

# META-LEARNING FOR STOCHASTIC GRADIENT MCMC

Wenbo Gong<sup>1\*</sup>, Yingzhen Li<sup>2\*†</sup>, José Miguel Hernández-Lobato<sup>1,2</sup>

<sup>1</sup>University of Cambridge <sup>2</sup>Microsoft Research Cambridge

{wg242, jmh233}@cam.ac.uk, Yingzhen.Li@microsoft.com

## ABSTRACT

Stochastic gradient Markov chain Monte Carlo (SG-MCMC) has become increasingly popular for simulating posterior samples in large-scale Bayesian modeling. However, existing SG-MCMC schemes are not tailored to any specific probabilistic model, even a simple modification of the underlying dynamical system requires significant physical intuition. This paper presents the first meta-learning algorithm that allows automated design for the underlying continuous dynamics of an SG-MCMC sampler. The learned sampler generalizes Hamiltonian dynamics with state-dependent drift and diffusion, enabling fast traversal and efficient exploration of energy landscapes. Experiments validate the proposed approach on learning tasks with Bayesian fully connected neural networks, Bayesian convolutional neural networks and Bayesian recurrent neural networks, showing that the learned sampler out-performs generic, hand-designed SG-MCMC algorithms, and generalizes to different datasets and larger architectures.

## 1 INTRODUCTION

There is a resurgence of research interests in Bayesian deep learning (Graves, 2011; Blundell et al., 2015; Hernández-Lobato & Adams, 2015; Hernandez-Lobato et al., 2016; Gal & Ghahramani, 2016; Ritter et al., 2018), which applies Bayesian inference to neural networks for better uncertainty estimation. It is crucial for e.g. better exploration in reinforcement learning (Deisenroth & Rasmussen, 2011; Depeweg et al., 2017), resisting adversarial attacks (Feinman et al., 2017; Li & Gal, 2017; Louizos & Welling, 2017) and continual learning (Nguyen et al., 2018). A popular approach to performing Bayesian inference on neural networks is stochastic gradient Markov chain Monte Carlo (SG-MCMC), which adds properly scaled Gaussian noise to a stochastic gradient ascent procedure (Welling & Teh, 2011). Recent advances in this area further introduced optimization techniques such as pre-conditioning (Ahn et al., 2012; Patterson & Teh, 2013), annealing (Ding et al., 2014) and adaptive learning rates (Li et al., 2016a; Chen et al., 2016). All these efforts have made SG-MCMC highly scalable to many deep learning tasks, including shape and texture modeling in computer vision (Li et al., 2016b) and language modeling with recurrent neural networks (Gan et al., 2017). However, inventing novel dynamics for SG-MCMC requires significant mathematical work to ensure the sampler’s stationary distribution is the target distribution, which is less friendly to practitioners. Furthermore, many of these algorithms are designed as a generic sampling procedure, and the associated physical mechanism might not be best suited for sampling neural network weights.

This paper aims to automate the SG-MCMC proposal design by introducing *meta-learning* techniques (Schmidhuber, 1987; Bengio et al., 1992; Naik & Mammone, 1992; Thrun & Pratt, 1998). The general idea is to train a *learner* on one or multiple tasks in order to acquire common knowledge that generalizes to future tasks. Recent applications of meta-learning include learning to transfer knowledge to unseen few-shot learning tasks (Santoro et al., 2016; Ravi & Larochelle, 2017; Finn et al., 2017), and learning algorithms such as gradient descent (Andrychowicz et al., 2016; Li & Malik, 2017; Wichrowska et al., 2017), Bayesian optimization (Chen et al., 2017) and reinforcement learning (Duan et al., 2016; Wang et al., 2016). Unfortunately, these advances cannot be directly transferred to the world of MCMC samplers, as a naive neural network parameterization of the transition kernel does not guarantee the posterior distribution to be the stationary distribution of the sampler.

\*Equal contribution

†Work done at the University of Cambridge

We present to the best of our knowledge the first attempt towards meta-learning an SG-MCMC algorithm. Concretely, our contribution includes:

An SG-MCMC sampler that extends Hamiltonian dynamics with *learnable* diffusion and curl matrices. Once trained, the sampler can generalize to different datasets and architectures. Extensive evaluation of the proposed sampler on Bayesian fully connected neural networks, Bayesian convolutional neural networks and Bayesian recurrent neural networks, with comparisons to popular SG-MCMC schemes based on e.g. Hamiltonian Monte Carlo (Chen et al., 2014) and pre-conditioned Langevin dynamics (Li et al., 2016a).

## 2 BACKGROUND: A COMPLETE FRAMEWORK FOR SG-MCMC

Consider sampling from a target density  $\pi(\boldsymbol{\theta})$  that is defined by an *energy function*:  $U(\boldsymbol{\theta})$ ,  $\boldsymbol{\theta} \in \mathbb{R}^D$ ,  $\pi(\boldsymbol{\theta}) \propto \exp(-U(\boldsymbol{\theta}))$ . In this paper, we focus on this sampling task with a Bayesian modeling set-up, i.e. given observed data  $D = \{ \mathbf{o}_n \}_{n=1}^N$ , we define a probabilistic model  $p(D, \boldsymbol{\theta}) = \prod_{n=1}^N p(\mathbf{o}_n | \boldsymbol{\theta}) p(\boldsymbol{\theta})$ , and we want samples from the target density defined as *posterior distribution*  $\pi(\boldsymbol{\theta}) = p(\boldsymbol{\theta} | D)$ . We use Bayesian neural networks as an illustrating example, in this case,  $\mathbf{o}_n = (\mathbf{x}_n, \mathbf{y}_n)$ , the prior  $p(\boldsymbol{\theta})$  is a Gaussian  $\mathcal{N}(\boldsymbol{\theta}; \mathbf{0}, \lambda^{-1} \mathbf{I})$ , and the energy function is defined as

$$U(\boldsymbol{\theta}) = \sum_{n=1}^N \log p(\mathbf{y}_n | \mathbf{x}_n, \boldsymbol{\theta}) \quad \log p(\boldsymbol{\theta}) = \sum_{n=1}^N \ell(\mathbf{y}_n, \text{NN}(\mathbf{x}_n)) + \lambda \|\boldsymbol{\theta}\|_2^2, \quad (1)$$

with  $\ell(\mathbf{y}, \hat{\mathbf{y}})$  usually defined as the  $\ell_2$  loss for regression or the cross-entropy loss for classification. A typical MCMC sampler constructs a Markov chain with a *transition kernel*, and corrects the proposed samples with Metropolis-Hastings (MH) rejection steps. Some of these methods, e.g. Hamiltonian Monte Carlo (HMC) (Duane et al., 1987; Neal et al., 2011), further augment the state space as  $\mathbf{z} = (\boldsymbol{\theta}, \mathbf{r})$  with auxiliary variables  $\mathbf{r}$ , and sample from the augmented distribution  $\pi(\mathbf{z}) \propto \exp(-H(\mathbf{z}))$ , with the *Hamiltonian*  $H(\mathbf{z}) = U(\boldsymbol{\theta}) + g(\mathbf{r})$  such that  $\int \exp(-g(\mathbf{r})) d\mathbf{r} = C$ . Thus, marginalizing out the auxiliary variable  $\mathbf{r}$  will not affect the stationary distribution  $\pi(\boldsymbol{\theta}) \propto \exp(-U(\boldsymbol{\theta}))$ .

For deep learning tasks, the observed dataset  $D$  often contains thousands, if not millions, of instances, making MH rejection steps computationally prohibitive. Fortunately this is mitigated by SG-MCMC, whose transition kernel is implicitly defined by a stochastic differential equation (SDE) that leaves the target density invariant (Welling & Teh, 2011; Ahn et al., 2012; Patterson & Teh, 2013; Chen et al., 2014; Ding et al., 2014). Such a Markov process is called *Itô diffusion* governed by the continuous-time SDEs:

$$d\mathbf{z} = \mathbf{f}(\mathbf{z}) dt + \sqrt{2\mathbf{D}(\mathbf{z})} d\mathbf{W}(t), \quad (2)$$

with  $\mathbf{f}(\mathbf{z})$  the deterministic *drift*,  $\mathbf{W}(t)$  the Wiener process, and  $\mathbf{D}(\mathbf{z})$  the *diffusion* matrix. As a simple example, Langevin dynamics considers  $\mathbf{z} = \boldsymbol{\theta}$ ,  $\mathbf{f}(\boldsymbol{\theta}) = -\nabla U(\boldsymbol{\theta})$  and  $\mathbf{D}(\boldsymbol{\theta}) = \mathbf{I}$ . Then using *forward Euler discretization* with step-size  $\eta$  the update rule of the parameters is

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \nabla U(\boldsymbol{\theta}_t) + \sqrt{2\eta} \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (3)$$

Stochastic gradient Langevin dynamics (SGLD, Welling & Teh, 2011) proposed an approximation to (3), by replacing the exact gradient  $-\nabla U(\boldsymbol{\theta})$  with an estimate using a mini-batch of datapoints:

$$-\nabla U(\boldsymbol{\theta}) \approx \frac{N}{M} \sum_{m=1}^M \nabla \log p(\mathbf{o}_m | \boldsymbol{\theta}) - \nabla \log p(\boldsymbol{\theta}), \quad \mathbf{o}_1, \dots, \mathbf{o}_M \sim D. \quad (4)$$

Therefore SGLD can be viewed as a stochastic gradient descent (SGD) that adds in a properly scaled Gaussian noise term. Similarly, SGHMC (Chen et al., 2014) is closely related to momentum SGD (see appendix). Furthermore, MH rejection steps are usually dropped in SG-MCMC when a carefully selected discretization step-size is in use. Therefore SG-MCMC has the same computational complexity as many stochastic optimization algorithms, making it highly scalable for sampling posterior distributions of neural network weights conditioned on big datasets.

Ma et al. (2015) derived a framework of SG-MCMC samplers using advanced statistical mechanics (Yin & Ao, 2006; Shi et al., 2012), which explicitly parameterizes the drift  $\mathbf{f}(\mathbf{z})$ :

$$\mathbf{f}(\mathbf{z}) = -[\mathbf{D}(\mathbf{z}) + \mathbf{Q}(\mathbf{z})] \nabla_{\mathbf{z}} H(\mathbf{z}) + \mathbf{z}, \quad \mathbf{i}(\mathbf{z}) = \sum_{j=1}^d \frac{\partial}{\partial \mathbf{z}_j} (\mathbf{D}_{ij}(\mathbf{z}) + \mathbf{Q}_{ij}(\mathbf{z})), \quad (5)$$

with  $\mathbf{Q}(\mathbf{z})$  the curl matrix,  $\mathbf{D}(\mathbf{z})$  the diffusion matrix and  $\mathbf{h}(\mathbf{z})$  a correction term. Remarkably Ma et al. (2015) showed the completeness of their framework:

1.  $\pi(\mathbf{z}) \propto \exp(-H(\mathbf{z}))$  is a stationary distribution of the SDE (2)+(5) for any pair of positive semi-definite matrix  $\mathbf{D}(\mathbf{z})$  and skew-symmetric matrix  $\mathbf{Q}(\mathbf{z})$ ;
2. for any Itô diffusion process that has the unique stationary distribution  $\pi(\mathbf{z})$ , under mild conditions there exist  $\mathbf{D}(\mathbf{z})$  and  $\mathbf{Q}(\mathbf{z})$  matrices such that the process is governed by (2)+(5).

As a consequence, the construction of an SG-MCMC algorithm reduces to defining the state-space  $\mathbf{z}$  and the  $\mathbf{D}(\mathbf{z})$ ,  $\mathbf{Q}(\mathbf{z})$  matrices. Indeed Ma et al. (2015) also cast existing SG-MCMC samplers within the framework, and proposed an improved version of SG-Riemannian-HMC. In general, an appropriate design of these two matrices leads to significant improvements in mixing as well as reduction of sample bias (Li et al., 2016a; Ma et al., 2015). However, this design has been historically based on strong physical intuitions from statistical mechanics (Duane et al., 1987; Neal et al., 2011; Ding et al., 2014). Therefore, it can still be difficult for practitioners to understand and engineer the sampling method that is best suited to their machine learning tasks.

In the next section, we will describe our recipe on meta-learning an SG-MCMC sampler of the form (2)+(5). Before the presentation, we emphasize that the completeness result of the framework is beneficial for our meta-learning task. On the one hand, as meta-learning searches the best algorithm for a given set of tasks, it is crucial that the search space is large enough to contain many useful candidates. On the other hand, some form of “correctness” guarantee is often required to achieve better generalization to test tasks that might not be very similar to the training tasks. Ma et al. (2015)’s completeness result indicates that our proposed method searches SG-MCMC samplers over the subset of all Itô diffusion processes such that each instance is a posterior sampler. Therefore, the proposed meta-learning algorithm has the best from both worlds, indeed our experiments show that the learned sampler is superior to a number of other baseline SG-MCMC methods.

### 3 META-LEARNING FOR SG-MCMC

This section presents a meta-learning approach to learn an SG-MCMC proposal from data. Our aim is to design an appropriate parameterization  $\mathbf{D}(\mathbf{z})$  and  $\mathbf{Q}(\mathbf{z})$ , so that the sampler can be trained on simple tasks with a meta-learning procedure, and generalize to more complicated densities. For simplicity, we only augment the state-space by one extra variable, the momentum (Duane et al., 1987; Neal et al., 2011), although generalization to e.g. thermostat variable (Ding et al., 2014) is straightforward. Thus, the augmented state-space is  $\mathbf{z} = (\mathbf{s}; \mathbf{p})$  (i.e.  $\mathbf{r} = \mathbf{p}$ ), and the Hamiltonian is defined as  $H(\mathbf{z}) = U(\mathbf{s}) + \frac{1}{2}\mathbf{p}^T \mathbf{p}$  with identity mass matrix.

#### 3.1 EFFICIENT PARAMETERIZATION OF DIFFUSION AND CURL MATRICES

For neural networks, the dimensionality of  $\mathbf{z}$  can be at least tens of thousands. Thus, training and applying full  $\mathbf{D}(\mathbf{z})$  and  $\mathbf{Q}(\mathbf{z})$  matrices can cause a huge computational burden, let alone gradient computations required by  $\mathbf{z}$ . To address this, we define the preconditioning matrices as follows:

$$\mathbf{Q}(\mathbf{z}) = \begin{pmatrix} \mathbf{0} & \mathbf{Q}_f(\mathbf{z}) \\ \mathbf{Q}_f(\mathbf{z}) & \mathbf{0} \end{pmatrix}; \quad \mathbf{D}(\mathbf{z}) = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_f(\mathbf{z}) \end{pmatrix}; \quad (6)$$

$$\mathbf{Q}_f(\mathbf{z}) = \text{diag}[\mathbf{f}_o(\mathbf{z})]; \quad \mathbf{D}_f(\mathbf{z}) = \text{diag}[\mathbf{f}_o(\mathbf{z}) + \mathbf{f}_d(\mathbf{z}) + c]; \quad ; c > 0:$$

Here  $\mathbf{f}_d$  and  $\mathbf{f}_o$  are neural network parameterized functions that will be detailed in section 3.2, and  $c$  is a small positive constant. We choose  $\mathbf{D}_f$  and  $\mathbf{Q}_f$  to be diagonal for fast computation, although future work can explore low-rank matrix solutions. From Ma et al. (2015), our design has the stationary distribution  $\pi(\mathbf{z}) \propto \exp(-U(\mathbf{s}))$  if  $\mathbf{f}_d$  is non-negative for all  $\mathbf{z}$ .

We discuss the role of each precondition matrix for better intuition. The curl matrix  $\mathbf{Q}(\mathbf{z})$  in (2) mainly controls the deterministic drift forces introduced by energy gradient  $-\nabla U(\mathbf{s})$  (as seen in many HMC-like procedures and in eq. (5)). Usually, we only have access to the stochastic gradient  $\mathbf{r} \sim \nabla U(\mathbf{s})$  through data sub-sampling, and  $\mathbf{h}(\mathbf{z})$  acts as the friction to counter for the associated noise that mainly affects the momentum. This explains the design of the diffusion matrix  $\mathbf{D}(\mathbf{z})$  that uses  $\mathbf{D}_f(\mathbf{z})$  to control the amount of friction and injected noise to the momentum. Furthermore,

(a) SGHMC (b) the meta-learned sampler (NNSGHMC)

Figure 1: Comparing the computation graphs of SGHMC and the meta-learned sampler (with modified forward Euler discretization). Here  $\mathbf{f}_o$  and  $\mathbf{G}$  transformations are defined by Eq. (7).

$\mathbf{D}_f(\mathbf{z})$  should also account for the pre-conditioning effect introduced by  $\mathbf{Q}_f(\mathbf{z})$ , e.g, when the magnitude of  $\mathbf{Q}_f(\mathbf{z})$  is large, we need higher friction correspondingly. This explains the squared term  $\mathbf{f}_o(\mathbf{z}) \mathbf{f}_o(\mathbf{z})$  in  $\mathbf{D}_f(\mathbf{z})$  design. The positive scaling constants is heuristically selected following (Chen et al., 2014; Ma et al., 2015) (see appendix). Finally, the extra term  $\mathbf{z} = [\mathbf{z}; \mathbf{p}(\mathbf{z})]$  is responsible for compensating the changes introduced by preconditioning matrices  $\mathbf{Q}_f(\mathbf{z})$  and  $\mathbf{D}_f(\mathbf{z})$ .

The discretized dynamics of the state  $(\mathbf{z}; \mathbf{p})$  with step-size  $\epsilon$  and stochastic gradient  $\mathcal{U}(\cdot)$  are

$$\begin{aligned} \mathbf{p}_{t+1} &= (1 - \epsilon \mathbf{D}_f(\mathbf{z}_t)) \mathbf{p}_t - \epsilon \mathbf{Q}_f(\mathbf{z}_t) \mathcal{U}(\mathbf{z}_t) + \epsilon \mathbf{p}(\mathbf{z}_t) + \epsilon \mathbf{N}(\mathbf{0}; 2\epsilon \mathbf{D}_f(\mathbf{z}_t)); \\ \mathbf{z}_{t+1} &= \mathbf{z}_t + \epsilon \mathbf{Q}_f(\mathbf{z}_t) \mathbf{p}_{t+1} + \epsilon \mathbf{z}(\mathbf{z}_t); \quad \mathbf{z}_t = [\mathbf{z}_t; \mathbf{p}_t]; \quad \mathbf{z}_t = [\mathbf{z}_t; \mathbf{p}_{t+1}]; \end{aligned} \quad (7)$$

We use a modified forward Euler discretization (Neal et al., 2011) here, and the computation graph of eq. (7) is visualized in the right part of Figure 1 (see appendix for SGHMC discretized updates). Again we see that  $\mathbf{Q}_f(\mathbf{z})$  is responsible for the acceleration of and from the  $(1 - \epsilon \mathbf{D}_f(\mathbf{z}))$  term in the update equation of  $\mathbf{p}$ , we see that  $\mathbf{D}_f(\mathbf{z})$  controls the friction introduced to the momentum. Note that in the big-data setting, the noisy gradient is approximately Gaussian distributed with mean  $\mathcal{U}(\cdot)$  and variance  $\mathbf{V}(\cdot)$ . Observing this, Ma et al. (2015) further suggested a correction scheme to counter for stochastic gradient noise, which samples the Gaussian noise  $\mathbf{N}(\mathbf{0}; 2\epsilon \mathbf{D}_f(\mathbf{z})^{-2} \mathbf{B}(\cdot))$  with an empirical estimate of the variance  $\mathbf{B}(\cdot) = \mathbf{Q}_f(\mathbf{z}) \mathbf{V}(\cdot) \mathbf{Q}_f^T(\mathbf{z})$  instead. These corrections can be dropped when the discretization step-size is small, therefore, we do not consider them in our experiments.

### 3.2 CHOICES OF INPUTS TO THE NEURAL NETWORKS

We now present detailed functional forms for  $\mathbf{f}_o$  and  $\mathbf{f}_D$ . When designing these, our goal was to achieve a good balance between generalization power and computational efficiency. Recall that the curl matrix  $\mathbf{Q}(\mathbf{z})$  mainly controls the drift of the dynamics, and the desired behavior is the fast traverse through low-density regions. One useful source of information to identify this is the energy function  $U(\cdot)$ .<sup>1</sup> We also include the momentum  $\mathbf{p}$  to the inputs of  $\mathbf{f}_o$ , allowing the  $\mathbf{Q}(\mathbf{z})$  matrix to observe the velocity information of the  $\mathbf{p}$ . We further add an offset to  $\mathbf{Q}(\mathbf{z})$  to prevent the vanishing of this matrix. Putting all of them together, we define the element of  $\mathbf{f}_o$  as

$$\mathbf{f}_{o;i}(\mathbf{z}) = \mathbf{f}_{o,i}(U(\cdot); \mathbf{p}_i); \quad (8)$$

The corresponding  $(\mathbf{z})$  term requires both  $\mathbf{f}_{o,i}(U(\cdot); \mathbf{p}_i)$  and  $\mathbf{p}_i \mathbf{f}_{o,i}(U(\cdot); \mathbf{p}_i)$ . The energy gradient  $\mathbf{g}_i U(\cdot)$  also appears in (7), so it remains to compute  $\mathbf{f}_{o,i}$ , which, along with  $\mathbf{p}_i \mathbf{f}_{o,i}(U(\cdot); \mathbf{p}_i)$ , can be obtained by automatic differentiation (Abadi et al., 2015).

Matrix  $\mathbf{D}(\mathbf{z})$  is responsible for the friction and the stochastic gradient noise, which are crucial for better exploration around high-density regions. Therefore, we also add the energy gradient  $\mathbf{g}_i U(\cdot)$  to the inputs, meaning that the element of  $\mathbf{f}_D$  is

$$\mathbf{f}_{D;i}(\mathbf{z}) = \mathbf{f}_{D,i}(U(\cdot); \mathbf{p}_i; \mathbf{g}_i U(\cdot)); \quad (9)$$

By the construction of the  $\mathbf{D}(\mathbf{z})$  matrix, the  $\mathbf{p}$  vector only requires  $\mathbf{p}_i \mathbf{D}_f$  without computing any higher order information.

<sup>1</sup>The energy gradient  $\mathbf{g}_i U(\cdot)$  is also informative here, however, it requires expensive computation of the diagonal Hessian for  $(\mathbf{z})$ . For similar reasons we do not consider other higher order derivatives as inputs.

In practice, both  $U(\cdot)$  and  $\mathbb{U}(\cdot)$  are replaced by their stochastic estimates  $u_t$  and  $\mathbb{u}_t$ , respectively. To keep the scale of the inputs roughly the same across tasks, we rescale all the inputs using statistics computed by simulating the sampler with randomly initialized  $\theta_0$ . When the computational budget is limited, we replace the exact gradient computation required by finite difference approximations. We refer the reader to the appendix for details.

### 3.3 LOSS FUNCTION DESIGN FOR METALEARNING

Another challenge is to design a meta-learning procedure for the sampler to encourage faster convergence and low bias on test tasks. To achieve these goals we propose two loss functions that we named as the cross-chain loss and the in-chain loss. From now on we consider the discretized dynamics and denote  $q_t(\cdot|D)$  as the marginal distribution of the random variable at time  $t$ .

**Cross-chain loss** We introduce a cross-chain loss that encourages the sampler to converge faster. Since the sampler is guaranteed to have the unique stationary distribution  $\exp(-U(\cdot))$ , fast convergence means that  $\mathbb{L}[q_t, \pi]$  is close to zero when  $t$  is small. Therefore this KL-divergence becomes a sensible objective to minimize, which is equivalent to maximizing the variational lower-bound (or ELBO):  $L_{VI}^t(q_t) = \mathbb{E}_{q_t}[U(\cdot)] + H[q_t]$  (Jordan et al., 1999; Beal, 2003). We further make the objective doubly stochastic: (1) the energy term is further approximated by its stochastic estimates  $u_t(\cdot)$ ; (2) we use Monte Carlo variational inference (MCVI, Ranganath et al., 2014; Blundell et al., 2015) which estimates the lower-bound with samples  $q_t(\cdot|D); k = 1; \dots; K$ . These samples  $\{g_{k=1:T}^k\}$  are obtained by simulating  $K$  parallel Markov chains with the sampler, and the cross-chain loss is defined by accumulating the lower-bounds through time:

$$L_{\text{cross-chain}} = \frac{1}{T} \sum_{t=1}^T L_{VI}^t(f_{t:T}^k g_{k=1}^k); \quad L_{VI}^t(f_{t:T}^k g_{k=1}^k) = \frac{1}{K} \sum_{k=1}^K \mathbb{U}(u_t^k) + \log q_t(\cdot|D); \quad (10)$$

By minimizing this objective, we can improve the convergence of the sampler, especially at the early times of the Markov chain. The objective also takes the sampler bias into account because the two distributions will match when the KL-divergence is minimized.

**In-chain loss** For very big neural networks, simulating multiple Markov chains is prohibitively expensive. The issue is mitigated by thinning that collects samples for every  $c$  step (after burn-in), which effectively draws samples from the averaged distribution  $q_s(\cdot|D) = \frac{1}{bT-c} \sum_{s=1}^{bT-c} q_s(\cdot)$ . The in-chain loss is, therefore, defined as the ELBO evaluated at the averaged distribution which is then approximated by Monte Carlo with samples  $\{g_{s=1}^k\}_{s=1}^{bT-c}$  obtained by thinning:

$$L_{\text{in-chain}} = \frac{1}{K} \sum_{k=1}^K L_{VI}^k(g_{s=1}^k); \quad L_{VI}^k(g_{s=1}^k) = \frac{1}{bT-c} \sum_{s=1}^{bT-c} \mathbb{U}(u_s^k) + \log q_s(\cdot|D); \quad (11)$$

**Gradient approximation** We leverage the recently proposed Stein gradient estimator (Li & Turner, 2018) to estimate the intractable gradients  $\log q_t(\cdot)$  for cross-chain loss and  $\log q_s(\cdot)$  for in-chain loss. Precisely, by the chain rule, we have  $\log q_t(\cdot) = \mathbb{r}^T \log q(\cdot)$ , so it remains to estimate the gradient  $\mathbb{G} = (\mathbb{r}_1^T \log q(\cdot); \dots; \mathbb{r}_K^T \log q(\cdot))^T$  at the sampled locations  $f_{t:T}^k g_{k=1}^k$ . The recipe first constructs a kernel matrix  $\mathbf{K}$  with  $K_{ij} = \mathbb{K}(g_{t:T}^i; g_{t:T}^j)$ , and then estimates the gradients  $\mathbb{G} = (\mathbf{K} + \mathbb{I})^{-1} \mathbb{h}; \mathbf{K} \mathbf{i}$ , where  $\mathbb{h}; \mathbf{K} \mathbf{i} = \sum_{k=1}^K \mathbb{U}(g_{t:T}^k) \mathbb{K}(g_{t:T}^k; \cdot)$ . In our experiments, we use the RBF kernel, and the corresponding gradient estimator has a simple analytic form that can be computed efficiently in  $\mathcal{O}(K^2 D + K^3)$  time (usually  $K \ll D$ ).

## 4 RELATED WORK

Since the development of SGLD (Welling & Teh, 2011), SG-MCMC has been increasingly popular for posterior sampling on big data. In detail, Chen et al. (2014) scaled up HMC with stochastic gradients, Ding et al. (2014) further augmented the state space with an auxiliary temperature variable, and Springenberg et al. (2016) improved robustness through scale adaptation. The SG-MCMC

Figure 2: (Left) Sampler's bias measured by KL. (Middle) NNSGHMC trajectory plot on a 2D-Gaussian with manually injected gradient noise. (Right) SGHMC plot for the same settings.

extensions to Riemannian Langevin dynamics and HMC (Girolami & Calderhead, 2011) have also been proposed (Patterson & Teh, 2013; Ma et al., 2015). Our proposed sampler architecture further generalizes SG-Riemannian-HMC as it decouples the desired  $\mathbf{P}(\mathbf{z})$  and  $\mathbf{Q}(\mathbf{z})$  matrices, and the detailed functional form of these two matrices are also learned from data.

Our approach is closely related to the recent line of work on learning optimization algorithms. Specifically, Andrychowicz et al. (2016) trained a recurrent neural network (RNN) based optimizer that transfers to similar tasks with supervised learning. Later Chen et al. (2017) generalized this approach to Bayesian optimization (Brochu et al., 2010; Snoek et al., 2012) which is gradient-free. We do not use RNNs in our approach as it cannot be represented within the framework of Ma et al. (2015). We leave the combination of learnable RNN proposals to future work. Also Li & Turner (2018) presented an initial attempt to meta-learn an approximate inference algorithm, which simply combined the stochastic gradient and the Gaussian noise with a neural network. Thus the stationary distribution of that sampler (if it exists) is only an approximation to the exact posterior. On the other hand, the proposed sampler (with  $\beta = 0$ ) is guaranteed to be correct by the complete framework (Ma et al., 2015). Very recently Wu et al. (2018) discussed that short-horizon meta-objectives for learning optimizers can cause a serious issue for long-time generalization. We found this bias is less severe in our approach, again due to the fact that the learned sampler is provably correct.

Recent research also considered improving HMC with a trainable transition kernel. Salimans et al. (2015) improved upon vanilla HMC by introducing a trainable re-sampling distribution for the momentum. Song et al. (2017) parameterized the HMC transition kernel with a trainable invertible transformation called non-linear independent components estimation (NICE, Dinh et al., 2014), and train it with Wasserstein adversarial training (Arjovsky et al., 2017). Levy et al. (2018) generalized HMC by augmenting the state space with a binary direction variable, and they parameterized the transition kernel with a non-volume preserving invertible transformation that is inspired by the real-valued non-volume preserving (RealNVP) flows (Dinh et al., 2017). The sampler is trained with the expected squared jump distance (Pasarica & Gelman, 2010). We note that adversarial training is less reliable for high dimensional data, thus it is not considered in this paper. Also, the jump distance does not explicitly take the sampling bias and convergence speed into account. More importantly, the purpose of these approaches is to directly improve the HMC-like sampler on the target distribution, and with NICE/RealNVP parametrization it is difficult to generalize the sampler to densities of different dimensions. In contrast, our goal is to learn an SG-MCMC sampler that can later be transferred to sample from different Bayesian neural network posterior distributions, which will typically have different dimensionality and include tens of thousands of random variables.

## 5 EXPERIMENTS

We evaluate the meta-learned SG-MCMC sampler, which is referred to as NNSGHMC or the meta sampler in the following. Detailed test set-ups are reported in the appendix. The code is available at <https://github.com/WenboGong/MetaSGMCMC>.

Figure 3: Learning curves on test error (top) and negative test LL (bottom).

Table 1: The final performance for the samplers on MNIST, averaged over 10 independent runs.

Sampler	NT Err.	NT+AF Err	NT+Data Err	NT NLL	NT+AF NLL	NT+Data NLL
NNSGHMC	98.36 <b>0:02%</b>	97.72 <b>0:02%</b>	98.62 <b>0:02%</b>	640 6:25	875 <b>3:19</b>	230 3:23
SGHMC	98.21 0:01%	97.72 <b>0:01%</b>	98.52 0:03%	705 3:44	929 2:95	246 5:43
SGLD	98.27 0:02%	97.62 0:02%	98.54 0:01%	631 3:15	905 2:36	232 1:93
PSGLD	98.31 0:02%	97.67 0:02%	98.60 0:02%	610 <b>2:93</b>	975 4:41	224 <b>1:97</b>

### 5.1 SYNTHETIC EXAMPLE: SAMPLING GAUSSIAN VARIABLES WITH NOISY GRADIENTS

We first consider sampling Gaussian variables to demonstrate fast convergence and low bias of the meta sampler. To mimic stochastic gradient settings, we manually inject Gaussian noise with unit variance to the gradient as suggested by (Chen et al., 2014). The training density is a 10D Gaussian with randomly generated diagonal covariance matrix, and the test density is a 20D Gaussian. For evaluation, we simulate  $k = 50$  parallel chains for  $t = 12;000$  steps. Then we follow Ma et al. (2015) to evaluate the sampler’s bias measured by the KL divergence from the empirical estimate to the ground truth. Results are visualized on the left panel of Figure 2, showing that the meta sampler both converges much faster and achieves lower bias compared to SGHMC. The effective sample size for SGHMC and NNSGHMC are 22 and 59, again indicating better efficiency of the meta sampler. For illustration purposes, we also plot in the other two panels the trajectory of samples by simulating NNSGHMC (middle) and SGHMC (right) on a 2D Gaussian for a fixed amount of time. This confirms that the meta sampler explores more efficiently and is less affected by the injected noise.

### 5.2 BAYESIAN FEEDFORWARD NEURAL NETWORKS

Next, we consider Bayesian neural network classification on MNIST data with three generalization tests: network architecture generalization (NT), activation function generalization (AF) and dataset generalization (Data). In all tests, the sampler is trained with a 1-hidden layer multi-layer perceptron (MLP) (20 units, ReLU activation) as the underlying model for the target distribution. We also report long-time horizon generalization results, meaning that the simulation time steps in test time are much longer than that of training (cf. Andrychowicz et al., 2016). Algorithms in comparison include SGLD (Welling & Teh, 2011), SGHMC (Chen et al., 2014) and preconditioned SGLD (PSGLD, Li et al., 2016a). Note that PSGLD uses RMSprop-like preconditioning techniques (Tieleman & Hinton, 2012) that require moving average estimates of the gradient’s second moments. Therefore the underlying dynamics of PSGLD cannot be represented within our framework (6). Thus we mainly focus on comparisons with SGLD and SGHMC, and leave the PSGLD results as reference. The discretization step-sizes for the samplers are tuned on the validation dataset for each task.

**Architecture generalization (NT)** In this test we use the trained sampler to draw samples from the posterior distribution of a 2-hidden layer MLP with 40 units and ReLU activations. Figure 3 shows the learning curves of test error and negative test log-likelihood (NLL) for 100 epochs, where the final performance is reported in Table 1. Overall NNSGHMC achieves the fastest convergence even when compared with PSGLD. It has the lowest test error compared to SGLD and SGHMC. NNSGHMC’s final test LL is on par with SGLD and slightly worse than PSGLD, but it is still better than SGHMC.

<sup>2</sup>Implementation follows the ESS function in the BEAST package <http://beast.community>.

Figure 4: (Left) The contour plot of function  $f_{\theta}$  (Middle) The contour plot for  $f_{\theta}$  for dimension 1 and 2 with fixed  $r = U(\cdot)$  (Right) The same plot for  $f_{\theta}$  for dimension 2 and 3 with fixed energy.

**Architecture + Activation function generalization (NT+AF)** Next we replace NT’s test network’s activation function with sigmoid and re-run the same test as before. Again results in Figure 3 and Table 1 show that NNSGHMC converges faster than others for both test error and NLL. It also achieves the best NLL results among all samplers, and the same test error as SGHMC.

**Architecture + Dataset generalization (NT+Data)** In this test we split MNIST into training task (classifying digits 0-4) and test task (digits 5-9). The meta sampler is trained with the smaller MLP, and it is evaluated on the task with the larger MLP with NT’s architecture. Thus, the meta sampler is trained without any knowledge of the test task’s training and test data. From Figure 3, we see that NNSGHMC, although a bit slower at the start, catches up quickly and proceeds to lower error. The difference between these samplers NLL results is marginal, and NNSGHMC is on par with PSGLD.

**Learned strategies** For better intuition, we visualize in Figure 4 the contours of  $f_{\theta}$  and  $f_{\theta}$ . From the left panel,  $f_{\theta}$  has learned a nearly linear strategy w.r.t. the energy and small variations w.r.t. the momentum. This enables the sampler for fast traversal through low density (high energy) regions and better exploration at high density (low energy) area.

The strategy learned for the diffusion matrix  $D(\mathbf{z})$  is rather interesting. Recall that  $D(\mathbf{z})$  is parametrized by both  $f_{\theta}$  and  $f_{\theta}$  (eq. (6)). Since Figure 4 (left) indicates that  $f_{\theta}$  is large in high energy regions, the amount of friction is adequate, and tends to reduce its output to maintain momentum (see the middle plot). By contrast, at low energy regions,  $f_{\theta}$  increases to add friction in order to prevent divergence. The right panel visualizes the interactions between the momentum and the mean gradient  $r = U(\cdot)$  at a fixed energy level. This indicates that the meta sampler has learned a strategy to prevent overshoot by producing large friction, and returns large values when the signs of the momentum and the gradient differ.

### 5.3 BAYESIAN CONVOLUTIONAL NEURAL NETWORKS

Following the setup of BNN MNIST experiments, we also test our algorithm on convolutional neural networks (CNNs) for CIFAR-10 (Krizhevsky, 2009) classification, again with three generalization tasks (NT, AF and Data). The meta sampler is trained using a smaller CNN with two convolutional layers  $(3 \times 3 \times 8)$  and  $(3 \times 3 \times 8 \times 8)$  and one fully connected (fc) layer (50 hidden units). ReLU activations, and max-pooling operators of size 2 are applied after each convolutional layer. The meta sampler is trained using 100 “meta-epochs”, where each “meta-epoch” has 5 data epochs. At the beginning of each “meta-epoch”, a “replay” technique inspired by experience replay (Lin, 1993) is utilized (see appendix). The discretization step-sizes are tuned on a validation dataset for each task.

**Architecture generalization (NT)** The test CNN has two convolutional layers  $(3 \times 3 \times 16)$  and  $(3 \times 3 \times 16 \times 16)$  and one fc layer (100 hidden units), resulting in roughly a dimensionality of 100. Figure 5 shows that the meta sampler achieves the fastest learning at the first 10 epochs, and continues to have better performance in both test accuracy and NLL. Interestingly, PSGLD slows down quickly after 3 epochs, and it converges to a worse answer. The best performance over 200 epochs is shown in Table 2, where the meta sampler is a clear winner in both accuracy and NLL. This demonstrates that our sampler indeed converges faster and has found a better posterior mode.



Figure 5: Learning curves on test error (top) and negative test LL/100 (bottom).

Table 2: The best performance on CIFAR-10 over 200 epoch, averaged over 5 independent runs. All the samplers achieved the best performance after around 190 epochs.

Methods	NT Err.	NT+AF Err	NT+Data Err	NT NLL/100	NT+AF NLL/100	NT+Data NLL/100
NNSGHMC	78.12 <b>0:035%</b>	74.41 <b>0:11%</b>	89.97 <b>0:04%</b>	68.88 <b>0:15</b>	79.55 <b>0:057</b>	15.66 0:28
SGHMC	77.63 0:068%	73.68 0:17%	90.11 <b>0:15%</b>	70.39 0:27	82.08 0:48	15.26 0:07
SGLD	76.38 0:085%	73.83 0:013%	89.34 0:06%	73.39 0:26	80.30 0:14	16.36 0:07
PSGLD	77.46 0:05%	73.16 0:1%	89.78 0:08%	69.89 0:2	83.70 0:21	15.72 0:04
Adam	70.94 0:10%	69.73 0:11%	85.44 0:14%	86.77 1:01	87.60 0:44	22.35 0:47
SGD-M	68.06 0:27%	68.76 0:17%	84.86 0:48%	99.12 1:06	90.35 0:29	23.64 0:73

Architecture + Activation function generalization (NT+AF) We use the same CNN architecture as in NT but replace the ReLU activations with sigmoid. Figure 5 and Table 2 show that the meta sampler again has better convergence speed and the best final performance.

Architecture + Dataset generalization (NT+Data) We split CIFAR-10 according to labels 0-4 as the training task and 5-9 as the test task. We also used the same CNN architecture as in NT. From Figure 5 and Table 2, the meta sampler consistently achieves the fastest convergence speed. It also achieves similar accuracy as SGHMC, but it has slightly worse test NLL compared to SGHMC.

#### 5.4 BAYESIAN RECURRENT NEURAL NETWORKS

Lastly, we consider a more challenging setup: sequence modeling with Bayesian RNNs. Here a single datum is a sequence  $\mathbf{x} = \mathbf{x}_1^T, \dots, \mathbf{x}_n^T$  and the log-likelihood is defined as  $\log p(\mathbf{x}) = \sum_{t=1}^n \log p(\mathbf{x}_t | \mathbf{x}_{1:t-1})$ , with each of the conditional densities produced by a gated recurrent unit (GRU) network (Cho et al., 2014). We consider four polyphonic music datasets for this task: Piano-midi (Piano) as training data, and Nottingham (Nott), MuseData (Muse) and JSB chorales (JSB) for evaluation. The meta sampler is trained on a small GRU with 100 hidden states. At test time, we follow Chen et al. (2016) and set the step-size to 0.001. We found SGLD significantly under-performs, so instead, we report the performances of two optimizers, Adam (Kingma & Ba, 2014) and Santa (Chen et al., 2016), taken from Chen et al. (2016). Again, these two optimizers use moving average schemes which are out of the scope of our framework, so we mainly compare the meta sampler with SGHMC and leave the others as references.

The meta sampler is tested on the four datasets using 200 unit GRU. So for Piano this corresponds to architecture generalization only. From Figure 6 we see that the meta sampler achieves faster convergence compared to SGHMC, at the same time it achieves similar speed as Santa at early stages. All the samplers achieve best results close to Santa on Piano. The meta sampler successfully generalizes to the other three datasets, demonstrating faster convergence than SGHMC consistently, and better final performance on Muse. Interestingly, the meta sampler's final results on Nott and JSB are slightly worse than other samplers. Presumably, these two datasets are very different from Muse and Piano, therefore, the energy landscape is less similar to the training density (see appendix). Specifically, JSB is a dataset with much shorter sequences. And in this case, SGHMC also exhibits over-fitting but to a smaller degree. Therefore, we further test the meta sampler on JSB without the offset in  $\mathbf{f}_0$  to reduce the acceleration (denoted as NNSGHMC-s). Surprisingly, NNSGHMC-s

Method	Piano	Muse	Nott	JSB
NNSGHMC	7.66	7.27	3.37	8.49
SGHMC	7.65	7.33	3.35	8.40
PSGLD	7.67	7.48	3.28	8.42
Santa	7.6	7.2	3.39	8.46
Adam	8	7.56	3.7	8.51

Figure 6: Test NLL learning curve (with zoom-in for sampling methods) and the best performance. Santa and Adam results are from Chen et al. (2016)

convergences in similar speeds as the original one, but with less amount of over-fitting and better final test NLL 8.40

## 6 CONCLUSIONS AND FUTURE WORK

We have presented a meta-learning algorithm that can learn an SG-MCMC sampler on simpler tasks and generalizes to more complicated densities in high dimensions. Experiments on Bayesian MLPs, Bayesian CNNs and Bayesian RNNs confirmed the strong generalization of the trained sampler to the long-time horizon as well as across datasets and network architectures. Future work will focus on better designs for both the sampler and the meta-learning procedure. For the former, temperature variable augmentation as well as moving average estimation will be explored. For the latter, better loss functions will be proposed for faster training, e.g. by reducing the unrolling steps of the sampler during training. Finally, the automated design of generic MCMC algorithms that might not be derived from continuous Markov processes remains an open challenge.

### ACKNOWLEDGMENTS

We thank Shixiang Gu, Mark Rowland and Cheng Zhang for comments on the manuscript. We also appreciate Changyou Chen for providing the experiment results of Bayesian RNN (Chen et al., 2016). Wenbo Gong is supported by the CSC-Cambridge Trust Scholarship.

### REFERENCES

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Sungjin Ahn, Anoop Korattikara, and Max Welling. Bayesian posterior sampling via stochastic gradient descent. arXiv preprint arXiv:1206.6380, 2012.

- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. *Advances in Neural Information Processing Systems*, pp. 3981–3989, 2016.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pp. 214–223, 2017.
- Matthew James Beal. Variational algorithms for approximate Bayesian inference. PhD thesis, University College London, 2003.
- Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. On the optimization of a synaptic learning rule. In *Conference on Optimality in Biological and Artificial Networks*, 1992.
- Mikołaj Bińkowski, Dougal J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pp. 1613–1622, 2015.
- Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- Changyou Chen, David Carlson, Zhe Gan, Chunyuan Li, and Lawrence Carin. Bridging the gap between stochastic gradient MCMC and stochastic optimization. *Artificial Intelligence and Statistics*, pp. 1051–1060, 2016.
- Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *International Conference on Machine Learning*, pp. 1683–1691, 2014.
- Yutian Chen, Matthew W Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Timothy P Lillicrap, Matt Botvinick, and Nando Freitas. Learning to learn without gradient descent by gradient descent. In *International Conference on Machine Learning*, pp. 748–756, 2017.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Kacper Chwiałkowski, Heiko Strathmann, and Arthur Gretton. A kernel test of goodness of fit. In *International Conference on Machine Learning*, 2016.
- Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *International Conference on machine learning*, pp. 465–472, 2011.
- Stefan Depeweg, José Miguel Hernández-Lobato, Finale Doshi-Velez, and Steffen Udluft. Learning and policy search in stochastic dynamical systems with Bayesian neural networks. *International Conference on Learning Representations*, 2017.
- Nan Ding, Youhan Fang, Ryan Babbush, Changyou Chen, Robert D Skeel, and Hartmut Neven. Bayesian sampling using stochastic gradient thermostats. *Advances in Neural Information Processing Systems*, pp. 3203–3211, 2014.
- Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=HkpbhH9lx>.
- Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Physics letters B* 195(2):216–222, 1987.

- Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. Detecting adversarial samples from artifacts. arXiv preprint arXiv:1703.00410, 2017.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic meta-learning for fast adaptation of deep networks. *International Conference on Machine Learning*, pp. 1126–1135, 2017.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *International Conference on Machine Learning*, pp. 1050–1059, 2016.
- Zhe Gan, Chunyuan Li, Changyou Chen, Yunchen Pu, Qinliang Su, and Lawrence Carin. Scalable Bayesian learning of Recurrent neural networks for language modeling. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) volume 1*, pp. 321–331, 2017.
- Mark Girolami and Ben Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2): 123–214, 2011.
- Jackson Gorham and Lester Mackey. Measuring sample quality with Stein's method. *Advances in Neural Information Processing Systems*, pp. 226–234, 2015.
- Alex Graves. Practical variational inference for neural networks. *Advances in Neural Information Processing Systems*, pp. 2348–2356, 2011.
- Jose Hernandez-Lobato, Yingzhen Li, Mark Rowland, Thang Bui, Daniel Hernandez-Lobato, and Richard Turner. Black-box Alpha divergence minimization. *International Conference on Machine Learning*, pp. 1511–1520, 2016.
- José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of Bayesian neural networks. *International Conference on Machine Learning*, pp. 1861–1869, 2015.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Daniel Levy, Matt D. Hoffman, and Jascha Sohl-Dickstein. Generalizing Hamiltonian Monte Carlo with neural networks. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=B1n8LexRZ>.
- Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. Mmd gan: Towards deeper understanding of moment matching networks. *Advances in Neural Information Processing Systems*, pp. 2203–2213, 2017.
- Chunyuan Li, Changyou Chen, David Carlson, and Lawrence Carin. Preconditioned stochastic gradient Langevin dynamics for deep neural networks. *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pp. 1788–1794. AAAI Press, 2016a.
- Chunyuan Li, Andrew Stevens, Changyou Chen, Yunchen Pu, Zhe Gan, and Lawrence Carin. Learning weight uncertainty with stochastic gradient MCMC for shape classification. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5666–5675, 2016b.
- Ke Li and Jitendra Malik. Learning to optimize. *International Conference on Learning Representations*, 2017.
- Yingzhen Li and Yarin Gal. Dropout inference in Bayesian neural networks with Alpha-divergences. *International Conference on Machine Learning*, pp. 2052–2061, 2017.

- Yingzhen Li and Richard E. Turner. Gradient estimators for implicit models. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SJi9WOeRb>.
- Long-Ji Lin. Reinforcement learning for robots using neural networks. Technical report, Carnegie-Mellon Univ Pittsburgh PA School of Computer Science, 1993.
- Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances In Neural Information Processing Systems* pp.2378–2386, 2016.
- Qiang Liu, Jason Lee, and Michael Jordan. A kernelized Stein discrepancy for goodness-of-fit tests. In *International Conference on Machine Learning* pp. 276–284, 2016.
- Christos Louizos and Max Welling. Multiplicative normalizing flows for variational Bayesian neural networks. In *International Conference on Machine Learning* pp. 2218–2227, 2017.
- Yi-An Ma, Tianqi Chen, and Emily Fox. A complete recipe for stochastic gradient MCMC. In *Advances in Neural Information Processing Systems* pp.2917–2925, 2015.
- Devang K Naik and RJ Mammone. Meta-neural networks that learn by learning. *Neural Networks*, 1992. *IJCNN*, International Joint Conference, volume 1, pp. 437–442. IEEE, 1992.
- Radford M Neal et al. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11), 2011.
- Cuong V. Nguyen, Yingzhen Li, Thang D. Bui, and Richard E. Turner. Variational continual learning. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BkQqq0gRb>.
- Cristian Pășărică and Andrew Gelman. Adaptively scaling the Metropolis algorithm using expected squared jumped distance. *Statistica Sinica* pp. 343–364, 2010.
- Sam Patterson and Yee Whye Teh. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems* pp.3102–3110, 2013.
- Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. *Artificial Intelligence and Statistics* pp. 814–822, 2014.
- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. *International Conference on Learning Representations*, 2017.
- Hippolyt Ritter, Aleksandar Botev, and David Barber. A scalable Laplace approximation for neural networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Skdvd2xAZ>.
- Tim Salimans, Diederik Kingma, and Max Welling. Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning* pp. 1218–1226, 2015.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International Conference on Machine Learning* pp. 1842–1850, 2016.
- Jürgen Schmidhuber. Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hood. PhD thesis, Technische Universität München, 1987.
- Jianghong Shi, Tianqi Chen, Ruoshi Yuan, Bo Yuan, and Ping Ao. Relation of a new interpretation of stochastic differential equations to its process. *Journal of Statistical Physics* 148(3):579–590, 2012.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems* pp.2951–2959, 2012.
- Jiaming Song, Shengjia Zhao, and Stefano Ermon. A-NICE-MC: Adversarial training for MCMC. In *Advances in Neural Information Processing Systems* pp.5146–5156, 2017.

- Jost Tobias Springenberg, Aaron Klein, Stefan Falkner, and Frank Hutter. Bayesian optimization with robust Bayesian neural networks. *Advances in Neural Information Processing Systems (NIPS)*, pp. 4134–4142, 2016.
- Charles M Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory*, pp. 583–602, 1972.
- Charles M Stein. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, pp. 1135–1151, 1981.
- Sebastian Thrun and Lior Pradel. *Learning to learn*. Springer Science & Business Media, 1998.
- Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012.
- Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. preprint arXiv:1611.05763, 2016.
- Max Welling and Yee W Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *International Conference on Machine Learning*, pp. 681–688, 2011.
- Olga Wichrowska, Niru Maheswaranathan, Matthew W Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Nando Freitas, and Jascha Sohl-Dickstein. Learned optimizers that scale and generalize. In *International Conference on Machine Learning*, pp. 3751–3760, 2017.
- Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial Intelligence and Statistics*, pp. 370–378, 2016.
- Yuhuai Wu, Mengye Ren, Renjie Liao, and Roger Grosse. Understanding short-horizon bias in stochastic meta-optimization. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=H1MczcgR->
- Lan Yin and Ping Ao. Existence and construction of dynamical potential in nonequilibrium processes without detailed balance. *Journal of Physics A: Mathematical and General*, 39(27):8593, 2006.

## A COMPARING MOMENTUM SGD AND SGHMC

Similar to the relationship between SGLD and SGD, SGHMC is closely related to SGD with momentum (SGD-M). First in HMC, the state space is augmented with an additional momentum variable denoted as  $\mathbf{p} \in \mathbb{R}^D$ . We assume an identity mass matrix associated with that momentum term. Then the corresponding drift  $\mathbf{f}(\mathbf{z}; \mathbf{p})$  and diffusion matrix  $\mathbf{D}$  are:

$$\mathbf{f}(\mathbf{z}; \mathbf{p}) = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{I} & \mathbf{C} \end{bmatrix} \begin{bmatrix} \nabla U(\mathbf{z}) \\ \mathbf{p} \end{bmatrix}; \quad \mathbf{D} = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{C} \end{bmatrix}; \quad (12)$$

where  $\mathbf{C}$  is a positive definite matrix called friction coefficient. Thus, HMC's continuous-time dynamics is governed by the following SDE:

$$\begin{aligned} d\mathbf{z} &= \mathbf{p} dt; \\ d\mathbf{p} &= -\nabla U(\mathbf{z}) dt - \mathbf{C} \mathbf{p} dt + \sqrt{\mathbf{C}} d\mathbf{W}(t); \end{aligned} \quad (13)$$

The discretized update rule (with simple Euler discretization) of HMC with step-size

$$\begin{aligned} \mathbf{z}_{t+1} &= \mathbf{z}_t + \mathbf{p}_t \Delta t; \\ \mathbf{p}_{t+1} &= (1 - \mathbf{C} \Delta t) \mathbf{p}_t - \nabla U(\mathbf{z}_t) \Delta t + \sqrt{\mathbf{C}} \Delta t \epsilon_t; \\ \epsilon_t &\sim \mathcal{N}(\mathbf{0}, 2\mathbf{C}): \end{aligned} \quad (14)$$

If stochastic gradient  $\hat{\nabla} U(\mathbf{z})$  is used, we need to replace the covariance matrix with  $2(\mathbf{C} + \hat{\mathbf{B}})$  where  $\hat{\mathbf{B}}$  is the variance estimation of the gradients.

On the other hand, the update equations of SGD with momentum (SGD-M) are the following:

$$\begin{aligned} \mathbf{z}_{t+1} &= \mathbf{z}_t + \mathbf{v}_t \Delta t; \\ \mathbf{v}_{t+1} &= \alpha \mathbf{v}_t - \eta \nabla U(\mathbf{z}_t); \end{aligned} \quad (15)$$

where  $\alpha$  and  $\eta$  are called momentum discount factor and learning rate respectively. Also we can rewrite the SGHMC update equations by setting  $\mathbf{p}_t = \mathbf{v}_t$ ,  $\alpha = 1 - \Delta t$ ,  $\eta = \Delta t$ ,  $\mathbf{C} = \mathbf{C}$ ,

$$\begin{aligned} \mathbf{z}_{t+1} &= \mathbf{z}_t + \mathbf{v}_t \Delta t; \\ \mathbf{v}_{t+1} &= \alpha \mathbf{v}_t - \eta \nabla U(\mathbf{z}_t) + \sqrt{2\mathbf{C}(1-\alpha)} \Delta t \epsilon_t; \\ \epsilon_t &\sim \mathcal{N}(\mathbf{0}, 2\mathbf{I}(1-\alpha)): \end{aligned} \quad (16)$$

Thus, the discretized SGHMC updates can be viewed as the SGD-M update injected with carefully controlled Gaussian noise. Therefore, the hyperparameter of SGHMC can be heuristically chosen based on the experience of SGD-M and vice versa.

Neal et al. (2011) showed that in practice, simple Euler discretization for HMC simulation might cause divergence, therefore advanced discretization schemes such as Leapfrog and modified Euler are recommended. We use modified Euler discretization in our implementation of SGHMC and the meta sampler, resulting in the following update:

$$\begin{aligned} \mathbf{p}_{t+1} &= (1 - \mathbf{C} \Delta t) \mathbf{p}_t - \nabla U(\mathbf{z}_t) \Delta t; \\ \mathbf{z}_{t+1} &= \mathbf{z}_t + \mathbf{p}_{t+1} \Delta t; \\ \epsilon_t &\sim \mathcal{N}(\mathbf{0}, 2\mathbf{C}): \end{aligned} \quad (17)$$

## B FINITE DIFFERENCE APPROXIMATION FOR THE GRADIENT VECTOR

The main computational burden is the gradient computation required by the vector. To simplify notations we will write  $\mathbf{D} = \mathbf{D}(\mathbf{z})$  and  $\mathbf{Q} = \mathbf{Q}(\mathbf{z})$ . From the parametrization of the  $\mathbf{C}$  and  $\mathbf{D}$  matrices in eq. (6), for  $\mathbf{z} \in \mathbb{R}^D$  we have  $\mathbf{z} = [\mathbf{z}; \mathbf{p}]$ . For the first term, we have

$$\mathbf{D}_{i,j} = \nabla_{\mathbf{p}} \mathbf{Q}_{i,j} = \frac{\partial \mathbf{Q}_{i,j}}{\partial \mathbf{p}}. \quad (18)$$

Due to the two-stage update of Euler integrator, at time  $t$ , we have  $\mathbf{z}_{t+1}^* = \mathbf{z}_{t+1} + \mathbf{p}_t^* \Delta t$ ,  $\mathbf{f}_{\mathbf{Q},i}^* = \mathbf{f}_{\mathbf{Q},i}(\mathbf{z}_{t+1}^*; \mathbf{p}_t^*)$  and  $\mathbf{f}_{\mathbf{D},i}^* = \mathbf{f}_{\mathbf{D},i}(\mathbf{z}_{t+1}^*; \mathbf{p}_t^*; \nabla_{\mathbf{z}} U(\mathbf{z}_{t+1}^*))$ . Thus a proper finite

difference method requires  $\frac{\partial f_{\alpha,i}}{\partial \mathbf{p}}(U(t); \mathbf{p}_i^{t-1})$ , which is not exactly the history from the previous time. Therefore we further approximate it using delayed estimate:

$$\frac{\partial f_{\alpha,i}}{\partial \mathbf{p}} \left( \frac{f_{\alpha,i}^{t-1}}{\mathbf{p}_i^{t-1}}, \frac{f_{\alpha,i}^t}{\mathbf{p}_i^t} \right) \approx \frac{\partial f_{\alpha,i}}{\partial \mathbf{p}} \left( \frac{\hat{Q}^{t-1}}{\mathbf{p}^t}, \frac{Q^{t-1}}{\mathbf{p}^{t-1}} \right). \quad (19)$$

Similarly, the  $\mathbf{p}$  term expands as

$$\begin{aligned} \mathbf{p}_{;i} &= \mathbf{r} \cdot [\mathbf{D} + \mathbf{Q}]_{;i}; \\ &= \frac{\partial f_{\alpha,i}}{\partial \mathbf{p}} + \frac{\partial f_{\beta,i}}{\partial \mathbf{p}} + 2 f_{\alpha,i} \frac{\partial f_{\alpha,i}}{\partial \mathbf{p}} \\ &= \frac{\partial f_{\alpha,i}}{\partial \mathbf{p}} \frac{\partial U}{\partial \mathbf{p}} + \frac{\partial f_{\beta,i}}{\partial \mathbf{p}} + 2 f_{\alpha,i} \frac{\partial f_{\alpha,i}}{\partial \mathbf{p}}. \end{aligned} \quad (20)$$

We further approximate  $\frac{\partial f_{\alpha,i}}{\partial U}$  by the following

$$\frac{\partial f_{\alpha,i}}{\partial U} \approx \frac{f_{\alpha,i}^t - f_{\alpha,i}^{t-1}}{U(t) - U(t-1)}. \quad (21)$$

This only requires the storage of previous  $\mathbf{Q}$  matrix. However,  $\frac{\partial f_{\beta,i}}{\partial \mathbf{p}}$  requires one further forward pass to obtain  $f_{\beta,i}^{t-1} = f_{\beta,i}(U(t); \mathbf{p}_i^{t-1}; \mathbf{r}_i(U(t)))$ , thus, we have

$$\begin{aligned} \frac{\partial f_{\beta,i}}{\partial \mathbf{p}} \left( \frac{f_{\beta,i}^t}{\mathbf{p}_i^t}, \frac{f_{\beta,i}^{t-1}}{\mathbf{p}_i^{t-1}} \right) &= \frac{\partial f_{\beta,i}}{\partial \mathbf{p}} \left( \frac{Q^t}{U(t)}, \frac{Q^{t-1}}{U(t-1)} \right) \mathbf{r}_i(U(t)) + \frac{f_{\beta,i}^t}{\mathbf{p}_i^t} \frac{\partial f_{\beta,i}}{\partial \mathbf{p}} \left( \frac{Q^{t-1}}{\mathbf{p}_i^{t-1}}, \frac{Q^{t-1}}{\mathbf{p}_i^{t-1}} \right). \end{aligned} \quad (22)$$

Therefore the proposed finite difference method only requires one more forward passes to compute  $f_{\beta,i}^{t-1}$  and instead, save 3 back-propagations. As back-propagation is typically more expensive than forward pass, our approach reduces running time drastically, especially when the sampler are applied to large neural network.

**Time complexity figures** Every SG-MCMC method (including the meta sampler) requires  $\mathcal{O}(HD)$ . The main burden is the forward pass and back-propagation through  $\mathbf{D}(\mathbf{z})$  and  $\mathbf{Q}(\mathbf{z})$  matrices, where the latter one has been replaced by the proposed finite difference scheme. The time complexity is  $\mathcal{O}(HD)$  for both forward pass and finite difference with the number of hidden units in the neural network of the meta sampler. Parallel computation with GPUs improves real-time speed, indeed in our MNIST experiment the meta sampler spends roughly 1.5x time when compared with SGHMC.

## C DETAILS OF THE STEIN GRADIENT ESTIMATOR

For a distribution  $q(\cdot)$  that is implicitly defined by a generative procedure, the density is often intractable. Li & Turner (2018) derived the Stein gradient estimator that estimates  $\mathbf{S} = (\mathbf{r}_1 \log q(\cdot^1); \dots; \mathbf{r}_K \log q(\cdot^K))^T$  on samples  $\mathbf{z}^1, \dots, \mathbf{z}^K \sim q(\cdot)$ . There are two different ways to derive this gradient estimator, here we briefly introduce one of them, and refer the readers to Li & Turner (2018) for details.

We start by introducing Stein's identity (Stein, 1972; 1981; Gorham & Mackey, 2015; Liu et al., 2016). Let  $h: \mathbb{R}^{d^1} \rightarrow \mathbb{R}^{d^0}$  be a differentiable multivariate test function which maps a column vector  $h(\cdot) = [h_1(\cdot); h_2(\cdot); \dots; h_{d^0}(\cdot)]^T$ . One can use integration by parts to show the following Stein's identity when a boundary condition  $\lim_{\|\mathbf{z}\| \rightarrow \infty} h(\mathbf{z}) = 0$  is assumed for the test function:

$$\mathbb{E}_q[h(\cdot) \mathbf{r} \log q(\cdot)^T + \mathbf{r} \cdot h(\cdot)] = 0; \quad \mathbf{r} \cdot h(\cdot) = (\mathbf{r}_1 h_1(\cdot); \dots; \mathbf{r}_{d^0} h_{d^0}(\cdot))^T \in \mathbb{R}^{d^0 \times d^0}. \quad (23)$$

This boundary condition holds for almost any test function  $h$  if it has sufficiently fast-decaying tails (e.g. Gaussian tails). Li & Turner (2018) proposed Stein gradient estimator for  $\mathbf{r} \log q(\cdot)$  by



inverting a Monte Carlo (MC) version of Stein's identity (23):

$$\frac{1}{K} \mathbb{H} \mathbb{G} \overline{\mathbf{r} \cdot \mathbf{h}}; \quad \mathbb{H} = \mathbb{h}(\mathbf{x}^1); \quad \mathbb{h}(\mathbf{x}^k) \in \mathbb{R}^{d^0 \times K}; \quad \overline{\mathbf{r} \cdot \mathbf{h}} = \frac{1}{K} \sum_{k=1}^K \mathbf{r}^k \cdot \mathbf{h}(\mathbf{x}^k) \in \mathbb{R}^{d^0 \times d}.$$

Then  $\mathbb{G}$  is obtained by ridge regression (with  $\|\cdot\|_F$  the Frobenius norm of a matrix)

$$\hat{\mathbb{G}}_{\mathbb{V}}^{\text{Stein}} := \arg \min_{\mathbb{G} \in \mathbb{R}^{K \times d}} \|\mathbb{G} \overline{\mathbf{r} \cdot \mathbf{h}}\|_F + \frac{1}{K} \mathbb{H} \mathbb{G} \mathbb{H}^T + \frac{1}{K^2} \sum_{j,j'} \|\mathbb{G}_{jj'}\|_F^2; \quad \mathbb{0}; \quad (24)$$

which has an analytical solution

$$\hat{\mathbb{G}}_{\mathbb{V}}^{\text{Stein}} = (\mathbb{K} + \mathbb{I})^{-1} \mathbb{H} \mathbf{r}; \quad \mathbb{K}; \quad (25)$$

where

$$\begin{aligned} \mathbb{K} &:= \mathbb{H}^T \mathbb{H}; \quad \mathbb{K}_{ij} = \mathbb{K}(\mathbf{x}^i; \mathbf{x}^j) := \mathbf{h}(\mathbf{x}^i)^T \mathbf{h}(\mathbf{x}^j); \\ \mathbf{H} \mathbf{r}; \quad \mathbb{K} \mathbf{i} &:= \mathbb{K} \mathbb{H}^T \overline{\mathbf{r} \cdot \mathbf{h}}; \quad \mathbf{H} \mathbf{r}; \quad \mathbb{K} \mathbf{i}_{ij} = \sum_{k=1}^K \mathbf{r}^k(j) \mathbb{K}(\mathbf{x}^i; \mathbf{x}^k); \end{aligned}$$

Here  $k(j)$  denotes the  $j^{\text{th}}$  element of vector  $\mathbf{k}$ . One can show that the RBF kernel satisfies Stein's identity (Liu et al., 2016). In this case  $\mathbb{h}(\mathbf{x}) = \mathbb{K}(\mathbf{x}; \cdot)$ ;  $d^0 = d + 1$  and by the reproducing kernel property,  $\mathbf{h}(\mathbf{x})^T \mathbf{h}(\mathbf{x}^0) = \mathbb{K}(\mathbf{x}; \mathbf{x}^0)$ ;  $\mathbb{K}(\mathbf{x}^0; \cdot) \mathbb{H} = \mathbb{K}(\cdot; \mathbf{x}^0)$ ; Li & Turner (2018) also show that the Stein gradient estimator can be obtained by minimizing a Monte Carlo estimate of the kernelized Stein discrepancy (Chwialkowski et al., 2016; Liu et al., 2016).

The kernel choice It is well-known for kernel methods that a better choice of the kernel can greatly improve the performance. However, optimal kernels are often problem specific, and they are generally difficult to obtain. Recently, a popular approach for kernel design is to compose a simple kernel (e.g. RBF kernel) on features extracted from a deep neural network. Representative work includes deep kernel learning for Gaussian processes (Wilson et al., 2016), and adversarial approaches to learn kernel parameters (Li et al., 2017; Błkowski et al., 2018). Unfortunately, both approaches do not scale very well to our application as has at least tens of thousands of dimensions. Furthermore, they both considered kernel learning for observed data, while in our case latent variable to be inferred. Therefore it remains a research question on how to learn kernels on latent variables efficiently, and addressing this question is out of the scope of the paper. Instead, we follow Liu & Wang (2016); Li & Turner (2018) to use RBF kernel for the gradient estimator. Other kernels can be trivially adapted to our method. We expect even better performance if an optimal kernel is in use, but we leave the investigation to future work.

Time complexity During meta sampler training, the Stein gradient estimator requires the kernel matrix inversion which is  $\mathcal{O}(K^3)$  for cross-chain training. In practice, we only run a few parallel Markov chains  $K = 20 \sim 50$ , thus, this will not incur huge computation cost. For in-chain loss the computation can also be reduced with proper thinning schemes.

## D IMPLEMENTATION DETAILS OF THE TRAINING LOSS

We visualize on the left panel of Figure 7 the unrolled computation scheme. We apply truncated back-propagate through time (BPTT) to train the sampler. Specifically, we manually stop the gradient flow through the input of  $\mathbb{D}$  and  $\mathbb{Q}$  matrices to avoid computing higher order gradients.

We also illustrate cross-chain in-chain training on the right panel of Figure 7. Cross-chain training encourages both fast convergence and low bias, provided that the samples are taken from parallel chains. On the other hand, in-chain training encourages sample diversity inside a chain. In practice, we might consider thinning the chains when performing in-chain training. Empirically this improves the Stein gradient estimator's accuracy as the samples are spread out. Computationally, this also prevents inverting big matrices for the Stein gradient estimator, and reduces the number of back-propagation operations. Another trick we applied is parallel chain sub-sampling: if all the chains are used, then there is less encouragement of single chain mixing, since the parallel chain samples can be diverse enough already to give reasonable gradient estimate.

Figure 7: (Left) The unrolled scheme of the meta sampler updates. Stop gradient operations are applied to the dashed arrows. (Right) A visualization of cross-chain in-chain training. The grey area represents samples across multiple chains, and we compute the cross chain loss for every 5 time steps. The purple area indicates the samples taken across time with sub-sampled chains 1 and 3. In this visualization the initial 15 samples are discarded for burn-in, and the thinning length is 1 (effectively no thinning).

## E INPUT PRE-PROCESSING

One potential challenge is that for different tasks and problem dimensions, the energy function, momentum and energy gradient can have very different scales and magnitudes. This affects the meta sampler’s generalization, for example, if training and test densities have completely different energy scales, then the meta sampler is likely to produce wrong strategies. This is especially the case when the meta sampler is generalized to much bigger networks or to very different datasets.

To mediate this issue, we propose to pre-process the inputs to both  $f_{\theta}$  and  $f_{\phi}$  networks to make it at similar scale as those in training task. Recall that the energy function  $U(\mathbf{x}) = -\sum_{n=1}^N \log p(\mathbf{y}_n | \mathbf{x}_n; \theta) - \log p(\mathbf{x}; \phi)$  where the prior  $\log p(\cdot)$  is often an isotropic Gaussian distribution. Thus the energy function scale linearly w.r.t both the dimensionality and the total number of observations  $N$ . Often the energy function is further approximated using mini-batches of datapoints. Putting them together, we propose pre-processing the energy as

$$\overline{U}(\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M \log p(\mathbf{y}_m | \mathbf{x}_m; \theta) + \frac{D_{\text{train}}}{ND_{\text{test}}} \log p(\mathbf{x}; \phi) \quad (26)$$

where  $D_{\text{train}}$  and  $D_{\text{test}}$  are the dimensionality of in the training task and the test task, respectively.

Importantly, for RNNs  $N$  represents the total sequence length, namely  $\sum_{n=1}^{N_{\text{data}}} T_n$ , where  $N_{\text{data}}$  is the total number of sequences and  $T_n$  is the sequence length for a datum. We also define  $M$  accordingly. The momentum and energy gradient magnitudes are estimated by simulating a randomly initialized meta sampler for short iterations. With these statistics we normalize both the momentum and the energy gradient to have roughly zero mean and unit variance.

## F EXPERIMENT SETUP

### F.1 TOY EXAMPLE

We train our meta sampler on a 10D uncorrelated Gaussian with mean 3 and randomly generated covariance matrix. We do not set any offset and additional frictions, i.e.  $\beta = 0$  and  $\gamma = 0$ . The noise estimation matrices are set to be 0 for both meta sampler and SGHMC. To mimic stochastic gradient, we manually inject Gaussian noise with zero mean and unit variance into  $\nabla U(\mathbf{x}) = \nabla U(\mathbf{x}) + \mathbf{N}(\mathbf{0}; \mathbf{I})$ . The functions  $f_{\theta}$  and  $f_{\phi}$  are represented by 1-hidden-layer MLPs with 40 hidden units. For training task, the meta sampler step size is 0.01. The initial positions are drawn from  $\text{Uniform}([0; 6]^D)$ . We train our sampler for 100 epochs and each epoch consists 4 x 100 steps. For every 100 steps, we update the  $\mathbf{D}$  matrices using Adam optimizer with learning rate 0.0005. Then we continue the updated sampler with last position and momentum until 4

sub-epochs are finished. We re-initialize the momentum and position. We use both cross-chain and in-chain losses. The Stein Gradient estimator uses RBF kernel with bandwidth chosen to be 0.5 times the median-heuristic estimated value. We unroll the Markov Chain for 20 steps before we manually stop the gradient. For cross-chain training, we take sampler across chain for each 2 time steps. For in-Chain, we discard initial 50 points for burn-in and sub-sample the chain with batch size 5. We thin the samples for every 3 steps. For both training and evaluation, we run 50 parallel Markov Chains.

The test task is to draw samples from a 20D correlated Gaussian with mean  $\mu$  and randomly generated covariance matrix. The step size is 0.025 for both meta sampler and SGHMC. To stabilize the meta sampler we also clamp the output values within  $[-5, 5]$ . The friction matrix for SGHMC is selected as

## F.2 BAYESIAN MLP MNIST

In MNIST experiment, we apply input pre-processing on energy function as in (26) and scale energy gradient by 70. Also, we scale  $\mu$  by 50 to account for sum of stochastic noise. The offset is selected as  $\frac{0.01}{N}$  as suggested by Chen et al. (2014), where  $\frac{\eta}{N}$  with  $\eta$  the per-batch learning rate. We also turn off the off-set and noise estimation, i.e.  $\beta = 0$  and  $\mathbf{B} = 0$ . We run 20 parallel chains for both training and evaluation. We only adopt the cross chain training with thinning samplers of 5 times step. We also use the finite difference technique during evaluation to speed-up computations.

### F.2.1 ARCHITECTURE GENERALIZATION (NT)

We train the meta sampler on a smaller BNN with architecture 784-20-10 and ReLU activation function, then test it on a larger one with architecture 784-40-40-10. In both cases the batch size is 500 following Chen et al. (2014). Both  $\mu$  and  $\mathbf{f}$  are parameterized by 1-hidden-layer MLPs with 10 units. The per-batch learning rate is 0.007. We train the sampler for 100 epochs and each one consists of 7 sub-epochs. For each sub-epoch, we run the sampler for 100 steps. We re-initialize and momentum after each epoch. To stabilize the meta sampler in evaluation, we first run the meta sampler with small per-batch learning rate 0.005 for 3 data epochs and clamp the values. After, we increase the per-batch learning rate to 0.018 with clipped  $\mathbf{f}$ . The learning rate for SGHMC is 0.01 for all times. For SGLD and PSGLD, they are 0.2 and  $1.4 \times 10^{-3}$  respectively. These step-sizes are tuned on MNIST validation data.

### F.2.2 NT + ACTIVATION FUNCTION GENERALIZATION

We modify the test network's activation function to sigmoid. We use almost the same settings as in network generalization tests, except that the per-batch learning rates are tuned again on validation data. For the meta sampler and SGHMC, they are 0.018 and 0.15. For SGLD and PSGLD, they are 0.2 and  $1.3 \times 10^{-2}$ .

### F.2.3 NT + DATASET GENERALIZATION

We train the meta sampler on ReLU network with architecture 784-20-5 to classify images 0-4, and test the sampler on ReLU network 784-40-40-5 to classify images 5-9. The settings are mostly the same as in network architecture generalization for both training and evaluation. One exception is again the per-batch learning rate for PSGLD, which is tuned as  $1.3 \times 10^{-3}$ . Note that even though we use the same per-batch learning rate as before, the discretization step-size is now different due to smaller training dataset, thus, will be automatically adjusted accordingly.

## F.3 BAYESIAN CONVOLUTIONAL NEURAL NETWORK ON CIFAR-10

CIFAR-10 dataset contains 50,000 training images with 10 labels and 10,000 test images. We train our meta sampler using smaller CNN classifier with two convolutional layers (3  $\times$  3  $\times$  8 and 3  $\times$  3  $\times$  8, no padding) and one fc layer of 50 hidden units. Therefore the dimensionality of is 15; 768. The training sampler discretization step-size is  $\frac{0.0007}{50000}$  and scaling term is  $= \frac{0.005}{50000}$ . To make it analogous to optimization methods, we use 0.007 as per-batch learning rate and 0.005

friction coefficient. The  $f_{\theta}$  and  $f_{\phi}$  are defined by 2-layer MLPs with 10 hidden units. We set the offset values to 0 for both  $\theta$  and  $\phi$ . Further, we scale up the output  $Df(\mathbf{z})$  by 10 and its gradient input  $U(\cdot)$  by 100. We scale up the energy input  $U(\cdot)$  to both  $f_{\theta}$  and  $f_{\phi}$  by 5. We train our meta sampler using 100 “meta epoch” with 5 data epoch and 500 batch size. Within each “meta epoch”, we repeat the following computation for 10 times: we run 50 parallel chains using the meta sampler for 50 iterations (0.5 dataset epoch), compute the loss function, and update the meta sampler’s parameters using Adam. We manually stop the gradient after 20 iterations. Then we start the next sub-epoch using the last  $\theta$  and  $\phi$ . After we finish all sub-epoch, we re-initialize the  $\theta$  and  $\phi$  using replay techniques with probability 0.15. The sub-sample chain number for in-chain loss is set to 5.

### F.3.1 THE REPLAY TECHNIQUE

Experience replay (Lin, 1993) is a technique broadly used in reinforcement learning literature. Inspired by this, in Bayesian CNN experiments we train the meta sampler in a similar way, and we found this replay technique particularly useful for more complicated dataset like CIFAR-10.

At the beginning of each “meta epoch”, each chain is initialized either with a specific state randomly chosen from a replay pool, or with a random state sampled from a Gaussian distribution. We use a pre-defined replay probability to control the replay strategy. The replay pool is updated after each sub-epoch, and it has a queue-like data structure of constant size, so that the old states are replaced by the new ones. Therefore, this replay technique is useful for both short-time and long-time horizon generalization. On one hand, the meta sampler can continue with previous states, allowing it to accommodate long-time horizon behavior. On the other hand, due to non-zero probability of random restart, the meta sampler can learn a better strategy for fast convergence. Therefore with this replay technique, the sampler can observe both burn-in and roughly-converged behavior, and this balance is controlled by the replay probability.

### F.3.2 ARCHITECTURE GENERALIZATION (NT)

For architecture generalization, the test CNN has two convolutional layers (3 × 3 × 16 and 3 × 3 × 16 × 16, no padding) and one fully connected layer with 100 hidden units. Thus, the dimensionality of  $\mathbf{z}$  is 61,478, roughly 4 times of the training dimension. We run 20 parallel chains in test time. We split the 50,000 training images into 45,000 training and 5,000 validation images, and tune the discretization step-size of each sampling and optimization methods on the validation set for 80 epochs. For test, we run the tuned samplers/optimizers for 200 data epoch (roughly 40 times longer than training) to ensure convergence. For the meta sampler, the per-batch learning rate is 0.001. For SGHMC, the per-batch is also 0.003 with friction coefficient 0.1. For SGLD, the per-batch learning rate is  $3 \times 10^{-3}$  as learning rate and 0.99 as moving average term. For optimization methods, we use learning rate 0.002 for Adam and 0.003 for SGD-M. The momentum term is 0.9. To prevent overfitting, we use weight penalty with coefficient 0.01.

### F.3.3 NT + SIGMOID GENERALIZATION

The test CNN has same architecture as in NT, except that it replaces all ReLU activation functions with sigmoid activations. We fix all other parameters for sampling method and only re-tune the step sizes using same setup as in NT. The per-batch rate for meta sampler, SGHMC, SGLD and PSGLD are 0.1, 0.03, 0.5 and 0.005 respectively. For optimization methods, the step size for Adam and SGD-M are 0.002 and 0.03 respectively.

### F.3.4 NT + DATASET GENERALIZATION

We split the CIFAR-10 training and test dataset according to the labels. We use training data with labels 0-4 for meta sampler training, training data with labels 5-9 for test CNN training, and test data with labels 5-9 for test CNN evaluation. Thus, the meta sampler has no access to the test task’s training and test data during sampler training. We train our sampler using the same scaling terms as in NT but reduce the discretization step-size to 0.0005. The rest setup is the same as in NT.

We use the same test CNN architecture and ReLU activation as in NT, and tune the learning rate using validation data. The step size for the meta sampler, SGHMC, SGLD and PSGLD are 0.1, 0.005,

Table 3: The basic statistics for 4 RNN datasets, bold figure represents large difference compared to others. Size is the number of data points. Avg. Time is the averaged sequence length. Energy scale is the rough scale of the train NLL when sampler converges.

	Piano	Muse	Nott	JSB
Size:train	87	524	694	229
Size:test	25	124	170	77
Avg. Time:train	872	467	254	60
Avg. Time:test	761	518	261	61
Energy scale:train	7:2	7	<b>2:5</b>	7:8

0.2 and 0.0018, respectively. For optimization methods, we use learning rate 0.002 and 0.003 for Adam and SGD-M respectively.

#### F.4 BAYESIAN RNN

The Piano data is selected as the training task, which is further split into training, validation and test subsets. We use batch-size 1, meaning that the energy and the gradient are estimated on a single sequence. The meta sampler uses similar neural network architectures as in MNIST tests. The training and evaluation per-batch learning rate for all the samplers is set to be 0.001 following Chen et al. (2016). We train the meta sampler for 40 epochs with 7 sub-epochs with only cross chain loss. Each sub-epochs consists 70 iterations. We scale the output by 20 and set  $\beta = 0.002$ , where  $\beta$  is defined in the same way as before. We use zero offset during training,  $\beta = 0$ . We apply input pre-processing for both  $\mathbf{h}_0$  and  $\mathbf{f}_0$ . To prevent divergence of the meta sampler at early training stage. We also set the constant of 100 to the  $\mathbf{h}_0$ . For dataset generalization, we tune the off-set value based on Piano validation set and transfer the tuned setting 1:5 to the other three datasets. For Piano architecture generalization, we do not tune any hyper-parameters including  $\beta$  and use exactly same settings as training. Exact gradient is used in RNN experiments instead of computing finite differences.

## G RNN DATASET DESCRIPTION

We list some data statistics in Table 3 which roughly indicates the similarity between datasets. Piano dataset is the smallest in terms of data number, however, the averaged sequence length is the largest. Muse dataset is similar to Piano in sequence length and energy scale but much larger in terms of data number. On the other hand, Nott dataset has very different energy scale compared to the other three. This potentially makes the generalization much harder due to inconsistent energy scale  $\beta$  and  $\mathbf{h}_0$ . For JSB, we notice a very short sequence length on average, therefore the GRU model is more likely to over-fit. Indeed, some algorithms exhibits significant over-fitting behavior on JSB dataset compared to other data (Santa is particularly severe).

## H ADDITIONAL PLOTS

### H.1 SHORT HORIZON PERFORMANCE COMPARISONS

We also run the samplers using the same settings as in MNIST experiments for a short period of time (500 iterations). We also compare to other optimization methods including momentum SGD (SGD-M) and Adam. We use the same per-batch learning rate for SGD-M and SGHMC as in MNIST experiment. For Adam, we use 0.002 for ReLU and 0.01 for Sigmoid network.

The results are shown in Figure 8. Meta sampler and Adam achieves the fastest convergence speed. This again confirms the faster convergence of the meta sampler especially at initial stages. We also provide additional contour plots (Figure 9) for MNIST experiments to demonstrate the strategy learned by  $\mathbf{h}_0$  for reference.

