

# Variational Hamiltonian Monte Carlo via Score Matching

Cheng Zhang\*, Babak Shahbaba†, and Hongkai Zhao‡

**Abstract.** Traditionally, the field of computational Bayesian statistics has been divided into two main subfields: variational methods and Markov chain Monte Carlo (MCMC). In recent years, however, several methods have been proposed based on combining variational Bayesian inference and MCMC simulation in order to improve their overall accuracy and computational efficiency. This marriage of fast evaluation and flexible approximation provides a promising means of designing scalable Bayesian inference methods. In this paper, we explore the possibility of incorporating variational approximation into a state-of-the-art MCMC method, Hamiltonian Monte Carlo (HMC), to reduce the required expensive computation involved in the sampling procedure, which is the bottleneck for many applications of HMC in big data problems. To this end, we exploit the regularity in parameter space to construct a free-form approximation of the target distribution by a fast and flexible surrogate function using an optimized additive model of proper random basis, which can also be viewed as a single-hidden layer feedforward neural network. The surrogate function provides sufficiently accurate approximation while allowing for fast computation in the sampling procedure, resulting in an efficient approximate Bayesian inference algorithm. We demonstrate the advantages of our proposed method using both synthetic and real data problems.

**MSC 2010 subject classifications:** Primary 65C60; secondary 65C05.

**Keywords:** Markov Chain Monte Carlo, variational inference, free-form approximation.

## 1 Introduction

Bayesian inference has been successful in modern data analysis. Given a probabilistic model for the underlying mechanism of the observed data, Bayesian methods properly quantify uncertainty and reveal the landscape or global structure of parameter space. While conceptually simple, exact posterior inference in many Bayesian models is often intractable. Therefore, in practice, it is common to use approximation methods among which Markov chain Monte Carlo (MCMC) and variational Bayes (VB) are the two most popular choices.

The MCMC approach is based on drawing a series of samples by constructing a Markov chain with guaranteed convergence to the target distribution. Therefore, MCMC methods are asymptotically unbiased. Simple methods such as random-walk Metropolis (Metropolis et al., 1953), however, often suffer from slow mixing (due to their random

---

\*UC Irvine, [chengz4@uci.edu](mailto:chengz4@uci.edu)

†UC Irvine, [babaks@uci.edu](mailto:babaks@uci.edu)

‡UC Irvine, [zhao@math.uci.edu](mailto:zhao@math.uci.edu)

walk nature) when encountering complicated models with strong dependencies among parameters. Introducing an auxiliary momentum variable, Hamiltonian Monte Carlo (HMC) (Duane et al., 1987; Neal, 2011) reduces the random walk behavior by proposing states following a Hamiltonian flow which preserves the target distribution. By incorporating the geometric information of the target distribution, e.g., the gradient, HMC is able to generate distant proposals with high acceptance probabilities, enabling more efficient exploration of the parameter space than standard random-walk proposals.

A major bottleneck of HMC, however, is the computation of the gradient of the potential energy function in order to simulate the Hamiltonian flow. As the datasets involved in many practical tasks may have millions to billions of observations, such gradient computations are infeasible since they need full scans of the entire dataset. In recent years, many attempts have been made to develop scalable MCMC algorithms that can cope with very large data sets (Welling and Teh, 2011; Ahn et al., 2012; Chen et al., 2014; Ding et al., 2014). The key idea of these methods stems from stochastic optimization where noisy estimates of the gradient based on small subsets of the data are utilized to scale up the algorithms. The noise introduced by subsampling, however, could lead to non-ignorable loss of accuracy, which in turn hinders the exploration efficiency of standard MCMC approaches (Betancourt, 2015).

The main alternative to MCMC is variational Bayes (VB) inference (Jordan et al., 1999; Wainwright and Jordan, 2008). As a deterministic approach, VB transforms Bayesian inference into an optimization problem where a parametrized distribution is introduced to approximate the target posterior distribution by minimizing the Kullback–Leibler (KL) divergence with respect to the variational parameters. Compared to MCMC methods, VB introduces bias but is usually faster.

A natural question would be: can we combine both methods to mitigate the drawbacks and get the best of both worlds? The first attempt in this direction was proposed by de Freitas et al. (2001) where a variational approximation was used as proposal distribution in a block Metropolis-Hasting (MH) algorithm in order to capture high probability regions quickly, thus facilitating convergence. Recently, a new synthesis of variational inference and Markov chain Monte Carlo methods has been explored in Salimans et al. (2015) where one or more steps of MCMC are integrated into variational approximation. The extra flexibility from MCMC steps provides a rich class of distributions to approximate the exact posterior distribution.

In this work, we explore the possibility of utilizing variational approximation to speed up HMC for problems with large scale datasets. The key idea is to integrate fast variational approximations into the sampling procedure so that the overall computational complexity can be reduced. The main contributions are summarized as follows:

1. We train a fast neural network surrogate using extreme learning machine (ELM) to approximate the log-posterior (potential energy function). Neural networks with randomly-mapped features (with almost any nonlinear activation functions) have been proved to approximate a rich class of functions arbitrarily well (Huang et al., 2006a). After the fast training of a reasonably accurate neural network surrogate, we use its gradient to simulate the surrogate induced Hamiltonian flow in order to

generate proposals. The learning and training of the surrogate function provides an effective and flexible way to explore both regularity of the parameter space and redundancy in the data collectively. This can be viewed as an implicit subsampling. However, unlike the popular subsampling-based methods where the *leapfrog* step-size needs to be annealed or remains small enough to compensate for the resulting bias, the leapfrog stepsize in our approach can remain close to that of the standard HMC algorithm while keeping a comparable acceptance probability. As such, we can maintain the efficiency of HMC while reducing the computational cost.

2. A new training procedure is proposed for an efficient and general surrogate method. The neural network surrogate is trained by minimizing the squared distance between the gradient of the surrogate and the gradient of the target (log-posterior), a procedure that resembles the *score matching* approach (Hyvärinen, 2005). The training data are collected while the “modified” HMC sampler (based on the surrogate induced Hamiltonian flow) explores the parameter space. In Hyvärinen (2005), the training data are assumed to be sampled from the target distribution. In our method, by including the gradient of the target distribution in the learning process, the information from the target distribution is explicitly included in the surrogate function. Therefore, the restriction on the training data is relaxed.
3. To further improve speed, we also use the computationally fast surrogate in the Metropolis-Hastings correction step. As a result, the modified HMC sampler will no longer converge to the exact target distribution. Instead, it will converge to the best approximation from an exponential family with pre-specified random sufficient statistics (which we refer to as *free-form* variational approximation). This variational perspective distinguishes our approach from the existing surrogate methods on accelerating HMC. Moreover, compared to traditional fixed-form variational approximations, the free-form variational approximation is usually more flexible and can provide a better fit to more general target distributions.

Our paper is organized as follows. In Section 2, we introduce the two ingredients of our method: Hamiltonian Monte Carlo and fixed-form variational Bayes. Section 3 presents our method, termed Variational Hamiltonian Monte Carlo (VHMC). We demonstrate the efficiency of VHMC using a number of experiments in Section 4 and conclude in Section 5.

## 2 Background

### 2.1 Hamiltonian Monte Carlo

Consider a set of independent and identically distributed observations  $Y = \{y_1, \dots, y_N\}$ . We denote their underlying distribution as  $p(y|\theta)$  with unknown parameter  $\theta$ . Assuming a prior distribution,  $p(\theta)$ , the posterior distribution is obtained by Bayes rule as follows:

$$p(\theta|Y) = \frac{p(Y|\theta)p(\theta)}{p(Y)} \propto \prod_{n=1}^N p(y_n|\theta) \cdot p(\theta).$$

To construct the Hamiltonian dynamical system commonly used in HMC (Duane et al., 1987; Neal, 2011), the position-dependent potential energy function is defined as the negative log of unnormalized posterior density

$$U(\theta) = - \sum_{n=1}^N \log p(y_n|\theta) - \log p(\theta), \quad (1)$$

and the kinetic energy function is defined as a quadratic function of an auxiliary momentum variable  $r$

$$K(r) = \frac{1}{2} r^T M^{-1} r,$$

where  $M$  is a mass matrix, which is often set to the identity matrix,  $I$ . The fictitious Hamiltonian, therefore, is defined as the total energy function of the system

$$H(\theta, r) = U(\theta) + K(r).$$

HMC suppresses random walk behavior by simulating the Hamiltonian dynamical system to propose distant states with high acceptance probabilities. To this end, in order to sample from the posterior distribution  $p(\theta|Y)$ , it augments the parameter space and generates samples from the joint distribution of  $(\theta, r)$

$$\pi(\theta, r) \propto \exp(-U(\theta) - K(r)). \quad (2)$$

Notice that  $\theta$  and  $r$  are separated in (2), so that we can simply drop the momentum samples  $r$  and the  $\theta$  samples follow the marginal distribution which is exactly the target posterior distribution.

To generate proposals, HMC simulates the Hamiltonian flow governed by the following differential equations

$$\frac{d\theta}{dt} = \frac{\partial H}{\partial r} = M^{-1} r, \quad (3)$$

$$\frac{dr}{dt} = -\frac{\partial H}{\partial \theta} = -\nabla_{\theta} U(\theta). \quad (4)$$

Over a period  $t$ , also called the trajectory length, (3) and (4) together define a map  $\phi_t : (\theta_0, r_0) \mapsto (\theta^*, r^*)$  in the augmented parameter space, from the starting state to the end state. As implied by a Hamiltonian flow,  $\phi_t$  is reversible, volume-preserving, and also preserves the Hamiltonian  $H(\theta_0, r_0) = H(\theta^*, r^*)$ . These properties allow us to construct  $\pi$ -invariant Markov chains whose proposals will always be accepted. In practice, however, (3) and (4) are not analytically solvable and we need to resort to numerical integrators. As a symplectic integrator, the *leapfrog* scheme (see Algorithm 1) maintains reversibility and volume preservation and hence is a common practice in the HMC literature. Although discretization introduces bias which needs to be corrected in an Metropolis-Hasting (MH) step, we can control the stepsizes to maintain high acceptance probabilities even for distant proposals.

In recent years, many variants of HMC have been developed to make the algorithm more flexible and generally applicable in a variety of settings. For example, methods proposed in Hoffman and Gelman (2014) and Wang et al. (2013) enable automatic tuning of hyper-parameters such as the stepsize  $\epsilon$  and the number of leapfrog steps. Riemannian manifold HMC (Girolami and Calderhead, 2011) further improves standard HMC's

**Algorithm 1** Hamiltonian Monte Carlo

---

**Input:** Starting position  $\theta^{(1)}$  and step size  $\epsilon$   
**for**  $t = 1, 2, \dots, T$  **do**  
  Resample momentum  $r$   
   $r^{(t)} \sim \mathcal{N}(0, M)$   
   $(\theta_0, r_0) = (\theta^{(t)}, r^{(t)})$   
  Simulate discretization of Hamiltonian dynamics:  
  **for**  $l = 1$  **to**  $L$  **do**  
     $r_{l-1} \leftarrow r_{l-1} - \frac{\epsilon}{2} \nabla_{\theta} U(\theta_{l-1})$   
     $\theta_l \leftarrow \theta_{l-1} + \epsilon M^{-1} r_{l-1}$   
     $r_l \leftarrow r_l - \frac{\epsilon}{2} \nabla_{\theta} U(\theta_l)$   
  **end for**  
   $(\theta^*, r^*) = (\theta_L, r_L)$   
  Metropolis-Hasting correction:  
   $u \sim \text{Uniform}[0, 1]$   
   $\rho = \exp[H(\theta^{(t)}, r^{(t)}) - H(\theta^*, r^*)]$   
  **if**  $u < \min(1, \rho)$  **then**  
     $\theta^{(t+1)} = \theta^*$   
  **else**  
     $\theta^{(t+1)} = \theta^{(t)}$   
  **end if**  
**end for**

---

efficiency by automatically adapting to local structure using the Riemannian geometry of the parameter space. These adaptive techniques could be potentially combined with our proposed method which focuses on reducing the computational complexity.

## 2.2 Fixed-Form Variational Bayes

An alternative to running a Markov chain is to approximate the intractable posterior distribution with a more convenient and tractable distribution. A popular approach for obtaining such an approximation is fixed-form variational Bayes (Honkela et al., 2010; Saul and Jordan, 1996; Salimans and Knowles, 2013) where a parametrized distribution  $q_v(\theta)$  is proposed to approximate the target posterior  $p(\theta|Y)$  by minimizing the KL divergence

$$\begin{aligned} D_{KL}(q_v(\theta)||p(\theta|Y)) &= \int q_v(\theta) \log \left( \frac{q_v(\theta)}{p(\theta|Y)} \right) d\theta, \\ &= \log(p(Y)) + \int q_v(\theta) \log \left( \frac{q_v(\theta)}{p(\theta, Y)} \right) d\theta. \end{aligned} \quad (5)$$

Because  $\log(p(Y))$  is a constant (used extensively in model selection), it suffices to minimize the second term in (5). Usually,  $q_v(\theta)$  is chosen from the exponential family of distributions with the following canonical form:

$$q_v(\theta) = \exp[T(\theta)v - A(v)]\nu(\theta), \quad (6)$$

where  $T(\theta)$  is a row vector of sufficient statistics,  $A(v)$  is for normalization, and  $\nu(\theta)$  is a base measure. The column vector  $v$  is often called the natural parameter of the exponential family distribution  $q_v(\theta)$ . Taking this approach and substituting into (5), we will have a parametric optimization problem in  $v$ :

$$\hat{v} = \arg \min_v \mathbb{E}_{q_v(\theta)} [\log q_v(\theta) - \log p(\theta, Y)]. \quad (7)$$

The above optimization problem can be solved using gradient-based optimization or fix-point algorithms if the expectation and its derivatives with respect to  $v$  can be evaluated analytically. Without assuming posterior independence and requiring conjugate exponential models, posterior approximations of this type are usually much more accurate than a factorized approximation following the mean-field assumptions. However, the requirement for analytical evaluation of the expectation and its derivatives is also very restrictive. To address this issue, some recent works have focused on evaluating these quantities based on stochastic approximation of Monte Carlo integration (Paisley et al., 2012; Ranganath et al., 2014; Kingma and Welling, 2013). As an alternative, Salimans and Knowles (2013) proposed a new optimization algorithm which relates (7) to stochastic linear regression. Compared to classic stochastic approximations, the randomness in the stochastic linear regression approach tends to cancel out and therefore improves the efficiency of the algorithm. Furthermore, their method can be easily extended to flexible mixture and hierarchical approximations through exploiting existing conditional independence structures in the model.

In the next section, we will discuss how the variational Bayes approach can be actually utilized to accelerate HMC. For this, we construct a fast and accurate approximation for the computationally expensive potential energy function. The approximation is provided by variational Bayes and is incorporated in the simulation of Hamiltonian flow. The variational approximation is obtained similarly to the stochastic linear regression approach, but is based on a new “score matching” distance and allows for a rich family of flexible free-form approximate distributions.

### 3 Variational Hamiltonian Monte Carlo

Besides subsampling, an alternative approach that can save computation cost is to construct fast and accurate surrogate functions for the expensive potential energy functions (Liu, 2001; Neal, 2011). As one of the commonly used models for emulating expensive-to-evaluate functions, Gaussian processes (GP) are used in Rasmussen (2003) to approximate the potential energy and its derivatives based on true values of these quantities (training set) collected during an initial exploratory phase. However, a major drawback of GP-based surrogate methods is that inference time typically grows cubically in the size of training set due to the necessity of inverting a dense covariance matrix. This is especially crucial in high dimensional spaces, where large training sets are often needed before a reasonable level of approximation accuracy is achieved. Our goal, therefore, is to develop a method that can scale to large training sets while still maintaining a desired level of flexibility. For this purpose, we propose to use neural networks along with efficient training algorithms to construct surrogate functions. A typical single-hidden layer feedforward neural network (SLFN) with scalar output is defined as

$$z(\theta) = \sum_{i=1}^s v_i \sigma(\theta; \gamma_i), \quad (8)$$

where  $\gamma_i$  and  $v_i$  are the node parameter and output weight for the  $i$ th hidden neuron,  $\sigma$  is a nonlinear activation function. Given a training dataset, the estimates of those parameters can be obtained by minimizing the mean square error (MSE) cost function. To save training time, randomly assigned node parameters  $\{\gamma_i\}_{i=1}^s$  are suggested in Ferrari and Stengel (2005) and Huang et al. (2006b) where the optimization is reduced to a linear regression problem with randomly mapped features, which in turn can be solved efficiently using algebraic approaches. This approach can also be viewed as using an additive model based on random (adaptive) basis to approximate the target distribution. Despite its simple form, the random network architecture in (8) has been proven to approximate a rich class of functions arbitrarily well (Rahimi and Recht, 2008). As a result, using randomized (as opposed to optimized) networks can provide comparable approximation accuracy while dramatically reducing the required computational effort. Unlike a standard Gaussian process, the above neural network surrogate scales linearly in the size of training data, and cubically in the number of hidden neurons. This allows us to explicitly balance evaluation time and model capacity. Taking all into consideration, we use random networks in our surrogate construction.

### 3.1 Free-Form Variational Bayes

The correspondence between distributions and their potential energy functions builds a bridge between distribution approximation and function approximation. Viewed this way, each random neural network in (8) corresponds to a distribution in the exponential family

$$q_v(\theta) \propto \exp(-z(\theta)) = \exp\left[-\sum_{i=1}^s v_i \sigma(\theta; \gamma_i) - \Phi(v)\right], \quad (9)$$

where  $v = (v_i, i = 1, 2, \dots, s)$  is called the vector of canonical parameters, and the collection of randomly-mapped features  $\Psi = (\Psi_i, i = 1, 2, \dots, s)$ , where  $\Psi_i = -\sigma(\theta; \gamma_i)$  and  $i = 1, 2, \dots, s$ , is known as a sufficient statistic. The quantity  $\Phi$ , known as the log partition function, is defined by the following integral:

$$\Phi(v) = \log \int \exp(v \cdot \Psi(\theta)) d\theta.$$

Note that  $q_v(\theta)$  is properly normalized if and only if the integral is finite. Therefore, the canonical parameters  $v$  of interest belong to the set

$$\Omega := \{v \in \mathbb{R}^s \mid \Phi(v) < +\infty\}. \quad (10)$$

We call  $q_v(\theta)$  the induced distribution of neural network surrogate  $z(\theta)$  if  $v \in \Omega$ .

As our neural network surrogates approximate the true potential energy function  $U(\theta)$ , the underlying distribution  $q_v(\theta)$  approximates the target posterior distribution  $p(\theta|Y)$ . Because both the surrogate induced distribution,  $q_v(\theta)$ , and the target posterior distribution,  $p(\theta|Y)$ , are known up to a constant, we introduce the following expected squared distance between unnormalized densities

$$\tilde{D}_{SM}(p(\theta|Y)||q_v(\theta)) = \frac{1}{2} \int q_v(\theta) \|\nabla_{\theta} z(\theta) - \nabla_{\theta} U(\theta)\|^2 d\theta, \quad (11)$$

which is similar to the well known score matching squared distance (Hyvärinen, 2005) (see Section 3.3 for a more detailed explanation). By minimizing the expected squared distance  $\tilde{D}_{SM}$ , we arrive at a variational inference algorithm

$$\hat{v} = \arg \min_{v \in \Omega} \tilde{D}_{SM}(p(\theta|Y)||q_v(\theta)). \quad (12)$$

The surrogate induced approximation (9) enriches our choices for variational approximation from fixed-form tractable distributions (e.g., Gaussian or mixture of Gaussians) to fast and flexible intractable distributions. The integral in (11) is usually hard to evaluate in practice but can be approximated using samples from the surrogate induced distribution  $q_v(\theta)$  by the law of large numbers. Unlike the fixed-form approximation, the surrogate induced distribution generally does not allow for drawing samples directly. However, we can use MCMC methods to generate samples from this distribution.

Due to the random nature in approximation (9), we refer to (12) as *free-form* variational Bayes. By choosing a proper number of hidden neurons  $s$ , the free-form variational Bayes provides an implicit subsampling procedure that can effectively remove redundancy and noise in the data while striking a good balance between computation cost and approximation accuracy of the underlying distribution.

**Remark.** From an approximation point of view, each  $\sigma(\theta; \gamma)$  is a basis with a random configuration of  $\gamma$ , which specifies the orientation and scaling properties, within a profile of the activation function  $\sigma$ . In particular we choose the *softplus* function  $\sigma(\theta; w, d) = \log(1 + \exp(w \cdot \theta + d))$  to approximate the potential function, e.g., the Hamiltonian corresponding to the posterior distribution, for our free-form variational Bayes method. Note that as shown in Zhang et al. (2016), different bases can be used in our free-form variational Bayes formulation. In particular, if an exponential squared kernel is used as radial basis functions (RBF) centered at given data, the GP method is recovered. However, by using kernel functions centered at very data point, GP methods do not effectively exploit redundancy in the data or regularity in parameter space and hence may result in expensive computation cost as well as instability for large data sets. Using the exponential squared kernel at a smaller set of properly induced points leads to a more computationally tractable GP model known as “sparse GP” (Snelson and Ghahramani, 2006; Quinero-Candela and Rasmussen, 2005). As shown in Zhang et al. (2016), compared to kernel functions which are usually local, a random basis based on the softplus function provides a more compact and better approximation of potential functions with a good balance between local features and global behaviors. Also the parameter  $s$  (i.e., the number of basis functions) can be tuned to improve flexibility and accuracy while avoiding overfitting.

### 3.2 Surrogate Induced Hamiltonian Flow

To sample from the surrogate induced distribution  $q_v(\theta)$ , we generate proposals by simulating the corresponding surrogate induced Hamiltonian flow governed by the following equations

$$\frac{d\theta}{dt} = \frac{\partial \tilde{H}}{\partial r} = M^{-1}r, \quad (13)$$

$$\frac{dr}{dt} = -\frac{\partial \tilde{H}}{\partial \theta} = -\nabla_{\theta} z(\theta), \quad (14)$$

where the modified Hamiltonian is

$$\tilde{H}(\theta, r) = z(\theta) + K(r).$$

In practice, we can use the leapfrog scheme to solve (13) and (14) numerically. The end-point  $(\theta^*, r^*)$  of the trajectory starting at  $(\theta_0, r_0)$  is accepted with probability

$$\alpha_{\text{vhmc}} = \min[1, \exp(\tilde{H}(\theta_0, r_0) - \tilde{H}(\theta^*, r^*))]. \quad (15)$$

Similar to the true Hamiltonian flow, we have the following theorem

**Theorem 1.** *If we construct a Markov chain by simulating surrogate induced Hamiltonian flow (13) and (14) using leapfrog steps and Metropolis-Hastings correction with the acceptance probability according to (15), the equilibrium distribution of the chain is*

$$\tilde{\pi}(\theta, r) \propto \exp(-z(\theta) - K(r)).$$

See Supplementary Materials (Zhang et al., 2017) for a proof.

Theorem 1 implies that the marginal distribution of  $\theta$  is exactly the surrogate induced distribution  $q_v(\theta)$ . Therefore, we can run the Markov chain and collect the values of interest (e.g., the potential energy function and its derivatives) as additional training data to improve the surrogate approximation (solving (12)) on the fly. This way, our algorithm can be viewed as a generalized version of the stochastic approximation algorithm proposed in Salimans and Knowles (2013) for free-form variational Bayes.

### 3.3 Score Matching

A well-known strategy to estimate unnormalized models is score matching (Hyvärinen, 2005). Assuming that data come from a distribution with unknown density  $p_D(\cdot)$ , we want to find an approximation with a parametric unnormalized density model  $q_v(\cdot)$ , where  $v$  is an  $s$ -dimensional vector of parameters. Score matching estimates the model by minimizing the expected squared distance between the model score function  $\psi_v(\xi) = \nabla_{\xi} \log q_v(\xi)$  and the data score function  $\psi_D(\xi) = \nabla_{\xi} \log p_D(\xi)$

$$J(v) = \frac{1}{2} \int p_D(\xi) \|\psi_v(\xi) - \psi_D(\xi)\|^2 d\xi. \quad (16)$$

A simple trick was suggested in Hyvärinen (2005) to avoid the expensive estimation of the data score function  $\psi_D(\xi)$ . A similar idea can be applied here to train our neural network surrogate. Note, however, the data in this case are samples collected from the parameter space (i.e.,  $\xi = \theta$ ) for which the density of the posterior distribution is known up to some constant. The data score function then can be evaluated exactly

$\psi_D(\theta) = -\nabla_\theta U(\theta)$ . This allows us to estimate our density model  $q_v(\theta)$  without requiring samples from the posterior distribution. Since sampling from the posterior distribution is computationally costly, we sample from the simpler and cheaper surrogate induced distribution (our variational model) instead. The corresponding expected squared distance is

$$\tilde{J}(v) = \frac{1}{2} \int q_v(\theta) \|\psi_v(\theta) - \psi_D(\theta)\|^2 d\theta. \quad (17)$$

Note that the model score function is  $\psi_v(\theta) = -\nabla_\theta z(\theta)$ , and  $\tilde{J}(v)$  is exactly the expected squared distance  $\tilde{D}_{SM}$  we introduced in Section 3.1. In the case that our density model is not degenerate, we have local consistency similar to Hyvärinen (2005) as shown by the following theorem.

**Theorem 2.** *Assume that the data density  $p_D(\cdot)$  follows the model:  $p_D(\cdot) = q_v(\cdot)$  for some  $v^*$ . Further, assume that no other parameter value gives a probability density that is equal to  $q_{v^*}(\cdot)$ , and that  $q_v(\theta) > 0$  for all  $v, \theta$ . Then*

$$\tilde{J}(v) = 0 \Leftrightarrow v = v^*.$$

As such, minimizing the expected squared distance  $\tilde{D}_{SM}$  would be sufficient to estimate the model. For a proof, see Supplementary Materials.

**Remark.** Note that the data score function is

$$\psi_D(\theta) = \sum_{n=1}^N \nabla_\theta \log p(y_n|\theta) + \nabla_\theta \log p(\theta).$$

We may choose our surrogate function as

$$z(\theta) = \sum_{n=1}^N \log p(y_n|\theta) + \log p(\theta). \quad (18)$$

Then the data density  $p_D$  follows the model exactly. In order to reduce computational cost, our model is usually much simpler than (18) ( $s \ll N$ ). This allows us to explore and exploit the redundancy in the data from a function approximation perspective.

In practice, samples from the surrogate induced distribution can be collected as observations and our surrogate can be trained by minimizing the empirical version of  $\tilde{J}(v)$ . A regularization term could be included to improve numerical stability.

Now suppose that we have collected training data of size  $t$  from the Markov chain history

$$\mathcal{T}_s^{(t)} := \{(\theta_n, \nabla_\theta U(\theta_n))\}_{n=1}^t \in \mathbb{R}^d \times \mathbb{R}^d, \quad (19)$$

where  $\theta_n$  is the  $n$ -th sample. The estimator of the output weight vector (variational parameter) can be obtained by minimizing the empirical square distance between the gradient of the surrogate and the gradient of the potential (i.e., score function) plus an additional regularization term:

$$\hat{v} = \arg \min_v \frac{1}{2} \sum_{n=1}^t \|\nabla_\theta z(\theta_n) - \nabla_\theta U(\theta_n)\|^2 + \frac{\lambda}{2} \|v\|^2, \quad (20)$$

which has an online updating formula summarized in the following Proposition 1. (See Supplementary Materials for a detailed proof.) For simplicity, we use the additive nodes  $\sigma(\theta; \gamma_i) = \sigma(w_i \cdot \theta + d_i)$ <sup>1</sup>.

**Proposition 1.** *Suppose our current estimator of the output weight vector is  $v^{(t)}$  based on the current training dataset  $\mathcal{T}_s^{(t)} := \{(\theta_n, \nabla_{\theta} U(\theta_n))\}_{n=1}^t \in \mathbb{R}^d \times \mathbb{R}^d$  using  $s$  hidden neurons. Given a new training data point  $(\theta_{t+1}, \nabla_{\theta} U(\theta_{t+1}))$ , the updating formula for the estimator is given by*

$$v^{(t+1)} = v^{(t)} + W^{(t+1)}(\nabla_{\theta} U(\theta_{t+1}) - A_{t+1}v^{(t)}), \quad (21)$$

where

$$W^{(t+1)} = C^{(t)} A'_{t+1} \left[ I_d + A_{t+1} C^{(t)} A'_{t+1} \right]^{-1},$$

$$A_{t+1} = (A_1(\theta_{t+1}), \dots, A_s(\theta_{t+1})),$$

with  $A_i(\theta_{t+1}) := \sigma'(w_i \cdot \theta_{t+1} + d_i)w_i$ , and  $C^{(t)}$  can be updated by the Sherman-Morrison-Woodbury formula:

$$C^{(t+1)} = C^{(t)} - W^{(t+1)} A_{t+1} C^{(t)}. \quad (22)$$

The estimator and inverse matrix can be initialized at  $v^{(0)} = 0$ ,  $C^{(0)} = \frac{1}{\lambda} I_s$ . The online learning can be achieved by storing the inverse matrix  $C$  and performing the above updating formulas, which cost  $\mathcal{O}(d^3 + ds^2)$  computation and  $\mathcal{O}(s^2)$  storage independent of  $t$ .

### 3.4 Variational HMC in Practice

The neural network based surrogate is capable of approximating the potential energy function well when there is enough training data. However, the approximation could be poor when only few training data are available which is true in the early stage of the Markov chain simulations. To alleviate this issue, we propose to add an auxiliary regularizer which provides enough information for the sampler at the beginning and gradually diminishes as the surrogate becomes increasingly accurate. While, we use the Laplace's approximation to the potential energy function, any other fast approximation could be used in practice. The regularized surrogate approximation then takes the form

$$V_t(\theta) = \mu_t z_t(\theta) + \frac{1}{2}(1 - \mu_t)(\theta - \theta^L)' \nabla_{\theta}^2 U(\theta^L)(\theta - \theta^L),$$

where  $\theta^L$  is the maximum *a posteriori* (MAP) estimate and  $\mu_t \in [0, 1]$  is a smooth monotone function monitoring the transition from the Laplace's approximation to the surrogate approximation. Refining the surrogate approximation by acquiring training data from simulating the regularized surrogate induced Hamiltonian flow, we arrive

---

<sup>1</sup>Note that our method works for a general class of nonlinear nodes, including the radial basis functions (RBF)

**Algorithm 2** Variational Hamiltonian Monte Carlo

---

**Input:** Regularization coefficient  $\lambda$ , transition function  $\mu_t$ , number of hidden neurons  $s$ , starting position  $\theta^{(1)}$  and HMC parameters  
*Find the Maximum A Posterior  $\theta^L$  and compute the Hessian matrix  $\nabla_{\theta}^2 U(\theta^L)$*   
*Randomly assign the node parameters:  $\{\gamma_i\}_{i=1}^s$*   
**for**  $t = 1, 2, \dots, T$  **do**  
    Propose  $(\theta^*, r^*)$  with regularized surrogate induced Hamiltonian flow, using  $\nabla_{\theta} V_t(\theta)$   
  
    Perform Metropolis-Hasting step according to the underlying distribution  $\pi_t \sim \exp(-V_t(\theta) - K(r))$   
    **if** New state is accepted &  $t < t_0$  **then**  
        Acquire new training data point  $(\theta_{t+1}, \nabla_{\theta} U(\theta_{t+1}))$   
        Update the output weight estimate  $v^{(t+1)} \leftarrow$  (21) and the inverse matrix  $C^{(t+1)} \leftarrow$  (22)  
    **else**  
         $v^{(t+1)} = v^{(t)}, C^{(t+1)} = C^{(t)}$   
    **end if**  
**end for**

---

at an efficient approximate inference method: Variational Hamiltonian Monte Carlo (VHMC; Algorithm 2).

In practice, the surrogate approximation may achieve a sufficient level of accuracy so by early stopping we could avoid unnecessary updating of the output weights. In fact, the stopping time  $t_0$  serves as a knob to control the desired approximation quality. Before stopping, VHMC acts as a free-form variational Bayes method that keeps improving itself by collecting training data from the Markov chain. After stopping, VHMC performs as a standard HMC algorithm which samples from the surrogate induced distribution. VHMC successfully combines the advantages of both variational Bayes and Hamiltonian Monte Carlo, resulting in higher computational efficiency compared to HMC and better approximation compared to VB.

## 4 Experiments

In this section, four experiments are conducted to verify the effectiveness and efficiency of the proposed Variational HMC framework. In the first two examples, we demonstrate the performance of VHMC from a pure approximation perspective and compare it to state-of-the-art fixed-form variational Bayes methods. We then test the efficiency of VHMC on two machine learning problems with large datasets and compare its performance to that of SGLD (Welling and Teh, 2011), which is one of the state-of-the-art stochastic gradient MCMC methods. Note that compared to standard HMC, our proposed method introduces some additional tuning parameters, such as the number of random bases  $s$  and the transition monitor  $\mu_t$ . As a general guideline, we want our method to be as efficient as possible while maintaining good approximation. Therefore, we keep  $s$  small as long as a reasonable level of accuracy can be achieved; we also keep

$\mu_t$  relatively small to ensure enough data are collected for the training procedure. In our experiments, we found that  $\mu_t = 1 - \exp(-t/n_s)$  is a reasonable schedule where  $n_s$  can be adjusted to accommodate different transition speeds according to the complexity of the model. In practice, we can find the right setting by monitoring the acceptance rate using an initial run as in Zhang et al. (2016).

#### 4.1 A Beta-Binomial Model for Overdispersion

We first demonstrate the performance of our variational Hamiltonian Monte Carlo method on a simple example, which considers the problem of estimating the rates of death from stomach cancer for the largest cities in Missouri (Albert, 2009). The data consist of 20 pairs  $(n_j, y_j)$ , where  $n_j$  records the number of individuals that were at risk for cancer in city  $j$ , and  $y_j$  is the number of cancer deaths that occurred in that city. The counts  $y_j$  are overdispersed compared to what would be expected under a binomial model with a constant probability; therefore, Albert (2009) assumes a beta-binomial model with mean  $m$  and precision  $K$ :

$$p(y_j|m, K) = \binom{n_j}{y_j} \frac{B(Km + y_j, K(1 - m) + n_j - y_j)}{B(Km, K(1 - m))},$$

and assume the following improper prior:

$$p(m, K) \propto \frac{1}{m(1 - m)} \frac{1}{(1 + K)^2}.$$

The resulting posterior is extremely skewed (as shown in the bottom right corner in Figure 1); hence, the following reparameterization:

$$x_1 = \text{logit}(m), \quad x_2 = \text{logit}(K),$$

is used to ameliorate this issue.

We choose  $\mu_t = 1 - \exp(-t/200)$  as our transition schedule and set up the tuning parameters of HMC to achieve around 85% acceptance rate. We run the variational Hamiltonian Monte Carlo long enough so that we can estimate the full approximation quality of our surrogate. To demonstrate the approximation performance in terms of the number of hidden neurons, we train the neural network surrogate using different values of  $s$  and examine the resulting KL-divergence and score matching squared distance to the true posterior density. As we can see from Figures 1 and 2, the neural network surrogate model indeed offers a high quality approximation and becomes more accurate as  $s$  increases. The surrogate induced Hamiltonian flow effectively explores the parameter space and transfers information from the posterior to the surrogate.

#### 4.2 Bayesian Probit Regression

Next, we demonstrate the approximation performance of our Variational HMC algorithm relative to existing variational approaches on a simple Bayesian classification problem using binary probit regression. Given  $N$  observed data pairs,

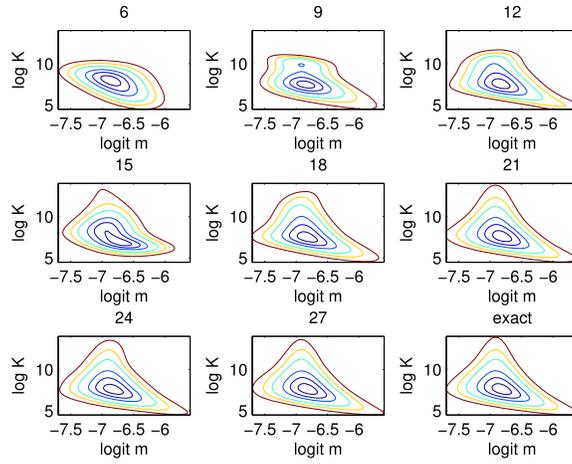


Figure 1: Approximate posteriors for a varying number of hidden neurons. Exact posterior at bottom right.

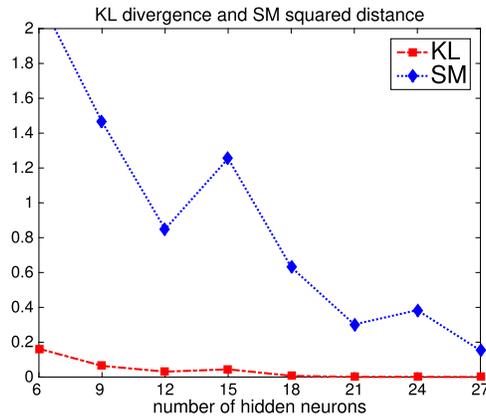


Figure 2: KL-divergence and score matching squared distance between the surrogate approximation and the exact posterior density using an increasing number of hidden neurons.

$$\{(y_n, x_n) | y_n \in \{0, 1\}, x_n \in \mathbb{R}^d\}_{n=1}^N,$$

we assume a probit likelihood function  $P(y_n = 1 | \theta) = \Phi(\theta^T x_n)$  and a Gaussian prior  $\mathcal{N}(0, 100)$  on regression parameters; here,  $\Phi$  is the standard Gaussian cdf. A full covariance multivariate normal approximation is used for all variational approaches. The synthetic data we use are simulated from the model, with  $N = 10000$  and  $d = 5$ . We compare our algorithm to Variational Bayesian Expectation Maximization (VBEM) (Beal and Ghahramani, 2002; Ormerod and Wand, 2010), and the best fixed-form variational

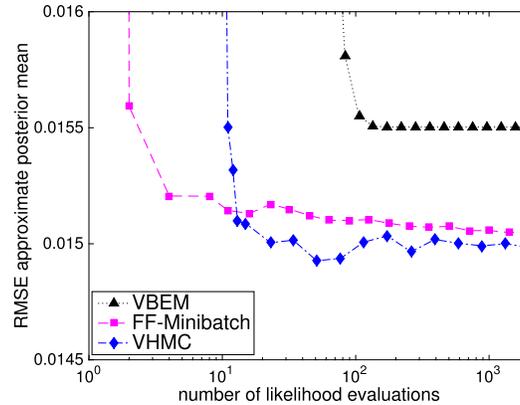


Figure 3: RMSE approximate posterior mean as a function of the number of likelihood evaluations for difference variational Bayesian approaches and our Variational HMC algorithm.

approximation of Salimans and Knowles (2013) that uses Hessian and subsampling (FF-Minibatch). VBEM and FF-Minibatch are both implemented using the code provided in Salimans and Knowles (2013) with the same learning rate and minibatch size. For all variational approaches, we initialize the posterior approximation to the prior. For our Variational HMC algorithm, we choose  $s = 100$  random hidden units for the surrogate and set the starting point to be the origin. The number of hidden units is chosen so that the surrogate is flexible enough to fit the target well and remain fast in computation. The HMC parameters are set to make the acceptance probability around 70%. The target density is almost Gaussian, and a fast transition  $\mu_t = 1 - \exp(-t/5)$  is enough to stabilize our algorithm. Since this experiment is based on synthetic data, we follow Salimans and Knowles (2013) to assess the approximation performance in terms of the root mean squared error (RMSE) between the estimate (variational mean for VB and sample mean for VHMC) and the true parameter that is used to generate the dataset.

Figure 3 shows the performance averaged over 50 runs for the three methods. As we can see from the graph, VHMC and the subsampling based fixed-form variational approach (FF-minibatch) achieve lower RMSE than the VBEM algorithm. This is due to the extra factorization assumptions made by VBEM when introducing the auxiliary variables (Ormerod and Wand, 2010). Even though Gaussian approximation is already sufficiently accurate on this simple example, VHMC can still arrive at a lower RMSE due to the extra flexibility provided by the free-form neural network surrogate function.

### 4.3 Bayesian Logistic Regression

Next, we test the efficiency of our Variational HMC method as a scalable sampling method by applying VHMC to a Bayesian logistic regression model. Given the  $i$ -th input vector  $x_i$ , the corresponding output (label)  $y_i = \{0, 1\}$  is assumed to follow the

probability  $p(y_i = 1|x_i, \beta) = 1/(1 + \exp(-x_i^T \beta))$  and a Gaussian prior  $\mathcal{N}(0, 100)$  is used for the model parameters  $\beta$ . We test our proposed algorithm on the `a9a` dataset (Lin et al., 2008). The original dataset, which is compiled from the UCI `adult` dataset, has 32561 observations and 123 features. We reduce the dimension to 50 by random projection of the original features. We choose  $s = 2000$  hidden units for the surrogate and set a transition schedule  $\mu_t = 1 - \exp(-t/500)$  for our VHMC algorithm. We then compare the algorithm to HMC (Duane et al., 1987; Neal, 2011) and to stochastic gradient Langevin dynamics (SGLD) (Welling and Teh, 2011). For HMC and VHMC, we set the leapfrog stepsize such that the acceptance rate is around 70%. For SGLD we choose batch size of 500 and use a range of fixed stepsizes.

To illustrate the sampling efficiency of all methods, we follow Ahn et al. (2012) to investigate the time normalized effective sample size (ESS)<sup>2</sup> averaged over the 51 parameters and compare this with the relative error after a fixed amount of computation time. The relative error of mean (REM) and relative error of covariance (REC) is defined as

$$\text{REM}_t = \frac{\sum_i |\bar{\beta}_i^t - \beta_i^o|}{\sum_i |\beta_i^o|}, \quad \text{REC}_t = \frac{\sum_{ij} |C_{ij}^t - C_{ij}^o|}{\sum_{ij} |C_{ij}^o|}, \quad (23)$$

where  $\bar{\beta}^t = \frac{1}{t} \sum_{t'=1}^t \beta_{t'}$ ,  $C^t = \frac{1}{t} \sum_{t'=1}^t (\beta_{t'} - \bar{\beta}^t)(\beta_{t'} - \bar{\beta}^t)^T$  are the sample mean and sample covariance up to time  $t$ , and the ground truth  $\beta^o$  and  $C^o$  are obtained using a long run ( $T = 500\text{K}$  samples) of HMC algorithm.

Figure 4 shows the relative error at time  $T = 300$ ,  $T = 3000$  as a function of the time normalized mean ESS, which is a measure of the mixing rate. The results for the mean are shown on the top, and those for the covariance are on the bottom. We run each algorithm with a different setting of parameters that control the mixing rate: number of leapfrog steps  $L = [50, 40, 30, 20, 10, 5, 1]$  for HMC and  $L = [50, 40, 30, 20, 10, 5]$  for VHMC, and stepsizes  $\epsilon = [2e-3, 1e-3, 5e-4, 1e-4, 5e-5, 1e-5]$  for SGLD.

As we decrease the stepsize, SGLD becomes less biased in the gradient approximation, resulting in smaller relative error. However, the exploration efficiency drops at the same time and sampling variance gradually dominates the relative error. In contrast, HMC uses a fixed leapfrog stepsize and therefore maintains high exploration efficiency in parameter space. The downside is the expensive computation of the full gradient and the possible turning back of the trajectories when the number of leapfrog steps is unnecessarily large. Adopting a flexible neural network surrogate, VHMC balances the computation cost and approximation quality much better than subsampling and achieves lower relative error with high mixing rates.

#### 4.4 Independent Component Analysis

Finally, we apply our method to sample from the posterior distribution of the unmixing matrix in Independent Component Analysis (ICA) (Hyvärinen and Oja, 2000). Given  $N$   $d$ -dimensional observations  $X = \{x_n \in \mathbb{R}^d\}_{n=1}^N$ , we model the data as

<sup>2</sup>Given  $B$  samples,  $\text{ESS} = B[1 + 2 \sum_{k=1}^K \gamma(k)]^{-1}$ , where  $\gamma(k)$  is the sample autocorrelation at lag  $k$  (Geyer, 1992)

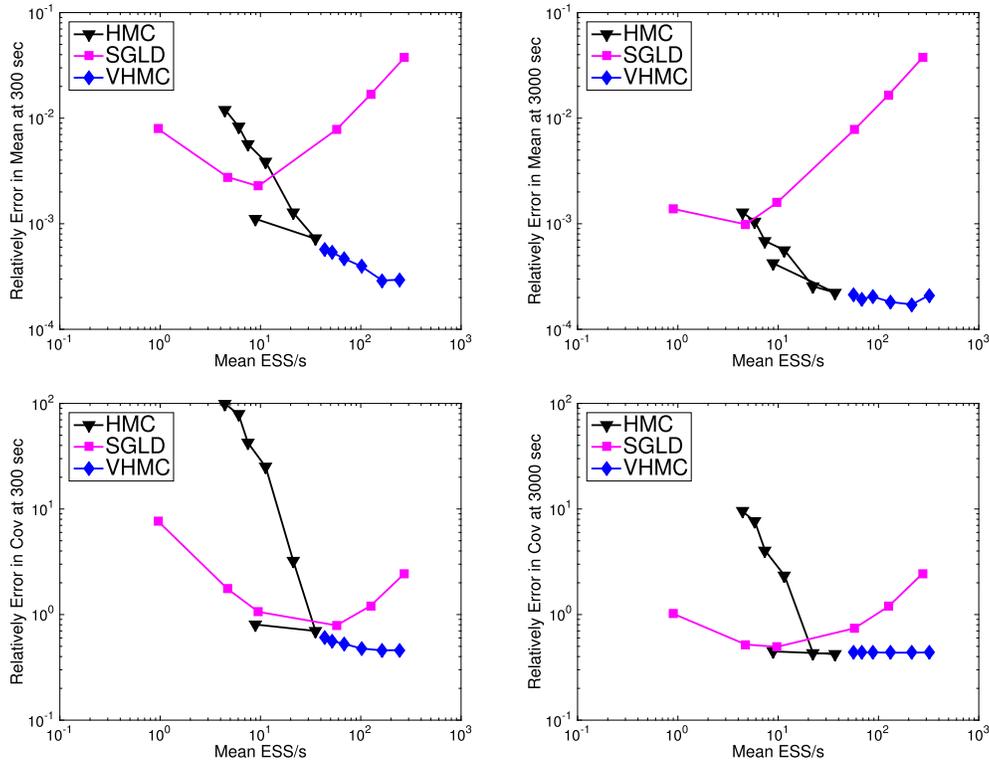


Figure 4: Final error of logistic regression at time  $T$  versus mixing rate for the mean (top) and covariance (bottom) estimates after 300 (left) and 3000 (right) seconds of computation. Each algorithm is run using different setting of parameters.

$$p(x|W) = |\det(W)| \prod_{i=1}^d p_i(w_i^T x), \quad (24)$$

where  $w_i$  is the  $i$ -th row of  $W$  and  $p_i$  is used to capture the true density of the  $i$ -th independent component. Following Welling and Teh (2011), we use a Gaussian prior over the unmixing matrix  $p(w_{ij}) = \mathcal{N}(0, \sigma)$  and assume  $p_i(y_i) = [4 \cosh(\frac{1}{2}y_i)]^{-1}$  with  $y_i = w_i^T x$ . We evaluate our method using the MEG dataset (Vigário et al., 1997), which has 122 channels and 17730 observations. We extract the first 5 channels for our experiment which leads to samples with 25 dimensions. We then compare our algorithm to standard HMC and SGLD (Welling and Teh, 2011). For SGLD, we use a natural gradient (Amari et al., 1996) which has been found to improve the efficiency of gradient descent significantly. We set  $\sigma = 100$  for the Gaussian prior. For HMC and Variational HMC, we set the leapfrog stepsize such that the acceptance ratio remains around 70% and set  $L = 40$  to allow an efficient exploration of the parameter space. For SGLD, we choose batch size of 500 and use stepsizes from a polynomial annealing schedule  $a(b+t)^{-\delta}$ , with  $a = 5 \times 10^{-3}$ ,  $b = 10^{-4}$  and  $\delta = 0.5$ . (This setting reduces the stepsize

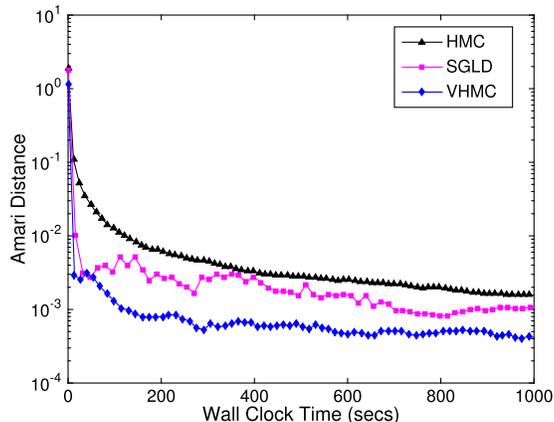


Figure 5: Convergence of Amari distance on the MEG data for HMC, SGLD and our Variational HMC algorithm.

from  $5 \times 10^{-5}$  to  $1 \times 10^{-6}$  during  $1e+7$  iterations.) For VHMC, we choose  $s = 1000$  hidden units and set the transition schedule  $\mu_t = 1 - \exp(-t/2000)$ . Note that the likelihood (24) is row-permutation invariant. To measure the convergence of all samplers on this ICA problem, we therefore use the Amari distance (Amari et al., 1996)  $d_A(\overline{W}, W_0)$ , where  $\overline{W}$  is the sample average and  $W_0$  is the true unmixing matrix estimated using a long run ( $T = 100K$  samples) of standard HMC algorithm.

In Figure 5, the Amari distance as a function of runtime is shown for each method. From the graph, we can see that SGLD converges faster than standard HMC. The noise introduced by subsampling is compensated by the fast exploration of the parameter space; this allows for an early arrival at the high probability domain for SGLD. However, when the variance dominates the bias, the convergence speed slows down since SGLD requires annealing (or small) stepsize that inevitably leads to low exploration efficiency. By maintaining efficient parameter exploration (same stepsize as HMC) while reducing the computation in simulating the Hamiltonian flow, VHMC outperforms SGLD by arriving at a lower Amari distance much more rapidly.

## 5 Conclusion

We have presented a novel approach, Variational Hamiltonian Monte Carlo, for approximate Bayesian inference. Our approach builds on the framework of HMC, but uses a flexible and efficient neural network surrogate function to approximate expensive gradient evaluations. The surrogate keeps refining its approximation by collecting training data while the sampler is exploring the parameter space. This way, our algorithm can be viewed as a free-form variational approach. Unlike subsampling-based MCMC methods, VHMC maintains the relatively high exploration efficiency of its MCMC counterparts while reducing the computation cost. Compared to fixed-form variational approximation, VHMC is more flexible and thus can approximate posterior distributions better.

As the complexity of the model increases (e.g., high dimensional models), the computational cost of our random network surrogate method would increase since more random bases might be needed for better approximation. In many cases, however, the free-style surrogate construction in VPMC can be adjusted to reduce the computation cost. To this end, we could exploit the underlying structure of a model to construct a more effective surrogate model. For example, if the structure results in conditional independence, we could use the corresponding graphical model to build a more structured surrogate function. Alternatively, we could train the full neural network with sparsity constraints imposed on the input weights.

## Supplementary Material

Variational Hamiltonian Monte Carlo via Score Matching – Supplementary Materials (DOI: [10.1214/17-BA1060SUPP](https://doi.org/10.1214/17-BA1060SUPP); .pdf).

## References

- Ahn, S., Korattikara, A., and Welling, M. (2012). “Bayesian Posterior Sampling via Stochastic Gradient Fisher Scoring.” In Langford, J. and Pineau, J. (eds.), *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, 1591–1598. New York, NY, USA: Omnipress. [486](#), [500](#)
- Albert, J. (2009). *Bayesian Computing with R*. Springer Science, New York. [MR2541616](#). [497](#)
- Amari, S. I., Cichocki, A., and Yang, H. H. (1996). “A New Learning Algorithm for Blind Signal Separation.” In Touretzky, D. S., Mozer, M. C., and Hasselmo, M. E. (eds.), *Advances in Neural Information Processing Systems 8*, 757–763. MIT Press. [501](#), [502](#)
- Beal, M. J. and Ghahramani, Z. (2002). “The variational Bayesian EM algorithm for incomplete data: with application to scoring graphical model structures.” In Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M. (eds.), *Bayesian Statistics 7: Proceedings of the 7th Valencia International Meeting*, 453–463. Oxford University Press, Oxford. [MR2003189](#). [498](#)
- Betancourt, M. (2015). “The Fundamental Incompatibility of Scalable Hamiltonian Monte Carlo and Naive Data Subsampling.” In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, 533–540. Lille, France: PMLR. [486](#)
- Chen, T., Fox, E., and Guestrin, C. (2014). “Stochastic Gradient Hamiltonian Monte Carlo.” In Xing, E. P. and Jebara, T. (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32, 1683–1691. Beijing, China: PMLR. [486](#)
- de Freitas, N., Højten-Sørensen, P., Jordan, M. I., and Russell, S. (2001). “Variational MCMC.” In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial*

- Intelligence*, 120–127. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. 486
- Ding, N., Fang, Y., Babbush, R., Chen, C., Skeel, R. D., and Neven, H. (2014). “Bayesian Sampling Using Stochastic Gradient Thermostats.” In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, 3203–3211. Curran Associates, Inc. 486
- Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). “Hybrid Monte Carlo.” *Physics Letters B*, 195(2): 216–222. 486, 488, 500
- Ferrari, S. and Stengel, R. F. (2005). “Smooth function approximation using neural networks.” *IEEE Transactions on Neural Networks*, 16(1): 24–38. 491
- Geyer, C. J. (1992). “Practical Markov Chain Monte Carlo.” *Statistical Science*, 7: 473–483. 500
- Girolami, M. and Calderhead, B. (2011). “Riemann manifold Langevin and Hamiltonian Monte Carlo methods (with discussion).” *Journal of the Royal Statistical Society*, 73(2): 123–214. MR2814492. doi: <https://doi.org/10.1111/j.1467-9868.2010.00765.x>. 488
- Hoffman, M. D. and Gelman, A. (2014). “The No-U-turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo.” *Journal of Machine Learning Research*, 15(1): 1593–1623. MR3214779. 488
- Honkela, A., Raiko, T., Kuusela, M., Tornio, M., and Karhunen, J. (2010). “Approximate Riemannian conjugate gradient learning for fixed-form variational Bayes.” *Journal of Machine Learning Research*, 11: 3235–3268. MR2746551. 489
- Huang, G. B., Chen, L., and Siew, C. K. (2006a). “Universal approximation using incremental constructive feedforward networks with random hidden nodes.” *IEEE Transactions on Neural Networks*, 17(4): 879–892. 486
- Huang, G. B., Zhu, Q. Y., and Siew, C. K. (2006b). “Extreme learning machine: Theory and applications.” *Neurocomputing*, 70(1–3): 489–501. 491
- Hyvärinen, A. (2005). “Estimation of non-normalized statistical models by score matching.” *Journal of Machine Learning Research*, 6: 695–709. MR2249836. 487, 492, 493, 494
- Hyvärinen, A. and Oja, E. (2000). “Independent component analysis: algorithms and applications.” *Neural Networks*, 13: 411–430. 500
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999). “An Introduction to Variational Methods for Graphical Models.” In *Machine Learning*, 183–233. MIT Press. 486
- Kingma, D. P. and Welling, M. (2013). “Auto-Encoding Variational Bayes.” In *The 2nd International Conference on Learning Representations (ICLR)*. 490

- Lin, C. J., Weng, R. C., and Keerthi, S. S. (2008). “Trust region Newton method for large-scale logistic regression.” *Journal of Machine Learning Research*, 9: 627–650. MR2417250. doi: <https://doi.org/10.1145/1273496.1273567>. 500
- Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer. MR1842342. 490
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). “Equation of State Calculations by Fast Computing Machines.” *The Journal of Chemical Physics*, 21(6): 1087–1092. 485
- Neal, R. M. (2011). “MCMC using Hamiltonian dynamics.” In Brooks, S., Gelman, A., Jones, G., and Meng, X. L. (eds.), *Handbook of Markov Chain Monte Carlo*, 113–162. Chapman and Hall/CRC. MR2858447. 486, 488, 490, 500
- Ormerod, J. T. and Wand, M. P. (2010). “Explaining Variational Approximations.” *The American Statistician*, 2(64): 140–153. MR2757005. doi: <https://doi.org/10.1198/tast.2010.09058>. 498, 499
- Paisley, J., Blei, D., and Jordan, M. (2012). “Variational Bayesian Inference with Stochastic Search.” In Langford, J. and Pineau, J. (eds.), *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, 1367–1374. New York, NY, USA: Omnipress. 490
- Quinonero-Candela, J. and Rasmussen, C. E. (2005). “A unifying view of sparse approximate Gaussian process regression.” *Journal of Machine Learning Research*, 6: 1939–1959. MR2249877. 492
- Rahimi, A. and Recht, B. (2008). “Uniform approximation of functions with random bases.” In *46th Annual Allerton Conference on Communication, Control, and Computing*, 555–561. 491
- Ranganath, R., Gerrish, S., and Blei, D. M. (2014). “Black Box Variational Inference.” In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics, AISTATS 2014, Reykjavik, Iceland*, 814–822. 490
- Rasmussen, C. E. (2003). “Gaussian Processes to Speed up Hybrid Monte Carlo for Expensive Bayesian Integrals.” *Bayesian Statistics*, 7: 651–659. MR2003529. 490
- Salimans, T., Kingma, D., and Welling, M. (2015). “Markov Chain Monte Carlo and Variational Inference: Bridging the Gap.” In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, 1218–1226. PMLR. 486
- Salimans, T. and Knowles, D. A. (2013). “Fixed-form variational posterior approximation through stochastic linear regression.” *Bayesian Analysis*, 8(4): 837–882. MR3150471. doi: <https://doi.org/10.1214/13-BA858>. 489, 490, 493, 499
- Saul, L. and Jordan, M. I. (1996). “Exploiting tractable substructures in intractable networks.” In Tesauro, G., Touretzky, D. S., and Leen, T. K. (eds.), *Advance in neural information processing systems 7 (NIPS 1996)*, 486–492. Cambridge, MA: MIT Press. 489

- Snelson, E. and Ghahramani, Z. (2006). “Sparse Gaussian Processes using Pseudo-inputs.” In Weiss, Y., Schölkopf, P. B., and Platt, J. C. (eds.), *Advances in Neural Information Processing Systems 18*, 1257–1264. MIT Press. [MR2441315](#). 492
- Vigário, R., Särelä, J., and Oja, E. (1997). “MEG data for studies using independent component analysis.” [http://research.ics.aalto.fi/ica/eegmeg/MEG\\_data.html](http://research.ics.aalto.fi/ica/eegmeg/MEG_data.html). 501
- Wainwright, M. and Jordan, M. (2008). “Graphical models, exponential families, and variational inference.” *Foundations and Trends in Machine Learning*, 1(1–2): 1–305. 486
- Wang, Z., Mohamed, S., and Freitas, N. (2013). “Adaptive Hamiltonian and Riemann Manifold Monte Carlo.” In Dasgupta, S. and McAllester, D. (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28, 1462–1470. Atlanta, Georgia, USA: PMLR. 488
- Welling, M. and Teh, Y. W. (2011). “Bayesian Learning via Stochastic Gradient Langevin Dynamics.” In Getoor, L. and Scheffer, T. (eds.), *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 681–688. New York, NY, USA: ACM. [MR3157685](#). doi: <https://doi.org/10.4310/CIS.2012.v12.n3.a3>. 486, 496, 500, 501
- Zhang, C., Shahbaba, B., and Zhao, H. (2016). “Hamiltonian Monte Carlo acceleration using surrogate functions with random bases.” *Statistics and Computing*, 1–18. 492, 497
- Zhang, C., Shahbaba, B., and Zhao, H. (2017). “Variational Hamiltonian Monte Carlo via Score Matching – Supplementary Materials.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/17-BA1060SUPP>. 493

#### Acknowledgments

This work is supported by NIH grant R01-AI107034 and NSF grants DMS-1622490 and DMS-1418422.