
Noise and Fluctuation of Finite Learning Rate Stochastic Gradient Descent

Kangqiao Liu^{*1} Liu Ziyin^{*1} Masahito Ueda¹²³

Abstract

In the vanishing learning rate regime, stochastic gradient descent (SGD) is now relatively well understood. In this work, we propose to study the basic properties of SGD and its variants in the non-vanishing learning rate regime. The focus is on deriving exactly solvable results and discussing their implications. The main contributions of this work are to derive the stationary distribution for discrete-time SGD in a quadratic loss function with and without momentum; in particular, one implication of our result is that the fluctuation caused by discrete-time dynamics takes a distorted shape and is dramatically larger than a continuous-time theory could predict. Examples of applications of the proposed theory considered in this work include the approximation error of variants of SGD, the effect of minibatch noise, the optimal Bayesian inference, the escape rate from a sharp minimum, and the stationary covariance of a few second-order methods including damped Newton’s method, natural gradient descent, and Adam.

1. Introduction

Behind the success of deep learning lies the simple optimization methods such as stochastic gradient descent (SGD) (Bottou, 1998; Sutskever et al., 2013; Dieuleveut et al., 2020; Mori and Ueda, 2020a) and its variants (Duchi et al., 2011; Flammarion and Bach, 2015; Kingma and Ba, 2015), which are used for neural network optimization. Despite the empirical efficiency of SGD, our theoretical understanding of SGD is still limited. Two types of noises for SGD are studied. When the noise is white, the dynamics is governed by the stochastic gradient Langevin dynamics (SGLD) (Welling

and Teh, 2011). When the noise is due to minibatch sampling, the noise is called the SGD noise or minibatch noise. So far, nearly all the theoretical attempts at understanding the noise in SGD have adopted the continuous-time approximation by assuming a vanishingly small learning rate (Mandt et al., 2017; Li et al., 2017; Jastrzebski et al., 2017; Chaudhari and Soatto, 2018; Zhu et al., 2019; Xie et al., 2021). This amounts to making an analogy to the diffusion theory in physics (Einstein, 1905; Van Kampen, 1992), and helps understand some properties of SGD and deep learning such as the flatness of the minima selected by training (Jastrzebski et al., 2017; Chaudhari and Soatto, 2018; Smith and Le, 2018; Xie et al., 2021). However, in reality, a large learning rate often leads to qualitatively distinct behavior, including reduced training time and potentially better generalization performance (Keskar et al., 2017; Li et al., 2019). The present work is motivated by the fact that the existing continuous theory is insufficient to describe and predict the properties and phenomena of learning in this large learning rate regime. In fact, we will show that the prediction by continuous theory may deviate arbitrarily from the experimental result.

In this work, we study the stationary-state solutions of discrete-time update rules of SGD. The result can be utilized to analyze SGD without invoking the unrealistic assumption of a vanishingly small learning rate. Specifically, our contributions are:

- We exactly solve the discrete-time update rules for SGD and its variants with momentum in a local minimum to obtain the analytic form of the covariance matrix of the model parameters at asymptotic time.
- We apply our results to various settings that have been studied in the continuous-time limit, such as finding the optimal learning rate in a Bayesian setting, understanding the escape from a sharp minimum, and the approximation error of various variants of SGD.
- Compared with the continuous-time theories, our work requires fewer assumptions and finds significantly improved agreement with experimental results.

In section 2, we present the background of this work. Related works are discussed in section 3. In section 4, we derive our main theoretical results for SGD and its momentum variant. In section 5, we verify our theoretical results

^{*}Equal contribution ¹Department of Physics, the University of Tokyo, Japan ²RIKEN CEMS, Japan ³Institute for Physics of Intelligence, the University of Tokyo, Japan. Correspondence to: Kangqiao Liu <kqliu@cat.phys.s.u-tokyo.ac.jp>, Liu Ziyin <zliu@cat.phys.s.u-tokyo.ac.jp>.

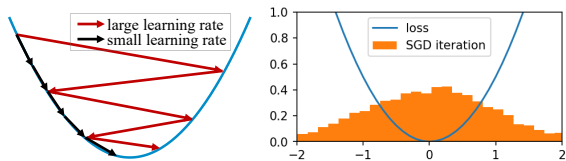


Figure 1: (Left) Schematic illustrations of the deterministic continuous-time evolution given by $w = w_0 e^{-\lambda k t}$ for small λ and the deterministic discrete-time evolution given by Eq. (3) for $1/k < \lambda < 2/k$. The black arrows represent each update according to the continuous-time update rule while the red ones represent the discrete-time updates. (Right) When noise exists, SGD iteration converges not to a point but a distribution.

experimentally. In section 6, we apply our solution to some well-known problems that have previously been investigated in the continuous-time limit. A summary of our results is given in Table 1.

2. Background

In the presence of noise, it is difficult to give a simple solution to the discrete dynamics. Under the assumptions of a constant Gaussian noise, a quadratic loss function, and an infinitesimal learning rate, theorists approximate the multidimensional update rule by a continuous-time stochastic differential equation (Mandt et al., 2017; Zhu et al., 2019), which is a multivariate Ornstein-Uhlenbeck process, $d\mathbf{w}_t = -\lambda K \mathbf{w}_t dt + \lambda C^{\frac{1}{2}} dW_t$, where C is the covariance matrix of the noise, W_t represents a standard multidimensional Brownian motion, and K is the Hessian of the local minimum¹. The stationary covariance matrix, $\Sigma := \lim_{t \rightarrow \infty} \text{cov}[w_t^T, w_t]$, is found to satisfy the following matrix equation (Van Kampen, 1992):

$$\Sigma K + K \Sigma = \frac{\lambda}{1 - \mu} C, \quad (1)$$

where μ is the momentum hyperparameter; when no momentum is used, $\mu = 0$. Despite the fact that a number of theoretical works has been performed on the basis of the above constant Gaussian noise and the continuous-time approximation (Chaudhari and Soatto, 2018; Jacot et al., 2018; Zhu et al., 2019; Lee et al., 2020; Xie et al., 2021), it is clear that the stationary distribution given by Eq. (1) substantially deviates from the true one obtained in experiments². The predictions based on these results are qualitatively ac-

¹In this work, we assume K to be full-rank; this is true in practice since, for applications, it is a standard practice to apply a very small L_2 regularization, which amounts to adding a small positive value to all the eigenvalues of K , making it full-rank. For more discussion, see Appendix B.

²For example, Figure 1 in Mandt et al. (2017) plots stationary distributions obtained from solving stochastic differential equations. These distributions generally deviate from experimental results. In Gitman et al. (2019), predictions also deviate far from experiments when the learning rate is large.

ceptable only in a small learning rate regime. For a large learning rate, the assumptions simply break down so that the theory becomes invalid.

To intuitively understand how a large learning rate makes a difference, we consider a model with a single parameter $w \in \mathbb{R}$ in a quadratic potential $L = \frac{1}{2} k w^2$ with $k > 0$. SGD obeys the dynamical equations as follows:

$$\begin{cases} g_t = k w_{t-1} + \eta_t; \\ w_t = w_{t-1} - \lambda g_t, \end{cases} \quad (2)$$

where λ is the learning rate and η_t is a normal random variable with zero mean and variance σ^2 . When $\sigma^2 = 0$, the dynamics is deterministic, and the common approach is to assume that $\lambda \ll 1/k$ such that one may take the continuous-time limit of this equation as $\dot{w} = -\lambda k w(t)$, which solves to give $w = w_0 e^{-\lambda k t}$. However, this continuous approximation fails when λ is large. To see this, we note that the deterministic discrete-time dynamics solves to give

$$w_t = (1 - \lambda k)^t w_0, \text{ for } t \in \mathbb{N}^0, \quad (3)$$

which is an exponential decay when $\lambda < \frac{1}{k}$, and, in this region, the standard continuous-time dynamics is valid with error $O((1 - \lambda k)^2)$. On the other hand, when $\lambda > \frac{2}{k}$, the learning is so large that the parameter w will diverge; therefore, the interesting region is when $\frac{1}{k} < \lambda < \frac{2}{k}$, where the dynamics is convergent, yet a simple continuous approximation fails. See a schematic illustration in Figure 1. It is therefore urgent to develop a theory that can handle SGD with a large learning rate.

3. Related Works

Large Learning Rate. Although continuous-time theory has been the dominant theoretical approach, Lewkowycz et al. (2020) took a step forward in understanding why a larger learning rate may generalize better. They characterized SGD into three regimes according to the learning rate. They conjectured that a rather large initial learning rate leads to a ‘‘catapult phase’’, which helps exploration and often leads to better generalization by converging to a flatter minimum. However, their work is mostly empirical and in the noise-free regime. There are more empirical works on large learning rate and generalization. LeCun et al. (2012) found that a large batch size (or a small learning rate) usually leads to reduced generalization performance. Keskar et al. (2017) proposed an empirical measure based on the sharpness of a minimum. They presented numerical evidence that a small learning rate prefers sharp minima that generalize poorly. Goyal et al. (2017) showed that setting the learning rate proportional to the minibatch size ensures good generalization, which is crucial for large scale training. There are also works about how noise, the batch size and the learning rate influence the generalization (Hoffer et al., 2017; Mori and Ueda, 2020a;b). The explanations

Noise and Fluctuation of Finite Learning Rate Stochastic Gradient Descent

Table 1: Comparison of the results of this work with previous results. For notational conciseness, we only show Σ when the noise matrix C commutes with the Hessian K . For the *approximation error* panel, the results apply to any K and C . *SGD: stochastic gradient descent. *SGDM: stochastic gradient descent with momentum. *QHM: quasi-hypobolic momentum. *DNM: damped newton’s method. *NGD: natural gradient descent.

	Previous Work	This Work
	Σ	Σ
SGD	$\frac{\lambda}{2} K^{-1} C$ (Mandt et al., 2017)	$\lambda [K(2I_D - \lambda K)]^{-1} C$
SGDM	$\frac{\lambda}{2(1-\mu)} K^{-1} C$ (Mandt et al., 2017)	$\lambda \left[\frac{K}{1+\mu} \left(2I_D - \frac{\lambda K}{1+\mu} \right) \right]^{-1} \frac{C}{1-\mu^2}$
QHM	$\frac{\lambda}{2(1-\mu)} K^{-1} C + O(\lambda^2)$ (Gitman et al., 2019)	$\lambda^2 h(K)^{-1} C$ [Eq. (90)]
DNM	-	$\frac{1+\mu}{1-\mu} \frac{\lambda}{2(1+\mu)-\lambda} K^{-2} C$
NGD	-	$\frac{1}{2} K^{-1} \left[Q + \frac{\lambda}{2(1+\mu)} I_D \right]$ [Eq. (107)]
Adam	-	$\frac{\lambda^2(1+c)}{4} I_D$
	Approximation Error	Approximation Error
SGD	$\frac{\lambda}{4} \text{Tr}[C] + \frac{\lambda^2}{8} \text{Tr}[KC] + O(\lambda^3)$ (Gitman et al., 2019)	$\frac{\lambda}{2} \text{Tr}[(2I_D - \lambda K)^{-1} C]$
SGDM	$\frac{\lambda}{4(1-\mu)} \text{Tr}[C] + \frac{\lambda^2}{8} \frac{1}{1-\mu^2} \text{Tr}[KC] + O(\lambda^3)$ (Gitman et al., 2019)	$\frac{\lambda}{2(1-\mu)} \text{Tr} \left[\left(2I_D - \frac{\lambda}{1+\mu} K \right)^{-1} C \right]$
QHM	$\frac{\lambda}{4} \text{Tr}[C] + \frac{\lambda^2}{8} \left[1 + \frac{2\mu\nu}{1-\mu} \left(\frac{2\mu\nu}{1+\mu} - 1 \right) \right] \text{Tr}[KC] + O(\lambda^3)$ (Gitman et al., 2019)	$\frac{\lambda^2}{2} \text{Tr}[h(K)^{-1} KC]$ [Eq. (89)]
DNM	-	$\frac{1+\mu}{1-\mu} \frac{\lambda}{4(1+\mu)-2\lambda} \text{Tr}[K^{-1} C]$
NGD	-	$\frac{1}{4} \text{Tr} \left[Q + \frac{\lambda}{2(1+\mu)} I_D \right]$
Adam	-	$\frac{\lambda^2(1+c)}{8} \text{Tr}[K]$

of why a flatter minimum generalizes better are given by some theories such as the minimum description length theory (Rissanen, 1983), a Bayesian view of learning (MacKay, 1992), and the Gibbs free energy (Chaudhari et al., 2019).

Escape from Sharp Minimum. Theoretically understanding why and how SGD converges to flat minima is crucially important (Hochreiter and Schmidhuber, 1997). In fact, among many complexity measures characterizing generalization (Dziugaite and Roy, 2017; Neyshabur et al., 2018a; Smith and Le, 2018; Chaudhari and Soatto, 2018; Keskar et al., 2017; Liang et al., 2019; Nagarajan and Kolter, 2019), the sharpness-based measures have been shown to be the best up to now (Jiang et al., 2019). The continuous approximation is usually adopted to understand how SGD chooses flat minima. Hu et al. (2019) used diffusion theory to show that escape is easier from a sharp minimum than from a flat one. Wu et al. (2018) studied the relationship among learning rate, batch size and generalization from the perspective of dynamical stability. Jastrzebski et al. (2017) used stochastic differential equations to prove that the higher the ratio of the learning rate to the batch size, the flatter minimum will be selected. Zhu et al. (2019) defined the escape efficiency for a minimum and obtained its explicit expression using diffusion theory. They show that anisotropic noise helps escape from sharp minima effectively. A recent work by Xie et al. (2021) calculated the escape rate from a minimum by adopting the formalism of the Kramers escape rate in physics (Kramers, 1940). It is shown that SGD with minibatch noise favors flatter minima. Our exact results for a large learning rate make it possible to study the selected flatness and the complexity measures more accurately and may enhance our understanding of deep learning.

Bayesian Inference. SGD has been used for Bayesian inference as well. In Bayesian inference, one assumes a probabilistic model $p(w, x)$ with data x and hidden parameter w . The goal is to approximate the posterior $p(w|x)$. Traditionally, stochastic gradient Markov Chain Monte Carlo (MCMC) methods have been used (Welling and Teh, 2011; Ma et al., 2015). A similarity between SGD and MCMC suggests the possibility of SGD being used as approximate Bayesian inference. Mandt et al. (2017) applied SGD to minimize the loss function defined as $-\ln p(w, x)$. They show that one can tune the learning rate such that the Kullback-Leibler (KL) divergence between the learned distribution by SGD and the posterior is minimized. This means that SGD can be regarded as an approximate Bayesian inference. However, for a large learning rate, their assumption is no more valid.

4. Theory of Discrete-Time SGD

We propose to deal with the discrete-time SGD directly. We use $\mathbf{w} \in \mathbb{R}^D$ to denote the weights of the model viewed as a vector, and the boldness is dropped when $D = 1$. We use capital $K \in \mathbb{R}^{D \times D}$ to denote the Hessian matrix of the quadratic loss function; when $D = 1$, we use the lower-case k . We use $\Sigma \in \mathbb{R}^{D \times D}$ to denote the asymptotic covariance matrix of \mathbf{w} , and $C \in \mathbb{R}^{D \times D}$ to denote the covariance matrix of a general type of noise $\eta \in \mathbb{R}^D$. When the learning rate λ is not a scalar but a matrix (sometimes called preconditioning matrix), we use the upper-case letter $\Lambda \in \mathbb{R}^{D \times D}$ instead of λ . L denotes the training loss function, and S the minibatch size. We use $\mathbb{E}[\cdot]$ to denote the expectation over the stationary distribution of the model parameters. The capital T is used as a superscript to denote matrix transpose

and lower case t is used to denote the time step of optimization. Due to space constraint, we leave derivations to Appendix E.1.

4.1. SGD

Consider a general loss function $\mathcal{L}(\mathbf{w}')$ for a general differentiable model with parameters $\mathbf{w} \in \mathbb{R}^D$; close to a local minimum, we may expand $\mathcal{L}(\mathbf{w}')$ up to second order in \mathbf{w}' . Therefore, close to a local minimum, the dynamics of SGD is governed by a general form of a quadratic potential:

$$\mathcal{L}(\mathbf{w}') \approx \frac{1}{2}(\mathbf{w}' - \mathbf{w}^*)^\top K(\mathbf{w}' - \mathbf{w}^*) := L(\mathbf{w}'), \quad (4)$$

where the Hessian K is a positive definite matrix, and \mathbf{w}^* is a constant vector. One can redefine $\mathbf{w}' - \mathbf{w}^* \rightarrow \mathbf{w}$ to obtain a simplified form: $L(\mathbf{w}) = \mathbf{w}^\top K \mathbf{w} / 2$. The SGD algorithm with momentum μ is defined by the update rule

$$\begin{cases} \mathbf{g}_t = \nabla L(\mathbf{w}_{t-1}) + \eta_{t-1} = K \mathbf{w}_{t-1} + \eta_{t-1}; \\ \mathbf{m}_t = \mu \mathbf{m}_{t-1} + \mathbf{g}_t; \\ \mathbf{w}_t = \mathbf{w}_{t-1} - \lambda \mathbf{m}_t, \end{cases} \quad (5)$$

where the noise η_t has a finite covariance C , and $\mu \in [0, 1)$ is the momentum hyperparameter. We first consider the case without momentum, i.e., $\mu = 0$. Let k^* be the largest eigenvalue of K . For any minimum with $\lambda k^* > 2$, the dynamics will diverge, and \mathbf{w} will escape from this minimum. Therefore, we focus on the range of $0 < \lambda < \frac{2}{k^*}$. This means that the absolute values of all the eigenvalues of $(I_D - \lambda K)$ are in the range of $(0, 1)$.

Theorem 1. (Model fluctuations of discrete SGD in a quadratic potential) Let \mathbf{w}_t be updated according to (5) with $\mu = 0$. The stationary covariance of \mathbf{w} satisfies

$$\Sigma K + K \Sigma - \lambda K \Sigma K = \lambda C. \quad (6)$$

Recall that we have shifted the underlying parameter \mathbf{w}' by \mathbf{w}^* , and this result translates to that $\mathbb{E}[\mathbf{w}'] = \mathbf{w}^*$, close to the local minimum \mathbf{w}^* .

Remark. The exact condition $\lambda \Sigma K + \lambda K \Sigma - \lambda^2 K \Sigma K = \lambda^2 C$ is different from the classical result obtained from a continuous Ornstein-Uhlenbeck process, which has $\lambda \Sigma K + \lambda K \Sigma = \lambda^2 C$. This suggests that the approximation of a discrete-time dynamics with a continuous-time one in Mandt et al. (2017) can be thought of as a λ -first-order approximation to the true dynamics. This approximation incurs an error of order $\mathcal{O}(\lambda^2)$. The error becomes significant or even dominant as λ gets large.

Remark. This theorem only requires the existence of a finite stationary noise covariance, in contrast to the unrealistic assumption of constant Gaussian noise by Mandt et al. (2017). In Appendix C, we show that the agreement of the theory is as good when the noises are non-gaussian.

4.2. SGD with Momentum

We now consider the case with arbitrary $\mu \in [0, 1)$ in Eq. (5).

Theorem 2. (Stationary distribution of discrete SGD with momentum) Let \mathbf{w}_t be updated according to (5) with arbitrary $\mu \in [0, 1)$. Then Σ is the solution to

$$\underbrace{-\frac{1+\mu^2}{1-\mu^2} \lambda^2 K \Sigma K + \frac{\mu}{1-\mu^2} \lambda^2 (K^2 \Sigma + \Sigma K^2)}_{\text{discrete-time}} + \underbrace{(1-\mu) \lambda (K \Sigma + \Sigma K)}_{\text{continuous-time}} = \lambda^2 C. \quad (7)$$

If the noise is Gaussian, then the stationary state distribution of \mathbf{w} is $\mathbf{w} \sim \mathcal{N}(0, \Sigma)$.

We examine the above result (7) with two limiting examples. First, if there is no momentum, namely $\mu = 0$, we recover the previous result (6) without momentum. Next, if $\lambda K \ll 1$, neglecting the $\mathcal{O}((\lambda K)^2)$ terms recovers the result of SGD with momentum under the continuous approximation (1) derived by Mandt et al. (2017). When C commutes with K , the above matrix equation can be solved explicitly.

Corollary 1. Let $[C, K] = 0$. Then

$$\Sigma = \left[\frac{\lambda K}{1+\mu} \left(2I_D - \frac{\lambda K}{1+\mu} \right) \right]^{-1} \frac{\lambda^2 C}{1-\mu^2}. \quad (8)$$

We can obtain a more general result when the learning rate is a matrix.

Theorem 3. If the learning rate is a positive definite preconditioning matrix Λ , then Σ satisfies

$$\begin{aligned} & -\frac{1+\mu^2}{1-\mu^2} \Lambda K \Sigma K \Lambda + \frac{\mu}{1-\mu^2} (\Lambda K \Lambda K \Sigma + \Sigma K \Lambda K \Lambda) \\ & + (1-\mu) (\Lambda K \Sigma + \Sigma K \Lambda) = \Lambda C \Lambda. \end{aligned} \quad (9)$$

Note that the matrix Λ does not necessarily commute with K . We consider an application of this general result to understanding second-order methods in section 6.4.

4.3. Two Typical Kinds of Noise

As mentioned, two specific types of noise are of particular interest for machine learning practices. The first is a multidimensional white noise: $\eta \sim \mathcal{N}(0, \sigma^2 I_D)$, where σ is a positive scalar. The second type of noise is of Gaussian type with a covariance proportional to the Hessian, which is approximately equal to the noise caused by a minibatch gradient descent algorithm (Zhu et al., 2019; Xie et al., 2021): $C = aK$, where a is a constant scalar, because

$$C(\mathbf{w}) \approx \frac{1}{S} \frac{1}{N} \sum_{n=1}^N \nabla L(\mathbf{w}, x_n) \nabla L(\mathbf{w}, x_n)^\top := \frac{1}{S} J(\mathbf{w}) \approx \frac{1}{S} K,$$

where the first approximation is due to the fact that noise dominates the dynamics in the vicinity of a minimum, and the second approximation is according to the similarity between the Fisher information $J(\mathbf{w})$ and the Hessian K near a minimum. This approximation is somewhat crude although it is often employed.

To be more general, one might wish to mix an isotropic Gaussian noise with the minibatch noise, namely $C = \sigma^2 I_D + aK$. The following theorem gives the distribution.

Theorem 4. *Let $C = \sigma^2 I_D + a(\lambda)K$ and $\mu = 0$. Then the stationary distribution of \mathbf{w} is*

$$\mathcal{N}\left(0, \lambda(\sigma^2 I_D + aK)[K(2I_D - \lambda K)]^{-1}\right). \quad (10)$$

Notice that, in this case, C commutes with K . From this result, we can derive the following two special cases by setting $\sigma^2 = 0$ or $a = 0$.

Corollary 2. *Let $\sigma^2 = 0$. Then $\Sigma = a\lambda(2I_D - \lambda K)^{-1}$.*

Corollary 3. *Let $a = 0$. Then $\Sigma = \sigma^2\lambda[K(2I_D - \lambda K)]^{-1}$.*

Interestingly, when $\lambda \ll 1$, the minibatch noise results in isotropic fluctuations, independent of the underlying geometry; the discrete time steps, however, causes fluctuations in the direction of large eigenvalues of the Hessian.

4.4. Approximation Error of SGD

We note that one important quantity for measuring the approximation error of SGD is $\text{Tr}[K\Sigma]$, because the expectation of a quadratic loss is

$$L_{\text{train}} := \mathbb{E}\left[\frac{1}{2}\mathbf{w}^T K \mathbf{w}\right] = \frac{1}{2}\text{Tr}[K\Sigma], \quad (11)$$

where the expectation is taken over the stationary distribution of \mathbf{w} .

Theorem 5. *(Approximation error for discrete SGD with momentum) Let the noise covariance C and Hessian K be any positive definite matrix. Then the training error for SGD with momentum is*

$$L_{\text{train}} = \frac{\lambda}{4(1-\mu)} \text{Tr}\left[\left(I_D - \frac{\lambda}{2(1+\mu)}K\right)^{-1} C\right]. \quad (12)$$

This result suggests that a larger eigenvalue in the Hessian causes larger training error. Also, compared with the continuous result: $L_{\text{train}} = \frac{\lambda}{4(1-\mu)}\text{Tr}[C]$, the discrete theory results in a larger approximation error.

4.5. Necessary Stability Condition

The main result in Theorem 2 also suggests a condition for the convergence of SGD. In order for a stationary distribution to exist at convergence, the covariance Σ needs to exist and be positive definite, and this is only possible when $\left(2I_D - \frac{\lambda}{1+\mu}K\right)$ is positive definite. This condition reflects the fact that discrete-time SGD becomes *ill-conditioned* as λ increases, and so the continuous-time approximation becomes less valid. Also, an important implication is that using momentum may mitigate the ill-conditioning of the large learning rate, but only up to a factor of $1 + \mu < 2$, before the momentum causes another divergence problem due to the term $\frac{1}{1-\mu}$. Therefore, when momentum is used, the necessary condition for convergence becomes $\lambda k^* \leq 2(1 + \mu) < 4$. We also comment that this is only a necessary condition for stability; the sufficient condition of stability is highly complicated and we leave this to a future work.

4.6. Regularization Effect of a Finite Learning Rate

For the continuous dynamics, the stationary distribution is known to obey the Boltzmann distribution, $P(\mathbf{w}) \sim \exp[-L(\mathbf{w})/T]$ for some scalar T determined by the noise strength (Landau and Lifshitz, 1980; Zhang et al., 2018). This implies that, close to a local minimum \mathbf{w}_* , the stationary distribution is approximated by

$$P(\mathbf{w}|\mathbf{w}^*) \sim \exp[-(\mathbf{w} - \mathbf{w}^*)^T K (\mathbf{w} - \mathbf{w}^*)/T]. \quad (13)$$

Comparing with the standard continuous-time solution (1), we see that this corresponds to an isotropic noise, namely $C = 2TI_D$, and $L_{\text{eff}} := -T \log P(\mathbf{w}|\mathbf{w}^*)$ may be defined as an effective loss function in analogy with an effective free energy in theoretical physics.

For discrete-time SGD, however, the stationary distribution has an additional term to leading order:

$$P_d(\mathbf{w}|\mathbf{w}^*) = \exp\left[-\frac{1}{T}(\mathbf{w} - \mathbf{w}^*)^T \left(K + \frac{\lambda}{2}K^2\right)(\mathbf{w} - \mathbf{w}^*) + O(\lambda^2)\right].$$

This implies a different form of the effective loss function:

$$L_{\text{deff}} \sim (\mathbf{w} - \mathbf{w}^*)^T K (\mathbf{w} - \mathbf{w}^*) + \frac{\lambda}{2}(\mathbf{w} - \mathbf{w}^*)^T K^2 (\mathbf{w} - \mathbf{w}^*),$$

where the first term is the same as the continuous-time loss function, while the second term emerges as a discrete-time contribution due to a large learning rate. In particular, it encourages \mathbf{w} to have a smaller norm around the minimum \mathbf{w}_* in the kernel K^2 . Therefore, to first order, the effect of the large learning rate can be understood as an effective weight decay term that encourages a smaller norm.

4.7. State-Dependent Noise

In reality, the minibatch noise depends on the current value of \mathbf{w} , and the noise is, therefore, not constant (Simsekli et al., 2019; 2020; Hodgkinson and Mahoney, 2020; Meng et al., 2020; Ziyin et al., 2021; Mori et al., 2021). This means that the noise covariance is a function of the parameters covariance, namely $C = C(\Sigma)$. The general Theorem 2 is still applicable to this setting, and Σ can be solved by setting $C = C(\Sigma)$, provided that Σ exists. One example of state-dependent noise is given in section 6.4 where the noise covariance is set to $C \approx K\Sigma K$. After the publication of this work, we notice that Ziyin et al. (2021) has applied our formalism to study the exact shape and strength of SGD noise when minibatch sampling is used, and we refer the readers to this work for a more detailed study of the state-dependent noise in SGD.

5. Experiments

Gaussian Noise. We first consider the case when $w \in \mathbb{R}$ is one-dimensional. The loss function is $L(w) = \frac{1}{2}kw^2$ with $k = 1$. In Figure 2(a), we plot the variance of w after 1000 training steps from 10^4 independent runs. We compare the

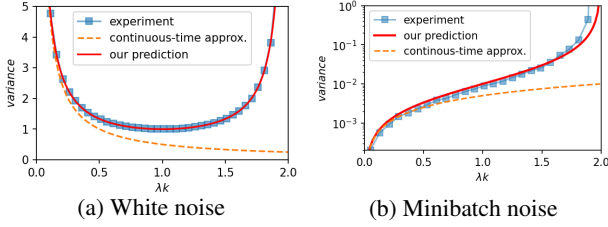


Figure 2: Comparison for a 1d example for predicting Σ . We see that the continuous prediction is only acceptable for $\lambda k < 1$, while the discrete theory holds well for all allowed λ . (a) White noise with strength $\sigma^2 = 1$. (b) Minibatch noise.

prediction of Corollary 1 with that of the continuous-time approximation in Mandt et al. (2017). We see that the proposed theory agrees excellently with the experiment, whereas the standard continuous-time approximation fails as λ increases. Moreover, the continuous-time approximation fails to predict the divergence of the variance at $\lambda k \rightarrow 2$, whereas the discrete theory captures this very well. Now, we consider a multidimensional case. We set the loss function to be $L(\mathbf{w}) = \frac{1}{2} \mathbf{w}^T K \mathbf{w}$. For visualization, we choose $D = 2$, and we set the eigenvalues of the Hessian matrix to be 1 and 0.1, and plot the fluctuation along the eigenvectors of this Hessian matrix. See Figures 3(a)-(c). As before, we compare the discrete-time results with the theoretical predictions of Mandt et al. (2017). After diagonalization, the continuous-time dynamics predicts $\Sigma = \text{diag}(\lambda/2, \lambda/0.2)$, whereas the discrete theory predicts $\Sigma = \text{diag}\left(\frac{\lambda}{2-\lambda}, \frac{\lambda}{0.1(2-0.1\lambda)}\right)$. The proposed theory exhibits no noticeable deviation from the experiment and successfully predicts a distortion along the direction with a large eigenvalue in the Hessian. In comparison, the continuous-time approximation always underestimates the fluctuation in the learning, and the discrepancy is larger as the learning rate gets larger; the prediction of the continuous-time theory can deviate arbitrarily far from the experiment as λ gets close to the divergence value.

Minibatch Noise. For minibatch noise, we solve a linear regression task with the loss function $L(w) = \frac{1}{N} \sum_i^N (\mathbf{w}^T x_i - y_i)^2$, where $N = 1000$ is the number of data points; for the 1d case, the data points x_i are sampled independently from a normal distribution $\mathcal{N}(0, 1)$; $y_i = \mathbf{w}^* x_i + \epsilon_i$ with a constant but fixed \mathbf{w}^* , ϵ_i are noises, also sampled from a normal distribution. For sampling of minibatches, we set the batch size $S = 100$. The theoretical noise covariance matrix is approximated by $C \approx K/S$. See Figure 2(b) for the 1d comparison. We see that the proposed theory agrees much better with the experiment than the continuous theory, both in trend and in magnitude. We also compare the predicted distribution for $D = 2$. Here, the data points x_i are sampled from $\mathcal{N}(0, \text{diag}(1, 0.1))$, which makes the expected Hessian equal to $\text{diag}(1, 0.1)$. See Figures 3(d)-(f) for the illustration. Again, an overall agreement with the experi-

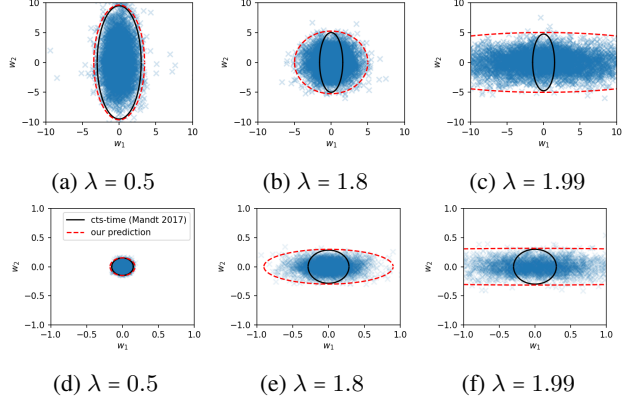


Figure 3: Comparison between our prediction and that of Mandt et al. (2017) at 3σ confidence interval, i.e., 99% of data points lie within the boundary of the theory. We see that our prediction agrees well with the experiments across all levels of the learning rate, whereas the prediction by Mandt et al. (2017) applies only at small λ . (a)-(c) White noise. (d)-(f) Minibatch noise.

ment is much better for the proposed theory. We notice that the prediction by Mandt et al. (2017) and the discrete theory agree well with each other when λ is small. Interestingly, the proposed theory slightly overestimates the variance of the parameters. This suggests the limitation of the commonly used approximation, $C \sim K$, of the minibatch noise. It is possible to treat the minibatch noise in discrete-time regime rigorously in the proposed framework (Ziyyin et al., 2021; Mori et al., 2021).

SGD with Momentum. For white noise, we set $L(w) = kw^2/2$ as before. In Figure 4(a), we plot the case with $\lambda k > 2$, where the dynamics will diverge if no momentum is present. The experiment does show this divergence at the value of $\mu \rightarrow \lambda k/2 - 1$ implied by the necessary stability condition, in contrast to the continuous-time theory that predicts no divergence. In Figure 4(b), we show the experiments with minibatch noise with the same λ . The loss is the same as that of the minibatch noise. The predicted theory agrees better than the continuous-time approximation. On the other hand, the agreement becomes worse as the fluctuation in w becomes large. This suggests the limitation of the commonly used approximation of minibatch noise, i.e., $C \sim H(w) = K$. More experimental results with a smaller learning rate are included in Appendix D.

6. Applications

In this section, we apply our exact solution to some important problems studied in the recent literature, such as the derivation of the Bayesian optimal learning rate in section 6.1, and escape from a sharp minimum³. Previously,

³In the existing machine learning literature, there are two definitions for the escape rate. The first was introduced by Zhu et al. (2019). We call it the ‘‘escape rate of the first kind’’ (in physics,

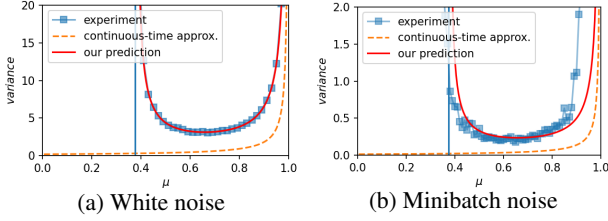


Figure 4: Comparison between the continuous-time theory and the discrete-time theory with momentum. $\lambda k = 2.75$ for both experiments. (a) White noise. The vertical line shows $\mu = \frac{\lambda k - 2}{2}$, where the variance is predicted to diverge. (b) Minibatch noise.

these problems have been solved with the continuous-time approximation in a quadratic loss. It is therefore of interest to investigate how the exact solution corrects the established results in the large- λ regime. The detailed derivations are included in Appendix E.2.

6.1. Optimal Bayesian Learning Rate

We follow the setting of Mandt et al. (2017) for analyzing SGD as an approximate Bayesian inference algorithm. We assume a probabilistic model $p(\mathbf{w}, \mathbf{x})$ with N -dimensional data \mathbf{x} . Our goal is to approximate the posterior

$$p(\mathbf{w}|\mathbf{x}) = \exp[\ln p(\mathbf{w}, \mathbf{x}) - \ln p(\mathbf{x})]. \quad (14)$$

The loss function is defined as $L(\mathbf{w}) := \frac{1}{N} \sum_{n=1}^N l_n(\mathbf{w})$, where $l_n(\mathbf{w}) := -\ln p(x_n|\mathbf{w}) - \frac{1}{N} \ln p(\mathbf{w})$. The posterior is approximately Gaussian:

$$p(\mathbf{w}) \propto \exp\left\{-\frac{N}{2} \mathbf{w}^T K \mathbf{w}\right\}. \quad (15)$$

Theorem 6. *If the noise covariance is $C = \frac{N-S}{NS} K$, the optimal learning rate for minimizing the KL divergence $D_{\text{KL}}(q||f)$ between the SGD stationary distribution $q(\mathbf{w})$ in Theorem 1 and the posterior (15) is the solution to*

$$\frac{N-2S}{S} \sum_{i=1}^D \frac{k_i}{2-\lambda k_i} + \frac{N-S}{S} \lambda \sum_{i=1}^D \frac{k_i^2}{(2-\lambda k_i)^2} = \frac{D}{\lambda}, \quad (16)$$

where k_i is the i -th eigenvalue of the Hessian K .

This relation constitutes a general solution to this problem, which can be solved numerically. The result by Mandt et al. (2017) can be seen as a solution to the above equation after ignoring non-linear terms in λ , which gives $\lambda_c^* = 2 \frac{S}{N} \frac{D}{\text{Tr}[K]}$ under the assumptions that $S \ll N$ and $\lambda k \ll 1$. With increasing λk , one requires increasingly higher-order corrections from Eq. (16).

6.2. Escape Rate of the First Kind

Now, we investigate the effect of discreteness on the escape rate from a sharp minimum. The first indicator for the

this rate is more properly called the ‘‘thermalization rate’’ rather than the ‘‘escape rate’’. The second type is the familiar ‘‘Kramers escape rate’’ (Kramers, 1940).

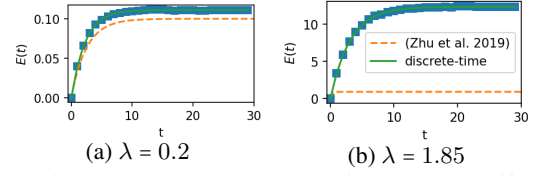


Figure 5: Theoretical prediction of the escape efficiency $E(t)$ vs. experiment. The continuous-time prediction follows from Zhu et al. (2019). The discrete-time prediction agrees very well with the experiment, whereas the continuous-time prediction shows significant deviations as λ increases and completely fails at λ is close to 2.

escape rate, called the escaping efficiency, is proposed by Zhu et al. (2019) as

$$E(t) := \mathbb{E}[L(\mathbf{w}_t) - L(\mathbf{w}_0)], \quad (17)$$

where \mathbf{w}_0 is the exact minimum and t is a fixed time. This indicator qualitatively characterizes the ability of escape from the minimum \mathbf{w}_0 . It is related to the escape probability via the Markov inequality $P(L(\mathbf{w}_t) - L(\mathbf{w}_0) \geq \delta) \leq \frac{E}{\delta}$, for $\delta > 0$, if $\delta = \Delta L$ is the height of the potential barrier.

Theorem 7. *(Escaping efficiency from a sharp minimum) Let the algorithm be updated according to Eq. (5) with $\mu = 0$. Then the escaping efficiency is*

$$E_d = \frac{\lambda}{4} \text{Tr} \left[\left(I_D - \frac{\lambda K}{2} \right)^{-1} [I_D - (I_D - \lambda K)^{2t}] C \right]. \quad (18)$$

The subscript d indicates discrete-time. In comparison, the escaping efficiency calculated from the continuous-time approximation is given by Zhu et al. (2019)

$$E_c = \frac{\lambda}{4} \text{Tr} [(I_D - e^{-2\lambda K t}) C], \quad (19)$$

where the subscript c indicates continuous-time. The two results can be easily compared in two limiting cases. First, in the short-time limit, the continuous-time theory predicts $E_c = \frac{t\lambda^2}{2} \text{Tr}[KC]$, which coincides with the single-step $t = 1$ result given by the discrete theory. Second, when $t \gg 1$, the continuous indicator becomes Hessian-independent: $E_c = \frac{\lambda}{4} \text{Tr}[C]$, whereas the discrete result still depends on the curvature: $E_d = \frac{\lambda}{2} \text{Tr} [(2I_D - \lambda K)^{-1} C]$. The conclusion that a flatter minimum relates to a smaller escaping efficiency still holds. If we take the small- λ limit, it recovers the trivial continuous result. In Figure 5, we compare the prediction of Eq. (18) with the continuous theory. We set $C = I_D$, $K = I_D$ and compare at two different levels of learning rate. The result is averaged over 50000 runs. We see that our solution agrees with the experiment perfectly, while the continuous theory significantly underestimates the escape rate and fails at a large learning rate.

The following corollary shows that the discrete theory predicts a higher escape probability than the continuous one.

Corollary 4. $\forall 0 < \lambda < 2/k^*$ and $t \geq 0$, $E_d \geq E_c$.

We then investigate the effect of anisotropic noise on the escape efficiency as in [Zhu et al. \(2019\)](#). We consider different structure of noise with the same magnitude $\text{Tr}[C]$. We define an *ill-conditioned* Hessian K as its descending ordered eigenvalues $k_1 \geq k_2 \cdots \geq k_D > 0$ satisfy $k_{l+1}, k_{l+2}, \dots, k_D < k_1 D^{-d}$ for some constant $l \ll D$ and $d > 1/2$. We assume that C is aligned with K . Let u_i be the corresponding unit eigenvector of eigenvalue k_i . For some coefficient $a > 0$, we have $u_1^T C u_1 \geq a k_1 \text{Tr}[C] / \text{Tr}[K]$. This is true if the maximal eigenvalues of C and K are aligned in proportion, namely $c_1 / \text{Tr}[C] \geq a_1 k_1 / \text{Tr}[K]$.

Theorem 8. *Under the conditions of an ill-conditioned Hessian and a noise covariance that aligns with the Hessian, the ratio between the escaping efficiencies of an anisotropic noise and its isotropic version $\bar{C} := \frac{\text{Tr}[C]}{D} I_D$ is*

$$\frac{\text{Tr}[KC]}{\text{Tr}[K\bar{C}]} = \mathcal{O}(aD^{2d-1}). \quad (20)$$

Remark. *This result shows that the previous understanding that the anisotropy in noise may help escape from a sharp minimum still holds in a discrete-time regime. Therefore, the qualitative result in [Zhu et al. \(2019\)](#) still holds when a large learning rate is used.*

6.3. Escape from Sharp Minima (Kramers Problem)

In physics, the Kramers escape problem ([Kramers, 1940](#)) concerns the approximate mean time for a particle confined in a local minimum of a potential $L(w)$ to escape across the potential barrier. For continuous-time dynamics, the standard approach to calculating this Kramers rate (or time) ([Hänggi et al., 1990](#); [Van Kampen, 1992](#)) is to employ the Fokker-Planck equation for the distribution $P_c(w, t)$ (c for continuous-time)

$$\frac{\partial P_c(w, t)}{\partial t} = -\nabla \cdot J(w, t), \quad (21)$$

where the probability current is defined as $J(w, t) := -\lambda P_c(w, t) \nabla L(w) - \mathcal{D} \nabla P_c(w, t)$. Here $\mathcal{D} := \frac{1}{2} C$ is the diffusion matrix and T is the effective ‘‘temperature’’. The mean escape rate is defined as

$$\gamma := \frac{P(w \in V_a)}{\int_{\partial a} J \cdot dS}, \quad (22)$$

where $P(w \in V_a)$ is the probability of a particle being inside the well a , and $\int_{\partial a} J \cdot dS$ is the probability flux through the boundary of the well a . For illustration, see [Figure 6\(a\)](#). In continuous theory, the probability inside the well a can be approximated by 1 because the distribution lies almost within the well. However, the discrete theory predicts larger fluctuations such that the distribution spreads out with large λ . By making proper approximations, we improve the result on the Kramers rate for the discrete SGD.

Theorem 9. *Let k_a and k_b be the local Hessian at the local minimum a and the barrier top b , respectively. Suppose l is a midpoint on the most probable escape path between a and*

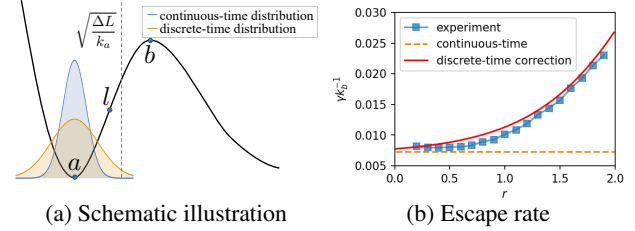


Figure 6: (a) Schematic illustration of the Kramers problem. The stationary distribution predicted by the discrete-time theory at minimum a is represented by the orange area while the blue area represents the continuous theory prediction. Because the discrete-time distribution spreads out of the valley, we approximate half of the width of the valley a by $\sqrt{\frac{\Delta L}{k_a}}$ as indicated by the dashed vertical line. (b) $\frac{\gamma}{|k_b|}$ versus r . The blue squares are taken from experiments. The orange dashed line shows the prediction of the continuous theory. It agrees with the experiments only for vanishingly small r . The red solid curve is our theoretical prediction multiplied by an empirical constant. We see that our prediction is consistent with the experimental data up to a constant coefficient, while the continuous theory underestimates the escape rate.

b such that $k(w) \approx k_a$ in the path $a \rightarrow l$ and $k(w) \approx k_b$ in $l \rightarrow b$. The approximate Kramers escape rate from a local minimum a for the discrete-time SGD is

$$\gamma \approx \frac{1}{2\pi} |k_b| \sqrt{\frac{2}{2 - \lambda k_a}} \text{erf} \left(\sqrt{\frac{S(2 - \lambda k_a) \Delta L}{\lambda k_a}} \right) \times \exp \left[-\frac{2S\Delta L}{\lambda} \left(\frac{l(1 - \lambda k_a/2)}{k_a} + \frac{1-l}{|k_b|} \right) \right], \quad (23)$$

where $\text{erf}(z)$ is the error function.

To compare, the mean escape rate obtained from the continuous-time theory ([Xie et al., 2021](#)) is

$$\gamma_c = \frac{1}{2\pi} |k_b| \exp \left[-\frac{2S\Delta L}{\lambda} \left(\frac{l}{k_a} + \frac{1-l}{|k_b|} \right) \right]. \quad (24)$$

In [Figure 6\(b\)](#) we plot the quantity $\frac{\gamma}{|k_b|}$ while rescaling the loss function by a factor r : $L \rightarrow rL$. The continuous theory predicts a constant escape rate by varying r , whereas the discrete theory expects a monotonic increase as r becomes large. Such monotonic increase is indeed observed in experiments. The theoretical curve is rescaled by a constant to make comparison easier. Our prediction is qualitatively consistent with the experiment, whereas the continuous theory is only valid in a rather limited range of small r . One surprising result here is that, the escape rate of continuous-time dynamics is invariant to the multiplication of the loss function by a constant r , while this is not true for discrete-time dynamics, where larger r leads to a larger escape rate from a sharp minimum.

6.4. Second-Order Methods

In this subsection, we show that the proposed theory can also be extended to analyze the stochastic versions of second-order methods. Here, we consider the stochastic version of the damped Newton’s method (DNM) (Nesterov et al., 2018) where $\Lambda = \lambda K^{-1}$ and the natural gradient descent (NGD) algorithm in which the learning rate matrix is defined as $\Lambda := \lambda J(\mathbf{w})^{-1}$ with $J(\mathbf{w}) := \mathbb{E}[\nabla L(\nabla L)^T]$ being the Fisher information. The NGD algorithm is known to be *Fisher-efficient* close to a local minimum (Amari, 1998; Amari and Nagaoka, 2007), namely, it is the fastest possible method to estimate a given statistical quantity, because the Fisher information is the “natural” metric in the probability space. The following corollaries give the model fluctuations of DNM and NGD under minibatch noise. A detailed discussion is given in Appendix E.2.7.

Corollary 5. *Let the DNM algorithm be updated with noise covariance being $C = \frac{N-S}{NS}K$. The model fluctuation is*

$$\Sigma = \frac{1+\mu}{1-\mu} \frac{\lambda}{2(1+\mu)-\lambda} \frac{N-S}{NS} K^{-1}. \quad (25)$$

Corollary 6. *Let the NGD algorithm be updated with noise covariance being $C = \frac{N-S}{NS}K\Sigma K$. Then,*

$$\Sigma = \lambda \frac{(1+\mu) \frac{N-S}{NS} + 1 - \mu}{2(1-\mu^2)} K^{-1}. \quad (26)$$

Remark. *For NGD, if $S \rightarrow N$, then $C \rightarrow 0$. In this situation, we have $\Sigma = \frac{\lambda}{2(1+\mu)} K^{-1}$. This means that the algorithm involves nonzero fluctuations even if the noise is absent! Moreover, the divergence caused by $\lambda k^* \rightarrow 2$ disappears for NGD. This suggests that, when the noise is due to minibatch sampling, NGD naturally corrects the ill-conditioned problem of discrete-time SGD.*

Remark. *Even more interestingly, both the DNM and the NGD algorithms induce fluctuations that are the same as those in the continuous-time SGD algorithm with Gaussian noise up to the coefficients, in the sense that the variance is proportional to K^{-1} which is the local geometry of the minimum. Intuitively, this means that DNM and NGD need no correction even in the discrete-time case; moreover, they are likely to generalize better because they better align with the underlying loss function.*

It would also be interesting to consider the stationary fluctuation of the adaptive gradient methods such as Adam (Kingma and Ba, 2014). Adam can be seen as an approximate second-order method with a diagonal preconditioning matrix $\hat{\Lambda} = \lambda/\sqrt{\mathbb{E}[\text{diag}(\mathbf{g}^2)]}$, where \mathbf{g}^2 denotes element-wise square. This can be approximated by a non-diagonal preconditioning matrix $\Lambda = \lambda/\sqrt{\mathbb{E}[\mathbf{g}\mathbf{g}^T]}$. Using $\mathbb{E}[\mathbf{g}\mathbf{g}^T]$. Such approximation is in fact not bad; as shown in Zhiyi and Ziyin (2021), the diagonal assumption of $\mathbb{E}[\mathbf{g}\mathbf{g}^T]$ leads to

surprisingly accurate predictions of the statistical properties of Adam for various modern neural architectures. Under this non-diagonal setting, we can solve the stationary covariance of Adam (without momentum).

Theorem 10. *(Stationary fluctuation of Adam) Let the preconditioning matrix of Adam be $\Lambda = \lambda/\sqrt{\mathbb{E}[\mathbf{g}\mathbf{g}^T]}$ and the noise be $C = cK\Sigma K$ with a positive c , then*

$$\Sigma_{\text{Adam}} = \frac{\lambda^2(1+c)}{4} I_D. \quad (27)$$

Remark. *Similar to SGD, Σ_{Adam} is isotropic. In general, we have $C \sim 1/S$. Two key differences exist: (1) that Σ_{SGD} vanishes when $S \gg 1$, while that of Adam converge to a finite constant, like NGD; (2) Σ_{Adam} scales as λ^2 , while that of SGD scales as λ , and so Adam has much smaller fluctuation than SGD.*

7. Concluding Remark

In this work, we have analyzed the SGD algorithm in a quadratic potential, which is a good approximation close to any local minimum and a common setting in the recent literature. Compared to the related works, our solution is exact, and relies on fewer assumptions than previous works, and, with the exact solutions, corrections to the previous results that were based on continuous-time approximation are obtained. In fact, we showed that even in the simplest settings, the prediction of the continuous-time solution may deviate significantly and eventually fails for a large learning rate. This suggests the fundamental limitation of making the continuous-time approximation in analyzing machine learning problems. Previous works have shown that, SGD leads to a flatter minimum due to the existence of anisotropic noise; this anisotropy is enhanced when the dynamics is discrete-time; this gives stronger mobility to model parameters along the sharper directions in the Hessian, and therefore, makes a flatter minimum more favorable. On the other hand, the distortion that the discrete-time SGD causes, in general, does not match the underlying landscape, indicating that using a large learning rate may cause a larger approximation error and worsened generalization. This tradeoff has been discussed only in a restricted setting in this work, and we hope the discussions here may stimulate further research in this direction.

Acknowledgment

We acknowledge Prof. Chikara Furusawa, Zeke Xie and Takashi Mori for valuable discussions. This work was supported by KAKENHI Grant No. JP18H01145 and a Grant-in-Aid for Scientific Research on Innovative Areas “Topological Materials Science (KAKENHI Grant No. JP15H05855) from the Japan Society for the Promotion of Science.

References

- Amari, S. and Nagaoka, H. (2007). *Methods of Information Geometry*. Translations of mathematical monographs. American Mathematical Society.
- Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276.
- Bottou, L. (1998). Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):142.
- Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. (2019). Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018.
- Chaudhari, P. and Soatto, S. (2018). Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–10. IEEE.
- Dieuleveut, A., Durmus, A., Bach, F., et al. (2020). Bridging the gap between constant step size stochastic gradient descent and markov chains. *Annals of Statistics*, 48(3):1348–1382.
- Duchi, J., Hazan, E., and Singer, Y. (2011). Adaptive sub-gradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7).
- Dziugaite, G. K. and Roy, D. M. (2017). Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*.
- Einstein, A. (1905). Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Annalen der physik*, 4.
- Flammarion, N. and Bach, F. (2015). From averaging to acceleration, there is only a step-size. In *Conference on Learning Theory*, pages 658–695. PMLR.
- Gitman, I., Lang, H., Zhang, P., and Xiao, L. (2019). Understanding the role of momentum in stochastic gradient methods. In *Advances in Neural Information Processing Systems*, pages 9633–9643.
- Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., and He, K. (2017). Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*.
- Hänggi, P., Talkner, P., and Borkovec, M. (1990). Reaction-rate theory: fifty years after kramers. *Reviews of modern physics*, 62(2):251.
- Higham, N. J. and Kim, H.-M. (2001). Solving a quadratic matrix equation by newton’s method with exact line searches. *SIAM Journal on Matrix Analysis and Applications*, 23(2):303–316.
- Hochreiter, S. and Schmidhuber, J. (1997). Flat minima. *Neural computation*, 9(1):1–42.
- Hodgkinson, L. and Mahoney, M. W. (2020). Multiplicative noise and heavy tails in stochastic optimization. *arXiv preprint arXiv:2006.06293*.
- Hoffer, E., Hubara, I., and Soudry, D. (2017). Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1729–1739.
- Hu, W., Li, C. J., Li, L., and Liu, J.-G. (2019). On the diffusion approximation of nonconvex stochastic gradient descent. *Annals of Mathematical Sciences and Applications*, 4(1).
- Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: convergence and generalization in neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 8580–8589.
- Jastrzebski, S., Kenton, Z., Arpit, D., Ballas, N., Fischer, A., Bengio, Y., and Storkey, A. (2017). Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*.
- Jiang, Y., Neyshabur, B., Mobahi, H., Krishnan, D., and Bengio, S. (2019). Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*.
- Keskar, N. S., Nocedal, J., Tang, P. T. P., Mudigere, D., and Smelyanskiy, M. (2017). On large-batch training for deep learning: Generalization gap and sharp minima. In *5th International Conference on Learning Representations, ICLR 2017*.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- Kramers, H. A. (1940). Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304.
- Landau, L. D. and Lifshitz, E. M. (1980). *Statistical physics, part 1*. Elsevier, third edition edition.

- LeCun, Y. A., Bottou, L., Orr, G. B., and Müller, K.-R. (2012). Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer.
- Lee, J., Xiao, L., Schoenholz, S. S., Bahri, Y., Novak, R., Sohl-Dickstein, J., and Pennington, J. (2020). Wide neural networks of any depth evolve as linear models under gradient descent. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12):124002.
- Lewkowycz, A., Bahri, Y., Dyer, E., Sohl-Dickstein, J., and Gur-Ari, G. (2020). The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint arXiv:2003.02218*.
- Li, Q., Tai, C., and Weinan, E. (2017). Stochastic modified equations and adaptive stochastic gradient algorithms. In *International Conference on Machine Learning*, pages 2101–2110. PMLR.
- Li, Y., Wei, C., and Ma, T. (2019). Towards explaining the regularization effect of initial large learning rate in training neural networks. *arXiv preprint arXiv:1907.04595*.
- Liang, T., Poggio, T., Rakhlin, A., and Stokes, J. (2019). Fisher-rao metric, geometry, and complexity of neural networks. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 888–896. PMLR.
- Ma, J. and Yarats, D. (2018). Quasi-hyperbolic momentum and adam for deep learning. In *International Conference on Learning Representations*.
- Ma, Y.-A., Chen, T., and Fox, E. (2015). A complete recipe for stochastic gradient mcmc. *Advances in Neural Information Processing Systems*, 28:2917–2925.
- MacKay, D. J. (1992). A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472.
- Mandt, S., Hoffman, M. D., and Blei, D. M. (2017). Stochastic gradient descent as approximate bayesian inference. *Journal of Machine Learning Research*, 18:1–35.
- Meng, Q., Gong, S., Chen, W., Ma, Z.-M., and Liu, T.-Y. (2020). Dynamic of stochastic gradient descent with state-dependent noise. *arXiv preprint arXiv:2006.13719*.
- Mori, T. and Ueda, M. (2020a). Improved generalization by noise enhancement. *arXiv preprint arXiv:2009.13094*.
- Mori, T. and Ueda, M. (2020b). Is deeper better? it depends on locality of relevant features. *arXiv preprint arXiv:2005.12488*.
- Mori, T., Ziyin, L., Liu, K., and Ueda, M. (2021). Logarithmic landscape and power-law escape rate of sgd. *arXiv preprint arXiv:2105.09557*.
- Nagarajan, V. and Kolter, J. Z. (2019). Generalization in deep networks: The role of distance from initialization. *arXiv preprint arXiv:1901.01672*.
- Nesterov, Y. et al. (2018). *Lectures on convex optimization*, volume 137. Springer.
- Neyshabur, B., Bhojanapalli, S., and Srebro, N. (2018a). A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. In *International Conference on Learning Representations*.
- Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y., and Srebro, N. (2018b). The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*.
- Pascanu, R. and Bengio, Y. (2013). Revisiting natural gradient for deep networks. *arXiv preprint arXiv:1301.3584*.
- Rao, C. R. (1992). Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in statistics*, pages 235–247. Springer.
- Rissanen, J. (1983). A universal prior for integers and estimation by minimum description length. *The Annals of statistics*, pages 416–431.
- Simsekli, U., Sagun, L., and Gurbuzbalaban, M. (2019). A tail-index analysis of stochastic gradient noise in deep neural networks. In *International Conference on Machine Learning*, pages 5827–5837. PMLR.
- Simsekli, U., Sener, O., Deligiannidis, G., and Erdogdu, M. A. (2020). Hausdorff dimension, heavy tails, and generalization in neural networks. *Advances in Neural Information Processing Systems*, 33.
- Smith, S. L. and Le, Q. V. (2018). A bayesian perspective on generalization and stochastic gradient descent. In *International Conference on Learning Representations*.
- Sutskever, I., Martens, J., Dahl, G., and Hinton, G. (2013). On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147. PMLR.
- Van Kampen, N. G. (1992). *Stochastic processes in physics and chemistry*, volume 1. Elsevier.
- Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML’11*, page 681–688, Madison, WI, USA. Omnipress.
- Wu, L., Ma, C., and E, W. (2018). How sgd selects the global minima in over-parameterized learning: A dynamical stability perspective. In *Proceedings of the 32nd*

International Conference on Neural Information Processing Systems, pages 8289–8298.

Xie, Z., Sato, I., and Sugiyama, M. (2021). A diffusion theory for deep learning dynamics: Stochastic gradient descent exponentially favors flat minima. In *International Conference on Learning Representations*.

Zhang, C., Liao, Q., Rakhlin, A., Miranda, B., Golowich, N., and Poggio, T. (2018). Theory of deep learning iib: Optimization properties of sgd. *arXiv preprint arXiv:1801.02254*.

Zhiyi, Z. and Ziyin, L. (2021). On the distributional properties of adaptive gradients.

Zhu, Z., Wu, J., Yu, B., Wu, L., and Ma, J. (2019). The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. In *International Conference on Machine Learning*, pages 7654–7663. PMLR.

Ziyin, L., Liu, K., Mori, T., and Ueda, M. (2021). On minibatch noise: Discrete-time sgd, overparametrization, and bayes. *arXiv preprint arXiv:2102.05375*.