$$\omega_0 = 1 + \frac{\eta}{s^2}, \quad \omega_1 = \frac{T_s(\omega_0)}{T'_s(\omega_0)}, \quad \mu_1 = \frac{\omega_1}{\omega_0}, \quad (14)$$

and for all $i = 2, \dots, s$,

$$\mu_i = \frac{2\omega_1 T_{i-1}(\omega_0)}{T_i(\omega_0)}, \quad \nu_i = \frac{2\omega_0 T_{i-1}(\omega_0)}{T_i(\omega_0)}, \quad \kappa_i = 1 - \nu_i. \quad (15)$$

Additionally, when $\nu \sim \mathcal{N}(0, I)$, we can further exploit the linear structure of the score function for a Gaussian distribution to design an efficient exponential integrator scheme. Below, we derive the corresponding update rules, distinguishing between the regimes $0 \leq \lambda_{\kappa t} < \tilde{\lambda}$ and $\tilde{\lambda} \leq \lambda_{\kappa t} \leq 1$.

Regime 1: $0 \leq \lambda_{\kappa t} < \tilde{\lambda}$. Define the modified schedule

$$\lambda'_{\kappa t} = \begin{cases} \lambda_\delta & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \lambda_{\kappa t} & \text{if } \lambda_\delta \leq \lambda_{\kappa t} \leq 1. \end{cases}$$

The slow-fast system (13) simplifies to

$$\begin{cases} dX_t = -\frac{X_t - Y_t}{1 - \lambda'_{\kappa t}} dt + \sqrt{2} dB_t \\ dY_t = \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda'_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda'_{\kappa t}}} \right) - \frac{Y_t - X_t}{1 - \lambda'_{\kappa t}} \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t. \end{cases}$$

The exponential integrator scheme is then expressed as

$$d \begin{pmatrix} X_t \\ Y_t \end{pmatrix} = -\frac{1}{1 - \lambda'_{\kappa t}} \begin{pmatrix} 1 & -1 \\ -\varepsilon^{-1} & \varepsilon^{-1} \end{pmatrix} \otimes I_d \begin{pmatrix} X_t \\ Y_t \end{pmatrix} dt - \frac{1}{\varepsilon \sqrt{\lambda'_{\kappa t}}} \left(\nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda'_{\kappa t}}} \right) \right) dt + \sqrt{2} \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{\varepsilon^{-1}} \end{pmatrix} \otimes I_d \begin{pmatrix} dB_t \\ d\tilde{B}_t \end{pmatrix},$$

where given a time discretisation $0 \leq T_0 < \dots < T_M$ of the corresponding time interval, we define $t_- = T_{l-1}$ when $t \in [T_{l-1}, T_l)$. The explicit update rule is then

$$\begin{pmatrix} X_{l+1} \\ Y_{l+1} \end{pmatrix} = A_0(T_l, T_{l+1}) \otimes I_d \begin{pmatrix} X_l \\ Y_l \end{pmatrix} - \frac{A_1(T_l, T_{l+1})}{\varepsilon} \otimes I_d \left(\frac{\mathbf{0}_d}{\nabla V_\pi \left(\frac{Y_l}{\sqrt{\lambda'_{T_l}}} \right)} \right) + A_2(T_l, T_{l+1}) \otimes I_d \xi_l,$$

where $\xi_l \sim \mathcal{N}(0, I_{2d})$ and

$$\begin{aligned} A_0(a, b) &= \exp \left(- \begin{pmatrix} 1 & -1 \\ -\varepsilon^{-1} & \varepsilon^{-1} \end{pmatrix} \int_a^b (1 - \lambda'_{\kappa u})^{-1} du \right) \\ A_1(a, b) &= \int_a^b \frac{1}{\sqrt{\lambda'_{\kappa u}}} A_0(u, b) du \\ A_2(a, b) &= \sqrt{2 \int_a^b A_0(u, b) \begin{pmatrix} 1 & 0 \\ 0 & \varepsilon^{-1} \end{pmatrix} A_0(u, b)^\top du}. \end{aligned}$$

Algorithm 2 MULTALMC sampler: overdamped version

Require: Schedule function λ_t , value for λ_δ and $\tilde{\lambda}$, number of sampling steps L , time discretisation $0 = T_0 < \dots < T_L = 1/\kappa$, step size $h_l = T_{l+1} - T_l$. Constants for SROCK step from (14), (15).

Initial samples $X_0 \sim \mathcal{N}(0, I)$, $Y_0 \sim \mathcal{N}(0, I)$. Define the schedule

$$\lambda'_{\kappa t} = \begin{cases} \lambda_\delta & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \lambda_{\kappa t} & \text{if } \lambda_\delta \leq \lambda_{\kappa t} \leq 1. \end{cases}$$

for $l = 0$ **to** L **do**

$$\xi_l^{(1)}, \xi_l^{(2)} \sim \mathcal{N}(0, I)$$

if $0 \leq \lambda_{\kappa T_l} < 1 - \lambda_\delta$ **then**

EM for slow dynamics

$$X_{l+1} = \begin{cases} X_l + \frac{1}{\sqrt{1-\lambda'_{\kappa T_l}}} \nabla \log \nu \left(\frac{X_l - Y_l}{\sqrt{1-\lambda'_{\kappa T_l}}} \right) h_l + \sqrt{2h_l} \xi_l^{(1)} & \text{if } \lambda'_{\kappa T_l} < \tilde{\lambda} \\ X_l - \frac{1}{\sqrt{\lambda'_{\kappa T_l}}} \nabla V_\pi \left(\frac{Y_l}{\sqrt{\lambda'_{\kappa T_l}}} \right) h_l + \sqrt{2h_l} \xi_l^{(1)} & \text{if } \lambda'_{\kappa T_l} \geq \tilde{\lambda} \end{cases}$$

SROCK for fast dynamics

$$\begin{cases} K_{l,0} = Y_l \\ K_{l,1} = K_{l,0} + \mu_1 \frac{h_l}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda'_{\kappa T_l}}} \nabla V_\pi \left(\frac{K_{l,0}}{\sqrt{\lambda'_{\kappa T_l}}} \right) + \frac{1}{\sqrt{1-\lambda'_{\kappa T_l}}} \nabla \log \nu \left(\frac{K_{l,0} - X_{l+1}}{\sqrt{1-\lambda'_{\kappa T_l}}} \right) \right) \\ \vdots \\ K_{l,i} = \mu_i \frac{h_l}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda'_{\kappa T_l}}} \nabla V_\pi \left(\frac{K_{l,i-1}}{\sqrt{\lambda'_{\kappa T_l}}} \right) + \frac{1}{\sqrt{1-\lambda'_{\kappa T_l}}} \nabla \log \nu \left(\frac{K_{l,i-1} - X_{l+1}}{\sqrt{1-\lambda'_{\kappa T_l}}} \right) \right) + \nu_i K_{l,i-1} + \kappa_i K_{l,i-2} \\ \vdots \\ Y_{l+1} = K_s + \sqrt{\frac{2h_l}{\varepsilon}} \xi_l^{(2)} \end{cases}$$

if $1 - \lambda_\delta \leq \lambda_{\kappa T_l} \leq 1$ **then**

$$X_{l+1} = X_l - h_l \nabla V_\pi(X_l) + \sqrt{2h_l} \xi_l^{(1)}$$

end for

return X_L

Regime 2: $\tilde{\lambda} \leq \lambda_{\kappa t} < 1 - \lambda_\delta$. Here, the dynamics become

$$\begin{cases} dX_t = -\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) dt + \sqrt{2} dB_t \\ dY_t = \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) - \frac{Y_t - X_t}{1 - \lambda_{\kappa t}} \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t. \end{cases}$$

In this case, the exponential integrator scheme can be formulated as follows

$$d \begin{pmatrix} X_t \\ Y_t \end{pmatrix} = -\frac{1}{1 - \lambda_{\kappa t}} \begin{pmatrix} 0 & 0 \\ -\varepsilon^{-1} & \varepsilon^{-1} \end{pmatrix} \otimes I_d \begin{pmatrix} X_t \\ Y_t \end{pmatrix} dt - \frac{1}{\sqrt{\lambda_{\kappa t}}} \begin{pmatrix} \nabla V_\pi \left(\frac{Y_{t-}}{\sqrt{\lambda_{t-}}} \right) \\ \frac{1}{\varepsilon} \nabla V_\pi \left(\frac{Y_{t-}}{\sqrt{\lambda_{t-}}} \right) \end{pmatrix} dt + \sqrt{2} \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{\varepsilon^{-1}} \end{pmatrix} \otimes I_d \begin{pmatrix} dB_t \\ d\tilde{B}_t \end{pmatrix},$$

The corresponding update rule is

$$\begin{pmatrix} X_{l+1} \\ Y_{l+1} \end{pmatrix} = \tilde{A}_0(T_l, T_{l+1}) \otimes I_d \begin{pmatrix} X_l \\ Y_l \end{pmatrix} - \tilde{A}_1(T_l, T_{l+1}) \otimes I_d \begin{pmatrix} \nabla V_\pi \left(\frac{Y_l}{\sqrt{\lambda_{T_l}}} \right) \\ \frac{1}{\varepsilon} \nabla V_\pi \left(\frac{Y_l}{\sqrt{\lambda_{T_l}}} \right) \end{pmatrix} + \tilde{A}_2(T_l, T_{l+1}) \otimes I_d \xi_l,$$

where $\xi_l \sim \mathcal{N}(0, I_{2d})$ and

$$\tilde{A}_0(a, b) = \exp \left(- \begin{pmatrix} 0 & 0 \\ -\varepsilon^{-1} & \varepsilon^{-1} \end{pmatrix} \int_a^b (1 - \lambda_{\kappa u})^{-1} du \right)$$

$$\begin{aligned}\tilde{A}_1(a, b) &= \int_a^b \frac{1}{\sqrt{\lambda_{\kappa u}}} \tilde{A}_0(u, b) du \\ \tilde{A}_2(a, b) &= \sqrt{2 \int_a^b \tilde{A}_0(u, b) \begin{pmatrix} 1 & 0 \\ 0 & \varepsilon^{-1} \end{pmatrix} \tilde{A}_0(u, b)^\top du}.\end{aligned}$$

B.1.2 Underdamped system

Similarly in the overdamped case, to mitigate numerical instabilities arising from the degeneracy of $\rho_{t,x}$ when $\lambda_{\kappa t} = 0$ and $\lambda_{\kappa t} = 1$, we consider the following modified dynamics.

$$\begin{cases} dX_t &= M^{-1}V_t dt \\ dV_t &= \begin{cases} \left(\frac{1}{\sqrt{1-\lambda_\delta}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1-\lambda_\delta}} \right) - \Gamma M^{-1} V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \left(\frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1-\lambda_{\kappa t}}} \right) - \Gamma M^{-1} V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } \lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda} \\ \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) - \Gamma M^{-1} V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } \tilde{\lambda} \leq \lambda_{\kappa t} < 1 - \lambda_\delta \\ \left(-\nabla V_\pi(X_t) - \Gamma M^{-1} V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } 1 - \lambda_\delta \leq \lambda_{\kappa t} \leq 1 \end{cases} \\ dY_t &= \begin{cases} \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_\delta}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_\delta}} \right) + \frac{1}{\sqrt{1-\lambda_\delta}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1-\lambda_\delta}} \right) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \frac{1}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{Y_t}{\sqrt{\lambda_{\kappa t}}} \right) + \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{Y_t - X_t}{\sqrt{1-\lambda_{\kappa t}}} \right) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t & \text{if } \lambda_\delta \leq \lambda_{\kappa t} \leq 1. \end{cases} \end{cases} \quad (16)$$

Note that when $\lambda_{\kappa t} \geq 1 - \lambda_\delta$, the dynamics in (16) reduce to standard underdamped Langevin dynamics with a warm start, which can be discretised using any preferred numerical scheme. Therefore, we focus on the regime $0 \leq \lambda_{\kappa t} < 1 - \lambda_\delta$. We propose to use a hybrid discretisation method that combines the OBABO splitting scheme [Monmarché, 2020] with an SROCK update [Abdulle et al., 2018] for the fast-dynamics. The algorithm is explicitly defined in Algorithm 3, where $\sqrt{\Gamma}$ denotes the matrix square root of the friction coefficient, which is well defined since Γ is a symmetric positive definite matrix. That is, $\sqrt{\Gamma}$ is any matrix satisfying $\sqrt{\Gamma} \sqrt{\Gamma}^\top = \Gamma$.

It is important to mention that for implementation, we set the mass parameter to $M = I$ and consider a time-dependent friction coefficient Γ_t to control the degree of acceleration. Additionally, the same discretisation scheme applies to the heavy-tailed diffusion, with a modified fast process defined by Eq. (7).

B.2 MULTCDIFF: Multiscale Controlled Diffusions

B.2.1 Overdamped system

Diffusion models typically consider a stochastic process $(\vec{X}_t)_{t \in [0, T]}$ constructed by initialising \vec{X}_0 at the target distribution π and then evolving according to the OU SDE

$$d\vec{X}_t = -\vec{X}_t dt + \sqrt{2} dB'_t. \quad (17)$$

Under mild regularity conditions [Anderson, 1982], the dynamics of the reverse process $(\overleftarrow{X}_t)_{t \in [0, T]}$ are described by the following SDE

$$d\overleftarrow{X}_t = \left(\overleftarrow{X}_t + 2\nabla \log q_{T-t}(\overleftarrow{X}_t) \right) dt + \sqrt{2} dB_t, \quad (18)$$

where q_t denotes the marginal distribution of the solution of the forward process (17), given by

$$q_t(x) = \int \pi(y) \rho_t(x|y) dy,$$

with $\rho_t(x|y)$ being the Gaussian transition density $\mathcal{N}(e^{-t}y, \sigma_t^2 I)$, where $\sigma_t^2 = 1 - e^{-2t}$. Using this, the score function $\nabla \log q_{T-t}(\overleftarrow{X}_t)$ can be expressed as

$$\nabla \log q_{T-t}(\overleftarrow{X}_t) = -\frac{\overleftarrow{X}_t - e^{-(T-t)} \mathbb{E}_{Y|\overleftarrow{X}_t}[Y]}{\sigma_{T-t}^2}.$$

Since the score term involves an intractable expectation, we approximate the averaged dynamics in (18) using a multiscale SDE system. The fast process of the system is modelled by an overdamped

Algorithm 3 MULTALMC sampler: accelerated version

Require: Schedule function λ_t , value for λ_δ and $\tilde{\lambda}$, friction coefficient Γ (scalar), mass parameter M (scalar), number of sampling steps L , time discretisation $0 = T_0 < \dots < T_L = 1/\kappa$, step size $h_l = T_{l+1} - T_l$. Constants for SROCK step from (14), (15).

Initial samples $X_0 \sim \mathcal{N}(0, I)$, $V_0 \sim \mathcal{N}(0, MI)$, $Y_0 \sim \mathcal{N}(0, I)$. Define the schedule

$$\lambda'_{\kappa t} = \begin{cases} \lambda_\delta & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \lambda_{\kappa t} & \text{if } \lambda_\delta \leq \lambda_{\kappa t} \leq 1. \end{cases}$$

for $l = 0$ **to** L **do**

$$\xi_l^{(1)}, \xi_l^{(2)}, \xi_l^{(3)} \sim \mathcal{N}(0, I)$$

if $0 \leq \lambda_{\kappa T_l} < 1 - \lambda_\delta$ **then**

Half-step for velocity component

$$V'_l = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right) V_l + \sqrt{h_l \Gamma} \xi_l^{(1)}$$

$$V''_l = \begin{cases} V'_l + \frac{h_l}{2\sqrt{1-\lambda'_{\kappa T_l}}} \nabla \log \nu \left(\frac{X_l - Y_l}{\sqrt{1-\lambda'_{\kappa T_l}}} \right) & \text{if } \lambda'_{\kappa T_l} < \tilde{\lambda} \\ V'_l - \frac{h_l}{2\sqrt{\lambda'_{\kappa T_l}}} \nabla V_\pi \left(\frac{Y_l}{\sqrt{\lambda'_{\kappa T_l}}} \right) & \text{if } \lambda'_{\kappa T_l} \geq \tilde{\lambda} \end{cases}$$

Full EM-step for position component

$$X_{l+1} = X_l + h_l M^{-1} V''_l$$

Full SROCK-step for fast component

$$\begin{cases} K_{l,0} = Y_l \\ K_{l,1} = K_{l,0} + \mu_1 \frac{h_l}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda'_{\kappa T_l}}} \nabla V_\pi \left(\frac{K_{l,0}}{\sqrt{\lambda'_{\kappa T_l}}} \right) + \frac{1}{\sqrt{1-\lambda'_{\kappa T_l}}} \nabla \log \nu \left(\frac{K_{l,0} - X_{l+1}}{\sqrt{1-\lambda'_{\kappa T_l}}} \right) \right) \\ \vdots \\ K_{l,i} = \mu_i \frac{h_l}{\varepsilon} \left(-\frac{1}{\sqrt{\lambda'_{\kappa T_l}}} \nabla V_\pi \left(\frac{K_{l,i-1}}{\sqrt{\lambda'_{\kappa T_l}}} \right) + \frac{1}{\sqrt{1-\lambda'_{\kappa T_l}}} \nabla \log \nu \left(\frac{K_{l,i-1} - X_{l+1}}{\sqrt{1-\lambda'_{\kappa T_l}}} \right) \right) + \nu_i K_{l,i-1} + \kappa_i K_{l,i-2} \\ \vdots \\ Y_{l+1} = K_s + \sqrt{\frac{2h}{\varepsilon}} \xi_l^{(2)} \end{cases}$$

Half-step for velocity component

$$V'''_l = \begin{cases} V''_l + \frac{h_l}{2\sqrt{1-\lambda'_{\kappa T_l}}} \nabla \log \nu \left(\frac{X_{l+1} - Y_{l+1}}{\sqrt{1-\lambda'_{\kappa T_l}}} \right) & \text{if } \lambda'_{\kappa T_l} < \tilde{\lambda} \\ V''_l - \frac{h_l}{2\sqrt{\lambda'_{\kappa T_l}}} \nabla V_\pi \left(\frac{Y_{l+1}}{\sqrt{\lambda'_{\kappa T_l}}} \right) & \text{if } \lambda'_{\kappa T_l} \geq \tilde{\lambda} \end{cases}$$

$$V_{l+1} = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right) V'''_l + \sqrt{h_l \Gamma} \xi_l^{(3)}$$

if $1 - \lambda_\delta \leq \lambda_{\kappa T_l} \leq 1$ **then**

Half-step for velocity component

$$V'_l = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right) V_l + \sqrt{h_l \Gamma} \xi_l^{(1)}$$

$$V''_l = V'_l - \frac{h_l}{2} \nabla V_\pi (X_l)$$

Full EM-step for position component

$$X_{l+1} = X_l + h_l M^{-1} V''_l$$

Half-step for velocity component

$$V'''_l = V''_l - \frac{h_l}{2} \nabla V_\pi (X_{l+1})$$

$$V_{l+1} = \left(1 - \frac{h_l}{2}\Gamma M^{-1}\right) V'''_l + \sqrt{h_l \Gamma} \xi_l^{(3)}$$

end for

return (X_L, V_L)

Langevin SDE targetting the conditional distribution $Y|\overleftarrow{X}_t$, which takes the form

$$dY_t = \frac{1}{\varepsilon} \left(\nabla \log \pi(Y_t) + \nabla_{Y_t} \log \rho_{T-t}(\overleftarrow{X}_t|Y_t) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t.$$

where

$$\nabla_{Y_t} \log \rho_{T-t}(\overleftarrow{X}_t, \overleftarrow{V}_t|Y_t) = -\frac{Y_t - e^{T-t}\overleftarrow{X}_t}{\sigma_{T-t}^2} e^{-2(T-t)}.$$

Putting everything together, we obtain the following multiscale system of SDEs

$$\begin{cases} d\overleftarrow{X}_t = \left(\overleftarrow{X}_t - \frac{2}{\sigma_{T-t}^2} (\overleftarrow{X}_t - e^{-(T-t)}Y_t) \right) dt + \sqrt{2} dB_t \\ dY_t = \frac{1}{\varepsilon} \left(\nabla \log \pi(Y_t) - \frac{e^{-2(T-t)}}{\sigma_{T-t}^2} (Y_t - e^{T-t}\overleftarrow{X}_t) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t. \end{cases} \quad (19)$$

This system of SDEs can be discretised using either an exponential integrator scheme or a numerical scheme that combines Euler-Maruyama for the slow dynamics with the SROCK method [Abdulle et al., 2018] for the fast component, similarly to the discretisation of MULTALMC in the overdamped case (see Appendix B.1.1).

B.2.2 Underdamped system

When using an OU noising process, the forward underdamped diffusion has the following form

$$\begin{cases} d\overrightarrow{X}_t = M^{-1}\overrightarrow{V}_t dt \\ d\overrightarrow{V}_t = (-\overrightarrow{X}_t - \Gamma M^{-1}\overrightarrow{V}_t) dt + \sqrt{2\Gamma} dB_t', \end{cases}$$

which can be written compactly as

$$d\overrightarrow{Z}_t = A \otimes I_d \overrightarrow{Z}_t dt + B \otimes I_d dB_t'',$$

where $\overrightarrow{Z}_t = (\overrightarrow{X}_t, \overrightarrow{V}_t)$ and

$$A = \begin{pmatrix} 0 & M^{-1} \\ -1 & -\Gamma M^{-1} \end{pmatrix}, \quad B = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{2\Gamma} \end{pmatrix},$$

where, as in Appendix B.1.2, $\sqrt{\Gamma}$ denotes the matrix square root of the friction coefficient Γ , which is a symmetric and positive definite. The solution of this SDE is given by [Karatzas and Shreve, 1991]

$$\overrightarrow{Z}_t = e^{At} \otimes I_d \overrightarrow{Z}_0 + \int_0^t (e^{A(t-s)} B) \otimes I_d dB_s''.$$

As discussed in the main text, there is a crucial balance between the mass M and friction coefficient Γ . We focus on the critically damping regime by setting $\Gamma^2 = 4M$. This provides an ideal balance between the Hamiltonian and the Ohnstein-Uhlenbeck components of the dynamics, which leads to faster convergence without oscillations [McCall, 2010]. Under this setting, the matrix A can be written as

$$A = \begin{pmatrix} 0 & 4\Gamma^{-2} \\ -1 & -4\Gamma^{-1} \end{pmatrix}.$$

Then, the matrix exponential e^{At} has the following form

$$e^{At} = \begin{pmatrix} 1 + 2\Gamma^{-1}t & 4\Gamma^{-2}t \\ -t & 1 - 2\Gamma^{-1}t \end{pmatrix} e^{-2\Gamma^{-1}t} = \begin{pmatrix} m_{1,t} & m_{2,t} \\ m_{3,t} & m_{4,t} \end{pmatrix}.$$

On the other hand, the reverse SDE is given by

$$\begin{cases} d\overleftarrow{X}_t = -4\Gamma^{-2}\overleftarrow{V}_t dt \\ d\overleftarrow{V}_t = \left(\overleftarrow{X}_t + 4\Gamma^{-1}\overleftarrow{V}_t + 2\Gamma \nabla_{V_t} \log q_{T-t}(\overleftarrow{X}_t, \overleftarrow{V}_t) \right) dt + \sqrt{2\Gamma} dB_t, \end{cases}$$

where q_t is the marginal distribution of the solution of the forward noising process which has the following expression

$$q_t(x, v) = \int \pi(y) \varphi(v_0) \rho_t(x, v|y, v_0) dy dv_0, \quad (20)$$

with $\varphi = \mathcal{N}(0, I)$ and conditional distribution $\rho_t(x, v|y, v_0)$ given by a Normal distribution with mean

$$m_t(y, v_0) = e^{At} \otimes I_d \begin{pmatrix} y \\ v_0 \end{pmatrix} = \begin{pmatrix} m_{1,t}y + m_{2,t}v_0 \\ m_{3,t}y + m_{4,t}v_0 \end{pmatrix}$$

and covariance matrix

$$\Sigma_t \otimes I_d = \begin{pmatrix} \Sigma_{11,t} & \Sigma_{12,t} \\ \Sigma_{21,t} & \Sigma_{22,t} \end{pmatrix} \otimes I_d = \left(\int_0^t e^{A(t-s)} B B^\top [e^{A(t-s)}]^\top ds \right) \otimes I_d,$$

which results into

$$\begin{aligned} \Sigma_{11,t} &= \left(e^{4t\Gamma^{-1}} - 1 - 4t\Gamma^{-1} - 8t^2\Gamma^{-2} \right) e^{-4t\Gamma^{-1}} \\ \Sigma_{12,t} &= \Sigma_{21,t} = 4t^2\Gamma^{-1} e^{-4t\Gamma^{-1}} \\ \Sigma_{22,t} &= \left(\frac{\Gamma^2}{4} \left(e^{4t\Gamma^{-1}} - 1 \right) + t\Gamma - 2t^2 \right) e^{-4t\Gamma^{-1}}. \end{aligned}$$

We note that v_0 in Eq. (20) can be integrated out analytically. Since both $\rho_t(x, v|y, v_0)$ and $\varphi(v_0)$ are Gaussian, it follows that $\rho_t(x, v|y)$ is also Gaussian with mean

$$\hat{m}(y) = \begin{pmatrix} m_{1,t}y \\ m_{3,t}y \end{pmatrix}$$

and covariance matrix

$$\tilde{\Sigma}_t \otimes I_d = \begin{pmatrix} \tilde{\Sigma}_{11,t} & \tilde{\Sigma}_{12,t} \\ \tilde{\Sigma}_{21,t} & \tilde{\Sigma}_{22,t} \end{pmatrix} \otimes I_d = \left[\begin{pmatrix} \Sigma_{11,t} & \Sigma_{12,t} \\ \Sigma_{21,t} & \Sigma_{22,t} \end{pmatrix} + \begin{pmatrix} m_{2,t}^2 & m_{2,t}m_{4,t} \\ m_{2,t}m_{4,t} & m_{4,t}^2 \end{pmatrix} \right] \otimes I_d.$$

We can rewrite Eq. (20) in terms of $\rho_t(x, v|y)$ as

$$q_t(x, v) = \int \pi(y) \rho_t(x, v|y) dy,$$

We note that the conditional distribution $\rho_t(v|x, y)$ remains Gaussian, with mean

$$\tilde{m}_t(x, y) = m_{3,t}y + \tilde{\Sigma}_{12,t}\tilde{\Sigma}_{11,t}^{-1}(x - m_{1,t}y),$$

which is linear in y and covariance $\sigma_t^2 I$, where

$$\sigma_t^2 = \tilde{\Sigma}_{22,t} - \tilde{\Sigma}_{12,t}\tilde{\Sigma}_{11,t}^{-1}\tilde{\Sigma}_{12,t}.$$

This provides the following expression for the score

$$\nabla_v \log q_t(x, v) = \nabla_v \log q_t(v|x) = -\frac{v - \mathbb{E}_{Y|x,v}[\tilde{m}_t(x, Y)]}{\sigma_t^2} = -\frac{v - f_t(x, \mathbb{E}_{Y|x,v}[Y])}{\sigma_t^2},$$

where

$$f_t(x, \mathbb{E}_{Y|x,v}[Y]) = m_{3,t}\mathbb{E}_{Y|x,v}[Y] + \tilde{\Sigma}_{12,t}\tilde{\Sigma}_{11,t}^{-1}(x - m_{1,t}\mathbb{E}_{Y|x,v}[Y]).$$

Using the derivations above, we define a multiscale SDE system that enables sampling from the target distribution via the reverse SDE, without requiring prior estimation of the denoiser $\mathbb{E}_{Y|x,v}[Y]$. We consider a fast process governed by an overdamped Langevin SDE targeting the conditional distribution of Y given the current state $(\tilde{X}_t, \tilde{V}_t)$, which takes the form

$$dY_t = \frac{1}{\varepsilon} \left(\nabla \log \pi(Y_t) + \nabla_{Y_t} \log \rho_{T-t}(\tilde{X}_t, \tilde{V}_t|Y_t) \right) dt + \sqrt{\frac{2}{\varepsilon}} dB_t.$$

where

$$\nabla_{Y_t} \log \rho_{T-t}(\tilde{X}_t, \tilde{V}_t|Y_t) = - \left[(m_{1,T-t}Y_t, m_{3,T-t}Y_t) - (\tilde{X}_t, \tilde{V}_t) \right] \tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix}.$$

Combining altogether, we arrive at the following multiscale system of SDEs

$$\begin{cases} d\tilde{X}_t = -4\Gamma^{-2}\tilde{V}_t dt \\ d\tilde{V}_t = \left(\tilde{X}_t + 4\Gamma^{-1}\tilde{V}_t - \frac{2\Gamma}{\sigma_{T-t}^2} \left(\tilde{V}_t - (m_{3,T-t}Y_t + \tilde{\Sigma}_{12,T-t}\tilde{\Sigma}_{11,T-t}^{-1}(\tilde{X}_t - m_{1,T-t}Y_t)) \right) \right) dt + \sqrt{2\Gamma} dB_t \\ dY_t = \frac{1}{\varepsilon} \left(\nabla \log \pi(Y_t) + \nabla_{Y_t} \log \rho_{T-t}(\tilde{X}_t, \tilde{V}_t|Y_t) \right) dt + \sqrt{\frac{2}{\varepsilon}} dB_t. \end{cases} \quad (21)$$

Inspired by Dockhorn et al. [2022b], we propose a novel discretisation scheme which leverages symmetric splitting techniques [Leimkuhler and Matthews, 2013] justified by the symmetric Trotter splitting and the Baker–Campbell–Hausdorff formula [Strang, 1968, Trotter, 1959, Tuckerman, 2010]. The l -th iteration of the proposed numerical scheme consists of the following steps, with h_l denoting the step size.

1. Evolve the following system of SDEs exactly for half-step $h_l/2$

$$\begin{cases} d\overleftarrow{X}_t = -4\Gamma^{-2}\overleftarrow{V}_t dt, \\ d\overleftarrow{V}_t = \left(\overleftarrow{X}_t + 4\Gamma^{-1}\overleftarrow{V}_t\right) dt + \sqrt{2\Gamma} dB_t. \end{cases} \quad (22)$$

2. Evolve the velocity component (full step) under the following ODE using Euler's method

$$d\overleftarrow{V}_t = -\frac{2\Gamma}{\sigma_{T-t}^2} \left(\overleftarrow{V}_t - \left(m_{3,T-t} Y_t + \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} \left(\overleftarrow{X}_t - m_{1,T-t} Y_t \right) \right) \right) dt.$$

3. Evolve the fast dynamics (full step) using SROCK method

$$dY_t = \frac{1}{\varepsilon} \left(\nabla \log \pi(Y_t) + \nabla_{Y_t} \log \rho_{T-t}(\overleftarrow{X}_t, \overleftarrow{V}_t | Y_t) \right) dt + \sqrt{\frac{2}{\varepsilon}} d\tilde{B}_t.$$

4. Evolve the system in Eq. (22) exactly for half-step $h_l/2$.

We summarise the steps above in the form of a concise algorithm. To do so, we first derive the exact solution of the SDE system (22) used in Step 1. Define the matrices

$$\hat{A} = \begin{pmatrix} 0 & -4\Gamma^{-2} \\ 1 & 4\Gamma^{-1} \end{pmatrix}, \quad \hat{B} = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{2\Gamma} \end{pmatrix}.$$

The matrix exponential $e^{\hat{A}t}$ takes the form

$$e^{\hat{A}t} = \begin{pmatrix} 1 - 2\Gamma^{-1}t & -4\Gamma^{-2}t \\ t & 1 + 2\Gamma^{-1}t \end{pmatrix} e^{2\Gamma^{-1}t}.$$

The solution to the system (22) at time t , given an initial condition (X_0, V_0) , can be expressed as

$$(X_t, V_t) \sim \mathcal{N}(\hat{m}_t(X_0, V_0), \hat{\Sigma}_t \otimes I_d)$$

where the mean and covariance matrix are defined by

$$\hat{m}_t(X_0, V_0) = e^{\hat{A}t} \otimes I_d \begin{pmatrix} X_0 \\ V_0 \end{pmatrix}, \quad \hat{\Sigma}_t = \begin{pmatrix} \hat{\Sigma}_{11,t} & \hat{\Sigma}_{12,t} \\ \hat{\Sigma}_{21,t} & \hat{\Sigma}_{22,t} \end{pmatrix} = \int_0^t e^{\hat{A}(t-s)} \hat{B} \hat{B}^\top \left[e^{\hat{A}(t-s)} \right]^\top ds. \quad (23)$$

The explicit expressions for the entries of the covariance matrix $\hat{\Sigma}_t$ are

$$\begin{aligned} \hat{\Sigma}_{11,t} &= \left(-e^{4t\Gamma^{-1}} + 1 - 4t\Gamma^{-1} + 8t^2\Gamma^{-2} \right) e^{4t\Gamma^{-1}} \\ \hat{\Sigma}_{12,t} &= \hat{\Sigma}_{21,t} = 4t^2\Gamma^{-1} e^{4t\Gamma^{-1}} \\ \hat{\Sigma}_{22,t} &= \left(\frac{\Gamma^2}{4} \left(1 - e^{4t\Gamma^{-1}} \right) + t\Gamma + 2t^2 \right) e^{4t\Gamma^{-1}}. \end{aligned}$$

The full algorithm is presented in Algorithm 4.

Challenges for heavy-tailed controlled diffusions. When using a heavy-tailed noising process given by the convolution with a Student's t distribution with tail index α , the forward underdamped diffusion takes the form

$$\begin{cases} d\overrightarrow{X}_t = M^{-1}\overrightarrow{V}_t dt \\ d\overrightarrow{V}_t = \left(-\frac{\alpha+d}{\alpha} \frac{\overrightarrow{X}_t}{1 + \frac{\|\overrightarrow{X}_t\|^2}{\alpha}} - \Gamma M^{-1}\overrightarrow{V}_t \right) dt + \sqrt{2\Gamma} dB_t. \end{cases}$$

The non-linear drift term in the velocity dynamics prevents us from obtaining analytic solutions via standard techniques for linear SDEs. As a result, we cannot directly reproduce the computations used in the case of an OU noising process. A detailed investigation of this regime is left for future work.

C Theoretical guarantees

In this section, we provide detailed proofs of the theoretical results presented in the paper and discuss the additional challenges posed by heavy-tailed diffusions.

Algorithm 4 MULTCDIFF sampler

Require: Friction coefficient Γ (scalar), $M = \Gamma^2/4$, number of sampling steps L , time discretisation $0 = T_0 < \dots < T_L = T$, step size $h_l = T_{l+1} - T_l$. Constants for SROCK step from (14), (15). Functions $\hat{m}_t(\cdot, \cdot)$ and $\hat{\Sigma}_t$ from (23).

Initial samples $X_0 \sim \mathcal{N}(0, I)$, $V_0 \sim \mathcal{N}(0, MI)$, $Y_0 \sim \mathcal{N}(0, I)$.

for $l = 0$ **to** L **do**

$$(X_{l+1/2}, V_{l+1/2}) \sim \mathcal{N}(\hat{m}_{h_l/2}(X_l, V_l), \hat{\Sigma}_{h_l/2})$$

$$V'_{l+1/2} = V_{l+1/2} - \frac{2\Gamma}{\sigma_{T-T_l}^2} \left(V_{l+1/2} - \left(m_{3,T-T_l} Y_l + \tilde{\Sigma}_{12,T-T_l} \tilde{\Sigma}_{11,T-T_l}^{-1} (X_{l+1/2} - m_{1,T-T_l} Y_l) \right) \right) \frac{h_l}{2}$$

$$\begin{cases} K_{l,0} &= Y_l \\ K_{l,1} &= K_{l,0} + \mu_1 \frac{h_l}{\varepsilon} \left(-\nabla V_\pi(K_{l,0}) + \nabla_Y \log \rho_{T-T_l}(X_{l+1/2}, V'_{l+1/2} | K_{l,0}) \right) \\ &\vdots \\ K_{l,i} &= \mu_i \frac{h_l}{\varepsilon} \left(-\nabla V_\pi(K_{l,i-1}) + \nabla_Y \log \rho_{T-T_l}(X_{l+1/2}, V'_{l+1/2} | K_{l,i-1}) \right) + \nu_i K_{l,i-1} + \kappa_i K_{l,i-2} \\ &\vdots \\ Y_{l+1} &= K_s + \sqrt{\frac{2h}{\varepsilon}} \xi_l, \quad \xi_l \sim \mathcal{N}(0, I) \end{cases}$$

$$(X_{l+1/2}, V_{l+1/2}) \sim \mathcal{N}(\hat{m}_{h_l/2}(X_{l+1}, V'_{l+1/2}), \hat{\Sigma}_{h_l/2})$$

end for

return (X_L, V_L)

C.1 Convergence of the slow-fast system to the averaged dynamics

Based on the stochastic averaging results from [Liu et al., 2020] presented in Appendix A, we restate our convergence result, Theorem 4.1, and provide the proof.

Theorem C.1. [Theorem 4.1 restated] *Let the base distribution $\nu \sim \mathcal{N}(0, I)$. Suppose the target distribution π and the schedule function λ_t satisfy assumptions A1 and A2, respectively. Then, for any $\varepsilon \in (0, \varepsilon_0)$, where ε_0 is specified in the proof, and any given initial conditions, there exists unique solutions $\{(X_t^\varepsilon, Y_t^\varepsilon), t \geq 0\}$, $\{(X_t^\varepsilon, V_t^\varepsilon, Y_t^\varepsilon), t \geq 0\}$ and $\{(\bar{X}_t^\varepsilon, \bar{V}_t^\varepsilon, Y_t^\varepsilon), t \geq 0\}$ to the slow-fast stochastic systems (4), (6) and (10), respectively. Furthermore, for any $p > 0$, it holds that*

$$\lim_{\varepsilon \rightarrow 0} \mathbb{E} \left(\sup_{t \in [0, T]} \|X_t^\varepsilon - \bar{X}_t\|^p \right) = 0,$$

where \bar{X}_t denotes the solution of the averaged system and X_t^ε is the corresponding slow component of the multiscale systems (4), (6) and (10).

Proof. If the coefficients of the proposed slow-fast systems (4), (6) and (10) satisfy assumptions A3 and A4, then the result follows directly from Theorems A.1 and A.2. Thus, we verify below that these assumptions hold in all cases.

It is important to mention that although the theorem specifically addresses the case where $\nu \sim \mathcal{N}(0, I)$, assumption A3 only requires that $\nabla \log \nu$ is Lipschitz continuous with constant L_ν and has a maximiser x^* . Without loss of generality, we assume $x^* = 0$, i.e., $\nabla \log \nu(0) = 0$. Therefore, we show that assumption A3 holds for general base distributions ν satisfying these properties.

Besides, for the systems based on annealed Langevin diffusion (4) and (6), we instead consider their modified versions (13) and (16), respectively, as these reflect the dynamics used in practice during implementation. We observed in Appendix B.1 that these systems reduced to

- A multiscale system with time-independent coefficients for $0 \leq \lambda_{\kappa t} < \lambda_\delta$.
- A multiscale system with time-dependent coefficients for $\lambda_\delta \leq \lambda_{\kappa t} < 1 - \lambda_\delta$. Here, we distinguish between two different dynamics depending on whether $\lambda_{\kappa t} < \tilde{\lambda}$ or $\lambda_{\kappa t} \geq \tilde{\lambda}$.
- Standard Langevin dynamics for $1 - \lambda_\delta \leq \lambda_{\kappa t} \leq 1$.

We analyse the first two cases, as standard Langevin dynamics is known to converge to its invariant distribution when initialised with a warm start [Chewi et al., 2021, Lee et al., 2021a, Wu et al., 2022]. See Appendix B.1 for further discussion.

We denote by m^* a minimiser of the target potential V_π .

Overdamped MULTALMC (4) or (13) The coefficients of the slow-fast system for $\lambda_{\kappa t} \in [0, 1 - \lambda_\delta]$ have the following expressions

$$b(t, x, y) = \begin{cases} \frac{1}{\sqrt{1-\lambda_\delta}} \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_\delta}} \right) & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) & \text{if } \lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda} \\ -\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) & \text{if } \tilde{\lambda} \leq \lambda_{\kappa t} < 1 - \lambda_\delta, \end{cases}$$

$$\sigma(t, x) = \sqrt{2},$$

$$f(t, x, y) = \begin{cases} \frac{1}{\sqrt{1-\lambda_\delta}} \nabla \log \nu \left(\frac{y-x}{\sqrt{1-\lambda_\delta}} \right) - \frac{1}{\sqrt{\lambda_\delta}} \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_\delta}} \right) & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{y-x}{\sqrt{1-\lambda_{\kappa t}}} \right) - \frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) & \text{if } \lambda_\delta \leq \lambda_{\kappa t} < 1 - \lambda_\delta, \end{cases} \quad (24)$$

$$g(t, x, y) = \sqrt{2}.$$

We consider two regimes based on the value of $\lambda_{\kappa t}$: Regime 1 $\lambda_{\kappa t} \in [0, \tilde{\lambda}]$ and Regime 2 $\lambda_{\kappa t} \in [\tilde{\lambda}, 1 - \lambda_\delta]$. We analyse the convergence within each regime separately.

Substituting the expressions of the coefficients, it follows that assumption **A3** holds with $\theta_1 = 0$, $\theta_2 = \theta_3 = \theta_5 = \theta_6 = 1$, $\theta_4 = 2$ and $Z_T = \|m^*\| \sqrt{\tilde{\lambda}}/2$ is a constant random variable in both regimes, as shown below.

$$\begin{aligned} & 2\|b(t, x_1, y) - b(t, x_2, y)\| \|x_1 - x_2\| + \|\sigma(t, x_1) - \sigma(t, x_2)\|^2 \\ & \leq \begin{cases} \mathbb{1}_{\{0 \leq \lambda_{\kappa t} < \lambda_\delta\}} \frac{2L_\nu}{1-\lambda_\delta} \|x_1 - x_2\|^2 + \mathbb{1}_{\{\lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \frac{2L_\nu}{1-\lambda_{\kappa t}} \|x_1 - x_2\|^2 & \text{Regime 1} \\ 0 & \text{Regime 2} \end{cases} \\ & \leq \frac{2L_\nu}{1-\lambda} \|x_1 - x_2\|^2, \end{aligned}$$

$$\begin{aligned} \|b(t, x, y_1) - b(t, x, y_2)\| & \leq \begin{cases} \mathbb{1}_{\{0 \leq \lambda_{\kappa t} < \lambda_\delta\}} \frac{L_\nu}{1-\lambda_\delta} \|y_1 - y_2\| + \mathbb{1}_{\{\lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \frac{L_\nu}{1-\lambda_{\kappa t}} \|y_1 - y_2\| & \text{Regime 1} \\ \mathbb{1}_{\{\tilde{\lambda} \leq \lambda_{\kappa t} < 1 - \lambda_\delta\}} \frac{L_\pi}{\lambda_{\kappa t}} \|y_1 - y_2\| & \text{Regime 2} \end{cases} \\ & \leq \max \left\{ \frac{L_\nu}{1-\lambda}, \frac{L_\pi}{\tilde{\lambda}} \right\} \|y_1 - y_2\|. \end{aligned}$$

For the following, we can assume without loss of generality that $s \leq t$

$$\begin{aligned} & \|b(t, x, y) - b(s, x, y)\| \\ & \leq \begin{cases} \mathbb{1}_{\{0 \leq \lambda_{\kappa s} < \lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \left| \frac{1}{\sqrt{1-\lambda_{\kappa t}}} - \frac{1}{\sqrt{1-\lambda_\delta}} \right| \left(L_\nu \frac{\|x\| + \|y\|}{\sqrt{1-\tilde{\lambda}}} + \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) \right\| \right) \\ + \mathbb{1}_{\{\lambda_\delta \leq \lambda_{\kappa t}, \lambda_{\kappa s} < \tilde{\lambda}\}} \left| \frac{1}{\sqrt{1-\lambda_{\kappa t}}} - \frac{1}{\sqrt{1-\lambda_{\kappa s}}} \right| \left(L_\nu \frac{\|x\| + \|y\|}{\sqrt{1-\tilde{\lambda}}} + \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) \right\| \right) & \text{Regime 1} \\ \mathbb{1}_{\{\tilde{\lambda} \leq \lambda_{\kappa t}, \lambda_{\kappa s} < 1 - \lambda_\delta\}} \left| \frac{1}{\sqrt{\lambda_{\kappa t}}} - \frac{1}{\sqrt{\lambda_{\kappa s}}} \right| \left(\frac{L_\pi \|y\|}{\sqrt{\tilde{\lambda}}} + \left\| \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) \right\| \right) & \text{Regime 2} \end{cases} \\ & \leq \begin{cases} \frac{C|t-s|^{\gamma_1}}{(1-\tilde{\lambda})^2} L_\nu (\|x\| + \|y\|) & \text{Regime 1} \\ \frac{C|t-s|^{\gamma_1}}{\tilde{\lambda}^2} L_\pi \left(\|y\| + \frac{\|m^*\| \sqrt{\tilde{\lambda}}}{2} \right) & \text{Regime 2} \end{cases} \\ & \leq C \max \left\{ \frac{L_\nu}{(1-\tilde{\lambda})^2}, \frac{L_\pi}{\tilde{\lambda}^2} \right\} |t-s|^{\gamma_1} \left(\|y\| + \|x\| + \frac{\|m^*\| \sqrt{\tilde{\lambda}}}{2} \right), \end{aligned}$$

$$2\langle x, b(t, x, y) \rangle$$

$$\begin{aligned}
&\leq \begin{cases} \mathbb{1}_{\{0 \leq \lambda_t < \lambda_\delta\}} \frac{2\|x\|}{\sqrt{1-\lambda_\delta}} \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_\delta}} \right) \right\| + \mathbb{1}_{\{\lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \frac{2\|x\|}{\sqrt{1-\lambda_{\kappa t}}} \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) \right\| & \text{Regime 1} \\ \mathbb{1}_{\{\tilde{\lambda} \leq \lambda_{\kappa t} < 1-\lambda_\delta\}} \frac{2\|x\|}{\sqrt{\lambda_{\kappa t}}} \left\| \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) \right\| & \text{Regime 2} \end{cases} \\
&\leq \begin{cases} \frac{L_\nu}{1-\tilde{\lambda}} (2\|x\|^2 + \|y\|^2) & \text{Regime 1} \\ \frac{L_\pi}{\tilde{\lambda}} (2\|x\|^2 + \|y\|^2 + \tilde{\lambda}\|m^*\|^2) & \text{Regime 2} \end{cases} \\
&\leq 2 \max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\} \max\{1, \tilde{\lambda}\|m^*\|^2\} (1 + \|x\|^2) + \max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\} \|y\|^2,
\end{aligned}$$

which leads to the constant λ_1 in assumption **A3** being given by $\lambda_1 = \max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\}$.

$$\begin{aligned}
\|b(t, x, y)\| &\leq \begin{cases} \frac{L_\nu(\|x\| + \|y\|)}{1-\tilde{\lambda}} & \text{Regime 1} \\ \frac{L_\pi}{\tilde{\lambda}} (\|y\| + \sqrt{\tilde{\lambda}}\|m^*\|) & \text{Regime 2} \end{cases} \\
&\leq \frac{L_\pi}{\sqrt{\tilde{\lambda}}} \|m^*\| + \max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\} (\|x\| + \|y\|),
\end{aligned}$$

$$\|\sigma(t, x)\|^2 = 2.$$

On the other hand, to verify that assumption **A4** holds, we recall that for a standard Normal distribution $\nabla \log \nu(x) = -x$. Besides, note that we can express the drift term of the fast process, $f(t, x, y)$, in a compact form by rewriting (24) as

$$f(t, x, y) = \frac{1}{\sqrt{1-\lambda'_{\kappa t}}} \nabla \log \nu \left(\frac{y-x}{\sqrt{1-\lambda'_{\kappa t}}} \right) - \frac{1}{\sqrt{\lambda'_{\kappa t}}} \nabla V_\pi \left(\frac{y}{\sqrt{\lambda'_{\kappa t}}} \right),$$

where the schedule $\lambda'_{\kappa t}$ is defined by

$$\lambda'_{\kappa t} = \begin{cases} \lambda_\delta & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_\delta \\ \lambda_{\kappa t} & \text{if } \lambda_\delta \leq \lambda_{\kappa t} < 1 - \lambda_\delta. \end{cases}$$

Therefore, we adopt this compact form in the following analysis and omit explicit distinctions between regimes to simplify notation. Using this, we have

$$\begin{aligned}
&2\langle f(t, x, y_1) - f(t, x, y_2), y_1 - y_2 \rangle + \|g(t, x, y_1) - g(t, x, y_2)\|^2 \\
&\leq -\frac{2}{1-\lambda'_{\kappa t}} \|y_1 - y_2\|^2 - \frac{2}{\sqrt{\lambda'_{\kappa t}}} \left\langle \nabla V_\pi \left(\frac{y_1}{\sqrt{\lambda'_{\kappa t}}} \right) - \nabla V_\pi \left(\frac{y_2}{\sqrt{\lambda'_{\kappa t}}} \right), y_1 - y_2 \right\rangle \\
&\leq -\left(\frac{2}{1-\lambda'_{\kappa t}} + \frac{2M_\pi}{\lambda'_{\kappa t}} \right) \|y_1 - y_2\|^2 \leq -\beta_\delta \|y_1 - y_2\|^2,
\end{aligned}$$

$$\begin{aligned}
\|f(t, x_1, y) - f(s, x_2, y)\| &\leq \left| \frac{1}{1-\lambda'_{\kappa t}} - \frac{1}{1-\lambda'_{\kappa s}} \right| (\|y\| + \|x_2\|) + \frac{\|x_1 - x_2\|}{1-\lambda'_{\kappa t}} \\
&\quad + \left| \frac{1}{\sqrt{\lambda'_{\kappa t}}} - \frac{1}{\sqrt{\lambda'_{\kappa s}}} \right| \left(\frac{L_\pi\|y\|}{\sqrt{\lambda'_{\kappa s}}} + \left\| \nabla V_\pi \left(\frac{y}{\sqrt{\lambda'_{\kappa t}}} \right) \right\| \right) \\
&\leq C_{1,\delta} |t-s|^{\gamma_1} (\|y\| + \|x_2\|) + C_{2,\delta} \|x_1 - x_2\| + C_{3,\delta} |t-s|^{\gamma_1} \|y\| \\
&\leq C_\delta (|t-s|^{\gamma_1} + \|x_1 - x_2\|) (1 + \|x_2\| + \|y\|),
\end{aligned}$$

$$\|f(t, x, y)\| \leq \max \left\{ \frac{L_\pi}{\lambda'_{\kappa t}(1-\lambda'_{\kappa t})}, \frac{1}{1-\lambda'_{\kappa t}}, \frac{L_\pi\|m^*\|}{\lambda'_{\kappa t}} \right\} (1 + \|x\| + \|y\|) \leq C_\delta (1 + \|x\| + \|y\|),$$

$$\|g(t, x_1, y_1) - g(s, x_2, y_2)\| = 0, \quad \|g(t, x, y)\| = \sqrt{2}.$$

Since the potential V_π is strongly convex with constant M_π , then V_π satisfies the dissipativity inequality $\langle \nabla V_\pi(x), x \rangle \geq a_\pi \|x\|^2 - b_\pi$. Therefore, we have that for any $k \geq 2$

$$\begin{aligned}
2\langle y, f(t, x, y) \rangle + (k-1)\|g(t, x, y)\|^2 &\leq -2 \left\langle y, \frac{y-x}{1-\lambda'_{\kappa t}} + \frac{1}{\sqrt{\lambda'_{\kappa t}}} \nabla V_\pi \left(\frac{y}{\sqrt{\lambda'_{\kappa t}}} \right) \right\rangle + 2(k-1) \\
&\leq -\frac{2}{1-\lambda'_{\kappa t}} (\|y\|^2 - \langle y, x \rangle) - \frac{2a_\pi}{\lambda'_{\kappa t}} \|y\|^2 + b_\pi + 2(k-1)
\end{aligned}$$

$$\begin{aligned} &\leq - \left(\frac{1}{1 - \lambda'_{\kappa t}} + \frac{2a_{\pi}}{\lambda'_{\kappa t}} \right) \|y\|^2 + \frac{1}{1 - \lambda'_{\kappa t}} \|x\|^2 + b_{\pi} + 2(k-1) \\ &\leq -\lambda_2 \|y\|^2 + C_{k,\delta} (\|x\|^2 + 1), \end{aligned}$$

where λ_2 is given by

$$\lambda_2 = \frac{1}{1 - \lambda_{\delta}} + \frac{2a_{\pi}}{\lambda_{\delta}}.$$

This concludes the proof for the overdamped MULTALMC, where the constant ε_0 in the theorem statement is explicitly given by

$$\varepsilon_0 = \frac{\lambda_2}{\lambda_1} = \frac{\frac{1}{1 - \lambda_{\delta}} + \frac{2a_{\pi}}{\lambda_{\delta}}}{\max \left\{ \frac{L_{\nu}}{1 - \lambda}, \frac{L_{\pi}}{\lambda} \right\}}.$$

Underdamped MULTALMC (6) or (16) In this case, since we introduce an auxiliary velocity variable, the slow component consists of both position and velocity, (x, v) . Accordingly, the coefficients of the slow-fast system are given by the following expressions

$$b(t, (x, v), y) = \begin{cases} \begin{pmatrix} M^{-1}v \\ \frac{1}{\sqrt{1 - \lambda_{\delta}}} \nabla \log \nu \left(\frac{x-y}{\sqrt{1 - \lambda_{\delta}}} \right) - \Gamma M^{-1}v \end{pmatrix} & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_{\delta} \\ \begin{pmatrix} M^{-1}v \\ \frac{1}{\sqrt{1 - \lambda_{\kappa t}}} \nabla \log \nu \left(\frac{x-y}{\sqrt{1 - \lambda_{\kappa t}}} \right) - \Gamma M^{-1}v \end{pmatrix} & \text{if } \lambda_{\delta} \leq \lambda_{\kappa t} < \tilde{\lambda} \\ \begin{pmatrix} M^{-1}v \\ -\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_{\pi} \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) - \Gamma M^{-1}v \end{pmatrix} & \text{if } \tilde{\lambda} \leq \lambda_{\kappa t} < 1 - \lambda_{\delta}, \end{cases}$$

$$\sigma(t, (x, v)) = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{2\Gamma} \end{pmatrix} \otimes I_d,$$

$$f(t, (x, v), y) = \begin{cases} \begin{pmatrix} \frac{1}{\sqrt{1 - \lambda_{\delta}}} \nabla \log \nu \left(\frac{y-x}{\sqrt{1 - \lambda_{\delta}}} \right) - \frac{1}{\sqrt{\lambda_{\delta}}} \nabla V_{\pi} \left(\frac{y}{\sqrt{\lambda_{\delta}}} \right) \\ \frac{1}{\sqrt{1 - \lambda_{\kappa t}}} \nabla \log \nu \left(\frac{y-x}{\sqrt{1 - \lambda_{\kappa t}}} \right) - \frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_{\pi} \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) \end{pmatrix} & \text{if } 0 \leq \lambda_{\kappa t} < \lambda_{\delta} \\ & \text{if } \lambda_{\delta} \leq \lambda_{\kappa t} < 1 - \lambda_{\delta}, \end{cases},$$

$$g(t, (x, v), y) = \sqrt{2}.$$

Since the coefficients of the fast process are the same as in the overdamped setting, then assumption **A4** also holds in the underdamped case. To verify assumption **A3**, we adopt the same two-regime framework defined in the overdamped case and analyse convergence within each regime separately.

$$\begin{aligned} &2 \|b(t, (x_1, v_1), y) - b(t, (x_2, v_2), y)\| \| (x_1, v_1) - (x_2, v_2) \| + \|\sigma(t, x_1) - \sigma(t, x_2)\|^2 \\ &\leq \begin{cases} 2 \left\| \begin{pmatrix} M^{-1}(v_1 - v_2) \\ \mathbb{1}_{\{0 \leq \lambda_{\kappa t} < \lambda_{\delta}\}} \frac{1}{\sqrt{1 - \lambda_{\delta}}} \left(\nabla \log \nu \left(\frac{x_1 - y}{\sqrt{1 - \lambda_{\delta}}} \right) - \nabla \log \nu \left(\frac{x_2 - y}{\sqrt{1 - \lambda_{\delta}}} \right) \right) - \Gamma M^{-1}(v_1 - v_2) \\ + \left(\mathbb{1}_{\{\lambda_{\delta} \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \frac{1}{\sqrt{1 - \lambda_{\kappa t}}} \left(\nabla \log \nu \left(\frac{x_1 - y}{\sqrt{1 - \lambda_{\kappa t}}} \right) - \nabla \log \nu \left(\frac{x_2 - y}{\sqrt{1 - \lambda_{\kappa t}}} \right) \right) \right) \end{pmatrix} \right\| \| (x_1, v_1) - (x_2, v_2) \| & \text{Regime 1} \\ 0 & \text{Regime 2} \end{cases} \\ &\leq 2 \sqrt{M^{-2} (1 + 2\Gamma^2) \|v_1 - v_2\|^2 + \frac{2L_{\nu}^2}{(1 - \tilde{\lambda})^2} \|x_1 - x_2\|^2} \| (x_1, v_1) - (x_2, v_2) \| \\ &\leq 2 \max \left\{ M^{-1} \sqrt{1 + 2\Gamma^2}, \frac{\sqrt{2}L_{\nu}}{1 - \tilde{\lambda}} \right\} \| (x_1, v_1) - (x_2, v_2) \|^2 \\ &\quad \|b(t, (x, v), y_1) - b(t, (x, v), y_2)\| \\ &\leq \begin{cases} \left\| \begin{pmatrix} 0 \\ \mathbb{1}_{\{0 \leq \lambda_{\kappa t} < \lambda_{\delta}\}} \left(\nabla \log \nu \left(\frac{x-y_1}{\sqrt{1 - \lambda_{\delta}}} \right) - \nabla \log \nu \left(\frac{x-y_2}{\sqrt{1 - \lambda_{\delta}}} \right) \right) \\ 0 \\ + \left(\mathbb{1}_{\{\lambda_{\delta} \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \left(\nabla \log \nu \left(\frac{x-y_1}{\sqrt{1 - \lambda_{\kappa t}}} \right) - \nabla \log \nu \left(\frac{x-y_2}{\sqrt{1 - \lambda_{\kappa t}}} \right) \right) \right) \end{pmatrix} \right\| & \text{Regime 1} \\ \left\| \begin{pmatrix} 0 \\ \frac{1}{\sqrt{\lambda_{\kappa t}}} \left(\nabla V_{\pi} \left(\frac{y_1}{\sqrt{\lambda_{\kappa t}}} \right) - \nabla V_{\pi} \left(\frac{y_2}{\sqrt{\lambda_{\kappa t}}} \right) \right) \end{pmatrix} \right\| & \text{Regime 2} \end{cases} \end{aligned}$$

$$\leq \max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\} \|y_1 - y_2\|.$$

For the following, we can assume without loss of generality that $s \leq t$

$$\begin{aligned} & \|b(t, (x, v), y) - b(s, (x, v), y)\| \\ & \leq \begin{cases} \mathbb{1}_{\{0 \leq \lambda_{\kappa s} < \lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \left| \frac{1}{\sqrt{1-\lambda_{\kappa t}}} - \frac{1}{\sqrt{1-\lambda_\delta}} \right| \left(L_\nu \frac{\|x\| + \|y\|}{\sqrt{1-\tilde{\lambda}}} + \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) \right\| \right) \\ \quad + \mathbb{1}_{\{\lambda_\delta \leq \lambda_{\kappa t}, \lambda_{\kappa s} < \tilde{\lambda}\}} \left| \frac{1}{\sqrt{1-\lambda_{\kappa t}}} - \frac{1}{\sqrt{1-\lambda_\kappa s}} \right| \left(L_\nu \frac{\|x\| + \|y\|}{\sqrt{1-\tilde{\lambda}}} + \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) \right\| \right) \end{cases} \quad \text{Regime 1} \\ & \leq \begin{cases} \mathbb{1}_{\{\tilde{\lambda} \leq \lambda_{\kappa t}, \lambda_{\kappa s} < 1-\lambda_\delta\}} \left| \frac{1}{\sqrt{\lambda_{\kappa t}}} - \frac{1}{\sqrt{\lambda_{\kappa s}}} \right| \left(L_\pi \frac{\|y\|}{\sqrt{\tilde{\lambda}}} + \left\| \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) \right\| \right) \end{cases} \quad \text{Regime 2} \\ & \leq \begin{cases} \frac{C|t-s|^{\gamma_1}}{(1-\tilde{\lambda})^2} L_\nu (\|x\| + \|y\|) \\ \frac{C|t-s|^{\gamma_1}}{\tilde{\lambda}^2} L_\pi \left(\|y\| + \frac{\|m^*\| \sqrt{\tilde{\lambda}}}{2} \right) \end{cases} \quad \text{Regime 1} \\ & \leq \begin{cases} \frac{L_\nu}{(1-\tilde{\lambda})^2} \frac{L_\pi}{\tilde{\lambda}^2} |t-s|^{\gamma_1} \left(\|y\| + \|x\| + \frac{\|m^*\| \sqrt{\tilde{\lambda}}}{2} \right), \end{cases} \end{aligned}$$

$$2\langle (x, v), b(t, (x, v), y) \rangle$$

$$\begin{aligned} & \leq \begin{cases} 2M^{-1} \langle x, v \rangle - 2\Gamma M^{-1} \|v\|^2 + \mathbb{1}_{\{0 \leq \lambda_{\kappa t} < \lambda_\delta\}} \frac{2\|v\|}{\sqrt{1-\lambda_\delta}} \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_\delta}} \right) \right\| \\ \quad + \mathbb{1}_{\{\lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \frac{2\|v\|}{\sqrt{1-\lambda_{\kappa t}}} \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) \right\| \end{cases} \quad \text{Regime 1} \\ & \leq \begin{cases} 2M^{-1} \langle x, v \rangle - 2\Gamma M^{-1} \|v\|^2 + \frac{2\|v\|}{\sqrt{\lambda_{\kappa t}}} \left\| \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) \right\| \end{cases} \quad \text{Regime 2} \\ & \leq \begin{cases} M^{-1} (\|x\|^2 + \|v\|^2) - 2\Gamma M^{-1} \|v\|^2 + \frac{L_\nu}{1-\tilde{\lambda}} (2\|v\|^2 + \|x\|^2 + \|y\|^2) \\ M^{-1} (\|x\|^2 + \|v\|^2) - 2\Gamma M^{-1} \|v\|^2 + \frac{L_\pi}{\tilde{\lambda}} (2\|v\|^2 + \|y\|^2 + \tilde{\lambda} \|m^*\|^2) \end{cases} \quad \text{Regime 1} \\ & \leq \max \left\{ M^{-1}, \frac{2L_\nu}{1-\tilde{\lambda}}, \frac{2L_\pi}{\tilde{\lambda}} \right\} \max \{1, \tilde{\lambda} \|m^*\|^2\} (1 + \|x\|^2 + \|v\|^2) + \max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\} \|y\|^2, \end{aligned}$$

$$\|b(t, (x, v), y)\|$$

$$\begin{aligned} & \leq \begin{cases} \sqrt{M^{-2} (1 + 2\Gamma^2) \|v\|^2 + \mathbb{1}_{\{0 \leq \lambda_{\kappa t} < \lambda_\delta\}} \frac{2}{1-\tilde{\lambda}} \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_\delta}} \right) \right\|^2 + \mathbb{1}_{\{\lambda_\delta \leq \lambda_{\kappa t} < \tilde{\lambda}\}} \frac{2}{1-\tilde{\lambda}} \left\| \nabla \log \nu \left(\frac{x-y}{\sqrt{1-\lambda_{\kappa t}}} \right) \right\|^2} \\ \sqrt{M^{-2} (1 + 2\Gamma^2) \|v\|^2 + \frac{2}{\tilde{\lambda}} \left\| \nabla V_\pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right) \right\|^2} \end{cases} \quad \text{Regime 1} \\ & \leq \begin{cases} \max \left\{ M^{-1} \sqrt{1 + 2\Gamma^2}, \frac{2L_\nu}{1-\tilde{\lambda}} \right\} \|(x, v)\| + \frac{2L_\pi \|m^*\|}{\sqrt{\tilde{\lambda}}} + 2 \max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\} \|y\| \\ \leq \frac{2L_\pi}{\sqrt{\tilde{\lambda}}} \|m^*\| + \max \left\{ M^{-1} \sqrt{1 + 2\Gamma^2}, \frac{2L_\nu}{1-\tilde{\lambda}}, \frac{2L_\pi}{\tilde{\lambda}} \right\} (\|(x, v)\| + \|y\|), \end{cases} \quad \text{Regime 2} \end{aligned}$$

$$\|\sigma(t, x)\|^2 = 2\Gamma.$$

This concludes that assumption **A3** holds with $\theta_1 = 0$, $\theta_2 = \theta_3 = \theta_5 = \theta_6 = 1$, $\theta_4 = 2$ and $Z_T = \|m^*\| \sqrt{\tilde{\lambda}}/2$ is a constant random variable in both regimes. Besides, the constant ε_0 in the theorem statement is given by

$$\varepsilon_0 = \frac{\lambda_2}{\lambda_1} = \frac{\frac{1}{1-\lambda_\delta} + \frac{2a_\pi}{\lambda_\delta}}{\max \left\{ \frac{L_\nu}{1-\tilde{\lambda}}, \frac{L_\pi}{\tilde{\lambda}} \right\}}.$$

Underdamped MULTCDIFF (10) or (21) From Appendix B.2, the coefficients of the slow-fast system are given by the following expressions

$$b(t, (x, v), y) = \begin{pmatrix} -4\Gamma^{-2}v \\ x + 4\Gamma^{-1}v - \frac{2\Gamma}{\sigma_{T-t}^2} \left(v - \left(m_{3,T-t}y + \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} (x - m_{1,T-t}y) \right) \right) \end{pmatrix}$$

$$\sigma(t, (x, v)) = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{2\Gamma} \end{pmatrix} \otimes I_d,$$

$$f(t, (x, v), y) = \nabla \log \pi(y) - [(m_{1,T-t}y, m_{3,T-t}y) - (x, v)] \tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix},$$

$$g(t, (x, v), y) = \sqrt{2}.$$

where $m_{1,T-t}, m_{3,T-t}, \tilde{\Sigma}_{T-t}$ are defined in Appendix B.2. Substituting these coefficients, it follows that assumption **A3** holds with $\theta_1 = 0, \theta_2 = \theta_3 = \theta_5 = \theta_6 = 1, \theta_4 = 2$ and Z_T is a constant random variable, as shown below.

$$\begin{aligned} & 2\|b(t, (x_1, v_1), y) - b(t, (x_2, v_2), y)\| \|(x_1, v_1) - (x_2, v_2)\| + \|\sigma(t, (x_1, v_1)) - \sigma(t, (x_2, v_2))\|^2 \\ & \leq 2 \left\| \left((x_1 - x_2) \left(1 + \frac{2\Gamma}{\sigma_{T-t}^2} \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} \right) + (v_1 - v_2) \left(4\Gamma^{-1} - \frac{2\Gamma}{\sigma_{T-t}^2} \right) \right) \right\| \|(x_1, v_1) - (x_2, v_2)\| \\ & \leq 2 \sqrt{2 \left(1 + \frac{2\Gamma}{\sigma_{T-t}^2} \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} \right)^2 \|x_1 - x_2\|^2 + \left(16\Gamma^{-4} + 2 \left(4\Gamma^{-1} - \frac{2\Gamma}{\sigma_{T-t}^2} \right)^2 \right) \|v_1 - v_2\|^2} \|(x_1, v_1) - (x_2, v_2)\| \\ & \leq C_t \|(x_1, v_1) - (x_2, v_2)\|^2, \end{aligned}$$

$$\begin{aligned} & \|b(t, (x, v), y_1) - b(t, (x, v), y_2)\| \\ & \leq \left\| \begin{pmatrix} 0 \\ \frac{2\Gamma}{\sigma_{T-t}^2} \left(m_{3,T-t} + \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t} \right) (y_1 - y_2) \end{pmatrix} \right\| \leq C_t \|y_1 - y_2\|, \end{aligned}$$

$$\begin{aligned} & \|b(t, (x, v), y) - b(s, (x, v), y)\| \\ & \leq 2\Gamma \left[\left| \frac{1}{\sigma_{T-t}^2} - \frac{1}{\sigma_{T-s}^2} \right| \|v\| + \left| \frac{\tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1}}{\sigma_{T-t}^2} - \frac{\tilde{\Sigma}_{12,T-s} \tilde{\Sigma}_{11,T-s}^{-1}}{\sigma_{T-s}^2} \right| \|x\| \right. \\ & \quad \left. + \left| \frac{m_{3,T-t} - \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t}}{\sigma_{T-t}^2} - \frac{m_{3,T-s} - \tilde{\Sigma}_{12,T-s} \tilde{\Sigma}_{11,T-s}^{-1} m_{1,T-s}}{\sigma_{T-s}^2} \right| \|y\| \right] \\ & \leq C_T |t - s|^{\tilde{\gamma}_1} (\|y\| + \|(x, v)\|), \end{aligned}$$

where we have used the expressions in Appendix B.2, which imply the existence of a constant $\tilde{\gamma}_1 \in (0, 1]$ such that

$$\begin{aligned} & \left| \frac{1}{\sigma_{T-t}^2} - \frac{1}{\sigma_{T-s}^2} \right|, \left| \frac{\tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1}}{\sigma_{T-t}^2} - \frac{\tilde{\Sigma}_{12,T-s} \tilde{\Sigma}_{11,T-s}^{-1}}{\sigma_{T-s}^2} \right|, \\ & \left| \frac{m_{3,T-t} - \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t}}{\sigma_{T-t}^2} - \frac{m_{3,T-s} - \tilde{\Sigma}_{12,T-s} \tilde{\Sigma}_{11,T-s}^{-1} m_{1,T-s}}{\sigma_{T-s}^2} \right| \leq C_T |t - s|^{\tilde{\gamma}_1}. \end{aligned}$$

$$\begin{aligned} & 2\langle (x, v), b(t, (x, v), y) \rangle \\ & \leq 2 \left(1 - 4\Gamma^{-2} + \frac{2\Gamma}{\sigma_{T-t}^2} \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} \right) \langle x, v \rangle + 2 \left(4\Gamma^{-1} - \frac{2\Gamma}{\sigma_{T-t}^2} \right) \|v\|^2 \\ & \quad + \frac{4\Gamma}{\sigma_{T-t}^2} \left(m_{3,T-t} - \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t} \right) \langle y, v \rangle \\ & \leq \max \left\{ \left(1 - 4\Gamma^{-2} + \frac{2\Gamma}{\sigma_{T-t}^2} \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} \right), 2 \left(4\Gamma^{-1} - \frac{2\Gamma}{\sigma_{T-t}^2} \right), \frac{2\Gamma}{\sigma_{T-t}^2} \left(m_{3,T-t} - \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t} \right) \right\} \|(x, v)\|^2 \\ & \quad + \frac{2\Gamma}{\sigma_{T-t}^2} \left(m_{3,T-t} - \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t} \right) \|y\|^2 \leq C_t \|(x, v)\|^2 + \lambda_1 \|y\|^2, \end{aligned}$$

$$\begin{aligned} & \|b(t, (x, v), y)\| \\ & \leq \sqrt{3 \left(1 + \frac{2\Gamma}{\sigma_{T-t}^2} \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} \right)^2 \|x\|^2 + \left(16\Gamma^{-4} + 3 \left(4\Gamma^{-1} - \frac{2\Gamma}{\sigma_{T-t}^2} \right)^2 \right) \|v\|^2 + \frac{12\Gamma^2}{\sigma_{T-t}^4} \left(m_{3,T-t} - \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t} \right)^2 \|y\|^2} \\ & \leq \max \left\{ \sqrt{3} \left| 1 + \frac{2\Gamma}{\sigma_{T-t}^2} \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} \right|, 4\Gamma^{-2} + \sqrt{3} \left| 4\Gamma^{-1} - \frac{2\Gamma}{\sigma_{T-t}^2} \right| \right\} \|(x, v)\| + \frac{2\sqrt{3}\Gamma}{\sigma_{T-t}^2} \left| m_{3,T-t} - \tilde{\Sigma}_{12,T-t} \tilde{\Sigma}_{11,T-t}^{-1} m_{1,T-t} \right| \|y\| \\ & \leq C_T (\|(x, v)\| + \|y\|), \end{aligned}$$

$$\|\sigma(t, x)\|^2 = 2\Gamma.$$

We finally check that assumption **A4** holds.

$$\begin{aligned} & 2\langle f(t, (x, v), y_1) - f(t, (x, v), y_2), y_1 - y_2 \rangle + \|g(t, (x, v), y_1) - g(t, (x, v), y_2)\|^2 \\ & \leq -2\langle \nabla V_\pi(y_1) - \nabla V_\pi(y_2), y_1 - y_2 \rangle - 2(m_{1,T-t}, m_{3,T-t})\tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} \|y_1 - y_2\|^2 \\ & \leq -2\left(M_\pi + (m_{1,T-t}^2 + m_{3,T-t}^2)\lambda_{\min}(\tilde{\Sigma}_{T-t}^{-1})\right)\|y_1 - y_2\|^2 \leq -\beta\|y_1 - y_2\|^2, \end{aligned}$$

where we have used that, for all t , the precision matrix $\tilde{\Sigma}_{T-t}^{-1}$ is positive definite and $\lambda_{\min}(\tilde{\Sigma}_{T-t}^{-1})$ denotes its smallest eigenvalue.

$$\begin{aligned} \|f(t, (x_1, v_1), y) - f(s, (x_2, v_2), y)\| & \leq \left| (m_{1,T-t}, m_{3,T-t})\tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} - (m_{1,T-s}, m_{3,T-s})\tilde{\Sigma}_{T-s}^{-1} \begin{pmatrix} m_{1,T-s} \\ m_{3,T-s} \end{pmatrix} \right| \|y\| \\ & + \|(x_1, v_1) - (x_2, v_2)\| \left\| \tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} \right\| + \|(x_2, v_2)\| \left\| \tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} - \tilde{\Sigma}_{T-s}^{-1} \begin{pmatrix} m_{1,T-s} \\ m_{3,T-s} \end{pmatrix} \right\| \\ & \leq C_T|t-s|^{\tilde{\gamma}_2} + C_T\|(x_1, v_1) - (x_2, v_2)\| + C_T\|(x_2, v_2)\||t-s|^{\tilde{\gamma}_2} \\ & \leq C_T\left(|t-s|^{\tilde{\gamma}_2} + \|(x_1, v_1) - (x_2, v_2)\|\right)(1 + \|(x_2, v_2)\|), \end{aligned}$$

where the existence of $\tilde{\gamma}_2 \in (0, 1]$ follows from the expressions for $m_{1,T-t}, m_{3,T-t}, \tilde{\Sigma}_{T-t}$ provided in Appendix B.2.

$$\begin{aligned} \|f(t, (x, v), y)\| & \leq \|\nabla \log \pi(y)\| + (m_{1,T-t}, m_{3,T-t})\tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} \|y\| + \left\| (x, v)\tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} \right\| \\ & \leq L_\pi(\|y\| + \|m^*\|) + \sqrt{m_{1,T-t}^2 + m_{3,T-t}^2} \lambda_{\max}(\tilde{\Sigma}_{T-t}^{-1}) \left(\sqrt{m_{1,T-t}^2 + m_{3,T-t}^2} \|y\| + \|(x, v)\| \right) \\ & \leq C_T(1 + \|y\| + \|(x, v)\|), \end{aligned}$$

$$\|g(t, (x_1, v_1), y_1) - g(s, (x_2, v_2), y_2)\| = 0, \quad \|g(t, (x, v), y)\| = \sqrt{2}.$$

Recalling that the strong convexity of V_π implies a dissipativity inequality, we have that for any $k \geq 2$

$$\begin{aligned} 2\langle y, f(t, (x, v), y) \rangle + (k-1)\|g(t, (x, v), y)\|^2 & = -2\langle y, \nabla V_\pi(y) \rangle - 2(m_{1,T-t}, m_{3,T-t})\tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} \|y\|^2 \\ & + 2\left\langle (x, v)\tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix}, y \right\rangle + 2(k-1) \\ & \leq -2(a_\pi\|y\|^2 - b_\pi) - (m_{1,T-t}, m_{3,T-t})\tilde{\Sigma}_{T-t}^{-1} \begin{pmatrix} m_{1,T-t} \\ m_{3,T-t} \end{pmatrix} \|y\|^2 + \|(x, v)\|^2 \left\| \tilde{\Sigma}_{T-t}^{-1/2} \right\|^2 + 2(k-1) \\ & \leq -\lambda_2\|y\|^2 + C_{T,k}(\|(x, v)\|^2 + 1), \end{aligned}$$

which concludes the proof. \square

C.2 Challenges of extension to heavy-tailed diffusions

When considering heavy-tailed diffusions, that is, when the base distribution ν is a Student's t distribution, the stochastic slow-fast systems (4) and (6) could still be used to derive sampling algorithms. However, Theorem 4.1 showing convergence to the averaged dynamics does not apply in this setting, as it relies on exponential ergodicity of the frozen process. In the case where the fast process follows overdamped Langevin dynamics targetting a heavy-tailed distribution, convergence is sub-exponential or polynomial rather than exponential [Wang, 2006, Chapter 4], and thus the frozen process fails to meet the required condition.

As discussed in Section 3.1, instead of using overdamped Langevin dynamics for the fast process, an alternative is to consider different diffusion processes that target the conditional distribution $\rho_{t,x}$ more efficiently. Motivated by the work of He et al. [2024a], we propose employing a natural Itô diffusion that arises in the context of the weighted Poincaré inequality, which has the following expression

$$dY_t = -\frac{1}{\varepsilon}(\alpha + d - 1)\nabla U_{\hat{\rho}_{t,x}}(y)dt + \sqrt{\frac{2U_{\hat{\rho}_{t,x}}}{\varepsilon}}dB_t, \quad (25)$$

$$\hat{\rho}_{t,X_t}(y) \propto \left(\sqrt{1 + \frac{\|x - y\|^2}{\alpha(1 - \lambda_{\kappa t})}} \pi \left(\frac{y}{\sqrt{\lambda_{\kappa t}}} \right)^{-\frac{1}{\alpha+d}} \right)^{-(\alpha+d)} = U_{\hat{\rho}_{t,X_t}}(y)^{-(\alpha+d)},$$

where α denotes the tail index of the noising distribution ν . This results into the following slow-fast system in the underdamped setting

$$\begin{cases} dX_t = M^{-1}V_t dt \\ dV_t = \begin{cases} \left(\frac{1}{\sqrt{1-\lambda_{\kappa t}}} \nabla \log \nu \left(\frac{X_t - Y_t}{\sqrt{1-\lambda_{\kappa t}}} \right) - \Gamma M^{-1}V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } \lambda_{\kappa t} < \tilde{\lambda} \\ \left(-\frac{1}{\sqrt{\lambda_{\kappa t}}} \nabla V_{\pi} \left(\frac{Y_t}{\sqrt{1-\lambda_{\kappa t}}} \right) - \Gamma M^{-1}V_t \right) dt + \sqrt{2\Gamma} dB_t & \text{if } \lambda_{\kappa t} \geq \tilde{\lambda} \end{cases} \\ dY_t = -\frac{1}{\varepsilon}(\alpha + d - 1) \nabla U_{\hat{\rho}_{t,X_t}}(y) dt + \sqrt{\frac{2U_{\hat{\rho}_{t,X_t}}}{\varepsilon}} d\tilde{B}_t. \end{cases} \quad (26)$$

However, although the modified Itô diffusion (25) offers improved convergence properties [He et al., 2024a], it does not guarantee exponential ergodicity of the frozen process under mild conditions. As a result, assumption A4 is not satisfied in this setting. We leave the analysis of the convergence of (26) to the averaged dynamics (5) for future work.

C.3 Bias of the annealed Langevin dynamics

We formally quantify the bias introduced by the underdamped averaged system (5) relative to the true diffusion path which is given by $(\hat{p}_t)_{t \in [0,1/\kappa]}$ with

$$\hat{p}_t(x, v) \propto \exp \left(-\frac{1}{2} v^T M^{-1} v + \log \hat{\mu}_t(x) \right). \quad (27)$$

Theorem C.2 (Theorem 4.2 restated). *Let $\mathbb{Q}_{\text{U-ALD}} = (q_{t,\text{U-ALD}})_{t \in [0,1/\kappa]}$ be the path measure of the diffusion annealed Langevin dynamics (5), and $\mathbb{Q} = (\hat{p}_t)_{t \in [0,1/\kappa]}$ that of a reference SDE such that the marginals at each time have distribution \hat{p}_t defined in (27). If $q_{0,\text{U-ALD}} = \hat{p}_0$, the KL divergence between the path measures is given by*

$$\text{KL}(\mathbb{Q} || \mathbb{Q}_{\text{U-ALD}}) = \frac{\kappa}{4} \mathcal{A}(p) = \frac{\kappa}{4} \mathcal{A}(\mu).$$

Proof. Let \mathbb{Q} be the path measure corresponding to the following reference SDE

$$\begin{cases} d\bar{Z}_t = (M^{-1}\bar{U}_t + \hat{v}_{1,t}(\bar{Z}_t, \bar{U}_t)) dt \\ d\bar{U}_t = (\nabla \log \hat{\mu}_t(\bar{Z}_t) - \Gamma M^{-1}\bar{U}_t + \hat{v}_{2,t}(\bar{Z}_t, \bar{U}_t)) dt + \sqrt{2\Gamma} dB_t, \end{cases} \quad \text{for } t \in [0, 1/\kappa].$$

The vector field $\hat{v} = ((\hat{v}_{1,t}, \hat{v}_{2,t}))_{t \in [0,1/\kappa]}$ is designed such that $(\bar{Z}_t, \bar{U}_t) \sim \hat{p}_t$ for all $t \in [0, 1/\kappa]$. Using the Fokker-Planck equation, we have that

$$\begin{aligned} \partial_t \hat{p}_t &= -\nabla_x \cdot (\hat{p}_t(M^{-1}v + \hat{v}_{1,t})) - \nabla_v \cdot (\hat{p}_t(\nabla \log \hat{\mu}_t - \Gamma M^{-1}v + \hat{v}_{2,t})) + \Gamma \Delta_v \hat{p}_t \\ &= -\nabla_x \cdot (\hat{p}_t \hat{v}_{1,t}) - \nabla_v \cdot (\hat{p}_t \hat{v}_{2,t}), \quad t \in [0, 1/\kappa]. \end{aligned}$$

This implies that $\hat{v}_t = (\hat{v}_{1,t}, \hat{v}_{2,t})$ satisfies the continuity equation and hence generates the curve of probability measures $(\hat{p}_t)_t$. Leveraging Lemma A.3, we choose \hat{v} to be the one that minimises the $L^2(\hat{p}_t)$ norm, resulting in $\|\hat{v}_t\|_{L^2(\hat{p}_t)} = |\dot{\hat{p}}|_t$ being the metric derivative. Using the form of Girsanov's theorem given in Lemma A.4 we have

$$\begin{aligned} \text{KL}(\mathbb{Q} || \mathbb{Q}_{\text{U-ALD}}) &= \frac{1}{4} \mathbb{E}_{\mathbb{Q}} \left[\int_0^{1/\kappa} \|\hat{v}_t(\bar{X}_t, \bar{V}_t)\|^2 dt \right] = \frac{1}{4} \int_0^{1/\kappa} \|\hat{v}_t(\bar{X}_t, \bar{V}_t)\|_{L^2(\hat{p})}^2 dt = \frac{1}{4} \int_0^{1/\kappa} |\dot{\hat{p}}|_t^2 dt \\ &= \frac{\kappa}{4} \int_0^1 |\dot{\hat{p}}|_t^2 dt = \frac{\kappa}{4} \mathcal{A}(p), \end{aligned}$$

where we have used that $|\dot{\hat{p}}|_t = \kappa |\dot{p}|_t$ and the change of variable formula.

To conclude, we note from (27) that the position x and the velocity v variable are independent, and that the marginal distribution of the velocity v remains constant along the diffusion path $(\hat{p}_t)_{t \in [0,1]}$. As a result, the metric derivative simplifies to

$$|\dot{p}|_t = \lim_{\delta \rightarrow 0} \frac{W_2(p_{t+\delta}, p_t)}{|\delta|} = \lim_{\delta \rightarrow 0} \frac{W_2(\mu_{t+\delta}, \mu_t)}{|\delta|} = |\dot{\mu}|_t.$$

Consequently, the action $\mathcal{A}(p)$ satisfies

$$\mathcal{A}(p) = \mathcal{A}(\mu),$$

where $\mu = (\mu_t)_{t \in [0,1]}$ is the diffusion path defined in (1). \square

Finally note that from the data processing inequality, it follows that the KL divergence between the marginals at final time is bounded by $\text{KL}(\hat{p}_{1/\kappa} \parallel q_{1/\kappa, \text{U-ALD}}) \leq \text{KL}(\mathbb{Q} \parallel \mathbb{Q}_{\text{U-ALD}})$.

By Cordero-Encinar et al. [2025, Lemmas 3.3 and 4.2], the action $\mathcal{A}(\mu)$ is bounded when the base distribution ν is either a Gaussian or a Student's t distribution, provided that the target distribution π has finite second order moment and the annealing schedule λ_t satisfies the following assumption.

A5. *Let $\lambda_t : \mathbb{R}^+ \rightarrow [0, 1]$ be non-decreasing in t and weakly differentiable, such that if $\nu \sim \mathcal{N}(0, I)$ there exists a constant C_λ satisfying either of the following conditions*

$$\max_{t \in [0, T]} |\partial_t \log \lambda_t| \leq C_\lambda$$

or

$$\max_{t \in [0, T]} \left| \frac{\partial_t \lambda_t}{\sqrt{\lambda_t(1 - \lambda_t)}} \right| \leq C_\lambda, \quad (28)$$

or if ν follows a Student's t distribution, then condition (28) holds for some constant C_λ .

D Experiment details

The code to reproduce our experiments is available at https://github.com/paulaoak/sampling_by_averaging.git. All experiments were implemented using JAX.

D.1 Target distributions

Mixture of Gaussians (MoG) and mixture of Student's t (MoS) distributions The 8 mixture of Gaussians distribution (8-MoG) consists of 8 equally weighted Gaussian distributions with mean $m_i = 10 \times (1 + \cos(2\pi i/8), 1 + \sin(2\pi i/8))$ for $i \in \{9, \dots, 7\}$ and covariance $0.7I_2$. We have shifted the distribution instead of considering the usual benchmark centred at 0 to make sampling more challenging. We note that some of the baselines get stuck in modes close to the initial standard Gaussian distribution, unlike our methods.

For the mixture of 40 Gaussians in dimensions 2 and 50, the modes are equally weighted, with means sampled uniformly from a hypercube of side length 40, and all the covariance matrices set to the identity matrix I .

The mixture of Student's t distributions consists of 10 standard Student's t distributions, each with 2 degrees of freedom (t_2). Following [Blessing et al., 2024, Chen et al., 2025], the mean of each component is sampled uniformly from a hypercube with a side length of 10. We evaluate this benchmark in dimensions 2, 10 and 50.

Rings This distribution is defined by the inverse polar reparameterisation of a distribution p_z which has itself a decomposition into two univariate marginals: p_r and p_θ . The radial component p_r is a mixture of 4 Gaussian distributions $\mathcal{N}(i + 1, 0.15^2)$ for $i \in \{0, \dots, 3\}$ describing the radial positions. The angular component p_θ is a uniform distribution over the interval $[0, 2\pi]$.

Funnel The density of this distribution is given by $\pi(x) \propto \mathcal{N}(x_1; 0, \eta^2) \prod_{i=2}^d \mathcal{N}(x_i; 0, e^{x_1})$ for $x = (x_i)_{i=1}^{10} \in \mathbb{R}^{10}$ with $\eta = 3$ [Neal, 2003].

Double well potential (DW) The unnormalised density of the d -dimensional DW is given by $\pi \propto \exp(-\sum_{i=1}^m (x_i^2 - \delta)^2 - \sum_{i=m+1}^d x_i^2)$ with $m \in \mathbb{N}$ and a separation parameter $\delta \in (0, \infty)$. This distribution has $2m$ modes, and larger values of δ make sampling more challenging due to higher energy barriers. Ground truth samples are obtained using rejection sampling with a Gaussian mixture proposal distribution [Midgley et al., 2023].

Bayesian logistic regression examples We consider binary classification problems on two benchmark datasets: Ionosphere (dimension 35) and Sonar (dimension 61). Given a training dataset $\mathcal{D} = \{(x_j, y_j)\}_{j=1}^N$ where $x_j \in \mathbb{R}^d$ and y_j , the posterior distribution over the model parameters (w, b) , $w \in \mathbb{R}^d$ and $b \in \mathbb{R}$, is defined by

$$p(w, b | \mathcal{D}) \propto p(w, b) \prod_{j=1}^N \text{Bernoulli}(y_j, \sigma(w^\top x_j + b)),$$

where σ denotes the sigmoid function. Following [Grenioux et al., 2024], we place independent Gaussian priors on the parameters $p(w, b) = \mathcal{N}(w; 0, I_d) \mathcal{N}(b; 0, 2.5^2)$.

Statistical physics model: ϕ^4 distribution The goal is to sample metastable states of the stochastic Allen-Cahn equation ϕ^4 model (dimension 100) [Albergo et al., 2019, Gabrie et al., 2022]. The ϕ^4 model is a continuous relaxation of the Ising model which is used to study phase transitions in statistical mechanics. Following [Albergo et al., 2019, Gabrie et al., 2022, Grenioux et al., 2024], we consider a version of the model discretised on a 1-dimensional grid of size $d = 100$. Each configuration of the model is represented by a d -dimensional vector $(\phi_i)_{i=1}^d$. We clip the field to 0 at the boundaries by defining $\phi_0 = \phi_{d+1} = 0$. The negative log-density of the distribution is defined as

$$\ln \pi_h(\phi) = -\beta \left(\frac{ad}{2} \sum_{i=1}^{d+1} (\phi_i - \phi_{i+1})^2 - \frac{1}{4ad} \sum_{i=1}^d (1 - \phi_i^2)^2 + h\phi_i \right).$$

We fix the parameters $a = 0.1$ and $\beta = 20$ to ensure a bimodal regime, and vary the value of h . We denote by w_+ the statistical occurrence of configurations with $\phi_{d/2} > 0$ and w_- the statistical occurrence of configurations with $\phi_{d/2} < 0$. At $h = 0$, the distribution is invariant under the symmetry $\phi \rightarrow -\phi$, so we expect $w_+ = w_-$. For $h > 0$, the negative mode becomes dominant.

When the dimension d is large, the relative probabilities of the two modes can be estimated by Laplace approximations at 0-th and 2-nd order. Let ϕ_+^h and ϕ_-^h denote the local maxima of the distribution, these approximations yield respectively

$$\frac{w_-}{w_+} \approx \frac{\pi_h(\phi_-^h)}{\pi_h(\phi_+^h)}, \quad \frac{w_-}{w_+} \approx \frac{\pi_h(\phi_-^h) \times |\det H_h(\phi_-^h)|^{-1/2}}{\pi_h(\phi_+^h) \times |\det H_h(\phi_+^h)|^{-1/2}},$$

where H_h is the Hessian of the function $\phi \rightarrow \ln \pi_h(\phi)$.

For this last benchmark, we only compare our method with SLIPS [Grenioux et al., 2024], as it is the only baseline capable of accurately recovering samples from the target distribution while preserving the correct relative mode weights.

D.2 Evaluation metrics

To evaluate the quality of the generated samples, we consider the entropy regularised Wasserstein-2 distance [Peyré and Cuturi, 2019], with regularisation parameter $\varepsilon = 0.05$, for distributions where true samples are available. This metric can be efficiently computed in JAX using the OTT library [Cuturi et al., 2022].

For the Funnel benchmark, we assess the sample quality using the Kolmogorov-Smirnov distance. Specifically, we use the sliced version introduced in Grenioux et al. [2023, Appendix D.1].

In the Bayesian logistic regression tasks, performance is measured via the mean predictive log-likelihood, computed as $p(w, b | \mathcal{D}_{\text{test}})$, where $\mathcal{D}_{\text{test}}$ is a held-out test dataset not used during training.

D.3 Algorithms and hyperparameters

All algorithms are initialised with samples drawn from a standard Gaussian distribution to ensure a fair comparison. This contrasts with the approach in Grenioux et al. [2024], where the initial samples are drawn using prior knowledge of the target distribution via a quantity denoted as R_π , which is upper bounded by the scalar variance of the target π . While we observe that initialising from this distribution can improve performance for our methods, such information is typically unavailable in real-world scenarios. Nevertheless, for the ϕ^4 model, we consider the same implementation for SLIPS as described in Grenioux et al. [2024], given the complexity of this benchmark.

Besides, to ensure consistency, all algorithms use the same number of energy evaluations, as detailed for each benchmark distribution in Table 3.

Table 3: Number of energy evaluations for each benchmark

	8-MoG ($d = 2$)	40-MoG ($d = 2$)	40-MoG ($d = 50$)	Rings ($d = 2$)	Funnel ($d = 10$)
Energy evaluations	3×10^5	3×10^5	5×10^6	3×10^5	5×10^5
	5-dim DW ($m = 5, \delta = 4$)	10-dim DW ($m = 5, \delta = 3$)	50-dim DW ($m = 5, \delta = 2$)	Ionosphere ($d = 35$)	Sonar ($d = 61$)
Energy evaluations	3×10^5	5×10^5	5×10^6	1×10^6	3×10^6
	ϕ^4 model ($d = 100$)	10-MoS ($d = 2$)	10-MoS ($d = 10$)	10-MoS ($d = 50$)	
Energy evaluations	1×10^8	3×10^5	5×10^5	1×10^6	

Hyperparameter selection For each baseline algorithm, we perform a grid search over a predefined set of hyperparameter values. Selection is based on the corresponding performance metric, computed using 4096 samples. The selected hyperparameters for each algorithm are summarised below.

- The SMC and AIS algorithms define a sequence of annealed distributions μ_k for $k \in [0, K]$. At each intermediate distribution, we perform $n = 64$ MCMC steps. The parameter n_{MCMC} is selected from the interval $[20, 100]$ using a grid step of 4. The total number of distributions K is chosen such that the product $n \times K$ matches the number of target evaluations specified in Table 3.
- For ULMC algorithm, we tune three hyperparameters: the mass M , the friction coefficient Γ and the step size h . In all benchmarks, we fix the mass to the identity matrix $M = I$. The step size is chosen within the following grid $h \in \{0.001, \dots, 0.009, 0.01, \dots, 0.09, 0.10, 0.11, \dots, 0.20\}$. The selected values for the step sizes h are provided in Table 4. Additionally, we use a non-constant friction coefficient. Specifically, when running the algorithm for L steps, the friction coefficient remains constant at $\Gamma_{\min} = 1$ for the first $L/2$ steps and then increases linearly from Γ_{\min} to $\Gamma_{\max} = 5$ during the remaining steps. The values of Γ_{\min} and Γ_{\max} are chosen within the grid $\{0.001, 0.005, 0.01, \dots, 0.09, 0.1, \dots, 0.9, 1, 1.5, \dots, 10\}$.
- For the PT algorithm, the number of chains K is selected from the interval $[1, 20]$ with a grid step of 2. The corresponding temperatures are chosen to be equally spaced on a logarithmic scale, with the minimum temperature fixed at 1 and the maximum temperature chosen from the grid $\{10^2, 10^3, 10^4, 10^5, 10^6, 10^7\}$. In our experiments, we use $K = 5$ parallel chains with temperatures $\{1.0, 5.6, 31.6, 177.8, 1000.0\}$. Each chain employs a Hamiltonian Monte Carlo sampler with n leapfrog steps per sample and a step size h specified in Table 4. The grid for the step size is the same as that of ULMC. The number of leapfrog steps n is chosen so that the total number of target evaluations matches the computational budget defined in Table 3.

Table 4: Selected step sizes for ULMC and PT algorithm across experiments.

	8-MoG	40-MoG 2-dim	40-MoG 50-dim	Rings	Funnel
ULMC	0.05	0.05	0.01	0.05	0.03
PT	0.10	0.10	0.03	0.10	0.05
	5-dim DW	10-dim DW	50-dim DW	Ionosphere	Sonar
ULMC	0.05	0.05	0.01	0.04	0.04
PT	0.10	0.05	0.05	0.05	0.03

- The DiGS algorithm uses K noise levels ranging from α_K to α_1 . The number of noise levels K is selected within the interval $[1, 20]$ with a grid step of 2. The maximum and minimum noise levels, α_1 and α_K , respectively, are chosen from the grid $\{0.05, \dots, 0.09, 0.1, 0.2, \dots, 1\}$. At each noise level, we perform n_{Gibbs} Gibbs sweeps, each consisting of n_{MALA} denoising steps using MALA with step size h . The number of Gibbs sweeps n_{Gibbs} is selected from the interval $[50, 500]$ using a grid step of 50, while the grid for the step size is the same as that of ULMC. The final values of K ,

α_K , α_1 , n_{Gibbs} , and h are provided in Table 5 for each experiment. The number of MALA steps, n_{MALA} , is determined based on the computational budget.

Table 5: Selected hyperparameters for DiGS algorithm across experiments.

	8-MoG	40-MoG 2-dim	40-MoG 50-dim	Rings	Funnel
K	3	1	5	1	3
α_K	0.1	0.1	0.1	0.1	0.1
α_1	0.5	0.1	0.9	0.1	0.5
n_{Gibbs}	200	300	100	300	200
h	0.05	0.10	0.01	0.10	0.03
	5-dim DW	10-dim DW	50-dim DW	Ionosphere	Sonar
K	1	3	5	3	3
α_K	0.1	0.1	0.1	0.1	0.1
α_1	0.1	0.5	0.8	0.5	0.5
n_{Gibbs}	300	200	100	200	200
h	0.06	0.01	0.005	0.03	0.01

- The RDMC algorithm has a hyperparameter T , corresponding to the final time of the OU process, its value is provided in Table 6. All other parameters follow the implementation detailed in [Grenioux et al., 2024].

Table 6: Selected hyperparameters for RDMC algorithm across experiments.

	8-MoG	40-MoG 2-dim	40-MoG 50-dim	Rings	Funnel
T	$-\log(0.80)$	$-\log(0.75)$	$-\log(0.70)$	$-\log(0.85)$	$-\log(0.90)$
	5-dim DW	10-dim DW	50-dim DW	Ionosphere	Sonar
T	$-\log(0.70)$	$-\log(0.70)$	$-\log(0.75)$	$-\log(0.95)$	$-\log(0.95)$

- For the SLIPS algorithm, we adopt the implementation and hyperparameters provided in the original paper [Grenioux et al., 2024]. To ensure a fair comparison with other algorithms initialised from a standard Gaussian distribution, we set the scalar variance parameter $R_\pi = \sqrt{d}$, yielding $\sigma = R_\pi/\sqrt{d} = 1$. While this choice may degrade SLIPS' performance compared to using an optimally tuned R_π , estimating such a parameter in practical scenarios is often non trivial. Nevertheless, for completeness, we report below the performance of SLIPS with its optimal parameters alongside our algorithms. In this setting, the results are comparable. However, a key advantage of our method is that it does not require estimation of the scalar variance. Moreover, if we initialise our algorithms using the same informed choice R_π by setting the base distribution $\nu \sim \mathcal{N}(0, \sigma^2 I)$, with $\sigma = R_\pi/\sqrt{d}$, we observe a performance improvement, particularly in the overdamped regime.
- For all numerical experiments in the main text, we use the underdamped version of our algorithms, as it yields improved performance. For MULTALMC, the required hyperparameters include the schedule function λ_t , the values of λ_δ and $\tilde{\lambda}$, the mass matrix M , ε , the friction coefficient Γ , the step size h , and the number of SROCK steps s . The grid for the step size is the same as that of ULMC. The scale separation parameter ε is chosen within the grid $\{0.001, 0.005, 0.01, \dots, 0.09, 0.10, 0.11, \dots, 0.20\}$. The parameter $\tilde{\lambda}$ is chosen from the interval $[0.3, 0.7]$ using a grid step of 0.05. In all experiments, we select $\lambda_\delta = 0.01$, $\tilde{\lambda} = 0.6$, $M = I$, and $s = 5$. Additionally, we consider a time-dependent friction coefficient. Specifically, when running the algorithm for L steps, the friction coefficient remains constant at Γ_{\min} for the first $L/2$ steps and then increases linearly from Γ_{\min} to Γ_{\max} during the remaining steps. The values of Γ_{\min} and Γ_{\max} are selected within the grid $\{0.001, 0.005, 0.01, \dots, 0.09, 0.1, \dots, 0.9, 1, 1.5, \dots, 10\}$. The mass matrix is set to the identity. It is worth noting that the number of SROCK steps can be reduced, which may result in a small compromise in performance. This effect is further analysed in Appendix E.2. Lastly, the number of iterations is determined based on the computational budget specified in Table 3.

Table 7: Metrics for different benchmarks averaged across 30 seeds. The metric for the mixture of Gaussian (MoG), Rings and the double well potential (DW) is the entropy regularised Wasserstein-2 distance (with regularisation parameter 0.05), the metric for the Funnel is the sliced Kolmogorov-Smirnov distance and the metric for the Bayesian logistic regression on Ionosphere and Sonar datasets is the average predictive posterior log-likelihood on a test dataset. We compare the performance of our algorithms (initialised with a standard Gaussian distribution) with that of SLIPS with an optimal value of the parameter R_π , refer to as SLIPS (R_π).

Algorithm	8-MoG (↓) ($d = 2$)	40-MoG (↓) ($d = 2$)	40-MoG (↓) ($d = 50$)	Rings (↓) ($d = 2$)	Funnel (↓) ($d = 10$)
SLIPS (R_π)	0.66 ± 0.11	0.98 ± 0.06	20.85 ± 0.63	0.19 ± 0.02	0.029 ± 0.007
MULTALMC	0.65 ± 0.07	0.91 ± 0.04	20.58 ± 0.71	0.18 ± 0.02	0.032 ± 0.005
MULTCDIFF	0.62 ± 0.10	0.93 ± 0.06	20.13 ± 0.59	0.19 ± 0.03	0.031 ± 0.005
Algorithm	5-dim DW (↓) ($m = 5, \delta = 4$)	10-dim DW (↓) ($m = 5, \delta = 3$)	50-dim DW (↓) ($m = 5, \delta = 2$)	Ionosphere (↑) ($d = 35$)	Sonar (↑) ($d = 61$)
SLIPS (R_π)	2.05 ± 0.30	4.45 ± 0.36	14.99 ± 0.59	-86.72 ± 0.10	-109.38 ± 0.12
MULTALMC	2.08 ± 0.16	4.48 ± 0.40	14.03 ± 0.63	-86.85 ± 0.09	-109.05 ± 0.13
MULTCDIFF	1.95 ± 0.24	4.23 ± 0.37	13.98 ± 0.56	-86.33 ± 0.10	-109.60 ± 0.21

Table 8: Selected hyperparameters for MULTALMC algorithm across experiments.

	8-MoG	40-MoG 2-dim	40-MoG 50-dim	Rings	Funnel
λ_t	Linear	Linear	Linear	Linear	Cosine-like
ε	0.10	0.05	0.05	0.10	0.07
$\Gamma_{\min}, \Gamma_{\max}$	0.07, 0.5	0.01, 0.5	0.01, 0.5	0.1, 0.5	0.01, 0.5
h	0.005	0.005	0.001	0.005	0.003
	5-dim DW	10-dim DW	50-dim DW	Ionosphere	Sonar
λ_t	Cosine-like	Cosine-like	Cosine-like	Cosine-like	Cosine-like
ε	0.10	0.07	0.05	0.10	0.10
$\Gamma_{\min}, \Gamma_{\max}$	0.01, 0.5	0.01, 0.5	0.01, 0.5	0.1, 0.5	0.1, 0.5
h	0.005	0.001	0.001	0.005	0.005
	ϕ^4 model	10-MoS 2-dim	10-MoS 10-dim	10-MoS 50-dim	
λ_t	Cosine-like	Linear	Linear	Linear	
ε	0.05	0.10	0.10	0.05	
$\Gamma_{\min}, \Gamma_{\max}$	0.01, 0.5	0.1, 0.5	0.1, 0.5	0.1, 0.5	
h	0.001	0.005	0.005	0.001	

On the other hand, for MULTCDIFF, the following hyperparameters need to be specified: the friction coefficient Γ , the scale separation parameter ε , the final time T of the OU process, the time discretisation $0 = T_0 < \dots < T_L = T$, and the number of SROCK steps s . The values of Γ , ε and s are set to be the same as in the MULTALMC algorithm. The time discretisation is chosen such that the difference $\lambda_{T_{l+1}} - \lambda_{T_l}$ is constant, where λ_t denotes the OU schedule. The number of iterations L is determined based on the computational budget for each benchmark.

D.4 Computation time

All experiments were conducted on a GPU server consisting of eight Nvidia GeForce RTX 3090 Ti GPU cards, 896 GB of memory and 14TB of local on-server data storage. Each GPU has 10496 cores as well as 24 GB of memory.

Recall that to ensure a fair comparison, we fixed the number of energy evaluations across all algorithms. As a result, the runtime of ULMC, DiGS, RDMC, SLIPS, and our proposed methods, MULTALMC and MULTCDIFF, are broadly similar. In contrast, algorithms such as SMC, AIS, and PT exhibit longer runtimes due to their accept/reject steps. This is demonstrated in Figure 2,

which shows average computation times (over 30 random seeds) for each method on the 40-MoG benchmark in 50 dimensions.

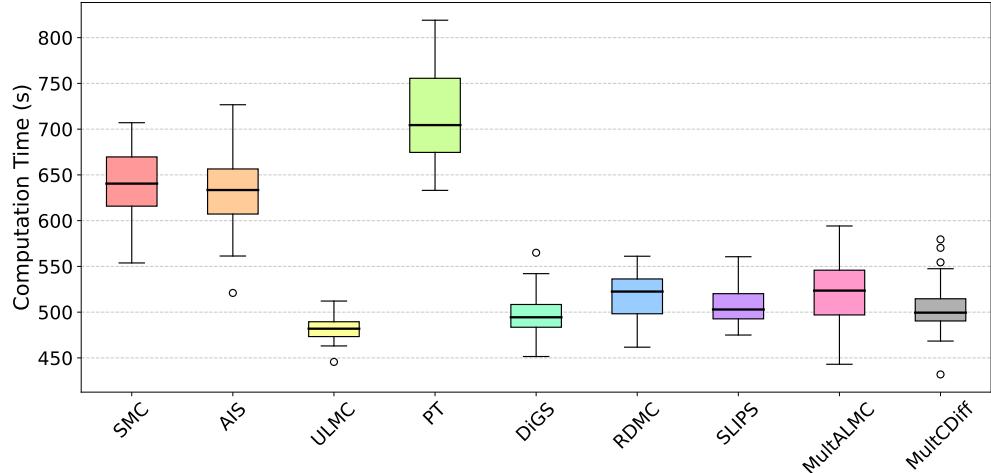


Figure 2: Boxplots of computation times for different algorithms on the 40-MoG benchmark in 50 dimensions, averaged over 30 random seeds.

In addition, we analyse how the runtime of our methods scales with dimensionality. To that end, Table 9 includes performance metrics, number of energy evaluations, and computation times for the 40-MoG benchmark across different dimensions, averaged over 30 random seeds. The performance metric used is the entropy-regularised Wasserstein-2 distance, W_2^δ , with regularisation parameter $\delta = 0.05$.

Table 9: Number of energy evaluations, entropy regularised Wasserstein-2 distance (with regularisation parameter 0.05) and computations times for our sampling methods, MULTALMC (Algorithm 3) and MULTCDIFF (Algorithm 4) evaluated on the 40-MoG benchmark across different dimensions. The results are averaged over 30 random seeds.

		d = 2	d = 5	d = 10	d = 20	d = 50
	Energy evaluations	3×10^5	6×10^5	1×10^6	2×10^6	5×10^6
W_2^δ	MULTALMC	0.91 ± 0.04	1.64 ± 0.12	4.10 ± 0.27	8.76 ± 0.34	20.58 ± 0.71
	MULTCDIFF	0.93 ± 0.06	1.59 ± 0.20	3.92 ± 0.36	8.68 ± 0.41	20.13 ± 0.59
Time (s)	MULTALMC	35 ± 5	57 ± 9	111 ± 13	223 ± 28	517 ± 28
	MULTCDIFF	33 ± 6	58 ± 7	104 ± 15	231 ± 21	502 ± 30

E Additional numerical experiments

E.1 Comparison of overdamped and underdamped dynamics

As noted in the main text, the overdamped version of MULTALMC requires a small value of κ , which corresponds to a slowly varying dynamics driven by $\nabla \log \hat{\mu}_t$ to perform well in practice. However, this leads to a large number of discretisation steps since we use a small step size with respect to this slowly changing dynamics resulting in high computational costs. In contrast, the underdamped version achieves better performance without the need for such small step sizes, thanks to the faster convergence properties of underdamped dynamics [Eberle and Lörler, 2024]. To further analyse this, we compare in Table 10 the performance and number of energy evaluations of the overdamped and underdamped versions of MULTALMC, as specified in Algorithms 2 and 3, respectively, on the 40-MoG benchmark across different dimensions. The results show that the overdamped version

requires approximately an order of magnitude more energy evaluations (and hence time steps) to achieve performance comparable to the underdamped version.

Table 10: Performance metrics on the 40-MoG benchmark across varying dimensions, averaged over 30 runs with different random seeds, along with the number of energy evaluations for the overdamped and underdamped versions of MULTALMC, as defined in Algorithms 2 and 3, respectively.

MULTALMC		d = 2	d = 5	d = 10	d = 20	d = 50
Overdamped	W_2^δ	1.12 ± 0.10	1.88 ± 0.22	4.46 ± 0.39	9.23 ± 0.67	21.33 ± 0.98
	# evaluations	1×10^6	8×10^6	2×10^7	6×10^7	1×10^8
Underdamped	W_2^δ	0.91 ± 0.04	1.64 ± 0.12	4.10 ± 0.27	8.76 ± 0.34	20.58 ± 0.71
	# evaluations	3×10^5	6×10^5	1×10^6	2×10^6	5×10^6

E.2 Analysis of the impact of the number of SROCK steps in the discretisation

The number of SROCK steps used in our algorithms—MULTALMC and MULTCDIFF—as implemented in Algorithms 2, 3, and 4, plays a critical role in balancing computational efficiency and numerical performance. Increasing the number of inner SROCK steps generally leads to better performance [Abdulle et al., 2018]; however, this comes at the cost of greater computational overhead. Ideally, we aim to identify the minimum number of steps required to maintain strong performance and computational efficiency.

In the numerical experiments presented in Section 5, we use five SROCK steps. We further analyse the sensitivity of our accelerated methods, MULTALMC (Algorithm 3) and MULTCDIFF (Algorithm 4), to the number of SROCK steps. To this end, we report the entropy regularised Wasserstein-2 distance for the 40-MoG benchmark in 50 dimensions (Table 11) when running our methods with varying numbers of SROCK steps while keeping the number of sampling steps L for the slow process fixed.

Table 11: Performance on the 40-MoG benchmark in 50 dimensions using varying numbers of SROCK steps in the discretisation of MULTALMC and MULTCDIFF, as proposed in Algorithms 3 and 4, respectively. The number of sampling steps L for the slow process is held fixed. The standard deviation is computed over 30 runs with different random seeds. The regularisation parameter for the entropy regularised Wasserstein-2 distance W_2^δ is set to $\delta = 0.05$.

SROCK steps		3	5	7	9	11
W_2^δ	MULTALMC	21.11 ± 0.74	20.58 ± 0.71	20.29 ± 0.68	20.08 ± 0.50	19.99 ± 0.45
	MULTCDIFF	20.92 ± 0.63	20.13 ± 0.59	20.01 ± 0.57	19.96 ± 0.48	19.89 ± 0.40

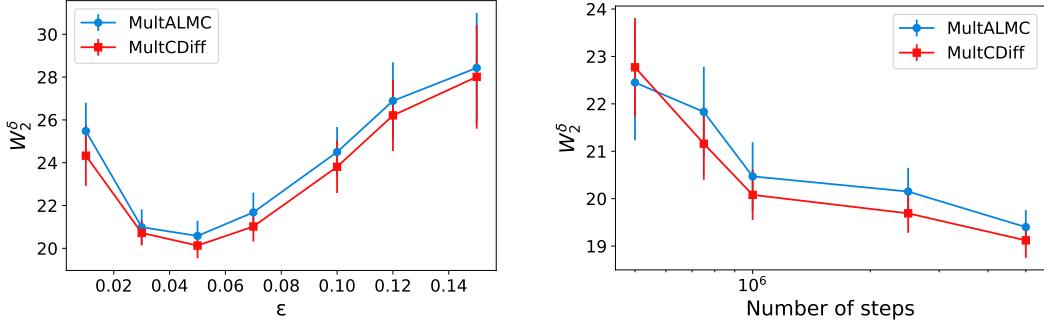
E.3 Stability across hyperparameter values

Our algorithms demonstrate robustness and stability across a broad range of hyperparameter values, as shown in Figures 3 and 4. These results suggest that while some tuning is beneficial, our methods do not require highly sensitive or exhaustive hyperparameter optimisation.

F Limitations and future work

We elaborate here on the limitations of our method discussed in Section 6.

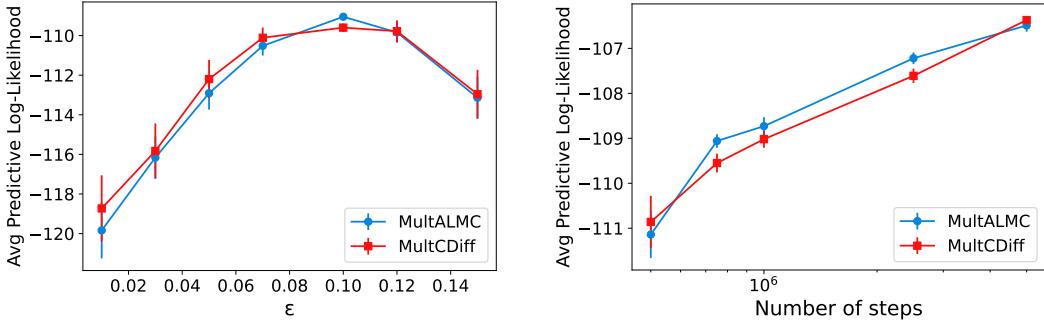
First, the theoretical guarantees presented in Section 4 rely on stringent assumptions, such as strong convexity of the potential V_π and Lipschitz continuity of the score function of the target $\nabla \log \pi$. Relaxing these assumptions, e.g., to strong convexity outside a compact region or to satisfying weak functional inequalities, could broaden the applicability of our analysis. The former would accommodate target distributions that are multimodal within a compact region and Gaussian-like in the tails, while the latter could capture heavy-tailed distributions.



(a) Performance against scale separation parameter ϵ .

(b) Performance against number of steps.

Figure 3: Ablation results on the 40-component Mixture of Gaussians benchmark in 50 dimensions. Results are averaged over 30 random seeds.



(a) Performance against scale separation parameter ϵ .

(b) Performance against number of steps.

Figure 4: Ablation results on the Sonar benchmark. Results are averaged over 30 random seeds. The performance metric is the average predictive posterior log-likelihood on a test dataset.

Additionally, the current method requires manual tuning of hyperparameters such as step size δ , scale separation parameters ϵ and friction coefficient Γ . Automating this tuning process, similar to the approach in Blessing et al. [2025], is a valuable direction for improving usability and robustness.

Further research could explore extending the controlled diffusion framework to heavy-tailed target distributions, which pose additional challenges as explained in Appendix B.2.2, or developing more efficient numerical schemes for implementing the proposed multiscale samplers.

G Broader impact

This work introduces MULTALMC and MULTCDIFF, two training-free, multiscale diffusion samplers that enable efficient, provably accurate sampling from complex distributions. These methods have the potential to significantly reduce the computational and environmental footprint of generative modelling and Bayesian inference. Potential societal benefits include faster scientific discovery, improved uncertainty quantification, and broader participation in high-impact ML research. As with all general-purpose algorithms, misuse is possible, e.g. lowering the cost of harmful synthetic content generation or accelerating harmful molecule design.