



Geometry-informed irreversible perturbations for accelerated convergence of Langevin dynamics

Benjamin J. Zhang¹ · Youssef M. Marzouk¹ · Konstantinos Spiliopoulos²

Received: 24 August 2021 / Accepted: 3 September 2022 / Published online: 19 September 2022
© The Author(s) 2022

Abstract

We introduce a novel geometry-informed irreversible perturbation that accelerates convergence of the Langevin algorithm for Bayesian computation. It is well documented that there exist perturbations to the Langevin dynamics that preserve its invariant measure while accelerating its convergence. Irreversible perturbations and reversible perturbations (such as Riemannian manifold Langevin dynamics (RMLD)) have separately been shown to improve the performance of Langevin samplers. We consider these two perturbations simultaneously by presenting a novel form of irreversible perturbation for RMLD that is informed by the underlying geometry. Through numerical examples, we show that this new irreversible perturbation can improve estimation performance over irreversible perturbations that do not take the geometry into account. Moreover we demonstrate that irreversible perturbations generally can be implemented in conjunction with the stochastic gradient version of the Langevin algorithm. Lastly, while continuous-time irreversible perturbations cannot impair the performance of a Langevin estimator, the situation can sometimes be more complicated when discretization is considered. To this end, we describe a discrete-time example in which irreversibility increases both the bias and variance of the resulting estimator.

Keywords Monte Carlo sampling · Stochastic gradient Langevin dynamics · Riemannian manifold Langevin dynamics · Geometry-informed irreversibility · Bayesian computation

1 Introduction

Bayesian inference often requires estimating expectations with respect to non-Gaussian distributions. To solve this problem, particularly in high dimensions, one frequently resorts to sampling methods. A commonly used class of sampling methods is based on the Langevin dynamics (LD), which uses the gradient of the log-target density to specify a stochastic differential equation (SDE) whose invariant distribution is the target (e.g., posterior) distribution of interest.

BJZ and YMM acknowledge support from the Air Force Office of Scientific Research, Analysis and Synthesis of Rare Events (ANSRE) MURI. KS was partially supported by the National Science Foundation (DMS 1550918, DMS 2107856) and Simons Foundation Award 672441.

✉ Benjamin J. Zhang
bjz@mit.edu

¹ Department of Aeronautics and Astronautics, Center for Computational Science and Engineering, Massachusetts Institute of Technology, Cambridge, USA

² Department of Mathematics and Statistics, Boston University, Boston, USA

Long term averages over a single trajectory of the SDE can be then used to estimate expectations of interest by appealing to the ergodicity of the stochastic process. Other LD-based approaches that reduce the mean squared error (MSE) of such estimators include the Metropolis-adjusted Langevin algorithm (MALA) (Roberts and Tweedie 1996; Girolami and Calderhead 2011), the stochastic gradient Langevin dynamics (SGLD) (Welling and Teh 2011), and their variants.

It is also known that certain perturbations to the LD can accelerate convergence of the dynamics to the stationary distribution. In Rey-Bellet and Spiliopoulos (2015) the authors show that suitable reversible and irreversible perturbations to diffusion processes can decrease the spectral gap of the generator, as well as increase the large deviations rate function and decrease the asymptotic variance of the estimators. One widely celebrated choice of *reversible* perturbation is the Riemannian manifold Langevin dynamics (Girolami and Calderhead 2011), in which one defines a Riemannian metric to alter the way distances and gradients are computed. The use of *irreversible* perturbations to accelerate convergence has also been well studied in a variety of contexts and general settings (Rey-Bellet and Spiliopoulos 2015; Hwang et al.

2005; Rey-Bellet and Spiliopoulos 2015, 2016); see also (Franke et al. 2010; Hwang et al. 1993; Bierkens 2016; Diaconis et al. 2010) and for linear systems, Lelievre et al. (2013). The authors of Ma et al. (2015) find general conditions on the drift and diffusion coefficients of an SDE so that a specified measure is the SDE's invariant measure—without, however, exploring how different choices of these coefficients impact sampling quality. Existing literature shows that augmenting the drift of the LD with a vector field that is orthogonal to the gradient of the log-target density will leave the invariant measure unchanged while decreasing the spectral gap. A convenient choice is simply to add the vector field induced by a skew-symmetric matrix applied to the gradient of the log posterior.

At the same time, traditional sampling methods for Bayesian inference are often intractable for extremely large datasets. While Langevin dynamics-based sampling methods only require access to the unnormalized posterior density, they need many evaluations of this unnormalized density and its gradient. When the dataset is extremely large, each evaluation of the density may be computationally intractable, as it requires the evaluation of the likelihood over the entire dataset. In the past decade the *stochastic gradient* Langevin dynamics (SGLD) has been introduced and analyzed (Welling and Teh 2011; Teh et al. 2016) to address the problem posed by large datasets. Rather than evaluating the likelihood over the entire dataset, SGLD subsamples a portion of the data (either with or without replacement) and uses the likelihood evaluated at the sampled data to estimate the true likelihood. The resulting chain can then be used to estimate ergodic averages.

In this paper we present a *state-dependent irreversible* perturbation of Riemannian manifold Langevin dynamics that is informed by the *geometry* of the manifold. This departs from existing literature, as the vector field of the resulting perturbation is *not* orthogonal to the original drift term. This geometry-informed irreversible perturbation accelerates convergence and, if desired, can be used in combination with the SGLD algorithm to exploit the computational savings of a stochastic gradient.

We demonstrate this approach on a variety of examples: a simple anisotropic Gaussian target, a posterior on the mean and variance parameters of a normal distribution, Bayesian logistic regression, and Bayesian independent component analysis (ICA). Generally, we observe that the geometry-informed irreversible perturbation improves the convergence rate of LD compared to a standard irreversible perturbation. The improvement tends to be more pronounced as the target distribution deviates from Gaussianity. Our numerical studies also show that introducing irreversibility can reduce the MSE of the resulting long-term average estimator, mainly by reducing variance. In many cases this reduction can be significant, e.g., 1–2 orders of magnitude.

One must, however, also take the effects of discretization into account. In the continuous-time setting, it is known theoretically that irreversible perturbations can at worst only leave the spectral gap fixed. In borderline cases, though—i.e., in cases where the continuous-time theoretical improvement is nearly zero—after accounting for discretization, stiffness can actually cause the resulting estimator to perform worse than if no irreversibility were applied at all. Indeed, we will describe in Appendix A an illustrative Gaussian example in which the standard Langevin algorithm performs better than the algorithm with the standard irreversible perturbation—that is, an example in which additional irreversibility leads to increased bias and variance of the long term average estimator (see Remark 3.2 and Remark A.1 for a theoretical explanation). Along similar lines, the idea of applying irreversible perturbations to SGLD has recently been studied in the context of nonconvex stochastic optimization (Hu et al. 2020). The authors also note that while irreversibility increases the rate of convergence, it increases the discretization error and amplifies the variance of the gradient, compared to a non-perturbed system with the same step size; see also (Brosse et al. 2018) for a related discussion on the relation of SGLD to SGD and convergence properties. This reflects the increased stiffness of irreversible SGLD relative to standard SGLD.

The rest of the paper is organized as follows. In Section 2 we review reversible and irreversible perturbations of the overdamped Langevin dynamics that may improve the efficiency of sampling from equilibrium. Then, in Section 2.3, we present our new geometry-informed irreversible perturbation. In Section 3 we present simulation studies that demonstrate the good performance of this geometric perturbation, relative to a variety of other standard reversible and irreversible choices. In several of these examples, we also demonstrate the use of stochastic gradients. Section 4 summarizes our results and outlines directions for future work. Appendix A details the simple Gaussian example showing that in “borderline” cases—i.e., when continuous-time analysis does not predict improvements from irreversible perturbations—the stiffness created by an irreversible perturbation can, after discretization, lead to poorer performance than the unperturbed case.

2 Improving the performance of Langevin samplers

We begin by recalling some relevant background on Langevin samplers, Riemannian manifold Langevin dynamics, perturbations of Langevin dynamics, and the stochastic gradient Langevin dynamics algorithm. Let $f(\theta)$ be a test function on state space $E \subset \mathbb{R}^d$ and let $\pi(\theta)$ be some unnormalized target density on E . In our experiments, $\pi(\theta)$ arises as a posterior density of the form $\pi(\theta) \propto L(\theta; X)\pi_0(\theta)$, where

$L(\theta; X)$ is the likelihood model, X are the data, and $\pi_0(\theta)$ is the prior density. Define $\{\theta(t)\}$ as a Langevin process that has invariant density $\pi(\theta)$:

$$d\theta(t) = \beta \nabla \log \pi(\theta(t)) dt + \sqrt{2\beta} dW(t), \quad (1)$$

where $\beta > 0$ denotes the temperature, $W(t)$ is a standard Brownian motion in $\mathbb{R}^d$, and the initial condition may be arbitrary. By ergodicity, we may compute expectations with respect to the posterior by the long term average of $f(\theta)$ over a single trajectory:

$$\mathbb{E}_\pi[f(\theta)] = \int_E f(\theta) \pi(\theta) d\theta = \lim_{t \rightarrow \infty} \frac{1}{T} \int_0^T f(\theta(t)) dt. \quad (2)$$

For practical computations, we must approximate (2) by discretizing the Langevin dynamics and choosing a large but finite T . Applying the Euler-Maruyama method to (1) with step size h yields the following recurrence relation,

$$\theta_{k+1} = \theta_k + h\beta \nabla \log \pi(\theta_k) dt + \sqrt{2\beta h} \xi_{k+1} \quad (3)$$

where ξ_k are independent standard normal random variables. The total number of steps is equal to $K = T/h$. The resulting estimator for (2) is

$$\mathbb{E}_\pi[f(\theta)] \approx \frac{1}{K} \sum_{k=0}^{K-1} f(\theta_k). \quad (4)$$

This estimator is the *unadjusted Langevin algorithm* (ULA), which has found renewed interest in the context of high-dimensional machine learning problems (Durmus and Moulines 2019). Discretization and truncation, however, introduce bias into the estimator. Moreover, there are noted examples in which the continuous-time process and the discretized version do not have the same invariant distribution no matter the choice of the fixed, but nonzero, discretization step h ; see Ganguly and Sundar (2021) for a related discussion. Certain Markov chain Monte Carlo (MCMC) methods such as MALA circumvent these issues by using the dynamics to propose new points, but accepting or rejecting them according to some rule so that the resulting discrete-time Markov chain has the target distribution as its invariant distribution (Roberts and Tweedie 1996; Girolami and Calderhead 2011).

Many different SDEs can have the same invariant distribution. Therefore, there has been much study into how the standard Langevin dynamics of some target distribution can be altered to increase its rate of convergence. Some examples of this can be found in the work of Hwang et al. (2005), Rey-Bellet and Spiliopoulos (2016) and others. The standard Langevin dynamics is a reversible Markov process, meaning that the process satisfies detailed balance. The work of Rey-Bellet and Spiliopoulos (2016) studies, in general terms, how

reversible and irreversible perturbations to reversible processes decrease the spectral gap, increase the large deviations rate function, and decrease the asymptotic variance. Yet how to *choose* such perturbations to most efficiently accelerate convergence is yet to be thoroughly studied in settings beyond linear diffusion processes (Lelievre et al. 2013). Also, with the exception of a few examples—see for instance (Duncun et al. 2017; Jianfeng and Spiliopoulos 2018)—these perturbations have mainly been studied in the continuous-time setting.

2.1 Reversible perturbations and Riemannian manifold Langevin dynamics

We only review relevant aspects of reversible perturbations and RMLD in this section. For a detailed review of RMLD and its related Monte Carlo methods, we refer the reader to Girolami and Calderhead (2011), Livingstone and Girolami (2014), Xifara et al. (2014). Let $\mathbf{B}(\theta)$ be a $d \times d$ symmetric positive definite matrix. A reversible perturbation on LD (1) is an SDE with multiplicative noise:

$$d\theta(t) = \beta [\mathbf{B}(\theta) \nabla \log \pi(\theta(t)) + \nabla \cdot \mathbf{B}(\theta)] dt + \sqrt{2\beta \mathbf{B}(\theta)} dW(t). \quad (5)$$

Here, the i -th component of $\nabla \cdot \mathbf{B}(\theta)$ is $\sum_{j=1}^d \partial_{\theta_j} \mathbf{B}_{ij}(\theta)$. This is equivalent to Langevin dynamics defined on a Riemannian manifold, where the metric is $\mathbf{G}(\theta) = \mathbf{B}(\theta)^{-1}$ (Xifara et al. 2014). A straightforward calculation shows that (5) with $\mathbf{B}(\theta)$ being any symmetric positive-definite matrix admits the same invariant distribution, π . The improved rate of convergence depends on the choice of the underlying metric. The work of Girolami and Calderhead (2011) argues that choosing the expected Fisher information matrix plus the Hessian of the log-prior to be the metric improves the performance of the resulting manifold MALA method. Meanwhile, Rey-Bellet and Spiliopoulos (2016) shows that under certain regularity conditions, if $\mathbf{B}(\theta)$ is chosen such that $\mathbf{B}(\theta) - \mathbf{I}$ is positive definite, then the resulting estimator is expected to have improved performance in terms of the asymptotic variance, the spectral gap, and the large deviations rate function.

2.2 Irreversible perturbations

Consider the following Langevin dynamics

$$d\theta(t) = [\beta \nabla \log \pi(\theta(t)) + \gamma(\theta(t))] dt + \sqrt{2\beta} dW(t). \quad (6)$$

When $\gamma(\theta) \equiv 0$, the process is reversible and has $\pi(\theta)$ as its invariant distribution. If $\gamma \neq 0$, then the resulting process will, in general, be time-irreversible unless $\gamma(\theta)$ can be written as a multiple of $\nabla \log \pi(\theta)$; see for example (Pavliotis 2014). However, an irreversible perturbation can still

preserve the invariant distribution of the system. By considering the Fokker-Planck equation, one can show that if $\gamma(\theta)$ is chosen such that $\nabla \cdot (\gamma\pi) = 0$, then π will still be the invariant distribution. A frequently used choice in the literature is $\gamma(\theta) = \mathbf{J}\nabla \log \pi(\theta)$, where $\mathbf{J}$ is a constant skew-symmetric matrix, i.e., $\mathbf{J} = -\mathbf{J}^T$. The computational advantage of this choice is clear since only one additional matrix-vector multiply is needed to implement it. The optimal choice of irreversible perturbation to a linear system (i.e., that which yields fastest convergence) was completely analyzed in Lelievre et al. (2013).

The advantages of using irreversible perturbations is widely noted. The main result of Hwang et al. (2005) is that under certain conditions, the spectral gap, i.e., the difference between the leading two eigenvalues of the generator of the Markov semigroup, increases when $\gamma \neq 0$. In Rey-Bellet and Spiliopoulos (2015, 2016), the large deviations rate function is introduced as a measure of performance in the context of sampling from the equilibrium, and upon connecting it to the asymptotic variance of the long term average estimator, it is proven that adding an appropriate perturbation γ not only increases the large deviations rate function but also decreases the asymptotic variance of the estimator. The use of irreversible proposals in the MALA was studied in Ottobre et al. (2019).

2.3 Irreversible perturbations for RMLD

In this section, we will introduce our novel geometry-informed irreversible perturbation to Langevin dynamics. Suppose that we are given a diffusion process as in (5), and we want to study how to choose an irreversible perturbation that leaves the invariant distribution fixed. Indeed, our previous choice of irreversible perturbation remains valid for this system, that is, adding $\gamma(\theta) = \mathbf{J}\nabla \log \pi(\theta)$ for a constant skew-symmetric matrix $\mathbf{J}$ to the drift term of (5) will preserve the invariant density. This choice yields the following SDE:

$$d\theta(t) = \left[(\beta \mathbf{B}(\theta(t)) + \mathbf{J})\nabla \log \pi(\theta(t)) + \beta \nabla \cdot \mathbf{B}(\theta(t)) \right] dt + \sqrt{2\beta \mathbf{B}(\theta)} dW(t) \quad (7)$$

We refer to this system as Riemannian manifold Langevin with an additive irreversible perturbation (RMirr). This choice, however, does not take into account the relevant features that the reversible perturbation may provide when constructing an irreversible perturbation.

The reversible perturbation leads to a positive definite matrix (a metric, in the terminology of Riemannian geometry) that is state-dependent. In contrast, the skew-symmetric matrix $\mathbf{J}$ is fixed in the irreversible perturbation. The skew-symmetric matrix need not be constant, however, as an

irreversible perturbation $\gamma(\theta)$ only needs to satisfy $\nabla \cdot (\gamma(\theta)\pi(\theta)) = 0$. In fact, if $\gamma(\theta) = \mathbf{C}(\theta)\nabla \log \pi(\theta) + \nabla \cdot \mathbf{C}(\theta)$ for $\mathbf{C}(\theta) = -\mathbf{C}(\theta)^T$, then this irreversible perturbation will also leave the invariant density intact. Noting that $\mathbf{C}_{ii}(\theta) = 0$ and that $\mathbf{C}_{ij} = -\mathbf{C}_{ji}$, observe that

$$\begin{aligned} \nabla \cdot (\gamma(\theta)\pi(\theta)) &= \nabla \cdot (\mathbf{C}(\theta)\nabla \pi(\theta) + (\nabla \cdot \mathbf{C}(\theta))\pi(\theta)) \\ &= \sum_{i,j=1}^d \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_i} \frac{\partial \pi(\theta)}{\partial \theta_j} + \mathbf{C}_{ij}(\theta) \frac{\partial^2 \pi(\theta)}{\partial \theta_i \partial \theta_j} \\ &\quad + \frac{\partial^2 \mathbf{C}_{ij}(\theta)}{\partial \theta_i \partial \theta_j} \pi(\theta) + \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_j} \frac{\partial \pi(\theta)}{\partial \theta_i} \\ &= \sum_{i>j,i=1}^d \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_i} \frac{\partial \pi(\theta)}{\partial \theta_j} + \mathbf{C}_{ij}(\theta) \frac{\partial^2 \pi(\theta)}{\partial \theta_i \partial \theta_j} \\ &\quad + \frac{\partial^2 \mathbf{C}_{ij}(\theta)}{\partial \theta_i \partial \theta_j} \pi(\theta) + \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_j} \frac{\partial \pi(\theta)}{\partial \theta_i} \\ &\quad + \frac{\partial \mathbf{C}_{ji}(\theta)}{\partial \theta_j} \frac{\partial \pi(\theta)}{\partial \theta_i} + \mathbf{C}_{ji}(\theta) \frac{\partial^2 \pi(\theta)}{\partial \theta_j \partial \theta_i} \\ &\quad + \frac{\partial^2 \mathbf{C}_{ji}(\theta)}{\partial \theta_j \partial \theta_i} \pi(\theta) + \frac{\partial \mathbf{C}_{ji}(\theta)}{\partial \theta_i} \frac{\partial \pi(\theta)}{\partial \theta_j} \\ &= \sum_{i>j,i=1}^d \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_i} \frac{\partial \pi(\theta)}{\partial \theta_j} + \mathbf{C}_{ij}(\theta) \frac{\partial^2 \pi(\theta)}{\partial \theta_i \partial \theta_j} \\ &\quad + \frac{\partial^2 \mathbf{C}_{ij}(\theta)}{\partial \theta_i \partial \theta_j} \pi(\theta) + \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_j} \frac{\partial \pi(\theta)}{\partial \theta_i} \\ &\quad - \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_j} \frac{\partial \pi(\theta)}{\partial \theta_i} - \mathbf{C}_{ij}(\theta) \frac{\partial^2 \pi(\theta)}{\partial \theta_j \partial \theta_i} \\ &\quad - \frac{\partial^2 \mathbf{C}_{ij}(\theta)}{\partial \theta_j \partial \theta_i} \pi(\theta) - \frac{\partial \mathbf{C}_{ij}(\theta)}{\partial \theta_i} \frac{\partial \pi(\theta)}{\partial \theta_j} = 0. \end{aligned}$$

We seek an irreversible perturbation that takes the reversible perturbation into account, with the possibility that $\mathbf{C}(\theta)$ is not a constant matrix, and investigate if it leads to any performance improvements of the long term average estimator. Note that in the literature, the above condition $\nabla \cdot (\gamma\pi) = 0$ is typically rewritten into the following sufficient conditions: $\nabla \cdot \gamma(\theta) = 0$ and $\gamma(\theta) \cdot \nabla \pi(\theta) = 0$ (Rey-Bellet and Spiliopoulos 2016). One can check, however, that when $\mathbf{C}$ is not constant, these conditions are not met, yet $\gamma(\theta)$ is still a valid irreversible perturbation. A simple choice of $\mathbf{C}(\theta)$ that incorporates $\mathbf{B}(\theta)$ is

$$\mathbf{C}(\theta) = \frac{1}{2} \mathbf{J} \mathbf{B}(\theta) + \frac{1}{2} \mathbf{B}(\theta) \mathbf{J}, \quad (8)$$

where $\mathbf{J}$ is a constant skew-symmetric matrix. The $\frac{1}{2}$ factor is introduced so that if $\mathbf{B}(\theta) = \mathbf{I}$, i.e., if there is no reversible perturbation, then this perturbation reverts to the standard irreversible perturbation (Irr). We arrive at the following system:

$$\begin{aligned} d\theta(t) = & [(\beta \mathbf{B}(\theta(t)) + \mathbf{C}(\theta(t))) \nabla \log \pi(\theta(t)) \\ & + \nabla \cdot (\beta \mathbf{B}(\theta(t)) + \mathbf{C}(\theta(t)))] dt + \sqrt{2\beta \mathbf{B}(\theta(t))} dW(t). \end{aligned} \quad (9)$$

We call this choice of perturbation the *geometry-informed irreversible perturbation* (GiIrr). Indeed, while there are infinitely many valid choices for $\mathbf{C}(\theta)$, we will investigate the choice in (8) in the numerical examples. Since we will have already explicitly constructed $\mathbf{B}(\theta)$ and $\mathbf{J}$ for the other systems, the additional computational cost of computing their product will be marginal. Furthermore, as mentioned earlier, this choice reduces to Irr when $\mathbf{B}(\theta) = \mathbf{I}$.

One may wonder when does GiIrr result in improved performance over standard irreversible perturbations such as in Equation 7. Based on the numerical results and intuition, we will argue that GiIrr results in better performance if the underlying reversible perturbation already improves the sampling. Namely, if one knows that RMLD leads to improved sampling on a given problem, then employing GiIrr is expected to improve sampling even further. As we mentioned earlier, the choice of GiIrr that is made in this paper is not unique, and a further investigation of its theoretical properties is left for future work; see also the discussion in the Section 4. The goal of this paper is to present this new class of irreversible perturbations and investigate it numerically in a number of representative computational studies.

2.4 Stochastic gradient Langevin dynamics

In certain Bayesian inference problems, the data are conditionally independent of each other given the parameter value. Therefore, the likelihood model can often be factorized and the posterior density can be written as follows:

$$\pi(\theta) \propto \pi_0(\theta) \prod_{i=1}^N \pi_i(X_i|\theta) \quad (10)$$

where $\pi(X_i|\theta)$ is the likelihood function for data point X_i . When the dataset is extremely large, i.e., when $N \gg 1$, however, ULA becomes exceedingly expensive as it requires repeatedly evaluating the likelihood over the entire dataset for each step of the trajectory. To mitigate this challenge, the stochastic gradient Langevin dynamics was presented to reduce the computational cost of evaluating the posterior density by only evaluating the likelihood over *subsets* of the data at each step. The true likelihood is estimated based on the likelihood function evaluated at the subsampled data (Welling and Teh 2011). Specifically, the gradient is estimated using a stochastic gradient

$$\begin{aligned} \nabla \log \pi(\theta|X) & \approx \nabla \log \widehat{\pi}(\theta|X) \\ & = \log \pi_0(\theta) + \frac{N}{n} \sum_{i=1}^n \log \pi(X_{\tau_i}|\theta) \end{aligned} \quad (11)$$

where τ is a random subset of $\{1, \dots, N\}$ of size n drawn with or without replacement. Depending on the choice of n , this approach cuts down on the computational costs dramatically with some additional variance incurred by the random sub-sampling of the data. The original version of this algorithm made the step size variable, approaching zero as the number of steps taken K became large. SGLD applied with a variable and shrinking step size was proven to be consistent: that is, the invariant distribution of the discretized system is equivalent to that of the continuous system (Teh et al. 2016). Having a decreasing step size counteracts the cost savings provided by computing the stochastic gradient, and therefore a version where the step size is fixed was presented in Vollmer et al. (2016), where theoretical characterizations of the asymptotic and finite-time bias and variance are also developed. In most of our numerical results, we use stochastic gradient version of the Langevin algorithm with fixed step size to demonstrate that SGLD can be used together with irreversible perturbations.

3 Numerical examples

In the following examples, we always apply the stochastic gradient version of each Langevin system unless otherwise stated. We fix $\beta = 1/2$ for all examples. The efficacy of the GiIrr perturbation does not change whether or not the stochastic gradient is used. We illustrate this explicitly in Section 3.4, where we report the results of all perturbations both with and without the stochastic gradient, for comparison.

3.1 Evaluating sample quality

Here we discuss the measures we use to evaluate sample quality for each Langevin sampler. For each example, we estimate the bias, variance, mean-squared error, and asymptotic variance of the estimators of the expectations of two observables: $\phi_1(\theta) = \sum_{l=1}^d \theta^{(l)}$ and $\phi_2(\theta) = \sum_{l=1}^d |\theta^{(l)}|^2$, where $\theta^{(l)}$ denotes the l th component of θ . Let $\bar{\phi}^K = \frac{1}{K} \sum_{k=0}^{K-1} \phi(\theta_k)$ denote the estimator of $\mathbb{E}_\pi[\phi(\theta)]$ obtained with a chain of length K .

$$\begin{aligned} \text{Bias}(\bar{\phi}^K) & = \mathbb{E}[\bar{\phi}^K] - \mathbb{E}_\pi[\phi(\theta)] \\ & = \mathbb{E}\left[\frac{1}{K} \sum_{k=0}^{K-1} \phi(\theta_k)\right] - \mathbb{E}_\pi[\phi(\theta)] \end{aligned}$$

$$\approx \frac{1}{M} \sum_{i=1}^M \left[\frac{1}{K} \sum_{k=0}^{K-1} \phi([\theta_k]_i) \right] - \mathbb{E}_{\pi} [\phi(\theta)], \quad (12)$$

where $[\theta_k]_i$ is the state of the i^{th} chain at iteration k . Here, we estimate $\mathbb{E}_{\pi} [\phi(\theta)]$ by applying the unadjusted Langevin algorithm with a very long simulated trajectory and small discretization step. The expected value of the estimator is computed by averaging over $M = 1000$ independent chains for the examples in Sections 3.2 and 3.3, and $M = 100$ independent chains for the examples in Sections 3.4 and 3.5. The variance of each estimator is defined and estimated as follows

$$\begin{aligned} \text{Var}(\bar{\phi}^K) &= \mathbb{E} \left[(\bar{\phi}^K)^2 \right] - \left(\mathbb{E} [\bar{\phi}^K] \right)^2 \\ &\approx \frac{1}{M} \sum_{i=1}^M \left(\frac{1}{K} \sum_{k=0}^{K-1} \phi([\theta_k]_i) \right)^2 \\ &\quad - \left(\frac{1}{M} \sum_{i=1}^M \left[\frac{1}{K} \sum_{k=0}^{K-1} \phi([\theta_k]_i) \right] \right)^2. \end{aligned} \quad (13)$$

The mean-squared error (MSE) of each estimator $\bar{\phi}^K$ follows analogously.

We also evaluate the asymptotic variance of the estimator of each observable, defined as

$$\sigma^2(\phi) = \lim_{t \rightarrow \infty} t \text{Var} \left(\frac{1}{t} \int_0^t \phi(\theta_t) dt \right) \approx \lim_{K \rightarrow \infty} K h \text{Var}(\bar{\phi}^K). \quad (14)$$

To compute these asymptotic variances, we use the batch means method in Asmussen and Glynn (2007). After the burn-in period, we evaluate the observable over each chain. Each observable chain is then batched into twenty separate chains, and their means are evaluated. The asymptotic variance is estimated by computing the empirical variance of those means and then multiplying by the length of each of the subsampled trajectories.

In addition to measuring the performance of estimators of specific observables, for each sampler we also evaluate overall sample quality by computing the recently-proposed kernelized Stein discrepancy (KSD) (Gorham and Mackey 2017). The KSD is a computable expression that can approximate certain integral probability metrics (IPMs) for a certain class of functions defined through the action of the Stein operator on a reproducing kernel Hilbert space. Let $\hat{\pi}_K = \frac{1}{K} \sum_{k=0}^{K-1} \delta_{\theta_k}$ be the empirical approximation to π based on samples $\{\theta_k\}_{k=0}^{K-1}$ produced by some Langevin algorithm. The IPM is defined as

$$d_{\mathcal{H}}(\hat{\pi}_K, \pi) := \sup_{h \in \mathcal{H}} |\mathbb{E}_{\hat{\pi}_K} [h(Z)] - \mathbb{E}_{\pi} [h(X)]|, \quad (15)$$

for some function space $\mathcal{H}$, where $h \in \mathcal{H}$ are functions from $\mathbb{R}^d$ to $\mathbb{R}$, and $Z \sim \hat{\pi}_K$ and $X \sim \pi$ are random variables. When $\mathcal{H}$ is large enough, $d_{\mathcal{H}}(\hat{\pi}_K, \pi) \rightarrow 0$ holds only if $\hat{\pi}_K \rightarrow \pi$ in distribution; see Gorham and Mackey (2017) and the references therein. This result implies that a better empirical approximation (i.e., better-quality samples) corresponds with a lower IPM value. Practically computing the IPM directly is intractable as it requires exactly knowing expectations with respect to the target distribution π (which is, after all, our original goal in sampling). By judiciously choosing the function space $\mathcal{H}$, however, we can estimate the IPM by computing the KSD, using only the samples that comprise the empirical distribution $\hat{\pi}$.

As the expectation with respect to the target distribution π is intractable to compute, one seeks an $\mathcal{H}$ such that for all $h \in \mathcal{H}$, $\mathbb{E}_{\pi} [h(X)] = 0$. Such a function space is found through *Stein's identity*, which states that $\mathbb{E}_{\pi} [\mathcal{A}g(X)] = 0$ for all $g : \mathbb{R}^d \rightarrow \mathbb{R}$ in some function space $\mathcal{G}$, where $\mathcal{A}$ is the Stein operator,

$$\mathcal{A}g(x) = \frac{1}{\pi(x)} \nabla \cdot (\pi(x)g(x)) = g(x) \cdot \nabla \log \pi(x) + \nabla \cdot g(x),$$

and $\nabla \cdot$ denotes the divergence operator. The space $\mathcal{G}$ is defined by specifying a reproducing kernel Hilbert space (RKHS) $\mathcal{R}_r$ of functions from $\mathbb{R}^d$ to $\mathbb{R}$, with a user-defined kernel $r(x, y)$ that is twice-continuously differentiable (Gorham and Mackey 2017; Izzatullah et al. 2020). Each function $g \in \mathcal{G}$ is made up of components $g_j \in \mathcal{R}_r$, for $j = 1, \dots, d$, such that the vector $(\|g_1\|_{\mathcal{R}_r}, \dots, \|g_d\|_{\mathcal{R}_r})$ has unit norm in the dual space ℓ^2 , where $\|\cdot\|_{\mathcal{R}_r}$ is the norm of the RKHS (Gorham and Mackey 2017). Hence, if we set $h = \mathcal{A}g$ for $g \in \mathcal{G}$, then $\mathbb{E}_{\pi} [h(X)] = 0$. Proposition 1 in Gorham and Mackey (2017) demonstrates that such a $\mathcal{G}$ is an appropriate domain for $\mathcal{A}$ so that one indeed has $\mathbb{E}_{\pi} [\mathcal{A}g(X)] = 0$ for all $g \in \mathcal{G}$. Then the space $\mathcal{H}$ is defined as $\mathcal{H} = \mathcal{A}\mathcal{G}$, i.e., the space of functions resulting from the Stein operator applied to functions in $\mathcal{G}$. With this $\mathcal{H}$, we define the corresponding KSD of a measure μ as $S(\mu) = d_{\mathcal{H}}(\mu, \pi) = \sup_{h \in \mathcal{H}} |\mathbb{E}_{\mu} [h(X)]|$.

Proposition 2 in Gorham and Mackey (2017) shows that the KSD admits a closed form. Indeed, for any $j = 1, \dots, d$ and letting $b_j(x) = \partial_{x_j} \log \pi(x)$, define the Stein kernel

$$\begin{aligned} r_0^j(x, y) &= b_j(x)b_j(y)r(x, y) + b_j(x)\nabla_{y_j}r(x, y) \\ &\quad + b_j(y)\nabla_{x_j}r(x, y) + \nabla_{x_j}\nabla_{y_j}r(x, y). \end{aligned} \quad (16)$$

Then, by Proposition 2 of Gorham and Mackey (2017), if $\sum_{j=1}^d \mathbb{E}_{\mu} \left[\sqrt{r_0^j(X, X)} \right] < \infty$, we have the equality $S(\mu) = \|w\|_2$, where for $j = 1, \dots, d$ we have

$$w_j = \sqrt{\mathbb{E}_{\mu \times \mu} [r_0^j(X, \tilde{X})]}, \quad X, \tilde{X} \stackrel{\text{i.i.d.}}{\sim} \mu.$$

Also, Gorham and Mackey (2017) emphasizes the importance of choosing the kernel $r(x, y)$ carefully so that convergence of the KSD to zero (for a sequence of empirical distributions) implies convergence in distribution to the target measure. As suggested by the results of Gorham and Mackey (2017), we use the inverse multiquadric kernel $r(x, y) = (c^2 + \|x - y\|_2^2)^\beta$ with $\beta = -1/2$ and $c = 1$. This choice has also been used in Izzatullah et al. (2020).

In our examples, $\mu = \hat{\pi}_K$ is a discrete distribution, so the KSD can be computed by evaluating the kernels over all pairs of sample points (Gorham and Mackey 2017). Namely, if $\mu = \hat{\pi}_K = \frac{1}{K} \sum_{k=0}^{K-1} \delta_{\theta_k}$, then

$$w_j = \sqrt{\sum_{k,k'=1}^K r_0^j(x_k, x_{k'})}. \quad (17)$$

We use the KSD to evaluate the quality of every Langevin sampler—effectively, sampler performance over a large class of observables—to complement our bias/variance/MSE computations for particular choices of observable. For each example below and for each Langevin sampler, we compute the KSD for 25 independent chains. The mean KSDs of the chains for each sampler are then plotted as a function of the number of steps of the chain. We comment on the quality of each sampler given these KSD plots. We refer the reader to Gorham et al. (2019), Gorham and Mackey (2015, 2017), Izzatullah et al. (2020) for further details on these choices and for further literature review.

3.2 Linear Gaussian example

Suppose we have data $\{X_i\}_{i=1}^N \subset \mathbb{R}^d$ generated from a multivariate normal distribution with mean $\theta \in \mathbb{R}^d$ and known precision matrix $\Gamma_X \in \mathbb{R}^{d \times d}$. From the data, we infer the value of θ . Endow θ with a normal prior with mean zero and precision $\Gamma_\theta \in \mathbb{R}^{d \times d}$. Then the posterior distribution is Gaussian with mean and precision

$$\mu_p = (\Gamma_\theta + N\Gamma_X)^{-1} \Gamma_X \sum_{i=1}^N X_i \quad \text{and} \quad \Gamma_p = (\Gamma_\theta + N\Gamma_X), \quad (18)$$

respectively. The Euler-Maruyama discretization with constant step size h applied to the corresponding Langevin dynamics is

$$\theta_{k+1} = (\mathbf{I} - \bar{\mathbf{A}}h) \theta_k + \bar{\mathbf{D}}_k h + \sqrt{h} \xi_k \quad (19)$$

where

$$\bar{\mathbf{A}} = \frac{1}{2}(\Gamma_\theta + N\Gamma_X), \quad \bar{\mathbf{D}}_k = \frac{1}{2}\Gamma_X \sum_{i=1}^N X_i, \quad \xi_k \sim \mathcal{N}(0, \mathbf{I}).$$

Using stochastic gradients yields the same recurrence above except with

$$\bar{\mathbf{D}}_k = \frac{1}{2}\Gamma_X \frac{N}{n} \sum_{i=1}^n X_{\tau_i^k} \quad (20)$$

where $n \leq N$ and $\tau_i^k \in \{1, \dots, N\}$ is randomly sampled (with or without replacement) (Welling and Teh 2011). Expectations with respect to the posterior are approximated by a long term average of the observable over the course of a trajectory. It has been shown that despite subsampling the data at each step in the dynamics, this estimator has comparable performance as the estimator produced by the regular Langevin dynamics with the full likelihood or MALA (Welling and Teh 2011; Vollmer et al. 2016).

Now, we consider the case where the dynamics are perturbed by an irreversible term that preserves the invariant distribution of the dynamics. We demonstrate that this leads to a lower MSE than standard SGLD or Langevin dynamics. In this case, we replace $\bar{\mathbf{A}}$ and $\bar{\mathbf{D}}_k$ with $\mathbf{A}$ and $\mathbf{D}_k$, which are

$$\mathbf{A} = \frac{1}{2}(\mathbf{I} + \mathbf{J})(\Gamma_\theta + N\Gamma_X), \quad \mathbf{D}_k = \frac{1}{2}(\mathbf{I} + \mathbf{J})\Gamma_X \frac{N}{n} \sum_{i=1}^n X_{\tau_i^k}. \quad (21)$$

and $\mathbf{J}$ is a skew-symmetric matrix.



For the numerical experiments, we choose $d = 3$, $N = 10$, where the mini-batches are of size $n = 2$. The initial condition is the zero vector. We have $\Gamma_X = 0.25\mathbf{I}$, Γ_θ is a precision matrix with eigenvalues 0.2, 0.01, 0.05 and eigenvectors that are randomly generated, and $h = 0.005$. Note that these matrices were chosen so that the resulting reversible perturbation has eigenvalues greater than one. To construct the perturbations, we choose $\mathbf{B} = \Gamma_p^{-1}$ and $\mathbf{J}$ to be

$$\mathbf{J} = \delta \begin{bmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{bmatrix} \quad (22)$$

for $\delta \in \mathbb{R}$. We consider the five different SDE systems presented in Table 1 and investigate how the MSE, bias, and variance differs for each case. For this example, since a constant metric is used, the geometry-informed irreversible perturbation simply produces a different constant skew-symmetric matrix than the other irreversible perturbations. Each system is simulated for $K = 10^5$ steps with step size $h = 5 \times 10^{-3}$. In Figures 1 and 2, we plot the MSE of the running average for each case when the observables are

Table 1 Summary of the five SDEs that share the same invariant density $\pi(\theta)$. Stochastic gradients can be considered instead of the deterministic gradients. All systems are of the form $d\theta_t = b(\theta_t)dt + \sigma(\theta_t)dW_t$. The term β denotes the temperature

	$b(\theta)$	$\sigma(\theta)$
LD	$\beta \nabla \log \pi(\theta)$	$\sqrt{2\beta \mathbf{I}}$
RM	$\beta \mathbf{B}(\theta) \nabla \log \pi(\theta) + \beta \nabla \cdot \mathbf{B}(\theta)$	$\sqrt{2\beta \mathbf{B}(\theta)}$
Irr	$(\beta \mathbf{I} + \mathbf{J}) \nabla \log \pi(\theta)$	$\sqrt{2\beta \mathbf{I}}$
RMirr	$(\beta \mathbf{B}(\theta) + \mathbf{J}) \nabla \log \pi(\theta) + \beta \nabla \cdot \mathbf{B}(\theta)$	$\sqrt{2\beta \mathbf{B}(\theta)}$
GiIrr	$(\beta \mathbf{B}(\theta) + \frac{1}{2} \mathbf{J} \mathbf{B}(\theta) + \frac{1}{2} \mathbf{B}(\theta) \mathbf{J}) \nabla \log \pi(\theta) + \nabla \cdot (\beta \mathbf{B}(\theta) + \frac{1}{2} \mathbf{J} \mathbf{B}(\theta) + \frac{1}{2} \mathbf{B}(\theta) \mathbf{J})$	$\sqrt{2\beta \mathbf{B}(\theta)}$

**Fig. 1** MSE of the running average for the first moment. Stochastic gradients are computed**Fig. 2** MSE of the running average for the second moment. Stochastic gradients are computed

the sums of the first and second moments. In Table 2, we report the asymptotic variance of the estimator for each system. We see that irreversible perturbations definitely improve the performance of the estimators, although the improvement provided by the geometry-informed irreversible perturbation seems marginal over RMirr when estimating the second moments.

When the reversible perturbation is chosen such that the drift matrix is exactly the identity (for example, when the matrix is chosen to be the covariance matrix of the posterior), additional irreversibility cannot widen the spectral gap of the system. This fact can be deduced from the results of Lelièvre et al. (2013). The improved performance of the geometry-informed irreversible perturbation is mostly due to the fact that the norm of the corresponding skew-symmetric matrix is greater than that of simple irreversibility. Even though one


can scale the skew-symmetric matrix for the other two cases to observe similar performance as geometry-informed irreversibility, GiIrr accomplishes that in a more systematic way.

In Figure 3 we plot the kernelized Stein discrepancy (KSD) for the linear-Gaussian example. We see that irreversible perturbations typically have smaller KSD compared to reversible perturbations and that in all cases the theoretical slope of $K^{-1/2}$ (see Liu et al. (2016)) is achieved.

Remark 3.1 Note that in Figure 3 and in all subsequent KSD-related figures, K ranges up to 10^4 , rather than up to 10^5 or higher (as in the bias/variance/MSE plots, e.g., Figures 1 and 2 and analogous figures in subsequent examples). The reason is that KSD is expensive to compute, and we find that evaluating it up to $K = 10^4$ is sufficient to draw conclusions.

Table 2 Asymptotic variance estimates for the linear Gaussian example

	$\mathbb{E}[\text{AVar}_{\phi_1}]$	$\text{Std}[\text{AVar}_{\phi_1}]$	$\mathbb{E}[\text{AVar}_{\phi_2}]$	$\text{Std}[\text{AVar}_{\phi_2}]$
LD	37.75	11.94	209.4	84.38
RM	20.09	6.420	132.8	49.21
Irr	15.72	5.008	135.4	47.91
RMirr	12.36	3.937	115.9	40.10
GiIrr	7.444	2.336	103.7	36.78


Fig. 3 Kernelized Stein discrepancy plot for the Gaussian example. Black line has slope $-1/2$, which denotes the expected convergence rate

Remark 3.2 While it is known that irreversible perturbations can, at worst, maintain the same performance as standard Langevin in the continuous-time setting (Rey-Bellet and Spiliopoulos 2016), when considering discretization and in borderline cases (i.e., when one does not expect much or any improvement in continuous time), irreversibility may actually harm the performance of the estimator as it introduces additional stiffness into the system without resulting in faster convergence to the invariant density. A detailed exploration of this effect is presented in Appendix A, in which we compute the bias and variance of the long term average estimator for a simple linear Gaussian problem where the posterior precision is a scalar multiple of the identity matrix. As further discussed in Remark A.1, in this case, the irreversible perturbation is not expected to lead to improvement in the sampling properties from the equilibrium. Hence, the stiffness induced upon discretization has a more profound impact on the practical performance of the irreversible perturbation.

In the current numerical study, the posterior precision is diagonal, but not a scalar multiple of the identity matrix. The eigenvalues of the resulting drift matrix are therefore distinct, and by the theory in Lelièvre et al. (2013), irreversible perturbations are able to reduce the spectral gap and result in improved performance. This is in contrast with the example studied in Appendix A. In the example of the appendix, the drift matrix is taken to be proportional to the identity matrix, and as explained in Remark A.1, the

spectral gap therefore cannot be widened. In that case, irreversibility leads to increased stiffness of the system, which then leads to increased bias and variance in the resulting estimator.

3.3 Parameters of a normal distribution

This example is identical to that used in Girolami and Calderhead (2011, Section 5) to demonstrate the performance of RMLD. Given a dataset of $\mathbb{R}$ -valued data $\mathbf{X} = \{X_i\}_{i=1}^N \sim \mathcal{N}(\mu, \sigma^2)$, we infer the parameters μ, σ . To be clear, in this example the state is $\theta = [\mu, \sigma]^T$. The prior on μ, σ is chosen to be flat (and, therefore, improper). The log-posterior is

$$\log p(\mu, \sigma | \mathbf{X}) = \frac{N}{2} \log 2\pi - N \log \sigma - \sum_{i=1}^N \frac{(X_i - \mu)^2}{2\sigma^2}. \quad (23)$$

The gradient is

$$\nabla \log p(\mu, \sigma | \mathbf{X}) = \begin{bmatrix} m_1(\mu)/\sigma^2 \\ -N/\sigma + m_2(\mu)/\sigma^3 \end{bmatrix} \quad (24)$$

where $m_1(\mu) = \sum_{i=1}^N (X_i - \mu)$, and $m_2(\mu) = \sum_{i=1}^N (X_i - \mu)^2$. In Girolami and Calderhead (2011), the authors propose using the geometry of the manifold defined by the parameter space of the posterior distribution to accelerate the resulting Metropolis-adjusted Langevin algorithm. The authors in Girolami and Calderhead (2011) suggest using the expected Fisher information matrix to define the Riemannian metric, which in the context of reversible diffusions (Rey-Bellet and Spiliopoulos 2016), is equivalent to choosing $\mathbf{B}(\mu, \sigma)$ to be the inverse of the sum of the expected Fisher information matrix and the negative Hessian of the log-prior. Straightforward computations yield

$$\mathbf{B} = \frac{\sigma^2}{N} \begin{bmatrix} 1 & 0 \\ 0 & 1/2 \end{bmatrix}, \quad \sqrt{\mathbf{B}} = \frac{\sigma}{\sqrt{N}} \begin{bmatrix} 1 & 0 \\ 0 & 1/\sqrt{2} \end{bmatrix}, \quad \nabla \cdot \mathbf{B} = \begin{bmatrix} 0 \\ \sigma/N \end{bmatrix}. \quad (25)$$

As for the geometry-informed irreversible perturbation, let $\mathbf{J} = \delta \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$, for $\delta \in \mathbb{R}$. Then the relevant quantities are

$$\frac{1}{2} \mathbf{J} \mathbf{B} + \frac{1}{2} \mathbf{B} \mathbf{J} = \frac{3\sigma^2}{4N} \mathbf{J}, \quad \frac{1}{2} \nabla \cdot (\mathbf{J} \mathbf{B} + \mathbf{B} \mathbf{J}) = \frac{3\delta\sigma}{2N} \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \quad (26)$$

In the experiments, we have $N = 30$, $h = 10^{-3}$, $\delta = 2$, and simulate $M = 1000$ independent trajectories of each system up to $T = 1000$ for a total of $K = 10^6$ steps. The initial condition is chosen to be $\mu = 5$ and $\sigma = 20$,

Fig. 4 Trajectory burn-in: each trajectory is run for $T = 2.5$. Left: single trajectories, right: mean paths. The gradients are computed exactly here





Fig. 5 Observable: $\phi_1(\mu, \sigma) = \mu + \sigma$, $\delta = 2$. Stochastic gradients are computed




Fig. 6 Observable: $\phi_2(\mu, \sigma) = \mu^2 + \sigma^2$, $\delta = 2$. Stochastic gradients are computed

which is consistent with the choice in Girolami and Calderhead (2011). The data are subsampled at a rate of $n = 6$ per stochastic gradient computation. Each trajectory is allotted a burn-in time of $T_b = 10$. The dataset is generated by drawing samples from a normal distribution with $\mu_{true} = 0$ and $\sigma_{true} = 10$. The observables we study are $\phi_1(\mu, \sigma) = \mu + \sigma$, and $\phi_2(\mu, \sigma) = \mu^2 + \sigma^2$. We plot the MSE, squared bias, and variance of resulting estimators for each observable in Figures 5 and 6. Moreover, in Table 3 we report the asymptotic variance of the estimators of each of the five systems. We plot the kernelized Stein discrepancy in Figure 7. Notice that the

irreversibly perturbed systems reach the $K^{-1/2}$ convergence rate (see Liu et al. (2016)) faster than the reversibly perturbed system. The main takeaway is that an irreversible perturbation that is adapted to the existing reversible perturbation performs much better than if the irreversible perturbation were applied without regard to the underlying geometry. Notice that the reversible perturbation considered here still improves the performance of the long term average estimator despite the fact that $\mathbf{B} - \mathbf{I}$ is not positive definite on the state space. Indeed, while $\mathbf{B} - \mathbf{I}$ being positive definite is a sufficient condition to obtain improved performance, it is not a

Table 3 Asymptotic variance estimates for the parameters of a normal distribution example. Stochastic gradients are employed

	$\mathbb{E}[\text{AVar}_{\phi_1}]$	$\text{Std}[\text{AVar}_{\phi_1}]$	$\mathbb{E}[\text{AVar}_{\phi_2}]$	$\text{Std}[\text{AVar}_{\phi_2}]$
LD	55.29	21.52	8332	4359
RM	20.63	6.019	4034	1378
Irr	5.791	2.638	2169	1072
RMirr	6.512	2.226	1729	631.2
GiIrr	1.400	0.4697	479.4	170.8

Table 4 Asymptotic variance estimates for the parameters of a normal distribution example. The gradients are computed exactly

	$\mathbb{E}[\text{AVar}_{\phi_1}]$	$\text{Std}[\text{AVar}_{\phi_1}]$	$\mathbb{E}[\text{AVar}_{\phi_2}]$	$\text{Std}[\text{AVar}_{\phi_2}]$
LD (no SG)	48.51	17.53	7339	3707
RM (no SG)	20.91	6.445	3855	1406
Irr (no SG)	5.658	2.108	2265	1191
RMirr (no SG)	6.276	2.075	1648	565.1
GiIrr (no SG)	1.363	0.4223	492.9	183.8




Fig. 7 Kernelized Stein discrepancy plot for the parameters of a normal distribution example. Black line has slope $-1/2$, which denotes the expected convergence rate

necessary one (Rey-Bellet and Spiliopoulos 2016). The reason for the reduced asymptotic variance we observed here is because the reversible perturbation $\mathbf{B}$ has eigenvalues larger than one where the bulk of the posterior distribution lies.

Figure 4 show single and mean trajectories of the burn-in period of trajectories from each of the five systems. The plot shows that the geometry-informed irreversible perturbation is able to find the bulk of the distribution sooner than the other systems without incurring additional errors due to stiffness.

To show that the GiIrr perturbation is not intimately tied to the stochastic gradient, we also report the results for each system when the gradients are computed exactly in Table 4. We see that there is little meaningful difference in the results compared to when stochastic gradients are used.

In Figure 7 we plot the KSD for the parameters-of-a-normal-distribution example. We see that GiIrr yields lower KSD than all other perturbations.

3.4 Bayesian logistic regression

Next we consider Bayesian logistic regression. Given data $\{(\mathbf{x}_i, t_i)\}_{i=1}^N$, where $\mathbf{x}_i \in \mathbb{R}^d$, and $t_i \in \{0, 1\}$, we seek a logistic function, parameterized by weights $\mathbf{w} \in \mathbb{R}^d$, that best fits the data. The weights are obtained in a Bayesian fashion, in which we endow the weights with a prior and seek to characterize its posterior distribution via sampling. Define $\varphi(y)$ to be the logistic function $\varphi(y) = (1 + \exp(-y))^{-1}$.

The log-likelihood function is

$$l(\mathbf{w}) = \sum_{i=1}^N t_i \mathbf{x}_i^T \mathbf{w} - \sum_{i=1}^N \log(1 + \exp(\mathbf{x}_i^T \mathbf{w})). \quad (27)$$

The prior for the weights is normally distributed with mean zero and covariance $\alpha^{-1} \mathbf{I}$. The gradient of the log-posterior is

$$\nabla_{\mathbf{w}} \log \pi(\mathbf{w}|\mathbf{X}) = -\alpha \mathbf{w} + \sum_{i=1}^N t_i \mathbf{x}_i - \sum_{i=1}^N \varphi(\mathbf{x}_i^T \mathbf{w}) \mathbf{x}_i. \quad (28)$$

This term is used in the drift part of the Langevin dynamics that fully computes the gradient of the log-likelihood at every step. If the data are subsampled as in SGLD, we instead compute

$$\nabla_{\mathbf{w}} \log \tilde{\pi}(\mathbf{w}|\mathbf{X}) = -\alpha \mathbf{w} + \frac{N}{n} \sum_{i=1}^n t_{\tau_i} \mathbf{x}_{\tau_i} - \frac{N}{n} \sum_{i=1}^n \varphi(\mathbf{x}_{\tau_i}^T \mathbf{w}) \mathbf{x}_{\tau_i}. \quad (29)$$

We use the `german` data set described in Gershman et al. (2012) for the numerical experiments. In this problem, there are 20 weight parameters to be learned. The training dataset is of size $N = 400$ and we choose to subsample at a rate of $n = 10$ per likelihood computation. The time step we choose is $h = 10^{-4}$ and $K = 4 \times 10^5$ steps. The initial condition is chosen to be the zero vector. We generate the skew-symmetric matrix by constructing a lower triangular matrix with entries randomly drawn from $\{1, -1\}$ and then subtracting its transpose. The diagonal is then set to zero and the matrix is scaled to have norm one.

As for the Riemannian manifold Langevin dynamics, in Girolami and Calderhead (2011) the authors use the expected Fisher information matrix plus the negative Hessian of the log-prior as the underlying metric, which in this case is equal to

$$\mathbf{G}(\mathbf{w}) = \alpha^{-1} \mathbf{I} + \mathbf{X} \mathbf{\Lambda}(\mathbf{w}) \mathbf{X}^T \quad (30)$$




Fig. 8 Observable: $\phi_1(w) = \sum_i w_i$. Bayesian logistic regression. Here, $d = 20$




Fig. 9 Observable: $\phi_2(w) = \sum_i w_i^2$. Bayesian logistic regression. Here, $d = 20$

where Λ is a diagonal matrix with entries $\Lambda_{ii}(w) = (1 - \varphi(\mathbf{x}_i^T w))\varphi(\mathbf{x}_i^T w)$ and $\mathbf{x}_i$ is the i -th column of $\mathbf{X}$. The resulting reversible perturbation uses the inverse of $\mathbf{G}(w)$. This perturbation, however, does not lead to accelerated convergence to the invariant measure since the eigenvalues of $\mathbf{G}$ are large. This implies that the eigenvalues of $\mathbf{G}^{-1}$ are less than one and so $\mathbf{G}^{-1}(w) - \mathbf{I}$ is not positive definite, a condition that needs to be satisfied to guarantee accelerated convergence (Rey-Bellet and Spiliopoulos 2015). To alleviate this issue, we consider the reversible perturbation $\mathbf{B}(w) = \mathbf{I} + \mathbf{G}^{-1}(w)$. This guarantees $\mathbf{B}(w)$ to be positive definite for all w , but the drawback is that computing the square root of $\mathbf{B}(w)$ requires explicitly computing or at least approximating the inverse of $\mathbf{G}(w)$ repeatedly in the simulation (and not just computing the action of the inverse). This additional computational cost is incurred for all examples that consider a geometry-informed perturbation, both reversible and irreversible. We show the result of this state-dependent perturbation in Figures 8 and 9 and report the asymptotic variance in Table 5. The geometry-informed irreversible perturbation does provide improvement over all other perturbations. We observe that the asymptotic variance is reduced by half over RM, with only little additional computational effort. Most of the computational cost of applying GiIrr is due to the evaluation of the reversible perturbation. Therefore we emphasize that if one is already applying the reversible perturbation to the

Table 5 Asymptotic variance estimates for the Bayesian logistic regression example with a state-dependent metric

	$\mathbb{E}[\text{AVar}_{\phi_1}]$	$\text{Std}[\text{AVar}_{\phi_1}]$	$\mathbb{E}[\text{AVar}_{\phi_2}]$	$\text{Std}[\text{AVar}_{\phi_2}]$
LD	1.967	0.9995	23.77	12.52
RM	1.328	0.6538	15.35	7.348
Irr	1.163	0.5698	14.84	7.738
RMIrr	0.8775	0.4228	10.68	5.306
GiIrr	0.7148	0.3450	8.798	4.490

Langevin dynamics, the marginal cost of applying the GiIrr perturbation is negligible.

In Figure 10 we plot the KSD for the Bayesian logistic regression example. We see that GiIrr has slightly lower KSD than the other perturbations, but the differences are small. The theoretical slope of $K^{-1/2}$ is still realized.

3.5 Independent component analysis

Our last example considers the problem of blind signal separation addressed in Welling and Teh (2011) and Amari et al. (1996). This problem yields a posterior that is strongly non-Gaussian and multi-modal, and we show that GiIrr has substantially better sampling performance over standard reversible and irreversible perturbations. Suppose there are m




Fig. 10 Kernelized Stein discrepancy plot for Bayesian logistic regression example. Black line has slope $-1/2$, which denotes the expected convergence rate

separate unknown independent signals $s^i(t)$ for $i = 1, \dots, m$ that are mixed by mixing matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$. Suppose we can observe the mixed signals $X(t) = \mathbf{M}s(t)$ for N instances in time. The goal of independent component analysis is to infer a de-mixing matrix $\mathbf{W}$ such that the m signals are recovered up to a nonzero constant and permutation. As such, this problem is generally ill-posed, but is suitable to be considered in a Bayesian context. The ICA literature states that, based on real-world data, it is best to assume a likelihood model with large kurtosis. Following (Welling and Teh 2011; Amari et al. 1996), let $p(y_i) = \frac{1}{4} \text{sech}^2(\frac{1}{2}y_i)$. The prior on the weights $\mathbf{W}_{ij}$ is Gaussian with zero mean and precision λ . The posterior is equal to

$$p(\mathbf{W}|X) \propto |\det \mathbf{W}| \prod_{i=1}^m p(\mathbf{w}_i^T \mathbf{x}) \prod_{ij} \mathcal{N}(\mathbf{W}_{ij}; 0, \lambda^{-1}). \quad (31)$$

The gradient of the log posterior with respect to the matrix $\mathbf{W}$ is then

$$f(\mathbf{W}) = \nabla_{\mathbf{W}} \log p(\mathbf{W}|X) = \left(N(\mathbf{W}^T)^{-1} - \sum_{n=1}^N \tanh\left(\frac{1}{2}\mathbf{y}_n\right) \mathbf{x}_n^T \right) - \lambda \mathbf{W}. \quad (32)$$

It is suggested in Amari et al. (1996) that the natural gradient should be used instead of the gradient we see here above to account for the information geometry of the problem. Specifically, Teh et al. (2016), Amari et al. (1996) post-multiply the gradient by $\mathbf{W}^T \mathbf{W}$ and arrive at the so-called natural gradient of the system

$$\mathcal{D}_{\mathbf{W}} := \left(N\mathbf{I}_d - \sum_{n=1}^N \tanh\left(\frac{1}{2}\mathbf{y}_n\right) \mathbf{y}_n^T \right) \mathbf{W} - \lambda \mathbf{W} \mathbf{W}^T \mathbf{W}. \quad (33)$$

In the context of RMLD, this is equivalent to perturbing the system with a reversible perturbation with $\tilde{\mathbf{B}}(\mathbf{W}) = \mathbf{W}^T \mathbf{W} \otimes \mathbf{I}_d$ pre-multiplied in front of the vectorized gradient. That is, we have

$$\text{vec}(f(\mathbf{W})\mathbf{W}^T \mathbf{W}) = (\mathbf{W}^T \mathbf{W} \otimes \mathbf{I}_d) \text{vec} f(\mathbf{W}).$$

This choice of reversible perturbation, however, may not be sufficient for accelerating the convergence of Langevin dynamics as $\tilde{\mathbf{B}}(\mathbf{W}) - \mathbf{I}_{d^2}$ is not positive definite throughout the state space (Rey-Bellet and Spiliopoulos 2016). Instead, we choose the reversible perturbation $\mathbf{B}(\mathbf{W}) = \mathbf{I}_{d^2} + (\mathbf{W}^T \mathbf{W} \otimes \mathbf{I}_d) = ((\mathbf{I}_d + \mathbf{W}^T \mathbf{W}) \otimes \mathbf{I}_d)$.

We construct the GiIrr term as follows. To take advantage of the matrix structure of the reversible perturbation, we choose the skew-symmetric matrix such that it acts within the computation of the natural gradient. We choose $\mathbf{J} = (\mathbf{I}_d \otimes \mathbf{C}_0) + (\mathbf{C}_0 \otimes \mathbf{I}_d)$ where $\mathbf{C}_0$ has the same sign pattern as (22) but such that $\mathbf{J}$ has matrix norm equal to 1. Then the geometry-informed irreversible perturbation is

$$\begin{aligned} \frac{1}{2} \mathbf{B}(\mathbf{W}) \mathbf{J} + \frac{1}{2} \mathbf{J} \mathbf{B}(\mathbf{W}) &= ((\mathbf{I}_d + \mathbf{W}^T \mathbf{W}) \otimes \mathbf{C}_0) + (\mathbf{C}_0 \otimes \mathbf{I}_d) \\ &\quad + \frac{1}{2} (\mathbf{W}^T \mathbf{W} \mathbf{C}_0 \otimes \mathbf{I}_d) \\ &\quad + \frac{1}{2} (\mathbf{C}_0 \mathbf{W}^T \mathbf{W} \otimes \mathbf{I}_d). \end{aligned}$$

To simulate the RM and GiIrr systems, correction terms (such as $\nabla \cdot \mathbf{B}(\theta)$) need to be computed. The correction terms are derived using the symbolic algebra toolbox in MATLAB. Since the perturbations are vectors of polynomials, the symbolic algebra toolbox can easily derive and efficiently evaluate the correction terms.

For the numerical experiments, we synthetically generate $m = 3$ signals, one of which is Laplace distributed, and two are distributed according to the squared hyperbolic secant distribution. The posterior distribution is $d = 9$ dimensional, there are a total of $N = 400$ data points, and the gradient is approximated by subsampling $n = 40$ data points per estimate. The initial condition here is chosen to be a diagonal matrix with either $+1$ or -1 entries, which are chosen randomly. Since the posterior is nine-dimensional and highly multimodal, it is difficult to evaluate its marginal densities directly, i.e., without sampling. Instead, we establish a baseline reference density by simulating the standard Langevin dynamics with exact computation of the likelihood over all the data over $T = 10000$ with $h = 10^{-4}$. One- and two-dimensional marginals of this baseline posterior distribution are plotted in Figure 11. The two-dimensional marginals highlight the challenges of sampling from this posterior. In Figure 12, we plot trace plots of the $\mathbf{W}_{11}$ variable for each system. By visual inspection, we see that that mixing is




Fig. 11 Posterior distribution sampled with standard Langevin with a deterministic gradient with $T = 10000$ and $h = 10^{-4}$. Notice that the system is very multimodal and non-Gaussian

best for the geometry-informed irreversibly perturbed system. One can intuitively expect that with better mixing, the geometry-informed irreversibility should yield better estimation performance than the other systems. We assess this quantitatively below.

As in the previous example, we simulate the five systems and compute the asymptotic variances of two observables for each system. Each system is simulated independently 100 times up to time $T = 2000$ with $h = 2 \times 10^{-5}$. The smaller step size is to account for the additional stiffness irreversible perturbations introduce. Since the true mean of the posterior distribution is unknown, and because standard sampling methods fail to adequately sample from the posterior distribution to get a reasonable estimate for the mean, we only plot the variances of the selected observables with respect to K in Figure 13. To compute the asymptotic variance, we allot a burn-in time of $T_b = 20$. The observables we estimate are $\phi_1(\mathbf{W}) = \sum_{i,j} \mathbf{W}_{ij}$, $\phi_2(\mathbf{W}) = \sum_{i,j} \mathbf{W}_{ij}^2$, and $\phi_3(\mathbf{W}) = \left(\sum_{i,j} \mathbf{W}_{ij}\right)^2$. The asymptotic variance numbers confirm that

the faster mixing observed in the geometry-adapted irreversible perturbation does lead to a better sampling method. The values of the asymptotic variance are reported in Table 6. The results for the asymptotic variance and variance of ϕ_2 are somewhat noisy, which is why GiIrr may appear to perform similarly as the other sampling methods. When estimating the posterior mean and an observable (ϕ_3) that includes cross-moments, the geometry-informed irreversible perturbation outperforms standard irreversibility applied to the reversible perturbation.

In Figure 14 we plot the convergence of KSD for the ICA example. GiIrr yields lower KSD than the other perturbations in this example, and the theoretical slope of $K^{-1/2}$ is also realized.

4 Conclusion

We presented a novel irreversible perturbation, GiIrr, that accelerates the convergence of Langevin dynamics. By intro-

Fig. 12 Trace plots of the W_{11} marginal






Fig. 13 Variance of running average estimators. For the second moment (middle plot), there is less difference among the samplers as the distribution is quite symmetric, and one can properly estimate the second

moment even if the samplers are stuck in a single mode. GiIrr is able to estimate the observable with cross-moments (right plot) better than the other samplers

Table 6 Asymptotic variance estimates for the ICA example

	$\mathbb{E}[\text{AVar}_{\phi_1}]$	$\text{Std}[\text{AVar}_{\phi_1}]$	$\mathbb{E}[\text{AVar}_{\phi_2}]$	$\text{Std}[\text{AVar}_{\phi_2}]$	$\mathbb{E}[\text{AVar}_{\phi_3}]$	$\text{Std}[\text{AVar}_{\phi_3}]$
LD	80.40	23.61	9.445×10^{-4}	3.616×10^{-4}	50.17	17.52
RM	53.01	12.22	5.489×10^{-4}	1.607×10^{-4}	26.75	8.442
Irr	52.38	17.27	6.854×10^{-4}	2.035×10^{-4}	27.02	9.134
RMIrr	39.20	10.36	5.794×10^{-4}	1.873×10^{-4}	19.47	6.086
GiIrr	15.50	4.441	1.253×10^{-3}	3.800×10^{-4}	6.381	1.777

**Fig. 14** Kernelized Stein discrepancy plot for the ICA example. Black line has slope $-1/2$, which denotes the expected convergence rate

ducing an irreversible perturbation that incorporates any given underlying reversible perturbation, which can also be interpreted as defining a Riemannian metric, we have shown through numerical examples that geometry-informed irreversible perturbations outperform those that are not informed as such. In the examples, we found that **GiIrr** seems to perform best when the target distribution is highly non-Gaussian.

Most of our numerical examples used stochastic gradients to cut down on computational effort in sampling each trajectory. This demonstrates that SGLD can be used in conjunction with irreversibility for practical computations.

We also provided some analysis on how irreversibility interacts with discretization of the SDE systems. Irreversibility introduces additional stiffness into the system, which may lead to additional bias or variance in the estimator. For practical purposes, one can simply choose a small enough step size so that the asymptotic bias and variance are sufficiently small. At the same time, we note an example (see Appendix A) where the introduction of the irreversible term, once discretized, leads to no improvement in the long term average estimator.

Future work could study the use of novel integrators which circumvent stiffness. For example, Jianfeng and Spiliopoulos (2018) uses a multiscale integrator, but it is not readily adapted to the data-driven setting of Bayesian inference. Another direction for future work is to theoretically characterize the performance of the geometry-informed irreversible perturbation and to compare it with that of other perturbations. A starting point for such an analysis could be

the general results of Rey-Bellet and Spiliopoulos (2016), in particular the large deviations Theorem 1 together with Propositions 2–4 therein. Preliminary investigation of this direction showed that it is a promising avenue for a theoretical investigation, but non-trivial work and a finer analysis are needed to demonstrate the effects of this class of irreversible perturbations. We leave this for future work, as such an analysis is outside the scope of this paper. Our goal in this paper has been to introduce the perturbation and showcase its potential through simulation examples.

Funding BJZ and YMM acknowledge support from the Air Force Office of Scientific Research, Analysis and Synthesis of Rare Events (ANSRE) MURI. KS was partially supported by the National Science Foundation (DMS 1550918, DMS 2107856) and Simons Foundation Award 672441.

Declarations

Conflicts of interest Not applicable.

Availability of data and material Not applicable.

Code availability Code can be made available upon request.

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

Appendix A The effects of discretization

In this section, we study the effects of discretization in the setting of an irreversibly perturbed Langevin system. Results in full generality are, as yet, elusive; therefore we only consider a Gaussian example, as it still provides insight into how irreversibility interacts with discretization in impacting the

asymptotic and finite sample bias and variance of the long term average estimator. While we do not present the results when a stochastic gradient is used, we note that the results are similar and can be easily extended based on what we present here. Recall that,

$$\mathbf{A} = \frac{1}{2}(\mathbf{I} + \mathbf{J})(\mathbf{\Gamma}_\theta + N\mathbf{\Gamma}_X),$$

$$\mathbf{D} = \frac{1}{2}(\mathbf{I} + \mathbf{J}) \left(\mathbf{\Gamma}_X \sum_{i=1}^N X_i \right), \text{ where } \mathbf{J} = \delta \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

For this analysis, all precision matrices are 2×2 scalar matrices. That is, we assume $\mathbf{\Gamma}_\theta = \sigma_\theta^{-2}\mathbf{I}$, $\mathbf{\Gamma}_X = \sigma_X^{-2}\mathbf{I}$. This is distinct from the example in Section 3.2, since the precision matrices there are diagonal but not scalar. Let $b = \frac{1}{2\sigma_X^2}$ and $S_X = \sum_{k=1}^N X_i$, so that $\mathbf{D} = b(\mathbf{I} + \mathbf{J})S_X$.

We summarize our findings here. For fixed discretization size h and scalar precision matrices as defined above, and introducing the irreversible perturbation scaled by δ , we find the following:

- The asymptotic bias for linear observables is zero, that is, $\mathbb{E}[\theta_\infty] = \mu_p$;
- The asymptotic variance for linear observables increases. We found that

$$\text{Tr Var}[\theta_\infty] = \frac{2}{2a - ha^2(1 + \delta^2)} \quad (34)$$

where $a = 0.5(1/\sigma_\theta^2 + N/\sigma_X^2)$;

- The finite time estimator for the observable $\phi(\theta) = \theta_1 + \theta_2$ has lower bias and variance;
- The finite time estimator for the observable $\phi(\theta) = \|\theta\|^2$ has higher bias and variance.

We focus on the finite time results and omit the asymptotic results, since the the former case is of more practical interest. The computations related to both are similar.

Finite time analysis: bias for linear observables. We study how the magnitude of the irreversibility, characterized by δ , impacts the mean-squared error $\text{MSE} = \mathbb{E}[\|\bar{\theta}_K - \mu_p\|^2]$ where $\bar{\theta}_K = \frac{1}{K} \sum_{k=0}^{K-1} \theta_k$. We approach this quantity via its bias-variance decomposition:

$$\text{MSE} = \|\mathbb{E}[\bar{\theta}_K] - \mu_p\|^2 + \text{Tr Var}(\bar{\theta}_K). \quad (35)$$

First, we compute the expected value of the sample average $\mathbb{E}[\bar{\theta}_K] = \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\theta_k]$. For simplicity, we assume that the initial condition is always $\theta_0 = \mathbf{0}$. For any k , we have

$$\mathbb{E}[\theta_k] = (\mathbf{I} - h\mathbf{A})\mathbb{E}[\theta_{k-1}] + h\mathbf{D}$$

$$= (\mathbf{I} - h\mathbf{A})^k \theta_0 + h \sum_{n=0}^{k-1} (\mathbf{I} - h\mathbf{A})^n \mathbf{D}$$

$$= h(\mathbf{A}h)^{-1}(\mathbf{I} - (\mathbf{I} - h\mathbf{A})^k) \mathbf{D}$$

$$= \mathbf{A}^{-1} \mathbf{D} - \mathbf{A}^{-1}(\mathbf{I} - h\mathbf{A})^k \mathbf{D}.$$

This yields

$$\mathbb{E}[\bar{\theta}_K] = \frac{1}{K} \sum_{k=0}^{K-1} (\mathbf{A}^{-1} \mathbf{D} - \mathbf{A}^{-1}(\mathbf{I} - h\mathbf{A})^k \mathbf{D})$$

$$= \mathbf{A}^{-1} \mathbf{D} - \frac{1}{K} \mathbf{A}^{-1}(\mathbf{A}h)^{-1}(\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) \mathbf{D}.$$

Since $\mu_p = \mathbf{A}^{-1} \mathbf{D}$, the bias is

$$\text{bias} = -\frac{1}{Kh} \mathbf{A}^{-2}(\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) \mathbf{D}.$$

The norm of the bias can in fact be computed. Note that $\mathbf{A}^2 = (1 + \delta^2)a^2\mathbf{I}$ and we have

$$\|\text{bias}\|^2$$

$$= \frac{1}{K^2 h^2} \mathbf{D}^T (\mathbf{I} - (\mathbf{I} - h\mathbf{A}^T)^K) \mathbf{A}^{-2T} \mathbf{A}^{-2} (\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) \mathbf{D}$$

$$= \frac{1}{K^2 h^2 a^4 (1 + \delta^2)^2} \mathbf{D}^T (\mathbf{I} - (\mathbf{I} - h\mathbf{A}^T)^K) (\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) \mathbf{D}$$

$$= \frac{b^2}{K^2 h^2 a^4 (1 + \delta^2)^2} S_X^T$$

$$(\mathbf{I} + \mathbf{J})^T (\mathbf{I} - (\mathbf{I} - \mathbf{A}^T h)^K) (\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) (\mathbf{I} + \mathbf{J}) S_X.$$

The inner matrix can be computed. Since each matrix above is simultaneously diagonalizable, we only need to consider the eigenvalues of each of the above matrices. Note that $\mathbf{I} + \mathbf{J}$ is a normal matrix, so we may write the eigenvalue decomposition $\mathbf{I} + \mathbf{J} = \mathbf{P} \mathbf{D} \mathbf{P}^*$, where $*$ denotes conjugate transpose, $\mathbf{Q} = \text{diag}(1 + i\delta, 1 - i\delta)$, and

$$\mathbf{P} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ i & -i \end{bmatrix}$$

is orthogonal. Now note that

$$\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K$$

$$= \mathbf{P} \begin{bmatrix} 1 - (1 - ah(1 + i\delta))^K & 0 \\ 0 & 1 - (1 - ah(1 - i\delta))^K \end{bmatrix} \mathbf{P}^*,$$

which implies

$$(\mathbf{I} - (\mathbf{I} - h\mathbf{A}^T)^K) (\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) = |1 - (1 - ah(1 + i\delta))^K|^2 \mathbf{I}.$$

Using the fact that $(\mathbf{I} + \mathbf{J})^T(\mathbf{I} + \mathbf{J}) = (1 + \delta^2)\mathbf{I}$, and we have the following

$$\|\text{bias}\|^2 = \frac{b^2}{K^2 h^2 a^4 (1 + \delta^2)} |1 - (1 - a(1 + i\delta)h)^K|^2 \|S_X\|. \quad (36)$$

To simplify further, we write $1 - a(1 + i\delta)h = re^{i\theta}$ where $r^2 = (1 - ah)^2 + \delta^2 a^2 h^2$, and $\tan \theta = \delta ah / (1 - ah)$. Then we obtain

$$\begin{aligned} \|\text{bias}\|^2 &= \frac{b^2}{K^2 h^2 a^4 (1 + \delta^2)} |1 - r^K e^{i\theta K}|^2 \|S_X\|^2 \\ &= \frac{b^2}{K^2 h^2 a^4 (1 + \delta^2)} (1 + r^{2K} - 2r^K \cos K\theta) \|S_X\|^2. \end{aligned}$$

We know that $r < 1$, since otherwise, the numerical scheme would be unstable. It is easy to see that for large, but not infinite, K , the bias decays as $\mathcal{O}(1/(Kh\sqrt{1 + \delta^2}))$, so the introduction of irreversibility decreases the constant in front of the expression and therefore slightly improves the convergence of the bias.

Finite time analysis: variance for linear observables. For simplicity, we assume $\theta_0 = 0$. We compute $\text{Tr Var}(\bar{\theta}_K)$. We begin with

$$\text{Tr Var}(\bar{\theta}_K) = \text{Tr } \mathbb{E}[\bar{\theta}_K \bar{\theta}_K^T] - \text{Tr } \mathbb{E}[\bar{\theta}_K] \mathbb{E}[\bar{\theta}_K]^T$$

and compute these terms separately. It is difficult to surmise a relationship between δ and $\text{Tr Var}(\bar{\theta}_K)$ even with exact formulas, so we appeal to plots of the expressions to see that the variance decreases with irreversibility. We computed $\mathbb{E}[\bar{\theta}_K]$ in the previous section.

With the observation that

$$\mathbf{A}^{-2}(\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) \mathbb{E}[\mathbf{D}] = \frac{b}{a^2} \mathbf{PQ}^* \mathbf{P}^* \mathbb{E}[S_X]$$

where

$$\mathbf{Q}' = \begin{bmatrix} \frac{1 - (1 - ah(1 + i\delta))^K}{1 + i\delta} & 0 \\ 0 & \frac{1 - (1 - ah(1 - i\delta))^K}{1 - i\delta} \end{bmatrix},$$

and $\mathbf{P}$ is defined in the previous section. We compute that

$$\begin{aligned} \text{Tr } \mathbb{E}[\bar{\theta}_K] \mathbb{E}[\bar{\theta}_K]^T &= \|\mu_p\|^2 + \|\text{bias}\|^2 - \frac{2b^2}{Kha^3(1 + \delta^2)} \\ &\quad \text{Re}\{(1 - i\delta)(1 - ah(1 + i\delta))^K\} \|\mathbb{E}[S_X]\|^2. \end{aligned} \quad (37)$$

The other term is more complicated and needs to be approached more gingerly. Observe that

$$\begin{aligned} \text{Tr } \mathbb{E}[\bar{\theta}_K \bar{\theta}_K^T] &= \frac{1}{K^2} \sum_{i,j=1}^K \text{Tr } \mathbb{E}[\theta_i \theta_j^T] = \frac{1}{K^2} \\ &\quad \left(\sum_{i=0}^{K-1} \text{Tr } \mathbb{E}[\theta_i \theta_i^T] + 2 \sum_{i < j=0}^{K-1} \text{Tr } \mathbb{E}[\theta_i \theta_j^T] \right). \end{aligned}$$

We take each term individually. To compute $\mathbb{E}[\theta_k \theta_k^T]$, it is actually better to consider the covariance matrix of θ_k , $\Sigma_k = \mathbb{E}[\theta_k \theta_k^T] - \mathbb{E}[\theta_k] \mathbb{E}[\theta_k]^T$.

We first compute

$$\begin{aligned} \mathbb{E}[\theta_k \theta_k^T] &= (\mathbf{I} - h\mathbf{A}) \mathbb{E}[\theta_{k-1} \theta_{k-1}^T] (\mathbf{I} - h\mathbf{A})^T + h^2 \mathbf{D} \mathbf{D}^T + h\mathbf{I} \\ &\quad + (\mathbf{I} - h\mathbf{A}) \mathbb{E}[\theta_{k-1}] \mathbf{D}^T h + h\mathbf{D} \mathbb{E}[\theta_{k-1}^T] (\mathbf{I} - h\mathbf{A})^T \\ \mathbb{E}[\theta_k] \mathbb{E}[\theta_k]^T &= (\mathbf{I} - h\mathbf{A}) \mathbb{E}[\theta_{k-1}] \mathbb{E}[\theta_{k-1}]^T (\mathbf{I} - h\mathbf{A})^T \\ &\quad + \mathbf{D} \mathbb{E}[\theta_{k-1}]^T (\mathbf{I} - h\mathbf{A})^T \\ &\quad + (\mathbf{I} - h\mathbf{A}) \mathbb{E}[\theta_{k-1}] \mathbf{D} + h^2 \mathbf{D} \mathbf{D}^T \end{aligned}$$

which imply the following recurrence relation. Assuming $\Sigma_0 = \mathbf{0}$, we have

$$\begin{aligned} \Sigma_k &= (\mathbf{I} - h\mathbf{A}) \Sigma_{k-1} (\mathbf{I} - h\mathbf{A})^T + h\mathbf{I} \\ &= h \sum_{n=0}^{k-1} \left((\mathbf{I} - h\mathbf{A}) (\mathbf{I} - h\mathbf{A})^T \right)^n \\ &= h \sum_{n=0}^{k-1} (\mathbf{I} - (\mathbf{A} + \mathbf{A}^T)h + h^2 \mathbf{A} \mathbf{A}^T)^n \\ &= ((\mathbf{A} + \mathbf{A}^T) - h^2 \mathbf{A} \mathbf{A}^T)^{-1} (\mathbf{I} - (\mathbf{I} - (\mathbf{A} + \mathbf{A}^T)h + \mathbf{A} \mathbf{A}^T h^2)^k). \end{aligned}$$

Let $s = 1 - 2ah + h^2 a^2 (1 + \delta^2)$, then, by recalling that $\mathbf{A} + \mathbf{A}^T = 2a\mathbf{I}$ and $\mathbf{A} \mathbf{A}^T = a^2 (1 + \delta^2) \mathbf{I}$, the above sum is equal to $\frac{1-s^k}{1-s} h\mathbf{I}$. Therefore,

$$\text{Tr } \Sigma_k = \frac{2h(1 - s^k)}{1 - s}. \quad (38)$$

Meanwhile note that

$$\mathbb{E}[\theta_k] = \mu_p - \mathbf{A}^{-1} (\mathbf{I} - h\mathbf{A})^k \mathbf{D}.$$

Therefore,

$$\begin{aligned} \text{Tr } \mathbb{E}[\theta_k] \mathbb{E}[\theta_k]^T &= \mathbb{E}[\theta_k]^T \mathbb{E}[\theta_k] \\ &= \|\mu_p\|^2 + \mathbf{D}^T (\mathbf{I} - h\mathbf{A}^T)^k \mathbf{A}^{-T} \mathbf{A}^{-1} \\ &\quad (\mathbf{I} - h\mathbf{A})^k \mathbf{D} - 2\mu_p^T \mathbf{A}^{-1} (\mathbf{I} - h\mathbf{A})^k \mathbf{D} \\ &= \|\mu_p\|^2 + \frac{s^k b^2}{a^2 (1 + \delta^2)} \|S_X\|^2 - 2\mu_p^T \mathbf{A}^{-1} (\mathbf{I} - h\mathbf{A})^k \mathbf{D}. \end{aligned}$$

We now take the sum for each expression from $k = 0$ to $K - 1$. We have

$$\begin{aligned}\sum_{i=0}^{K-1} \text{Tr} \Sigma_i &= 2h \left(\frac{K}{1-s} - \sum_{i=0}^{K-1} \frac{s^i}{1-s} \right) \\ &= 2h \left(\frac{K}{1-s} - \frac{1-s^K}{(1-s)^2} \right)\end{aligned}$$

and

$$\begin{aligned}\sum_{i=0}^{K-1} \mathbb{E}[\theta_i]^T \mathbb{E}[\theta_i] &= K \|\mu_p\|^2 + \frac{(1-s^K)b^2}{a^2(1+\delta^2)(1-s)} \\ &\quad \|S_X\|^2 - 2\mu_p^T \mathbf{A}^{-1} (h\mathbf{A})^{-1} (\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) \mathbf{D} \\ &= K \|\mu_p\|^2 + \frac{(1-s^K)b^2}{a^2(1+\delta^2)(1-s)} \\ &\quad \|S_X\|^2 - 2\mu_p^T h^{-1} \mathbf{A}^{-2} (\mathbf{I} - (\mathbf{I} - h\mathbf{A})^K) \mathbf{D}.\end{aligned}$$

For the cross-terms, observe that we may write

$$\sum_{i < j=0}^{K-1} \theta_i \theta_j^T = \sum_{i=0}^{K-1} \theta_i \sum_{j=i+1}^{K-1} \theta_j^T \quad (39)$$

which can be simplified further. First note that

$$\begin{aligned}\theta_j &= (\mathbf{I} - h\mathbf{A})\theta_{j-1} + \mathbf{D}h + \sqrt{h}\xi_{j-1} \\ &= (\mathbf{I} - h\mathbf{A})^{j-i}\theta_i + h \sum_{n=0}^{j-1-i} (\mathbf{I} - h\mathbf{A})^n \mathbf{D} \\ &\quad + \sqrt{h} \sum_{n=0}^{j-i-1} (\mathbf{I} - h\mathbf{A})^n \xi_{j-1-n}.\end{aligned}$$

Plugging this expression into the double sum above, we have

$$\begin{aligned}\sum_{i=0}^{K-1} \theta_i \sum_{j=i+1}^{K-1} \left[\theta_i^T (\mathbf{I} - \mathbf{A}^T h)^{j-i} + h \sum_{n=0}^{j-1-i} \mathbf{D}^T (\mathbf{I} - \mathbf{A}^T h)^n \right. \\ \left. + \sqrt{h} \sum_{n=0}^{j-i-1} \xi_{j-1-n}^T (\mathbf{I} - \mathbf{A}^T h)^n \right]. \quad (40)\end{aligned}$$

Taking expectations, we have

$$\begin{aligned}\sum_{i=0}^{K-1} \sum_{j=i+1}^{K-1} \left[\mathbb{E}[\theta_i \theta_i^T] (\mathbf{I} - h\mathbf{A}^T)^{j-i} + h \mathbb{E}[\theta_i] \mathbf{D}^T \sum_{n=0}^{j-1-i} (\mathbf{I} - h\mathbf{A}^T)^n \right. \\ \left. + h \mathbb{E}[\theta_i] \mathbf{D}^T (\mathbf{A}^T h)^{-1} (\mathbf{I} - (\mathbf{I} - \mathbf{A}^T h)^{j-i}) \right].\end{aligned}$$




Fig. 15 Variance for different δ , fixed h

Carrying out the computation for the first term, we have

$$\begin{aligned}F &= \sum_{i=0}^{K-1} \mathbb{E}[\theta_i \theta_i^T] \sum_{j=i+1}^{K-1} (\mathbf{I} - \mathbf{A}^T h)^{j-i} \\ &= \sum_{i=0}^{K-1} \mathbb{E}[\theta_i \theta_i^T] (\mathbf{I} - \mathbf{A}^T h) (\mathbf{A}^T h)^{-1} (\mathbf{I} - (\mathbf{I} - \mathbf{A}^T h)^{K-1-i}).\end{aligned}$$

For the second term we have,

$$\begin{aligned}\sum_{i=0}^{K-1} \mathbb{E}[\theta_i] \mu_p^T \sum_{j=i+1}^{K-1} (\mathbf{I} - (\mathbf{I} - \mathbf{A}^T h)^{j-i}) \\ = \sum_{i=0}^{K-1} \mathbb{E}[\theta_i] \mu_p^T \left[(K-1-i) \mathbf{I} - (\mathbf{I} - \mathbf{A}^T h) (\mathbf{A}^T h)^{-1} \right. \\ \left. (\mathbf{I} - (\mathbf{I} - \mathbf{A}^T h)^{K-1-i}) \right].\end{aligned}$$

The summations are difficult to compute precisely, so we compute them by direct evaluation instead. For simplicity, we assume that $\mu_p = \mathbb{E}[S_X] = [0, 0]^T$, σ_X and σ_θ are chosen such that $a = 1$. For this scenario, the bias is zero and only the variance contributes to the MSE. The variance is

$$\text{Tr}[\text{Var} \bar{\theta}_K] = \frac{1}{K^2} \left(2h \left(\frac{K}{1-s} - \frac{1-s^K}{(1-s)^2} \right) + 2\text{Tr} F \right)$$

where $\mathbb{E}[\theta_i \theta_i^T] = \Sigma_i = \frac{1-s^i}{1-s} h \mathbf{I}$.

In Figure 15 we plot the variance for varying choices of δ . In this plots, $h = 0.001$, $K = 2 \times 10^5$, and δ varies between zero and ten. We can clearly see that strengthening the irreversible perturbation leads to improvement of the squared bias and variance of the long term average estimator.

Finite sample analysis for the quadratic observable $\phi(\theta) = \|\theta\|^2$. The previous finite sample results for the observable $\phi(\theta) = \theta_1 + \theta_2$ suggests that both the bias and variance of the long term average estimator goes down with a larger irreversible term. In this section, we show that this is

actually a special case, and that when the observable is not linear, then the bias and variance may increase. We analyze the bias and variance of the long term average estimator of the observable $\phi(\theta) = \|\theta\|^2$. Define

$$\bar{\phi} = \int \phi(\theta)\pi(\theta)d\theta, \quad \bar{\phi}_K = \frac{1}{K} \sum_{k=0}^{K-1} \phi(\theta_k). \quad (41)$$

As before, we assume that $\mu_p = [0, 0]^T$, $\mathbb{E}[S_X] = \mathbf{0}$, and σ_x and σ_θ are chosen such that $a = 1$. We compute $|\mathbb{E}\bar{\phi}_K - \bar{\phi}|^2$ and $\text{Var}\bar{\phi}_K$ and see how they vary with δ . From previous computations, we can show that

$$\mathbb{E}\bar{\phi}_K = 2h \left(\frac{1}{1-s} - \frac{1-s^K}{K(1-s)^2} \right), \quad (42)$$

where $s = 1 - 2ah + a^2h^2(1 + \delta^2)$. Given this, the only term left to compute is the variance of the second moment of this observable:

$$\begin{aligned} \mathbb{E}[(\bar{\phi}_K)^2] &= \frac{1}{K^2} \mathbb{E} \left[\left(\sum_{k=0}^{K-1} \theta_k^T \theta_k \right)^2 \right] \\ &= \frac{1}{K^2} \sum_{k=0}^{K-1} \mathbb{E}[(\theta_k^T \theta_k)^2] \\ &\quad + \frac{2}{K^2} \sum_{k=0}^{K-1} \sum_{l=k+1}^{K-1} (\theta_k^T \theta_k)(\theta_l^T \theta_l). \end{aligned} \quad (43)$$

To compute the first sum, consider the following:

$$\begin{aligned} \theta_k^T \theta_k &= \theta_{k-1}^T (\mathbf{I} - h\mathbf{A})^T (\mathbf{I} - h\mathbf{A}) \theta_{k-1} + h\xi_{k-1}^T \xi_{k-1} \\ &\quad + 2\sqrt{h}\theta_{k-1}^T (\mathbf{I} - h\mathbf{A})^T \xi_{k-1} \end{aligned}$$

and so we have

$$\begin{aligned} (\theta_k^T \theta_k)^2 &= s^2(\theta_{k-1}^T \theta_{k-1})^2 + h^2(\xi_{k-1}^T \xi_{k-1})^2 \\ &\quad + 4h(\xi_{k-1}^T (\mathbf{I} - h\mathbf{A}) \theta_{k-1})^2 + 2s\theta_{k-1}^T \theta_{k-1} h\xi_{k-1}^T \xi_{k-1} \\ &\quad + 4s\sqrt{h}(\theta_{k-1}^T \theta_{k-1})\theta_{k-1}^T (\mathbf{I} - h\mathbf{A})^T \xi_{k-1} \\ &\quad + 4h^{3/2}(\xi_{k-1}^T \xi_{k-1})\theta_{k-1}^T (\mathbf{I} - h\mathbf{A})^T \xi_{k-1}. \end{aligned}$$

Taking the expectation, we have

$$\begin{aligned} \mathbb{E}[(\theta_k^T \theta_k)^2] &= s^2 \mathbb{E}[(\theta_{k-1}^T \theta_{k-1})^2] + h^2 \mathbb{E}[(\xi_{k-1}^T \xi_{k-1})^2] \\ &\quad + 4h \mathbb{E}[(\xi_{k-1}^T (\mathbf{I} - h\mathbf{A}) \theta_{k-1})^2] \\ &\quad + 2sh \mathbb{E}[(\theta_{k-1}^T \theta_{k-1})(\xi_{k-1}^T \xi_{k-1})]. \end{aligned}$$

After simplifying, we arrive at the following recurrence relation:

$$\mathbb{E}[(\theta_k^T \theta_k)^2] = s^2 \mathbb{E}[(\theta_{k-1}^T \theta_{k-1})^2] + 8h^2 + 8sh \mathbb{E}[\theta_{k-1}^T \theta_{k-1}]. \quad (44)$$

Let $\beta_k = \mathbb{E}[(\theta_k^T \theta_k)^2]$, $\zeta_k = 8sh \mathbb{E}[\theta_k^T \theta_k]$, and $\kappa = 8h^2$. We have the following recurrence, which we solve

$$\begin{aligned} \beta_k &= s^2 \beta_{k-1} + \zeta_{k-1} + \kappa \\ &= s^{2k} \beta_0 + \sum_{n=0}^{k-1} s^{2n} \zeta_{k-n-1} + \sum_{n=0}^{k-1} s^{2n} \kappa. \end{aligned}$$

From previous for the term ζ_k , we have

$$\begin{aligned} \beta_k &= \sum_{n=0}^{k-1} s^{2n} 8sh \cdot 2h \frac{1-s^{k-n-1}}{1-s} + \kappa \frac{1-s^{2k}}{1-s^2} \\ &= \frac{16sh^2}{1-s} \sum_{n=0}^{k-1} (s^{2n} - s^{k+n-1}) + \frac{8h^2(1-s^{2k})}{1-s^2} \\ &= \frac{16sh^2}{1-s} \left(\frac{1-s^{2k}}{1-s^2} - s^{k-1} \frac{1-s^k}{1-s} \right) + \frac{8h^2(1-s^{2k})}{1-s^2}. \end{aligned}$$

Next we compute the summation of the cross terms. Define R_k such that

$$\sum_{k=0}^{K-1} R_k = \sum_{k=0}^{K-1} \sum_{l=k+1}^{K-1} \mathbb{E}[(\theta_k^T \theta_k)(\theta_l^T \theta_l)]. \quad (45)$$

We write

$$\begin{aligned} \theta_l^T \theta_l &= s\theta_{l-1}^T \theta_{l-1} + h\xi_{l-1}^T \xi_{l-1} + 2\sqrt{h}\xi_{l-1}^T (\mathbf{I} - h\mathbf{A}) \theta_{l-1} \\ &= s^{l-k} \theta_k^T \theta_k \\ &\quad + \sum_{n=0}^{l-k-1} h s^n \xi_{l-n-1}^T \xi_{l-n-1} + 2\sqrt{h} s^n \xi_{l-n-1}^T (\mathbf{I} - h\mathbf{A}) \theta_{l-n-1}. \end{aligned}$$

This implies that

$$\begin{aligned} R_k &= \sum_{l=k+1}^{K-1} \left(s^{l-k} \mathbb{E}[(\theta_k^T \theta_k)^2] + \sum_{n=0}^{l-k-1} 2hs^n \mathbb{E}[\theta_k^T \theta_k] \right) \\ &= \sum_{l=k+1}^{K-1} \beta_k s^{l-k} + 2h \mathbb{E}[\theta_k^T \theta_k] \sum_{n=0}^{l-k-1} s^n \\ &= \sum_{l=k+1}^{K-1} \beta_k s^{l-k} + 2h \mathbb{E}[\theta_k^T \theta_k] \frac{1-s^{l-k}}{1-s} \\ &= \beta_k \sum_{l=k+1}^{K-1} s^{l-k} + \frac{2h \mathbb{E}[\theta_k^T \theta_k]}{1-s} \sum_{l=k+1}^{K-1} 1-s^{l-k} \end{aligned}$$




Fig. 16 (Squared) bias and variance of $\phi(\theta) = \|\theta\|^2$ for varying levels of irreversibility

$$= \beta_k \frac{s - s^{K-k}}{1-s} + \frac{2h\mathbb{E}[\theta_k^T \theta_k]}{1-s} \left(K - 1 - k - \frac{s - s^{K-k}}{1-s} \right).$$

To summarize, we have

$$\mathbb{E}[(\bar{\phi}_K)^2] = \frac{1}{K^2} \sum_{k=0}^{K-1} (\beta_k + 2R_k), \quad (46)$$

the squared bias is $(\mathbb{E}\bar{\phi}_K - 1)^2$ and the variance is $\mathbb{E}[(\bar{\phi}_K)^2] - \mathbb{E}[\bar{\phi}_K]^2$. These expressions are not simplifiable easily, so we plot these expressions and study their trends. In Figure 16, we plot the squared bias and variance for fixed h and K and varying δ . In these plots, $h = 0.001$, $K = 2 \times 10^5$, and δ varies between zero and ten. Notice that for these choices, both the squared bias and variance increases as δ grows, showing that for irreversibility provides no benefit, and in fact, harms the performance of the standard estimator.

Remark A.1 When discretization is considered, sampling properties in this simple example in which the drift is proportional to the state cannot be improved upon using irreversibility. Let us now explain this phenomenon from a theoretical point of view. In Rey-Bellet and Spiliopoulos (2016), the authors show that adding an irreversible perturbation to the generator of the diffusion process may decrease the spectral gap, and will never increase it. They further prove that in the continuous case, decreasing the spectral gap then decreases the asymptotic variance. However this improvement is not strict, that is, irreversibility is only guaranteed to not increase the spectral gap.

Meanwhile, in Lelievre et al. (2013), the authors consider irreversibility only in the context of linear systems, and rigorously study optimal irreversible perturbations that accelerate convergence to the invariant distribution. Their results show that when the drift matrix is proportional to the identity matrix, the spectral gap cannot be widened. Proposition 4 in Lelievre et al. (2013) shows that the leading nonzero eigenvalue of the irreversibly perturbed drift matrix is bounded

above by the leading nonzero eigenvalue of the original drift matrix and below by the trace of the original drift matrix over the dimension of the state space. The lower bound is then the optimal spectral gap. For a drift matrix that is a multiple of the identity, the upper and lower bounds are the same, which implies that the spectral gap can never decrease from its original value in the continuous case. After factoring in discretization, the irreversible perturbation increases stiffness of the system, which contributes to increased bias and variance in the resulting estimator.

References

- Amari, Shun-ichi, Cichocki, Andrzej, Yang, Howard Hua: A new learning algorithm for blind signal separation. In: *Advances in Neural Information Processing Systems*, pages 757–763. Morgan Kaufmann Publishers, (1996)
- Asmussen, Søren., Glynn, Peter W.: *Stochastic simulation: algorithms and analysis*, vol. 57. Springer Science & Business Media, Germany (2007)
- Bierkens, Joris: Non-reversible Metropolis-Hastings. *Stat. Comput.* **26**, 1213–1228 (2016)
- Brosse, Nicolas, Durmus, Alain, Moulines, Éric: The promises and pitfalls of stochastic gradient Langevin dynamics. In: *NeurIPS 2018 (Advances in Neural Information Processing Systems 2018)*, (2018)
- Diaconis, P., Holmes, S., Neal, R.: Analysis of a nonreversible Markov chain sampler. *Ann. Appl. Probab.* **10**, 726–752 (2010)
- Duncan, A.B., Pavliotis, G.A., Zygalkakis, K.C.: Nonreversible Langevin samplers: Splitting schemes, analysis and implementation. *arXiv preprint arXiv:1701.04247*, (2017)
- Durmus, Alain, Moulines, Eric: High-dimensional Bayesian inference via the unadjusted Langevin algorithm. *Bernoulli* **25**(4A), 2854–2882 (2019)
- Franke, Brice, Hwang, C.-R., Pai, H.-M., Sheu, S.-J.: The behavior of the spectral gap under growing drift. *Trans. Am. Math. Soc.* **362**(3), 1325–1350 (2010)
- Ganguly, Arnab, Sundar, P.: Inhomogeneous functionals and approximations of invariant distribution of ergodic diffusions: error analysis through central limit theorem and moderate deviation asymptotics. *Stoch. Proc. Appl.* **133**(C), 74–110 (2021)
- Gershman, Samuel J., Hoffman, Matthew D., Blei, David M.: Non-parametric variational inference. In: *Proceedings of the 29th*

- International Conference on International Conference on Machine Learning, ICML'12, pages 235–242, Madison, WI, USA, (2012). Omnipress
- Girolami, Mark, Calderhead, Ben: Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)* **73**(2), 123–214 (2011)
- Gorham, Jackson, Duncan, Andrew B., Vollmer, Sebastian J., Mackey, Lester: Measuring sample quality with diffusions. *Ann. Appl. Probab.* **29**(5), 2884–2928 (2019)
- Gorham, Jackson, Mackey, Lester: Measuring sample quality with stein's method. *Advances in Neural Information Processing Systems*, 28, (2015)
- Gorham, Jackson, Mackey, Lester: Measuring sample quality with kernels. In: *International Conference on Machine Learning*, pages 1292–1301. PMLR, (2017)
- Hu, Yuanhan, Wang, Xiaoyu, Gao, Xuefeng, Gürbüzbalaban, Mert, Zhu, Lingjiong: Non-convex optimization via non-reversible stochastic gradient Langevin dynamics. *arXiv preprint arXiv:2004.02823*, (2020)
- Hwang, Chii-Ruey, Hwang-Ma, Shu-Yin, Sheu, Shuenn-Jyi: Accelerating Gaussian diffusions. *The Annals of Applied Probability*, pages 897–913, (1993)
- Hwang, Chii-Ruey, Hwang-Ma, Shu-Yin, Sheu, Shuenn-Jyi: Accelerating diffusions. *Ann. Appl. Probab.* **15**(2), 1433–1444 (2005)
- Izzatullah, Muhammad, Baptista, Ricardo, Mackey, Lester, Marzouk, Youssef, Peter, Daniel: Bayesian seismic inversion: Measuring langevin mcmc sample quality with kernels. In: *SEG International Exposition and Annual Meeting, OnePetro*, (2020)
- Lelièvre, Tony, Nier, Francis, Pavliotis, Grigorios A.: Optimal non-reversible linear drift for the convergence to equilibrium of a diffusion. *J. Stat. Phys.* **152**(2), 237–274 (2013)
- Liu, Q., Lee, J., Jordan, M.: A kernelized Stein discrepancy for goodness-of-fit tests. *Proc. of 33rd ICML* **48**, 276–284 (2016)
- Livingstone, Samuel, Girolami, Mark: Information-geometric Markov chain Monte Carlo methods using diffusions. *Entropy* **16**(6), 3074–3102 (2014)
- Jianfeng, Lu., Spiliopoulos, Konstantinos: Analysis of multiscale integrators for multiple attractors and irreversible Langevin samplers. *Mult. Model. Simul.* **16**(4), 1859–1883 (2018)
- Ma, Yi.-An., Chen, Tianqi, Fox, Emily B.: A complete recipe for stochastic gradient MCMC. *NIPS'15: Proceedings of the 28th International Conference on Neural Information Processing Systems* **2**, 2917–2925 (2015)
- Ottobre, Michela, Pillai, Natesh S., Spiliopoulos, Konstantinos: Optimal scaling of the MALA algorithm with irreversible proposals for Gaussian targets. *Stochastics and Partial Differential Equations: Analysis and Computations*, pages 1–51, (2019)
- Pavliotis, Grigorios A.: *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, vol. 60. Springer, Germany (2014)
- Rey-Bellet, Luc, Spiliopoulos, Konstantinos: Irreversible Langevin samplers and variance reduction: a large deviations approach. *Nonlinearity* **28**(7), 2081 (2015)
- Rey-Bellet, Luc, Spiliopoulos, Konstantinos: Variance reduction for irreversible Langevin samplers and diffusion on graphs. *Electronic Communications in Probability*, 20, (2015)
- Rey-Bellet, Luc, Spiliopoulos, Konstantinos: Improving the convergence of reversible samplers. *J. Stat. Phys.* **164**(3), 472–494 (2016)
- Roberts, Gareth O., Tweedie, Richard L.: Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2**(4), 341–363 (1996)
- Teh, Yee Whye, Thiery, Alexandre H., Vollmer, Sebastian J.: Consistency and fluctuations for stochastic gradient Langevin dynamics. *Journal of Machine Learning Research*, 17, (2016)
- Vollmer, Sebastian J., Zygalakis, Konstantinos C., Teh, Yee Whye: Exploration of the (non-) asymptotic bias and variance of stochastic gradient Langevin dynamics. *J. Mach. Learn. Res.* **17**(1), 5504–5548 (2016)
- Welling, Max, Teh, Yee W.: Bayesian learning via stochastic gradient Langevin dynamics. In: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688. Citeseer, (2011)
- Xifara, Tatiana, Sherlock, Chris, Livingstone, Samuel, Byrne, Simon, Girolami, Mark: Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Stat. Probab. Lett.* **91**, 14–19 (2014)

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