HAAR-WEAVE-METROPOLIS KERNEL

KAMATANI, Kengo
Institute of Statistical Mathematics

SONG, Xiaolin
Graduate School of Engineering Science, Osaka University

<https://doi.org/10.5109/4755997>

出版情報 : Bulletin of informatics and cybernetics. 54 (1), pp.1-31, 2022. 統計科学研究会
バージョン :
権利関係 :

HAAR-WEAVE-METROPOLIS KERNEL

by

Kengo KAMATANI and Xiaolin SONG

*Reprinted from the Bulletin of Informatics and Cybernetics
Research Association of Statistical Sciences, Vol.54, No. 1*

FUKUOKA, JAPAN
2022

HAAR-WEAVE-METROPOLIS KERNEL

By

Kengo KAMATANI* and Xiaolin Song†

Abstract

Recently, many Markov chain Monte Carlo methods have been developed with deterministic reversible transform proposals inspired by the Hamiltonian Monte Carlo method. The deterministic transform is relatively easy to reconcile with the local information (gradient etc.) of the target distribution. However, as the ergodic theory suggests, these deterministic proposal methods seem to be incompatible with robustness and lead to poor convergence, especially in the case of target distributions with heavy tails. On the other hand, the Markov kernel using the Haar measure is relatively robust since it learns global information about the target distribution introducing global parameters. However, it requires a density preserving condition, and many deterministic proposals break this condition. In this paper, we carefully select deterministic transforms that preserve the value of the density function and create a Markov kernel, the Weave-Metropolis kernel, using the deterministic transforms. By combining with the Haar measure, we also introduce the Haar-Weave-Metropolis kernel. In this way, the Markov kernel can employ the local information of the target distribution using the deterministic proposal, and thanks to the Haar measure, it can employ the global information of the target distribution. Finally, we show through numerical experiments that the performance of the proposed method is superior to other methods in terms of effective sample size and mean square jump distance per second.

Key Words and Phrases: Markov chain, Hamiltonian Monte Carlo, Haar measure, Bayesian inference.

1. Introduction

The fundamental object of Bayesian statistics is the posterior distribution, and all statistical inference is performed by integrating a given quantity with respect to the posterior distribution. Therefore, the evaluation of the integral is the main obstacle in Bayesian statistics, and Bayesian computational methods have been developed for this purpose. The Markov chain Monte Carlo method, especially, the Metropolis–Hastings kernel has been the gold standard for Bayesian computation for the last thirty years. However, its efficiency seems to be diminishing due to the arise of modern big data with complex dependent models. Most Bayesian computational methods work efficiently for high-dimensional models in the 1990 sense, but have difficulty in the modern era. Bayesian computation needs a breakthrough to keep pace with change.

* Institute of Statistical Mathematics, 10-3 Midoricho, Tachikawa, Tokyo 190-0014, Japan

† Corresponding Author: Graduate School of Engineering Science, Osaka University 1-3 Machikaneyama-cho, Toyonaka, Osaka 560-0043 Japan. songxl@sigmath.es.osaka-u.ac.jp

Recently, the Metropolis–Hastings kernel based on deterministic reversible transform have been developed such as Murray et al. (2010), Song et al. (2017), Levy et al. (2017), Zhang (2018). Typical, and the earliest example is the Hamilton Monte Carlo method by Duane et al. (1987). Reversible transform-based methods can use local information, that is, the gradients of the target probability density function. In addition, these methods can train the transition kernels, e.g., through the neural network strategy using the generated random sequences. However, due to the lack of global information, reversible transform-based methods can lead to poor performance because the local information is sometimes less informative, e.g., for heavy-tailed target distributions.

On the other hand, it is possible to use global information of the target distribution by introducing a global parameter into the Metropolis–Hastings kernel. The global parameters are estimated from random samples. If there is a sufficient information, we can plug in this information to the estimator. However, this is not a robust strategy because a poor estimate leads directly to poor convergence. It is advisable to set a non-informative prior as it leads to robust choice. In this purpose, it is natural to use the Haar measure for the prior distribution of the parameter. The usefulness of Haar measure has been analysed such as Liu et al. (1999), Liu et al. (2000), Nishihara et al. (2014), Shariff et al. (2015), Kamatani (2017), Kamatani (2018). This strategy sometimes improves the performance of Markov chain Monte Carlo drastically. The improvement is theoretically proved for a specific kernel in terms of ergodic property by Kamatani (2017) and high-dimensional convergence speed by Kamatani (2018). In this paper, we will provide a general theory (Theorem 3.3) that partly explain the benefit of the use of Haar measure.

Therefore, on the one hand, there is an efficient method based on the reversible transform, which takes into account the local information of the target distribution. On the other hand, there is a method based on the Haar measure that uses global information. It is natural to consider a combination of local and global information methods to solve complicated problems. However, this is not an easy task because, at least in our framework, the Haar measure-based methods require that the transform is measure preserving respect to a probability measure. In other words, it should satisfy a density preserving condition (see the paragraph after Definition 4.2). The locally-informed reversible transform usually destroys this structure. In this paper, we carefully select reversible transforms that satisfy the density preserving condition.

As in the Hamiltonian Monte Carlo kernel, the reversible transform is defined for the variable (x, v) , where x is the state variable and v is an auxiliary variable. If we focus on the behaviour of x , then the path of x evolves with the elliptic motion induced by the Circle transform. After a certain time, the bounce forces x to follow another ellipse. Because of this property, the path looks like a weaving behaviour. For this reason, we call this transform the Weaving transform. See Figure 1 for a typical behaviour of the path of the Weave transform.

Thanks to the Bounce transform going in the opposite direction to the gradient, the Weaving transform does not change the corresponding potential energy (negative log-likelihood) as much. This is similar to the Hamiltonian flow, which does not change the value of the Hamiltonian as much. In addition, the Weave transform does not change the distance from the origin. This avoids an unfavourable behaviour for the super-light target distribution. See, e.g. Theorem 4.2 of Roberts et al. (1996). The

Figure 1: Both the left and right panels show the trace of the Weave transform with step number $L = 40$. The left panel is target on a 2-dimensional Student t -distribution and the right panel is toward on 3-dimensional Student t -distribution.

Weave transform itself is not sufficient to move throughout the state space. This local behaviour, reinforced by the Haar motion, traverses the state space quite well.

1.1. Paper organisation

The paper is organised as follows. Section 2 gives a brief overview of the augmented Metropolis–Hastings kernel (Definition 2.3) as a combination of the Metropolis kernel (Definition 2.1) and the flipped reversible transform (Definition 2.2). The Bounce transform (Definition 2.6) is an example of the flipped reversible transform, and Examples 2.4-2.7 are examples of the augmented Metropolis–Hastings kernel. These definitions and examples are more or less well known.

Our main contribution begins in Section 3. In Section 3, we introduce the Haar mixture Metropolis kernel (Definition 3.1) with Example 3.2 as an example to illustrate how to construct a Haar mixture Metropolis kernel. In addition, we explain the advantage of using heavy-tailed reference measure for Metropolis kernels (Theorem 3.3). As far as we are aware, this is the first general result for the uniform ergodicity of the Metropolis kernel. Note here that the definition of the Metropolis kernel is wider than the usual one (see Definition 2.1).

In Section 4, we define augmented Metropolis kernel (Definition 4.1) as Metropolis counterparts of Definition 2.3. Also we introduce augmented Haar–Metropolis kernel (Definition 4.2) as combination of Definition 3.1 and Definition 4.1. Then by introducing the Circle transform (Definition 4.3) as a flipped reversible transform, we propose the Weave-Metropolis kernel (Definition 4.4) as an augmented Haar–Metropolis kernel using a combination of Circle and Bounce transform as flipped reversible transform. Also we define Haar-Weave-Metropolis (Definition 4.5) as the Haar–Metropolis counterpart. Some theoretical properties are summarised in Propositions 4.7 and 4.8 using a result of reversible transform (Proposition 4.6). The proposed algorithms (Weave-Metropolis and Haar-Weave-Metropolis) and related algorithms in this paper are summarised in

Table 1.

Section 5 studies the limit behaviour of the Weave transform (Proposition 5.2 and Theorem 5.3) that illustrates a good insight into how to evolve the Weave transform.

In Section 6, we compare the proposed algorithms with other Markov chain Monte Carlo methods through some numerical experiments.

Table 1: Summary of the algorithms in this paper. HMC stands for the Hamiltonian Monte Carlo and $\mathcal{N}_d$ stands for the d -dimensional normal distribution. Also, Π is the target probability measure. The reference measures in this table are those of the augmented Metropolis(–Hastings) kernel.

	Transform	Reference measure	Augmented Metropolis(–Hastings)	Augmented Haar–Metropolis
Sec. 2.	Hamiltonian	$\Pi \otimes \mathcal{N}_d$	HMC	–
Ex.2.4	Leap-frog	$\text{Leb}^d \otimes \text{Leb}^d$	Leap-frog approximated HMC	–
Ex.2.5	(A version of) Leap-frog	$\text{Leb}^d \otimes \text{Leb}^d$	Infinite dimensional HMC	–
Ex.2.7	Hug	$\text{Leb}^d \otimes \text{Leb}^d$	Hug Metropolis	–
Def.4.4, Def.4.5	Weave	$\mathcal{N}_d \otimes \mathcal{N}_d$	Weave-Metropolis	Haar-Weave- Metropolis

1.2. Notation

Write $\mathcal{N}_d(\mu, \Sigma)$ for the d -dimensional normal distribution with mean μ and variance Σ . Write $\mathcal{N}(\mu, \Sigma)$ when $d = 1$. Similarly, we write $\mathcal{G}(\nu, \alpha)$ for the gamma distribution with shape parameter ν and rate parameter α , and we write $\mathcal{B}e(\alpha_1, \alpha_2)$ for the Beta distribution with shape parameters α_1 and α_2 . Let $\mathcal{T}_d(\nu, \mu, \Sigma)$ be a d -dimensional multivariate Student t -distribution with shift parameter μ , scale parameter Σ and degree of freedom $\nu > 0$.

For a vector $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, $|x| = (\sum_{i=1}^d x_i^2)^{1/2}$ is the Euclidean norm. If A is a $d \times d$ -square matrix, $|A| = \sup_{|x|=1, x \in \mathbb{R}^d} |Ax|$. I_d is the $d \times d$ -identity matrix.

2. Reversible transform based Metropolis–Hastings kernel

A Markov kernel Q on $(E, \mathcal{E})$ is called μ -reversible for a measure μ if

$$\int_A \mu(dx) Q(x, B) = \int_B \mu(dx) Q(x, A)$$

for any $A, B \in \mathcal{E}$. The Metropolis kernel, defined as follows, is Π -reversible where $\Pi(dx) = \exp(-U(x))\mu(dx)$ is a probability measure, where $U : E \rightarrow \mathbb{R}$. The definition here is somewhat broader than usual, which is useful for our purpose. We want to define the Metropolis kernel as a subclass of the Metropolis–Hastings kernel such that the acceptance rate can be written as a fraction of the target densities.

DEFINITION 2.1 METROPOLIS KERNEL. Let μ be a σ -finite measure. For a μ -reversible Markov kernel Q on $(E, \mathcal{E})$, a Markov kernel P defined by

$$P(x, A) = \delta_x(A) \left(1 - \int_{y \in E} Q(x, dy) \alpha(x, y) \right) + \int_A Q(x, dy) \alpha(x, y)$$

is called the Metropolis kernel, where δ_x is the Dirac measure on $x \in E$, and

$$\alpha(x, y) = \min \{1, \exp(-U(y) + U(x))\}. \quad (1)$$

The Markov kernel Q is called the proposal kernel, the function α is called the acceptance probability, and the measure μ is called the reference measure for the Metropolis kernel.

We call any one-to-one measurable mapping Φ from E to itself a transform. Furthermore, we call a transform gradient-based if it uses gradient information of the target density. For $L \in \mathbb{N}$, we write

$$\Phi^L(x) = \overbrace{\Phi \circ \dots \circ \Phi}^L(x).$$

We introduce a Metropolis kernel based on a transform. We call Φ μ -measure preserving if

$$\mu(\{x \in E : \Phi(x) \in A\}) = \mu(A)$$

for $A \in \mathcal{E}$. We also call Φ μ -reversible transform if

$$\mu(\{x \in A : \Phi(x) \in B\}) = \mu(\{x \in B : \Phi(x) \in A\}) \quad (2)$$

for $A, B \in \mathcal{E}$. Note that every μ -reversible transform Φ is μ -measure preserving, and if the μ -measure preserving transform Φ satisfies the condition $\Phi \circ \Phi(x) = x$, then it is μ -reversible. For the construction of a Metropolis kernel based on μ -measure preserving transform, we need a slightly more general version of reversibility, flipped reversibility, which is sometimes called skew reversibility.

DEFINITION 2.2 FLIPPED REVERSIBLE TRANSFORM. Let μ be a probability measure on $(E, \mathcal{E})$. Let $\kappa : E \rightarrow E$ be a μ -measure preserving transform such that $\kappa \circ \kappa(x) = x$ ($x \in E$). We call $\Phi : E \rightarrow E$ (μ, κ) -reversible transform if $\kappa \circ \Phi$ is μ -reversible.

The properties of the flipped reversible transform are explored in Section 4.3. We will design a Metropolis–Hastings kernel, that uses a flipped reversible transform. Let us consider an extended state space $(E^2, \mathcal{E}^{\otimes 2})$ so that the construction of the flipped reversible transform becomes simpler in practice. Let

$$\mu^{\otimes 2}(dx, dv) = \mu(dx)\mu(dv), \quad \mu \otimes \lambda(dx, dv) = \mu(dx)\lambda(dv).$$

DEFINITION 2.3 AUGMENTED METROPOLIS–HASTINGS KERNEL. Let $\kappa(x, v) = (x, -v)$. Let $\Phi : E^2 \rightarrow E^2$ be $(\mu^{\otimes 2}, \kappa)$ -reversible transform. Let $\Pi(dx) = \exp(-U(x))\mu(dx)$ and $\lambda(dv) = \exp(-K(v))\mu(dv)$ be probability measures on $(E, \mathcal{E})$ such that $K(v) = K(-v)$. The augmented Metropolis–Hastings kernel is a Markov kernel on $(E, \mathcal{E})$ defined by

$$\begin{aligned} P(x, A) = & \delta_x(A) \left(1 - \int_{v \in E} \lambda(dv) \alpha((x, v), \Phi(x, v)) \right) \\ & + \int_{v \in E} 1(\{\Phi(x, v) \in A \times E\}) \lambda(dv) \alpha((x, v), \Phi(x, v)) \end{aligned} \quad (3)$$

with the acceptance probability

$$\alpha((x, v), (x^*, v^*)) = \min \{1, \exp(-H(x^*, v^*) + H(x, v))\}$$

where the function $H(x, v)$, the so-called Hamiltonian, is defined by

$$H(x, v) = U(x) + K(v).$$

We call μ the reference measure.

See Section 4.3. for the proof of Π -reversibility of P . The Hamiltonian Monte Carlo kernel belongs to this class. In this case μ is the Lebesgue measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ and λ is a normal distribution. For $h > 0$, the Hamiltonian transform $\Phi(x, v) = (x_h, v_h)$ is defined by the solution $(x_t, v_t)_{t \geq 0}$ of

$$\frac{dx}{dt} = \frac{\partial H}{\partial v}, \quad \frac{dv}{dt} = -\frac{\partial H}{\partial x}.$$

The Hamiltonian flow does not change the value of the Hamiltonian. Thanks to this property together with the (Leb^d, κ) -reversible property, the Hamiltonian transform is $(\Pi \otimes \lambda, \kappa)$ -reversible and the acceptance probability is always 1 (see Proposition 4.6). In practice, however, we need a discretisation, and the leap-frog approximation is commonly used.

EXAMPLE 2.4 LEAP-FROG APPROXIMATED HAMILTONIAN MONTE CARLO KERNEL.

Let $h > 0$. Suppose $U : \mathbb{R}^d \rightarrow \mathbb{R}$ and $K : \mathbb{R}^d \rightarrow \mathbb{R}$ are differentiable. The leap-frog approximated transform $\phi_h(x, v) = (x_h, v_h)$ is defined by

$$\begin{cases} v_{h/2} &= v - \frac{h}{2} \nabla U(x) \\ x_h &= x + h \nabla K(v_{h/2}) \\ v_h &= v_{h/2} - \frac{h}{2} \nabla U(x_h). \end{cases}$$

In practice, the iterated transform $\Phi(x, v) = (\phi_h)^L(x, v)$ is used, where $L \in \mathbb{N}$. The approximated transform no longer preserves the Hamiltonian and is therefore not $(\Pi \otimes \lambda, \kappa)$ -reversible. However, it is still (Leb^d, κ) -reversible. Therefore, we can construct an augmented Metropolis–Hastings kernel where the underlying measure μ is the Lebesgue measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$.

EXAMPLE 2.5 INFINITE DIMENSIONAL HAMILTONIAN MONTE CARLO KERNEL. A version of the leap-frog transform, which replaces the middle step with the Circle transform (Definition 4.3) and replaces $U(x)$ by $U(x) + |x|^2/2$, is also (Leb^d, κ) -reversible. The corresponding Metropolis–Hastings kernel is called the infinite dimensional Hamiltonian Monte Carlo kernel. See Neal (2011), Beskos et al. (2011). Because of the similarity to the kernel we will introduce, we will focus on this kernel later in Section 6.

Now let us look at another closely related kernel, the Hug kernel introduced in Ludkin et al. (2019). The core for the kernel is the Bounce transform. Let $\xi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a vector-valued function. Let $\mathcal{N} = \{x \in \mathbb{R}^d : \xi(x) = 0\}$. For any $x \in \mathcal{N}^c$, let

$$\bar{\xi}(x) = \frac{\xi(x)}{|\xi(x)|}. \quad (4)$$

The function $\xi(x)$ is indeterminate if $x \in \mathcal{N}$.

DEFINITION 2.6. (**Bounce transform**) For any $x, v \in \mathbb{R}^d$, consider an operation

$$R(x) v = v - 2(\bar{\xi}(x)^\top v)\bar{\xi}(x) = (I - 2\bar{\xi}(x)\bar{\xi}(x)^\top)v \quad (5)$$

if $x \notin \mathcal{N}$ and set $R(x)v = -v$ if $x \in \mathcal{N}$. We call $\Phi_{\text{bounce}}(x, v) = (x, R(x)v)$ the Bounce transform. More generally, for a d -dimensional vector M with a symmetric positive definite matrix Σ , we consider an operation

$$R(x \mid M, \Sigma) v = M + \left(I - 2 \frac{\Sigma \xi(x) \xi(x)^\top}{\xi(x)^\top \Sigma \xi(x)} \right) (v - M)$$

and set $\Phi_{\text{bounce}}(x, v \mid M, \Sigma) = (x, R(x \mid M, \Sigma)v)$. We have $\xi(x)^\top (R(x \mid M, \Sigma)v - M) = -\xi(x)^\top (v - M)$. In particular, $\xi(x)^\top R(x)v = -\xi(x)^\top v$.

The Bounce transform reflects the velocity v with respect to the hyper-plane orthogonal to $\xi(x)$. The Bounce transform has been used in the Monte Carlo literature, for example in Neal (2003), Peters et al. (2012), Sherlock and Thiery (2021), Ludkin et al. (2019). Note that the general bounce operation $R(x \mid M, \Sigma)$ and the corresponding transform $\Phi_{\text{bounce}}(x, v \mid M, \Sigma)$ are used only in the simulation section. In the rest of the paper, for simplicity, we consider only $R(x)$ and Φ_{bounce} .

EXAMPLE 2.7 HUG KERNEL. Let $h > 0$. Assume that U and K are differentiable. The Hug transform $\phi_h(x, v) = (x_h, v_h)$ is defined by

$$\begin{cases} x_{h/2} &= x + \frac{h}{2} \nabla K(v) \\ v_h &= R(x_{h/2}) v \\ x_h &= x_{h/2} + \frac{h}{2} \nabla K(v_h). \end{cases}$$

This transform is (Leb^d, κ) -reversible. The augmented Metropolis–Hastings kernel using $\Phi = (\phi_h)^L$ will be referred to as the Hug Metropolis kernel for $L \in \mathbb{N}$.

Under suitable conditions, the transform Φ changes the potential energy $U(x)$ on the order of h^2 when $h \rightarrow 0$. See Ludkin et al. (2019). On the other hand, the transform lacks the ability to efficiently traverse the state space, so it is slow to converge to equilibrium. In Ludkin et al. (2019), a different Markov kernel, called the Hop kernel, was introduced to support traveling in the state space.

3. Haar mixture Metropolis kernel

3.1. Haar measure and Haar mixture kernel

In this section, we recall the Haar mixture Metropolis kernel discussed in Kamatani et al. (2020) that is a trick improving convergence of Metropolis kernels. Let $(G, \times)$ be a locally compact topological group with a Borel σ -algebra. We also assume that the topological group is unimodular, that is, the left Haar measure and the right Haar measure coincide up to a multiplicative constant. Let ν be the Haar measure. The Haar measure satisfies $\nu(H) = \nu(gH)$ for every $g \in G$ and every Borel set H of G . Let E be an abelian topological group with a Borel σ -algebra $\mathcal{E}$. We assume that E is a G -module, i.e., there is an operation $(g, x) \mapsto gx$ such that $ga + gb = g(a + b)$. A typical

example of the pair G and E is $\mathbb{R}$ and $\mathbb{R}^d$ with a scalar multiplication $(a, x) \mapsto ax$ as the left group action.

Let μ be a probability measure and let $Q(x, dy)$ be a μ -reversible Markov kernel on $(E, \mathcal{E})$. Let $\mu_g(A) = \mu(gA)$ and let $Q_g(x, A) = Q(gx, gA)$. Let

$$\mu_*(A) = \int \mu_g(A) \nu(dg).$$

Assume that μ_* is σ -finite. Observe that μ_* is a left invariant measure. Let

$$K(x, dg) = \frac{d\mu_g}{d\mu_*}(x) \nu(dg).$$

By Fubini's lemma, $K(x, G) = 1$ for μ_* -a.s. $x \in E$. But for simplicity, we assume $K(x, G) = 1$ ($x \in E$).

DEFINITION 3.1 HAAR MIXTURE METROPOLIS KERNEL. Let $\Pi(dx) = \exp(-U(x))\mu_*(dx)$. The Haar mixture Metropolis kernel is a Metropolis kernel with proposal kernel

$$Q_*(x, A) = \int_{g \in G} K(x, dg) Q_g(x, A)$$

with acceptance probability (1). The Markov kernel Q_* is μ_* -reversible.

EXAMPLE 3.2 AUTOREGRESSIVE KERNEL. Let $h \in [0, 2\pi)$. Consider a Markov kernel

$$Q(x, \cdot) = \mathcal{N}_d(\cos h \, x, (\sin h)^2 I_d)$$

on $\mathbb{R}^d$. Simple calculation yields that it is $\mathcal{N}_d(0, I_d)$ -reversible.

Let $G = (\mathbb{R}_+, \times)$ and set $(g, x) \mapsto g^{1/2}x$. Then the Haar measure is $\nu(dg) \propto g^{-1}dg$. A simple calculation yields $\mu_g = \mathcal{N}_d(0, g^{-1}I_d)$ and $Q_g(x, \cdot) = \mathcal{N}_d(\cos h \, x, g^{-1}(\sin h)^2 I_d)$. Also, $\mu_*(dx) \propto |x|^{-d}dx$ and $K(x, dg) = \mathcal{G}(d/2, |x|^2/2)$. We have a closed form (up to a constant) of expression of $Q_*(x, \cdot)$ as follows:

$$Q_*(x, dy) \propto \left[1 + \frac{|y - \cos h \, x|^2}{(\sin h)^2 |x|^2} \right]^{-d} dx.$$

See Kamatani et al. (2020) for the detail.

In the above example, defined on $E = \mathbb{R}^d$, the mixed distribution μ_* is heavier than μ in the sense that a μ_* -integrability implies μ -integrability. In particular, μ has any order of moment, but μ_* does not. In general a class of functions h such that $h(x) \leq h(gx) + h(g^{-1}x)$ ($g \in G, x \in E$) we have

$$\int_E h(x) \mu(dx) \leq \int_{g \in G} \int_{x \in E} (h(gx) + h(g^{-1}x)) \mu(dx) \nu(dg) = 2 \int_E h(x) \mu_*(dx).$$

Thus μ_* -integrability implies μ -integrability for the above wide class of functions. Therefore, μ_* is expected to be heavier when E is a Euclidean space.

3.2. Benefit of using heavy-tailed reference measure

In this section, we provide a result that explains in part the importance of the heavy-tailed reference measure (Definition 2.1). More precisely, for uniform ergodicity, we show that μ must be heavier than the target distribution Π so that the density is bounded from above. As explained in Section 3.1., Haar mixture can make the reference measure heavier. This explains in part why Haar mixture improves performance.

Suppose that the measurable space $(E, \mathcal{E})$ is countably generated. By this condition, there exists a small set of positive Π -probability. See Section 5.2 of Meyn et al. (1993). A Markov kernel P is uniformly ergodic if there exists a probability measure Π such that

$$\sup_{x \in E} \sup_{A \in \mathcal{E}} |P^m(x, A) - \Pi(A)| \xrightarrow{m \rightarrow \infty} 0.$$

The following is a generalisation of Theorem 2.1 of Mengersen et al. (1996). In their work, they investigated a sufficient condition for uniform (and exponential) ergodicity for the independent type Metropolis kernel. We present a similar result for general Metropolis kernels. It is worth explaining that for the independent Metropolis kernel, uniform ergodicity is equivalent to exponential ergodicity. In general, however, there is a large gap between the two, and we focus here on uniform ergodicity.

THEOREM 3.3. *Suppose that μ is an atomless σ -finite measure, and $Q(x, dy) = q(x, y)\mu(dy)$ is a μ -reversible Markov kernel with a jointly measurable density $q(x, y)$. Suppose that a probability measure Π has a density $\pi(x)$ with respect to μ such that*

$$\mu\text{-ess sup } \pi(x) = +\infty.$$

Then the Metropolis kernel P is not uniformly ergodic.

PROOF. Suppose now that P is uniformly ergodic. Then Markov kernel is Π -irreducible and also aperiodic by Theorem 5.4.4 of Meyn et al. (1993). Also, there exists $c > 0$ and a probability measure ξ such that

$$P^M(x, dy) \geq c \xi(dy) \quad (x \in E) \tag{6}$$

for some $M \in \mathbb{N}$ by Theorem 16.0.2 of Meyn et al. (1993). By construction, the Metropolis kernel P satisfies

$$P(x, dy) \leq I(x, dy) + Q(x, dy)\alpha(x, y)$$

where $I(x, A) = \delta_x(A)$ and $\alpha(x, y) = \min\{1, \pi(y)/\pi(x)\}$. Observe that $\alpha(x, y)\alpha(y, z) \leq \alpha(x, z)$. Therefore

$$\begin{aligned} P^M(x, dy) &\leq (I + Q\alpha)^M(x, dy) \\ &= I(x, dy) + \sum_{m=1}^M \binom{M}{m} (Q\alpha)^m(x, dy) \\ &\leq I(x, dy) + \sum_{m=1}^M \binom{M}{m} Q^m(x, dy)\alpha(x, y) =: I(x, dy) + Q_*(x, dy)\alpha(x, y). \end{aligned}$$

Observe that $Q_*(x, E) = \sum_{m=1}^M \binom{M}{m} = 2^M - 1$. For each $x \in E$, $P^M(x, dy)$ has a unique atom at the current state x and $P^M(x, dz)$ and $P^M(y, dz)$ does not share an atom if $x \neq y$. Therefore, ξ does not have an atom, and we have

$$Q_*(x, dy)\alpha(x, y) \geq c \xi(dy). \quad (7)$$

Since $Q^m(x, dy)$ is absolutely continuous with respect to μ for each $x \in E$, the probability measure ξ is also absolutely continuous and it has a density function $h(x) = d\xi/d\mu(x)$. Observe that without loss of generality, we can assume that $h(x)$ is bounded above since (6) holds for a truncated version of h , that is, $P^M(x, dy) \geq c_K \xi^K(dy)$ where the probability measure ξ^K is $\xi^K(dy) = \min\{K, h(y)\}\mu(dy) / \int \min\{K, h(z)\}\mu(dz)$ for $c_K = c \int \min\{K, h(z)\}\mu(dz)$ for some $K > 0$ such that $c_K > 0$.

Let $q_*(x, y)$ be the density of $Q_*(x, dy)$ with respect to μ . The left-hand side of (7) is

$$q_*(x, y)\mu(dy)\alpha(x, y) = q_*(x, y)\Pi(dy) \min\left\{\frac{1}{\pi(y)}, \frac{1}{\pi(x)}\right\}.$$

For any $N > 0$, set $N(x) = \{y \in E : q_*(x, y) \leq N\}$. By assumption, there exists $x_n \in E$ ($n = 1, 2, \dots$) such that $\pi(x_n) \rightarrow \infty$. Then

$$\int_{y \in N(x_n)} q_*(x_n, y)\Pi(dy) \min\left\{\frac{1}{\pi(y)}, \frac{1}{\pi(x_n)}\right\} \leq \frac{N}{\pi(x_n)} \rightarrow_{n \rightarrow \infty} 0.$$

Therefore, by the inequality (7),

$$\xi(N(x_n)) \rightarrow_{n \rightarrow \infty} 0 \implies \xi(N(x_n)^c) \rightarrow_{n \rightarrow \infty} 1. \quad (8)$$

On the other hand, by Markov's inequality, we have

$$\mu(N(x_n)^c) \leq \int_E \frac{q_*(x_n, y)}{N} \mu(dy) = \frac{Q_*(x_n, E)}{N} = \frac{2^M - 1}{N}$$

and hence

$$\xi(N(x_n)^c) = \int_{N(x_n)^c} h(y)\mu(dy) \leq \|h\|_\infty \mu(N(x_n)^c) \leq \|h\|_\infty \frac{2^M - 1}{N}$$

where $\|h\|_\infty = \sup_{x \in E} |h(x)|$. This contradicts the convergence (8) when N is greater than $\|h\|_\infty (2^M - 1)$. Therefore, P is not uniformly ergodic.

We do not investigate further ergodic properties of our new kernels that will be explained in Section 4.2. The exponential ergodicity of the new kernels is an interesting topic, but it exceeds the scope of our paper.

4. Weave kernels as combination of transform-based and Haar mixture kernels

4.1. Augmentation and the Haar measure

In this section, we introduce a new kernel that combines a flipped reversible transform and the Haar mixture Metropolis kernel. The new kernel is intended to have a locally and globally informed. To this end, we consider a slightly simpler version of Definition 2.3. We call this simpler version Metropolis kernel, not Metropolis–Hastings kernel, since the acceptance rate depends only on the fraction of target densities.

DEFINITION 4.1 AUGMENTED METROPOLIS KERNEL. Let μ be a probability measure. Let $\Phi : E^2 \rightarrow E^2$ be a $(\mu^{\otimes 2}, \kappa)$ -reversible transform. Let $\Pi(dx) = \exp(-U(x))\mu(dx)$ be a probability measure on $(E, \mathcal{E})$. The augmented Metropolis kernel is a Metropolis kernel on $(E, \mathcal{E})$ with the proposal kernel

$$Q(x, A) = \mu(\{w \in E : \Phi(x, w) \in A \times E\}) \quad (x \in E, A \in \mathcal{E}) \quad (9)$$

with acceptance probability (1). We call μ the reference measure.

Based on this Metropolis kernel, we construct a Haar mixture version. Let $(G, \times), \nu$ and μ_* be as in Section 3.1.

DEFINITION 4.2 AUGMENTED HAAR-METROPOLIS KERNEL. Let μ be a probability measure, and let $\Phi : E^2 \rightarrow E^2$ be a $(\mu_g^{\otimes 2}, \kappa)$ -reversible transform for any $g \in G$. Let $\Pi(dx) = \exp(-U(x))\mu_*(dx)$ be a probability measure on $(E, \mathcal{E})$. The augmented Haar metropolis kernel is a Metropolis kernel on $(E, \mathcal{E})$ with the proposal kernel

$$Q_*(x, A) = \int_G K(x, dg) \mu_g(\{w \in E : \Phi(x, w) \in A \times E\}) \quad (x \in E, A \in \mathcal{E}) \quad (10)$$

with acceptance probability (1). We call μ_* the reference measure.

In Definitions 4.1 and 4.2, the transform Φ should be measure preserving with respect to a probability measure, say, μ . This assumption is crucial. There are many transforms that are Lebesgue measure preserving. These transforms can be μ -measure preserving if the transform does not change the value of the density (see Proposition 4.6). We say that Lebesgue-measure preserving transform satisfies the **density preserving condition** if it does not change the value of the density. However, most gradient-based transforms do not satisfy this condition. An important exception is the Bounce transform, whose kernel will form the basis for our new Metropolis kernels.

4.2. Weave-Metropolis kernel

We are in the position to introduce two new kernels, which use a bounce transform, and a circle transform introduced by the following definition.

DEFINITION 4.3. (**Circle transform**) For $x \in \mathbb{R}^d$, $v \in \mathbb{R}^d$, and $h \in [0, 2\pi)$, the Circle transform denoted by $(x^*, v^*) = \Phi_{\text{circle}}(x, v)$ is defined by

$$\begin{aligned} x^* &\leftarrow x \cos h + v \sin h \\ v^* &\leftarrow -x \sin h + v \cos h. \end{aligned}$$

More generally, for a d -dimensional vector M , a circle transform with parameter M is denoted by $(x^*, v^*) = \Phi_{\text{circle}}(x, v \mid M)$ and defined by

$$\begin{aligned} x^* &\leftarrow M + (x - M) \cos h + (v - M) \sin h \\ v^* &\leftarrow M - (x - M) \sin h + (v - M) \cos h. \end{aligned}$$

The Circle transform is $(\mathcal{N}_d(0, \Sigma)^{\otimes 2}, \kappa)$ -reversible for $\kappa(x, v) = (x, -v)$ and has been used as a proposal kernel of Metropolis-Hastings kernels. See Neal (1999),

Beskos et al. (2008), Murray et al. (2010), Bierkens et al. (2020). The transform naturally fits the Gaussian prior distribution used in many statistical problems.

Based on the Circle transform, we would like to construct an efficient transform that uses local information about the potential energy. It is possible to introduce global information into the Circle transform, as in Law (2014), Rudolf et al. (2018), Cui et al. (2016). However, the Circle transform is blind to the local information of potential energy. Therefore, we need to combine another transform to fulfill our purpose, i.e., to connect with the local information.

The leap-frog approximated transform is commonly used to introduce local information. However, we do not use this strategy because it breaks the $(\mathcal{N}_d(0, \Sigma)^{\otimes 2}, \kappa)$ -reversibility of the Circle transform. More precisely, the transform does not satisfy the density preserving condition for the normal distribution because it changes the value of the normal density in general.

In this work, we use the Bounce transform to employ the local information of the target distribution, and introduce the Weave transform:

$$\phi_h(z) = (\Phi_{\text{circle}} \circ \Phi_{\text{bounce}} \circ \Phi_{\text{circle}})(z). \quad (11)$$

The transform is similar to the Hug transform in Example 2.7. The only difference is that it replaces the shift transform with the Circle transform. Thanks to this difference, the Weave transform always preserves the distance from the origin. This transform satisfies the density preserving condition and also keeps the value of the potential energy approximately constant. The latter assertion is described in Section 5. The **Weave-Metropolis kernel** is the Metropolis kernel using the Weave transform and is considered a discrete-time version of the boomerang sampler proposed in Bierkens et al. (2020).

DEFINITION 4.4 WEAVE-METROPOLIS KERNEL. Let $L \in \mathbb{N}$ and $h > 0$, and let $\kappa(x, v) = (x, -v)$. Let $\Pi(dx) = \exp(-U(x))\mu(dx)$ be a probability measure on $\mathbb{R}^d$ with $\mu = \mathcal{N}_d(0, I_d)$. The Weave-Metropolis kernel is an augmented Metropolis kernel with $\Phi = (\phi_h)^L$.

We also introduce the Haar mixture version of the Weave-Metropolis kernel. We will see that the Haar mixture version, the **Haar-Weave-Metropolis kernel**, performs much better for many target probability distributions.

DEFINITION 4.5 HAAR-WEAVE-METROPOLIS KERNEL. Let $L \in \mathbb{N}$ and $h > 0$, $\kappa(x, v) = (x, -v)$, and let $\mu = \mathcal{N}_d(0, I_d)$. Let $\Pi(dx) = \exp(-U(x))\mu_*(dx)$ be a probability measure on $\mathbb{R}^d$ with $\mu_*(dx) = |x|^{-d}dx$. The Haar-Weave-Metropolis kernel is an augmented Haar-Metropolis kernel on $\mathbb{R}^d$ with $\Phi = (\phi_h)^L$. In this case, $K(x, dg) = \mathcal{G}(d/2, |x|^2/2)$.

The step-by-step description of the algorithms using Weave-Metropolis and Haar-Weave-Metropolis kernels are described in Algorithms 1 and 2.

4.3. Reversibility of Metropolis kernels

In this section, we further investigate the flipped reversible transforms. We would like to remind the reader here that, if the absolute value of the Jacobian determinant of $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is 1 and one-to-one, then ψ is Lebesgue measure preserving. It follows that, the Bounce transform (Definition 2.6) and the Circle transform (Definition 4.3) are Lebesgue measure preserving.

PROPOSITION 4.6. *Let μ and ν be σ -finite measures on $(E, \mathcal{E})$ such that $\nu \ll \mu$. Let $p(x) = d\nu/d\mu(x)$. Let $\kappa : E \rightarrow E$ be a μ -reversible transform such that $\kappa \circ \kappa(x) = x$.*

1. *If the μ -measure preserving transform ψ satisfies the condition $\kappa \circ \psi \circ \kappa \circ \psi(x) = x$ in μ -a.s., then ψ is a (μ, κ) -reversible transform.*
2. *If ψ is (μ, κ) -reversible, then ψ^L is (μ, κ) -reversible for $L \in \mathbb{N}$.*
3. *If ψ_1, ψ_2 are (μ, κ) -reversible transforms, then $\psi_1 \circ \psi_2 \circ \psi_1$ is also (μ, κ) -reversible.*
4. *If ψ is (μ, κ) -reversible and if $p(x) = p(\kappa(x))$ and $p(x) = p(\psi(x))$ in μ -a.s., then ψ is also (ν, κ) -reversible.*

PROOF. 1. By substituting $\kappa \circ \psi$ for ψ , it suffices to show that $\psi \circ \psi(x) = x$ implies the μ -reversibility of ψ when ψ is μ -measure preserving. However, for any $A, B \in \mathcal{E}$ we have.

$$\begin{aligned} \mu(\{x \in A : \psi(x) \in B\}) &= \mu(\{\psi \circ \psi(x) \in A : \psi(x) \in B\}) \\ &= \mu(\{\psi(x) \in A : x \in B\}). \end{aligned}$$

2. For any transform ϕ with a (μ, κ) -reversible transform ψ , the following holds:

$$\begin{aligned} \mu(\{x \in A : \kappa \circ \phi \circ \psi(x) \in B\}) &= \mu(\{x \in A : \kappa \circ \phi \circ \kappa \circ \kappa \circ \psi(x) \in B\}) \\ &= \mu(\{\kappa \circ \psi(x) \in A : \kappa \circ \phi \circ \kappa(x) \in B\}) \\ &= \mu(\{\kappa \circ \psi \circ \kappa(x) \in A : \kappa \circ \phi(x) \in B\}). \end{aligned}$$

Applying this equation sequentially, we obtain

$$\begin{aligned} \mu(\{x \in A : \kappa \circ \psi^L(x) \in B\}) &= \mu(\{\kappa \circ \psi^L \circ \kappa(x) \in A : \kappa(x) \in B\}) \\ &= \mu(\{\kappa \circ \psi^L(x) \in A : x \in B\}). \end{aligned}$$

3. A similar argument as above proves the assertion.
4. By construction we have

$$\begin{aligned} \nu(\{x \in A : \kappa \circ \psi(x) \in B\}) &= \int_E 1(\{x \in A : \kappa \circ \psi(x) \in B\}) p(x) \mu(dx) \\ &= \int_E 1(\{\kappa \circ \psi(x) \in A : x \in B\}) p(\kappa \circ \psi(x)) \mu(dx) \\ &= \int_E 1(\{\kappa \circ \psi(x) \in A : x \in B\}) p(x) \mu(dx) \\ &= \nu(\{\kappa \circ \psi(x) \in A : x \in B\}). \end{aligned}$$

PROPOSITION 4.7. *Augmented Metropolis–Hastings kernel P in Definition 2.3 is Π -reversible.*

PROOF. It suffices to show that $\widehat{P}$ is Π -reversible, where

$$\widehat{P}(x, A) := \int_{v \in E} 1(\{\Phi(x, v) \in A \times E\}) \lambda(dv) \alpha((x, v), \Phi(x, v)).$$

By construction we have $\Pi(dx) \lambda(dv) = e^{-H(z)} \mu^{\otimes 2}(dz)$ and

$$\begin{aligned} \int_A \Pi(dx) \widehat{P}(x, B) &= \int_{(x,v) \in A \times E} \Pi(dx) 1(\{\Phi(x, v) \in B \times E\}) \lambda(dv) \alpha((x, v), \Phi(x, v)) \\ &= \int \mu^{\otimes 2}(dz) \beta(z, \Phi(z)) \end{aligned}$$

where

$$\beta(z, z') = 1(\{z \in A \times E, z' \in B \times E\}) \min\{e^{-H(z)}, e^{-H(z')}\}.$$

Since $K(v) = K(-v)$, we have $\beta(z, \Phi(z)) = \beta(z, \kappa \circ \Phi(z))$. On the other hand, by $(\mu^{\otimes 2}, \kappa)$ -reversibility of Φ ,

$$\begin{aligned} \int_A \Pi(dx) \widehat{P}(x, B) &= \int \mu^{\otimes 2}(dz) \beta(z, \kappa \circ \Phi(z)) \\ &= \int \mu^{\otimes 2}(dz) \beta(\kappa \circ \Phi(z), z) \\ &= \int \mu^{\otimes 2}(dz) \beta(\Phi(z), z) = \int_B \Pi(dx) \widehat{P}(x, A). \end{aligned}$$

Thus the claim follows.

REMARK. Let $\phi = \Phi_{\text{bounce}}$ or $\phi = \Phi_{\text{circle}}$. For $\kappa(x, v) = (x, -v)$, we have $\kappa \circ \phi \circ \kappa \circ \phi(x, v) = (x, v)$. Since ϕ is Lebesgue measure preserving map, it is (Leb^d, κ) -reversible flow. Also, since $|z| = |\kappa(z)| = |\phi(z)|$, it is $(\mu_g^{\otimes 2}, \kappa)$ -reversible for $\mu_g = \mathcal{N}_d(0, g^{-1}I_d)$. Finally, ϕ_h in (11) is $(\mu_g^{\otimes 2}, \kappa)$ -reversible.

PROPOSITION 4.8. *The proposal kernels Q and Q_* defined in Definitions 4.1 and 4.2 are μ and μ_* -reversible respectively. In particular, the augmented Metropolis and augmented Haar-Metropolis kernels are Π -reversible.*

PROOF. The proof of Q is essentially the same as that of Proposition 4.7. For Q_* , by the identity $\mu_*(dx) K(x, dg) = \nu(dg) \mu_g(dx)$, we have

$$\begin{aligned} \int_A \mu_*(dx) Q_*(x, B) &= \int_A \mu_*(dx) \int_G K(x, dg) \mu_g(\{v \in E : \Phi(x, v) \in B \times E\}) \\ &= \int_G \nu(dg) \int_A \mu_g(dx) \mu_g(\{v \in E : \Phi(x, v) \in B \times E\}) \\ &= \int_G \nu(dg) \mu_g^{\otimes 2}(\{(x, v) \in A \times E, \Phi(x, v) \in B \times E\}). \end{aligned}$$

The rest of the proof is also similar, since Φ is $(\mu_g^{\otimes 2}, \kappa)$ -reversible. The assertion for the Metropolis kernel is obvious.

5. Limit of the Weave transform

We will take a closer look at the Weave transform ϕ_h defined in (11). More precisely, we provide a short-time expansion of ϕ_h together with the limit process induced by the transform. These results give a good insight into the behaviour of the Weave-Metropolis kernel, which is different from similar Markov kernels such as the Hamiltonian Monte Carlo kernel.

Let $\xi = \nabla U$. For $\alpha \in \mathbb{R}^d$, let

$$\mathcal{M}_\alpha = \{x \in \mathbb{R}^d : U(x) = \alpha\}$$

be the level set of U . We will assume that $\mathcal{N} \cap \mathcal{M}_\alpha = \emptyset$ where $\mathcal{N} = \{x : \xi(x) = 0\}$. In this case $\mathcal{M}_\alpha$ is a $d - 1$ -dimensional regular submanifold. For $\epsilon > 0$ we introduce an ϵ -perturbation of $\mathcal{M}_\alpha$ by

$$\mathcal{M}_\alpha^\epsilon = \{x \in \mathbb{R}^d : \exists y \text{ s.t. } |x - y| \leq \epsilon, |U(y) - \alpha| \leq \epsilon\}.$$

Let $B_r = \{x \in \mathbb{R}^d : |x| \leq r\}$. Consider a projection of v to the tangential space of $\xi(x)$:

$$P(x)v = (I - \bar{\xi}(x)\bar{\xi}(x)^\top)v$$

for $v \in \mathbb{R}^d$ for $x \notin \mathcal{N}$. Moreover, let $Q(x) = I - P(x) = \bar{\xi}(x)\bar{\xi}(x)^\top$. From this notation we have

$$R = P - Q, \quad I = P + Q \quad \rightsquigarrow \quad R - I = -2Q, \quad R + I = 2P \quad (12)$$

where $R(x)$ is as in Definition 2.6. Also,

$$P(x)\bar{\xi}(x) = 0, \quad Q(x)\bar{\xi}(x) = \bar{\xi}(x). \quad (13)$$

Thus $P(x)Q(x) = Q(x)P(x) = 0$. Observe that if ξ is twice differentiable and if $x \notin \mathcal{N}$, then

$$\partial \bar{\xi}(x)[u, v] = \sum_{i,j} \partial_{x_j} \bar{\xi}_i(x) u_i v_j = \frac{\partial \xi(x)[P(x)u, v]}{|\xi(x)|}. \quad (14)$$

First we show that the transform ϕ_h does not change the value of $U(x)$ so much. As in Theorem 1 of Ludkin et al. (2019), the Hug transform has the same property for the potential with respect to the Lebesgue measure, i.e. the Hug transform does not change the value of $\tilde{U}(x)$ so much, where $\tilde{U}(x)$ is defined by $\Pi(dx) = \exp(-\tilde{U}(x))dx$. This lemma and Theorem 1 of Ludkin et al. (2019) illustrate the similarity of the Weave transform and the Hug transform, both of which use the Bounce transform. Let $p_X(x, v) = x$.

LEMMA 5.1. *Let $\alpha \in \mathbb{R}, r, T > 0$. Suppose that ξ is continuously differentiable. Suppose that $z = (x, v) \in \mathbb{R}^d \times \mathbb{R}^d$, satisfies $|z| \leq r$. Then*

$$|U(p_X(\phi_h(z))) - U(x)| \leq h^2 C(r)$$

for any $h > 0$, where $C(r) = r \sup_{x \in B_r} |\xi(x)| + r^2 \sup_{x \in B_r} |\partial \xi(x)|$.

PROOF. Let us introduce the temporary notation x_0, x_i, v_i, v_i^* ($i = \pm 1$) defined by the following chain relation:

$$z = (x_{-1}, v_{-1}) \xrightarrow{\Phi_{\text{circle}}} (x_0, v_{-1}^*) \xrightarrow{\Phi_{\text{bounce}}} (x_0, v_1^*) \xrightarrow{\Phi_{\text{circle}}} (x_1, v_1) = \phi_h(z).$$

We can rewrite the both ends of the chain (x_i, v_i) ($i = \pm 1$) with the intermediate state x_0 and the velocities v_i^* ($i = \pm 1$):

$$\begin{aligned} \begin{pmatrix} x_{-1} \\ v_{-1} \end{pmatrix} &= \Phi_{\text{circle}}^{-1}(x_0, v_{-1}^*) = \begin{pmatrix} x_0 \cos h - v_{-1}^* \sin h \\ x_0 \sin h + v_{-1}^* \cos h \end{pmatrix}, \\ \begin{pmatrix} x_1 \\ v_1 \end{pmatrix} &= \Phi_{\text{circle}}(x_0, v_1^*) = \begin{pmatrix} x_0 \cos h + v_1^* \sin h \\ -x_0 \sin h + v_1^* \cos h \end{pmatrix}. \end{aligned} \quad (15)$$

Since $v_1^* = R(x_0)v_{-1}^*$, from (12), the following simple relations are obtained for the two intermediate velocities:

$$\begin{aligned} v_1^* - v_{-1}^* &= (R(x_0) - I)v_{-1}^* = -2Q(x_0)v_{-1}^* \\ v_1^* + v_{-1}^* &= (R(x_0) + I)v_{-1}^* = 2P(x_0)v_{-1}^*. \end{aligned} \quad (16)$$

We need to estimate the difference

$$U(x_1) - U(x_{-1}) = \sum_{i=\pm 1} i(U(x_i) - U(x_0)) = \sum_{i=\pm 1} i(U(x_0 \cos h + iv_i^* \sin h) - U(x_0)). \quad (17)$$

We assume that $x_0 \notin \mathcal{N}$ since if $x_0 \in \mathcal{N}$, then $v_{-1}^* = -v_1^*$ and hence $x_{-1} = x_1$, i.e. $U(x_{-1}) = U(x_1)$. Since $\xi = \nabla U$, the right-hand side of the equation (17) is

$$\begin{aligned} \sum_{i=\pm 1} i(U(x_i) - U(x_0)) &= \left\{ - \sum_{i=\pm 1} i \int_0^h \xi(x_0 \cos \theta + iv_i^* \sin \theta)^\top x_0 \sin \theta \, d\theta \right\} \\ &\quad + \left\{ \sum_{i=\pm 1} \int_0^h \xi(x_0 \cos \theta + iv_i^* \sin \theta)^\top v_i^* \cos \theta \, d\theta \right\}. \end{aligned}$$

The absolute value of the first term in the right-hand side is dominated above by $\sup_{x \in B_r} |\xi(x)| r h^2$ since $x_0 \cos \theta + iv_i^* \sin \theta \in B_r$ and $|\sin \theta| \leq |\theta|$. By (13) and (16), we have $\sum_{i=\pm 1} \xi(x)v_i^* = 0$. Thus the second term is

$$\sum_{i=\pm 1} \int_0^h \{ \xi(x_0 \cos \theta + iv_i^* \sin \theta) - \xi(x_0) \}^\top v_i^* \cos \theta \, d\theta$$

whose absolute value is dominated above by $\sup_{x \in B_r} |\partial \xi(x)| r^2 h^2$. Thus the claim follows.

Let $z = (x, v)$. We focus on a projection

$$\mathbf{p}(z) = \begin{pmatrix} x \\ P(x)v \end{pmatrix}.$$

We have the following decomposition of z using $\mathbf{p}(z)$:

$$z = \begin{pmatrix} 0 \\ Q(x)v \end{pmatrix} + \begin{pmatrix} x \\ P(x)v \end{pmatrix} = \begin{pmatrix} 0 \\ Q(x)v \end{pmatrix} + \mathbf{p}(z).$$

The first part in the right-hand side corresponds the sign flip, and it is the fast move of the sequence $\{(\phi_h)^n(z) : n = 0, 1, \dots\}$. The second part is tangential to the sign flip, and it evolves slowly. We are interested into the second part since the first move will

be canceled out in the long run. A short term expansion of the transform $\mathbf{p}(\phi_h(z))$ is composed by

$$a(x, v) = 2 \begin{pmatrix} P(x)v \\ -P(x)x \end{pmatrix}, \quad (18)$$

and

$$b(x, v) = \partial P(x)[\cdot, v, v] + \partial P(x)[\cdot, R(x)v, R(x)v] \in \mathbb{R}^d \quad (19)$$

where ∂P is defined on $\mathcal{N}^c$ as

$$\partial P(x)[u, v, w] = \sum_{i,j,k} u_i \partial_{x_k} P_{ij}(x) v_j w_k = -(\bar{\xi}(x)^\top v) \partial \bar{\xi}(x)[u, w] - (\bar{\xi}(x)^\top u) \partial \bar{\xi}(x)[v, w]. \quad (20)$$

Here, $P_{i,j}(x)$ denotes the (i, j) entry of the matrix $P(x)$. Note here that $|a(z)| \leq 2|z|$ since $|P(x)v| \leq |v|$ and $|P(x)x| \leq |x|$.

PROPOSITION 5.2. *Let $\alpha \in \mathbb{R}$, $r > 0$ and $v \in \mathbb{R}^d$. Suppose that ξ is twice continuously differentiable. Suppose that $\mathcal{M}_\alpha \cap \mathcal{N}^c \neq \emptyset$. Choose $\epsilon > 0$ such that $\mathcal{M}_\alpha^\epsilon \subset \mathcal{N}^c$. Then there is a constant $C(\epsilon, \alpha, r)$ such that for $z = (x, v)$ with $|z| \leq r$ and $x \in \mathcal{M}_\alpha^{\epsilon/2}$,*

$$\left| \mathbf{p}(\phi_h(z)) - \mathbf{p}(z) - h \begin{pmatrix} a \\ b \end{pmatrix}(z) \right| \leq h^2 C(\epsilon, \alpha, r) \quad (21)$$

for $0 < h < \min\{r^{-1}\epsilon, (\epsilon/2C(r))^{1/2}\}$, where $C(r)$ is defined in Lemma 5.1.

PROOF. For a function $f(x)$, we use a generic notation $f(x) = R(\epsilon, \alpha, r)$ if $\sup_{x \in B_r \cap \mathcal{M}_\alpha^\epsilon} |f(x)| < \infty$. By Lemma 5.1, with the same temporary notation, for $0 \leq \theta \leq h$, we have

$$|U(x_1) - U(x_{-1})| \leq h^2 C(r) \leq \epsilon/2 \quad \rightsquigarrow \quad |U(x_{\pm 1}) - \alpha| \leq \epsilon$$

by the triangle inequality, since $x_{-1} \in \mathcal{M}_\alpha^{\epsilon/2}$. On the other hand, since $x_{\pm 1} = x_0 \cos h \pm v_{\pm 1}^* \sin h$ by (15), we have

$$|x_0 \cos \theta \pm v_{\pm 1}^* \sin \theta - x_{\pm 1}| \leq r\theta \leq \epsilon.$$

By the two inequalities, we have

$$x_0 \cos \theta \pm v_{\pm 1}^* \sin \theta \in \mathcal{M}_\alpha^\epsilon \quad (0 \leq \theta \leq h). \quad (22)$$

We now proceed to show (21). We have

$$\begin{aligned} \begin{pmatrix} x_1 - x_{-1} \\ P(x_0)(v_1 - v_{-1}) \end{pmatrix} &\stackrel{(15)}{=} \begin{pmatrix} (v_1^* + v_{-1}^*) \sin h \\ -2P(x_0)x_0 \sin h + P(x_0)(v_1^* - v_{-1}^*) \cos h \end{pmatrix} \\ &\stackrel{(13), (16)}{=} \begin{pmatrix} 2P(x_0) v_{-1}^* \sin h \\ -2P(x_0) x_0 \sin h \end{pmatrix} \\ &\stackrel{(18)}{=} a(x_0, v_{-1}^*) \sin h. \end{aligned} \quad (23)$$

Let

$$\hat{b}(x_0, v_{-1}^*) = h^{-1} \{(P(x_1) - P(x_0))v_1 - (P(x_{-1}) - P(x_0))v_{-1}\}.$$

From (23), we obtain

$$\mathbf{p}(\phi_h(z)) - \mathbf{p}(z) = \begin{pmatrix} x_1 - x_{-1} \\ P(x_1)v_1 - P(x_{-1})v_{-1} \end{pmatrix} = a(x_0, v_{-1}^*) \sin h + \begin{pmatrix} 0 \\ \widehat{b}(x_0, v_{-1}^*) \end{pmatrix} h.$$

Thus it is sufficient to show two inequalities

$$\begin{aligned} |a(x_0, v_{-1}^*) \sin h - a(z)h| &\leq h^2 R(\epsilon, \alpha, r), \\ |\widehat{b}(x_0, v_{-1}^*) - b(z)| &\leq h R(\epsilon, \alpha, r). \end{aligned} \quad (24)$$

(a) Since $|a(z)| \leq 2|z| \leq 2r$, we have

$$|a(x_0, v_{-1}^*) \sin h - a(x_0, v_{-1}^*)h| = |a(x_0, v_{-1}^*)| |\sin h - h| \leq 2r \frac{h^3}{3!}.$$

Also, by (15) and (22), we have

$$|a(x_0, v_{-1}^*) - a(x_{-1}, v_{-1})| \leq rh \sup_{x \in B_r \cap \mathcal{M}_\alpha^\epsilon, v \in B_r} |a'(x, v)| = h R(\epsilon, \alpha, r).$$

Thus the first part of (24) follows since $z = (x_{-1}, v_{-1})$.

(b) Let $l_i(\theta) = x_0 \cos \theta + i v_i^* \sin \theta$. Note that $l_i'' = -l_i$. By (15), we have

$$\begin{aligned} \widehat{b}(x_0, v_{-1}^*) &= h^{-1} \sum_{i=\pm 1} i(P(x_i) - P(x_0))v_i \\ &= h^{-1} \sum_{i=\pm 1} (P(l_i(h)) - P(l_i(0)))l_i'(h) \\ &= h^{-1} \sum_{i=\pm 1} \int_0^h \partial P(l_i(\theta))[\cdot, l_i'(\theta), l_i'(\theta)] d\theta \\ &\quad - h^{-1} \sum_{i=\pm 1} \int_0^h (P(l_i(\theta)) - P(l_i(0)))l_i(\theta) d\theta. \end{aligned}$$

The last equality follows from the integration form of the second equation with respect to h , i.e. if we write the second line as $h^{-1} \sum_{i=\pm 1} \xi_i(h)$, the third line is $h^{-1} \sum_{i=\pm 1} \int_0^h \xi_i'(\theta) d\theta$. The second term in the right-hand side is on the order of h since P is differentiable on $B_r \cap \mathcal{M}_\alpha^\epsilon$. Similarly, since P is twice differentiable on $B_r \cap \mathcal{M}_\alpha^\epsilon$, the first term is

$$\sum_{i=\pm 1} \partial P(x_0)[\cdot, v_i^*, v_i^*] + h R(\epsilon, \alpha, r) = b(x_0, v_{-1}^*) + h R(\epsilon, \alpha, r).$$

On the other hand, by Taylor's expansion, we have (recall that $z = (x_{-1}, v_{-1})$)

$$|b(x_{-1}, v_{-1}) - b(x_0, v_{-1}^*)| \leq h \sup_{x \in B_r \cap \mathcal{M}_\alpha^\epsilon, v \in B_r} |b'(x, v)| = h R(\epsilon, \alpha, r).$$

Thus the second part of (24) follows.

We study limit behaviour of the sequence $\{(\phi_h)^n(z) : n = 0, 1, 2, \dots\}$ as $h \rightarrow 0$. Let $z^h(t) = (x^h(t), v^h(t)) = (\phi_h)^n(z)$ if $nh \leq t < (n+1)h$ for $n = 0, 1, \dots$ be the continuous-time process version of the sequence $\{(\phi_h)^n(z) : n = 0, 1, \dots\}$. We have considered a continuous-time version here because it is easier to compare with the limit process, which is a continuous-time process. We also define $w^h(t) = P(x^h(t))v^h(t)$ that is the tangential component of $v^h(t)$ to $\bar{\xi}(x^h(t))$.

THEOREM 5.3. *Suppose that ξ is thrice differentiable and $\inf_{|x| \rightarrow \infty} U(x) = +\infty$. Let $\alpha \in \mathbb{R}$ such that $\mathcal{M}_\alpha \subset \mathcal{N}^c$ and $\mathcal{M}_\alpha \neq \emptyset$. Choose $x \in \mathcal{M}_\alpha$ $v \in \mathbb{R}^d$. For $T > 0$, consider the process $(x^h(t), w^h(t))$ such that $z^h(0) = (x, v)$. Let $r = |z^h(0)|$. Then there exists $H > 0$ such that for $0 < h < H$, we have*

$$e_T := \sup_{0 \leq t \leq T} |\mathbf{p}(z^h(t)) - \zeta(t)| \leq h C(\epsilon, \alpha, r, T)$$

for some $C(\epsilon, \alpha, r, T) > 0$ where $\zeta(t) = (x(t), w(t))$ is the solution of

$$\zeta'(t) = \mathbf{a}(\zeta(t)) + \begin{pmatrix} 0 \\ \mathbf{b}(\zeta(t)) \end{pmatrix}, \quad \zeta(0) = (x(0), w(0)) = (x, P(x)v) \quad (25)$$

where

$$\begin{aligned} \mathbf{a}(x, w) &= 2(w, -P(x)x)^\top, \\ \mathbf{b}(x, w) &= -2(r^2 - |w|^2 - |x|^2)\partial\bar{\xi}(x)[\cdot, \bar{\xi}(x)] - 2\partial\bar{\xi}(x)[w, w]\bar{\xi}(x). \end{aligned} \quad (26)$$

PROOF. First, we consider the existence and uniqueness of the limit process $\zeta(t)$. Since $\inf_{|x| \rightarrow \infty} U(x) = +\infty$, the level set $\mathcal{M}_\alpha$ is a compact set. Since ξ is thrice differentiable, the functions a and b are differentiable and the derivatives are bounded on $\mathcal{M}_\alpha \times K$ where K is any compact set of $\mathbb{R}^d$. Therefore, by the Picard–Lindelöf theorem, there exists a unique solution $\zeta(t) = (x(t), w(t))$ ($0 \leq t \leq h_1$) for some $h_1 > 0$. By Remark 5., $x(h) \in \mathcal{M}_\alpha$. Therefore, by the same argument replacing $t = 0$ by $t = h_1$, there exists a unique solution $\zeta(t) = (x(t), w(t))$ ($0 \leq t \leq h_2$) for some $h_2 > h_1$. By iterating this argument, we obtain the unique solution $\zeta(t)$ ($0 \leq t \leq T$).

By (20) and with $R(x)\bar{\xi}(x) = -\bar{\xi}(x)$, the vector-valued function b can be written as

$$\begin{aligned} w^\top b(x, v) &\stackrel{(19)}{=} \partial P(x)[w, v, v] + \partial P[w, R(x)v, R(x)v] \\ &\stackrel{(20)}{=} -(\bar{\xi}(x)^\top v)\partial\bar{\xi}(x)[w, v] - (\bar{\xi}(x)^\top w)\partial\bar{\xi}(x)[v, v] \\ &\quad + (\bar{\xi}(x)^\top v)\partial\bar{\xi}(x)[w, R(x)v] - (\bar{\xi}(x)^\top w)\partial\bar{\xi}(x)[v, R(x)v] \\ &\stackrel{(12)}{=} -2(\bar{\xi}(x)^\top v)^2\partial\bar{\xi}(x)[w, \bar{\xi}(x)] - 2(\bar{\xi}(x)^\top w)\partial\bar{\xi}(x)[v, P(x)v], \end{aligned} \quad (27)$$

where we used the fact that $\partial\bar{\xi}[v, \cdot] = \partial\bar{\xi}[R(x)v, \cdot]$ by (13) and (14). Let

$$\begin{aligned} c(x, v) &= a(x, v) + (0, b(x, v))^\top, \\ \mathbf{c}(x, w) &= \mathbf{a}(x, w) + (0, \mathbf{b}(x, w))^\top. \end{aligned}$$

Observe that $a(z) = \mathbf{a}(\mathbf{p}(z))$. Also, if $x^2 + v^2 = r^2$, then $r^2 - |P(x)v|^2 - |x|^2 = |v|^2 - |P(x)v|^2 = (\bar{\xi}(x)^\top v)^2$ and hence

$$\begin{aligned} \mathbf{b}(\mathbf{p}(z)) &\stackrel{(26)}{=} -2(r^2 - |P(x)v|^2 - |x|^2) \partial \bar{\xi}(x)[\cdot, \bar{\xi}(x)] - 2\partial \bar{\xi}(x)[P(x)v, P(x)v] \bar{\xi}(x) \\ &\stackrel{(13)}{=} -2(\bar{\xi}(x)^\top v)^2 \partial \bar{\xi}(x)[\cdot, \bar{\xi}(x)] - 2\partial \bar{\xi}(x)[v, P(x)v] \bar{\xi}(x) \\ &\stackrel{(27)}{=} b(z). \end{aligned}$$

Thus

$$c(z) = \mathbf{c}(\mathbf{p}(z)). \quad (28)$$

Let $N = \lceil T/h \rceil$. The process z^h always remains on a sphere, i.e. there exists $r > 0$ such that

$$|z^h(t)|^2 = r^2 \quad (0 \leq t \leq T).$$

We choose $\epsilon > 0$ so that $\mathcal{M}_\alpha^\epsilon \subset \mathcal{N}^c$ and set

$$H = C(r)^{-1} T^{-1} (\epsilon/2).$$

From Lemma 5.1, for $n = 1, \dots, N$ and for $0 < h < H$,

$$\begin{aligned} |U(x^h(nh)) - U(x^h((n-1)h))| &\leq h^2 C(r) \quad \rightsquigarrow \quad |U(x^h(nh)) - U(x)| \leq \frac{\epsilon}{2} \\ &\rightsquigarrow \quad x^h(nh) \in \mathcal{M}_\alpha^{\epsilon/2}. \end{aligned}$$

The inequality (21) holds for every $z^h(nh)$ by Proposition 5.2. Thus, for

$$\delta^h(t) := \mathbf{p}(z^h(t)) - \mathbf{p}(z^h(0)) - \int_0^t c(z^h(s)) ds,$$

we have

$$\sup_{0 \leq t \leq T} |\delta^h(t)| \leq h T C(\epsilon, \alpha, r) =: hTC \quad (29)$$

for $0 < h < \min\{H, r^{-1}\epsilon, (\epsilon/2C(r))^{1/2}\}$. At the same time, we have

$$\delta(t) := \zeta(t) - \zeta(0) - \int_0^t \mathbf{c}(\zeta(s)) ds = 0 \quad (0 \leq t \leq T).$$

Therefore, according to (28), for $0 < t \leq T$, we have

$$\begin{aligned} e_t &= \sup_{0 \leq s \leq t} |\mathbf{p}(z^h(s)) - \zeta(s)| \\ &= \sup_{0 \leq s \leq t} \left| \delta^h(t) - \delta(t) + \int_0^s \mathbf{c}(\mathbf{p}(z^h(u))) - \mathbf{c}(\zeta(u)) du \right| \\ &\leq htC + \sup_{0 \leq s \leq t} \left| \int_0^s \mathbf{c}(\mathbf{p}(z^h(u))) - \mathbf{c}(\zeta(u)) du \right| \\ &\leq htC + \left(\sup_{x \in B_r \cap \mathcal{M}_\alpha^c, v \in B_r} |\mathbf{c}'(x, v)| \right) \int_0^t e_u du. \end{aligned}$$

The claim follows from Grönwall's lemma.

REMARK. The process $w(t)$ is always tangential to $\bar{\xi}(x(t))$, that is,

$$\bar{\xi}(x(t))^\top w(t) = 0. \quad (30)$$

This equation is true when $t = 0$, and if $t > 0$, we have

$$\begin{aligned} \frac{d}{dt}(\bar{\xi}(x(t))^\top w(t)) &= \left\{ \frac{d}{dt} \bar{\xi}(x(t)) \right\}^\top w(t) + \bar{\xi}(x(t))^\top \left\{ \frac{d}{dt} w(t) \right\} \\ &\stackrel{(25)}{=} 2\partial \bar{\xi}(x(t))[w(t), w(t)] - 2\bar{\xi}(x(t))^\top P(x(t))x(t) + \bar{\xi}(x(t))^\top b(x(t), w(t)) \\ &= 0 \end{aligned}$$

since $P(x(t))\bar{\xi}(x) = 0$ and $\partial \bar{\xi}(x)[\bar{\xi}(x), \cdot] = 0$ by (13) and (25). Thus (30) holds for $t \geq 0$. Due to the property, the process $x(t)$ does not change the value of U since

$$\frac{dU(x(t))}{dt} = \xi(x(t))^\top \frac{dx(t)}{dt} \stackrel{(25)}{=} 2\xi(x(t))^\top w(t) \stackrel{(30)}{=} 0.$$

Note that the process $\zeta(t) = (x(t), w(t))$ does not stay on a sphere, since it is the limit of $(x^h(t), P(x^h(t))v^h(t))$, not of $(x^h(t), v^h(t))$, where the latter remains on the sphere of radius $r = (|x^h(0)|^2 + |v^h(0)|^2)^{1/2}$, but the former is inside the sphere of radius r (see Figure 2). The process $x(t)$ always stay on the manifold $\mathcal{M}_\alpha$, while $x^h(t)$ will move along some ellipses which are bounded and around the trace of $x(t)$ according to Theorem 5.3.

Figure 2: Illustration of the path of the process $x^h(t)$. The surface represents the manifold $\mathcal{M}_\alpha$. The red dashed line denotes $x(t) = p_X(\zeta(t))$ and the blue line is the trace of $x^h(t)$, the circles are the points where bounce events occur.

6. Simulation

6.1. Simulation settings

In this section, we compare the performance of eight proposed kernels as described in Table 2. The kernels are divided into two classes. The first class consists of four gradient-free, information-blind kernels. This class includes the random-walk Metropolis kernel along with the preconditioned Crank–Nicolson kernel (PCN), the mixed preconditioned Crank–Nicolson kernel by Kamatani (2017) (MPCN) and the guided mixed

preconditioned Crank–Nicolson kernel by Kamatani et al. (2020) (GMPCN). The PCN kernel is the Metropolis kernel with normal reference measure, and MPCN is its Haar mixture counterpart, and GMPCN is the non-reversible version of MPCN. The second class consists of four gradient-based kernels. This class includes the infinite dimensional Hamiltonian Monte Carlo kernel (Example 2.5) (∞ -HMC), the Hug-and-Hop kernel by Ludkin et al. (2019) (HH), and the two proposed algorithms, the Weave-Metropolis kernel and the Haar-Weave-Metropolis kernel. We do not show the results of the Hamiltonian Monte Carlo, since the infinite dimensional Hamiltonian Monte Carlo always performs better in our experiments. All kernels except RWM and HH are Gaussian-based in the sense that the proposal distributions are invariant to the normal distribution or its mixture.

Table 2: Markov kernels in Section 6.. The first four algorithms are gradient-free, information blind kernels. The last four algorithms are gradient based, informed kernels.

RWM	Random-walk Metropolis
PCN	Preconditioned Crank–Nicolson
MPCN	Mixed preconditioned Crank–Nicolson
GMPCN	Δ -guided mixed preconditioned Crank–Nicolson
∞ -HMC	Infinite dimensional Hamiltonian Monte Carlo
HH	Hug-and-Hop
WM	Weave-Metropolis
HWM	Haar-Weave-Metropolis

Table 3: Abbreviations in tables in Section 6..

AR	The acceptance ratio
ESS-min	The minimum of the effective sample sizes for the coordinates
ESSL	The effective sample size for the likelihood
MSJD	The mean square jump distance

The efficiency of the Markov kernels is compared using the effective sample size and the mean square jump distance. See p. 126 of Liu (2001) and Roberts et al. (1997) for these statistics. Let $x_1, x_2, \dots$ be a sequence of one-dimensional stationary Markov chain with exponentially ergodic Markov kernel. The effective sample size is defined by

$$N \frac{\lambda^2}{\sigma^2}$$

where λ^2 is the variance of the invariant measure and

$$\sigma^2 = \lambda^2 + 2 \sum_{n=1}^{\infty} \text{cov}(x_1, x_{n+1})$$

is the asymptotic variance of the Markov kernel corresponding to the function of interest $f(x) = x$. In our case, the Markov chain is multidimensional. Therefore, we evaluate the (approximation of the) effective sample size for each coordinate of the Markov chain and take its minimum. We also consider the effective sample size of the log likelihood

with respect to the Lebesgue measure. The coordinate processes and the log likelihood processes are not Markov chains, but we expect these statistics to illustrate the efficiency of the Markov kernels. In this work, all effective sample sizes are computed using the *R* package *mcmcse* developed by Flegal et al. (2021). On the other hand, the mean square jump distance is defined as

$$\frac{1}{N} \sum_{n=0}^{N-1} |x_{n+1} - x_n|^2.$$

All statistics are summarised in Table 3. Note that two acceptance probabilities are given for the Hug-and-Hop kernel. This is because the Hug-and-Hop procedure consists of two types of Metropolis–Hastings procedures, resulting in two acceptance probabilities.

In the experiment, the tuning parameters of the Markov kernels are divided into three categories. The first category is the global parameters. More precisely the shift parameter M and the scale parameter Σ of the target distribution. These parameters are estimated using the adaptive Metropolis kernel, independently of the kernels in Table 2, to allow a fair comparison. The number of iterations is fixed at 10^5 . These parameters are inserted into the normal distributions of the Gaussian-based kernels. The scale parameter is also used in RWM, and both parameters are used in HH. For the Haar motion methods (MPCN, GMPCN, HWM), not all global parameters are fully determined at this stage because one global parameter g is updated by the iteration.

The second category is the local parameters. This class includes the local scale parameter s in $Q(x, \cdot) = \mathcal{N}_d(x, s\Sigma)$ of RWM and the parameter h in the Circle transform. These parameters are tuned using an acceptance rate criterion. For RWM, the parameters are tuned so that the acceptance rate is approximately 25%. See Roberts et al. (1997). For preconditioned Crank–Nicolson kernels, there is no known optimal criterion, but we choose acceptance rates of 30% to 50%, which leads to better performance in our experiments. As suggested by Beskos et al. (2013), the acceptance probability of ∞ -HMC is chosen around 65%. For the Weave kernels, the acceptance rates are chosen around 60%.

For the remaining parameters, rather artificial values are set. Since we consider relatively high-dimensional problems, the gradient-based methods have high computational costs in our experiments. Therefore, we set the number of transforms per iteration to $L = 1$ so that they are comparable to the well-tuned non-gradient-based methods. In HH, there are two types of transforms. Therefore, we take a single step $L = 1$ for each of them. For other tuning parameters of HH, we follow the tuning recommendations described in Ludkin et al. (2019). The kernel had difficulty for our experiments since it is designed for light or super-light target distributions. See Ludkin et al. (2019) for the detail.

We performed all experiments on a desktop computer with 6 cores Intel i7-5930K (3.50 GHz) CPU. All algorithms use the package *RcppArmadillo* version 0.9.850.1.0. See Eddelbuettel et al. (2014). The code for all experiments is available in the online repository at the link <https://github.com/Xiaolin-Song/WM>.

Finally, we describe some theoretical conclusions related to the simulation. The posterior distributions in Sections 6.2. and 6.4. are heavy-tailed. Therefore, the algorithms based on the normal distribution (PCN, HH, WM) have difficulties, as suggested by Theorem 3.3. On the other hand, algorithms based on the mixture of the normal distri-

bution (MPCN, GMPCN, HWM) are expected to have better performance. This assertion is confirmed by comparing the results in Sections 6.2. and 6.4. and those in Section 6.3. On the other hand, as proved in Theorem 5.3, the Weave transform remains approximately on the level set of the likelihood. For this reason, WM has a poor ESSL/S. The property of the Weave transform is efficient when combined with the Haar measure as in HWM.

6.2. Logistic regression

Posterior inference of Bayesian logistic regression with the d -dimensional Cauchy prior distribution $\propto (1 + |x|^2)^{-(d+1)/2} dx$ is considered. The model is applied to two data sets, one is the cancer dataset and the other is the sonar dataset from Dua et al. (2017). The cancer dataset contains 569 observations and 31 variables, and the sonar dataset contains 208 observations and 60 explanatory variables. As suggested in Gelman et al. (2008), all non-binary variables are scaled to have a mean of 0 and a standard deviation of 0.5. We run each algorithm for 1×10^6 iterations and take the first 1×10^5 iterations as burn-in. The results are shown in Tables 4 and 5. The technical difficulty arises from the fact that both data sets contain a large number of highly correlated variables. This makes the likelihood surface complicated.

For both tables, in terms of effective sample size per second, the Haar motion methods and ∞ -HMC perform better than the others. This suggests that the Haar motion effectively explores this heavy-tailed distribution and ∞ -HMC effectively explores the complicated likelihood surface. However, Table 5 shows the advantage of the gradient information more clearly, which is related to the fact that the sonar dataset has a larger number of explanatory variables.

The important point to note here is that when comparing two normal reference type kernels (PCN, WM) and the three normal mixture reference type kernels (MPCN, GMPCN, HWM), the Weave methods perform better in each of the statistics. This indicates the advantage of the Bounce transform.

Another important observation is the comparison between ∞ -HMC and WM. As we will see in other experiments, ∞ -HMC always outperforms WM in terms of effective sample size. This is natural since WM does not change the likelihood as much as described in Remark 5. On the other hand, the mean square jump per second illustrates the limitation of ∞ -HMC. This is probably due to the fact that for a large $|x|$, the kernel behaves like the random-walk Metropolis kernel, since the contribution of the gradient is relatively small when x is far from the origin. On the other hand, the WM effectively uses the directional information even when $|x|$ is large. While WM has difficulty updating the likelihood, HWM shows the best performance for all statistics, even better than ∞ -HMC, thanks to Haar motion.

6.3. The stochastic volatility model

Next, we consider the stochastic volatility model. In this example, we consider the sampling of the latent variables in the following stochastic volatility model. For a

Table 4: The performance for the logistic regression result with Cancer data

Methods	ESSL	ESS-min	MSJD	ESSL/s	ESS- min/s	MSJD/s	Time	AR
RWM	1906.38	5001.71	108.56	33.06	86.74	1.88	57.66	0.20
PCN	1970.02	3554.88	4542.60	32.85	59.28	75.75	59.97	0.31
MPCN	18412.10	33675.57	5505.95	311.18	569.15	93.05	59.17	0.35
GMPCN	26227.38	46693.52	5507.35	433.63	772.00	91.06	60.48	0.35
∞ -HMC	46175.15	86752.44	5062.57	437.12	821.25	47.93	105.63	0.63
HH	22263.42	49639.42	2115.66	79.62	177.53	7.57	279.62	0.83/0.25
WM	8987.86	36601.77	14302.22	89.23	363.36	141.99	100.73	0.62
HWM	48826.31	140611.82	20164.62	471.93	1359.08	194.90	103.46	0.64

Table 5: The performance for the logistic regression result with Sonar data

Methods	ESSL	ESS-min	MSJD	ESSL/s	ESS- min/s	MSJD/s	Time	AR
RWM	790.29	1782.22	0.84	26.88	60.61	0.03	29.40	0.23
PCN	1366.79	2552.64	23.79	42.85	80.03	0.75	31.90	0.28
MPCN	6364.51	7294.80	49.71	209.53	240.16	1.64	30.38	0.26
GMPCN	11819.67	18009.25	50.17	361.32	550.53	1.53	32.71	0.27
∞ -HMC	38026.66	60822.82	160.52	623.98	998.05	2.63	60.94	0.66
HH	17923.04	35698.99	52.98	64.98	129.43	0.19	275.81	0.83/0.34
WM	6076.32	16777.28	318.16	102.96	284.29	5.39	59.01	0.68
HWM	41157.92	89597.69	442.81	690.57	1503.32	7.43	59.60	0.65

Table 6: The performance for the stochastic volatility model

Methods	ESSL	ESS-min	MSJD	ESSL/s	ESS-min/s	MSJD/s	Time	AR
RWM	865.27	964.13	33.46	46.42	51.73	1.80	18.64	0.20
PCN	1128.64	890.82	35.00	60.38	47.66	1.87	18.69	0.37
MPCN	3159.22	1783.71	76.93	164.08	92.64	4.00	19.25	0.34
GMPCN	3747.38	1975.54	100.17	160.97	84.86	4.30	23.28	0.28
∞ -HMC	17861.48	14655.91	2105.46	465.30	381.80	54.84	38.39	0.63
HH	11104.22	18234.27	3415.83	18.21	29.90	5.60	609.77	0.52/0.23
WM	1772.11	8027.00	3206.11	46.07	208.66	83.34	38.47	0.62
HWM	22531.37	18314.68	3871.59	581.06	472.31	99.83	38.78	0.59

positive integer T , for $t = 1, \dots, T$, let

$$\begin{aligned}
x_t &= \phi x_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2) \\
y_t &= \exp\left(\frac{x_t}{2}\right) w_t, \quad w_t \sim \mathcal{N}(0, 1) \\
x_0 &= \frac{\phi}{1 - \phi^2} w_0, \quad w_0 \sim \mathcal{N}(0, \sigma^2)
\end{aligned}$$

where $\epsilon_1, \epsilon_2, \dots$ and w_0, w_1 are independent. The variables $y_1, y_2, \dots$ are observed, while $x_0, x_1, \dots$ are not observed. The parameters of interest are the mean inversion parameter $\phi \sim \mathcal{B}e(2, 5)$ and the standard deviation $\sigma \sim \mathcal{G}(5, 0.2)$. In our simulation, the number of observations is $T = 100$ with $\phi = 0.5$ and $\sigma = 10$.

The result is shown in Table 6. It seems that gradient-based methods are better for this example, which is probably due to the fact that the likelihood in this case is relatively cheap to evaluate. We would like to note that when comparing MPCN and HWM, the Haar motion itself is not sufficient to explore the likelihood surface. On the other hand, the Bounce transform alone is not sufficient when comparing ∞ -HMC and some bounce kernels (HH, WM). The combination is effective, as shown by the performance of HWM.

6.4. Discrete observation of the diffusion process

Finally, we consider statistical inference for the stochastic process driven by the Wiener process. The motivation for this model is to study a scenario with a very complicated high-dimensional target distribution. In this case, the gradient-based methods are likely to be less effective.

Let k, d and N be positive integers, T be a positive number, and let $\alpha \in \mathbb{R}^k$. Suppose $(X_t)_{t \in [0, T]}$ is a solution process of a stochastic differential equation

$$dX_t = a(X_t, \alpha) dt + b(X_t) dW_t; X_0 = x_0$$

where $(W_t)_{t \in [0, T]}$ is the d -dimensional standard Wiener process and $a : \mathbb{R}^d \times \mathbb{R}^k \rightarrow \mathbb{R}^d$ and $b : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ are the drift and diffusion coefficients, respectively. We have only one discrete observation from the path $(X_t)_{t \in [0, T]}$. For simplicity, we assume equally time spaced observations $X_0, X_h, X_{2h}, \dots, X_{Nh}$, where $h = T/N$.

In general, the likelihood is not available for the solution process of the stochastic differential equation. The local Gaussian approximation has been extensively

Table 7: The performance of the stochastic differential equation

Methods	ESSL	ESS-min	MSJD	ESSL/s	ESS-min/s	MSJD/s	Time	AR
RWM	2034.12	2799.71	310.46	6.49	8.94	0.99	313.18	0.31
PCN	12897.57	12785.26	850.10	40.99	40.63	2.70	314.68	0.43
MPCN	16548.68	21805.46	11392.07	52.12	68.67	35.88	317.53	0.30
GMPCN	23184.34	22797.25	11445.11	71.48	70.28	35.28	324.37	0.30
∞ -HMC	15252.22	13420.56	3383.18	16.25	14.30	3.60	938.74	0.69
HH	10061.25	29574.45	12093.45	3.93	11.56	4.73	2557.61	0.52/0.20
WM	14260.80	45389.57	26949.13	15.54	49.47	29.37	917.47	0.64
HWM	44554.04	61870.96	41298.56	48.31	67.09	44.78	922.19	0.63

studied in the past, such as Prakasa Rao et al. (1983), Prakasa Rao et al. (1988), Florens-zmirou (1989), Yoshida (1992). We follow this approach and consider a Bayesian inference using an approximated likelihood.

For our experiment, we set $d = 50$, $N = 10^2$, and $T = 5$. For simplicity, we consider the case of a constant diffusion coefficient, i.e. $b \equiv 1$. The drift coefficient is

$$a(x, \alpha) = -\frac{1}{2}\nabla V(x - \alpha);$$

where $V(x) = 27.5 \log(1 + x^\top \Sigma^{-1}x)$. The symmetric positive definite matrix Σ is generated from the Wishart distribution with 50 degrees of freedom and the identity matrix as the scale matrix. The prior distribution of the parameter α is the multivariate Student t -distribution $\mathcal{T}_{50}(3, 0, 10I_{50})$.

As mentioned earlier, gradient evaluation is expensive. Therefore, the non-gradient based methods show good performance, with the exception of the random-walk Metropolis kernel. In particular, MPCN shows the best performance for the effective sample size per second. The HWM is the best at the mean square jump distance thanks to the Haar motion. The WM kernel and ∞ -HMC have similar performance for effective sample size of log likelihood, but the WM performs better for other statistics per second.

7. Discussion

We introduced a new algorithm, the Weave-Metropolis kernel, which is based on the Weave transform and combines Circle and Bounce transforms. The Haar motion lifts the kernel more efficiently, especially for heavy-tailed target distributions. The Weave transform has similarities with Hamiltonian flow. The former does not change the potential energy as much (see Remark 5.) and the latter does not change the Hamiltonian. However, the former can be combined with the Haar motion, while the latter cannot. The combined Haar-Weave-Metropolis kernel is expected to work well and robust for a wide class of target distributions. On the other hand, the benefit may be lower if the target distribution has light-tailed as in Section 6.3.

We would like to remark about parameter tuning. The effects of the number of iterations per step L for the Weave kernels are different from the Hamiltonian Monte Carlo kernels. For the Hamiltonian Monte Carlo kernel, a large L usually leads to a large mean square jump distance and a small acceptance probability. However, due to the

Circle transform, the effect is non-monotonic for the Weave kernels. To obtain the best performance, L and also the step size h should be carefully tuned. Second, it is advisable to randomise the step size of the Circle transform, since the transform may be periodic. This does not mean that the Markov kernel is reducible, since v is refreshed at each iteration. However, this periodicity can slow down convergence. Similar phenomena can occur in the Hamiltonian Monte Carlo kernel. See p.127 in Brooks et al. (2011).

The parameters of the reference measure must also be chosen properly. The parameters can be estimated from the current samples by treating them as if they had been generated from the reference measure. If it gives poor estimates, then our proposed methods may be inefficient.

Finally, we would like to discuss possible extensions of the results. One approach is to combine various transforms with the Haar measure. For example, in Kamatani et al. (2020) they studied the Beta-Gamma Haar mixture on $\mathbb{R}_+^d$, which can be combined with gradient information. Also, due to the similarity with the Hamiltonian Monte Carlo kernel, we can use the same or similar techniques that can improve the performance of the Hamiltonian Monte Carlo kernel, e.g., parallel computation in Yang et al. (2018), delayed acceptance in Park et al. (2020). This technique can further improve performance.

Acknowledgement

KK and XS were supported by JST CREST Grant number JPMJCR14D7. KK was supported by JSPS KAKENHI Grant number 20H04149 and 21K18589. XS was supported by the Ichikawa International Scholarship Foundation.

A Algorithms

Algorithm 1 Weave-Metropolis

For target distribution $\Pi(dx) = \exp(-U(x))\mu(dx)$ where $\mu = \mathcal{N}_d(M, \Sigma)$.

Input: current state x , pre-condition matrix Σ , step number L , step size h

- 1: Generate $v \sim \mu$
- 2: $(x_0, v_0) = (x, v)$
- 3: **for** $l = 1, 2, \dots, L$ **do**
- 4: $(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}}) = \Phi_{\text{circle}}(x_{l-1}, v_{l-1} \mid M)$ ▷ Circle transform
- 5: $(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}}^*) = \Phi_{\text{bounce}}(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}} \mid M, \Sigma)$ ▷ Bounce transform
- 6: $(x_l, v_l) = \Phi_{\text{circle}}(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}}^* \mid M)$ ▷ Circle transform
- 7: **end for**
- 8: Generate $w \sim \mathcal{U}[0, 1]$
- 9: Compute the acceptance rate $\alpha(x, x_L) = \min\{1, \exp(-U(x_L) + U(x_0))\}$
- 10: **if** $w \leq \alpha(x, x_L)$ **then**
- 11: $x_{\text{new}} = x_L$
- 12: **else**
- 13: $x_{\text{new}} = x$
- 14: **end if**

Output: new state x_{new}

Algorithm 2 Haar-Weave-Metropolis

The target distribution is $\Pi(dx) = \exp(-U(x))\mu_*(dx)$, where $\mu_*(dx) = (\Delta x)^{-d/2}dx$ with $\Delta x = (x - M)^\top \Sigma^{-1}(x - M)$.

Input: current state x , pre-condition matrix Σ , step number L , step size h

- 1: Generate $g \sim \mathcal{G}(d/2, \Delta x/2)$, $v \sim \mathcal{N}_d(M, g^{-1} \Sigma)$ ▷ Haar motion
- 2: $(x_0, v_0) = (x, v)$
- 3: **for** iteration $l = 1, 2, \dots L$ **do**
- 4: $(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}}) = \Phi_{\text{circle}}(x_{l-1}, v_{l-1} \mid M)$ ▷ Circle transform
- 5: $(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}}^*) = \Phi_{\text{bounce}}(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}} \mid M, \Sigma)$ ▷ Bounce transform
- 6: $(x_l, v_l) = \Phi_{\text{circle}}(x_{l-\frac{1}{2}}, v_{l-\frac{1}{2}}^* \mid M)$ ▷ Circle transform
- 7: **end for**
- 8: Generate $w \sim \mathcal{U}[0, 1]$
- 9: Compute the acceptance rate $\alpha(x, x_L) = \min\{1, \exp(-U(x_L) + U(x_0))\}$
- 10: **if** $w \leq \alpha(x, x_L)$ **then**
- $x_{\text{new}} = x_L$
- 11: **else**
- $x_{\text{new}} = x_0$
- 12: **end if**

Output: new state x_{new}

References

- Beskos, A., Pillai, N., Roberts, G., Sanz-Serna, J.-M., and Stuart, A. (2013). Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501–1534.
- Beskos, A., Pinski, F. J., Sanz-Serna, J. M., and Stuart, A. M. (2011). Hybrid Monte Carlo on Hilbert spaces. *Stochastic Process. Appl.*, 121(10):2201–2230.
- Beskos, A., Roberts, G., Stuart, A., and Voss, J. (2008). MCMC methods for diffusion bridges. *Stoch. Dyn.*, 8(3):319–350.
- Bierkens, J., Grazi, S., Kamatani, K., and Roberts, G. (2020). The Boomerang Sampler. In *International conference on machine learning*, pages 908–918. PMLR.
- Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. (2011). *Handbook of Markov chain Monte Carlo*. CRC press.
- Cui, T., Law, K. J., and Marzouk, Y. M. (2016). Dimension-independent likelihood-informed MCMC. *Journal of Computational Physics*, 304:109 – 137.
- Dua, D. and Graff, C. (2017). UCI machine learning repository.
- Duane, S., Kennedy, A., Pendleton, B. J., and Roweth, D. (1987). Hybrid Monte Carlo. *Physics Letters B*, 195(2):216 – 222.
- Eddelbuettel, D. and Sanderson, C. (2014). Rcpparmadillo: Accelerating R with high-performance c++ linear algebra. *Computational Statistics and Data Analysis*, 71:1054–1063.
- Flegal, J. M., Hughes, J., Vats, D., Dai, N., Gupta, K., and Maji, U. (2021). *mcmcse: Monte Carlo Standard Errors for MCMC*. Riverside, CA, and Kanpur, India. R package version 1.5-0.

- Florens-zmirou, D. (1989). Approximate discrete-time schemes for statistics of diffusion processes. *Statistics*, 20(4):547–557.
- Gelman, A., Jakulin, A., Pittau, M. G., and Su, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *Ann. Appl. Stat.*, 2(4):1360–1383.
- Kamatani, K. (2017). Ergodicity of Markov chain Monte Carlo with reversible proposal. *J. Appl. Probab.*, 54(2):638–654.
- Kamatani, K. (2018). Efficient strategy for the Markov chain Monte Carlo in high-dimension with heavy-tailed target probability distribution. *Bernoulli*, 24(4B):3711–3750.
- Kamatani, K. and Song, X. (2020). Non-reversible guided Metropolis kernel. *arXiv preprint arXiv:2005.05584*.
- Law, K. (2014). Proposals which speed up function-space MCMC. *Journal of Computational and Applied Mathematics*, 262:127–138. Selected Papers from NUMDIFF-13.
- Levy, D., Hoffman, M. D., and Sohl-Dickstein, J. (2017). Generalizing Hamiltonian Monte Carlo with Neural Networks. *arXiv preprint arXiv:1711.09268*.
- Liu, J. and Sabatti, C. (2000). Generalised Gibbs sampler and multigrid Monte Carlo for Bayesian computation. *Biometrika*, 87(2):353–369.
- Liu, J. S. (2001). *Monte Carlo strategies in scientific computing*, volume 10. Springer.
- Liu, J. S. and Wu, Y. N. (1999). Parameter expansion for data augmentation. *Journal of the American Statistical Association*, 94(448):1264–1274.
- Ludkin, M. and Sherlock, C. (2019). Hug and Hop: a discrete-time, non-reversible Markov chain Monte-Carlo algorithm. *arXiv preprint arXiv:1907.13570*.
- Mengersen, K. L. and Tweedie, R. L. (1996). Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Statist.*, 24(1):101–121.
- Meyn, S. P. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer.
- Murray, I., Adams, R., and MacKay, D. (2010). Elliptical slice sampling. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 541–548. JMLR Workshop and Conference Proceedings.
- Neal, R. M. (1999). Regression and classification using Gaussian process priors. In *Bayesian statistics, 6 (Alcoceber, 1998)*, pages 475–501. Oxford Univ. Press, New York.
- Neal, R. M. (2003). Slice sampling. *Ann. Statist.*, 31(3):705–767. With discussions and a rejoinder by the author.
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In *Handbook of Markov chain Monte Carlo*, Chapman & Hall/CRC Handb. Mod. Stat. Methods, pages 113–162. CRC Press, Boca Raton, FL.
- Nishihara, R., Murray, I., and Adams, R. P. (2014). Parallel MCMC with generalized elliptical slice sampling. *J. Mach. Learn. Res.*, 15:2087–2112.
- Park, J. and Atchadé, Y. (2020). Markov chain Monte Carlo algorithms with sequential proposals. *Statistics and Computing*, 30(5):1325–1345.

- Peters, E. A. J. F. and de With, G. (2012). Rejection-free Monte Carlo sampling for general potentials. *Physical Review E*, 85(2).
- Prakasa Rao, B. L. S. (1983). Asymptotic theory for non-linear least squares estimator for diffusion processes. *Series Statistics*, 14(2):195–209.
- Prakasa Rao, B. L. S. (1988). Statistical inference from sampled data for stochastic processes. *Statistical inference from stochastic processes (Ithaca, NY, 1987)*, volume 80 of *Contemp. Math.*, pages 249–284. Amer. Math. Soc., Providence, RI.
- Roberts, G. O., Gelman, A., and Gilks, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Appl. Probab.*, 7(1):110–120.
- Roberts, G. O. and Tweedie, R. L. (1996). Exponential convergence of Langevin diffusions and their discrete approximations. *Bernoulli*, 2:341–363.
- Rudolf, D. and Sprungk, B. (2018). On a generalization of the preconditioned Crank-Nicolson Metropolis algorithm. *Found Comput Math* 18, 309–343 (2018).
- Shariff, R., György, A., and Szepesvári, C. (2015). Exploiting Symmetries to Construct Efficient MCMC Algorithms With an Application to SLAM. In *Artificial Intelligence and Statistics*, pages 866–874. PMLR.
- Sherlock, C. and Thiery, A. H. (2021). A discrete bouncy particle sampler. *Biometrika*.
- Song, J., Zhao, S., and Ermon, S. (2017). A-NICE-MC: Adversarial Training for MCMC. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Yang, S., Chen, Y., Bernton, E., and Liu, J. S. (2018). On parallelizable Markov chain Monte Carlo algorithms with waste-recycling. *Statistics and Computing*, 28(5):1073–1081.
- Yoshida, N. (1992). Estimation for diffusion processes from discrete observation. *Journal of Multivariate Analysis*, 41(2):220 – 242.
- Zhang, Y. (2018). The Theory and Algorithm of Ergodic Inference. *arXiv preprint arXiv:1811.07192*.

Received: November 12, 2021

Revised: February 10, 2022

Accept: February 10, 2022