

# UC Irvine

## UC Irvine Previously Published Works

### Title

Neural network gradient Hamiltonian Monte Carlo

### Permalink

<https://escholarship.org/uc/item/1148n4q0>

### Journal

Computational Statistics, 34(1)

### ISSN

0943-4062

### Authors

Li, Lingge  
Holbrook, Andrew  
Shahbaba, Babak  
[et al.](#)

### Publication Date

2019-03-01

### DOI

10.1007/s00180-018-00861-z

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Neural Network Gradient Hamiltonian Monte Carlo

Lingge Li, Andrew Holbrook, Babak Shahbaba, Pierre Baldi  
University of California, Irvine

November 16, 2017

## Abstract

Hamiltonian Monte Carlo is a widely used algorithm for sampling from posterior distributions of complex Bayesian models. It can efficiently explore high-dimensional parameter spaces guided by simulated Hamiltonian flows. However, the algorithm requires repeated gradient calculations, and these computations become increasingly burdensome as data sets scale. We present a method to substantially reduce the computation burden by using a neural network to approximate the gradient. First, we prove that the proposed method still maintains convergence to the true distribution though the approximated gradient no longer comes from a Hamiltonian system. Second, we conduct experiments on synthetic examples and real data sets validate the proposed method.

## 1 Introduction

Hamiltonian Monte Carlo (HMC) uses local geometric information provided by the log-posterior gradient to explore the high posterior density regions of the parameter space [10]. Compared to the Metropolis-Hastings random walk algorithm, HMC has high acceptance probability and low sample auto-correlation even when the parameter space is high-dimensional. That said, the advantages of HMC come at a computational cost that limits its application to smaller data sets. The gradient calculation involves the entire data set and scales linearly with the number of observations. As HMC needs to calculate the gradient multiple times within every single step, performing HMC on millions of observations requires an enormous computational budget. Allowing HMC to scale to large data sets would help tackle the double challenge of big data and big models.

There have been two main approaches to scaling HMC to larger data sets. The first is stochastic gradient HMC, which calculates the gradient on subsets of the data. [12] implemented a stochastic gradient version of Langevin Dynamics, which may be viewed as single-step HMC. [3] introduced stochastic gradient HMC with “friction” to counterbalance the inherently noisy gradient. However,

these methods may not be optimal because subsampling inevitably reduces the acceptance probability of HMC [2].

The second approach relies on a surrogate function, the gradient of which is less expensive to calculate. [11][7] used Gaussian process (GP) to produce satisfactory results in lower dimensions. However, training a GP is itself computationally expensive and training points must be chosen with great care. More recently, [13] implemented neural network surrogate with random bases. It outperforms GP surrogate in their experiments but fails in parameter spaces of moderate dimensionality.

In this paper, we develop a third approach, neural network gradient HMC (NNgHMC), by using a neural network to directly approximate the gradient instead of using it as a surrogate. We also train all the neural network weights through backpropagation rather than having random weights [13]. Compared to existing methods, the proposed method can emulate Hamiltonian flows accurately even when dimensionality increases. In Section 3, details of our method and proof of convergence are presented. Section 4 includes experiments to validate our method and comparisons with previous methods on synthetic and real datasets.

## 2 Background

### 2.1 Hamiltonian Monte Carlo

Let  $x \sim \pi(x|q)$  denote a probabilistic model with  $p$  a probability density function and  $q$  its corresponding parameter. We also make  $q$  a random variable by giving the parameter a prior distribution  $\pi(q)$ . The integration constant of the posterior distribution

$$\pi(q|x) = \frac{\pi(x|q)\pi(q)}{\int \pi(x|q)\pi(q) dq} \quad (1)$$

is usually analytically intractable, but the distribution can be numerically simulated using MCMC. The Metropolis-Hastings algorithm constructs a Markov chain that randomly proposes a new state  $q'$  from current state  $q$  based on transition distribution  $g(q'|q)$  and moves from  $q$  to  $q'$  with probability  $\min\{1, \frac{\pi(q'|X)g(q|q')}{\pi(q|X)g(q'|q)}\}$ . Unfortunately, in a higher dimensional space, the probability of randomly moving to  $q'$  drops dramatically as the unit ball around  $q$  grows exponentially in volume. Therefore, the MH algorithm has trouble exploring the posterior efficiently in higher dimensions.

The idea of HMC is to explore a frictionless landscape induced by potential energy function  $U$  and kinetic energy function  $K$  where potential energy  $U(q) = -\log \pi(x|q)\pi(q)$  is proportional to the negative log posterior. HMC introduces an auxiliary Gaussian momentum  $p$ , and  $K(p)$  is the negative log density of  $p$ . Potential energy  $U$  tends to convert to kinetic energy  $K$  so  $q$  will likely move to a position with higher posterior density. More formally, the Hamiltonian system

is defined by the following equations.

$$H(q, p) = U(q) + K(p) = -\left(\log \pi(q) + \sum_{i=1}^N \log \pi(x_i|q)\right) + \frac{1}{2}p^T p, \quad (2)$$

$$\frac{dq}{dt} = \frac{\partial H}{\partial p} = \frac{\partial K}{\partial p} = p \quad (3)$$

$$\frac{dp}{dt} = -\frac{\partial H}{\partial q} = -\frac{\partial U}{\partial q} = \nabla_q \left(\log \pi(q) + \sum_{i=1}^N \log \pi(x_i|q)\right). \quad (4)$$

In theory, convergence of HMC is guaranteed by the time reversibility of the Hamiltonian dynamics which, in turn, renders the Markov chain transitions reversible, thus ensuring detailed balance. By conservation of the Hamiltonian, HMC has acceptance probability 1 and can travel arbitrarily long trajectories along energy level contours. In practice, the Hamiltonian dynamics is simulated with the leapfrog algorithm which adds numerical errors. To ensure convergence to the posterior, a Metropolis correction step is done at the end of each trajectory.

Within each simulated trajectory, the leapfrog algorithm iterates back and forth between Equations (3) and (4), the latter of which features a summation over the log-likelihood evaluated at *each* separate data point. For large data sets, this repeated evaluation of the gradient becomes infeasible. In Section 3, we show how to greatly speed up HMC using neural network approximations to this gradient term, but first we introduce an important predecessor to our method, the surrogate HMC class of algorithms.

## 2.2 Surrogate HMC

Two methods for approximating the log-posterior in the HMC context have already been advanced. The first uses a Gaussian Process prior to model the log-posterior, the second uses a neural network. We refer to the latter as neural network surrogate HMC (NNsHMC). It is natural that both models would be used in such a capacity: Cybenko [4] showed that neural networks can provide universal function approximation, and Neal [9] showed that certain probabilistic neural networks converge to Gaussian processes as the number of hidden units goes to infinity. In this section, we focus on NNsHMC, since it is more closely related to our method (Section 3).

NNsHMC approximates the potential energy  $U$  with a neural network surrogate  $\hat{U}$  and uses  $\nabla \hat{U}$  during leapfrog steps. The surrogate neural network has one hidden layer with softplus activation:

$$\hat{U}(q) = W_2 \ln(1 + \exp(W_1 q + b_1)) + b_1 \quad (5)$$

where  $W_1, W_2$  and  $b_1, b_2$  are weight matrices and bias vectors, respectively. Under this formulation, one can explicitly calculate the gradient

$$\nabla \hat{U} = W_1^T \text{diag}(W_2) \frac{1}{1 + \exp(-(W_1 q + b_1))} \quad (6)$$

and represent  $\nabla\widehat{U}$  with another neural network, which is just the backpropagation graph of  $\widehat{U}$ . Therefore, we can view neural network surrogate as using a constrained network with tied weights and local connections to approximate the gradient.

For training the neural network, Zhang et al [13] uses extreme learning machine (ELM) [5]. ELM is a simple algorithm that randomly projects the input to the hidden layer and only trains the weights from the hidden layer to the output. Random projection is widely used in machine learning but backpropagation is the “default” training method for most neural networks with its optimality theoretically explained by Baldi et al [1]. Moreover, since the goal is to improve computational efficiency, we want to make the surrogate neural network as small as possible. From this point of view, large hidden layers often seen in ELMs are less than optimal.

### 3 Neural network gradient HMC

---

**Algorithm 1** Neural network gradient HMC

---

```

Initialize  $q^{(0)}$ , leapfrog step number  $L$  and step size  $\epsilon$ 
for  $t = 1, 2, \dots, T$  do
     $q_0 = q^{(t-1)}$ 
    Sample momentum  $p_0 \sim \mathcal{N}(0, I)$ 
     $p_0 = p_0 - \frac{\epsilon}{2} \widehat{\nabla U}(q_t)$   $\triangleright$  Leapfrog steps with approximated gradient  $\widehat{\nabla U}$ 
    instead of  $\nabla U$ 
    for  $i = 1, 2, \dots, L$  do
         $q_i = q_{i-1} + \epsilon p_{i-1}$ 
         $p_i = p_{i-1} - \epsilon \widehat{\nabla U}(q_i)$ 
    end for
     $p_L = p_L - \frac{\epsilon}{2} \widehat{\nabla U}(q_L)$ 
     $r = \exp(H(q_L, p_L) - H(q_0, p_0))$ ,  $u \sim Uniform(0, 1)$ 
    if  $u < \min(1, r)$  then  $\triangleright$  Metropolis accept/reject based on  $H = U + K$ 
         $q^{(t)} = q_L$ 
    else
         $q^{(t)} = q_0$ 
    end if
end for

```

---

In contrast to previous work, NNgHMC does not use a surrogate function for  $U$  but fits a neural network to approximate  $\nabla U$  directly with backpropagation. Training data  $(q, \nabla U(q))$  for the neural network are collected during the early period of HMC shortly after convergence. Once the approximate gradient is learned, the algorithm is exactly the same as classical HMC, but with neural network gradient  $\widehat{\nabla U}$  replacing  $\nabla U$ . Details are given in Algorithm 1.

One benefit of our method occurs as early as the data collection process. Since we approximate the gradient  $\nabla U$  and not  $U$ , we can collect more training

data faster: surrogate HMC must reach the end of a leapfrog trajectory before obtaining a single (functional evaluation) training sample; the same leapfrog trajectory renders a new (gradient evaluation) training sample for each leapfrog step, and the number of such steps in a single trajectory can range into the hundreds.

Suppose that there are  $N$  data points  $x_n$  and that the parameter space is  $d$ -dimensional. In this case, gradient calculations involve  $d$  partial derivatives

$$\frac{\partial U}{\partial q_j} = -\frac{\partial}{\partial q_j} \log \left( \pi(q) \prod_{i=1}^N \pi(x_i|q) \right) = -\frac{\partial}{\partial q_j} \log \pi(q) - \sum_{i=1}^N \frac{\partial}{\partial q_j} \log \pi(x_i|q), \quad (7)$$

each of which involves a summation over the  $N$  data points. On the other hand, performing a forward pass in a shallow neural network is proportional only to the hidden layer size  $s \ll N$ . Once the neural network is trained on burn-in samples, posterior sampling with approximated gradient is orders of magnitude faster.

Although the neural network gradient approximation  $\widehat{\nabla}U(q)$  is not the same as  $\nabla U(q)$ , the method nonetheless samples from the true posterior. If one were able to simulate the Hamiltonian system directly, i.e. without numerical integration, then all the benefits of HMC would be preserved in the limit, as the gradient field approximates the true gradient field to arbitrary degree. On the other hand, the NNgHMC transition kernel—characterized by the approximate gradient leapfrog integrator combined with the Metropolis accept-reject step—leaves the posterior distribution invariant. We formalize the relevant results here and defer proofs to the appendix.

An important litmus test for the validity of our method is that it should leave the Hamiltonian invariant in the limit as step-sizes and gradient approximation errors approach zero. In turn, this result will imply high acceptance probabilities when the system is simulated from numerically, and when gradient approximations are good.

**Proposition 1.** *When the system induced by the approximate gradient field is simulated directly, changes in the Hamiltonian  $H(q, p) = U(q) + K(p)$  converge in probability to 0 as the approximate gradient converges pointwise to the true gradient. That is, for a sequence of approximate gradient fields  $\{\widehat{\nabla}_q^n U\}_{n=1}^\infty$  converging to the true gradient field  $\nabla_q U$ , the change in Hamiltonian values satisfies*

$$\left( \frac{dH}{dt} \right)_n = o_p(1). \quad (8)$$

*Proof.* Following [4], assume we are able to construct a sequence of approximate gradients  $\widehat{\nabla}_q^n H$  satisfying

$$\nabla_q H = \widehat{\nabla}_q^n H + E_n(q), \quad E_n(q) \in B_{1/n}(0), \quad (9)$$

where  $B_{1/n}(0)$  is the ball around the origin of radius  $1/n$ . In this case, the vector field given by the approximate gradient induces a new system of equations:

$$\begin{aligned}\frac{dq_i}{dt} &= \frac{\partial H}{\partial p_i} \\ \frac{dp_i}{dt} &= -\frac{\widehat{\partial H}}{\partial q_i}.\end{aligned}\tag{10}$$

Then it follows that

$$\begin{aligned}\frac{dH}{dt} &= \sum_{i=1}^d \left[ \frac{dq_i}{dt} \frac{\partial H}{\partial q_i} + \frac{dp_i}{dt} \frac{\partial H}{\partial p_i} \right] \\ &= \sum_{i=1}^d \left[ \frac{\partial H}{\partial p_i} \left( \frac{\widehat{\partial H}}{\partial q_i} + E_{n,i}(q) \right) - \frac{\widehat{\partial H}}{\partial q_i} \frac{\partial H}{\partial p_i} \right] \\ &= \sum_{i=1}^d \frac{\partial H}{\partial p_i} E_{n,i} \\ &= p^T E_n \sim N(0, E_n^T E_n).\end{aligned}\tag{11}$$

This last line implies  $p^T E_n$  is  $O_p(\sqrt{E_n^T E_n})$ , and hence that  $\frac{dH}{dt}$  is  $o_p(1)$ .  $\square$

We note that Proposition 1 is a local result, and that local deviations from the true Hamiltonian flow will accrue to larger global deviations in general. While this may seem disconcerting, NNghMC maintains remarkably high acceptance rates in practice. To help understand why this is the case, we present local and global error analyses for the dynamics of the ordinary differential equation initial value problem

$$\frac{d}{dt} z = f(z), \quad z(t_0) = z^0 \in \mathbb{R}^k,\tag{12}$$

approximated with function  $\widehat{f} \approx f$ . These results will then be related back to NNghMC by specifying  $z = (q, p)$  and

$$z = (q, p)^T, \quad f(q, p) = \left( p, -\frac{\partial H}{\partial q} \right)^T, \quad \text{and} \quad \widehat{f}(q, p) = \left( p, -\frac{\widehat{\partial H}}{\partial q} \right)^T.\tag{13}$$

The general form of the following proofs follows after Section 2.1.2 of [8].

**Proposition 2.** (*Local error bounds*) Let  $z^0 = z(0)$  be the initial value, let  $z(\Delta t)$  be the value of the exact, true trajectory after traveling for time  $\Delta t$ , and let  $z^1$  be the value of the computed trajectory using Euler's method applied to the approximated gradient field. Finally, assume that the exact solution is twice continuously differentiable. Then the local error  $\epsilon^1 = z(\Delta t) - z^1$  has the following bounds:

$$\|\epsilon^1\| \leq \Delta t \delta + O(\Delta t^2),\tag{14}$$

where  $\delta = \|f(z_0) - \widehat{f}(z_0)\|$  measures the difference between the true, exact trajectory and the approximated trajectory at point  $z_0$ .

*Proof.* The proof follows from the Taylor expansion of both  $z(\Delta t)$  and  $z^1$ :

$$\begin{aligned}\epsilon^1 &= (z_0 + \Delta t \dot{z}(0) + \frac{1}{2} \Delta t^2 \ddot{z}(\tau)) - (z_0 + \Delta t \widehat{f}(z_0)) \\ &= \Delta t (f(z_0) - \widehat{f}(z_0)) + \frac{1}{2} \Delta t^2 \ddot{z}(\tau),\end{aligned}\quad (15)$$

where  $\tau \in [0, \Delta t]$ . The result follows from the triangular inequality.  $\square$

From the above result, it follows that the local error rate approaches the  $O(\Delta t^2)$  error rate of Euler's method using the true gradient field as  $\delta = \|f(z_0) - \widehat{f}(z_0)\| = \|\frac{\partial H}{\partial q}(z_0) - \widehat{\frac{\partial H}{\partial q}}(z_0)\|$  goes to 0. The same approach can be used to obtain global error bounds.

**Proposition 3.** (*Global error bounds*) We adopt the same notation as above with the addition of the error at iteration  $n$ ,  $\epsilon^n = z(n\Delta t) - z^n$ , where  $z^n$  is the value after  $n$  Euler updates using the approximate gradient field. Again we assume that the exact solution is twice differentiable, and we further assume that it is Lipschitz with constant  $L$ . Then the following bounds on  $\epsilon^n$  hold:

$$\|\epsilon^n\| \leq (e^{n\Delta t L} - 1) \left( \frac{\delta}{L} + O(\Delta t) \right), \quad \text{for } \delta = \max \|f(z(j\Delta t)) - \widehat{f}(z^j)\|, \quad (16)$$

and  $j = 0, \dots, n$ .

*Proof.* The proof proceeds by recursion. Assume that we have obtained  $\epsilon^n = z(t_n) - z^n$ . Letting  $\tau \in [t_n, t_{n+1}]$ , a Taylor's expansion gives:

$$\begin{aligned}\epsilon^{n+1} &= \left( z(t_n) + \Delta t \dot{z}(t_n) + \frac{1}{2} \Delta t^2 \ddot{z}(\tau) \right) - \left( z^n + \Delta t \widehat{f}(z^n) \right) \\ &= \left( z(t_n) + \Delta t f(z(t_n)) + \frac{1}{2} \Delta t^2 \ddot{z}(\tau) \right) - \left( z^n + \Delta t \widehat{f}(z^n) \right) \\ &= (z(t_n) - z^n) + \Delta t (f(z(t_n)) - \widehat{f}(z^n)) + \frac{1}{2} \Delta t^2 \ddot{z}(\tau).\end{aligned}\quad (17)$$

But  $\ddot{z}$  is continuous by assumption, so we can bound it on  $[t_n, t_{n+1}]$  by a constant  $M$ . Furthermore, the Lipschitz assumption combined with the triangle inequality give:

$$\begin{aligned}\|\epsilon^{n+1}\| &\leq \|\epsilon^n\| + \Delta t \left( \|f(z(t_n)) - f(z^n)\| + \|f(z^n) - \widehat{f}(z^n)\| \right) + \frac{\Delta t^2 M}{2} \\ &\leq (1 + \Delta t L) \|\epsilon^n\| + \Delta t \delta + \frac{\Delta t^2 M}{2}\end{aligned}\quad (18)$$

Next we make use of the following recursion relationship:

$$a_{n+1} \leq C a_n + D \implies a_n \leq C^n a_0 + \frac{C^n - 1}{C - 1} D \quad (19)$$



for  $C = (1 + \Delta t L)$  and  $D = \Delta t \delta + \Delta t^2 M/2$ . Noting that  $a_0 = \epsilon^0 = 0$  gives

$$\|\epsilon^n\| \leq (e^{t_n L} - 1) \left( \frac{\delta}{L} + \frac{\Delta t M}{2L} \right), \quad (20)$$

and the result follows.  $\square$

The above result suggests that the usual numerical error caused by a large Lipschitz constant  $L$  can overpower the effects of gradient approximation error  $\delta$ .

The preservation of volume entailed by both the theoretical Hamiltonian flow and the leapfrog integrator is important for HMC. The latter fact implies there is no need for Jacobian corrections within the accept-reject step. It turns out that the NNgHMC dynamics also preserve volume, both for direct and for leapfrog simulation.

**Lemma 1.** *Both for infinitesimal and finite step sizes, the NNgHMC trajectory preserves volume.*

*Proof.* For the finite case, the leapfrog algorithm iterates between shear transformations and so preserves volume [10]. For the case of direct simulation, we use the fact that the Hamiltonian vector field induced by the approximate gradient field has zero divergence (Liouville's Theorem, [10]). Using the notation of Proposition 1:

$$\begin{aligned} \sum_{i=1}^d \left[ \frac{\partial}{\partial q_i} \frac{dq_i}{dt} + \frac{\partial}{\partial p_i} \frac{dp_i}{dt} \right] &= \sum_{i=1}^d \left[ \frac{\partial}{\partial q_i} \frac{\partial H}{\partial p_i} - \frac{\partial}{\partial p_i} \widehat{\frac{\partial H}{\partial q_i}} \right] \\ &= \sum_{i=1}^d \left[ \frac{\partial}{\partial q_i} \frac{\partial H}{\partial p_i} - \frac{\partial}{\partial p_i} \left( \frac{\partial H}{\partial q_i} - E_i \right) \right] \\ &= \sum_{i=1}^d \frac{\partial}{\partial p_i} E_i = 0. \end{aligned} \quad (21)$$

$\square$

Not only does the NNgHMC trajectory preserve volume, it is reversible as well. This easy fact is shown in the proof of Proposition 2.

**Theorem 1.** *The NNgHMC transition kernel leaves the canonical distribution  $\exp\{-H(q, p)\}$  invariant.*

*Proof.* Since leapfrog integration preserves volume and since the Metropolis acceptance probability is the same as for classical HMC, all we need to show is that the leapfrog integration is reversible. This fact follows in the exact same way as for HMC, despite the use of an approximate gradient field to generate

the dynamics:

$$\begin{aligned}
 p_i(t + \epsilon/2) &= p_i(t) - (\epsilon/2) \frac{\widehat{\partial U}}{\partial q_i}(q(t)) \\
 q_i(t + \epsilon) &= q_i(t) + \epsilon p_i(t + \epsilon/2) \\
 p_i(t + \epsilon) &= p_i(t + \epsilon/2) - (\epsilon/2) \frac{\widehat{\partial U}}{\partial q_i}(q(t + \epsilon)).
 \end{aligned}
 \tag{22}$$

These are the same equations as in [10] except with  $\frac{\widehat{\partial U}}{\partial q_i}$  replacing  $\frac{\partial U}{\partial q_i}$ . Hence, just as in [10], the NNgHMC leapfrog equations are symmetric and thus reversible: to reverse a sequence of leapfrog dynamics, negate  $p$ , take the same number of steps, and negate  $p$  again. It follows that the NNgHMC transition kernel leaves the canonical distribution invariant and is an asymptotically exact method for sampling from the posterior distribution.  $\square$

## 4 Experiments

### 4.1 Gaussian with ill-conditioned covariance

First, we illustrate the proposed method on a multivariate Gaussian distribution  $q \sim \mathcal{N}_{30}(0, \Sigma)$  where  $\Sigma$  is a diagonal matrix with smallest value 0.1, largest value 1000 and other values uniformly drawn between 1 and 100. As the distribution is on very disparate scales in different dimensions, HMC needs accurate gradient information to move accordingly. For HMC, the leapfrog step size is set to be 0.5 and the number of steps is set to be 100 so that acceptance probability is around 0.7. If the step size is too big, HMC would miss the high density region in the narrowest dimension. Without a sufficiently long trajectory, HMC would fail to explore the elongated tails in the widest dimension.

We collect sample gradients until 50 iterations after convergence to train the neural network. The network architecture is very simple: 30-100-30 with ReLU units in the hidden layer. With the same tuning parameters as HMC, NNgHMC has acceptance probability around 0.5. Hence, there is little downside to using the approximate gradient field. As shown in Figure 1, NNgHMC converges to the true posterior as HMC. Moreover, Figure 2 compares two trajectories in the true and approximated gradient fields.

### 4.2 Gaussian process comparison

Gaussian process (GP) is characterized by the kernel function  $\mathcal{K}(\cdot)$ . Here we use the squared exponential kernel parametrized by  $\mathcal{K}(x, x') = \exp - \frac{(x-x')^2}{2l^2}$  with added white noise  $\sigma^2 I$  to the diagonal of the covariance matrix. The squared exponential kernel, the default choice, is infinitely differentiable and gives rise to another Gaussian process as the derivative. Given observations  $X$  and  $Y$ , we can explicitly write down the mean of the derivative at  $x^*$ .




Figure 1: NNgHMC posterior (bottom) captures the highly elongated shape of the Gaussian distribution in the two most extreme dimensions ( $\sigma_1^2 = 0.1, \sigma_{30}^2 = 1000$ ) as well as the HMC posterior (top). Note that the x- and y-axes are on very different scales.




Figure 2: With the same initial position and momentum, the leapfrog trajectory in the same dimensions as in Figure 1 using approximated gradient (blue) faithfully resembles the one using true gradient (red) despite heavy oscillation on the energy level contour. The periodic nature of the Hamiltonian flow reflects the fact that the Hamiltonian is that of the harmonic oscillator, i.e. both potential and kinetic energies are quadratic. The vectors and trajectories are slightly jittered for plotting.

$$E \frac{\partial f}{\partial x^*} = \frac{\partial}{\partial x^*} E f = \frac{\partial}{\partial x^*} \mathcal{K}(x^*, X) \mathcal{K}(X, X)^{-1} Y \quad (23)$$

The white noise is often added for numerical stability and its magnitude affects the “wiggleness.” Here we estimate both the length scale  $l$  in the squared exponential kernel and the white noise parameter  $\sigma$  jointly with maximum likelihood. The estimation requires inverting the observed covariance matrix, which is  $\mathcal{O}(n^3)$  where  $n$  is the number of observations.

We compare with GP surrogate method on multivariate Gaussian distributions with covariance  $I_d$  where  $d$  varies from 10 to 40. For each Gaussian, we generate  $n$  training data points to train the GP surrogate and neural network. The neural network has 100 hidden units and is trained for 10 epochs. After training, both methods are used to draw 1000 samples.

Table 1 compares acceptance probability for the two methods. We can see that the neural network predicted gradient provides better approximation over-

all than the gradient of GP as indicated by higher acceptance probability. This advantage is more pronounced in higher dimension.

Table 1: Acceptance probability when sampling from multivariate Gaussian

Method	Dimension / Training	500	1000	2000
Gaussian process	10	0.65	0.61	0.57
	20	0.64	0.65	0.62
	40	0.31	0.32	0.32
Neural network gradient	10	0.95	0.96	0.97
	20	0.82	0.87	0.91
	40	0.61	0.75	0.87

### 4.3 Stochastic gradient HMC comparison

Naïve stochastic gradient HMC using mini-batches of data is problematic as the noisy gradient can push the sampler away from the target region. Recent more advanced stochastic gradient method uses a friction term and is shown to sample from the true posterior asymptotically. The formulation of SGHMC is given by:

$$d\theta = M^{-1}r dt \tag{24}$$

$$dr = -\nabla U(\theta)dt - BM^{-1}r dt + N(0, 2Bdt) \tag{25}$$

where  $N(0, 2Bdt)$  is the noise added to the gradient by subsampling. In practice, the friction term  $BM^{-1}r dt$  is set arbitrarily.

To further improve speed, SGHMC does not perform Metropolis-Hastings correction and uses very small step sizes. The SGHMC posterior is dependent on the choice of step size; however, *a priori* one would not know the optimal step size. Here we want to show that while SGHMC provides fast approximation of the true posterior when data are abundant, the SGHMC posterior may not be suitable for inference.

In our experiment, the Cover Type data from UCI machine learning repository is used. Standard HMC is run for 4000 iterations with 50 leapfrog steps and step size 0.002. SGHMC is also run for 4000 iterations with default parameters and varying step sizes from  $5e - 6$  to  $5e - 8$ .

Figure 3 illustrates the main issue with SGHMC. For these two marginal distributions, the SGHMC posteriors have roughly the same location but completely different shapes. On the other hand, NNgHMC posteriors agree with the standard HMC posteriors almost exactly.




Figure 3: Histograms of marginal posteriors of logistic regression model coefficients with Laplace prior on Cover Type data. Blue: standard HMC; Red: stochastic gradient HMC; Green: neural network gradient HMC

#### 4.4 Speed evaluation on real data

Similar to other surrogate methods, NNgHMC has three stages: training data collection, training and sampling. We have demonstrated that using a neural network can provide accurate approximation of the gradient, however, the effectiveness of our method still needs to be evaluated by time. If the neural network requires too much training data, then it would not save computation time. Here we first run standard HMC to draw a desired number of samples (10000) and record time as benchmark. Then we collect different amounts of training data points (10%, 15% and 20% of total number) and use NNgHMC to draw remaining samples.

The News Popularity data set has 39,797 observations with 60 features derived from Mashable articles and a quantitative outcome, which is the number of shares. As some features are highly skewed and require additional pre-processing, we only keep 44 of the features and dichotomize the outcome at its median for logistic regression.

HMC tuning parameters are chosen with target acceptance probability 0.8 in mind. The number of leapfrog steps is 10 and step size is 0.01. Convergence of HMC is determined by visual inspection of trace plots and takes about 400 iterations. The two neural network training hyperparameters, number of hidden units and number of epochs, are fixed at 100 and 50.

As shown in Table 2, 10% of training data is sufficient for the neural network to learn the gradient and gives the most speed-up. While adding more training data increases the quality of gradient approximation, the computation cost outweighs the benefit of higher acceptance probability.

NNgHMC is implemented in Keras and uses the default Adam optimizer [6] during training. All experiments are performed on a 3.4 GHz Intel Quad-Core CPU and our code is available at: <https://github.com/linggeli7/hamiltonian>.

Table 2: Experiment results on News data

Method	AP	ESS	CPU time	Median ESS/s	Speed-up
Standard	0.77	(777, 2021, 5929)	3607s	0.66	1
NNg (10%)	0.61	(605, 1416, 4865)	502s	2.82	4.27
NNg (15%)	0.64	(620, 1382, 5500)	678s	2.04	3.09
NNg (20%)	0.68	(700, 1731, 5397)	854s	2.03	3.08

AP: acceptance probability

ESS: effective sample size (min, median, max) after removing 10% burn-in

## 5 Discussion

Whereas HMC is helpful for computing large Bayesian models, its repeated gradient evaluations become overly costly for big data analysis. We have presented a method that circumvents the costly gradient evaluations, not by subsampling data batches but by learning an approximate gradient that is functionally free of the data. We found that multi-output, feedforward neural networks were ripe for this application: our NNgHMC was able to handle models of comparatively large dimensionality.

The NNgHMC algorithm is an important paradigm shift away from the class of surrogate function approximate HMC algorithms, but this shift leaves many open questions. Much work is needed to extend NNgHMC to an on-line, adaptive methodology: what measures of approximation error will be useful criteria for ending the training regime of the algorithm, and are there benefits to iterating between training and sampling regimes? Are there any valid second-order extensions to the NNgHMC algorithm à la Riemannian HMC? Finally—and most interestingly—can the representational power of deep neural networks be leveraged for more accurate approximations to the Hamiltonian flow?

## References

- [1] Pierre Baldi and Peter Sadowski. A theory of local learning, the learning channel, and the optimality of backpropagation. *Neural Networks*, 83:51–74, 2016.
- [2] MJ Betancourt. The fundamental incompatibility of hamiltonian monte carlo and data subsampling. *arXiv preprint arXiv:1502.01510*, 2015.
- [3] Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient hamiltonian monte carlo. In *International Conference on Machine Learning*, pages 1683–1691, 2014.
- [4] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems (MCSS)*, 2(4):303–314, 1989.
- [5] Guang-Bin Huang, Qin-Yu Zhu, and Chee-Kheong Siew. Extreme learning machine: a new learning scheme of feedforward neural networks. In *Neural Networks, 2004. Proceedings. 2004 IEEE International Joint Conference on*, volume 2, pages 985–990. IEEE, 2004.
- [6] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [7] Shiwei Lan, Tan Bui-Thanh, Mike Christie, and Mark Girolami. Emulation of higher-order tensors in manifold monte carlo methods for bayesian inverse problems. *Journal of Computational Physics*, 308:81–101, 2016.
- [8] Benedict Leimkuhler and Sebastian Reich. *Simulating hamiltonian dynamics*, volume 14. Cambridge university press, 2004.
- [9] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [10] Radford M Neal et al. Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2:113–162, 2011.
- [11] Carl Edward Rasmussen, JM Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, and M West. Gaussian processes to speed up hybrid monte carlo for expensive bayesian integrals. In *Bayesian Statistics 7*, pages 651–659, 2003.
- [12] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688, 2011.
- [13] Cheng Zhang, Babak Shahbaba, and Hongkai Zhao. Hamiltonian monte carlo acceleration using surrogate functions with random bases. *Statistics and Computing*, pages 1–18, 2015.