
A Unified Particle-Optimization Framework for Scalable Bayesian Sampling

Changyou Chen

University at Buffalo (cchangyou@gmail.com)

Ruiyi Zhang, Wenlin Wang, Bai Li, Liqun Chen

Duke University

Abstract

There has been recent interest in developing scalable Bayesian sampling methods such as stochastic gradient MCMC (SG-MCMC) and Stein variational gradient descent (SVGD) for big-data analysis. A standard SG-MCMC algorithm simulates samples from a discrete-time Markov chain to approximate a target distribution, thus samples could be highly correlated, an undesired property for SG-MCMC. In contrary, SVGD directly optimizes a set of particles to approximate a target distribution, and thus is able to obtain good approximations with relatively much fewer samples. In this paper, we propose a principle particle-optimization framework based on Wasserstein gradient flows to unify SG-MCMC and SVGD, and to allow new algorithms to be developed. Our framework interprets SG-MCMC as particle optimization on the space of probability measures, revealing a strong connection between SG-MCMC and SVGD. The key component of our framework is several particle-approximate techniques to efficiently solve the original partial differential equations on the space of probability measures. Extensive experiments on both synthetic data and deep neural networks demonstrate the effectiveness and efficiency of our framework for scalable Bayesian sampling.

1 INTRODUCTION

Bayesian methods have been playing an important role in modern machine learning, especially in unsupervised learning (Kingma and Welling, 2014; Li et al., 2017), and recently in deep reinforcement learning (Houthoofd et al., 2016; Liu et al., 2017). When dealing with big data, two lines of research directions have been developed to scale

up Bayesian methods, *e.g.*, variational-Bayes-based and sampling-based methods. Stochastic gradient Markov chain Monte Carlo (SG-MCMC) is a family of scalable Bayesian learning algorithms designed to efficiently sample from a target distribution such as a posterior distribution (Welling and Teh, 2011; Chen et al., 2014; Ding et al., 2014; Chen et al., 2015). In principle, SG-MCMC generates samples from a Markov chain, which are used to approximate a target distribution. Under a standard setting, samples from SG-MCMC are able to match a target distribution exactly with an infinite number of samples (Teh et al., 2016; Chen et al., 2015). However, this is practically infeasible, as only a finite number of samples are obtained. Although nonasymptotic approximation bounds w.r.t. the number of samples have been investigated (Teh et al., 2016; Vollmer et al., 2016; Chen et al., 2015), there are no theory/algorithms to guide learning an optimal set of fixed-size samples/particles. This is an undesirable property of SG-MCMC, because in practice one often seeks to learn the optimal samples of a finite size that best approximate a target distribution.

A remedy for this issue is to adopt the idea of particle-based sampling methods, where a set of particles (or samples) are initialized from some simple distribution, followed by iterative updates to better approximate a target distribution. The updating procedure is usually done by optimizing some metrics such as a distance measure between the target distribution and the current approximation. There is not much work in this direction for large-scale Bayesian sampling, with an outstanding representative being the Stein variational gradient descent (SVGD) (Liu and Wang, 2016a). In SVGD, the update of particles are done by optimizing the KL-divergence between the empirical particle distribution and a target distribution, thus the samples are designed to be updated optimally to reduce the KL-divergence in each iteration. Because of this property, SVGD is found to perform better than SG-MCMC when the number of samples used to approximate a target distribution is limited, and has been

applied to other problems such as deep generative models (Feng et al., 2017) and deep reinforcement learning (Liu et al., 2017; Haarnoja et al., 2017; Zhang et al., 2018b).

Though often achieving comparable performance in practice, little work has been done on investigating connections between SG-MCMC and SVGD, and on developing particle-optimization schemes for SG-MCMC. In this paper, adopting ideas from Wasserstein-gradient-flow literature, we propose a unified particle-optimization framework for scalable Bayesian sampling. The idea of our framework is to work directly on the evolution of a density functions on the space of probability measures, *e.g.*, the Fokker-Planck equation in SG-MCMC. To make the evolution solution computationally feasible, particle approximations are adopted for densities, where particles can be optimized during the evolution process. Both SG-MCMC and SVGD are special cases of our framework, and are shown to be highly related. Notably, sampling with SG-MCMC becomes a deterministic particle-optimization problem as SVGD on the space of probability measures, overcoming the aforementioned correlated-sample issue. Furthermore, we are able to develop new *unified* particle-optimization algorithms by combing SG-MCMC and SVGD, which is less prone to high-dimension space and thus obtains better performance for large-scale Bayesian sampling. We conduct extensive experiments on both synthetic data and Bayesian learning of deep neural networks, verifying the effectiveness and efficiency of our proposed framework.

2 PRELIMINARIES

In this section, we review related concepts and algorithms for SG-MCMC, SVGD, and Wasserstein gradient flows (WGF) on the space of probability measures.

2.1 Stochastic gradient MCMC

Diffusion-based sampling methods Generating random samples from a distribution (*e.g.*, a posterior distribution) is one of the fundamental problems in Bayesian statistics, which has many important applications in machine learning. Traditional Markov Chain Monte Carlo methods (MCMC), such as the Metropolis–Hastings algorithm (Metropolis et al., 1953) produces unbiased samples from a desired distribution when the density function is known up to a normalizing constant. However, most of these methods are based on random walk proposals which suffer from high dimensionality and often lead to highly correlated samples. On the other hand, dynamics-based sampling methods such as the Metropolis adjusted Langevin algorithm (MALA) (Xifara et al., 2014) avoid this high degree of correlation by combining dynamical

systems with the Metropolis step. In fact, these dynamical systems are derived from a more general mathematical technique called diffusion process, or more specifically, Itô diffusion (Øksendal, 1985).

Specifically, our objective is to generate random samples from a posterior distribution $p(\boldsymbol{\theta} | \mathbf{X}) \propto p(\mathbf{X} | \boldsymbol{\theta})p(\boldsymbol{\theta})$, where $\boldsymbol{\theta} \in \mathbb{R}^r$ represents the model parameter, and $\mathbf{X} \triangleq \{\mathbf{x}_i\}_{i=1}^N$ represents the data. The canonical form is $p(\boldsymbol{\theta} | \mathbf{X}) = (1/Z) \exp(U(\boldsymbol{\theta}))$, where $U(\boldsymbol{\theta}) =$

$$\log p(\mathbf{X} | \boldsymbol{\theta}) + \log p(\boldsymbol{\theta}) \triangleq \sum_{i=1}^N \log p(\mathbf{x}_i | \boldsymbol{\theta}) + \log p(\boldsymbol{\theta})$$

is referred to as the potential energy based on an i.i.d. assumption of the model, and Z is the normalizing constant. In Bayesian sampling, the posterior distribution corresponds to the (marginal) stationary distribution of a (continuous-time) Itô diffusion, defined as a stochastic differential equation of the form:

$$d\boldsymbol{\Theta}_t = F(\boldsymbol{\Theta}_t)dt + g(\boldsymbol{\Theta}_t)d\mathcal{W}_t, \quad (1)$$

where t is the time index; $\boldsymbol{\Theta}_t \in \mathbb{R}^p$ represents the full variables in a dynamical system, and $\boldsymbol{\Theta}_t \supseteq \boldsymbol{\theta}_t$ (thus $p \geq r$) is potentially an augmentation of model parameter $\boldsymbol{\theta}$; $\mathcal{W}_t \in \mathbb{R}^p$ is p -dimensional Brownian motion. Functions $F : \mathbb{R}^p \rightarrow \mathbb{R}^p$ and $g : \mathbb{R}^p \rightarrow \mathbb{R}^p \times \mathbb{R}^p$ are assumed to satisfy the Lipschitz continuity condition (Ghosh, 2011). By Fokker-Planck equation (or the forward Kolmogorov equation) (Kolmogoroff, 1931; Risken, 1989), when appropriately designing the diffusion-coefficient functions $F(\cdot)$ and $g(\cdot)$, the stationary distribution of the corresponding Itô diffusion equals the posterior distribution of interest, $p(\boldsymbol{\theta} | \mathbf{X})$. For example, the 1st-order Langevin dynamic defines $\boldsymbol{\Theta} = \boldsymbol{\theta}$, and $F(\boldsymbol{\Theta}_t) = \frac{1}{2} \nabla_{\boldsymbol{\theta}} U(\boldsymbol{\theta})$, $g(\boldsymbol{\Theta}_t) = \mathbf{I}_r$; the 2nd-order Langevin diffusion defines $\boldsymbol{\Theta} = (\boldsymbol{\theta}, \mathbf{q})$, and $F(\boldsymbol{\Theta}_t) = \begin{pmatrix} \mathbf{q} \\ -B\mathbf{q} - \nabla_{\boldsymbol{\theta}} U(\boldsymbol{\theta}) \end{pmatrix}$, $g(\boldsymbol{\Theta}_t) = \sqrt{2B} \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_r \end{pmatrix}$ for a scalar $B > 0$; \mathbf{q} is an auxiliary variable known as the momentum (Chen et al., 2014; Ding et al., 2014).

Let the density of $\boldsymbol{\Theta}_t$ be μ_t , it is known μ_t is characterized by the Fokker-Planck (FP) equation (Risken, 1989):

$$\frac{\partial \mu_t}{\partial t} = -\nabla_{\boldsymbol{\Theta}} \cdot (\mu_t F(\boldsymbol{\Theta}_t)) + \nabla_{\boldsymbol{\Theta}} \nabla_{\boldsymbol{\Theta}} : (\mu_t \Sigma(\boldsymbol{\Theta}_t)) \quad (2)$$

where $\Sigma(\boldsymbol{\Theta}_t) \triangleq g(\boldsymbol{\Theta}_t)g^\top(\boldsymbol{\Theta}_t)$, $\mathbf{a} \cdot \mathbf{b} \triangleq \mathbf{a}^\top \mathbf{b}$ for vectors \mathbf{a} and \mathbf{b} , $\mathbf{A} : \mathbf{B} \triangleq \text{trace}(\mathbf{A}^\top \mathbf{B})$ for matrices \mathbf{A} and \mathbf{B} . The FP equation is the key to develop our particle-optimization framework for SG-MCMC. In the following, we focus on the simplest case of 1st-order Langevin dynamics if not stated explicitly, though the derivations apply to other variants.

Stochastic gradient MCMC SG-MCMC algorithms are discretized numerical approximations of the Itô dif-

fusion (I). They mitigate the slow mixing and non-scalability issues encountered in traditional MCMC algorithms by *i*) adopting gradient information of the posterior distribution, *ii*) using minibatches of the data in each iteration of the algorithm to generate samples, and *iii*) ignoring the rejection step as in standard MCMC. To make the algorithms scalable in a big-data setting, three developments will be implemented based on the Itô diffusion: *i*) define appropriate functions F and g in the Itô-diffusion formula so that the (marginal) stationary distributions coincide with the target posterior distribution $p(\boldsymbol{\theta} | \mathbf{X})$; *ii*) replace F or g with unbiased stochastic approximations to reduce the computational complexity, *e.g.*, approximating F with a random subset of the data instead of using the full data. For example, in the 1st-order Langevin dynamics, $\nabla_{\boldsymbol{\theta}} U(\boldsymbol{\theta})$ could be *approximated* by an unbiased estimator with a subset of data:

$$\nabla_{\boldsymbol{\theta}} \tilde{U}(\boldsymbol{\theta}) \triangleq \nabla \log p(\boldsymbol{\theta}) + \frac{N}{n} \sum_{i=1}^n \nabla_{\boldsymbol{\theta}} \log p(\mathbf{x}_{\pi_i} | \boldsymbol{\theta}) \quad (3)$$

where π is a size- n random subset of $\{1, 2, \dots, N\}$, leading to the first SG-MCMC algorithm in machine learning – stochastic gradient Langevin dynamics (SGLD) (Welling and Teh, 2011); and *iii*) solve the generally intractable continuous-time Itô diffusions with a numerical method, *e.g.*, the Euler method (Chen et al., 2015). For example, this leads to the following update in SGLD:

$$\boldsymbol{\theta}_{\ell} = \boldsymbol{\theta}_{\ell-1} + \nabla_{\boldsymbol{\theta}} \tilde{U}(\boldsymbol{\theta}_{\ell-1})h + \sqrt{2h} \boldsymbol{\delta}_{\ell},$$

where h means the stepsize, ℓ indexes the samples, $\boldsymbol{\delta}_{\ell} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is a random sample from an isotropic normal distribution. After running the algorithm for L steps, the collection of samples $\{\boldsymbol{\theta}_{\ell}\}_{\ell=1}^L$ are used to approximate the unknown posterior distribution $\frac{1}{2} e^{U(\boldsymbol{\theta})}$.

2.2 Stein variational gradient descent

Different from SG-MCMC, SVGD initializes a set of particles which are iteratively updated so that the empirical particle distribution approximates the posterior distribution. Specifically, we consider a set of particles $\{\boldsymbol{\theta}^{(i)}\}_{i=1}^M$ drawn from some distribution q . SVGD tries to update these particles by doing gradient descent on the interactive particle system via

$$\boldsymbol{\theta}^{(i)} \leftarrow \boldsymbol{\theta}^{(i)} + h\phi(\boldsymbol{\theta}^{(i)}), \quad \phi = \arg \max_{\phi \in \mathcal{F}} \left\{ \frac{\partial}{\partial h} \text{KL}(q_{[h\phi]} || p) \right\}$$

where ϕ is a function perturbation direction chosen to minimize the KL divergence between the updated density $q_{[h\phi]}$ estimated by the particles and the posterior $p(\boldsymbol{\theta} | \mathbf{X})$ (p for short). Since $\text{KL}(q || p)$ is convex in q , global optimum of $q = p$ can be guaranteed. SVGD considers \mathcal{F} as the unit ball of a vector-valued reproducing kernel Hilbert

space (RKHS) \mathcal{H} associated with a kernel $\kappa(\boldsymbol{\theta}, \boldsymbol{\theta}')$. In such as setting, Liu and Wang (2016b) shown:

$$-\frac{\partial}{\partial h} \text{KL}(q_{[h\phi]} || p)|_{h=0} = \mathbb{E}_{\boldsymbol{\theta} \sim q} [\text{trace}(\Gamma_p \phi(\boldsymbol{\theta}))], \quad (4)$$

with $\Gamma_p \phi(\boldsymbol{\theta}) \triangleq \nabla_{\boldsymbol{\theta}} U(\boldsymbol{\theta})^{\top} \phi(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \cdot \phi(\boldsymbol{\theta})$,

where Γ_p is called the Stein operator. Assuming that the update function $\phi(\boldsymbol{\theta})$ is in a RKHS with kernel $\kappa(\cdot, \cdot)$, it was shown in (Liu and Wang, 2016b) that (4) is maximized with:

$$\phi(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta} \sim q} [\kappa(\boldsymbol{\theta}, \boldsymbol{\theta}') \nabla_{\boldsymbol{\theta}} U(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \kappa(\boldsymbol{\theta}, \boldsymbol{\theta}')]. \quad (5)$$

When approximating the expectation $\mathbb{E}_{\boldsymbol{\theta} \sim q}[\cdot]$ with empirical particle distribution and adopting stochastic gradients, we arrive at the following updates for the particles (ℓ denotes the iteration number): $\boldsymbol{\theta}_{\ell+1}^{(i)} = \boldsymbol{\theta}_{\ell}^{(i)} +$

$$\frac{h}{M} \sum_{j=1}^M \left[\kappa(\boldsymbol{\theta}_{\ell}^{(j)}, \boldsymbol{\theta}_{\ell}^{(i)}) \nabla_{\boldsymbol{\theta}_{\ell}^{(j)}} \tilde{U}(\boldsymbol{\theta}_{\ell}^{(j)}) + \nabla_{\boldsymbol{\theta}_{\ell}^{(j)}} \kappa(\boldsymbol{\theta}_{\ell}^{(j)}, \boldsymbol{\theta}_{\ell}^{(i)}) \right] \quad (6)$$

SVGD applies updates (6) repeatedly, moving the samples to a target distribution p .

2.3 Wasserstein Gradient Flows

For a better motivation of WGF, we start from gradient flows defined on the Euclidean space.

Gradient flows on the Euclidean space For a smooth function $E : \mathbb{R}^r \rightarrow \mathbb{R}$, and a starting point $\boldsymbol{\theta}_0 \in \mathbb{R}^r$, the gradient flow of $E(\boldsymbol{\theta})$ is defined as the solution of the differential equation: $\frac{d\boldsymbol{\theta}}{dt} = -\nabla E(\boldsymbol{\theta}(t))$, s.t. $\boldsymbol{\theta}(0) = \boldsymbol{\theta}_0$. This is a standard Cauchy problem (Rulla, 1996), endowed with a unique solution if ∇E is Lipschitz continuous. When E is non-differentiable, the gradient is replaced with its subgradient, defined as $\partial E(\boldsymbol{\theta}) \triangleq \{\mathbf{p} \in \mathbb{R}^r : F(\boldsymbol{\theta}') \geq F(\boldsymbol{\theta}) + \mathbf{p} \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta}), \forall \boldsymbol{\theta}' \in \mathbb{R}^r\}$. Note $\partial E(\boldsymbol{\theta}) = \{\nabla E(\boldsymbol{\theta})\}$ if E is differentiable at $\boldsymbol{\theta}$. In this case, the gradient flow formula above is replaced with: $\frac{d\boldsymbol{\theta}}{dt} \in -\partial E(\boldsymbol{\theta}(t))$.

Wasserstein gradient flows Let $\mathcal{P}(\Omega)$ denote the space of probability measures on $\Omega \subset \mathbb{R}^r$. WGF is an extension of gradient flows in Euclidean space by lifting the definition onto the space of probability measures. Formally, let $\mathcal{P}(\Omega)$ be endowed with a Riemannian geometry induced by the 2nd-order Wasserstein distance, *i.e.*, the curve length between two elements (two distributions) is defined as:

$$W_2^2(\mu, \nu) \triangleq \inf_{\gamma} \left\{ \int_{\Omega \times \Omega} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2 d\gamma(\boldsymbol{\theta}, \boldsymbol{\theta}') : \gamma \in \Gamma(\mu, \nu) \right\}$$

where $\Gamma(\mu, \nu)$ is the set of joint distributions over $(\boldsymbol{\theta}, \boldsymbol{\theta}')$ such that the two marginals equal μ and ν , respectively. The Wasserstein distance can be explained as an optimal-transport problem, where one wants to transform elements

in the domain of μ to ν with a minimum cost (Villani 2008). The term $\|\theta - \theta'\|_2^2$ represents the cost to transport θ in μ to θ' in ν , and can be replaced by a general metric $c(\theta, \theta')$ in a metric space. If μ is absolutely continuous w.r.t. the Lebesgue measure, there is a unique optimal transport plan from μ to ν , i.e., a mapping $\mathcal{T} : \mathbb{R}^r \rightarrow \mathbb{R}^r$ pushing elements in the domain of μ onto ν satisfying $\mathcal{T}_\# \mu = \nu$. Here $\mathcal{T}_\# \mu$ denotes the pushforward measure of μ . The Wasserstein distance is equivalently reformulated as: $W_2^2(\mu, \nu) \triangleq \inf_{\mathcal{T}} \left\{ \int_{\Omega} c(\theta, \mathcal{T}(\theta)) d\mu(\theta) \right\}$.

Let $\{\mu_t\}_{t \in [0,1]}$ be an absolutely continuous curve in $\mathcal{P}(\Omega)$ with finite second-order moments. We consider to define the change of μ_t 's by investigating $W_2^2(\mu_t, \mu_{t+h})$. Motivated by the Euclidean-space case, this is reflected by a vector field, $\mathbf{v}_t(\theta) \triangleq \lim_{h \rightarrow 0} \frac{\mathcal{T}(\theta_t) - \theta_t}{h}$ called the *velocity of the particle*. A gradient flow can be defined on $\mathcal{P}(\Omega)$ correspondingly (Ambrosio et al., 2005).

Lemma 1 *Let $\{\mu_t\}_{t \in [0,1]}$ be an absolutely-continuous curve in $\mathcal{P}(\Omega)$ with finite second-order moments. Then for a.e. $t \in [0, 1]$, the above vector field \mathbf{v}_t defines a gradient flow on $\mathcal{P}(\Omega)$ as: $\partial_t \mu_t + \nabla \cdot (\mathbf{v}_t \mu_t) = 0$.*

The gradient flow describes the evolution of a functional E , which is a lifted version of the function in the case of Euclidean space in Section 2.3 to the space of probability measures. E maps a probability measure μ to a real value, i.e., $E : \mathcal{P}(\Omega) \rightarrow \mathbb{R}$. We will focus on the case where E is convex in this paper, which is enough considering gradient flows for SG-MCMC and SVGD, though the theory applies to a more general λ -convex energy functional setting (Ambrosio et al., 2005). It can be shown that \mathbf{v}_t in Lemma 1 has the form $\mathbf{v}_t = -\nabla \frac{\delta E}{\delta \mu_t}(\mu_t)$ (Ambrosio et al., 2005), where $\frac{\delta E}{\delta \mu_t}$ is called the *first variation* of E at μ_t . Based on this, gradient flows on $\mathcal{P}(\Omega)$ can be written

$$\partial_t \mu_t = -\nabla \cdot (\mathbf{v}_t \mu_t) = \nabla \cdot \left(\mu_t \nabla \left(\frac{\delta E}{\delta \mu_t}(\mu_t) \right) \right). \quad (7)$$

Remark 1 *Intuitively, an energy functional E characterizes the landscape structure (appearance) of the corresponding manifold in $\mathcal{P}(\Omega)$, and the gradient flow (7) defines a geodesic path on this manifold. Usually, by choosing appropriate E , the landscape is convex, e.g., for the cases of both SG-MCMC and SVGD described below. This provides a theoretical guarantee on the optimal convergence of a gradient flow.*

3 PARTICLE-OPTIMIZATION-BASED SAMPLING

In this section, we interpret the continuous versions of both SG-MCMC and SVGD as WGFs, followed by several techniques for particle optimization in the next section. In the following, μ_t denotes the distribution of θ_t .

3.1 SVGD as WGF

The continuous-time and infinite-particle limit of SVGD with full gradients, denoted as SVGD $^\infty$, is known to be a special instance of the Vlasov equation in nonlinear partial-differential-equation literature (Liu, 2017):

$$\partial_t \mu_t = \nabla \cdot ((\mathbf{W} * \mu_t) \mu_t), \quad (8)$$

where $(\mathbf{W} * \mu_t)(\theta) \triangleq \int \mathbf{W}(\theta - \theta') \mu_t(\theta') d\theta'$ is the convolutional operator applied for some function $\mathbf{W} : \mathbb{R}^r \rightarrow \mathbb{R}$. To specify SVGD $^\infty$, we generalize the convolutional operator, and consider \mathbf{W} as a function with two input arguments, i.e.,

$$(\mathbf{W} * \mu_t)(\theta) \triangleq \int \mathbf{W}(\theta, \theta') \mu_t(\theta') d\theta'.$$

Under this setting, we can specify the function $\mathbf{W}(\cdot, \cdot)$ for SVGD $^\infty$ as

$$\begin{aligned} \mathbf{W}(\theta, \theta') &\triangleq \nabla_{\theta'} \log p(\theta' | \mathbf{X}) \kappa(\theta', \theta) + \nabla_{\theta'} \kappa(\theta', \theta) \\ &= \nabla_{\theta'} [p(\theta' | \mathbf{X}) \kappa(\theta, \theta')] / p(\theta'). \end{aligned} \quad (9)$$

As will be shown in Section 4, \mathbf{W} in (9) naturally leads to the SVGD algorithm, without the need to derive from an RKHS perspective.

Proposition 2 *The stationary distribution of (8) is $\lim_{t \rightarrow \infty} \mu_t \triangleq \mu = p(\theta | \mathbf{X})$.*

To interpret SVGD $^\infty$ as a WGF, we need to specify two quantities, the energy functional and an underlying metric to measure distances between density functions.

Energy functional and distance metric of SVGD $^\infty$

There are two ways to derive energy functionals for SVGD $^\infty$, depending on the underlying metrics for probability distributions. When adopting the WGF framework where W_2 is used as the underlying metric, according to (7), the energy functional E_s must satisfy

$$\begin{aligned} \nabla_{\theta} \left(\frac{\delta E_s}{\delta \mu_t}(\mu_t) \right) &= \mathbf{W}(\theta, \theta') * \mu_t \\ &= \mathbb{E}_{\theta' \sim \mu_t} [\nabla_{\theta'} [p(\theta' | \mathbf{X}) K(\theta, \theta')] / p(\theta' | \mathbf{X})]. \end{aligned} \quad (10)$$

In general, there is no close-form solution for the above equation. Alternatively, Liu (2017) proved another form of the energy functional by defining a different distance metric on the space of probability measures, called \mathcal{H} -Wasserstein distance:

$$W_{\mathcal{H}}(q_1, q_2) \triangleq \inf_{\phi_t, \mu_t} \left\{ \int_0^1 \|\phi_t\|_{\mathcal{H}} dt, \text{ s.t. } \mu_t = -\nabla_{\theta} \cdot (\phi_t \mu_t), \mu_0 = q_1, \mu_1 = q_2 \right\}, \quad (11)$$

where $\phi_t \triangleq \mathbf{W} * \mu_t$, and $\|\cdot\|_{\mathcal{H}}$ is the norm in the Hilbert space induced by $\kappa(\cdot, \cdot)$. Under this metric, the underlying energy functional is proved to be the standard KL-divergence between μ_t and p , e.g.,

$$E_s = \text{KL}(\mu_t, p(\cdot | \mathbf{X})).$$

As can be seen in Section 4, this interpretation allows one to derive SVGD, a particle-optimization-based algorithm to approximate the continuous-time equation (8).

3.2 SG-MCMC as WGF

The continuous-time limit of SG-MCMC, when considering gradients to be exact, corresponds to standard Itô diffusions. We consider the Itô diffusion of SGLD for simplicity, *e.g.*,

$$d\theta_t = \frac{1}{2} \nabla U(\theta_t) dt + dW. \quad (12)$$

Energy functional The energy functional for SG-MCMC is easily seen by noting that the corresponding FP equation (2) is in the gradient-flow form of (7). Specifically, the energy functional E is defined as:

$$E(\mu) \triangleq \underbrace{- \int U(\theta) \mu(\theta) d\theta}_{E_1} + \underbrace{\int \mu(\theta) \log \mu(\theta) d\theta}_{E_2} \quad (13)$$

Note E_2 is the energy functional of a pure Brownian motion (*e.g.*, $U(\theta) = 0$ in (12)). We can verify (13) by showing that it satisfies that FP equation. According to (7), the first variation of E_1 and E_2 is calculated as

$$\frac{\delta E_1}{\delta \mu} = -U, \quad \frac{\delta E_2}{\delta \mu} = \log \mu + 1. \quad (14)$$

Substituting (14) into (7) recovers the FP equation (2) for the Itô diffusion (12).

4 PARTICLE OPTIMIZATION

An efficient way to solve the generally infeasible WGF formula (7) is to adopt numerical methods with particle approximation. With a little abuse of notation but for conciseness, we do not distinguish subscripts t and ℓ for the particle θ , *i.e.*, θ_t denotes the continuous-time version of the particle, while θ_ℓ denotes the discrete-time version. We develop several techniques to approximate different types of WGF for SG-MCMC and SVGD. In particle approximation, the continuous density μ_t is approximated by a set of M particles $(\theta_t^{(i)})_{i=1}^M$ that evolve over time t with weights $(m_i)_{i=1}^M$ such that $\sum_{i=1}^M m_i = 1$, *i.e.*, $\mu_t \approx \sum_{i=1}^M m_i \delta(\theta_t^{(i)})$, where $\delta(\theta_t^{(i)}) = 1$ when $\theta = \theta_t^{(i)}$ and 0 otherwise. Typically m_i 's are chosen at the beginning and fixed over time, thus we assume $m_i = \frac{1}{M}$ and rewrite $\mu_t \approx \frac{1}{M} \sum_{i=1}^M \delta(\theta_t^{(i)})$ in the following for simplicity. We investigate two types of particle-approximation methods in the following, discrete gradient flows and by blob methods.

Particle approximation by discrete gradient flows

Denote $\mathcal{P}_s(\mathbb{R}^r)$ be the space of probability measures with finite 2nd-order moments. Define the following optimization problem with stepsize h :

$$J_h(\mu) \triangleq \arg \min_{\nu \in \mathcal{P}_s(\mathbb{R}^d)} \left\{ \frac{1}{2h} W_2^2(\mu, \nu) + E(\nu) \right\}. \quad (15)$$

A discrete gradient flow of the continuous one in (7) up to time T is the composition of a sequence of the solutions $(\tilde{\mu}_\ell)_{\ell=1}^{T/h}$ of (15), *i.e.*,

$$\tilde{\mu}_\ell \triangleq J_h(\tilde{\mu}_{\ell-1}) = J_h(J_h(\cdots \mu_0)) \triangleq J_h^\ell \mu_0. \quad (16)$$

One can show that when $h \rightarrow 0$, the discrete gradient flow (16) converges to the true flow (7) for all ℓ . Specifically, let $\partial E(\mu)$ be the set of Wasserstein subdifferential of E at μ , *i.e.*, $\xi \in \partial E(\mu)$ if $\partial_t \mu = \xi$ is satisfied. Define $|\partial E|(\mu) = \min\{\|\xi\|_{L^2(\mu)} : \xi \in \partial E(\mu)\}$ to be the minimum norm of the elements in $\partial E(\mu)$. We have

Lemma 3 (Craig (2014)) *Assume E is proper, coercive and lower semicontinuous (specify in Section B of the Supplementary Material (SM)). For an μ_0 and $t \geq 0$, as $\frac{T}{h} \rightarrow \infty$, the discrete gradient sequence $\tilde{\mu}_{T/h} \triangleq J_h^{T/h} \mu_0$ converge uniformly in t to a compact subset of $[0, +\infty)$, and $W_2^2(\tilde{\mu}_{T/h}, \mu_T) \leq \sqrt{3} |\partial E|(\mu) \sqrt{Th}$.*

Lemma 3 suggests the discrete gradient flow can approximate the original WGF arbitrarily well if a small enough stepsize h is adopted. Consequently, one solves (16) through a sequence of optimization procedures to update the particles. We will derive a particle-approximation method for the W_2 term in (15), which allows us to solve SG-MCMC efficiently. However, this technique is not applicable to SVGD, as we neither have an explicit form of the energy functional in (10) when adopting the W_2 metric, nor have an explicit form for the metric $W_{\mathcal{H}}$ in (11) when adopting the KL-divergence as the energy functional. Fortunately, this can be solved by the second approximation method called blob methods.

Particle approximation by blob methods The name of blob methods comes from the classical fluids literature, where instead of evolving the density in (7), one evolves all particles on a grid with time-spacing h (Carrillo et al., 2017). Specifically, note the function \mathbf{v}_t in (7) represents velocity of particles via transportation map \mathcal{T} , thus solving a WGF is equivalent to evolving the particles along their velocity in each iteration. Formally, one can prove

Proposition 4 (Craig and Bertozzi (2016)) *Let $\mu_0 \approx \frac{1}{M} \sum_{i=1}^M \delta(\theta_0^{(i)})$. Assume \mathbf{v}_t in (7) is well-defined and continuous w.r.t. each $\theta_t^{(i)}$ at time t . Then solving the PDE (7) reduces to solving a system of ordinary differential equations for the locations of the Dirac masses:*

$$d\theta_t^{(i)} / dt = -\mathbf{v}_t(\theta_t^{(i)}). \quad (17)$$

Proposition 4 suggests evolving each particle along the directions defined by \mathbf{v}_t , eliminating the requirement to know an explicit form of the energy functional. In the following, we apply the above particle-optimization techniques to derive algorithms for SVGD and SG-MCMC.

4.1 A particle-optimization algorithm for SVGD

As mentioned above, discrete-gradient-flow approximation does not apply to SVGD. We thus rely on the blob method. From Section 3.1, \mathbf{v}_t in SVGD is defined as $\mathbf{v}_t(\boldsymbol{\theta}) = (\mathbf{W} * \mu_t)(\boldsymbol{\theta})$. When $\mu_t(\boldsymbol{\theta})$ is approximated by particles, $\mathbf{v}_t(\boldsymbol{\theta}_t^{(i)})$ is simplified as:

$$\mathbf{v}_t(\boldsymbol{\theta}_t^{(i)}) = \frac{1}{M} \sum_{j=1}^M \mathbf{W}(\boldsymbol{\theta}_t^{(i)}, \boldsymbol{\theta}_t^{(j)}).$$

As a result, with the definition of \mathbf{W} in (9), updating $\{\boldsymbol{\theta}_t^{(i)}\}$ by time discretizing (17) recovers the update equations for standard SVGD in (6).

4.2 Particle-optimization algorithms for SG-MCMC

Both the discrete-gradient-flow and the blob methods can be applied for SG-MCMC, which are detailed below.

Particle optimization with discrete gradient flows

We first specify Lemma 3 in the case of SG-MCMC in Lemma 5, which is known as the Jordan-Kinderlehrer-Otto scheme (Jordan et al., 1998).

Lemma 5 (Jordan et al. (1998)) *Assume that $p(\boldsymbol{\theta}_t | \mathbf{X}) \leq C_1$ is infinitely differentiable, and $\|\nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{\theta} | \mathbf{X})\| \leq C_2 (1 + C_1 - \log p(\boldsymbol{\theta} | \mathbf{X}))$ ($\forall \boldsymbol{\theta}$) for some constants $\{C_1, C_2\}$. Let $T = hK$ with K the number of iterations, $\tilde{\mu}_0$ be an arbitrary distribution with same support as $p(\boldsymbol{\theta} | \mathbf{X})$, and $\{\tilde{\mu}_k\}_{k=1}^K$ be the solution of the functional optimization problem:*

$$\tilde{\mu}_k = \arg \min_{\mu \in \mathcal{P}_s(\mathbb{R}^r)} \text{KL}(\mu \| p) + \frac{1}{2h} W_2^2(\tilde{\mu}_{k-1}, \mu). \quad (18)$$

Then $\tilde{\mu}_K$ converges to μ_T in the limit of $h \rightarrow 0$, i.e., $\lim_{h \rightarrow 0} \tilde{\mu}_K = \mu_T$, where μ_T is the solution of the FP equation (2) at time T .

According to Lemma 5, it is apparent that SG-MCMC can be implemented by iteratively solving the optimization problem in (18). However, particle approximations for both terms in (18) are challenging. In the following, we develop efficient techniques to solve the problem.

First, rewrite the optimization problem in (18) as

$$\min_{\mu \in \mathcal{P}_s(\mathbb{R}^r)} \underbrace{-\mathbb{E}_{\mu}[\log p(\boldsymbol{\theta} | \mathbf{X})]}_{F_1} + \underbrace{\mathbb{E}_{\mu}[\log \mu] + \frac{1}{2h} W_2^2(\tilde{\mu}_{k-1}, \mu)}_{F_2}$$

We aim at deriving gradient formulas for both the F_1 and F_2 terms under a particle approximation in order to perform gradient descent for the particles. Let $\mu \approx \frac{1}{M} \sum_{i=1}^M \delta(\boldsymbol{\theta}^{(i)})$. The gradient of F_1 is easily approximated as

$$\frac{\partial F_1}{\partial \boldsymbol{\theta}^{(i)}} \approx -\nabla_{\boldsymbol{\theta}^{(i)}} \log p(\boldsymbol{\theta}^{(i)} | \mathbf{X}). \quad (19)$$

To approximate the gradient for F_2 , let p_{ij} denote the joint distribution of the particle-pair $(\boldsymbol{\theta}^{(i)}, \boldsymbol{\theta}_{k-1}^{(j)})$. Note $\mathbb{E}_{\mu}[\log \mu]$ is minimized when the particles $\{\boldsymbol{\theta}^{(i)}\}$ are uniformly distributed. In other words, the marginal distribution vector $(\sum_j p_{ij})_i$ is a uniform distribution. Combining $\mathbb{E}_{\mu}[\log \mu]$ with the definition of W_2 , calculating F_2 is equivalent to solving the following optimization problem:

$$P \triangleq \{p_{ij}\} = \arg \min_{p_{i,j}} \sum_{i,j} p_{ij} d_{ij} \quad (20)$$

$$s.t. \quad \sum_j p_{ij} = \frac{1}{M}, \quad \sum_i p_{ij} = \frac{1}{M},$$

where $d_{ij} \triangleq \|\boldsymbol{\theta}^{(i)} - \boldsymbol{\theta}_{k-1}^{(j)}\|^2$. We can further enforce the joint distribution $\{p_{ij}\}$ to have maximum entropy by introducing a regularization term $\mathbb{E}_{p_{ij}}[\log p_{ij}]$, which is stronger than the regularizer enforced for the marginal distribution above. After introducing Lagrangian multipliers $\{\alpha_i, \beta_j\}$ to deal with the constraints in (20), we arrive at the dual problem:

$$\begin{aligned} \max \mathcal{L}^D(\{p_{ij}\}, \{\alpha_i\}, \{\beta_j\}) &= \lambda \sum_{i,j} p_{ij} \log p_{ij} + p_{ij} d_{ij} \\ &+ \sum_i \alpha_i \left(\sum_j p_{ij} - \frac{1}{M} \right) + \sum_j \beta_j \left(\sum_i p_{ij} - \frac{1}{M} \right), \end{aligned}$$

where λ is the weight for the regularizer. The optimal p_{ij} 's can be obtained by applying KKT conditions to set the derivative w.r.t. p_{ij} to be zero, ending up with the following form:

$$p_{ij}^* = u_i e^{-d_{ij}/\lambda} v_j,$$

where $u_i \triangleq e^{-\frac{1}{2} - \frac{\alpha_i}{\lambda}}$, $v_j \triangleq e^{-\frac{1}{2} - \frac{\beta_j}{\lambda}}$. As a result, the particle gradients on F_2 can be approximated as

$$\begin{aligned} \frac{\partial F_2}{\partial \boldsymbol{\theta}^{(i)}} &\approx - \frac{\sum_j u_i v_j d_{ij} e^{-d_{ij}/\lambda}}{\partial \boldsymbol{\theta}^{(i)}} \\ &= \sum_j 2u_i v_j \left(\frac{d_{ij}}{\lambda} - 1 \right) e^{-d_{ij}/\lambda} (\boldsymbol{\theta}^{(i)} - \boldsymbol{\theta}_{k-1}^{(j)}). \end{aligned} \quad (21)$$

Theoretically, we need to adaptively update $\{u_i, v_j\}$ as well to ensure the constraints in (20). In practice, however, we use a fixed scaling factor γ to approximate $u_i v_j$ for the sake of simplicity.

Particle gradients are obtained by combining (19) and (21), which are then used to update the particles $\{\boldsymbol{\theta}^{(i)}\}$ by standard gradient descent. Intuitively, (19) encourages particles move to local modes while (21) regularizes particle interactions. Different from SVGD, our scheme imposes both attractive and repulsive forces for the particles. Specifically, by inspecting (21), we can conclude that: *i)* When $\boldsymbol{\theta}^{(i)}$ is far from a previous particle $\boldsymbol{\theta}_k^{(j)}$, i.e., $\frac{d_{ij}}{\lambda} > 1$, $\boldsymbol{\theta}^{(i)}$ is pulled close to $\{\boldsymbol{\theta}_k^{(j)}\}$ with force proportional to $(\frac{d_{ij}}{\lambda} - 1)e^{-d_{ij}/\lambda}$; *ii)* when $\boldsymbol{\theta}^{(i)}$ is close enough to a previous particle $\boldsymbol{\theta}_k^{(j)}$, i.e., $\frac{d_{ij}}{\lambda} < 1$, $\boldsymbol{\theta}^{(i)}$ is pushed away, preventing it from collapsing to $\boldsymbol{\theta}_k^{(j)}$.

Particle optimization with blob methods The idea of blob methods can also be applied to particle approximation for SG-MCMC, which require the velocity vector field \mathbf{v}_t . According to (13), this is calculated as: $\mathbf{v}_t(\boldsymbol{\theta}) = -\nabla_{\boldsymbol{\theta}} \frac{\delta(E_1+E_2)}{\delta\mu} = -\nabla_{\boldsymbol{\theta}} U - \nabla_{\boldsymbol{\theta}} \mu / \mu$. Unfortunately, direct application of particle approximation is infeasible because the term $\nabla_{\boldsymbol{\theta}} \mu$ is undefined with discrete μ . To tackle this problem, we adopt the idea in Carrillo et al. (2017) to approximate the energy functional E_2 in (13) as: $E_2 \approx \int \mu(\boldsymbol{\theta}) \log(\mu * K)(\boldsymbol{\theta}) d\boldsymbol{\theta}$, where $K(\cdot, \cdot)$ is another kernel function to smooth out μ . Consequently, based on Carrillo et al. (2017), the velocity \mathbf{v}_t can be calculated as (details in Section C of the SM):

$$\mathbf{v}_t(\boldsymbol{\theta}) = -\nabla_{\boldsymbol{\theta}} U - \sum_{j=1}^n \nabla_{\boldsymbol{\theta}_t^{(j)}} K(\boldsymbol{\theta}, \boldsymbol{\theta}_t^{(j)}) / \sum_k K(\boldsymbol{\theta}_t^{(j)}, \boldsymbol{\theta}_t^{(k)}) - \sum_{j=1}^n \nabla_{\boldsymbol{\theta}_t^{(j)}} K(\boldsymbol{\theta}, \boldsymbol{\theta}_t^{(j)}) / \sum_{k=1}^n K(\boldsymbol{\theta}, \boldsymbol{\theta}_t^{(k)}) \quad (22)$$

Given \mathbf{v}_t , particle updates can be obtained by solving (17) numerically as in SVGD. By inspecting the formula of \mathbf{v}_t in (22), the last two terms both act as repulsive forces. Interestingly, the mechanism is similar to SVGD, but with adaptive force between different particle pairs.

5 THE GENERAL RECIPE

Based on the above development, a more general particle-optimization framework is proposed by combining the PDEs of both SG-MCMC and SVGD. As a result, we propose the following PDE to drive evolution of densities

$$\frac{\partial \mu_t}{\partial t} = -\nabla_{\boldsymbol{\theta}} \cdot (\mu_t F(\boldsymbol{\theta}_t)) + \lambda_1 \nabla_{\boldsymbol{\theta}} \cdot ((\mathbf{W} * \mu_t) \mu_t) + \lambda_2 \nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}} : (\mu_t g(\boldsymbol{\theta}_t) g^{\top}(\boldsymbol{\theta}_t)) \quad (23)$$

where λ_1 and λ_2 are two constants. It is easily seen that to ensure the stationary distribution of (23) to be equal to $p(\boldsymbol{\theta} | \mathbf{X})$, the following condition must be satisfied:

$$\nabla_{\boldsymbol{\theta}} \cdot (p(\boldsymbol{\theta} | \mathbf{X}) F(\boldsymbol{\theta})) = \lambda_1 \nabla_{\boldsymbol{\theta}} \cdot ((\mathbf{W} * p(\boldsymbol{\theta} | \mathbf{X})) p(\boldsymbol{\theta} | \mathbf{X})) + \lambda_2 \nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}} : (p(\boldsymbol{\theta} | \mathbf{X}) g(\boldsymbol{\theta}) g^{\top}(\boldsymbol{\theta})) \quad (24)$$

There are many feasible choices for the functions and parameters $\{F(\boldsymbol{\theta}), \mathbf{W}, g(\boldsymbol{\theta}), \lambda_1, \lambda_2\}$ to satisfy (24). However, the verification procedure might be complicated given the present of a convolutional term in (24). We recommend the following choices for simplicity:

- $F(\boldsymbol{\theta}) = \frac{1}{2}U(\boldsymbol{\theta})$, $\mathbf{W} = 0$, $g(\boldsymbol{\theta}) = \mathbf{I}$ and $\lambda_2 = 1$: this reduces to the Wasserstein-based SGLD with particle optimization. Specifically, when the discrete-gradient-flow approximation is adopted, the algorithm is denoted as w -SGLD; whereas when the blob method is adopted, it is denoted as w -SGLD-B.

- $F(\boldsymbol{\theta}) = 0$, $g(\boldsymbol{\theta}) = 0$, \mathbf{W} is defined as (9): this reduces to standard SVGD.

- $F(\boldsymbol{\theta}) = \frac{1}{2}U(\boldsymbol{\theta})$, $g(\boldsymbol{\theta}) = \mathbf{I}$, \mathbf{W} is defined as (9), and $\lambda_2 = 1$: this is the combination of SGLD and SVGD, and is called particle interactive SGLD, denoted as PI-SGLD or π -SGLD.

It is easy to verify that condition (24) is satisfied for all the above three particle-optimization algorithms. Furthermore, particle updates are readily developed by applying either the discrete-gradient-flow or blob-based methods.

6 RELATED PARTICLE-BASED MCMC METHODS

There have been related particle-based MCMC algorithms. Representative methods are sequential Monte Carlo (SMC) (Moral et al. 2006), particle MCMC (PMCMC) (Andrieu et al. 2010) and many variants. In SMC, particles are sample from a proposal distribution, and the corresponding weights are updated by a resampling step. PMCMC extends SMC by sampling from an extended distribution interacted with a MH-rejection step. Compared to our framework, their proposal distributions are typically hard to choose; furthermore, optimality of the particles from both methods can not be guaranteed. Furthermore, the methods are typically much more computationally expensive. Recently, Dai et al. (2016) proposed a particle-based MCMC algorithm by approximating a target distribution with weighted kernel density estimator, which updates particle weights based on likelihoods of the corresponding particles. This approach is theoretically sound but lacks an underlying geometry interpretation. Finally, we note that w -SGLD has been successfully applied to reinforcement learning recently for improved policy optimization (Zhang et al. 2018a).

7 EXPERIMENTS

We verify our framework on a set of experiments, including a number of toy experiments and applications to Bayesian sampling of deep neural networks (DNNs).

7.1 Demonstrations

Toy Distributions We compare various sampling methods on multi-mode toy examples, *i.e.*, SGLD, SVGD, w -SGLD, w -SGLD-B and π -SGLD. We aim to sample from four unnormalized 2D densities $p(z) \propto \exp\{U(z)\}$, with detailed functional form provided in the SM. We optimize/sample 2000 particles to approximate target distributions. The results are shown in Figure 1. It can be seen from Figure 1 that though SGLD maintains good asymptotic properties, it is inaccurate to approximate distributions with only a few samples; in some case, the

samples cannot even cover all the modes. Interestingly, all other particle-optimization-based algorithms successfully find all the modes and fit the distributions well. w -SGLD is good at finding modes, but worse at modeling the correct variance due to difficulty of controlling the balance between attractive and repulsive forces between particles. w -SGLD-B is better than w -SGLD at modeling the distribution variance, performing similarly to SVGD and π -SGLD. Even though, we note that w -SGLD is very useful when the number of particles is small, which fits a distribution better, as shown in Section E of the SM.




Figure 1: Illustration of different algorithms on toy distributions. Each column is a distribution case. 1st row: Ground truth; 2nd row: standard SGLD; 3rd row: w -SGLD; 4th row: w -SGLD-B; 5th row: SVGD; 6th row: π -SGLD.

Bayesian Logistic regression We next compare the three variants of our framework (*i.e.* SVGD, w -SGLD and w -SGLD-B) on a simple logistic-regression task with quantitative evaluations. We use the same model, data and experimental settings as Liu and Wang (2016a). The Covertypes dataset contains 581,012 data points and 54 features. We perform 5 runs for each setting and report the mean of testing accuracies/log-likelihoods. Figure 2 plots both test accuracies and test log-likelihoods w.r.t. the




Figure 2: Test accuracies (left) and log-likelihoods (right) v.s. iterations for SVGD, w -SGLD and w -SGLD-B.

number of training iterations. It is clearly that while all methods converge to the same accuracy/likelihood level, both w -SGLD and w -SGLD-B converge slightly faster than SVGD. In addition, w -SGLD and w -SGLD-B have similar convergence behaviors, thus we only use w -SGLD in the DNN experiments below.

Parameter Sensitivity

Now we study the role of hyper-parameters in π -SGLD: the number of particles M and the scaling factor γ to




Figure 3: Impact of W_2^2 factor γ and particle number M .

replace the $u_i v_j$ -term in (21). We use the same dataset and model as the above experiment. Figure 3 plots test accuracies along with different parameter settings. As expected, the best performance is achieved with appropriate scale of W_2^2 . The performance keep improving with increasing particles. Interestingly, the Wasserstein regularization is more important when the number of particles is small, demonstrating the superiority when approximate distributions with very few particles.

7.2 Applications on deep neural networks

We conduct experiments for Bayesian learning of DNNs. Different from traditional optimization for DNNs, we are interested in modeling weight uncertainty of neural networks, an important topic that has been well explored (Hernández-Lobato and Adams, 2015; Blundell et al., 2015a; Li et al., 2016; Louizos and Welling, 2016). We assign priors to the weights, which are simple isotropic Gaussian priors in our case, and perform posterior sampling with the proposed particle-optimization-based algorithms, as well as other standard algorithms such as SGLD and SGD. We use the RMSprop optimizer for feed-forward networks (FNN), and Adam for for convolutional neural networks (CNNs) and recurrent neural

networks (RNNs). For all methods, we use a RBF kernel $K(\theta, \theta') = \exp(-\|\theta - \theta'\|_2^2/h)$, with the bandwidth set to $h = \text{med}^2/\log M$. Here med is the median of the pairwise distance between particles. All experiments are conducted on a single TITAN X GPU.

Feed-forward Neural Networks We perform the classification tasks on the standard MNIST dataset. A two-layer model 784-X-X-10 with ReLU activation function is used, with X being the number of hidden units for each layer. The training epoch is set to 100. The test errors are reported in Table 1. Not surprisingly, Bayesian methods generally perform better than their optimization counterparts. The new π -SGLD which combines w -SGLD and SVGD improves both methods with little computational overhead. In addition, w -SGLD seems to perform better than SVGD in this case, partially due to a better asymptotic property mentioned in (Liu, 2017). Furthermore, standard SGLD which is based on MCMC obtains higher test errors compared to particle-optimization-based algorithms, partially due to the correlated-sample issue discussed in the introduction. See (Blundell et al., 2015b) for details on the other methods in Table 1.

Table 1: Classification error of FNN on MNIST.

| Method | Test Error | |
|---------------|--------------|--------------|
| | 400-400 | 800-800 |
| π -SGLD | 1.36% | 1.30% |
| w -SGLD | 1.44% | 1.37% |
| SVGD | 1.53% | 1.40% |
| SGLD | 1.64% | 1.41% |
| RMSprop | 1.59% | 1.43% |
| RMSspectral | 1.65% | 1.56% |
| SGD | 1.72% | 1.47% |
| BPB, Gaussian | 1.82% | 1.99% |
| SGD, dropout | 1.51% | 1.33% |

Convolution Neural Networks We use the CIFAR-10 dataset to test our framework on CNNs. We adopt a CNN of three convolution layers, using 3×3 filter size with C64-C128-C256 channels, and 2×2 max-pooling after each convolution layer. Our implementation adopts batch normalization, drop out and data augmentation to improve the performance. Training losses and test accuracies are presented in Table 2. Consistently, π -SGLD outperforms all other algorithms in terms of test accuracy. ADAM obtains a better training loss but worse test accuracy, indicating worse generalization ability of the optimization-based methods compared to Bayesian methods.

Recurrent Neural Networks For RNNs, we run standard language models. Experiments are presented on three publicly available corpora: APNEWS, IMDB and BNC. APNEWS is a collection of Associated Press news articles from 2009 to 2016. IMDB is a set of movie re-

Table 2: Classification error of CNN on CIFAR-10.

| Method | Training Loss | Test Accuracy |
|-------------|---------------|---------------|
| ADAM | 23.80 | 86.76% |
| SVGD | 30.57 | 88.72% |
| SGLD | 28.52 | 88.64% |
| w -SGLD | 31.26 | 88.80% |
| π -SGLD | 25.06 | 89.52% |

views collected by (Maas et al., 2011), and BNC (BNC Consortium, 2007) is the written portion of the British National Corpus, which contains excerpts from journals, books, letters, essays, memoranda, news and other types of text. These datasets can be downloaded from Github*.

Table 3: Perplexity of language model on three corpora.

| Method | APNEWS | IMDB | BNC |
|-------------|--------------|--------------|--------------|
| SGD | 64.13 | 72.14 | 102.89 |
| SGLD | 63.01 | 68.12 | 95.13 |
| SVGD | 61.64 | 69.25 | 94.99 |
| w -SGLD | 61.22 | 67.41 | 93.68 |
| π -SGLD | 59.83 | 67.04 | 92.33 |

We follow the standard set up as (Wang et al., 2017). Specifically, we lower case all the word tokens and filter out word tokens that occur less than 10 times. All the datasets are divided into training, development and testing sets. For the language model set up, we consider a 1-layer LSTM model with 600 hidden units. The sequence length is fixed to be 30. In order to alleviate overfitting, dropout with a rate of 0.4 is used in each LSTM layer. Results in terms of test perplexities are presented in Table 3. Again, we see that π -SGLD performs best among all algorithms, and w -SGLD is slightly better than SVGD, both of which are better than other algorithms.

8 CONCLUSION

We propose a unified particle-optimization framework for large-scale Bayesian sampling. Our framework defines gradient flows on the space of probability measures, and uses particles to approximate the corresponding densities. Consequently, solving gradient flows reduces to optimizing particles on the parameter space. Our framework includes the standard SVGD as a special case, and also allows us to develop efficient particle-optimization algorithms for SG-MCMC, which is highly related to SVGD. Extensive experiments are conducted, demonstrating the effectiveness and efficiency of our proposed framework. Interesting future work includes designing more practically efficient variants of the proposed particle-optimization framework, and developing theory to study general convergence behaviors of the algorithms, in addition to the asymptotic results presented in (Liu, 2017).

*<https://github.com/jhlau/topically-driven-language-model>

REFERENCES

- Ambrosio, L., Gigli, N., and Savaré, G. (2005). *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics ETH Zürich.
- Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015a). Weight uncertainty in neural networks. In *ICML*.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015b). Weight uncertainty in neural networks. In *ICML*.
- BNC Consortium, B. (2007). The british national corpus, version 3 (bnc xml edition). *Distributed by Bodleian Libraries, University of Oxford, on behalf of the BNC Consortium*. URL:<http://www.natcorp.ox.ac.uk/>.
- Carrillo, J. A., Craig, K., and Patacchini, F. S. (2017). A blob method for diffusion. (arXiv:1709.09195).
- Chen, C., Ding, N., and Carin, L. (2015). On the convergence of stochastic gradient MCMC algorithms with high-order integrators. In *NIPS*.
- Chen, T., Fox, E. B., and Guestrin, C. (2014). Stochastic gradient Hamiltonian Monte Carlo. In *ICML*.
- Craig, K., editor (2014). *The Exponential Formula for the Wasserstein Metric*. PhD thesis, The State University of New Jersey.
- Craig, K. and Bertozzi, A. L. (2016). A blob method for the aggregation equation. *Mathematics of Computation*, 85(300):1681–1717.
- Dai, B., He, N., Dai, H., and Song, L. (2016). Provable Bayesian inference via particle mirror descent. In *AISTATS*.
- Ding, N., Fang, Y., Babbush, R., Chen, C., Skeel, R. D., and Neven, H. (2014). Bayesian sampling using stochastic gradient thermostats. In *NIPS*.
- Feng, Y., Wang, D., and Liu, Q. (2017). Learning to draw samples with amortized stein variational gradient descent. In *UAI*.
- Ghosh, A. P. (2011). *Backward and Forward Equations for Diffusion Processes*. Wiley Encyclopedia of Operations Research and Management Science.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. (2017). Reinforcement learning with deep energy-based policies. In *ICML*.
- Hernández-Lobato, J. M. and Adams, R. P. (2015). Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *ICML*.
- Houthoofd, R., Chen, X., Duan, Y., Schulman, J., De Turck, F., and Abbeel, P. (2016). Vime: Variational information maximizing exploration. In *NIPS*.
- Jordan, R., Kinderlehrer, D., and Otto, F. (1998). The variational formulation of the Fokker-Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational Bayes. In *ICLR*.
- Kolmogoroff, A. (1931). Some studies in machine learning using the game of checkers. *Mathematische Annalen*, 104(1):415–458.
- Li, C., Chen, C., Carlson, D., and Carin, L. (2016). Preconditioned stochastic gradient Langevin dynamics for deep neural networks. In *AAAI*.
- Li, C., Liu, H., Chen, C., Pu, Y., Chen, L., Heno, R., and Carin, L. (2017). ALICE: Towards understanding adversarial learning for joint distribution matching. In *NIPS*.
- Liu, Q. (2017). Stein variational gradient descent as gradient flow. In *NIPS*.
- Liu, Q. and Wang, D. (2016a). Stein variational gradient descent: A general purpose Bayesian inference algorithm. In *NIPS*.
- Liu, Q. and Wang, D. (2016b). Stein variational gradient descent: A general purpose bayesian inference algorithm. In *NIPS*.
- Liu, Y., Ramachandran, P., Liu, Q., and Peng, J. (2017). Stein variational policy gradient. In *UAI*.
- Louizos, C. and Welling, M. (2016). Structured and efficient variational deep learning with matrix Gaussian posteriors. In *ICML*.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. (2011). Learning word vectors for sentiment analysis. In *ACL*.
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., and Teller, E. (1953). Equation of State Calculations by Fast Computing Machines. *Journal of Chemical Physics*, 21:1087–1092.
- Moral, P. D., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B*, 68(3):411–436.
- Øksendal, B., editor (1985). *Stochastic Differential Equations*. Springer-Verlag, Berlin.
- Risken, H. (1989). *The Fokker-Planck equation*. Springer-Verlag, New York.
- Rulla, J. (1996). Error analysis for implicit approximations to solutions to Cauchy problems. *SIAM Journal on Numerical Analysis*, 33(1):68–87.
- Teh, Y. W., Thiery, A. H., and Vollmer, S. J. (2016). Consistency and fluctuations for stochastic gradient Langevin dynamics. *JMLR*, 17(1):193–225.
- Villani, C. (2008). *Optimal transport: old and new*. Springer Science & Business Media.
- Vollmer, S. J., Zygalakis, K. C., and Teh, Y. W. (2016). (exploration of the (Non-)asymptotic bias and variance of stochastic gradient Langevin dynamics. *JMLR*, 1:1–48.
- Wang, W., Gan, Z., Wang, W., Shen, D., Huang, J., Ping, W., Satheesh, S., and Carin, L. (2017). Topic compositional neural language model. *arXiv preprint arXiv:1712.09783*.
- Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient Langevin dynamics. In *ICML*.
- Xifara, T., Sherlock, C., Livingstone, S., Byrne, S., and Girolami, M. (2014). Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Statistics & Probability Letters*, 91:14–19.
- Zhang, R., Chen, C., Li, C., and Carin, L. (2018a). Policy optimization as wasserstein gradient flows. In *ICML*.
- Zhang, R., Li, C., Chen, C., and Carin, L. (2018b). Learning structural weight uncertainty for sequential decision-making. In *AISTATS*.