
Magnetic Hamiltonian Monte Carlo

Nilesh Tripuraneni¹ Mark Rowland² Zoubin Ghahramani^{2,3} Richard Turner²

Abstract

Hamiltonian Monte Carlo (HMC) exploits Hamiltonian dynamics to construct efficient proposals for Markov chain Monte Carlo (MCMC). In this paper, we present a generalization of HMC which exploits *non-canonical* Hamiltonian dynamics. We refer to this algorithm as magnetic HMC, since in 3 dimensions a subset of the dynamics map onto the mechanics of a charged particle coupled to a magnetic field. We establish a theoretical basis for the use of non-canonical Hamiltonian dynamics in MCMC, and construct a symplectic, leapfrog-like integrator allowing for the implementation of magnetic HMC. Finally, we exhibit several examples where these non-canonical dynamics can lead to improved mixing of magnetic HMC relative to ordinary HMC.

1. Introduction

Probabilistic inference in complex models generally requires the evaluation of intractable, high-dimensional integrals. One powerful and generic approach to inference is to use Markov chain Monte Carlo (MCMC) methods to generate asymptotically exact (but correlated) samples from a posterior distribution for inference and learning. Hamiltonian Monte Carlo (HMC) (Duane et al., 1987; Neal, 2011) is a state-of-the-art MCMC method which uses gradient information from an absolutely continuous target density to encourage efficient sampling and exploration. Crucially, HMC utilizes proposals inspired by Hamiltonian dynamics (corresponding to the classical mechanics of a point particle) which can traverse long distances in parameter space. HMC, and variants like NUTS (which eliminates the need to hand-tune the algorithm’s hyperparameters), have been successfully applied to a large class of probabilistic inference problems where they are often the gold standard for

(asymptotically) exact inference (Neal, 1996; Hoffman & Gelman, 2014; Carpenter et al., 2016).




Figure 1. Example sample paths for standard HMC (left) and MHMC (right) for an isotropic Gaussian target distribution.

In this paper, we first review important properties of Hamiltonian dynamics, namely energy-preservation, symplecticity, and time-reversibility, and derive a more general class of dynamics with these properties which we refer to as *non-canonical* Hamiltonian dynamics. We then discuss the relationship of non-canonical Hamiltonian dynamics to well-known variants of HMC and propose a novel extension of HMC. We refer to this method as magnetic HMC (see Algorithm 1) since it corresponds to a particular subset of the non-canonical dynamics that in 3 dimensions map onto to the mechanics of a charged particle coupled to a magnetic field – see Figure 1 for an example of these dynamics. Furthermore, we construct an explicit, symplectic, leapfrog-like integrator for magnetic HMC which allows for an efficient numerical integration scheme comparable to that of ordinary HMC. Finally, we evaluate the performance of magnetic HMC on several sampling problems where we show how its non-canonical dynamics can lead to improved mixing. The proofs of all results in this paper are presented in the corresponding sections of the Appendix.

2. Markov chain Monte Carlo

Given an unnormalized target density $\rho(\theta)$ defined on \mathbb{R}^d , an MCMC algorithm constructs an ergodic Markov chain $(\Theta_n)_{n \in \mathbb{N}}$ such that the distribution of Θ_n converges to ρ (e.g. in total variation) (Robert & Casella, 2004). Often, the transition kernel of such a Markov chain is specified by the Metropolis-Hastings (MH) algorithm which (i) given the

¹UC Berkeley, USA ²University of Cambridge, UK
³Uber AI Labs, USA. Correspondence to: Nilesh Tripuraneni
<nileshtrip@gmail.com>.

current state $\Theta_n = \theta$, proposes a new state $\tilde{\theta}$ by sampling from a proposal distribution $Q(\cdot|\theta)$, and (ii) sets $\Theta_{n+1} = \theta$ with probability $\min\left(1, \frac{\rho(\tilde{\theta})Q(\theta|\tilde{\theta})}{\rho(\theta)Q(\tilde{\theta}|\theta)}\right)$ and $\Theta_{n+1} = \theta$ otherwise. The role of the acceptance step is to enforce reversibility (or detailed balance) of the Markov chain with respect to ρ – which implies ρ is a stationary distribution of the transition kernel.

Heuristically, a good MH algorithm should have low inter-sample correlation while maintaining a high acceptance ratio. Hamiltonian Monte Carlo provides an elegant mechanism to do this by simulating a particle moving along the contour lines of a dynamical system, constructed from the target density, to use as a MCMC proposal.

2.1. Hamiltonian Monte Carlo

In Hamiltonian Monte Carlo, the target distribution is augmented with “momentum” variables \mathbf{p} which are independent of the θ variables but of equal dimension. For the remainder of the paper, we take the distribution over the momentum variables to be Gaussian, as is common in the literature (indeed, there is evidence that in many cases, the choice of a Gaussian distribution may be optimal (Betancourt, 2017)). The joint target distribution is therefore:

$$\rho(\theta, \mathbf{p}) \propto e^{-U(\theta) - \mathbf{p}^\top \mathbf{p}/2} \equiv e^{-H(\theta, \mathbf{p})}. \quad (1)$$

Crucially, this augmentation allows Hamiltonian dynamics to be used as a proposal for an MCMC algorithm over the space (θ, \mathbf{p}) , where we interpret θ (resp., \mathbf{p}) as position (resp., momentum) coordinates of a physical particle with total energy $H(\theta, \mathbf{p})$, given by the sum of its potential energy $U(\theta)$ and kinetic energy $\mathbf{p}^\top \mathbf{p}/2$. We briefly review the Markov chain construction below; see (Neal, 2011) or (Duane et al., 1987) for a more detailed description. Given the Markov chain state (θ_n, \mathbf{p}_n) at time n , the new state for time $n + 1$ is obtained by first resampling momentum $\mathbf{p}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and then proposing a new state according to the following steps: (i) Simulate the deterministic Hamiltonian flow defined by the differential equation

$$\begin{aligned} \frac{d}{dt} \begin{bmatrix} \theta(t) \\ \mathbf{p}(t) \end{bmatrix} &= \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} \nabla_{\theta} H(\mathbf{p}(t), \theta(t)) \\ \nabla_{\mathbf{p}} H(\mathbf{p}(t), \theta(t)) \end{bmatrix} \\ &\equiv \begin{bmatrix} \mathbf{p}(t) \\ -\nabla_{\theta} U(\theta(t)) \end{bmatrix}. \end{aligned} \quad (2)$$

for time τ , with initial condition (θ_n, \mathbf{p}_n) , to obtain $(\theta'_n, \mathbf{p}'_n) = \Phi_{\tau, H}(\theta_n, \mathbf{p}_n)$ ¹; (ii) Flip the resulting momentum component with the map $\Phi_{\mathbf{p}}(\theta, \mathbf{p}) = (\theta, -\mathbf{p})$ to ob-

tain $(\tilde{\theta}_{n+1}, \tilde{\mathbf{p}}_{n+1}) = \Phi_{\mathbf{p}}(\theta'_n, \mathbf{p}'_n) = \tilde{\Phi}_{\tau, H}(\theta_n, \mathbf{p}_n)$; (iii) Apply a MH-type accept/reject step to enforce detailed balance with respect to the target distribution; (iv) Flip the momentum again with $\Phi_{\mathbf{p}}$ so it points in the original direction.

Note that because the map $\Phi_{\tau, H}$ is time-reversible (in the sense that if the path $(\theta(t), \mathbf{p}(t))$ is a solution to (2) then the path with negated momentum traversed in reverse $(\theta(-t), -\mathbf{p}(-t))$ is also a solution), the map $\tilde{\Phi}_{\tau, H}$ is self-inverse. From this, the acceptance ratio in step (iii) enforcing detailed balance can be shown (see e.g. (Green, 1995)) to have the form:

$$\min\left(1, \frac{\exp(-H(\tilde{\theta}_{n+1}, \tilde{\mathbf{p}}_{n+1}))}{\exp(-H(\theta_n, \mathbf{p}_n))} \left| \det \nabla_{\theta, \mathbf{p}} \tilde{\Phi}_{\tau, H}(\theta_n, \mathbf{p}_n) \right| \right). \quad (3)$$

Note that the Hamiltonian flow & momentum flip operator $\tilde{\Phi}_{\tau, H}$ is volume-preserving², which immediately yields that the Jacobian term in the acceptance ratio (3) is simply 1. The acceptance probability therefore reduces to $\min(1, \exp(H(\theta_n, \mathbf{p}_n) - H(\tilde{\theta}_{n+1}, \tilde{\mathbf{p}}_{n+1})))$. Furthermore, since the Hamiltonian flow defined in (2) is energy-preserving (i.e. $H(\tilde{\theta}_{n+1}, \tilde{\mathbf{p}}_{n+1}) = H(\theta_n, \mathbf{p}_n)$) – the acceptance ratio is identically 1. Moreover, the momentum resampling in (i) and momentum flip in (iv) both leave the joint distribution invariant.

While the momentum resampling ensures the Markov chain explores the joint (θ, \mathbf{p}) space, the proposals inspired by Hamiltonian dynamics can traverse long distances in parameter space θ , reducing the random-walk behavior of MH that often results in highly correlated samples (Neal, 2011).

2.2. Symplectic Numerical Integration

Unfortunately, it is rarely possible to integrate the flow defined in (2) analytically; instead an efficient numerical integration scheme must be used to generate a proposal for the MH-type accept/reject test. Typically, the leapfrog (Störmer-Verlet) integrator is used since it is an explicit method that is both symplectic and time-reversible (Neal, 2011). One elegant way to motivate this integrator is by decomposing the Hamiltonian into a symmetric splitting:

$$H(\theta, \mathbf{p}) = \underbrace{U(\theta)/2}_{H_1(\theta)} + \underbrace{\mathbf{p}^\top \mathbf{p}/2}_{H_2(\mathbf{p})} + \underbrace{U(\theta)/2}_{H_1(\theta)} \quad (4)$$

and then defining $\Phi_{\epsilon, H_1(\theta)}$ and $\Phi_{\epsilon, H_2(\mathbf{p})}$ to be the exactly-integrated flows for the sub-Hamiltonians $H_1(\theta)$ and

Hamiltonian flow associated with H for time τ . In addition, $\tilde{\Phi}_{\tau, H}$ denotes the composition of $\Phi_{\tau, H}$ with the momentum flip map $\Phi_{\mathbf{p}}$.

²In fact the Hamiltonian flow satisfies the stronger condition of symplecticity with respect to the \mathbf{A} matrix $([\nabla_{\theta, \mathbf{p}} \Phi_{\tau, H}(\theta, \mathbf{p})]^\top \mathbf{A}^{-1} [\nabla_{\theta, \mathbf{p}} \Phi_{\tau, H}(\theta, \mathbf{p})]) = \mathbf{A}^{-1}$ which immediately implies it is volume-preserving by taking determinants of this relation.

¹Throughout this paper, we use $\Phi_{\tau, H}$ to denote the map that takes a given position-momentum pair as initial conditions for the

$H_2(\mathbf{p})$, respectively. These updates (which are equivalent to Euler translations) can be written:

$$\begin{aligned}\Phi_{\epsilon, H_1(\theta)} \begin{bmatrix} \theta \\ \mathbf{p} \end{bmatrix} &= \begin{bmatrix} \theta \\ \mathbf{p} - \frac{\epsilon}{2} \nabla_{\theta} U(\theta) \end{bmatrix} \\ \Phi_{\epsilon, H_2(\mathbf{p})} \begin{bmatrix} \theta \\ \mathbf{p} \end{bmatrix} &= \begin{bmatrix} \theta + \epsilon \mathbf{p} \\ \mathbf{p} \end{bmatrix}\end{aligned}\quad (5)$$

since the Hamilton equations (2) for the sub-Hamiltonians $H_1(\theta)$ and $H_2(\mathbf{p})$ are linear, and hence analytically integrable. One leapfrog step is then defined as:

$$\Phi_{\epsilon, H(\theta, \mathbf{p})}^{\text{frog}} = \Phi_{\epsilon, H_1(\theta)} \circ \Phi_{\epsilon, H_2(\mathbf{p})} \circ \Phi_{\epsilon, H_1(\theta)} \quad (6)$$

with the overall proposal given by L leapfrog steps, followed by the momentum flip operator $\Phi_{\mathbf{p}}$ as before:

$$\tilde{\Phi}_{L, \epsilon, H}^{\text{frog}} = \Phi_{\mathbf{p}} \circ \left(\Phi_{\epsilon, H(\theta, \mathbf{p})}^{\text{frog}} \right)^L.$$

As each of the flows $\Phi_{\epsilon, H_1(\theta)}$, $\Phi_{\epsilon, H_2(\mathbf{p})}$ exactly integrates a sub-Hamiltonian, they inherit the symplecticity, volume-preservation, and time-reversibility of the exact dynamics. Moreover, since the composition of symplectic flows is also symplectic and the splitting scheme is symmetric (implying the composition of time-reversible flows is also time-reversible), the Jacobian term in the acceptance probability (3) is exactly 1 as in the case of perfect simulation.

The leapfrog scheme will not exactly preserve the Hamiltonian H , so the remaining acceptance ratio $\exp(H(\theta_n, \mathbf{p}_n) - H(\tilde{\theta}_{n+1}, \tilde{\mathbf{p}}_{n+1}))$ must be calculated. However, the leapfrog integrator has error $\mathcal{O}(\epsilon^3)$ in one leapfrog step (Hairer et al., 2006). This error scaling will lead to good energy conservation properties (and thus high acceptance rates in the MH step), even when simulating over long trajectories.

3. Non-Canonical Hamiltonian Monte Carlo

In Section 2, we noted the role time-reversibility, volume-preservation, and energy conservation of canonical Hamiltonian dynamics play in making them useful candidates for MCMC. In this section, we develop the properties of a general class of flows we refer to as *non-canonical* Hamiltonian systems that parallel these properties, we use to construct our method magnetic HMC (see Algorithm 1):

Lemma 1. *The map $\Phi_{\tau, H}^{\mathbf{A}}(\theta, \mathbf{p})$ defined by integrating the non-canonical Hamiltonian system*

$$\frac{d}{dt} \begin{bmatrix} \theta(t) \\ \mathbf{p}(t) \end{bmatrix} = \mathbf{A} \nabla_{\theta, \mathbf{p}} H(\theta(t), \mathbf{p}(t)) \quad (7)$$

with initial conditions (θ, \mathbf{p}) for time τ , where $\mathbf{A} \in \mathcal{M}_{2n \times 2n}$ is any invertible, antisymmetric matrix induces a flow on the coordinates (θ, \mathbf{p}) that is still energy-conserving ($\partial_{\tau} H(\Phi_{\tau, H}^{\mathbf{A}}(\theta, \mathbf{p})) = 0$) and symplectic with

respect to \mathbf{A} ($[\nabla_{\theta, \mathbf{p}} \Phi_{\tau, H}(\theta, \mathbf{p})]^{\top} \mathbf{A}^{-1} [\nabla_{\theta, \mathbf{p}} \Phi_{\tau, H}(\theta, \mathbf{p})] = \mathbf{A}^{-1}$) which also implies volume-preservation of the flow.

Within the formal construction of classical mechanics, it is known that any Hamiltonian flow defined on the cotangent bundle (θ, \mathbf{p}) of a configuration manifold, which is equipped with an arbitrary symplectic 2-form, will preserve its symplectic structure and admit the corresponding Hamiltonian as a first integral invariant (Arnold, 1989). The statement of Lemma 1 is simply a restatement of this fact grounded in a coordinate system. Similar arbitrary, antisymmetric terms have also appeared in the study of MCMC algorithms based on diffusion processes; such samplers often do not enforce detailed balance with respect to the target density and are often implemented as discretizations of stochastic differential equations (Rey-Bellet & Spiliopoulos, 2015; Ma et al., 2015), in contrast to the approach taken here.

Our second observation is that the dynamics in (7) are *not* time-reversible in the traditional sense. Instead, if we consider the parametrization of \mathbf{A} as:

$$\mathbf{A} = \begin{bmatrix} \mathbf{E} & \mathbf{F} \\ -\mathbf{F}^{\top} & \mathbf{G} \end{bmatrix} \quad (8)$$

where \mathbf{E} , \mathbf{G} are antisymmetric and \mathbf{F} is taken to be general such that \mathbf{A} is invertible, then the *non-canonical* dynamics have a (pseudo) time-reversibility symmetry:

Lemma 2. *If $(\theta(t), \mathbf{p}(t))$ is a solution to the non-canonical dynamics:*

$$\frac{d}{dt} \begin{bmatrix} \theta(t) \\ \mathbf{p}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{E} & \mathbf{F} \\ -\mathbf{F}^{\top} & \mathbf{G} \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} \nabla_{\theta} H(\theta(t), \mathbf{p}(t)) \\ \nabla_{\mathbf{p}} H(\theta(t), \mathbf{p}(t)) \end{bmatrix} \quad (9)$$

then $(\tilde{\theta}(t), \tilde{\mathbf{p}}(t)) = (\theta(-t), -\mathbf{p}(-t))$ is a solution to the modified non-canonical dynamics:

$$\frac{d}{dt} \begin{bmatrix} \tilde{\theta}(t) \\ \tilde{\mathbf{p}}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} -\mathbf{E} & \mathbf{F} \\ -\mathbf{F}^{\top} & -\mathbf{G} \end{bmatrix}}_{\tilde{\mathbf{A}}} \begin{bmatrix} \nabla_{\tilde{\theta}} H(\tilde{\theta}(t), \tilde{\mathbf{p}}(t)) \\ \nabla_{\tilde{\mathbf{p}}} H(\tilde{\theta}(t), \tilde{\mathbf{p}}(t)) \end{bmatrix} \quad (10)$$

if $H(\theta, \mathbf{p}) = H(\theta, -\mathbf{p})$. In particular if $\mathbf{E} = \mathbf{G} = 0$ then $\mathbf{A} = \tilde{\mathbf{A}}$, which reduces to the traditional time-reversal symmetry of canonical Hamiltonian dynamics.

Lemma 1 suggests a generalization of HMC that can utilize an arbitrary invertible antisymmetric \mathbf{A} matrix in its dynamics; however Lemma 2 indicates the non-canonical dynamics lack a traditional time-reversibility symmetry which poses a potential difficulty to satisfying detailed balance. In particular, we cannot compose $\Phi_{\mathbf{p}}$ with an exact/approximate simulation of $\Phi_{\tau, H}^{\mathbf{A}}$ to make $\tilde{\Phi}_{\tau, H}^{\mathbf{A}} = \Phi_{\mathbf{p}} \circ \Phi_{\tau, H}^{\mathbf{A}}$ self-inverse.

Our solution to obtaining a time-reversible proposal is simply to flip the elements of the \mathbf{E} and \mathbf{G} matrices just as ordinary HMC flips the auxiliary variable \mathbf{p} i) at the end of Hamiltonian flow in the proposal and ii) once again after the MH acceptance step to return \mathbf{p} to its original direction. In this vein, we view the parameters \mathbf{E} and \mathbf{G} as auxiliary variables in the state space, and simultaneously flip \mathbf{p} , \mathbf{E} , and \mathbf{G} after having simulated the dynamics, rendering the proposal time-reversible according to Lemma 2 – see Section 2 in the Appendix for full details of this construction. This ensures that detailed balance is satisfied for this entire proposal. To avoid “random walk” behaviour in the resulting Markov chain, we can apply a sign flip to \mathbf{E} and \mathbf{G} , in addition to \mathbf{p} , to return them to their original directions after the MH acceptance step.

The validity of this construction relies on equipping \mathbf{E} and \mathbf{G} with symmetric auxiliary distributions. For the remainder of this paper, we further restrict to binary symmetric auxiliary distributions supported on a given antisymmetric matrix \mathbf{V}_0 and its sign flip $-\mathbf{V}_0$ – see Appendix 1.1 for full details. This restriction is not necessary, but gives rise to a simple and interpretable class of algorithms, which is in spirit closest to using fixed parameters \mathbf{E} and \mathbf{G} , whilst ensuring the proposal satisfies detailed balance. This construction is also reminiscent of lifting constructions prevalent in the discrete Markov chain literature (Chen et al., 1999); heuristically, the signed variables \mathbf{E} and \mathbf{G} favour proposals in opposing directions.

3.1. Symplectic Numerical Integration for Non-Canonical Dynamics

As with standard HMC, exactly simulating the flow $\Phi_{\tau,H}^{\mathbf{A}}$ is rarely tractable, and a numerical integrator is required to approximate the flow. It is not absolutely necessary to use an explicit, symplectic integration scheme; indeed implicit integrators are used in Riemannian HMC to maintain symplecticity of the proposal which comes at a greater complexity and computational cost (Girolami et al., 2009). However explicit, symplectic integrators are simple, have good energy-conservation properties, and are volume-preserving/time-reversible (Hairer et al., 2006), so for the present discussion we restrict our attention to investigating leapfrog-like schemes.

We begin, as in Section 2.2, by considering the symmetric splitting (4), yielding the sub-Hamiltonians $H_1(\theta) = U(\theta)/2$, $H_2(\mathbf{p}) = \mathbf{p}^\top \mathbf{p}/2$. The corresponding non-canonical dynamics for the sub-Hamiltonians $H_1(\theta)$ and $H_2(\mathbf{p})$ are:

$$\frac{d}{dt} \begin{bmatrix} \theta \\ \mathbf{p} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{E} & \mathbf{F} \\ -\mathbf{F}^\top & \mathbf{G} \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} \nabla_\theta U(\theta)/2 \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{E} \nabla_\theta U(\theta)/2 \\ -\mathbf{F}^\top \nabla_\theta U(\theta)/2 \end{bmatrix}$$

and:

$$\frac{d}{dt} \begin{bmatrix} \theta \\ \mathbf{p} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{E} & \mathbf{F} \\ -\mathbf{F}^\top & \mathbf{G} \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} \mathbf{0} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{F} \mathbf{p} \\ \mathbf{G} \mathbf{p} \end{bmatrix}.$$

We denote the corresponding flows by $\Phi_{\epsilon,H_1(\theta)}^{\mathbf{A}}$ and $\Phi_{\epsilon,H_2(\mathbf{p})}^{\mathbf{A}}$ respectively. The flow $\Phi_{\epsilon,H_1(\theta)}^{\mathbf{A}}$ is generally not explicitly tractable unless we take $\mathbf{E} = \mathbf{0}$ – in which case it is solved by an Euler translation as before. Crucially, the flow in $\Phi_{\epsilon,H_2(\mathbf{p})}^{\mathbf{A}}$ is a *linear* differential equation and hence analytically integrable. If \mathbf{G} is invertible (and $\mathbf{F} = \mathbf{I}$) then:

$$\Phi_{\epsilon,H_2(\mathbf{p})}^{\mathbf{A}} \begin{bmatrix} \theta \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \theta + \mathbf{G}^{-1}(\exp(\mathbf{G}\epsilon) - \mathbf{I})\mathbf{p} \\ \exp(\mathbf{G}\epsilon)\mathbf{p} \end{bmatrix}. \quad (11)$$

See the Appendix for a detailed derivation which also handles the general case where \mathbf{G} is not invertible. Thus when $\mathbf{E} = \mathbf{0}$, the flows $\Phi_{\epsilon,H_1(\theta)}^{\mathbf{A}}$ and $\Phi_{\epsilon,H_2(\mathbf{p})}^{\mathbf{A}}$ are analytically tractable and will inherit the generalized symplecticity and (pseudo) time-reversibility of the exact dynamics in (7). Therefore if we use the symmetric splitting (4) to construct a leapfrog-like step:

$$\Phi_{\epsilon,H(\theta,\mathbf{p})}^{\text{frog},\mathbf{A}} = \Phi_{\epsilon,H_1(\theta)}^{\mathbf{A}} \circ \Phi_{\epsilon,H_2(\mathbf{p})}^{\mathbf{A}} \circ \Phi_{\epsilon,H_1(\theta)}^{\mathbf{A}} \quad (12)$$

we can construct a total proposal that consists of several leapfrog steps, followed by a flip of the momentum and \mathbf{G} , $\Phi_{\mathbf{p}} \circ \Phi_{\mathbf{G}}$, which will be a volume-preserving, self-inverse map:

$$\tilde{\Phi}_{\epsilon,H(\theta,\mathbf{p})}^{\text{frog},\mathbf{A}} = \Phi_{\mathbf{p}} \circ \Phi_{\mathbf{G}} \circ (\Phi_{\epsilon,H_1(\theta)}^{\mathbf{A}} \circ \Phi_{\epsilon,H_2(\mathbf{p})}^{\mathbf{A}} \circ \Phi_{\epsilon,H_1(\theta)}^{\mathbf{A}})^L. \quad (13)$$

Henceforth we will always take $\mathbf{E} = \mathbf{0}$ when we use $\Phi_{\epsilon,H(\theta,\mathbf{p})}^{\text{frog},\mathbf{A}}$ to generate leapfrog proposals, which interestingly corresponds to a magnetic dynamics as discussed in Lemma 4. A full description of the magnetic HMC algorithm using this numerical integrator is described in Section 5.

4. Special Cases

Here, we describe several tractable subcases of the general formulation of non-canonical Hamiltonian dynamics since these they have interesting physical interpretations.

4.1. Mass Preconditioned Dynamics

One simple variant of HMC is preconditioned HMC where $\mathbf{p} \sim \mathcal{N}(0, \mathbf{M})$ (Neal, 2011), and can be implemented nearly identically to ordinary HMC. We note that preconditioning can be recovered within our framework using a simple form for the non-canonical \mathbf{A} matrix:

Lemma 3. *i) Preconditioned HMC with momentum variable $\mathbf{p} \sim \mathcal{N}(0, \mathbf{M})$ in the (θ, \mathbf{p}) coordinates is exactly equivalent to simulating non-canonical HMC with $\mathbf{p}' = \mathbf{M}^{-1/2}\mathbf{p} \sim \mathcal{N}(0, \mathbf{I})$ and the non-canonical matrix $\mathbf{A} =$*

$\begin{bmatrix} \mathbf{0} & \mathbf{M}^{1/2} \\ -(\mathbf{M}^{1/2})^\top & \mathbf{0} \end{bmatrix}$, and then transforming back to (θ, \mathbf{p}) coordinates using $\mathbf{p} = \mathbf{M}^{1/2}\mathbf{p}'$. Here $\mathbf{M}^{1/2}$ is a Cholesky factor for \mathbf{M} .

ii) On the other hand if we apply a change of basis (via an invertible matrix \mathbf{F}) to our coordinates $\theta' = \mathbf{F}^{-1}\theta$, simulate HMC in the (θ', \mathbf{p}) coordinates, and transform back to the original basis using \mathbf{F} , this is exactly equivalent to non-canonical HMC with $\mathbf{A} = \begin{bmatrix} \mathbf{0} & \mathbf{F} \\ -\mathbf{F}^\top & \mathbf{0} \end{bmatrix}$.

Lemma 3 illustrates a fact alluded to in (Neal, 2011); using a mass preconditioning matrix \mathbf{M} and a change of basis given by matrix $-(\mathbf{M}^{-1/2})^\top$ acting on θ leaves the HMC algorithm invariant.

4.2. Magnetic Dynamics

The primary focus of this paper is to investigate the subcase of the dynamics where:

$$\mathbf{A} = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{G} \end{pmatrix} \quad (14)$$

for two important reasons³. Firstly for this choice of \mathbf{A} we can construct an explicit, symplectic, leapfrog-like integration scheme which is important for developing an efficient HMC sampler as discussed in Section 3.1. Secondly, the dynamics have an interesting physical interpretation quite distinct from mass preconditioning and other HMC variants like Riemannian HMC (Girolami et al., 2009):

Lemma 4. *In 3 dimensions the non-canonical Hamiltonian dynamics corresponding to the Hamiltonian $H(\theta, \mathbf{p}) = U(\theta) + \frac{1}{2}\mathbf{p}^\top \mathbf{p}$ and \mathbf{A} matrix as in (14) are equivalent to the Newtonian mechanics of a charged particle (with unit mass and charge) coupled to a magnetic field $\vec{\mathbf{B}}$ (given by a particular function of \mathbf{G} - see Appendix): $\frac{d^2\theta}{dt^2} = -\nabla_\theta U(\theta) + \frac{d\theta}{dt} \times \vec{\mathbf{B}}$.*

This interpretation is perhaps surprising since Hamiltonian formulations of classical magnetism are uncommon, although the quantum mechanical treatment naturally incorporates a Hamiltonian framework. However, in light of Lemma 3 we might wonder if by a clever rewriting of the Hamiltonian we can reproduce this system of ODEs using the canonical \mathbf{A} matrix (i.e. $\mathbf{E} = \mathbf{G} = \mathbf{0}$, $\mathbf{F} = \mathbf{I}$). This is not the case:

Lemma 5. *The non-canonical Hamiltonian dynamics with magnetic \mathbf{A} and Hamiltonian $H(\theta, \mathbf{p}) = U(\theta) + \frac{1}{2}\mathbf{p}^\top \mathbf{p}$ cannot be obtained using canonical Hamiltonian dynamics for any choice of smooth Hamiltonian. (See Appendix).*

³Note that the effect of a non-identity \mathbf{F} matrix can be achieved by simply composing these magnetic dynamics with a coordinate-transformation as suggested in Lemma 3.

5. The Magnetic HMC (MHMC) Algorithm

Using the results discussed in Section 3 and Section 3.1 we can now propose Magnetic HMC – see Algorithm 1.

Algorithm 1 Magnetic HMC (MHMC)

Input: $H, \mathbf{G}, L, \epsilon$

Initialize (θ_0, \mathbf{p}_0) , and set $\mathbf{G}_0 \leftarrow \mathbf{G}$

for $n = 1, \dots, N$ **do**

 Resample $\mathbf{p}_{n-1} \sim N(\mathbf{0}, \mathbf{I})$

 Set $(\tilde{\theta}_n, \tilde{\mathbf{p}}_n) \leftarrow \text{LF}(H, L, \epsilon, (\theta_{n-1}, \mathbf{p}_{n-1}, \mathbf{G}_{n-1}))$ with $\Phi_{\epsilon, H_2(\mathbf{p})}^{\mathbf{A}}$ as in (11)

 Flip momentum $(\tilde{\theta}_n, \tilde{\mathbf{p}}_n) \leftarrow (\tilde{\theta}_n, -\tilde{\mathbf{p}}_n)$ and set $\tilde{\mathbf{G}}_n \leftarrow -\mathbf{G}_{n-1}$

if $\text{Unif}([0, 1]) < \min(1, \exp(H(\theta_{n-1}, \mathbf{p}_{n-1}) - H(\tilde{\theta}_n, \tilde{\mathbf{p}}_n)))$ **then**

 Set $(\theta_n, \mathbf{p}_n, \mathbf{G}_n) \leftarrow (\tilde{\theta}_n, \tilde{\mathbf{p}}_n, \tilde{\mathbf{G}}_n)$

else

 Set $(\theta_n, \mathbf{p}_n, \mathbf{G}_n) \leftarrow (\theta_{n-1}, \mathbf{p}_{n-1}, \mathbf{G}_{n-1})$

end if

 Flip momentum $\mathbf{p}_n \leftarrow -\mathbf{p}_n$ and flip $\mathbf{G}_n \leftarrow -\mathbf{G}_n$

end for

Output: $(\theta_n)_{n=0}^N$

One further remark is that by construction the integrator for magnetic HMC is expected to have similarly good energy conservation properties to the integrator of standard HMC:

Lemma 6. *The symplectic leapfrog-like integrator for magnetic HMC will have the same local ($\sim \mathcal{O}(\epsilon^3)$) and global ($\sim \mathcal{O}(\epsilon^2)$) error scaling (over $\tau \sim \frac{L}{\epsilon}$ steps), as the canonical leapfrog integrator of standard HMC if the Hamiltonian is separable. (See Appendix).*

It is worthwhile to contrast the algorithmic differences between magnetic HMC and ordinary HMC. Intuitively, the role of the flow $\Phi_{\epsilon, H_2(\mathbf{p})}^{\mathbf{A}}$ – which reduces to the standard Euler translation update of ordinary HMC when $\mathbf{G} = \mathbf{0}$ – is to introduce a *rotation* into the momentum space of the flow. In particular, a non-zero element \mathbf{G}_{ij} will allow momentum to periodically flow between p_i and p_j . If we regard \mathbf{G} as an element in the Lie algebra of antisymmetric matrices, which can be thought of as infinitesimal rotations, then the exponential map $\exp(\mathbf{G}\epsilon)$ will project this transformation into the Lie group of real orthogonal linear maps.

With respect to computational cost, although magnetic HMC requires matrix exponentiation/diagonalization to simulate $\Phi_{\epsilon, H_2(\mathbf{p})}^{\mathbf{A}}$, this only needs to be computed once upfront for $\pm\mathbf{G}$ and cached; moreover, as $\pm\mathbf{G}$ is diagonalizable, the exact exponential can be calculated in $\mathcal{O}(d^3)$ time. Additionally, there is an $\mathcal{O}(d^2)$ cost for the matrix-vector products needed to implement the flow $\Phi_{\epsilon, H_2(\mathbf{p})}^{\mathbf{A}}$ as with preconditioning. However, it is possible to design sparsi-

fied matrix representations of \mathbf{A} which will translate into sparsified rotations if we only wish to "curl" in a specific subspace of dimension d_0 – which will translate into a computational cost of $\mathcal{O}(d_0^3)$ and $\mathcal{O}(d_0^2)$ respectively.

An important problem to address is the selection of the \mathbf{G} matrix, which affords a great deal of flexibility to MHMC relative to HMC; this point is further discussed in the Experiments section, where we argue that in certain cases intuitive heuristics can be used to select the \mathbf{G} matrix.

6. Experiments

Here we investigate the performance of magnetic HMC against standard HMC in several examples; in each case commenting on our choice of the magnetic field term \mathbf{G} . Step sizes (ϵ) and number of leapfrog steps (L) were tuned to achieve an acceptance rate between .7 – .8, after which the norm of the non-zero elements in \mathbf{G} was set to $\sim .1 - .2$ which was found to work well.

In the Appendix we also display illustrations of different MHMC proposals across several targets in order to provide more intuition for MHMC's dynamics. Further experimental details and an additional experiment on a Gaussian funnel target are also provided in the Appendix.

6.1. Multiscale Gaussians

We consider two highly ill-conditioned Gaussians similar to as in (Sohl-Dickstein et al., 2014) to illustrate a heuristic for \mathbf{G} matrix selection and demonstrate properties of the magnetic dynamics. In particular we consider a centered, uncorrelated 2D Gaussian with covariance eigenvalues of 10^6 and 1 as well as a centered, uncorrelated 10D Gaussian with two large covariance eigenvalues of 10^6 and remaining eigenvalues of 1. We denote their coordinates as $\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$ and $\mathbf{x} = (x_1, \dots, x_{10}) \in \mathbb{R}^{10}$ respectively. HMC will have difficulty exploring the directions of




Figure 2. Averaged Autocorrelation of HMC vs MHMC on a 10D ill-conditioned Gaussian (left) and Averaged Autocorrelation of HMC vs MHMC on a 2D ill-conditioned Gaussian.

large marginal variance since its exploration will often be limited by the smaller variance directions. Accordingly, in order to induce a periodic momentum flow between the directions of small and large variance, we introduce nonzero components \mathbf{G}_{ij} into the subspaces spanned directly between the large and small eigenvalues. Indeed, we find that magnetic \mathbf{G} term is encouraging more efficient exploration as we can see from the averaged autocorrelation of samples generated from the HMC/MHMC chains – see Figure 2. Further, by running the 50 parallel chains for 10^7 timesteps, we computed both the bias and Monte Carlo standard errors (MCSE) of the estimators of the target moments as shown in Table 1 and Table 2.

Table 1. Absolute Bias \pm MCSE for 2D ill-conditioned Gaussian moments for HMC vs. MHMC. Note that $\mathbb{E}[x_1^2] = 10^6$ and $\mathbb{E}[x_2^2] = 1$.

algorithm	x_1^2 (Bias \pm MCSE)	x_2^2 (Bias \pm MCSE)
HMC	.947 \pm 41.7	.00108 \pm .0247
MHMC	.657 \pm 22.5	.000280 \pm .00438

Table 2. Absolute Bias \pm MCSE for 10D ill-conditioned Gaussian moments for HMC vs. MHMC. Note that $\mathbb{E}[x_1^2] = 10^6$ and $\mathbb{E}[x_{10}^2] = 1$.

algorithm	x_1^2 (Bias \pm MCSE)	x_{10}^2 (Bias \pm MCSE)
HMC	24.7 \pm 46.6	0.00376 \pm 0.0249
MHMC	9.68 \pm 32.7	0.00127 \pm 0.0132

6.2. Mixture of Gaussians

We compare MHMC vs. HMC on a simple, but interesting, 2D density over $\mathbf{x} = (x, y) \in \mathbb{R}^2$ comprised of an evenly weighted mixture of isotropic Gaussians: $p(\mathbf{x}) = \frac{1}{2}\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \Sigma) + \frac{1}{2}\mathcal{N}(\mathbf{x}; -\boldsymbol{\mu}, \Sigma)$ for $\sigma_x^2 = \sigma_y^2 = 1$, $\rho_{xy} = 0$ and $\boldsymbol{\mu} = (2.5, -2.5)$. This problem is challenging for HMC because the gradients in canonical Hamiltonian dynamics force it to one of the two modes. We tuned HMC to achieve an acceptance rate of $\sim .75$ and used the same ϵ, L for MHMC, generating 15000 samples from both HMC and MHMC with these settings. The addition of the magnetic field term \mathbf{G} – which has one degree of freedom in this case – introduces an asymmetric "curl" into the dynamics that pushes the sampler across the saddlepoint to the other mode allowing it to efficiently mix around both modes and between them – see Figure 3. The maximum mean discrepancy between exact samples generated from the target density and samples generated from both HMC and MHMC chains was also estimated for various magnitudes of \mathbf{G} , using a quadratic ker-




Figure 3. Left: 500 samples from HMC; Middle: 500 samples from MHMC; Right: MMD between HMC/MHMC samples for various magnitudes of the non-zero component of the magnetic field – denoted g . Note $g = 0$ corresponds to standard HMC.

nel $k(\mathbf{x}, \mathbf{x}') = (1 + \langle \mathbf{x}, \mathbf{x}' \rangle)^2$ and averaged over 100 runs of the Markov chains (Borgwardt et al., 2006). Here, we clearly see that for various values of the nonzero component of \mathbf{G} , denoted g , the samples generated by MHMC more faithfully reflect the structure of the posterior. As before, we ran 50 parallel chains for 10^7 timesteps to compute both the bias and Monte Carlo standard errors (MCSE) of the estimators of the target moments as shown in Table 3. Additional experiments over a range of ϵ , L (and corre-

Table 3. Bias \pm MCSE for 2D Mixture of Gaussians for HMC vs. MHMC with $g = 0.1$. Note that $\mathbb{E}[x] = 0$ and $\mathbb{E}[x^2] = 7.25$.

algorithm	x (Bias \pm MCSE)	x^2 (Bias \pm MCSE)
HMC	.0132 \pm .0644	0.00264 \pm 0.0114
MHMC	.00239 \pm .012	0.000596 \pm 0.00365

sponding acceptance rates) and details are included in the Appendix for this example, demonstrating similar behavior.

6.3. FitzHugh-Nagumo model

Finally, we consider the problem of Bayesian inference over the parameters of the FitzHugh-Nagumo model (a set of nonlinear ordinary differential equations, originally developed to model the behavior of axial spiking potentials in neurons) as in (Ramsay et al., 2007; Girolami & Calderhead, 2011). The FitzHugh-Nagumo model is a dynamical system $(V(t), R(t))$ defined by the following coupled differential equations:

$$\begin{aligned} \dot{V}(t) &= c(V(t) - V(t)^3/3 + R(t)) \\ \dot{R}(t) &= -(V(t) - a + bR(t))/c \end{aligned} \quad (15)$$

We consider the problem where the initial conditions

$(V(0), R(0))$ of the system (15) are known, and a set of noise-corrupted observations $(\tilde{V}(t_n), \tilde{R}(t_n))_{n=0}^N = (V_{a,b,c}(t_n) + \epsilon_n^V, R_{a,b,c}(t_n) + \epsilon_n^R)_{n=0}^N$ at discrete time points $0 = t_0 < t_1 < \dots < t_N$, are available - note that we illustrate dependence of the trajectories on the model parameters explicitly via subscripts. It is not possible to recover the true parameter values of the model from these observations, but we can obtain a posterior distribution over them by specifying a model for the observation noise and a prior distribution over the model parameters.

Similar to (Ramsay et al., 2007; Girolami & Calderhead, 2011), we assume that the observation noise variables $(\epsilon_n^V)_{n=0}^N$ and $(\epsilon_n^R)_{n=0}^N$ are iid $\mathcal{N}(0, 0.1^2)$, and take an independent $\mathcal{N}(0, 1)$ prior over each parameter a , b , and c . This yields a posterior distribution of the form

$$p(a, b, c) \propto \mathcal{N}(a; 0, 1) \mathcal{N}(b; 0, 1) \mathcal{N}(c; 0, 1) \times \prod_{n=0}^N \mathcal{N}(\tilde{V}(t_n); V_{a,b,c}(t_n), 0.1^2) \quad (16)$$

Importantly, the highly non-linear dependence of the trajectory on the parameters a , b and c yields a complex posterior distribution - see Figure 4. Full details of the model set-up can be found in (Ramsay et al., 2007; Girolami & Calderhead, 2011).

For our experiments, we used fixed parameter settings of $a = 0.2, b = 0.2, c = 3.0$ to generate 200 evenly-spaced noise-corrupted observations over the time interval $t = [0, 20]$ (as in (Ramsay et al., 2007; Girolami & Calderhead, 2011)). We performed inference over the posterior distribution of parameters (a, b, c) with this set of observations using both the HMC and MHMC algorithms, which was perturbed with a magnetic field in each of the 3 axial planes of parameters – along the ab , ac , and bc axes with magnitude $g = 0.1$. The chains were run to gener-

ate 1000 samples over 100 repetitions with settings of $\epsilon = 0.015$, $L = 10$, which resulted in an average acceptance rate of $\sim .8$. The effective sample size of each of the chains normalized per unit time was then computed for each chain.

Since each query to the posterior log-likelihood or posterior gradient log-likelihood requires solving an augmented set of differential equations as in (15), the computation time (~ 238 s) of all the methods was nearly identical. Moreover,




Figure 4. Marginal posterior density contour plot over a , b , with $c = 3$.

Table 4. HMC vs. MHMC performance targeting the Fitzhugh-Nagumo posterior parameters

algorithm	ESS	ESS(a, b, c)/time (s)
HMC	1000 318 606	4.20, 1.33, 2.55
MHMC ab	1000 349 658	4.20, 1.47, 2.77
MHMC ac	1000 336 628	4.20, 1.42, 2.64
MHMC bc	1000 326 649	4.20, 1.37, 2.73

note that all methods achieved nearly perfect mixing over the first coordinate so their effective sample size were truncated at 1000 for the a coordinate. In this example, we can see that all magnetic perturbations slightly increase the mixing rate of the sampler over each of the (b, c) coordinates with the ab perturbation performing best.

7. Discussion and Conclusion

We have investigated a framework for MCMC algorithms based on non-canonical Hamiltonian dynamics and have given a construction for an explicit, symplectic integrator that is used to implement a generalization of HMC we refer to as magnetic HMC. We have also shown several examples where the non-canonical dynamics of MHMC can improve upon the sampling performance of standard HMC. Important directions for further research include finding more automated, adaptive mechanisms to set the matrix \mathbf{G} , as well as investigating positionally-dependent magnetic field components, similar to how Riemannian HMC corresponds to local preconditioning. We believe that exploiting more general deterministic flows (such as also maintaining a non-zero \mathbf{E} in the top left-block of a general \mathbf{A} matrix) could form a fruitful area for further research on MCMC methods.

Acknowledgements

The authors thank John Aston, Adrian Weller, Maria Lomeli, Yarin Gal and the anonymous reviewers for helpful comments. MR acknowledges support by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/L016516/1 for the University of Cambridge Centre for Doctoral Training, the Cambridge Centre for Analysis. RET thanks EPSRC grants EP/M0269571 and EP/L000776/1 as well as Google for funding.

References

- Arnold, Vladimir. *Mathematical Methods of Classical Mechanics*. 1989. ISBN 0387968903.
- Betancourt, Michael. A Conceptual Introduction to Hamiltonian Monte Carlo. *arXiv:1701.02434*, 2017.
- Borgwardt, Karsten M., Gretton, Arthur, Rasch, Malte J., Kriegel, Hans Peter, Scholkopf, Bernhard, and Smola, Alex J. Integrating structured biological data by Kernel Maximum Mean Discrepancy. In *Bioinformatics*, volume 22, 2006.
- Carpenter, Bob, Gelman, Andrew, Hoffman, Matt, Lee, Daniel, Goodrich, Ben, Betancourt, Michael, Brubaker, Marcus A, Li, Peter, and Riddell, Allen. Journal of Statistical Software Stan : A Probabilistic Programming Language. *Journal of Statistical Software*, VV(Ii), 2016.
- Chen, Fang, Lovász, László, and Pak, Igor. Lifting markov chains to speed up mixing. In *Proceedings of the Thirty-first Annual ACM Symposium on Theory of Computing, STOC '99*, pp. 275–281, New York, NY, USA, 1999. ACM. ISBN 1-58113-067-8. doi: 10.1145/301250.301315.
- Duane, Simon, Kennedy, A. D., Pendleton, Brian J., and Roweth, Duncan. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216 – 222, 1987.
- Girolami, Mark and Calderhead, Ben. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, March 2011.
- Girolami, Mark, Calderhead, Ben, and Chin, Siu A. Riemannian Manifold Hamiltonian Monte Carlo. *Physics*, 73(2):35, 2009.
- Green, Peter J. Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination. *Biometrika*, 82(4):711–732, 1995.
- Hairer, Ernst, Hochbruck, Marlis, Iserles, Arieh, and Lubich, Christian. Geometric Numerical Integration. *Oberwolfach Reports*, pp. 805–882, 2006.

- Hoffman, Matt and Gelman, Andrew. The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(2008):1–31, 2014.
- Ma, Yi-an, Chen, Tianqi, and Fox, Emily B. A Complete Recipe for Stochastic Gradient MCMC. *NIPS*, pp. 1–19, 2015.
- Neal, Radford. *Bayesian Learning for Neural Networks*, volume 118. 1996.
- Neal, Radford M. MCMC using Hamiltonian Dynamics. *Handbook of Markov Chain Monte Carlo*, pp. 113–162, 2011.
- Ramsay, J. O., Hooker, G., Campbell, D., and Cao, J. Parameter estimation for differential equations: A generalized smoothing approach. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, 2007.
- Rey-Bellet, Luc and Spiliopoulos, Konstantinos. Irreversible Langevin Samplers and Variance Reduction: a Large Deviations Approach. *Nonlinearity*, 28(7):2081, 2015.
- Robert, Christian P and Casella, George. *Monte Carlo Statistical Methods*, volume 95. 2004.
- Sohl-Dickstein, J, Mudigonda, Mayur, and DeWeese, M. Hamiltonian Monte Carlo Without Detailed Balance. *Proceedings of The 31st International Conference on Machine Learning*, 32:9, 2014.