

POLYNOMIAL CONVERGENCE OF RIEMANNIAN DIFFUSION MODELS

Xingyu Xu, Ziyi Zhang, Yorie Nakahira, Guannan Qu, Yuejie Chi

Department of Electrical and Computer Engineering

Carnegie Mellon University

Pittsburgh, PA 15213, USA

{xingyuxu, ziyizhan, ynakahir, gqu, yuejiec}@andrew.cmu.edu

ABSTRACT

Diffusion models have demonstrated remarkable empirical success in recent years and are considered one of the state-of-the-art generative models in modern AI. These models consist of a forward process, which gradually diffuses the data distribution to a noise distribution spanning the whole space, and a backward process, which inverts this transformation to recover the data distribution from noise. Most of the existing literature assumes that the underlying space is Euclidean. However, in many practical applications, the data are constrained to lie on a submanifold of Euclidean space. Addressing this setting, [De Bortoli et al. \(2022\)](#) introduced Riemannian diffusion models and proved that using an exponentially small step size yields a small sampling error in the Wasserstein distance, provided the data distribution is smooth and strictly positive, and the score estimate is L_∞ -accurate. In this paper, we greatly strengthen this theory by establishing that, under L_2 -accurate score estimate, a *polynomially small stepsize* suffices to guarantee small sampling error in the total variation distance, without requiring smoothness or positivity of the data distribution. Our analysis only requires mild and standard curvature assumptions on the underlying manifold. The main ingredients in our analysis are Li-Yau estimate for the log-gradient of heat kernel, and Minakshisundaram-Pleijel parametrix expansion of the perturbed heat equation. Our approach opens the door to a sharper analysis of diffusion models on non-Euclidean spaces.

1 INTRODUCTION

Initially introduced by [Sohl-Dickstein et al. \(2015\)](#) and later advanced by [Song & Ermon \(2019\)](#); [Ho et al. \(2020\)](#); [Dhariwal & Nichol \(2021\)](#), diffusion model has become one of the bedrocks in generative modeling across a variety of application domains such as vision, video, speech, and many others. On a high level, diffusion models generate samples from a target distribution by operating on two stochastic processes:

1. A *forward* process

$$X_0 \xrightarrow{\text{add noise}} X_1 \xrightarrow{\text{add noise}} \dots \xrightarrow{\text{add noise}} X_T,$$

where X_0 is sampled from the target distribution p_0 in \mathbb{R}^d , and X_T resembles pure noise.

2. A *reverse* process

$$Y_T \xrightarrow{\text{denoise}} Y_{T-1} \xrightarrow{\text{denoise}} \dots \xrightarrow{\text{denoise}} Y_0,$$

where Y_T starts from pure noise, and gradually removes the noise, so that at Y_0 , we recover a new sample from a distribution close to p_0 .

The reverse process is built to recover the target data distribution by step-wise reversing the forward process, with a goal of matching the probabilities $Y_t \approx X_t$ in distribution for $t \in \{T, \dots, 1\}$. Leveraging the theory of backward stochastic differential equations (SDE) ([Anderson, 1982](#); [Haussmann & Pardoux, 1986](#)), this can be formally achieved as soon as the as long as the *score function*, i.e.,

the log-gradient of the the marginal density of the forward process, becomes available, which can be estimated via score matching (Hyvärinen & Dayan, 2005).

Tremendous recent progresses have been made in understanding the convergence of diffusion models in the Euclidean space, e.g. Lee et al. (2023); Chen et al. (2023); Benton et al. (2024); Li et al. (2024), which establish near-tight polynomial iteration complexities of discrete-time samplers under L_2 -accurate score estimates and mild assumptions of the data distribution. These convergence results provide strong justifications to the empirical success of diffusion models for generating from complex multi-modal distributions.

Many scientific domains, however, are intrinsically *non-Euclidean*; examples include orientations on $\text{SO}(3)$, directions on spheres, toroidal angles, articulated poses, and symmetric positive definite (SPD) matrices are naturally modeled on Riemannian manifolds (Piggott & Solo, 2016; Muniz et al., 2022). Recently, there has been an increasing interest in effectively sampling from distributions supported on manifolds and providing theoretical guarantees (Girolami & Calderhead, 2011; Gutmry & Vempala, 2022; Li & Erdogdu, 2023; Guan et al., 2025). Although sampling on manifolds has been studied extensively (Cheng et al., 2023), extending diffusion models to manifolds requires careful treatments to incorporate the manifold constraints into both the time-inhomogeneous forward and reverse processes, with selected attempts in De Bortoli et al. (2022); Huang et al. (2022); Lou et al. (2023); Liu et al. (2023); Fishman et al. (2023).

One notable development is De Bortoli et al. (2022), who introduced *Riemannian Score-Based Generative Models* (RSGMs) with convergence guarantees in the Wasserstein distance. Specifically, they established a time-reversal diffusion process for geometric Brownian motion on manifolds, which can be similarly learned via score matching (Hyvärinen & Dayan, 2005). While groundbreaking, their convergence bound suffers from a few caveats: 1) it requires an exponentially small stepsize, leading to a possibly exponential iteration complexity in some of the manifold parameters; 2) it requires L_∞ -accurate score estimates, which are impractical in deep learning; and 3) the data distribution is required to be smooth and strictly positive on compact manifolds. This naturally raises the following questions:

Can we achieve polynomial iteration complexity for manifold diffusion models using L_2 -accurate score estimates under mild data assumptions?

1.1 OUR CONTRIBUTION

We provide a discrete-time analysis of the RSGM sampler in De Bortoli et al. (2022), assuming L_2 -accurate score estimates. Under mild geometric conditions of the manifold without assuming smooth or strictly positive data densities, we establish that *polynomial* stepsizes suffice for accurate sampling on manifolds in *total variation* (TV). More precisely, for some $\varepsilon > 0$, the TV error between the distribution of the output Y_0 and p_δ , that is, an approximation to p_0 with early-stopping time $\delta > 0$, obeys

$$\text{TV}(p_\delta, \text{Law}(Y_0)) \lesssim \varepsilon + \varepsilon_{\text{score}} + \sqrt{hT} \text{poly}(d, \delta^{-1}),$$

as long as the horizon satisfies $T \gtrsim \lambda_1^{-1}(d \log d + \log(d/\varepsilon))$. Here, d is the dimension of \mathcal{M} , λ_1 is the spectral gap of $-\Delta_M$, $h > 0$ is the stepsize, and $\varepsilon_{\text{score}}$ is the L_2 score estimation error.¹ This bound suggests that a polynomial stepsize is sufficient: take $T \asymp \lambda_1^{-1}(d \log d + \log(d/\varepsilon))$ and $h = \frac{\varepsilon^2}{\text{poly}(d, \delta^{-1})T}$, then the TV error is bounded by $\varepsilon + \varepsilon_{\text{score}}$ after an iteration complexity of

$$N = T/h \asymp \text{poly}(d, \delta^{-1})/(\lambda_1 \varepsilon)^2.$$

This conveys a much more benign message about the efficiency of Riemannian diffusion models, compared with the iteration complexity in De Bortoli et al. (2022) that scales exponentially with the dimension d , under relaxed assumptions on both the data distribution and the score estimates.

Techniques. Our proof highlights three ingredients: (i) high-probability Li-Yau gradient bounds for the manifold heat kernel together with early stopping to control $\|\nabla \log p_t\|$ without assuming positivity/smoothness of p_0 ; (ii) a localization scheme that “freezes” drifts across nearby tangent

¹For simplicity, we omitted polynomial dependence on manifold geometry parameters and poly-log dependence on ε . The complete version can be found in Theorem 1.

Work	Structure	Metric	Iteration complexity	Data distribution
Benton et al. (2024)	Euclidean	TV	$\tilde{O}(d/\varepsilon^2)$	bounded moment
Li et al. (2024)	Euclidean	TV	$\tilde{O}(\text{poly}(d)/\varepsilon)$	bounded support
Li & Yan (2025)	Euclidean	TV	$\tilde{O}(d/\varepsilon)$	bounded moment
De Bortoli et al. (2022)	Manifold	W_p	$\tilde{O}(\exp(O(d))/\varepsilon^{-1/\lambda_1})^{\textcolor{red}{a}}$	smooth, strictly positive
This work	Manifold	TV	$\tilde{O}\left(\frac{\text{poly}(d)}{\lambda_1^2 \varepsilon^2}\right)$	None (early stopping)

^aThe original W_p error in De Bortoli et al. (2022) is stated in the form of $\tilde{O}(Ce^{-\lambda_1 T} + e^T \sqrt{h})$, where C is defined in Proposition C.6 therein, which in turn is specified by Urakawa (2006, Proposition 2.6) as the supremum of $t^{d/2} H(t, x, y)$, where H is the heat kernel on the manifold. In general, the best estimate for this is due to Li-Yau (Li & Yau, 1986), which gives $C \leq e^{O(d)}$. To achieve ε -error, we must set $T = \lambda_1^{-1}(\Omega(d) + \log \varepsilon^{-1})$ and $h = e^{-2T} \varepsilon^2$, then the iteration complexity T/h has the claimed form.

Table 1: Comparison of the current theoretical guarantees on diffusion probabilistic models on Euclidean spaces and manifolds. Here, $\lambda_1 > 0$ is the spectral gap of the Laplace-Beltrami operator.

spaces but preserves continuous Brownian motion (BM), to separate the effects of discretizing scores and BM; and (iii) a quantitative estimates for Minakshisundaram–Pleijel parametrix that controls one-step deviations between the manifold heat flow and its discretized proxy. These components allow us to handle the discretization errors sharply to avoid exponential dependence.

1.2 RELATED WORKS

Non-asymptotic convergence for Euclidean diffusion models. Early convergence analyses of diffusion models require L_∞ -accurate score estimates (De Bortoli et al., 2021). For stochastic samplers such as DDPM (Ho et al., 2020), early bounds under Lipschitz/smoothness assumptions of the data distribution admit an $O(T^{-\frac{1}{2}})$ iteration complexity in the total variation distance assuming L_2 -accurate score estimates (Chen et al., 2023), with subsequent analyses relaxing the Lipschitz assumption yet retaining the same complexity (Lee et al., 2023; Benton et al., 2024; Li et al., 2024). More recently, Li & Yan (2025) has improved the iteration complexity to $\tilde{O}(T^{-1})$. For deterministic samplers, Chen et al. (2023) established polynomial convergence with exact scores, and Li et al. (2024) established a convergence rate of $O(T^{-1})$ under L_2 -accurate scores. See Beyler & Bach (2025); Liang et al. (2024); Li & Jiao (2024) for additional analyses that established convergence in the Wasserstein distance and improved discrete-time rates. Several works (Li & Yan, 2024; Liang et al., 2025; Huang et al., 2024; Potapchik et al., 2024) also developed non-asymptotic convergence rates of diffusion models under the manifold hypothesis, suggesting diffusion models are adaptive to low-dimensional structures. This line of work should not be confused with ours, where the diffusion process is designed specifically to be constrained on the manifold.

Sampling on Riemannian manifold. Cheng et al. (2022; 2023) analyzed the geometric Euler–Maruyama (EM) discretization for time-homogeneous SDEs, and proved a polynomial complexity guarantee under dissipative-distant geometric assumptions on the manifold. See also Bharath et al. (2025) for follow-ups. Guan et al. (2025) proposed a Riemannian proximal sampler with convergence guarantees under the log-Sobolev inequality. Various sampling algorithms are also studied for a related problem known as sampling from constrained spaces (Srinivasan et al., 2024; Ahn & Chewi, 2021). Nonetheless, convergence analyses of Riemannian diffusion models under general data distributions remain highly limited, with De Bortoli et al. (2022) being the only prior work with non-asymptotic convergence rates.

2 BACKGROUNDS

2.1 DIFFUSION MODELS ON EUCLIDEAN SPACE

We briefly recall diffusion processes on \mathbb{R}^d . Let $(W_t)_{t \geq 0}$ be a standard Brownian motion in \mathbb{R}^d .

Forward SDE and Fokker–Planck. Given a drift term $b_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$, the forward process $(X_t)_{t \in [0, T]}$ solves the Itô SDE

$$dX_t = b_t(X_t) dt + dW_t, \quad X_0 \sim p_0.$$

Let p_t denote the law of X_t , then p_t satisfies the Fokker–Planck equation $\partial_t p_t = -\nabla(b_t p_t) + \frac{1}{2} \Delta p_t$. In the driftless setting where $b_t \equiv 0$, the marginal $p_t = p_0 * \varphi_t$ is a Gaussian smoothing of p_0 with kernel $\varphi_t(z) = (2\pi t)^{-d/2} \exp(-\|z\|^2/(2t))$.

Score and reverse process. The *score* of the forward process at time t is defined as $s_t(x) := \nabla \log p_t(x)$. The time-reversal identity (Anderson, 1982) yields a reverse-time process $(Y_t)_{t \in [0, T]}$ whose marginals match those of $(X_t)_{t \in [0, T]}$ which solves the reverse-time SDE (note that t flows from T to 0):

$$dY_\tau = [-b_\tau(Y_\tau) + s_\tau(Y_\tau)] dt + dW_t, \quad \tau = T - t, \quad Y_T \sim p_{X_T}.$$

Discretization with approximate score. Let $0 = t_0 < t_1 < \dots < t_N = T$, where $t_i - t_{i-1} =: h$ be a time grid. In practice, exact score function is often unavailable. Instead, we use an approximation \hat{s}_t trained via score matching (Hyvärinen & Dayan, 2005). The Euler–Maruyama discretization of the reverse-time SDE above is given by

$$y_{k-1} = y_k + h[-b_{t_k}(y_k) + \hat{s}_{t_k}(y_k)] + \sqrt{h}g_k, \quad g_k \sim \mathcal{N}(0, I_d).$$

In our driftless setting, this reduces to

$$y_{k-1} = y_k + h\hat{s}_{t_k}(y_k) + \sqrt{h}g_k, \quad g_k \sim \mathcal{N}(0, I_d).$$

2.2 GEOMETRY AND NOTATION

We assume some familiarity with Riemannian geometry, and make use of standard notation. Please refer to Jost (2017); Petersen (2006) for a more in-depth treatment. In particular, we use $\alpha, \beta, \xi, \zeta$, etc., to index *coordinate representation* of tensors, and assume Einstein’s summation convention. Let (\mathcal{M}, g) be a connected, compact d -dimensional Riemannian manifold, with geodesic distance $\rho(\cdot, \cdot)$ and volume measure μ . We assume $\mu(\mathcal{M}) = 1$. The Levi–Civita connection is denoted by ∇ , and the Laplace–Beltrami operator by

$$\Delta_{\mathcal{M}} f := \nabla_\alpha \nabla^\alpha f.$$

We use $T_x \mathcal{M}$ for the tangent space at x and use $\exp_x : T_x \mathcal{M} \rightarrow \mathcal{M}$ for the exponential map and \log_x for its local inverse on the normal neighborhood of x . Furthermore, we define *geodesic diameter* of (\mathcal{M}, g) is

$$\text{Diam}(\mathcal{M}) := \sup_{x, y \in \mathcal{M}} \rho(x, y),$$

where $\rho(\cdot, \cdot)$ is the geodesic distance induced by g . We further denote Rm as the Riemannian curvature tensor. Geodesic ball centered at x with radius r is denoted $B_x(r)$.

We use the *total variation* (TV) and the *Kullback–Leibler* (KL) distance to measure the discrepancy between two distributions p, q :

$$\text{TV}(p, q) = \int_{\mathcal{M}} |dp - dq|, \quad \text{KL}(p \parallel q) = \int_{\mathcal{M}} \left(\log \frac{dp}{dq} \right) dp.$$

2.3 HEAT FLOW, BROWNIAN MOTION, AND DIFFUSION ON \mathcal{M}

We also recall the setup for SDE and diffusion processes on Riemannian manifolds introduced in De Bortoli et al. (2022); Cheng et al. (2023). Let $(W_t)_{t \geq 0}$ be a standard Brownian motion in \mathbb{R}^d and $U_x : \mathbb{R}^d \rightarrow T_x \mathcal{M}$ any orthonormal frame at x . The *Geometric Brownian motion* solves

$$dX_t = U_{X_t} \circ dW_t,$$

where \circ denotes Stratonovich integral, and its transition density $p_t(x, y)$ with respect to μ solves the heat equation

$$\partial_t p_t(\cdot, y) = \frac{1}{2} \Delta_{\mathcal{M}} p_t(\cdot, y).$$

Algorithm 1 Riemannian Score-Based Generative Models (RSGM)

```

1: Manifold  $(\mathcal{M}, g)$ ; score  $\hat{s}_t(x)$ ; early stopping time  $\delta > 0$ ; reverse time grid  $\delta = t_0 < t_1 < \dots < t_N = T$ ; step size  $h = t_k - t_{k-1}$ ; initial  $x_N \sim \mu$  (uniform distribution);
2: for  $k \in \{N, \dots, 1, 0\}$  do
3:   Choose an orthonormal frame  $U_k$  at  $Y_k$ , which is a linear map from  $\mathbb{R}^d$  to  $T_{Y_k} \mathcal{M}$ .
4:    $\xi_k \sim \mathcal{N}(0, I_d)$  in  $\mathbb{R}^d$ ;  $G_k \leftarrow U_k \xi_k \in T_{Y_k} \mathcal{M}$ .
5:    $b_k \leftarrow \hat{s}_{t_k}(Y_k) \in T_{Y_k} \mathcal{M}$ 
6:    $\Delta_k \leftarrow h b_k + \sqrt{h} G_k \in T_{Y_k} \mathcal{M}$ 
7:   if  $\|\Delta_k\| \leq h^{1/4}$  then
8:      $Y_{k-1} \leftarrow \exp_{Y_k}(\Delta_k)$ 
9:   else
10:     $Y_{k-1} \sim \mu$ 
11: return  $Y_0$ 

```

Equivalently, Brownian motion can be defined abstractly as the solution to the martingale problem for the operator $\frac{1}{2}\Delta_{\mathcal{M}}$. Concretely, for any $f \in C^\infty([0, \infty) \times \mathcal{M})$, the process

$$M_t^f := f(t, X_t) - f(0, X_0) - \int_0^t \left(\partial_s + \frac{1}{2}\Delta_{\mathcal{M}} \right) f(s, X_s) ds$$

is a martingale with respect to the natural filtration of X . More generally, a forward diffusion process with drift is given by

$$dX_t = b_t(X_t) dt + U_{X_t} \circ dW_t,$$

with Fokker–Planck equation $\partial_t p_t = -\nabla(b_t p_t) + \frac{1}{2}\Delta_{\mathcal{M}} p_t$. Note that in this setting, the following process is a martingale for smooth f :

$$M_t^f := f(t, X_t) - f(0, X_0) - \int_0^t \left(\partial_s f + \langle b_t, \nabla f \rangle + \frac{1}{2}\Delta_{\mathcal{M}} f \right) (s, X_s) ds. \quad (1)$$

Let p_t denote the density of X_t w.r.t. μ , and define the *score* $s_t := \nabla \log p_t$. The time-reversal identity on manifolds yields a reverse SDE:

$$d\tilde{X}_\tau = (-b_t(\tilde{X}_\tau) + \nabla \log p_\tau(\tilde{X}_\tau)) dt + U_{\tilde{X}_\tau} \circ dW_t, \quad \tau = T - t, \quad \tilde{X}_T \sim p_{X_T}.$$

In practice, the score $\nabla \log p_t$ is approximated by a trained neural network $\hat{s}_t(x)$.

Last but not least, note that on compact manifolds, $-\Delta_{\mathcal{M}}$ admits a spectral gap $\lambda_1 > 0$. Any initial distribution mixes to the uniform distribution μ along the heat flow with rate $e^{-\lambda_1 t}$.

3 MAIN RESULT

In this section, for completeness, we first introduce the RSGM algorithm in [De Bortoli et al. \(2022\)](#). Then, we offer our polynomial convergence guarantee in [Theorem 1](#). For simplicity, we use a driftless forward process:

$$dX_t = U_{X_t} \circ dW_t, \quad X_0 \sim p_0.$$

The time-reversal identity yields the *reverse-time SDE*

$$dY_t = \nabla \log p_\tau(Y_\tau) dt + U_{Y_\tau} \circ dW_t, \quad \tau = T - t, \quad Y_T \sim p_T. \quad (2)$$

In [Algorithm 1](#), we provide an outline of discretized reverse-time SDE on Riemannian manifold, modified from [De Bortoli et al. \(2022\)](#). In each reverse step $k \in \{N, \dots, 1, 0\}$, we select an orthonormal frame U_k at y_k , then sample Gaussian noise ξ_k and lift it to the tangent space $T_{y_k} \mathcal{M}$ using the orthonormal frame, obtaining $G_k \in T_{y_k} \mathcal{M}$. Afterwards, we propose a tangent update $\Delta_k = h \hat{s}_{t_k}(y_k) + \sqrt{h} G_k$ and the project to the manifold using the exponential map. To prevent the update from exiting the injective radius, we perform a rejection sampling step that rejects exceedingly large update. The algorithm terminates at $k = 0$ and returns the final iterate y_0 . In this way, we ensure every update is well-defined in normal coordinates during the algorithm.

Before presenting the main theorem, we formalize the assumptions needed for the convergence guarantee.

Assumption 1 (Regularity). *Let (\mathcal{M}, g) be a connected, compact d -dimensional Riemannian manifold. We assume the following conditions on \mathcal{M} :*

(A1) **Positive injectivity radius:** *there exists some constant $K \geq 1$ such that the injective radius $\geq 1/K$.*

(A2) **Uniform curvature bounds:** *for the same constant K (which can be enlarged if necessary), we have*

$$\max \left\{ \text{Diam}(\mathcal{M}), \|\text{Rm}\|_{L^\infty}, \|\nabla \text{Rm}\|_{L^\infty}, \|\nabla^2 \text{Rm}\|_{L^\infty} \right\} \leq K.$$

(A3) **Regularity of score estimates:** *there exists a polynomial $\text{poly}(d, K)$, such that*

$$\|\widehat{s}_{t_k}(x)\| \leq \text{poly}(d, K) (\|\nabla \log p_{t_k}(x)\| + t_k^{-1}), \quad \forall x \in \mathcal{M}.$$

In Assumption 1, we made the standard ‘‘bounded geometry’’ assumption; similar assumptions also occur in [Cheng et al. \(2022\)](#); [De Bortoli et al. \(2022\)](#). A positive injective radius ensures that we have sufficient room to operate on the tangent spaces as a proxy of operating on manifolds, since for every $x \in \mathcal{M}$, the exponential map \exp_x is a diffeomorphism on the geodesic ball within injective radius. Bounds on Riemannian tensors rule out pathological cases, which helps to control the error propagation along the reverse diffusion. Lastly, compactness ensures a positive spectral gap of $\Delta_{\mathcal{M}}$ with $\lambda_1 > 0$, which is necessary to guarantee that the forward process mixes. The mild assumption (A3) on the score estimates avoids excessively large drifts in diffusion, and can be implemented easily in practice by clipping. In addition to the above, we also need a standard assumption on the score estimation error ([Chen et al., 2023](#)).

Assumption 2 (Score estimation error). *There exists $\varepsilon_{\text{score}} > 0$ such that*

$$\sum_{k=1}^N (t_k - t_{k-1}) \mathbb{E} \|\widehat{s}_{t_k}(Y_{t_k}) - \nabla \log p_{t_k}(Y_{t_k})\|^2 \leq \varepsilon_{\text{score}}^2.$$

With the above assumptions, we are now ready to present our main convergence guarantee for RSGM, as outlined in the following TV-accuracy bound.

Theorem 1. *Assume Assumptions 1 and 2 hold. There exists some universal constant $C, C' > 0$ such that the following holds. If $T \geq \frac{C}{\lambda_1} (d \log(Kd) + K + \log(\frac{N}{\varepsilon}))$, then the output Y_0 of Algorithm 1 obeys*

$$\text{TV}(p_\delta, \text{Law}(Y_0)) \leq \varepsilon + C' \varepsilon_{\text{score}} + \sqrt{hT} \text{poly}(d, K, \delta^{-1}),$$

where h is the discretization step size, $\lambda_1 > 0$ is the mixing rate of the geometric Brownian Motion on \mathcal{M} , i.e., the smallest eigenvalue of $-\Delta_{\mathcal{M}}$ in $L^2(\mu)$.

A few remarks are in order.

Iteration complexity. The error bound decomposes cleanly into three terms: ε results from mixing of the heat semigroup at the spectral gap λ_1 , $\varepsilon_{\text{score}}$ captures error from imperfect score estimation, and $\sqrt{hT} \text{poly}(d, K, \delta^{-1})$ is the discretization error controlled by the step size and curvature. Consequently, choosing $T \asymp \lambda_1^{-1} (d \log d + \log(d/\varepsilon))$ and $h = \frac{\varepsilon^2}{\text{poly}(d, K, \delta^{-1})T}$, then the TV error is bounded by $\varepsilon + \varepsilon_{\text{score}}$ after polynomially many iterations

$$N = T/h \asymp \frac{\text{poly}(d, K, \delta^{-1})}{(\lambda_1 \varepsilon)^2}.$$

Compared to prior convergence rates in the Wasserstein metric ([De Bortoli et al., 2022](#)), which require exponential complexity, we achieve polynomial convergence of Riemannian diffusion models for the first time. Nonetheless, we emphasize that TV and Wasserstein distances are incomparable with each other in general, and our guarantee complements prior Wasserstein results ([De Bortoli et al., 2022](#)) by ensuring distributional closeness in a different notion with a much smaller number of iterations.

Possible improvements. We note that the bound established in Theorem 1 holds under very mild geometric assumptions, requiring only constraints on the injective radius and Riemannian curvature. The purpose of this study is to demonstrate that, in the manifold setting, the exponential blow-up in T can be avoided and polynomial complexity can be achieved. To keep the exposition as simple as possible and to clearly highlight the key ideas, we have not attempted to optimize the current bound on the degree of the polynomial. Potential approaches for sharper bounds include: (i) a better design of discretization schedule, possibly adaptive to the manifold geometry, and a more careful computation of discretization error, such as those in Li & Jiao (2024), Benton et al. (2024) (notably, the dependence on δ might be improved to poly-logarithmic in this way); (ii) a tailored analysis for TV error that does not rely on Pinsker’s inequality, like those in Li & Yan (2025), may also be extended to manifolds; (iii) a tighter version of our Minakshisundaram-Pleijel parametrix bound. We leave these improvements as future work.

4 PROOF OUTLINE

Throughout the proof, we assume that

$$h \leq \frac{1}{\text{poly}(d, K, \delta^{-1})}, \quad (3)$$

since otherwise the bound in Theorem 1 would be trivial (recall that TV distance is always bounded by 2). We start by recalling the sequence considered in RSGM. Let $(Y_k)_{k \in \{0, \dots, N\}}$ be given by $Y_N \sim \mu$ and for any $k \in \{0, \dots, N-1\}$:

$$Y_{k-1} = \begin{cases} \exp_{Y_k} \left[h \hat{s}_{t_k}(Y_k) + \sqrt{h} G_k \right], & \|h \hat{s}_{t_k}(Y_k) + \sqrt{h} G_k\| \leq h^{1/4}, \\ \text{drawn from } \mu, & \text{otherwise.} \end{cases}$$

This defines a sequence of probability transition kernels $\hat{K}_{t_k, t_{k-1}}$. For simplicity, we denote this by \hat{K}_k . Let q_k be the law of Y_k . We have

$$q_0 = q_N \hat{K}_N \hat{K}_{N-1} \cdots \hat{K}_1.$$

Similarly, the probability transition kernel from time t_k to t_{k-1} in (2) is denoted by $K_{t_k, t_{k-1}}$ or K_k in short. We have

$$p_0 = p_N K_N K_{N-1} \cdots K_1.$$

Our goal would be to bound $\text{TV}(p_0, q_0)$ as in Theorem 1, by decomposing the total error into four components:

$$(\text{initialization error}) + (\text{score error}) + (\text{drift discretization error}) + (\text{BM simulation error}).$$

More concretely:

- **Initialization error** arises from initializing Y_N with μ instead of the true marginal p_N ;
- **Score error** arises from imperfect score estimation;
- **Drift discretization error** arises from approximating the continuous-time drift $\hat{s}_t(Y_t)$ by its “time-frozen” counterpart $\hat{s}_{t_k}(Y_{t_k})$;
- **Brownian motion (BM) simulation error** is a distinctive feature of the manifold setting. Unlike in Euclidean space — where the transition kernel of Brownian motion over $[t_k, t_{k-1}]$ is exactly Gaussian with variance $(t_k - t_{k-1})$ — the transition kernel of manifold-valued Brownian motion cannot be simulated exactly by any discrete-time process, even after time discretization. This inherent inexactness gives rise to this final error term.

The first two components are relatively easier to bound using well-established tools: mixing rate bounds of heat flow (Urakawa, 2006) and Girsanov transform (Chen et al., 2023). For the drift discretization error, recent techniques developed in the Euclidean setting (Benton et al., 2024) can also be adapted with modifications that account for the manifold curvature. However, the last component — the Brownian motion simulation error — represents the core challenge in the manifold setting, which fundamentally denies a direct extension of Euclidean analysis.

Step I. Constructing auxiliary kernels via localization. In view of this, we first introduce an intermediate random process that separates the drift discretization error from the BM simulation error. Constructing such a process, however, involves additional technicality. In particular, the frozen drift $\hat{s}_{t_k}(Y_{t_k})$ is a vector in the tangent space $T_{Y_k}\mathcal{M}$, and is therefore only well-defined at the fixed point Y_k . This poses a compatibility issue: as Brownian motion evolves continuously on the manifold, it immediately departs from Y_k , rendering the frozen drift ill-defined. Careful geometric considerations are thus required to reconcile the piecewise-constant drift approximation with the intrinsic curvature of the manifold.

In our analysis, this is handled using localization by the construction of an auxiliary sequence of transition kernels K_k^{aux} . These kernels do not appear in the algorithm itself; they serve solely as an analytical tool to facilitate the proof. These kernels expose the behavior of the time-reverse SDE (2) when the estimated score \hat{s}_t is frozen to be a constant vector field in between discretization steps, meanwhile keeping the continuous Brownian motion.

Let $\eta : [0, \infty) \rightarrow [0, 1]$ be a smooth cutoff function, i.e., η is decreasing, $\eta|_{[0,1]} \equiv 1$ and $\eta|_{[4,\infty)} \equiv 0$. Such a function can be chosen such that $|\eta'| + |\eta''| + |\eta'''| \leq 100$. Recall that the injective radius of \mathcal{M} is lower bounded by $1/K$, and the curvature is upper bounded by K . Define

$$\omega := \frac{c_\omega}{Kd^4}, \quad \eta_\omega(r) = \eta\left(\frac{4r^2}{\omega^2}\right), \quad r \geq 0, \quad (4)$$

where $c_\omega > 0$ is a small universal constant. We have $\eta_\omega|_{[0, \frac{\omega}{2}]} \equiv 1$ and $\eta_\omega|_{[\omega, \infty)} \equiv 0$. For $t > 0$, $x, y \in \mathcal{M}$, define the following vector field on \mathcal{M} :

$$\mathcal{S}_{t,x}(y) = (\text{d exp}_x)_{\log_x y}(\eta_\omega(\rho(x, y)) \cdot \hat{s}_t(x)) \in T_y\mathcal{M}.$$

Intuitively speaking, $\mathcal{S}_{t,x}(\cdot)$ is the ‘‘constant’’ velocity field $\hat{s}_t(x)$ in normal coordinates, which represents our idea of freezing the drift term for a time period. The d exp_x in the formula is responsible for identifying $T_y\mathcal{M}$ with $T_x\mathcal{M}$.² On the other hand, the cut-off function η_ω is necessary to keep all our discussions restricted to the injective radius, so as to avoid pathologies of cut locus.

With this in mind, we are ready to define K_k^{aux} as the transition kernel from time t_k to t_{k-1} of the reverse-time SDE

$$\text{d}Y_\tau = \mathcal{S}_{t_k, Y_{t_k}}(Y_\tau) \text{d}t + U_{Y_\tau} \circ \text{d}W_t, \quad \tau = T - t, \quad \tau \in [t_{k-1}, t_k], \quad (5)$$

and in addition,

$$p_k^{\text{aux}} = p_N K_N^{\text{aux}} K_{N-1}^{\text{aux}} \cdots K_{k+1}^{\text{aux}}, \quad k = N, N-1, \dots, 0.$$

Step II. Decomposing different sources of error. We now decompose

$$\text{TV}(p_0, q_0) \leq \text{TV}(p_0, p_0^{\text{aux}}) + \text{TV}(p_0^{\text{aux}}, q_0) \leq \sqrt{2\text{KL}(p_0 \parallel p_0^{\text{aux}})} + \text{TV}(p_0^{\text{aux}}, q_0),$$

where the last inequality used Pinsker’s inequality. To control $\text{KL}(p_0 \parallel p_0^{\text{aux}})$, we further introduce the counterpart of $\mathcal{S}_{t,x}$ using the exact score function $\nabla \log p_t$:

$$\mathcal{S}_{t,x}^*(y) = (\text{d exp}_x)_{\log_x y}(\eta_\omega(\rho(x, y)) \cdot \nabla \log p_t(x)) \in T_y\mathcal{M}.$$

We apply Girsanov’s theorem (Hsu, 2002) to compare (5) with (2), in a way that is standard in recent literature (Chen et al., 2023; De Bortoli et al., 2022). Denote the path law of the solution of (2) by $\text{Law}(Y)$, and the path law of the solution of (5) by $\text{Law}(Y^{\text{aux}})$. Girsanov’s theorem asserts that the KL divergence $\text{KL}(\text{Law}(Y) \parallel \text{Law}(Y^{\text{aux}}))$ is upper bounded by the expectation of the squared norm of the difference between the drift terms in the two SDEs.³ More concretely,

$$\text{KL}(\text{Law}(Y) \parallel \text{Law}(Y^{\text{aux}})) \leq \sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \left\| \nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}(Y_t) \right\|^2 \text{d}t.$$

²Generally speaking, it is more natural to use parallel transport to identify different tangent spaces. However, this would later lead to a more complicated treatment of the perturbed heat equation with variable drifts. We choose to use parallelism in normal coordinates instead for simplicity.

³In its classical form, Girsanov’s theorem requires integrability such as Novikov’s condition to hold. In our setting, this can be bypassed with a localization argument as in Chen et al. (2023).

Since p_0 and p_0^{aux} are marginals of $\text{Law}(Y)$ and $\text{Law}(Y^{\text{aux}})$ respectively at time $t = t_0$, by post-processing inequality, we have

$$\begin{aligned} \text{KL}(p_0 \parallel p_0^{\text{aux}}) &\leq \sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \left\| \nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}(Y_t) \right\|^2 dt \\ &\leq 2 \underbrace{\sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \|\nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 dt}_{\text{drift discretization}} + 2 \underbrace{\sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \|\mathcal{S}_{t_k, Y_{t_k}}(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 dt}_{\text{score matching}}. \end{aligned} \quad (6)$$

It remains to decompose $\text{TV}(p_0^{\text{aux}}, q_0)$. To isolate the initialization error, we introduce

$$q_0^* = p_N \hat{K}_N \hat{K}_{N-1} \cdots \hat{K}_1.$$

By triangle inequality and post-processing inequality, we have

$$\text{TV}(p_0^{\text{aux}}, q_0) \leq \text{TV}(p_0^{\text{aux}}, q_0^*) + \text{TV}(q_0^*, q_0) \leq \underbrace{\text{TV}(p_0^{\text{aux}}, q_0^*)}_{\text{BM simulation}} + \underbrace{\text{TV}(p_N, q_N)}_{\text{initialization}}.$$

Step III. Controlling initialization and score errors. By our design, $q_N = \mu$, and $\text{TV}(p_N, q_N) = \text{TV}(p_N, \mu)$. This is known as the mixing rate of heat flow in total variation norm, and has well-established bounds, e.g., [Urakawa \(2006\)](#). The score-matching error, on the other hand, can be controlled with an analysis on the distortion on the Riemannian metric in normal coordinates. We compile the bounds into the following lemma.

Lemma 1. *There exists a universal constant $C > 0$, such that whenever $T \geq 1$, we have*

$$\text{TV}(p_N, q_N) \leq e^{C(K+d \log d)} e^{-\frac{\lambda_1}{2}(T-\frac{1}{2})},$$

and

$$\sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \|\mathcal{S}_{t_k, Y_{t_k}}(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 dt \leq 2\epsilon_{\text{score}}^2.$$

Step IV. Controlling drift discretization error with Itô/Stratonovich calculus and Li-Yau estimates. The drift discretization error defined in (6) has a similar form to the discretization error for the Euclidean setting ([Benton et al., 2024](#)), though additional complications arise due to non-constant $\mathcal{S}_{t_k, Y_{t_k}}^*$. The idea is to study the time derivative of $\mathbb{E} \|\nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2$, which in view of $\partial_\tau \log p_t = -\frac{1}{2} \Delta_M p_t$ (negative sign due to reverse time) involves space derivatives of $\log p_t$ up to third order. Fortunately, after applying Itô/Stratonovich calculus to simplify the expression, a key property in the proof of the Euclidean setting carries over: third-order derivatives of $\log p_t$ cancel out. The remaining first and second-order derivatives can be controlled by Li-Yau estimates on the log-gradient of the heat kernel. We obtain

Lemma 2. *Under the assumptions in Theorem 1, there is a universal constant $C > 0$ such that*

$$\sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \|\nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 dt \leq \frac{Cd^6 K^8}{\delta^3} h^2 N.$$

Step V. Controlling BM simulation error using parametrix estimates. Our approach is inspired by the following consequence of post-processing inequality and Pinsker's inequality:

$$\text{TV}(p_0^{\text{aux}}, q_0^*) \leq \sqrt{2 \text{KL}(p_0^{\text{aux}} \parallel q_0^*)} \leq \sqrt{2 \sum_{k=1}^N \text{KL}(p_k^{\text{aux}} K_k^{\text{aux}} \parallel p_k^{\text{aux}} \hat{K}_k)}.$$

This leads us to compare the kernel K_k^{aux} and \hat{K}_k . In normal coordinates, the Fokker-Planck equation shows that these two are the solutions of the heat equations with the Euclidean Laplacian and with the manifold Laplace-Beltrami operator. We utilize the Minakshisundaram-Pleijel parametrix theory ([Berline et al., 2003](#)) in geometric analysis for this comparison, and establish a quantitative bound in polynomially small radius and polynomially short time (cf. Lemma 20).

5 CONCLUSION

We developed a discrete-time theory for Riemannian diffusion models showing that a polynomial stepsize suffices for TV-accurate sampling under mild geometric conditions. In particular, our results show that choosing a stepsize polynomially small in manifold parameters achieves any prescribed TV target without exponential blow-ups in dimension or curvature. This complements prior Wasserstein-type guarantees which require exponentially many steps. Several important future directions remain open.

- *Sharper bounds.* For simplicity, we did not attempt to establish sharp bounds for the error terms in our analysis, and it is likely that the degree of the polynomial in the bound could be improved significantly by refining our analysis, and some of the polynomial dependencies can be improved to logarithmic ones (Benton et al., 2024; Li & Jiao, 2024).
- *Analysis of deterministic samplers.* We focused on DDPM-style stochastic samplers in our analysis. For practical purpose, it is also tempting to develop an analogous theory for DDIM-style deterministic samplers (Song et al., 2021; Li et al., 2024).
- *Conditional sampling.* Our theory was for unconditional diffusion models. Applications like solving inverse problems require conditional sampling, which calls for both new algorithm design and new theoretical analysis (Xu & Chi, 2024).

6 ACKNOWLEDGMENT

The work of X. Xu and Y. Chi is supported in part by Air Force Office of Scientific Research under FA9550-25-1-0060, and by National Science Foundation under ECCS-2126634/2537078. Z. Zhang, Y. Nakahira, and G. Qu are supported in part by NSF CAREER Award 2339112, NSF Award 2512805, NSF Award 2442948, and the Pennsylvania Infrastructure Technology Alliance. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

REFERENCES

Kwangjun Ahn and Sinho Chewi. Efficient constrained sampling via the mirror-Langevin algorithm. *Advances in Neural Information Processing Systems*, 34:28405–28418, 2021.

Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. ISSN 0304-4149. doi: [https://doi.org/10.1016/0304-4149\(82\)90051-5](https://doi.org/10.1016/0304-4149(82)90051-5).

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

Nicole Berline, Ezra Getzler, and Michele Vergne. *Heat kernels and Dirac operators*. Springer Science & Business Media, 2003.

Eliot Beyler and Francis Bach. Convergence of deterministic and stochastic diffusion-model samplers: A simple analysis in Wasserstein distance. *arXiv preprint arXiv:2508.03210*, 2025.

Karthik Bharath, Alexander Lewis, Akash Sharma, and Michael V Tretyakov. Sampling and estimation on manifolds using the langevin diffusion. In *Advances in Neural Information Processing Systems*, 2025.

Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R. Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:2209.11215*, 2023.

Xiang Cheng, Jingzhao Zhang, and Suvrit Sra. Efficient sampling on Riemannian manifolds via Langevin MCMC. *Advances in Neural Information Processing Systems*, 35:5995–6006, 2022.

Xiang Cheng, Jingzhao Zhang, and Suvrit Sra. Theory and algorithms for diffusion processes on riemannian manifolds, 2023.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34:17695–17709, 2021.

Valentin De Bortoli, Emile Mathieu, Michael John Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022.

Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021.

Nic Fishman, Leo Klarner, Valentin De Bortoli, Emile Mathieu, and Michael Hutchinson. Diffusion models for constrained domains. *arXiv preprint arXiv:2304.05364*, 2023.

Khashayar Gatmiry and Santosh S. Vempala. Convergence of the riemannian langevin algorithm. *ArXiv*, abs/2204.10818, 2022.

Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73:123 – 214, 03 2011. doi: 10.1111/j.1467-9868.2010.00765.x.

Alfred Gray. *Tubes*, volume 221. Springer Science & Business Media, 2003.

Yunrui Guan, Krishnakumar Balasubramanian, and Shiqian Ma. Riemannian proximal sampler for high-accuracy sampling on manifolds. *arXiv preprint arXiv:2502.07265*, 2025.

Richard S Hamilton. Matrix harnack estimate for the heat equation. *Communications in analysis and geometry*, 1(1):113–126, 1993.

Qing Han and Qi S Zhang. An upper bound for hessian matrices of positive solutions of heat equations. *The Journal of Geometric Analysis*, 26(2):715–749, 2016.

Ulrich G. Haussmann and Étienne Pardoux. Time reversal of diffusions. *Annals of Probability*, 14: 1188–1205, 1986.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020.

E.P. Hsu. *Stochastic Analysis on Manifolds*. Graduate studies in mathematics. American Mathematical Society, 2002. ISBN 9780821808023.

Chin-Wei Huang, Milad Aghajohari, Joey Bose, Prakash Panangaden, and Aaron C Courville. Riemannian diffusion models. *Advances in Neural Information Processing Systems*, 35:2750–2761, 2022.

Zhihan Huang, Yuting Wei, and Yuxin Chen. Denoising diffusion probabilistic models are optimally adaptive to unknown low dimensionality. *arXiv preprint arXiv:2410.18784*, 2024.

Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.

Jürgen Jost. *Riemannian Geometry and Geometric Analysis*. Universitext. Springer, Cham, 7 edition, 2017. ISBN 978-3-319-61859-3. doi: 10.1007/978-3-319-61860-9.

Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. In *Proceedings of the 34th International Conference on Algorithmic Learning Theory*, volume 201 of *PMLR*, pp. 946–985, 2023.

Gen Li and Yuchen Jiao. Improved convergence rate for diffusion probabilistic models. In *The Thirteenth International Conference on Learning Representations*, 2024.

Gen Li and Yuling Yan. Adapting to unknown low-dimensional structures in score-based diffusion models. *Advances in Neural Information Processing Systems*, 37:126297–126331, 2024.

Gen Li and Yuling Yan. O(d/t) convergence theory for diffusion probabilistic models under minimal assumptions. In *The Thirteenth International Conference on Learning Representations*, 2025.

Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. Towards non-asymptotic convergence for diffusion-based generative models. In *International Conference on Learning Representations (ICLR)*, 2024.

Junfang Li and Xiangjin Xu. Differential harnack inequalities on riemannian manifolds i: Linear heat equation. *Advances in Mathematics*, 226(5):4456–4491, 2011. ISSN 0001-8708. doi: <https://doi.org/10.1016/j.aim.2010.12.009>.

Mufan Bill Li and Murat A. Erdogdu. Riemannian langevin algorithm for solving semidefinite programs. *Bernoulli*, 29:3093 – 3113, 2023.

Peter Li and Shing Yau. On the parabolic kernel of the schödinger operator. *Acta Mathematica*, 156: 153–201, 07 1986. doi: 10.1007/BF02399203.

Jiadong Liang, Zhihan Huang, and Yuxin Chen. Low-dimensional adaptation of diffusion models: Convergence in total variation. *arXiv preprint arXiv:2501.12982*, 2025.

Yang Liang, Peng Ju, Yingbin Liang, and Ness Shroff. Non-asymptotic convergence of discrete-time diffusion models: New approach and improved rate. *arXiv preprint arXiv:2402.13901*, 2024.

Guan-Horng Liu, Tianrong Chen, Evangelos Theodorou, and Molei Tao. Mirror diffusion models for constrained and watermarked generation. *Advances in Neural Information Processing Systems*, 36:42898–42917, 2023.

Aaron Lou, Minkai Xu, Adam Farris, and Stefano Ermon. Scaling Riemannian diffusion models. *Advances in Neural Information Processing Systems*, 36:80291–80305, 2023.

Michelle Muniz, Matthias Ehrhardt, Michael Günther, and Renate Winkler. Higher strong order methods for linear \hat{t} ô sdes on matrix lie groups. *BIT Numerical Mathematics*, 62, 01 2022. doi: 10.1007/s10543-021-00905-9.

Peter Petersen. *Riemannian Geometry*, volume 171 of *Graduate Texts in Mathematics*. Springer, New York, 2 edition, 2006. ISBN 978-0-387-29246-5. doi: 10.1007/978-0-387-29403-2.

M. Piggott and V. Solo. Geometric euler–maruyama schemes for stochastic differential equations in $so(n)$ and $se(n)$. *SIAM Journal on Numerical Analysis*, 54:2490–2516, 08 2016. doi: 10.1137/15M1019726.

Yury Polyanskiy and Yihong Wu. *Information theory: From coding to learning*. Cambridge university press, 2025.

Peter Potaptchik, Iskander Azangulov, and George Deligiannidis. Linear convergence of diffusion models under the manifold hypothesis. *arXiv preprint arXiv:2410.09046*, 2024.

Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293. Springer Science & Business Media, 2013.

Richard M. Schoen, , and Shing-Tung Yau. Lectures on differential geometry. *Conference proceedings and lecture notes in geometry and topology*, 1, 1994.

Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2256–2265, Lille, France, 07–09 Jul 2015. PMLR.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021.

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

Vishwak Srinivasan, Andre Wibisono, and Ashia Wilson. Fast sampling from constrained spaces using the Metropolis-adjusted mirror Langevin algorithm. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 4593–4635. PMLR, 2024.

Hajime Urakawa. Convergence rates to equilibrium of the heat kernels on compact riemannian manifolds. *Indiana University mathematics journal*, pp. 259–288, 2006.

Xingyu Xu and Yuejie Chi. Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction. *Advances in Neural Information Processing Systems*, 37:36148–36184, 2024.

A NUMERICAL EXPERIMENTS

In this section, we verify the results in Theorem 1 on compact manifolds by measuring the exit probability in the reverse steps of Algorithm 1, which is extensively used in the proof of Lemma 19 to ensure the convergence of Algorithm 1, and the TV distance between the target distribution (a Gaussian mixture) and the recovered distribution.

Reset probability on \mathbb{S}^2 and \mathbb{T}^2 . We start by examining the total reset probability on the unit 2-sphere \mathbb{S}^2 and on the 2-torus \mathbb{T}^2 . We run the backward process in Algorithm 1 with different stepsizes h in each setting with p_0 being Gaussian mixture (see below for definition), and record the fraction of trials whose tangent update $\Delta_k = h s_{t_k}(Y_k) + \sqrt{h} G_k$ violates $\|\Delta_k\| \leq h^{1/4}$ at least once among all steps. On both manifolds, Figure 1a shows a clear linear trend of the logarithm of reset probability of against $h^{-1/2}$, which can be obtained by Gaussian tails. This confirms that the rejection sampling has no practical impact on the performance of Algorithm 1.

High-dimensional torus. We extend this experiment to the d -dimensional flat torus \mathbb{T}^d for different stepsizes. Figure 1b reports the logarithm of the reset probability versus $h^{-1/2}$ for $d \in \{2, 4, 8\}$. Increasing d raises the baseline reset rate, yet the slope of the decay remains essentially unchanged—resets remain exponentially rare as $h \downarrow 0$, which aligns with our analysis.

TV accuracy on \mathbb{T}^d with warped Gaussian mixture. Finally, we assess the distributional accuracy in TV for a *warped Gaussian mixture* target on \mathbb{T}^d , $d \in \{1, 2, 3\}$. Here, the warped Gaussian distribution is defined to be the push-forward of the Gaussian distribution by the universal covering $\mathbb{R}^d \rightarrow \mathbb{T}^d$ given by $(x_1, \dots, x_d) \mapsto (e^{i2\pi x_1}, \dots, e^{i2\pi x_d})$, and warped Gaussian mixture is similarly the push-forward of Gaussian mixture in \mathbb{R}^d . The result is depicted in Figure 2, which confirms that the total variation decays fast with the increase of the number of steps.

B PRELIMINARIES

We first introduce some tools we use in the rest of the proof.

Lemma 3 (Pinsker’s inequality, [Polyanskiy & Wu \(2025\)](#)). *For any two probability distributions p, q on \mathcal{M} , we have*

$$\text{TV}(p, q) \leq \sqrt{2\text{KL}(p \parallel q)}.$$

Lemma 4. *Let v be a vector field on \mathcal{M} . In a local coordinate on \mathcal{M} , we have*

$$\partial_\alpha v_\beta = (\nabla_\alpha v)^\beta - \Gamma_{\alpha\gamma}^\beta v_\gamma.$$

Here $\Gamma_{\alpha\gamma}^\beta$ is the Christoffel symbol, defined as

$$\Gamma_{\alpha\gamma}^\beta = \frac{1}{2} g^{\beta\delta} (\partial_\alpha g_{\gamma\delta} + \partial_\gamma g_{\alpha\delta} - \partial_\delta g_{\alpha\gamma}).$$

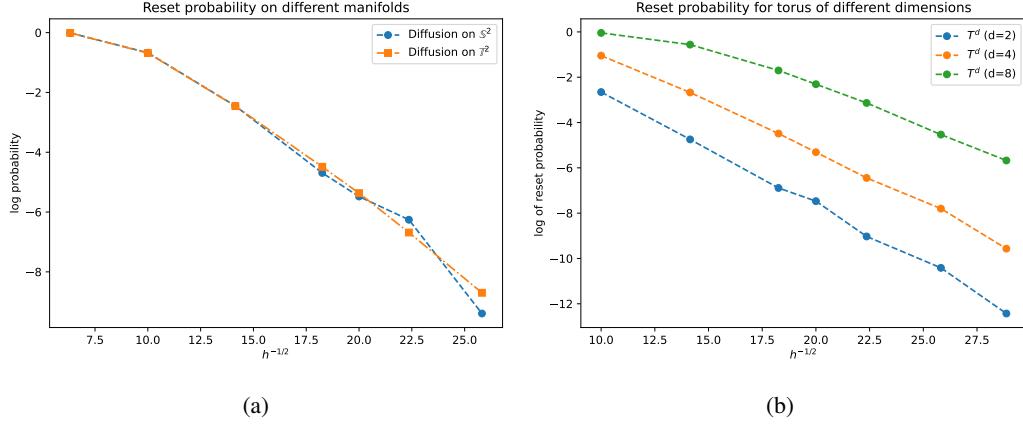


Figure 1: Reset probabilities on spheres and tori. In Figure 1a, we examine the relationship between $h^{-1/2}$ and the log of the reset probability of Algorithm 1 on both sphere \mathbb{S}^2 and torus \mathbb{T}^2 under the reset rules of Algorithm 1. In both cases, we see that the reset probability decays exponentially, confirming the conclusion in (43). In Figure 1b, we examine the same statistics on high-dimensional tori, and we find increasing d only shifts the curves to the right but leaves the exponential decay rate in $h^{-1/2}$ unchanged.

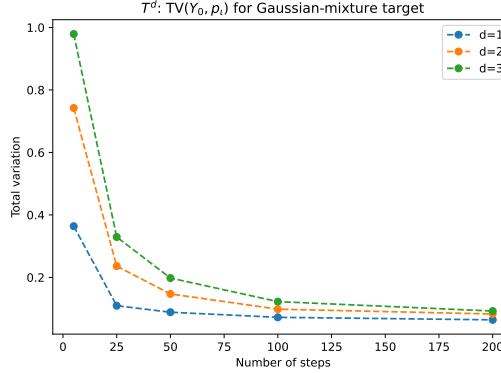


Figure 2: TV distance on \mathbb{T}^d with a warped Gaussian-mixture target. The total variation is estimated with a kernel density estimator.

Lemma 5 (Metric distortion in normal coordinates). *There exist coefficients $c, C > 0$ polynomial in d and constant in other parameters, such that the following holds. Let $x \in \mathcal{M}$. In the normal coordinates (∂_α) at x , for any $y \in \mathcal{M}$ such that $\rho(x, y) \leq c/K$, we have*

$$\begin{aligned} \|g(y) - I\| &\leq CKd^2(x, y), \\ \|\partial_\alpha g_{\beta\gamma}\| &\leq CK\rho(x, y), \\ \|\partial_{\alpha\beta} g_{\gamma\xi}\| &\leq CK. \end{aligned}$$

Proof. This is a quantitative version of the well-known Taylor expansion of g in normal coordinates (cf. [Berline et al. \(2003, Proposition 1.28\)](#)):

$$g_{\alpha\beta}(\exp_x(u)) = \delta_{\alpha\beta} - \frac{1}{3}R_{\alpha\gamma\beta\xi}(x)u^\gamma u^\xi + O((\|\text{Rm}\| + \|\nabla\text{Rm}\|)\|u\|^3), \quad \|u\| \leq c/K.$$

Let $(e_\alpha)_{\alpha=1}^d$ be an orthonormal basis of $T_x\mathcal{M}$; normal coordinates at x are defined by identifying a point $\exp_x(z^\alpha e_\alpha)$ with its coordinate vector (z^α) . Denote by ∂_α the coordinate vector fields.

Step I: Representation by Jacobi fields. Define the geodesic segment with *unit parameter* $s \in [0, 1]$:

$$\gamma(s) := \exp_x(sv), \quad \gamma(0) = x, \quad \gamma(1) = y, \quad \dot{\gamma}(s) = \frac{d}{ds}\gamma(s).$$

Then $\nabla_s \dot{\gamma} = 0$ and $|\dot{\gamma}(s)| \equiv |v| = r$.

Fix an orthonormal basis $(e_\alpha)_{\alpha=1}^d$ of $T_x \mathcal{M}$ and parallel transport it along γ to obtain an orthonormal frame $(E_\alpha(s))$ along γ :

$$\nabla_s E_\alpha(s) = 0, \quad E_\alpha(0) = e_\alpha.$$

Let $J_\beta(s)$ be the Jacobi field along γ corresponding to varying the initial point in direction e_β in the normal coordinate chart, i.e.

$$J_\beta(s) := d(\exp_x)_{sv}(se_\beta) \in T_{\gamma(s)} \mathcal{M}.$$

Equivalently, J_β is the unique Jacobi field solving

$$\nabla_s^2 J_\beta + \text{Rm}(J_\beta, \dot{\gamma})\dot{\gamma} = 0, \quad J_\beta(0) = 0, \quad \nabla_s J_\beta(0) = e_\beta. \quad (7)$$

Note that in normal coordinates, $\partial_\beta|_x = e_\beta$ and the geodesic variation $\exp_x(s(v + \varepsilon e_\beta))$ yields (7).

Write $J_\beta(s)$ in the parallel frame:

$$J_\beta(s) = \sum_{\alpha=1}^d \mathbf{J}_{\alpha\beta}(s) E_\alpha(s),$$

and let $\mathbf{J}(s) \in \mathbb{R}^{d \times d}$ be the matrix with entries $\mathbf{J}_{\alpha\beta}(s)$. Since E_α is parallel, (7) becomes the *matrix Jacobi equation*

$$\mathbf{J}''(s) + \mathbf{R}(s) \mathbf{J}(s) = 0, \quad \mathbf{J}(0) = 0, \quad \mathbf{J}'(0) = I, \quad (8)$$

where the curvature matrix $\mathbf{R}(s)$ is defined by

$$(\mathbf{R}(s) u)_\alpha := \left\langle \text{Rm} \left(\sum_\mu u_\mu E_\mu(s), \dot{\gamma}(s) \right) \dot{\gamma}(s), E_\alpha(s) \right\rangle.$$

From $\|\text{Rm}\| \leq K$ and $|\dot{\gamma}| = r$ we have

$$\|\mathbf{R}(s)\| \leq K |\dot{\gamma}(s)|^2 = Kr^2, \quad s \in [0, 1]. \quad (9)$$

In normal coordinates, the coordinate vector fields at $y = \gamma(1)$ are

$$\partial_\beta|_y = d(\exp_x)_v(e_\beta) = J_\beta(1).$$

Since the frame at $s = 1$ is orthonormal, the metric coefficients are

$$g_{\beta\gamma}(y) = \langle \partial_\beta|_y, \partial_\gamma|_y \rangle = \langle J_\beta(1), J_\gamma(1) \rangle = \sum_{\alpha=1}^d \mathbf{J}_{\alpha\beta}(1) \mathbf{J}_{\alpha\gamma}(1) = (\mathbf{J}(1)^\top \mathbf{J}(1))_{\beta\gamma}.$$

Hence, as matrices,

$$g(y) = \mathbf{J}(1)^\top \mathbf{J}(1). \quad (10)$$

Step II: Control of $\mathbf{J}(1) - I$ via Grönwall inequality. From (8), integrating twice and using $\mathbf{J}(0) = 0$, $\mathbf{J}'(0) = I$, we get the exact Volterra equation

$$\mathbf{J}(s) = sI - \int_0^s (s - \tau) \mathbf{R}(\tau) \mathbf{J}(\tau) d\tau, \quad s \in [0, 1]. \quad (11)$$

Taking operator norms and using (9) gives for $s \in [0, 1]$:

$$\|\mathbf{J}(s)\| \leq s + Kr^2 \int_0^s (s - \tau) \|\mathbf{J}(\tau)\| d\tau.$$

A standard Grönwall argument yields

$$\|\mathbf{J}(s)\| \leq Cs \quad \text{and} \quad \|\mathbf{J}'(s)\| \leq C \quad \text{for all } s \in [0, 1], \text{ provided } r \leq c/\sqrt{K}. \quad (12)$$

Now subtract sI in (11):

$$J(s) - sI = - \int_0^s (s - \tau) R(\tau) J(\tau) d\tau.$$

Using (9) and (12),

$$\|J(1) - I\| \leq \int_0^1 (1 - \tau) \|R(\tau)\| \|J(\tau)\| d\tau \leq C \int_0^1 (1 - \tau) (Kr^2) \tau d\tau \leq CKr^2.$$

Combine with (10):

$$g(y) - I = (J(1)^\top J(1) - I) = (J(1) - I)^\top + (J(1) - I) + (J(1) - I)^\top (J(1) - I),$$

so

$$\|g(y) - I\| \leq C\|J(1) - I\| \leq CKr^2. \quad (13)$$

Step III: Control of first derivatives. We first control $\partial_\alpha J(1)$ as a function of the coordinate v . Let $v \mapsto J_v(s)$ denote the Jacobi matrix for the geodesic $\gamma_v(s) = \exp_x(sv)$. Differentiate the ODE (8) w.r.t. v^α :

$$\partial_\alpha J'' + R \partial_\alpha J = -(\partial_\alpha R) J, \quad \partial_\alpha J(0) = 0, \quad \partial_\alpha J'(0) = 0. \quad (14)$$

Bound on $\partial_\alpha R$. Recall $R(s)$ represents the operator $u \mapsto \text{Rm}(u, \dot{\gamma})\dot{\gamma}$ in the parallel frame. Varying v changes both γ and $\dot{\gamma}$; the corresponding variation field $V_\alpha(s) := \partial_\alpha \gamma_v(s)$ along γ is itself a Jacobi field with $V_\alpha(0) = 0$, $\nabla_s V_\alpha(0) = e_\alpha$, hence by the same estimate as (12)

$$\|V_\alpha(s)\| \leq Cs, \quad \|\nabla_s V_\alpha(s)\| \leq C. \quad (15)$$

Using the product rule on $\text{Rm}(\cdot, \dot{\gamma})\dot{\gamma}$ and our Assumption 1, together with $|\dot{\gamma}| = r$ and (15), one obtains the uniform operator bound

$$\|\partial_\alpha R(s)\| \leq C \left(\|\nabla \text{Rm}\| \|V_\alpha(s)\| |\dot{\gamma}|^2 + \|\text{Rm}\| |\dot{\gamma}| \|\partial_\alpha \dot{\gamma}(s)\| \right) \leq C(K \cdot s \cdot r^2 + K \cdot r \cdot 1) \leq CKr, \quad (16)$$

for all $s \in [0, 1]$ (since $s \leq 1$). Here we used $\partial_\alpha \dot{\gamma} = \nabla_s V_\alpha$.

Now solve (14) by the same Duhamel principle: integrating twice with zero initial data gives

$$\partial_\alpha J(s) = - \int_0^s (s - \tau) \left(R(\tau) \partial_\alpha J(\tau) + (\partial_\alpha R)(\tau) J(\tau) \right) d\tau. \quad (17)$$

Using (9), (16), and (12), we obtain

$$\|\partial_\alpha J(s)\| \leq Kr^2 \int_0^s (s - \tau) \|\partial_\alpha J(\tau)\| d\tau + CKr \int_0^s (s - \tau) \|J(\tau)\| d\tau \leq Kr^2 \int_0^s (s - \tau) \|\partial_\alpha J(\tau)\| d\tau + CKr s^3.$$

Apply the same Grönwall comparison as before (now with a forcing term $CKr s^3$) to conclude, for $r \leq c/\sqrt{K}$,

$$\|\partial_\alpha J(1)\| \leq CKr. \quad (18)$$

Finally differentiate $g = J^\top J$:

$$\partial_\alpha g = (\partial_\alpha J)^\top J + J^\top (\partial_\alpha J),$$

so by (12) (at $s = 1$) and (18),

$$\|\partial_\alpha g(y)\| \leq 2\|\partial_\alpha J(1)\| \|J(1)\| \leq C(Kr) \cdot 1 \leq CKr. \quad (19)$$

Step IV: Control of second derivatives. Differentiate (14) once more:

$$\partial_{\alpha\beta} J'' + R \partial_{\alpha\beta} J = -(\partial_{\alpha\beta} R) J - (\partial_\alpha R) \partial_\beta J - (\partial_\beta R) \partial_\alpha J, \quad (20)$$

with zero initial data at $s = 0$.

Bound on $\partial_{\alpha\beta} R$. Under Assumption 1, $\partial_{\alpha\beta} R$ can be bounded uniformly by CK on $[0, 1]$ as follows: expanding the second parameter derivative of $\text{Rm}(\cdot, \dot{\gamma})\dot{\gamma}$ produces terms of the schematic form

$$(\nabla \text{Rm})(V) \cdot \dot{\gamma} \cdot (\partial \dot{\gamma}), \quad \text{Rm}(\cdot, \partial \dot{\gamma}) \cdot (\partial \dot{\gamma}), \quad (\nabla \text{Rm})(\partial V) \cdot \dot{\gamma} \cdot \dot{\gamma},$$

and also terms involving $\nabla_s(\partial V)$, all of which are controlled using $\|V\| \lesssim 1$, $\|\nabla_s V\| \lesssim 1$ and the fact that each appearance of $\dot{\gamma}$ contributes a factor r . Concretely, one shows (using (15) for both V_α, V_β and the same Jacobi estimates for their derivatives) that

$$\|\partial_{\alpha\beta} R(s)\| \leq CK \quad \text{for all } s \in [0, 1]. \quad (21)$$

Now apply Duhamel's principle to (20) with zero initial data:

$$\partial_{\alpha\beta} J(s) = - \int_0^s (s-\tau) \left(R \partial_{\alpha\beta} J + (\partial_{\alpha\beta} R) J + (\partial_\alpha R) \partial_\beta J + (\partial_\beta R) \partial_\alpha J \right) (\tau) d\tau.$$

Take norms and use (9), (21), (12), (16), (18):

$$\|\partial_{\alpha\beta} J(s)\| \leq Kr^2 \int_0^s (s-\tau) \|\partial_{\alpha\beta} J(\tau)\| d\tau + CK \int_0^s (s-\tau) \|J(\tau)\| d\tau + C(Kr)(Kr) \int_0^s (s-\tau) d\tau.$$

Since $\|J(\tau)\| \leq C\tau$, the second integral is bounded by CKs^3 , and the third is bounded by $CK^2r^2s^2 \leq CKs^2$ provided $r \leq c/\sqrt{K}$. Thus for $s \leq 1$,

$$\|\partial_{\alpha\beta} J(s)\| \leq Kr^2 \int_0^s (s-\tau) \|\partial_{\alpha\beta} J(\tau)\| d\tau + CK.$$

Applying Grönwall argument once more yields

$$\|\partial_{\alpha\beta} J(1)\| \leq CK. \quad (22)$$

Finally differentiate $g = J^\top J$ twice:

$$\partial_{\alpha\beta} g = (\partial_{\alpha\beta} J)^\top J + J^\top (\partial_{\alpha\beta} J) + (\partial_\alpha J)^\top (\partial_\beta J) + (\partial_\beta J)^\top (\partial_\alpha J).$$

Hence by (12), (18), (22) (and $K^2r^2 \leq CK$ for $r \leq c/\sqrt{K}$),

$$\|\partial_{\alpha\beta} g(y)\| \leq C\|\partial_{\alpha\beta} J(1)\| \cdot \|J(1)\| + C\|\partial_\alpha J(1)\| \|\partial_\beta J(1)\| \leq CK + C(Kr)^2 \leq CK. \quad (23)$$

This completes the proof. \square

Lemma 6. Fix $x \in \mathcal{M}$. Define

$$J(x, u) := |\det d\exp_x(u)| = \sqrt{\det g_{ij}(\exp_x u)}.$$

There exist universal constants $c, C > 0$, such that for $u \in T_x \mathcal{M}$ with $\|u\| \leq \frac{c}{Kd}$, we have the following bound on $J(x, u)$:

$$|J(x, u) - 1| \leq CdK\|u\|^2. \quad (24)$$

In particular, we have

$$\frac{1}{2} \leq J(x, u) \leq 2, \quad \|u\| \leq \frac{c}{Kd}.$$

Proof. Work in normal coordinates at x so that $\exp_x : B_{\text{euc}}(0, 1/K) \subset T_x \mathcal{M} \rightarrow B_{\text{geo}}(x, 1/K)$ is a diffeomorphism and $g_{ij}(0) = \delta_{ij}$, $\Gamma_{ij}^k(0) = 0$.

From Lemma 5, we know that $\|g(\exp_x u) - I\| \leq CK\|u\|^2$. In the region $\|u\| \leq \frac{c}{Kd}$, we have

$$\|g(\exp_x u) - I\| \leq \frac{c}{d}.$$

Therefore, by Taylor expansion of determinants, we know

$$\begin{aligned} |\det g(\exp_x u) - 1| &= |\det(I + g(\exp_x u) - I) - 1| \\ &\leq C \text{tr}(g(\exp_x u) - I) \\ &\leq Cd \cdot \|g(\exp_x u) - I\| \\ &\leq Cd \cdot CK\|u\|^2. \end{aligned}$$

This concludes the proof by adjusting C if necessary. \square

The metric distortion bound implies that geodesic is almost a straight line, in a sufficiently small normal neighborhood. The following quantitative bound shall be useful.

Lemma 7 (Geodesics are almost straight in small balls). *There exist coefficients $c, C > 0$ polynomial in d and constant in other parameters, such that the following holds. Fix any $x \in \mathcal{M}$ and let $0 < r \leq c/K$. Let $y, z \in B_x(r)$ and γ be the unit-speed geodesic connecting y to z . Write*

$$y(s) := \exp_x^{-1}(\gamma(s)) \in T_x \mathcal{M} \simeq \mathbb{R}^d$$

for its representation in normal coordinates at x . Then:

(i) (Almost constant velocity)

$$\sup_{s \in [0, \ell]} |\dot{y}(s) - \dot{y}(0)| \leq CKr^2.$$

(ii) (Almost linear trajectory)

$$\sup_{s \in [0, \ell]} |y(s) - y(0) - s \dot{y}(0)| \leq CKr^3.$$

In words, in normal coordinates at x , any geodesic segment contained in $B_x(r)$ deviates from the Euclidean line segment connecting its endpoints by at most $O(Kr^3)$ in position and $O(Kr^2)$ in direction.

Proof. Work in normal coordinates at x . By bounded geometry and the choice of r , the metric coefficients satisfy, in view of Lemma 5, that

$$\|g(y) - I\| \leq CK|y|^2, \quad \|\partial g(y)\| \leq CK|y|, \quad |y| \leq r,$$

which implies the Christoffel symbols obey

$$|\Gamma(y)| \leq CK|y| \leq CKr.$$

The coordinate representation $y(s)$ of the geodesic satisfies the geodesic equation

$$\ddot{y}^k(s) + \Gamma_{ij}^k(y(s)) \dot{y}^i(s) \dot{y}^j(s) = 0.$$

Since γ is unit-speed and $g(y)$ is uniformly equivalent to the Euclidean metric on $|y| \leq r$, we have $|\dot{y}(s)| \asymp 1$. Consequently,

$$|\ddot{y}(s)| \leq CKr \quad \text{for all } s \in [0, \ell].$$

Integrating once gives

$$|\dot{y}(s) - \dot{y}(0)| \leq \int_0^s |\ddot{y}(u)| du \leq CKr s \leq CKr^2,$$

proving (i). Integrating again yields

$$|y(s) - y(0) - s \dot{y}(0)| \leq \int_0^s \int_0^u |\ddot{y}(w)| dw du \leq CKr s^2 \leq CKr^3,$$

which proves (ii). \square

We spell out the constants in a few classical inequalities in geometric analysis.

Lemma 8 (Schoen et al. (1994, Thm. 4.6)). *Let (\mathcal{M}, g) be a complete Riemannian manifold with $\text{Ric}(\mathcal{M}) \geq -K$ for some $K \geq 0$. Let $H(x, y, t)$ be the heat kernel, i.e., the fundamental solution of $(\Delta - \frac{\partial}{\partial t})u(x, t) = 0$. Then, for every $\delta_{\text{Sch}} > 0$ and $\alpha > 1$,*

$$H(t, x, y) \leq C(\delta_{\text{Sch}}, d, \alpha) V_x(\sqrt{t})^{-1/2} V_y(\sqrt{t})^{-1/2} \exp\left[-\frac{r^2(x, y)}{(4 + \delta_{\text{Sch}})t} + C_1 \delta_{\text{Sch}} Kt\right],$$

where $V_x(R) = \mu(B_x(R))$, $C(\delta_{\text{Sch}}, d, \alpha) = (1 + \delta_{\text{Sch}})^{d\alpha} \exp(\frac{1+\alpha}{\delta_{\text{Sch}}})$, and $C_1 = \frac{\alpha d}{\alpha - 1}$.

To unleash the power of Lemma 8, we need the following lower bound on volume of geodesic balls assuming bounded geometry.

Lemma 9 (Günther's comparison theorem). *Under Assumption 1, we have*

$$V_x(r) \geq \frac{(2\pi)^{d/2}}{\Gamma(d/2)} \int_0^r \left(\frac{\sin(t\sqrt{K})}{\sqrt{K}} \right)^{d-1} dt, \quad 0 \leq r \leq 1/K.$$

Here Γ is the Gamma function. In particular, when $r \leq c/K$ for some small universal constant $c > 0$, we have

$$V_x(r) \geq \frac{\pi^{d/2}}{d\Gamma(d/2)} r^d \geq \frac{1}{d^{d/2}} r^d.$$

Proof. We observe that $\|\text{Rm}\| \leq K$ implies that the sectional curvature is upper bounded by K , since by definition $\text{sec}(u, v) = \text{Rm}(u, v, u, v)$. The first inequality follows from the classical form of Günther's comparison theorem, see for example [Gray \(2003, Theorem 3.17\)](#). The second inequality follows from the elementary bound that $\sin(x) \geq \frac{1}{2}x$ for $x \in [0, c]$, where c is a small universal constant, and the crude bound $\Gamma(x) \leq x^{x-1}$ for $x \geq 1$. \square

A complementary lower bound to Lemma 8 is as follows.

Lemma 10 ([Li & Xu \(2011, Thm. 1.5\)](#)). *Let (\mathcal{M}, g) be complete, possibly with $\text{Ric}(\mathcal{M}) \geq -K$. For the (Neumann) heat kernel $H(x, y, t)$ and all $x, y \in \mathcal{M}, t > 0$,*

$$\begin{aligned} H(t, x, y) &\geq (4\pi t)^{-d/2} \frac{(2Kt)^{d/2}}{(e^{2Kt} - 2Kt - 1)^{d/4}} \exp\left[-\frac{\rho(x, y)^2}{4t}\left(1 + \frac{Kt \coth(Kt) - 1}{Kt}\right)\right], \\ H(t, x, y) &\geq (4\pi t)^{-d/2} \exp\left[-\frac{\rho(x, y)^2}{4t}\left(1 + \frac{1}{3}Kt\right) - \frac{d}{4}Kt\right]. \end{aligned} \quad (25)$$

The above bounds for heat kernel translates seamlessly to p_t , since p_t is a convolution of p_0 with $H(t, \cdot, \cdot)$. We formalize this in the following lemma.

Lemma 11. *We have*

$$\inf_{x, y \in \mathcal{M}} H(t, x, y) \leq \inf_{x \in \mathcal{M}} p_t(x) \leq \sup_{x \in \mathcal{M}} p_t(x) \leq \sup_{x, y \in \mathcal{M}} H(t, x, y).$$

Proof. This follows from taking infimum and supremum respectively in the formula (Duhamel principle)

$$p_t(y) = \int_{\mathcal{M}} p_0(x) H(t, x, y) \mu(dx).$$

\square

The following lemma compiles a few follow-ups of Li-Yau estimates ([Hamilton, 1993; Han & Zhang, 2016](#)) with constants made explicit.

Lemma 12. *Under Assumption 1, we have Han-Zhang's inequality*

$$\frac{\nabla^2 p_t}{p_t} \preceq C_{\text{HZ}} \left(\frac{1}{t} + K \right) \left(1 + \log \frac{\sup p_{t/2}}{p_t} \right).$$

On the other hand, we also have Hamilton's Harnack inequality

$$\nabla^2 \log p_t = \frac{\nabla^2 p_t}{p_t} - \frac{(\nabla p_t)(\nabla p_t)^\top}{p_t^2} \succeq -\frac{1}{2t} g - C_{\text{Ham}} \left(1 + \log \frac{\sup p_{t/2}}{p_t} \right) g$$

and

$$\|\nabla \log p_t\|^2 = \frac{\|\nabla p_t\|^2}{p_t^2} \leq C \left(\frac{1}{t} + K \right) \log \frac{\sup p_{t/2}}{p_t}.$$

Here $C > 0$ is a universal constant, $C_{\text{HZ}} = CdK$, $C_{\text{Ham}} = CdK^2$.

Proof. The last inequality follows from [Hamilton \(1993, Theorem 1.1\)](#). The proof of the rest two inequalities require tracing the proofs of [Hamilton \(1993\)](#); [Han & Zhang \(2016\)](#). The details would be too tedious to reproduce here, so we leave pointers to relevant proofs for interested readers.

To prove the second inequality, we trace the proof of [Hamilton \(1993, Theorem 4.3\)](#) to see that if $A > 0$ is such that

$$\frac{\Delta_{\mathcal{M}} p_t}{p_t} \leq \frac{A}{t} \left(d + \log \frac{\sup p_{t/2}}{p_t} \right)$$

and

$$\frac{\|\nabla p_t\|^2}{p_t^2} \leq \frac{A}{t} \left(d + \log \frac{\sup p_{t/2}}{p_t} \right),$$

then

$$\frac{\nabla^2 p_t}{p_t} - \frac{(\nabla p_t)(\nabla p_t)^\top}{p_t^2} \succeq -\frac{1}{2t}g - CA(\|\text{Rm}\| + \|\nabla \text{Rm}\|)g.$$

Here $C > 0$ is an absolute constant. By Assumption 1, this implies

$$C_{\text{Ham}} \leq CKC_{\text{HZ}}.$$

Tracing the proof the main theorem in [Han & Zhang \(2016, Page 9\)](#), we can see

$$C_{\text{HZ}} \leq C_{\text{LY}}(1 + K) = CdK,$$

where C_{LY} is the maximum of the coefficients before t^{-1} and K in [Li & Yau \(1986, Theorem 1.2\)](#). This was explicitly defined as Cd therein, by setting $\alpha = 2$ there. This completes the proof. \square

Lemma 13. *For any three points $x, y, z \in \mathcal{M}$, we have*

$$\frac{\rho(x, z)^2}{1-t} + \frac{\rho(z, y)^2}{t} \geq \rho(x, y)^2, \quad \forall t \in (0, 1).$$

The equality is attainable at some point z_ on the minimum-length geodesic from x to y . Moreover, if x, y, z are within $\iota \leq 1/\text{poly}(d, K)$ distance to each other, then the function*

$$\psi(z) := \frac{\rho(x, z)^2}{1-t} + \frac{\rho(z, y)^2}{t} - \rho(x, y)^2$$

is $\frac{1-Cd^2K^2\iota}{t(1-t)}$ -strongly convex in the normal coordinates at x (or y), where $C > 0$ is a universal constant.

Proof. The first inequality follows from Cauchy-Schwarz inequality and triangle inequality:

$$(1-t+t) \left(\frac{\rho(x, z)^2}{1-t} + \frac{\rho(z, y)^2}{t} \right) \geq (\rho(x, z) + \rho(z, y))^2 \geq \rho(x, y)^2.$$

Let $\gamma : [0, 1] \rightarrow \mathcal{M}$ be the constant-speed, minimum-length geodesic from x to y . It is then straightforward to check that

$$\frac{\rho(x, z_*)^2}{1-t} + \frac{\rho(z_*, y)^2}{t} = \rho(x, y)^2, \quad z_* := \gamma(\lambda_*),$$

where $\lambda_* \in (0, 1)$ solves the quadratic equation

$$\frac{\lambda^2}{1-t} + \frac{(1-\lambda)^2}{t} = 1.$$

Proving the strong convexity requires Lemma 5 and Lemma 6, which implies that $\rho(x, z)^2$ and $\rho(z, y)^2$ are both $(1 - Cd^2K^2\iota)$ -strongly convex in the normal coordinates. The desired conclusion then follows from

$$\frac{1}{1-t} + \frac{1}{t} = \frac{1}{t(1-t)}.$$

The proof is completed. \square

C INITIALIZATION ERROR

We require the following result from [Urakawa \(2006, Proposition 2.6\)](#).

Lemma 14. *Denote $H(t, x, y)$ the heat kernel on \mathcal{M} . Assume*

$$A := \sup_{t \leq 1} \sup_{x \in \mathcal{M}} t^{d/2} H(t, x, x).$$

Then for any probability distribution p_0 , its evolution along heat flow $\partial_t p_t = \frac{1}{2} \Delta_{\mathcal{M}}$ satisfies

$$\text{TV}(p_t, \mu) \leq \sqrt{A} e^{-\frac{1}{2\lambda_1}(t-\frac{1}{2})}, \quad t \geq 1.$$

We combine this with the Li-Yau upper bound (Lemma 8) to obtain

Lemma 15 (First part of Lemma 1). *Under Assumption 1, there exists a universal constant $C > 0$ such that*

$$\text{TV}(p_N, \mu) \leq e^{C(K+d \log(Kd))} e^{-\frac{1}{2\lambda_1}(T-\frac{1}{2})}.$$

Proof. Plug the bound in Lemma 9 into Lemma 8 and use the fact that $\sup_{x,y} H(t, x, y)$ is decreasing in t (by convolution inequality), we obtain

$$\begin{aligned} A &\leq (Cd/K)^d e^{CK} \\ &\leq \exp(C'K + C'd \log(Kd)), \end{aligned}$$

for some universal constants $C, C' > 0$. Note that we absorbed $d \log K$ into $K + d \log d$. The desired claim follows. \square

D SCORE MATCHING ERROR

We now prove the second inequality in Lemma 1.

Lemma 16. *Under the same assumptions as in Theorem 1, we have*

$$\sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \|\mathcal{S}_{t_k, Y_{t_k}}(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 dt \leq 2\varepsilon_{\text{score}}^2.$$

Proof. This is relatively straightforward. Notice that in normal coordinates, by Lemma 5, we have

$$\begin{aligned} \|\mathcal{S}_{t_k, Y_{t_k}}(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 &= g_{\alpha\beta}(Y_t) (\widehat{s}_t - \nabla \log p_t)^\alpha (\widehat{s}_{t_k}(Y_{t_k}) - \nabla \log p_{t_k}(Y_{t_k}))^\beta \\ &\leq \|g(Y_t)\| \cdot \|\widehat{s}_{t_k}(Y_{t_k}) - \nabla \log p_{t_k}(Y_{t_k})\|^2 \\ &\leq 2\|\widehat{s}_{t_k}(Y_{t_k}) - \nabla \log p_{t_k}(Y_{t_k})\|^2 \end{aligned}$$

for $\rho(Y_t, Y_{t_k}) \leq c/K$, and is 0 otherwise due to our cutoff η_ω . Therefore

$$\mathbb{E} \|\mathcal{S}_{t_k, Y_{t_k}}(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 \leq 2\mathbb{E} \|\widehat{s}_{t_k}(Y_{t_k}) - \nabla \log p_{t_k}(Y_{t_k})\|^2,$$

and consequently,

$$\begin{aligned} &\sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \|\mathcal{S}_{t_k, Y_{t_k}}(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 dt \\ &\leq 2 \sum_{k=1}^N (t_k - t_{k-1}) \mathbb{E} \|\widehat{s}_{t_k}(Y_{t_k}) - \nabla \log p_{t_k}(Y_{t_k})\|^2 = 2\varepsilon_{\text{score}}^2. \end{aligned}$$

This proves the claim, as desired. \square

E DISCRETIZATION ERROR

Lemma 17. *Under the same assumptions as in Theorem 1 and assuming (3) without loss of generality, there is a universal constant $C > 0$ such that for $t_k - h \leq t \leq t_k$, we have*

$$\mathbb{E}\|\nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 \leq \frac{Cd^6 K^8}{t^3}(t_k - t).$$

Proof. For convenience, set the reverse time

$$\tau := t_k - t.$$

The main challenge is that Li-Yau estimates provide sharp uniform control up to second-order derivatives of $\log p_t$, but a naïve calculation of the difference $\nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)$ involves third-order derivatives. More precisely, a straightforward Taylor expansion will introduce a factor of $\partial_\tau \nabla \log p_t$, which, by reverse-time heat equation $\partial_\tau p_t = -\frac{1}{2} \Delta_{\mathcal{M}} p_t$, contains third-order derivatives of p_t . We bypass this difficulty by making use of Itô’s calculus to show that third-order derivatives cancel out; this is inspired by [Benton et al. \(2024\)](#), where a similar strategy was employed to the Euclidean setting.

Step I. Applying Itô/Stratonovich formula. We first compute $\partial_\tau \nabla \log p_t$. Since the forward heat equation is $\partial_t p_t = \frac{1}{2} \Delta_{\mathcal{M}} p_t$, we have

$$\partial_\tau \nabla \log p_t = -\partial_t \nabla \log p_t = -\nabla \left(\frac{\partial_t p_t}{p_t} \right) = -\frac{1}{2} \nabla \left(\frac{\Delta_{\mathcal{M}} p_t}{p_t} \right).$$

Use the manifold quotient rule and the identity $\Delta_{\mathcal{M}} \log p_t = \frac{\Delta_{\mathcal{M}} p_t}{p_t} - \|\nabla \log p_t\|^2$ to rewrite

$$\frac{\Delta_{\mathcal{M}} p_t}{p_t} = \Delta_{\mathcal{M}} \log p_t + \|\nabla \log p_t\|^2.$$

Therefore, we have

$$\begin{aligned} \partial_\tau \nabla \log p_t &= -\frac{1}{2} \nabla \left(\frac{\Delta_{\mathcal{M}} p_t}{p_t} \right) = -\frac{1}{2} \nabla \Delta_{\mathcal{M}} \log p_t - \frac{1}{2} \nabla \|\nabla \log p_t\|^2 \\ &= -\frac{1}{2} \nabla \Delta_{\mathcal{M}} \log p_t - \nabla^2 \log p_t \cdot \nabla \log p_t. \end{aligned}$$

On the other hand, it is straightforward to calculate

$$\nabla^2 \log p_t = \frac{\nabla^2 p_t}{p_t} - \frac{(\nabla p_t)(\nabla p_t)^\top}{p_t^2}.$$

Now, from Itô’s formula, we know

$$\begin{aligned} d\nabla \log p_t(Y_t) &= (\partial_\tau \nabla \log p_t)(Y_t) d\tau + \nabla^2 \log p_t(Y_t) (\nabla \log p_t(Y_t) d\tau + U_{Y_t} \circ dW_t) + \frac{1}{2} \Delta_{\mathcal{M}} \nabla \log p_t(Y_t) d\tau \\ &= -\frac{1}{2} \nabla \Delta_{\mathcal{M}} \log p_t d\tau - \nabla^2 \log p_t \cdot \nabla \log p_t d\tau + \nabla^2 \log p_t \cdot \nabla \log p_t d\tau \\ &\quad + \nabla^2 \log p_t \cdot U_{Y_t} \circ dW_t + \frac{1}{2} \Delta_{\mathcal{M}} \nabla \log p_t d\tau \\ &= \frac{1}{2} (\Delta_{\mathcal{M}} \nabla - \nabla \Delta_{\mathcal{M}}) \log p_t d\tau + \nabla^2 \log p_t \cdot U_{Y_t} \circ dW_t \\ &= \frac{1}{2} \text{Ric}^\sharp(\nabla \log p_t, \cdot) d\tau + \nabla^2 \log p_t \cdot U_{Y_t} \circ dW_t, \end{aligned}$$

where the last line follows from Bochner’s identity $(\Delta_{\mathcal{M}} \nabla - \nabla \Delta_{\mathcal{M}})f = \text{Ric}^\sharp(\nabla f, \cdot)$, and Ric^\sharp denotes the $(1, 1)$ -tensor obtained by raising one index in Ricci curvature. Notice here the cancellation of third-order derivatives.

On the other hand,

$$d\mathcal{S}_{t_k, Y_{t_k}}^*(Y_t) = \nabla \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t) \cdot (\nabla \log p_t d\tau + U_{Y_t} \circ dW_t) + \frac{1}{2} \Delta_{\mathcal{M}} \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t) d\tau.$$

In normal coordinates, $\mathcal{S}_{t_k, Y_{t_k}}^*$ is a constant vector field inside $B(0, \omega/3)$, therefore we have (cf. Lemma 4):

$$\nabla_\alpha \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)^\beta = \Gamma_{\alpha\gamma}^\beta \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)^\gamma = \Gamma_{\alpha\gamma}^\beta \nabla^\gamma \log p_{t_k}(Y_{t_k}), \quad \rho(Y_t, Y_{t_k}) \leq \omega/3,$$

and similarly, when $\rho(Y_t, Y_{t_k}) \leq \omega/3$, we have

$$\Delta_{\mathcal{M}} \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)^\alpha = -\text{Ric}^\alpha_\beta \nabla^\beta \log p_{t_k}(Y_{t_k}) - g^{\beta\gamma} \left(\partial_\gamma \Gamma_{\beta\xi}^\alpha + \Gamma_{\gamma\xi}^\alpha \Gamma_{\beta\xi}^\zeta + \Gamma_{\zeta\xi}^\alpha \Gamma_{\beta\gamma}^\zeta \right) \nabla^\xi \log p_{t_k}(Y_{t_k}).$$

Step II. Bounding the coefficients. Combine the above formulas with the estimates given in Lemma 5, we obtain for some universal constant $C > 0$:

$$\begin{aligned} \left\| \text{Ric}^\sharp(\nabla \log p_t, \cdot) \right\| &\leq CK \|\nabla \log p_t(Y_t)\|, \\ \left\| \nabla \mathcal{S}_{t_k, Y_{t_k}}^* \right\| &\leq CK \|\nabla \log p_{t_k}(Y_{t_k})\|, \quad \rho(Y_t, Y_{t_k}) \leq \omega/2, \\ \left\| \Delta_{\mathcal{M}} \mathcal{S}_{t_k, Y_{t_k}}^* \right\| &\leq CK^2 \|\nabla \log p_{t_k}(Y_{t_k})\|, \quad \rho(Y_t, Y_{t_k}) \leq \omega/2. \end{aligned}$$

Outside the geodesic ball $B_{Y_{t_k}}(\omega/2)$, the field $\mathcal{S}_{t_k, Y_{t_k}}^*$ is non-zero only inside $B_{Y_{t_k}}(\omega)$. Between these two balls, we have to take into account the radial derivative of the cutoff function η_ω , whose first order derivative is bounded by $C\omega^{-1}$ and second order derivative by $C\omega^{-2}$ by our construction of η_ω (recall that $|\eta'| + |\eta''| \leq 100$). Apply Lemma 4 and Lemma 5 again, this time we bound

$$\begin{aligned} \left\| \nabla \mathcal{S}_{t_k, Y_{t_k}}^* \right\| &\leq CK\omega^{-1} \|\nabla \log p_{t_k}(Y_{t_k})\|, \\ \left\| \Delta_{\mathcal{M}} \mathcal{S}_{t_k, Y_{t_k}}^* \right\| &\leq CK^2\omega^{-2} \|\nabla \log p_{t_k}(Y_{t_k})\|. \end{aligned}$$

We now apply Itô's formula on manifold (i.e., take expectation and invoke the martingale property in (1)) and collect the above bounds to obtain (cf. Benton et al. (2024))

$$\begin{aligned} &\left| \frac{d}{d\tau} \mathbb{E} \|\nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 \right| \\ &\leq CK^2 (\mathbb{E} \|\nabla \log p_t(Y_t)\|^2 + \mathbb{E} \|\nabla \log p_{t_k}(Y_{t_k})\|^2 + \mathbb{E} \|\nabla^2 \log p_t(Y_t)\|^2) \\ &\quad + CK^2\omega^{-2} \mathbb{E} \left[(\|\nabla \log p_t(Y_t)\|^2 + \|\nabla \log p_{t_k}(Y_{t_k})\|^2 + \|\nabla^2 \log p_t(Y_t)\|^2) \mathbf{1}_{\rho(Y_t, Y_{t_k}) > \omega/3} \right] \\ &\leq CK^2 (\mathbb{E} \|\nabla \log p_t(Y_t)\|^2 + \mathbb{E} \|\nabla \log p_{t_k}(Y_{t_k})\|^2 + \mathbb{E} \|\nabla^2 \log p_t(Y_t)\|^2) \\ &\quad + CK^2\omega^{-2} \sqrt{\mathbb{E} \|\nabla \log p_t(Y_t)\|^4 + \mathbb{E} \|\nabla \log p_{t_k}(Y_{t_k})\|^4 + \mathbb{E} \|\nabla^2 \log p_t(Y_t)\|^4} \sqrt{\mathbb{P}(\rho(Y_t, Y_{t_k}) > \omega/3)} \\ &\leq CK^2 d \left(\frac{1}{t} + dK^2 \right)^2 \sup_{t \leq s \leq t_k} \sqrt{\mathbb{E} \log^4 \frac{\sup p_{s/2}}{p_s(Y_s)}} \cdot \left(1 + \omega^{-2} \sqrt{\mathbb{P}(\rho(Y_t, Y_{t_k}) > \omega/3)} \right). \quad (26) \end{aligned}$$

Here the last line follows from Lemma 12.

Step III. Controlling expectations via Chebyshev and Li-Yau estimates. To bound $\mathbb{E} \log^4 \frac{\sup p_{s/2}}{p_s(Y_s)}$, we note that

$$\mathbb{E} \left(\frac{1}{p_t(Y_t)} \right) = \int \frac{1}{p_t} p_t d\mu = \int d\mu = 1.$$

By Chebyshev's inequality, we have

$$\mathbb{P} \left(\frac{1}{p_t(Y_t)} \geq \lambda \right) \leq \lambda^{-1}, \quad \lambda > 0,$$

and then

$$\mathbb{P} \left(\log^4 \frac{1}{p_t(Y_t)} \geq \lambda \right) \leq e^{-\sqrt[4]{\lambda}}, \quad \lambda \geq 0.$$

Integrate with respect to λ , we see

$$\mathbb{E} \log^4 \frac{1}{p_t(Y_t)} \leq C.$$

We then apply Li-Yau's estimate (Lemma 8) combined with Lemma 9, Lemma 11 to obtain $\sup \log p_{s/2} \leq \sup \log H(s/2, x, y) \lesssim d \log \frac{d}{s} + Ks$, where the first inequality follows from $p_{s/2}$ being the convolution of p_0 with $H(s/2, x, y)$. These together shows

$$\mathbb{E} \log^4 \frac{\sup p_{s/2}}{p_s(Y_s)} \leq \left(Cd \log \frac{d}{s} + CKs \right)^4 \leq \frac{Cd^5 K^4}{\delta}.$$

Here we used $\log \frac{d}{s} \leq C \left(\frac{d}{s} \right)^{1/4}$, and $s \geq t \geq \delta$.

Step IV. Controlling exit probability via stopping time. It remains to bound the probability $\mathbb{P}(\rho(Y_t, Y_{t_k}) > \omega/3)$. This would follow from a stopping time argument. We claim that given $t_k - t \leq h$, we have

$$\mathbb{P}(\rho(Y_t, Y_{t_k}) > \omega/3) \leq \exp \left(-\frac{c\omega^2}{t_k - t} \right) \leq \omega^4. \quad (27)$$

where the last inequality follows from (3). Plug this back into the desired conclusion of the lemma is proved.

We now prove (27). Let σ be the largest $t \leq t_k$ such that $\rho(Y_t, Y_{t_k}) > \omega/3$. We have

$$\mathbb{P}(\rho(Y_t, Y_{t_k}) > \omega/3) \leq \mathbb{P}(\sigma \geq t).$$

In the interval $[\sigma, t_k]$, Y_t stays in the geodesic ball $B_{Y_{t_k}}(\omega/3)$, and follows the SDE (2). In normal coordinates, this can be spelled out explicitly:

$$\begin{aligned} dY_t^\alpha &= \nabla^\alpha \log p_t(Y_t) dt + A_\beta^\alpha(Y_t) \circ dW_t^\beta \\ &= \left(\nabla^\alpha \log p_t(Y_t) + \frac{1}{2} (\partial_\gamma A_\beta^\alpha) A^{\beta\gamma} \right) dt + A_\beta^\alpha dW_t^\beta, \end{aligned} \quad (28)$$

where A is the square root of the matrix representing the coefficients of the Laplace-Beltrami operator

$$\frac{1}{\sqrt{\det g}} \partial_\alpha (\sqrt{\det g} \cdot g^{\alpha\beta} \partial_\beta).$$

It can be checked with the help of Lemma 5 that $\|A - I\| \leq CdK\omega^2$, and $\|\partial_\alpha A\| \leq Cd^2K\omega$; we omit the computation that has a similar pattern as many of the previous arguments. Furthermore, we have $\|\nabla \log p_t(Y_t)\| \lesssim (\delta^{-1} + K) \log \frac{\sup p_{t/2}}{p_t}$ by the same argument via Lemma 12 as before. This time we combine the uniform bound provided by Lemma 10 with Lemma 8 to conclude $\log \frac{\sup p_{t/2}}{p_t} \lesssim (\delta^{-1} + K + d \log d)^2 \text{Diam}(\mathcal{M})^2 \lesssim (\delta^{-1} + K + d \log d)^2 K^2$ by Assumption 1. This shows

$$\sup \|\nabla \log p_t\| \lesssim (\delta^{-1} + K + d \log d)^3 K^2. \quad (29)$$

Therefore, in view of Lemma 5 to convert the above bound to normal coordinates, and together with the aforementioned bound for A and $\partial_\alpha A$, we see that the drift term up to time σ will not exceed

$$\sup \left\| \nabla^\alpha \log p_t(Y_t) + \frac{1}{2} (\partial_\gamma A_\beta^\alpha) A^{\beta\gamma} \right\| \cdot (t_k - \sigma) \leq C(\delta^{-1} + K + d \log d)^3 K^2 (t_k - \sigma) \leq \frac{\omega}{12}, \quad (30)$$

where the last inequality used (3). On the other hand, the bound on A implies that the quadratic variation of the martingale part does not exceed

$$\int_\sigma^{t_k} A_\gamma^\alpha A_\beta^\gamma dt \leq 2(t_k - \sigma) I.$$

By Burkholder-Davis-Gundy inequality (Revuz & Yor, 2013), the tail of $\int_\sigma^{t_k} A_\beta^\alpha dW_t^\beta$ is $O(1)$ -subgaussian, thus we have

$$\mathbb{P} \left(\sigma \geq t, \left\| \int_\sigma^{t_k} A_\beta^\alpha dW_t^\beta \right\| > \frac{\omega}{12} \right) \leq \exp \left(\frac{-c(\omega - 2\sqrt{\rho(t_k - t)})^2}{t_k - t} \right).$$

In view of (3), combine this with (30) and (28), we have proved (27) as claimed. \square

Lemma 18. *Under the same assumptions as in Theorem 1 and assuming (3) without loss of generality, the discretization error obeys the following upper bound:*

$$\sum_{k=1}^N \int_{t_{k-1}}^{t_k} \mathbb{E} \|\nabla \log p_t(Y_t) - \mathcal{S}_{t_k, Y_{t_k}}^*(Y_t)\|^2 dt \leq \frac{Cd^6 K^8}{\delta^3} h^2 N,$$

where $C > 0$ is a universal constant.

Proof. This follows directly from Lemma 17. \square

F BROWNIAN MOTION SIMULATION ERROR

In this section, we handle the Brownian motion simulation error using the machinery of Minakshisundaram-Pleijel parametrix. A complete introduction to this heavy machinery would require establish a whole system of notations and lemmas in geometric analysis, which is unduly burdensome. We instead refer the interested reader to [Berline et al. \(2003\)](#) for a comprehensive treatment, and point to results there whenever needed.

F.1 OVERVIEW

Our aim is to prove the following lemma.

Lemma 19. *Under the same assumptions as in Theorem 1, and assuming (3) without loss of generality, we have*

$$\text{TV}(p_0^{\text{aux}}, q_0^*) \leq \sqrt{hT} \text{ poly}(d, K, \delta^{-1}).$$

To better explain the idea of the proof, we ignore the rejection sampling procedure in the construction of \hat{K}_k temporarily. Our starting point is the observation that by Fokker-Planck equation, \hat{K}_k is the heat kernel associated to the Euclidean Laplacian with drift $\mathcal{S}_{t_k, Y_{t_k}}$, in normal coordinates. On the other hand, K_k^{aux} is also a heat kernel with the same drift, but associated to the manifold Laplace-Beltrami operator. The following lemma shows that the two solutions coincide up to first order in time, at least in a polynomially small neighborhood of initial point and in a polynomially short time.

Lemma 20. *Let $F^{\mathcal{H}}(t, x, y)$ be the (generalized) heat kernel for the operator $\mathcal{H} = \frac{1}{2}\Delta_{\mathcal{M}} + \langle \mathcal{S}_{t_k, Y_{t_k}}, \nabla \rangle$. Define the Euclidean density*

$$\varphi_t(u; x) := \frac{1}{(2\pi t)^{d/2}} \exp\left(-\frac{\|u - \mathcal{S}_{t_k, Y_{t_k}}(x)t\|^2}{2t}\right), \quad u \in T_x \mathcal{M},$$

and let $\Phi(t, x, y)$ be the density of the push-forward by \exp_x of $\eta_{\omega} \varphi_t(\cdot; x)$ with respect to the volume measure, where η_{ω} is the cutoff function defined in (4). Then there exists polynomial $\text{poly}(d, K)$ with universally constant coefficients, such that for all $0 < t \leq \frac{1}{\text{poly}(d, K, \delta^{-1})}$ and for all $\rho(y, Y_{t_k}) \leq t^{5/12}$, we have

$$\left| \frac{F^{\mathcal{H}}(t, Y_{t_k}, y)}{\Phi(t, Y_{t_k}, y)} - 1 \right| \leq \text{poly}(d, K, \delta^{-1}) t.$$

With Lemma 20 in hand, it is tempting to calculate the KL error with the following heuristic:

$$\begin{aligned} \text{KL}(p_k^{\text{aux}} K_k^{\text{aux}} \| p_k^{\text{aux}} \hat{K}_k) &\lesssim -\mathbb{E} \int F^{\mathcal{H}}(h, Y_{t_k}, \cdot) \log \frac{\Phi(h, Y_{t_k}, \cdot)}{F^{\mathcal{H}}(h, Y_{t_k}, \cdot)} \\ &\lesssim \mathbb{E} \int F^{\mathcal{H}}(h, Y_{t_k}, \cdot) \left(\frac{F^{\mathcal{H}}(h, Y_{t_k}, \cdot)}{\Phi(h, Y_{t_k}, \cdot)} - 1 \right)^2 \\ &\lesssim \text{poly}(d, K, \delta^{-1}) h^2, \end{aligned}$$

where we ignore the fact that Lemma 20 holds only in a small neighborhood; the first line is post-processing inequality, and the second line stems from the fact that for two distributions p, q , we have

$$\int p \log \frac{q}{p} = \int p \log \left(1 + \frac{q-p}{p} \right)$$

$$\begin{aligned} &\geq \int p \left(\frac{q-p}{p} - C \frac{(q-p)^2}{p^2} \right) \\ &= -C \int p \left(\frac{q}{p} - 1 \right)^2, \end{aligned}$$

given $\frac{q}{p} - 1$ is sufficiently small, where the last line follows from $\int p = \int q = 1$. From this, we conclude that the accumulated error along N steps is bounded by $\text{poly}(d, K)h^2N = \text{poly}(d, K)hT$, and the desired bound follows from Pinsker's inequality.

Apart from Lemma 20, the above computation is the essence of this proof. The rest of this section is mainly devoted to proving Lemma 20, and then formalizing the above computation by handling exceptional events of exiting the polynomially small neighborhood.

F.2 PROOF OF LEMMA 20: A PARAMETRIX ESTIMATE

We begin the proof of Lemma 20. For simplicity, denote by v^α the normal coordinate representation of $\hat{s}_{t_k}(Y_{t_k})$. Naturally, our initial test solution is the drifted heat kernel, as simulated by our discretized process:

$$\varphi_t(u) := \frac{1}{(2\pi t)^{d/2}} \exp\left(-\frac{\|u - vt\|^2}{2t}\right), \quad u \in \mathbb{R}^d.$$

Before we compare this with the manifold heat kernel, there is one subtlety we need to keep in mind. The density φ_t is with respect to the *Lebesgue* measure on $T_x\mathcal{M}$, not with respect to the *volume* on \mathcal{M} . We compute and define the corresponding density on \mathcal{M} as follows:

$$\Phi(t, x, y) = \varphi_t(\log_x y) \sqrt{\Delta(x, y)}, \quad \Delta(x, y) := \frac{|\det d\log_x y|^2}{\det g(y)}, \quad y \in B_x(\omega).$$

Here all quantities are computed in normal coordinates. The factor $\Delta(x, y)$ is known as the van Vleck-Morette determinant. From Lemma 5 and Lemma 6, we know that

$$\frac{1}{2} \leq \Delta(x, y) \leq 2, \quad \text{if } \rho(x, y) \leq \frac{c}{Kd}. \quad (31)$$

We consider the generalized Laplacian

$$\mathcal{H} := \frac{1}{2} \Delta_{\mathcal{M}} + \langle \mathcal{S}_{t_k, Y_{t_k}}, \nabla \rangle.$$

As in Lemma 20, denote by $F^{\mathcal{H}}$ the heat kernel of \mathcal{H} at time $t_k - t_{k-1}$. We also propose an approximation of $F^{\mathcal{H}}$ by

$$\Psi(t, x, y) := G(t, x, y) \exp(\psi(x, y)) \sqrt{\Delta(x, y)}, \quad G(t, x, y) := \frac{1}{(2\pi t)^{d/2}} \exp\left(-\frac{d^2(x, y)}{2t}\right),$$

where for any two point $x, y \in \mathcal{M}$, letting $\gamma : [0, 1] \rightarrow$ be a constant-speed geodesic connecting x to y , we define

$$\psi(x, y) := \int_0^1 \left\langle \mathcal{S}_{t_k, Y_{t_k}}(\gamma(s)), \frac{d}{ds} \gamma(s) \right\rangle_g ds.$$

The auxiliary function Ψ bridges $F^{\mathcal{H}}$ and Φ in the following sense. On the one hand, we relate Φ and Ψ with the following lemma:

Lemma 21. *There exists a polynomial $\text{poly}(d, K)$ with universally constant coefficients such that the following holds. For any $0 < r \leq \frac{1}{\text{poly}(d, K, \delta^{-1})}$, $0 < t \leq \frac{1}{\text{poly}(d, K, \delta^{-1})}$ and for all $x, y \in B_{Y_{t_k}}(r)$, we have*

$$\left| \frac{\Phi(t, x, y)}{\Psi(t, x, y)} - 1 \right| \leq \text{poly}(d, K, \delta^{-1})(r^3 + t).$$

On the other hand, we have the following asymptotic expansion:

$$F^{\mathcal{H}}(t, x, y) = \Psi(t, x, y) \cdot \left(1 + \sum_{i=1}^{\infty} t^i u_i(x, y) \right), \quad t \rightarrow 0^+,$$

where u_i are smooth functions that can be computed explicitly via a recursive formula (Berline et al., 2003). We will not need the formula here, but instead require u_1 and the remainder terms to be bounded properly. Such bounds have been well-established, which we wrap up into the following lemma. Recall the cutoff function η_{ω} with radius ω defined in (4). It is clear we can replace ω with any $\iota > 0$ to define a cutoff η_{ι} of radius ι .

Lemma 22 (adapted from Berline et al. (2003)). *Fix a positive $\iota \leq 1/\text{poly}(d, K)$. There exists a smooth function $u_1(x, y)$ on $\mathcal{M} \times \mathcal{M}$ such that*

$$\|u_1\|_{\infty} + \|\nabla_y u_1\|_{\infty} \leq \text{poly}(d, K, \delta^{-1}),$$

and for all $0 < t \leq 1/\text{poly}(d, K)$, $y \in B_x(\iota)$, we have

$$|(\partial_t - \mathcal{H}) [\eta_{\iota}(\rho(x, y)) \Psi(t, x, y) (1 + tu_1(x, y))]| \leq r_{\eta}(t, x, y) + r_{\psi}(t, x, y), \quad (32)$$

where

$$r_{\eta}(t, x, y) \leq \frac{1}{\iota t} \text{poly}(d, K) \mathbb{1}_{\frac{\iota}{2} \leq \rho(x, y) \leq \iota} G(t, x, y), \quad (33a)$$

$$r_{\psi}(t, x, y) \leq t \cdot \text{poly}(d, K, \delta^{-1}) \mathbb{1}_{\rho(x, y) \leq \iota} G(t, x, y). \quad (33b)$$

Proof. The inequality on u_1 follows from Theorem 2.26 in Berline et al. (2003), with \mathcal{H} the same as our \mathcal{H} and therefore $F = \langle \mathcal{S}_{t_k, Y_{t_k}}, \nabla \rangle$. Note that all the coefficients in \mathcal{H} are bounded in C^2 by $\text{poly}(d, K)(1 + \|\mathcal{S}_{t_k, Y_{t_k}}\|_{C^2(\mathcal{M})})$, which is further bounded by $\text{poly}(d, K, \delta^{-1})$ as we will show momentarily. In fact, by Lemma 5 and Assumption 1, (A3), we have

$$\|\mathcal{S}_{t_k, Y_{t_k}}\|_{C^2(\mathcal{M})} = \|\mathcal{S}_{t_k, Y_{t_k}}\|_{\infty} + \|\nabla \mathcal{S}_{t_k, Y_{t_k}}\|_{\infty} + \|\nabla^2 \mathcal{S}_{t_k, Y_{t_k}}\|_{\infty} \leq \text{poly}(d, K) \omega^{-2} (1 + \|\nabla \log p_{t_k}(Y_{t_k})\|),$$

and then we control $\|\nabla \log p_{t_k}(Y_{t_k})\| \leq \text{poly}(d, K, \delta^{-1})$ via (29), yielding the claimed bound (recall that $\omega^{-1} = \text{poly}(d, K)$ by definition).

We proceed to prove (32). We follow the proof of Theorem 2.29, item (iii) in Berline et al. (2003), and choose the cutoff function ψ there to be η_{ι} as defined in (4), and with the differential operator B defined in Berline et al. (2003), we have

$$|(\partial_t - \mathcal{H}) [\Psi(t, x, y) (1 + tu_1(x, y))]| \leq \underbrace{\frac{1}{\iota t} \text{poly}(d) G(t, x, y)}_{=: r_{\eta}, \nabla \eta_{\iota} \text{ related terms}} + \underbrace{t \cdot G(t, x, y) \cdot |(B_y u_2)(x, y)|}_{=: r_{\psi}}.$$

Here B_y can be viewed as a coordinate-transformed version of \mathcal{H} , applied to the variable y (precise definition can be found in the reference), and u_2 is the second order term in the expansion. Similar to the argument we used to bound $\|u_1\|$, in virtue of Theorem 2.26 in Berline et al. (2003), we have

$$\|\mathcal{B}_x u_2\|_{L^{\infty}(\mu)} \leq \text{poly}(d, K, \delta^{-1}).$$

The claimed bound follows from combining the above inequalities. \square

We now state the Volterra series representation of heat kernel.

Lemma 23 (Volterra series, Theorem 2.23 in Berline et al. (2003)). *Fix a $\iota > 0$. Let*

$$\begin{aligned} \Psi_1(t, x, y) &:= \eta_{\iota}(\rho(x, y)) \Psi(t, x, y) (1 + tu_1(x, y)), \\ r_1(t, x, y) &:= (\partial_t - \mathcal{H}) \Psi_1(t, x, y). \end{aligned}$$

Define the time-space convolution operator $*$ as

$$(f * g)(t, x, y) = \int_0^t \int f(t-s, x, z) g(s, z, y) \mu(dz) ds.$$

Then we have

$$F^{\mathcal{H}} = \Psi_1 + \sum_{k=1}^{\infty} (-1)^k \Psi_1 * r_1^{*k}, \quad \text{where} \quad r_1^{*k} := \underbrace{r_1 * \dots * r_1}_{k \text{ times}},$$

on any domain such that the series on the right hand side converges absolutely uniformly.

Lemma 24 (Iterative bounds for Volterra series). *There exists a polynomial $\text{poly}(d, K, \delta^{-1})$ with universally constant coefficients such that the following holds. Assume $0 < t \leq 1/\text{poly}(d, K, \delta^{-1})$, take $\iota = 4t^{5/12}d$ in the definition of Ψ . Then we have, for all $\rho(x, y) \leq t^{5/12} = \iota/(4d)$, that*

$$\sum_{k=1}^{\infty} |\Psi_1 * r_1^{*k}|(t, x, y) \leq t \cdot \text{poly}(d, K, \delta^{-1}) G(t, x, y).$$

Proof. Recall Lemma 22, and denote by P the polynomial factor $\text{poly}(d, K)$ therein. Denote by $\lambda\Delta^k$ the dilated standard simplex

$$\lambda\Delta^k = \{(s_1, \dots, s_{k+1}) : s_i \geq 0, \sum_{i=1}^{k+1} s_i = \lambda\}, \quad \lambda > 0.$$

Fix some $k \geq 1$. Set $z_0 = x$ and $z_k = y$, we have

$$\begin{aligned} & |\Psi_1 * r_1^{*k}|(t, x, y) \\ & \leq t^k P^k \int_{t\Delta^{k-1}} \mathrm{d}s \int_{\mathcal{M}^{k-1}} \left(\prod_{i=1}^k G(s_i, z_{i-1}, z_i) \left(\mathbb{1}_{\rho(z_{i-1}, z_i) \leq \iota} + \frac{1}{\iota s_i} \mathbb{1}_{\rho(z_{i-1}, z_i) > \iota/2} \right) \right) \mu^{\otimes(k-1)}(\mathrm{d}z). \end{aligned} \quad (34)$$

We split the integral in (34) into a *local* part and an *outlier* part. Define a small “local” region

$$\mathcal{R} := \{(z_1, \dots, z_{k-1}) : \rho(z_i, x) \leq 2\iota, \rho(z_{i-1}, z_i) \leq \iota/2, i = 1, \dots, k\}.$$

We further define

$$\begin{aligned} I_{\text{loc}}(s) &:= \int_{\mathcal{R}} \left(\prod_{i=1}^k G(s_i, z_{i-1}, z_i) \right) \mu^{\otimes(k-1)}(\mathrm{d}z), \\ I_{\text{out}}(s) &:= \int_{\mathcal{R}^c} \left(\prod_{i=1}^k G(s_i, z_{i-1}, z_i) \mathbb{1}_{\rho(z_{i-1}, z_i) \leq \iota} \left(1 + \frac{1}{\iota s_i} \mathbb{1}_{\rho(z_{i-1}, z_i) > \iota/2} \right) \right) \mu^{\otimes(k-1)}(\mathrm{d}z). \end{aligned}$$

It is clear that

$$|\Psi_1 * r_1^{*k}|(t, x, y) \leq t^k P^k \int_{t\Delta^{k-1}} (I_{\text{loc}}(s) + I_{\text{out}}(s)) \mathrm{d}s. \quad (35)$$

We will establish bounds for I_{loc} and I_{out} respectively.

Bounding the local integral. For ease of understanding, we begin by computing the first integral in I_{loc} with respect to z_1 . Extracting the factors containing z_1 , we need to calculate

$$\int_{\{\rho(z_1, x) \leq 2\iota\}} \frac{1}{(2\pi s_1)^{d/2}} \frac{1}{(2\pi s_2)^{d/2}} \exp \left(-\frac{\rho(x, z_1)^2}{2s_1} - \frac{\rho(z_1, z_2)^2}{2s_2} \right) \mu(\mathrm{d}z_1). \quad (36)$$

To proceed, we will invoke Lemma 13. Let z_* be a minimizer of

$$V(z) = \frac{\rho(x, z)^2}{s_1(s_1 + s_2)^{-1}} + \frac{\rho(z, z_2)^2}{s_2(s_1 + s_2)^{-1}} - \rho(x, z_2)^2.$$

By Lemma 13, $V(z_*) = 0$. Since $V(z) > 0$ whenever $\rho(x, z) \geq \rho(x, z_2)$, we know that $z_* \in B_x(2\iota)$. Moreover, Lemma 13 and Lemma 5 together imply

$$V(z) \geq (1 - Cd^2K^2\iota) \cdot \frac{(s_1 + s_2)^2}{s_1 s_2} \rho(z, z_*)^2, \quad \forall z \in B_x(4\iota).$$

Here, the second inequality follows from strong convexity given by Lemma 13 and a comparison of geometric distance and Euclidean distance in normal coordinates fueled by Lemma 5. Denote for the moment that

$$\theta := 1 - Cd^2K^2\iota.$$

Plugging this back into (36), we obtain

$$\begin{aligned}
& \int_{\{\rho(z_1, x) \leq 2\iota\}} \frac{1}{(2\pi s_1)^{d/2}} \frac{1}{(2\pi s_2)^{d/2}} \exp\left(-\frac{\rho(x, z_1)^2}{2s_1} - \frac{\rho(z_1, z_2)^2}{2s_2}\right) \mu(dz) \\
&= \int_{\{\rho(z_1, x) \leq 2\iota\}} \frac{1}{(2\pi s_1)^{d/2}} \frac{1}{(2\pi s_2)^{d/2}} \exp\left(-\frac{V(z_1)}{2(s_1 + s_2)}\right) \mu(dz_1) \\
&\leq \int_{\{\rho(z_1, x) \leq 2\iota\}} \frac{1}{(2\pi(t-s))^{d/2}} \frac{1}{(2\pi s)^{d/2}} \exp\left(-\frac{\theta(s_1 + s_2)\rho(z_1, z_*)^2}{2s_1 s_2} - \frac{\rho(x, z_2)^2}{2(s_1 + s_2)}\right) \mu(dz_1) \\
&\leq 4 \exp\left(-\frac{\rho(x, z_2)^2}{2(s_1 + s_2)}\right) \cdot \int \frac{1}{(2\pi s_1)^{d/2}} \frac{1}{(2\pi s_2)^{d/2}} \exp\left(-\frac{\theta(s_1 + s_2)\|Z\|^2}{2s_1 s_2}\right) dZ \\
&= 4(2\pi(s_1 + s_2))^{d/2} G(s_1 + s_2, x, z_2) \cdot \frac{1}{(2\pi s_1)^{d/2}} \frac{1}{(2\pi s_2)^{d/2}} \left(2\pi \cdot \frac{s_1 s_2}{\theta(s_1 + s_2)}\right)^{d/2} \\
&\leq 4\theta^{-d/2} G(s_1 + s_2, x, z_2),
\end{aligned}$$

where the third-to-last line follows from change of variable to normal coordinates at z_* and from using (31) to bound the determinant; the penultimate line follows from Gaussian integration. Now, we note that

$$\theta^{-d/2} = (1 - Cd^2 K^2 \iota)^{-d/2} \leq \exp(2Cd^3 K^2 \iota) \leq 2,$$

give $\iota \leq \frac{1}{100Cd^3K^2}$. Putting these pieces together, we proved

$$\int_{\{\rho(z_1, x) \leq 2\iota\}} \frac{1}{(2\pi s_1)^{d/2}} \frac{1}{(2\pi s_2)^{d/2}} \exp\left(-\frac{\rho(x, z_1)^2}{2s_1} - \frac{\rho(z_1, z_2)^2}{2s_2}\right) \mu(dz_1) \leq 8G(s_1 + s_2, x, z_2).$$

Iterate the above argument for the integration over z_2, \dots, z_{k-1} to obtain

$$I_{\text{loc}}(s) \leq 8^k \cdot G\left(\sum_{i=1}^k s_i, x, y\right) = 8^k \cdot G(t, x, y). \quad (37)$$

Bounding the outlier integral. Next, we show how to control I_{out} . We first write

$$\prod_{i=1}^k G(s_i, z_{i-1}, z_i) = \frac{1}{\prod_{i=1}^k (2\pi s_i)^{d/2}} \exp\left(-\sum_{i=1}^k \frac{\rho(z_{i-1}, z_i)^2}{2s_i}\right). \quad (38)$$

We claim that for any $z \in \mathcal{R}^c$, we have

$$T := \prod_{i=1}^k \exp\left(-\sum_{i=1}^k \frac{\rho(z_{i-1}, z_i)^2}{2s_i(d+1)}\right) \mathbb{1}_{\rho(z_{i-1}, z_i) \leq \iota} \left(1 + \frac{1}{\iota s_i} \mathbb{1}_{\rho(z_{i-1}, z_i) > \iota/2}\right) \leq \exp\left(-\frac{\iota^2}{16td}\right). \quad (39)$$

The claim is proved at the end of this proof. It is tempting to plug this back into (38), and argue that when t is polynomially small, the integrand in I_{out} becomes exponentially small. However, this would not work since it does not resolve the singular factors $\prod_{i=1}^k s_i^{-d/2}$ in the integrand. For this purpose, we need the following crucial ‘‘freezing’’ trick, which follows trivially from $1 = \frac{1}{1+d^{-1}} + \frac{1}{d+1}$:

$$\begin{aligned}
& \prod_{i=1}^k G(s_i, z_{i-1}, z_i) \mathbb{1}_{\rho(z_{i-1}, z_i) \leq \iota} \left(1 + \frac{1}{\iota s_i} \mathbb{1}_{\rho(z_{i-1}, z_i) > \iota/2}\right) \\
&= \left(\prod_{i=1}^k (1+d^{-1})^{d/2} G((1+d^{-1})s_i, z_{i-1}, z_i) \mathbb{1}_{\rho(z_{i-1}, z_i) \leq \iota} \right) T.
\end{aligned}$$

The idea is to keep the Gaussian behavior to resolve the $s_i^{-d/2}$ factors, and only single out a very small proportion to demonstrate exponential smallness. Plug (39) into the above identity to obtain

$$I_{\text{out}}(s) \leq \exp\left(-\frac{\iota^2}{16td}\right) \cdot \int_{\mathcal{R}^c} \left(\prod_{i=1}^k (1+d^{-1})^{d/2} G((1+d^{-1})s_i, z_{i-1}, z_i) \mathbb{1}_{\rho(z_{i-1}, z_i) \leq \iota} \right) \mu^{\otimes(k-1)}(dz).$$

Integrate successively for each variable, convert to normal coordinates, and apply Lemma 5, Eqn. (31), and Gaussian integration as we did in bounding I_{loc} , we obtain

$$\int_{\mathcal{R}^c} \left(\prod_{i=1}^k G((1+d^{-1})s_i, z_{i-1}, z_i) \mathbf{1}_{\rho(z_{i-1}, z_i) \leq \iota} \right) \mu^{\otimes(k-1)}(dz) \leq 8^k (1+d^{-1})^{dk/2} (ct)^{-d/2} \leq 32^k (ct)^{-d/2}.$$

Therefore

$$I_{\text{out}}(s) \leq 32^k (ct)^{-d/2} \cdot \exp\left(-\frac{\iota^2}{16td}\right) \leq 32^k G(t, x, y), \quad (40)$$

where in the last inequality we used the assumption $\rho(x, y) \leq t^{5/12} = \iota/(4d)$ and $t \leq 1/\text{poly}(d, K)$, so that $\exp(-\iota^2/(32td)) \leq \exp(-\rho(x, y)^2/(2t))$ and $\exp(-\iota^2/(32td)) = \exp(-\frac{1}{2}dt^{-1/6}) \leq (2\pi t)^{-d/2}$.

Putting things together. We plug the bounds (37) and (40) into (35) to obtain

$$|\Psi_1 * r_1^{*k}|(t, x, y) \leq t^k P^k \int_{t\Delta^{k-1}} (8^k + 32^k) G(t, x, y) ds \leq \frac{40}{(k-1)!} (40t)^{2k-1} P^k G(t, x, y).$$

The desired conclusion of Lemma 24 follows from the above inequality by summing over k and taking $t \leq 1/\text{poly}(d, K)$.

Proof of Claim (39). For $z \in \mathcal{R}^c$, let

$$J = \{i : \rho(z_{i-1}, z_i) > \iota/2\}.$$

By definition of \mathcal{R}^c , either J is nonempty, or there is i_0 such that $\rho(x, z_{i_0}) > 2\iota$. For $i \in J$, we note that

$$\exp\left(-\frac{\rho(z_{i-1}, z_i)^2}{2s_i d}\right) \left(1 + \frac{1}{\iota s_i}\right) \leq \exp\left(-\frac{\iota^2}{8s_i d}\right) \left(1 + \frac{1}{\iota s_i}\right) \leq \exp\left(-\frac{\iota^2}{16td}\right), \quad (41)$$

where the last inequality follows from $s_i \leq t$, $\iota = 4t^{5/12}d$, and that $t \leq 1/\text{poly}(d, K)$. When J is nonempty, we readily deduce (39) as all the other factors are ≤ 1 .

When J is empty, let i_0 be such that $\rho(x, z_{i_0}) > 2\iota$. We apply Cauchy-Schwarz to obtain

$$\sum_{i=1}^k \frac{\rho(z_{i-1}, z_i)^2}{s_i} \geq \frac{1}{\sum_{i=1}^k s_i} \left(\sum_{i=1}^k \rho(z_{i-1}, z_i) \right)^2 = \frac{1}{t} \left(\sum_{i=1}^k \rho(z_{i-1}, z_i) \right)^2.$$

Then, by triangle inequalities, we have

$$\sum_{i=1}^{i_0} \rho(z_{i-1}, z_i) \geq \rho(z_0, z_{i_0}) = \rho(x, z_{i_0}) > 2\iota, \quad \text{therefore} \quad \sum_{i=1}^k \frac{\rho(z_{i-1}, z_i)^2}{s_i} \geq \frac{2\iota^2}{t}.$$

The desired claim (39) follows immediately, given that J is empty. \square

We now have all the ingredients to prove Lemma 20.

Proof of Lemma 20. This follows immediately from Lemma 21, Lemma 23 and Lemma 24. Note that in applying Lemma 21, we used $r \leq t^{5/12}$, thus $r^3 \leq t^{5/4} \leq t$. \square

F.3 PROOF OF LEMMA 21

By definition, we can compute

$$\frac{\Phi(t, x, y)}{\Psi(t, x, y)} = \exp((\log_x y) \cdot v - \psi(x, y)) \exp(-\|v\|^2 t/2).$$

Note that $\|v\| \leq C\|\nabla \log p_{t_k}(Y_{t_k})\|$ by Lemma 5, which in turn is bounded by $\text{poly}(d, K, \delta^{-1})$ by (29). When $t \leq \frac{1}{\text{poly}(d, K, \delta^{-1})} \leq \frac{1}{4\|v\|^2}$, we have $|\exp(-\|v\|^2 t/2) - 1| \leq \|v\|^2 t \leq \text{poly}(d, K, \delta^{-1})t$. Therefore, it suffices to show

$$|(\log_x y) \cdot v - \psi(x, y)| \leq \text{poly}(d, K, \delta^{-1})r^3.$$

Recall the definition of ψ . Note that since $r \leq \frac{1}{\text{poly}(d, K, \delta^{-1})}$, when the polynomial $\text{poly}(d, K, \delta^{-1})$ is sufficiently large, the geodesic γ from x to y is unique and is inside $B_{Y_{t_k}}(r)$. In the normal coordinate on Y_{t_k} within radius r , the vector field $\mathcal{S}_{t_k, Y_{t_k}}$ is represented by the constant vector v . We also recognize that in normal coordinate, $\gamma(1) - \gamma(0) = \log_x y$. We thus have

$$\begin{aligned} |(\log_x y) \cdot v - \psi(x, y)| &= \left| \delta_{\alpha\beta} v^\alpha (\gamma^\beta(1) - \gamma^\beta(0)) - \int_0^1 g_{\alpha\beta}(\gamma(s)) v^\alpha \frac{d}{ds} \gamma^\beta(s) ds \right| \\ &= \left| \int_0^1 (g_{\alpha\beta}(\gamma(s)) - \delta_{\alpha\beta}) v^\alpha \frac{d}{ds} \gamma^\beta(s) ds \right| \\ &\leq C \int_0^1 \|g(\gamma(s)) - I\| \cdot \|v\| \cdot \left\| \frac{d}{ds} \gamma(s) \right\| ds \\ &\leq C \|v\| \cdot \rho(x, y) \int_0^1 CK(s\rho(x, y))^2 ds \\ &\leq \text{poly}(d, K, \delta^{-1}) \rho(x, y)^3, \end{aligned}$$

as desired. Here the penultimate line follows from Lemma 5.

F.4 PROOF OF LEMMA 19: HANDLING EXCEPTIONAL EVENTS

Proof of Lemma 19. Recall that $p_0^{\text{aux}} = p_N K_N^{\text{aux}} K_{N-1}^{\text{aux}} \cdots K_1^{\text{aux}}$ and $q_0^* = p_N \hat{K}_N \hat{K}_{N-1} \cdots \hat{K}_1$. We need to compare the kernel K_k^{aux} with \hat{K}_k , $k = 1, \dots, N$. To apply Lemma 20, we define two auxiliary kernels \tilde{K}_k^{aux} , \tilde{K}_k that are “localized” version of K_k^{aux} and \hat{K}_k . We show the auxiliary kernels are close to K_k^{aux} and \hat{K}_k respectively in total variation, and establish bound on $\text{KL}(\tilde{K}_k^{\text{aux}} \parallel \tilde{K}_k)$. Denote

$$\mathcal{R}_x := \{y \in \mathcal{M} : \rho(x, y) \leq h^{5/12}\}, \quad \mathcal{R}_x^c := \mathcal{M} \setminus \mathcal{R}_x.$$

Recall the notation Φ in the proof of Lemma 24. To distinguish the kernels at different step, we denote Φ_k as the corresponding Φ at step k . Define \tilde{K}_k^{aux} , \tilde{K}_k by

$$\begin{aligned} \tilde{K}_k^{\text{aux}}(x, dy) &= K_k^{\text{aux}}(x, dy) \mathbb{1}_{\mathcal{R}_x}(y) + \frac{K_k^{\text{aux}}(x, \mathcal{R}_x^c)}{\mu(\mathcal{R}_x^c)} \mu(dy) \mathbb{1}_{\mathcal{R}_x^c}, \\ \tilde{K}_k(x, dy) &= \Phi_k(h, x, y) \mathbb{1}_{\mathcal{R}_x}(y) + \frac{\int_{\mathcal{R}_x^c} \Phi_k(h, x, z) \mu(dz)}{\mu(\mathcal{R}_x^c)} \mu(dy) \mathbb{1}_{\mathcal{R}_x^c}, \end{aligned}$$

By converting to normal coordinate and invoking Gaussian integration in the same way as in the proof of Lemma 24, we obtain

$$\exp\left(-\frac{2}{h^{1/6}}\right) \leq \int_{\mathcal{R}_x^c} \Phi_k(h, x, z) \mu(dz) \leq \exp\left(-\frac{1}{16h^{1/6}}\right). \quad (42)$$

When $h \leq 1/\text{poly}(d, K, \delta^{-1})$, it is apparent (e.g., follows from Gromov’s volume comparison theorem) that $\mu(\mathcal{R}_x^c) \leq 1/2$, thus

$$\frac{1}{2} \leq \mu(\mathcal{R}_x^c) \leq \mu(\mathcal{M}) = 1.$$

We observe that \hat{K}_k differs from \tilde{K}_k by a rejection sampling with radius $h^{1/4}$. With the above bounds and the same Gaussian integration technique, we see that the probability of rejection is bounded by

$$\mathbb{P}(\text{rejection at step } k) \leq \exp\left(-\frac{(h^{1/4})^2}{16h}\right) \leq \exp\left(-\frac{1}{16h^{1/2}}\right). \quad (43)$$

Summing up, We readily obtain

$$\text{TV}(\widehat{K}_k, \widetilde{K}_k) \leq \exp\left(-\frac{1}{16h^{1/6}}\right). \quad (44)$$

On the other hand, by using the stopping time argument as in the proof of (27), we have

$$K_k^{\text{aux}}(x, \mathcal{R}_x^c) \leq \exp\left(-\frac{1}{16h^{1/6}}\right). \quad (45)$$

Therefore, the following TV bound is obvious:

$$\text{TV}(K_k^{\text{aux}}, \widetilde{K}_k^{\text{aux}}) \leq \exp\left(-\frac{1}{16h^{1/6}}\right), \quad (46)$$

Now we compute $\text{KL}(\widetilde{K}_k^{\text{aux}}(x, \cdot) \parallel \widetilde{K}_k(x, \cdot))$. By definition, we have

$$\text{KL}(\widetilde{K}_k^{\text{aux}}(x, \cdot) \parallel \widetilde{K}_k(x, \cdot)) = \underbrace{\int_{\mathcal{R}_x} \left(\log \frac{\widetilde{K}_k^{\text{aux}}(x, dy)}{\widetilde{K}_k(x, dy)} \right) \widetilde{K}_k^{\text{aux}}(x, dy)}_{=:T_1} + \underbrace{\left(\log \frac{K_k^{\text{aux}}(x, \mathcal{R}_x^c)}{\int_{\mathcal{R}_x^c} \Phi_k(h, x, z) \mu(dz)} \right) \widetilde{K}_k(x, \mathcal{R}_x^c)}_{=:T_2}.$$

We control the two terms separately.

Controlling T_1 . We invoke Lemma 20 to see

$$\left| \frac{\widetilde{K}_k^{\text{aux}}(x, dy)}{\widetilde{K}_k(x, dy)} - 1 \right| \leq \text{poly}(d, K, \delta^{-1})h.$$

Therefore, we use the elementary fact that $\log(1+x) \geq x - 2x^2$ for $x \in [-1/2, 1/2]$ to obtain

$$\begin{aligned} \log \frac{\widetilde{K}_k^{\text{aux}}(x, dy)}{\widetilde{K}_k(x, dy)} &= -\log \frac{\widetilde{K}_k(x, dy)}{\widetilde{K}_k^{\text{aux}}(x, dy)} \\ &\leq 1 - \frac{\widetilde{K}_k(x, dy)}{\widetilde{K}_k^{\text{aux}}(x, dy)} + 2 \left(\frac{\widetilde{K}_k(x, dy)}{\widetilde{K}_k^{\text{aux}}(x, dy)} - 1 \right)^2 \\ &\leq 1 - \frac{\widetilde{K}_k(x, dy)}{\widetilde{K}_k^{\text{aux}}(x, dy)} + \text{poly}(d, K, \delta^{-1})h^2, \end{aligned}$$

provided $h \leq 1/\text{poly}(d, K, \delta^{-1})$. Integrate with respect to $\widetilde{K}_k^{\text{aux}}(x, dy)$ over $y \in \mathcal{R}_x$ to obtain

$$\begin{aligned} T_1 &\leq \widetilde{K}_k^{\text{aux}}(x, \mathcal{R}_x) - \widetilde{K}_k(x, \mathcal{R}_x) + \text{poly}(d, K, \delta^{-1})h^2 \\ &\leq 2 \exp\left(-\frac{1}{16h^{1/6}}\right) + \text{poly}(d, K, \delta^{-1})h^2 \\ &\leq \text{poly}(d, K, \delta^{-1})h^2, \end{aligned}$$

where the second line follows from (42) and (45), and the last line follows from $h \leq 1/\text{poly}(d, K, \delta^{-1})$ so that the exponential term is sufficiently small.

Controlling T_2 . This is straightforward given (45) and (42). We obtain in the same way as above that

$$T_2 \leq \exp\left(-\frac{1}{32h^{1/6}}\right) \leq \text{poly}(d, K, \delta^{-1}) \leq h^2.$$

Summarizing the above, we have shown that

$$\text{KL}(\widetilde{K}_k^{\text{aux}}(x, \cdot) \parallel \widetilde{K}_k(x, \cdot)) \leq \text{poly}(d, K, \delta^{-1})h^2.$$

Accumulate the error over all N steps using post-processing inequality and apply Pinsker's inequality, we obtain

$$\text{TV}(p_0^{\text{aux}} \parallel q_0^*) \leq \sqrt{\text{poly}(d, K, \delta^{-1})h^2 N} \leq \sqrt{hT} \text{poly}(d, K, \delta^{-1}),$$

since $hN = T - \delta \leq T$, as claimed. \square

G PROOF OF MAIN RESULTS

Proof of Lemma 1. This follows from combining Lemma 15 and Lemma 16. \square

Proof of Lemma 2. This follows from Lemma 18 and our choice of schedule $hN = T - \delta \leq T$. \square

Proof of Theorem 1. This follows from Lemma 1, Lemma 2, and Lemma 19. \square