Sampling from multi-modal distributions with polynomial query complexity in fixed dimension via reverse diffusion

Anonymous Author(s)
Affiliation
Address
email

Abstract

Sampling from multi-modal distributions is challenging, even in low dimensions. We provide the first sampling algorithm for a broad class of distributions — including all Gaussian mixtures — with a query complexity that is polynomial in the parameters governing multi-modality, assuming fixed dimension. Our sampling algorithm simulates a time-reversed diffusion process, using a specific Monte Carlo estimator of the intermediate score functions. Unlike previous works, it avoids metastability, requires no prior knowledge of the mode locations, and does not rely on restrictive smoothness assumptions that exclude general Gaussian mixtures.

1 Introduction

Sampling from a distribution is a fundamental problem in statistics. Formally, given some potential $V : \mathbb{R}^d \to \mathbb{R}$ such that $\int e^{-V(x)} \mathrm{d}x < \infty$, the sampling problem consists in obtaining a sample from some distribution p such that p is ϵ-close to $\mu \propto e^{-V}$ with respect to some divergence while maintaining the complexity, a.k.a. the number of queries to V and possibly to its derivatives, as low as possible. Depending on the shape of the distribution, the typical complexity of existing sampling algorithms can significantly differ.

Log-concave distributions As in Euclidean optimization, a common assumption in the sampling literature is to assume that μ is log-concave and log-smooth or equivalently, that V is convex and smooth. Specifically, when the potential V is assumed to be α-strongly convex and to have an L-Lipschitz gradient, the popular Unajusted Langevin Algorithm (ULA) is known to achieve fast convergence [Durmus and Moulines, 2017, Dalalyan and Karagulyan, 2017]. In the more general case where μ satisfies an α^{-1}-log-Sobolev inequality (which is true in particular for α-strongly convex potentials V, but is more general), ULA achieves ϵ-error in Kullback-Leibler (KL) divergence in $\tilde{O}(L^2\alpha^{-2}d\epsilon^{-1})$ queries to ∇V [Vempala and Wibisono, 2019]. While these guarantees are appealing, strongly convex potentials are uni-modal thus do not apply to real-world settings where distributions, such as Bayesian posteriors, are often multi-modal. Moreover, truly multi-modal target distributions generally have poor functional inequalities, thus leading to weak convergence guarantees for ULA. For instance, for Gaussian mixtures, the log-Sobolev constant can grow exponentially with the distance between modes [Chen et al., 2021]. This is not simply a theoretical issue—practitioners are well aware that when dealing with multi-modal distributions, ULA-based algorithms suffer *metastability*, where they get stuck in local modes, leading to slow convergence [Deng et al., 2020].

Multi-modal distributions Designing sampling algorithms for multi-modal distributions is an active area of research. However, most existing sampling algorithms are limited in at least two of the following ways:

1. The query complexity is exponential in the parameters of the problem.
2. Guarantees are obtained under a smoothness assumption that is violated by canonical multi-modal distributions such as Gaussian mixtures.
3. The algorithm requires explicit knowledge on the target distribution, such as the localization of its modes, that are unknown generally in practice.

To illustrate these limitations, consider the previous example, where the potential is no longer α-strongly convex globally, but only outside of a centered ball of radius R in which the multi-modality is confined. First, the complexity ULA measured in KL degrades to an *exponential* function of the parameters of the problem $\tilde{O}(e^{16LR^2}L^2\alpha^{-2}d\epsilon^{-1})$ [Ma et al., 2019, Prop. 2]. Second, we show in this work that the smoothness assumption is *violated* by general Gaussian mixtures. Third, the complexity result was assuming prior knowledge of the smoothness constant which is used to choose a "small enough" step size.

Generally speaking, finding a sampling algorithm that addresses even the first limitation is not possible. Recent results provide lower-bounds on the complexity of sampling from a multi-modal distribution: they are exponential in the dimension, as shown in Lee et al. [2018, Th K.1], Chak [2024, Th 3.3] and He and Zhang [2025, Th. 3]. However, when the dimensionality is fixed, the question of whether the three limitations can be addressed remains. *In fixed dimension, can we sample from a broad class of multi-modal distributions with a polynomial number of queries in the parameters of the problem, (e.g. L, R, α) and without prior knowledge on these parameters*?

Our contributions Our work answers this question positively. We provide a sampling algorithm that addresses the three limitations outlined above. First, we show that our algorithm has *polynomial* complexity in all the problem parameters but the dimension, thus enabling to efficiently sample from highly multi-modal distributions in fixed dimension. Second, we show that unlike most existing results, our guarantees hold under assumptions that cover *general* Gaussian mixtures. Third, our algorithm yield guarantees *without* prior knowledge on the parameters of distribution.

Of independent interest, we highlight that our sampling algorithm is based on simulating a time-reversed diffusion. Such processes are popular for sampling from distributions when access to samples is available, such as in generative modeling [Ho et al., 2020, Song et al., 2021]. Recently, there has been growing interest in adapting these methods to settings where only an unnormalized density of the target distribution is known [Huang et al., 2024a,b, He et al., 2024]. While several algorithms have been proposed, complexity guarantees remain scarce. Our work contributes to this line of research.

This paper is structured as follows. First, we survey related work on multi-modal sampling in Section 2 and present our main result in Section 3. Then, we detail our sampling algorithm in Section 4 and provide a sketch of proof in Section 5.

2 Related work

Before detailing our sampling algorithm in Section 4, we review the main alternatives to ULA when sampling from a multi-modal distribution and the guarantees they offer. We show that existing approaches suffer from at least two of three drawbacks outlined in the introduction: exponential query complexity, a restrictive smoothness assumption that excludes general Gaussian mixtures (as shown later in Sec. 3.2), and unavailable prior knowledge of the distribution.

Proposal-based algorithms At a high level, proposal-based algorithms use a proposal distribution that is easy to sample from and that can either be directly used as a proxy for the target, or, alternatively, whose samples will be rejected (resp. re-weighted) to obtain approximate samples from the target; the corresponding algorithm for the latter case is the well-known rejection sampling (resp. importance sampling) scheme. Guarantees for these methods typically assume that the target distribution is log-smooth. Furthermore, they must use carefully designed proposals that require prior knowledge of the target distribution that is often unavailable in practice.

For instance, when the target is L-log-smooth and with second moment m_2, one can design and sample from a proxy distribution that is ϵ-close in TV to the target using $\tilde{O}((Lm_2\epsilon^{-1})^{O(d)})$ queries to the potential function V and its gradient ∇V [He and Zhang, 2025]. While this bound is indeed polynomial in fixed dimension, the design of the proxy requires an ϵ-approximation of the global minimum of V which can only be achieved if m_2, therefore the location of the mass, is explicitly available; this is rarely the case in practice.

Similarly, if we further assume that the target is α-strongly log-concave outside of a ball, one can achieve an ϵ-precise approximation of the target distribution in χ^2 divergence with polynomial number of queries to ∇V when d is fixed via an importance sampling scheme [Chak, 2024, Th. 2.3]. In this case, the proposal is such that it coincides with the target when the potential of the target is above some cutoff and is flat elsewhere. While there exists a cutoff value ensuring the log-concavity of the proposal, thus allowing its efficient sampling, this value depends on the unknown constants L, R and α [Chak, 2024, Prop. 5.1].

Tempering-based algorithms Instead of directly sampling from $\mu \propto e^{-V}$, these algorithms start by sampling from a flattened version of μ given by $\mu^\beta \propto e^{-\beta V}$ for small β, and gradually increase β to 1, a strategy sometimes referred to also as annealing. When μ is assumed to be a finite mixture of the same shifted α-strongly log-concave and L-log-smooth distribution, Lee et al. [2018] proved that for a well-chosen (stochastic) sequence of flattened distributions μ^{β_i}, sampling up to precision ϵ in KL can be achieved in $poly(L, \alpha^{-1}, w_{min}^{-1}, d, \epsilon^{-1}, R)$ queries to ∇V, where R is the location of the furthest mode and $w_{\min}$ the minimum weight in the mixture. However, this setting is quite restrictive: for Gaussian mixtures for instance, it only handles the case where all the covariance matrices are identical. Furthermore the algorithm requires explicit knowledge on R which is unavailable.

Föllmer Flows Instead of directly sampling from the target distribution $\mu \propto e^{-V}$, Föllmer Flows start from a simple (e.g. Gaussian or Dirac mass) distribution that is progressively interpolated to the target using a Schrödinger bridge. When the initial distribution is a Dirac mass at 0, the time marginals of this bridge solve a closed-from SDE [Wang et al., 2021, Theorem 3] which can thus be discretized to generate samples from the target. In the context of Bayesian inference, where one only has access to the unnormalized density, Vargas et al. [2022] estimated the drift of the resulting SDE via neural methods; in particular no guarantees on the sampling quality are provided. In Huang et al. [2021], Jiao et al. [2021], Ruzayqat et al. [2023], Ding et al. [2023], the drift is estimated via a Monte-Carlo method. While in appearance, these works provide strong polynomial guarantees, a closer look shows that these guarantees only hold if the function $f(x) = e^{-V(x)+\|x\|^2/2}$ is Lipschitz, smooth and bounded from below; we show in Appendix A.1 that this assumption is quite restrictive.

Diffusion-based methods Over the past few years, diffusion-based algorithms, and especially reverse diffusion, that we shall review in details in Sec. 3, have emerged as solid candidates for multi-modal sampling. In essence, they allow to transfer the sampling problem into the problem of estimating the scores of intermediate distributions that are given by the convolution of the initial distribution with increasing levels of Gaussian noise. Under an ϵ-oracle of these intermediate scores, it has been shown that diffusion-based methods could yield an ϵ-approximate sample of the target in poly(ϵ^{-1}) time under milder and milder assumptions [Chen et al., 2023b,a, Benton et al., 2024, Li et al., 2024, Conforti et al., 2025, Gentiloni-Silveri and Ocello, 2025, Cordero-Encinar et al., 2025]. This framework has been applied with tremendous empirical success in generative modeling, where numerous samples of the target are already available and one seeks to produce new samples from the target, for several years now [Song and Ermon, 2019, Ho et al., 2020, Bortoli et al., 2021]. However, it was only quite recently that this framework has been applied in a Bayesian context, where only an unnormalized density of the target, instead of samples, is available [Huang et al., 2024a,b, He et al., 2024, Grenioux et al., 2024, Akhound-Sadegh et al., 2024]. In a work closely related to ours, Huang et al. [2024b] showed that if the intermediate scores remain L-log-smooth for any noise level, which in particular implies that the target itself is L-log-smooth, their algorithm could reach a complexity of $O(e^{L^3 \log^3((Ld+m_2)/\epsilon)})$ with m_2 the second order moment of the target. In He et al. [2024], the smoothness assumption is relaxed with a sub-quadratic growth assumption $V(x) - V(x^*) \leq L\|x - x^*\|^2$, where x^* is any global minimizer of V, yet at the price of an oracle access to $V(x^*)$; under these assumptions, the authors manage to obtain a complexity that is at best $O(L^{d/2}\epsilon^{-d}e^{L\|x^*\|^2+\|x_N\|^2})$ where x_N is the final output sample. Under the reasonable (and desirable) assumption that $\mathbb{E}[\|x_N^2\|] = m_2$, Jensen's inequality yield an overall $O(L^{d/2}\epsilon^{-d}e^{L\|x^*\|^2+m_2})$ complexity. In particular, both these works suffer at least two of the limitations mentioned in the introduction, making them ill-suited candidates for multi-modal sampling.

3 Presentation of the main result and application to Gaussian Mixtures

3.1 Main result

As in the works mentioned above, we rely on the recent advances on reverse diffusion and focus on the task of estimating the intermediate scores. Using an estimator that is described in Sec. 4, we recover polynomial sampling guarantees for densities verifying the assumptions described hereafter.

Assumption 1 (Semi-log-convexity) *We assume that $\mu \propto e^{-V}$ is such that $\log(\mu)$ is C^2 and verifies $\nabla^2 \log(\mu) \succeq -\beta I_d$ or equivalently $\nabla^2 V \preceq \beta I_d$ for some $\beta \geq 0$.*

This assumption shall be referred to as *semi-log-convexity*, by analogy with the functional analysis literature [Mikulincer and Shenfeld, 2023, Theorem 3]. Note that is has also been referred to as 1-sided Lipschitzness for ∇V [Gentiloni-Silveri and Ocello, 2025]. It relaxes the classical log-smoothness assumption which implies the additional lower bound $\nabla^2 V \succeq -\beta I_d$. In particular, unlike the latter, a mixture of semi-log-convex densities remains semi-log-convex Marshall et al. [1979, Chap. 16.B]; we provide a quantitative version of this statement in Sec. 3.2 in the case of Gaussian mixtures.

Assumption 2 (Dissipativity) *We assume that $\mu \propto \exp^{-V}$ is such that there exists $a > 0, b \geq 0$ for which its potential satisfies $\langle \nabla V(x), x \rangle \geq a\|x\|^2 - b$ for all $x \in \mathbb{R}^d$.*

This assumption is referred to as *dissipativity* as common in the sampling and optimization literature [Raginsky et al., 2017, Zhang et al., 2017, Erdogdu and Hosseinzadeh, 2021]. Note that this assumption relaxes strong convexity outside of a ball which can be equivalently re-written as $\langle \nabla V(x) - \nabla V(y), x - y \rangle \geq a\|x - y\|^2 - b$ for all pairs $(x, y) \in \mathbb{R}^d$ (also referred to as *strong dissipativity* [Erdogdu et al., 2022]). Unlike the latter, it does not imply a local behavior on ∇V and *a fortiori* a control on $\nabla^2 V$ but only allows for an asymptotical control on ∇V. In particular, unlike strong-convexity outside of a ball, dissipativity is stable for mixtures. We provide a quantitative version of this statement in Sec. 3.2.

Theorem 1 *[Main result, informal] Suppose that Assumption 1 and 2 hold. Then, for all $\epsilon > 0$, there exists a stochastic algorithm whose parameters only depend on ϵ (and not on the parameters on the problem), that outputs a sample $X \sim \hat{p}$ such that $\mathbb{E}[\mathrm{KL}(\mu, \hat{p})] \lesssim \epsilon \beta^{d+3}(b + d)/a^2$ in $O(\mathrm{poly}(\epsilon^{-d}))$ queries to V with $\lesssim$ hiding a universal constant as well as log quantities in $d, \epsilon^{-1}, \beta, a, b$.*

In particular, when d is fixed, our algorithm can output a sample from a distribution that is ϵ-close to μ in expected KL in $\mathrm{poly}(b/a, \beta, \epsilon^{-1})$ running time.

Our algorithm addresses the three limitations outlined in the sections above. When the dimension d is fixed, we obtain a *polynomial* query complexity. This guarantee does *not* assume log-smoothness and applies to general Gaussian mixtures, as will be shown in the next subsection. Moreover, this guarantee does *not* require running the algorithm with any explicit knowledge of the target distribution's constants a, b, β.

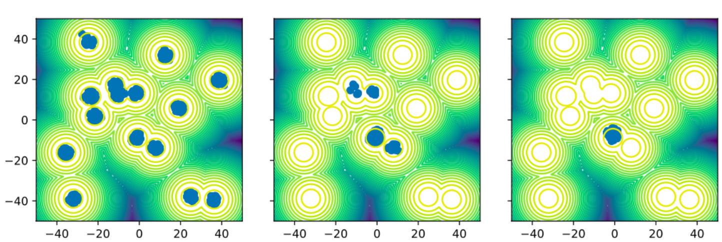

Figure 1: From left to right: our algorithm vs. ULA vs. Huang et al. [2024a].

Overall comparison We summarize in Table 1 how our algorithm compares to previous approaches in terms of assumptions, oracles required to run the algorithm and the resulting complexity. Along with the work of He and Zhang [2025], our algorithm is the only one that is polynomial in the parameters of the distribution when d is fixed. Furthermore, while our dissipativity assumption is stronger than bounded second moment, we relax the log-smoothness assumption by semi-log-convexity which notably covers general Gaussian mixtures. Finally, as mentioned above, because their algorithm requires an ϵ-approximation of the global minimum of V, they require an explicit upper-bound on the second order moment of μ which may not be available in practice.

Algorithm	Assumptions	Oracle	Complexity
ULA [Ma et al., 2019]	L-log-smoothness, α-strong log concavity outside B_R	∇V, L	$O(e^{16LR^2}(L/\alpha)^2 d\epsilon^{-1})$
Proposal-based [He and Zhang, 2025]	L-log-smoothness, $m_2 \leq M$	V, ∇V, M, L	$O((LM\epsilon^{-1})^{O(d)})$
RD + ULA [Huang et al., 2024b]	L-log-smoothness along the forward flow, finite m_2	∇V	$e^{O(L^3 \log^3((Ld+m_2/\epsilon))}$
RD + Rejection Sampling [He et al., 2024]	L-sub quadratic growth, finite m_2	V, V^*	$O(L^{d/2}\epsilon^{-d}e^{L\|x^*\|^2+m_2})$
RD + Self-normalized IS (ours)	β-semi-log-convexity, (a,b)-dissipativity	V	$O(d\beta^{d+2}(b+d)/a^2\epsilon^{-(2d+8)})$

Table 1: Complexity of sampling algorithms. We denote by x^* a global minimizer of V and by $V^* = V(x^*)$ the global minimum of V. In "RD + ULA", RD refers to a Reverse Diffusion algorithm and ULA refers to how the intermediate scores are estimated.

We illustrate our results numerically in Figure 1, where we compare our algorithm against Unadjusted Langevin Algorithm (ULA) and to the reverse diffusion algorithm of Huang et al. [2024a]. We also implemented the zeroth-order method of He et al. [2024], but it failed to converge. The corresponding task is standard in the literature: sampling from a mixture of 16 equally weighted Gaussians with unit variance and centers uniformly distributed in $[-40, 40]^2$ [Midgley et al., 2023].

ULA was initialized from $\mathcal{N}(0, I_2)$ and run for 5×10^4 steps. All three methods used the same discretization step size $h = 0.01$. Our reverse diffusion algorithm and that of Huang et al. [2024a] were both run with 500 reverse diffusion steps and used 100 samples to estimate the scores of the intermediate diffusion sequence. For Huang et al. [2024a], the samples used to estimate the score are drawn from an auxiliary distribution that are generated using 100 steps of ULA. While ULA and our algorithm took approximately one minute to run on a computer locally, the method of Huang et al. [2024a] required over an hour. We observe that unlike the two others, our algorithm successfully recovers all the modes.

3.2 Application to Gaussian mixtures

We now apply our results to derive provable sampling guarantees for general Gaussian mixtures. Apart from the very recent exception of Lytras and Mertikopoulos [2025] that we discuss below, note that despite their wide popularity, no sampling guarantees were yet derived for general Gaussian mixtures . Indeed, unless the covariance matrices verify some specific assumptions such as identical covariance matrices [Cordero-Encinar et al., 2025, Lemma B.1], general Gaussian mixtures do not satisfy the log-smoothness assumption, thus they do not fit the framework of many previously discussed works. In fact, their gradient may not even be Hölder continuous: consider the simple counter-example of a two-dimensional mixture $\mu = 0.5\mathcal{N}(0, \Sigma_1) + 0.5\mathcal{N}(0, \Sigma_2)$ with covariances $\Sigma_1 = diag(1, 0.5)$ and $\Sigma_2 = diag(0.5, 1)$. On the diagonal $x = y$, the score is $\nabla \log(\mu)(x, x) = -3/2(x, x)$, while near the diagonal, right above it for instance, the score behaves asymptotically as $\nabla \log(\mu) \sim_{x\to+\infty} -(2x, x)$; we provide a rigorous analysis in Appendix A.2. Fortunately though, Gaussian mixtures do verify Assumptions 1-2, as shown in the following proposition.

Proposition 2 *Let $\mu = \sum_{i=1}^p w_i \mathcal{N}(\mu_i, \Sigma_i)$ and denote $\lambda_{min} > 0$ (resp. λ_{max}) the minimum (resp. maximum) eigenvalue of the covariance matrices Σ_i. It holds that $-\nabla^2 \log(\mu) \preceq I_d/\lambda_{min}$ and that for all $x \in \mathbb{R}^d$, $\langle -\nabla \log(\mu)(x), x\rangle \geq \|x\|^2/(2\lambda_{max}) - \lambda_{max} \max_i(\|\mu_i\|/\lambda_{min})^2$.*

The proof is deferred to Appendix A.3. Combined with Theorem 1, this shows that we can sample any Gaussian mixture with average precision ϵ in KL in $O(\mathrm{poly}(\kappa, R, d, \lambda_{min}^{-d}, \epsilon^{-d}))$ queries to V where $R = \max_i \|\mu_i\|$ and $\kappa = \lambda_{max}/\lambda_{min}$. While there exists a relatively new literature seeking to relax the log-smoothness assumption [Erdogdu and Hosseinzadeh, 2021, Mou et al., 2022, Nguyen et al., 2023, Lehec, 2023], the only reference we know of that applies to general Gaussian mixtures is the work of Lytras and Mertikopoulos [2025], who used a regularized version of Langevin. We show in Appendix A.4 that for an appropriate regularization, their algorithm provides a sample with precision ϵ in $O(\mathrm{poly}(e^{(\kappa R)^2}, \lambda_{min}^{-d}, d, \epsilon^{-1}))$ queries to V, hence it is exponential and not polynomial in the problem parameters. In particular, in the highly multi-modal case $R \gg 1$ for instance, our algorithm is more competitive.

4 Our sampling algorithm

In this section, we introduce the reverse diffusion framework and how it allows to reduce the sampling problem to the problem of estimating the scores along the forward Ornstein-Uhlenbeck (OU) process.

4.1 Reverse diffusion: from sampling to score estimation

Reverse diffusion methods emerged as an alternative to Langevin-based samplers in order to overcome metastability and were first introduced to the ML community in Song and Ermon [2019]. They rely on the so-called forward process

$$\begin{cases} \mathrm{d}X_t = -X_t \mathrm{d}t + \sqrt{2}\mathrm{d}B_t \\ X_0 \sim \mu\,, \end{cases} \tag{1}$$

which corresponds to the standard OU process initialized at μ, that is a specific case of a Langevin diffusion targeting a standard Gaussian, that we will denote π. Note that since the target of this process, the standard Gaussian, is 1-strongly log-concave, the resulting process converges exponentially fast to the equilibrium. In order to sample from μ, reverse diffusion algorithms rely on the semi-discretized *backward process*: given a horizon T that we discretize as $0 = t_0 \leq t_1 \leq \cdots t_{N-1} \leq t_N = T$, the latter writes

$$\mathrm{d}Y_t = Y_t \mathrm{d}t + 2\nabla \log\big(p_{t_{N-k}}\big)(Y_t)\mathrm{d}t + \sqrt{2}\mathrm{d}B_t\,, \quad t \in]T - t_{N-k}, T - t_{N-(k+1)}]\,, \tag{2}$$

with $Y_0 \sim p_T$ and where p_t is the distribution of the forward process Eq. 1 at time t. Note that this reverse process cannot be readily implemented for two reasons: first, it requires the knowledge of the intermediate scores $\nabla \log(p_{t_k})$ which are not available in closed form. Second, it requires sampling from the distribution p_T. Nevertheless, if one can access a proxy s_{t_k} of the scores $\nabla \log(p_{t_k})$, and considering T large enough so that $p_T \approx \pi$, we can implement instead

$$\mathrm{d}Y_t = Y_t \mathrm{d}t + 2 s_{t_k}(Y_k)\mathrm{d}t + \sqrt{2}\mathrm{d}B_t\,, \quad t \in]T - t_{N-k}, T - t_{N-(k+1)}]\,, \tag{3}$$

with $Y_0 \sim \pi$ and where all iterations can be solved in closed form. Because the forward process Eq. 1 converges exponentially fast, we can expect the initialization error $Y_0 \sim \pi$ instead of $Y_0 \sim p_T$ to be small after a short time T. Furthermore, if the proxies s_{t_k} are sufficiently accurate, one can expect that the process output by the approximate scheme Eq. 3 has a distribution that is close to the target μ. Over the past three years, several works provided quantitative bounds of the error induced by the discretization, the use of an approximate score and the initialization error with respect to different divergences and under various assumptions [Bortoli et al., 2021, Lee et al., 2022, Chen et al., 2023b,a, Conforti et al., 2025]. Yet, we shall rely exclusively on the following theorem as it is the most suited to our framework.

Theorem 3 (Conforti et al. [2025]) *Assume that $\mu \propto e^{-V}$ has finite Fisher-information w.r.t. π the standard gaussian density in $\mathbb{R}^d$:*

$$\mathcal{I}(\mu, \pi) := \int \|x - \nabla V(x)\|^2 \mathrm{d}\mu(x) < +\infty\,.$$

Then, for the constant step-size discretization $t_k = kT/N$, denoting p the distribution of the sample Y_T output by Eq. 3, it holds that

$$\mathrm{KL}(\mu, p) \lesssim (d + m_2)e^{-T} + \frac{1}{N}\sum_{k=1}^{N} \|\nabla \log(p_{t_k}) - s_{t_k}\|^2_{L^2(p_{t_k})} + \frac{T}{N}\mathcal{I}(\mu, \pi)\,,$$

where m_2 is the second order moment of μ and where $\lesssim$ hides a universal constant.

The previous theorem shows that under mild assumptions that notably allow for multi-modality, the problem of sampling from μ can be transferred into a score approximation problem along the forward process. In the next subsection, we build an estimator for these intermediate scores that is tractable given the knowledge of the unnormalized density $\mu \propto e^{-V}$, and show how it differs from previous works.

4.2 Construction of our estimator

The key observation to derive an estimator of the intermediate scores is that the forward process Eq. 1 is nearly available in closed form. However, this closed form may be interpreted in different manners.

Previous estimators of the scores Consider Eq. 1 integrates to $X_t = e^{-t}X_0 + B_{1-e^{-2t}}$, where B_s is the standard d-dimensional Brownian motion evaluated at time s. In particular, conditionally on X_0, X_t has a normal distribution $X_t|X_0 \sim \mathcal{N}(e^{-t}X_0, (1-e^{-2t})I_d)$. By Bayes rule, the distribution p_t^V of the forward process X_t initialized at $\mu \propto e^{-V}$ thus integrates up to a normalizing constant to

$$p_t^V(z) \propto \int e^{-\frac{\|e^{-t}x - z\|^2}{2(1-e^{-2t})}} e^{-V(x)} \mathrm{d}x\,.$$

One thus recovers that the score of the forward process re-writes as

$$\nabla \log\big(p_t^V\big)(z) = \mathbb{E}_{Y\sim q_{t,z}}\left[\frac{z - e^{-t}Y}{1-e^{-2t}}\right], \tag{4}$$

which requires samples from an auxiliary distribution $q_{t,z}(x) \propto e^{-\frac{\|e^{-t}x-z\|^2}{2(1-e^{-2t})}} e^{-V(x)}$. As noted in Huang et al. [2024a,b], He et al. [2024], Grenioux et al. [2024], if samples $y_i \sim q_{t,z}$ are available, one can recover an empirical estimator of the score given by $\frac{-1}{n(1-e^{-2t})}\sum_{i=1}^n y_i$ ready to be plugged in the reverse diffusion process Eq. 3. However, by doing so, the problem of sampling from e^{-V} is simply transferred to the problem of sampling from $q_{t,z}$ which progressively becomes as hard as t grows; in particular, the standard algorithms described in the sections above will still fail to accurately sample from these intermediate distributions. Since Huang et al. [2024a] use ULA and He et al. [2024] use rejection sampling to generate these intermediate samples, unsurprisingly, they recover poor guarantees.

Our self-normalized estimator of the scores Instead, we observe that after applying the change of variable $y_t = z - e^{-t}x$ both on the numerator and on the denominator, we obtain a score formula which re-writes as a *ratio* of expectations under standard Gaussians

$$\nabla \log\big(p_t^V\big)(z) = \frac{-1}{1-e^{-2t}}\frac{\mathbb{E}[Y_t e^{-V(e^t(z-Y_t))}]}{\mathbb{E}[e^{-V(e^t(z-Y_t))}]}\,,$$

with $Y_t \sim \mathcal{N}(0, (1-e^{-2t})I_d)$. Unlike the previous identity, which involves an expectation under $q_{t,z}$, the expectations above are computationally cheap as we can easily sample from Y_t. While conventional statistical wisdom would suggest both the numerator and the denominator to be estimated with independent samples, we voluntary choose to correlate them and implement instead

$$\hat{s}_{t,n}(z) := \frac{-1}{1-e^{-2t}}\frac{\sum_{i=1}^n y_i e^{-V(e^t(z-y_i))}}{\sum_{i=1}^n e^{-V(e^t(z-y_i))}}\,, \tag{5}$$

where the y_i are independent Gaussians such that $y_i \sim \mathcal{N}(0, (1-e^{1-2t})I_d)$; we shall refer to this estimator as *self-normalized* as common in the sampling literature [Agapiou et al., 2017]. Depending on the locations of the modes of V we expect that this estimator behaves reasonably well in some regions of the space and poorly in others. Yet, the key property of self-normalized estimators is that they remain nearly bounded: in our case, it holds uniformly in z, t that

$$\|\hat{s}_{t,n}(z)\| \leq \frac{\max_i \|y_i\|}{1-e^{-2t}} \sim \sqrt{d\log(n)}\,.$$

This boundedness will allow us to derive a non-asymptotical control on the quadratic error that we present in the next section.

Note that the estimator Eq. 5 was already considered in [Saremi et al., 2024] yet no estimation guarantees were provided. A similar estimator was also already considered in the context of Föllmer Flows [Jiao et al., 2021, Ruzayqat et al., 2023, Ding et al., 2023]. Yet as previously discussed, they were derived under very restrictive assumptions. As a consequence, these works avoid most of the technical difficulties we were confronted with.

5 Sketch of proof

The proof is decomposed in three steps: (i) we derive a non-asymptotic error of the estimator presented in Eq. 5 (ii) we show that under Assumptions 1-2, the integrated error of this estimator (that appears in the bound of Theorem 3) can be fully controlled by the zeroth and second order moments of the ratio $\Phi_t = p_t^{2V}/p_t^V$ with p_t^{2V} (resp. p_t^V) the density of the forward process defined in Eq. 1 initialized at μ^2 (resp.μ) (iii) we provide a quantitative bound on these moments as well as other relevant quantities and we conclude with Theorem 3.

Proposition 4 (Non-asymptotic bound on the quadratic error) *For all $z \in \mathbb{R}^d$, n and $t > 0$, denoting p_t^{2V} (resp. p_t^V) the density of the forward process defined in Eq. 1 initialized at $\mu^2 \propto e^{-2V}$ (resp.$\mu \propto e^{-V}$), Z_{2V} (resp. Z_V) the normalizing constant of e^{-2V} (resp. e^{-V}) and π the density of the standard Gaussian, it holds that*

$$\mathbb{E}\left[\|\hat{s}_{t,n}(z) - \nabla\log\big(p_t^V\big)(z)\|^2\right] \leq \frac{32e^2(d+\log(n))}{n(1-e^{-2t})}\left(1+\theta_t(z)\frac{p_t^{2V}(z)Z_{2V}e^{td}}{(p_t^V(z)Z_V)^2}\right),$$

with $\theta_t(z) = \Delta\log\big(p_t^{2V}\big)(z) - \frac{\Delta\log(\pi)(z)}{1-e^{-2t}} + \|\nabla\log\big(p_t^V\big)(z)\|^2 + \|\nabla\log\big(p_t^{2V}\big)(z)\|^2 + 1.$

The complete proof is left to Appendix B yet we briefly sketch the main arguments. We split the expectation on the event A where the empirical denominator $\hat{D}$ of Eq. 5 (resp. the empirical numerator $\hat{N}$) is not too small (resp. not too large) with respect to its expectation D (resp. $\|N\|$) and on the complementary $\bar{A}$. Across A, we use a Taylor expansion of order 2 to make the variance of the numerator and of the denominator appear and we both compute them explicitly. Conversely, the estimator remains almost bounded on $\bar{A}$. We use Chebyshev's inequality to upper-bound $\mathbb{P}(\bar{A})$ and make the variances of the numerator and of the denominator appear again, which concludes the proof.

Remark 5 *Generic bounds on self-normalized estimators were derived in Agapiou et al. [2017, Theorem 2.3]. Yet, if used in our context, they would involve an additional $1/(p_t^V)^4$ term which would cause the integrated error $\int\mathbb{E}\left[\|\hat{s}_{t,n}(z) - \nabla\log\big(p_t^V\big)(z)\|^2\right]\mathrm{d}p_t^V(z)$ to diverge.*

Now, we need to control the Laplacian of the forward processes as well as their gradient. As mentioned in the previous section, the intermediate scores can be re-written as

$$\nabla\log\big(p_t^V\big)(z) = \frac{z - e^{-t}\mathbb{E}_{q_{t,z}}[Y]}{1-e^{-2t}}\,,$$

with $q_{t,z}(x) \propto e^{-V(x)}e^{-\frac{\|e^{-t}x-z\|^2}{2(1-e^{-2t})}}$. Now we note that if $\mu \propto e^{-V}$ is dissipative, so is $q_{t,z}$; in particular we can quantitatively bound its second order moment and *a fortiori*, upper-bound $\|\nabla\log\big(p_t^V\big)(z)\|^2$.

Proposition 6 (Regularity bounds on the forward) *Suppose that Assumption 2 holds. Then, for all $t > 0$, it holds that*

$$\begin{cases} \|\nabla\log\big(p_t^V\big)(z)\|^2 \leq \frac{\|z\|^2}{(1-e^{-2t})^2}\left(2+\frac{e^{-2t}}{a(1-e^{-2t})}\right) + \frac{2e^{-2t}(2b+d)}{a(1-e^{-2t})^2}\,, \\ \Delta\log\big(p_t^V\big) \leq \frac{e^{-2t}}{(1-e^{-2t})^2}\left(\frac{\|z\|^2}{2a(1-e^{-2t})} + \frac{2b+d}{a}\right). \end{cases}$$

The complete proof is left in Appendix C.2. This result implies that the $\theta(z)$ term defined in Proposition 4 is at most of order of order $\sim (1 + \|z\|^2)\frac{p_t^{2V}(z)}{p_t^V(z)^2}$ w.r.t. z. In particular, the average integrated error $\mathbb{E}[\|\hat{s}_{t,n}(z) - \nabla \log(p_t^V)\|^2_{L^2(p_t^V)}]$ can be upper-bounded with respect to the zeroth and the second order moments of the ratio $\Phi_t = p_t^{2V}/p_t^V$. In the next proposition, we show that these moments are bounded under the semi-log-convexity assumption. By slight abuse of notations, we shall denote $m_i(\Phi_t) = \int \|x\|^i \Phi_t(z)\mathrm{d}z$ the i-th moment of Φ_t.

Lemma 7 (Bound on the moments of the ratio) *Assume that $\mu \propto e^{-V}$ has finite second moment m_2 and that Assumption 1 holds. Then,*

$$\begin{cases} m_0(\Phi_t) \leq \frac{(Z_V)^2}{Z_{2V}} e^{td}(\beta(1-e^{-2t}) + e^{-2t})^d \,, \\ m_2(\Phi_t) \leq \frac{2e^{t(d+2)}(Z_V)^2}{Z_{2V}}(\beta(1-e^{-2t}) + e^{-2t})^{d+2}\left[m_2 + d(\beta+1) + 2d\log\left(\frac{\beta\mu(0)^{-2/d}}{2\pi}\right)\right] \,, \end{cases}$$

where Z_{2V} (resp. Z_V) is the normalization constant of e^{-2V} (resp. e^{-V}).

The proof of this lemma is deferred in Appendix D. It relies on a key result in Mikulincer and Shenfeld [2023, Lemma 5] where it is shown that for β-semi-log convex distributions, the following bound holds:

$$\nabla^2 \log(\pi) - \nabla^2 \log(p_t^V) \preceq \frac{(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1} I_d \,.$$

In particular, the right-hand-side remains bounded w.r.t. β whenever $t > 0$ which allows to avoid an exponential dependence in β in our final bounds. It suffices now to control $\mathcal{I}(\mu,\pi), m_2$ and $\mu(0)$ in order to apply Theorem 3 and conclude.

Lemma 8 *Assume that $\mu \propto e^{-V}$ is such that $\nabla^2 V \preceq \beta I_d$ and such that $\langle \nabla V(x), x\rangle \geq a\|x\|^2 - b$ for some $a > 0, b, \beta \geq 0$. It holds that*

$$\begin{cases} m_2 \leq (b+2d)/a \,, \\ \mathcal{I}(\mu,\pi) \leq 2(b+2d)/a + 2\beta d \,, \\ \log(\mu(0)^{-2/d}) \leq 4\beta b/(da) + 2\pi + \log(2/a) \,. \end{cases}$$

The proof is left in Appendix C.1. We can now state our main result.

Theorem 9 *Under Assumptions 1-2, if we run algorithm Eq. 3 with $T = \log(1/\epsilon)$, $N = 1/\epsilon$, $t_k = kT/N$ and with the stochastic score estimators $\hat{s}_{n_k,t_k}$ defined in Eq. 5 with $n_k = d^2\max(\epsilon^{-2(d+1)+1}, \epsilon^{-5})$, then, denoting $\hat{p}$ the stochastic distribution of the output Y_N, it holds that*

$$\mathbb{E}[\mathrm{KL}(\mu,\hat{p})] \lesssim \epsilon\beta^{d+3}(b+d)/a^2 \,,$$

where $\lesssim$ hides a universal constant as well as log factors with respect to $d, \epsilon^{-1}, a, b, \beta$. In particular, the error above can is achieved in $\sum_{k=1}^N n_k = d^2\max(\epsilon^{-2(d+2)}, \epsilon^{-6})$ queries to V.

The proof is deferred in Appendix E and is mainly an application of the results collected above.

6 Conclusion

In this article, we successfully apply the reduction from sampling to intermediate score estimation initiated over the past three years by Chen et al. [2023b,a], Conforti et al. [2025], Benton et al. [2024] to the problem of low dimensional multi-modal sampling. Using our self-normalized estimator of the scores, our results provide *polynomial* query complexity guarantees, apply to *general* Gaussian mixtures, and do *not* require prior knowledge of the target distribution's constants.

References

Sergios Agapiou, Omiros Papaspiliopoulos, Daniel Sanz-Alonso, and Andrew M Stuart. Importance sampling: Intrinsic dimension and computational cost. *Statistical Science*, 2017.

Tara Akhound-Sadegh, Jarrid Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, et al. Iterated denoising energy matching for sampling from Boltzmann densities. In *ICML*, 2024.

Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and geometry of Markov diffusion operators*. Springer, 2014.

Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Nearly d-Linear convergence bounds for diffusion models via stochastic localization. In *ICLR*, 2024.

Valentin De Bortoli, James Thornton, Jeremy Heng, and A. Doucet. Diffusion Schrödinger Bridge with applications to score-based generative modeling. In *NeurIPS*, 2021.

Martin Chak. On theoretical guarantees and a blessing of dimensionality for nonconvex sampling. *arXiv preprint arXiv:2411.07776*, 2024.

Hong-Bin Chen, Sinho Chewi, and Jonathan Niles-Weed. Dimension-free log-Sobolev inequalities for mixture distributions. *Journal of Functional Analysis*, 2021.

Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *ICML*, 2023a.

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 *ICLR*, 2023b.

Sinho Chewi and Austin J Stromme. The ballistic limit of the log-sobolev constant equals the polyak-$\{\backslash\mathrm{L}\}$ ojasiewicz constant. *arXiv preprint arXiv:2411.11415*, 2024.

Giovanni Conforti, Alain Durmus, and Marta Gentiloni Silveri. Score diffusion models without early stopping: finite Fisher information is all you need. *SIAM Journal on Mathematics of Data Science (SIMODS)*, 2025.

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.

SAM CRAIG. The Poincaré inequality on convex domains. 2021.

Arnak S Dalalyan and Avetik G Karagulyan. User-friendly guarantees for the Langevin Monte Carlo with inaccurate gradient. *arXiv preprint arXiv:1710.00095*, 2017.

Wei Deng, Guang Lin, and Faming Liang. A contour stochastic gradient Langevin dynamics algorithm for simulations of multi-modal distributions. In *NeurIPS*, 2020.

Zhao Ding, Yuling Jiao, Xiliang Lu, Zhijian Yang, and Cheng Yuan. Sampling via Föllmer Flow. *arXiv preprint arXiv:2311.03660*, 2023.

Alain Durmus and Eric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 2017.

Murat A Erdogdu and Rasa Hosseinzadeh. On the convergence of Langevin monte carlo: The interplay between tail growth and smoothness. In *COLT*, 2021.

Murat A. Erdogdu, Rasa Hosseinzadeh, and Shunshi Zhang. Convergence of Langevin Monte Carlo in chi-squared and Rényi divergence. In *AISTATS*, 2022.

Marta Gentiloni-Silveri and Antonio Ocello. Beyond log-concavity and score regularity: Improved convergence bounds for score-based generative models in W2-distance. *arXiv preprint arXiv:2501.02298*, 2025.

Louis Grenioux, Maxence Noble, Marylou Gabrié, and Alain Oliviero Durmus. Stochastic localization via iterative posterior sampling. In *ICML*, 2024.

Ye He, Kevin Rojas, and Molei Tao. Zeroth-order sampling methods for non-log-concave distributions: Alleviating metastability by denoising diffusion. In *NeurIPS*, 2024.

Yuchen He and Chihao Zhang. On the query complexity of sampling from non-log-concave distributions. *arXiv preprint arXiv:2502.06200*, 2025.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

Jian Huang, Yuling Jiao, Lican Kang, Xu Liao, Jin Liu, and Yanyan Liu. Schrödinger-Föllmer sampler: Sampling without ergodicity. *arXiv preprint arXiv:2106.10880*, 2021.

Xunpeng Huang, Hanze Dong, Hao Yifan, Yi-An Ma, and Tong Zhang. Reverse diffusion Monte Carlo. In *ICLR*, 2024a.

Xunpeng Huang, Difan Zou, Hanze Dong, Yi-An Ma, and Tong Zhang. Faster sampling without isoperimetry via diffusion-based Monte Carlo. In *COLT*, 2024b.

Yuling Jiao, Lican Kang, Yanyan Liu, and Youzhou Zhou. Convergence analysis of Schrödinger-Föllmer sampler without convexity. *arXiv preprint arXiv:2107.04766*, 2021.

Holden Lee, Andrej Risteski, and Rong Ge. Beyond Log-concavity: Provable Guarantees for Sampling Multi-modal Distributions using Simulated Tempering Langevin Monte Carlo. In *NeurIPS*, 2018.

Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. In *NeurIPS*, 2022.

Joseph Lehec. The Langevin Monte Carlo algorithm in the non-smooth log-concave case. *The Annals of Applied Probability*, 2023.

Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards non-asymptotic convergence for diffusion-based generative models. In *ICLR*, 2024.

Iosif Lytras and Panayotis Mertikopoulos. Tamed Langevin sampling under weaker conditions. In *AISTATS*, 2025.

Yi-An Ma, Yuansi Chen, Chi Jin, Nicolas Flammarion, and Michael I Jordan. Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 2019.

Albert W Marshall, Ingram Olkin, and Barry C Arnold. *Inequalities: theory of majorization and its applications*. Springer, 1979.

Laurence Illing Midgley, Vincent Stimper, Gregor N.C. Simm, Bernard Schölkopf, and José Miguel Hernández-Lobato. Flow annealed importance sampling bootstrap. In *ICLR*, 2023.

Dan Mikulincer and Yair Shenfeld. *On the Lipschitz Properties of Transportation Along Heat Flows*. Springer International Publishing, 2023.

Wenlong Mou, Nicolas Flammarion, Martin J Wainwright, and Peter L Bartlett. An efficient sampling algorithm for non-smooth composite potentials. *Journal of Machine Learning Research*, 2022.

Dao Nguyen, Xin Dang, and Yixin Chen. Unadjusted Langevin algorithm for non-convex weakly smooth potentials. *Communications in Mathematics and Statistics*, 2023.

Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. In *COLT*, 2017.

Hamza Ruzayqat, Alexandros Beskos, Dan Crisan, Ajay Jasra, and Nikolas Kantas. Unbiased estimation using a class of diffusion processes. *Journal of Computational Physics*, 2023.

Saeed Saremi, Ji Won Park, and F. Bach. Chain of log-concave markov chains. In *ICLR*, 2024.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *NeurIPS*, 2019.

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

Francisco Vargas, Andrius Ovsianas, David Fernandes, Mark Girolami, Neil D. Lawrence, and Nikolas Nüsken. Bayesian learning via neural Schrödinger-Föllmer flows. In *AABI*, 2022.

Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices. In *NeurIPS*, 2019.

Gefei Wang, Yuling Jiao, Qian Xu, Yang Wang, and Can Yang. Deep generative learning via schrödinger bridge. In *ICLR*, 2021.

Yuchen Zhang, Percy Liang, and Moses Charikar. A hitting time analysis of stochastic gradient Langevin dynamics. In *COLT*, 2017.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes] , [No] , or [NA] .
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes] " is generally preferable to "[No] ", it is perfectly acceptable to answer "[No] " provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No] " or "[NA] " is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

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 and introduction list our contributions, i.e. novel complexity guarantees for a sampling algorithm that simulates a time-reversed diffusion with a specific estimator of the intermediate scores.

   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 discuss the limitations of our theoretical results in the Related Work section.

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: For each theoretical result, we worked with justified assumptions, making all the dependencies of the problem clear. We provide clear and detailed proofs in the appendix.

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: Our experiments are simple and illustrative of our theory. We provide details on how to reproduce them.

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: [NA]

   Justification: our experiments are illustrative, we do not use specific datasets.

   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: our experiments are simple and we detail their setup.

   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: [NA]

   Justification: our experiments illustrate a theoretical result on the convergence rate of an algorithm. The goal is not investigate the stochasticity of results across different runs.

   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: Our experiments are simple. We detail how long they take to run on a local computer.

   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: In our opinion, this paper does not address societal impact directly, and considers the generic problem of sampling from a distribution.

   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: [NA]

    Justification: In our opinion the paper does not have direct positive or negative social impact.

    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: Our paper does not present such risks.

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: [NA]

    Justification: we coded our our own experiments that are illustrative of our theoretical result.

    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: [Yes]

    Justification: we document the setup of our experiments.

    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: our experiments do not involve crowdsourcing.

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: our study does not involve risk for participants.

    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: LLMs are not a part of this research project.

    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 Section 3

A.1 Discussion assumptions Föllmer Flows

In the works of Huang et al. [2021], Jiao et al. [2021], Ruzayqat et al. [2023], the authors all make the assumption that $f := \frac{\mathrm{d}\mu}{\mathrm{d}\mathcal{N}(0,I_d)}$ is Lipschitz, has Lipschitz gradient and is bounded from below. We show that even if the second assumption is dropped, μ cannot be a Gaussian mixture.

Write $\mu = \sum_i^p w_i \mathcal{N}(\mu_i, \Sigma_i^{-1})$ so that f re-writes

$$f(x) = \sum_{i=1}^{p} \tilde{w}_i e^{-x^\top Q_i x + \tilde{\mu}_i^\top x}\,,$$

where $\tilde{w}_i = w_i \det\bigl(\Sigma_i^{-1}\bigr)^{1/2} e^{-\|\mu_i\|^2/2}$, $\tilde{\mu}_i = 2\Sigma_i^{-1}\mu_i$ and $Q_i = (\Sigma_i^{-1} - I_d)/2$. Now let us consider the smallest eigenvalue of the (Q_i): $\lambda = \inf_{\|u\|=1} \min_i u^\top Q_i u$ and a corresponding eigen value u. We denote $\mathcal{I} = \{i \mid u^\top Q_i u = \lambda\}$ the set of indexes of the matrices whose eigenvalue λ is associated with eigenvector u. Remark that for $i \notin \mathcal{I}$, $u^\top Q_i u > \lambda$. We prove that $\lambda \geq 0$. For $x = \alpha u$ with $\alpha \to \infty$, it holds that

$$\begin{aligned}\nabla f(x) &= \sum_{i\in\mathcal{I}} \tilde{w}_i(-2\alpha u + \tilde{\mu}_i) e^{-\alpha^2\lambda + \alpha\tilde{\mu}_i^\top u} + \sum_{i\notin\mathcal{I}} \tilde{w}_i e^{-\alpha^2 u^\top Q_i u + \alpha\tilde{\mu}_i^\top u} \\ &\sim -2\alpha u e^{-\lambda\alpha^2} \sum_{i\in\mathcal{I}} \tilde{w}_i e^{\alpha\tilde{\mu}_i^\top u}\,,\end{aligned}$$

In particular, for f to be Lipschitz, we must have $\lambda \geq 0$. This implies for all i that $\Sigma_i^{-1} \succeq I_d$. Furthermore, for f to be bounded from below, we need to have for some index i that $Q_i \preceq I_d$ which thus implies $Q_i = I_d$. We note $\mathcal{J}$ the set of indexes j such that $Q_j \succ I_d$. For all $j \in \mathcal{J}$, there exists $\|u_j\| = 1$ such that $u_j^\top Q_j u_j > 0$. Let us define $\bar{u} = \sum_{j\in\mathcal{J}} u_j$ consider $x = \alpha \sum_{j\in\mathcal{J}} u_j$ for $\alpha \to \infty$. As previously, it holds that

$$\nabla f(x) \sim -2u\alpha \sum_{j\notin\mathcal{J}} \tilde{w}_j e^{\alpha\tilde{\mu}_j^\top \bar{u}}\,.$$

To ensure that f is Lipschitz, it must hold that $\tilde{\mu}_j^\top \bar{u} \geq 0$. Similarly, considering $-\bar{u}$, it must hold that $\tilde{\mu}_j^\top \bar{u} \leq 0$ hence we eventually recover

$$\nabla f(x) \sim -2\alpha u \sum_{j\notin\mathcal{J}} \tilde{w}_j\,,$$

which contradicts f Lipschitz. Hence, there exits no Gaussian mixtures such that f is bounded from below and Lispchitz. More broadly, the complexity bounds in Huang et al. [2021], Jiao et al. [2021], Ruzayqat et al. [2023], Ding et al. [2023] depend polynomially on the lower-bound, on the Lispchitz and smoothness constants of f which are exponential in the gap $V - \frac{\|x\|^2}{2}$ making them very sensitive to deviation from purely Gaussian distributions. As noted in Vargas et al. [2022], this limitation is not only a theoretical artifact; in practice, these methods are very unstable and fail to convergence on simple examples.

A.2 Details on the non-smooth example

Consider the mixture $\mu = 0.5\mathcal{N}(0,\Sigma_1) + 0.5\mathcal{N}(0,\Sigma_2)$ with $\Sigma_1 = diag(1, 0.5)$ and $\Sigma_1 = diag(0.5, 1)$. For $(x,y) \in \mathbb{R}^2$, it holds that

$$-\nabla \log(\mu)(x,y) = \frac{(x,2y)e^{-x^2/2-y^2} + (2x,y)e^{-x^2-y^2/2}}{e^{-x^2/2-y^2} + e^{-x^2-y^2/2}}\,.$$

Hence, when $x = y$, we have $-\nabla\log(\mu)(x,y) = 3/2(x,x)$. Now for $y = x + \eta$, the score reads

$$\begin{aligned}-\nabla\log(\mu)(x,x+\eta) &= \frac{(x,2x+2\eta)e^{-x^2/2-x^2-2\eta x-\eta^2} + (2x,x+\eta)e^{-x^2-x^2/2-\eta x-\eta^2/2}}{e^{-x^2/2-x^2-2\eta x-\eta^2} + e^{-x^2-x^2/2-\eta x-\eta^2/2}} \\ &= \frac{(x,2x+2\eta)e^{-\eta x-\eta^2/2} + (2x,x+\eta)}{e^{-\eta x-\eta^2/2} + 1} \\ &\sim_{x\to+\infty} (2x,x)\,.\end{aligned}$$

In particular, $-\nabla \log(\mu)$ is not Hölder.

A.3 Proof of Proposition 2

Recall we denote $\lambda_{\min} := \min_i \lambda_{\min}(\Sigma_i)$, $\quad \lambda_{\max} := \max_i \lambda_{\max}(\Sigma_i)$. We write:

$$\mu(x) = \sum_{i=1}^{p} \tilde{w}_i \exp\left(-\frac{1}{2}(x-\mu_i)^\top \Sigma_i^{-1}(x-\mu_i)\right) := \sum_{i=1}^{p} \tilde{w}_i \phi_i(x)$$

with $\tilde{w}_i = w_i (2\pi)^{-d/2} \det(\Sigma_i)^{-1/2}$.

Bound on the Hessian. The Hessian of μ writes:

$$\nabla^2 \log \mu(x) = \frac{1}{\mu(x)} \nabla^2 \mu(x) - \frac{1}{\mu(x)^2} \nabla \mu(x) \nabla \mu(x)^\top$$

Since

$$\begin{aligned}
\nabla \phi_i(x) &= -\Sigma_i^{-1}(x-\mu_i)\, \phi_i(x) \\
\nabla^2 \phi_i(x) &= \left[\Sigma_i^{-1}(x-\mu_i)(x-\mu_i)^\top \Sigma_i^{-1} - \Sigma_i^{-1}\right] \phi_i(x)
\end{aligned}$$

We have

$$\begin{aligned}
\nabla \mu(x) &= \sum_i \tilde{w}_i \nabla \phi_i(x) = -\sum_i \tilde{w}_i \phi_i(x) \Sigma_i^{-1}(x-\mu_i) \\
\nabla^2 \mu(x) &= \sum_{i=1}^{p} \tilde{w}_i \nabla^2 \phi_i(x) = \sum_{i=1}^{p} \tilde{w}_i \phi_i(x) \left[\Sigma_i^{-1}(x-\mu_i)(x-\mu_i)^\top \Sigma_i^{-1} - \Sigma_i^{-1}\right]
\end{aligned}$$

Denoting $\gamma_i(x) := \frac{\tilde{w}_i \phi_i(x)}{\mu(x)}$ $\quad$ and $s_i(x) = -\Sigma_i^{-1}(x-\mu_i)$ we get

$$\begin{aligned}
\nabla^2 \log \mu(x) &= \sum_i \gamma_i(x) \left[s_i(x) s_i(x)^\top - \Sigma_i^{-1}\right] - \left(\sum_i \gamma_i(x) s_i(x)\right)\left(\sum_j \gamma_j(x) s_j(x)\right)^\top \\
&= \mathrm{Cov}_{\gamma(x)}[s_i(x)] - \sum_i \gamma_i(x) \Sigma_i^{-1}
\end{aligned}$$

where the first term is a covariance matrix of the vectors $s_i(x)$ under the weights $\gamma_i(x)$, hence which is a positive semi-definite matrix. Therefore:

$$-\nabla^2 \log \mu(x) \preceq \sum_i \gamma_i(x) \Sigma_i^{-1} \preceq \sum_i \gamma_i(x) \cdot \frac{I_d \cdot}{\lambda_{\min}} = \frac{I_d \cdot}{\lambda_{\min}}.$$

Bound on the drift. For the (negative) score of the mixture we have:

$$-\nabla \log \mu(x) = \sum_i \gamma_i(x) \Sigma_i^{-1}(x-\mu_i)$$

Hence

$$\begin{aligned}
\langle -\nabla \log \mu(x), x \rangle &= \sum_{i=1}^{p} \gamma_i(x) \langle \Sigma_i^{-1}(x-\mu_i), x \rangle \\
&= \sum_{i=1}^{p} \gamma_i(x) \left(x^\top \Sigma_i^{-1} x - \mu_i^\top \Sigma_i^{-1} x\right) \geq \sum_{i=1}^{p} \gamma_i(x) \left(\frac{\|x\|^2}{\lambda_{\max}} - \mu_i^\top \Sigma_i^{-1} x\right),
\end{aligned}$$

where we have used $\Sigma_i^{-1} \succeq \frac{1}{\lambda_{\max}} I_d$ for the first term. Now for the second term, using Cauchy-Schwartz and Young's inequality:

$$\mu_i^\top \Sigma_i^{-1} x \leq \|\mu_i\| \cdot \|\Sigma_i^{-1} x\| \leq \frac{\|\mu_i\| \|x\|}{\lambda_{\min}} \leq \frac{1}{2\lambda_{\min}} (\|x\|^2 + \|\mu_i\|^2).$$

Hence, using that $\sum_i \gamma_i(x) = 1$ and that $\sum_i \gamma_i(x)\|\mu_i\|^2 \leq \max_i \|\mu_i\|^2$ we have:

$$\begin{aligned}
\langle -\nabla \log \mu(x), x\rangle &\geq \sum_{i=1}^{p} \gamma_i(x)\left(\frac{\|x\|^2}{\lambda_{\max}} - \frac{1}{2\lambda_{\min}}(\|x\|^2 + \|\mu_i\|^2)\right) \\
&= \left(\frac{1}{\lambda_{\max}} - \frac{1}{2\lambda_{\min}}\right)\|x\|^2 - \frac{1}{2\lambda_{\min}}\sum_{i=1}^{p}\gamma_i(x)\|\mu_i\|^2. \\
&\geq \left(\frac{1}{\lambda_{\max}} - \frac{1}{2\lambda_{\min}}\right)\|x\|^2 - \frac{1}{2\lambda_{\min}}\max_i \|\mu_i\|^2 \\
&\geq \frac{\|x\|^2}{2\lambda_{\max}} - \lambda_{\max}\cdot \max_i \left(\frac{\|\mu_i\|}{\lambda_{\min}}\right)^2.
\end{aligned}$$

A.4 Discussion on [Lytras and Mertikopoulos, 2025]

To justify our claims at the end of Section 3.2, we first need some preliminary background on Poincaré constants. Let $q \in \mathcal{P}_{\mathrm{ac}}(\mathbb{R}^d)$. We say that q satisfies the *Poincaré inequality* with constant $C_P \geq 0$ if for all $f \in \mathcal{C}_c^\infty(\mathbb{R}^d)$ ($C^\infty(\mathbb{R}^d)$ functions with compact support):

$$\int f^2(x)dq(x) - \left(\int f(x)dq(x)\right)^2 = \mathrm{Var}_q(f) \leqslant C_P\|\nabla f\|^2_{L^2(q)}, \tag{6}$$

and let $C_P(q)$ be the best constant in Eq. 6, or $+\infty$ if it does not exist.

Proposition 10 *Let $V \in \mathcal{C}^4(\mathbb{R}^d)$, satisfying Assumption 1 and 2. Then the measure $\mu \propto e^{-V}$ satisfies a Poincaré inequality with constant*

$$C_P(\mu) \leq \textit{poly}\left(\frac{b+d}{a}\right) e^{\beta\frac{3b+2d}{a}}.$$

Proof. We will rely on Chewi and Stromme [2024, Lemma 10] which is itself adapted from Bakry et al. [2014, Theorem 4.6.2]. Following their notation, we define the operator $L_\mu := \Delta + \langle \nabla \log \mu, \nabla\rangle = \Delta + \langle -\nabla V, \nabla\rangle$ and a Lyapunov function $W(x) = \|x\|^2 + 1$.

The Lyapunov function verifies the following inequality

$$\begin{aligned}
-L_\mu(W)(x) &= -\Delta W(x) - \langle -\nabla V(x), 2x\rangle = -2d + 2\langle \nabla V(x), x\rangle & (7)\\
&\geq -2d + 2(a\|x\|^2 - b) = 2aW(x) - 2(a+b+d) & (8)
\end{aligned}$$

where we used Assumption 2. This yields

$$\frac{-L(W)(x)}{W(x)} \geq 2a - 2\frac{a+b+d}{W(x)} \tag{9}$$

In general case, by noting that $W(x) \geq 1$, we can lower-bound this as

$$\frac{-L(W)(x)}{W(x)} \geq -2(a+b+d). \tag{10}$$

In the special case when x satisfies $\|x\|^2 \geq 1 + 2\frac{b+d}{a}$, or equivalently, when x is located outside the ball B_R with radius $R = \sqrt{1 + 2\frac{b+d}{a}}$, then we can refine the lower bound and get

$$\frac{-L(W)(x)}{W(x)} \geq a. \tag{11}$$

We can summarize these two cases in one equation, as

$$\frac{-L(W)(x)}{W(x)} \geq a\Big(1 - \big(3 + \frac{2b}{a} + \frac{2d}{a}\big)\mathbb{I}_{B_R}\Big). \tag{12}$$

By applying Chewi and Stromme [2024, Lemma 10], we obtain that

$$C_P(\mu) \leq \frac{1}{a} + \big(3 + \frac{2b}{a} + \frac{2d}{a}\big) C_P(\mu_{|B_R}). \tag{13}$$

We are thus left to control $C_P(\mu_{|B_R})$, which we will obtain by comparing $\mathrm{Var}_{\mu_{|B_R}}(f)$ and $\|\boldsymbol{\nabla} f\|^2_{L^2(\mu_{|B_R})}$. Define $\mu_{|B_R}(f) = \int f(x) d\mu_{|B_R}(x)$ and let $\alpha(x)$ be the Lebesgue measure. For any test function f, we have

$$\mathrm{Var}_{\mu_{|B_R}}(f) = \int (f - \mu_{|B_R}(f))^2 d\mu_{|B_R}(x) \leq \int (f - \alpha_{|B_R}(f))^2 d\mu_{|B_R}(x) \tag{14}$$

given that the mean $\mu_{|B_R}(f)$ minimizes the quadratic error between f and a scalar. Then, we have

$$\mathrm{Var}_{\mu_{|B_R}}(f) \leq \int (f - \alpha_{|B_R}(f))^2 \frac{e^{-V(x)}}{\int_{B_R} e^{-V(x)} dx} dx \tag{15}$$

$$\leq \frac{\sup_{x \in B_R} e^{-V(x)}}{\int_{B_R} e^{-V(x)} dx} \int_{B_R} (f - \alpha_{|B_R}(f))^2 dx \tag{16}$$

We recall that the uniform (Lebesgue) measure restricted to B_R satisfies $C_P(\alpha_{|B_R}) \leq \frac{4R^2}{\pi^2}$, see [CRAIG, 2021, Eq (1.1)]), hence

$$\mathrm{Var}_{\mu_{|B_R}}(f) \leq \frac{\sup_{x \in B_R} e^{-V(x)}}{\int_{B_R} e^{-V(x)} dx} \frac{4R^2}{\pi^2} \int_{B_R} \|\nabla f(x) - \nabla \alpha_{|B_R}(f)\|^2 dx \tag{17}$$

$$= \frac{\sup_{x \in B_R} e^{-V(x)}}{\int_{B_R} e^{-V(x)} dx} \frac{4R^2}{\pi^2} \int_{B_R} \|\nabla f(x)\|^2 dx \tag{18}$$

$$\leq \frac{\sup_{x \in B_R} e^{-V(x)}}{\int_{B_R} e^{-V(x)} dx} \frac{4R^2}{\pi^2} \int \|\nabla f(x)\|^2 \frac{e^{-V(x)}}{\inf_{x \in B_R} e^{-V(x)}} \tag{19}$$

$$= \frac{\sup_{x \in B_R} e^{-V(x)}}{\inf_{x \in B_R} e^{-V(x)}} \frac{4R^2}{\pi^2} \int \|\nabla f(x)\|^2 d\mu_{|B_R}(x). \tag{20}$$

From the above, we read that the Poincaré constant must satisfy

$$C_P(\mu_{|B_R}) \leq \frac{4R^2}{\pi^2} \cdot \frac{\sup_{x \in B_R} e^{-V(x)}}{\inf_{x \in B_R} e^{-V(x)}} \tag{21}$$

$$= \frac{4R^2}{\pi^2} \cdot \sup_{x \in B_R} e^{V(x^*) - V(x)} \tag{22}$$

$$= \frac{4(1 + 2\frac{b+d}{a})^2}{\pi^2} \cdot \sup_{x \in B_R} e^{V(x^*) - V(x)} \tag{23}$$

where $x^* \in \arg\max_{x \in B_R} V(x)$. There remains to bound $\sup_{x \in B_R} e^{V(x^*) - V(x)}$.

Using the smoothness assumption, $\nabla^2 V \preceq \beta I_d$, we can write

$$0 \leq V(x*) - V(x) \leq \frac{\beta}{2} \|x - x^*\|^2 \leq \frac{\beta}{2} \left(\|x\|^2 + \|x^*\|^2 \right). \tag{24}$$

We can now bound each of the two rightmost terms. Using the dissipativity assumption $\langle \nabla V(x*), x* \rangle = 0 \geq a\|x*\|^2 - b$, we obtain $\|x^*\|^2 \leq \frac{b}{a}$. Furthermore, because x is in B_R, we have $\|x\|^2 \leq 1 + 2\frac{b+d}{a}$. Plugging this in, we get

$$0 \leq V(x*) - V(x) \leq \frac{\beta}{2} \big(1 + \frac{3b + 2d}{a}\big) \tag{25}$$

and

$$\exp(V(x^*) - V(x)) \leq \exp\big(\frac{\beta}{2}\big(1 + \frac{3b + 2d}{a}\big)\big). \tag{26}$$

Finally, we have

$$C_P(\mu_{|B_R}) \leq \frac{4(1+2\frac{b+d}{a})^2}{\pi^2}\exp\big(\frac{\beta}{2}\big(1+\frac{3b+2d}{a}\big)\big) \tag{27}$$

and

$$\begin{aligned} C_P(\mu) &\leq \frac{1}{a} + \big(3+\frac{2b}{a}+\frac{2d}{a}\big)C_P(\mu_{|B_R}) && (28)\\ &\leq \frac{1}{a} + \big(3+\frac{2b}{a}+\frac{2d}{a}\big)\frac{4(1+2\frac{b+d}{a})^2}{\pi^2}\exp\big(\frac{\beta}{2}\big(1+\frac{3b+2d}{a}\big)\big) && (29)\end{aligned}$$

□

In Lytras and Mertikopoulos [2025, Theorem 3], the best complexity obtained is of the form poly$(C_P(\mu), \epsilon^{-1})$. By applying the above bound with a, b, and β as in Proposition 10 we recover the expected result.

B Proof of Proposition 4

Before starting the proof, we recall the following identities.

Proposition 11 (Tweedie's formulas) *Denoting p_t^V the density of the forward process X_t initialized at $\mu \propto e^{-V}$, and $Y_t \sim \mathcal{N}(0, (1-e^{-2t})I_d)$, it holds for all $z \in \mathbb{R}^d$ that*

$$p_t^V(z) = \frac{e^{td}}{Z_V}\mathbb{E}[e^{-V((z-Y_t)e^t)}]\,, \tag{30}$$

that

$$\nabla \log\big(p_t^V\big)(z) = \frac{\mathbb{E}[-Y_t e^{-V((z-Y_t)e^t)}]}{(1-e^{-2t})\mathbb{E}[e^{-V((z-Y_t)e^t)}]}\,, \tag{31}$$

and that

$$\nabla^2 \log\big(p_t^V\big)(z) = \frac{\mathbb{E}[Y_t Y_t^\top e^{-V((z-Y_t)e^t)}]}{(1-e^{-2t})^2\mathbb{E}[e^{-V((z-Y_t)e^t)}]} - \frac{I_d}{(1-e^{-2t})} - (\nabla \log\big(p_t^V\big)(z)))(\nabla \log\big(p_t^V\big)(z))^\top\,. \tag{32}$$

Proof. Recall that p_t^V is the law of the variable

$$X_t = e^{-t}X_0 + B_{1-e^{-2t}}\,,$$

with $X_0 \sim \mu$ and B_s the standard Brownian motion evaluated at time s. Hence, using Bayes formula, we have

$$p_t^V(z) = \int p_t(z|x)dp_0(x) = \frac{1}{Z^V(1-e^{-2t})^{d/2}(2\pi)^{d/2}}\int_{\mathbb{R}^d} e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}}e^{-V(x)}dx\,.$$

After taking the logarithm and differentiating with respect to z, we obtain

$$\nabla \log\big(p_t^V\big)(z) = \frac{\int_{\mathbb{R}^d} -(z-xe^{-t})e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}}e^{-V(x)}dx}{(1-e^{-2t})\int_{\mathbb{R}^d} e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}}e^{-V(x)}dx}\,.$$

To obtain the Hessian, we differentiate the formula above. The Jacobian of the numerator is given by

$$-I_d\int_{\mathbb{R}^d} e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}}e^{-V(x)}dx + \frac{1}{1-e^{-2t}}\int_{\mathbb{R}^d}(z-xe^{-t})(z-xe^{-t})^\top e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}}e^{-V(x)}dx\,,$$

from which we can deduce

$$
\begin{aligned}
\nabla^2 \log\big(p_t^V\big)(z) = &-\frac{I_d}{(1-e^{-2t})} + \frac{\int_{\mathbb{R}^d} (z-xe^{-t})(z-xe^{-t})^\top e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}} e^{-V(x)} dx}{(1-e^{-2t})^2 \int_{\mathbb{R}^d} e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}} e^{-V(x)} dx} \\
&- \frac{(\int_{\mathbb{R}^d} (z-xe^{-t}) e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}} e^{-V(x)} dx)(\int_{\mathbb{R}^d} (z-xe^{-t}) e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}} e^{-V(x)} dx)^\top}{(1-e^{-2t})^2 (\int_{\mathbb{R}^d} e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}} e^{-V(x)} dx)^2} \\
= &-\frac{I_d}{(1-e^{-2t})} + \frac{\int_{\mathbb{R}^d} (z-xe^{-t})(z-xe^{-t})^\top e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}} e^{-V(x)} dx}{(1-e^{-2t})^2 \int_{\mathbb{R}^d} e^{-\frac{\|z-xe^{-t}\|^2}{2(1-e^{-2t})}} e^{-V(x)} dx} \\
&- (\nabla \log\big(p_t^V\big)(z))(\nabla \log\big(p_t^V\big)(z))^\top .
\end{aligned}
$$

In order to rewrite the quantities above as expectations, we make the change of variable $y = z - xe^{-t}$ so that $x = (z-y)e^t$ and we obtain for the density p_t^V :

$$
p_t^V(z) = \frac{e^{td}}{Z_V} \mathbb{E}[e^{-V((z-Y_t)e^t)}] ,
$$

where $Y_t \sim \mathcal{N}(0, I_d(1-e^{-2t}))$. Conversely, the score rewrites as

$$
\nabla \log\big(p_t^V\big)(z) = \frac{\mathbb{E}[-Y_t e^{-V((z-Y_t)e^t)}]}{(1-e^{-2t})\mathbb{E}[e^{-V((z-Y_t)e^t)}]} ,
$$

and the Hessian rewrites as

$$
\nabla^2 \log\big(p_t^V\big)(z) = \frac{\mathbb{E}[Y_t Y_t^\top e^{-V((z-Y_t)e^t)}]}{(1-e^{-2t})^2 \mathbb{E}[e^{-V((z-Y_t)e^t)}]} - \frac{I_d}{(1-e^{-2t})} - (\nabla \log\big(p_t^V\big)(z)))(\nabla \log\big(p_t^V\big)(z))^\top .
$$

□

For the rest of the proof we shall drop the dependence in z and write $\hat{s}_{t,n}(z) = \frac{\hat{N}}{\hat{D}}$ with the empirical numerator $\hat{N} = -\frac{\sum_{i=1}^n y_i e^{-V(e^t(z-y_i))}}{1-e^{-2t}}$ and denominator $\hat{D} = \sum_{i=1}^n e^{-V(e^t(z-y_i))}$ where we recall $y_i \sim \mathcal{N}(0, (1-e^{-2t})I_d)$. In the following proposition, we explicitly compute the variances of $\hat{N}$ and $\hat{D}$ using the formulas above. In what follows, we shall denote $N = \mathbb{E}[N]$ and $D = \mathbb{E}[D]$

Proposition 12 (Variance of estimators) *Let $y_1, \dots, y_n$ i.i.d. distributed as $\mathcal{N}(0, (1-e^{-2t})I_d)$. Denote by π a standard normal density, and by $\hat{N}(z)$ and $\hat{D}(z)$ the numerator and denominator of the estimator defined in Eq. 5. We have:*

$$
\begin{aligned}
\mathbb{E}[\|\hat{N} - N\|^2] &\le \frac{p_t^{2V} Z_{2V} e^{-td}}{n} \left(\Delta \log\big(p_t^{2V}\big)(z) - \frac{\Delta \log(\pi)(z)}{1-e^{-2t}} + \|\nabla \log\big(p_t^{2V}\big)\|^2(z) \right) , \\
\mathbb{E}[\|\hat{D} - D\|^2] &\le \frac{p_t^{2V}(z) Z_{2V} e^{-td}}{n} .
\end{aligned}
$$

Proof. For the numerator, we have

$$
\hat{N} - N = \frac{-1}{n} \sum_{i=1}^n \frac{y_i e^{-V((z-y_i)e^t)}}{1-e^{-2t}} + N ,
$$

hence, since the $y_i,\ i = 1, \dots, n$ are i.i.d. distributed as $Y_t \sim \mathcal{N}(0, (1-e^{-2t})I_d)$,

$$
\mathbb{E}[\|\hat{N} - N\|^2] = \frac{1}{n} \mathbb{E}\left[\left\| \frac{Y_t e^{-V((z-Y_t)e^t)}}{1-e^{-2t}} - N(z) \right\|^2 \right] \le \frac{1}{n} \mathbb{E}\left[\frac{\|Y_t\|^2 e^{-2V((z-Y_t)e^t)}}{(1-e^{-2t})^2} \right] .
$$

Taking in the trace in the log hessian identity in Proposition 11 yields

$$\mathbb{E}\left[\frac{\|Y_t\|^2 e^{-2V((z-Y_t)e^t)}}{(1-e^{-2t})^2}\right] = \left(\Delta \log\!\left(p_t^{2V}\right) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log\!\left(p_t^{2V}\right)\|^2\right) p_t^{2V} Z_{2V} e^{-td} \,.$$

Similarly, we have

$$\hat{D} - D = \frac{1}{n}\sum_{i=1}^{n} e^{-V((z-y_i)e^t)} - D \,,$$

hence we get using again Proposition 11,

$$\mathbb{E}[(\hat{D}-D)^2] = \frac{1}{n}\left(\mathbb{E}[e^{-2V((z-Y_t)e^t)}] - \mathbb{E}[e^{-2((z-Y_t)e^t)}]^2\right) \leq \frac{p_t^{2V}(z) Z_{2V} e^{-td}}{n} \,. \qquad (33)$$

□

We can now prove Proposition 4.

Proof. Define the event $A = (\hat{D} \geq \eta D) \cap (\|\hat{N}\| \leq \kappa \|N\|)$ where $\eta \leq 1, \kappa \geq 1$ are positive scalars to be chosen later. We start to decompose the quadratic error as:

$$\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2\right] = \mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2 \mathbb{1}_A\right] + \mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2 \mathbb{1}_{\bar{A}}\right] .$$

We now separately analyze the first and the second term. For the first term, define

$$\begin{cases} \theta : \mathbb{R}^d \times \mathbb{R}^* \to \mathbb{R} \\ (x,p) \mapsto \left\|\frac{x}{p} - \frac{N}{D}\right\|^2 . \end{cases}$$

The gradient and Hessian of θ are given by

$$\begin{cases} \nabla\theta(x,p) = \frac{-2}{p}\left(\frac{N}{D} - \frac{x}{p}, \left\|\frac{x}{p}\right\|^2 - \left\langle \frac{x}{p}, \frac{N}{D}\right\rangle\right) , \\ \nabla^2\theta(x,p) = \frac{-2}{p^2}\begin{pmatrix} -I_d & (\frac{2x}{p} - \frac{N}{D})^\top \\ (\frac{2x}{p} - \frac{N}{D}) & -3\left\|\frac{x}{p}\right\|^2 + 2\langle \frac{x}{p}, \frac{N}{D}\rangle \end{pmatrix} . \end{cases}$$

We thus make a Taylor expansion of order 2 of $\theta(\hat{N}, \hat{D})$ around (N, D): there exists (a random) $\hat{t} \in [0,1]$ such that

$$\begin{aligned} \theta(\hat{N}, \hat{D}) &= \theta(N, D) + \nabla\theta(N,D)^\top(\hat{N} - N, \hat{D} - D) \\ &\quad + \frac{1}{2}(\hat{N} - N, \hat{D} - D)^\top \nabla^2\theta(\hat{N}_{\hat{t}}, \hat{D}_{\hat{t}})(\hat{N} - N, \hat{D} - D) \,. \end{aligned}$$

where we denoted $\hat{N}_{\hat{t}} = \hat{t}\hat{N} + (1-\hat{t})N$ and $\hat{D}_{\hat{t}} = \hat{t}\hat{D} + (1-\hat{t})D$. The two first terms in the expansion are null and we are left with

$$\begin{aligned} \theta(\hat{N}, \hat{D}) &= \frac{1}{\hat{D}_{\hat{t}}^2}\left(\|\hat{N} - N\|^2 - 2(\langle 2\frac{\hat{N}_{\hat{t}}}{\hat{D}_{\hat{t}}} - \frac{N}{D}, \hat{N} - N\rangle(\hat{D} - D)) + \left(3\frac{\|\hat{N}_{\hat{t}}\|^2}{\hat{D}_{\hat{t}}^2} - 2\langle \frac{\hat{N}_{\hat{t}}}{\hat{D}_{\hat{t}}}, \frac{N}{D}\rangle\right)(\hat{D} - D)^2\right) \\ &\leq \frac{1}{\hat{D}_{\hat{t}}^2}\left(\|\hat{N} - N\|^2 + 2\|2\frac{\hat{N}_{\hat{t}}}{\hat{D}_{\hat{t}}} - \frac{N}{D}\|\|\hat{N} - N\||\hat{D} - D| + \left(3\frac{\|\hat{N}_{\hat{t}}\|^2}{\hat{D}_{\hat{t}}^2} + 2\left\|\frac{\hat{N}_{\hat{t}}}{\hat{D}_{\hat{t}}}\right\|\left\|\frac{N}{D}\right\|\right)(\hat{D} - D)^2\right) . \end{aligned}$$

Hence, almost surely over A

$$\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2 \leq \frac{1}{\eta^2 D^2}\left(\|\hat{N} - N\|^2 + \frac{6\kappa}{\eta}\left\|\frac{N}{D}\right\|\|\hat{N} - N\||\hat{D} - D| + 5\left\|\frac{N}{D}\right\|^2\left(\frac{\kappa}{\eta}\right)^2(\hat{D} - D)^2\right) .$$

Hence, after taking the expectation and applying Cauchy-Schwarz, we obtain

$$\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}}-\frac{N}{D}\right\|^2 \mathbb{1}_A\right] \leq \frac{1}{\eta^2 D^2}(\mathbb{E}[\|\hat{N}-N\|^2] + \frac{6\kappa}{\eta}\left\|\frac{N}{D}\right\|\mathbb{E}[\|\hat{N}-N\|^2]^{1/2}\mathbb{E}[(\hat{D}-D)^2]^{1/2}$$
$$+ 5\left\|\frac{N}{D}\right\|^2\left(\frac{\kappa}{\eta}\right)^2\mathbb{E}[(\hat{D}-D)^2])\,.$$

Now recall that $D = p_t^V Z_V e^{-td}$ and that $\frac{N}{D} = \nabla\log\big(p_t^V\big)$ which, combined with Proposition 12 yields for the first term:

$$\frac{\mathbb{E}[\|\hat{N}-N\|^2]}{\eta^2 D^2} \leq \frac{p_t^{2V} Z_{2V} e^{td}}{\eta^2 (p_t^V)^2 (Z_V)^2 n}\left(\Delta\log\big(p_t^{2V}\big) - \frac{\Delta\log(\pi)}{1-e^{-2t}} + \|\nabla\log\big(p_t^{2V}\big)\|^2\right)\,,$$

for the second term:

$$\frac{6\kappa}{\eta^3 D^2}\left\|\frac{N}{D}\right\|\|\hat{N}-N\||\hat{D}-D| \leq \frac{6\kappa p_t^{2V} Z_{2V} e^{td}}{\eta^3 (p_t^V)^2 (Z_V)^2 n}\|\nabla\log\big(p_t^V\big)\|\left(\Delta\log\big(p_t^{2V}\big) - \frac{\Delta\log(\pi)}{1-e^{-2t}} + \|\nabla\log\big(p_t^{2V}\big)\|^2\right)^{1/2}$$

and for the last term:

$$5\left\|\frac{N}{D}\right\|^2\left(\frac{\kappa}{\eta^2 D}\right)^2(\hat{D}-D)^2 \leq \frac{5\kappa^2 p_t^{2V} Z_{2V} e^{td}}{\eta^4 (p_t^V)^2 (Z_V)^2 n}\|\nabla\log\big(p_t^V\big)\|^2\,.$$

Hence we finally obtain

$$\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}}-\frac{N}{D}\right\|^2 \mathbb{1}_A\right] \leq \frac{p_t^{2V} Z_{2V} e^{td}}{n\eta^2 (p_t^V Z_V)^2}\Bigg(\frac{6\kappa}{\eta}\|\nabla\log\big(p_t^V\big)\|\left[\Delta\log\big(p_t^{2V}\big) - \frac{\Delta\log(\pi)}{1-e^{-2t}} + \|\nabla\log\big(p_t^{2V}\big)\|^2\right]^{1/2}$$
$$+ \Delta\log\big(p_t^{2V}\big) - \frac{\Delta\log(\pi)}{1-e^{-2t}} + 6\left(\frac{\kappa}{\eta}\right)^2\|\nabla\log\big(p_t^V\big)\|^2\Bigg)\,. \tag{34}$$

Let us now handle the quadratic error of the estimator on the complementary $\bar{A}$. We have, using Young's inequality $||a-b||^2 \leq 2(||a||^2 + ||b||^2)$,

$$\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}}-\frac{N}{D}\right\|^2 \mathbb{1}_{\bar{A}}\right] \leq 2\left\|\frac{N}{D}\right\|^2\mathbb{P}(\bar{A}) + 2\mathbb{E}\left[\mathbb{1}_{\bar{A}}\frac{\max_i\|y_i\|^2}{(1-e^{-2t})^2}\right]\,.$$

Now, recall that $X_i = ||y_i||^2(1-e^{-2t})^{-1}$ are n independent variables such that for all i, $X_i \sim \chi^2(d)$. Using Hölder inequality for some $p \geq 1$, the second term can be upper-bounded as

$$\mathbb{E}\left[\mathbb{1}_{\bar{A}}\frac{\max_i\|y_i\|^2}{(1-e^{-2t})^2}\right] \leq \frac{1}{1-e^{-2t}}\mathbb{E}[\max_i X_i^p]^{1/p}\mathbb{P}(\bar{A})^{1-1/p}$$
$$\leq \frac{1}{1-e^{-2t}}(n\mathbb{E}[X_1^p])^{1/p}\mathbb{P}(\bar{A})^{1-1/p} \leq \frac{1}{1-e^{-2t}}n^{1/p}(d+2p)\mathbb{P}(\bar{A})^{1-1/p}\,,$$

where we used in the penultimate inequality that the max is smaller than the sum, and in the last one that $\mathbb{E}[X_1^p] = \prod_{i=0}^{p-1}(d+2i)$ when $X_1 \sim \chi^2(d)$ combined with the fact that the geometric mean is lower than the arithmetic mean.

We now upper bound the probability of the event $\bar{A} = (\hat{D} < \eta D)\cup(\|\hat{N}\| > \kappa\|N\|)$. By Chebyshev's inequality, using $\eta < 1$, it holds that

$$\mathbb{P}(\hat{D} < \eta D) \leq \frac{\mathbb{E}[(\hat{D}-D)^2]}{D^2(\eta-1)^2} \leq \frac{p_t^{2V} Z_{2V} e^{td}}{n(p_t^V Z_V)^2(\eta-1)^2} := \frac{U}{n(\eta-1)^2}\,.$$

Similarly, recalling that $||N|| = D||\nabla\log\big(p_t^V\big)||$, we have

$$\mathbb{P}(\|\hat{N}\| > \kappa\|N\|) \leq \frac{\mathbb{E}[||\hat{N}-N||^2]}{||N||^2(\kappa-1)^2}$$
$$\leq \frac{U}{n\|\nabla\log\big(p_t^V\big)\|^2(\kappa-1)^2}\left(\Delta\log\big(p_t^{2V}\big) - \frac{\Delta\log(\pi)}{1-e^{-2t}} + \|\nabla\log\big(p_t^V\big)\|^2\right)\,.$$

We now make a disjunction of cases: if $\|\nabla \log(p_t^V)\| \geq 1$, we pick $\eta = 1/2$ and $\kappa = 3/2$ so we recover

$$\begin{aligned}
\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2 \mathbb{1}_{\bar{A}}\right] &\leq \frac{8U}{n}\left[\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + 2\|\nabla \log(p_t^V)\|^2\right] \\
&\quad + \frac{n^{1/p}(d+2p)}{1-e^{-2t}}\left[\frac{8U}{n}\left(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^V)\|^2 + 1\right)\right]^{1-1/p} \\
&\leq \frac{8U}{n}\left[\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + 2\|\nabla \log(p_t^V)\|^2\right] \\
&\quad + \frac{n^{1/p}(d+2p)}{1-e^{-2t}}\left[\frac{8U}{n}\left(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^V)\|^2 + 1\right) + 1\right].
\end{aligned}$$

We thus pick $p = \log(n)$ to get

$$\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2 \mathbb{1}_{\bar{A}}\right] \leq \frac{16e^2(d+2\log(n))}{n(1-e^{-2t})}\left[U\left(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^V)\|^2 + 1\right) + 1\right].$$

Combining this with the bound Eq. 34 eventually yields

$$\delta^2 \leq \frac{32e^2(d+2\log(n))}{n(1-e^{-2t})}\left[U\left(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^V)\|^2 + \|\nabla \log(p_t^{2V})\|^2 + 1\right) + 1\right],$$

where we used the inequality

$$\begin{aligned}
&\|\nabla \log(p_t^V)\|\left[\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^{2V})\|^2\right]^{1/2} \\
&\leq \frac{1}{2}\left(\|\nabla \log(p_t^V)\|^2 + \Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^{2V})\|^2\right).
\end{aligned}$$

In the case where $\|\nabla \log(p_t)\|^2 < 1$, we instead pick $\eta = 1/2$ and $\kappa = 1 + \frac{1}{2\|\nabla \log(p_t^V)\|}$. We obtain that

$$\mathbb{P}(\bar{A}) \leq \frac{4U}{n}\left(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^V)\|^2 + 1\right)$$

and as previously, for $p = \log(n)$ we get

$$\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2 \mathbb{1}_{\bar{A}}\right] \leq \frac{16e^2(d+2\log(n))}{n(1-e^{-2t})}\left[U\left(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^V)\|^2 + 1\right) + 1\right].$$

For this choice of κ, η, the bound on A becomes

$$\begin{aligned}
\mathbb{E}\left[\left\|\frac{\hat{N}}{\hat{D}} - \frac{N}{D}\right\|^2 \mathbb{1}_{A}\right] &\leq \frac{4U}{n}\left(6\left[\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^{2V})\|^2\right]^{1/2} + \Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + 6\right) \\
&\leq \frac{4U}{n}\left(7(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}}) + 6(\|\nabla \log(p_t^{2V})\|^2 + 1)\right).
\end{aligned}$$

Thus, we obtain as previously

$$\delta^2 \leq \frac{32e^2(d+2\log(n))}{n(1-e^{-2t})}\left[U\left(\Delta \log(p_t^{2V}) - \frac{\Delta \log(\pi)}{1-e^{-2t}} + \|\nabla \log(p_t^V)\|^2 + \|\nabla \log(p_t^{2V})\|^2 + 1\right) + 1\right],$$

$\square$

C Proof of Proposition 6 and Lemma 8

Before starting the proofs, we recall the following usefull lemma that bounds the second order moment of dissipative distributions.

Lemma 13 *Let V be such that $\langle \nabla V(x), x\rangle \geq a\|x\|^2 - b$. Then for $\mu \propto e^{-V}$, denoting m_2 the second moment of μ, it holds that $m_2 \leq \frac{b+d}{a}$.*

Proof. Define the Laplacian of μ as $L(f) = \Delta f - \langle \nabla V, \nabla f\rangle$ for f sufficiently smooth. By integration by parts, it holds that

$$\int L(f)(x)\mathrm{d}\mu(x) = 0\,.$$

In particular, for $f(x) = \|x\|^2$, we recover $\int \langle \nabla V(x), \nabla f(x)\rangle \mathrm{d}\mu(x) = 2d$. For V dissipative, it implies

$$a\int \|x\|^2 \mathrm{d}\mu(x) \leq 2d + b\,,$$

or equivalently $m_2 \leq (b+2d)/a$. □

C.1 Proof of Lemma 8

Proof. The first inequality was shown in the Lemma above. For the Fisher information, it holds that

$$\begin{aligned}
\mathcal{I}(\mu, \pi) &= \int \|\nabla V(x) - x\|^2 \mathrm{d}\mu(x)\,,\\
&\leq 2\int \langle \nabla V(x), \nabla V(x) e^{-V(x)}\rangle / Z_V \mathrm{d}x + 2m_2\,,\\
&= 2\int \Delta V \mathrm{d}\mu(x) + 2m_2\,,\\
&\leq 2\beta d + 2m_2\,.
\end{aligned}$$

There remains to lower-bound $\mu(0)$. Denote x^* a global minimizer of V. By dissipativity, it must hold that $\|x^*\|^2 \leq b/a$. Now observe that

$$\begin{aligned}
V(x^*) - V(0) &= \int_0^1 (x^*)^\top \nabla^2 V(tx^*) x^* \mathrm{d}t\,,\\
&\leq \beta\|x^*\|^2\,.
\end{aligned}$$

Combined with the fact that for all $x \in \mathbb{R}^d$, it holds that $V(x) - V(x^*) \geq 0$, be recover that $V(x) - V(0) \geq -\beta\|x^*\|^2 \geq -\beta b/a$. Furthermore, for $0 < \delta < 1$, we have

$$\begin{aligned}
V(x) - V(0) &= \int_0^1 \langle \nabla V(tx), x\rangle \mathrm{d}t\,,\\
&= \int_\delta^1 \langle \nabla V(tx), x\rangle \mathrm{d}t + V(\delta x) - V(0)\,,\\
&\geq \int_\delta^1 \langle \nabla V(tx), tx\rangle / t\mathrm{d}t - \beta b/a\,,\\
&\geq \int_\delta^1 at\|x\|^2 - b/t\mathrm{d}t\,,\\
&= a\|x\|^2/2 - a\delta^2\|x\|^2/2 + b\log(\delta) - \beta b/a\,.
\end{aligned}$$

In particular, for $\delta = 1/\sqrt{2}$, we obtain

$$V(x) - V(0) \geq a\|x\|^2/4 - b(\beta/a + \log(2)/2)\,.$$

Hence, we have

$$\begin{aligned}\mu(0) &= \frac{e^{-V(0)}}{\int e^{-V(x)}\mathrm{d}x}\,,\\ &= \frac{1}{\int e^{-(V(x)-V(0))}\mathrm{d}x}\,,\\ &\leq \frac{1}{\int e^{-a\|x\|^2/4+b(\beta/a+\log(2)/2)}\mathrm{d}x}\\ &= e^{-b(\beta/a+\log(2)/2)}(a/2)^{d/2}(2\pi)^{-d/2}\,.\end{aligned}$$

Since $\beta/a \geq 1$ and $\log(2)/2 \leq 1$, we recover that $\log\big(\mu(0)^{-2/d}\big) \leq 4\beta b/ad + 2\pi + \log(2/a)$.

$\square$

C.2 Proof of Proposition 6

Recall that the intermediate scores read

$$\nabla \log\big(p_t^V\big)(z) = \frac{z - e^{-t}\mathbb{E}_{q_{t,z}}[Y]}{1-e^{-2t}}\,,$$

with $q_{t,z}(x) \propto e^{-V(x)}e^{-\frac{\|e^{-t}x-z\|^2}{2(1-e^{-2t})}}$. In particular, if V is dissipative with constant a, b then it holds that

$$\begin{aligned}\langle -\nabla \log(q_{t,z})(x), x\rangle &= \langle \nabla V(x) + \frac{e^{-t}(e^{-t}x - z)}{1-e^{-2t}}, x\rangle\,,\\ &\geq a\|x\|^2 - b + \frac{e^{-t}}{1-e^{-2t}}(e^{-t}\|x\|^2 - \langle z, x\rangle).\end{aligned}$$

Now recall that $e^{-t}\|x\|^2 - \langle z, x\rangle \geq -e^{t}\|z\|^2/4$ which yields

$$\langle -\nabla \log(q_{t,z})(x), x\rangle \geq a\|x\|^2 - (b + \|z\|^2/(4(1-e^{-2t})))\,.$$

Hence, using Lemma 13, it holds that

$$\mathbb{E}_{q_{t,z}}[\|Y\|^2] \leq \frac{2b + \|z\|^2/(2(1-e^{-2t})) + d}{a}\,.$$

Hence we recover that

$$\begin{aligned}\|\nabla \log\big(p_t^V\big)(z)\|^2 &\leq \frac{2\|z\|^2}{(1-e^{-2t})^2} + \frac{2e^{-2t}\mathbb{E}_{q_{t,z}}[\|Y\|^2]}{(1-e^{-2t})^2}\,,\\ &\leq \frac{2\|z\|^2}{(1-e^{-2t})^2} + \frac{e^{-2t}}{(1-e^{-2t})^2}\left(\frac{2(2b+d)}{a} + \frac{\|z\|^2}{(1-e^{-2t})a}\right),\\ &= \frac{\|z\|^2}{(1-e^{-2t})^2}\left(2 + \frac{e^{-2t}}{a(1-e^{-2t})}\right) + \frac{2e^{-2t}(2b+d)}{a(1-e^{-2t})^2}\,.\end{aligned}$$

Similarly, recall that

$$\nabla^2 \log\big(p_t^V\big)(z) - \nabla^2 \log(\pi)(z) = \frac{e^{-2t}}{1-e^{-2t}}\left(\frac{\mathrm{Cov}_{q_{t,z}}(X)}{1-e^{-2t}} - I_d\right),$$

from which we can deduce $\Delta \log\big(p_t^{2V}\big) \leq \frac{e^{-2t}}{(1-e^{-2t})^2}\mathbb{E}_{q_{t,z}}[\|Y\|^2]$ which yields again

$$\Delta \log\big(p_t^{2V}\big) \leq \frac{e^{-2t}}{(1-e^{-2t})^2}\left(\frac{\|z\|^2}{2a(1-e^{-2t})} + \frac{2b+d}{a}\right).$$

D Proof of Lemma 7

Before starting the proof, we recall the result of Mikulincer and Shenfeld [2023].

Proposition 14 *Let μ be a β-semi-log-convex probability distribution. Then, denoting p_t^V the distribution of the forward process in Eq. 1 and π the density of the standard Gaussian, it holds for all $z \in \mathbb{R}^d$ that*

$$\nabla^2 \log(\pi)(z) - \nabla^2 \log\big(p_t^V\big)(z) \preceq \frac{(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1} I_d \,. \tag{35}$$

Proof. Define the Ornstein-Uhlenbeck semi-group Q_t as

$$\begin{aligned} Q_t(g)(z) &= \int g(ze^{-t} + \sqrt{1-e^{-2t}}y)e^{-\frac{\|y\|^2}{2}}(2\pi)^{-d/2}\mathrm{d}y \\ &= \frac{1}{\sqrt{1-e^{-2t}}} \int g(u)e^{-\frac{\|u-e^{-t}z\|^2}{2(1-e^{-2t})}}(2\pi)^{-d/2}\mathrm{d}u \,, \end{aligned}$$

for all function g integrable w.r.t. the standard Gaussian measure. Taking g as $f = \frac{\mathrm{d}\mu}{\mathrm{d}\pi}$ with π the standard Gaussian, we obtain that

$$\begin{aligned} Q_t(f)(z) &= \frac{1}{Z_V\sqrt{1-e^{-2t}}} \int e^{-V(u)}e^{\frac{\|u\|^2}{2}}e^{-\frac{\|u-e^{-t}z\|^2}{2(1-e^{-2t})}}\mathrm{d}u \\ &= \frac{1}{Z_V\sqrt{1-e^{-2t}}} \int e^{-V(u)}e^{\frac{\|u\|^2(1-e^{-2t})-\|u\|^2+\langle z, ue^{-t}\rangle - e^{-2t}\|z\|^2}{2(1-e^{-2t})}}\mathrm{d}u \\ &= \frac{1}{Z_V\sqrt{1-e^{-2t}}} \int e^{-V(u)}e^{\frac{-\|ue^{-t}-z\|^2+\|z\|^2-e^{-2t}\|z\|^2}{2(1-e^{-2t})}}\mathrm{d}u \\ &= \frac{e^{\frac{\|z\|^2}{2}}}{Z_V\sqrt{1-e^{-2t}}} \int e^{-V(u)}e^{-\frac{\|ue^{-t}-z\|^2}{2(1-e^{-2t})}}\mathrm{d}u \,. \end{aligned}$$

In particular, we remark that $\nabla \log(Q_t(f)) = \nabla \log\big(p_t^V\big) - \nabla \log(\pi)$. Now, the quantity $\nabla \log(Q_t(f))$ was studied in Mikulincer and Shenfeld [2023] and they prove in Lemma 5 that for all z

$$\nabla^2 \log(Q_t(f))(z) \succeq \frac{(1-\beta)e^{-2t}}{(1-e^{-2t})(\beta-1)+1} I_d \,,$$

which is equivalent to

$$\nabla^2 \log(\pi)(z) - \nabla^2 \log\big(p_t^V\big)(z) \preceq \frac{(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1} I_d \,.$$

□

Before proving Lemma 7, we introduce this preliminary result on the evolution of Φ_t.

Lemma 15 (Evolution of the ratio) *Let $t > 0$, it holds that*

$$\partial_t \Phi_t = \Phi_t \left(\Delta \log(\Phi_t) - \langle \nabla \log(\Phi_t), \nabla \log(\pi)\rangle + \|\nabla \log(\Phi_t)\|^2 + 2\langle \nabla \log\big(p_t^V\big), \nabla \log(\Phi_t)\rangle\right) \,.$$

Proof. Recall that the log-density $\log\big(p_t^V\big)$ evolves as

$$\partial_t \log\big(p_t^V\big) = \Delta \log\big(p_t^V\big) + \|\nabla \log\big(p_t^V\big)\|^2 - \langle \nabla \log\big(p_t^V\big), \nabla \log(\pi)\rangle - \Delta \log(\pi) \,.$$

Hence, we deduce that $\log(\Phi_t)$ evolves as

$$\partial_t \log(\Phi_t) = \Delta \log(\Phi_t) - \langle \nabla \log(\Phi_t), \nabla \log(\pi)\rangle + \|\nabla \log\big(p_t^{2V}\big)\|^2 - \|\nabla \log\big(p_t^V\big)\|^2 \,.$$

The difference of quadratic terms can be expressed as

$$\begin{aligned} \|\nabla \log\big(p_t^{2V}\big)\|^2 - \|\nabla \log\big(p_t^V\big)\|^2 &= \|\nabla \log(\Phi_t) + \nabla \log\big(p_t^V\big)\|^2 - \|\nabla \log\big(p_t^V\big)\|^2 \\ &= \|\nabla \log(\Phi_t)\|^2 + 2\langle \nabla \log\big(p_t^V\big), \nabla \log(\Phi_t)\rangle \,, \end{aligned}$$

which allows to recover

$$\partial_t \log(\Phi_t) = \Delta \log(\Phi_t) - \langle \nabla \log(\Phi_t), \nabla \log(\pi)\rangle + \|\nabla \log(\Phi_t)\|^2 + 2\langle \nabla \log\big(p_t^V\big), \nabla \log(\Phi_t)\rangle \,.$$

□

We now provide the proof of Lemma 7.

Proof. Until the rest of the proof, the dependence on z of the integrand shall be implied unless expressed explicitly. We start by differentiating $m_0(\Phi_t)$ with respect to t:

$$\begin{aligned}\partial_t m_0(\Phi_t) &= \int \partial_t \Phi_t \mathrm{d}z \\ &= \int \Phi_t \left(\Delta \log(\Phi_t) - \langle \nabla \log(\Phi_t), \nabla \log(\pi)\rangle + \|\nabla \log(\Phi_t)\|^2 + 2\langle \nabla \log\big(p_t^V\big), \nabla \log(\Phi_t)\rangle\right) \mathrm{d}z\,,\end{aligned}$$

where we used Lemma 15 to compute $\partial_t \Phi_t$. Using integration by parts, the first term reads

$$\int \Phi_t \Delta \log(\Phi_t)\mathrm{d}z = -\int \langle \nabla \Phi_t, \nabla \log(\Phi_t)\rangle \mathrm{d}z = -\int \Phi_t \|\nabla \log(\Phi_t)\|^2 \mathrm{d}z,$$

hence the first and the third terms cancel and we recover

$$\partial_t m_0(\Phi_t) = 2\int \langle \nabla \log\big(p_t^V\big), \nabla \Phi_t\rangle \mathrm{d}z - \int \langle \nabla \log(\pi), \nabla \Phi_t\rangle \mathrm{d}z\,.$$

Using integration by parts again, we recover

$$\begin{aligned}\partial_t m_0(\Phi_t) &= \int \Delta \log(\pi)\Phi_t \mathrm{d}z - 2\int \Delta \log\big(p_t^V\big)\Phi_t \mathrm{d}z \\ &= 2\int (\Delta \log(\pi) - \Delta \log\big(p_t^V\big))\Phi_t \mathrm{d}z - \int \Delta \log(\pi)\Phi_t \mathrm{d}z \\ &= d m_0(\Phi_t) + 2\int (\Delta \log(\pi) - \Delta \log\big(p_t^V\big))\Phi_t \mathrm{d}z\,.\end{aligned}$$

Using Proposition 14, since μ is β-semi-log-convex, the term $(\Delta \log(\pi) - \Delta \log\big(p_t^V\big))$ can be upper-bounded uniformly by $\frac{d(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1}$ so we eventually get

$$\partial_t m_0(\Phi_t) \leq d m_0(\Phi_t)\left(1 + \frac{2(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1}\right)\,.$$

Hence we can use Gronwall's lemma which yields

$$m_0(\Phi_t) \leq m_0(\Phi_0)\exp\left(d\int_0^t \left(1 + \frac{2(\beta-1)e^{-2s}}{(1-e^{-2s})(\beta-1)+1}\right)\mathrm{d}s\right)\,.$$

Denoting by Z_V (resp. Z_{2V}) the normalizing constant of e^{-V} (resp. e^{-2V}), the term $m_0(\Phi_0)$ reads

$$m_0(\Phi_0) = \int \frac{p_0^{2V}(z)}{p_0^V(z)}\mathrm{d}z = \frac{Z_V}{Z_{2V}}\int \frac{e^{-2V(z)}}{e^{-V(z)}}\mathrm{d}z = \frac{(Z_V)^2}{Z_{2V}}\,.$$

Finally, let us compute the integral above. Making the change of variable $u = e^{-2s}(\beta-1)$ we have $\mathrm{d}u = -2(\beta-1)e^{-2s}\mathrm{d}s$ which yields

$$\begin{aligned}\int_0^t \frac{2(\beta-1)e^{-2s}}{(1-e^{-2s})(\beta-1)+1}\mathrm{d}s &= -\int_{\beta-1}^{(\beta-1)e^{-2t}} \frac{1}{\beta-u}\mathrm{d}u \\ &= [\log(\beta-u)]_{\beta-1}^{(\beta-1)e^{-2t}} \\ &= \log\big(\beta-(\beta-1)e^{-2t}\big) \\ &= \log\big(\beta(1-e^{-2t})+e^{-2t}\big)\,.\end{aligned}$$

Hence we recover

$$m_0(\Phi_t) \leq \frac{e^{td}(Z_V)^2}{Z_{2V}}(\beta(1-e^{-2t})+e^{-2t})^d\,.$$

$\square$

In order to recover a bound on the second moment of Φ_t, we need several intermediate results. We first prove that the maximum of the ratio decreases through time.

Lemma 16 (Decrease of the maximum of Φ) *The maximum of the ratio Φ_t decreases with t.*

Proof. Let z_t be a point where Φ_t attains its maximum and denote $M_t = \log(\Phi_t)(z_t)$. By the implicit function theorem, z_t is differentiable hence we can compute $\partial_t M_t$ as

$$\begin{aligned}\partial_t M_t &= \partial_t \log(\Phi_t)(z_t) + \langle \partial_t z_t, \nabla \log(\Phi_t)(z_t)\rangle \\ &= \Delta \log(\Phi_t)(z_t)\,.\end{aligned}$$

Since z_t is a maximum, we have in particular $\Delta \log \Phi_t(z_t) \leq 0$ which implies that M_t decreases. □

We then derive an upper-bound on the maximum of a log-smooth distribution.

Proposition 17 (Upper-bound of the maximum) *If μ is β-semi-log-convex then it holds that $\frac{\mathrm{d}\mu}{\mathrm{d}z} \leq \left(\frac{\beta}{2\pi}\right)^{\frac{d}{2}}$.*

Proof. Recall that the density of μ can be re-written as

$$\frac{\mathrm{d}\mu}{\mathrm{d}z} = \frac{e^{-(V(z)-V_*)}}{\int_z e^{-(V(z)-V_*)}\mathrm{d}z}\,,$$

where V_* the minimum of V attained for some z_*. By definition $e^{-(V(z)-V_*)} \leq 1$ for all z. Furthermore, since V verifies $\nabla^2 V \preceq \beta I_d$, we are ensured that

$$V(z) - V_* \leq \beta \frac{\|z - z_*\|^2}{2}\,,$$

which implies in particular that

$$\frac{1}{\int_z e^{-(V(z)-V_*)}\mathrm{d}z} \leq \left(\frac{\beta}{2\pi}\right)^{\frac{d}{2}}\,.$$

□

Using the previous result, we can derive an upper-bound on the integrated squared gradient at 0.

Lemma 18 (Upper-bound integrated gradient) *Let $\mu \propto e^{-V}$ be a β-semi-log-convex measure. Denoting $\mu(0)$ the density of μ with respect to the Lebesgue measure at 0, it holds that*

$$\int_0^t \|\nabla \log\big(p_s^V\big)\|^2(0)\mathrm{d}s \leq -\log\left(\mu(0)\right) + \frac{d}{2}\log\left(\frac{\beta}{2\pi}\right)\,.$$

Proof. Denoting π the density of the standard d dimensional Gaussian, recall that the density p_t^V evolves as

$$\partial_t p_t^V = \nabla \cdot \left(p_t^V \nabla \log\left(\frac{p_t^V}{\pi}\right)\right)\,,$$

which can also be re-written as

$$\begin{aligned}\partial_t \log\big(p_t^V\big) &= \frac{\Delta p_t^V}{p_t^V} - \langle \nabla \log\big(p_t^V\big), \nabla \log(\pi)\rangle - \Delta \log(\pi) \\ &= \Delta \log\big(p_t^V\big) + \|\nabla \log\big(p_t^V\big)\|^2 - \langle \nabla \log\big(p_t^V\big), \nabla \log(\pi)\rangle - \Delta \log(\pi)\,.\end{aligned}$$

In particular, for $z = 0$ this yields

$$\partial_t \log\big(p_t^V\big)(0) = \Delta \log\big(p_t^V\big)(0) + \|\nabla \log\big(p_t^V\big)\|^2(0) - \Delta \log(\pi)(0)\,,$$

which implies

$$\int_0^t \|\nabla \log\big(p_s^V\big)\|^2(0)\mathrm{d}s = \log\big(p_t^V\big)(0) - \log\big(p_0^V\big)(0) + \int_0^t \Delta \log(\pi)(0) - \Delta \log\big(p_t^V\big)(0)\mathrm{d}s\,.$$

Using the uniform upper-bound of Proposition 14, the second term can upper-bounded as

$$\begin{aligned}\int_0^t \Delta\log(\pi)(0) - \Delta\log\big(p_t^V\big)(0)\mathrm{d}s &\leq d\int_0^t \frac{(\beta-1)e^{-2s}}{(1-e^{-2s})(\beta-1)+1}\mathrm{d}s \\ &= \frac{d}{2}(\beta-1)\int_{e^{-2t}}^1 \frac{1}{\beta-u(\beta-1)}\mathrm{d}u \\ &= \frac{d}{2}\log\big(\beta(1-e^{-2t})+e^{-2t}\big)\,.\end{aligned}$$

Furthermore, Proposition 14 shows that $-\nabla^2\log\big(p_t^V\big) \preceq \frac{\beta}{\beta(1-e^{-2t})+e^{-2t}}$. Thus, using Proposition 17, we recover that $\log\big(p_t^V\big)(0) \leq \frac{d}{2}\log\left(\frac{\beta}{\beta(1-e^{-2t})+e^{-2t}}\right) - \frac{d}{2}\log(2\pi)$. In particular, we recover

$$\int_0^t \|\nabla\log\big(p_s^V\big)\|^2(0)\mathrm{d}s \leq -\log\left(\frac{\mathrm{d}\mu}{\mathrm{d}x}(0)\right) + \frac{d}{2}\log(\beta) - \frac{d}{2}\log(2\pi)\,.$$

□

We can now bound the first order moment of Φ_t

Lemma 19 *Let μ be a β-semi-log-convex measure with finite second moment m_2. It holds that*

$$m_1(\Phi_t) \leq \frac{(Z_V)^2}{Z_{2V}}e^{t(d+1)}(\beta(1-e^{-2t})+e^{-2t})^d\left(\sqrt{m_2}+\sqrt{-2\log\left(\mu(0)\right)+d\log\left(\frac{\beta}{2\pi}\right)}+2\sqrt{d\beta}\right)\,.$$

Proof. We differentiate $m_1(\Phi_t)$ and we recover bou

$$\partial_t m_1(\Phi_t) = \int \Phi_t\left(\Delta\log(\Phi_t) - \langle\nabla\log(\Phi_t),\nabla\log(\pi)\rangle + \|\nabla\log(\Phi_t)\|^2 + 2\langle\nabla\log\big(p_t^V\big),\nabla\log(\Phi_t)\rangle\right)\|z\|\mathrm{d}z\,.$$

Integration by parts of the first term yields

$$\int \Phi_t\Delta\log(\Phi_t)\|z\|\mathrm{d}z = -\int \Phi_t\|\nabla\log(\Phi_t)\|^2\|z\| + \langle\nabla\Phi_t, \frac{z}{\|z\|}\rangle\mathrm{d}z\,,$$

hence the squared gradients terms cancel and we recover

$$\partial_t m_1(\Phi_t) = \int\langle 2\nabla\log\big(p_t^V\big) - \nabla\log(\pi),\nabla\Phi_t\rangle\|z\|\mathrm{d}z - \int\langle\nabla\Phi_t,\frac{z}{\|z\|}\rangle\mathrm{d}z\,.$$

Let us denote by A the first term above. Integration by parts yields:

$$\begin{aligned}A &= \int \Phi_t(\Delta\log(\pi) - 2\Delta\log\big(p_t^V\big))\|z\|\mathrm{d}z + \int \Phi_t\langle\nabla\log(\pi) - 2\nabla\log\big(p_t^V\big),\frac{z}{\|z\|}\rangle\mathrm{d}z \\ &= -\int \Phi_t\Delta\log(\pi)\|z\|\mathrm{d}z + 2\int\Phi_t(\Delta\log(\pi) - \Delta\log\big(p_t^V\big))\|z\|\mathrm{d}z \\ &\quad - \int\Phi_t\langle\nabla\log(\pi),\frac{z}{\|z\|}\rangle\mathrm{d}z + 2\int\Phi_t\langle\nabla\log(\pi) - \nabla\log\big(p_t^V\big),\frac{z}{\|z\|}\rangle\mathrm{d}z \\ &= (d+1)m_1(\Phi_t) + 2\int\Phi_t(\Delta\log(\pi) - \Delta\log\big(p_t^V\big))\|z\|\mathrm{d}z + 2\int\Phi_t\langle\nabla\log(\pi) - \nabla\log\big(p_t^V\big),\frac{z}{\|z\|}\rangle\mathrm{d}z\,.\end{aligned}$$

Using the upper bound given in Proposition 14, we get $2\int\Phi_t(\Delta\log(\pi) - \Delta\log\big(p_t^V\big))\|z\|\mathrm{d}z \leq 2d\frac{(\beta-1)e^{-2t}}{(\beta-1)(1-e^{-2t})+1}m_1(\Phi_t)$. Similarly, we re-write the second term as

$$\begin{aligned}2\int\Phi_t\langle\nabla\log(\pi) - \nabla\log\big(p_t^V\big),\frac{z}{\|z\|}\rangle\mathrm{d}z &= 2\int\Phi_t\langle\nabla\log\left(\frac{\pi}{p_t^V}\right)(z) - \nabla\log\left(\frac{\pi}{p_t^V}\right)(0),\frac{z}{\|z\|}\rangle\mathrm{d}z \\ &\quad + 2\int\Phi_t\langle\nabla\log\left(\frac{\pi}{p_t^V}\right)(0),\frac{z}{\|z\|}\rangle\mathrm{d}z \\ &\leq 2m_1(\Phi_t)\frac{(\beta-1)e^{-2t}}{(\beta-1)(1-e^{-2t})+1} + 2\|\nabla\log\big(p_t^V\big)(0)\|m_0(\Phi_t)\,.\end{aligned}$$

Let us now handle the term $B = -\int \langle \nabla \Phi_t, \frac{z}{\|z\|} \rangle \mathrm{d}z$. In one dimension, $B = 2\Phi_t(0) \leq \max(\Phi_t)$ and for $d \geq 2$, we have

$$
\begin{aligned}
B &= \int \frac{\Phi_t(d-1)}{\|z\|} \mathrm{d}z \\
&= \int_{B_R} \frac{\Phi_t(d-1)}{\|z\|} \mathrm{d}z + \int_{\overline{B_R}} \frac{\Phi_t(d-1)}{\|z\|} \mathrm{d}z \\
&\leq \max(\Phi_t) \int_{B_R} \frac{d-1}{\|z\|} \mathrm{d}z + \frac{d-1}{R} m_0(\Phi_t) \\
&= \max(\Phi_t) \frac{2\pi^{d/2}}{\Gamma(d/2)} R^{d-1} + \frac{d-1}{R} m_0(\Phi_t)\,.
\end{aligned}
$$

Using Lemma 16 and Proposition 17, we have that $\max(\Phi_t) \leq \max(\Phi_0) = \frac{(Z_V)^2}{Z_{2V}} \max(\frac{\mathrm{d}\mu}{\mathrm{d}x}) \leq \frac{(Z_V)^2}{Z_{2V}} (\frac{\beta}{2\pi})^{d/2}$. Hence, if we pick $R = \left(Z_{2V} m_0(\Phi_t) \Gamma(d/2)/Z_V^2\right)^{1/d} \beta^{-1/2} 2^{1/2-1/d}$, we get as an upper-bound for B:

$$
\begin{aligned}
B &\leq \frac{Z_V^2}{Z_{2V}} d\sqrt{\beta} 2^{1/d-1/2} (m_0(\Phi_t) Z_{2V}/Z_V^2)^{\frac{d-1}{d}} \Gamma(d/2)^{-1/d} \\
&\leq \frac{2Z_V^2}{Z_{2V}} \sqrt{d\beta} (m_0(\Phi_t) Z_{2V}/Z_V^2)^{\frac{d-1}{d}}\,.
\end{aligned}
$$

In particular, we recover that

$$
\partial_t m_1(\Phi_t) \leq (d+1) \left(1 + 2\frac{(\beta-1)e^{-2t}}{(\beta-1)(1-e^{-2t})+1}\right) m_1(\Phi_t) + 2\|\nabla \log\left(p_t^V\right)(0)\| m_0(\Phi_t) + \frac{2Z_V^2}{Z_{2V}} \sqrt{\beta d} (m_0(\Phi_t) Z_{2V}/Z_V^2)^{\frac{d-1}{d}}\,,
$$

hence using Gronwall lemma, we have that

$$
\begin{aligned}
m_1(\Phi_t) &\leq e^{t(d+1)} (\beta(1-e^{-2t}) + e^{-2t})^{d+1} m_1(\Phi_0) \\
&\quad + 2\int_0^t \left[\|\nabla \log\left(p_s^V\right)(0)\| m_1(\Phi_s) + \sqrt{d\beta} (m_0(\Phi_s) Z_{2V}/Z_V^2)^{\frac{d-1}{d}}\right] \frac{e^{t(d+1)} (\beta(1-e^{-2t}) + e^{-2t})^{d+1}}{e^{s(d+1)} (\beta(1-e^{-2s}) + e^{-2s})^{d+1}} \mathrm{d}s\,.
\end{aligned}
$$

Using Lemma 7, we have that $m_0(\Phi_s) \leq \frac{Z_V^2}{Z_{2V}} e^{sd} (\beta(1-e^{-2s}) + e^{-2s})^d$ hence the first term of the integral is upper-bounded as:

$$
\int_0^t \frac{e^{-s(d+1)} m_0(\Phi_s) \|\nabla \log\left(p_s^V\right)(0)\|}{(\beta(1-e^{-2s}) + e^{-2s})^{d+1}} \mathrm{d}s \leq \frac{Z_V^2}{Z_{2V}} \int_0^t \|\nabla \log\left(p_s^V\right)(0)\| \frac{e^{-s}}{\beta(1-e^{-2s}) + e^{-2s}} \mathrm{d}s\,.
$$

By Cauchy-Schwarz it holds that

$$
\int_0^t \frac{\|\nabla \log\left(p_s^V\right)(0)\| e^{-s}}{L(1-e^{-2s}) + e^{-2s}} \mathrm{d}s \leq \sqrt{\int_0^t \|\nabla \log(p_s^V)(0)\|^2 \mathrm{d}s} \sqrt{\int_0^t \frac{e^{-2s}}{(\beta(1-e^{-2s}) + e^{-2s})^2} \mathrm{d}s}\,.
$$

The integral term is given by

$$
\begin{aligned}
\int_0^t \frac{e^{-2s}}{(\beta(1-e^{-2s}) + e^{-2s})^2} \mathrm{d}s &= \frac{1}{2} \int_{e^{-2t}}^1 \frac{1}{(\beta(1-u) + u)^2} \mathrm{d}u \\
&= \frac{1}{2(1-\beta)} [-\frac{1}{\beta(1-u) + u}]_{e^{-2t}}^1 \\
&= \frac{1}{2(1-\beta)} (\frac{1}{\beta(1-e^{-2t}) + e^{-2t}} - 1) \\
&= \frac{1 - e^{-2t}}{2(\beta(1-e^{-2t}) + e^{-2t})}
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
\int_0^t m_0(\Phi_s)^{\frac{d-1}{d}} \frac{e^{-s(d+1)}}{(\beta(1-e^{-2s}) + e^{-2s})^{d+1}} \mathrm{d}s &\leq \int_0^t \frac{e^{-2s}}{(\beta(1-e^{-2s}) + e^{-2s})^2} \mathrm{d}s \\
&= \frac{1 - e^{-2t}}{2(\beta(1-e^{-2t}) + e^{-2t})}
\end{aligned}
$$

Hence, we obtain

$$m_1(\Phi_t) \leq \frac{Z_V^2}{Z_{2V}} e^{t(d+1)}(\beta(1-e^{-2t})+e^{-2t})^{d+1} \left(\frac{Z_{2V}}{Z_V^2} m_1(\Phi_0) + \sqrt{-2\log(\mu(0)) + d\log\left(\frac{\beta}{2\pi}\right)} + \sqrt{d\beta} \right).$$

Finally, $m_1(\Phi_0) = \frac{(Z_V)^2}{Z_{2V}} \int \|z\| \frac{e^{-V(z)}}{\int e^{-V}\mathrm{d}z}\mathrm{d}z = \frac{(Z_V)^2}{Z_{2V}} \sqrt{m_2(\mu)}$. $\square$

We can derive our upper-bound on $m_2(\Phi_t)$.

Proof. We start by differentiating $m_2(\Phi_t)$:

$$\begin{aligned}
\partial_t m_2(\Phi_t) &= \int \|z\|^2 \partial_t \Phi_t \mathrm{d}z \\
&= \int \|z\|^2 (\mathrm{div}(\nabla\Phi_t) - \langle \nabla\Phi_t, \nabla\log(\pi)\rangle + 2\langle \nabla \log(p_t^V), \nabla\Phi_t\rangle)\mathrm{d}z \\
&= -\int 2\langle z, \nabla\Phi_t\rangle \mathrm{d}z + \int (\Delta\log(\pi) - 2\Delta\log(p_t^V))\|z\|^2\Phi_t\mathrm{d}z + 2\int \langle \nabla\log(\pi) - 2\log(p_t^V), z\rangle \Phi_t \mathrm{d}z \\
&= -2\int \Delta\log(\pi)\Phi_t\mathrm{d}z - \int \Delta\log(\pi)\|z\|^2\Phi_t\mathrm{d}z - 2\int \langle\nabla\log(\pi), z\rangle\Phi_t\mathrm{d}z \\
&\quad + 2\int (\Delta\log(\pi) - \Delta\log(p_t^V))\|z\|^2\Phi_t\mathrm{d}z + 4\int\langle \nabla\log(\pi) - \nabla\log(p_t^V), z\rangle\Phi_t\mathrm{d}z \\
&= 2dm_0(\Phi_t) + dm_2(\Phi_t) + 2m_2(\Phi_t) + 2\int (\Delta\log(\pi) - \Delta\log(p_t^V))\|z\|^2\Phi_t\mathrm{d}z + 4\int \langle \nabla\log(\pi) - \nabla\log(p_t^V), z\rangle
\end{aligned}$$

The first term $\int (\Delta\log(\pi) - \Delta\log(p_t^V))\|z\|^2\Phi_t\mathrm{d}z$ is upper-bounded by $\frac{(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1} dm_2(\Phi_t)$ and for the second term we have

$$\begin{aligned}
\int \langle \nabla\log(\pi) - \nabla\log(p_t^V), z\rangle\Phi_t\mathrm{d}z &= \int \left\langle \nabla\log\left(\frac{\pi}{p_t^V}\right)(z) - \nabla\log\left(\frac{\pi}{p_t^V}\right)(0), z\right\rangle\Phi_t\mathrm{d}z \\
&\quad + \int \left\langle \nabla\log\left(\frac{\pi}{p_t^V}\right)(0), z\right\rangle\Phi_t\mathrm{d}z \\
&\leq \int \frac{(L-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1}\|z\|^2\Phi_t\mathrm{d}z + \|\log(p_t^V)(0)\| \int \|z\|\Phi_t\mathrm{d}z \\
&= \frac{(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1} m_2(\Phi_t) + \|\nabla\log(p_t^V)(0)\| m_1(\Phi_t).
\end{aligned}$$

Hence we recover

$$\partial_t m_2(\Phi_t) \leq (d+2)\left(1 + \frac{2(\beta-1)e^{-2t}}{(1-e^{-2t})(\beta-1)+1}\right) m_2(\Phi_t) + 4\|\nabla\log(p_t^V)(0)\| m_1(\Phi_t) + 2dm_0(\Phi_t).$$

We now use the Gronwall lemma to obtain

$$m_2(\Phi_t) \leq e^{t(d+2)}(\beta(1-e^{-2t})+e^{-2t})^{d+2}\left(m_2(\Phi_0) + 2\int_0^t \frac{e^{-s(d+2)}(2m_1(\Phi_s)\|\nabla\log(p_s^V)(0)\| + dm_0(\Phi_s))}{(\beta(1-e^{-2s})+e^{-2s})^{(d+2)}}\mathrm{d}s\right)$$

Recalling the upper-bound $m_1(\Phi_t) \leq e^{t(d+1)}(\beta(1-e^{-2t})+e^{-2t})^{d+1}C$ where C is defined in Lemma 19, we upper-bound the integral term as

$$\begin{aligned}
\int_0^t \frac{e^{-s(d+2)}m_1(\Phi_s)\|\nabla\log(p_s^V)(0)\|}{(\beta(1-e^{-2s})+e^{-2s})^{(d+2)}}\mathrm{d}s &\leq C\int_0^t \frac{e^{-s}}{\beta(1-e^{-2s})+e^{-2s})}\|\nabla\log(p_s^V)(0)\|\mathrm{d}s \\
&\leq C\sqrt{\int_0^t \frac{e^{-2s}}{(\beta(1-e^{-2s})+e^{-2s})^2}\mathrm{d}s}\sqrt{\int_0^t \|\nabla\log(p_t^V)(0)\|^2\mathrm{d}s} \\
&= C\sqrt{\frac{1-e^{-2t}}{2}}\sqrt{\int_0^t \|\nabla\log(p_t^V)(0)\|^2\mathrm{d}s}.
\end{aligned}$$

Similarly,

$$\int_0^t \frac{e^{-s(d+2)}}{(\beta(1-e^{-2s})+e^{-2s})^{(d+2)}} m_0(\Phi_s)\mathrm{d}s \leq \int_0^t \frac{e^{-2s}}{(\beta(1-e^{-2s})+e^{-2s})^2}$$
$$= \frac{1-e^{-2t}}{2}\,.$$

Hence we recover that

$$m_2(\Phi_t) \leq e^{t(d+2)}(\beta(1-e^{-2t})+e^{-2t})^{d+2}\left(m_2(\Phi_0)+2C\sqrt{-2\log\left(\mu(0)\right)+d\log\left(\frac{\beta}{2\pi}\right)}+d\right)\,.$$

Using the expression of C, we recover eventually that

$$m_2(\Phi_t) \leq \frac{2e^{t(d+2)}Z_V^2}{Z_{2V}}(\beta(1-e^{-2t})+e^{-2t})^{d+2}\left[m_2(\mu)+d(\beta+1)-4\log\left(\mu(0)\right)+2d\log\left(\frac{\beta}{2\pi}\right)\right]\,.$$

□

E Proof of Theorem 9

Proof. As Proposition 6 shows, the average error of the estimator can be upper-bounded as

$$\mathbb{E}[\|\hat{s}_{t,n}(z)-\nabla\log\left(p_t^V\right)(z)\|^2] \lesssim \frac{d\log(n)e^{td}Z_{2V}p_t^{2V}(z)}{na(Z_V)^2(p_t^V(z))^2}\left[\frac{\|z\|^2}{(1-e^{-2t})^4}+\frac{b+d}{(1-e^{-2t})^3}\right]\,.$$

Hence, the average integrated error reads

$$\mathbb{E}[\int\|\hat{s}_{t,n}(z)-\nabla\log\left(p_t^V\right)(z)\|^2\mathrm{d}p_t(z)] \lesssim \frac{d\beta^{d+2}\log(n)e^{2t(d+1)}}{na(1-e^{-2t})^4}\left[\frac{Z_{2V}m_2(\Phi_t)e^{-t(d+2)}}{(Z_V)^2}+\frac{Z_{2V}(b+d)m_0(\Phi_t)e^{-t(d+2)}}{(Z_V)^2}\right]\,.$$

We then apply the upper-bounds in Lemma 7 and in Lemma 8 and we recover

$$\mathbb{E}[\int\|\hat{s}_{t,n}(z)-\nabla\log\left(p_t^V\right)(z)\|^2\mathrm{d}p_t(z)] \lesssim \frac{d\beta^{d+3}\log(n)e^{2t(d+1)}(b+d)}{na^2(1-e^{-2t})^4}$$

We thus set n as $n=d^2\max(\epsilon^{-2(d+1)+1},\epsilon^{-5})$ and we eventually get for $\epsilon\leq t\leq\log(1/\epsilon)$ that

$$\mathbb{E}[\int\|\hat{s}_{t,n}(z)-\nabla\log\left(p_t^V\right)(z)\|^2\mathrm{d}p_t(z)] \lesssim \frac{\epsilon\beta^{d+3}(b+d)}{a^2}\,.$$

Hence, plugging again the bounds of Lemma 8 in Theorem 3, we recover the desired result. □