Sampling from Non-Log-Concave Distributions via Stochastic Variance-Reduced Gradient Langevin Dynamics

Difan Zou
Department of Computer Science
University of California,
Los Angeles

Pan Xu
Department of Computer Science
University of California,
Los Angeles

Quanquan Gu
Department of Computer Science
University of California,
Los Angeles

Abstract

We study stochastic variance reduction-based Langevin dynamic algorithms, SVRG-LD and SAGA-LD (Dubey et al., 2016), for sampling from non-log-concave distributions. Under certain assumptions on the log density function, we establish the convergence guarantees of SVRG-LD and SAGA-LD in 2-Wasserstein distance. More specifically, we show that both SVRG-LD and SAGA-LD require $\widetilde{O}\big(n+n^{3/4}/\epsilon^2+n^{1/2}/\epsilon^4\big)\cdot\exp\big(\widetilde{O}(d+\gamma)\big)$ stochastic gradient evaluations to achieve ϵ-accuracy in 2-Wasserstein distance, which outperforms the $\widetilde{O}\big(n/\epsilon^4\big)\cdot\exp\big(\widetilde{O}(d+\gamma)\big)$ gradient complexity achieved by Langevin Monte Carlo Method (Raginsky et al., 2017). Experiments on synthetic data and real data back up our theory.

1 INTRODUCTION

In the past decade, there has been an increasing interest in applying gradient based Markov Chain Monte Carlo (MCMC) methods for sampling from posterior distributions in Bayesian machine learning (Neal et al., 2011; Welling and Teh, 2011; Ahn et al., 2012; Chen et al., 2014; Ma et al., 2015; Cheng et al., 2017). In detail, this class of MCMC methods is based on the Langevin dynamics, which is described by the following stochastic differential equation (SDE)

$$\mathrm{d}\boldsymbol{X}(t) = -\nabla F(\boldsymbol{X}(t))\mathrm{d}t + \sqrt{2/\gamma}\mathrm{d}\boldsymbol{B}(t), \tag{1.1}$$

where $\gamma > 0$ is the inverse temperature parameter and $\{\boldsymbol{B}(t)\}_{t\geq 0}$ is the standard Brownian motion in $\mathbb{R}^d$.

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89.

Under certain assumptions on the drift term $\nabla F(\mathbf{x})$, the distribution of $\boldsymbol{X}(t)$ can be described by Fokker-Planck equation, and is able to converge to an invariant stationary distribution $\pi \propto \exp\big(-\gamma F(\mathbf{x})\big)$ (Chiang et al., 1987). In Bayesian inference, one aims to sample the target distribution with the form $\pi \propto \exp\big(-\gamma F(\mathbf{x})\big)$, and a typical way is to apply Euler-Maruyama discretization (Kloeden and Platen, 1992) to (1.1), which gives rise to the celebrated Langevin Monte Carlo (LMC) method (Roberts and Tweedie, 1996a),

$$\boldsymbol{X}_{k+1} = \boldsymbol{X}_k - \nabla F(\boldsymbol{X}_k)\eta + \sqrt{2\eta/\gamma}\boldsymbol{\epsilon}_k, \tag{1.2}$$

where $\boldsymbol{\epsilon}_k$ follows a standard multivariate normal distribution, and $\eta > 0$ denotes the step size. When the target distribution is strongly log-concave, i.e., function $F(\mathbf{x})$ is strongly convex, the convergence property of LMC has been widely studied based on total variation (Dalalyan, 2014; Durmus and Moulines, 2015, 2016) and 2-Wasserstein (Dalalyan, 2017; Dalalyan and Karagulyan, 2017) distances. On the other hand, for many machine learning problems involving extremely large amount of data, the function $F(\mathbf{x})$ on the drift term of (1.1) can be written as an average of n component functions

$$F(\mathbf{x}) = \frac{1}{n}\sum_{i=1}^{n} f_i(\mathbf{x}),$$

where $f_i(\mathbf{x}) : \mathbb{R}^d \to \mathbb{R}$ is the negative log likelihood function on the i-th example. When the data sample size n is enormous, the computation of the full gradient $\nabla F(\boldsymbol{X})$ in LMC is often very expensive. To overcome this computational burden, one resorts to using stochastic gradient to approximate the drift term in (1.1), which gives rise to the celebrated stochastic gradient Langevin dynamics (SGLD) method (Welling and Teh, 2011). In practice, the SGLD algorithm has achieved great success in Bayesian learning (Welling and Teh, 2011; Ahn et al., 2012) and Bayesian deep learning (Chaudhari et al., 2016; Ye et al., 2017). However, the SGLD algorithm requires more iteration steps

to achieve a high sampling precision compared with LMC due to the large variance of stochastic gradients. In order to alleviate this issue as well as save the gradient computation, Dubey et al. (2016) incorporated the idea of variance reduction (Johnson and Zhang, 2013; Reddi et al., 2016) into SGLD, and proposed two types of stochastic variance reduced algorithms based on gradient Langevin dynamics, namely SVRG-LD and SAGA-LD. Recently, Chatterji et al. (2018) proved the convergence rate of SVRG-LD and SAGA-LD in 2-Wasserstein distance when the target distribution is strongly log-concave, which characterizes the feasible regime where SVRG-LD and SAGA-LD outperform LMC and SGLD. The convergence rate of SVRG-LD was further improved by Zou et al. (2018b) recently. However, the current convergence analyses (Chatterji et al., 2018; Zou et al., 2018b) of stochastic variance-reduced gradient Langevin dynamics are mostly restricted to the strongly log-concave distributions, except Dubey et al. (2016); Chen et al. (2017). Nevertheless, Dubey et al. (2016); Chen et al. (2017) only investigated the mean square error of the sample path average. It is of more interest to establish the nonasymptotic convergence guarantee in terms of certain distance between the target distribution and that of the current iterate, which provides a fine-grained characterization of the sampling algorithms.

In this paper, we provide convergence analyses of SVRG-LD and SAGA-LD in 2-Wasserstein distance for non-log-concave target distributions. Different from the analysis of sampling from strongly log-concave distributions, the contraction property of 2-Wasserstein distance along the Langevin diffusion (1.1) no longer holds, which poses a great challenge for our analysis and makes existing proof techniques (Chatterji et al., 2018) for strongly log-concave distribution not applicable to our case. To address this challenge, we provide a new proof technique by extending the idea of Raginsky et al. (2017) for analyzing SGLD in nonconvex optimization. More specifically, our proof technique is based on a coupled Brownian motion between the discrete-time Markov chain and a continuous-time Markov chain generated by (1.1) and decomposes the 2-Wasserstein distance between the target distribution and that of the current iterate into two parts: the 2-Wasserstein distance between distributions of the current iterate and the corresponding continuous-time Markov Chain, and the distance between the distribution of the position in the coupled Markov chain and its stationary distribution, i.e., the target distribution π.

Our Contributions The major contributions of this paper are highlighted as follows.

- We study the SVRG-LD and SAGA-LD methods for sampling from non-log-concave distributions and prove their nonasymptotic convergence to the target distribution in terms of 2-Wasserstein distance. Specifically, we show that both SVRG-LD and SAGA-LD require $\widetilde{O}\big(n + n^{3/4}/\epsilon^2 + n^{1/2}/\epsilon^4\big) \cdot \exp\big(\widetilde{O}(d+\gamma)\big)$ stochastic gradient evaluations to achieve ϵ-accuracy, where n is the number of samples, γ is the inverse temperature and d is the problem dimension, which outperforms the gradient complexities of LMC and SGLD.
- We conduct experiments on both synthetic and real-world data to compare different first-order Langevin methods (SVRG-LD, SAGA-LD, SGLD, LMC) for sampling from non-log-concave distributions. The comparison suggests that the SVRG-LD and SAGA-LD have similar performance, and attain faster mixing time and perform better than their counterparts even when the target distribution is non-log-concave.

Notation We denote a deterministic vector by lower case bold symbol $\mathbf{x}$ and a random vector by upper case italicized bold symbol $\boldsymbol{X}$. We also use $\boldsymbol{X}_k$ (with subscript k) to denote the iterate of a discrete-time algorithm and $\boldsymbol{X}(t)$ (with index t in a parenthesis) to denote the continuous-time random process.For a vector $\mathbf{x} \in \mathbb{R}^d$, we denote by $\|\mathbf{x}\|_2$ the Euclidean norm. For a matrix $\mathbf{X}$, we denote $\|\mathbf{X}\|_F$ as the Frobenius norm.For a random vector $\boldsymbol{X} \in \mathbb{R}^d$, we denote its probability distribution function by $P(\boldsymbol{X})$. We denote by $\mathbb{E}_u(\boldsymbol{X})$ the expectation of $\boldsymbol{X}$ under probability measure u. We denote the 2-Wasserstein distance between two probability measures u and v as

$$\mathcal{W}_2^2(u,v) = \inf_{\zeta \in \Gamma(u,v)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\boldsymbol{X}_u - \boldsymbol{X}_v\|_2^2 \mathrm{d}\zeta(\boldsymbol{X}_u, \boldsymbol{X}_v),$$

where the infimum is over all joint distributions ζ with u and v being its marginal distributions. We denote by $\mathrm{KL}(p_1||p_2)$ the KL-divergence between probability measures p_1 and p_2. We use $a_n = O(b_n)$ to denote that $a_n \leq C b_n$ for some universal constant $C > 0$, and use $a_n = \widetilde{O}(b_n)$ to hide some logarithmic terms of b_n. We also use $a \wedge b$ to denote $\min\{a, b\}$.

2 RELATED WORK

In this section, we briefly review the literature on generic Langevin dynamics based algorithms.

Langevin Monte Carlo (LMC) methods, which are derived from discretizing Langevin dynamics, have been widely used for approximate sampling in Bayesian inference. For example, Dalalyan (2014) proved that the distribution of the last step in LMC converges to the

stationary distribution in $O(d/\epsilon^2)$ iterations in terms of total variation distance. Durmus and Moulines (2015) improved the results by showing the same result holds for any starting point and established similar bounds for the Wasserstein distance. Recently Dalalyan (2017) improved the existing results in terms of the 2-Wasserstein distance and provide further insights on the close relation between approximate sampling and optimization. Bubeck et al. (2015) analyzed sampling from log-concave distributions with compact support via projected LMC. Brosse et al. (2017) proposed a proximal LMC algorithm. Note that the Euler discretization on SDEs introduces a bias, and might fail to converge to the target distribution (Roberts and Tweedie, 1996a,b). An effective way to address this issue is incorporating the metropolis hasting correction step (Hastings, 1970) into LMC, which gives rise to metropolis adjusted Langevin algorithm (MALA) (Roberts and Rosenthal, 1998). Following this line of research, Bou-Rabee and Hairer (2012) provided nonasymptotic bounds on the mixing time of MALA, but the explicit dependence on the dimension d and target accuracy remains implicit. Eberle et al. (2014) established a clearer mixing time bound of MALA in terms of a modified Wasserstein distance for log-concave densities. Dwivedi et al. (2018) investigated MALA for strongly log-concave densities, and proved a linear rate of convergence in total variation distance.

Due to the increasing amount of data in modern machine learning problems, stochastic gradient Langevin dynamics (SGLD) (Welling and Teh, 2011; Ahn et al., 2012; Ma et al., 2015) has received extensive attentions in Bayesian learning. Vollmer et al. (2016) analyzed the nonasymptotic bias and variance of the SGLD algorithm by using Poisson equations. Dalalyan and Karagulyan (2017) proved $\widetilde{O}(d\sigma^2/\epsilon^2)$ convergence rate for SGLD in 2-Wasserstein distance when the target distribution is strongly log-concave. Moreover, Neal et al. (2011) introduced fictitious momentum term in Hamilton dynamics, which gives rise to Hamiltonian Monte Carlo (HMC) method. Similar to SGLD, stochastic gradient Hamiltonian Monte Carlo (SGHMC) (Chen et al., 2014) was proposed to overcome the limitation of gradient evaluation over extremely large datasets, and demonstrated better performance in learning Bayesian neural networks and online Bayesian matrix factorization (Chen et al., 2014). Under a similar framework, Chen et al. (2014) studied the stochastic MCMC method with higher-order integrator in terms of the MSE of the average sample path. Cheng et al. (2017) proposed underdamped MCMC method and proved the convergence guarantee in 2-Wasserstein distance when the target distribution is strongly log-concave. Although SGLD and SGHMC have acheived great success in Bayesian learning, the large variance of stochastic gradient may lead to unavoidable bias due to the lack of metropolis hasting (MH) correction. To overcome this, Whye Teh et al. (2014) proposed to apply decreasing step size to alleviate the bias, and proved the asymptotic rate of SGLD in terms of MSE. Betancourt (2015) pointed out that stochastic HMC may also lead to poor sampling performance, and there exists a tradeoff between the step size selection and acceptance probability in MH correction. This issue has been addressed by Dang et al. (2017) where they proposed a modified HMC algorithm that uses a subset of data to estimate both the dynamics and the subsequent MH acceptance probability.

The other direction to alleviate the variance of stochastic gradient and further save gradient computation is applying variance-reduction technique. Dubey et al. (2016) proposed a variance-reduced stochastic gradient Langevin dynamics for Bayesian posterior inference, and proved that their method improves the mean square error upon SGLD. Since that, many attempts have been made to incorporate variance reduction technique into Langevin based algorithms. Baker et al. (2017) applied zero variance control variates to stochastic MCMC method, and showed that such technique is able to reduce the computational cost of stochastic gradient Langevin dynamics to $O(1)$. Chatterji et al. (2018) studied two variants of variance-reduced stochastic Langevin dynamics proposed in Dubey et al. (2016), and proved their convergence guarantee for strongly log-concave distributions. Moreover, by replacing the full gradient in the outer loop of SVRG-LD with a subsampled one, Chen et al. (2017) and Zou et al. (2018b) studied the convergence rate of subsampled SVRG-LD method in MSE and 2-Wasserstein distance, respectively. The variance-reduced Hamiltonian Monte Carlo has also been investigated recently in Zou et al. (2018a); Li et al. (2018).

In Table 1, we summarize the gradient complexity[1] of LMC, SGLD, SVRG-LD and SAGA-LD in 2-Wasserstein distance for sampling from strongly log-concave and non-log-concave densities. To the best of our knowledge, there is no convergence result in 2-Wasserstein distance for sampling from general log-concave densities using Langevin dynamics based algorithms. It should be noted that for sampling from a non-log-concave the dependence on dimension d is inevitably exponential. In fact, it is proved in Bovier et al. (2004) that the lower bound of metastable exit time of SDE is exponential in d when the nonconvex function F in (1.2) has multiple local minima and saddle points.

[1]Gradient complexity is defined as the number of stochastic gradient evaluations.

Table 1: Gradient complexities to converge to the stationary distribution in 2-Wasserstein distance. Note that Raginsky et al. (2017) shows that SGLD dose not converge in 2-Wasserstein distance for non-log-concave densities.

	Strongly log-concave[2]	Non-log-concave
LMC	$\widetilde{O}\Big(\frac{nd^{1/2}}{\epsilon}\Big)$(Dalalyan, 2017)	$\widetilde{O}\big(\frac{n}{\epsilon^4}\big)\cdot e^{\widetilde{O}(d)}$(Raginsky et al., 2017)
SGLD	$\widetilde{O}\Big(\frac{d}{\epsilon^2}\Big)$(Dalalyan, 2017)	–
SVRG-LD	$\widetilde{O}\Big(n+\frac{n^{1/2}d^{1/2}}{\epsilon}\Big)$(Zou et al., 2018b)	$\widetilde{O}\Big(n+\frac{n^{3/4}}{\epsilon^2}+\frac{n^{1/2}}{\epsilon^4}\Big)\cdot e^{\widetilde{O}(d)}$**(This paper)**
SAGA-LD	$\widetilde{O}\Big(n+\frac{n^{1/2}d^{1/2}}{\epsilon}\Big)$ (Chatterji et al., 2018)	$\widetilde{O}\Big(n+\frac{n^{3/4}}{\epsilon^2}+\frac{n^{1/2}}{\epsilon^4}\Big)\cdot e^{\widetilde{O}(d)}$**(This paper)**

Algorithm 1 Stochastic Variance-Reduced Gradient Langevin Dynamics (SVRG-LD)

1: **input:** step size $\eta>0$; batch size B; epoch length m; inverse temperature parameter $\gamma>0$
2: **initialization:** $\boldsymbol{X}_0=\boldsymbol{0}$, $\widetilde{\boldsymbol{X}}^{(0)}=\boldsymbol{X}_0$
3: **for** $s=0,1,\ldots,(K/m)$ **do**
4: $\quad\widetilde{\boldsymbol{G}}=\nabla F(\widetilde{\boldsymbol{X}}^{(s)})$
5: $\quad$**for** $\ell=0,\ldots,m-1$ **do**
6: $\quad\quad k=sm+\ell$
7: $\quad\quad$randomly pick a subset I_k from $\{1,\ldots,n\}$ of size $|I_k|=B$; randomly draw $\boldsymbol{\epsilon}_k\sim N(\boldsymbol{0},\mathbf{I}_{d\times d})$
8: $\quad\quad\widetilde{\nabla}_k=\frac{1}{B}\sum_{i_k\in I_k}\big(\nabla f_{i_k}(\boldsymbol{X}_k)-\nabla f_{i_k}(\widetilde{\boldsymbol{X}}^{(s)})+\widetilde{\boldsymbol{G}}\big)$
9: $\quad\quad\boldsymbol{X}_{k+1}=\boldsymbol{X}_k-\eta\widetilde{\nabla}_k+\sqrt{2\eta/\gamma}\boldsymbol{\epsilon}_k$
10: $\quad$**end for**
11: $\quad\widetilde{\boldsymbol{X}}^{(s+1)}=\boldsymbol{X}_{(s+1)m}$
12: **end for**

Algorithm 2 Stochastic Average Gradient Langevin Dynamics (SAGA-LD)

1: **input:** step size $\eta>0$; batch size B; epoch length m; inverse temperature parameter $\gamma>0$
2: **initialization:** $\boldsymbol{X}_0=\boldsymbol{0}$, $\widetilde{\boldsymbol{G}}=[\nabla f_1(\boldsymbol{X}_0),\ldots,\nabla f_n(\boldsymbol{X}_0)]$
3: **for** $k=0,1,\ldots,K$ **do**
4: $\quad\widetilde{\boldsymbol{g}}_k=n^{-1}\sum_{i=1}^n\widetilde{\boldsymbol{G}}_i$, where $\widetilde{\boldsymbol{G}}_i$ denotes the i-th column of Matrix $\widetilde{\boldsymbol{G}}$
5: $\quad$randomly pick a subset I_k from $\{1,\ldots,n\}$ of size $|I_k|=B$; randomly draw $\boldsymbol{\epsilon}_k\sim N(\boldsymbol{0},\mathbf{I}_{d\times d})$
6: $\quad\widetilde{\nabla}_k=\frac{1}{B}\sum_{i_k\in I_k}\big(\nabla f_{i_k}(\boldsymbol{X}_k)-\widetilde{\boldsymbol{G}}_{i_k}+\widetilde{\boldsymbol{g}}_k\big)$
7: $\quad\boldsymbol{X}_{k+1}=\boldsymbol{X}_k-\eta\widetilde{\nabla}_k+\sqrt{2\eta/\gamma}\boldsymbol{\epsilon}_k$
8: $\quad\widetilde{\boldsymbol{G}}_{i_k}=\nabla f_{i_k}(\boldsymbol{X}_k)$ for $i_k\in I_k$
9: **end for**

3 REVIEW OF SVRG-LD AND SAGA-LD

In this section, we review the SVRG-LD and SAGA-LD algorithms, which incorporates the variance reduction technique into the Langevin based algorithm.

Algorithm 1 displays the detail of SVRG-LD, which consists of multiple epochs. In the beginning of the s-th epoch, we compute the full gradient of $F(\widetilde{\boldsymbol{X}}^{(s)})$ by scanning all samples

$$\widetilde{G}=\nabla F(\widetilde{\boldsymbol{X}}^{(s)})=\frac{1}{n}\sum_{i=1}^n\nabla f_i(\widetilde{\boldsymbol{X}}^{(s)}).$$

Regarding the l-th inner iteration in the s-th epoch (the k-th update in the total iteration sequence), the semi-stochastic gradient $\widetilde{\nabla}_k$ is computed based on the snapshot gradient $\widetilde{G}$ and a new minibatch of samples I_k, which yields

$$\widetilde{\nabla}_k=\frac{1}{B}\sum_{i_k\in I_k}\big(\nabla f_{i_k}(\boldsymbol{X}_k)-\nabla f_{i_k}(\widetilde{\boldsymbol{X}}^{(s)})+\widetilde{\boldsymbol{G}}\big),$$

where i_k is uniformly sampled from $[n]=\{1,2,\ldots,n\}$, and $|I_k|=B$ denotes the minibatch size. Then we perform the following update based on the semi-stochastic gradient with an injected Gaussian noise $\boldsymbol{\epsilon}_k$,

$$\boldsymbol{X}_{k+1}=\boldsymbol{X}_k-\eta\widetilde{\nabla}_k+\sqrt{2\eta/\gamma}\boldsymbol{\epsilon}_k.$$

At the end of the epoch, we use the last iterate as the starting point of the next epoch, i.e., $\widetilde{\boldsymbol{X}}^{(s+1)}=\boldsymbol{X}_{(s+1)m}$.

Now we present SAGA-LD in Algorithm 2. Compared with SVRG-LD, SAGA-LD requires higher memory cost, since it explicitly stores n stochastic gradients in memory, which formulates n columns of a matrix $\widetilde{\boldsymbol{G}}$. $\widetilde{\boldsymbol{G}}$ is initialized as $[\nabla f_1(\boldsymbol{X}_0),\ldots,\nabla f_n(\boldsymbol{X}_0)]$. In the k-th update, we first compute the average of the column vectors in $\widetilde{\boldsymbol{G}}$, i.e., $\widetilde{\boldsymbol{g}}_k=n^{-1}\sum_{i=1}^n\widetilde{\boldsymbol{G}}_i$ as a snapshot gradient, where $\widetilde{\boldsymbol{G}}_i$ is the i-th column of $\widetilde{\boldsymbol{G}}$. Then an index set I_k is uniformly generated from $[n]$ to compute the following approximated gradient

$$\widetilde{\nabla}_k=\frac{1}{B}\sum_{i_k\in I_k}\big(\nabla f_{i_k}(\boldsymbol{X}_k)-\widetilde{\boldsymbol{G}}_{i_k}+\widetilde{\boldsymbol{g}}_k\big),$$

where $B=|I_k|$ is the size of index set I_k. Then we apply such approximated gradient to perform one-step update on the iterate $\boldsymbol{X}_k$, as shown in the line 7 of Algorithm 2. At the end of each iteration, we update

[2]LMC, SVRG-LD and SAGA-LD require hessian Lipschitz assumption in the strongly log-concave regime.

the columns in $\widetilde{\boldsymbol{G}}$ whose indexes belong to I_k with the stochastic gradients computed in the current iteration, i.e., we set $\widetilde{\boldsymbol{G}}_{i_k} = \nabla f_{i_k}(\boldsymbol{X}_k)$ for all $i_k \in I_k$.

Algorithms 1 and 2 stem from Dubey et al. (2016). However, they only analyzed the mean square error of averaged the sample path based on all iterates $\{\boldsymbol{X}_k\}_{k=0}^K$, while we aim at developing a nonasymptotic analyses of SVRG-LD and SAGA-LD in terms of 2-Wasserstein distance and Algorithms 1 and 2 only require the last iterate $\boldsymbol{X}_K$.

4 MAIN THEORY

In this section, we present our main theoretical results, which characterize the convergence rates of SVRG-LD and SAGA-LD for sampling from non-log-concave distributions. We first lay out the assumptions that are necessary for our theory.

Assumption 4.1 (Smoothness)**.** The function $f_i(\mathbf{x})$ is M-smooth with $M > 0$, i.e., for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $i = 1, \ldots, n$, we have

$$\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\|_2 \leq M\|\mathbf{x} - \mathbf{y}\|_2.$$

The smoothness assumption is also known as gradient Lipschitzness in the literature.

Assumption 4.2 (Dissipative)**.** There exist constants $a, b > 0$, such that for all $\mathbf{x} \in \mathbb{R}^d$ we have

$$\langle \nabla F(\mathbf{x}), \mathbf{x} \rangle \geq b\|\mathbf{x}\|_2^2 - a.$$

It is worthy noting that the smoothness assumption is made on all component function $f_i(\mathbf{x})$, while the dissipative assumption is only made on the average of the component functions. Assumption 4.2 is a typical assumption for the ergodicity analysis of stochastic differential equations (SDE) and diffusion approximation (Mattingly et al., 2002; Vollmer et al., 2016; Raginsky et al., 2017; Zhang et al., 2017). It means that, starting from a position that is sufficiently far away from the origin, the Markov process defined by (1.1) moves towards the origin on average. Note that the class of distribution satisfying dissipative assumption covers many densities of interest such as Gaussian mixture model (Ge et al., 2017).

4.1 Convergence Guarantee for SVRG-LD

Now we present our main theoretical results on the nonasymptotic convergence of SVRG-LD.

Theorem 4.3. Under Assumptions 4.1 and 4.2, consider $\{\boldsymbol{X}_k\}_{k=0,1,\ldots,K}$ generated by Algorithm 1 with initial point $\boldsymbol{X}_0 = \mathbf{0}$. The 2-Wasserstein distance between the distribution of $\boldsymbol{X}_k$ and the target distribution π is bounded by

$$\begin{aligned}
&\mathcal{W}_2\big(P(\boldsymbol{X}_k), \pi\big) \\
&\leq D_1\bigg[D_2\Big(\frac{m^2}{B} + 1\Big)k\eta^3 + D_3\Big(\frac{m}{B} + 1\Big)k\eta^2\bigg]^{1/4} \\
&\quad + D_4 e^{-\frac{k\eta}{\gamma D_5}}, \qquad (4.1)
\end{aligned}$$

where the parameters are defined as

$$\begin{aligned}
D_1 &= 4\sqrt{3/2 + (2b + d/\gamma)k\eta}, \\
D_2 &= 3\gamma M^2\big(2M^2(1 + 1/b)(a + G^2 + d/\gamma) + G^2\big), \\
D_3 &= M^2 d,
\end{aligned}$$

and $G = \max_{i\in[n]} \|f_i(0)\|_2$. Moreover, B is the batch size, m is inner loop length of Algorithm 1, and parameters D_4, D_5 are both in the order of $\exp(O(d + \gamma))$.

Based on Theorem 4.3, we are able to characterize the gradient complexity of Algorithm 1 as well as the choices of hyper parameters including η, m and B. We state these results in the following corollary.

Corollary 4.4. Under the identical assumptions in Theorem 4.3, in order to guarantee that the target accuracy satisfies $\mathcal{W}_2\big(P(\mathbf{x}_k), \pi\big) \leq \epsilon$, we set $mB = O(n)$, $\eta = \widetilde{O}(\epsilon^2 B^{3/2}/n^2 \wedge \epsilon^4 B^2/n) \cdot \exp\big(-\widetilde{O}(\gamma + d)\big)$. Then the gradient complexity of Algorithm 1 is

$$T_g = \widetilde{O}\Big(\frac{nB^{-1/2}}{\epsilon^2} + \frac{n/B + B}{\epsilon^4} + n\Big) \cdot \exp\big(\widetilde{O}(\gamma + d)\big).$$

Moreover, if we set $B = O\big(n^{1/2}\big)$ and $\eta = \widetilde{O}\big(\epsilon^2/n^{1/4} \wedge \epsilon^4\big) \cdot \exp\big(-\widetilde{O}(\gamma + d)\big)$, the gradient complexity is

$$T_g = \widetilde{O}\Big(n + \frac{n^{3/4}}{\epsilon^2} + \frac{n^{1/2}}{\epsilon^4}\Big) \cdot \exp\big(\widetilde{O}(\gamma + d)\big).$$

Remark 4.5. Under identical assumptions in Theorem 4.3, LMC achieves ϵ-accuracy in 2-Wasserstein distance after $T_g = \widetilde{O}\big(n/\epsilon^4\big) \cdot \exp\big(\widetilde{O}(d + \gamma)\big)$ stochastic gradient evaluations (Raginsky et al., 2017). It is obvious that SVRG-LD requires less stochastic gradient evaluations to achieve ϵ-accuracy than LMC.

4.2 Convergence Guarantee for SAGA-LD

Next, we present the following theorem that spells out the convergence rate of SAGA-LD.

Theorem 4.6. Under Assumptions 4.1 and 4.2, consider $\{\boldsymbol{X}_k\}_{k=0,1,\ldots,K}$ generated by Algorithm 2 with initial point $\boldsymbol{X}_0 = \mathbf{0}$. The 2-Wasserstein distance between the distribution of $\boldsymbol{X}_k$ and the target distribu-

tion π is bounded by

$$\begin{aligned}&\mathcal{W}_2\big(P(\boldsymbol{X}_k),\pi\big)\\&\leq D_1\bigg[D_2\bigg(\frac{48n^2}{B^3}+1\bigg)k\eta^3+D_3\bigg(\frac{4n}{B^2}+1\bigg)k\eta^2\bigg]^{1/4}\\&\quad+D_4e^{-\frac{k\eta}{\gamma D_5}},\end{aligned}\tag{4.2}$$

where $G=\max_{i\in[n]}\|f_i(0)\|_2$, B is the batch size, parameters D_1, D_2, D_3, D_4 and D_5 are identical to those in Theorem 4.3.

Based on Theorem 4.6, we present the gradient complexity of SAGA-LD in the following corollary.

Corollary 4.7. Under the same assumptions as in Theorem 4.6, in order to guarantee that the target accuracy satisfies $\mathcal{W}_2\big(P(\mathbf{x}_k),\pi\big)\leq\epsilon$, we set $\eta=\widetilde{O}(\epsilon^2B^{3/2}/n^2\wedge\epsilon^4B^2/n)\cdot\exp\big(-\widetilde{O}(\gamma+d)\big)$, and the gradient complexity of Algorithm 2 is

$$T_g=\widetilde{O}\Big(\frac{nB^{-1/2}}{\epsilon^2}+\frac{n/B+B}{\epsilon^4}+n\Big)\cdot\exp\big(\widetilde{O}(\gamma+d)\big).$$

Moreover, if we set $B=O\big(n^{1/2}\big)$ and $\eta=\widetilde{O}\big(\epsilon^2/n^{1/4}\wedge\epsilon^4\big)\cdot\exp\big(-\widetilde{O}(\gamma+d)\big)$, the gradient complexity becomes

$$T_g=\widetilde{O}\Big(n+\frac{n^{3/4}}{\epsilon^2}+\frac{n^{1/2}}{\epsilon^4}\Big)\cdot\exp\big(\widetilde{O}(d+\gamma)\big).$$

Remark 4.8. It can be clearly observed that the gradient complexities of SVRG-LD and SAGA-LD are essentially identical when we set $mB=O(n)$ in SVRG-LD. This observation also matches the result in Dubey et al. (2016) and Zou et al. (2018b), where the former focuses on the mean squared error of sample path average and the latter only establishes the convergence guarantees for strongly log-concave densities.

Remark 4.9. It is worth noting that our analyses on SVRG-LD and SAGA-LD do not imply the convergence rate of SGLD. However, the convergence rate of SGLD in 2-Wasserstein distance is similar to Equation (3.2) in Raginsky et al. (2017). Based on the argument in Raginsky et al. (2017), the SGLD algorithm cannot be guaranteed to converge to the target distribution if the batch size is not carefully specified. However, empirical study shows that SGLD converges in most cases, which indicates a gap between the theory and the experiment. In particular, we found that SGLD actually converge to the target distribution in our experiment, even when the batch size is set to be 1, and enjoys faster rate than LMC.

5 EXPERIMENTS

In order to explore the behavior of SVRG-LD and SAGA-LD for sampling from non-log-concave densities, we carry out numerical experiments on both synthetic and real dataset in this section. Specifically, we compare the SVRG-LD and SAGA-LD algorithms with LMC and SGLD for sampling from non-log-concave density, independent component analysis (ICA) and Bayesian logistic regression.

5.1 Sampling for Gaussian Mixture Distribution

We first compare the performances of SVRG-LD, SAGA-LD, LMC and SGLD on synthetic data. In particular, we consider the target distribution with form $\pi\propto\exp(-F(\mathbf{x}))=\exp\big(-\sum_{i=1}^n f_i(\mathbf{x})/n\big)$, where each component $\exp(-f_i(\mathbf{x}))$ is defined as

$$\exp(-f_i(\mathbf{x}))=e^{-\|\mathbf{x}-\mathbf{a}_i\|_2^2/2}+e^{-\|\mathbf{x}+\mathbf{a}_i\|_2^2/2},\quad\mathbf{a}_i\in\mathbb{R}^d.$$

It is easy to verify that $\exp(-f_i(\mathbf{x}))$ is proportion to the PDF of a Gaussian mixture distribution. The function $f_i(\mathbf{x})$ and its gradient can be further simplified as

$$\begin{aligned}f_i(\mathbf{x})&=\frac{1}{2}\|\mathbf{x}-\mathbf{a}_i\|_2^2-\log\big(1+\exp(-2\mathbf{x}^\top\mathbf{a}_i)\big),\\\nabla f_i(\mathbf{x})&=\mathbf{x}-\mathbf{a}_i+\frac{2\mathbf{a}_i}{1+\exp(2\mathbf{x}^\top\mathbf{a}_i)}.\end{aligned}$$

According to Dalalyan (2016); Dwivedi et al. (2018), when the parameter $\mathbf{a}_i$ is chosen such that $\|\mathbf{a}_i\|_2^2>1$, function $f_i(\mathbf{x})$ defined as above is nonconvex. Moreover, it can be seen that

$$\begin{aligned}\langle\nabla f_i(\mathbf{x}),\mathbf{x}\rangle&=\|\mathbf{x}\|_2^2+\frac{1-\exp(2\mathbf{x}^\top\mathbf{a}_i)}{1+\exp(2\mathbf{x}^\top\mathbf{a}_i)}\langle\mathbf{a}_i,\mathbf{x}\rangle\\&\geq\frac{1}{2}\|\mathbf{x}\|_2^2-\frac{1}{2}\|\mathbf{a}_i\|_2^2,\end{aligned}$$

which suggests that function $f_i(\mathbf{x})$ satisfies Dissipative Assumption 4.2 with $b=1/2$ and $a=\|\mathbf{a}_i\|_2^2/2$ and further implies that $F(\mathbf{x})$ is also dissipative. Then we set sample size $n=500$ and dimension $d=10$, and randomly generate parameters $\mathbf{a}_i\sim N(\mu,\boldsymbol{\Sigma})$ with $\mu=(2,\ldots,2)^\top$ and $\boldsymbol{\Sigma}=\mathbf{I}_{d\times d}$. Since it takes a large number of samples to characterize the distribution, which makes repeated experiments computationally expensive, we instead follow Bardenet et al. (2017) to use iterates along one Markov chain to visualize the distribution of iterates obtained by MCMC algorithms. Specifically, we run all four algorithms for 2×10^4 data passes, and make use of the iterates in the last 10^4 data passes to visualize distributions, where the batch sizes for SGLD, SVRG-LD and SAGA-LD are all set to be 10. In Figures 1(a) - 1(d), We compare the distributions generated by LMC, SGLD, SVRG-LD and SAGA-LD while using MCMC with Metropolis-Hasting correction as a reference. It can be observed that both SVRG-LD and SAGA-LD can well approximate the target distribution within

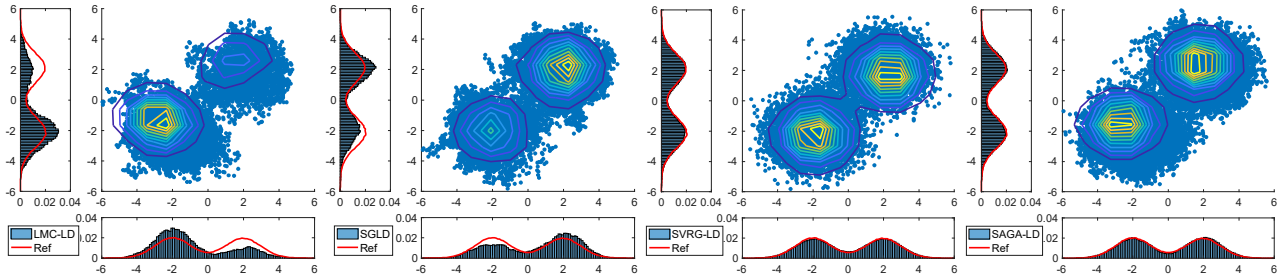

(a) LMC (b) SGLD (c) SVRG-LD (d) SAGA-LD

Figure 1: 2D projection of the kernel densities of random samples generated after 10^4 data passes. (a) - (d) represent 4 different algorithms.

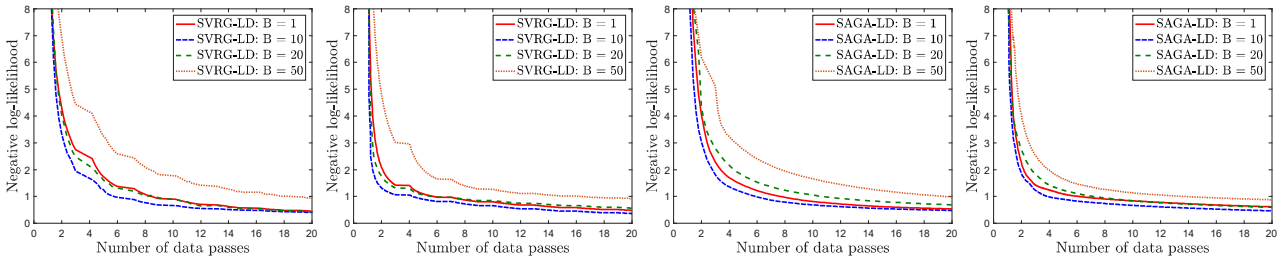

(a) $n = 500$ (b) $n = 5000$ (c) $n = 500$ (d) $n = 5000$

Figure 2: Experiment results for independent components analysis, where x axis indicates the number of data pass and y axis shows the negative log-likelihood on the test data. (a)-(b) Experiment results for SVRG-LD with different batch size. (c)-(d) Experiment results for SAGA-LD with different batch size.

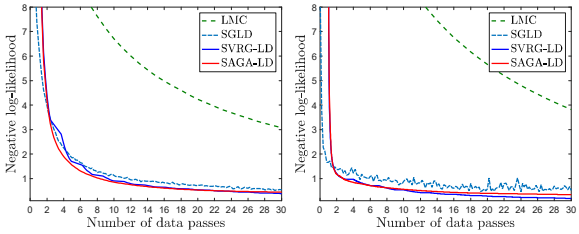

(a) $n = 500$ (b) $n = 5000$

Figure 3: Experiment results of ICA for different algorithms.

2×10^4 datapass, while the distributions generated by LMC and SGLD have obvious deviation from the true one. This suggests that SVRG-LD and SAGA-LD enjoy faster convergence rate than LMC and SGLD, which verifies our theory. However, if we run SGLD and LMC for more iterations, SGLD and LMC can both well approximate the target distribution. More interestingly, we find that SGLD actually requires less gradient evaluations than LMC to well approximate the target distribution, which does not well align with the existing theory.

5.2 Independent Components Analysis

We further apply the SVRG-LD and SAGA-LD algorithms to a Bayesian Independent Component Analysis (ICA) model, and compare their performance with LMC and SGLD. In the ICA model, we are given a dataset with n examples $\mathbf{X} = \{\mathbf{x}_i\}_{i=1,\ldots,n}$. The probability of samples $\mathbf{x}_i$ given the model matrix $\mathbf{W}$ can be written as follows (Welling and Teh, 2011; Dubey et al., 2016)

$$p(\mathbf{x}_i|\boldsymbol{W}) = |\det(\mathbf{W})| \prod_i p(\mathbf{w}_i^\top \mathbf{x}_i),$$

where $p(\mathbf{w}_i^\top \mathbf{x}_i) = 1/(4\cosh^2(\mathbf{w}_i^\top \mathbf{x}_i/2))$. We consider Gaussian prior over $\mathbf{W}$, i.e., $p(\mathbf{W}) \sim \mathcal{N}(0, \lambda^{-1}\mathbf{I})$. Then we formulate the log-posterior as the average of n component functions, i.e., $\sum_{i=1}^n f_i(\mathbf{W})/n$, where

$$f_i(\mathbf{W}) = -n\big[\log(|\det(\mathbf{W})|) + 2\sum_{i=1}^{d} \log\big(\cosh(\mathbf{w}_i^\top \mathbf{x}_i/2)\big)\big] + \lambda\|\mathbf{W}\|_F^2.$$

We perform the ICA algorithm on EEG dataset[3], which contains 125337 samples with 34 channels. In this experiment, we consider two regimes with different sample size n. To achieve this, we extract two subsets with size 500 and 5000 from the original dataset, and extract 5000 samples from the rest dataset for test. Follow the same procedures in Welling and Teh (2011); Chen et al. (2014); Zou et al. (2018a), we discard the first 50 iterates as burnin and compute the sample path average to estimate the model matrix parameter $\mathbf{W}$. We first run SVRG-LD and SAGA-LD with different batch sizes $B = 1$, $B = 10$, $B = 20$ and $B = 50$ (the epoch length is set to be $m = 2n/B$ for SVRG-LD),

[3] https://mmspg.epfl.ch/cms/page-58322.html

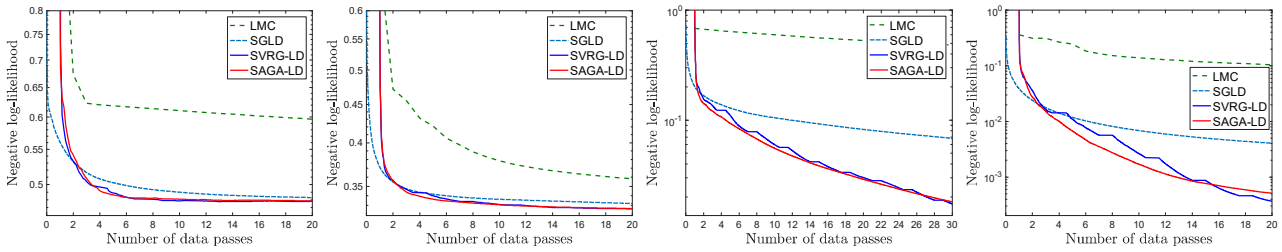

(a) *pima* (b) *a3a* (c) *gisette* (d) *mushroom*

Figure 4: Experiment results for Bayesian logistic regression, where x axis indicates the number of data pass and y axis shows the negative log-likelihood on the test dataset. (a) - (d) represent 4 different datasets.

and plot the negative log-likelihood on test dataset with respect to the number of effective data pass in Figures 2(a)-2(d). It can be seen that both SVRG-LD and SAGA-LD algorithms have the best performance when the batch size is $B = 10$. Next, we set batch size to be $B = 10$ for both SVRG-LD and SAGA-LD, and compare their convergence performances with those of LMC and SGLD, which are displayed in Figures 3(a)-3(b). It should be noted that in the first epoch, SVRG-LD and SAGA-LD compute the full gradient using all n samples, thus the curves of SVRG-LD and SAGA-LD should start from the first data pass. Moreover, we observe that SVRG-LD and SAGA-LD have comparable performance and both converge faster than SGLD and LMC, this supports our theory.

5.3 Bayesian Logistic Regression

We also apply LMC, SGLD, SVRG-LD and SAGA-LD to a Bayesian logistic regression problem. In this problem, n i.i.d samples $\{\mathbf{x}_i, y_i\}_{i=1,\ldots,n}$ are observed, where $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i \in \{-1, 1\}$ denote the feature and the corresponding label of the i-th sample. In Bayesian logistic model, the likelihood function takes the form $p(y_i|\mathbf{x}_i, \boldsymbol{\beta}) = 1/\big(1+\exp(-y_i\mathbf{x}_i^\top\boldsymbol{\beta})\big)$ where $\boldsymbol{\beta}$ is the regression parameter that requires to be trained. In order to evaluate the performance of SVRG-LD and SAGA-LD when dealing with non-log-concave densities, we consider Gamma prior $p(\boldsymbol{\beta}) \propto \|\boldsymbol{\beta}\|_2^{-\lambda}\exp(-\theta\|\boldsymbol{\beta}\|_2)$. Then we formulate the logarithmic posterior distribution as follows,

$$\log\big[p(\boldsymbol{\beta}|\mathbf{x}_1,\ldots,\mathbf{x}_n;y_1,\ldots,y_n)\big] \propto -\frac{1}{n}\sum_{i=1}^{n} f_i(\boldsymbol{\beta}),$$

where $f_i(\beta) = n\log\big(1+e^{-y_i\mathbf{x}_i^\top\boldsymbol{\beta}}\big) + \lambda\log(\|\boldsymbol{\beta}\|_2) + \theta\|\boldsymbol{\beta}\|_2$. We compare SVRG-LD and SAGA-LD with the baseline algorithms on four datasets from UCI[4] and Libsvm[5] libraries, which are *pima*, *a3a*, *gisette*, and *mushroom*. Since *pima* and *mushroom* do not have test data, we manually split the whole dataset into training and test parts. Again, we compute the sample path average to estimate the regression parameter $\boldsymbol{\beta}$. The comparison between different algorithms for different datasets are displayed in Figure 4(a) - 4(d). Similarly, SVRG-LD and SAGA-LD start from the first data pass. It can be observed that the performances of SVRG-LD and SAGA-LD are quite similar, and both converge faster than another two baseline algorithms, which suggests that the SVRG-LD and SAGA-LD methods serve as better choices for Bayesian logistic regression with non-log-concave prior compared with LMC and SGLD.

[4] `https://archive.ics.uci.edu/ml/`

[5] `https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/`

6 CONCLUSIONS AND FUTURE WORK

We studied the SVRG-LD and SAGA-LD methods for sampling from non-log-concave densities, and proved the corresponding convergence rate as well as the gradient complexity when the sampling error is measured as 2-Wasserstein distance. Experimental results showed that SVRG-LD and SAGA-LD achieve similar performance, and converge faster than LMC and SGLD when the target distribution is non-log-concave, which is consistent with our theory.

There are many possible future directions that demand to be explored, such as the convergence rate of SGLD in Wasserstein distance when the target distribution is non-log-concave. In addition, it is also of interest to investigate whether the metropolis hasting step can be applied to further improve the current results.

Acknowledgements

We would like to thank the anonymous reviewers for their helpful comments. This research was sponsored in part by the National Science Foundation IIS-1906169 and BIGDATA IIS-1855099. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

Ahn, S., Korattikara, A. and Welling, M. (2012). Bayesian posterior sampling via stochastic gradient fisher scoring. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*.

Baker, J., Fearnhead, P., Fox, E. B. and Nemeth, C. (2017). Control variates for stochastic gradient mcmc. *arXiv preprint arXiv:1706.05439* .

Bakry, D., Barthe, F., Cattiaux, P., Guillin, A. et al. (2008). A simple proof of the poincaré inequality for a large class of probability measures. *Electronic Communications in Probability* **13** 60–66.

Bakry, D., Gentil, I. and Ledoux, M. (2013). *Analysis and geometry of Markov diffusion operators*, vol. 348. Springer Science & Business Media.

Bardenet, R., Doucet, A. and Holmes, C. (2017). On markov chain monte carlo methods for tall data. *The Journal of Machine Learning Research* **18** 1515–1557.

Betancourt, M. (2015). The fundamental incompatibility of scalable hamiltonian monte carlo and naive data subsampling. In *International Conference on Machine Learning*.

Bolley, F. and Villani, C. (2005). Weighted csiszár-kullback-pinsker inequalities and applications to transportation inequalities. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, vol. 14.

Bou-Rabee, N. and Hairer, M. (2012). Nonasymptotic mixing of the mala algorithm. *IMA Journal of Numerical Analysis* **33** 80–110.

Bovier, A., Eckhoff, M., Gayrard, V. and Klein, M. (2004). Metastability in reversible diffusion processes i: Sharp asymptotics for capacities and exit times. *Journal of the European Mathematical Society* **6** 399–424.

Brosse, N., Durmus, A., Moulines, É. and Pereyra, M. (2017). Sampling from a log-concave distribution with compact support with proximal langevin monte carlo. *arXiv preprint arXiv:1705.08964* .

Bubeck, S., Eldan, R. and Lehec, J. (2015). Sampling from a log-concave distribution with projected langevin monte carlo. *arXiv preprint arXiv:1507.02564* .

Cattiaux, P., Guillin, A. and Wu, L.-M. (2010). A note on talagrands transportation inequality and logarithmic sobolev inequality. *Probability theory and related fields* **148** 285–304.

Chatterji, N. S., Flammarion, N., Ma, Y.-A., Bartlett, P. L. and Jordan, M. I. (2018). On the theory of variance reduction for stochastic gradient monte carlo. *arXiv preprint arXiv:1802.05431* .

Chaudhari, P., Choromanska, A., Soatto, S. and LeCun, Y. (2016). Entropy-sgd: Biasing gradient descent into wide valleys. *arXiv preprint arXiv:1611.01838* .

Chen, C., Wang, W., Zhang, Y., Su, Q. and Carin, L. (2017). A convergence analysis for a class of practical variance-reduction stochastic gradient mcmc. *arXiv preprint arXiv:1709.01180* .

Chen, T., Fox, E. and Guestrin, C. (2014). Stochastic gradient hamiltonian monte carlo. In *International Conference on Machine Learning*.

Cheng, X., Chatterji, N. S., Bartlett, P. L. and Jordan, M. I. (2017). Underdamped langevin mcmc: A non-asymptotic analysis. *arXiv preprint arXiv:1707.03663* .

Chiang, T.-S., Hwang, C.-R. and Sheu, S. J. (1987). Diffusion for global optimization in rˆn. *SIAM Journal on Control and Optimization* **25** 737–753.

Dalalyan, A. S. (2014). Theoretical guarantees for approximate sampling from smooth and log-concave densities. *arXiv preprint arXiv:1412.7392* .

Dalalyan, A. S. (2016). Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* .

Dalalyan, A. S. (2017). Further and stronger analogy between sampling and optimization: Langevin monte carlo and gradient descent. *arXiv preprint arXiv:1704.04752* .

Dalalyan, A. S. and Karagulyan, A. G. (2017). User-friendly guarantees for the langevin monte carlo with inaccurate gradient. *arXiv preprint arXiv:1710.00095* .

Dang, K.-D., Quiroz, M., Kohn, R., Tran, M.-N. and Villani, M. (2017). Hamiltonian monte carlo with energy conserving subsampling. *arXiv preprint arXiv:1708.00955* .

Dubey, K. A., Reddi, S. J., Williamson, S. A., Poczos, B., Smola, A. J. and Xing, E. P. (2016). Variance reduction in stochastic gradient langevin dynamics. In *Advances in Neural Information Processing Systems*.

Durmus, A. and Moulines, E. (2015). Nonasymptotic convergence analysis for the unadjusted langevin algorithm. *arXiv preprint arXiv:1507.05021* .

Durmus, A. and Moulines, E. (2016). Sampling from strongly log-concave distributions with the unadjusted langevin algorithm. *arXiv preprint arXiv:1605.01559* .

Dwivedi, R., Chen, Y., Wainwright, M. J. and Yu, B. (2018). Log-concave sampling: Metropolis-hastings algorithms are fast! *arXiv preprint arXiv:1801.02309* .

Eberle, A. et al. (2014). Error bounds for metropolis–hastings algorithms applied to perturbations of gaussian measures in high dimensions. *The Annals of Applied Probability* **24** 337–377.

Ge, R., Lee, H. and Risteski, A. (2017). Beyond log-concavity: Provable guarantees for sampling multi-modal distributions using simulated tempering langevin monte carlo. *arXiv preprint arXiv:1710.02736* .

Gyöngy, I. (1986). Mimicking the one-dimensional marginal distributions of processes having an itô differential. *Probability theory and related fields* **71** 501–516.

Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications .

Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*.

Kloeden, P. E. and Platen, E. (1992). Higher-order implicit strong numerical schemes for stochastic differential equations. *Journal of statistical physics* **66** 283–314.

Li, Z., Zhang, T. and Li, J. (2018). Stochastic gradient hamiltonian monte carlo with variance reduction for bayesian inference. *arXiv preprint arXiv:1803.11159* .

Ma, Y.-A., Chen, T. and Fox, E. (2015). A complete recipe for stochastic gradient mcmc. In *Advances in Neural Information Processing Systems*.

Mattingly, J. C., Stuart, A. M. and Higham, D. J. (2002). Ergodicity for sdes and approximations: locally lipschitz vector fields and degenerate noise. *Stochastic processes and their applications* **101** 185–232.

Neal, R. M. et al. (2011). Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo* **2** 113–162.

Raginsky, M., Rakhlin, A. and Telgarsky, M. (2017). Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. *arXiv preprint arXiv:1702.03849* .

Reddi, S. J., Hefny, A., Sra, S., Poczos, B. and Smola, A. (2016). Stochastic variance reduction for nonconvex optimization. *arXiv preprint arXiv:1603.06160* .

Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **60** 255–268.

Roberts, G. O. and Tweedie, R. L. (1996a). Exponential convergence of langevin distributions and their discrete approximations. *Bernoulli* 341–363.

Roberts, G. O. and Tweedie, R. L. (1996b). Geometric convergence and central limit theorems for multidimensional hastings and metropolis algorithms. *Biometrika* **83** 95–110.

Vollmer, S. J., Zygalakis, K. C. and Teh, Y. W. (2016). Exploration of the (non-) asymptotic bias and variance of stochastic gradient langevin dynamics. *Journal of Machine Learning Research* **17** 1–48.

Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning*.

Whye Teh, Y., Thiéry, A. and Vollmer, S. (2014). Consistency and fluctuations for stochastic gradient langevin dynamics. *arXiv preprint arXiv:1409.0578* .

Xu, P., Chen, J., Zou, D. and Gu, Q. (2018). Global convergence of langevin dynamics based algorithms for nonconvex optimization. In *Advances in Neural Information Processing Systems*.

Ye, N., Zhu, Z. and Mantiuk, R. K. (2017). Langevin dynamics with continuous tempering for high-dimensional non-convex optimization. *arXiv preprint arXiv:1703.04379* .

Zhang, Y., Liang, P. and Charikar, M. (2017). A hitting time analysis of stochastic gradient langevin dynamics. *arXiv preprint arXiv:1702.05575* .

Zou, D., Xu, P. and Gu, Q. (2018a). Stochastic variance-reduced Hamilton Monte Carlo methods. In *Proceedings of the 35th International Conference on Machine Learning*.

Zou, D., Xu, P. and Gu, Q. (2018b). Subsampled stochastic variance-reduced gradient langevin dynamics. In *UAI*.