

Kinetic energy choice in Hamiltonian/hybrid Monte Carlo

BY S. LIVINGSTONE, M. F. FAULKNER

*School of Mathematics, University of Bristol, University Walk,
Bristol BS8 1TW, U.K.*

samuel.livingstone@bristol.ac.uk michael.faulkner@bristol.ac.uk

5

AND G. O. ROBERTS

Department of Statistics, University of Warwick, Coventry CV4 7AL, U.K.

gareth.o.roberts@warwick.ac.uk

SUMMARY

We consider how different choices of kinetic energy in Hamiltonian Monte Carlo affect algorithm performance. To this end, we introduce two quantities which can be easily evaluated, the composite gradient and the implicit noise. Results are established on integrator stability and geometric convergence, and we show that choices of kinetic energy that result in heavy-tailed momentum distributions can exhibit an undesirable negligible moves property, which we define. A general efficiency-robustness trade off is outlined, and implementations which rely on approximate gradients are also discussed. Two numerical studies illustrate our theoretical findings, showing that the standard choice which results in a Gaussian momentum distribution is not always optimal in terms of either robustness or efficiency.

10

15

Some key words: Markov chain Monte Carlo; Bayesian Inference; Hamiltonian Monte Carlo; Bayesian computation; Hybrid Monte Carlo; MCMC.

20

1. INTRODUCTION

Hamiltonian Monte Carlo is a Markov chain Monte Carlo method which is now both widely used in Bayesian inference, and increasingly studied and developed. The idea of the approach is to use the deterministic measure-preserving dynamics of Hamiltonian flow to promote fast exploration of a parameter space of interest. To achieve this the space is augmented with a momentum variable, and the Markov chain evolves by switching between re-sampling this momentum and solving Hamilton's equations for a prescribed amount of time. Typically the equations cannot be solved exactly and so a time-reversible and volume-preserving numerical integrator is used, with discretisation errors controlled for using a Metropolis step. Comprehensive reviews are given in Neal (2011) and Betancourt et al. (2017).

25

30

The free variables in Hamiltonian Monte Carlo are the time for which Hamilton's equations should be solved between momentum refreshments, the choice of numerical integrator and step-size, and the choice of distribution for the momentum. Typically the Störmer–Verlet or leapfrog numerical integrator is used, which gives a reasonable balance between energy preservation and computational cost. Guidelines for step-size choice are given in Beskos et al. (2013), which is typically tuned to reach a 65-80% Metropolis acceptance rate. Heuristics exist for the integration time (Hoffman & Gelman, 2014), and there is some justification for a stochastic choice given in Bou-Rabee et al. (2017).

35

This paper is concerned with the choice of momentum distribution $\nu(\cdot)$, addressing a question which has been raised previously (Barthelmé & Chopin, 2011; Stephens, 2011). We will assume throughout that $\nu(\cdot)$ possesses a density which is proportional to $\exp\{-K(p)\}$ for some $K(p)$, which we call the kinetic energy of the system. The standard choice is $K(p) = p^T p/2$, with the resulting $\nu(\cdot)$ a Gaussian distribution. The general requirements are simply that $K(p)$ is differentiable, symmetric about zero and that $\nu(\cdot)$ can be sampled from. Alternative options to the Gaussian have recently been suggested (Lu et al., 2017; Zhang et al., 2016). Here we consider how such a choice affects the algorithm as a whole, in terms of stability and convergence, and develop guidelines for practitioners.

Our key findings are that the robustness and efficiency of the method can to some degree be controlled through the quantity $\nabla K \circ \nabla U(x)$, notation defined below, which we term the composite gradient. In particular, we propose that balancing the tails of the kinetic energy with those of the potential to make this approximately linear in x should give good performance, and that no faster than linear growth is necessary for algorithm stability. When $\pi(\cdot)$ is very light-tailed, this can be done by choosing $\nu(\cdot)$ to be heavier-tailed. There are, however, serious disadvantages to choosing any heavier tails than those of a Laplace distribution, which can result in the sampler moving very slowly in certain regions of the space. We introduce a negligible moves property for Markov chains to properly characterize this behaviour. We also find that in practice considering the distribution of $\nabla K(p)$, which we term the implicit noise, is important, since this governs the behaviour of the sampler in regions where $\|\nabla U(x)\|$ is small. We suggest various choices for controlling these two quantities, and support our findings with a numerical study. We also comment on how changing the kinetic energy can affect behaviour when $\nabla U(x)$ is approximated, either through subsampling or within an exact-approximate Monte Carlo scheme.

1.1. Setup, notation and assumptions

Throughout we denote the distribution from which expectations are desired by $\pi(\cdot)$. We assume that $\pi(\cdot)$ possesses a density which is proportional to $\exp\{-U(x)\}$, for some $C^1(\mathbb{R}^d)$ potential $U : \mathbb{R}^d \rightarrow [0, \infty)$, where $x \in \mathbb{R}^d$. The Hamiltonian is formed as $H(x, p) = U(x) + K(p)$, where $p \in \mathbb{R}^d$ and $K : \mathbb{R}^d \rightarrow [0, \infty)$ is a continuously differentiable kinetic energy, and the corresponding Hamiltonian dynamics are given by

$$\dot{x} = \nabla K(p), \quad \dot{p} = -\nabla U(x). \quad (1)$$

We restrict ourselves to separable Hamiltonians, meaning the distribution for p does not depend on x , as in this case general purpose and well-studied explicit numerical integrators are available (Leimkuhler & Reich, 2004). We also focus on the leapfrog numerical integrator (Neal, 2011) and a choice of integration time $T = L\varepsilon$, where ε is the integrator step-size and L is the number of leapfrog steps. We assume that L is drawn from some distribution $\Psi(\cdot)$, which is independent of the current x and p , and denote the result of solving (1) using this integrator for T units of time given the starting position and momentum (x, p) as $\tilde{\varphi}_T(x, p)$. With these tools, the Hamiltonian Monte Carlo method is described as Algorithm 1 below.

We work on the Borel space $(\mathbb{R}^d, \mathcal{B})$ and product space \mathbb{R}^{2d} with the product σ -field. Throughout let $\|x\|$, $\|x\|_1$ and $\|x\|_\infty$ be the Euclidean, L_1 and L_∞ norms of $x \in \mathbb{R}^d$, $B_r(x) = \{y \in \mathbb{R}^d : \|y - x\| < r\}$ be the Euclidean ball of radius r centred at x , and $\partial_i = \partial/\partial x_i$. For functions $f, g : \mathbb{R}^d \rightarrow \mathbb{R}^d$, let $f \circ g(x) = f(g(x))$. Whenever a distribution is referred to as having a density then this with respect to the relevant Lebesgue measure. We write x^i for the i th point in a Markov chain produced by Hamiltonian Monte Carlo, and $x_{i\varepsilon}$ for the i th step of the leapfrog integrator within a single iteration of the method.

Algorithm 1. Hamiltonian Monte Carlo

Require x^0, ε
 For $i = 1$ to $i = n$
 Draw $p \sim \nu(\cdot)$
 Draw $L \sim \Psi(\cdot)$ and set $T = L\varepsilon$
 Propose $(x^*, p^*) \leftarrow \tilde{\varphi}_T(x^{i-1}, p)$
 Set $x^i \leftarrow x^*$ with probability $1 \wedge \exp\{H(x^{i-1}, p) - H(x^*, p^*)\}$,
 otherwise set $x^i \leftarrow x^{i-1}$
 Output $x = (x^0, \dots, x^n)$

2. GENERAL OBSERVATIONS

The effects of changing the proposal input noise distribution are more complex in Hamiltonian Monte Carlo than in the random walk Metropolis and the Metropolis-adjusted Langevin algorithm, two methods for which this question has been studied previously (Jarner & Roberts, 2007; Stramer & Tweedie, 1999). For both of these cases the input noise is combined linearly with some deterministic function of the current point of the Markov chain. As a result, larger values of this noise will typically result in larger proposed moves, which may be advantageous if $\pi(\cdot)$ has heavy tails or possesses multiple modes. By contrast, in Hamiltonian Monte Carlo the choice of kinetic energy alters both the input noise distribution and the Hamiltonian dynamics, and so the impact of different choices is not so transparent.

If we consider how a single proposal in Hamiltonian Monte Carlo perturbs the current position x^{i-1} at iteration i , then setting $x_0 = x^{i-1}$, $p_0 \sim \nu(\cdot)$ and solving (1) for T units of time gives

$$x^* = x_0 + \int_0^T \nabla K \left\{ p_0 - \int_0^s \nabla U(x_u) du \right\} ds. \quad (2)$$

An explicit numerical method typically involves the approximation $\int_0^h \nabla U(x_u) du \approx \nabla U(x_0)h$ for some suitably chosen h . In the case of the leapfrog integrator the approximate solution over a single leapfrog step becomes

$$x_\varepsilon = x_0 + \varepsilon \nabla K \left\{ p_0 - \frac{\varepsilon}{2} \nabla U(x_0) \right\}, \quad (3)$$

with the corresponding momentum update

$$p_\varepsilon = p_0 - \frac{\varepsilon}{2} \{ \nabla U(x_0) + \nabla U(x_\varepsilon) \}. \quad (4)$$

Recalling that $K(p)$ is an even function, meaning $\nabla K(p)$ is odd, then it will turn out that in many cases the function that intuitively governs the speed of the x -dynamics when $\|\nabla U(x)\|$ is large is

$$\nabla K \circ \nabla U(x), \quad (5)$$

which we refer to from this point as the composite gradient. Although in practice the choice of ε will influence sampler behaviour, we will demonstrate that a qualitative understanding can be developed by considering only (5). Similarly, when $\|\nabla U(x)\|$ is small, x_ε will resemble a random walk proposal with perturbation ε multiplied by $\nabla K(p_0)$, which we refer to as the implicit noise. We argue that an appropriate choice of kinetic energy is one for which these two quantities are both suitably optimized over, and in the next sections we examine both when this can be done and what a suitable choice for each should be.

3. THE COMPOSITE GRADIENT

3.1. Geometric convergence and long term stability

We begin with some general results relating the choice of kinetic energy to the composite gradient, and potential consequences for the convergence rate of the method.

DEFINITION 1. We call a distribution which has a density $f(x) \propto \exp\{-g(x)\}$ heavy-tailed if $\lim_{\|x\| \rightarrow \infty} \|\nabla g(x)\| = 0$, and light-tailed if $\lim_{\|x\| \rightarrow \infty} \|\nabla g(x)\| = \infty$.

Distributions which do not fall into either category are those for which the tails are heavy in some directions and light in others, for which the sampler will behave very differently in distinct regions of the space. A detailed study of these is beyond the scope of this work, though we conjecture that the least favourable region would dictate convergence. We also do not include distributions with exponentially decaying densities, which are often treated as a special case when analysing Metropolis–Hastings methods (Jarner & Hansen, 2000).

Recall that a π -invariant Markov chain with transition kernel P is called geometrically ergodic if for some positive constants $C < \infty$, $\rho < 1$ and a Lyapunov function $V : \mathbb{R}^d \rightarrow [1, \infty)$ the bound

$$\|P^n(x, \cdot) - \pi(\cdot)\|_{TV} \leq CV(x)\rho^n \quad (6)$$

can be constructed, where $\|\mu(\cdot) - \nu(\cdot)\|_{TV}$ denotes the total variation distance between two measures $\mu(\cdot)$ and $\nu(\cdot)$. See for example Roberts & Rosenthal (2004) for more details.

PROPOSITION 1. In either of the following cases Algorithm 1 will not produce a geometrically ergodic Markov chain:

- (i) $\pi(\cdot)$ is heavy-tailed
- (ii) $\pi(\cdot)$ is light-tailed and the following conditions hold:

1. The composite gradient satisfies

$$\lim_{\|x\| \rightarrow \infty} \frac{\|\nabla K \circ \nabla U(x)\|}{\|x\|} = \infty, \quad (7)$$

2. There is a strictly increasing unbounded function $\phi : [0, \infty) \rightarrow [1, \infty)$ and $m_1, m_2 < \infty$ with $m_1 \geq 1$ such that for every $M \geq m_1$ and $\|x\| \geq m_2$

$$\|y\| \geq M\|x\| \implies \|\nabla U(y)\| \geq \phi(M)\|\nabla U(x)\|. \quad (8)$$

3. Setting $\Delta H(x_0, p_0) = H(x_{L\varepsilon}, p_{L\varepsilon}) - H(x_0, p_0)$ and $\Delta(x_0, p_0) = \|x_{L\varepsilon}\| - \|x_0\| + \|p_{L\varepsilon}\| - \|p_0\|$,

$$\Delta(x_0, p_0) \rightarrow \infty \implies \Delta H(x_0, p_0) \rightarrow \infty. \quad (9)$$

4. For all $\delta < \infty$, there is a $C_\delta > 0$ such that for all $x \in \mathbb{R}^d$

$$\frac{\|\nabla K\{\delta \nabla U(x)\}\|}{\|\nabla K \circ \nabla U(x)\|} \geq C_\delta \quad (10)$$

5. Either $K(p)$ satisfies

$$\|p\| \geq \|y\| \implies \|\nabla K(p)\| \geq \|\nabla K(y)\|, \quad (11)$$

or there is a $k : \mathbb{R} \rightarrow \mathbb{R}$ such that $K(p) := \sum_{i=1}^d k(p_i)$ and

$$|p_i| \geq |y_i| \implies |k'(p_i)| \geq |k'(y_i)|. \quad (12)$$

for all i .

Remark 1. Condition 4 in part (ii) will be satisfied if, for example, ∇K is a homogeneous function. It is satisfied for all choices considered in this work. 140

For the standard quadratic choice of kinetic energy the Hamiltonian Monte Carlo method will not produce a geometrically ergodic Markov chain when $\pi(\cdot)$ has heavy-tails (Livingstone et al., 2016). Part (i) of the above states that no choice of kinetic energy which is independent of x can rectify this. We explain why in subsection 5.3. Part (ii) states that provided the density is sufficiently regular that (9) holds, and that the degree of tail oscillation in $\|\nabla U(x)\|$ is restricted through (8), then the composite gradient should not grow faster than linearly in $\|x\|$ for geometric ergodicity. Within the class of kinetic energies for which this requirement is met, we motivate an optimal choice in subsection 3.3. 145

The composite gradient was deduced based on studying the leapfrog integrator over a single step. Next, however, we argue that controlling this quantity in fact induces stability over multiple leapfrog steps. 150

PROPOSITION 2. *If for some $q > 0$ and $A, B, C, D < \infty$*

$$\|\nabla U(x)\| \leq A\|x\|^q + B, \quad \|\nabla K(p)\| \leq C\|p\|^{\frac{1}{q}} + D, \quad (13)$$

then $\limsup_{\|x\| \rightarrow \infty} \|\nabla K \circ \nabla U(x)\|/\|x\| < \infty$, and furthermore if $\|p_0\| \leq E_0\|x_0\|^q + F_0$ for some $E_0, F_0 < \infty$, then writing $x_{L\varepsilon}$ and $p_{L\varepsilon}$ as functions of x_0 and p_0 , it holds that

$$\|x_{L\varepsilon}(x_0, p_0)\| \leq E_L\|x_0\| + F_L, \quad \|p_{L\varepsilon}(x_0, p_0)\| \leq \mathbf{G}_L\|x_0\|^q + \mathbf{H}_L.$$

for some $E_L, F_L, \mathbf{G}_L, \mathbf{H}_L < \infty$ which do not depend on x_0 . 155

As a contrast, if the assumption on $\nabla K(p)$ was replaced with $\|\nabla K(p)\| \leq C\|p\| + D$, then the bound would read $\|x_{L\varepsilon}(x_0, p_0)\| = E_L\|x_0\|^{1+qL} + F_L$, which grows exponentially in L when $q > 1$.

3.2. Kinetic energies that induce negligible moves

Next we consider a heavy-tailed choice of $\nu(\cdot)$. In this instance the composite gradient will in fact decay as $\|\nabla U(x)\|$ grows, meaning that if $\pi(\cdot)$ is light-tailed the resulting Markov chain will move very slowly in some regions of the space. We formalize this intuition below. 160

DEFINITION 2. *We say that a Markov chain on \mathbb{R}^d with transition kernel P possesses the negligible moves property if for every $\delta > 0$*

$$\lim_{\|x\| \rightarrow \infty} P\{x, B_\delta(x)\} = 1. \quad (14)$$

PROPOSITION 3. *If $\pi(\cdot)$ is light-tailed and $\nu(\cdot)$ is heavy-tailed then a Markov chain produced by the Hamiltonian Monte Carlo method possesses the negligible moves property.* 165

In general this property seems undesirable, the implication being that if the current point x has large norm then the chain is unlikely to move very far in any appreciable number of steps. It does not, however, preclude geometric ergodicity, as often such steps can be close to deterministic in a desirable direction. We give a simple example to illustrate the phenomenon. 170

Example 1. Consider the Markov chain with transition

$$x^i = \max(x^{i-1} - 1 + \xi^i, 1), \quad \xi^i \sim N(0, 1),$$

which is geometrically ergodic as shown in section 16.1.3 of Meyn & Tweedie (1993). Then the transformed Markov chain $(y^i)_{i \geq 1}$, where each $y^i = \log(x^i)$, is also geometrically ergodic, but possesses the negligible moves property.

175 Although a geometrically decaying convergence bound can still be established in such cases, the slow movement of the chain when $\|x\|$ is large is accounted for in the Lyapunov function $V(x)$ present in the bound (6).

PROPOSITION 4. *If (14) holds for a Hamiltonian Monte Carlo method then any Lyapunov function $V : \mathbf{X} \rightarrow [1, \infty)$ cannot satisfy either of the following:*

180 (i) *$\log V$ is uniformly lower semi-continuous, meaning for all $\epsilon' > 0$, there is a $\delta' < \infty$ such that if $\|y - x\| \leq \delta'$, uniformly in x , then $\log V(y) \geq \log V(x) - \epsilon'$.*

(ii) *There is an $s \in (0, \infty)$ such that $\liminf_{\|x\| \rightarrow \infty} V(x)e^{-s\|x\|} = c$ for some $c \in (0, \infty)$.*

The above implies that Lyapunov functions will typically exhibit faster than exponential growth in $\|x\|$, meaning the penalty in the bound (6) for starting the chain in the tails of the distribution will be very large. This reflects the fact that the chain will take a long time to move an appreciable distance and hence approach equilibrium. So although the negligible moves property does not preclude geometric convergence, we regard it as an undesirable feature.

3.3. Illustrative example

Part (ii) of Proposition 1 relates to the numerical solution of (1) using the leapfrog method. If 190 the composite gradient grows faster than linearly in $\|x\|$ then the dynamical system exhibits stiffness, meaning gradients change very quickly and hence numerical approximations in which they are assumed to be constant over small time periods will become unstable. Two possible remedies are to use a more complex implicit numerical method or adaptive step-size control, as discussed in Okudo & Suzuki (2016). With the standard choice of kinetic energy $\nabla K \circ \nabla U(x) = \nabla U(x)$, 195 meaning these instabilities occur when $\|\nabla U(x)\|$ grows faster than linearly in $\|x\|$. Since the composite gradient dictates this phenomenon in general, fast growth in $\nabla U(x)$ can be counteracted by choosing a kinetic energy which is subquadratic in $\|p\|$. Physically, choosing a slowly growing kinetic energy slows down the Hamiltonian dynamics, meaning the system (1) is no longer stiff in the x -coordinate and hence simple numerical schemes like the leapfrog method 200 become stable. We can further use the same idea when $\nabla U(x)$ exhibits relatively slow growth by choosing a fast-growing kinetic energy which will speed up the flow allowing the chain to explore the state space more efficiently.

To demonstrate this phenomenon more concretely we consider the family of separable Hamiltonians $H : \mathbb{R}^2 \rightarrow [0, \infty)$ given by

$$H(x, p) = \alpha^{-1}|x|^\alpha + \beta^{-1}|p|^\beta, \quad (15)$$

205 for suitable choices $\alpha, \beta > 1$, which is necessary to ensure continuity of the derivative. Hamilton's equations are

$$\begin{aligned} \dot{x} &= \operatorname{sgn}(p)|p|^{\beta-1} \\ \dot{p} &= -\operatorname{sgn}(x)|x|^{\alpha-1} \end{aligned} \quad (16)$$

The flow is periodic since the contours of H are closed curves in \mathbb{R}^2 . These kinetic energies correspond to choosing $\nu(\cdot)$ from the exponential power family of Box & Tiao (1973), for which 210 direct sampling schemes are provided in Mineo (2014). Zhang et al. (2016) recommend choosing this class of momentum distributions in Hamiltonian Monte Carlo. The following proposition shows how the the period length depends on the initial value of the Hamiltonian.

PROPOSITION 5. For the class of Hamiltonians given by (15) and dynamics given by the exact solutions to (16), the period length

215

$$\mathcal{P}(E) = \inf_{t>0} \{t : (x_t, p_t) = (x_0, p_0)\}$$

is related to the initial value of the Hamiltonian $E = H(x_0, p_0)$ through the expression

$$\mathcal{P}(E) = c_{\alpha,\beta} E^\eta, \quad \eta = \frac{1 - (\beta - 1)(\alpha - 1)}{\alpha\beta},$$

where $c_{\alpha,\beta} < \infty$ does not depend on E . In particular, if $\alpha = 1 + \gamma$ and $\beta = 1 + \gamma^{-1}$ for some $\gamma \in (0, \infty)$, then $\nabla K \circ \nabla U(x) = x$ and $\eta = 0$, meaning the period length does not depend on E .

The above result characterizes precisely how the speed of the Hamiltonian flow depends on the growth rates of the potential and kinetic energies. If $\eta < 0$ the flow will be very fast at higher energies, and so explicit numerical methods will typically become unstable when $|x|$ is large. If $\eta > 0$ then the period length will increase with $|x|$, so numerical methods will be more stable, but the sampler may take some time to move from the tails back into the centre of the space.

220

The ideal choice in terms of finding the right balance between efficiency and numerical stability would therefore appear to be setting $\eta = 0$, which for this class of models can be achieved by choosing the kinetic energy as the L_p dual of the potential. This corresponds to setting $\alpha = 1 + 1/(\beta - 1)$, which produces an algorithm for which the flow behaves in an auto-regressive manner when $|x|$ is large, since the composite gradient will be similar in magnitude to the current point. In this case tuning the method should also be more straightforward, as a choice of integrator step-size and integration time made in the transient phase should still be reasonable at equilibrium. When $\eta < 0$, for any fixed choice of step-size $\varepsilon > 0$ there is an $M_\varepsilon < \infty$ such that when $|x| > M_\varepsilon$ the algorithm will become numerically unstable, and when $\eta > 0$ then it would be desirable in the transient phase to choose a comparatively large value for ε and take more leapfrog steps, which may result in many proposals being rejected and unnecessary computational expense when the chain reaches equilibrium.

225

230

235

4. THE IMPLICIT NOISE

There are likely to be regions in the centre of the space in which $\|\nabla U(x_0)\|$ is small, meaning the increment $\varepsilon \nabla K \{p_0 - \varepsilon \nabla U(x_0)/2\} \approx \varepsilon \nabla K(p_0)$. This can occur when the density $\pi(x)$ contains flat regions, and we show one such case in section 6. In these regions Hamiltonian Monte Carlo resembles a random walk Metropolis with proposal $x_{L\varepsilon} \approx x_0 + L\varepsilon \nabla K(p_0)$. It would seem sensible to choose $K(p)$ such that the distribution of $\nabla K(p)$ would be an appropriate choice in a random walk Metropolis here.

240

For the standard quadratic choice $\nabla K(p) = p$, which is hence also Gaussian. Different choices of kinetic energy, which may seem sensible by analysing the composite gradient, can however result in very different distributions for the implicit noise, which can either become bi-modal or exhibit very light tails. We plot histograms of $\nabla K(p)$ for some natural choices of momentum distribution in figure 1.

245

The optimum choice of implicit noise will typically be problem dependent, but it would seem logical to follow the standard choices used in the random walk Metropolis. Bi-modal distributions have recently been suggested for this method (Yang & Rodríguez, 2013), but we found that something which resembles a Gaussian distribution for small values of p performed favourably in experiments.

250

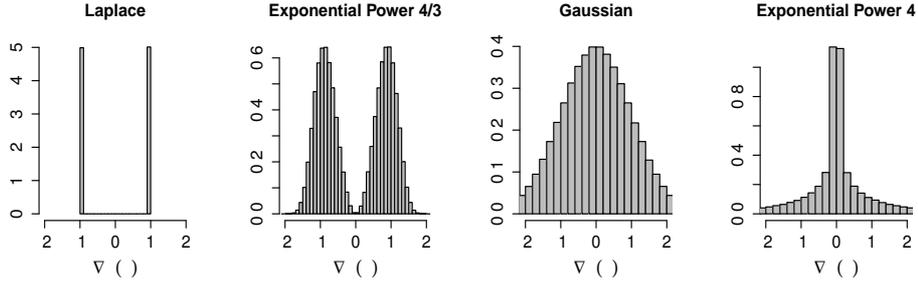


Fig. 1. Histograms showing how different choices of the momentum distribution $\nu(\cdot)$ can affect typical values of $\nabla K(p)$.

5. GENERAL GUIDELINES AND THE EFFICIENCY-ROBUSTNESS TRADE OFF

255 The above results allow us to offer some general strategies and observations for choosing a kinetic energy, dependent on the objectives of the user.

If we deem the goal to be to maximize the speed of the flow when $\|x\|$ is large whilst retaining numerical stability, then an appropriate choice of kinetic energy for a given potential U is one for which when $\|\nabla U(x)\|$ is large $\nabla K\{p_0 - \varepsilon \nabla U(x_0)/2\} \approx -c_\varepsilon x$ for some $c_\varepsilon > 0$. This can
 260 be achieved when $\nabla U(x)$ is invertible through the choice

$$\nabla K(p) \approx \nabla U^{-1}(p), \quad (17)$$

which implies that $\nabla K \circ \nabla U(x) \approx \nabla U^{-1} \circ \nabla U(x) = x$. Of course often $\nabla U(x)$ will not possess an inverse, in which case the considering the leading order term in $\|\nabla U(x)\|$ is of greatest importance. Below we give some examples of how such a choice for $K(p)$ can be constructed.

Example 2. If $x \sim N(0, \Sigma)$ then $\nabla U(x) = \Sigma^{-1}x$, meaning $-(\nabla U)^{-1}(-p) = \Sigma p$, implying
 265 the choice $p \sim N(0, \Sigma^{-1})$ and corresponding kinetic energy $K(p) = p^T \Sigma p/2$.

Remark 2. This result is intuitive, as if $K(p) = p^T M^{-1}p/2$, then the proposal after a single leapfrog step will be $x_\varepsilon = x_0 - \varepsilon^2 M^{-1} \nabla U(x)/2 + \varepsilon M^{-1/2} z$ with $z \sim N(0, I)$. This is equivalent to a Metropolis-adjusted Langevin algorithm proposal with pre-conditioning matrix M^{-1} , and in that setting it is known that this should be made equal to Σ for optimal sampling (Roberts
 270 & Rosenthal, 2001; Dalalyan, 2017).

Example 3. Consider a distribution with density $\pi(x) \propto e^{-U(x)}$ where

$$U(x) = C(\alpha x^4 - \beta x^2),$$

for some $C, \alpha, \beta > 0$. Such double well potential models are important in many areas of physics, including quantum mechanical tunnelling (Liang & Müller-Kirsten, 1992). In this case $\nabla U(x) = 4C\alpha x^3 - 2C\beta x$, which is typically not invertible, but for large $|x|$ we see that $\nabla U(x)/x^3 \approx$
 275 $4C\alpha$, which motivates the choice $\nabla K(p) = p^{1/3}/(4C\alpha)$ for large $|p|$.

Example 4. Park & Casella (2008) suggest a Bayesian version of the bridge estimator for regression of Frank & Friedman (1993). In this case

$$U(\beta) = L(\beta) + \lambda \sum_{j=1}^d |\beta_j|^q,$$

for some $\lambda > 0$, where $L(\beta)$ is a suitable loss function. If $q > 1$ and the asymptotic growth of $L(\beta)$ is at most linear in $\|\beta\|$, then the choice $K(p) = \sum_i |p_i|^\alpha$ with $\alpha = 1 + 1/(q - 1)$ for large $\|p\|$ would maximize the speed of convergence during the transient phase of the algorithm. Logistic regression, the Huberized loss of Park & Casella (2008) and the asymmetric pinball loss (Yu & Moyeed, 2001) are possible choices for $L(\beta)$ which fit this criterion.

280

Example 5. Neal (2003) proposes the hierarchical model $x_i|\nu \sim N(0, e^\nu)$ for $i = 1, \dots, 9$, with $\nu \sim N(0, 3^2)$, resulting in a joint distribution that resembles a ten-dimensional funnel. Here $U(x_1, \dots, x_9, \nu) = \nu^2/18 + 9\nu/2 + e^{-\nu} \sum_i x_i^2/2$, meaning

285

$$\nabla U(x_1, \dots, x_9, \nu) = \left(2x_1 e^{-\nu}, \dots, 2x_9 e^{-\nu}, \frac{\nu}{2} + \frac{9}{2} - \frac{e^{-\nu}}{2} \sum_i x_i^2 \right)^T.$$

When $\nu \ll 0$ these gradients can be exponentially large, so a choice of $\nabla K(p)$ that is at most logarithmic in p would be needed, which could arise from the choice of kinetic energy $K(p_i) \approx \{1 + \log(1 + |p_i|)\}|p_i|$.

A general appreciation for the consequences of different kinetic energy choices can be gained by considering a formal Taylor series expansion

290

$$\nabla K \left\{ p_0 - \frac{\varepsilon}{2} \nabla U(x_0) \right\} = \nabla K \left\{ -\frac{\varepsilon}{2} \nabla U(x_0) \right\} + \nabla^2 K \left\{ -\frac{\varepsilon}{2} \nabla U(x_0) \right\} p_0 + \dots$$

In the Gaussian case $\nabla^2 K(p) = I$, meaning the refreshed momentum p_0 interacts linearly with the composite gradient term regardless of the current position in the state space. If however, the momentum distribution is chosen to have heavier than Gaussian tails, then each element of $\nabla^2 K(p)$ will become negligibly small when $\|\nabla U(x_0)\|$ is large, reducing the influence of p_0 . A sampler with a flow of this type will move towards a high-density region of the space in an almost deterministic fashion when $\|x\|$ is large. We formalize this intuition below.

295

DEFINITION 3. A Hamiltonian Monte Carlo sampler is called deterministic in the tails if as $\|x\| \rightarrow \infty$

$$\left\| \nabla^2 K \left\{ \frac{\varepsilon}{2} \nabla U(x) - p \right\} - \nabla^2 K \left\{ \frac{\varepsilon}{2} \nabla U(x) \right\} \right\| \rightarrow 0 \tag{18}$$

a.s., where $p \sim \nu(\cdot)$.

Given the above intuition, the following result is immediate.

300

PROPOSITION 6. A Hamiltonian Monte Carlo sampler will be deterministic in the tails if $\pi(\cdot)$ is light-tailed and $K(p)$ is twice differentiable with

$$\lim_{\|p\| \rightarrow \infty} \|\nabla^2 K(p)\| = 0. \tag{19}$$

Remark 3. Typically if a choice of momentum distribution with heavier than Gaussian tails is made then (19) will hold.

As alluded to in subsection 3.3, there is a trade off to be made between efficiency and step-size robustness. Choosing a lighter-tailed momentum distribution will result in a faster flow, and potentially a faster mixing Markov chain as a result, but the risk of numerical instabilities is higher, and so the algorithm will be very sensitive to the choice of ε . Choosing a heavier-tailed momentum distribution will result in slower flows, meaning algorithm performance will be less sensitive to step-size choice, but the best-case mixing time under the optimal step-size

305

310

may be slower. We give a numerical example illustrating this in the supplementary material. If the tails are too heavy then the algorithm will also exhibit the negligible moves property of section 3.2. A kinetic energy choice with Laplacian tails, for example, would result in a sampler which may be comparatively slow in the tails but is guaranteed to be numerically stable, since in this case the composite gradient is $\text{sign}\{p_0 - \varepsilon \nabla U(x_0)/2\}$, which is bounded from above and below regardless of how fast $\|\nabla U(x_0)\|$ grows.

Remark 4. It has recently been shown that by choosing $\nu(\cdot)$ to be a Laplace distribution one can perform Hamiltonian Monte Carlo on problems for which the standard method is not applicable, such as problems with discontinuous likelihoods or involving discrete parameters. We do not comment further on this here, but see Nishimura et al. (2017) for details.

5.1. Relativistic power kinetic energies

From the above discussion it would seem sensible to design a kinetic energy function that looks similar to the Gaussian choice when its argument is small, but also allows a robust and efficient choice of composite gradient. The relativistic kinetic energy

$$K(p) = \sum_i m_i c_i^2 \left(1 + \frac{p_i^2}{m_i^2 c_i^2} \right)^{\frac{1}{2}}$$

suggested in Lu et al. (2017), has many such desirable features, since it behaves as a quadratic function for small p and is linear when p is large. We slightly generalize this choice to allow for different tail behaviours, resulting in the family of relativistic power kinetic energies

$$K(p) = \sum_i \beta^{-1} (1 + \gamma_i^{-1} p_i^2)^{\beta/2}, \quad (20)$$

where $\beta \geq 1$ and each $\gamma_i > 0$. When each $p_i^2 \ll \gamma_i$ then

$$K(p) \approx \sum_i \beta^{-1} (1 + \gamma_i^{-1} p_i^2) \implies \partial_i K(p) \approx 2(\beta \gamma_i)^{-1} p_i$$

and when each $p_i^2 \gg \gamma_i$

$$K(p) \approx \sum_i \beta^{-1} \gamma_i^{-\beta/2} |p_i|^\beta \implies \partial_i K(p) \approx \gamma_i^{-\beta/2} \text{sgn}(p) |p_i|^{\beta-1}.$$

Drawing samples from this family is a requirement for implementation. In our experiments we used adaptive rejection sampling (Gilks & Wild, 1992), which we found to be reasonably efficient. However, since the distribution is fairly regular in shape it is likely that specialized methods can be designed to be even more efficient. In addition, other strategies such a dependent sampling using Metropolis–Hastings could be employed, as in Akhmatskaya & Reich (2008).

5.2. Approximate gradients and doing away with Metropolis–Hastings

One example in which the robustness-efficiency trade off outlined above can be informative is when estimates are used in place of $\nabla U(x)$, which may be either intractable or very expensive to compute. Examples are given in Chen et al. (2014); Strathmann et al. (2015); Lindsten & Doucet (2016).

If we assume that at each iteration the approximate gradient $\tilde{\nabla} U(x)$ satisfies

$$\tilde{\nabla} U(x) = \nabla U(x) + \eta_x,$$

where η_x follows a distribution that may depend on x , but such that $E(\eta_x) = 0$ for all x , then a similar Taylor expansion gives

$$\nabla K \left\{ p_0 - \frac{\varepsilon}{2} \tilde{\nabla} U(x_0) \right\} = \nabla K \left\{ -\frac{\varepsilon}{2} \nabla U(x_0) \right\} + \nabla^2 K \left\{ -\frac{\varepsilon}{2} \nabla U(x_0) \right\} (p_0 + \eta_x) + \dots$$

Therefore, if the momentum distribution is chosen to have heavy enough tails that for large $\|x\|$ the last term becomes negligibly small, then the effects of such an approximation are mitigated and the resulting approximate composite gradient will closely resemble $\nabla K \left\{ -\frac{\varepsilon}{2} \nabla U(x_0) \right\}$. 345

In some of the above mentioned approximate implementations, the Metropolis step is also omitted to reduce computational costs. When the gradient term $\|\nabla U(x)\|$ grows at a faster than linear rate, then often the resulting Markov chains become transient when this is done, as shown for the unadjusted Langevin algorithm in Roberts & Tweedie (1996). In such cases it is difficult to give long time guarantees on the level of bias induced from a finite simulation. Ensuring that the composite gradient grows no faster than linearly should mean that such transience is averted, meaning a stronger grasp of the degree of approximation can be established. We leave a detailed exploration of this for future work. 350

5.3. Heavy-tailed models

Based on the composite gradient intuition, it may seem desirable when $\pi(\cdot)$ is heavy-tailed to choose a kinetic energy for which $\|\nabla K(p)\| \rightarrow \infty$ as $\|p\| \rightarrow 0$. In this case when $\|x\|$ is large the potential gradient will be very small, so choosing the kinetic energy in this way one could still make $\nabla K \circ \nabla U(x)$ linear in $\|x\|$. However, the actual leapfrog perturbation is 355

$$\varepsilon \nabla K \left\{ p_0 - \frac{\varepsilon}{2} \nabla U(x_0) \right\}.$$

When $\pi(\cdot)$ is light-tailed the p_0 term becomes safe to ignore when $\|x\|$ is large, but if it is heavy-tailed then this is no longer the case as for large $\|x\|$ typical proposals will be of the form $x^* \approx x + L\varepsilon \nabla K(p)$, where $p \sim \nu(\cdot)$, which resembles a random walk. It is because of this that, as shown in Part (i) of Proposition 1, no choice of kinetic energy which is independent of x can produce a geometrically ergodic chain. 360

It will likely be the case that choosing a heavier-tailed momentum distribution will be of benefit in this scenario in terms of rates of convergence, since this is true of the random walk Metropolis, as shown in Jarner & Roberts (2007). However, since it is the implicit noise $\nabla K(p)$ which drives the sampler, then care must also be taken with its form. In the case where $\nu(\cdot)$ is a Cauchy distribution, for example, then $\nabla K(p) = 2p/(1 + p^2)$, which is lighter-tailed than p itself, so here the benefits of choosing heavier tails are not as strong as those for the random walk algorithm. 365
370

6. EXAMPLES

6.1. Quantile regression

We consider a Bayesian quantile regression as introduced in Yu & Moyeed (2001). The goal is to estimate the τ th quantile of a response $y \in \mathbb{R}$ conditioned on a collection of covariates $x \in \mathbb{R}^m$, written $\mu(x) = F^{-1}(\tau|x)$. Given n data points $\{(x_i, y_i)\}_{i=1}^n$, we take the natural choice $\mu(x_i, \beta) = \sum_{j=1}^m x_{ij} \beta_j$, and follow the approach of Fasiolo et al. (2017) by estimating a posterior distribution for β as $\pi(\beta|y, x) \propto \exp \left\{ -\sum_{i=1}^n L(\beta, x_i, y_i) \right\} \pi_0(\beta)$, using the general Bayesian updating framework of Bissiri et al. (2016), with $L(\beta, x_i, y_i)$ a smoothed version of 375

the pinball loss introduced in Fasiolo et al. (2017), given by

$$L(\beta, x, y) = (\tau - 1) \left\{ \frac{y - \mu(x, \beta)}{\sigma} \right\} + \xi \log \left[1 + \exp \left\{ \frac{y - \mu(x, \beta)}{\xi \sigma} \right\} \right] + g(\xi, \sigma, \tau).$$

380 Here $g(\xi, \sigma, \tau) = \log \{ \xi \sigma \text{Beta}[\xi(1 - \tau), \xi \tau] \}$. As $\xi \rightarrow 0$ then the non-smooth pinball loss is recovered. This is linear in β , meaning that if L_q priors are chosen as in bridge regression (Park & Casella, 2008) then the corresponding potential is

$$U(\beta) = \sum_{i=1}^n L(\beta, x_i, y_i) + \lambda \sum_{j=1}^d |\beta_j|^q,$$

for $1 < q \leq 2$. The dominant term in $U(\beta)$ is the contribution from the prior, meaning that each $|\partial_j U(\beta)| = O(|\beta_j|^{q-1})$ for $j = 1, \dots, m$.

385 We performed two studies on 20 simulated data points with $m = 2$, fixing $\sigma = 1$, $\lambda = 1$ and $\xi = 0.01$. In the first we set $q = 2$ for each β_j , and in the second $q = 1.5$. Four different momentum distributions were tested: an exponential power family with shape parameter 3, Gaussian, Laplacian, and a t distribution with 4 degrees of freedom. Although the set up is simple it still enables a demonstration of both the negligible moves property and how the composite gradient
390 dictates sampler performance. The number of leapfrog steps was set to 1 and the step size was tuned based on achieving a 65-75% acceptance rate at equilibrium. The samplers were then initialised far from the region of high probability and convergence speed was assessed, and also initialised at equilibrium to assess mixing in favourable areas. Results are shown in Figure 2. In the Gaussian prior study the exponential power choice is not shown since the resulting sampler
395 did not move, as explained by Proposition 1. As expected, the Gaussian choice, which results in a linearly growing composite gradient, reaches equilibrium quicker than the others. The Laplacian choice converges in a straight-line fashion, as the size of the proposed jump is always $\varepsilon\sqrt{2}$ regardless of the current position. The Student's t choice exhibits the negligible moves property as outlined in Proposition 3, with convergence visibly slow during the first 5000 iterations. When
400 the heavier-tailed prior is chosen, the exponential power choice is now numerically stable and produces faster convergence than the Gaussian. Since the potential growth is $O(|\beta|^{1.5})$ then this option results in a linear composite gradient, so the performance is to be expected. The speed of the remaining choices is dictated by how the composite gradient grows, with the Student's t distribution again performing the worst. In this study the difference is not so pronounced, as in
405 probabilistic terms the sampler is initialised in a region that is not so far into the tails. The slower convergence overall compared to the Gaussian study reflects the fact that $\pi(\cdot)$ is heavier-tailed. In both studies all samplers mixed similarly well regardless of the kinetic energy when initialised at equilibrium, with slightly worse performance observed for the Student's t distribution. As discussed in Section 4, the optimal choice of implicit noise is likely to be problem specific, with the
410 Gaussian shape sensible when no other information is available.

6.2. The Ginzburg–Landau model

As a second example we take the model of phase transitions in condensed-matter physics proposed by Ginzburg and Landau (Goldenfeld, 1992, Chapter 5). We consider a three-dimensional 10^3 lattice, where each site represents a random variable $\psi_{ijk} \in \mathbb{R}$. The corresponding potential
415 for the 1000-dimensional parameter ψ is

$$U(\psi) = \sum_{i,j,k} \left\{ \frac{(1 - \tau)}{2} \psi_{ijk}^2 + \frac{\tau \alpha}{2} \|\tilde{\nabla} \psi_{ijk}\|^2 + \frac{\tau \lambda}{4} \psi_{ijk}^4 \right\},$$

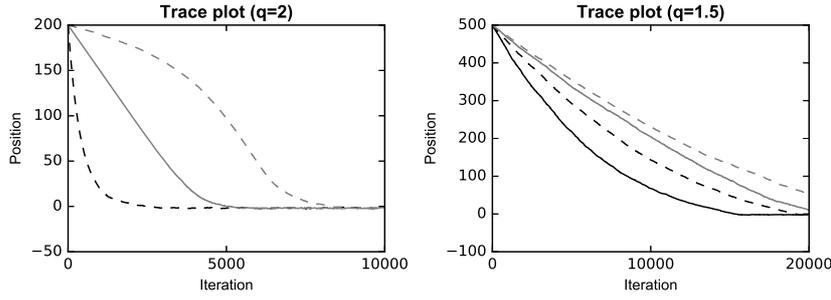


Fig. 2. Trace plots from the quantile regression studies. The solid black line is the exponential power momentum choice, dashed black lines are the Gaussian choice, solid grey lines are the Laplacian choice and dashed grey lines are the t distribution with 4 degrees of freedom.

where $\alpha, \lambda, \tau > 0$ and $\tilde{\nabla}\psi_{ijk} = (\psi_{i_+j_+k} - \psi_{ijk}, \psi_{ij_+k} - \psi_{ijk}, \psi_{ijk_+} - \psi_{ijk})$, where $i_+ = (i + 1) \bmod (10)$, and j_+, k_+ are defined analogously. The modular structure imposes periodic boundary conditions on the lattice. When the parameter $\tau < 1$, the conditional distribution for each ψ_{ijk} looks Gaussian in the centre of the space but with lighter tails. The model exhibits a phase transition at $\tau = 1$, so that when $\tau > 1$ each conditional distribution is bi-modal. The parameter α controls the strength of correlation between neighbouring lattice sites, with larger values making the sampling problem more challenging. When $\tau < 1$, each ψ_{ijk} is likely to be close to zero, while in the bi-modal phase they are more likely to be non-negligible in magnitude. When α is large, neighbouring parameters are likely to have the same sign. The bi-modal phase therefore represents the system in its ordered state, whereas the system is disordered for $\tau < 1$.

The inference problem is to estimate expectations with respect to the distribution with density proportional to $\exp\{-U(\psi)\}$. We generate samples using Hamiltonian Monte Carlo with four different choices of kinetic energy: the standard quadratic choice, the relativistic choice of Lu et al. (2017), the exponential power choice of Zhang et al. (2016) and the relativistic power choice introduced in section 5.1. In the latter two cases we set the shape parameter $\beta = 4/3$, which results in a composite gradient that is linear when $\|\psi\|$ is large.

We perform two studies. In the first we initialize samplers in the centre of the space and compute the effective sample sizes from 10,000 iterations. In the second we initialize each $\psi_{ijk} \sim U[-10, 10]$, and compute the number of iterations until $\max|\psi_{ijk}| \leq 2$. Results are shown in Table 1. The kinetic energies with Gaussian implicit noise produce larger effective sample sizes in this example. The reason is that there is a sharp drop in density when the quartic term begins to dominate the potential. As the exponential power choice results in a bi-modal implicit noise, it is more likely to propose larger jumps which fall into this low density region, and are hence rejected. To prevent this from happening too often a smaller step-size must be chosen than for the other methods. Among the other choices the Gaussian has the highest effective sample size, though this cannot necessarily be relied upon as the sampler is not geometrically ergodic. The slower speed for larger values of $\|\psi\|$ enforced in the relativistic and to a lesser degree relativistic power cases appears to slightly reduce efficiency here, which is sacrificed in favour of robustness. We simply set each tuning parameter $\gamma_i = 1$ in our experiments. This could be tuned in order to improve performance, but we preferred to use default values to limit the extra tuning required for this choice. Sensitivity checks with $\gamma_i = 2$ and 3 did not yield appreciably different results. In terms of iterations to the centre, the Gaussian choice performs poorly as expected. The remaining samplers perform similarly. The relativistic power choice outperforms

Table 1. *Hamiltonian Monte Carlo on the Ginzburg–Landau model*

	Effective sample size			Iterations to centre
	Minimum	Mean	Maximum	
Gaussian	6,251	8,748	10,000	N/A
Relativistic Power $\beta = 4/3$	5,253	6,777	8,271	4.2
Relativistic	3,591	4,639	5,525	8.6
Exponential Power $\beta = 4/3$	810	1,108	1,303	11.9

Effective sample sizes at equilibrium and iterations until $\max_i |\psi_i| < 2$ for Hamiltonian Monte Carlo on the Ginzburg–Landau model with $\alpha = 0.1$, $\lambda = 0.5$ and $\tau = 2$. Step-sizes were chosen based on optimizing effective sample sizes. In each simulation the number of leapfrog steps was set to 10, and results are averaged over 10 runs.

the exponential power choice because a smaller step-size was needed in the latter case for adequate mixing at equilibrium. The relativistic choice, though slower, is not critically so, because 10 is not very large in absolute terms, meaning the differences between a drift which is $O(1)$ and $O(\|\psi\|)$, while visible, is not substantial here.

7. DISCUSSION

We have described how changing the kinetic energy in Hamiltonian Monte Carlo affects performance. In practice, several other strategies could be considered, such as mixtures of kinetic energies, or employing a delayed rejection approach as in Mira et al. (2001). In addition, other stochastic processes that utilise Hamiltonian dynamics, such as underdamped Langevin diffusions, can also be considered (Stoltz & Trstanova, 2016).

In this work kinetic energies of the form $K(p) = \sum_i K_i(p_i)$ were used. The reason for this is that $\partial_i K(p) = \partial_i K_i(p_i)$, meaning that when $\pi(x)$ is of the product form the Hamiltonian flows become independent. If a different choice were made, for example $K(p) = \beta^{-1}(1 + \gamma^{-1}p^T p)^{\beta/2}$, then $\partial_i K(p) = (1 + \gamma^{-1}p^T p)^{\beta/2-1} p_i$, so this does not always happen, here only when $\beta = 2$. It is unclear which of these approaches is preferable from a theoretical perspective, but evaluating $p^T p$ is certainly a computational burden, and previous empirical evidence favours our approach (Lu et al., 2017).

Outside the realm of separable Hamiltonians, in Girolami & Calderhead (2011) the choice $K(x, p) = \{\log |G(x)| + p^T G^{-1}(x)p\}/2$ is advised. Since an implicit integrator is required, the composite gradient intuition no longer directly applies. Here instead

$$x_\varepsilon = x_0 + \frac{\varepsilon}{2} \left[\{G^{-1}(x_0) + G^{-1}(x_\varepsilon)\} \left\{ p_0 - \frac{\varepsilon}{2} \nabla U(x_0) - \frac{\varepsilon}{2} \nabla f(x_0, p_{\frac{\varepsilon}{2}}) \right\} \right],$$

where $G(x)$ is some Riemannian metric and $2f(x, p) = \log |G(x)| + p^T G^{-1}(x)p$. Key drivers of the dynamics are the terms $G^{-1}(x_0)\nabla U(x_0)$ and $G^{-1}(x_\varepsilon)\nabla U(x_0)$, which are often called the natural gradient (Amari & Nagaoka, 2007). Given recent results (Taylor, 2015), this approach may be advantageous when $\pi(\cdot)$ has heavy-tails. The log-determinant term can also be beneficial (Betancourt & Girolami, 2015).

Another line of further study is to assess how kinetic energy choice affects dynamic implementations such as the No-U-Turn Sampler (Hoffman & Gelman, 2014). Results in simple cases suggest that such implementations are favourable when $\pi(\cdot)$ is heavy-tailed (Livingstone et al., 2016).

ACKNOWLEDGEMENT

We thank the reviewers for helping to improve the paper. SL thanks Michael Betancourt for sharing a calculation related to Proposition 5, Matteo Fasiolo for advice on quantile regression, and Gabriel Stoltz, Zofia Trstanova and Paul Fearnhead for useful discussions. The authors acknowledge support from the Engineering and Physical Sciences Research Council for support through grants EP/K014463/1, i-like (SL and GOR) and EP/D002060/1, CRiSM (GOR).

SUPPLEMENTARY MATERIAL

Supplementary material available at *Biometrika* online includes proofs of propositions along with some intermediate technical results and a numerical example of the efficiency-robustness trade-off.

REFERENCES

- AKHMATSKAYA, E. & REICH, S. (2008). GSHMC: An efficient method for molecular simulation. *Journal of Computational Physics* **227**, 4934–4954.
- AMARI, S.-I. & NAGAOKA, H. (2007). *Methods of information geometry*, vol. 191. American Mathematical Soc.
- BARTHELMÉ, S. & CHOPIN, N. (2011). Discussion on Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B* **73**, 163–164.
- BESKOS, A., PILLAI, N., ROBERTS, G., SANZ-SERNA, J.-M., STUART, A. et al. (2013). Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli* **19**, 1501–1534.
- BETANCOURT, M., BYRNE, S., LIVINGSTONE, S. & GIROLAMI, M. (2017). The geometric foundations of Hamiltonian Monte Carlo. *Bernoulli* **23**, 2257–2298.
- BETANCOURT, M. & GIROLAMI, M. (2015). Hamiltonian Monte Carlo for hierarchical models. *Current trends in Bayesian methodology with applications* **79**, 30.
- BISSIRI, P. G., HOLMES, C. & WALKER, S. G. (2016). A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- BOU-RABEE, N., SANZ-SERNA, J. M. et al. (2017). Randomized hamiltonian monte carlo. *The Annals of Applied Probability* **27**, 2159–2194.
- BOX, G. E. & TIAO, G. C. (1973). *Bayesian inference in statistical analysis*. Reading, Mass: Addison-Wesley.
- CHEN, T., FOX, E. B. & GUESTRIN, C. (2014). Stochastic Gradient Hamiltonian Monte Carlo. In *ICML*.
- DALALYAN, A. S. (2017). Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **79**, 651–676.
- FASIOLO, M., GOUDE, Y., NEDELLEC, R. & WOOD, S. N. (2017). Fast calibrated additive quantile regression. *arXiv preprint arXiv:1707.03307*.
- FRANK, L. E. & FRIEDMAN, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35**, 109–135.
- GILKS, W. R. & WILD, P. (1992). Adaptive rejection sampling for Gibbs sampling. *Applied Statistics*, 337–348.
- GIROLAMI, M. & CALDERHEAD, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**, 123–214.
- GOLDENFELD, N. (1992). *Lectures on phase transitions and the renormalization group*. Addison-Wesley, Advanced Book Program, Reading.
- HOFFMAN, M. D. & GELMAN, A. (2014). The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* **15**, 1593–1623.
- JARNER, S. F. & HANSEN, E. (2000). Geometric ergodicity of Metropolis algorithms. *Stochastic processes and their applications* **85**, 341–361.
- JARNER, S. F. & ROBERTS, G. O. (2007). Convergence of Heavy-tailed Monte Carlo Markov Chain Algorithms. *Scandinavian Journal of Statistics* **34**, 781–815.
- LEIMKUEHLER, B. & REICH, S. (2004). *Simulating hamiltonian dynamics*, vol. 14. Cambridge University Press.
- LIANG, J.-Q. & MÜLLER-KIRSTEN, H. (1992). Periodic instantons and quantum-mechanical tunneling at high energy. *Phys. Rev. D* **46**, 4685–4690.
- LINDSTEN, F. & DOUCET, A. (2016). Pseudo-Marginal Hamiltonian Monte Carlo. *arXiv preprint arXiv:1607.02516*.
- LIVINGSTONE, S., BETANCOURT, M., BYRNE, S. & GIROLAMI, M. (2016). On the Geometric Ergodicity of Hamiltonian Monte Carlo. *arXiv preprint arXiv:1601.08057*.

- 530 LU, X., PERRONE, V., HASENCLEVER, L., TEH, Y. W. & VOLLMER, S. J. (2017). Relativistic Monte Carlo. *AISTATS 2017*.
- MEYN, S. P. & TWEEDIE, R. L. (1993). *Markov chains and stochastic stability*. London: Springer-Verlag.
- MINEO, A. M. (2014). *normalp: Routines for Exponential Power Distribution*. R package version 0.7.0.
- MIRA, A. et al. (2001). On Metropolis-Hastings algorithms with delayed rejection. *Metron* **59**, 231–241.
- 535 NEAL, R. M. (2003). Slice sampling. *Annals of statistics*, 705–741.
- NEAL, R. M. (2011). MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo* **2**, 113–162.
- NISHIMURA, A., DUNSON, D. & LU, J. (2017). Discontinuous Hamiltonian Monte Carlo for sampling discrete parameters. *arXiv preprint arXiv:1705.08510*.
- OKUDO, M. & SUZUKI, H. (2016). Hamiltonian Monte Carlo with Explicit, Reversible, and Volume-preserving Adaptive Step Size Control. *Technical report*.
- 540 PARK, T. & CASELLA, G. (2008). The bayesian lasso. *Journal of the American Statistical Association* **103**, 681–686.
- ROBERTS, G. O. & ROSENTHAL, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical science* **16**, 351–367.
- ROBERTS, G. O. & ROSENTHAL, J. S. (2004). General state space Markov chains and MCMC algorithms. *Probability Surveys* **1**, 20–71.
- 545 ROBERTS, G. O. & TWEEDIE, R. L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 341–363.
- STEPHENS, D. (2011). Discussion on Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B* **73**, 197.
- 550 STOLTZ, G. & TRSTANOVA, Z. (2016). Stable and accurate schemes for Langevin dynamics with general kinetic energies. *arXiv preprint arXiv:1609.02891*.
- STRAMER, O. & TWEEDIE, R. (1999). Langevin-type models ii: self-targeting candidates for MCMC algorithms. *Methodology and Computing in Applied Probability* **1**, 307–328.
- STRATHMANN, H., SEJDINOVIC, D., LIVINGSTONE, S., SZABO, Z. & GRETTON, A. (2015). Gradient-free Hamiltonian Monte Carlo with efficient kernel exponential families. In *Advances in Neural Information Processing Systems*.
- 555 TAYLOR, K. (2015). *Exact Algorithms for simulation of diffusions with discontinuous drift and robust Curvature Metropolis-adjusted Langevin algorithms*. Ph.D. thesis, University of Warwick.
- YANG, Z. & RODRÍGUEZ, C. E. (2013). Searching for efficient Markov chain Monte Carlo proposal kernels. *Proceedings of the National Academy of Sciences* **110**, 19307–19312.
- 560 YU, K. & MOYEED, R. A. (2001). Bayesian quantile regression. *Statistics & Probability Letters* **54**, 437–447.
- ZHANG, Y., WANG, X., CHEN, C., HENAO, R., FAN, K. & CARIN, L. (2016). Towards Unifying Hamiltonian Monte Carlo and Slice Sampling. In *Advances in Neural Information Processing Systems*.