
Non-convex Learning via Replica Exchange Stochastic Gradient MCMC

Wei Deng¹ Qi Feng^{*2} Liyao Gao^{*1} Faming Liang¹ Guang Lin¹

Abstract

Replica exchange Monte Carlo (reMC), also known as parallel tempering, is an important technique for accelerating the convergence of the conventional Markov Chain Monte Carlo (MCMC) algorithms. However, such a method requires the evaluation of the energy function based on the full dataset and is not scalable to big data. The naïve implementation of reMC in mini-batch settings introduces large biases, which cannot be directly extended to the stochastic gradient MCMC (SGMCMC), the standard sampling method for simulating from deep neural networks (DNNs). In this paper, we propose an adaptive replica exchange SGMCMC (reSGMCMC) to automatically correct the bias and study the corresponding properties. The analysis implies an acceleration-accuracy trade-off in the numerical discretization of a Markov jump process in a stochastic environment. Empirically, we test the algorithm through extensive experiments on various setups and obtain the state-of-the-art results on CIFAR10, CIFAR100, and SVHN in both supervised learning and semi-supervised learning tasks.

1. Introduction

The increasing concern for AI safety problems draws our attention to MCMC, which is known for its asymptotic uncertainty quantification (Chen et al., 2015; Teh et al., 2016), and guarantees in non-convex optimizations (Zhang et al., 2017; Raginsky et al., 2017). Traditional MCMC methods have achieved tremendous success. However, the efficient sampling algorithm in DNNs was not well studied until the invention of stochastic gradient Langevin dynamics (SGLD) (Welling and Teh, 2011), which scales up the computation

in DNNs by injecting noises to stochastic gradients. Since then, various high-order SGMCMC algorithms have been proposed, which incorporate strategies such as Hamiltonian dynamics (Chen et al., 2014; Ma et al., 2015; Ding et al., 2014), Hessian approximation (Li et al., 2016; Şimşekli et al., 2016), and high-order numerical schemes (Chen et al., 2015; Li et al., 2019) to improve the convergence.

In addition to the high-order algorithms, we can also follow traditional MCMC algorithms combined with simulated annealing (Kirkpatrick et al., 1983), simulated tempering (Marinari and Parisi, 1992), dynamical weighting (Wong and Liang, 1997) or replica exchange Monte Carlo (Swendsen and Wang, 1986; Earl and Deem, 2005). Among these advancements, simulated annealing SGMCMC (Mangoubi and Vishnoi, 2018) and simulated tempering SGMCMC (Lee et al., 2018) show how dynamical temperatures speed up the convergence. However, simulated annealing is very sensitive to the fast-decaying temperatures, and simulated tempering requires a lot on the approximation of the normalizing constant. For the latter, the replica exchange Monte Carlo is easier to analyze and implement and is suitable for parallelism. Specifically, the replica exchange Langevin diffusion utilizes multiple diffusion processes with different temperatures and proposes to swap the processes while training. Intuitively, the high-temperature process acts as a bridge to connect the various modes. As such, the acceleration effect can be theoretically quantified (Dupuis et al., 2012; Chen et al., 2019). However, despite these advantages, a proper replica exchange SGMCMC (reSGMCMC) has long been missing in the deep learning community.

A bottleneck that hinders the development of reSGMCMC is the naïve extension of the acceptance-rejection criterion that fails in mini-batch settings. Various attempts (Bardenet et al., 2017; Korattikara et al., 2014) were proposed to solve this issue. However, they introduce biases even with the ideal normality assumption on the noise. Some unbiased estimators (Bhanot and Kennedy, 1985; Beskos et al., 2006) have ever been presented, but the large variance leads to inefficient inference. To remove the bias while maintaining efficiency, Ceperley and Dewing (1999) proposed a corrected criterion under normality assumptions, and Seita et al. (2017); Quiroz et al. (2019) further analyzed the model errors with the asymptotic normality assumptions. However, the above algorithms fail when the required corrections are

^{*}Equal contribution ¹Purdue University, West Lafayette, IN, USA. ²University of Southern California, Los Angeles, CA, USA.. Correspondence to: Wei Deng <weideng056@gmail.com>, Guang Lin <guanglin@purdue.edu>, Faming Liang <fm-liang@purdue.edu>.

time-varying and much larger than the energies as shown in Fig.3(a-b). Consequently, an effective algorithm with the potential to adaptively estimate the corrections and balance between acceleration and accuracy is in great demand.

In this paper, we propose an adaptive replica exchange SGMCMC algorithm via stochastic approximation (SA) (Robbins and Monro, 1951; Liang et al., 2007; Deng et al., 2019), a standard method in adaptive sampling to estimate the latent variable: the unknown correction. The adaptive algorithm not only shows the asymptotic convergence in standard scenarios but also gives a good estimate when the corrections are time-varying and excessively large. We theoretically analyze the discretization error for reSGMCMC in mini-batch settings and show the accelerated convergence in 2-Wasserstein distance. Such analysis sheds light on the use of biased estimates of unknown corrections to obtain a trade-off between acceleration and accuracy. In summary, this algorithm has three main contributions:

1. We propose a novel reSGMCMC to speed up the computations of SGMCMC in DNNs with theoretical guarantees. The theory shows the potential of using biased corrections and a large batch size to obtain better performance.
2. We identify the problem of time-varying corrections in DNNs and propose to adaptively estimate the time-varying corrections, with potential extension to a variety of time-series prediction techniques.
3. We test the algorithm through extensive experiments using various models. It achieves the state-of-the-art results in both supervised learning and semi-supervised learning tasks on CIFAR10, CIFAR100, and SVHN datasets.

2. Preliminaries

A standard sampling algorithm is the Langevin diffusion, which is a stochastic differential equation (SDE) as follows:

$$d\beta_t^{(1)} = -\nabla U(\beta_t^{(1)})dt + \sqrt{2\tau_1}d\mathbf{W}_t^{(1)}, \quad (1)$$

where $\beta_t^{(1)} \in \mathbb{R}^d$, $U(\cdot)$ is the energy function, $\mathbf{W}_t^{(1)} \in \mathbb{R}^d$ is the Brownian motion, and $\tau_1 > 0$ is the temperature.

Under mild growth conditions on U , the Langevin diffusion $\{\beta_t^{(1)}\}_{t \geq 0}$ converges to the unique invariant Gibbs distribution $\pi_{\tau_1}(\beta^{(1)}) \propto \exp(-\frac{U(\beta^{(1)})}{\tau_1})$, where the temperature τ_1 is crucial for both optimization and sampling of the non-convex energy function U . On the one hand, a high-temperature τ_1 achieves the *exploration* effect: the convergence to the flattened Gibbs distribution of the whole domain is greatly facilitated. However, the flattened distribution is less concentrated around the global optima (Raginsky et al., 2017), and the geometric connection to the global minimum is affected (Zhang et al., 2017). On the other hand, a low-temperature τ_1 leads to the *exploitation* effect:

the solutions explore the local geometry rapidly, but they are more likely to get trapped in local optima, leading to a slow convergence in both optimization and sampling. Therefore, the potential of using a fixed temperature is quite limited.

A powerful algorithm called replica exchange Langevin diffusion (reLD), also known as parallel tempering Langevin diffusion, has been proposed to accelerate the convergence of the SDE as shown in Fig.1.

reLD proposes to simulate a high-temperature particle for *exploration* and a low-temperature particle for *exploitation* and allows them to swap simultaneously. Now consider the following coupled processes with a higher temperature $\tau_2 > \tau_1$ and $\mathbf{W}^{(2)}$ independent of $\mathbf{W}^{(1)}$:

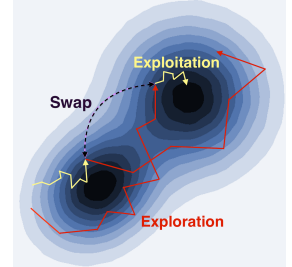


Figure 1. Paths of reLD.

$$\begin{aligned} d\beta_t^{(1)} &= -\nabla U(\beta_t^{(1)})dt + \sqrt{2\tau_1}d\mathbf{W}_t^{(1)} \\ d\beta_t^{(2)} &= -\nabla U(\beta_t^{(2)})dt + \sqrt{2\tau_2}d\mathbf{W}_t^{(2)}. \end{aligned} \quad (2)$$

Eq.(2) converges to the invariant distribution with density

$$\pi(\beta^{(1)}, \beta^{(2)}) \propto e^{-\frac{U(\beta^{(1)})}{\tau_1} - \frac{U(\beta^{(2)})}{\tau_2}}. \quad (3)$$

By allowing the two particles to swap, the positions are likely to change from $(\beta_t^{(1)}, \beta_t^{(2)})$ to $(\beta_{t+dt}^{(2)}, \beta_{t+dt}^{(1)})$ with a swapping rate $rS(\beta_t^{(1)}, \beta_t^{(2)})dt$, where the constant $r \geq 0$ is the swapping intensity, and $S(\cdot, \cdot)$ satisfies

$$S(\beta_t^{(1)}, \beta_t^{(2)}) := e^{\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right)(U(\beta_t^{(1)}) - U(\beta_t^{(2)}))}. \quad (4)$$

In such a case, reLD is a Markov jump process, which is reversible (Chen et al., 2019) and leads to the same invariant distribution (3).

3. Replica Exchange Stochastic Gradient Langevin Dynamics

The wide adoption of the replica exchange Monte Carlo in traditional MCMC algorithms motivates us to design replica exchange stochastic gradient Langevin dynamics for DNNs, but the straightforward extension of reLD to replica exchange stochastic gradient Langevin dynamics is highly non-trivial (Chen et al., 2014; Ma et al., 2015; Şimşekli et al., 2016). In this section, we will first show that naïve extensions of replica exchange Monte Carlo to SGLD (naïve reSGLD) lead to large biases. Afterward, we will present an adaptive replica exchange stochastic gradient Langevin dynamics (reSGLD) that will automatically adjust the bias and yield a good approximation to the correct distribution.

3.1. Naïve reSGLD

We denote the entire data by $\mathcal{D} = \{\mathbf{d}_i\}_{i=1}^N$, where \mathbf{d}_i is a data point. Given the model parameter $\tilde{\beta}$, we consider the following energy function (negative log-posterior)

$$L(\tilde{\beta}) = -\log p(\tilde{\beta}) - \sum_{i=1}^N \log P(\mathbf{d}_i | \tilde{\beta}). \quad (5)$$

where $p(\tilde{\beta})$ is a proper prior and $\sum_{i=1}^N \log P(\mathbf{d}_i | \tilde{\beta})$ is the complete data log-likelihood. When the number of data points N is large, it is expensive to evaluate $L(\tilde{\beta})$ directly. Instead, we propose to approximate the energy function with a mini-batch of data $\mathcal{B} = \{\mathbf{d}_{s_i}\}_{i=1}^n$, where $s_i \in \{1, 2, \dots, N\}$. We can easily check that if \mathcal{B} is sampled randomly with or without replacement, we obtain the following unbiased estimator of the energy function

$$\tilde{L}(\tilde{\beta}) = -\log p(\tilde{\beta}) - \frac{N}{n} \sum_{i=1}^n \log P(\mathbf{d}_{s_i} | \tilde{\beta}). \quad (6)$$

Let $\tilde{\beta}_k$ denote the estimate of $\tilde{\beta}$ at the k -th iteration. Next, SGLD proposes the following iterations:

$$\tilde{\beta}_{k+1} = \tilde{\beta}_k - \eta_k \nabla \tilde{L}(\tilde{\beta}_k) + \sqrt{2\eta_k \tau_1} \xi_k, \quad (7)$$

where η_k is the learning rate, the stochastic gradient $\nabla \tilde{L}(\tilde{\beta}_k)$ is the unbiased estimator of the exact gradient $\nabla L(\tilde{\beta}_k)$, ξ is a standard d -dimensional Gaussian vector with mean $\mathbf{0}$ and identity covariance matrix. It is known that SGLD asymptotically converges to a unique invariant distribution $\pi(\tilde{\beta}) \propto \exp(-L(\tilde{\beta})/\tau_1)$ (Teh et al., 2016) as $\eta_k \rightarrow 0$. If we simply replace gradients with stochastic gradients in the replica exchange gradient Langevin dynamics, we have

$$\begin{aligned} \tilde{\beta}_{k+1}^{(1)} &= \tilde{\beta}_k^{(1)} - \eta_k \nabla \tilde{L}(\tilde{\beta}_k^{(1)}) + \sqrt{2\eta_k \tau_1} \xi_k^{(1)} \\ \tilde{\beta}_{k+1}^{(2)} &= \tilde{\beta}_k^{(2)} - \eta_k \nabla \tilde{L}(\tilde{\beta}_k^{(2)}) + \sqrt{2\eta_k \tau_2} \xi_k^{(2)}. \end{aligned} \quad (8)$$

Furthermore, we swap the Markov chains in (8) with the naïve stochastic swapping rate $r\mathbb{S}(\tilde{\beta}_{k+1}^{(1)}, \tilde{\beta}_{k+1}^{(2)})\eta_k$ [§]:

$$\mathbb{S}(\tilde{\beta}_{k+1}^{(1)}, \tilde{\beta}_{k+1}^{(2)}) = e^{\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \left(\tilde{L}(\tilde{\beta}_{k+1}^{(1)}) - \tilde{L}(\tilde{\beta}_{k+1}^{(2)})\right)}. \quad (9)$$

Apparently, the unbiased estimators $\tilde{L}(\tilde{\beta}_{k+1}^{(1)})$ and $\tilde{L}(\tilde{\beta}_{k+1}^{(2)})$ in $\mathbb{S}(\tilde{\beta}_{k+1}^{(1)}, \tilde{\beta}_{k+1}^{(2)})$ do not provide an unbiased estimator of $S(\tilde{\beta}_{k+1}^{(1)}, \tilde{\beta}_{k+1}^{(2)})$ after a non-linear transformation as shown in (9), which leads to a large bias.

3.2. Replica Exchange Stochastic Gradient Langevin Dynamics with Correction

A viable MCMC algorithm requires the approximately unbiased estimators of the swapping rates to “satisfy” the

[§]In the implementations, we fix $r\eta_k = 1$ by default.

detailed balance property (Ceperley and Dewing, 1999; Andrieu and Roberts, 2009; Nicholls et al., 2012) and the weak solution of a Markov jump process with unbiased stochastic coefficients has also been studied in Gyöngy (1986); Bentata and Cont (2012). When we make normality assumption on the stochastic energy $\tilde{L}(\tilde{\beta}) \sim \mathcal{N}(L(\tilde{\beta}), \sigma^2)$, it follows

$$\tilde{L}(\tilde{\beta}^{(1)}) - \tilde{L}(\tilde{\beta}^{(2)}) = L(\tilde{\beta}^{(1)}) - L(\tilde{\beta}^{(2)}) + \sqrt{2}\sigma W_1, \quad (10)$$

where W_1 follows the standard normal distribution and can be viewed as a Brownian motion at $t = 1$. Consider the evolution of the stochastic swapping rate $\{\tilde{S}_t\}_{t \in [0,1]}$ in each swap as a geometric Brownian motion:

$$\begin{aligned} \tilde{S}_t &= e^{\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \left(\tilde{L}(\tilde{\beta}^{(1)}) - \tilde{L}(\tilde{\beta}^{(2)}) - \left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \sigma^2 t\right)} \\ &= e^{\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \left(L(\tilde{\beta}^{(1)}) - L(\tilde{\beta}^{(2)}) - \left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \sigma^2 t + \sqrt{2}\sigma W_t\right)}. \end{aligned} \quad (11)$$

Set $\tau_\delta = \frac{1}{\tau_1} - \frac{1}{\tau_2}$ and take the partial derivatives of \tilde{S}_t

$$\frac{d\tilde{S}_t}{dt} = -\tau_\delta^2 \sigma^2 \tilde{S}_t, \quad \frac{d\tilde{S}_t}{dW_t} = \sqrt{2}\tau_\delta \sigma \tilde{S}_t, \quad \frac{d^2\tilde{S}_t}{dW_t^2} = 2\tau_\delta^2 \sigma^2 \tilde{S}_t.$$

Itô’s lemma shows that

$$d\tilde{S}_t = \left(\frac{d\tilde{S}_t}{dt} + \frac{1}{2} \frac{d^2\tilde{S}_t}{dW_t^2} \right) dt + \frac{d\tilde{S}_t}{dW_t} dW_t = \sqrt{2}\tau_\delta \sigma \tilde{S}_t dW_t.$$

Notice that $\{\tilde{S}_t\}_{t \in [0,1]}$ is a Martingale and yields the same expectation for $\forall t \in [0, 1]$. By fixing $t = 1$ in (11), we have

$$\tilde{S}_1 = e^{\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \left(\tilde{L}(\tilde{\beta}^{(1)}) - \tilde{L}(\tilde{\beta}^{(2)}) - \left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \sigma^2\right)}, \quad (12)$$

where the stochastic swapping rate \tilde{S}_1 is an unbiased estimator of $\tilde{S}_0 = e^{\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \left(L(\tilde{\beta}^{(1)}) - L(\tilde{\beta}^{(2)})\right)}$, and the correction term $\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \sigma^2$ aims to remove the bias from the swaps.

An advantage of interpreting the correction term from the perspective of geometric Brownian motion is that we can naturally extend it to geometric Lévy process (Applebaum, 2004), which is more suitable for the heavy-tailed energy noise (Şimşekli et al., 2019). Admittedly, the estimation of the tail-index of extreme-value distributions and the correction under Lévy process go beyond the scope of this paper, so we leave it for future works.

3.3. Adaptive Replica Exchange Stochastic Gradient Langevin Dynamics

In reality, the exact variance σ^2 is hardly known and subject to estimation. The normality assumption may be violated and even no longer time-independent.

3.3.1. FIXED VARIANCE σ^2

We use stochastic approximation (SA) to adaptively estimate the unknown variance while sampling from the posterior. In each SA step, we obtain an unbiased sample variance $\tilde{\sigma}^2$ for the true σ^2 and update the adaptive estimate $\hat{\sigma}_{m+1}^2$ through

$$\hat{\sigma}_{m+1}^2 = (1 - \gamma_m)\hat{\sigma}_m^2 + \gamma_m\tilde{\sigma}_{m+1}^2, \quad (13)$$

where γ_m is the smoothing factor at the m -th SA step. The SA step is updated less frequently than the standard sampling to reduce the computational cost. When the normality assumption holds, we notice that $\hat{\sigma}_m^2 = \sum_{i=1}^m \tilde{\sigma}_i^2 / m$ when $\gamma_m = \frac{1}{m}$. Following central limit theorem (CLT), we have that $\hat{\sigma}_m^2 - \sigma^2 = \mathcal{O}(\frac{1}{\sqrt{m}})$. Inspired by theorem 2 from [Chen et al. \(2015\)](#), we expect that the weak convergence of the adaptive sampling algorithm holds since the bias decreases sufficiently fast ($\frac{1}{m} \sum_{l=1}^m \mathcal{O}(\frac{1}{\sqrt{l}}) \rightarrow 0$ as $m \rightarrow \infty$).

In practice, the normality assumption is likely to be violated when we use a small batch size n , but the unknown distribution asymptotically approximates the normal distribution as $n \rightarrow \infty$ and yield a bias $\mathcal{O}(\frac{1}{n})$ in each SA step. Besides, the mini-batch setting usually introduces a very large noise on the estimator of the energy function, which requires a large correction term and leads to *almost-zero swapping rates*.

To handle this issue, we introduce a correction factor F to reduce the correction term from $(\frac{1}{\tau_1} - \frac{1}{\tau_2})\hat{\sigma}^2$ to $\frac{(\frac{1}{\tau_1} - \frac{1}{\tau_2})\hat{\sigma}^2}{F}$. We note that a large $F > 1$ introduces some bias, but may significantly increase the acceleration effect, giving rise to an acceleration-accuracy trade-off in finite time. Now, we show the algorithm in [Alg. 1](#). In addition to simulations of multi-modal distributions, our algorithm can be also combined with simulated annealing ([Li et al., 2009](#); [Martino et al., 2016](#)) to accelerate the non-convex optimization and increase the hitting probability to the global optima ([Mangoubi and Vishnoi, 2018](#)).

3.3.2. TIME-VARYING VARIANCE σ^2

In practice, the variance σ^2 usually varies with time, resulting in time-varying corrections. For example, in the optimization of residual networks on CIFAR10 and CIFAR100 datasets, we notice from [Fig.3\(a-b\)](#) that the corrections are time-varying. As such, we cannot use a fixed correction anymore to deal with the bias. The treatment for the time-varying corrections includes standard methods for time-series data, and a complete recipe for modeling the data goes beyond our scope. We still adopt the method of stochastic approximation and choose a fixed smoothing factor γ so that

$$\hat{\sigma}_{m+1}^2 = (1 - \gamma)\hat{\sigma}_m^2 + \gamma\tilde{\sigma}_{m+1}^2. \quad (14)$$

Such a method resembles the simple exponential smoothing and acts as robust filters to remove high-frequency noise.

Algorithm 1 Adaptive Replica Exchange Stochastic Gradient Langevin Dynamics Algorithm. For *sampling purposes*, we fix the temperatures τ_1 and τ_2 ; for *optimization purposes*, we keep annealing τ_1 and τ_2 during each epoch. Empirically, a larger γ_m tracks the dynamics better but is less robust. The intensity r and η are omitted in the corrected swaps.

repeat

Sampling Step

$$\begin{aligned} \tilde{\beta}_{k+1}^{(1)} &= \tilde{\beta}_k^{(1)} - \eta_k^{(1)} \nabla \tilde{L}(\tilde{\beta}_k^{(1)}) + \sqrt{2\eta_k^{(1)}\tau_1} \xi_k^{(1)} \\ \tilde{\beta}_{k+1}^{(2)} &= \tilde{\beta}_k^{(2)} - \eta_k^{(2)} \nabla \tilde{L}(\tilde{\beta}_k^{(2)}) + \sqrt{2\eta_k^{(2)}\tau_2} \xi_k^{(2)}, \end{aligned}$$

SA Step

Obtain an unbiased estimate $\tilde{\sigma}_{m+1}^2$ for σ^2 .

$$\hat{\sigma}_{m+1}^2 = (1 - \gamma_m)\hat{\sigma}_m^2 + \gamma_m\tilde{\sigma}_{m+1}^2,$$

Swapping Step

Generate a uniform random number $u \in [0, 1]$.

$$\hat{S}_1 = e^{\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right) \left(\tilde{L}(\tilde{\beta}_{k+1}^{(1)}) - \tilde{L}(\tilde{\beta}_{k+1}^{(2)}) - \frac{(\frac{1}{\tau_1} - \frac{1}{\tau_2})\hat{\sigma}_{m+1}^2}{F} \right)}.$$

if $u < \hat{S}_1$ **then**

Swap $\tilde{\beta}_{k+1}^{(1)}$ and $\tilde{\beta}_{k+1}^{(2)}$.

end if

until $k = k_{\max}$

It can be viewed as a special case of autoregressive integrated moving average (ARIMA) (0,1,1) model but often outperforms the ARIMA equivalents because it is less sensitive to the model selection error ([Bossons, 1966](#)). From the regression perspective, this method can be viewed as a zero-degree local polynomial kernel model ([Gijbels et al., 1999](#)), which is robust to distributional assumptions.

4. Convergence Analysis

We theoretically analyze the acceleration effect and the accuracy of reSGLD in terms of 2-Wasserstein distance between the Borel probability measures μ and ν on \mathbb{R}^d

$$\mathcal{W}_2(\mu, \nu) := \inf_{\Gamma \in \text{Couplings}(\mu, \nu)} \sqrt{\int \|\beta_\mu - \beta_\nu\|^2 d\Gamma(\beta_\mu, \beta_\nu)},$$

where $\|\cdot\|$ is the Euclidean norm, and the infimum is taken over all joint distributions $\Gamma(\beta_\mu, \beta_\nu)$ with μ and ν being its marginal distributions.

Our analysis begins with the fact that reSGLD in [Algorithm.1](#) tracks the replica exchange Langevin diffusion (2). For ease of analysis, we consider a fixed learning rate η for both chains. reSGLD can be viewed as a special discretization of the continuous-time Markov jump process. In

particular, it differs from the standard discretization of the continuous-time Langevin algorithms (Chen et al., 2019; Yin and Zhu, 2010; Raginsky et al., 2017; Sato and Nakagawa, 2014) in that we need to consider the discretization of the Markov jump process in a stochastic environment. To handle this issue, we follow Dupuis et al. (2012) and view the swaps of positions as swaps of the temperatures, which have been proven equivalent in distribution.

Lemma 1 (Discretization Error). *Given the smoothness and dissipativity assumptions in the appendix, and a small learning rate η , we have that*

$$\mathbb{E}[\sup_{0 \leq t \leq T} \|\beta_t - \tilde{\beta}_t^\eta\|^2] \leq \tilde{O}(\eta + \max_i \mathbb{E}[\|\phi_i\|^2] + \max_i \sqrt{\mathbb{E}[\|\psi_i\|^2]}),$$

where $\tilde{\beta}_t^\eta$ is the continuous-time interpolation for reSGLD, $\phi := \nabla \tilde{U} - \nabla U$ is the noise in the stochastic gradient, and $\psi := \tilde{S} - S$ is the noise in the stochastic swapping rate.

Then we quantify the evolution of the 2-Wasserstein distance between ν_t and the invariant distribution π , where ν_t is the probability measure associated with reLD at time t . The key tool is the exponential decay of entropy when π satisfies the log-Sobolev inequality (LSI) (Bakry et al., 2014). To justify LSI, we first verify LSI for reLD without swaps, which is a direct result (Cattiaux et al., 2010) given the Lyapunov function criterion and the Poincaré inequality (Chen et al., 2019). Then we verify LSI for reLD with swaps by analyzing the corresponding Dirichlet form, which is strictly larger than the Dirichlet form associated with reLD without swaps. Finally, the exponential decay of the 2-Wasserstein distance follows from the Otto-Villani theorem (Bakry et al., 2014) by connecting 2-Wasserstein distance with the relative entropy.

Lemma 2 (Accelerated exponential decay of \mathcal{W}_2). *Under the smoothness and dissipativity assumptions, we have that the replica exchange Langevin diffusion converges exponentially fast to the invariant distribution π :*

$$\mathcal{W}_2(\nu_t, \pi) \leq D_0 e^{-k\eta(1+\delta_S)/c_{LS}}, \quad (15)$$

where $\delta_S := \inf_{t>0} \frac{\mathcal{E}_S(\sqrt{\frac{d\nu_t}{d\pi}})}{\mathcal{E}(\sqrt{\frac{d\nu_t}{d\pi}})} - 1$ is the very **acceleration effect** depending on the swapping rate S , \mathcal{E} and \mathcal{E}_S are the Dirichlet forms defined in the appendix, c_{LS} is the constant in the log-Sobolev inequality, $D_0 = \sqrt{2c_{LS}D(\nu_0|\pi)}$.

Finally, combining the definition of Wasserstein distance and the triangle inequality, we have that

Theorem 1 (Convergence of reSGLD). *Let the smoothness and dissipativity assumptions hold. For the distribution $\{\mu_k\}_{k \geq 0}$ associated with the discrete dynamics $\{\tilde{\beta}_k\}_{k \geq 1}$, we have the following estimates for $k \in \mathbb{N}^+$:*

$$\begin{aligned} \mathcal{W}_2(\mu_k, \pi) &\leq D_0 e^{-k\eta(1+\delta_S)/c_{LS}} \\ &\quad + \tilde{O}(\eta^{\frac{1}{2}} + \max_i (\mathbb{E}[\|\phi_i\|^2])^{\frac{1}{2}} + \max_i (\mathbb{E}[\|\psi_i\|^2])^{\frac{1}{4}}), \end{aligned}$$

where $D_0 = \sqrt{2c_{LS}D(\mu_0|\pi)}$, $\delta_S := \min_i \frac{\mathcal{E}_S(\sqrt{\frac{d\mu_i}{d\pi}})}{\mathcal{E}(\sqrt{\frac{d\mu_i}{d\pi}})} - 1$.

Ideally, we want to boost the acceleration effect δ_S by using a larger swapping rate S and increase the accuracy by reducing the mean squared errors $\mathbb{E}[\|\phi_i\|^2]$ and $\mathbb{E}[\|\psi_i\|^2]$. One possible way is to apply a *large enough batch size*, which may be yet inefficient given a large dataset. Another way is to *balance* between acceleration and accuracy by tuning the correction factor F . In practice, a larger F leads to a larger acceleration effect and also injects more biases.

5. Experiments

5.1. Simulations of Gaussian Mixture Distributions

In this group of experiments, we evaluate the acceleration effects and the biases for reSGLD on multi-modal distributions based on different assumptions on the estimators for the energy function. As a comparison, we choose SGLD and the naïve reSGLD without corrections as baselines. The learning rates $\eta^{(1)}$ and $\eta^{(2)}$ are both set to 0.03, and the temperatures τ_1 and τ_2 are set to 1 and 10, respectively. In particular, SGLD uses the learning rate $\eta^{(1)}$ and the temperature τ_1 . We simulate 100,000 samples from each distribution and propose to estimate the correction every 100 iterations. The correction estimator is calculated based on the variance of 10 samples of $\tilde{U}_1(x)$. The initial correction is set to 100 and the step size γ_m for stochastic approximation is chosen as $\frac{1}{m}$. The correction factor F is 1 in the first two examples.

We first demonstrate reSGLD on a simple Gaussian mixture distribution $e^{-U_1(x)} \sim 0.4\mathcal{N}(-3, 0.7^2) + 0.6\mathcal{N}(2, 0.5^2)$, where $U_1(x)$ is the energy function. We assume we can only obtain the unbiased energy estimator $\tilde{U}_1(x) \sim \mathcal{N}(U_1(x), 2^2)$ and the corresponding stochastic gradient at each iteration. From Fig.2(a,b), we see that SGLD suffers from the local trap problem and takes a long time to converge. The naïve reSGLD algorithm alleviated the local trap problem, but is still far away from the ground truth distribution without a proper correction. The naïve reSGLD converges faster than reSGLD in the early phase due to a higher swapping rate, but ends up with a large bias when the training continues. By contrast, reSGLD successfully identifies the right correction through adaptive estimates and yields a close approximation to the ground truth distribution. The high-temperature chain serves as a bridge to facilitate the movement, and the local trap problem is greatly reduced.

In the second example, we relax the normality assumption to a heavy-tail distribution. Given a Gaussian mixture distribution $e^{-U_2(x)} \sim 0.4\mathcal{N}(-4, 0.7^2) + 0.6\mathcal{N}(3, 0.5^2)$, we assume that we can obtain the stochastic energy estimator $\tilde{U}_2(x) \sim U_2(x) + t(\nu = 5)$, where $t(\nu = 5)$ denotes the Student's t-distribution with degree of freedom 5. We see

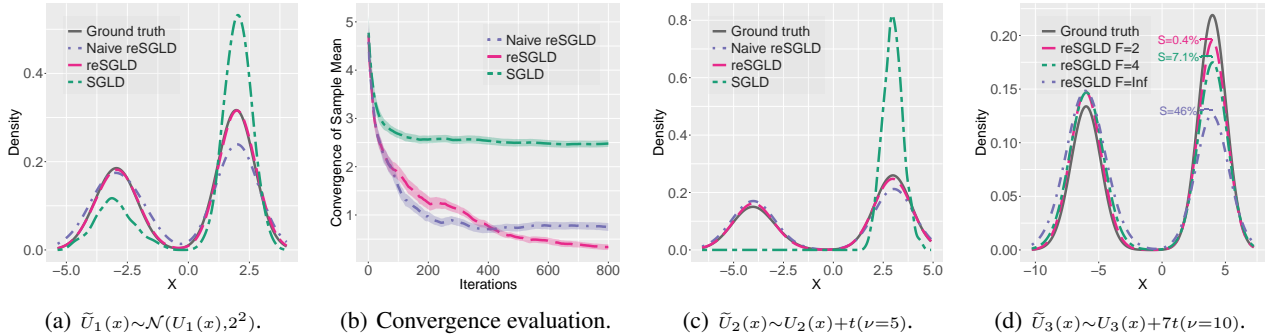


Figure 2. Evaluation of reSGLD on Gaussian mixture distributions, where the naïve reSGLD doesn’t make any corrections and reSGLD proposes to adaptively estimate the unknown corrections. In Fig.1(d), we omit SGLD because it gets stuck in a single mode.

from Fig.2(c) that reSGLD still gives a good approximation to the ground true distribution while the others don’t.

In the third example, we show a case when the correction factor F is useful. We sample from a Gaussian mixture distribution $e^{-U_3(x)} \sim 0.4\mathcal{N}(-6, 0.7^2) + 0.6\mathcal{N}(4, 0.5^2)$. We note that the two modes are far away from each other and the distribution is more difficult to simulate. More interestingly, we assume $\tilde{U}_3(x) \sim U_3(x) + 7t(\nu = 10)$, which requires a large correction term and ends up with no swaps in the finite 100,000 iterations at $F = 1$. In such a case, the unbiased algorithm behaves like the ordinary SGLD as in Fig.2(c) and still suffers from the local trap problems. To achieve larger acceleration effects to avoid local traps and maintain accuracy, we try F at 2, 4 and ∞ (Inf), where the latter is equivalent to the naïve reSGLD. We see from Fig.2(d) that $F = 2$ shows the best approximations, despite that the swapping rate S is only 0.4%. Further increases on the acceleration effect via larger correction factors F give larger swapping rates (7.1% and 46%) and potentially accelerate the convergence in the beginning. However, the biases become more significant as we increase F and lead to larger errors in the end.

5.2. Optimization of Supervised Learning

We evaluate the adaptive replica exchange Monte Carlo on CIFAR10 and CIFAR100, which consist of 50,000 32×32 RGB images for training and 10,000 images for testing. CIFAR10 and CIFAR100 have 10 classes and 100 classes, respectively. We adopt the well-known residual networks (ResNet) and wide ResNet (WRN) as model architectures. We use the 20, 32, 56-layer ResNet (denoted as ResNet-20, et al.), WRN-16-8 and WRN-28-10, where, for example, WRN-16-8 denotes a ResNet that has 16 layers and is 8 times wider than the original. Inspired by the popular momentum stochastic gradient descent, we use stochastic gradient Hamiltonian Monte Carlo (SGHMC) as the baseline sampling algorithm and use the numerical method proposed by Saatici and Wilson (2017) to reduce the tuning cost. We refer to the momentum stochastic gradient descent algorithm

as M-SGD and the adaptive replica exchange SGHMC algorithm as reSGHMC.

We first run several experiments to study the ideal corrections for the optimization of deep neural networks based on the fixed temperatures $\tau_1 = 0.01$ and $\tau_2 = 0.05$. We observe from Fig.(3)(a, b) that the corrections are thousands of times larger than the energy losses, which implies that an exact correction leads to no swaps in practice and no acceleration can be achieved. The desire to obtain more acceleration effects drives us to manually shrink the corrections by increasing F to increase the swapping rates, although we have to suffer from some model bias.

We study the model performance by applying different correction factors F . We choose batch-size 256 and run the experiments within 500 epochs. We first tune the optimal hyperparameters for M-SGD, SGHMC and the low-temperature chain of reSGHMC: we set the learning rate $\eta_k^{(1)}$ to $2e-6$ in the first 200 epochs and decay it afterward by a factor of 0.984 every epoch; the low temperature follows an annealing schedule $\tau_1 = \frac{0.01}{1.02^k}$ to accelerate the optimization; the weight decay is set to 25. Then, as to the high-temperature chain of reSGHMC, we use a larger learning rate $\eta_k^{(2)} = 1.5\eta_k^{(1)}$ and a higher temperature $\tau_2 = 5\tau_1$. Following the dynamic temperatures, we fix $F_k = F_0\alpha^{\mathbb{N}_k}1.02^k$, where \mathbb{N}_k denotes the number of swaps in the first k epochs and $\alpha = 0.8$. The variance estimator is updated each epoch based on the variance of 10 samples of the stochastic energies and the smoothing factor is set to $\gamma = 0.3$ in (14). Consequently, the computations only increase by less than 5%. In addition, we use a thinning factor 200 and report all the results based on Bayesian model averaging. We repeat every experiment five times to obtain the mean and 2 standard deviations.

We see from Fig.3(c,d) that both datasets rely on a very large initial correction factor F_0 to yield good performance and the optimal initial correction factor \hat{F}_0 is achieved at $3e5$. Empirically, we notice that the first five swaps provide the largest marginal improvement in acceleration. A larger F_0

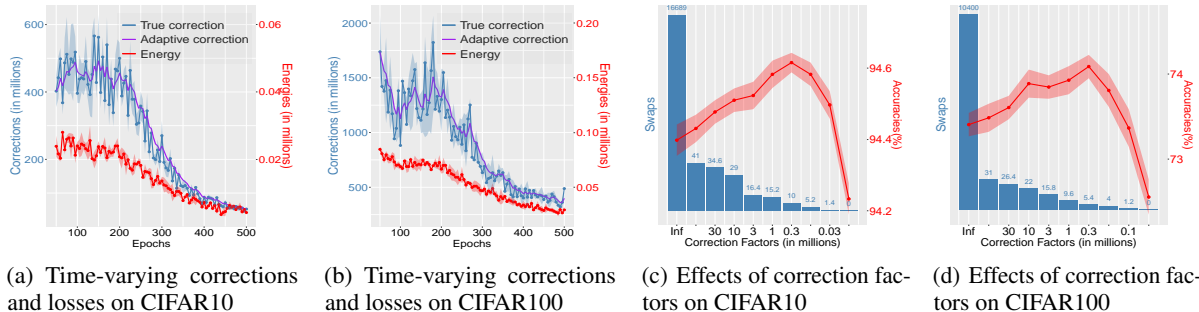


Figure 3. Time-varying variances of the stochastic energy based on batch-size 256 on CIFAR10 & CIFAR100 datasets.

Table 1. PREDICTION ACCURACIES (%) WITH DIFFERENT ARCHITECTURES ON CIFAR10 AND CIFAR100.

MODEL	CIFAR10			CIFAR100		
	M-SGD	SGHMC	RESGHMC	M-SGD	SGHMC	RESGHMC
RESNET-20	94.21±0.16	94.22±0.12	94.62±0.18	72.45±0.20	72.49±0.18	74.14±0.22
RESNET-32	95.15±0.08	95.18±0.06	95.35±0.08	75.01±0.22	75.14±0.28	76.55±0.30
RESNET-56	96.01±0.08	95.95±0.10	96.12±0.06	78.96±0.32	79.04±0.30	80.14±0.34
WRN-16-8	96.71±0.06	96.73±0.08	96.87±0.06	81.70±0.26	82.07±0.22	82.95±0.30
WRN-28-10	97.33±0.08	97.32±0.06	97.42±0.06	83.79±0.18	83.76±0.14	84.38±0.18

than \hat{F}_0 leads to a larger swapping rate with more swaps and thus a larger acceleration effect, however, the performance still decreases as we increase F_0 , implying that the biases start to dominate the error and the diminishing marginal improvement on the acceleration effect is no longer significant. We note that there is only one extra hyper-parameter, namely, the correction factor F , required to tune, and it is independent of the standard SGHMC. This shows that the tuning cost is acceptable.

Table 2. PREDICTION ACCURACIES (%) WITH DIFFERENT BATCH SIZES ON CIFAR10 & CIFAR100 USING RESNET-20.

BATCH	M-SGD	SGHMC	RESGHMC
CIFAR10			
256	94.21±0.16	94.22±0.12	94.62±0.18
1024	94.49±0.12	94.57±0.14	95.01±0.16
CIFAR100			
256	72.45±0.20	72.49±0.18	74.14±0.21
1024	73.31±0.18	73.23±0.20	75.11±0.26

To obtain a comprehensive evaluation of reSGHMC, we use the optimal correction factor for reSGHMC and test it on ResNet20, 32, 56, WRN-16-8 and WRN-28-10. From Table.1, we see that reSGHMC consistently outperforms SGHMC and M-SGD on both datasets, showing the robustness of reSGHMC to various model architectures. For CIFAR10, our method works better with ResNet-20 and ResNet-32 and improves the prediction accuracy by 0.4% and 0.2%, respectively. Regarding the other model architectures, it still slightly outperforms the baselines by roughly 0.1%-0.2%, although this dataset is highly optimized. Specifically, reSGHMC achieves the state-of-the-art 97.42% accuracy with WRN-28-10 model. For CIFAR100,

reSGHMC works particularly well based on various model architectures. It outperforms the baseline by as high as 1.5% using ResNet-20 and ResNet-32, and around 1% based on the other architectures. It also achieves the state-of-the-art 84.38% based on WRN-28-10 on CIFAR100.

We also conduct the experiments using larger batch sizes with ResNet-20 and report the results in Table.2. We run the same iterations and keep the other setups the same. We find that a larger batch size significantly boosts the performance of reSGHMC by as much as 0.4% accuracies on CIFAR10 and 1% on CIFAR100, which shows the *potential of using a large batch size* in the future.

5.3. Bayesian GANs for Semi-supervised Learning

Semi-supervised learning (SSL) is an economic machine learning task because it doesn't require all the data to have pricey labels and still shows promising results. However, the multi-modal problem is more severe in the training of SSLs, such as Bayesian GANs (Saatici and Wilson, 2017), which motivates us to utilize a more powerful algorithm for multi-modal sampling. Therefore, we further evaluate reSGHMC in SSL on CIFAR10, CIFAR100 and the StreetView House Numbers dataset (SVHN) [†] using Bayesian GANs and study how swaps boost the performance.

Regarding the Bayesian GANs used for SSLs, we transform the ordinary discriminator into a K+1-class classifier, where K is the number of classes in each dataset, and $K = 10$ for CIFAR10 and SVHN and $K = 100$ for CIFAR100. During training, a five-layer Bayesian deconvolutional GAN

[†]SVHN consists of 73,257 10-class images for training and 26,032 images for testing.

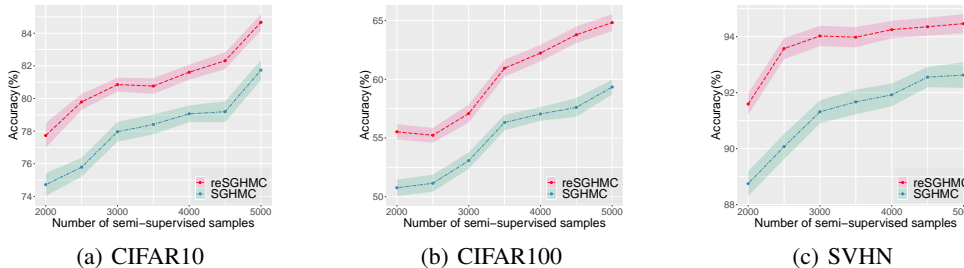


Figure 4. reSGHMC versus SGHMC on benchmark datasets in semi-supervised learning.

Table 3. SEMI-SUPERVISED LEARNING ON CIFAR10, CIFAR100 AND SVHN BASED ON DIFFERENT NUMBER OF LABELS.

N_s	CIFAR10		CIFAR100		SVHN	
	SGHMC	RESGHMC	SGHMC	RESGHMC	SGHMC	RESGHMC
2000	74.72±0.39	77.73±0.31	50.76±0.71	55.53±0.64	88.75±0.44	91.59±0.38
3000	77.96±0.32	80.85±0.23	53.07±0.71	57.09±0.77	91.32±0.41	94.03±0.36
4000	79.06±0.29	81.61±0.24	57.05±0.59	62.23±0.69	91.92±0.41	94.25±0.31
5000	81.74±0.36	84.67±0.28	59.34±0.64	64.83±0.72	92.63±0.46	94.33±0.34

is used as the generator to increase the performance of the discriminator. After training, we discard the generator and use the discriminator for predictions. Since SGHMC is the standard baseline method to simulate from Bayesian GAN, we only compare reSGHMC with SGHMC and no longer report the results based on M-SGD. Following Saatici and Wilson (2017), we take 10 Monte Carlo (MC) samples for the generator and just 1 MC sample for the discriminator. We also simulate 2 samples from SGHMC or reSGHMC for each MC sample. Moreover, we only use 2 chains in the reSGHMC algorithm.

We study the model performance based on different number of labeled data N_s in semi-supervised learning: $N_s = \{2000, 2500, \dots, 5000\}$. Similar to the experiments of supervised learning in Section 5.2, a large correction factor F tends to decrease the performance by injecting more biases, and a smaller correction factor F leads to smaller acceleration effects. Therefore, a good correction factor F is required to balance between the acceleration effects and accuracies. We choose batch-size 64 and detail the other hyper-parameter settings in the Appendix. We decay the learning rate while training, but no longer decrease the temperatures. We report each result using Bayesian model average and repeat each experiment five times to get the average and 2 standard deviations.

As shown in Fig.4 and Table.3, we observe that a larger number of labeled images leads to better performance for all the three datasets. In particular, the 3000 additional labeled images boost the prediction accuracies on CIFAR10, CIFAR100, and SVHN by 7%, 9%, and 3%, respectively. CIFAR10 and CIFAR100 are more sensitive to the labeled images and show larger marginal improvements given a smaller number of labeled images.

Compared to SGHMC, reSGHMC shows a significantly pronounced difference in performance. The consistent improvements are nearly 3% for CIFAR10, 5% for CIFAR100 and 2% for SVHN, respectively. From Table.1 and Table.3, the large improvement in SSL indicates that the multi-modal problem is more severe in Bayesian GANs and the high-temperature chain facilitates the low-temperature chain to jump over distinct modes for the exploration of rich multi-modal distributions. In the end, the low-temperature chain obtains both the exploration ability to traverse the whole domain and the exploitation ability to explore the local geometry, which greatly avoids the mode collapse problems and enables the state-of-the-art performance in SSL.

6. Conclusion and Future Work

We propose the adaptive replica exchange SGMCMC algorithm and prove the accelerated convergence in terms of 2-Wasserstein distance. The theory implies an accuracy-acceleration trade-off and guides us to tune the correction factor F to obtain the optimal performance. We support our theory with extensive experiments and obtain significant improvements over the vanilla SGMCMC algorithms on CIFAR10, CIFAR100, and SVHN.

For future works, it is promising to relax the asymptotic normality assumption to the heavy-tailed generalization of Lévy-stable distribution (Şimşekli et al., 2019) and apply Itô’s lemma to geometric Lévy process to analyze the bias from fat-tailed noises with small batch sizes. Besides, variance reduction (Xu et al., 2018) of the stochastic noise to obtain a larger acceleration effect is also appealing in both theory and practice. From the computational perspective, it is also interesting to study parallel multi-chain reSGMCMC in larger machine learning tasks.

Acknowledgements

We thank the reviewers for their suggestions. We acknowledge the support from the Bilsland Dissertation Fellowship (Deng), the National Science Foundation DMS-1555072, DMS-1736364, DMS-1821233 (Lin) and DMS-1818674 (Liang) and the GPU grant program from NVIDIA.

References

- Christophe Andrieu and Gareth O. Roberts. The Pseudo-Marginal Approach for Efficient Monte Carlo Computations. *The Annals of Statistics*, 37(2):697–725, 2009.
- David Applebaum. *Lévy Processes and Stochastic Calculus*. Cambridge University Press, 2004.
- Dominique Bakry, Ivan Gentil, and Michel Ledoux. Analysis and Geometry of Markov Diffusion Operators. *Springer*, 2014.
- Rémi Bardenet, Arnaud Doucet, and Chris Holmes. On Markov Chain Monte Carlo Methods for Tall Data. *Journal of Machine Learning Research*, 18:1–43, 2017.
- Amel Bentata and Rama Cont. Mimicking the Marginal Distributions of a Semimartingale. *arXiv:0910.3992v5*, 2012.
- Alexandros Beskos, Omiros Papaspiliopoulos, Gareth O. Roberts, and Paul Fearnhead. Exact and Computationally Efficient Likelihood-Based Estimation for Discretely Observed Diffusion Processes. *Journal of the Royal Statistical Society, Series B*, 68:333–382, 2006.
- Gyan V. Bhanot and Anthony D. Kennedy. Bosonic Lattice Gauge Theory with Noise. *Physics Letters B*, 157B:70–76, 1985.
- John Bossons. The Effects of Parameter Misspecification and Non-stationary on The Applicability of Adaptive Forecasts. *Management Science*, 58:659–669, 1966.
- Patrick Cattiaux, Arnaud Guillin, and Li-Ming Wu. A Note on Talagrand’s Transportation Inequality and Logarithmic Sobolev Inequality. *Prob. Theory and Rel. Fields*, 148:285–334, 2010.
- David Ceperley and Mark Dewing. The Penalty Method for Random Walks with Uncertain Energies. *The Journal of Chemical Physics*, 110:9812–9820, 1999.
- Changyou Chen, Nan Ding, and Lawrence Carin. On the Convergence of Stochastic Gradient MCMC Algorithms with High-order Integrators. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2278–2286, 2015.
- Tianqi Chen, Emily B. Fox, and Carlos Guestrin. Stochastic Gradient Hamiltonian Monte Carlo. In *Proc. of the International Conference on Machine Learning (ICML)*, 2014.
- Yi Chen, Jinglin Chen, Jing Dong, Jian Peng, and Zhaoran Wang. Accelerating Nonconvex Learning via Replica Exchange Langevin Diffusion. In *Proc. of the International Conference on Learning Representation (ICLR)*, 2019.
- Umut Şimşekli, Roland Badeau, A. Taylan Cemgil, and Gaë Richard. Stochastic Quasi-Newton Langevin Monte Carlo. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 642–651, 2016.
- Umut Şimşekli, Levent Sagun, and Mert Gürbüzbalaban. A Tail-Index Analysis of Stochastic Gradient Noise in Deep Neural Networks. In *Proc. of the International Conference on Machine Learning (ICML)*, 2019.
- Wei Deng, Xiao Zhang, Faming Liang, and Guang Lin. An Adaptive Empirical Bayesian Method for Sparse Deep Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Nan Ding, Youhan Fang, Ryan Babbush, Changyou Chen, Robert D. Skeel, and Hartmut Neven. Bayesian Sampling Using Stochastic Gradient Thermostats. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3203–3211, 2014.
- Paul Dupuis, Yufei Liu, Nuria Plattner, and J. D. Doll. On the Infinite Swapping Limit for Parallel Tempering. *SIAM J. Multiscale Modeling & Simulation*, 10, 2012.
- David J. Earl and Michael W. Deem. Parallel Tempering: Theory, Applications, and New Perspectives. *Phys. Chem. Chem. Phys.*, 7:3910–3916, 2005.
- Irène Gijbels, Alun Pope, and Matt P. Wand. Understanding Exponential Smoothing via Kernel Regression. *Journal of the Royal Statistical Society, Series B*, 61:39–50, 1999.
- I. Gyöngy. Mimicking the One-dimensional Marginal Distributions of Processes Having an Ito Differential. *Probability Theory and Related Fields*, 71:501–516, 1986.
- Scott Kirkpatrick, D. Gelatt Jr, and Mario P. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598): 671–680, 1983.
- Anoop Korattikara, Yutian Chen, and Max Welling. Austerity in MCMC Land: Cutting the Metropolis-Hastings Budget. In *Proc. of the International Conference on Machine Learning (ICML)*, 2014.
- Holden Lee, Andrej Risteski, and Rong Ge. Beyond Log-concavity: Provable Guarantees for Sampling

- Multi-modal Distributions using Simulated Tempering Langevin Monte Carlo. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Chunyu Li, Changyou Chen, David Carlson, and Lawrence Carin. Preconditioned Stochastic Gradient Langevin Dynamics for Deep Neural Networks. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, pages 1788–1794, 2016.
- Xuechen Li, Denny Wu, Lester Mackey, and Murat A. Erdogdu. Stochastic Runge-Kutta Accelerates Langevin Monte Carlo and Beyond. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 7746–7758, 2019.
- Yaohang Li, Vladimir A. Protopopescu, Nikita Arnold, Xinyu Zhang, and Andrey Gorin. Hybrid Parallel Tempering and Simulated Annealing Method. *Applied Mathematics and Computation*, 212(1):216–228, 2009.
- Faming Liang, Chuanhai Liu, and Raymond J. Carroll. Stochastic Approximation in Monte Carlo Computation. *Journal of the American Statistical Association*, 102:305–320, 2007.
- Yi-An Ma, Tianqi Chen, and Emily B. Fox. A Complete Recipe for Stochastic Gradient MCMC. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- Oren Mangoubi and Nisheeth K. Vishnoi. Convex Optimization with Unbounded Nonconvex Oracles using Simulated Annealing. In *Proc. of Conference on Learning Theory (COLT)*, 2018.
- E Marinari and G Parisi. Simulated Tempering: A New Monte Carlo Scheme. *Europhysics Letters (EPL)*, 19(6): 451–458, 1992.
- L. Martino, V. Elvira, D. Luengo, J. Corander, and F. Louzada. Orthogonal Parallel MCMC Methods for Sampling and Optimization. *Digital Signal Processing*, 58:64–84, 2016.
- Geoff K. Nicholls, Colin Fox, and Alexis Muir Watt. Coupled MCMC with a Randomized Acceptance Probability. *ArXiv e-prints*, 2012.
- Matias Quiroz, Robert Kohn, Mattias Villani, and Minh-Ngoc Tran. Speeding Up MCMC by Efficient Data Sub-sampling. *Journal of the American Statistical Association*, 114:831–843, 2019.
- Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex Learning via Stochastic Gradient Langevin Dynamics: a Nonasymptotic Analysis. In *Proc. of Conference on Learning Theory (COLT)*, June 2017.
- Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- Yunus Saatci and Andrew G Wilson. Bayesian GAN. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3622–3631, 2017.
- Issei Sato and Hiroshi Nakagawa. Approximation Analysis of Stochastic Gradient Langevin Dynamics by using Fokker-Planck Equation and Itô Process. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 982–990, 2014.
- Daniel Seita, Xinlei Pan, Haoyu Chen, and John Canny. An Efficient Minibatch Acceptance Test for Metropolis-Hastings. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.
- Robert H. Swendsen and Jian-Sheng Wang. Replica Monte Carlo Simulation of Spin-Glasses. *Phys. Rev. Lett.*, 57: 2607–2609, 1986.
- Yee Whye Teh, Alexandre Thiéry, and Sebastian Vollmer. Consistency and Fluctuations for Stochastic Gradient Langevin Dynamics. *Journal of Machine Learning Research*, 17:1–33, 2016.
- Max Welling and Yee Whye Teh. Bayesian Learning via Stochastic Gradient Langevin Dynamics. In *Proc. of the International Conference on Machine Learning (ICML)*, pages 681–688, 2011.
- Wing Hung Wong and Faming Liang. Dynamic Weighting in Monte Carlo and Optimization. *Proc. Natl. Acad. Sci.*, 94:14220–14224, 1997.
- Pan Xu, Jinghui Chen, Difan Zou, and Quanquan Gu. Global Convergence of Langevin Dynamics Based Algorithms for Nonconvex Optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- George Yin and Chao Zhu. *Hybrid Switching Diffusions: Properties and Applications*. Springer, 2010.
- Yuchen Zhang, Percy Liang, and Moses Charikar. A Hitting Time Analysis of Stochastic Gradient Langevin Dynamics. In *Proc. of Conference on Learning Theory (COLT)*, pages 1980–2022, 2017.