

The Boomerang Sampler

Bierkens, J.; Grazzi, S.; Kamatani, K.; Roberts, G. O.

Publication date

2020

Document Version

Final published version

Published in

Proceedings of the 37th International Conference on Machine Learning

Citation (APA)

Bierkens, J., Grazzi, S., Kamatani, K., & Roberts, G. O. (2020). The Boomerang Sampler. In *Proceedings of the 37th International Conference on Machine Learning* (pp. 908-918)
<http://proceedings.mlr.press/v119/bierkens20a.html>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

The Boomerang Sampler

Joris Bierkens¹ Sebastiano Grazzi¹ Kengo Kamatani² Gareth Roberts³

Abstract

This paper introduces the Boomerang Sampler as a novel class of continuous-time non-reversible Markov chain Monte Carlo algorithms. The methodology begins by representing the target density as a density, e^{-U} , with respect to a prescribed (usually) Gaussian measure and constructs a continuous trajectory consisting of a piecewise elliptical path. The method moves from one elliptical orbit to another according to a rate function which can be written in terms of U . We demonstrate that the method is easy to implement and demonstrate empirically that it can out-perform existing benchmark piecewise deterministic Markov processes such as the bouncy particle sampler and the Zig-Zag. In the Bayesian statistics context, these competitor algorithms are of substantial interest in the large data context due to the fact that they can adopt data subsampling techniques which are exact (ie induce no error in the stationary distribution). We demonstrate theoretically and empirically that we can also construct a control-variate subsampling boomerang sampler which is also exact, and which possesses remarkable scaling properties in the large data limit. We furthermore illustrate a factorised version on the simulation of diffusion bridges.

1. Introduction

Markov chain Monte Carlo remains the gold standard for asymptotically exact (ie bias-free) Bayesian inference for complex problems in Statistics and Machine Learning; see for example (Brooks et al., 2011). Yet a major impediment to its routine implementation for large data sets is the need to evaluate the target density (and possibly other related functionals) at each algorithm iteration.

¹Technische Universiteit Delft, Netherlands ²Graduate School of Engineering Science, Osaka University, Japan ³Department of Statistics, University of Warwick, United Kingdom. Correspondence to: Joris Bierkens <joris.bierkens@tudelft.nl>.

Partly motivated by this, in recent years there has been a surge in the development of innovative piecewise deterministic Monte Carlo methods (PDMC, most notably the Bouncy Particle Sampler (BPS) (Bouchard-Côté et al., 2017) and the Zig-Zag Sampler (ZZ) (Bierkens et al., 2019)), as a competitor for classical MCMC algorithms such as Metropolis-Hastings and Gibbs sampling. We refer to (Fearnhead et al., 2018) for an accessible introduction to the PDMC setting. The primary benefits of these methods are the possibility of *exact subsampling* and *non-reversibility*. Exact subsampling refers to the possibility of using only a subset of the full data set (or even just a single observation) at each iteration of the algorithm, without introducing bias in the output of the algorithm (Fearnhead et al., 2018). Non-reversibility is a property of MCMC algorithms related to a notion of direction of the algorithm, reducing the number of backtracking steps, thus reducing the diffusivity of the algorithm and reducing the asymptotic variance; as analyzed e.g. in (Diaconis et al., 2000; Andrieu & Livingstone, 2019).

The current key proponents BPS and ZZ of the PDMC paradigm share the following description of their dynamics. The process moves continuously in time according to a constant velocity over random time intervals, which are separated by ‘switching events’. These switching events occur at stochastic times at which the velocity, or a component of it, is either reflected, or randomly refreshed. The direction of a reflection, and the random time at which it occurs, is influenced by the target probability distribution.

In this paper we explore the effect of modifying the property of constant velocity. By doing so we introduce the Boomerang Sampler which has dynamics of the simple form $\frac{dx}{dt} = v$, $\frac{dv}{dt} = -x$. Similar ideas were introduced in (Vanetti et al., 2017) and termed Hamiltonian-BPS, a method which can be seen as a special case of our approach. We generalise the Hamiltonian-BPS algorithm in three important ways.

1. We relax a condition which restricts the covariance function of the auxiliary velocity process to be isotropic. This generalisation is crucial to ensure good convergence properties of the algorithm.
2. Furthermore we extend the Boomerang Sampler to allow for exact subsampling (as introduced above),

thus permitting its application efficiently for large data sets.

3. We also introduce a factorised extension of the sampler which has important computational advantages in the common situation where the statistical model exhibits suitable conditional dependence structure.

Our method also has echoes of the elliptical slice sampler (Murray et al., 2010) which has been a successful discrete-time MCMC method especially within machine learning applications. Both methods are strongly motivated by Hamiltonian dynamics although there are also major differences in the two approaches. Finally we mention some other PDMP methods with non-linear dynamics such as Randomized HMC (Bou-Rabee & Sanz-Serna, 2017; Deligiannidis et al., 2018), and others (Markovic & Sepehri, 2018; Terenin & Thorngren, 2018).

We shall study the Boomerang Sampler and two subsampling alternatives theoretically by analysing the interaction of Bayesian posterior contraction, data size (n) and subsampling schemes in the regular (smooth density) case. We shall show that no matter the rate of posterior contraction, a suitably constructed subsampled Boomerang sampler achieves an $\mathcal{O}(n)$ advantage over non-subsampled algorithms.

At the same time, we show that for the (non-subsampled) Boomerang Sampler, the number of switching events, and thus the computational cost, can be reduced by factor $\mathcal{O}(1/d)$ (where d is the number of dimensions) relative to other piecewise deterministic methods, thanks to the deterministic Hamiltonian dynamics of the Boomerang Sampler.

We illustrate these analyses with empirical investigations in which we compare the properties of Boomerang samplers against other PDMC benchmarks demonstrating the superiority of subsampled Boomerang for sufficiently large data size for any fixed dimension in the setting of logistic regression. We shall also give an empirical study to compare the Boomerang Sampler with its competitors as dimension increases. Finally, as a potentially very useful application we describe the simulation of diffusion bridges using the Factorised Boomerang Sampler, demonstrating substantial computational advantages compared to its natural alternatives.

Notation

For $\mathbf{a} \in \mathbb{R}^d$ and Σ a positive definite matrix in $\mathbb{R}^{d \times d}$ we write $\mathcal{N}(\mathbf{a}, \Sigma)$ for the Gaussian distribution in \mathbb{R}^d with mean \mathbf{a} and covariance matrix Σ . Let $\langle \cdot, \cdot \rangle$ denote the Euclidean inner product in \mathbb{R}^d . We write $(a)_+ := \max(a, 0)$ for the positive part of $a \in \mathbb{R}$, and we write $\langle \cdot, \cdot \rangle_+ := (\langle \cdot, \cdot \rangle)_+$ for the positive part of the inner product.

2. The Boomerang Sampler

The Boomerang Sampler is a continuous time, piecewise deterministic Markov process (PDMP) with state space $S = \mathbb{R}^d \times \mathbb{R}^d$. The two copies of \mathbb{R}^d will be referred to as the *position space* and the *velocity space*, respectively. Our primary interest is in sampling the position coordinate, for which the auxiliary velocity coordinate is a useful tool for us.

Let μ_0 denote a Gaussian measure on S specified by $\mu_0 = \mathcal{N}(\mathbf{x}_*, \Sigma) \otimes \mathcal{N}(\mathbf{0}, \Sigma)$, where Σ is a positive definite matrix in $\mathbb{R}^{d \times d}$. Often we take $\mathbf{x}_* = \mathbf{0}$ to shorten expressions, which can be done without loss of generality by a shift in the position coordinate. The measure μ_0 will be referred to as the *reference measure*. The Boomerang Sampler is designed in such a way that it has stationary probability distribution μ with density $\exp(-U(\mathbf{x}))$ relative to μ_0 . Equivalently, it has density

$$\exp\left(-U(\mathbf{x}) - \frac{1}{2}(\mathbf{x} - \mathbf{x}_*)^\top \Sigma^{-1}(\mathbf{x} - \mathbf{x}_*) - \frac{1}{2}\mathbf{v}^\top \Sigma^{-1}\mathbf{v}\right)$$

relative to the Lebesgue measure $d\mathbf{x} \otimes d\mathbf{v}$ on $\mathbb{R}^d \times \mathbb{R}^d$. We assume that this density has a finite integral. The marginal distribution of μ with respect to \mathbf{x} is denoted by Π .

The Boomerang process moves along deterministic trajectories $(\mathbf{x}_t, \mathbf{v}_t) \in \mathbb{R}^d \times \mathbb{R}^d$ which change direction at random times. The deterministic trajectories satisfy the following simple ordinary differential equation:

$$\frac{d\mathbf{x}_t}{dt} = \mathbf{v}_t, \quad \frac{d\mathbf{v}_t}{dt} = -(\mathbf{x}_t - \mathbf{x}_*), \quad (1)$$

with explicit solution $\mathbf{x}_t = \mathbf{x}_* + (\mathbf{x}_0 - \mathbf{x}_*) \cos(t) + \mathbf{v}_0 \sin(t)$, $\mathbf{v}_t = -(\mathbf{x}_0 - \mathbf{x}_*) \sin(t) + \mathbf{v}_0 \cos(t)$. Note that $(\mathbf{x}, \mathbf{v}) \mapsto \langle \mathbf{x} - \mathbf{x}_*, \mathbf{Q}(\mathbf{x} - \mathbf{x}_*) \rangle + \langle \mathbf{v}, \mathbf{Q}\mathbf{v} \rangle$ is invariant with respect to the flow of (1) for any symmetric matrix \mathbf{Q} . In particular the flow of (1) preserves the Gaussian measure μ_0 on S .

Given an initial position $(\mathbf{x}_0, \mathbf{v}_0) \in S$, the process moves according to the motion specified by (1), resulting in a trajectory $(\mathbf{x}_t, \mathbf{v}_t)_{t \geq 0}$, until the first event occurs. The distribution of the first reflection event time T is specified by

$$\mathbb{P}(T \geq t) = \exp\left(-\int_0^t \lambda(\mathbf{x}_s, \mathbf{v}_s) ds\right),$$

where $\lambda : S \rightarrow [0, \infty)$ is the *event rate* and is specified as

$$\lambda(\mathbf{x}, \mathbf{v}) = \langle \mathbf{v}, \nabla U(\mathbf{x}) \rangle_+. \quad (2)$$

For $\mathbf{x} \in \mathbb{R}^d$ we define the *contour reflection* $\mathbf{R}(\mathbf{x})$ to be the linear operator from \mathbb{R}^d to \mathbb{R}^d given, for $(\mathbf{x}, \mathbf{v}) \in S$, by

$$\mathbf{R}(\mathbf{x})\mathbf{v} = \mathbf{v} - \frac{2\langle \nabla U(\mathbf{x}), \mathbf{v} \rangle}{|\Sigma^{1/2}\nabla U(\mathbf{x})|^2} \Sigma \nabla U(\mathbf{x}). \quad (3)$$

Importantly the reflection satisfies

$$\langle \mathbf{R}(\mathbf{x})\mathbf{v}, \nabla U(\mathbf{x}) \rangle = -\langle \mathbf{v}, \nabla U(\mathbf{x}) \rangle \quad (4)$$

and

$$|\Sigma^{-1/2} \mathbf{R}(\mathbf{x}) \mathbf{v}| = |\Sigma^{-1/2} \mathbf{v}|, \quad (5)$$

which are key in establishing that the resulting Boomerang Sampler has the correct stationary distribution.

At the random time T at which a switch occurs, we put $\mathbf{v}_T := \mathbf{R}(\mathbf{x}_{T-}) \mathbf{v}_{T-}$, where we use the notation $\mathbf{y}_{t-} := \lim_{s \uparrow t} \mathbf{y}_s$. The process then starts afresh according to the dynamics (1) from the new position $(\mathbf{x}_T, \mathbf{v}_T)$. Additionally, at random times generated by a homogeneous Poisson process with rate $\lambda_{\text{refr}} > 0$ the velocity is refreshed, i.e. at such a random time T we independently draw $\mathbf{v}_T \sim \mathcal{N}(\mathbf{0}, \Sigma)$. This additional input of randomness guarantees that the Boomerang Sampler can visit the full state space and is therefore ergodic, as is the case for e.g. BPS (Bouchard-Côté et al., 2017).

In Section 1 of the Supplement we define the generator of the Boomerang Sampler, which can in particular be used to prove that μ is a stationary distribution for the Boomerang process, and which can be used in subsequent studies to understand its probabilistic properties.

Remark 2.1 (On the choice of the reference measure). In principle we can express *any* probability distribution $\Pi(d\mathbf{x}) \propto \exp(-E(\mathbf{x})) d\mathbf{x}$ as a density relative to μ_0 by defining

$$U(\mathbf{x}) = E(\mathbf{x}) - \frac{1}{2}(\mathbf{x} - \mathbf{x}_*)^\top \Sigma^{-1}(\mathbf{x} - \mathbf{x}_*). \quad (6)$$

As mentioned before we can take μ_0 to be identical to a Gaussian prior measure in the Bayesian setting. Alternatively, and this is an approach which we will adopt in this paper, we may choose μ_0 to be a Gaussian approximation of the measure Π which may be obtained at relatively small computational cost in a preconditioning step.

2.1. Factorised Boomerang Sampler

As a variation to the Boomerang Sampler introduced above we introduce the Factorised Boomerang Sampler (FBS), which is designed to perform well in situations where the conditional dependencies in the target distribution are sparse. For simplicity we restrict to the case with a diagonal reference covariance $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$.

The deterministic dynamics of the FBS are identical to those of the standard Boomerang Sampler, and given by (1). The difference is that every component of the velocity has its own switching intensity. This is fully analogous with the difference between BPS and ZZ, where the latter can be seen as a factorised Bouncy Particle Sampler. In the current case, this means that as switching intensity for the i -th component of the velocity we take

$$\lambda_i(\mathbf{x}, \mathbf{v}) = (v_i \partial_i U(\mathbf{x}))_+,$$

and once an event occurs, the velocity changes according to the operator $\mathbf{F}_i(\mathbf{v})$ given by

$$\mathbf{F}_i(\mathbf{v}) = (v_1, \dots, v_{i-1}, -v_i, v_{i+1}, \dots, v_d)^\top.$$

Also, the velocity of each component is refreshed according to $v_i \sim \mathcal{N}(0, \sigma_i^2)$ at rate $\lambda_{\text{refr},i} > 0$.

Note that the computation of the reflections has a computational cost of $\mathcal{O}(1)$, compared to the reflections in (3) being at least of $\mathcal{O}(d)$, depending upon the sparsity of Σ . The sparse conditional dependence structure implies that the individual switching intensities $\lambda_i(\mathbf{x})$ are in fact functions of a subset of the components of \mathbf{x} , contributing to a fast computation. This feature can be exploited by an efficient ‘local’ implementation of the FBS algorithm which reduces the number of Poisson times simulated by the algorithm (similar in spirit to the local Bouncy Particle Sampler (Bouchard-Côté et al., 2017) and the local Zig-Zag sampler in (Bierkens et al., 2020)). In Section 3.2 we will briefly comment on the dimensional scaling of FBS. As an illustration of a realistic use, FBS will be applied to the simulation of diffusion bridges in Section 4.2.

2.2. Subsampling with control variates

Let $E(\mathbf{x})$ be the energy function, i.e., negative log density of Π with respect to the Lebesgue measure. Consider the setting where $E(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n E^i(\mathbf{x})$, as is often the case in e.g. Bayesian statistics or computational physics. (Let us stress that n represents a quantity such as the number of interactions or the size of the data, and *not* the dimensionality of \mathbf{x} , which is instead denoted by d .) Using this structure, we introduce a subsampling method using the Gaussian reference measure as a tool for the efficient construction of the Monte Carlo method.

Relative to a Gaussian reference measure with covariance Σ centred at \mathbf{x}_* , the negative log density is given by (6). Let us assume

$$\Sigma = [\nabla^2 E(\mathbf{x}_*)]^{-1} \quad (7)$$

for a reference point \mathbf{x}_* . In words, the curvature of the reference measure will agree around \mathbf{x}_* with the curvature of the target distribution. We can think of \mathbf{x}_* as the mean or mode of an appropriate Gaussian approximation used for the Boomerang Sampler. Note however that we shall not require that $\nabla E(\mathbf{x}_*) = \mathbf{0}$ for the sampler and its subsampling alternatives to work well, although some restrictions will be imposed in Section 3.1. In this setting it is possible to employ a subsampling method which is exact, in the sense that it targets the correct stationary distribution. This is an extension of methodology used for subsampling in other piecewise deterministic methods, see e.g. (Fearnhead et al., 2018) for an overview.

Assume for notational convenience that $\mathbf{x}_* = \mathbf{0}$. As an

unbiased estimator for the log density gradient of U we could simply take

$$\widehat{\nabla U(\mathbf{x})} = \nabla E^I(\mathbf{x}) - \nabla^2 E(\mathbf{0})\mathbf{x}, \quad (8)$$

where I is a random variable with uniform distribution over $\{1, \dots, n\}$. We shall see in Proposition 3.1 that this estimator will lead to weights which increase with n and therefore we shall consider a control variate alternative.

Therefore also consider the control variate gradient estimator $\widehat{\nabla U(\mathbf{x})} = \mathbf{G}^I(\mathbf{x})$, where, for $i = 1, \dots, n$,

$$\mathbf{G}^i(\mathbf{x}) = \nabla E^i(\mathbf{x}) - \nabla^2 E^i(\mathbf{0})\mathbf{x} - \nabla E^i(\mathbf{0}) + \nabla E(\mathbf{0}). \quad (9)$$

Taking the expectation with respect to I ,

$$\begin{aligned} \mathbb{E}_I \widehat{\nabla U(\mathbf{x})} &= \frac{1}{n} \sum_{i=1}^n \{ \nabla E^i(\mathbf{x}) - \nabla^2 E^i(\mathbf{0})\mathbf{x} - \nabla E^i(\mathbf{0}) + \nabla E(\mathbf{0}) \} \\ &= \nabla E(\mathbf{x}) - \nabla^2 E(\mathbf{0})\mathbf{x} = \nabla U(\mathbf{x}), \end{aligned}$$

so that $\widehat{\nabla U(\mathbf{x})}$ is indeed an unbiased estimator for $\nabla U(\mathbf{x})$.

In Section 3 we shall show that $\widehat{\nabla U(\mathbf{x})}$ has significantly superior scaling properties for large n than $\nabla U(\mathbf{x})$.

Remark 2.2. In various situations we can find a reference point \mathbf{x}_* such that $\nabla E(\mathbf{x}_*) = \mathbf{0}$, in which case the final term in (9) vanishes. We include the term here so that it can accommodate the general situation in which $\nabla E(\mathbf{x}_*) \neq \mathbf{0}$.

Upon reflection, conditional on the random draw I , we reflect according to

$$\mathbf{R}^I(\mathbf{x})\mathbf{v} = \mathbf{v} - \frac{2\langle \mathbf{G}^I(\mathbf{x}), \mathbf{v} \rangle}{|\Sigma^{1/2} \mathbf{G}^I(\mathbf{x})|^2} \Sigma \mathbf{G}^I(\mathbf{x}).$$

The Boomerang Sampler that switches at the random rate $\widehat{\lambda(\mathbf{x}, \mathbf{v})} = \langle \mathbf{v}, \widehat{\nabla U(\mathbf{x})} \rangle_+$, and reflects according to \mathbf{R}^I will preserve the desired target distribution in analogy to the argument found in (Bierkens et al., 2019).

2.3. Simulation

The implementation of the Boomerang Sampler depends crucially on the ability to simulate from a nonhomogeneous Poisson process with a prescribed rate. In this section we will make a few general comments on how to achieve these tasks for the Boomerang Sampler and for the Subsampled Boomerang Sampler.

Suppose we wish to generate the first event according to a switching intensity $\lambda(\mathbf{x}_t, \mathbf{v}_t)$ where $(\mathbf{x}_t, \mathbf{v}_t)$ satisfy (1). This is challenging because it is non-trivial to generate points according to time inhomogeneous Poisson process,

but also the function $\lambda(\mathbf{x}_t, \mathbf{v}_t)$ may be expensive to evaluate. It is customary in simulation of PDMPs to employ the technique of *Poisson thinning* to generate an event according to a deterministic rate function $\bar{\lambda}(t) \geq 0$, referred to as *computational bound*, such that $\lambda(\mathbf{x}_t, \mathbf{v}_t) \leq \bar{\lambda}(t)$ for all $t \geq 0$. The function $\bar{\lambda}(t)$ should be suitable in the sense that we can explicitly simulate T according to the law

$$\mathbb{P}(T \geq t) = \exp\left(-\int_0^t \bar{\lambda}(s) ds\right).$$

After generating T from this distribution, we accept T as a true switching event with probability $\lambda(\mathbf{x}_T, \mathbf{v}_T)/\bar{\lambda}(T)$. As a consequence of this procedure, the first time T that gets accepted in this way is a Poisson event with associated intensity $\lambda(\mathbf{x}_t, \mathbf{v}_t)$.

In this paper we will only consider bounds $\bar{\lambda}(t)$ of the form $\bar{\lambda}(t; \mathbf{x}_0, \mathbf{v}_0) = a(\mathbf{x}_0, \mathbf{v}_0) + tb(\mathbf{x}_0, \mathbf{v}_0)$. We will call the bound *constant* if $b(\mathbf{x}, \mathbf{v}) = 0$ for all (\mathbf{x}, \mathbf{v}) , and *affine* otherwise. As a simple example, consider the situation in which $|\nabla U(\mathbf{x})| \leq m$ for all \mathbf{x} . In this case we have

$$\lambda(\mathbf{x}, \mathbf{v}) = \langle \mathbf{v}, \nabla U(\mathbf{x}) \rangle_+ \leq m|\mathbf{v}| \leq m\sqrt{|\mathbf{x}|^2 + |\mathbf{v}|^2}.$$

Since the final expression is invariant under the dynamics (1), we find that

$$\lambda(\mathbf{x}_t, \mathbf{v}_t) \leq m\sqrt{|\mathbf{x}_0|^2 + |\mathbf{v}_0|^2}, \quad t \geq 0,$$

which gives us a simple constant bound.

In the case of subsampling the switching intensity $\widehat{\lambda(\mathbf{x}, \mathbf{v})}$ is random. Still, the bound $\bar{\lambda}(t; \mathbf{x}_0, \mathbf{v}_0)$ needs to be an upper bound for all random realizations of $\widehat{\lambda(\mathbf{x}, \mathbf{v})}$. In the case we use the unbiased gradient estimator $\widehat{\nabla U(\mathbf{x})} = \mathbf{G}^I$ of (9), we can bound e.g.

$$\widehat{\lambda(\mathbf{x}, \mathbf{v})} \leq \sup_{i, \mathbf{x}} |\mathbf{G}^i(\mathbf{x})| |\mathbf{v}| \leq \sup_{i, \mathbf{x}} |\mathbf{G}^i(\mathbf{x})| \sqrt{|\mathbf{x}|^2 + |\mathbf{v}|^2},$$

assuming all gradient estimators \mathbf{G}^i are globally bounded. We will introduce different bounds in detail in Section 2 of the Supplement.

3. Scaling for large data sets and large dimension

3.1. Robustness to large n

In this section, we shall investigate the variability of the rates induced by the Boomerang Sampler and its subsampling options. The size of these rates is related to the size of the upper bounding rate Poisson process used to simulate them. Moreover, the rate of the upper bounding Poisson rate is proportional to the number of density evaluations, which

in turn is a sensible surrogate for the computing cost of running the algorithm.

As in Section 2.2, we describe E as a sum of n constituent negative log-likelihood terms: $E(\mathbf{x}) = -\sum_{i=1}^n \ell_i(\mathbf{x})$. (In the notation above we are just setting $\ell_i(\mathbf{x}) = -nE_i(\mathbf{x})$.) Under suitable regularity conditions, the target probability measure Π satisfies posterior contraction around $\mathbf{x} = \mathbf{0}$ at the rate η , that is for all ϵ there exists $\delta > 0$ such that $\Pi(B_{n^{-\eta}\delta}(\mathbf{0})) > 1 - \epsilon$ where $B_r(\mathbf{0})$ denotes the ball of radius r centred at $\mathbf{0}$. As a result of this, we typically have velocities of order $n^{-\eta}$ ensuring that the dynamics in (1) circles the state space in $\mathcal{O}(1)$ time.

The various algorithms will have computational times roughly proportional to the number of likelihood evaluations, which in turn depends on the event rate (and its upper bound). Therefore we shall introduce explicitly the subsampling bounce rates corresponding to the use of the unbiased estimators in (8) and (9).

$$\widetilde{\lambda(\mathbf{x}, \mathbf{v})} = \langle \mathbf{v}, \widetilde{\nabla U(\mathbf{x})} \rangle_+; \quad \widehat{\lambda(\mathbf{x}, \mathbf{v})} = \langle \mathbf{v}, \widehat{\nabla U(\mathbf{x})} \rangle_+.$$

To simplify the arguments below, we also assume that ℓ_i has all its third derivatives uniformly bounded, implying that all third derivative terms of E are bounded by a constant multiple of n . This allows us to write down the expansion

$$\begin{aligned} \nabla U(\mathbf{x}) &= \nabla E(\mathbf{0}) + \nabla^2 E(\mathbf{0})\mathbf{x} - \Sigma^{-1}\mathbf{x} + \mathcal{O}(n|\mathbf{x}|^2) \\ &= \nabla E(\mathbf{x}) \\ &= \nabla E(\mathbf{0}) + \mathcal{O}(n|\mathbf{x}|^2). \end{aligned} \quad (10)$$

Similarly we can write

$$\begin{aligned} \widehat{\nabla U(\mathbf{x})} &= n\nabla \ell_I(\mathbf{x}) - n\nabla^2 \ell_I(\mathbf{0})\mathbf{x} - n\ell_I(\mathbf{0}) \\ &\quad + \nabla E(\mathbf{0}) - \Sigma^{-1}\mathbf{x} \\ &= \nabla E(\mathbf{0}) + \mathcal{O}(n|\mathbf{x}|^2). \end{aligned} \quad (11)$$

using the same Taylor series expansion.

We can now use this estimate directly to obtain bounds on the event rates. We summarise this discussion in the following result.

Proposition 3.1. *Suppose that $\mathbf{x}, \mathbf{v} \in B_{n^{-\eta}\delta}(\mathbf{0})$ for some δ , and under the assumptions described above, we have that*

$$\lambda(\mathbf{x}, \mathbf{v}) \leq \mathcal{O}(n^{-\eta}(|\nabla E(\mathbf{0})| + n^{1-2\eta})) \quad (12)$$

$$\widetilde{\lambda(\mathbf{x}, \mathbf{v})} \leq \mathcal{O}(|\nabla E(\mathbf{0})|) + \mathcal{O}(n) \quad (13)$$

$$\widehat{\lambda(\mathbf{x}, \mathbf{v})} \leq \mathcal{O}(n^{-\eta}(|\nabla E(\mathbf{0})| + n^{1-2\eta})) \quad (14)$$

Thus the use of $\widehat{\nabla U(\mathbf{x})}$ does not result in an increased event rate (in order of magnitude). There is therefore an $\mathcal{O}(n)$ computational advantage obtained from using subsampling due to each target density valuation being $\mathcal{O}(n)$ quicker.

Proposition 3.1 shows that as long as the reference point \mathbf{x}^* (chosen to be $\mathbf{0}$ here for convenience) is chosen to be sufficiently close to the mode so that $|\nabla E(\mathbf{0})|$ is at most $\mathcal{O}(n^{1-2\eta})$, then we have that

$$\lambda(\mathbf{x}, \mathbf{v}) = \widehat{\lambda(\mathbf{x}, \mathbf{v})} = \mathcal{O}(n^{1-3\eta}).$$

Note that this rate can go to 0 when $\eta > 1/3$. In particular in the regular case where Bernstein von-Mises theorem holds, we have $\eta = 1/2$. In this case the rate of jumps for the Boomerang can recede to 0 at rate $n^{-1/2}$ so long as $|\nabla E(\mathbf{0})|$ is at most $\mathcal{O}(1)$.

3.2. Scaling with dimension

In this section, we will discuss how the Boomerang Sampler has an attractive scaling property for high dimension. This property is qualitatively similar to the preconditioned Crank-Nicolson algorithm (Neal, 1999; Beskos et al., 2008) and the elliptical slice sampler (Murray et al., 2010) which take advantage of the reference Gaussian distribution.

The dimensional complexity of BPS and ZZ was studied in (Bierkens et al., 2018; Deligiannidis et al., 2018; Andrieu et al., 2018). For the case of an isotropic target distribution, the rate of reflections per unit of time is constant for BPS and proportional to d for ZZ with unit speeds in all directions. On the other hand, the time until convergence is of order d for the BPS and 1 for ZZ. Therefore, the total number of reflections required for convergence of these two algorithms is of the same order which grows linearly with dimension.

For the Boomerang Sampler we consider the following setting. Consider reference measures $\mu_{0,d}(d\mathbf{x}, d\mathbf{v}) = \mathcal{N}(\mathbf{0}, \Sigma_d) \otimes \mathcal{N}(\mathbf{0}, \Sigma_d)$ for increasing dimension d , where for every $d = 1, 2, \dots$, Σ_d is a d -dimensional positive definite matrix. Relative to these reference measures we consider a sequence of potential functions $U_d(\mathbf{x})$. Thus relative to Lebesgue measure our target distribution $\Pi_d(d\mathbf{x})$ has density $\exp(-E_d(\mathbf{x}))$, where $E_d(\mathbf{x}) = U_d(\mathbf{x}) + \frac{1}{2}\langle \mathbf{x}, \Sigma_d^{-1}\mathbf{x} \rangle$. Let \mathbb{E}_d denote expectation with respect to $\Pi_d(d\mathbf{x}) \otimes \mathcal{N}(\mathbf{0}, \Sigma_d)(d\mathbf{v})$. We assume that the sequence (U_d) satisfies

$$\sup_{d=1,2,\dots} \mathbb{E}_d[|\Sigma_d^{1/2}\nabla U_d(\mathbf{x})|^2] < \infty, \quad (15)$$

The condition (15) arises naturally for instance in the context of Gaussian regression, spatial statistics, Bayesian inverse problems as well as the setting of the diffusion bridge simulation example described in detail in Section 4.2.

Furthermore we assume that the following form of the Poincaré inequality holds,

$$\mathbb{E}_d[f_d(\mathbf{x})^2] \leq \frac{1}{C^2} \mathbb{E}_d[|\Sigma_d^{1/2}\nabla f_d(\mathbf{x})|^2] \quad (16)$$

with constant $C > 0$ independent of dimension, and where $f_d : \mathbb{R}^d \rightarrow \mathbb{R}$ is any mean zero differentiable function. A sufficient condition for (16) to hold is

$$C^2 I \preceq \Sigma_d^{1/2} \nabla^2 E_d(\mathbf{x}) \Sigma_d^{1/2} = \Sigma_d^{1/2} \nabla^2 U_d(\mathbf{x}) \Sigma_d^{1/2} + I$$

by the classical Brascamp-Lieb inequality (Brascamp & Lieb, 1976; Bakry et al., 2014); note that it may also hold in the non-convex case, see e.g. (Lorenzi & Bertoldi, 2007), Section 8.6.

Under the stated assumptions we argue that (i) the expected number of reflections per unit time scales as $\mathcal{O}(1)$ with respect to dimension, and (ii) within a continuous time interval that scales as $\mathcal{O}(1)$, the Boomerang Sampler mixes well. Claims (i) and (ii) are provided with a heuristic motivation in the Supplement. A rigorous proof for this claim remains part of our future work.

In the ideal but non-sparse scenario, the computational cost of the event time calculation for the Boomerang Sampler is thus expected to be a factor d smaller compared to BPS and ZZ assuming that the cost per event is the same for these algorithms. However, this comparison is unrealistic since in general we can not simulate reflections directly. In practice, we need to use the thinning method as discussed in Section 2.3. The thinning method introduces a significant amount of shadow events (which are rejected after inspection), and the true events usually represent a small portion relative to the number of shadow events. As a result there can be a high cost for calculating shadow events even when the number of true events is small.

For the FBS, the expected number of events per unit of time is $\sum_{i=1}^d \mathbb{E}_d[(v_i \partial_i U(\mathbf{x}))_+]$. Under the hypothesis above, this is of $\mathcal{O}(d^{1/2})$. Thus, the number of events is much bigger than that of the Boomerang. However, as in the case of ZZ, under a sparse model assumption, the cost of calculation per jump is of constant order whereas it is of the order of d for the Boomerang Sampler. Therefore, the Factorised Boomerang Sampler should outperform the Boomerang Sampler for this sparse setup.

4. Applications and experiments

4.1. Logistic regression

As a suitable test bed we consider the logistic regression inference problem. Given predictors $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n)}$ in \mathbb{R}^d , and outcomes $z^{(1)}, \dots, z^{(n)}$ in $\{0, 1\}$, we define the log likelihood function as

$$\ell(\mathbf{x}) = - \sum_{i=1}^n \left\{ \log(1 + e^{\mathbf{x}^\top \mathbf{y}^{(i)}}) - z^{(i)} \mathbf{x}^\top \mathbf{y}^{(i)} \right\}.$$

Furthermore we impose a Gaussian prior distribution over \mathbf{x} which for simplicity we keep fixed to be a standard normal

distribution throughout these experiments. As a result we arrive at the negative log target density

$$E(\mathbf{x}) = \sum_{i=1}^n \left\{ \log(1 + e^{\mathbf{x}^\top \mathbf{y}^{(i)}}) - z^{(i)} \mathbf{x}^\top \mathbf{y}^{(i)} \right\} + \frac{1}{2} \mathbf{x}^\top \mathbf{x}.$$

As a preprocessing step when applying the Boomerang Sampler, and all subsampled methods, we find the mode \mathbf{x}_* of the posterior distribution and define Σ by (7). We apply the Boomerang Sampler, with and without subsampling. These samplers are equipped with an affine computational bound and a constant computational bound respectively, both discussed in Section 2 of the Supplement (the affine bound is usually preferred over a constant bound, but a useful affine bound is not available in the subsampling case).

We compare the Boomerang to both BPS and ZZ with and without subsampling. In all subsampling applications we employ appropriate control variance techniques to reduce the variability of the random switching intensities, as discussed in Section 2.2. Furthermore in the dimension dependent study we include the Metropolis adjusted Langevin algorithm (MALA) for comparison. Throughout these experiments we use Effective Sample Size (ESS) per second of CPU time as measure of the efficiency of the methods used. ESS is estimated using the Batch Means method, where we take a fixed number of 50 batches for all our estimates. ESS is averaged over the dimensions of the simulation and then divided by the runtime of the algorithm to obtain ‘‘average ESS per second’’ (other ESS summaries could also have been used). The time horizon is throughout fixed at 10,000 (with 10,000 iterations for MALA). For ZZ and BPS the magnitude of the velocities is rescaled to be comparable on average with Boomerang, to avoid unbalanced runtimes of the different algorithms. In Figures 1 and 2 the boxplots are taken over 20 randomly generated experiments, where each experiment corresponds to a logistic regression problem with a random (standard normal) parameter, based on randomly generated data from the model.¹ The refreshment rates for BPS and the Boomerang Samplers are taken to be 0.1.

The Boomerang Sampler is seen to outperform the other algorithms, both in terms of scaling with dimension as with respect to an increase in the number of observations. For a fixed dimension, the subsampling algorithms will clearly outperform the non-subsampling algorithms as number of observations n grows. In particular, the ESS/sec stays fixed for the subsampled algorithms, and decreases as $\mathcal{O}(n)$ for the non-subsampled versions. In this case, we did not include the MALA algorithm since we observed its complexity strongly deteriorating as the number of observations

¹The code used to carry out all of the experiments of this paper may be found online at <https://github.com/jbierkens/ICML-boomerang>.




Figure 1. Scaling of Boomerang Sampler compared to other PDMC methods for the logistic regression problem of Section 4.1 as a function of the number of observations. Here $d = 2$.

increases. For a large number of observations ($n \geq 10,000$, $d = 2$) we see that the Boomerang Sampler (with and without subsampling) accepts almost none of the proposed switches. This means that effectively we are sampling from the Gaussian reference measure. This observed behaviour is in line with the scaling analysis in Section 3.1.

In the second experiment we let the dimension d grow for a fixed number of observations. The subsampling algorithms currently do not scale as well as the non-sampled versions. For practical purposes we therefore only consider non-sampled algorithms for the comparison with respect to dimensional dependence. For the dimensions $d \leq 32$ we tested the Boomerang outperforms MALA, but it seems empirically that MALA has a better scaling behaviour with dimension. Note that MALA needs careful tuning to exhibit this good scaling. We remark that the beneficial scaling properties of the underlying Boomerang Process as discussed in Section 3.2 may be adversely affected by suboptimal computational bounds. We are optimistic that the dimensional scaling of subsampled algorithms can be further improved by designing better computational bounds.

In all cases the necessary preprocessing steps can be done very quickly. In particular the plots are not affected by including (or excluding) the preprocessing time in the computation of ESS/sec.




Figure 2. Scaling of Boomerang Sampler compared to other PDMC methods and MALA for the logistic regression problem of Section 4.1 as a function of the number of dimensions. Here the number of observations is $n = 1,000$.

4.2. Diffusion bridges

In (Bierkens et al., 2020) the authors introduce a framework for the simulation of diffusion bridges (diffusion processes conditioned to hit a prescribed endpoint) taking strong advantage of the use of factorised piecewise deterministic samplers. This invites the use of the Factorised Boomerang Sampler (FBS). We consider time-homogeneous one-dimensional conditional diffusion processes (diffusion bridges) of the form

$$dX_t = b(X_t)dt + dW_t, \quad X_0 = u, \quad X_T = v$$

where W is a scalar Brownian motion and b satisfies some mild regularity conditions (see (Bierkens et al., 2020) for details). This simulation problem is an essential building block in Bayesian analysis of non-linear diffusion models with low frequency observations (Roberts & Stramer, 2001).

We consider the approach of (Bierkens et al., 2020) where the diffusion path on $[0, T]$ is expanded with a truncated Faber Schauder basis as

$$X_t^N = \bar{\phi}(t)u + \bar{\phi}(t)v + \sum_{i=0}^N \sum_{j=0}^{2^i-1} \phi_{i,j}(t)x_{i,j}.$$

Here,

$$\bar{\phi}(t) = t/T, \quad \bar{\phi}(t) = 1 - t/T,$$

$$\phi_{0,0}(t) = \sqrt{T}((t/T)\mathbf{1}_{[0,T/2]}(t) + (1 - t/T)\mathbf{1}_{(T/2,T]}(t)),$$

$$\phi_{i,j}(t) = 2^{-i/2}\phi_{0,0}(2^i t - jT) \quad i \geq 0, \quad 0 \leq j \leq 2^i - 1,$$

are the Faber-Schauder functions and N is the truncation of the expansion. In (Bierkens et al., 2020), ZZ is used to sam-

ple the corresponding coefficients $\mathbf{x} := (x_{0,0}, \dots, x_{N,2^N-1})$ which have a density measure written with respect to a standard Gaussian reference measure (corresponding to a Brownian bridge measure in the path space, see (Bierkens et al., 2020) for details). In particular we have that

$$\frac{d\mu}{d\mu_0}(\mathbf{x}, \mathbf{v}) \propto \exp \left\{ -\frac{1}{2} \int_0^T (b^2(X_s^N) + b'(X_s^N)) ds \right\} \quad (17)$$

where b' is the derivative of b and $\mu_0 = \mathcal{N}(\mathbf{0}, \mathbf{I}) \otimes \mathcal{N}(\mathbf{0}, \mathbf{I})$ with \mathbf{I} the $2^{N+1} - 1$ dimensional identity matrix. The measure given by (17) has a remarkable conditional independence property (Proposition 2, (Bierkens et al., 2020)) and the coefficients $x_{i,j}$, for i large, responsible for the local behaviour of the process, are approximately independent standard Gaussian, reflecting the fact that, locally, the process behaves as a Brownian motion.

In (Bierkens et al., 2020) the authors devise a local implementation of ZZ which optimally exploits the sparse conditional independence structure of the target distribution, alleviating the computational costs in high dimensional setting (e.g. of a high truncation level N). Since the Girsanov density (17) is expressed relative to a standard normal distribution on the coefficients \mathbf{x} , the Factorised Boomerang Sampler is a natural candidate for a further reduction in computational cost, by reducing the required number of simulated events, in particular at the higher levels where the coefficients have approximately a Gaussian distribution. This will allow for a further increase of the truncation level N and/or faster computations at a fixed truncation levels.

We consider the the class of diffusion bridges with drift equal to

$$b(x) = \alpha \sin(x), \quad \alpha \geq 0. \quad (18)$$

The higher α , the stronger is the attraction of the diffusion paths to the stable points $(2k - 1)\pi, k \in \mathbb{N}$ while for $\alpha = 0$ the process reduces to a Brownian bridge with $\mu = \mu_0$. Equivalently to (Bierkens et al., 2020), we use subsampling as the gradient of the log density in (17) involves a time integral that cannot be solved analytically in most of the cases. The unbiased estimator for $\partial_{x_{i,j}} U(\mathbf{x})$ is the integrand evaluated at a uniform random point multiplied by the range of the integral. The Poisson bounding rates used for the subsampling can be found in the Supplement, Section 5.

Figure 3 shows the resulting bridges for $\alpha = 1$, starting at $u = -\pi$ and hitting $v = 3\pi$ at final time $T = 50$ after running the FBS with clock $T^* = 20000$, as simulated on a standard desktop computer. The refreshment rate relative to each coefficient $x_{i,j}$ is fixed to $\lambda_{\text{refr},i,j} = 0.01$ and the truncation of the expansion is $N = 6$.

In Figure 4, we compare the performances of the Boomerang Sampler and ZZ by computing the average number of reflections (y -axis on a log-scale) for the coefficients $x_{i,j}$ at




Figure 3. 1000 diffusion bridges with drift equal to (18) with $\alpha = 1$, $u = -\pi$, $v = 3\pi$, $T = 50$, $L = 6$ sampled with the FBS with time horizon $T^* = 20,000$ and refreshment rates $\lambda_{\text{refr},i} = 0.01$ for all i . The straight horizontal lines are the attraction points.

each level (x -axis). The number of reflections is understood as a measure of complexity of the algorithm. We repeat the experiment for $\alpha = 0.5$ and $\alpha = 0$ (where $\mu = \mu_0$) and fix the truncation level to be $N = 10$ which corresponds to a $2047 + 2047$ dimensional space for (\mathbf{x}, \mathbf{v}) . For a fair comparison we set the expected ℓ^1 norm of the velocities and the time horizon of the two samplers to be the same. In both cases, the average number of reflections converges to the average number of reflections under μ_0 (dashed lines) indicating that the coefficients at high levels are approximately standard normally distributed but while ZZ requires a fixed number of reflections for sampling from $\mu = \mu_0$, the Boomerang does not, allowing to high resolutions of the diffusion bridges at lower computational cost.




Figure 4. Average number of reflections (on a log-scale) for the coefficients $x_{i,j}$ at the level $i = 0, 1, \dots, 10$ for the diffusion bridge given by (18) with $\alpha = 0.5$ (solid lines) and $\alpha = 0.0$ (dashed lines) for the Zig-Zag Sampler (blue lines) and the Factorised Boomerang Sampler (red lines) with $T^* = 2,000$ and Boomerang refreshment rates $\lambda_{\text{refr},i} = 0.01$ for all i .

4.3. Dependence upon reference measure

In a final experiment we investigate the dependence of the performance of the Boomerang Sampler upon the choice of reference measure. For this we consider a simple setting in which the target distribution is a standard normal distribution in d dimensions. However, instead of using the standard normal distribution as reference measure, we perturb it in two ways: (i) we vary the component-wise variance σ^2 of the reference measure, and (ii) we vary the mean \mathbf{x}_* of the reference measure. Specifically, we choose a reference measure $\mathcal{N}(\mathbf{x}_*, \Sigma) \otimes \mathcal{N}(\mathbf{0}, \Sigma)$, which we choose in case (i) to be $\mathbf{x}_* = \mathbf{0}$, $\Sigma = \sigma^2 \mathbf{I}$, and in case (ii), $\mathbf{x}_* = \alpha(1, \dots, 1)^\top$, $\Sigma = \mathbf{I}$. As performance measure we use the ESS per second for the quantity $|\mathbf{x}|^2$. We use refreshment rate 0.1 for Boomerang, and we compare to the Bouncy Particle Sampler, with refreshment 1.0, with both samplers run over a time horizon of 10,000. In Figure 5 the results of this experiment are displayed for varying σ^2 , and in Figure 6 the results are displayed for varying \mathbf{x}_* . The box plots are taken over 20 experiments of the Boomerang Sampler, which are compared to a single run of the Bouncy Particle Sampler (dashed line).




Figure 5. Effect of perturbing the variance of the reference measure. As reference measure we choose $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) \otimes \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, where σ^2 is varied from 0.5 to 2.0.

In this setting, the Boomerang Sampler significantly outperforms the BPS, although the performance is seen to depend upon the choice of reference measure. Note however that the dependencies of Σ on σ^2 and of \mathbf{x}_* upon α scale as $\text{trace } \Sigma = \sigma^2 d$ and $\|\mathbf{x}_*\| = \alpha d^{1/2}$ respectively, so that in high dimensional cases the sensitivity on \mathbf{x}_* and Σ may be more moderate than might appear from Figures 5 and 6.




Figure 6. Effect of perturbing the mean of the reference measure. As reference measure we choose $\mathcal{N}(\alpha \mathbf{1}, \mathbf{I}) \otimes \mathcal{N}(\mathbf{0}, \mathbf{I})$, where α is varied from 0.0 to 2.0.

5. Conclusion

We presented the Boomerang Sampler as a new and promising methodology, outperforming other piecewise deterministic methods in the large n , moderate d setting, as explained theoretically and by performing a suitable benchmark test. The theoretical properties of the underlying Boomerang Sampler in high dimension are very good. However currently a large computational bound and therefore a large number of rejected switches are hampering the efficiency. We gave a numerical comparison which demonstrates that Boomerang performs well against its natural competitors; however one should be cautious about drawing too many conclusions about the performance of the Boomerang without a more comprehensive simulation study. Further research is required to understand in more detail the dependence upon e.g. reference covariance Σ , centering position \mathbf{x}_* , refreshment rate, computational bounds and the choice of efficiency measure.

We furthermore introduced the Factorised Boomerang Sampler and illustrated its ability to tackle a challenging simulation problem using an underlying sparse structure.

An important direction for further research is the improvement of the computational bounds, in particular with the aim of having good scaling with dimension of subsampled algorithms. Related to this it is important to gain a better understanding of the relative optimality of subsampled versus non-subsampled algorithms in the large n , large d case.

Acknowledgements

We thank the anonymous reviewers and metareviewer for their helpful comments, which have helped to improve the paper. JB and SG acknowledge funding under the research programme ‘Zig-zagging through computational barriers’ with project number 016.Vidi.189.043, which is financed by the Netherlands Organisation for Scientific Research (NWO). KK is supported by JSPS KAKENHI Grant Number 20H04149 and JST CREST Grant Number JPMJCR14D7. GR is supported by EPSRC grants EP/R034710/1 (CoSInES) and EP/R018561/1 (Bayes for Health), and by Wolfson merit award WM140096.

References

- Andrieu, C. and Livingstone, S. Peskun-Tierney ordering for Markov chain and process Monte Carlo: beyond the reversible scenario. *preprint arXiv:1906.06197*, 2019. URL <http://arxiv.org/abs/1906.06197>.
- Andrieu, C., Durmus, A., Nüsken, N., and Roussel, J. Hypocoercivity of Piecewise Deterministic Markov Process-Monte Carlo. *preprint arXiv:1808.08592*, aug 2018. URL <http://arxiv.org/abs/1808.08592>.
- Bakry, D., Gentil, I., and Ledoux, M. Analysis and geometry of markov diffusion operators. *Grundlehren der mathematischen Wissenschaften*, 2014. ISSN 2196-9701. doi: 10.1007/978-3-319-00227-9. URL <http://dx.doi.org/10.1007/978-3-319-00227-9>.
- Beskos, A., Roberts, G. O., Stuart, A., and Voss, J. MCMC methods for diffusion bridges. *Stoch. Dyn.*, 8(3):319–350, 2008. URL <https://doi.org/10.1142/S0219493708002378>.
- Bierkens, J., Kamatani, K., and Roberts, G. O. High-dimensional scaling limits of piecewise deterministic sampling algorithms. *preprint arXiv:1807.11358*, 2018. URL <http://arxiv.org/abs/1807.11358>.
- Bierkens, J., Fearnhead, P., and Roberts, G. O. The Zig-Zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Statist.*, 47(3):1288–1320, 2019. doi: 10.1214/18-AOS1715. URL <https://doi.org/10.1214/18-AOS1715>.
- Bierkens, J., Grazi, S., van der Meulen, F., and Schauer, M. A piecewise deterministic Monte Carlo method for diffusion bridges. *preprint arXiv:2001.05889*, 2020. URL <http://arxiv.org/abs/2001.05889>.
- Bou-Rabee, N. and Sanz-Serna, J. M. Randomized Hamiltonian Monte Carlo. *Ann. Appl. Probab.*, 27(4):2159–2194, 2017. doi: 10.1214/16-AAP1255. URL <https://doi.org/10.1214/16-AAP1255>.
- Bouchard-Côté, A., Vollmer, S. J., and Doucet, A. The Bouncy Particle Sampler: A Non-Reversible Rejection-Free Markov Chain Monte Carlo Method. *Journal of the American Statistical Association*, 2017. doi: 10.1080/01621459.2017.1294075. URL <http://www.tandfonline.com/doi/abs/10.1080/01621459.2017.1294075>.
- Brascamp, H. J. and Lieb, E. H. On extensions of the Brunn-Minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation. *J. Functional Analysis*, 22(4):366–389, 1976. doi: 10.1016/0022-1236(76)90004-5. URL [https://doi.org/10.1016/0022-1236\(76\)90004-5](https://doi.org/10.1016/0022-1236(76)90004-5).
- Brooks, S., Gelman, A., Jones, G. L., and Meng, X.-L. (eds.). *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL, 2011. ISBN 978-1-4200-7941-8. doi: 10.1201/b10905. URL <https://doi.org/10.1201/b10905>.
- Deligiannidis, G., Paulin, D., and Doucet, A. Randomized Hamiltonian Monte Carlo as Scaling Limit of the Bouncy Particle Sampler and Dimension-Free Convergence Rates. *preprint arXiv:1808.04299*, aug 2018. URL <http://arxiv.org/abs/1808.04299>.
- Diaconis, P., Holmes, S., and Neal, R. M. Analysis of a non-reversible Markov chain sampler. *Ann. Appl. Probab.*, 10(3):726–752, 2000. ISSN 1050-5164. doi: 10.1214/aoap/1019487508. URL <https://doi.org/10.1214/aoap/1019487508>.
- Fearnhead, P., Bierkens, J., Pollock, M., and Roberts, G. O. Piecewise Deterministic Markov Processes for Continuous-Time Monte Carlo. *Statist. Sci.*, 33(3):386–412, 2018. doi: 10.1214/18-STS648. URL <https://doi.org/10.1214/18-STS648>.
- Lorenzi, L. and Bertoldi, M. *Analytical methods for Markov semigroups*, volume 283 of *Pure and Applied Mathematics (Boca Raton)*. Chapman & Hall/CRC, Boca Raton, FL, 2007. ISBN 978-1-58488-659-4; 1-58488-659-5.
- Markovic, J. and Sepehri, A. Bouncy hybrid sampler as a unifying device. *preprint arXiv:1802.04366*, 2018. URL <http://arxiv.org/abs/1802.04366>.
- Murray, I., Adams, R., and MacKay, D. Elliptical slice sampling. In Teh, Y. W. and Titterton, M. (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 541–548, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL <http://proceedings.mlr.press/v9/murray10a.html>.

Neal, R. M. Regression and classification using Gaussian process priors. In *Bayesian statistics, 6 (Alcoceber, 1998)*, pp. 475–501. Oxford Univ. Press, New York, 1999.

Roberts, G. O. and Stramer, O. On inference for partially observed nonlinear diffusion models using the metropolis-hastings algorithm. *Biometrika*, 88(3):603–621, 2001. ISSN 00063444. URL <http://www.jstor.org/stable/2673434>.

Terenin, A. and Thorngren, D. A Piecewise Deterministic Markov Process via (r, θ) swaps in hyperspherical coordinates. *preprint arXiv:1807.00420*, 2018. URL <http://arxiv.org/abs/1807.00420>.

Vanetti, P., Bouchard-Côté, A., Deligiannidis, G., and Doucet, A. Piecewise-Deterministic Markov Chain Monte Carlo. *preprint arXiv:1707.05296*, 2017. URL <http://arxiv.org/abs/1707.05296>.