



Efficient Random Walks on Riemannian Manifolds

Simon Schwarz¹ · Michael Herrmann³ · Anja Sturm¹ · Max Wardetzky²

Received: 16 September 2022 / Revised: 17 October 2023 / Accepted: 31 October 2023
© The Author(s) 2023

Abstract

According to a version of Donsker's theorem, geodesic random walks on Riemannian manifolds converge to the respective Brownian motion. From a computational perspective, however, evaluating geodesics can be quite costly. We therefore introduce approximate geodesic random walks based on the concept of retractions. We show that these approximate walks converge in distribution to the correct Brownian motion as long as the geodesic equation is approximated up to second order. As a result, we obtain an efficient algorithm for sampling Brownian motion on compact Riemannian manifolds.

Keywords Brownian motion · Riemannian manifolds · Retractions · Geodesic random walks

Mathematics Subject Classification 65C30 · 60H35 · 58J65

Communicated by Carola-Bibiane Schönlieb.

✉ Simon Schwarz
simon.schwarz@uni-goettingen.de
Michael Herrmann
michael.herrmann@tu-braunschweig.de
Anja Sturm
anja.sturm@mathematik.uni-goettingen.de
Max Wardetzky
wardetzky@math.uni-goettingen.de

¹ Institute for Mathematical Stochastics, University of Göttingen, Goldschmidtstraße 7, 37077 Göttingen, Germany

² Institute for Numerical and Applied Mathematics, University of Göttingen, Lotzestr. 16-18, 37083 Göttingen, Germany

³ Institute for Partial Differential Equations, Technische Universität Braunschweig, Universitätsplatz 2, 38106 Brunswick, Germany

1 Introduction

Probabilistic models in continuous time with geometric constraints lead to Brownian motion and stochastic differential equations (SDEs) on Riemannian manifolds (M, g) . The theory of Brownian motion and SDEs on manifolds has been extensively studied (see, e.g., [16, 20])—and many results from the Euclidean setting can be generalized to manifold-valued SDEs. One example is Donsker’s theorem. Here, one considers a geodesic random walk for which the update step at the current location x_i is obtained by uniformly sampling a unit tangent vector $\bar{v} \sim \text{Unif}(S_{x_i})$ with $S_x := \{u \in T_x M : \|u\|_g = 1\}$ and following a geodesic path by setting $x_{i+1} := \text{Exp}_{x_i}(\varepsilon \bar{v})$ for some small enough parameter ε , where $v := \sqrt{m} \bar{v}$, m denotes the dimension of the manifold, and $\text{Exp}_x : T_x M \rightarrow M$ is the exponential map. We will discuss the reason for scaling by \sqrt{m} below. Jørgensen proved that on complete Riemannian manifolds that satisfy mild assumptions (which in particular hold for compact manifolds), the scaling limit of such a geodesic random walk is the Brownian motion on (M, g) , see [21]. This result has recently been generalized to the setting of Finsler manifolds, see [26].

In principle, it is thus possible to simulate Brownian motion on Riemannian manifolds by computing geodesic random walks—analogously to how one would use Donsker’s theorem in the Euclidean case. Unfortunately, this procedure is not computationally efficient on manifolds since the solution of a nonlinear second-order geodesic ODE has to be computed in every step. With an eye on applications, we are interested in efficient and simple, yet convergent, methods in order to approximate Brownian motion on Riemannian manifolds.

We introduce a class of algorithms that yield approximate geodesic random walks by making use of *retractions* that approximate the exponential map and have been introduced in the context of numerical optimization on manifolds, see [1–3, 8]. We show that retractions which approximate the geodesic ODE up to second order converge to Brownian motion in the Skorokhod topology, see Sect. 3. We consider two prevalent scenarios, where random walks based on such second-order retractions can be computed efficiently: the case of parameterized manifolds and the case of implicitly defined manifolds, see Sect. 2. Our approach generalizes the setting of approximate geodesic walks on manifolds with positive sectional curvatures presented in [27]. Moreover, our method works for arbitrary dimensions, and its cost only depends on the evaluation of the respective retraction.

A retraction in our sense can be understood as both, an approximation of the geodesic ODE, or, equivalently, as a 2-jet in the sense of [5], which allows for treating coordinate-free SDEs. Armstrong and King developed a 2-jet scheme for simulating SDEs on manifolds in [6], which uses an unbounded stepsize in each iteration. Different from our retraction-based approach, which is based on *bounded* stepsizes, the case of unbounded stepsizes hampers efficiency (for the case of local charts) or might even prohibit an implementation (for the case of implicitly defined submanifolds when using a step-and-project approach). Moreover, our random walk method is based on a Donsker-type theorem, while their approach improves Euler–Maruyama

approximations—resulting in a different type of convergence in terms of the respective topologies.

Outside the setting of 2-jets, one typically considers SDEs on embedded manifolds $M \subset \mathbb{R}^n$ and first solves them in ambient Euclidean space, followed by projecting back to M using Lagrange multipliers, see, e.g., [7, 25]. Projection-based approaches, however, do not immediately extend to the case of parameterized manifolds. Within our unified framework based on retractions, we cover both, the case of embedded manifolds and the case of parameterized manifolds, see Sects. 2.1 and 2.2. In particular, our approach leads to an efficient and convergent numerical treatment of Brownian motion on Riemannian manifolds (including drift—see Remark 2.4).

A different scenario where geodesic random walks are commonly used is the problem of sampling the volume measure of polytopes (defined as the convex hull resulting from certain linear constraints). Originally, Euclidean random walk methods had been considered in this context, see [30]. Although the Euclidean perspective is natural in this situation, small stepsizes are required in order to reach points near the boundary, which hampers efficiency in higher dimensions. To circumvent this issue, Lee and Vempala [24] introduced a Riemannian metric based on the Hessian of a log-barrier—to the effect of rescaling space as the walk approaches the boundary. In their implementation, they used collocation schemes for solving the requisite ODEs in order to simulate geodesic random walks. We leave for future research the question whether such approaches can be reformulated in terms of retractions.

Furthermore, the problem of sampling on manifolds is commonly addressed by Markov chain Monte Carlo (MCMC) methods, see, e.g., [10–12, 31]. Different from our focus, MCMC methods are typically concerned with obtaining samples from a given probability distribution on M . Our algorithm, however, yields approximate sample *paths* of a Brownian motion on a manifold. We are not aware of any results concerned with proving that an MCMC method samples the correct dynamics of a given SDE. Although not the main focus of our exposition, we show that our method also correctly recovers the stationary measure of a geodesic random walk in the limit of stepsize tending to zero, see Theorem 3.2.

2 Retraction-Based Random Walks

Throughout this exposition, we consider m -dimensional compact and orientable Riemannian manifolds (M, g) without boundary. We use *retractions* to approximate the exponential map on such manifolds.

Definition 2.1 Let $\text{Ret} : TM \rightarrow M$ be a smooth map, and denote the restriction to the tangent space $T_x M$ by Ret_x for any $x \in M$. Ret is a *retraction* if the following two conditions are satisfied for all $x \in M$ and all $v \in T_x M$:

1. $\text{Ret}_x(0) = x$, where 0 is the zero element in $T_x M$ and
2. $\frac{d}{d\tau} \text{Ret}_x(\tau v) \Big|_{\tau=0} = v$ (where we identify $T_0 T_x M \simeq T_x M$).

A retraction is a *second-order retraction* if it additionally satisfies that for all $x \in M$ and for all $v \in T_x M$ one has that

$$\frac{D}{d\tau} \left(\frac{d}{d\tau} \text{Ret}_x(\tau v) \right) \Big|_{\tau=0} = \frac{D}{d\tau} \left(\frac{d}{d\tau} \text{Exp}_x(\tau v) \right) \Big|_{\tau=0} = 0, \tag{1}$$

where $\frac{D}{d\tau} \left(\frac{d}{d\tau} \gamma(\tau) \right)$ denotes covariant differentiation of the tangent vector field $\dot{\gamma}(\tau) = \frac{d}{d\tau} \gamma(\tau)$ along the curve γ (following standard notation, see, e.g., [15]). If M is a submanifold of a Euclidean space, Eq. (1) is equivalent to

$$\frac{d^2}{d\tau^2} \text{Ret}_x(\tau v) \Big|_{\tau=0} \in \mathcal{N}_x M,$$

where $\mathcal{N}_x M$ denotes the normal bundle of M . Consequently, for the case of submanifolds, one has

$$\text{Ret}_x(\tau v) = \text{Exp}_x(\tau v) + \mathcal{O}(\tau^3)$$

for all $x \in M$ and $v \in T_x M$ as $\tau \rightarrow 0$.

Clearly, the exponential map is itself a retraction. The main benefit of retractions is, however, that they can serve as computationally efficient *approximations* of the exponential map. Using retractions yields Algorithm 1 for simulating random walks on a Riemannian manifold.

Algorithm 1 Retraction-based random walk

- 1: **Input:** Retraction Ret ; iterations N ; stepsize ε ; dimension m ; initial position x_0 ; set $x_0^\varepsilon := x_0$
 - 2: **for** $1 \leq i \leq N$ **do**
 - 3: Sample $\bar{v} \in S_{x_{i-1}^\varepsilon} := \{u \in T_{x_{i-1}^\varepsilon} M : \|u\|_g = 1\}$ uniformly wrt. the Riemannian metric g
 - 4: Set $v := \sqrt{m} \bar{v}$
 - 5: $x_i^\varepsilon := \text{Ret}_{x_{i-1}^\varepsilon}(\varepsilon v)$
 - 6: **end for**
 - 7: **Return:** $(x_j^\varepsilon)_{0 \leq j \leq N}$
-

Notice that the uniform sampling in Step 3 of the above algorithm is performed with respect to the Riemannian metric restricted to the *unit sphere* in the attendant tangent space. We provide details for this step in Sects. 2.1 and 2.2. Notice furthermore that ε takes the role of a *spatial* stepsize, while the total physical *time* simulated by Algorithm 1 is given by $\varepsilon^2 N$; see also the time rescaling in Theorem 2.2. Step 4 of the algorithm is necessary since uniform sampling on the unit sphere in a space of dimension m results in a covariance matrix $\frac{1}{m} \text{Id}$, while convergence to Brownian motion requires the identity as covariance matrix. The normalizing factor \sqrt{m} in Step 4 precisely ensures the latter. Indeed, without the renormalization in Step 4 the random walk constructed in Algorithm 1 converges to a time-changed Brownian motion.

The next theorem shows convergence of the retraction-based random walk resulting from Algorithm 1 and provides an explicit expression for the generator of the resulting limit process:

Theorem 2.2 Consider the sequence of random variables $(X_i^\varepsilon)_{i \in \mathbb{N}}$ constructed in Algorithm 1 (with $N = \infty$). The continuous-time process $X^\varepsilon := (X_{\lfloor \varepsilon^{-2}t \rfloor}^\varepsilon)_{t \geq 0}$ converges in distribution to the stochastic process with generator

$$(Lf)(x) = \frac{1}{2}(\Delta_g f)(x) + \frac{m}{2\omega_m} \int_{S_x} df|_x \left(\frac{D}{d\tau} \frac{d}{d\tau} \text{Ret}_x(\tau \bar{v}) \right) \Big|_{\tau=0} d\bar{v} \tag{2}$$

in the Skorokhod topology (see Sect. 3 for details). Here, S_x denotes the unit sphere in $T_x M$, ω_m is the volume of this sphere, and Δ_g is the Laplace–Beltrami operator on (M, g) .

We defer the proof to Sect. 3. Notice that Theorem 2.2 immediately implies a necessary and sufficient condition for a retraction-based random walk to converge to Brownian motion on a Riemannian manifold (M, g) :

Corollary 2.3 Using the same notation as in Theorem 2.2, the continuous-time process X^ε converges in distribution to the Brownian motion on M if and only if the Laplacian (also known as the tension) of $\text{Ret}_x : T_x M \rightarrow M$ vanishes at $0 \in T_x M$ for every $x \in M$. In particular, this holds for second-order retractions.

Proof By definition, the Laplacian of $\text{Ret}_x : T_x M \rightarrow M$ is defined as

$$\Delta \text{Ret}_x := \text{trace } \nabla d \text{Ret}_x,$$

where $\nabla d \text{Ret}_x$ is the Hessian of the C^∞ -map Ret_x . Clearly, the Laplacian of Ret_x vanishes at $0 \in T_x M$ if and only if

$$0 = \int_{S_x} \frac{D}{d\tau} \frac{d}{d\tau} \text{Ret}_x(\tau \bar{v}) \Big|_{\tau=0} d\bar{v}.$$

Applying Theorem 2.2 proves the first claim. If Ret_x is a second-order retraction, the second claim immediately follows since $\frac{D}{d\tau} \frac{d}{d\tau} \text{Ret}_x(\tau v) \Big|_{\tau=0} = 0$ for all $v \in T_x M$. □

Thus, random walks generated by second-order retractions converge to the Brownian motion on (M, g) .

Remark 2.4 Drift can seamlessly be incorporated into our approach. Indeed, when modeling drift by a vector field X , then Step 4 of Algorithm 1 needs to be altered according to $v := \sqrt{m}\bar{v} + \varepsilon X|_{x_{t-1}^\varepsilon}$. The stochastic process resulting from this altered version of Algorithm 1 converges to the process with generator $Lf + df(X)$, where Lf is the generator given by (2). The proof of this fact follows the proof of Theorem 2.2 with obvious modifications.

In the next two subsections, we consider two concrete examples of second-order retractions that are *computationally efficient*.

2.1 Retractions Based on Local Parameterizations

Our first example of computationally efficient retractions is based on local parameterizations of compact m -dimensional Riemannian manifolds (M, g) . Consider an atlas consisting of finitely many charts $\{(U_i, \phi_i)\}_{i \in I}$, where every $U_i \subset \mathbb{R}^m$ is open, $\phi_i : U_i \rightarrow M$ is a diffeomorphism onto its image in M , and where the relatively open sets $\{\phi_i(U_i)\}$ cover M . In concrete applications, such parameterizations often arise naturally, e.g., for isometrically embedded manifolds $M \subset \mathbb{R}^n$.

For a given $x \in M$, $v \in T_x M$, and (U_i, ϕ_i) with $x \in \phi_i(U_i)$, let $\tilde{x} = \phi_i^{-1}(x)$ and $\tilde{v} = (d\phi_i^{-1})|_x v$. We assume that for some fixed $\varepsilon > 0$ and any $x \in M$, there exists a chart (U_k, ϕ_k) with $\tilde{x} \in U_k$ such that

$$\tilde{x} + \tilde{v} - \frac{1}{2} \Gamma_{ij} \tilde{v}^i \tilde{v}^j \in U_k$$

for all $\tilde{v} = (d\phi_k^{-1})|_x(\varepsilon\sqrt{m}\bar{v})$ with $\bar{v} \in S_x$, where (Γ_{ij}) denotes the Christoffel symbols of the respective chart. This assumption is readily satisfied for any compact M , provided that ε is chosen small enough. In order to find such a sufficiently small ε in Algorithm 1, we propose to restart the algorithm with stepsize $\frac{\varepsilon}{2}$ if the condition is violated. Then, we define the retraction

$$\text{p-Ret}_x(v) := \phi_k \left(\tilde{x} + \tilde{v} - \frac{1}{2} \Gamma_{ij} \tilde{v}^i \tilde{v}^j \right). \tag{3}$$

Notice that Christoffel symbols can be computed numerically. For low-dimensional manifolds, symbolic differentiation accelerates their computation.

Lemma 2.5 *The retraction defined by Eq. (3) is of second order.*

Proof Let $(\gamma(\tau))_{\tau \geq 0}$ denote the geodesic satisfying $\gamma(0) = x$ and $\dot{\gamma}(0) = v$. In the parameter domain, γ satisfies the ordinary differential equation

$$\ddot{\gamma}^k + \Gamma_{ij}^k \dot{\gamma}^i \dot{\gamma}^j = 0. \tag{4}$$

The Taylor expansion of γ reads

$$\begin{aligned} \gamma^k(\tau) &= \gamma^k(0) + \tau \dot{\gamma}^k(0) + \frac{\tau^2}{2} \ddot{\gamma}^k(0) + \mathcal{O}(\tau^3) \\ &= \gamma^k(0) + \tau \dot{\gamma}^k(0) - \frac{\tau^2}{2} \Gamma_{ij}^k \dot{\gamma}^i(0) \dot{\gamma}^j(0) + \mathcal{O}(\tau^3) \\ &= \text{p-Ret}_x^k(v) + \mathcal{O}(\tau^3). \end{aligned}$$

Since $\text{Exp}_x(\tau v) = \gamma(\tau)$, this shows that

$$\text{p-Ret}_x(\tau v) = \text{Exp}_x(\tau v) + \mathcal{O}(\tau^3)$$

as $\tau \rightarrow 0$, which proves that p-Ret is indeed a second-order retraction. □

Table 1 Comparison between geodesic and retraction-based random walks

Dimension	Sampling (SVD)	Computing geodesics	Computing retractions
$m = 2$	0.107s	6.112s	0.083s (74×)
$m = 5$	0.193s	13.151s	0.190s (69×)
$m = 10$	0.322s	43.581s	0.675s (65×)
$m = 15$	0.516s	125.081s	1.994s (63×)
$m = 20$	0.873s	319.442s	4.987s (64×)

Computation times for geodesic and retraction-based random walks on m -dimensional tori with a non-diagonal metric $(g_{ij}) = B^T(\tilde{g}_{ij})B$, where $\tilde{g}_{ii} = 1.5 + \cos(x_{m-(i-1)})$ for $i = 1, \dots, m$, $\tilde{g}_{ij} = 0$ for $i \neq j$, and B is a random orthogonal matrix

In each case, 10,000 steps were calculated in varying dimensions with stepsize $\varepsilon = 0.1$. All computations were performed on a standard laptop in MATLAB, using MATLAB’s method ode45 for solving the geodesic equation. The Christoffel symbols were computed numerically

Another aspect of parameterization-based approaches concerns the computation of tangent vectors \tilde{v} in the parameter domain that are *uniformly distributed with respect to the Riemannian metric* g . An efficient procedure for this task is presented in Algorithm 2.

Algorithm 2 Uniform sampling of \tilde{v} w.r.t. the Riemannian metric

- 1: **Input:** Chart (U, ϕ) ; current position $\tilde{x} \in U \subset \mathbb{R}^m$
- 2: Sample m standard normal variates w_1, \dots, w_m .
- 3: Compute

$$z := \left(\sum_{i=1}^m w_i^2 \right)^{-\frac{1}{2}} (w_1, \dots, w_m)^T. \tag{5}$$

- 4: Compute the singular value decomposition (SVD) of the pullback metric $(g_{ij}) = \phi^*g$ at point \tilde{x} :

$$(g_{ij})|_{\tilde{x}} = \mathbf{V}\Sigma\mathbf{V}^T$$

with diagonal matrix Σ and orthogonal matrix \mathbf{V} .

- 5: Compute

$$\tilde{v} := \sum_{i=1}^m \frac{z_i}{\sqrt{\Sigma_{ii}}} \mathbf{V}_i,$$

where \mathbf{V}_i the i th column vector of \mathbf{V} .

- 6: **Return:** \tilde{v}

Notice that Eq. (5) is well defined since $w_k \neq 0$ a.s. for all $1 \leq k \leq m$ and that z is a sample of a uniformly distributed random variable on \mathbb{S}^{m-1} . It is straightforward to verify the following claim:

Lemma 2.6 *The vector \tilde{v} constructed in Algorithm 2 is a unit tangent vector that is uniformly sampled with respect to the Riemannian metric (g_{ij}) .*

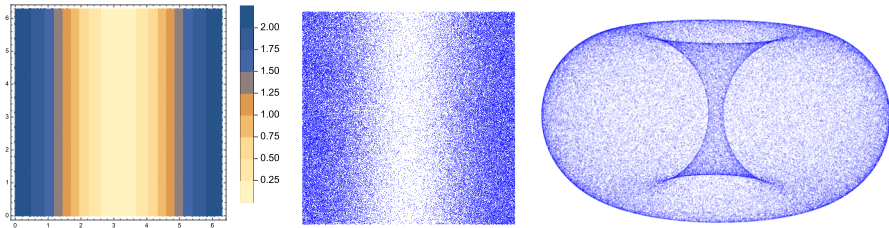


Fig. 1 Results for a parametric surface defined as $\phi(s, t) = ((1.1 + \cos(s)) \cos(t), (1.1 + \cos(s)) \sin(t), \sin(s))$. Left: Contour plot for the density of the stationary measure (i.e., the volume form of the Riemannian metric) in the parameter domain. Middle and right: Result of an approximate geodesic random walk based on p-Ret (using a single chart with periodic boundary conditions) with 100.000 steps and stepsize $\varepsilon = 0.5$, depicted in the parameter domain and in \mathbb{R}^3 , respectively. Using an implementation in MATHEMATICA on a standard laptop, the computation takes a few seconds. For a discussion of the relationship between the stationary measure and the empirical measure resulting from the random walk, see Sect. 2.3

In summary, using \tilde{v} as computed in Algorithm 2, setting $\bar{v} = \tilde{v}$ in Step 3 of Algorithm 1, and using the retraction p-Ret from (3), provides an efficient method for computing random walks for parameterized manifolds. See Table 1 for a comparison between geodesic random walks and retraction-based random walks, showing that the latter are significantly faster than the former. Notice that the attendant algorithm might require to update the respective local chart during the simulated random walk. Results of a simulated random walk on a 2-torus (using a single chart with periodic boundary conditions) are presented in Fig. 1.

2.2 Projection-Based Retractions

Our second example of computationally efficient retractions is based on a step-and-project method for compact manifolds that are given as zero-level sets of smooth functions. More specifically, we consider the setting where

$$f : \mathbb{R}^n \rightarrow \mathbb{R}^k$$

is a smooth function, $k < n$, and the m -dimensional manifold in question, with $m = (n - k)$, is given by

$$M = \{x \in \mathbb{R}^n : f(x) = 0\}. \tag{6}$$

We assume that 0 is a regular value of f , i.e., that the differential df has full rank along M . In this case, a retraction can be defined via

$$\pi\text{-Ret}_x(v) = \pi_M(x + v),$$

where $x \in M$, $v \in T_x M$, and $\pi_M(y) = \arg \min_{p \in M} \|p - y\|$ denotes the (closest point) projection to M . Notice that $\pi_M(y)$ is well defined as long as the distance $d(y, M)$ is less than the so-called *reach* of M . The reach of M is the distance of M

to its medial axis, which in turn is defined as the set of those points in \mathbb{R}^n that do not have a unique closest point on M . Smoothly embedded compact manifolds always have positive reach.

The following result is important in our setup. For a proof, see, e.g., Theorem 4.10 in [3].

Lemma 2.7 *π -Ret is a second-order retraction.*

It remains to specify the computation of a random unit direction $v \in T_x M$ as well as the implementation of the projection step. As for the former, a uniformly sampled unit tangent direction at some $x \in M$ can be computed by first sampling a randomly distributed unit vector $u \in \mathbb{S}^{n-1}$ in ambient space, and denoting its projection to $\text{im}(df|_x)$ by \tilde{u} . Then,

$$\bar{v} = \frac{u - \tilde{u}}{\|u - \tilde{u}\|} \tag{7}$$

yields the requisite uniformly distributed unit tangent vector. Finally, for computing the closest point projection of some given point $y \in \mathbb{R}^n$ to M , consider the Lagrangian

$$\mathcal{L}_y(z, \lambda) = \frac{1}{2} \|z - y\|^2 - \lambda^T f(z),$$

where $\lambda \in \mathbb{R}^k$ denotes the (vector-valued) Lagrange multiplier. Then, Newton’s method offers an efficient implementation for solving the Euler–Lagrange equations resulting from \mathcal{L}_y . In our implementation, we let Newton’s method run until the threshold $|f(z)|/\|df|_z\| < \varepsilon^3$ is satisfied. Notice that efficiency results from the fact that the point $y = x + \tau v$ is close to M for $x \in M$ and τ small enough.

In summary, using \bar{v} as computed in (7) for Step 3 of Algorithm 1, together with the retraction π -Ret from (6), provides an efficient method for computing random walks for implicitly given submanifolds. Results of a simulated random walk on the 2-torus are presented in Fig. 2.

Remark 2.8 Our projection-based retraction is an instance of retractions induced by so-called *normal foliations*. The latter are defined as foliations in a local neighborhood of $M = f^{-1}(0)$ whose leaves have dimension k and intersect M orthogonally. As shown in [32], retractions induced by normal foliations are always of second order and therefore provide alternative approaches for sampling Brownian motion on implicit manifolds. Examples of retractions induced by normal foliations include the gradient flow along $-\nabla f$ and Newton retractions based on the update step $\delta z = -(df)^\dagger(z) f(z)$, where $(df)^\dagger$ is the Moore–Penrose pseudo-inverse of the Jacobian df .

2.3 Stationary Measure

For second-order retractions, we show below that the stationary measure of the process in Algorithm 1 converges to the stationary measure of the Brownian motion, which is

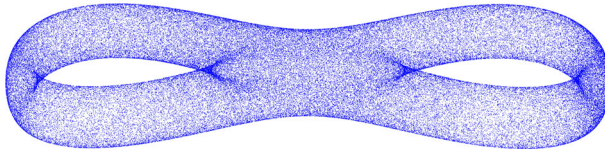


Fig. 2 Result of a random walk using π -Ret for a genus two surface given as the zero-level set of $f : \mathbb{R}^3 \rightarrow \mathbb{R}$, $f(x, y, z) = (x^2(1 - x^2) - y^2)^2 + z^2 - 0.01$ with 100.000 steps and stepsize $\varepsilon = 0.1$. Using an implementation in MATHEMATICA on a standard laptop, the computation takes less than a minute

the normalized Riemannian volume form, see Theorem 3.2. Nonetheless, the *empirical* approximation $\frac{1}{N} \sum_{i=1}^N \delta_{x_i^\varepsilon}$ of the Riemannian volume form using random walks—e.g., in order to obtain results as in Figs. 1 and 2—requires the stepsize ε and the number of steps N in Algorithm 1 to be chosen appropriately. An indicator at which time the empirical approximation is close to the stationary measure is the so-called δ -cover time, i.e., the time needed to hit any δ -ball. Hence, as a rough heuristic choice of ε and N , one may utilize [13, 14]. For a Brownian motion on a 2-dimensional manifold M , the δ -cover time C_δ of M satisfies [14, Theorem 1.3]

$$\lim_{\delta \rightarrow 0} \frac{C_\delta}{(\log \delta)^2} = \frac{2A}{\pi} \quad \text{a.s.,}$$

where A denotes the Riemannian area of M . Based on the observation that a second-order retraction converges to Brownian motion in the limit of vanishing stepsize ε , and using that the physical time of the walk resulting from Algorithm 1 is given by $\varepsilon^2 N$, one may set $\varepsilon^2 N = C_\delta$. If δ is small enough, one may then choose N according to

$$N \geq \frac{2A}{\pi} \left(\frac{\log \delta}{\varepsilon} \right)^2.$$

For dimension $m > 2$, the corresponding result in [13], Theorem 1.1 yields the heuristic choice

$$N \geq m\kappa_M \frac{-\delta^{2-m} \log \delta}{\varepsilon^2},$$

where

$$\kappa_M = \frac{2}{(m - 2)\omega_m} V(M),$$

$V(M)$ denotes the volume of M , and ω_m the volume of the unit sphere in \mathbb{R}^m .

Such heuristics, however, do not quantify the error between the *empirical measure* obtained by simulating a random walk and the Riemannian volume measure. In the setting of Euclidean random walks $(Y_n)_{n \in \mathbb{N}}$ with invariant distribution ν , it has been shown in [23] that under some conditions on Y the expectation of the 1-Wasserstein distance between the empirical and the invariant measure asymptotically decreases according to the law

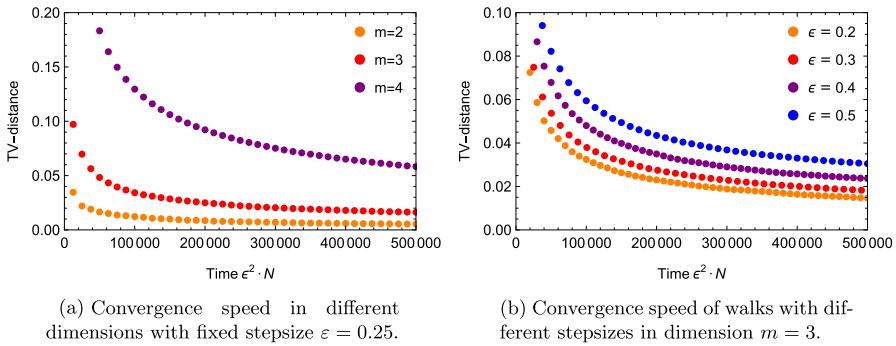


Fig. 3 Convergence speed of retraction-based random walks to the invariant measures on m -dimensional tori parametrized by $[0, 2\pi]^m$ with periodic boundary conditions and equipped with the diagonal metric $g_{ii} = 1.5 + \cos(x_{m-(i-1)})$ for $i = 1, \dots, m$ and $g_{ij} = 0$ for $i \neq j$. Results are shown for varying dimension m and fixed stepsize ϵ (left) as well as fixed dimension m and varying stepsize ϵ (right). Simulations were performed 10 times for each combination of dimension and stepsize. The points display the arithmetic mean of the total variation distances between the invariant measure and the empirical measures resulting from a box count using 20^m equally sized cubes

$$\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \delta_{Y_i} - \nu \right\|_{W_1} \lesssim \alpha \frac{(\log(\vartheta n))^\beta}{(\vartheta n)^\gamma} \tag{8}$$

as $N \rightarrow \infty$ with constants $\alpha, \beta, \gamma > 0$ depending on the dimension m and $\vartheta \in (0, 1)$ depending on the process Y . Our simulations, see Fig. 3, followed by a simple fit to the model given by the right-hand side of Eq. (8) suggest that a similar qualitative behavior might be present for retraction-based random walks on Riemannian manifolds. We leave a deeper analysis for future research.

3 Convergence

This section is devoted to the proof of Theorem 2.2, which yields convergence in the Skorokhod topology of the process constructed in Algorithm 1 as the stepsize $\epsilon \rightarrow 0$. The Skorokhod space $\mathcal{D}_M[0, \infty)$ is the space of all càdlàg, i.e., right-continuous functions with left-hand limits from $[0, \infty)$ to M , equipped with a metric turning $\mathcal{D}_M[0, \infty)$ into a complete and separable metric space, and hence a Polish space (see, e.g., [17], Chapter 3.5). The topology induced by this metric is the Skorokhod topology, which generalizes the topology of local uniform convergence for continuous functions $[0, \infty) \rightarrow M$ in a natural way. For local uniform convergence, we have uniform convergence on any compact time interval $K \subset [0, \infty)$. In the Skorokhod topology, functions with discontinuities converge if the times and magnitudes of the discontinuities converge (in addition to local uniform convergence of the continuous parts of the functions). Thus, $\mathcal{D}_M[0, \infty)$ is a natural and—due to its properties—also a convenient space for convergence of discontinuous stochastic processes.

With reference to the notation used in Theorem 2.2, notice that the process $X^\epsilon = (X_{\lfloor \epsilon^{-2}t \rfloor}^\epsilon)_{t \geq 0}$ has deterministic jump times and is therefore not Markov as a process

in continuous time. In order to apply the theory of Markov processes, let $(\eta_t)_{t \geq 0}$ be a Poisson process with rate ε^{-2} and define the pseudo-continuous process $Z^\varepsilon = (Z_t^\varepsilon)_{t \geq 0}$ by $Z_t^\varepsilon := X_{\eta_t}^\varepsilon$ for any $t \geq 0$, which is Markov since the Poisson process is Markov. Then, the convergence of Z^ε implies convergence of X^ε to the same limit by the law of large numbers (see [22], Theorem 19.28). First we restate Theorem 2.2, using an equivalent formulation based on the process Z^ε :

Theorem 3.1 *Consider the sequence of random variables $(X_i^\varepsilon)_{i \in \mathbb{N}}$ obtained by the construction from Algorithm 1. The process $Z^\varepsilon = (X_{\eta_t}^\varepsilon)_{t \geq 0}$ converges in distribution in the Skorokhod topology to the L-diffusion Z, i.e., the stochastic process with generator L defined in (2). This means that for any continuous and bounded function $f : \mathcal{D}_M[0, \infty) \rightarrow \mathbb{R}$ we have*

$$\mathbb{E}(f((Z_t^\varepsilon)_{t \geq 0})) \rightarrow \mathbb{E}(f((Z_t)_{t \geq 0})) \quad \varepsilon \rightarrow 0.$$

Certainly one cannot, in general, deduce convergence of stationary measures from time-local convergence statements about paths. However, standard arguments from the theory of Feller processes indeed allow to do so. This yields the following statement regarding stationary measures.

Theorem 3.2 *Let μ_ε be the stationary measure of the retraction-based random walk with stepsize ε . Then, the weak limit of $(\mu_\varepsilon)_{\varepsilon > 0}$ as $\varepsilon \rightarrow 0$ is the stationary measure of the L-diffusion, with generator L defined in (2). In the case of second-order retractions, the limit is the Riemannian volume measure.*

Remark 3.3 Notice that due to the approximation of the exponential map, the process Z^ε is not in general reversible and the generator of Z^ε is not in general self-adjoint. This is relevant since the stationary measure corresponds to the kernel of the adjoint of the generator.

Our proof of Theorems 3.1 and 3.2 (presented below) hinges on convergence of generators that describe the infinitesimal evolution of Markov processes, see, e.g., [17, Chapter 4]. We show that the generator of Z^ε converges to the generator L defined in (2). The generator of the process Z^ε is spelled out in Lemma 3.4, and the convergence of this generator to L is treated in Lemma 3.5.

The following result is standard for transition kernels U^ε on compact state spaces, see, e.g., [17, Chapter 8.3]:

Lemma 3.4 *The process $(Z_t^\varepsilon)_{t \geq 0}$ is Feller, and its generator is given by*

$$L^\varepsilon f = \frac{1}{\varepsilon^2}(U^\varepsilon f - f) \tag{9}$$

for all $f \in C(M)$, the continuous real-valued functions on M. Here, $(U^\varepsilon f)(x) = \frac{1}{\omega_m} \int_{S_x} f(\text{Ret}_x(\varepsilon\sqrt{m}\bar{v}))d\bar{v}$ and ω_m is the volume of the unit $(m - 1)$ -sphere.

Notice that Lemma 3.4 additionally states that the process Z^ε satisfies the Feller property, i.e., that its semigroup is a contraction semigroup that also fulfills a right-continuity property. Therefore, in particular, the process Z^ε is Markov.

In the following, $C^{(2,\alpha)}(M)$, $0 < \alpha \leq 1$ denotes the space of two times differentiable functions on M with α -Hölder continuous second derivative.

Lemma 3.5 For any $f \in C^{(2,\alpha)}(M)$ with $0 < \alpha \leq 1$,

$$\|L^\varepsilon f - Lf\|_\infty \rightarrow 0$$

as $\varepsilon \rightarrow 0$, where L^ε is the operator from Lemma 3.4 and L as defined in (2).

Proof Let P^ε be the transition kernel of the geodesic random walk, i.e.,

$$(P^\varepsilon f)(x) = \frac{1}{\omega_m} \int_{S_x} f(\text{Exp}_x(\varepsilon\sqrt{m}\bar{v}))d\bar{v}.$$

For $x \in M$ and $v \in T_x M$, define

$$(Gf)(x, v) := \frac{1}{2}df|_x \left(\frac{D}{d\tau} \frac{d}{d\tau} \text{Ret}_x(\tau v) \right) \Big|_{\tau=0}$$

and

$$(\tilde{G}f)(x) := \frac{m}{\omega_m} \int_{S_x} (Gf)(x, \bar{v})d\bar{v}.$$

Then, since $Lf = \frac{1}{2}\Delta_g f + \tilde{G}f$,

$$\|L^\varepsilon f - Lf\|_\infty \leq \left\| \frac{1}{\varepsilon^2}(U^\varepsilon f - P^\varepsilon f) - \tilde{G}f \right\|_\infty + \left\| \frac{1}{\varepsilon^2}(P^\varepsilon f - f) - \frac{1}{2}\Delta_g f \right\|_\infty.$$

Proposition 2.3 in [21] proves that the second summand on the right-hand side converges to 0 as $\varepsilon \rightarrow 0$. Let $\tilde{\varepsilon} := \sqrt{m}\varepsilon$. Then,

$$\begin{aligned} & \left| \frac{1}{\varepsilon^2}(U^\varepsilon f(x) - P^\varepsilon f(x)) - (\tilde{G}f)(x) \right| \\ & \leq \frac{m}{\omega_m} \int_{S_x} \left| \frac{1}{\tilde{\varepsilon}^2}[f(\text{Ret}_x(\tilde{\varepsilon}\bar{v})) - f(\text{Exp}_x(\tilde{\varepsilon}\bar{v}))] - (Gf)(x, \bar{v}) \right| d\bar{v}. \end{aligned}$$

Now, consider the Taylor expansions of the functions $f \circ \text{Ret}_x$ and $f \circ \text{Exp}_x$ in $\tilde{\varepsilon}$ at 0. To this end, notice that for any $x \in M$ and $v \in T_x M$, one has

$$\frac{d}{d\tau} f(\text{Ret}_x(\tau v)) = df|_{\text{Ret}_x(\tau v)}(v_\tau),$$

where $v_\tau = \frac{d}{d\tau} \text{Ret}_x(\tau v)$. Hence, we obtain that

$$R_{\text{Ret}}^{x,v}(\tau) := \frac{d^2}{d\tau^2} f(\text{Ret}_x(\tau v)) = \text{Hess } f(v_\tau, v_\tau) + df|_{\text{Ret}_x(\tau v)}\left(\frac{D}{d\tau} v_\tau\right),$$

where $\text{Hess } f = \nabla df$ is the Riemannian Hessian of f . Using that $v_0 = v$, the desired Taylor expansion takes the form

$$f(\text{Ret}_x(\tilde{\varepsilon}v)) = f(x) + \tilde{\varepsilon}df|_x(v) + \frac{1}{2}\tilde{\varepsilon}^2 R_{\text{Ret}}^{x,v}(\xi_1),$$

where $\xi_1 \in [0, \tilde{\varepsilon}]$. Defining $u_\tau = \frac{d}{d\tau} \text{Exp}_x(\tau v)$, we similarly obtain that

$$f(\text{Exp}_x(\tilde{\varepsilon}v)) = f(x) + \tilde{\varepsilon}df|_x(v) + \frac{1}{2}\tilde{\varepsilon}^2 R_{\text{Exp}}^{x,v}(\xi_2)$$

with $\xi_2 \in [0, \tilde{\varepsilon}]$ and

$$R_{\text{Exp}}^{x,v}(\tau) = \text{Hess } f(u_\tau, u_\tau),$$

where for the last equality we used the fact that $\frac{d}{d\tau} u_\tau = 0$. Since $R_{\text{Ret}}^{x,v}(0) = R_{\text{Exp}}^{x,v}(0) + 2(Gf)(x, v)$, we obtain that

$$\begin{aligned} &|f(\text{Ret}_x(\tilde{\varepsilon}v)) - f(\text{Exp}_x(\tilde{\varepsilon}v)) - \tilde{\varepsilon}^2(Gf)(x, v)| \\ &= \frac{\tilde{\varepsilon}^2}{2} |R_{\text{Ret}}^{x,v}(\xi_1) - R_{\text{Ret}}^{x,v}(0) - R_{\text{Exp}}^{x,v}(\xi_2) + R_{\text{Exp}}^{x,v}(0)| \\ &\leq \frac{\tilde{\varepsilon}^2}{2} \left(|R_{\text{Ret}}^{x,v}(\xi_1) - R_{\text{Ret}}^{x,v}(0)| + |R_{\text{Exp}}^{x,v}(\xi_2) - R_{\text{Exp}}^{x,v}(0)| \right). \end{aligned}$$

Since $\text{Hess } f$ is Hölder-continuous by assumption and Exp and Ret are smooth, both R_{Exp} and R_{Ret} are Hölder-continuous as products and sums of Hölder-continuous functions. Therefore, there exists a constant $C > 0$ such that one has the uniform bound

$$|f(\text{Ret}_x(\tilde{\varepsilon}v)) - f(\text{Exp}_x(\tilde{\varepsilon}v)) - \tilde{\varepsilon}^2(Gf)(x, v)| \leq \frac{C}{2} (\xi_1^\alpha + \xi_2^\alpha) \tilde{\varepsilon}^2 \leq C \tilde{\varepsilon}^{2+\alpha}.$$

Hence,

$$\sup_{x \in M} \left| \frac{1}{\varepsilon^2} (U^\varepsilon f(x) - P^\varepsilon f(x)) - (\tilde{G}f)(x) \right| \leq Cm^{1+\frac{\alpha}{2}} \varepsilon^\alpha \xrightarrow{\varepsilon \rightarrow 0} 0,$$

which proves the lemma. □

Using the above results, we can finish the proof of Theorem 3.1.

Proof of Theorem 3.1 By the proof of Theorem 2.1 and the remark on page 38 in [21], the space $C^{(2,\alpha)}(M)$ is a core for the differential operator L . For a sequence of Feller processes $(Z^\varepsilon)_{\varepsilon>0}$ with generators $L^\varepsilon, \varepsilon > 0$ and another Feller process Z with generator L and core D , the following are equivalent [22, Theorem 19.25]:

- (i) If $f \in D$, there exists a sequence of $f_\varepsilon \in \text{Dom}(L^\varepsilon)$ such that $\|f_\varepsilon - f\|_\infty \rightarrow 0$ and $\|L^\varepsilon f_\varepsilon - Lf\|_\infty \rightarrow 0$ as $\varepsilon \rightarrow 0$.

- (ii) If the initial conditions satisfy $Z_0^\varepsilon \rightarrow Z_0$ in distribution as $\varepsilon \rightarrow 0$ in M , then $Z^\varepsilon \rightarrow Z$ as $\varepsilon \rightarrow 0$ in distribution in the Skorokhod space $\mathcal{D}_M([0, \infty))$.

In our case, $D = C^{(2,\alpha)}(M)$ and $f_\varepsilon = f \in D$ satisfy (i) due to Lemma 3.5. This concludes the proof of Theorem 3.1 since the L -diffusion is also a Feller process on a compact manifold. \square

From an analytical point of view, it is perhaps not surprising that convergence of generators, see Lemma 3.5, implies convergence of stationary measures as stated in Theorem 3.2. For the proof, we here follow standard arguments of the theory of Feller processes.

Proof of Theorem 3.2 The family of (unique) stationary measures $(\mu_\varepsilon)_{\varepsilon>0}$ is a tight family of measures since all measures are supported on the same compact manifold M . By Prokhorov's theorem, which provides equivalence of tightness and relative compactness, any subsequence of the family $(\mu_{\varepsilon_n})_{n \in \mathbb{N}}$ with $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$ has a convergent subsubsequence. But the uniform convergence of the generators that we showed in Lemma 3.5, see also (i) in the proof of Theorem 3.1, is also equivalent to the convergence of the associated semigroups with respect to the supremum norm in space, see again [22, Theorem 19.25]. This implies [17, Chapter 4, Theorem 9.10] that all subsequential limits must be the unique stationary measure of the L -diffusion, and therefore, all subsubsequences converge to the same measure. A standard subsubsequence argument then proves the theorem. \square

Retractions on Sub-Riemannian Manifolds

A possible extension of our work concerns random walks in *sub-Riemannian* geometries. Indeed, sub-Riemannian structures can be used to model low-dimensional noise that lives in a higher dimensional state space [19]. A sub-Riemannian structure on a smooth manifold M consists of a smoothly varying positive definite quadratic form on a sub-bundle E of the tangent bundle TM . Similar to the Riemannian setting, so-called *normal* sub-Riemannian geodesics arise from a first-order ODE on the cotangent bundle [28]. Such geodesics are uniquely determined by an initial position $x \in M$ and a 1-form $\alpha \in T_x^*M$ and can therefore be approximated by a second-order Taylor expansion of the respective ODE on the cotangent bundle. Such an approximation would algorithmically resemble our retraction-based approach, when attempting to efficiently simulate a sub-Riemannian geodesic random walk. However, in general, there exists no canonical choice for *sampling* a requisite 1-form α in every step in order to construct a sub-Riemannian geodesic random walk (or an approximation thereof). Indeed, one requires an additional choice of a measure for such sampling, and different choices may lead to different limiting processes, see [4, 9]. Notice that once such a choice has been made, one can seamlessly adapt our algorithm to the sub-Riemannian setting. We leave for future work a respective convergence analysis. Moreover, there might exist canonical choices of such measures for special sub-Riemannian structures—e.g., when considering the frame bundle of a smooth manifold M in order to model anisotropic Brownian motion and general diffusion processes [18, 29].

Acknowledgements We acknowledge support by the Collaborative Research Center 1456 funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation). We thank Pit Neumann for pointing out an error in a previous version of the manuscript and the anonymous referees for improving the quality of our manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Absil, P.-A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton, New Jersey (2008)
2. Absil, P.-A., Malick, J.: Projection-like Retractions on Matrix Manifolds. *SIAM J. Optim.* **22**(1), 135–158 (2012). <https://doi.org/10.1137/100802529>
3. Adler, R.L., Dedieu, J.-P., Margulies, J.Y., Martens, M., Shub, M.: Newton's method on Riemannian manifolds and a geometric model for the human spine. *IMA J. Numer. Anal.* **22**, 359–390 (2002). <https://doi.org/10.1093/imanum/22.3.359>
4. Agrachev, A., Boscain, U., Neel, R., Rizzi, L.: Intrinsic random walks in Riemannian and sub-Riemannian geometry via volume sampling. *ESAIM Control Optim. Calc. Var.* **24**(3), 1075–1105 (2018). <https://doi.org/10.1051/cocv/2017037>
5. Armstrong, J., Brigo, D.: Intrinsic stochastic differential equations as jets. *Proc. R. Soc. A* **474**, 20170559 (2018). <https://doi.org/10.1098/rspa.2017.0559>
6. Armstrong, J., King, T.: Curved schemes for stochastic differential equations on, or near, manifolds. *Proc. R. Soc. A* **478**, 20210785 (2022). <https://doi.org/10.1098/rspa.2021.0785>
7. Averina, T.A., Rybakov, K.A.: A Modification of Numerical Methods for Stochastic Differential Equations with First Integrals. *Numer. Anal. Appl.* **12**, 203–218 (2019). <https://doi.org/10.1134/S1995423919030017>
8. Bonnabel, S.: Stochastic gradient descent on Riemannian manifolds. *IEEE Trans. Automat. Contr.* **58**(9), 2217–2229 (2013). <https://doi.org/10.1109/TAC.2013.2254619>
9. Boscain, U., Neel, R., Rizzi, L.: Intrinsic random walks and sub-Laplacians in sub-Riemannian geometry. *Adv. Math.* **314**, 124–184 (2017). <https://doi.org/10.1016/j.aim.2017.04.024>
10. Brubaker, M., Salzmann, M., Urtasun, R.: A Family of MCMC Methods on Implicitly Defined Manifolds. In: Lawrence, N.D., Girolami, M. (eds.) Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, vol. 22, pp. 161–172 (2012)
11. Byrne, S., Girolami, M.: Geodesic Monte Carlo on Embedded Manifolds. *Scand. J. Stat.* **40**(4), 825–845 (2013). <https://doi.org/10.1111/sjso.12036>
12. Cantarella, J., Shonkwiler, C.: The Symplectic Geometry of Closed Equilateral Random Walks in 3-space. *Ann. Appl. Probab.* **26**(1), 549–596 (2016). <https://doi.org/10.1214/15-aap1100>
13. Dembo, A., Peres, Y., Rosen, J.: Brownian Motion on Compact Manifolds: Cover Time and Late Points. *Electron. J. Probab.* **8**, 1–14 (2003). <https://doi.org/10.1214/EJP.v8-139>
14. Dembo, A., Peres, Y., Rosen, J., Zeitouni, O.: Cover times for Brownian motion and random walks in two dimensions. *Ann. Math.* **160**, 433–464 (2004). <https://doi.org/10.4007/annals.2004.160.433>
15. do Carmo, M.P.: Riemannian Geometry. Birkhäuser, Boston (1992)
16. Emery, M.: Stochastic Calculus in Manifolds. Springer, Berlin Heidelberg (1989)
17. Ethier, S.N., Kurtz, T.G.: Markov Processes. John Wiley & Sons, New York (1986)

18. Grong, E., Sommer, S.: Most Probable Paths for Anisotropic Brownian Motions on Manifolds. *Found. Comput. Math.* (2022). <https://doi.org/10.1007/s10208-022-09594-4>
19. Habermann, K.: Geometry of sub-Riemannian diffusion processes. PhD thesis, University of Cambridge (2017)
20. Hsu, E.P.: *Stochastic Analysis on Manifolds*. Graduate Studies in Mathematics, vol. 38. American Mathematical Society, Providence, RI (2002)
21. Jørgensen, E.: The Central Limit Problem for Geodesic Random Walks. *Z. Wahrscheinlichkeitstheorie verw. Gebiete* **32**, 1–64 (1975)
22. Kallenberg, O.: *Foundations of Modern Probability*. Springer, New York (2002)
23. Kloeckner, B.R.: Empirical measures: regularity is a counter-course to dimensionality. *ESAIM: Prob. Stat.* **24**, 408–434 (2020). <https://doi.org/10.1051/ps/2019025>
24. Lee, Y.T., Vempala, S.: Geodesic Walks in Polytopes. *SIAM J. Comput.* **51**(2), 400–488 (2022). <https://doi.org/10.1137/17M1145999>
25. Leimkuhler, B., Matthews, C.: *Molecular Dynamics*. Springer, Cham (2015)
26. Ma, T., Matveev, V.S., Pavlyukevich, I.: Geodesic Random Walks, Diffusion Processes and Brownian Motion on Finsler Manifolds. *J. Geom. Anal.* **31**, 12446–12484 (2021). <https://doi.org/10.1007/s12220-021-00723-z>
27. Mangoubi, O., Smith, A.: Rapid mixing of geodesic walks on manifolds with positive curvature. *Ann. Appl. Probab.* **28**(4), 2501–2543 (2018). <https://doi.org/10.1214/17-AAP1365>
28. Montgomery, R.: *A Tour of Subriemannian Geometries, Their Geodesics and Applications*. Mathematical Surveys and Monographs, vol. 91. American Mathematical Society, Providence, RI (2002)
29. Sommer, S., Svane, A.M.: Modelling anisotropic covariance using stochastic development and sub-Riemannian frame bundle geometry. *J. Geom. Mech.* **9**(3), 391–410 (2017). <https://doi.org/10.3934/jgm.2017015>
30. Vempala, S.: Geometric random walks: a survey. In: Goodman, J.E., Pach, J., Welzl, E. (eds.) *Combinatorial and Computational Geometry*. Mathematical Sciences Research Institute Publications, vol. 52, pp. 573–612 (2005)
31. Yanush, V., Kropotov, D.: Hamiltonian Monte-Carlo for Orthogonal Matrices (2019) [arXiv:1901.08045](https://arxiv.org/abs/1901.08045) [stat.ML]
32. Zhang, R.: Newton retraction as approximate geodesics on submanifolds (2020) [arXiv:2006.14751](https://arxiv.org/abs/2006.14751) [math.NA]

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.