



Randomized time Riemannian Manifold Hamiltonian Monte Carlo

Peter A. Whalley¹ · Daniel Paulin¹ · Benedict Leimkuhler¹

Received: 14 August 2022 / Accepted: 26 September 2023 / Published online: 24 November 2023
© The Author(s) 2023

Abstract

Hamiltonian Monte Carlo (HMC) algorithms, which combine numerical approximation of Hamiltonian dynamics on finite intervals with stochastic refreshment and Metropolis correction, are popular sampling schemes, but it is known that they may suffer from slow convergence in the continuous time limit. A recent paper of Bou-Rabee and Sanz-Serna (*Ann Appl Prob*, 27:2159–2194, 2017) demonstrated that this issue can be addressed by simply randomizing the duration parameter of the Hamiltonian paths. In this article, we use the same idea to enhance the sampling efficiency of a constrained version of HMC, with potential benefits in a variety of application settings. We demonstrate both the conservation of the stationary distribution and the ergodicity of the method. We also compare the performance of various schemes in numerical studies of model problems, including an application to high-dimensional covariance estimation.

Keywords Piecewise deterministic Markov process · Sampling · Riemannian manifold · High dimensional inference

1 Introduction and motivation

Efficient sampling of high dimensional probability distributions is required for Bayesian inference and is a challenge in many fields including biological modelling (Wilkinson 2007), economic modelling (Greenberg 2012), machine learning with large data sets (Pakman et al. 2017; Barber 2012) and molecular dynamics (Perez et al. 2015). A popular approach is Markov chain Monte Carlo, which defines a Markov chain $X_{i+1} \sim p(\cdot | X_i)$ with invariant measure μ and from which we may estimate expected values from the relation $\mathbb{E}_{X \sim \mu} f(X) \approx \frac{1}{N} \sum_{i=1}^N f(X_i)$; however convergence of such averages can be slow for high dimensional and multimodal distributions (see e.g. Quiroz et al. (2018)). Recent attempts to address this problem include the local bouncy particle sampler of Bouchard-Côté et al. (2018) and the Zig-Zag process of Bierkens et al. (2019). These methods can be viewed as piecewise deterministic Markov processes (PDMPs), see (Vanetti et al. 2018). The Randomized Hamil-

tonian Monte Carlo (RHMC), proposed in Bou-Rabee and Sanz-Serna (2017) and further studied in Deligiannidis et al. (2021), evolves a Hamiltonian flow for a duration drawn from an exponential distribution. In standard HMC the choice of integration time is a challenging task (see Hoffman and Gelman 2014) and mixing can be inefficient for some choices of integration time. By contrast, RHMC does not suffer from this problem as randomization of the duration prevents periodicities. This strategy has been studied from both analytic and numerical perspectives in Bou-Rabee and Sanz-Serna (2017). Other recent algorithms have been proposed which build on this idea (for example Riou-Durand and Vogrinc (2022) and Kleppe (2022)). We remark that RHMC is a special case of Andersen dynamics which is popular in the Molecular dynamics literature (see (Bou-Rabee and Eberle 2022)[Remark 2.2] and Andersen (1980)). Andersen dynamics has been studied in Bou-Rabee and Eberle (2022); Weinan and Li (2008) and Li (2007).

The algorithms discussed above are targeted to sampling from distributions in Euclidean space. The need to work with Riemannian manifolds is motivated by applications where constraints are imposed from modelling considerations or are introduced in order to restrict sampling to a relevant subdomain derived from statistical analysis (see (Brubaker et al. 2012)). Examples of manifolds include products of spheres or orthogonal matrices which arise in applications in protein configuration modelling with the Fisher-Bingham

✉ Peter A. Whalley
p.a.whalley@sms.ed.ac.uk

Daniel Paulin
dpaulin@ed.ac.uk

Benedict Leimkuhler
b.leimkuhler@ed.ac.uk

¹ Department of Mathematics, University of Edinburgh,
Edinburgh EH9 3FD, UK

distribution (Hamelryck et al. 2006), texture analysis using distributions over rotations (Kunze and Schaeben 2004) and fixed-rank matrix factorization for collaborative filtering (Meyer et al. 2011; Salakhutdinov and Mnih 2008). Methods that sample from probability distributions on manifolds have been considered and studied in Hartmann (2008); Brubaker et al. (2012); Byrne and Girolami (2013); Girolami and Calderhead (2011); Lelièvre et al. (2012); Zappa et al. (2018); Diaconis et al. (2013); Lee and Vempala (2018); Lelièvre et al. (2019) and Laurent and Vilmart (2021). In this article, we focus on manifolds defined by algebraic constraints. In order to maintain the constraints, in practice one needs to perform projections at each step of the algorithm, an additional overhead compared to Euclidean MCMC algorithms.

In this paper we propose the Randomized Time Riemannian Manifold Hamiltonian Monte Carlo (RT-RMHMC) method, an RHMC scheme for Riemannian manifolds. We establish invariance under a compactness assumption of the desired measure in the (small stepsize limit) continuous-time PDMP version of our method, where the algorithm is rejection free. Our approach to proving invariance is based on the PDMP framework established in Durmus et al. (2021), we construct an approximation of RT-RMHMC by truncating the velocity distribution consistently and prove all the conditions needed in Durmus et al. (2021) with the approximation; we have not found such a construction technique in the literature. Further, we demonstrate the invariance of the discretized method with Metropolis-Hastings adjustment and prove ergodicity of the discretized method with Metropolis-Hastings adjustment. We show in numerical experiments that this method has improved robustness, demonstrating for example that the convergence rate is relatively flat in the choice of mean time parameter; these results mirror those obtained for the Euclidean version of the method. Moreover, we compare RT-RMHMC to a constrained underdamped Langevin integrator (g-BAOAB) introduced in Leimkuhler and Matthews (2016).

To our knowledge, there is no theoretical or numerical treatment of RHMC in the manifold setting and there has been no theoretical treatment of Riemannian Hamiltonian Monte Carlo methods in the continuous time setting. We provide a first result to establish invariance of a continuous time Riemannian Hamiltonian Monte Carlo method in the compact setting. A biased RHMC method was recently introduced (see (Kleppe 2022)) which has event rates which depend on the position in the state space, these state dependent event rates can be incorporated into our RHMC Riemannian framework when the framework is unadjusted. We note that in the appendix of that article, a version of RHMC is introduced in the setting of adapting the metric for sampling on Euclidean space but not for working on a Riemannian manifold.

The remainder of this article is organized as follows. In the next section we describe the algorithm and provide invariance

in the continuous time setting under a compactness assumption. Section 3 considers the numerical implementation with and without Metropolis test. Section 4 provides conservation of the stationary distribution of the discretized algorithm and the ergodicity of the method with Metropolis-Hastings adjustment. Section 5 discusses numerical experiments and Section 6 gives some thoughts on future developments. We include several appendices addressing the generator, the invariance of the target measure and the irreducibility of the scheme, from which ergodicity necessarily follows.

2 Algorithm

Let (\mathcal{M}, g) be a d -dimensional Riemannian manifold and $T\mathcal{M}$ denote its tangent bundle. Let $G(x)$ denote the positive definite matrix associated to the metric g at $x \in \mathcal{M}$. Consider a target distribution on \mathcal{M} with density

$$\pi_{\mathcal{H}}(x) = \frac{1}{Z_{\mathcal{M}}} \exp(-U_{\mathcal{H}}(x)),$$

with respect to $\sigma_{\mathcal{M}}(dx)$, the surface measure (Hausdorff measure) of \mathcal{M} defined by $\sigma_{\mathcal{M}}(dx) = \sqrt{\det G(x)} dx$ and $Z_{\mathcal{M}} = \int_{\mathcal{M}} \exp(-U_{\mathcal{H}}(x)) \sigma_{\mathcal{M}}(dx)$, which we assume to be finite. Consider an extension of the distribution to $T\mathcal{M}$ as

$$\mu(dz) = \frac{1}{Z_{T\mathcal{M}}} \exp(-H(x, v)) \lambda_{T\mathcal{M}}(dz), \quad (1)$$

where $\lambda_{T\mathcal{M}}(dz)$ is the Liouville measure of $T\mathcal{M}$, H is defined by

$$\begin{aligned} H(x, v) &= U_{\mathcal{H}}(x) + \frac{1}{2} v^T G(x)^{-1} v \\ &= U(x) + \frac{1}{2} \log \{(2\pi)^d \det G(x)\} + \frac{1}{2} v^T G(x)^{-1} v, \end{aligned} \quad (2)$$

for $(x, v) \in T\mathcal{M}$ and $Z_{T\mathcal{M}} = \int_{T\mathcal{M}} \exp(-H(x, v)) \lambda_{T\mathcal{M}}(dz)$, which is finite when $Z_{\mathcal{M}}$ is. We have that

$$\mu(dz) = \pi_{\mathcal{H}}(x) \sigma_{\mathcal{M}}(dx) \psi(x)(dv),$$

where $\psi(x)(dv)$ is simply the Gaussian measure on $T_x\mathcal{M}$ given by

$$\begin{aligned} \psi(x)(dv) &= \frac{1}{\sqrt{\{(2\pi)^d \det G(x)\}}} \\ &\times \exp\left(-\frac{1}{2} v^T G(x)^{-1} v\right) \sigma_{T_x\mathcal{M}}(dv), \end{aligned}$$

in local coordinates and $\sigma_{T_x\mathcal{M}}(dv)$ is the Lebesgue measure on $T_x\mathcal{M}$. In particular we have that μ has marginal distribution $\pi_{\mathcal{H}}$ with respect to the Hausdorff measure (Girolami and

Calderhead 2011; Byrne and Girolami 2013; Lelievre et al. 2010[Section 3.3.2]).

We will define a stochastic process which is a Riemannian version of the Randomized Hamiltonian Monte Carlo of Bou-Rabee and Sanz-Serna (2017). The stochastic process follows constrained Hamiltonian dynamics for an time duration t sampled from $t \sim \exp(\lambda)$ for some rate $\lambda > 0$ before an event. This event is a random velocity refreshment under the distribution $\psi(x)$.

Algorithm 1 defines Randomized time Riemannian Manifold Hamiltonian Monte Carlo (RT-RMHMC) with rate parameter $\lambda > 0$, and Hamiltonian dynamics governed by the Hamiltonian $H(x, v) = U_{\mathcal{H}}(x) + \frac{1}{2}v^T G(x)^{-1}v$ defined on $T\mathcal{M}$. This stochastic process has invariant measure $\mu(z) = \exp(-H(z))$ with respect to the Liouville measure on $T\mathcal{M}$.

Algorithm 1: RT-RMHMC

- Initialize x_0 arbitrarily on \mathcal{M} and sample $v_0 \sim \psi(x_0)$ on $T_{x_0}\mathcal{M}$ such that $(x_0, v_0) \in T\mathcal{M}$.
- Initialize $t_0 = 0$.
- for $k = 1, 2, \dots$ do
 - Update time via $t_k = t_{k-1} + \delta t$, where $\delta t \sim \exp(\lambda)$.
 - Evolve over $[t_{k-1}, t_k]$ Hamilton’s equations with initial condition $(x(t_{k-1}), v(t_{k-1})) = (x_{t_{k-1}}, v_{t_{k-1}})$.
 - Set $(x_s, v_s) = (x(s), v(s))$ for $s \in [t_{k-1}, t_k]$.
 - Set $x_{t_k} = x(t_k)$ and sample $v_{t_k} \sim \psi(x_{t_k})$ such that $(x_{t_k}, v_{t_k}) \in T\mathcal{M}$.

To sample from a distribution π with respect to the Hausdorff measure we define $U = -\log \pi$ under the assumption that π is integrable on \mathcal{M} .

We can define the generator for this stochastic process as

$$\mathcal{L}f(z) = X_H(f(z)) + \lambda(Qf(z) - f(z)), \tag{3}$$

where

$$Qf(x, v) := \frac{1}{\sqrt{\{(2\pi)^d \det G(x)\}}} \int_{T_x\mathcal{M}} \exp\left(-\frac{1}{2}\xi^T G(x)^{-1}\xi\right) f(x, \xi) d\xi,$$

is the transition kernel for a completely randomized velocity refreshment according to a Gaussian distribution on the tangent space $T_x\mathcal{M}$ and X_H is the Hamiltonian vector field associated to H , which may be defined by $X_H = \left(\frac{\partial H}{\partial v_i}, -\frac{\partial H}{\partial x_i}\right)$ in local coordinates. In the Appendix, we will prove that this is the generator of this stochastic process in Section A and invariance of the measure in Section B under a compactness assumption. Our main theoretical result about Algorithm 1 is the following.

Corollary 1 (Invariant measure for RT-RMHMC) *Let $(P_t)_{t \geq 0}$ be the transition semigroup of a simulation of Algorithm 1 with characteristics (φ, λ, Q) on $T\mathcal{M}$ and Hamiltonian $H \in C^2(T\mathcal{M})$, where (\mathcal{M}, g) is a compact smooth Riemannian manifold and φ is the Hamiltonian flow associated to the Hamiltonian. Let μ be the measure on $(T\mathcal{M}, \mathcal{B}(T\mathcal{M}))$ given by*

$$\mu(dz) \propto e^{-H(x,v)} d\lambda_{T\mathcal{M}}(z),$$

where $d\lambda_{T\mathcal{M}}$ is the Liouville measure of $T\mathcal{M}$. Then μ is invariant for RT-RMHMC.

3 Constrained symplectic integrator and metropolis-hastings adjustment

In this section, we will state some more broadly implementable versions of Algorithm 1 that are applicable when the Hamiltonian dynamics cannot be solved exactly. We start with a brief introduction to Lagrangian and Hamiltonian dynamics with constraints based on Lee et al. (2017)[Chapter 3].

Consider manifolds \mathcal{M} embedded in \mathbb{R}^d that can be described by algebraic equations

$$\mathcal{M} := \{x \in \mathbb{R}^d \mid c_i(x) = 0, i = 1, \dots, m\} \subset \mathbb{R}^d,$$

where $c_i : \mathbb{R}^d \rightarrow \mathbb{R} \ i = 1, \dots, m$ are continuously differentiable functions with linearly independent gradient functions for all $x \in \mathcal{M}$.

We refer to such a submanifold as an algebraic constraint manifold. We can express the Euler-Lagrange equations as an orthogonal projection of the Euler-Lagrange equations in \mathbb{R}^d onto the constraint manifold, hence we have

$$\frac{d}{dt} \left(\frac{\partial L(x, \dot{x})}{\partial \dot{x}} \right) - \frac{\partial L(x, \dot{x})}{\partial x} + \sum_{i=1}^m \lambda_i \frac{\partial c_i(x)}{\partial x} = 0,$$

where λ_i are Lagrange multipliers for each of the constraints. We can then define an augmented Lagrangian function $L^a : T^*\mathcal{M} \times \mathbb{R}^m \rightarrow \mathbb{R}$ by $L^a(x, \dot{x}, \lambda) = L(x, \dot{x}) - \sum_{i=1}^m \lambda_i c_i(x)$. Then the Euler-Lagrange equations can be expressed as

$$\frac{d}{dt} \left(\frac{\partial L^a(x, \dot{x}, \lambda)}{\partial \dot{x}} \right) - \frac{\partial L^a(x, \dot{x}, \lambda)}{\partial x} = 0,$$

and the augmented Hamiltonian function $H^a : T^*\mathcal{M} \times \mathbb{R}^m \rightarrow \mathbb{R}$ as $H^a(x, \mu, \lambda) = \mu \cdot \dot{x} - L^a(x, \dot{x}, \lambda)$, and we therefore obtain Hamilton’s equations (see Hartmann 2007)

$$\dot{x} = \frac{\partial H^a(x, \mu, \lambda)}{\partial \mu} \quad \dot{\mu} = -\frac{\partial H^a(x, \mu, \lambda)}{\partial x}.$$

We next introduce a new formulation of RT-RMHMC for constraint manifolds which we will use for numerical simulation. Note that a constraint manifold Hamiltonian Monte Carlo method was introduced in Brubaker et al. (2012), but with a deterministic duration parameter. We will use the same notation as that used in Brubaker et al. (2012) to introduce randomized time into this algorithm.

Let us denote our constraints $c(x) := (c_1(x), \dots, c_m(x))^T$ and let $C(x) = \frac{\partial c}{\partial x}$ denote the Jacobian of the constraints, which we assume to have full rank everywhere. Define a Hamiltonian of the constrained system as $H(x, v) = U_{\mathcal{H}}(x) + K(v)$, where $K(v) = \frac{1}{2}v^T G(x)^{-1}v$ is the kinetic energy and v lies in the cotangent space, $\mathcal{T}_x^* \mathcal{M} = \{v \mid C(x) \frac{\partial H}{\partial v}(x, v) = 0\}$. The dynamics of the constrained system in terms of the Hamiltonian is thus given by

$$\dot{v} = -\frac{\partial H}{\partial x} + C(x)^T \lambda, \quad \dot{x} = \frac{\partial H}{\partial v}, \quad \text{such that } c(x) = 0,$$

where we remark that we can naturally identify the tangent and cotangent spaces and bundles.

We let $\pi_{\mathcal{H}}$ be our target measure with respect to the Hausdorff measure and $U_{\mathcal{H}}(x) = -\log \pi_{\mathcal{H}}(x)$ be the potential energy of our constrained system. We can then simulate the constrained Hamiltonian dynamics using the RATTLE scheme (Andersen 1983). However, if we know $\pi_{\mathcal{H}}$ explicitly we can avoid computation of the metric tensor by assuming our system is isometrically embedded in Euclidean space. Under this assumption we can then consider Algorithm 2, which is an explicit algorithm for simulation of Randomized time constrained Hamiltonian Monte Carlo (RT-CHMC). We will discuss and justify the embedding assumption further in Sect. 3.1.

Algorithm 2: RT-CHMC

- Initialize x_0 arbitrarily on \mathcal{M} and sample $v_0 \sim \mathcal{N}(0, I \mid C(x_0)v_0 = 0)$.
 - Initialize $t_0 = 0$.
 - for $k = 1, 2, \dots$ do
 - Update time via $t_k = t_{k-1} + \delta t$, where $\delta t \sim \exp(1/\lambda)$.
 - Evolve over $[t_{k-1}, t_k]$ Hamilton’s equations subject to constraints $c(\cdot)$ with initial condition $(x(t_{k-1}), v(t_{k-1})) = (x_{t_{k-1}}, v_{t_{k-1}})$ and Hamiltonian $H(x, v) = -\log \pi_{\mathcal{H}}(x) + \frac{1}{2}v^T v$.
 - Set $(x_s, v_s) = (x(s), v(s))$ for $s \in [t_{k-1}, t_k]$
 - Set $x_{t_k} = x(t_k)$ and sample $v_{t_k} \sim \mathcal{N}(0, I \mid C(x_{t_k})v_{t_k} = 0)$
-

In Algorithm 2 we sample the Gaussian distribution on the tangent space at a point on \mathcal{M} . We can do this by sampling a vector whose components are independent standard normal random variables and then projecting this orthogonally. To orthogonally project a momentum vector onto $T^* \mathcal{M}$ and

correctly resample the momentum in Algorithm 2 at $x \in \mathcal{M}$ we apply the projector

$$P_{\mathcal{M}}(x) := I - C(x)^T (C(x)C(x)^T)^{-1} C(x).$$

Proposition 1 *If $v' \sim \mathcal{N}(0, I)$ then $v = P_{\mathcal{M}}(x)v'$ is distributed according to $v \sim \mathcal{N}(0, I \mid C(x)v = 0)$*

Proof Can be found in Graham et al. (2022). □

3.1 Embedded manifolds

We next introduce the theory of manifold embeddings as it was presented in Byrne and Girolami (2013) to show that numerical simulation of RT-CHMC is in fact simulation of RT-RMHMC on constraint manifolds.

If we know the form of the distribution $\pi_{\mathcal{H}}$ with respect to the Hausdorff measure, then we can avoid the computation of the metric tensor and the lack of a global coordinate system (Byrne and Girolami 2013). We achieve this using isometric embeddings, remarking that every Riemannian manifold can be isometrically embedded in Euclidean space due to the Nash embedding theorem (Nash 1956). If we have an isometric embedding $\xi : \mathcal{M} \rightarrow \mathbb{R}^n$, then considering a path $q(t)$ on \mathcal{M} , the path $x(t) = \xi(q(t))$ is such that $\dot{x}_i(t) = \sum_j \frac{\partial x_i}{\partial q_j} \dot{q}_j(t)$. The phase space (q, p) , where $\dot{q} = G^{-1}p$, can then be transformed to the embedded phase space (x, v) , where

$$v = \dot{x} = XG(q)^{-1}p = X(X^T X)^{-1}p, \quad \text{where } X_{ij} = \frac{\partial x_i}{\partial q_j},$$

since $G = X^T X$ due to the fact that the embedding is isometric and preserves inner products (see (Byrne and Girolami 2013)). Now the Hamiltonian (Eq. 2) is

$$H(x, v) = -\log \pi_{\mathcal{H}}(x) + \frac{1}{2}v^T v,$$

in terms of coordinates (x, v) . When considering sampling of the velocities in Algorithm 1 and Algorithm 2, since $p \sim \mathcal{N}(0, G(q))$, we have

$$v \sim \mathcal{N}(0, X(X^T X)^{-1}X^T),$$

where $X(X^T X)^{-1}X^T$ is the orthogonal projection onto the tangent space of the embedded manifold (Byrne and Girolami 2013). Therefore we can sample from $\mathcal{N}(0, I)$ and project onto the tangent space to obtain a necessary sample. The Hamiltonian is thus expressed in a form which is independent of the metric (provided we know the density with respect to the Hausdorff measure). We next introduce the RATTLE

scheme for numerical integration on a manifold (Leimkuhler and Reich 2004)[Chapter 7]:

$$\begin{aligned} x_{n+1} &= x_n + \Delta t v_{n+1/2} \\ v_{n+1/2} &= v_n - \frac{\Delta t}{2} \nabla_x U(x_n) - \frac{\Delta t}{2} C(x_n)^T \lambda_{(r)}^n \\ c(x_{n+1}) &= 0 \\ v_{n+1} &= v_{n+1/2} - \frac{\Delta t}{2} \nabla_x U(x_{n+1}) - \frac{\Delta t}{2} C(x_{n+1})^T \lambda_{(v)}^{n+1} \\ C(x_{n+1})v_{n+1} &= 0, \end{aligned}$$

where we solve for $\lambda_{(r)}^n$ and $\lambda_{(v)}^{n+1}$ at each iteration so that the iterates lie in the tangent bundle. We solve for $\lambda_{(r)}^n$ (a non-linear system of equations) by cycling through the constraints, adjusting one multiplier at each iteration. Denote by C_i the i th row of C and we first initialize

$$Q := \bar{x}_{n+1} = x_n + \Delta t v_n - \frac{\Delta t^2}{2} \nabla_x U(x_n).$$

Next we cycle through the list of constraints one after another as follows: for each $i = 1, \dots, m$ compute

$$\Delta \Lambda_i := \frac{c_i(Q)}{C_i(Q)C_i(x_n)},$$

and update Q by $Q := Q - C_i(x_n)^T \Delta \Lambda_i$ until $c_i(Q) < tol$ for all $i = 1, \dots, m$, where tol is a certain prescribed tolerance. Then we set $x_{n+1} = Q$ and have $x_{n+1} \in \mathcal{M}$ within the tolerance. (Note that other stopping criteria could be used (see (Ortega and Rheinboldt 2000)).) We solve for $\lambda_{(v)}^{n+1}$ by solving the linear system:

$$\left(C(x_n)C(x_n)^T \right) \lambda_{(v)}^n = C(x_n) \left(\frac{2}{\Delta t} v_{n-1/2} - \nabla_x U(x_n) \right).$$

Once the linear system has been solved we obtain $(x_{n+1}, v_{n+1}) \in T^* \mathcal{M}$.

Theorem 2 *Let \mathcal{M} be a constraint manifold. Let $H \in C^2(T\mathcal{M})$, the RATTLE numerical integrator of the Hamiltonian system defined by H in $T\mathcal{M}$ is symmetric, symplectic and of order 2. Further it respects the manifold constraints.*

Proof Given in Leimkuhler and Skeel (1994). □

3.2 Metropolis hastings adjustment

Let $\Psi_{\Delta t}^L : T\mathcal{M} \rightarrow T\mathcal{M}$ be the numerical integrator defined by L steps of RATTLE with stepsize Δt . This integrator approximates the Hamiltonian dynamics. For theoretical purposes we will also define the map $N : T\mathcal{M} \rightarrow T\mathcal{M}$ which negates the momentum term i.e. $N(x, v) \equiv (x, -v)$. Note that this leaves the Hamiltonian invariant and due to the fact

that the momentum is resampled this has no affect on the samples from $\pi_{\mathcal{H}}$. We will define the following Metropolized RT-RMHMC, where we sample $T \sim \exp(\lambda)$ and fix a maximum time length Δt_{\max} below the stability threshold of the numerical integrator. Then we choose the number of leapfrog steps L to be $\lceil T/\Delta t_{\max} \rceil$. Having chosen L in this way, we set $\Delta t = T/L \leq \Delta t_{\max}$. At each step we perform L RATTLE steps with stepsize Δt . We propose this method of discretisation instead of purely randomizing the stepsize and fixing a number of leapfrog steps to avoid numerical instabilities in the numerical integrator. One could also propose fixing a stepsize within the numerical stability threshold of the integrator and simply sampling an integer number of leapfrog steps geometrically to randomize the time. However our proposed method closer relates to the continuous dynamics without the issues due to numerical instabilities.

Remark 1 For large choices of stepsize Δt it has been shown that $\Phi_{\Delta t}^L$ is not reversible where RATTLE is used to integrate on the manifold, see (Lelièvre et al. 2019; Zappa et al. 2018). In Lelièvre et al. (2019) they propose to combat this by adding a reversibility check incorporated into the Metropolis-Hastings adjustment, although in practice such checks may be neglected in favor of an implicit assumption that Δt is sufficiently small to avoid non-reversibility issues. We will investigate this further in Sect. 5.

In light of Remark 1, we include $\text{Rev}(\cdot)$ as a additional (optional) accept-reject condition which implements a reversibility check (following (Lelièvre et al. 2019)). In numerical experiments we examine the stepsize threshold where the reversibility condition fails (See Fig. 1).

Remark 2 Our framework can be adapted to handle inequality constraints by incorporating an additional rejection condition in the Metropolis-Hastings step, which rejects samples which aren't within the boundary.




Fig. 1 Ratio of samples out of 10^6 samples which don't satisfy reversibility check for different choices of Δt for the BVMF distribution with parameters $A = \text{diag}(-1000, 0, 1000)$ and $c = (100, 0, 0)$

Algorithm 3: RT-CHMC with Metropolis-Hastings step

- Initialize Δt_{\max} within the stability threshold.
- Initialize x_0 arbitrarily on \mathcal{M} and sample $v_0 \sim \mathcal{N}(0, I \mid C(x_0)v_0 = 0)$.
- Initialise $t_0 = 0$.
- for $k = 1, 2, \dots$ do
 1. Sample $v_{k-1} \sim \mathcal{N}(0, I \mid C(x_{k-1})v_{k-1} = 0)$.
 2. – Sample $T \sim \exp(\lambda)$.
 - Set $L = \lceil T/\Delta t_{\max} \rceil$ and $\Delta t = T/L$.
 - Set $(x^*, v^*) = \Psi(x_{k-1}, v_{k-1}) = \text{Rev}(N(\Psi_{\Delta t}^L(x_{k-1}, v_{k-1})))$.
 - Accept (x^*, v^*) with probability $\min\{1, \exp\{H(x, v) - H(x^*, v^*)\}\}$ and set $(x_k, v_k) = (x^*, v^*)$.
 - Otherwise $(x_k, v_k) = (x_{k-1}, v_{k-1})$.

This will be used in our application in Section 5.4 to impose a half-normal prior on some dimensions of our Bayesian model.

4 Ergodicity

We will now prove ergodicity and exact invariance of the desired measure of the discrete time algorithm with Metropolis-Hastings adjustment. We will provide ergodicity under two assumptions by the same technique as Brubaker et al. (2012) and restating some of their results.

Proposition 3 Assuming that Ψ for $\Delta t_{\max} > 0$ then μ is invariant with respect to the Markov kernel proposed in Algorithm 3.

Proof See Section C of the Appendix. □

Assumption 1 Let $\mathcal{M} \in \{x \in \mathbb{R}^n \mid c(x) = 0\}$ be Riemannian manifold which is connected, smooth and differentiable. We assume that $\partial c/\partial x$ is full rank everywhere.

Assumption 2 Let \mathcal{M} be a Riemannian manifold which satisfies Assumption 1. For $x \in \mathcal{M}$ we define $\mathcal{B}_r(x) = \{x' \in \mathcal{M} \mid d(x', x) \leq r\}$ to be the geodesic ball of radius r of x . We assume that there exists a $r > 0$ such that for every $x \in \mathcal{M}$ and $x' \in \mathcal{B}_r(x)$ there exists a unique choice of Lagrange multipliers and velocity $v \in T_x \mathcal{M}$, $v' \in T_{x'} \mathcal{M}$ for which $(v', x') = \Psi_{\Delta t}^L(v, x)$ for sufficiently small Δt .

Theorem 4 (Accessibility) Let $U \in C^2(\mathcal{M})$, and assuming Assumption 1. For any $x_0, x_1 \in \mathcal{M}$ and Δt sufficiently small, there exists finite $v_0 \in T\mathcal{M}$, $v_1 \in T\mathcal{M}$ and Lagrange multipliers λ_0, λ_1 such that $(v_1, x_1) = \Psi_{\Delta t}(v_0, x_0)$.

Proof Found in Brubaker et al. (2012)[Theorem 2] and is an extension of the results of Marsden and West (2001)[Theorem 2.1.1] and Hairer et al. (2006)[Theorem 5.6, Section IX.5.2]. □

Theorem 5 (μ -irreducible) Let $U \in C^2(\mathcal{M})$, and under Assumptions 1 and 2 we have that for any $x \in \mathcal{M}$, and measurable set $A \subset \mathcal{M}$ with positive measure. Then there exists an $n \in \mathbb{N}$ such that

$$K^n(x, A) > 0,$$

where K denotes the marginal transition kernel defined on \mathcal{M} of Algorithm 3.

Proof See Section C of the Appendix. □

Lemma 6 (Aperiodic) Let $U \in C^2(\mathcal{M})$ and under Assumptions 1 and 2 Algorithm 3 is aperiodic.

Proof Proof given in Brubaker et al. (2012)[Lemma 1]. □

Theorem 7 (Ergodicity) Let $U \in C^2(\mathcal{M})$ and under Assumptions 1 and 2 we have for μ -almost all starting values x

$$\lim_{t \rightarrow \infty} \int_{\mathcal{M}} |K^t(x, y) - \pi_{\mathcal{H}}(y)| \sigma_{\mathcal{M}}(dy) = 0.$$

Proof Since Algorithm 3 is μ -invariant by Theorem 3, μ -irreducible by Theorem 5 and aperiodic by Theorem 6, the required result holds by Tierney (1994)[Theorem 1]. □

5 Numerical results

We perform numerical simulations of the RT-CHMC algorithm and compare to the CHMC algorithm of Brubaker et al. (2012); Girolami and Calderhead (2011), specifically exploring the underlying dynamics of the two processes. MCMC schemes are used to approximate expected values of certain functions f over some distribution with probability density function π

$$\mathbb{E}_{\pi}(f) = \int f(x)\pi(x)dx,$$

where we can estimate this quantity using our MCMC scheme by

$$\bar{f} := \mathbb{E}_{\pi}(f) \approx \frac{1}{M} \sum_{i=1}^M f(X^i),$$

where X^i is the Markov chain from our MCMC method. We quantify the convergence rate associated to approximation of $\mathbb{E}_{\pi}(f)$ by considering the integrated autocorrelation function and essential sample size.

5.1 g-BAOAB

As a comparison method we implemented the g-BAOAB integrator of Leimkuhler and Matthews (2016), a numerical integrator for constrained underdamped Langevin dynamics. Constrained underdamped Langevin dynamics can be described by

$$\begin{aligned} \dot{x} &= v, \\ 0 &= C(x)v, \\ \dot{v} &= -\nabla_x U(x) - \gamma v + \sqrt{2\gamma}R(t) - C(x)^T \lambda, \\ 0 &= c(x), \end{aligned}$$

where γ is a friction coefficient and $R(t)$, is a vector-valued, stationary, zero-mean Gaussian process. The numerical integrator g-BAOAB is a splitting method for such dynamics, which uses similar constrained integrators as that of RT-CHMC. We note that g-BAOAB is a biased sampling algorithm due to the error in the numerical integrator. For a full description of g-BAOAB and a discussion of the sampling error we refer to Leimkuhler and Matthews (2016).

5.2 Computational cost

The cost of propagation steps are essentially the same as those of g-BAOAB. In situations where there are relatively few constraints compared to the dimension of the parameter space, the constraint cost should be affordable. This happens even in complex applications like molecular dynamics of proteins, where the constraint solver rarely introduces a cost greater than a few percent of the total cost of, say, long-ranged force evaluations (Xie et al. 2000). Obviously the practicality of the scheme will depend on the problem itself. The cost of solving the constraints using the semi-explicit SHAKE/RATTLE or geodesic integrator steps is expected to be lower than the cost to integrate the equations of CHMC in the general metric setting (using unconstrained ODEs), since the nonseparable Hamiltonian structure demands an implicit symplectic integration.

5.3 Test examples

We next provide examples of distributions on implicitly defined manifolds embedded in Euclidean space, with the distributions defined with respect to the Hausdorff measure of the manifold. We will consider two types of constraint manifolds: spheres and Stiefel manifolds.

Bingham-Von Mises-Fisher distribution on S^n

The first test case is the Bingham-Von Mises-Fisher (BVMF) distribution defined on the n -dimensional sphere embedded

in \mathbb{R}^{n+1} , that is $S^n := \{x \in \mathbb{R}^{n+1} \mid \sum_{i=1}^{n+1} x_i^2 = 1\}$. The BVMF distribution is the exponential family on $S^n \subset \mathbb{R}^{n+1}$ with density of the form

$$\pi_{\mathcal{H}}(x) \propto \exp \{c^T x + x^T A x\},$$

where $c \in \mathbb{R}^{n+1}$ and $A \in M_{n+1}(\mathbb{R})$ is a symmetric matrix.

We compare the integrated autocorrelation (IAC) of $-\log \pi_{\mathcal{H}}$ of the RT-CHMC method to that of the CHMC method introduced in Brubaker et al. (2012) for a number of distributions with parameters defined in the captions. We also compare the maximum IAC of x_i for $i = 1, \dots, n$ to compare the worst efficiency of the mixing in all dimensions. We compare methods by setting the event rate parameter λ of RT-CHMC to be the deterministic duration parameter of CHMC (running the dynamics for this duration before momentum randomization). We then compute the integrated autocorrelation of $-\log \pi_{\mathcal{H}}$ and x_i for $i = 1, \dots, n$ for the two methods for varying choices of λ by a Monte Carlo averaging procedure as described in Section Appendix E. Regarding the reversibility issue for large choices of stepsize (as discussed in Section 3.2), for the geometries and distributions chosen, this is shown to exhibit behaviour as in Fig. 1, where there is a dramatic change in reversibility failure for a small change in step-size. Before this point all samples generated satisfy reversibility conditions. We simply chose stepsizes which are below this threshold in our simulations.

The results are presented in Fig. 2. We choose the stepsize in RATTLE to be $\Delta t = 0.01$ and sample $N = 1,000,000$ events with a burn time of 10% of samples before we compute the Monte Carlo average. We also use lags of up to $M = N/50$, 2 percent of the number of samples used to estimate the IAC. As our choice of Δt is small, the acceptance rate is high so this process is close to the continuous version. The IAC compares the efficiency of the continuous processes.

Remark 3 Due to the symmetry in the x_3 -coordinate ($x_3 \mapsto -x_3$) in the BVMF distribution with parameters $A = \text{diag}(-1000, 0, 1000)$ and $c = (100, 0, 0)$, the distribution is bimodal. However in practice we consider this a unimodal distribution as the probability of transferring between modes is extremely small. This can be seen in Fig. 3, as the second mode is at $\theta = \pi$. Our dynamics stay near the mode at $(0, 0, 1)$ and do not visit the other mode in all our simulations.

In our first example and Fig. 2 we can see that the regularity of the quality of samples with respect to the duration parameter is poor when a deterministic duration parameter is used and nearly uniform across a wide interval for a randomized duration with the same expected value. This is illustrated in Fig. 3, where a small change in duration parameter causes the dynamics dramatically slows convergence and due to very slow mixing. The fact that CHMC behaves erratically for large mean duration parameters may not be very surprising

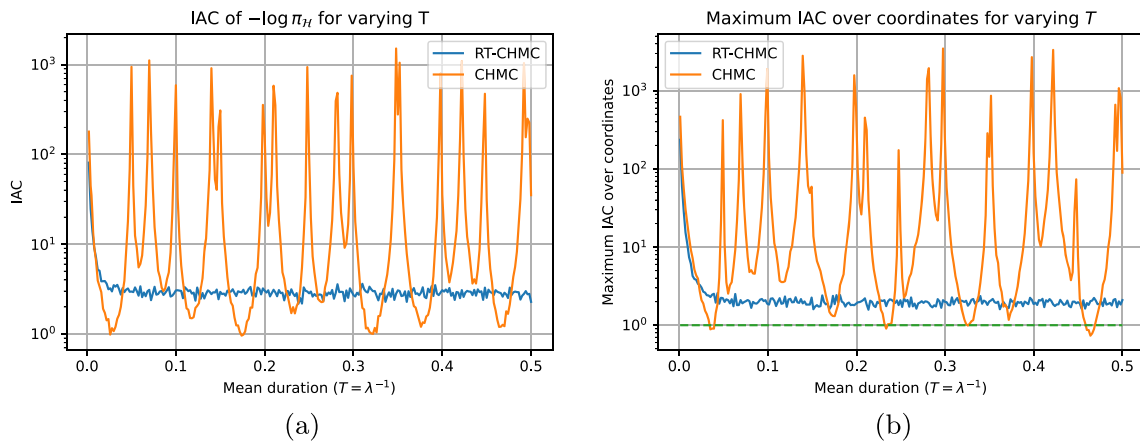


Fig. 2 IAC estimates for different choices of λ for the BVMF distribution with parameters $A = \text{diag}(-1000, 0, 1000)$ and $c = (100, 0, 0)$ averaged over 20 independent runs.. Left: IAC of $-\log \pi_{\gamma_{\mathcal{L}}}$. Right: Maximum IAC over x_1, x_2 and x_3

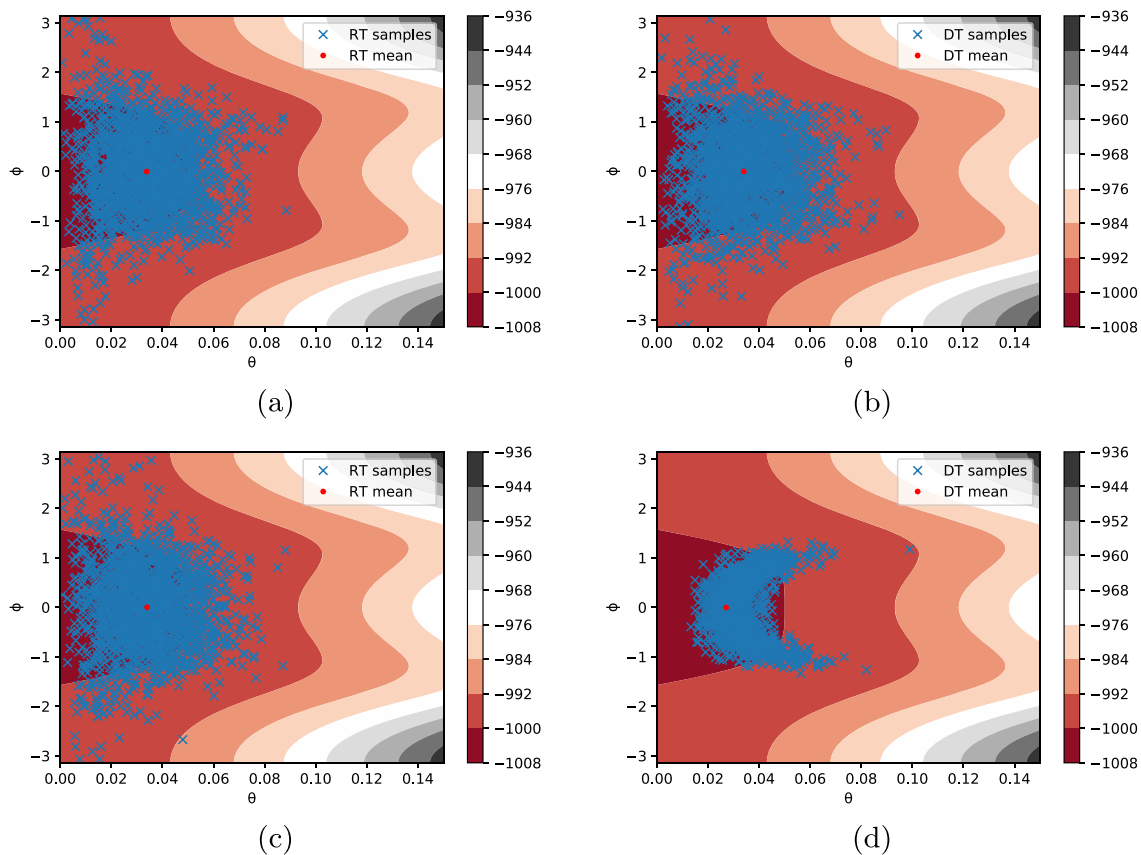


Fig. 3 Contour plot of $-\log \pi_{\gamma_{\mathcal{L}}}$ for the BVMF distribution with parameters $A = \text{diag}(-1000, 0, 1000)$ and $c = (100, 0, 0)$. The axis being a 2D parameterisation of S^2 . The points are 2000 samples after a 8000

sample burn in time. Upper left: RT-CHMC for $\lambda^{-1} = 0.09$. Upper right: CHMC for $\lambda^{-1} = 0.09$. Lower left: RT-CHMC for $\lambda^{-1} = 0.1$. Lower right: CHMC for $\lambda^{-1} = 0.1$

to some readers as the theoretical convergence bound for HMC without randomization requires a limit on the duration T (see (Mangoubi and Smith 2017)). This is due to the fact that when T is set too large, the coupling argument breaks down.

We next compare efficiencies using the metric gradient evaluations per effective sample size, which tells us the number of gradient evaluations needed for one independent sample in estimating our observables. We compare this metric for varying choices of step-size up to when the reversibility condition is broken and the numerical integra-

tor becomes unstable. Our observables will be $-\log \pi_{\mathcal{H}}$ and x_i for $i = 1, \dots, n$. We can see as in Figs. 4 and 5 CHMC (with deterministic time) exhibits the same behaviour as in Fig. 2 for all choices of step-size, which is not the case for RT-CHMC. We next compare the efficiency of the method with the g-BAOAB constrained Langevin integrator, we find in Figures 4 and 5 that g-BAOAB outperforms RT-CHMC for large choices of the friction parameter γ (for this example $\gamma = 50$). g-BAOAB has no Metropolis-Hastings adjustment and hence is a biased sampling method. The bias in the samples creates errors in computed observables. For large choices of γ , this bias is dramatically reduced, but use of high friction may slow convergence of metastable systems.

To explore this we next consider a bimodal distribution from Byrne and Girolami (2013) in Fig. 6. It is shown in Fig. 6 that g-BAOAB incurs bias for large stepsizes and convergence is slow for large choices of the friction parameter for this metastable system. The figure also shows that this is not the case for RT-CHMC. In Fig. 4 we choose step-sizes up to which the integrator is reversible and stable.

The efficiency of the methods with their optimal choices of parameters is comparable, but RT-CHMC is much less sensitive with respect to the choice of parameters (stepsize and

number of leapfrog steps) compared to CHMC, so RT-CHMC is more reliable from this point of view. This is important as it is hard to know an appropriate choice of parameters a priori and the integration length between samples might have to be arbitrarily small for CHMC to be efficient.

Von Mises-Fisher distribution on $\mathbb{V}_{d,p}$

Definition 1 A Stiefel manifold $\mathbb{V}_{d,p}$ is the set of $d \times p$ matrices X such that $X^T X = I$.

These arise in many statistical problems which are discussed in Byrne and Girolami (2013). Applications include dimensionality reduction such as is used in factor analysis, principal component analysis (Jolliffe 2002) and directional statistics (Mardia et al. 2000). These are a generalisation of orthogonal groups. The von Mises-Fisher distribution on Stiefel manifolds is defined by the density

$$p_{vMF}(X) \propto \exp(\text{Tr}(F^T X)) = \exp(\langle f_1, x_1 \rangle + \dots + \langle f_p, x_p \rangle),$$

where x_i and f_i are the columns of F and X . We simulate IAC estimates for two example distributions for varying duration

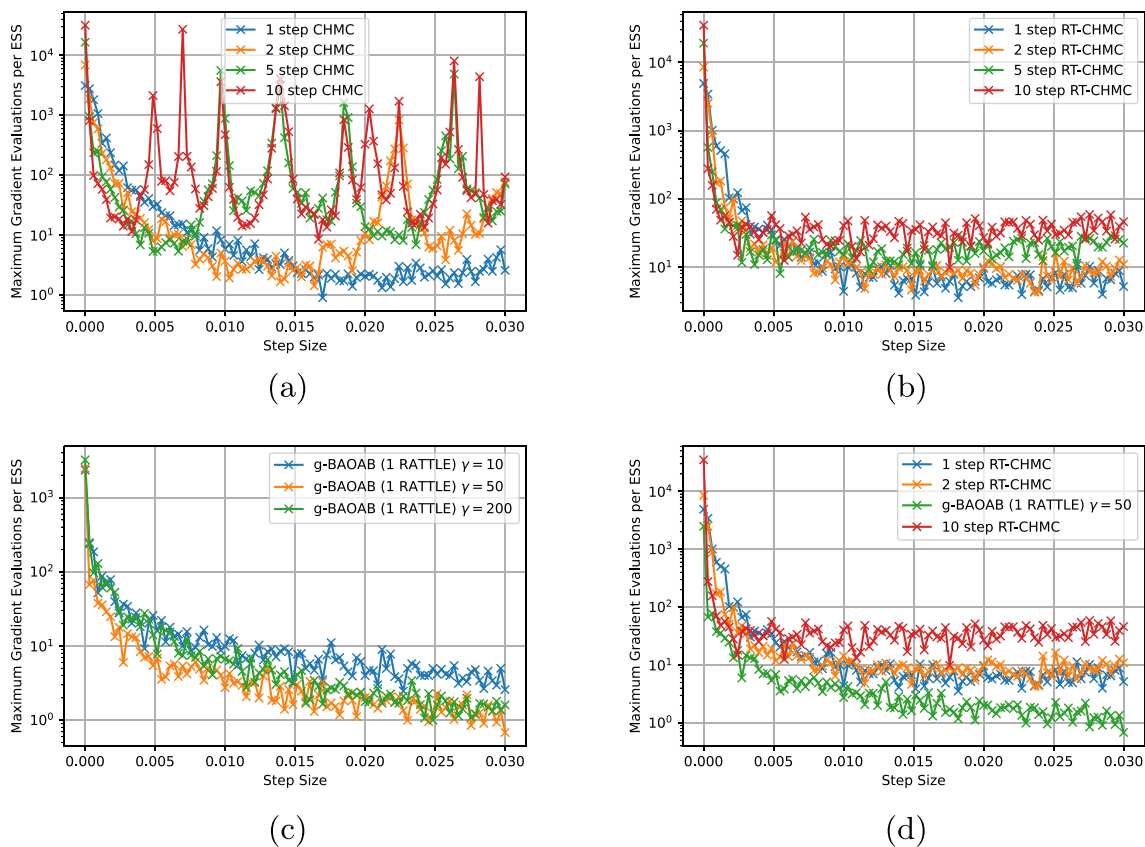


Fig. 4 Gradient evaluation per ESS estimates of $-\log \pi_{\mathcal{H}}$ for the BVMF distribution using 100,000 samples with parameters $A = \text{diag}(-1000, 0, 1000)$ and $c = (100, 0, 0)$ and for varying choices of

step-size. Upper left: CHMC. Upper right: RT-CHMC. Lower left: g-BAOAB. Lower right: g-BAOAB and RT-CHMC

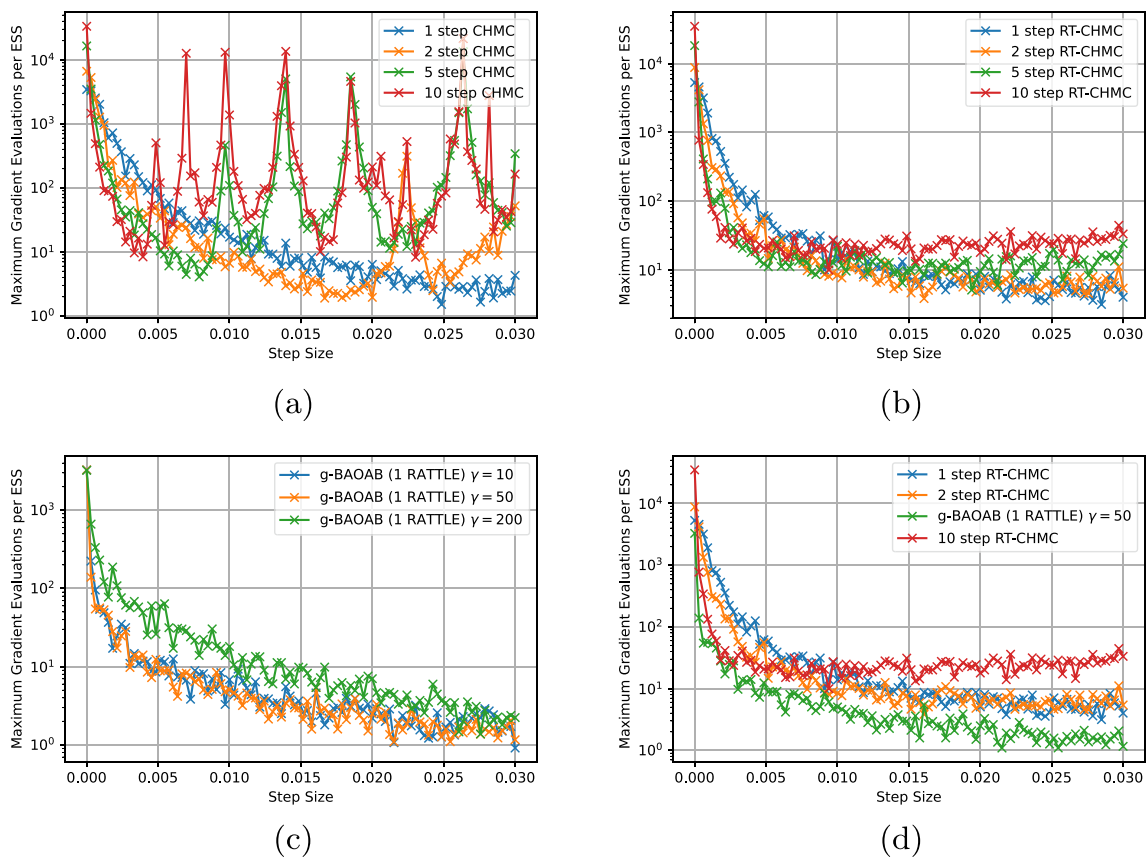


Fig. 5 Maximum gradient evaluation per ESS estimates over x_1, x_2 and x_3 for the BVMF distribution using 100,000 samples with parameters $A = \text{diag}(-1000, 0, 1000)$ and $c = (100, 0, 0)$ and for varying choices

of step-size. Upper left: CHMC. Upper right: RT-CHMC. Lower left: g-BAOAB. Lower right: g-BAOAB and RT-CHMC




Fig. 6 Monte Carlo average of $-\log \pi_H$ for the BVMF distribution with parameters $A = \text{diag}(-20, -10, 0, 10, 20)$ and $c = (40, 0, 0, 0, 0)$ with $N = 5 \times 10^6$ samples. Left: 10 step RT-CHMC. Middle: g-BAOAB with $\gamma = 2$. Right: g-BAOAB with $dt = 0.01$

parameters. In the simulations we use a stepsize of $\Delta t = 0.01$ and 100,000 samples in each IAC estimate. The results are shown in Fig. 7. We can see similar behaviour as the easier example on the Sphere. Both examples it is clear CHMC is much more sensitive to the mean duration and hence with respect to stepsize and number of leapfrog steps. We note that $\text{Skew}(2, -45, -4)$ denotes the 3 by 3 skew-symmetric matrix with up triangular entries 2, -45 and -4 .

5.4 High dimensional covariance estimation

In many statistical applications for analysing high dimensional data sets it is necessary to estimate sample covariances. This can be challenging when the number of dimensions is larger than the number of data points, as the sample covariance estimator does not work well in such cases. Lam (2020) provides a review of high-dimensional covariance estimation and applications in principal component analysis (Shen et al. 2016), cosmological data analysis (Joachimi 2017) and




Fig. 7 IAC estimates of $-\log \pi_T$ for different choices of λ for the VMF distribution on $O(3)$ and $V_{18,3}$ with parameters given in the captions averaged over 20 independent runs. Left: VMF distribution on

$O(3)$ with parameters $F = A := \text{Skew}(2, -45, -4)$. Right: VMF distribution on $V_{18,3}$ with parameters $F = [I, -A, I, -A, I, -A]^T$

finance (Lam 2016). Lam (2020) focuses on the setup where the matrix dimension is diverging or even larger than the sample size. In this setting one needs to estimate the population covariance matrix Σ of a set of n, p -dimensional data vectors, which we assume are drawn from an underlying distribution.

One estimator is the sample covariance matrix, which is defined by $\Sigma_S = 1/n \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T$, where $\bar{\mathbf{x}} = 1/n \sum_{k=1}^n \mathbf{x}_k$ is the sample mean. However this is a poor estimator of Σ when p is large compared to the sample size n (due to rank deficiency). A way to combat this is to consider regularized covariance matrix estimators which include structural assumptions on the covariance matrix Σ .

One such method which has been proposed assumes the structure of a low rank matrix plus a sparse matrix (see (Ross 1976) and Lam (2020)). This structure is known as a spiked covariance structure and has been studied in Bouchard et al. (2020); Lam (2020) and Cai et al. (2015). There have been interesting applications to finance (Fan et al. 2008), chemometrics (Kritchman and Nadler 2008), and astronomy (Joachim 2017). The covariance matrix Σ is assumed to be expressible in the form

$$\Sigma = X D_1 X^T + D_2,$$

where X is a Stiefel manifold of dimension $p \times m$ for $p \gg m$ and D_i for $i = 1, 2$ are diagonal matrices of dimensions $m \times m$ and $p \times p$ respectively. The motivation for this structure is that we assume that lower dimensional variables \mathbf{y}_i can describe the data \mathbf{x}_i such that $\mathbf{x}_i = X \mathbf{y}_i + \epsilon_i$, where X is a $p \times m$ matrix with orthogonal columns. We have that

$$\Sigma = X \Sigma_y X^T + \Sigma_\epsilon,$$

which we interpret as a low rank matrix (rank m) plus a sparse matrix. We take Σ_y and Σ_ϵ to be diagonal, which is

an approximation of the spiked covariance structure (Chamberlain and Rothschild 1982).

Assume a uniform prior on X with respect to the Hausdorff measure on X . Further assume a half-normal prior of the diagonal entries of D_i for $i = 1, 2$ to ensure positive definiteness. We also consider the following likelihood for the covariance estimation

$$\mathcal{L}(\Sigma | \mathbf{x}_1, \dots, \mathbf{x}_p) = (2\pi)^{-np/2} \prod_{i=1}^n \det(\Sigma)^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{x}_i - \bar{\mathbf{x}})^T \Sigma^{-1}(\mathbf{x}_i - \bar{\mathbf{x}})\right).$$

We introduce the posterior distribution $p(\Sigma | \mathbf{x}_1, \dots, \mathbf{x}_p) \propto \mathcal{L}(\Sigma | \mathbf{x}_1, \dots, \mathbf{x}_p) p(\Sigma)$, where $p(\Sigma) = p(X) p(D_1) p(D_2)$ for $X \sim \mathcal{U}(V_{p,m})$, $D_{1jj} \sim \mathcal{N}_+(0, \sigma_1^2)$ for $j = 1, \dots, m$ and $D_{2jj} \sim \mathcal{N}_+(0, \sigma_2^2)$ for $j = 1, \dots, p$ and \mathcal{N}_+ denotes the half-normal distribution. Define the potential $U : V_{p,m} \times \mathbb{R}^m \times \mathbb{R}^p \mapsto \mathbb{R}$ by $U(X, \mathbf{d}_1, \mathbf{d}_2) = -\log \mathcal{L}(\Sigma(X, \mathbf{d}_1, \mathbf{d}_2) | \mathbf{x}_1, \dots, \mathbf{x}_p) - \log p(X) - \log p(\mathbf{d}_1) - \log p(\mathbf{d}_2)$ with forces given by

$$\frac{\partial U}{\partial X_{ij}} = \frac{\partial U}{\partial \Sigma_{kl}} \frac{\partial \Sigma_{kl}}{\partial X_{ij}} = \left(\frac{1}{2} n (\Sigma^{-1})_{kl}^T + \frac{1}{2} \sum_{r=1}^n (\mathbf{x}_r - \bar{\mathbf{x}})^T B^{kl} (\mathbf{x}_r - \bar{\mathbf{x}}) \right) \frac{\partial \Sigma_{kl}}{\partial X_{ij}},$$

where $[B^{kl}]_{ij} = -(\Sigma^{-1})_{ik} (\Sigma^{-1})_{lj}$ and

$$\frac{\partial \Sigma_{kl}}{\partial X_{ij}} = \begin{cases} X_{kj} D_{jj}^1 & k \neq i, l = i, \\ D_{jj}^1 X_{lj} & k = i, l \neq i, \\ 2 X_{ij} D_{jj}^1 & k = i, l = i, \\ 0 & \text{otherwise.} \end{cases}$$

We also have that

$$\frac{\partial U}{\partial d_{ij}} = \frac{\partial U}{\partial \Sigma_{kl}} \frac{\partial \Sigma_{kl}}{\partial d_{ij}} + \frac{d_{ij}}{\sigma_1^2},$$

for $i = 1, 2$ and where

$$\frac{\partial \Sigma_{kl}}{\partial d_{1j_1}} = X_{kj_1} X_{lj_1} \quad \text{and} \quad \frac{\partial \Sigma_{kl}}{\partial d_{2j_2}} = \begin{cases} 1 & \text{if } k = l = j_2 \\ 0 & \text{otherwise,} \end{cases}$$

for $j_1 = 1, \dots, m$ and $j_2 = 1, \dots, p$. We will use the likelihood and its gradient for implementing our RT-CHMC algorithm for such models.

Covariance estimation for cosmological data

We consider an application of high dimensional covariance estimation in cosmological data analysis introduced in Joachimi (2017) and discussed in Lam (2020). The data is taken from Joachimi (2017) and consists of covariances of two-point correlation functions of cosmic weak lensing. It is simulated using coupled log-normal random fields from angular power spectra. For further information we refer the

reader to Joachimi (2017). We will test this method using $2n/3$ data vectors, where n is the dimension of the data vectors. Therefore we are in the setting where the dimension of the covariance matrix is larger than the number of samples. For our low rank plus sparse structure we choose $m = p/6 \ll p$ and to ensure fast convergence we will normalize the data entry-wise and initialize our Markov chain via a eigenvalue decomposition of the sample covariance. We initialize our Markov chain as $\Sigma_0 = X^T D_1 X + D_2 \approx \Sigma_S$, using the sample covariance matrix Σ_S and where D_2 is the diagonal of Σ_S and $X^T D_1 X$ corresponds to the eigenvalue decomposition of $\Sigma_S - D_2$, but with the p largest eigenvectors. After the covariance of the normalized data is estimated it can easily be rescaled to match the real data via entry-wise multiplication with the outer product of the entry-wise standard deviations. We compare our method to a maximum a posterior (MAP) estimate of the covariance matrix which uses a simple constrained gradient descent algorithm with Lagrange multipliers to ensure that the “low-rank plus sparse” structure is maintained. We compare the Bayesian and MAP approaches using a relative Frobenius norm and a covariance metric introduced by Förstner and Moonen (2003) which is defined by

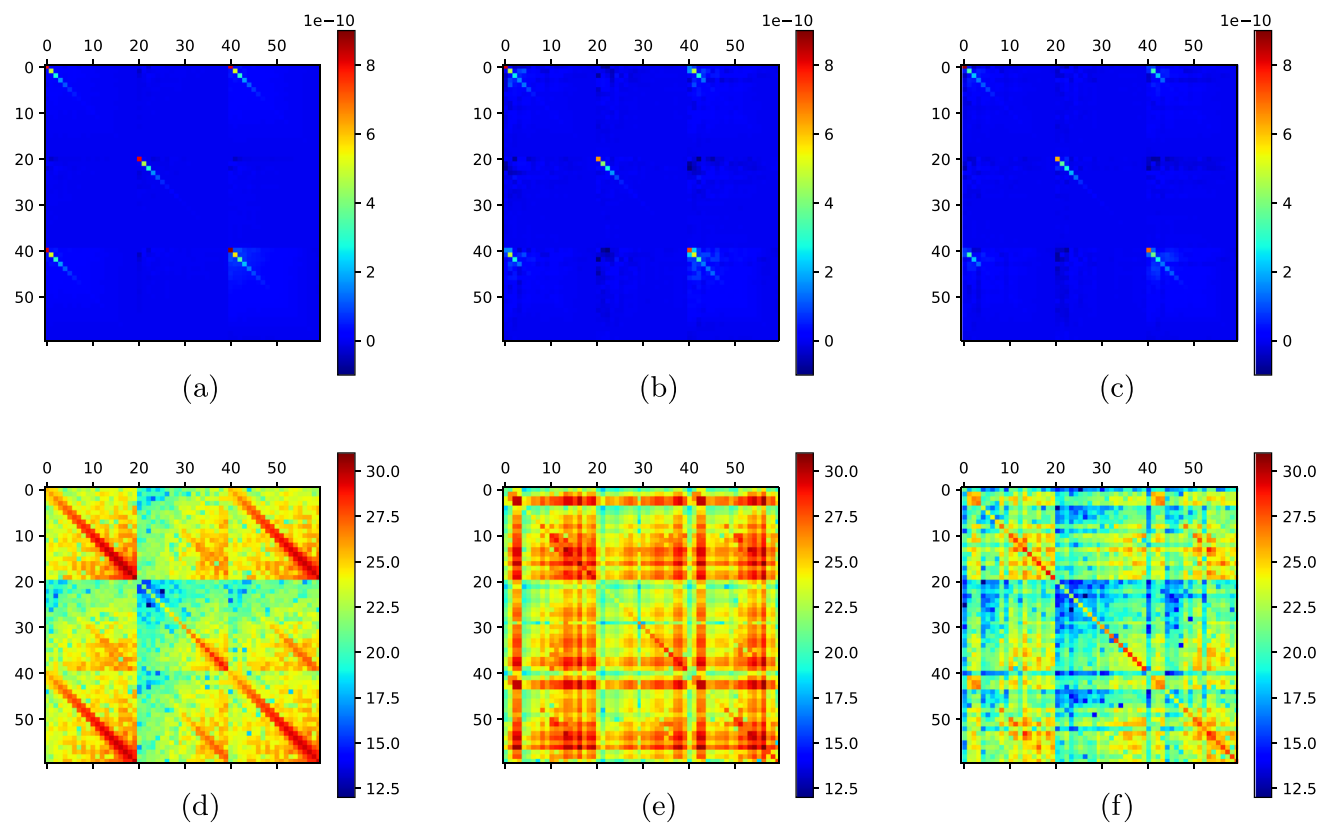


Fig. 8 Covariance estimates of astronomical data of 60-dimensional data vectors. Σ is the true covariance using all 2000 data points. Σ_{MAP} is a maximum a posteriori covariance estimate and $\hat{\Sigma}$ is a posterior expectation estimate of Σ using 40 data points. $\hat{\Sigma}^{-1}$ is a posterior expectation estimate of Σ^{-1} using 40 data points.

The posterior expectation estimate uses priors of $\sigma_1 = \sigma_2 = 2$ after normalisation and 500,000 samples from RT-CHMC. a): Σ . b): Σ_{MAP} . c): $\hat{\Sigma}$. d): $\ln |\Sigma^{-1}|$. e): $\ln |\Sigma_{MAP}^{-1}|$. f): $\ln |\hat{\Sigma}^{-1}|$

Table 1 A comparison using three metrics between the MAP estimate and the posterior expectation estimate using 500,000 samples from RT-CHMC

Metric	MAP	Posterior expectation
$\ \Sigma - \Sigma_{EST}\ _F / \ \Sigma\ _F$	0.4862	0.5155
$\ \Sigma^{-1} - \Sigma_{EST}^{-1}\ _F / \ \Sigma^{-1}\ _F$	1.8503	0.8946
$d(\Sigma, \Sigma_{EST})$	15.2515	14.7547

The relative Frobenius norm of the estimate and its inverse and the covariance metric introduced by Förstner and Moonen (2003)

$$d(A, B) = \sqrt{\sum_{i=1}^n \ln^2 \lambda_i(A, B)},$$

where A and B are covariance matrices and $\lambda_i(A, B)$ are the generalized eigenvalues from $\det(\lambda A - B) = 0$. As pointed out in Förstner and Moonen (2003), this covariance metric is affine invariant and invariant to inversion.

Note that in Fig. 8 we have not included the sample estimate with 40 samples because the sample covariance matrix is rank deficient and hence it is not possible to invert the matrix. We notice from Fig. 8 and Table 1 that the MAP and Posterior Expectation perform well when estimating the covariance matrix for only using 40 data points, but lose accuracy under inversion. The posterior expectation using a sampling method seems to retain more structure when inverted and provides a more accurate estimate according to the metric of Förstner and Moonen (2003).

It is clear from Table 1 that using the posterior means do not sacrifice accuracy compared to using the MAP estimators. An additional benefit of the Bayesian approach is that we can compute posterior standard deviations for each component of the covariance estimator, which gives error estimates. This is illustrated in Fig. 9. By comparing these standard deviations with the covariance estimates, we can get a sense of the relative error we are making. This information can be useful when deciding on the number of data samples we need to get a satisfactory level of accuracy in

Fig. 9 A comparison between the log of the absolute error and the log of the posterior standard deviations in each component. Left: $\ln|\Sigma - \hat{\Sigma}|$. Right: $\ln \Sigma_{SD}$

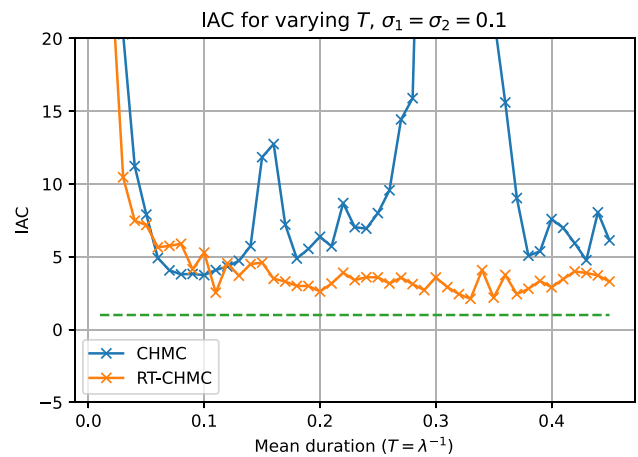



Fig. 10 IAC estimates for different choices of λ for the Bayesian posterior estimation of the upper left quadrant of the covariance matrix for CHMC and RT-CHMC averaged over 20 independent runs

estimating the covariance matrix. Since in practice we do not have access to the true covariance matrix, there is no straightforward way to compute error estimates based on the MAP estimator, and it is challenging to see whether we have reached sufficient accuracy.

Further it is clear from Fig. 10 that the CHMC still suffers from the robustness issue that were present in the lower-dimensional examples and RT-CHMC maintains robustness for this application.

6 Conclusion and future work

In this work we have introduced a Randomized Time Riemannian Manifold Hamiltonian Monte Carlo (RT-RMHMC), which is a robust alternative to Riemannian Manifold Hamiltonian Monte Carlo methods introduced by Girolami and Calderhead (2011) and Brubaker et al. (2012). We establish invariance of the desired measure under a compactness assumption in the continuous (small stepsize limit) setting. We provide an Metropolis adjusted version of RT-RMHMC in the discrete setting and prove invariance and ergodicity

of the adjusted discretized algorithm. We show that RT-RMHMC is a more robust method with respect to parameter choice on a number of numerical examples arising in applications and provide an example to demonstrate that our Riemannian manifold sampling method can be used for high-dimensional covariance estimation. We expect the stability with respect to choice of parameters is especially needed in poorly conditioned problems, where RMHMC would require very short time steps for stability but this may lead to some random walk behaviour and highly inefficient mixing in some principal directions.

In terms of future developments for RT-RMHMC, the next step would be to establish invariance of the measure in the non-compact setting and further to this establish (geometric) ergodicity of RT-RMHMC, which is already established in Bou-Rabee and Sanz-Serna (2017) for the Euclidean setting. Then one could find optimal choices of integration parameters and step-size. Another possibility would be to establish mixing time guarantees for RT-RMHMC by a coupling argument like Bou-Rabee et al. (2020) and Mangoubi and Smith (2018). In Mangoubi and Smith (2018) they establish rapid mixing guarantees for a geodesic walk algorithm on manifolds with positive curvature, which is RMHMC for the uniform distribution. One may be able to use a similar coupling argument to guarantee mixing times for RT-RMHMC for manifolds with positive curvature.

The C code for RT-CHMC and CHMC for each application is available at <https://github.com/PAWhalley/Randomized-Time-Riemannian-Manifold-Hamiltonian-Monte-Carlo>.

Acknowledgements The authors would like to thank the anonymous referees for their valuable feedback and suggestions, which have improved the quality and presentation of the paper. The authors acknowledge the support of the Engineering and Physical Sciences Research Council Grant EP/S023291/1 (MAC-MIGS Centre for Doctoral Training).

Author Contributions P.A.W., D.P. and B.L. developed the theory, conceptualized the algorithm and the simulation tests. P.A.W., D.P. and B.L. wrote the main manuscript text. D.P. and B.L. supervised the project.

Funding The research leading to these results received support of the Engineering and Physical Sciences Research Council Grant EP/S023291/1 (MAC-MIGS Centre for Doctoral Training).

Declarations

Conflicts of interest The authors have no competing interests to declare that are relevant to the content of this article.

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/>.

Appendix A: Generator of RT-RMHMC

To prove that the generator of this stochastic process takes the form of Equation (3) and that the measure (Equation (1)) is invariant under RT-RMHMC we use the framework of Durmus et al. (2021) viewing RT-RMHMC as a piecewise deterministic Markov process (PDMP) defined on $T\mathcal{M}$. For a general state space \mathcal{Z} , a \mathcal{Z} -valued continuous-time PDMP (φ, λ, Q) consists of the following components:

- a differential flow $\varphi : \mathbb{R}_+ \times \mathcal{Z} \mapsto \mathcal{Z}$ which satisfies the semi group property and is measurable. Moreover, is continuously differentiable with respect to time and a C^1 -diffeomorphism of \mathcal{Z} .
- an event rate $\lambda : \mathcal{Z} \rightarrow \mathbb{R}^+$, which is measurable and locally bounded.
- an inhomogeneous Markov transition kernel $Q : \mathbb{R}_+ \times \mathcal{Z} \times \mathcal{B}(\mathcal{Z}) \rightarrow [0, 1]$, such that for all $A \in \mathcal{B}(\mathcal{Z})$, $(t, z) \mapsto Q(t, z, A)$ is measurable and for all $(t, z) \in \mathbb{R}_+ \times \mathcal{Z}$, $Q(t, z, \cdot) \in \mathcal{P}(\mathcal{Z})$,

where $\mathcal{B}(\mathcal{Z})$ denotes the σ -algebra on the space \mathcal{Z} and $\mathcal{P}(\mathcal{Z})$ denotes the space of probability measures on the space \mathcal{Z} . For RT-RMHMC we consider $\mathcal{Z} = T\mathcal{M}$.

Definition 2 For a PDMP $Z = (Z_t)_{t \geq 0}$, we call $\tau_\infty(Z) = \inf\{t \geq 0 \mid Z_t = \infty\}$ the explosion time of the process $(Z_t)_{t \geq 0}$. A process $(Z_t)_{t \geq 0}$ is said to be non-explosive if $\tau_\infty(Z) = +\infty$ almost surely. PDMP characteristics are said to be non-explosive if for all initial distribution the associated PDMP is non-explosive.

Due to the event rate λ of RT-RMHMC being constant and bounded we have that RT-RMHMC is non-explosive. As RT-RMHMC is a non-explosive PDMP we can use the theory of Durmus et al. (2021)[Section 7 and 8] to establish the generator and invariance of the desired measure.

We shall use \mathcal{L} to denote the generator of a PDMP, we use $D(\mathcal{L})$ to denote the core of the generator (or domain), which informally is all the continuous functions which vanish at infinity and for which the generator is well-defined. For more information we refer you to Durmus et al. (2021). Under the assumption that the expected number of events in any unit time interval $[0, t]$ is finite, it is shown in Davis (1993)[Theorem 26.14] that for a non-explosive PDMP with generator

\mathcal{L} with domain $D(\mathcal{L})$ that all $f \in D(\mathcal{L})$ and $x \in T\mathcal{M}$,

$$\mathcal{A}f(x) = D_\varphi f(x) + \lambda(x)(Qf(x) - f(x)),$$

where D_φ is the associated vector field to ϕ .

If we let N_t denote the number of events in the interval $[0, t]$ then we have for RT-RMHMC $\mathbb{E}_x(N_t) = \lambda t < \infty$, \mathbb{E}_x denoting the expected value given the stochastic process starts with initial condition x . Therefore all the assumptions are satisfied of Davis (1993)[Theorem 26.14] and Durmus et al. (2021)[Section 7] and we have that the generator of RT-RMHMC is given by

$$\mathcal{L}f(x) = X_H(f) + \lambda(Qf(x) - f(x)),$$

where X_H is the Hamiltonian vector field and Q is the transition kernel for the Gaussian distribution induced by the metric $G(x)$ on the tangent space of $x \in \mathcal{M}$.

Appendix B: invariant measure

To prove that μ is an invariant measure of RT-RMHMC it is sufficient to show that $\int \mathcal{L}f(x)d\mu = 0$ for all $f \in D(\mathcal{L})$. As it is difficult to consider $D(\mathcal{L})$, one approach is to show that $C_c^1(T\mathcal{M})$ is a core of the generator and that $\int \mathcal{L}f(x)d\mu = 0$ for all $f \in C_c^1(T\mathcal{M})$, where $C_c^k(T\mathcal{M})$ denotes the space of k times differentiable functions $f : T\mathcal{M} \rightarrow \mathbb{R}$ with compact support.

Theorem 8 (Infinitesimal Invariance of RT-RMHMC) *Let \mathcal{M} be a smooth Riemannian manifold with metric g and let $(P_t)_{t \geq 0}$ be the semigroup of RT-RMHMC defined on \mathcal{M} with potential $U \in C^2(\mathcal{M})$ and Hamiltonian $H = U + K \in C^2(T\mathcal{M})$. Let μ be the measure on $(T\mathcal{M}, \mathcal{B}(T\mathcal{M}))$ defined by*

$$\mu(dz) \propto e^{-H(x,v)} d\lambda_{T\mathcal{M}}(z),$$

where $d\lambda_{T\mathcal{M}}$ is the Liouville measure of $T\mathcal{M}$. Then for all $f \in C_c^1(T\mathcal{M})$

$$\int_{T\mathcal{M}} \mathcal{L}f(x, v)\mu(dz) = 0,$$

where \mathcal{L} is the generator of RT-RMHMC.

Proof We have that

$$\begin{aligned} \int_{T\mathcal{M}} \mathcal{L}f(x, v)\mu(dz) &= \int_{T\mathcal{M}} X_H(f)d\mu \\ &+ \lambda \int_{T\mathcal{M}} (Q - I)f d\mu, \end{aligned}$$

where I is the identity operator. We will now consider these two integrals separately. Considering the first integral, due to the fact that μ is a Liouville measure, μ is invariant under the Hamiltonian flow by Liouville’s theorem (see for example Grothaus and Mertin (2022)) and hence the first integral is identically zero. Now considering the second integral we have

$$\begin{aligned} &\int [Qf(x, v) - f(x, v)]\mu(dz) \\ &= \int \int C(x)e^{-\frac{1}{2}\xi^T G(x)^{-1}\xi} (f(x, \xi) \\ &\quad - f(x, v))d\xi e^{-H(x,v)} d\lambda_{T\mathcal{M}}(z) \\ &= \int \int C(x)e^{-\frac{1}{2}\xi^T G(x)^{-1}\xi} e^{-U(x)-\frac{1}{2}v^T G(x)^{-1}v} (f(x, \xi) \\ &\quad - f(x, v))d\xi d\lambda_{T\mathcal{M}}(z) \\ &= \int C(x)e^{-U(x)} \int \int e^{-\frac{1}{2}\xi^T G(x)^{-1}\xi} e^{-\frac{1}{2}v^T G(x)^{-1}v} (f(x, \xi) \\ &\quad - f(x, v))d\xi dv d\sigma_{\mathcal{M}}(x) = 0, \end{aligned}$$

where $C(x)$ is varying depending on x and does not depend on v or ξ . The last equality is identically integral as it can be separated as the subtraction of identical integrals using a change of variables. Therefore we have that

$$\int_{T\mathcal{M}} \mathcal{L}f(x, v)\mu(dz) = 0,$$

and μ is an infinitesimally invariant measure. □

We will next demonstrate that $C_c^1(T\mathcal{M})$ is a core of $D(A)$ by showing that certain conditions established in Durmus et al. (2021) hold under the assumption that \mathcal{M} is compact. To show that $C_c^1(T\mathcal{M})$ is a core of $D(A)$ we use the approach of compactly approximating RT-RMHMC by a more well-behaved PDMP, which has PDMP characteristics $(\varphi, \lambda, Q^\epsilon)$ satisfying the Assumption A3 from Durmus et al. (2021) and has a Feller transition semigroup $(P_t)_{t \geq 0}$. We then use this approximation to show that RT-RMHMC is Feller and $C_c^1(T\mathcal{M})$ is a core of the strong generator of RT-RMHMC, whose transition semigroup $(P_t)_{t \geq 0}$ is seen as a semigroup on $C_0(T\mathcal{M})$. Note that $C_0(T\mathcal{M})$ denotes the space of continuous functions $f : T\mathcal{M} \rightarrow \mathbb{R}$ that vanish at infinity and $C_0(T\mathcal{M})$ is a Banach space when equipped with the $\|\cdot\|_\infty$ norm.

We first approximate our PDMP (φ, λ, Q) (RT-RMHMC) with the PDMP with characteristics $(\varphi, \lambda, Q^\epsilon)$ in the sense that

$$\begin{aligned} &\sup_{z \in T\mathcal{M}, A \in \mathcal{B}(T\mathcal{M})} \{ \lambda^\epsilon(z) \wedge \lambda(z) | Q^\epsilon(z, A) - Q(z, A) | \\ &\quad + | \lambda^\epsilon(z) - \lambda(z) | \} \leq \epsilon, \end{aligned}$$

where Q^ϵ is constructed as a Markov kernel corresponding to a consistently truncated Gaussian distribution on each tangent space as follows.

Define $G : \mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}$ by

$$G(x, a) = \int_{B(0,a)} \psi(x)(dv) - (1 - \epsilon/2\lambda),$$

where $\psi(x)$ denotes the probability density function of the Gaussian distribution on $T_x\mathcal{M}$ defined by $\psi(x)(dv) \propto \exp(-\frac{1}{2}v^T G(x)^{-1}v)dv$, known as the Maxwellian distribution. Then we have that $\partial G/\partial a \neq 0$ due to the fact that $G(x, \cdot)$ is strictly increasing. By the implicit function theorem there exists a unique continuously differentiable function $M : \mathcal{M} \rightarrow \mathbb{R}$ such that $G(x, M(x)) = 0$ for all $x \in \mathcal{M}$. We define the transition kernel as follows:

$$Q^\epsilon(z, dz') = \lambda \delta_x(dx') \psi_\epsilon(x)(dv'),$$

where

$$\psi_\epsilon(x)(dv') = \begin{cases} \frac{1}{1-\epsilon/2\lambda} \psi(x)(dv') & \text{for } |v'|_g \leq M(x) \\ 0 & \text{otherwise,} \end{cases}$$

is the truncated Maxwellian distribution. Then we have that for any $(x, v) \in T\mathcal{M}$ and $A \in \mathcal{B}(T\mathcal{M})$

$$\begin{aligned} & |Q^\epsilon((x, v), A) - Q((x, v), A)| \\ &= \lambda \left(\frac{1}{1-\epsilon/2\lambda} - 1 \right) \int_{A \cap B(0, M(x))} \psi(x)(dv') \\ &\quad - \int_{A \cap B(0, M(x))^c} \psi(x)(dv') \\ &\leq \lambda \left(\frac{1}{1-\epsilon/2\lambda} - 1 \right) \int_{A \cap B(0, M(x))} \psi(x)(dv') \\ &\quad + \lambda \int_{A \cap B(0, M(x))^c} \psi(x)(dv') \\ &\leq \lambda \left(\frac{1}{1-\epsilon/2\lambda} - 1 \right) (1 - \epsilon/2\lambda) + \lambda (1 - (1 - \epsilon/2\lambda)) \\ &= \epsilon. \end{aligned}$$

Lemma 9 (Continuity of Semigroup) *Let (\mathcal{M}, g) be a smooth Riemannian manifold, and let $U \in C^1(\mathcal{M})$ and hence $H \in C^1(T\mathcal{M})$. Let $(P_t)_{t \geq 0}$ be the transition semigroup of $(\varphi, \lambda, Q^\epsilon)$, then*

$$|P_t f - f| \rightarrow 0 \text{ as } t \rightarrow 0, \text{ for all } f \in C_0(T\mathcal{M}).$$

Proof Let $(Z_t)_{t \geq 0}$ denote a sample path of $(\varphi, \lambda, Q^\epsilon)$. We have that

$$\begin{aligned} |P_t f(z) - f(z)| &= |\mathbb{E}_z(f(Z_t) - f(z))| \\ &= |\mathbb{E}_z(f(Z_t)\mathbb{1}(S_1 \leq t) + f(Z_t)\mathbb{1}(S_1 > t)) - f(z)| \end{aligned}$$

$$\begin{aligned} &\leq \|f\|_\infty \mathbb{P}(S_1 \leq t) + |f(\varphi_t(z))e^{-\lambda t} - f(z)| \\ &= \|f\|_\infty (1 - e^{-\lambda t}) + |f(\varphi_t(z)) - f(z)| \rightarrow 0, \end{aligned}$$

where S_1 is the time of the first event and $\varphi_t(z)$ is the solution of the Hamiltonian flow. If H is continuously differentiable everywhere then $\varphi_t(z)$ is well defined for all $t > 0$, and $\varphi_t(z) \rightarrow z$ as $t \rightarrow 0$ (see for example Chicone (2006)[Theorem 1.186]). \square

Lemma 10 *Let (\mathcal{M}, g) be a compact Riemannian manifold, and let $U \in C^1(\mathcal{M})$. Let $(\varphi, \lambda, Q^\epsilon)$ be the PDMP approximation of RT-RMHMC defined above. The set of all possible sample paths of $(\varphi, \lambda, Q^\epsilon)$ with initial condition (X_0, V_0) is contained in a compact set.*

Proof Let $M(x)$ denote the continuous function in the definition of Q_ϵ which controls the truncation of the Gaussian distribution. $M(x)$ is a continuous function on a compact set and hence bounded by M_ϵ . We further choose M_ϵ such that $|V_0|_g \leq M_\epsilon$. Define the set

$$U_\epsilon = \{(x, v) \mid x \in \mathcal{M}, |v|_g \leq M_\epsilon\} \subset T\mathcal{M}.$$

Due to the fact that \mathcal{M} is compact it follows that U_ϵ is a compact subset of $T\mathcal{M}$ by Lemma 13. We have that H restricted to U_ϵ is bounded by M_H as it's continuous on a compact set. We also have that H is constant between event times of the PDMP, by the definition of Hamiltonian flow. Therefore the Hamiltonian defined on the PDMP (X_t, V_t) takes values which are defined by the image of (X_{t_i}, V_{t_i}) , for events $t_i, i = 1, 2, \dots$. At event time $t_i \sim \exp \lambda$, we have that $(X_{t_i}, V_{t_i}) \in U_\epsilon$, where $(X_{t_i}, V_{t_i}) \sim Q(X_{t_i-}, V_{t_i-}, \cdot)$. Therefore we can bound the Hamiltonian by M_H on $\{(X_t, V_t) \mid t \geq 0\}$. Now we have that

$$M_H \geq H(X_t, V_t) = U(X_t) + |V_t|_g \geq m_U + |V_t|_g,$$

for all t . Therefore

$$\begin{aligned} \{(X_t, V_t) \mid t \geq 0\} &\subset \{(x, v) \mid x \in \mathcal{M}, |v|_g \leq M_V \\ &:= M_H - m_U\}, \end{aligned}$$

which is compact by Lemma 13. \square

We have the following assumption from Durmus et al. (2021), which we use to establish Proposition 11.

Definition 3 Durmus et al. (2021)(Definition 16) We say that a homogeneous differential flow φ on $T\mathcal{M}$ and a homogeneous Markov kernel Q on $T\mathcal{M}$ are compactly compatible if for all compact sets $K \subset T\mathcal{M}$ and $T \geq 0$, there exists a compact set $\tilde{K} \subset T\mathcal{M}$ satisfying: for all $n \in \mathbb{N}^*$, $(t_i)_{i \in [1, n]} \in \mathbb{R}_+^n, \sum_{i=1}^n t_i \leq T$, there exists a sequence $(K_i)_{i \in [1, n]}$ of compact sets of $T\mathcal{M}$ such that, setting $K_0 = K$,

- for all $i \in \llbracket 1, n \rrbracket$, K_i only depends on $(t_j)_{j \in \llbracket 1, n \rrbracket}$ and $\cup_{i=0}^n K_i \subset \tilde{K}$;
- for all $i \in \llbracket 0, n-1 \rrbracket$, $s_{i+1} \in [0, t_{i+1}]$ and $s_{n+1} \in [0, T - \sum_{j=1}^n t_j]$,

$$\bigcup_{x \in K_i} \text{supp}\{Q(\varphi_{t_{i+1}}(x), \cdot)\} \subset K_{i+1}, \varphi_{s_{i+1}}(K_i) \subset \tilde{K},$$

$$\varphi_{s_{n+1}}(K_n) \subset \tilde{K}.$$

Assumption 3 Durmus et al. (2021)(A3) The homogeneous characteristics (φ, λ, Q) satisfy

- the flow φ and the Markov kernel Q are compactly compatible;
- $\lambda \in C^1(T\mathcal{M})$ and for all $f \in C^1(T\mathcal{M})$, $\lambda Q^\epsilon f \in C^1(T\mathcal{M})$ and there exists a locally bounded function $\Psi : T\mathcal{M} \rightarrow \mathbb{R}_+$ such that for all $x \in K$,

$$\lambda \|\nabla(Q^\epsilon f)(x)\| \leq \|\Psi\|_{\infty, K} \sup\{|f(y)| + \|\nabla f(y)\| : y \in \text{supp}\{Q^\epsilon(x, \cdot)\}\};$$

- $(t, x) \mapsto \varphi_t(x) \in C^1(\mathbb{R}_+ \times T\mathcal{M})$ and for all compact $K \subset T\mathcal{M}$ and $t \geq 0$,

$$\sup\{\|\nabla \varphi_s(x)\| \mid s \in [0, t], x \in K\} < +\infty.$$

Proposition 11 (Feller and Core of Generator) Let $(P_t)_{t \geq 0}$ be the transition semigroup of $(\varphi, \lambda, Q^\epsilon)$ on $T\mathcal{M}$, where (\mathcal{M}, g) is a compact smooth Riemannian manifold and φ is the Hamiltonian flow associated to the Hamiltonian $H \in C^2(T\mathcal{M})$. Then, $(P_t)_{t \geq 0}$ is Feller and $C_c^1(T\mathcal{M})$ is a core for the strong generator of $(P_t)_{t \geq 0}$ seen as a semigroup on $C_0(T\mathcal{M})$.

Proof If we prove that $(\varphi, \lambda, Q^\epsilon)$ satisfies Assumption 3, then from Durmus et al. (2021)[Theorem 17] $(P_t)_{t \geq 0}$ satisfies the Feller property. Once the Feller property is established by Lemma 9 and due to the fact that $T\mathcal{M}$ is a complete metric space we have by Böttcher et al. (2013)[Lemma 1.4] strong continuity of $(P_t)_{t \geq 0}$ and that $(P_t)_{t \geq 0}$ is Feller. Further to this $C_c^1(T\mathcal{M})$ is a core for the strong generator of $(P_t)_{t \geq 0}$ seen as a semigroup on $C_0(T\mathcal{M})$ is a consequence of Durmus et al. (2021)[Theorem 17] and Ethier and Kurtz (1986)[Proposition 3.3, Chapter 1]. We will now establish Assumption 3.

For any compact set $K \subset T\mathcal{M}$, as K is compact, $|v|_g \leq M_K$ for some constant $M_K \geq 0$ and for all v such that $(\cdot, v) \in K$. Then by the same argument to that of Lemma 10, but choosing M_ϵ larger than M_K we have that all PDMPs starting in K are contained in a compact set \tilde{K} . We can define $K_0 = K$ and $K_i = \tilde{K}$ for all $i \geq 1$. Then we have the flow φ and Q_ϵ are compactly compatible and hence Assumption 3i) holds.

We show Assumption 3ii) as follows. Trivially we have $\lambda \in C^1(T\mathcal{M})$. We have taken the metric to be smooth and

hence, as the truncated Gaussian distribution has a smooth transition kernel, we have that $Q^\epsilon f \in C^1(T\mathcal{M})$. Firstly we note that

$$\text{supp}\{Q^\epsilon(x, \cdot)\} = \{(x, v) \mid |v|_g \leq M(x)\},$$

which is compact by Lemma 13. For all continuously differentiable functions $f : T\mathcal{M} \rightarrow \mathbb{R}$, with $(x, y) \in T\mathcal{M}$, we define

$$A(x, y) := \lambda Q^\epsilon f(x, y) = \frac{\lambda}{1 - \epsilon/2\lambda} \int_{B(0, M(x))} f(x, y') \psi(x) (dy').$$

Therefore it is sufficient to show that for all compact sets $K \subset T\mathcal{M}$, and for all $(x, y) \in K$,

$$\|\nabla A(x, y)\| \leq \sup_{(w, z) \in K} \{\Psi(w, z)\} \sup\{|f(x, y')| + \|\nabla f(x, y')\| \mid |y'|_g \leq M(x)\},$$

where $\Psi : T\mathcal{M} \rightarrow \mathbb{R}$ is bounded on compact sets of $T\mathcal{M}$. Define $\|\cdot\|_{\infty, M(x)} \equiv \|\cdot\|_{\infty, B(0, M(x))}$. We have that for all $(x, y) \in T\mathcal{M}$, since all functions considered are C^1 and hence bounded on all compact sets of $T\mathcal{M}$ we have the following computation which uses the dominated convergence theorem, a Leibniz's integral rule and a spherical coordinate system:

$$\begin{aligned} \|\nabla A(x, y)\| &= \frac{\lambda}{1 - \epsilon/2\lambda} \|\nabla_x \int_{B(0, M(x))} f(x, y') \psi(x, y') dy'\| \\ &= C \|\nabla_x \int_0^{M(x)} \int_{S^{n-1}} f(x, r, \sigma) \psi(x, r, \sigma) r^{n-1} d\sigma dr\| \\ &= C \|\int_0^{M(x)} \int_{S^{n-1}} \nabla_x (f(x, r, \sigma) \psi(x, r, \sigma) r^{n-1}) d\sigma dr \\ &\quad + \int_{S^{n-1}} f(x, M(x), \sigma) \psi(x, M(x), \sigma) M(x)^{n-1} d\sigma \\ &\quad \times \cdot \nabla_x M(x)\| \\ &\leq C \|\int_{B(0, M(x))} \nabla_x (f(x, y') \psi(x, y')) dy'\| \\ &\quad + C \|\int_{S^{n-1}} f(x, M(x), \sigma) \psi(x, M(x), \sigma) d\sigma\| \\ &\quad \times \cdot \|M(x)^{n-1} \nabla_x M(x)\|_\infty \\ &\leq C_1 \|\int_{B(0, M(x))} \nabla_x (f(x, y')) \psi(x, y') dy'\| \\ &\quad + C_1 \|\int_{B(0, M(x))} f(x, y') \nabla_x (\psi(x, y')) dy'\| \\ &\quad + C_2 \|f(x, \cdot)\|_{\infty, M(x)} \\ &\leq C_1 \|\nabla_x f(x, \cdot)\|_{\infty, M(x)} + C_2 \|f(x, \cdot)\|_{\infty, M(x)} \end{aligned}$$

$$\begin{aligned}
 &+ C_3 \|f(x, \cdot)\|_{\infty, M(x)} \left\| \int_{B(0, M(x))} \nabla_x \psi(x, y') dy' \right\| \\
 \leq &C_1 \|\nabla_x f(x, \cdot)\|_{\infty, M(x)} + C_2 \|f(x, \cdot)\|_{\infty, M(x)} \\
 &+ C_3 \|f(x, \cdot)\|_{\infty, M(x)} (M(x))^n \\
 \leq &C_1 \|\nabla f(x, \cdot)\|_{\infty, M(x)} + C_2 \|f(x, \cdot)\|_{\infty, M(x)} \\
 \leq &(C_1 + C_2) \| \|\nabla f(x, \cdot)\| + |f(x, \cdot)| \|_{\infty, M(x)},
 \end{aligned}$$

where C, C_1, C_2 and C_3 are general constants carrying line by line and ∇_x denotes the differential operator with respect to position on \mathcal{M} and we have bounded $\nabla_x \psi$ universally on $\{(x, y) \mid x \in \mathcal{M}, |y|_g \leq M(x)\} \subset T\mathcal{M}$. Therefore we have the required result by setting $\Psi = C_1 + C_2$. Finally we have to show Assumption 3iii), where we use the fact that φ is continuously differentiable, when $U \in C^2(\mathcal{M})$ and for any compact set $K \subset T\mathcal{M}$ we have that $|v|_g \leq M_K$ for all $(x, v) \in T\mathcal{M}$. Then by the same argument as that of Lemma 10 we can define a larger constant such that all PDMPs starting in K have bounded velocity and hence are contained in a compact set \tilde{K} . Hence Assumption 3iii) holds by the fact that a continuous function on a compact set is bounded. \square

Theorem 12 (RT-RMHMC Feller and Core) *Let $(P_t)_{t \geq 0}$ be the transition semigroup of (φ, λ, Q) on $T\mathcal{M}$, where (\mathcal{M}, g) is a compact smooth Riemannian manifold and φ is the Hamiltonian flow associated to the Hamiltonian $H \in C^1(T\mathcal{M})$. Then, $(P_t)_{t \geq 0}$ is Feller and $C_c^1(T\mathcal{M})$ is a core for the strong generator of $(P_t)_{t \geq 0}$ seen as a semigroup on $C_0(T\mathcal{M})$.*

Proof By construction of Q^ϵ we have the property that

$$\begin{aligned}
 &\sup_{x \in \mathcal{M}, A \in \mathcal{B}(\mathcal{M})} \{ \lambda^\epsilon(x) \wedge \lambda(x) | Q^\epsilon(x, A) - Q(x, A) \\
 &+ | \lambda^\epsilon(x) - \lambda(x) | \} \leq \epsilon.
 \end{aligned}$$

Using Durmus et al. (2021)[Theorem 11], Proposition 11, Durmus et al. (2021)[Theorem 17] and the same argument as Durmus et al. (2021)[Theorem 21] we have the required result. \square

Corollary 2 (Invariant measure for RT-RMHMC) *Let $(P_t)_{t \geq 0}$ be the transition semigroup of (φ, λ, Q) on $T\mathcal{M}$, where (\mathcal{M}, g) is a compact smooth Riemannian manifold and φ is the Hamiltonian flow associated to the Hamiltonian $H \in C^2(T\mathcal{M})$. Let μ be the measure on $(T\mathcal{M}, \mathcal{B}(T\mathcal{M}))$ given by*

$$\mu(dz) \propto e^{-H(x,v)} d\lambda_{T\mathcal{M}}(z),$$

where $d\lambda_{T\mathcal{M}}$ is the Liouville measure of $T\mathcal{M}$. Then μ is invariant for RT-RMHMC.

Appendix C: Proof of invariance and μ -irreducibility for the Metropolized algorithm

Proof of Proposition 3 Let P_1 be the Markov kernel corresponding to the first step. It is clear that resampling from the Gaussian measure on the tangent space keeps $\pi_{\mathcal{H}}$ invariant as it has marginal $\phi(x)$, and therefore keeps μ invariant.

Let P_2 be the Markov Kernel corresponding the second step (the combination of the sampling the time duration, deterministic step by Ψ and the Metropolis-Hastings accept-reject step. Let L be an arbitrary number of RATTLE steps we will check that μ is reversible with respect to P_2 and hence also invariant.

P_2 is reversible with respect to μ if for every measurable bounded function $f : T\mathcal{M} \times T\mathcal{M} \rightarrow \mathbb{R}$

$$\begin{aligned}
 &\int \int f(z_1, z_2) \mu(dz_1) P_2(z_1, dz_2) \\
 &= \int \int f(z_1, z_2) \mu(dz_2) P_2(z_2, dz_1).
 \end{aligned}$$

For P_2 we have that $P_2(z_1, dz_2)$ is non-zero if and only if $z_2 = \Psi(z_1)$ or $z_2 = z_1$, hence we have that

$$\begin{aligned}
 &\int \int f(z_1, z_2) \mu(dz_1) P_2(z_1, dz_2) \\
 &= \int \int f(z_1, \Psi(z_1)) \min [1, \exp (H(z_1) \\
 &\quad - H(\Psi(z_1)))] \mu(dz_1) \\
 &+ \int \int f(z_1, z_1) (1 - \min [1, \exp (H(z_1) \\
 &\quad - H(\Psi(z_1)))] \mu(dz_1).
 \end{aligned}$$

Now let $z_2 = \Psi(z_1)$, then due to the momentum reversal map N , we have that $z_1 = \Psi(z_2) = \Psi(\Psi(z_1))$, and by the volume preserving property of Ψ (preserving the Liouville measure), we have that

$$\begin{aligned}
 \mu(dz_2) &= \mu(dz_1) \cdot \frac{\exp (-H(z_2))}{\exp (-H(z_1))} = \mu(dz_1) \cdot \exp (H(z_1) \\
 &\quad - H(z_2)),
 \end{aligned}$$

and using this property we have that the first part of the above sum can be written as

$$\begin{aligned}
 &\int \int f(z_1, \Psi(z_1)) \min [1, \exp (H(z_1) - H(\Psi(z_1)))] \mu(dz_1) \\
 &= \int \int f(\Psi(z_2), z_2) \min [1, \exp (H(\Psi(z_2)) - H(z_2))] \\
 &\quad \times \exp (H(z_2) - H(\Psi(z_2))) \mu(dz_2) \\
 &= \int \int f(\Psi(z_2), z_2) \min [1, \exp (H(z_2) \\
 &\quad - H(\Psi(z_2)))] \mu(dz_2).
 \end{aligned}$$

Now considering the second part of the sum, through the change of variables $z_1 = z_2$ and combining this with the above result we have the required equality. We therefore have that μ is reversible with respect to P_2 . Considering f to be an indicator function we have invariance P_2 with respect to μ . Due to the fact that this calculation was independent of time we have that μ is invariant with respect to the Markov kernel of this algorithm. \square

Proof of Theorem 5 Based on Brubaker et al. (2012)[Theorem 3]. Fix $\Delta t > 0$ sufficiently small such that our assumption holds. For a measurable set $A \subset \mathcal{M}$, we can say A is contained in a compact set K , which can be covered by $\{B_{r/2}(x) \mid x \in K\}$. Then we have that for some $x' \in K$, $B_{r/2}(x') \cap A$ has positive measure. We can connect x and x' by a sequence of points $x_0, \dots, x_i, \dots, x_n$ for $0 \leq i \leq n$, defined on the geodesic between $x_0 = x$ and $x_n = x'$ such that $d(x, x_n) \leq r/2$. We can find unique v_0, \dots, v_n such that $(x_{i+1}, v_{i+1}) = \Psi_{\Delta t}^L(x_i, v_i)$ by Theorem 4. We have that

$$K(x_i, x_{i+1}) > 0$$

due to the Theorem 4 and the fact that $\phi(x_i)(v_i) > 0$. Considering the final step we have due to the triangle inequality $|x_{n-1} - \tilde{x}| < r$ for all $\tilde{x} \in B_{r/2}(x') \cap A$. Hence by the same reasoning and Theorem 4 we have that $K(x_{n-1}, \tilde{x}) > 0$ for all $\tilde{x} \in B_{r/2}(x') \cap A$. Using the fact that $K(x_i, x_{i+1}) > 0$ for all $0 \leq i \leq n-2$, and $K(x_{n-1}, \tilde{x}) > 0$ for all $\tilde{x} \in B_{r/2}(x') \cap A$ we have that $K^n(x, \tilde{x}) > 0$ for all $\tilde{x} \in B_{r/2}(x') \cap A$ and

$$\begin{aligned} K^n(x, A) &\geq K^n(x, B_{r/2}(x') \cap A) \\ &= \int_{B_{r/2}(x') \cap A} K^n(x, y) \sigma_{\mathcal{M}}(dy) > 0. \end{aligned}$$

\square

Appendix D: Additional results

Lemma 13 Let (\mathcal{M}, g) be a smooth k -dimensional Riemannian manifold, let $K \subset \mathcal{M}$ be compact and let $R \in C^1(\mathcal{M})$ such that $R(x) > 0$ for all $x \in \mathcal{M}$, then the set

$$\tilde{K}_R := \{(x, v) \mid x \in K, v \in T_x \mathcal{M}, |v|_g \leq R(x)\}$$

is a compact subset of $T\mathcal{M}$.

Lemma 13 can be shown by showing that the embedding of \tilde{K}_R is closed and bounded. Closure can be established by showing that the limit of convergent sequences is contained in \tilde{K}_R , using the derivative of the local parametrisation as defined in Guillemin and Pollack (1974)[Page 50].

Appendix E: Integrated autocorrelation and ESS

If the MCMC method converges quickly, we have that the variance $\sigma^2(\bar{f})$ (the variance of the estimator) is small. From the central limit theorem we know that as $N \rightarrow \infty$,

$$\sqrt{N}(\bar{f} - \langle f \rangle) \sim \mathcal{N}(0, a^2),$$

and hence

$$\lim_{N \rightarrow \infty} \sigma^2(\bar{f}) = \frac{a^2}{N},$$

where the quantity a is known as the asymptotic variance. We have the following result

$$a^2 = \tau_f \sigma^2(f),$$

where $\sigma^2(f)$ is the variance of f under the distribution π and is independent of the MCMC scheme used (see (Casella 2004)[Chapter 12] for an in depth study).

We also have

$$\tau_f = 1 + 2 \sum_{i=1}^{\infty} \text{corr}(f(X^0), f(X^i)),$$

which is known as the integrated autocorrelation (IAC). If all samples are independent, then $\tau_f = 1$. Generally, MCMC schemes generate correlated samples, and the larger the value of τ_f the more correlated the samples are.

The IAC (τ_f) is a measure of how dependent the samples are and the closer this value is to 1, the higher the quality of the MCMC samples produced. Note that we will use X^\bullet to denote the random variables in a Markov chain and X_\bullet to denote the outputs of an MCMC scheme. In the following numerics we approximate the IAC by a Monte Carlo method, that is we create a finite chain $\{f_n\}_{n=1}^N = \{f(X_i)\}_{n=1}^N$ from the MCMC schemes we want to test. We estimate

$$\text{corr}(f(X^0), f(X^i)) \approx \frac{c_f(i)}{c_f(0)},$$

where

$$c_f(i) = \frac{1}{N-i} \sum_{n=1}^{N-i} (f_n - \mu_f)(f_{n+i} - \mu_f),$$

and

$$\mu_f = \frac{1}{N} \sum_{n=1}^N f_n.$$

We have that

$$\tau_f \approx 1 + 2 \sum_{i=1}^M \frac{c_f(i)}{c_f(0)},$$

for some large M such that $M \ll N$. Note that in practice one uses a fast Fourier transform method to calculate $c_f(\cdot)$ as it is much more computationally efficient.

We now define an additional metric of quality of samples known as effective sample size (ESS) which is defined as

$$N_{\text{eff}} = \frac{N}{\tau_f},$$

for a sample size of size N . This metric is used to say that a sample of size N of an MCMC algorithm has the efficiency of N_{eff} independent samples for computing the Monte Carlo average of f .

References

- Andersen, H.C.: Molecular dynamics simulations at constant pressure and/or temperature. *J. Chem. Phys.* **72**(4), 2384–2393 (1980)
- Andersen, H.C.: Rattle: a “velocity” version of the Shake algorithm for molecular dynamics calculations. *J. Comput. Phys.* **52**(1), 24–34 (1983)
- Barber, D.: *Bayesian Reasoning and Machine Learning*. Cambridge University Press, Cambridge (2012)
- Bierkens, J., Fearnhead, P., Roberts, G.: The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Stat.* **47**(3), 1288–1320 (2019)
- Böttcher, B., Schilling, R., Wang, J.: *Lévy Matters III. Lecture Notes in Mathematics*. Springer, Cham (2013)
- Bouchard, F., Breloy, A., Ginolhac, G., Pascal, F.: Riemannian framework for robust covariance matrix estimation in spiked models. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5979–5983 (2020). IEEE
- Bouchard-Côté, A., Vollmer, S.J., Doucet, A.: The bouncy particle sampler: a non-reversible rejection-free Markov chain Monte Carlo method. *J. Am. Stat. Assoc.* **113**(522), 855–867 (2018)
- Bou-Rabee, N., Eberle, A.: Couplings for Andersen dynamics. In: *Annales de l’Institut Henri Poincaré (B) Probabilités et Statistiques*, vol. 58, pp. 916–944 (2022). Institut Henri Poincaré
- Bou-Rabee, N., Sanz-Serna, J.M.: Randomized Hamiltonian Monte Carlo. *Ann. Appl. Probab.* **27**(4), 2159–2194 (2017)
- Bou-Rabee, N., Eberle, A., Zimmer, R.: Coupling and convergence for Hamiltonian Monte Carlo. *Ann. Appl. Probab.* **30**(3), 1209–1250 (2020)
- Brubaker, M., Salzmann, M., Urtasun, R.: A family of MCMC methods on implicitly defined manifolds. In: *Artificial Intelligence and Statistics*, pp. 161–172 (2012). PMLR
- Byrne, S., Girolami, M.: Geodesic Monte Carlo on embedded manifolds. *Scandinavian J. Stat.* **40**(4), 825–845 (2013)
- Cai, T., Ma, Z., Wu, Y.: Optimal estimation and rank detection for sparse spiked covariance matrices. *Probab. Theory Relat. Fields* **161**(3), 781–815 (2015)
- Casella, C.R.G.: *Monte Carlo Statistical Methods*, 2nd edn. Springer, New York (2004)
- Chamberlain, G., Rothschild, M.: *Arbitrage, Factor Structure, and Mean-variance Analysis on Large Asset markets*. National Bureau of Economic Research Cambridge, Cambridge (1982)
- Chicone, C.: *Ordinary Differential Equations with Applications*. Texts in Applied Mathematics. Springer, New York (2006)
- Davis, M.H.A.: *Markov Models & Optimization*. CRC Press, Boca Raton (1993)
- Deligiannidis, G., Paulin, D., Bouchard-Côté, A., Doucet, A.: Randomized Hamiltonian Monte Carlo as scaling limit of the bouncy particle sampler and dimension-free convergence rates. *Ann. Appl. Probab.* **31**(6), 2612–2662 (2021)
- Diaconis, P., Holmes, S., Shahshahani, M.: Sampling from a manifold. *Adv. Modern Stat. Theory Appl Festschrift Honor Morris L. Eaton* **10**, 102–125 (2013)
- Durmus, A., Guillin, A., Monmarché, P.: Piecewise deterministic Markov processes and their invariant measures. *Ann. de l’Institut Henri Poincaré, Probabilités et Statistiques* **57**(3), 1442–1475 (2021)
- Ethier, S.N., Kurtz, T.G.: *Markov Processes. Characterization and Convergence*. Wiley, Hoboken (1986)
- Fan, J., Fan, Y., Lv, J.: High dimensional covariance matrix estimation using a factor model. *J. Econ.* **147**(1), 186–197 (2008)
- Förstner, W., Moonen, B.: A metric for covariance matrices. In: *Geodesy—the Challenge of the 3rd Millennium*, pp. 299–309. Springer, Berlin (2003)
- Girolami, M., Calderhead, B.: Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. Royal Stat. Soc. Series B Stat. Method.* **73**(2), 123–214 (2011)
- Graham, M.M., Thiery, A.H., Beskos, A.: Manifold Markov chain Monte Carlo methods for Bayesian inference in diffusion models. *J. Royal Stat. Soc Series B (Stat. Methodol)* **84**(4), 1229–1256 (2022)
- Greenberg, E.: *Introduction to Bayesian Econometrics*, 2nd edn. Cambridge University Press, Cambridge (2012)
- Grothaus, M., Mertin, M.C.: Hypocoercivity of Langevin-type dynamics on abstract smooth manifolds. *Stoch. Process. Appl.* **146**, 22–59 (2022)
- Guillemin, V., Pollack, A.: *Differential Topology*. Prentice-Hall Inc, Englewood Cliffs, N.J. (1974)
- Hairer, E., Hochbruck, M., Iserles, A., Lubich, C.: Geometric numerical integration. *Oberwolfach Rep.* **3**(1), 805–882 (2006)
- Hamelryck, T., Kent, J.T., Krogh, A.: Sampling realistic protein conformations using local structural bias. *PLoS Comput. Biol.* **2**(9), 131 (2006)
- Hartmann, C.: *Model reduction in classical molecular dynamics*. Freie Universität Berlin (2007)
- Hartmann, C.: An ergodic sampling scheme for constrained Hamiltonian systems with applications to molecular dynamics. *J. Stat. Phys.* **130**(4), 687–711 (2008)
- Hoffman, M.D., Gelman, A.: The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**(1), 1593–1623 (2014)
- Joachimi, B.: Non-linear shrinkage estimation of large-scale structure covariance. *Mon. Notices Royal Astron. Soc. Lett.* **466**(1), 83–87 (2017)
- Jolliffe, I.: Generalizations and adaptations of principal component analysis. In: *Principal Component Analysis*. Springer Series in Statistics, 2nd edn. Springer, New York (2002)
- Kleppe, T.S.: Connecting the dots: numerical randomized Hamiltonian Monte Carlo with state-dependent event rates. *J. Comput. Graph. Stat.* **31**, 1–16 (2022)
- Kritchman, S., Nadler, B.: Determining the number of components in a factor model from limited noisy data. *Chemom. Intell. Lab. Syst.* **94**(1), 19–32 (2008)

- Kunze, K., Schauben, H.: The Bingham distribution of quaternions and its spherical radon transform in texture analysis. *Math. Geol.* **36**(8), 917–943 (2004)
- Lam, C.: Nonparametric eigenvalue-regularized precision or covariance matrix estimator. *Ann. Stat.* **44**(3), 928–953 (2016)
- Lam, C.: High-dimensional covariance matrix estimation. *Wiley Interdiscip Rev: Comput Stat* **12**(2), 1485 (2020)
- Laurent, A., Vilmart, G.: Order conditions for sampling the invariant measure of ergodic stochastic differential equations on manifolds. *Found. Comput. Math.* **22**, 1–47 (2021)
- Lee, Y.T., Vempala, S.S.: Convergence rate of Riemannian Hamiltonian Monte Carlo and faster polytope volume computation. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1115–1121 (2018)
- Lee, T., Leok, M., McClamroch, N.H.: *Global Formulations of Lagrangian and Hamiltonian Dynamics on Manifolds*. Springer, Cham (2017)
- Leimkuhler, B., Matthews, C.: Efficient molecular dynamics using geodesic integration and solvent-solute splitting. *Proc. Royal Soc. A Math. Phys. Eng. Sci.* **472**(2189), 20160138 (2016)
- Leimkuhler, B., Reich, S.: *Simulating Hamiltonian Dynamics*. Cambridge University Press, Cambridge (2004)
- Leimkuhler, B.J., Skeel, R.D.: Symplectic numerical integrators in constrained Hamiltonian systems. *J. Comput. Phys.* **112**(1), 117–125 (1994)
- Lelièvre, T., Rousset, M., Stoltz, G.: *Free Energy Computations: A Mathematical Perspective*. Imperial College Press, London (2010)
- Lelièvre, T., Rousset, M., Stoltz, G.: Langevin dynamics with constraints and computation of free energy differences. *Math. Comput.* **81**(280), 2071–2125 (2012)
- Lelièvre, T., Rousset, M., Stoltz, G.: Hybrid Monte Carlo methods for sampling probability measures on submanifolds. *Numerische Mathematik* **143**(2), 379–421 (2019)
- Li, D.: On the rate of convergence to equilibrium of the Andersen thermostat in molecular dynamics. *J. Stat. Phys.* **129**(2), 265–287 (2007)
- Mangoubi, O., Smith, A.: Rapid mixing of Hamiltonian Monte Carlo on strongly log-concave distributions. *arXiv preprint arXiv:1708.07114* (2017)
- Mangoubi, O., Smith, A.: Rapid mixing of geodesic walks on manifolds with positive curvature. *Ann. Appl. Probab.* **28**(4), 2501–2543 (2018)
- Mardia, K.V., Jupp, P.E., Mardia, K.: *Directional Statistics*. Wiley series in probability and statistics, Wiley, Chichester (2000)
- Marsden, J.E., West, M.: Discrete mechanics and variational integrators. *Acta Numerica* **10**, 357–514 (2001)
- Meyer, G., Bonnabel, S., Sepulchre, R.: Linear Regression under Fixed-Rank Constraints: A Riemannian Approach. In: *Proceedings of the 28th International Conference on Machine Learning*, pp. 545–552 (2011)
- Nash, J.: The imbedding problem for Riemannian manifolds. *Ann. Math.* **63**, 20–63 (1956)
- Ortega, J.M., Rheinboldt, W.C.: *Iterative Solution of Nonlinear Equations in Several Variables*, vol. 30. SIAM, Philadelphia (2000)
- Pakman, A., Gilboa, D., Carlson, D., Paninski, L.: Stochastic bouncy particle sampler. In: *International Conference on Machine Learning*, pp. 2741–2750 (2017). PMLR
- Perez, A., MacCallum, J.L., Dill, K.A.: Accelerating molecular simulations of proteins using Bayesian inference on weak information. *Proc. Natl. Acad. Sci.* **112**(38), 11846–11851 (2015)
- Quiroz, M., Kohn, R., Villani, M., Tran, M.-N.: Speeding up MCMC by efficient data subsampling. *J. Am. Stat. Assoc.* **114**, 831–843 (2018)
- Riou-Durand, L., Vogrinc, J.: Metropolis Adjusted Langevin Trajectories: a robust alternative to Hamiltonian Monte Carlo. *arXiv preprint arXiv:2202.13230* (2022)
- Ross, S.A.: The arbitrage theory of capital asset pricing. *J. Econ. Theory* **13**, 341–360 (1976)
- Salakhutdinov, R., Mnih, A.: Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In: *Proceedings of the 25th International Conference on Machine Learning*, pp. 880–887 (2008)
- Shen, D., Shen, H., Marron, J.: A general framework for consistency of principal component analysis. *J. Mach. Learn. Res.* **17**(1), 5218–5251 (2016)
- Tierney, L.: Markov Chains for exploring posterior distributions. *Ann. Stat.* **22**(4), 1701–1728 (1994)
- Vanetti, P., Bouchard-Côté, A., Deligiannidis, G., Doucet, A.: Piecewise-Deterministic Markov Chain Monte Carlo. *arXiv preprint arXiv:1707.05296* (2018)
- Weinan, E., Li, D.: The Andersen thermostat in molecular dynamics. *Commun. Pure Appl. Math.* **61**(1), 96–136 (2008)
- Wilkinson, D.J.: Bayesian methods in bioinformatics and computational systems biology. *Brief. Bioinf.* **8**(2), 109–116 (2007)
- Xie, D., Ridgway Scott, L., Schlick, T.: Analysis of the SHAKE-SOR algorithm for constrained molecular dynamics simulations. *Methods Appl. Anal.* **7**(3), 577–590 (2000)
- Zappa, E., Holmes-Cerfon, M., Goodman, J.: Monte Carlo on Manifolds: sampling densities and integrating functions. *Commun. Pure Appl. Math.* **71**(12), 2609–2647 (2018)

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