Irreversible Samplers from Jump and Continuous Markov Processes

Yi-An Ma · Emily B. Fox · Tianqi Chen · Lei Wu

Abstract In this paper, we propose irreversible versions of the Metropolis Hastings (MH) and Metropolis adjusted Langevin algorithm (MALA) with a main focus on the latter. For the former, we show how one can simply switch between different proposal and acceptance distributions upon rejection to obtain an irreversible jump sampler (I-Jump). The resulting algorithm has a simple implementation akin to MH, but with the demonstrated benefits of irreversibility. We then show how the previously proposed MALA method can also be extended to exploit irreversible stochastic dynamics as proposal distributions in the I-Jump sampler. Our experiments explore how irreversibility can increase the efficiency of the samplers in different situations.

Yi-An Ma
Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
E-mail: yianma@berkeley.edu

Emily B. Fox
Department of Statistics, University of Washington
E-mail: ebfox@stat.washington.edu

Tianqi Chen
Department of Computer Science and Engineering, University of Washington

Lei Wu
School of Mathematical Sciences, Peking University

1 Introduction

Markov chain Monte Carlo (MCMC) methods are the defacto tools for inference in Bayesian models [44, 62]. The Metropolis-Hastings (MH) algorithm is often used as the default approach because of its ease of implementation. One designs a proposal distribution to generate samples and uses an accept-reject procedure to ensure that the target distribution is maintained. A focus has been on developing clever proposals [66, 33, 44] to specify a Markov process with good mixing rates, but traditional methods are often strongly coupled to a specific challenge setting, like multimodal targets [66] or heavy tailed distributions [33]. In practice, one often does not know the structure of the target distribution, which might additionally exhibit a combination of these factors.

To address some of these limitations, the Metropolis adjusted Langevin algorithm (MALA) [63, 74] and Hamiltonian Monte Carlo (HMC) [21, 51] have been proposed. These methods are constructed to use the local gradient information of the target distribution in the proposal. In these approaches, the task of constructing good samplers is translated to finding continuous

Markov dynamics, in the form of differential equations (stochastic or deterministic), that can generate better MH proposals for general target distributions. A lot of recent attention has been on extensions to the Langevin and Hamiltonian dynamics, such as through use of second order information of the target distribution resulting in Riemann manifold variants [74, 25].

One restriction of all the above approaches is that regardless of the properties of the underlying continuous dynamics, direct use of the MH correction necessarily results in the whole sampler being reversible (the forward path is statistically indistinguishable from the backward one). This represents a serous limitation since irreversibility has been shown to increase the mixing rate of samplers in general [50, 19, 14, 12].

Our goal is to address this issue by defining easy-to-specify and computationally straightforward irreversible samplers. In particular, we focus in on defining an irreversible MALA algorithm (I-MALA); the construction of this algorithm consists of two parts. One part is to construct irreversible continuous Markov dynamics that leaves the target distribution invariant (see Sec. 2). Then, to correct for the discretization error involved in simulating the continuous Markov dynamics—while avoiding use of the reversible MH procedure—we develop an irreversible jump sampler (I-Jump). The I-Jump algorithm has an implementation similar to the simplicity of the MH algorithm, but allows for irreversibility (see Sec. 3).

We start in Sec. 2 by providing background on samplers using continuous Markov processes and present a general stochastic differential equation (SDE) based framework; a preliminary version of this work appeared in [46]. Within the framework, general irreversible continuous dynamics are defined through specifying two matrices: a positive semidefinite matrix and a skew-symmetric matrix. We prove that for any choice of these matrices, the continuous dynamics leave the target distribution invariant. We likewise prove that any continuous Markov process with the correct stationary distribution has a representation in this framework.

We then turn our attention in Sec. 3 to introducing the I-Jump sampler alternative to the MH algorithm. We first reparameterize the jump process and arrive at a straightforward set of constraints on the transition probabilities that ensures that the target distribution is

the stationary distribution. We then revise the MH algorithm to allow for irreversibility, while at the same time satisfying the derived constraints for correctness. The resulting sampler implementation has the ease and efficiency of the standard MH method. Importantly, our method outperforms existing approaches [48, 27] in terms of mixing rate versus runtime in a range of settings, from heavy-tailed to multimodal targets. We further demonstrate these performance gains in a set of challenging real world applications.

In Sec. 4, we then make use of the general irreversible stochastic dynamics framework to define better proposal distributions in the I-Jump sampler, leading to our sought-after I-MALA algorithm. Similar to the MALA or HMC algorithm, we use a discretization of the continuous dynamics to propose a sample, but then use our I-Jump sampler rather than the reversible MH correction. Importantly, local gradient information of the target distribution is taken into account in the SDE for better efficiency than the I-Jump algorithm with standard independent proposals. Furthermore, the simplicity of the MH algorithm is still inherited while the overall sampling process remains irreversible. We can view the benefits of this approach from two angles: (i) the SDE can provide an efficient proposal distribution for our I-Jump sampler or (ii) the accept-reject scheme allows us to correct for the bias introduced by sampling the continuous dynamics via a discretized SDE. This also opens up the possibility to combine, for example, the fast mixing of Langevin diffusion within local modes with the fast traversing of Hamiltonian dynamics. We examine the increased efficiency of our proposed I-MALA algorithm within the context of a Bayesian logistic regression model and stochastic volatility model.

2 Samplers Using Continuous Markov Processes

We start with the standard MCMC goal of drawing samples from an unnormalized target distribution $\pi(\theta)$. It is often customary to include auxiliary variables r (e.g. in the HMC algorithm) to facilitate the sampling process. Hence we write $\mathbf{z} = (\theta, r)$ to contain all variables being sampled according to the joint distribution $\pi(\mathbf{z})$. In this section, we focus on continuous Markov processes and discuss possible choices of continuous dynamics that

leave $\pi(\mathbf{z})$ invariant. If the stochastic process is further ergodic, then simulating the stationary stochastic dynamics equates with providing samples from $\pi(\mathbf{z})$. These processes will be deployed in the I-MALA algorithm of Sec. 4.

A realization of a continuous Markov process can be represented as the following stochastic differential equation (SDE):

$$\mathrm{d}\mathbf{z} = \mathbf{f}(\mathbf{z})\mathrm{d}t + \sqrt{2\mathbf{D}(\mathbf{z})}\mathrm{d}\mathbf{W}(t), \tag{1}$$

where $\mathbf{z}: [0, +\infty) \times \Omega \to \mathbb{R}^d$ is a real random vector defined on a probability space (Ω, Σ, P), parameterized by time t; $\mathbf{f} : \mathbb{R}^d \to \mathbb{R}^d$ is a real vector valued function; $\mathbf{D}(\mathbf{z})$ is a $d \times d$ positive semidefinite diffusion matrix; and $\mathbf{W}(t)$ is a d dimensional Wiener process. Following Itô's convention, (1) defines the following diffusion process (with an abuse of notation, we also use $\mathbf{z}$ to denote value of the random vector $\mathbf{z}$):

$$\frac{\partial}{\partial t}p(\mathbf{z};t) = \sum_{i,j=1}^{d} \frac{\partial^2}{\partial \mathbf{z}_i \partial \mathbf{z}_j}\big[\mathbf{D}_{ij}(\mathbf{z})p(\mathbf{z};t)\big]$$
$$- \sum_{i=1}^{d} \frac{\partial}{\partial \mathbf{z}_i}\big[\mathbf{f}_i(\mathbf{z})p(\mathbf{z};t)\big], \tag{2}$$

where $p(\mathbf{z};t)$ is the probability density function (assuming it exists for all t) of random vector $\mathbf{z}$ at time t. The continuous Markov process, (1), can be used to generate samples from $\pi(\mathbf{z})$ if $\pi(\mathbf{z})$ is a stationary solution to (2).

Although (1) provides a way to simulate the continuous dynamics and obtain samples from the Markov process, it is not clear which choices of $\mathbf{f}$ and $\mathbf{D}$ will result in a stationary distribution of (2) equal to the target distribution $\pi(\mathbf{z})$. For a given $\mathbf{f}$ and $\mathbf{D}$, (2) allows us to analyze this stationary distribution, but it is very challenging to define the set of $\mathbf{f}$ and $\mathbf{D}$ that yield a specified stationary distribution. Researchers have resorted to special cases such as overdamped Langevin [63, 72], underdamped Langevin [29, 13] and Nosé-Hoover [20, 64] dynamics in the statistical physics literature for inspiration.

2.1 Complete Recipe with Continuous Markov Dynamics

We propose using an alternative form for (1) specified via two matrices $\mathbf{D}$ and $\mathbf{Q}$, as well as the target distribution $\pi(\mathbf{z})$, as first considered in [46, 65]:

$$\mathrm{d}\mathbf{z} = \Big[-\big(\mathbf{D}(\mathbf{z}) + \mathbf{Q}(\mathbf{z})\big)\nabla H(\mathbf{z}) + \varGamma(\mathbf{z})\Big]\mathrm{d}t$$
$$+ \sqrt{2\mathbf{D}(\mathbf{z})}\mathrm{d}\mathbf{W}(t),$$
$$\varGamma_i(\mathbf{z}) = \sum_{j=1}^{d} \frac{\partial}{\partial \mathbf{z}_j}\big(\mathbf{D}_{ij}(\mathbf{z}) + \mathbf{Q}_{ij}(\mathbf{z})\big) \tag{3}$$

Here, $H(\mathbf{z}) = -\log(\pi(\mathbf{z}))$; $\mathbf{D}(\mathbf{z})$ is a positive semidefinite diffusion matrix and $\mathbf{Q}(\mathbf{z})$ a skew-symmetric matrix. As discussed in Appendix A.2, (3) decomposes into reversible and irreversible process. Matrix $\mathbf{D}(\mathbf{z})$ corresponds to the reversible part while matrix $\mathbf{Q}(\mathbf{z})$ determines the irreversible part. In the following, we use the Fokker-Planck equation associated with (3) to verify two properties of this representation: One is that (3) has $\pi(\mathbf{z})$ as its invariant distribution; the other is that any continuous Markov process with $\pi(\mathbf{z})$ as the invariant distribution can be written in the form of (3) (i.e., there exists a $\mathbf{D}(\mathbf{z})$ and $\mathbf{Q}(\mathbf{z})$ that place the process in this representation). Together, these two properties (i) allow us to very straightforwardly explore a set of valid samplers by specifying pairs $(\mathbf{D}(\mathbf{z}), \mathbf{Q}(\mathbf{z}))$ of positive semidefinite and skew-symmetric matrices, respectively, and (ii) ensure that as we span the space of all possible $(\mathbf{D}(\mathbf{z}), \mathbf{Q}(\mathbf{z}))$, we know we have covered all possible valid samplers. That is, our representation is *complete*. A preliminary version of this complete framework appeared in [46].

The Fokker-Planck equation (following Itô's convention) associated with (3) is (see Appendix A.1):

$$\frac{\partial}{\partial t}p(\mathbf{z};t) \tag{4}$$
$$= \nabla^T \cdot \bigg(\big[\mathbf{D}(\mathbf{z}) + \mathbf{Q}(\mathbf{z})\big]\big[p(\mathbf{z};t)\nabla H(\mathbf{z}) + \nabla p(\mathbf{z};t)\big] \bigg),$$

where $\nabla^T \cdot \mathbf{f}(\mathbf{z}) = \sum_{i=1}^{d} \frac{\partial f_i(\mathbf{z})}{\partial \mathbf{z}_i}$. It is straightforward to verify that $p^s(\mathbf{z}) \propto \pi(\mathbf{z}) = e^{-H(\mathbf{z})}$ is a stationary solution to (4). More significantly, Theorem 1 states that any process (1) and (2) with stationary distribution $\pi(\mathbf{z})$ has a representation in our framework.

Theorem 1 *Suppose* (2) *has stationary probability density function* $p^s(\mathbf{z}) \propto \pi(\mathbf{z})$. *Further assume that* $\left[\mathbf{f}_i(\mathbf{z})\pi(\mathbf{z}) - \sum_{j=1}^d \frac{\partial}{\partial \theta_j}\Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z})\Big)\right]$ *is Lebesgue integrable. Then, there exists a skew-symmetric matrix* $\mathbf{Q}(\mathbf{z})$ *such that the right hand side of* (2) *is equivalent to the right hand side of* (4).

For the reader's convenience, we include the constructive proof of Theorem 1 in Appendix A.1. To the best of our knowledge, the exact form of (3) was first presented in the statistical mechanics literature [75, 65]; however, the completeness of the representation of continuous Markov processes was made only later in our preliminary paper [46]. The proof of Theorem 1 is comprised of two sets of ideas stemming from different fields: In the study of continuous Markov processes, earlier work [59, 17, 58, 18, 70, 54] realized that diffusion processes with stationary probability density function $\pi(\mathbf{z})$ can be decomposed into reversible and irreversible parts preserving $\pi(\mathbf{z})$ as the invariant measure. In stochastic models of fluid dynamics and homogenization, earlier work [36] found that divergenceless vector fields can be represented as the divergence of an anti-symmetric matrix valued potential. Combining both ideas leads to the discovery of (4) that underlies the proof of Theorem 1. Similar structures have also been examined when one or both of $\mathbf{D}(\mathbf{z})$ and $\mathbf{Q}(\mathbf{z})$ are constant matrices [30, 39].

2.2 Continuous Markov Process Sampling Algorithm

We can simulate from (3) by using the following ϵ-discretization of the SDE:

$$
\begin{aligned}
\mathbf{z}_{t+1} \leftarrow & \mathbf{z}_t + \epsilon_t \left[-\big(\mathbf{D}(\mathbf{z}_t) + \mathbf{Q}(\mathbf{z}_t)\big)\nabla H(\mathbf{z}_t) + \Gamma(\mathbf{z}_t)\right] \\
& + \eta_t, \quad \eta_t \sim \mathcal{N}(0, 2\epsilon_t \mathbf{D}(\mathbf{z}_t)).
\end{aligned} \tag{5}
$$

Although (5) is in the form of the Euler–Maruyama method, higher order numerical schemes can be used for better accuracy [11, 9, 40]. For example, in HMC, since the diffusion matrix $\mathbf{D}$ is zero, a neutral integration scheme such as leap frog (a second order integration scheme) is often used for accuracy and stability of integration [51]. Other higher order numerical methods

---

Algorithm 1: Continuous Markov Process Sampling Algorithm

initialize $\mathbf{z}_0$
for $t = 0, 1, 2 \cdots N_{iter}$ **do**
  for $i = 1 \cdots n$ **do**
    $\Gamma_i(\mathbf{z}) = \sum_j \frac{\partial}{\partial \mathbf{z}_j}\left(\mathbf{D}_{ij}(\mathbf{z}) + \mathbf{Q}_{ij}(\mathbf{z})\right)$
  end
  sample $\eta_t \sim \mathcal{N}(0, 2\epsilon_t \mathbf{D}(\mathbf{z}_t))$
  $\mathbf{z}_{t+1} \leftarrow \mathbf{z}_t +$
  $\epsilon_t \left[-\big(\mathbf{D}(\mathbf{z}_t) + \mathbf{Q}(\mathbf{z}_t)\big)\nabla H(\mathbf{z}_t) + \Gamma(\mathbf{z}_t)\right] + \eta_t$
end

---

such as the splitting scheme [11, 9] and simple modifications to the Euler–Maruyama method [40] can also lead to higher order of accuracy in different scenarios.

The resulting algorithm according to (5) is outlined in Algorithm 1. Note that relying on a sample path from the discretized system of (3) typically leads to the introduction of bias due to discretization error. In these cases, the samples only provide unbiased estimates in the limit as $\epsilon_t \to 0$ unless further corrections are introduced. In Sec. 4, we use the dynamics of (5) as the proposal distribution inside an irreversible jump process to correct for any potential discretization error.

2.3 Previous Dynamics in MCMC Algorithms as Special Cases

We explicitly state how some previous continuous-dynamic-based MCMC methods fit within the proposed framework based on specific choices of $\mathbf{D}(\mathbf{z})$, $\mathbf{Q}(\mathbf{z})$ and $H(\mathbf{z})$.

Hamiltonian Dynamics   The key ingredient in HMC [21, 51] is Hamiltonian dynamics, which simulates the physical motion of an object with position θ, momentum r, and mass $\mathbf{M}$ on an frictionless surface as follows:

$$
\begin{cases}
\mathrm{d}\theta = \mathbf{M}^{-1} r \mathrm{d}t \\
\mathrm{d}r = -\nabla U(\theta)\mathrm{d}t.
\end{cases} \tag{6}
$$

It can be observed that (6) is a special case of the proposed framework with $\mathbf{z} = (\theta, r)$, $H(\theta, r) = U(\theta) + \frac{1}{2}r^T \mathbf{M}^{-1} r$, $\mathbf{Q}(\theta, r) = \begin{pmatrix} 0 & -\mathbf{I} \\ \mathbf{I} & 0 \end{pmatrix}$ and $\mathbf{D}(\theta, r) = \mathbf{0}$. See Sec. 4.1 for a more complete discussion.

Langevin Dynamics   The Langevin dynamics sampler [63, 72] proposes to use the following first order (no momentum) Langevin dynamics to generate samples

$$d\theta = -\mathbf{D}\nabla U(\theta)dt + \sqrt{2\mathbf{D}}d\mathbf{W}(t). \tag{7}$$

This algorithm corresponds to taking $\mathbf{z} = \theta$ with $H(\theta) = U(\theta)$, $\mathbf{D}(\theta) = \mathbf{D}$, and $\mathbf{Q}(\theta) = \mathbf{0}$.

Riemannian Langevin Dynamics   The Langevin dynamics sampler can be generalized to use an adaptive diffusion matrix $\mathbf{D}(\theta)$. Specifically, it is interesting to take $\mathbf{D}(\theta) = \mathbf{G}^{-1}(\theta)$, where $\mathbf{G}(\theta)$ is the Fisher information metric [74, 53]. The sampler iterates

$$d\theta = [-\mathbf{G}(\theta)^{-1}\nabla U(\theta) + \Gamma^{\mathbf{D}}(\theta)]dt$$
$$+ \sqrt{2\mathbf{G}(\theta)^{-1}}d\mathbf{W}(t). \tag{8}$$

We can cast this Riemannian Langevin dynamics sampler [53] into our framework taking $\mathbf{D}(\theta) = \mathbf{G}(\theta)^{-1}$, and $\mathbf{Q}(\theta) = \mathbf{0}$. From our framework, we know that here

$$\Gamma_i^{\mathbf{D}}(\theta) = \sum_j \frac{\partial \mathbf{D}_{ij}(\theta)}{\partial \theta_j}. \tag{9}$$

Interestingly, in earlier literature [25], $\Gamma_i^{\mathbf{D}}(\theta)$ was taken to be $2 |\mathbf{G}(\theta)|^{-1/2} \sum_j \frac{\partial}{\partial \theta_j} \left(\mathbf{G}_{ij}^{-1}(\theta)|\mathbf{G}(\theta)|^{1/2} \right)$. More recently, it was found that this correction term corresponds to the distribution function with respect to a non-Lebesgue measure [63]. For the Lebesgue measure, the revised $\Gamma_i^{\mathbf{D}}(\theta)$ was as determined by our framework [63]. This is an example of how our framework can provide guidance in devising correct samplers.

2.4 Irreversibility in continuous dynamics

In Appendix A.2, we show that the continuous stochastic dynamics (3) decompose into (i) general Riemannian Langevin dynamics and (ii) conserved, deterministic dynamics generalizing Hamiltonian dynamics. The first component is reversible and is determined by $\mathbf{D}(\mathbf{z})$ while the second part is irreversible and is determined by $\mathbf{Q}(\mathbf{z})$. The irreversible dynamics generates a circular motion that traverses through the state space.

It has been proven that incorporating this irreversible dynamics (parameterized by $\mathbf{Q}(\mathbf{z})$) can only increase the mixing of the Markov process [30, 31, 22, 61]. The intuition can be drawn from analyzing the nonzero part of the spectrum of the Fokker-Planck operator in (4). It can be shown that the traversing motion brings different eigenvalues closer to each other, making the overall spectrum (except zero) narrower [34]. When the spectral gap of the reversible dynamics is nonzero, introducing this narrowing effect only increases the spectral gap. For a Gaussian target distribution and taking $\mathbf{D}(\theta) = \mathbf{I}$, the optimal choice of a constant $\mathbf{Q}$ to increase the spectral gap has been studied [42, 73].

3 Irreversible Jump Sampler

Although irreversible continuous dynamics increase the mixing of the overall stochastic dynamics, the discretized algorithm of (5) typically leads to bias due to discretization error as mentioned in Sec. 2.2. If we use Metropolis-Hastings (MH) to correct for this error, the whole process becomes reversible again. Instead, in this section we propose an irreversible jump sampler (I-Jump) that can be used in place of MH to correct for the discretization error in (5) while maintaining irreversibility.

Although the focus of this paper is on using this I-Jump algorithm with continuous dynamics proposals—as explored in Sec. 4—the I-Jump algorithm can be used with more traditional proposal distributions as a generic replacement for MH. Hence, in this section we turn our attention to the jump processes and consider an equivalent representation that enables more ready analysis of the properties of the process, and the development of an efficient irreversible jump process sampler.

3.1 Irreversible Jump Processes

A Markov jump process can be defined by the Kolmogorov forward equation

$$\frac{\partial}{\partial t}p(\mathbf{z}|\mathbf{y};t) \tag{10}$$
$$= \int_{\mathbb{R}^d} d\mathbf{x}\Big[W(\mathbf{z}|\mathbf{x})p(\mathbf{x}|\mathbf{y};t) - W(\mathbf{x}|\mathbf{z})p(\mathbf{z}|\mathbf{y};t)\Big],$$

where, similar to (2), $p(\mathbf{z}|\mathbf{y};t)$ is the probability density function of random vector $\mathbf{z}$ (parameterized by t) taking value $\mathbf{z}$ at time t conditional on it taking value of $\mathbf{y}$ at time 0. For the process specified by a generic bivariate integrable function $W(\mathbf{x}|\mathbf{z}) : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$, it is challenging to determine which choice of $W(\mathbf{x}|\mathbf{z})$ leads to a jump process with the correct stationary distribution. Even if one can construct such a W, it can be challenging to use W to sample a realization of the jump process. Instead, one often restricts attention to reversible processes and uses MH.

We revisit Markov jump processes under an equivalent but alternative representation defined in terms of two bivariate functions S and $A : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$. A simple set of constraints on S and A ensures that the target distribution $\pi(\mathbf{z})$ is the stationary distribution of the jump process. In particular, we consider

$$\frac{\partial}{\partial t}p(\mathbf{z}|\mathbf{y};t) = \int_{\mathbb{R}^d} \big(S(\mathbf{x},\mathbf{z}) + A(\mathbf{x},\mathbf{z})\big)\frac{p(\mathbf{x}|\mathbf{y};t)}{\pi(\mathbf{x})}\mathrm{d}\mathbf{x}$$
$$- \int_{\mathbb{R}^d} S(\mathbf{x},\mathbf{z})\frac{p(\mathbf{z}|\mathbf{y};t)}{\pi(\mathbf{z})}\mathrm{d}\mathbf{x}, \qquad (11)$$

where $S(\mathbf{x},\mathbf{z}) = S(\mathbf{z},\mathbf{x})$ is symmetric representing the reversible part of the process and $A(\mathbf{x},\mathbf{z}) = -A(\mathbf{z},\mathbf{x})$ is anti-symmetric defining the irreversible part. Based on the form of (11), as shown in Appendix B.1, we simply have to satisfy the following constraints in order to ensure that $\pi(\mathbf{z})$ is the stationary distribution:

1. $S(\mathbf{x},\mathbf{z})\pi^{-1}(\mathbf{x})$ and $A(\mathbf{x},\mathbf{z})\pi^{-1}(\mathbf{x})$ are bounded and integrable
2. $S(\mathbf{x},\mathbf{z}) + A(\mathbf{x},\mathbf{z}) > 0$
3. $\int_{\mathbb{R}^d} A(\mathbf{x},\mathbf{z})\mathrm{d}\mathbf{x} = 0$.

Discretizing (11) with Δt step size gives the following update rule:

$$p(\mathbf{z}|\mathbf{y};\Delta t) = \frac{\Delta t}{\pi(\mathbf{y})}\big(S(\mathbf{y},\mathbf{z}) + A(\mathbf{y},\mathbf{z})\big) \qquad (12)$$
$$+ \left[1 - \frac{\Delta t}{\pi(\mathbf{y})}\int_{\mathbb{R}^d} S(\mathbf{y},\mathbf{x})\mathrm{d}\mathbf{x}\right]\delta(\mathbf{z} - \mathbf{y}),$$

which defines a Markov transition kernel entirely by functions S and A: $P(\mathbf{y},\mathrm{d}\mathbf{z}) = p(\mathbf{z}|\mathbf{y};\Delta t)\mathrm{d}\mathbf{z}$. Since the image of P is a probability, it follows from the first constraint of S and A (under Eq. (11)) that $\Delta t \leq 1/\max\{||(S+A)/\pi||_1, ||(S+A)/\pi||_\infty\}$.

Since the jump operator has $\pi(\mathbf{z})$ as the stationary distribution assuming the constraints of S and A are satisfied, the transition probability of (12) defines a valid procedure for drawing samples from the target $\pi(\mathbf{z})$. In particular, over time Δt, state $\mathbf{y}$ transitions to state $\mathbf{z}$ with probability $\Delta t(S(\mathbf{y},\mathbf{z}) + A(\mathbf{y},\mathbf{z}))/\pi(\mathbf{y})$, and state $\mathbf{y}$ remains unchanged with probability

$$\left[1 - \frac{\Delta t}{\pi(\mathbf{y})}\int_{\mathbb{R}^d} S(\mathbf{y},\mathbf{x})\mathrm{d}\mathbf{x}\right]. \qquad (13)$$

We can further derive from (12) that $\Delta t \cdot A$ is the difference between the probability of a forward path and the backward path in the update procedure:

$$A(\mathbf{x},\mathbf{z}) = \frac{1}{2\Delta t}\big(\pi(\mathbf{y})p(\mathbf{z}|\mathbf{y};\Delta t) - \pi(\mathbf{z})p(\mathbf{y}|\mathbf{z};\Delta t)\big). \qquad (14)$$

From this, we clearly see how A determines the irreversibility of the process. In Sec. 3.3, we examine a practical algorithm for efficiently implementing such a procedure based on an accept-reject scheme analogous to the MH algorithm outlined in Algorithm 2. The important challenge we conquer is handling the irreversibility of the process arising from $A \neq 0$.

3.2 Reversible Samplers as Special Cases ($A = 0$)

As with past continuous-dynamic-based samplers, we now cast a set of past jump-process-based samplers into our framework.

Direct resampling Methods that sample directly from $\pi(\mathbf{z})$ take $S(\mathbf{y},\mathbf{z}) = \frac{1}{\Delta t}\pi(\mathbf{y})\pi(\mathbf{z})$ and $A(\mathbf{y},\mathbf{z}) = 0$. We can verify this by substituting into (12).

Metropolis-Hastings The very popular MH algorithm (Algorithm 2) falls into our framework taking

$$S(\mathbf{y},\mathbf{z}) = \frac{1}{\Delta t}\min\big(\pi(\mathbf{y})q(\mathbf{z}|\mathbf{y}), \pi(\mathbf{z})q(\mathbf{y}|\mathbf{z})\big), \qquad (15)$$

where $q(\mathbf{z}|\mathbf{y})$ is the conditional probability of transiting from $\mathbf{y}$ to $\mathbf{z}$ and $A(\mathbf{y},\mathbf{z}) = 0$.

Algorithm 2: Metropolis-Hastings Algorithm

---

for $t = 0, 1, 2 \cdots N_{iter}$ **do**

    sample $u \sim \mathcal{U}_{[0,1]}$

    propose $\mathbf{z}(*) \sim q(\mathbf{z}(*)|\mathbf{z}(t))$

    $\alpha\left(\mathbf{z}^{(t)}, \mathbf{z}(*)\right) = \min \left\{ 1, \dfrac{\pi\left(\mathbf{z}(*)\right) q(\mathbf{z}(t)|\mathbf{z}(*))}{\pi\left(\mathbf{z}(t)\right) q(\mathbf{z}(*)|\mathbf{z}(t))} \right\}$

    if $u < \alpha\left(\mathbf{z}(t), \mathbf{z}(*)\right)$, $\mathbf{z}(t+1) = \mathbf{z}(*)$

    else $\mathbf{z}(t+1) = \mathbf{z}(t)$

end

Algorithmically, we can define

$$\alpha(\mathbf{y}, \mathbf{z}) = \Delta t \cdot S(\mathbf{y}, \mathbf{z})/(\pi(\mathbf{y})q(\mathbf{z}|\mathbf{y}))$$
$$= \min \left(1, \frac{\pi(\mathbf{z})q(\mathbf{y}|\mathbf{z})}{\pi(\mathbf{y})q(\mathbf{z}|\mathbf{y})} \right), \tag{16}$$

such that both $q(\mathbf{z}|\mathbf{y})$ and $\alpha(\mathbf{y}, \mathbf{z})$ are less than or equal to 1. Then the update rule can be expressed as [15]:

$$p(\mathbf{z}, \Delta t|\mathbf{y}) = q(\mathbf{z}|\mathbf{y})\alpha(\mathbf{y}, \mathbf{z}) \tag{17}$$
$$+ \left[1 - \int_{\mathbb{R}^d} q(\mathbf{z}|\mathbf{y})\alpha(\mathbf{y}, \mathbf{z})\mathrm{d}\mathbf{x} \right] \delta(\mathbf{z} - \mathbf{y}).$$

When in state $\mathbf{y}$ at time t, we propose to jump to state $\mathbf{z}$ at $t + \Delta t$ with conditional probability $q(\mathbf{z}|\mathbf{y})$, realized via a random number generator that has a distribution according to $q(\mathbf{z}|\mathbf{y})$; we accept this proposal with probability $\alpha(\mathbf{y}, \mathbf{z})$ to ensure that the target distribution will be preserved under this procedure. Hence, the total probability of transiting from state $\mathbf{y}$ to $\mathbf{z}$ is $q(\mathbf{z}|\mathbf{y})\alpha(\mathbf{y}, \mathbf{z})$. Otherwise, we stay in state $\mathbf{y}$. We see that MH restricts our attention to reversible cases as $A(\mathbf{y}, \mathbf{z})$ is always set to be 0.

Summary of past samplers In the previously mentioned algorithms, and a majority of those used in practice, only reversible Markov jump processes ($A(\mathbf{z}, \mathbf{y})$ being 0) are considered. In Sec. 3.3, we explore the case where the process is irreversible, i.e., $A(\mathbf{z}, \mathbf{y}) \neq 0$.

3.3 Construction of a Practical Irreversible Jump Sampler

Analogous to the discussion of Sec. 2, there are two issues with designing samplers using Markov jump processes. One is the construction of transition probabilities, a task that has been alleviated in part by the new formulation of (12) in terms of $S(\mathbf{y}, \mathbf{z})$ and $A(\mathbf{y}, \mathbf{z})$ with simple constraints, though we still have to construct such probability density functions. Another is simulating the Markov process of (12). In all but the simplest cases, we might not be able to sample from the transition probability $\Delta t \cdot (S(\mathbf{y}, \mathbf{z}) + A(\mathbf{y}, \mathbf{z}))/\pi(\mathbf{y})$. These two issues are often intertwined posing challenges to the design of samplers. As mentioned in Sec. 2, the MH algorithm is often resorted to due to its ease of implementation. It separates the process of proposing a sample into two simple steps: (1) proposing a candidate according to a known conditional probability distribution $q(\mathbf{z}|\mathbf{y})$ and (2) accepting or rejecting the candidate according to a certain probability. An important drawback of the vanilla MH sampler, however, is that the reversibility of the jump process being designed can greatly restrict possible ways to increase the mixing of the Markov chain.

There have been previous efforts to break the restriction of reversibility in different cases. For example, the *non-reversible MH* algorithm adds a vorticity function to the MH procedure [6] while the *lifting method* makes two replica of the original state space with a skew detailed balance condition to facilitate irreversibility [67, 71]. The authors have shown examples of sampling special distributions, but it is unclear how to generalize these previous methods to handle a broad set of target distributions. See Sec. 5 for a detailed discussion of these and other methods. Here, we show how we can devise a practical and efficient irreversible jump process algorithm analogous to MH that can be applied to general targets; this procedure implicitly defines valid functions $S(\mathbf{y}, \mathbf{z})$ and $A(\mathbf{y}, \mathbf{z})$. In particular, just as MH corresponds to restricting the class of functions $W(\mathbf{z}|\mathbf{y})$, our algorithm also focuses in on particular instances of $A(\mathbf{y}, \mathbf{z})$, but importantly allows $A(\mathbf{y}, \mathbf{z}) \neq 0$ (i.e., irreversible processes). The value of this in practice is demonstrated in the experiments of Sec. 6.

A naïve approach A straightforward approach to revise the MH algorithm to make the antisymmetric function $A(\mathbf{y}, \mathbf{z})$ nonzero is to utilize different proposal distributions $f(\mathbf{z}|\mathbf{y})$ and $g(\mathbf{z}|\mathbf{y})$, instead of a single $q(\mathbf{z}|\mathbf{y})$. That is, the transition function of the MH algorithm in

(15) is changed to

$$F(\mathbf{y}, \mathbf{z}) = S(\mathbf{y}, \mathbf{z}) + A(\mathbf{y}, \mathbf{z})$$
$$= \frac{1}{\Delta t} \min \big(\pi(\mathbf{y}) f(\mathbf{z}|\mathbf{y}), \pi(\mathbf{z}) g(\mathbf{y}|\mathbf{z}) \big). \quad (18)$$

Here we are considering jump processes with $A(\mathbf{y}, \mathbf{z}) = \frac{1}{2} \big(F(\mathbf{y}, \mathbf{z}) - F(\mathbf{z}, \mathbf{y}) \big) \neq 0$, in contrast to what we saw for MH. By adjusting f and g, faster mixing rates can possibly be attained while maintaining a simple sampling procedure akin to that of MH (see Algorithm 2, but with f in place of q in the numerator and g in place of q in the denominator of the α calculation). The more f and g differ, the more irreversibility effect is incorporated in the design of the sampler. Functions f and g can even be selected to have non-overlapping support in the state space (as is chosen in our experiments), so that new proposals are guided in certain directions until being rejected, encouraging the algorithm to explore farther states. The primary issue with this construction is that $\int_{\mathbb{R}^d} A(\mathbf{y}, \mathbf{z}) d\mathbf{y} \neq 0$ in general, rendering the stationary distribution *not* the $\pi(\mathbf{z})$ that we desire. The question is how to design the anti-symmetric function $A(\mathbf{y}, \mathbf{z})$, such that $\int_{\mathbb{R}^d} A(\mathbf{y}, \mathbf{z}) d\mathbf{y} = 0$.

Lifting for sampling when $d = 1$ A simple modified approach is to follow an adjoint Markov process after being rejected by the original one. This is inspired by the *lifting* idea in discrete spaces [67, 71]. Importantly, this approach has $\pi(\mathbf{z})$ as the stationary distribution (marginalized over the auxiliary variable).

Algorithmically, this process introduces a one-dimensional, uniformly distributed discrete auxiliary variable $\mathbf{y}^p \in \{-1, 1\}$. We then define

$$\widetilde{f}(\mathbf{z}, \mathbf{z}^p | \mathbf{y}, \mathbf{y}^p) = \big(\mathbf{1}_{\mathbf{y}^p \geq 0} f(\mathbf{z}|\mathbf{y}) + \mathbf{1}_{\mathbf{y}^p < 0} g(\mathbf{z}|\mathbf{y}) \big)$$
$$\widetilde{g}(\mathbf{z}, \mathbf{z}^p | \mathbf{y}, \mathbf{y}^p) = \big(\mathbf{1}_{\mathbf{y}^p < 0} f(\mathbf{z}|\mathbf{y}) + \mathbf{1}_{\mathbf{y}^p \geq 0} g(\mathbf{z}|\mathbf{y}) \big),$$
$$(19)$$

where $f(\mathbf{z}|\mathbf{y})$ and $g(\mathbf{z}|\mathbf{y})$ are different conditional probability distributions, and $\mathbf{1}_{\mathcal{A}}$ is the indicator function for the set $\mathcal{A}$.

We modify the MH algorithm as described in Algorithm 3, where we update state $\mathbf{y}$ and the auxiliary

---

Algorithm 3: One-Directional I-Jump Sampler

randomly pick z^p from $\{1, -1\}$ with equal probability
for $t = 0, 1, 2 \cdots N_{iter}$ **do**
  sample $u \sim \mathcal{U}_{[0,1]}$
  if $z^p > 0$ **then**
    sample $\mathbf{z}(*) \sim f(\mathbf{z}(*)|\mathbf{z}(t))$
    $\alpha(\mathbf{z}(t), \mathbf{z}(*)) =$
    $\min \left\{ 1, \dfrac{\pi(\mathbf{z}(*)) \, g(\mathbf{z}(t)|\mathbf{z}(*))}{\pi(\mathbf{z}(t)) \, f(\mathbf{z}(*)|\mathbf{z}(t))} \right\}$
  end
  else
    sample $\mathbf{z}(*) \sim g(\mathbf{z}(*)|\mathbf{z}(t))$
    $\alpha(\mathbf{z}(t), \mathbf{z}(*)) =$
    $\min \left\{ 1, \dfrac{\pi(\mathbf{z}(*)) \, f(\mathbf{z}(t)|\mathbf{z}(*))}{\pi(\mathbf{z}(t)) \, g(\mathbf{z}(*)|\mathbf{z}(t))} \right\}$
  end
  if $u < \alpha(\mathbf{z}(t), \mathbf{z}(*))$,
  $\mathbf{z}(t+1) = \mathbf{z}(*); z^p(t+1) = z^p(t)$
  else $\mathbf{z}(t+1) = \mathbf{z}(t); z^p(t+1) = -z^p(t)$
end

---

variable $\mathbf{y}^p$ according to the following transition probability (as in our recipe of (12)):

$$p(\mathbf{z}, \mathbf{z}^p | \mathbf{y}, \mathbf{y}^p; \Delta t)$$
$$= \frac{\Delta t}{\pi(\mathbf{y})\pi^p(\mathbf{y}^p)} \mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p) \delta(\mathbf{z}^p - \mathbf{y}^p) \quad (20)$$
$$+ \left(1 - \frac{\Delta t}{\pi(\mathbf{y})\pi^p(\mathbf{y}^p)} \int_{\mathbb{R}^d} \mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{x}, -\mathbf{z}^p) d\mathbf{x} \right)$$
$$\cdot \delta(\mathbf{z}^p + \mathbf{y}^p) \delta(\mathbf{z} - \mathbf{y}),$$

in which $\mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ is defined using $\widetilde{f}$ and $\widetilde{g}$:

$$\mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$$
$$= \min \big(\pi(\mathbf{y})\pi^p(\mathbf{y}^p) \widetilde{f}(\mathbf{z}, \mathbf{z}^p | \mathbf{y}, \mathbf{y}^p),$$
$$\pi(\mathbf{z})\pi^p(\mathbf{z}^p) \widetilde{g}(\mathbf{y}, \mathbf{y}^p | \mathbf{z}, \mathbf{z}^p) \big).$$

This update rule can be understood as follows. With probability $\mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)/(\pi(\mathbf{y})\pi^p(\mathbf{y}^p))$, state $\mathbf{y}$ becomes state $\mathbf{z}$ while the auxiliary state $\mathbf{y}^p$ remains the same. Alternatively, with probability

$$\left[1 - \frac{1}{\pi(\mathbf{y})\pi^p(\mathbf{y}^p)} \int_{\mathbb{R}^d} \mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{x}, -\mathbf{z}^p) d\mathbf{x} \right],$$ no new state $(\mathbf{x}, \mathbf{y}^p)$ is accepted conditioning on currently being at state $(\mathbf{y}, \mathbf{y}^p)$. Instead, state $(\mathbf{y}, \mathbf{y}^p)$ is directly changed to state $(\mathbf{y}, -\mathbf{y}^p)$, leading to a different jump process in $\mathbf{y}$. An illustration of the update rule is shown in Fig. 1. It is worth noting that in the lifted space of
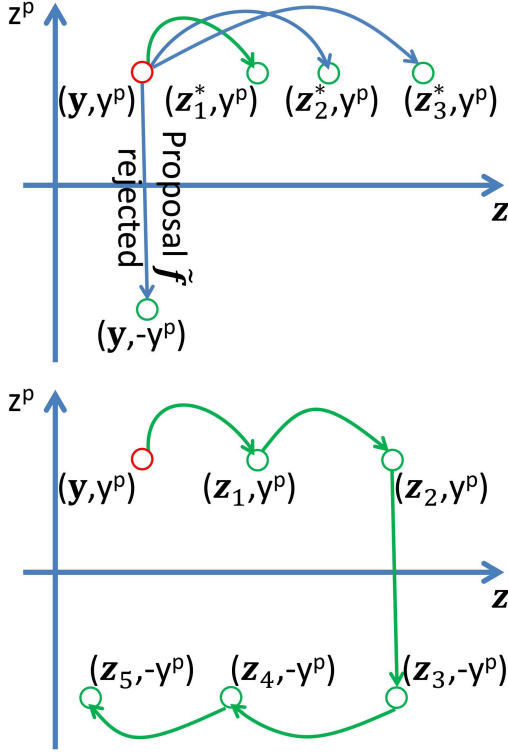
Fig. 1: Update rule starting from state $(\mathbf{y}, \mathbf{y}^p)$. *Left:* Several possible states $(\mathbf{z}^*, \mathbf{z}^p)$ that the algorithm could visit in the next step. Without resampling the auxiliary variables, $\mathbf{z}^p$ can only be $\mathbf{y}^p$ or $-\mathbf{y}^p$. *Right:* Assuming the algorithm visits $(\mathbf{z}_1, \mathbf{y}^p)$ as the next state to $(\mathbf{y}, \mathbf{y}^p)$ (indicated by the green arrow), a sample trajectory of states generated.

$(\mathbf{y}, \mathbf{y}^p)$, there are other choices for the probability of transition from $(\mathbf{y}, \mathbf{y}^p)$ to $(\mathbf{y}, -\mathbf{y}^p)$ (c.f. [71]). However, they all lead to a nonzero probability of staying in $(\mathbf{y}, \mathbf{y}^p)$, which reduces irreversibility.

From (14), we see that this proposed algorithm takes the anti-symmetric function $A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ to be

$$A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$$
$$= \frac{1}{2\Delta t}\Big(\pi(\mathbf{y})\pi^p(\mathbf{y}^p)p(\mathbf{z}, \mathbf{z}^p|\mathbf{y}, \mathbf{y}^p; \Delta t)$$
$$- \pi(\mathbf{z})\pi^p(\mathbf{z}^p)p(\mathbf{y}, \mathbf{y}^p|\mathbf{z}, \mathbf{z}^p; \Delta t)\Big) \qquad (21)$$

with $p(\mathbf{z}, \mathbf{z}^p|\mathbf{y}, \mathbf{y}^p; \Delta t)$ as in (20). To ensure correctness of the sampler, $A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ must satisfy (condition 3):

$$\int_{\mathbb{R}^{d+1}} A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p) \, \mathrm{d}\mathbf{y} \, \mathrm{d}\mathbf{y}^p = 0. \qquad (22)$$

In Appendix B.2, we prove that this is indeed the case for $A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ as in (21). The intuition is that the jump in the auxiliary variable introduces a circulative behavior to the whole process (see Fig. 1 for illustration). This circulation of probability flux is exactly balanced with the jumps in the original variable and the auxiliary variable. We also see in Fig. 1 that irreversibility introduces a directional effect (just like HMC introduces a direction of rotation). This algorithm is a generalization of the *guided walk Metropolis* method [26] and works well in one dimension, as we demonstrate in Sec. 6.2.1. In what follows, we generalize this idea to higher dimensions $d > 1$.

Moving to higher dimensions An irreversible sampler in $\mathbb{R}^d$ can be constructed as follows. We expand the state space by introducing a d^p-dimensional auxiliary variable $\mathbf{y}^p \in \mathbb{R}^{d^p}$ in the new state space $(\mathbf{y}, \mathbf{y}^p)$. The total probability can be designated as: $(\mathbf{y}, \mathbf{y}^p) \sim \pi(\mathbf{y})\pi^p(\mathbf{y}^p)$. We further impose symmetry on the auxiliary variables such that $\pi^p(\mathbf{y}^p) = \pi^p(-\mathbf{y}^p)$, and let

$$\widetilde{f}(\mathbf{z}, \mathbf{z}^p|\mathbf{y}, \mathbf{y}^p) \qquad (23)$$
$$= \prod_{i=1}^{n} \big(\mathbf{1}_{\mathbf{y}_i^p \geq 0} f_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p) + \mathbf{1}_{\mathbf{y}_i^p < 0} g_i(\mathbf{z}|\mathbf{y}, -\mathbf{y}^p)\big);$$
$$\widetilde{g}(\mathbf{z}, \mathbf{z}^p|\mathbf{y}, \mathbf{y}^p)$$
$$= \prod_{i=1}^{n} \big(\mathbf{1}_{\mathbf{y}_i^p < 0} f_i(\mathbf{z}|\mathbf{y}, -\mathbf{y}^p) + \mathbf{1}_{\mathbf{y}_i^p \geq 0} g_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)\big),$$

where n can be chosen by the user; and $f_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ and $g_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ are conditional probability distributions defined by the value of $\mathbf{y}^p$. Fitting this definition into the transition probability $p(\mathbf{z}, \mathbf{z}^p|\mathbf{y}, \mathbf{y}^p; \Delta t)$ in (20), the generalized update rule is defined and described in Algorithm 4. It is worth noting that the accept-reject step in the current setting is the same as in random-walk MH.

This definition of $\widetilde{f}$ and $\widetilde{g}$ is a direct generalization of the definition of (19) to the multi-dimensional case. Again, we have the anti-symmetric function $A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ as in (21). As

we prove in Appendix B.2, this construction has $\int_{\mathbb{R}^{d+d^p}} A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p) \, \mathrm{d}\mathbf{y} \, \mathrm{d}\mathbf{y}^p = 0$ even with our d^p-dimensional *continuous* auxiliary variables.

In summary, we can use (20) to devise a practical algorithm for sampling (I-Jump sampler of Algorithm 4). In particular, if we define $f_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ and $g_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ that are easy to sample from, then we can use the definitions of $\widetilde{f}$ and $\widetilde{g}$ in (23) to propose samples in the same way as the MH algorithm. Optionally, we can periodically resample $\mathbf{y}^p$ according to $\pi^p(\mathbf{y}^p)$.

One natural choice of $\widetilde{f}$ and $\widetilde{g}$ is to consider the random walk MH algorithm in our framework and add a nonzero $A(\mathbf{y}, \mathbf{z})$ to (15) while leaving $S(\mathbf{y}, \mathbf{z})$—the reversible part—unchanged. Then it can be proven that the resulting process will only provide a faster-mixing Markov chain [61, 34] (Note that this is in contrast to the naïve approach of (18) to modify MH). This can be achieved by choosing $n = 1$; $d^p = d$; $\pi^p(\mathbf{y}^p)$ equal to a restricted uniform distribution on $\{\mathbf{y}^p \mid ||\mathbf{y}^p||_2 = 1\}$; and f and g to be half-space Gaussian distributions. We write $f(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ as:

$$\mathbf{z} = \mathbf{y} + \boldsymbol{\eta} \cdot \mathrm{sgn}(<\boldsymbol{\eta}, \mathbf{y}^p>), \; \boldsymbol{\eta} \sim \mathcal{N}(0, \sigma^2 I),$$

$g(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ as:

$$\mathbf{z} = \mathbf{y} - \boldsymbol{\eta} \cdot \mathrm{sgn}(<\boldsymbol{\eta}, \mathbf{y}^p>), \; \boldsymbol{\eta} \sim \mathcal{N}(0, \sigma^2 I),$$

where $\mathrm{sgn}(x) = 2\mathbf{1}_{x \geq 0} - 1$. In other words,

$$f(\mathbf{z}|\mathbf{y}, \mathbf{y}^p) = \mathbf{1}_{<\mathbf{y}^p, \mathbf{z}-\mathbf{y}> \, \geq \, 0} \frac{2}{(2\pi\sigma)^{d/2}} e^{-\frac{1}{2\sigma}||\mathbf{z}-\mathbf{y}||_2^2},$$

$$g(\mathbf{z}|\mathbf{y}, \mathbf{y}^p) = \mathbf{1}_{<\mathbf{y}^p, \mathbf{z}-\mathbf{y}> \, < \, 0} \frac{2}{(2\pi\sigma)^{d/2}} e^{-\frac{1}{2\sigma}||\mathbf{z}-\mathbf{y}||_2^2}.$$

Intuitively, we are choosing a random direction $\mathbf{y}^p$ so that the proposal distribution is restricted in the half space aligned with $\mathbf{y}^p$. Then, $\mathbf{y}^p$ flips signs upon rejection. One can show that $S(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ for the above choice decomposed as $S(\mathbf{y}, \mathbf{z})S(\mathbf{y}^p, \mathbf{z}^p)$ and $S(\mathbf{y}, \mathbf{z})$ is indeed $\frac{1}{\Delta t} \min(\pi(\mathbf{y})q(\mathbf{z}|\mathbf{y}), \pi(\mathbf{z})q(\mathbf{y}|\mathbf{z}))$ for $q(\mathbf{z}|\mathbf{y})$ a Gaussian distribution centered at $\mathbf{y}$. That is, the reversible dynamic component determined by $S(\mathbf{y}, \mathbf{z})$ is exactly that of the random walk MH (See (15)). However, $A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p) \neq 0$ while satisfying (22). This construction maximizes the irreversibility without changing $S(\mathbf{y}, \mathbf{z})$. We test this setting in Sec. 6.3.

Alternatively, we can use one-sided distributions for proposals. For example, we can consider gamma distributions by setting $n = d^p = d$, and $f_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ as:

$$\mathbf{z}_i = \mathbf{y}_i + \gamma \cdot \mathbf{y}_i^p, \; \gamma \sim \Gamma(\alpha, \beta);$$

and $g_i(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ as:

$$\mathbf{z}_i = \mathbf{y}_i - \gamma \cdot \mathbf{y}_i^p, \; \gamma \sim \Gamma(\alpha, \beta);$$

and let $\pi^p(\mathbf{z}^p)$ to be a restricted uniform distribution on the set $\left\{ \mathbf{y}^p \middle| \frac{1}{N}||\mathbf{y}^p||_1 = 1 \right\}$. We will examine the benefits of this setting in synthetic experiments of Sec. 6.2.

It can be seen that in both settings there is always a direction of exploration, $\mathbf{y}^p$, that enjoys most of the benefits from irreversibility. In multiple dimensions, a favorable direction of exploration is often not clear. Hence, in experiments, we periodically resample the auxiliary variable $\mathbf{y}^p$ to explore all directions. Of course, there are still potential issues when d is large due to the fact that resampling of $\mathbf{y}^p$ may be inefficient for exploring the entire d dimensions. We leave this as a direction for future research. In the next Sec. 4, we focus on using the continuous Markov process to provide a favorable direction of exploration and take $n = d^p = 1$. Then $\mathbf{z}^p$ belongs to a binary set $\{-1, 1\}$, rendering Algorithm 4 the same as the simpler version, Algorithm 3, which is the continuous state space generalization of the *lifting* method [67, 71].

4 Irreversible Metropolis Adjusted Langevin Algorithm

In this section we discuss how to use the continuous dynamics of Sec. 2 as a proposal distribution in our I-Jump sampler of Sec. 3 (Algorithm 3), even when the continuous dynamics are not reversible. This procedure corrects for the discretization error introduced from simulating the continuous dynamics. Previously, methods such as the Metropolis-adjusted Langevin diffusion (MALA) and Riemannian Metropolis-adjusted Langevin diffusion (RMALA) [63, 74, 25] have only been proposed for reversible processes. These methods use one step integration of reversible SDEs to propose samples within an MH algorithm that accepts or rejects

---

Algorithm 4: I-Jump Sampler

---

for $t = 0, 1, 2 \cdots N_{iter}$ **do**

    optionally, periodically resample auxiliary variable $\mathbf{z}^p$ as $\mathbf{z}^p(t) \sim \pi^p(\mathbf{z}^p)$

    sample $u \sim \mathcal{U}_{[0,1]}$

    sample $\mathbf{z}(*) \sim \widetilde{f}(\mathbf{z}(*), \mathbf{z}^p(*)|\mathbf{z}(t), \mathbf{z}^p(t))$

    $\alpha\left(\mathbf{z}(t), \mathbf{z}^p(t), \mathbf{z}(*), \mathbf{z}^p(*)\right) = \min\left\{1, \dfrac{\pi\left(\mathbf{z}(*)\right)\pi^p\left(\mathbf{z}^p(*)\right)\widetilde{g}\left(\mathbf{z}(t), \mathbf{z}^p(t)|\mathbf{z}(*), \mathbf{z}^p(*)\right)}{\pi\left(\mathbf{z}(t)\right)\pi^p\left(\mathbf{z}^p(t)\right)\widetilde{f}\left(\mathbf{z}(*), \mathbf{z}^p(*)|\mathbf{z}(t), \mathbf{z}^p(t)\right)}\right\}$

    if $u < \alpha\left(\mathbf{z}(t), \mathbf{z}^p(t), \mathbf{z}(*), \mathbf{z}^p(*)\right),\quad (\mathbf{z}(t+1), \mathbf{z}^p(t+1)) = (\mathbf{z}(*), \mathbf{z}^p(t))$

    else $\quad(\mathbf{z}(t+1), \mathbf{z}^p(t+1)) = (\mathbf{z}(t), -\mathbf{z}^p(t))$

end

---

the proposal. We extend these methods to include proposals from any SDE in the form of (3) (any SDE with a mild integrability condition), without the requirement of reversibility. Furthermore, our I-Jump sampler allows the resulting overall process to still be irreversible.

In Sec. 4.1, we review the (reversible) MALA algorithm. We then delve into our proposed irreversible MALA (I-MALA) algorithm in two stages. In Sec. 4.2, we first discuss how one can use the continuous Markov processes of (3) as a proposal distribution in the I-Jump sampler of Sec. 3 and get an acceptance rate equal to 1 when the continuous Markov process is simulated exactly. In Sec. 4.3, we use a one-step discretized simulation of the continuous Markov process as the proposal distribution and specify the details of the resulting practical I-MALA algorithm. In the experiments of Sec. 6, we show that I-MALA can generate better results in terms of rapid and efficient exploration of a distribution than MALA or HMC.

4.1 Metropolis Adjusted Langevin Algorithm (MALA)

Since the MALA algorithm is a special case of the RMALA algorithm (with $\mathbf{D}(\mathbf{z})$ taken to be constant), we will simply introduce RMALA in this section. The RMALA algorithm takes $\mathbf{z} = \theta$ and constructs the proposal distribution $q(\theta(*)|\theta(t))$ in the MH algorithm (Algorithm 2) to be the discretized Riemannian Langevin dynamics:

$$\theta(*) \leftarrow \theta(t) + \eta(t)$$
$$+ \Delta t \cdot [-\mathbf{G}(\theta(t))^{-1}\nabla U(\theta(t)) + \Gamma^{\mathbf{D}}(\theta(t))],$$
$$\eta(t) \sim \mathcal{N}(0, 2\Delta t\mathbf{G}(\theta(t))^{-1}). \qquad (24)$$

Here, the diffusion matrix $\mathbf{D}(\theta(t)) = \mathbf{G}(\theta(t))^{-1}$ and $\Gamma_i^{\mathbf{D}}(\theta) = \sum_j \dfrac{\partial \mathbf{D}_{ij}(\theta)}{\partial \theta_j}$. In the original MALA algorithm [63], $\mathbf{G} = \mathbf{I}$, whereas in the RMALA algorithm [74, 25], $\mathbf{G}$ is taken to be the Fisher information metric. Therefore, the resulting transition probability $q(\theta(*)|\theta(t))$ in the MH Algorithm 2 is:

$$q(\theta(*)|\theta(t)) = \mathcal{N}\left\{\theta(*)\big|\mu(\theta(t), \Delta t), 2\Delta t\mathbf{G}(\theta(t))^{-1}\right\}, \qquad (25)$$

where

$$\mu(\theta(t), \Delta t)$$
$$= \theta(t) + \Delta t \cdot [-\mathbf{G}(\theta(t))^{-1}\nabla U(\theta(t)) + \Gamma^{\mathbf{D}}(\theta(t))].$$

This algorithm provides a sampling procedure to exactly simulate the reversible continuous Markov dynamics. And in doing so, gradient information is used to help the sampler efficiently explore the target distribution.

Use of the MH procedure inevitably restricts the sampler to be reversible. From (25), we also observe that only reversible Langevin dynamics are used in the MALA algorithm. Although irreversible dynamics can be used in (24), as will be discussed in the Sec. 4.2, the acceptance rate would decrease with the increase of irreversibility since irreversibility increases the difference between $q(\theta(*)|\theta(t))$ and the reverse proposal $q(\theta(t)|\theta(*))$. We explore this further in Sec. 4.2 (see (31)).

Comparison with HMC HMC expands the sampling space to $\mathbf{z} = (\theta, r)$ and simulates the Hamiltonian dynamics (6) discussed in Sec. 2.3 with a neutral integra-

tion scheme. When the mass matrix $\mathbf{M}$ in (6) is constant, a leap-frog step integration is often applied:

$$
\begin{cases}
r_{t+1/2} \leftarrow r_t - \epsilon_t/2\nabla U(\theta_t) \\
\theta_{t+1} \leftarrow \theta_t + \epsilon_t \mathbf{M}^{-1} r_{t+1/2} \\
r_{t+1} \leftarrow r_{t+1/2} - \epsilon_t \nabla U(\theta_{t+1}).
\end{cases}
\tag{26}
$$

When $\mathbf{M} = \mathbf{M}(\theta)$ is adaptive, a more involved symplectic integrator is often needed to preserve volume in $\mathbf{z}$. To obtain valid samples, the momentum variable r is resampled (an example of a jump process sampler described in "direct resampling" of Sec. 3.2) for ergodicity. An MH accept-reject step is applied with acceptance rate calculated using the volume preserving property of the Hamiltonian dynamics. While MALA quickly converges towards a mode and diffusively explores around it, HMC excels at deterministically traversing along level sets of the Hamiltonian

$$
H(\theta, r) = U(\theta) + \frac{1}{2}r^T \mathbf{M}^{-1} r.
$$

4.2 General SDE Proposals under Small Step Size Limit

Our ultimate goal is to use the stochastic dynamics of (3) to propose samples in the framework of Algorithm 3. In practice, we need to simulate from the discretized SDE of (5). Before analyzing this case, we first examine what would happen if we could *exactly* simulate the SDE of (3).

Here, we imagine using the transition probability density $P(\mathbf{z}|\mathbf{y}; t)$ of the continuous Markov process to construct a particular case of Algorithm 3. We take $f(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ in (23) to be equal to $P(\mathbf{z}|\mathbf{y}; t)$ defined via an exact solution to the SDE, starting at $\mathbf{y}$:

$$
\begin{aligned}
d\mathbf{z} = & \Big[-\big(\mathbf{D}(\mathbf{z}) + \mathbf{Q}(\mathbf{z})\big)\nabla H(\mathbf{z}) + \varGamma(\mathbf{z})\Big] dt \\
& + \sqrt{2\mathbf{D}(\mathbf{z})}d\mathbf{W}(t),
\end{aligned}
\tag{27}
$$

where $\varGamma_i(\mathbf{z}) = \sum_j \frac{\partial}{\partial \mathbf{z}_j}(\mathbf{D}_{ij}(\mathbf{z}) + \mathbf{Q}_{ij}(\mathbf{z}))$.

For the reverse proposal $g(\mathbf{z}|\mathbf{y}, \mathbf{y}^p)$ in (23), we use the adjoint process $P^\dagger(\mathbf{z}|\mathbf{y}; t)$, inverting the irreversible dynamics via $\mathbf{Q}(\mathbf{z}) \to -\mathbf{Q}(\mathbf{z})$ [47]:

$$
\begin{aligned}
d\mathbf{z} = & \Big[-\big(\mathbf{D}(\mathbf{z}) - \mathbf{Q}(\mathbf{z})\big)\nabla H(\mathbf{z}) + \widetilde{\varGamma}(\mathbf{z})\Big] dt \\
& + \sqrt{2\mathbf{D}(\mathbf{z})}d\mathbf{W}(t),
\end{aligned}
\tag{28}
$$

where $\widetilde{\varGamma}_i(\mathbf{z}) = \sum_j \frac{\partial}{\partial \mathbf{z}_j}(\mathbf{D}_{ij}(\mathbf{z}) - \mathbf{Q}_{ij}(\mathbf{z}))$.

Theorem 2 *For the Markov processes*

$$
P(\mathbf{z}|\mathbf{y}; t) \text{ and } P^\dagger(\mathbf{z}|\mathbf{y}; t)
$$

defined by the SDEs of (27) *and* (28) *through Itô integration, the following equality holds:*

$$
\frac{P(\mathbf{z}|\mathbf{y}; t)}{P^\dagger(\mathbf{y}|\mathbf{z}; t)} = \frac{\pi(\mathbf{z})}{\pi(\mathbf{y})}.
\tag{29}
$$

The proof is in Appendix C. Using Theorem 2, we have

$$
\begin{aligned}
\alpha(\mathbf{y}, \mathbf{z}) &= \min\left\{1, \frac{\pi(\mathbf{z})P^\dagger(\mathbf{y}|\mathbf{z}; t)}{\pi(\mathbf{y})P(\mathbf{z}|\mathbf{y}; t)}\right\} \\
&= 1.
\end{aligned}
\tag{30}
$$

Even though in Sec. 2.1 we saw that SDEs of the form in (3) have $\pi(\mathbf{z})$ as the invariant distribution, it is not immediately obvious that using this SDE as a proposal in Algorithm 3 would lead to an acceptance rate of 1. In fact, it might be tempting to directly plug $P(\mathbf{z}|\mathbf{y}; t)$ into the MH Algorithm 2. However, that would result in a MH acceptance rate:

$$
\begin{aligned}
\alpha^{MH}(\mathbf{y}, \mathbf{z}) &= \min\left\{1, \frac{\pi(\mathbf{z})P(\mathbf{y}|\mathbf{z}; t)}{\pi(\mathbf{y})P(\mathbf{z}|\mathbf{y}; t)}\right\} \\
&\neq 1.
\end{aligned}
\tag{31}
$$

And the more irreversibility is introduced, the less the acceptance rate α^{MH} will be in the MH algorithm. This gap between α^{MH} and 1 was first discovered in the statistical mechanics literature and relates to the "house keeping heat" [16, 28, 47].

(30) also gives us insight into the fact that using more accurate numerical integrators could lead to higher acceptance rates. In Sec. 4.3, we analyze the accept-reject scheme for the simple first-order integration of (5) with finite step size Δt.

4.3 I-MALA via I-Jump Correction

Since in practice we rely on finite step sizes $\Delta t > 0$, there will be numerical error resulting in $\frac{P(\mathbf{z}(*)|\mathbf{z}(t); \Delta t)}{P^\dagger(\mathbf{z}(t)|\mathbf{z}(*); \Delta t)}$ differing from $\frac{\pi(\mathbf{z}(*))}{\pi(\mathbf{z}(t))}$. We now

propose an irreversible generalization of the MALA algorithm to correct for these errors. We make use of Algorithm 3 and take a general SDE and its adjoint process defined in Sec. 4.2 to propose samples using a one-step numerical integration (as in MALA). Because we have the local gradient information in the SDEs to guide us, the direction of the exploration is determined. So, we simply use a 1-dimensional discrete auxiliary variable $\mathbf{y}^p$, and thus rely on Algorithm 3 instead of the more general Algorithm 4. We call the resulting algorithm the *I-MALA* method.

Assuming a one-step numerical integration using a Δt period of time, the discretization of the SDE of (27) leads to

$$P(\mathbf{z}|\mathbf{y}; \Delta t) = \mathcal{N}\{\mathbf{z}|\mu(\mathbf{y}, \Delta t), 2\Delta t \cdot \mathbf{D}(\mathbf{y})\}, \qquad (32)$$

where

$$\mu(\mathbf{y}, \Delta t) = \mathbf{y} + \left[-\big(\mathbf{D}(\mathbf{y}) + \mathbf{Q}(\mathbf{y})\big)\nabla H(\mathbf{y}) + \Gamma(\mathbf{y}) \right]\Delta t,$$

$$\Gamma_i(\mathbf{z}) = \sum_j \frac{\partial}{\partial \mathbf{z}_j}\big(\mathbf{D}_{ij}(\mathbf{z}) + \mathbf{Q}_{ij}(\mathbf{z})\big). \qquad (33)$$

Importantly, this allows us to compute $f(\mathbf{z}(*)|\mathbf{z}(t)) = P(\mathbf{z}(*)|\mathbf{z}(t); \Delta t)$ in Algorithm 3. The corresponding calculation for the adjoint process with the SDE in (28) is:

$$P^\dagger(\mathbf{z}|\mathbf{y}; \Delta t) = \mathcal{N}\{\mathbf{z}|\mu^\dagger(\mathbf{y}, \Delta t), 2\Delta t \cdot \mathbf{D}(\mathbf{y})\}, \qquad (34)$$

where

$$\mu^\dagger(\mathbf{y}, \Delta t) = \mathbf{y} + \left[-\big(\mathbf{D}(\mathbf{y}) - \mathbf{Q}(\mathbf{y})\big)\nabla H(\mathbf{y}) + \widetilde{\Gamma}(\mathbf{y}) \right]\Delta t,$$

$$\widetilde{\Gamma}_i(\mathbf{z}) = \sum_j \frac{\partial}{\partial \mathbf{z}_j}\big(\mathbf{D}_{ij}(\mathbf{z}) - \mathbf{Q}_{ij}(\mathbf{z})\big). \qquad (35)$$

This allows us to compute

$$g(\mathbf{z}(*)|\mathbf{z}(t)) = P^\dagger(\mathbf{z}(*)|\mathbf{z}(t); \Delta t).$$

The resulting I-MALA algorithm is summarized in Algorithm 5.

Algorithm 5: I-MALA

randomly pick z^p from $\{1, -1\}$ with equal probability
for $t = 0, 1, 2 \cdots N_{iter}$ **do**
  sample $u \sim \mathcal{U}_{[0,1]}$
  if $z^p > 0$ **then**
    sample $\eta_t \sim \mathcal{N}(0, 2\epsilon_t \mathbf{D}(\mathbf{z}_t))$
    $\mathbf{z}(*) \leftarrow \mathbf{z}_t -$
    $\epsilon_t \left[\big(\mathbf{D}(\mathbf{z}_t) + \mathbf{Q}(\mathbf{z}_t)\big)\nabla H(\mathbf{z}_t) + \Gamma(\mathbf{z}_t)\right] + \eta_t$
    $\alpha(\mathbf{z}(t), \mathbf{z}(*)) =$
    $\min\left\{1, \dfrac{\pi(\mathbf{z}(*))\, P^\dagger(\mathbf{z}(t)|\mathbf{z}(*); \Delta t)}{\pi(\mathbf{z}(t))\, P(\mathbf{z}(*)|\mathbf{z}(t); \Delta t)}\right\}$
  end
  else
    sample $\eta_t \sim \mathcal{N}(0, 2\epsilon_t \mathbf{D}(\mathbf{z}_t))$
    $\mathbf{z}(*) \leftarrow \mathbf{z}_t -$
    $\epsilon_t \left[\big(\mathbf{D}(\mathbf{z}_t) - \mathbf{Q}(\mathbf{z}_t)\big)\nabla H(\mathbf{z}_t) + \widetilde{\Gamma}(\mathbf{z}_t)\right] + \eta_t$
    $\alpha(\mathbf{z}(t), \mathbf{z}(*)) =$
    $\min\left\{1, \dfrac{\pi(\mathbf{z}(*))\, P(\mathbf{z}(t)|\mathbf{z}(*); \Delta t)}{\pi(\mathbf{z}(t))\, P^\dagger(\mathbf{z}(*)|\mathbf{z}(t); \Delta t)}\right\}$
  end
  if $u < \alpha(\mathbf{z}(t), \mathbf{z}(*))$,
  $\mathbf{z}(t+1) = \mathbf{z}(*);\ z^p(t+1) = z^p(t)$
  else $\quad \mathbf{z}(t+1) = \mathbf{z}(t);\ z^p(t+1) = -z^p(t)$
end

We know from Sec. 4.2 that in the small Δt limit,

$$\alpha(\mathbf{z}(t), \mathbf{z}(*)) = \min\left\{1, \frac{P^\dagger(\mathbf{z}(t)|\mathbf{z}(*))}{P(\mathbf{z}(*)|\mathbf{z}(t))} \cdot \frac{\pi(\mathbf{z}(*))}{\pi(\mathbf{z}(t))}\right\}$$
$$\to 1. \qquad (36)$$

From this result, we see that there seems to be a step-size/acceptance-rate tradeoff. As mentioned in Sec. 4.2, a higher-order numerical scheme could potentially increase the acceptance rate with the same step size [11, 9, 40]. We leave this as a direction for future research.

5 Related Work

There have been previous efforts to construct irreversible Markov processes for sampling. One example is using continuous dynamics to achieve this goal, which has been studied extensively. Theoretically, one can make use of Hamiltonian or generalized Hamiltonian dynamics to introduce irreversibility into the sampling procedure [30, 31, 60, 22]. There have been samplers that utilize specific irreversible continuous dy-

namics stemming from physical systems, such as underdamped Langevin [13] and Nosé-Hoover [20, 64, 41] dynamics and their generalizations [45, 43], although irreversibility was not the emphasis in these works. As described in Sec. 2.1, any dynamic process that has a nonzero $\mathbf{Q}$ matrix can be used to devise an irreversible sampler within our framework. The problem, however, is that simulating the continuous Markov processes using the discretized system typically leads to the introduction of bias due to discretization error (see Sec. 2.2). One option to correct for this bias is to introduce a MH step, but of course this yields the whole process reversible again. In all the previously mentioned works, no MH correction is used and instead a small, finite stepsize is used and some resulting bias tolerated. Our I-MALA method provides a mechanism for handling the discretization error while maintaining an overall irreversible process.

One alternative way of introducing irreversibility into samplers without causing additional bias is to combine Hamiltonian dynamics with an ergodic stochastic dynamics that can be exactly simulated. For Hamiltonian dynamics with quadratic kinetic energy, flipping the sign of the momentum variable is equivalent to following the adjoint process ($\mathbf{Q} \rightarrow -\mathbf{Q}$) in Algorithm 3 because of its special structure. This is a special case of the I-Jump algorithm where the proposal is via simulating Hamiltonian dynamics. Therefore, one can simulate the ergodic dynamics exactly, simulate the Hamiltonian dynamics, then change the sign of the momentum variable after rejection (instead of resampling it) to form an irreversible sampler. For example, in [29, 52], Langevin dynamics over the momentum variable is integrated exactly and combined with the Hamiltonian dynamics to form an irreversible sampler (SOL-HMC) [52]. During the review of this paper, it came to our attention that a recently released manuscript [56] proposes two new irreversible algorithms. The first uses a special case of I-MALA with discussions on novel implicit integration schemes. The second one (Hybrid MALA) uses MALA to simulate reversible dynamics exactly and combine with Hamiltonian dynamics to obtain an irreversible sampler. In this version of our manuscript, we compare to both SOL-HMC and hybrid MALA in the experiments of Sec. 6.3.

Only recently have researchers constructed irreversible jump processes that form valid sampling procedures. In the *non-reversible MH* algorithm [6], a vorticity function (or matrix) is added to the MH procedure. Then, the difficulty of construction is translated to defining a valid vorticity function, similar to the difficulty of defining the antisymmetric function $A(\mathbf{y}, \mathbf{z})$. For the multivariate Gaussian distribution, the author discretized an irreversible Ornstein-Uhlenbeck process to obtain a suitable vorticity function. The *lifting method* [67, 71] makes a replica of the original state space ($\mathbb{R}^d \times \{-1, 1\}$) to facilitate irreversibility in the sampling procedure. A skew detailed balance condition is imposed to ensure a valid antisymmetric function $A(\mathbf{y}, \mathbf{z})$ in the expanded state space. The authors showed an example of applying the method to spin models. For both the *non-reversible MH* and *lifting* methods, it has not been clear how to come up with practical, easy-to-construct algorithms to handle a broad set of target distributions. Our I-Jump sampler incorporates both ideas: lifting the state space (to $\mathbb{R}^d \times \mathbb{R}^{d^p}$) and using an irreversible accept-reject procedure similar to the non-reversible MH algorithm. Combination of the two ideas yields a simple MCMC procedure that generalizes to arbitrary target distributions.

The combined approach of using both continuous dynamics and jump processes has recently been proposed for constructing irreversible samplers. The *bouncy particle* [10] and *Zig-Zag* [8, 7] samplers use deterministic dynamics (irreversible in nature) combined with a Poisson process to create valid MCMC procedures. These two methods use continuous dynamics to guide a Poisson jump process with an inhomogeneous rate (or intensity) to ensure the invariance of the target distribution. In practice, a Poisson thinning step is often required to generate Poisson processes with inhomogeneous rates, hence posing further constraints on the target distribution (e.g., prior knowledge of global lower bounds for the norm of the gradient or Hessian of the negative log posterior). Our I-MALA algorithm avoids the difficulty of sampling from a Poisson process. Additionally, we end up with an algorithm that is a simple modification of vanilla MH, making it straightforward to use and plug in to existing algorithmic frameworks.

6 Experiments

In this section, we first examine the correctness and attributes of our I-Jump sampler (Algorithm 4). We consider various simulated scenarios, including the challenging cases of heavy tailed, multimodal, and correlated distributions.

We then explore the I-MALA algorithm (Algorithm 5) and compare it against numerous baselines.

6.1 Visual Comparison of Samplers

We first perform a qualitative comparison between the random walk MH algorithm, our I-Jump sampler, and the I-MALA algorithm to provide insights into their differences. It is demonstrated in Fig. 2 that the standard MH sampler jumps around randomly, but does so within a local region of the previous sample and irrespective of previous (directions of) jumps, leading to slow exploration of the distribution. In contrast, our irreversible counterpart (here using gamma proposals) more rapidly traverses the distribution by following the direction of the previous jump, until being rejected. Finally, the I-MALA algorithm provides an even smoother trajectory by using continuous dynamics in place of independent gamma proposals.

Having visually examined the differences between the samplers to gain intuition, in what follows we provide a more quantitative analysis of the proposed samplers.

6.2 Synthetic Experiments for I-Jump Sampler

In the following experiments we examine the I-Jump sampler on various challenging synthetic distributions. For this section, we consider gamma proposal. That is, as mentioned in Sec. 3.3, we take $f_i\left(\mathbf{z}(*)|\mathbf{z}(t), \mathbf{z}^p(t)\right)$ according to $\mathbf{z}_i(*) = \mathbf{z}_i(t) + \gamma \mathbf{z}_i^p(t)$ with $\gamma \sim \Gamma(\alpha, \beta)$ and $g_i\left(\mathbf{z}(*)|\mathbf{z}(t), \mathbf{z}^p(t)\right)$ according to $\mathbf{z}_i(*) = \mathbf{z}_i(t) - \gamma \mathbf{z}_i^p(t)$. We set $\pi(\mathbf{z}^p)$ to be a restricted uniform distribution on the set $\left\{\mathbf{z}^p | \frac{1}{N} ||\mathbf{z}^p||_1 = 1 \right\}$.

The hyperparameters α and β are chosen using a generic procedure. We take the gamma shape parameter to be $\alpha = 1.1$, and change the rate parameter β

approximately as $\beta \propto \sqrt{V}$ with V is the volume of the region we would like to explore. Further details are in Appendix D.1.

6.2.1 1D Heavy-tailed Distribution

We start by considering the task of sampling from 1D normal and log-normal distributions, the latter of which is a heavy-tailed distribution. The motivation for considering the simple 1D normal distribution is to validate the correctness of the sampler and to serve as a comparison relative to the heavy-tailed setting. We compare performance to a Gaussian random walk MH algorithm. The results are shown in Fig. 3. We also compare against an MH algorithm using a symmetrized gamma proposal distribution:

$$(\mathbf{z}(*) - \mathbf{z}(t)) \sim \frac{1}{2}\left(f\left(\mathbf{z}(*)|\mathbf{z}(t)\right) + g\left(\mathbf{z}(*)|\mathbf{z}(t)\right)\right)$$

to more closely match the gamma proposal of the I-Jump sampler.

We compute the decrease of total variation distance and autocorrelation function to compare the performance of the samplers in Fig. 3. The total variation distance is computed through discretizing the space and comparing the histogram of the samples and the true target distribution. We find the I-Jump sampler to have the best performance even in the simple normal target distribution case. In particular, I-Jump can decrease autocorrelation without increasing the rejection rate (the rejection rate of all three methods are similar). Intuitively, this result can be understood from Fig. 1: the irreversible algorithm leads to further exploration in one direction before circling back. (Also, see Fig. 2.)

For the heavy-tailed distribution, similar behavior is observed: I-Jump converges to the desired distribution faster because its samples decorrelate more rapidly as a function of runtime.

6.2.2 Multimodal Distributions

2D Bimodal distributions We use our I-Jump sampler to sample increasingly challenging bimodal distributions in 2D, $\pi(z_1, z_2) = 2(z_1^2 - \tau)^2 - 0.2z_1 - 5z_1^2 + 5z_2^2$ for $\tau = 0.5, 1.0,$ and 1.5, as displayed in Fig. 4. Based on the results of Sec. 6.2.1, we simply compare against
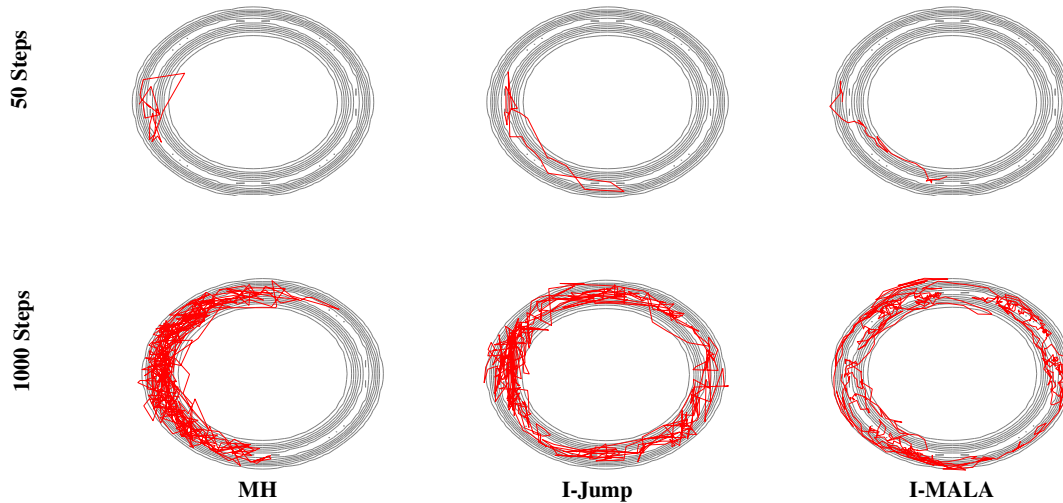
Fig. 2: *Top row:* Trajectory of first 50 steps of (left) MH algorithm using Gaussian random walk proposals, (middle) I-Jump algorithm with gamma proposals and (right) I-MALA algorithm. *Bottom row:* Similarly for the first 1000 steps of the algorithms.

the Gaussian random walk MH sampler and drop the symmetrized gamma proposal case. In Fig. 4 we see that the I-Jump sampler significantly outperforms the random walk MH algorithm. Intuitively, this is facilitated by the greater traversing ability of the I-Jump sampler so that with the same acceptance rate, the I-Jump sampler can explore more possible states than the reversible sampler, and have greater chance of transiting into another mode.

One way to quantify this effect is in terms of *escape time* from local modes, which we summarize in Table 1. We see that the I-Jump jump sampler has escape times orders of magnitude lower. Furthermore, these escape times increase at a much smaller rate as the local modes become more concentrated, indicating much more rapid mixing between modes.

2D Multimodal distributions We also tested our method against a recently considered multimodal setting [66]. In the first setting considered, the target distribution is highly multimodal in 2D with unevenly distributed modes. Furthermore, the high mass modes have smaller radii of variation. In the second setting

considered, these modes are highly concentrated and well separated, which is an extremely challenging setting for most samplers. See Figs. 5 and 6. In [66], a *repulsing-attracting Metropolis (RAM)* sampler was proposed with a structure specifically designed to efficiently handle these types of multimodal distributions. We use this as a gold-standard comparison, since this method was already shown to outperform parallel tempering and alternatives [38] in this setting.

We focus our performance analysis on the decay speed of the autocorrelation function (ACF). This can be understood by taking the Gaussian random walk MH algorithm as an example: Although the Gaussian random walk MH algorithm seems to perform well in terms of convergence of total variation distance, this effect is based on exploring one mode really well in a short period of time, instead of making more distant moves to explore other modes. In contrast, the ACF better characterizes the exploration of the samples through the whole space.

Our results are summarized in Figs. 5 and 6 for each of the two simulated multimodal scenarios. In the first scenario, our sampler outperforms both MH and
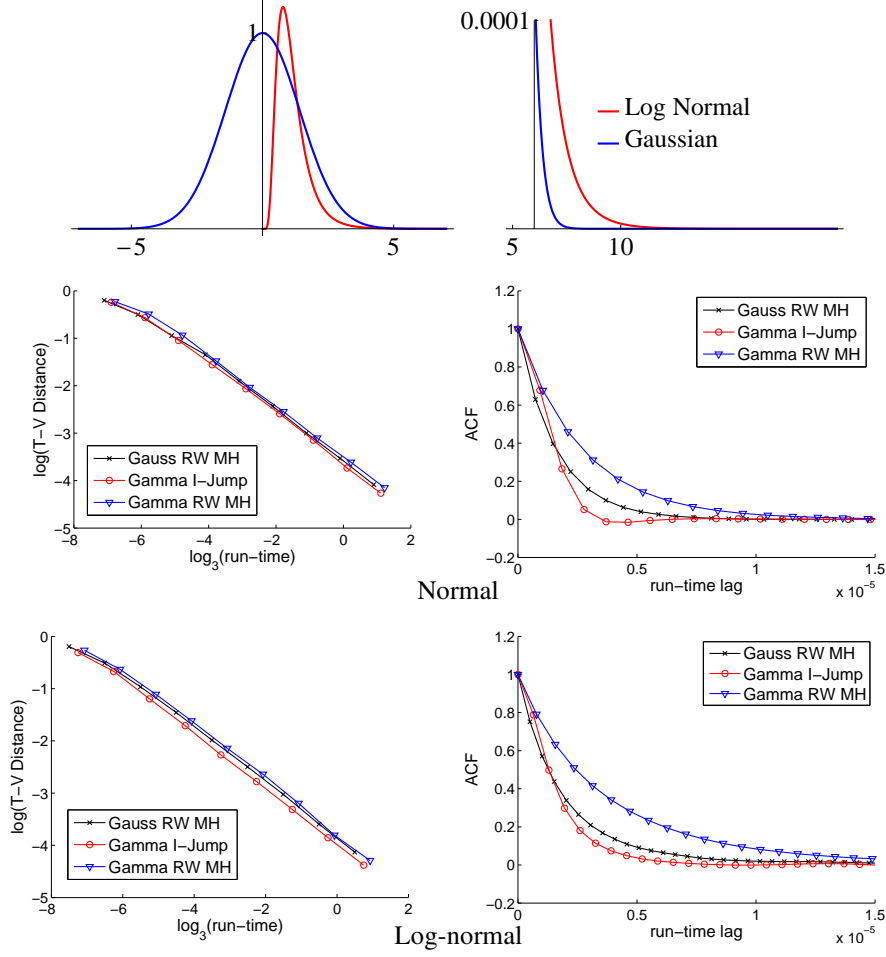
Fig. 3: *Top row:* (Left) Normal and log-normal target distributions, and (right) zoom in of the tails. *Middle row:* Results for normal target in terms of log total variation distance (T-V distance) vs. log runtime (left) and ACF vs. lag in runtime (right). *Bottom row:* Analogous plots for log normal target. Comparisons are made among the I-Jump sampler with gamma proposals (Gamma I-Jump), random walk MH algorithm with Gaussian proposals (Gauss RW MH), and random walk MH algorithm with symmetrized gamma proposals (Gamma RW MH). Runtime is measured in seconds.

RAM. In the second scenario, where we have highly concentrated and separated modes, the RAM method tailored to this scenario slightly outperforms our approach. Overall, however, the I-Jump sampler provides surprisingly good performance despite not having been designed specifically for this setting.

6.2.3 Correlated Distribution

We now test the correctness and attributes of our algorithm on a highly correlated (moon-shaped) target distribution, where $\pi(z_1, z_2) = z_1^4/10 + (4(z_2 + 1.2) - z_1^2)^2/2$. In terms of number of iterations, the I-Jump sampler decorrelates and converges to the target distribution faster. However, in terms of runtime, the I-Jump sampler does not perform as well as random walk MH,
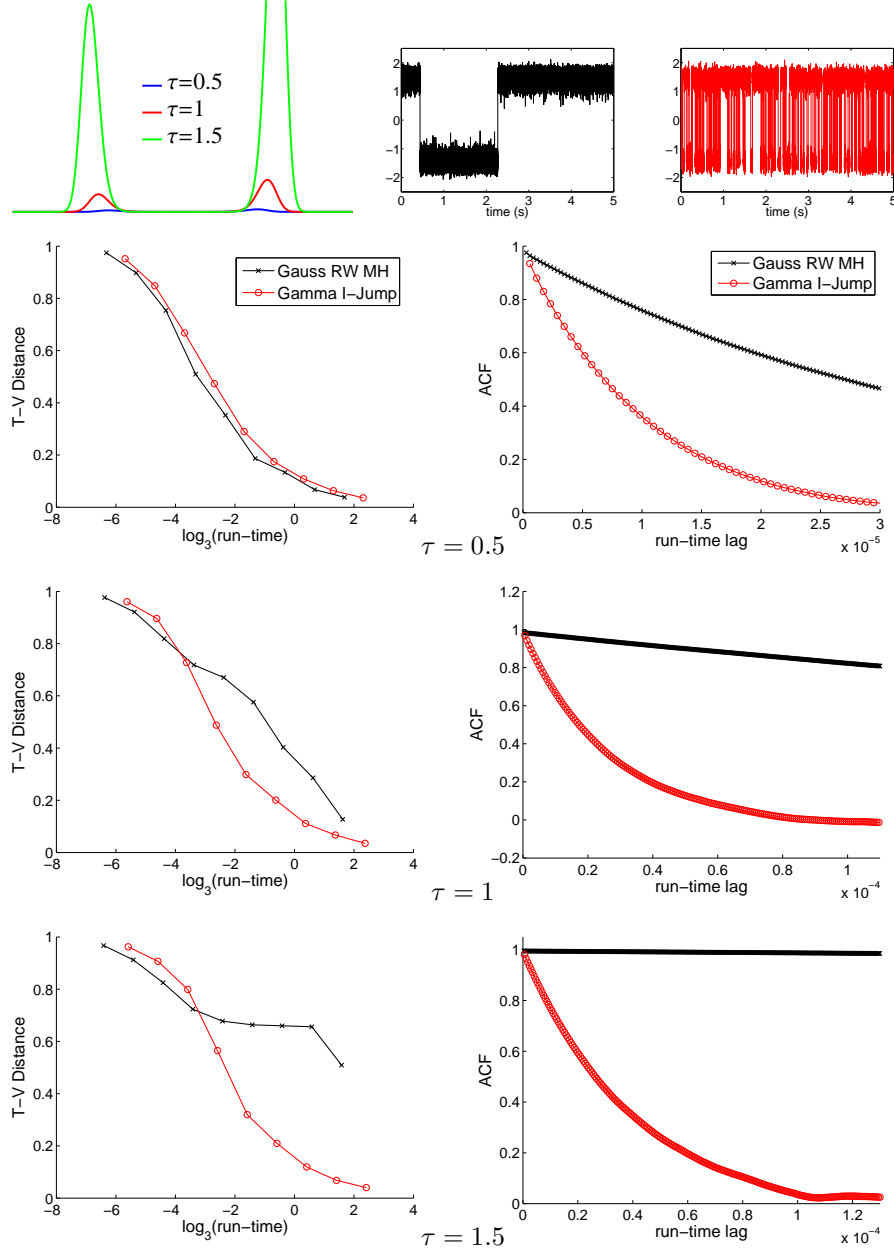
Fig. 4: *Top row:* (Left) Bimodal targets, $\pi(z_1, z_2) = 2(z_1^2 - \tau)^2 - 0.2z_1 - 5z_1^2 + 5z_2^2$, for various values of τ. Here we demonstrate a 1D cross section of the 2D distribution. (Middle) Sample state trajectories for Gauss RW MH and (right) Gamma I-Jump for $\tau = 1$. *Bottom rows:* Total variational distance vs. log runtime (left) and ACF vs. lag in runtime (right), with each row corresponding to a specific choice of τ.

τ	Avg. Escape Time for I-Jump Sampler	Avg. Escape Time for MH Sampler
0.5	1.94×10^2	1.06×10^3
1	4.64×10^2	2.47×10^4
1.5	9.06×10^2	7.89×10^5
2	2.41×10^3	N/A

Table 1: Comparison of average escape time from one local mode to another between Gamma I-Jump and Gauss RW MH. The distribution in 2D becomes more challenging with larger values of τ (plotted in Fig. 4). "N/A" in the last entry means that the escape time is so long that an accurate estimate is not available.
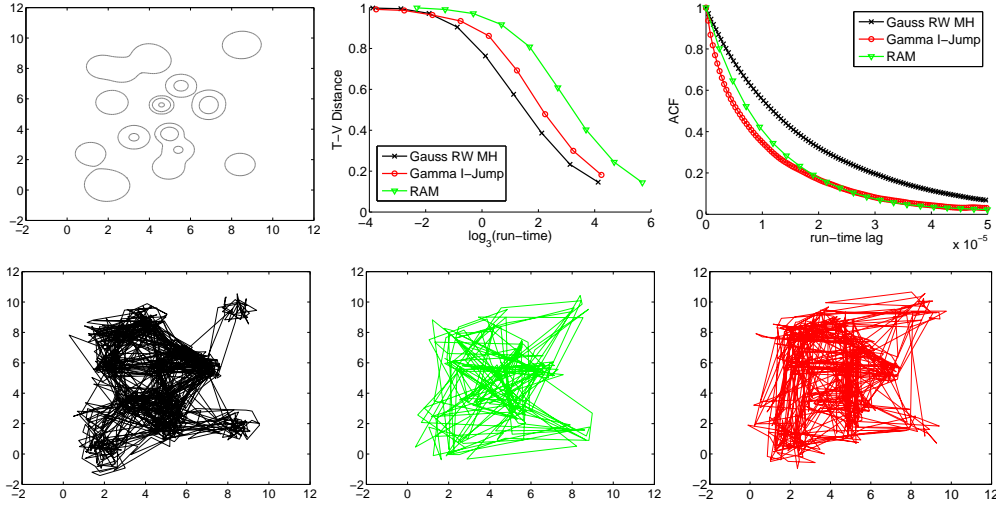

Fig. 5: *Top row:* (Left) Contour plot of a challenging multimodal probability density function; (middle) T-V distance and ACF comparisons among Gauss RW MH, Gamma I-Jump, and the recently proposed repulsing-attracting Metropolis (RAM) sampler [66]. *Bottom row:* A sample run of all three samplers, respectively.

as explored in Fig. 7. The reason is that the correlated distribution has a complex geometry. Faster exploration in random directions, as provided by our I-Jump sampler with independent proposals, only marginally increases the mixing effect in each step relative to the reversible independent proposals of MH. Since the calculation of the distribution is not demanding in this case, the small overhead of the irreversible sampler (generating gamma proposals and periodically resampling the direction of exploration) actually makes a difference and results in slightly worse performance in terms of runtime.

To improve the performance of our I-Jump sampler further in this correlated target case, it would be appealing to take the geometric information about the level sets—including the higher mass regions—into account.

Indeed, we are able to do this by replacing the independent gamma proposals with proposals from our continuous dynamics sampler—our I-MALA algorithm (Algorithm 5)—as described in Sec. 4. To demonstrate the effect of irreversibility, we choose $\mathbf{D}(\mathbf{z}) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ and $\mathbf{Q}(\mathbf{z}) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ in Eqs. (27), (28), (32), and (34). In this case, I-MALA significantly outperforms random walk MH. Because the target distribution has complex geometry, the continuous dynamics can provide guidance on locating the higher mass regions and exploring the contours rapidly with the gradient information. HMC and MALA also exploit this effect, so we provide comparisons to these methods as well. From these
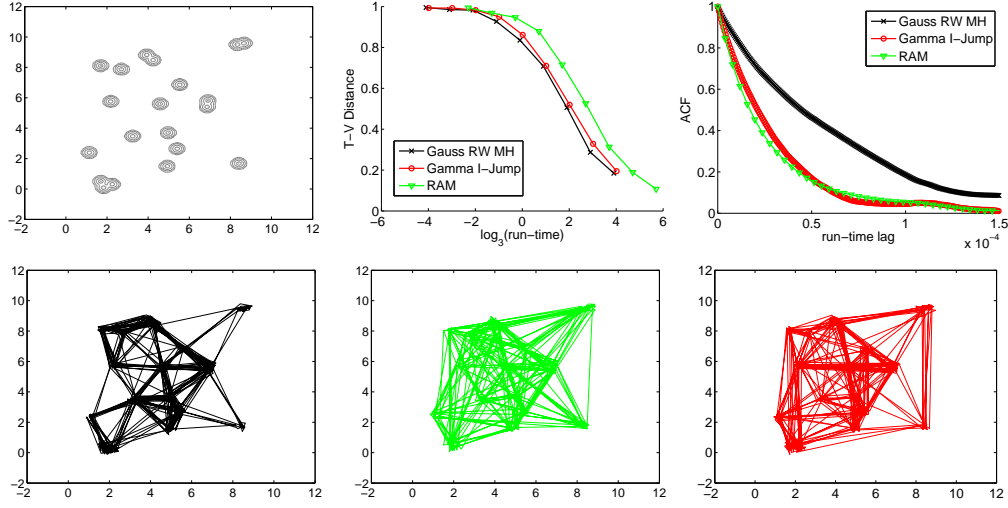
Fig. 6: Plots as in Fig. 5, but for an even more challenging multimodal case where the modes are very concentrated and well separated.
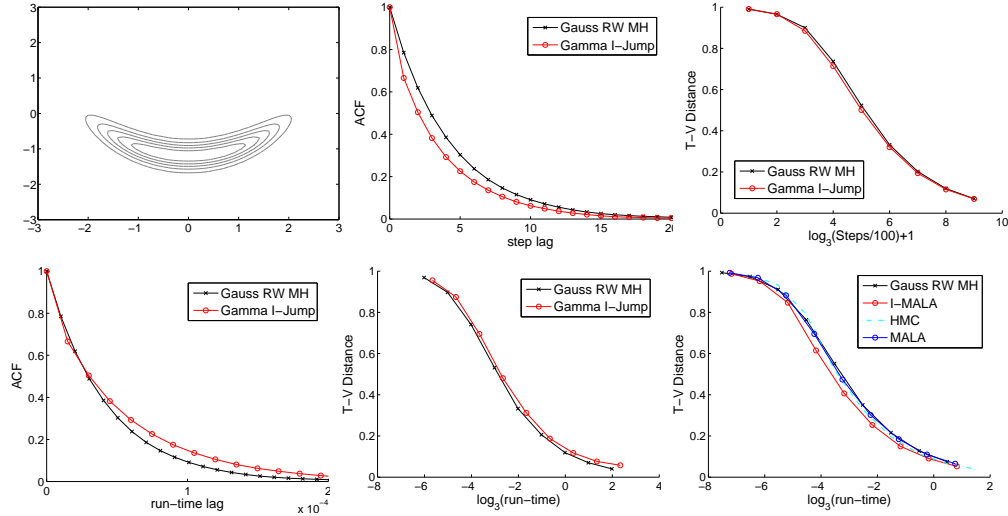

Fig. 7: *Top row:* Correlated distribution with complex geometry in 2D, $\pi(z_1, z_2) = z_1^4/10 + (4(z_2 + 1.2) - z_1^2)^2/2$ (left), ACF vs. lag in steps (middle), and T-V distance vs. log steps of Gamma I-Jump against Gauss RW MH (right). *Bottom row:* ACF vs. lag in runtime of Gamma I-Jump against Gauss RW MH (left), T-V distance vs. log runtime between Gauss RW MH and Gamma I-Jump (middle) and Gauss RW MH, I-MALA, HMC, and MALA (right).

comparisons, we see the benefits of the irreversibility itself of the I-MALA algorithm.

This experiment demonstrates the gains that are possible by combining our continuous dynamics and jump process frameworks, beyond what either can pro-

vide individually. We explore this I-MALA algorithm more extensively on a set of real-world scenarios in the following section.

6.3 Experiments for I-MALA

We test our I-MALA algorithm in sampling both Bayesian logistic regression and stochastic volatility models. In these experiments, we compare to numerous baselines. The comparison with I-Jump allow us to disentangle the importance of our SDE proposal in place of independent proposals. We used half-space Gaussian proposals since it is guaranteed to only increase the mixing of random walk MH. The comparison with MALA provides insights into the importance of an irreversible accept-reject correction to underlying irreversible SDE proposals. The vanilla random walk MH and HMC algorithms provide other popular choices to serve as baselines. We also provide comparison with some of the recently proposed irreversible samplers like Zig-Zag, SOL-HMC, and hybrid MALA algorithms.

6.3.1 Bayesian Logistic Regression

In this section, we demonstrate results from sampling a Bayesian logistic regression model. Similar to the setting in [25], we consider an $N \times D$ design matrix $\mathbf{X}$ comprised of N samples each with D covariates and a binary response variable $\mathbf{y} \in \{0, 1\}^N$. If we denote the logistic link function by $s(\cdot)$, a Bayesian logistic regression model of the binary response is obtained by introducing regression coefficients $\beta \in \mathbb{R}^D$ with an appropriate prior [23, 44]; for illustration we take $\beta \sim \mathcal{N}(\mathbf{0}, \alpha \mathbf{I})$, where $\alpha = 100$ in the experiments, employing a wide Gaussian prior.

We make use of three datasets available at the STATLOG project: https://archive.ics.uci.edu/ml/machine-learning-databases/statlog/. The first two datasets describe the connections between credit card approval and various attributes of the applicants in Australia and Germany. The third dataset is about the connection between the absence or presence of heart disease and various patient-specific covariates.

Performance metric We measure performance in terms of the per second effective sample size (ESS/runtime), where ESS is calculated as the number of steps N divided by the integrated autocorrelation time τ_{int}: $ESS = \dfrac{N}{\tau_{int}}$ [44, 62, 1, 49]. In [25], τ_{int} is estimated through the initial positive sequence estimator: $\tau_{int} = 1 + 2 \sum_k \gamma(k)$, where $\gamma(k)$ is the k-lagged autocorrelations and the sum is over the K monotone sample autocorrelations [24]. The initial positive sequence estimator assumes the Markov chain is reversible. Hence to include irreversible chains, we use the *Bartlett window estimator* for the integrated autocorrelation time τ_{int} [57, 24, 5] so that:

$$\widehat{ESS}^{BW} = \frac{N}{1 + 2\sum_{k=1}^{M}\left(1 - \dfrac{k}{M}\right)\gamma(k)}, \qquad (37)$$

where M is a large number (taken to be 3000 in the experiments). We also use a more robust estimator, the *multivariate batch mean estimator* [69, 68], to estimate the effective sample size:

$$\widehat{ESS}^{MBM} = K\left(\frac{|\widehat{\Lambda}|}{|\widehat{\Sigma}|}\right)^{D}, \qquad (38)$$

where $K = \sqrt{N}$ is the number of batches, $\widehat{\Sigma}$ is the estimated covariance of generated samples, $\widehat{\Lambda}$ is the estimated covariance of sample batch means, and $|\cdot|$ denotes determinant of the covariance matrices. (When the dimension D is very large, e.g., $D = 2000$ in the stochastic volatility model, we calculate $\widehat{ESS}^{MBM}$ as the median of the batch mean estimator over each dimension to circumvent the rank deficiency issue in $\widehat{\Lambda}$).

Optimal hyperparameters We select hyperparameters for each method via a grid search. For MH, we corroborate that the selected hyperparameters are indeed those obtained by tuning the acceptance rate between 20% and 40% at stationary. We also compare to the Zig-Zag sampler (without subsampling). One option is to use a lower bound on the Hessian of the negative log posterior (see Section 5), which in this case is $\frac{1}{4}\mathbf{X}\mathbf{X}^T + \alpha^{-1}\mathbf{I}$. However, as discussed in the experimental results below, we found this bound to be too loose to be of use. Instead, we combine the bound on

Sampler \ Dataset	Australian Credit ($D = 15$)	German Credit ($D = 25$)	Heart ($D = 14$)
MH	4.30, 4.29	2.82, 2.51	22.92, 10.76
I-Jump	4.93, 5.29	2.94, 2.93	24.32, 12.35
Zig-Zag	5.24, 5.65	3.09, 3.07	25.82, 11.68
MALA	9.05, 5.17	3.67, 3.65	29.26, 13.29
HMC	10.96, 4.36	3.73, 2.89	30.83, 18.04
SOL-HMC	13.29, 8.40	3.75, 2.92	32.71, 25.83
Hybrid MALA	11.01, 5.24	4.08, 3.63	39.54, 26.69
I-MALA	**15.95**, **8.86**	**4.47**, **4.03**	**41.23**, **27.27**

Table 2: Comparison of $\widehat{ESS}$ per second of runtime (*Left*: using $\widehat{ESS}^{BW}$ of (37) and *Right*: using $\widehat{ESS}^{MBM}$ of (38)) for various samplers on three datasets described in Sec. 6.3.1.

the gradient of the log likelihood (a constant term) with the Poisson rate for the Gaussian prior, which results in a tighter bound on the Poisson rate of the overall posterior. To derive the Poisson rate function bound, we used the recently proposed idea of superposition of Poisson rates for decomposable posteriors [10]. For the I-Jump sampler, we find that taking the same parameters as in MH already generates better performance than MH. Taking a slightly smaller variance in the half-space Gaussian proposal leads to even higher $\widehat{ESS}$ per second with an acceptance rate between 30% and 50%.

For HMC, we find that using 10 leapfrog steps to generate a sample is most efficient in terms of $\widehat{ESS}$ per second (as opposed to the commonly used 50 to 100 steps). The acceptance rate is around 90%. For MALA, an acceptance rate between 40% and 60% generates the best $\widehat{ESS}$ per second. For SOL-HMC and hybrid MALA, we use the optimal hyperparameters found for HMC and MALA to center our grid search. It seems that the hybrid MALA is more sensitive to the hyperparameters when numerical stability is concerned, possibly because of the combination of simulating Lagenvin and Hamiltonian dynamics. For the I-MALA algorithm, we take $\mathbf{D} = \mathbf{I}$, and $\mathbf{Q}(\mathbf{z}) = \begin{pmatrix} 0 & -\mathbf{I}_{d \times d} \\ \mathbf{I}_{d \times d} & 0 \end{pmatrix}$ where $d = \lfloor (D + 1)/2 \rfloor$, just as in Sec. 6.2.3 to combine benefits of Langevin diffusion and Hamiltonian dynamics. The acceptance rates are between 40% and 60%.

Experimental results Our results are summarized in Table 2. We see that I-Jump provides a gain over random walk MH in all datasets. The Zig-Zag sampler outperforms both MH and I-Jump. However, it is impor-

tant to note that having a tight bound for the transition rate function—as we were able to derive in this case—is critical to this observed performance. When using a looser bound (in this case, the bound on the Hessian of the negative log posterior), multiple steps are taken before a change of direction of exploration can happen. In practice, we found this lead to performance below that of MH and I-Jump.

HMC provides even better performance, improving over MALA and the I-Jump sampler on these examples. The irreversible variants, SOL-HMC and hybrid MALA, further improve upon HMC, achieving the highest $\widehat{ESS}$ *per sample* generated. However, the algorithms are slower per iteration than our proposed I-MALA. As a result, as seen in Table 2, I-MALA has by far the best performance across all datasets. The I-MALA algorithm combines the benefits of HMC's traversing and MALA's diffusion, plus our previously demonstrated benefits of irreversibility.

6.3.2 Stochastic Volatility Model

In this section, we follow [25, 44, 35] to study a stochastic volatility model. The daily returns y_t are modeled as $y_t = \epsilon_t \beta \exp(x_t/2)$, where $\epsilon_t \sim \mathcal{N}(0, 1)$ and the latent volatilities x_t follow an order-1 autoregressive process $x_{t+1} = \phi x_t + \eta_{t+1}$ with $x_1 \sim \mathcal{N}(0, \sigma^2/(1 - \phi^2))$. The joint probability is given by

$$p(\mathbf{y}, \mathbf{x}, \beta, \phi, \sigma) = \prod_{t=1}^{T} p(y_t | x_t, \beta) p(x_1) \pi(\beta)$$
$$\prod_{t=2}^{T} p(x_t | x_{t-1}, \phi, \sigma) \pi(\phi) \pi(\sigma),$$

with priors on the parameters chosen to be $\pi(\beta) \propto 1/\beta$, $\sigma^2 \sim \text{Inv-}\chi^2(10, 0.05)$, and $(\phi + 1)/2 \sim \text{Beta}(20, 1.5)$. We transform the constrained parameters ϕ and σ to the real line as $\sigma = \exp(\gamma)$ and $\phi = \tanh(\alpha)$, taking into account the Jacobian of the transformation. We iteratively sample over latent volatilities $\mathbf{x}$ and parameters (β, ϕ, σ) within a Gibbs sampling procedure (see [25] for further details).

Hyperparameters For MH, I-Jump and MALA, the stationary acceptance rates are similar to ones verified in Sec. 6.3.1 (between 20% and 40%, 30% and 50%, and 40% and 60%, respectively). For HMC, the leapfrog steps decorrelate quickly when sampling the latent variables, but are heavily correlated when sampling the model parameters. Hence we take 6 leapfrog steps for the former and 10 steps for the latter (more leapfrog steps can help HMC decorrelate, but increases runtime). Step sizes are then tuned to have acceptance rates for HMC between 80% and 90%.

For I-MALA, we expand the space as in the HMC algorithm and take $\mathbf{D}(\mathbf{z}) = \begin{pmatrix} a^{-1}\mathbf{I}_{d \times d} & 0 \\ 0 & a\mathbf{I}_{d \times d} \end{pmatrix}$ and $\mathbf{Q}(\mathbf{z}) = a \begin{pmatrix} 0 & -\mathbf{I}_{d \times d} \\ \mathbf{I}_{d \times d} & 0 \end{pmatrix}$, where $d = 2000$ for the latent variables and 3 for model parameters. The motivation for this choice is as follows. When a is larger, the dynamics of I-MALA are similar to HMC; when a is smaller, I-MALA becomes closer to MALA. Thus, the form of $\mathbf{D}(\mathbf{z})$, $\mathbf{Q}(\mathbf{z})$ specified above allows us to combine the benefits of HMC and MALA. Here, we leverage intuition from statistical mechanics: high-dimensional variables that are only pairwise correlated move approximately according to Newtonian mechanics perfectly described by the Hamiltonian dynamics while the low-dimensional parameters are similar to the summary quantities (or collective variables) and intrinsically follow stochastic Langevin dynamics. As such, we choose a to give HMC-like behavior for the latent state sequence sampling and MALA-like behavior for the parameter sampling by setting $a = 20$ for latent variables and $a = 5$ for model parameters.

Experimental results The results are summarized in Table 3. Again the I-Jump sampler improves upon MH. Furthermore, in this model, MALA provides the best baseline in terms of sampling the parameter space while HMC excels at sampling the latent variables. This behavior, further explored in Fig. 8, follows the pattern previously described by the statistical mechanics intuition. Hence, HMC provides faster traversing in the easy-to-explore high-dimensional latent variable space while MALA helps to diffuse faster in the lower-dimensional parameter space that exhibits complex dependencies.

SOL-HMC seems to be able to mitigate the slow mixing in model parameters. Likewise, hybrid MALA seems to be able to combine some of the benefits from HMC and MALA. Again, however, the I-MALA algorithm combines the benefits of HMC and MALA and exceeds them to provide the best overall performance. It appears from the experiments that in high dimensions, the vanilla I-Jump sampler with half-space Gaussian proposal has less improvements over MH than in low dimensional cases. We further verify this observation by the experimental results in Appendix D.2 on a standard normal distribution with doubling dimensions. However, the benefits of I-MALA over reversible MALA does not seem to have strong correlation with the dimension of the space, thus corroborating with our intuition that a direction of efficient exploration is important for the irreversible samplers. This is demonstrated in Table 3, where we notice that the improvements of I-MALA over MALA are still pronounced when going from $d = 3$ to $d = 2000$.

7 Conclusion

In this paper, we primarily focused on developing an irreversible MALA (I-MALA) algorithm (Algorithm 5). The construction of the I-MALA algorithm is comprised of two components: irreversible continuous dynamics as the proposal distribution, and an irreversible jump sampler (I-Jump) to correct for any numerical error in simulating the continuous dynamics.

Building on our preliminary work [46], we first provided a framework for specifying irreversible continuous dynamics defined via a positive semidefinite matrix $\mathbf{D}$ and a skew-symmetric matrix $\mathbf{Q}$. We prove that any choice of these two matrices leaves the target distribution invariant, and any continuous Markov process with

Sampler \ Variables	Latent Variables ($d = 2000$)	Parameters ($d = 3$)
MH	1.33, 3.24	3.86, 5.17
I-Jump	1.51, 3.38	4.15, 6.12
MALA	2.41, 3.67	11.72, 9.00
HMC	7.59, 6.21	0.96, 1.18
SOL-HMC	6.80, 6.92	2.06, 2.71
Hybrid MALA	5.65, 5.99	7.40, 9.83
I-MALA	**7.71, 7.64**	**18.27, 13.58**

Table 3: Comparison of $\widehat{ESS}$ per second of runtime (*Left*: using $\widehat{ESS}^{BW}$ of (37) and *Right*: using $\widehat{ESS}^{MBM}$ of (38)) for various algorithms sampling latent variables and parameters of a stochastic volatility model described in Sec. 6.3.2.
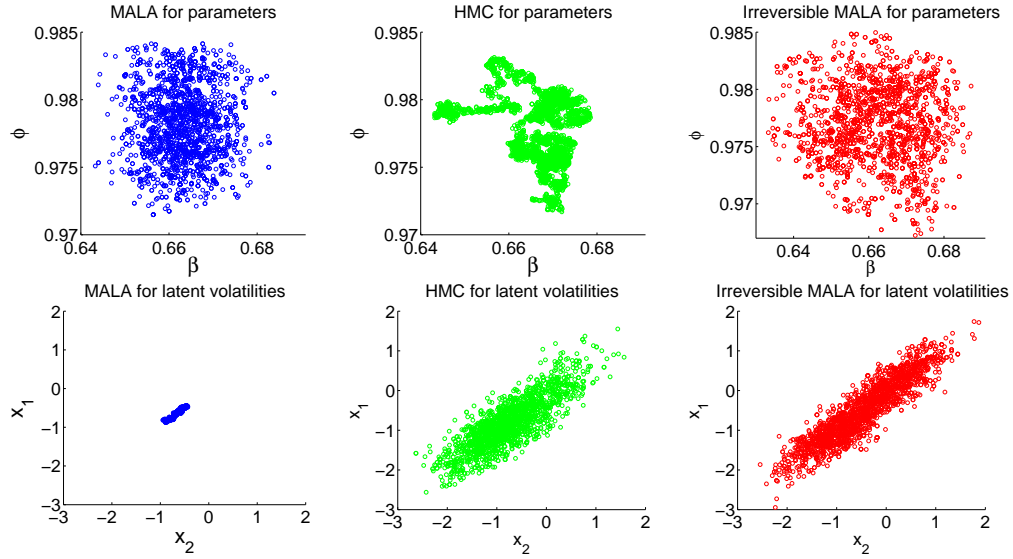

Fig. 8: $2,000$ samples of MALA, HMC, and I-MALA for parameters and latent volatilities.

the correct stationary distribution has a representation in our formulation with a certain **D** and **Q**.

We then designed the I-Jump sampler, which in itself serves as a general purpose MCMC algorithm and excels at various sampling tasks over the MH algorithm. We achieved this irreversibility by simply switching between different proposal and acceptance distributions upon rejection. The algorithm is nearly identical structurally and computationally to MH, and thus serves as a simple replacement.

Using these irreversible continuous dynamics as a proposal distribution and the I-Jump sampler to accept or reject the proposal, we arrive at our I-MALA algorithm. We demonstrated the benefits of I-MALA in sampling Bayesian logistic regression and stochastic volatility models on various real-world and synthetic datasets.

When computing target distributions that represent the posterior distribution of large datasets, scalability of the sampling algorithms becomes important. To cope with this scenario, stochastic gradient (SG) counterparts to the continuous dynamical samplers have been proposed (see [46] for more discussion). However, these SG-MCMC methods rely on the stepsize to asymptotically approach zero to converge to the correct target distribution. Otherwise a bias is introduced. One could use our proposed I-Jump sampler to correct this bias, but as with standard MH, the accept/reject cal-

culations involve touching the whole dataset, exactly what the SG methods seek to avoid. Instead, we could imagine extending the MH data-subsampling ideas of [37, 2, 3] to our I-Jump sampler to decrease the computational burden of the accept-reject step.

Another direction for future work is to examine the effect of incorporating second order information about the target distribution (similar to RMALA) into our I-MALA. It is well-studied that incorporating the Fisher information metric in $\mathbf{D}$ and $\mathbf{M}$ can speed up mixing rates in MALA and HMC [74, 25, 4]; the benefits can be explained by the fact that using an approximation of the Hessian of the target distribution to precondition the gradient makes the space locally "flat". However, in our case there are many ways in which such second order information can be incorporated in $\mathbf{D}$ and $\mathbf{Q}$. Exploring these options and the interplay between $\mathbf{D}$ and $\mathbf{Q}$ is a direction of future work.

Acknowledgments

A Samplers from Continuous Markov Processes

A.1 Proof of Theorem 1 (existence of $\mathbf{Q}(\mathbf{z})$ in (4))

Proof of Theorem 1 is comprised of two sets of ideas appearing in different fields. Here, we write the proof in two steps accordingly.

1. Plug the stationary solution $\pi(\mathbf{z})$ into (2) and observe that

$$\sum_i \frac{\partial}{\partial \mathbf{z}_i} \left\{ \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big) - \mathbf{f}_i \pi(\mathbf{z}) \right\} = 0.$$

2. Constructively prove that if entries of

$$\left\{ \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big) - \mathbf{f}_i \pi(\mathbf{z}) \right\}$$

belong to $L^1(\mathbb{R}^d)$, then there exists a matrix $\mathbf{Q}(\mathbf{z})$ with entries in $W^{1,1}(\pi)$, such that $\sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{Q}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big) = \mathbf{f}_i \pi(\mathbf{z}) - \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big)$.

Here we denote $L^1(\mathbb{R}^d)$ as the space of Lebesgue integrable functions on $\mathbb{R}^d$, and $W^{1,1}(\pi)$ as the Sobolev space of functions with weak derivatives and function values integrable with respect to the density π times the Lebesgue measure. Step 1 can be found in literatures of continuous Markov processes [59, 17, 58, 18, 70, 54]. Step 2 has been found in earlier works on stochastic models in fluid dynamics and homogenization [36].

Step 1 is accomplished by noting that the un-normalized density function $\pi(\mathbf{z}) \propto p^s(\mathbf{z})$ is also a stationary solution of (2). Hence,

$$\sum_i \frac{\partial}{\partial \mathbf{z}_i} \left\{ \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big) - \mathbf{f}_i \pi(\mathbf{z}) \right\} = 0. \qquad (39)$$

Step 2 provides a possible form of $\mathbf{Q}$ in terms of $\mathbf{D}$, $\mathbf{f}$, and π. First, compare the right hand sides of (2) and (4) and observe that they are equivalent if and only if there exists $\mathbf{Q}(\mathbf{z})$ such that:

$$\sum_j \frac{\partial}{\partial \mathbf{z}_j} \mathbf{Q}_{ij}(\mathbf{z})\pi(\mathbf{z}) = \mathbf{f}_i \pi(\mathbf{z}) - \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big). \qquad (40)$$

Since entries of $\left\{ \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big) - \mathbf{f}_i \pi(\mathbf{z}) \right\}$ belong to $L^1(\mathbb{R}^d)$, Fourier transform of it exists:

$$\hat{\mathbf{F}}_i(\mathbf{k}) = \int_{\mathbb{R}^d} d\mathbf{z} \left(\mathbf{f}_i(\mathbf{z})\pi(\mathbf{z}) - \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big) \right)$$
$$e^{-2\pi \mathrm{i}\, \mathbf{k}^T \mathbf{z}}.$$

Therefore, we can take $\mathbf{Q}(\mathbf{z})$ such that

$$\mathbf{Q}_{ij}(\mathbf{z})\pi(\mathbf{z}) = \int_{\mathbb{R}^d} \frac{\mathbf{k}_j \hat{\mathbf{F}}_i(\mathbf{k}) - \mathbf{k}_i \hat{\mathbf{F}}_j(\mathbf{k})}{(2\pi \mathrm{i}) \cdot \sum_l \mathbf{k}_l^2} e^{2\pi \mathrm{i} \sum_l \mathbf{k}_l \mathbf{x}_l} d\mathbf{k}, \quad (41)$$

where entries of $\mathbf{Q}(\mathbf{z})$ belong to $W^{1,1}(\pi)$.

Remark 1 It is worth noting that the condition of

$$\left\{ \sum_j \frac{\partial}{\partial \mathbf{z}_j} \Big(\mathbf{D}_{ij}(\mathbf{z})\pi(\mathbf{z}) \Big) - \mathbf{f}_i \pi(\mathbf{z}) \right\} \in L^1(\mathbb{R}^d)$$

can be rewritten for the vector field related to the SDE as: $\left\{ \mathbf{f}(\mathbf{z}) + \mathbf{D}(\mathbf{z})\nabla H(\mathbf{z}) - \Gamma^{\mathbf{D}}(\mathbf{z}) \right\} \in L^1(\pi)$, where $\Gamma_i^{\mathbf{D}}(\mathbf{z}) = \sum_{j=1}^d \frac{\partial}{\partial \mathbf{z}_j} \mathbf{D}_{ij}(\mathbf{z})$. We see that the condition is relatively mild for the purpose of constructing samplers.

A.2 Reversible and Irreversible Continuous Dynamics for Sampling

As has been previously noted [59, 17, 58, 18, 70, 54], the stochastic dynamics of Eqs. (4) and the corresponding (3) can be decomposed into a reversible Markov process, and an irreversible process. Formally, this can be elucidated by the infinitesimal generator $\mathcal{G}[\cdot]$ of the stochastic process (3) (see [54, 36] for more background). For ease of derivation, we work with the Hilbert space $L^2(\pi)$ of square integrable functions with respect to $\pi(\mathbf{z})$, equipped with inner product $< \phi, \varphi >_\pi = \int_{\mathbb{R}^d} \bar{\phi}(\mathbf{z})\varphi(\mathbf{z})\pi(\mathbf{z})d\mathbf{z}$. For $\phi(\mathbf{z}) \in W^{2,2}(\pi)$ (i.e., ϕ and its second order weak derivatives belong to $L^2(\pi)$),

$$\mathcal{G}[\phi(\mathbf{z})] = \nabla^T \Big(\big(\mathbf{D}(\mathbf{z}) - \mathbf{Q}(\mathbf{z}) \big) \nabla\phi(\mathbf{z}) \Big)$$
$$- \nabla H(\mathbf{z})^T \big(\mathbf{D}(\mathbf{z}) - \mathbf{Q}(\mathbf{z}) \big) \nabla\phi(\mathbf{z}). \qquad (42)$$

Then adjoint of $\mathcal{G}$ in $L^2(\pi)$ is:

$$\mathcal{G}^*[\phi(\mathbf{z})] = \nabla^T \Big(\big(\mathbf{D}(\mathbf{z}) + \mathbf{Q}(\mathbf{z}) \big) \nabla\phi(\mathbf{z}) \Big)$$
$$- \nabla H(\mathbf{z})^T \big(\mathbf{D}(\mathbf{z}) + \mathbf{Q}(\mathbf{z}) \big) \nabla\phi(\mathbf{z}). \qquad (43)$$

Therefore, $\mathcal{G}$ decomposes into a self-adjoint part $\mathcal{G}^S[\phi(\mathbf{z})] = \nabla^T \Big(\mathbf{D}(\mathbf{z})\nabla\phi(\mathbf{z}) \Big) - \nabla H(\mathbf{z})^T \mathbf{D}(\mathbf{z})\nabla\phi(\mathbf{z})$ and an anti-self-adjoint part $\mathcal{G}^A[\phi(\mathbf{z})] = \nabla^T \Big(\mathbf{Q}(\mathbf{z})\nabla\phi(\mathbf{z}) \Big) - \nabla H(\mathbf{z})^T \mathbf{Q}(\mathbf{z})\nabla\phi(\mathbf{z})$. The self-adjoint operator $\mathcal{G}^S$ corresponds to the reversible Markov process while the anti-self-adjoint operator $\mathcal{G}^A$ corresponds to the irreversible process.

It can be seen that the reversible process is determined solely by the diffusion matrix $\mathbf{D}$, where evolution of its probability den-

sity function is:

$$
\frac{\partial}{\partial t} p(\mathbf{z}; t) = \nabla^T \cdot \left(\mathbf{D}(\mathbf{z}) \left[p(\mathbf{z}; t) \nabla H(\mathbf{z}) + \nabla p(\mathbf{z}; t) \right] \right)
$$

$$
= \sum_{i,j=1}^{d} \frac{\partial^2}{\partial \mathbf{z}_i \partial \mathbf{z}_j} \left(\mathbf{D}_{ij}(\mathbf{z}) p(\mathbf{z}; t) \right) \tag{44}
$$

$$
+ \sum_{i=1}^{d} \frac{\partial}{\partial \mathbf{z}_i} \left(\left[\sum_j \mathbf{D}_{ij}(\mathbf{z}) \frac{\partial H(\mathbf{z})}{\partial \mathbf{z}_j} - \Gamma_i^{\mathbf{D}}(\mathbf{z}) \right] p(\mathbf{z}; t) \right)
$$

Here, $\Gamma_i^{\mathbf{D}}(\mathbf{z}) = \sum_{j=1}^{d} \frac{\partial}{\partial \mathbf{z}_j} \mathbf{D}_{ij}(\mathbf{z})$. According to Itô's convention, (44) corresponds to reversible Brownian motion in a potential force field on a Riemannian manifold specified by the diffusion matrix $\mathbf{D}(\mathbf{z})$: $d\mathbf{z} = \left[-\mathbf{D}(\mathbf{z}) \nabla H(\mathbf{z}) + \Gamma^{\mathbf{D}}(\mathbf{z}) \right] dt + \sqrt{2\mathbf{D}(\mathbf{z})} d\mathbf{W}(t)$. This is referred to as *Riemannian Langevin dynamics* [25, 74]. When $\mathbf{D}(\mathbf{z})$ is positive definite, the reversible Markov dynamics have nice statistical regularity and will drive the system to converge to the stationary distribution.

The irreversible process is determined solely by $\mathbf{Q}$, with its probability density function evolving according to:

$$
\frac{\partial}{\partial t} p(\mathbf{z}; t) = \nabla^T \cdot \left(\mathbf{Q}(\mathbf{z}) \left[p(\mathbf{z}; t) \nabla H(\mathbf{z}) + \nabla p(\mathbf{z}; t) \right] \right)
$$

$$
= \nabla^T \cdot \left(\left[\mathbf{Q}(\mathbf{z}) \nabla H(\mathbf{z}) - \Gamma^{\mathbf{Q}}(\mathbf{z}) \right] p(\mathbf{z}; t) \right). \tag{45}
$$

Here, $\Gamma_i^{\mathbf{Q}}(\mathbf{z}) = \sum_{j=1}^{d} \frac{\partial}{\partial \mathbf{z}_j} \mathbf{Q}_{ij}(\mathbf{z})$. The last line of (45) is a Liouville equation, which describes the density evolution of $p(\mathbf{z}; t)$ according to conserved, deterministic dynamics: $d\mathbf{z}/dt = -\mathbf{Q}(\mathbf{z}) \nabla H(\mathbf{z}) + \Gamma^{\mathbf{Q}}(\mathbf{z})$, with $\pi(\mathbf{z})$ its invariant measure.

Combining the dynamics of (44) and (45) leads to a general SDE, (3), with stationary distribution $\pi(\mathbf{z})$. Previous methods [63, 74, 72, 53, 51] have primarily focused on solely reversible or irreversible processes, respectively, as we make explicit in Sec. 2.3.

B Irreversible Jump Processes for MCMC

B.1 Equivalence of (10) and (11)

We introduce a symmetric bivariate function $S(\mathbf{x}, \mathbf{z}) = S(\mathbf{z}, \mathbf{x}) = \frac{1}{2} \left(W(\mathbf{z}|\mathbf{x}) \pi(\mathbf{x}) + W(\mathbf{x}|\mathbf{z}) \pi(\mathbf{z}) \right)$, and an anti-symmetric bivariate function $A(\mathbf{x}, \mathbf{z}) = -A(\mathbf{z}, \mathbf{x}) = \frac{1}{2} \left(W(\mathbf{z}|\mathbf{x}) \pi(\mathbf{x}) - W(\mathbf{x}|\mathbf{z}) \pi(\mathbf{z}) \right)$ for (10). A different form of the jump process (10) can be written according to $S(\mathbf{x}, \mathbf{z})$ and $A(\mathbf{x}, \mathbf{z})$ as

$$
\frac{\partial p(\mathbf{z}|\mathbf{y}; t)}{\partial t} = \int_{\mathbb{R}^d} d\mathbf{x} \left[S(\mathbf{x}, \mathbf{z}) \frac{p(\mathbf{x}|\mathbf{y}; t)}{\pi(\mathbf{x})} - S(\mathbf{x}, \mathbf{z}) \frac{p(\mathbf{z}|\mathbf{y}; t)}{\pi(\mathbf{z})} + A(\mathbf{x}, \mathbf{z}) \frac{p(\mathbf{x}|\mathbf{y}; t)}{\pi(\mathbf{x})} \right].
$$

Plugging $p(\mathbf{z}|\mathbf{y}; t) = \pi(\mathbf{z})$ into the above equation, we find that as long as $\int_{\mathbb{R}^d} A(\mathbf{x}, \mathbf{z}) d\mathbf{x} = 0$, $\pi(\mathbf{z})$ is a stationary solution to the equation. Since $\frac{S(\mathbf{x}, \mathbf{z}) + A(\mathbf{x}, \mathbf{z})}{\pi(\mathbf{x})}$ denotes a transition probability density, $S(\mathbf{x}, \mathbf{z}) + A(\mathbf{x}, \mathbf{z}) > 0$ for any $\mathbf{x}$ and $\mathbf{z}$. The restriction that $S(\mathbf{x}, \mathbf{z}) \pi^{-1}(\mathbf{x})$ and $A(\mathbf{x}, \mathbf{z}) \pi^{-1}(\mathbf{x})$ are bounded is imposed for the practical purpose of proposing samples in (12). We thereby notice that the requirement that $\pi(\mathbf{z})$ is a stationary distribution of the jump process is translated into simpler constraints.

B.2 Verifying condition 3 on $A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ in Section 3.3

The anti-symmetric function $A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)$ (expressed in (21)) of (20) can be written as:

$$
A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p)
$$

$$
= \frac{1}{2\Delta t} \Big(\pi(\mathbf{y}) \pi(\mathbf{y}^p) p(\mathbf{z}, \mathbf{z}^p | \mathbf{y}, \mathbf{y}^p; \Delta t)
$$

$$
- \pi(\mathbf{z}) \pi(\mathbf{z}^p) p(\mathbf{y}, \mathbf{y}^p | \mathbf{z}, \mathbf{z}^p; \Delta t) \Big)
$$

$$
= \frac{1}{2} \delta(\mathbf{z}^p - \mathbf{y}^p) \Big(\mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p) - \mathfrak{F}(\mathbf{z}, \mathbf{z}^p, \mathbf{y}, \mathbf{y}^p) \Big)
$$

$$
- \frac{1}{2} \delta(\mathbf{z}^p + \mathbf{y}^p) \delta(\mathbf{z} - \mathbf{y})
$$

$$
\cdot \int_{\mathbb{R}^d} \Big(\mathfrak{F}(\mathbf{y}, \mathbf{y}^p, \mathbf{x}, -\mathbf{z}^p) - \mathfrak{F}(\mathbf{z}, \mathbf{z}^p, \mathbf{x}, -\mathbf{y}^p) \Big) d\mathbf{x}.
$$

Below we prove that, as required,

$$
\int_{\mathbb{R}^{d+d^P}} A(\mathbf{y}, \mathbf{y}^p, \mathbf{z}, \mathbf{z}^p) \, d\mathbf{y} \, d\mathbf{y}^p = 0.
$$

Proof

$$
\int_{\mathbb{R}^{d+d^P}} A(\mathbf{x}, \mathbf{x}^p, \mathbf{z}, \mathbf{z}^p) \, d\mathbf{x} \, d\mathbf{x}^p
$$

$$
= \frac{1}{2} \int_{\mathbb{R}^d} \Big(\mathfrak{F}(\mathbf{y}, \mathbf{z}^p, \mathbf{z}, \mathbf{z}^p) - \mathfrak{F}(\mathbf{z}, \mathbf{z}^p, \mathbf{y}, \mathbf{z}^p) \Big) \, d\mathbf{y}
$$

$$
- \frac{1}{2} \int_{\mathbb{R}^d} \Big(\mathfrak{F}(\mathbf{z}, -\mathbf{z}^p, \mathbf{x}, -\mathbf{z}^p) - \mathfrak{F}(\mathbf{z}, \mathbf{z}^p, \mathbf{x}, \mathbf{z}^p) \Big) d\mathbf{x}.
$$

One can check that in (23) and (19),

$$
\widetilde{f}(\mathbf{z}, \cdot | \mathbf{y}, -\mathbf{y}^p) = \widetilde{g}(\mathbf{z}, \cdot | \mathbf{y}, \mathbf{y}^p).
$$

Hence,

$$
\mathfrak{F}(\mathbf{y}, -\mathbf{y}^p, \mathbf{z}, -\mathbf{z}^p) = \mathfrak{F}(\mathbf{z}, \mathbf{z}^p, \mathbf{y}, \mathbf{y}^p).
$$

Therefore

$$
\int_{\mathbb{R}^{d+d^P}} A(\mathbf{x}, \mathbf{x}^p, \mathbf{z}, \mathbf{z}^p) \, d\mathbf{x} \, d\mathbf{x}^p \tag{46}
$$

$$
= \frac{1}{2} \int_{\mathbb{R}^d} \Big(\mathfrak{F}(\mathbf{z}, -\mathbf{z}^p, \mathbf{y}, -\mathbf{z}^p) - \mathfrak{F}(\mathbf{z}, \mathbf{z}^p, \mathbf{y}, \mathbf{z}^p) \Big) \, d\mathbf{y}
$$

$$
- \frac{1}{2} \int_{\mathbb{R}^d} \Big(\mathfrak{F}(\mathbf{z}, -\mathbf{z}^p, \mathbf{x}, -\mathbf{z}^p) - \mathfrak{F}(\mathbf{z}, \mathbf{z}^p, \mathbf{x}, \mathbf{z}^p) \Big) d\mathbf{x} = 0.
$$

C Proof of Theorem 2 (relation between forward process and adjoint process)

We first prove that for the infinitesimal generators, the backward transition probability density following the adjoint process and the forward transition probability density are related as: $\pi(\mathbf{y})P(\mathbf{z}|\mathbf{y};\mathrm{d}t) = \pi(\mathbf{z})P^\dagger(\mathbf{y}|\mathbf{z};\mathrm{d}t)$. Taking path integrals with respect to the infinitesimal generators leads to the conclusion.

As is standard, we use two arbitrary smooth test functions $\psi(\mathbf{y})$ and $\phi(\mathbf{z})$. Then we use the definition of the infinitesimal generator of the process P and $P^\dagger$: $\mathcal{G}[\phi(\mathbf{y})]\mathrm{d}t = \int P(\mathbf{z}|\mathbf{y};\mathrm{d}t)\phi(\mathbf{z})\mathrm{d}\mathbf{z} - \phi(\mathbf{y})$ and obtain

$$\int\int \mathrm{d}\mathbf{y}\mathrm{d}\mathbf{z}\, P(\mathbf{z}|\mathbf{y};\mathrm{d}t)\pi(\mathbf{y})\,\psi(\mathbf{y})\phi(\mathbf{z})$$
$$= <\psi,(I+\mathrm{d}t\,\mathcal{G})[\phi]>_\pi,$$

and

$$\int\int \mathrm{d}\mathbf{y}\mathrm{d}\mathbf{z}\, P^\dagger(\mathbf{y}|\mathbf{z};\mathrm{d}t)\pi(\mathbf{z})\,\psi(\mathbf{y})\phi(\mathbf{z})$$
$$= <(I+\mathrm{d}t\,\mathcal{G}^*)[\psi],\phi>_\pi.$$

Since $\mathcal{G}$ and $\mathcal{G}^*$ are adjoint in $L^2(\pi)$: $<\psi,\mathcal{G}[\phi]>_\pi = <\mathcal{G}^*[\psi],\phi>_\pi$,

$$\pi(\mathbf{y})P(\mathbf{z}|\mathbf{y};\mathrm{d}t) = \pi(\mathbf{z})P^\dagger(\mathbf{y}|\mathbf{z};\mathrm{d}t).$$

Then we take path integrals over the forward path and the backward one. Using the Markov properties,

$$P(\mathbf{z}_N,t_N|\mathbf{z}_0,t_0)$$
$$= \int\cdots\int \prod_{i=1}^{N-1}\mathrm{d}\mathbf{z}_i \prod_{i=0}^{N-1} P(\mathbf{z}_{i+1},t_{i+1}|\mathbf{z}_i,t_i);$$

and

$$P^\dagger(\mathbf{z}_0,t_N|\mathbf{z}_N,t_0)$$
$$= \int\cdots\int \prod_{i=1}^{N-1}\mathrm{d}\mathbf{z}_i \prod_{i=0}^{N-1} P^\dagger(\mathbf{z}_i,t_{i+1}|\mathbf{z}_{i+1},t_i)$$
$$= \int\cdots\int \prod_{i=1}^{N-1}\mathrm{d}\mathbf{z}_i \prod_{i=0}^{N-1} \frac{\pi(\mathbf{z}_i)}{\pi(\mathbf{z}_{i+1})}P(\mathbf{z}_{i+1},t_{i+1}|\mathbf{z}_i,t_i)$$
$$= \int\cdots\int \prod_{i=1}^{N-1}\mathrm{d}\mathbf{z}_i \prod_{i=0}^{N-1} \frac{\pi(\mathbf{z}_i)}{\pi(\mathbf{z}_{i+1})}P(\mathbf{z}_{i+1},t_{i+1}|\mathbf{z}_i,t_i)$$
$$= \frac{\pi(\mathbf{z}_0)}{\pi(\mathbf{z}_N)}P(\mathbf{z}_N,t_N|\mathbf{z}_0,t_0).$$

Taking the time interval between t_i and t_{i+1} to be infinitesimal, we obtain that $\frac{P(\mathbf{z}^{(T)},T|\mathbf{z}^{(t)},t)}{P^\dagger(\mathbf{z}^{(t)},T|\mathbf{z}^{(T)},t)} = \frac{\pi(\mathbf{z}^{(T)})}{\pi(\mathbf{z}^{(t)})}$. The same conclusion can be reached by proving that the semigroups $e^{t\mathcal{G}}$ and $e^{t\mathcal{G}^*}$ generated by $\mathcal{G}$ and $\mathcal{G}^*$ are also adjoint with each other [32, 55, 56].

Dimensions	MH	I-Jump	I-Jump:MH
10	16683.95	20717.96	1.24
20	7101.00	9648.46	1.36
40	3464.73	3772.41	1.09
80	1555.40	1787.41	1.15
160	537.87	561.94	1.04
320	72.46	76.97	1.06
640	28.18	28.99	1.03

Table 4: Comparison of $\widehat{ESS}$ per second of runtime for I-Jump versus MH with different dimensions of Gaussian target distributions.

D Experiments

D.1 Parameter settings for the irreversible jump sampler

In the 1D experiments, we use $\beta = 1.2$ for the normal distribution case (where the length of the region of definition is 10) and $\beta = 0.8$ for the log-normal distribution (where the length of the region of definition is 5). The acceptance rate is around 50% in these cases. Due to the irreversibility of the sampler, a high acceptance rate can be maintained while reducing the autocorrelation time.

In the visual comparison of samplers of Fig. 2, we use $\beta = 0.15$ (where the lengths of the region of definition is 2×2). In the 2D bimodal experiments, we use $\beta = 0.4$ (where the lengths of the region of definition is 6×3). In the 2D multimodal experiments, we take $\beta = 1.5$ (where the lengths of the region of definition is 14×14). For the 2D correlated distribution, we take $\beta = 0.25$ (where the lengths of the region of definition is 5×2).

D.2 Effect of Dimensionality on I-Jump versus MH

In this appendix, we explore how the relative improvement of I-Jump over MH scales with dimensionality. We consider a standard normal distribution and double the dimension from one comparison to the next. For the I-Jump sampler, we use the vanilla half-space Gaussian proposal. It can be observed from Table 4 that the potential benefits of irreversibility in the I-Jump sampler slowly diminish with increasing dimensionality.

References

1. C. Andrieu and J. Thoms. A tutorial on adaptive MCMC. *Stat. Comput.*, 18:343–373, 2008.

2. R. Bardenet, A. Doucet, and C. Holmes. Towards scaling up Markov chain Monte Carlo: An adaptive subsampling approach. In *Proceedings of the 30th International Conference on Machine Learning (ICML'14)*, 2014.

3. R. Bardenet, A. Doucet, and C. Holmes. On Markov chain Monte Carlo methods for tall data. arXiv:1505.02827, 2015.

4. A. Barp, F.-X. Briol, A. D. Kennedy, and M. Girolami. Geometry and dynamics for Markov chain Monte Carlo. arXiv:1705.02891, 2017.

5. M. S. Bartlett. Smoothing periodograms from time-series with continuous spectra. *Nature*, 161:686–687, 1948.

6. J. Bierkens. Non-reversible Metropolis-Hastings. *Stat. Comput.*, pages 1–16, 2015.

7. J. Bierkens, P. Fearnhead, and G. Roberts. The Zig-Zag process and super-efficient sampling for Bayesian analysis of big data. arXiv:1607.03188, 2016.

8. J. Bierkens and G. Roberts. A piecewise deterministic scaling limit of Lifted Metropolis-Hastings in the Curie-Weiss model. arXiv:1509.00302, 2016.

9. N. Bou-Rabee and H. Owhadi. Long-run accuracy of variational integrators in the stochastic context. *SIAM J. Num. Anal.*, 48:278–297, 2010.

10. A. Bouchard-Côté, S. J. Vollmer, and A. Doucet. The bouncy particle sampler: A non-reversible rejection-free Markov chain Monte Carlo method. arXiv:1510.02451, 2016.

11. C. Chen, N. Ding, and L. Carin. On the convergence of stochastic gradient MCMC algorithms with high-order integrators. In *Advances in Neural Information Processing Systems 28 (NIPS'15)*, pages 2278–2286. 2015.

12. F. Chen, L. Lovász, and I. Pak. Lifting Markov chains to speed up mixing. In *Proceedings of the 31st annual ACM STOC*, pages 275–281. 1999.

13. T. Chen, E. B. Fox, and C. Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *Proceeding of 31st International Conference on Machine Learning (ICML'14)*, 2014.

14. T.-L. Chen and C.-R. Hwang. Accelerating reversible Markov chains. *Stat. Probabil. Lett.*, 83(9):1956–1962, 2013.

15. S. Chib and E. Greenberg. Understanding the Metropolis-Hastings algorithm. *Am. Stat.*, 49(4):327–335, 1995.

16. G. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E.*, 60:2721–2726, 1999.

17. A. Dembo and J.-D. Deuschel. Markovian perturbation, response and fluctuation dissipation theorem. *Ann. Inst. H. Poincar Probab. Statist.*, 46:822–852, 2010.

18. J. D. Deuschel and D. W. Stroock. *Large Deviations*. Amer. Math. Soc., 2001.

19. P. Diaconis, S. Holmes, and R. M. Neal. Analysis of a nonreversible Markov chain sampler. *Ann. Appl. Probab.*, 10:726–752, 2000.

20. N. Ding, Y. Fang, R. Babbush, C. Chen, R. D. Skeel, and H. Neven. Bayesian sampling using stochastic gradient thermostats. In *Advances in Neural Information Processing Systems 27 (NIPS'14)*. 2014.

21. S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 195(2):216 – 222, 1987.

22. A. B. Duncan, T. Lelièvre, and G. A. Pavliotis. Variance reduction using nonreversible Langevin samplers. *J. Stat. Phys.*, 163(3):457–491, 2016.

23. A. Gelman, J. B. Carhn, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall, 2004.

24. C. J. Geyer. Practical Markov chain Monte Carlo. *Statist. Sci.*, 7:473–483, 1992.

25. M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Series B (Stat. Methodol.)*, 73(2):123–214, 2011.

26. P. Gustafson. A guided walk Metropolis algorithm. *Stat. Comput.*, 8(4):357–364, 1998.

27. W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:pp 97–109, 1970.

28. T. Hatano and S.-I. Sasa. Steady-state thermodynamics of Langevin systems. *Phys. Rev. Lett.*, 86:3463–3466, 2001.

29. A. M. Horowitz. A generalized guided Monte Carlo algorithm. *Phys. Lett. B*, 268(2):247 – 252, 1991.

30. C.-R. Hwang, S.-Y. Hwang-Ma, and S.-J. Sheu. Accelerating Gaussian diffusions. *Ann. Appl. Probab.*, 3(3):897–913, 08 1993.

31. C.-R. Hwang, S.-Y. Hwang-Ma, and S.-J. Sheu. Accelerating diffusions. *Ann. Appl. Probab.*, 15(2):1433–1444, 05 2005.

32. S. Jansen and N. Kurt. On the notion(s) of duality for Markov processes. *Probab. Surveys*, 11:59–120, 2014.

33. S. F. Jarner and G. O. Roberts. Convergence of heavy-tailed Monte Carlo Markov chain algorithms. *Scand. J. Stat.*, 34(4):781–815, 2007.

34. Marcus Kaiser, Robert L. Jack, and Johannes Zimmer. Acceleration of convergence to equilibrium in Markov chains by breaking detailed balance. *J. Stat. Phys.*, 168(2):259–287, Jul 2017.

35. S. Kim, N. Shephard, and S. Chib. Stochastic volatility: likelihood inference and comparison with ARCH models. *Rev. Econ. Stud.*, 65:361–393, 1998.

36. T. Komorowski, C. Landim, and S. Olla. *Fluctuations in Markov Processes – Time Symmetry and Martingale Approximation*. Springer, Berlin-Heidelberg-New York, 2012.

37. A. Korattikara, Y. Chen, and M. Welling. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In *Proceedings of the 30th International Conference on Machine Learning (ICML'14)*, 2014.

38. S. C. Kou, Q. Zhou, and W. H. Wong. Discussion paper: Equi-energy sampler with applications in statistical inference and statistical mechanics. *Ann. Stat.*, 34(4):1581–1619, 2006.

39. C. Kwon, P. Ao, and D. J. Thouless. Structure of stochastic dynamics near fixed points. *Proc. Natl. Acad. Sci.*, 102(37):13029–13033, 2005.

40. B. Leimkuhler, C. Matthews, and M. Tretyakov. On the long-time integration of stochastic gradient systems. *Proc. R. Soc. A*, 470:20140120, 2014.

41. B. Leimkuhler and X. Shang. Adaptive thermostats for noisy gradient systems. *SIAM J. Sci. Comput*, 38(2):A712–A736, 2016.

42. T. Leliévre, F. Nier, and G.A. Pavliotis. Optimal non-reversible linear drift for the convergence to equilibrium of a diffusion. *J. Stat. Phys.*, 152:237–274, 2013.

43. C. Liu, J. Zhu, and Y. Song. Stochastic gradient geodesic MCMC methods. In *Advances in Neural Information Processing Systems 29 (NIPS'16)*, pages 3009–3017, 2016.

44. J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, 2001.

45. X. Lu, V. Perrone, L. Hasenclever, Y. W. Teh, and S. J. Vollmer. Relativistic Monte Carlo. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS'17)*. 2017.

46. Y.-A Ma, T. Chen, and E. Fox. A complete recipe for stochastic gradient MCMC. In *Advances in Neural Information Processing Systems 28 (NIPS'15)*, pages 2899–2907. 2015.

47. Y.-A Ma and H. Qian. Universal ideal behavior and macroscopic work relation of linear irreversible stochastic thermodynamics. *New J. Phys.*, 17(6):065013, 2015.

48. M. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and Teller. E. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:pp 1087–1092, 1953.

49. R. M. Neal. *Bayesian Learning for Neural Networks*. Springer, 1996.

50. R. M. Neal. Improving asymptotic variance of MCMC estimators: Non-reversible chains are better. arXiv:math/0407281, 2004.

51. R. M. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 54:113–162, 2010.

52. M. Ottobre, N. S. Pillai, F. J. Pinski, and A. M. Stuart. A function space HMC algorithm with second order Langevin diffusion limit. *Bernoulli*, 22(1):60–106, 02 2016.

53. S. Patterson and Y. W. Teh. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems 26 (NIPS'13)*. 2013.

54. G. A. Pavliotis. *Stochastic Processes and Applications*. Springer, 2014.

55. A. Pazy. *Semigroups of Linear Operators and Applications to Partial Differential Equations*. Springer-Verlag, 1983.

56. R. Poncet. Generalized and hybrid Metropolis-Hastings overdamped Langevin algorithms. arXiv:1701.05833, 2017.

57. M. B. Priestley. *Spectral Analysis and Time Series*. Academic, 1981.

58. H. Qian. A decomposition of irreversible diffusion processes without detailed balance. *J. Math. Phys.*, 54:053302, 2013.

59. H. Qian, M. Qian, and X. Tang. Thermodynamics of the general diffusion process: Time-reversibility and entropy production. *J. Stat. Phys.*, 107:1129, 2002.

60. L. Rey-Bellet and K. Spiliopoulos. Irreversible Langevin samplers and variance reduction: A large deviations approach. *Nonlinearity*, 28(7):2081, 2015.

61. L. Rey-Bellet and K. Spiliopoulos. Improving the convergence of reversible samplers. *J. Stat. Phys.*, 164(3):472–494, 2016.

62. C. Robert and G. Casella. *Monte Carlo statistical methods*. Springer Science & Business Media, 2nd edition, 2004.

63. G. O. Roberts and O. Stramer. Langevin diffusions and Metropolis-Hastings algorithms. *Methodol. Comput. Appl. Probab.*, 4:337–357, 2002.

64. X. Shang, Z. Zhu, B. Leimkuhler, and A. Storkey. Covariance-controlled adaptive Langevin thermostat for large-scale Bayesian sampling. In *Advances in Neural Information Processing Systems 28 (NIPS'15)*. 2015.

65. J. Shi, T. Chen, R. Yuan, B. Yuan, and P. Ao. Relation of a new interpretation of stochastic differential equations to Itô process. *J. Stat. Phys.*, 148(3):579–590, 2012.

66. H. Tak, X.-L. Meng, and D. A. van Dyk. A repulsive-attractive Metropolis algorithm for multimodality. arXiv:1601.05633, 2016.

67. K. S. Turitsyn, M. Chertkov, and M. Vucelja. Irreversible Monte Carlo algorithms for efficient sampling. *Physica D*, 240(4–5):410–414, 2011.

68. Flegal J. M. Vats, D. and G. L. Jones. Strong consistency of multivariate spectral variance estimators in Markov chain Monte Carlo. arXiv:1507.08266, 2016.

69. Flegal J. M. Vats, D. and G. L. Jones. Multivariate output analysis for Markov chain Monte Carlo. 2017.

70. C. Villani. *Hypocoercivity*. American Mathematical Soc., 2009.

71. M. Vucelja. Lifting – a nonreversible Markov chain Monte Carlo algorithm. arXiv:1412.8762, 2015.

72. M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML'11)*, pages 681–688, June 2011.

73. S.-J. Wu, C.-R. Hwang, and M. T. Chu. Attaining the optimal Gaussian diffusion acceleration. *J. Stat. Phys.*, 155(3):571–590, 2014.

74. T. Xifara, C. Sherlock, S. Livingstone, S. Byrne, and M. Girolami. Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Stat. Probabil. Lett.*, 91:14–19, 2014.

75. L. Yin and P. Ao. Existence and construction of dynamical potential in nonequilibrium processes without detailed balance. *J. Phys. A*, 39(27):8593, 2006.

Irreversible Jump