
Time-Independent Information-Theoretic Generalization Bounds for SGLD

Futoshi Futami*

Osaka University / RIKEN AIP

futami.futoshi.es@osaka-u.ac.jp

Masahiro Fujisawa*†

RIKEN AIP

masahiro.fujisawa@riken.jp

Abstract

We provide novel information-theoretic generalization bounds for stochastic gradient Langevin dynamics (SGLD) under the assumptions of smoothness and dissipativity, which are widely used in sampling and non-convex optimization studies. Our bounds are time-independent and decay to zero as the sample size increases, regardless of the number of iterations and whether the step size is fixed. Unlike previous studies, we derive the generalization error bounds by focusing on the time evolution of the Kullback–Leibler divergence, which is related to the stability of datasets and is the upper bound of the mutual information between output parameters and an input dataset. Additionally, we establish the first information-theoretic generalization bound when the training and test loss are the same by showing that a loss function of SGLD is sub-exponential. This bound is also time-independent and removes the problematic step size dependence in existing work, leading to an improved excess risk bound by combining our analysis with the existing non-convex optimization error bounds.

1 Introduction

Stochastic optimization, including stochastic gradient descent (SGD), is central to realizing practical large-scale or deep-learning models. There are currently considerable active discussions on accurately determining the generalization performance of models trained by SGD or its variants. In particular, stochastic gradient Langevin dynamics (SGLD) [13, 39, 29], a noisy variant of SGD, has garnered much attention in this type of study since it provides a useful theoretical framework for generalization error analysis based on the Langevin diffusion context [29]. Our study aims to contribute to a more accurate understanding and evaluation of the generalization performance for SGLD.

There are two main approaches to generalization analysis in SGLD. One is the *information-theoretic analysis* proposed by Russo and Zou [30] and Xu and Raginsky [40], by which a generalization error bound is derived using the mutual information (MI) between the learned parameters and the training dataset. Recently, some extensions using gradient information have been made to investigate the generalization properties of SGLD, for example, upper-bounding the MI with the norm of gradients [28] and the sum of gradient variances [25, 35, 37, 36]. Information-theoretic generalization bounds are applicable to a wide range of noisy iterative algorithms such as differentially private SGD [11] and stochastic gradient Hamiltonian Monte Carlo [5] modified to include a noisy momentum.

The other approach is *stability analysis*, by which the effects of changes in the learning algorithm due to the addition or removal of a single training data point on the generalization performance are investigated. Raginsky et al. [29] derived the non-asymptotic generalization and excess risk bound of SGLD via the exponential ergodicity of Langevin diffusion. Starting with the study by Raginsky

*Equal contribution.

†Corresponding author.

et al. [29], there have been many attempts to improve the generalization analysis in SGLD from the stability perspective, such as those by Zhang et al. [43], Mou et al. [23] and Li et al. [20].

Unfortunately, these existing generalization bounds are *time-dependent*; namely, they diverge with increasing number of iterations unless the step size is adjusted so that the order of bound values is $\mathcal{O}(n^{-1})$ or $\mathcal{O}(n^{-1/2})$, where n is the sample size (see Section 5 for details). Farghly and Rebeschini [10] attempted to avoid this problem by Wasserstein stability analysis through reflection coupling [9] under the smoothness and dissipativity [14] assumptions commonly used in sampling and non-convex optimization communities [29, 44]. Although their bounds bypass the divergence problem when taking a supremum over time, the geometry introduced for the reflection coupling yields an unnatural dependence on step size, resulting in a vacuous bound as the step size decreases (see Table 1).

In this paper, we provide novel generalization bounds for SGLD under smooth and dissipative loss functions obtained by the information-theoretic approach. We focus on the upper bound of the MI, namely, the Kullback–Leibler (KL) divergence between the distributions of parameters learned from different training datasets. We then analyze its time evolution caused by the update of the SGLD algorithm through the Fokker–Planck (FP) equation (Lemma 1). On the basis of this analysis, we obtain time-independent generalization error bounds that decay to zero as $n \rightarrow \infty$ regardless of the number of iterations or whether the step size is fixed (Theorem 4 and Corollary 1). Conventional information-theoretic generalization bounds [29, 28, 35, 36] are derived by bounding the MI between the parameters at all iterations and the training dataset. Therefore, these bounds grow linearly with the number of iterations, resulting in a time-dependent generalization error bound. Our analysis based on time evolution eliminates both this linearity issue and the unnatural dependence on step size (the inverse of step size) in the time-independent bound of Farghly and Rebeschini [10].

Another contribution is providing the first information-theoretic generalization bound and the excess risk bound when the same loss is used for training and the generalization performance evaluation. In the conventional information-theoretic approach, deriving generalization error bounds under this setting was challenging owing to the unknown tail behavior of a loss function of SGLD. We overcome this difficulty with our discovery that a smooth and dissipative loss function of SGLD is sub-exponential.

2 Preliminaries

2.1 Problem settings and stochastic gradient Langevin dynamics

We represent random variables in capital letters, such as X , and deterministic values in lowercase letters, such as x , and express the Euclidean inner product and distance as \cdot and $\|\cdot\|$. Let μ be an unknown generating distribution on the instance space \mathcal{Z} and $w \in \mathcal{W} \subseteq \mathbb{R}^d$ be the d -dimensional parameters such as weights of neural networks, where \mathcal{W} is the space of the parameters. We consider a loss function $l : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$ and the following optimization problem:

$$\min_{w \in \mathcal{W}} L_\mu(w) := \mathbb{E}_Z[l(w, Z)] = \int_{\mathcal{Z}} l(w, Z) d\mu(z),$$

which cannot be computed since μ is unknown. Instead, we typically minimize the empirical risk estimated using the dataset $S := \{Z_i\}_{i=1}^n$:

$$\min_{w \in \mathcal{W}} L_S(w) := \frac{1}{n} \sum_{i=1}^n l(w, Z_i),$$

where $\{Z_i\}_{i=1}^n$ are independent and identically distributed (i.i.d.) samples from μ , i.e., $Z_i \stackrel{\text{i.i.d.}}{\sim} \mu$.

Stochastic gradient Langevin dynamics. In this paper, we use the SGLD algorithm [39] to solve the empirical risk minimization. SGLD utilizes the gradient information of the loss function; however, some loss functions, such as the 0-1 loss, are not differentiable. In this case, it is common to use the differentiable surrogate loss function $f : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$ (e.g., the cross-entropy loss) and minimize the following empirical risk: $F_S(w) := \frac{1}{n} \sum_{i=1}^n f(w, Z_i)$. Given a mini-batch $B \subset [n] := \{1, \dots, n\}$ with $k = |B| \leq n$, we define its mini-batch version as

$$F(w, B) := \frac{1}{k} \sum_{i \in B} f(w, Z_i).$$

The SGLD algorithm updates the parameters using the following recursion:

$$W_{t+1} = W_t - \eta_t \nabla F(W_t, B_t) + \sqrt{2\beta_t^{-1}\eta_t} \xi_t, \quad W_0 \sim P_{W_0},$$

where P_{W_0} is a given initial distribution, $\nabla F(w, B)$ is a stochastic gradient, t is the number of iterations, η_t is the step size, β_t is the inverse temperature, and $(B_t)_{t=0}^\infty$ is an i.i.d. sequence of random variables distributed uniformly on $\{B \subset [n] : |B| = k\}$. In addition, $(\xi_t)_{t=0}^\infty$ is an i.i.d. sequence of standard Gaussian random variables, i.e., $\xi_t \sim \mathcal{N}(0, \mathbf{I}_d)$, where \mathbf{I}_d is the d -dimensional identity matrix. The output parameters $W \in \mathcal{W}$ obtained using SGLD can be seen as the samples from a conditional distribution $P_{W|S} : \mathcal{Z}^n \rightarrow \mathcal{W}$. We express the t -th output of the SGLD algorithm as W_t .

2.2 Expected generalization error and its bounds

The focus of this paper is the *expected generalization error*, defined as

$$\text{gen}(\mu, P_{W|S}; L) := \mathbb{E}_{S, W} [L_\mu(W) - L_S(W)], \quad (1)$$

where the expectation is taken over the joint distribution of (S, W) , i.e., $\mu^n \otimes P_{W|S}$.

Information-theoretic generalization bounds. Russo and Zou [30] and Xu and Raginsky [40] have shown that Eq. (1) can be bounded by the MI between the input dataset S and the output parameters W under the following sub-Gaussian assumption.

Assumption 1 (sub-Gaussian losses). *A loss function $l(w, Z)$ is sub-Gaussian under $Z \sim \mu$ for all $w \in \mathcal{W}$, that is, there is a positive constant σ_g^2 such that $\log \mathbb{E}_Z [\exp(\lambda(l(w, Z) - \mathbb{E}l(w, Z)))] \leq \lambda^2 \sigma_g^2 / 2$ for all constant $\lambda \in \mathbb{R}$.*

For example, bounded or Lipschitz-continuous loss functions satisfy this assumption. Assumptions regarding the tail behavior of the loss function distributions as in the above are necessary for the information-theoretic generalization error analysis. Bu et al. [3] have investigated information-theoretic generalization bounds with another tail-behavior assumption such as sub-exponential losses.

We introduce the following standard information-theoretic generalization bound.

Theorem 1 (Russo and Zou [30] and Xu and Raginsky [40]). *Suppose that Assumption 1 holds. Then, we have*

$$|\text{gen}(\mu, P_{W|S}; L)| \leq \sqrt{\frac{2\sigma_g^2}{n} I(W; S)}, \quad (2)$$

under a training dataset $S = \{Z_i\}_{i=1}^n$ and the algorithm's output W , where $I(W; S)$ is the MI between W and S .

In the SGLD context, $I(W, S)$ of Eq. (2) can be upper-bounded in a form that incorporates the gradient variance [28, 25, 35, 36]. Given the output of the T -th iterate of the SGLD algorithm, W_T , the following upper bound can be obtained.

Theorem 2 (Modified bound of Pensia et al. [28]). *Let $f(\cdot, z)$ be an L -Lipschitz continuous function, namely, there is a constant $L > 0$ such that $\|f(w, z) - f(\bar{w}, z)\| \leq L\|w - \bar{w}\|$ holds for all $w, \bar{w} \in \mathcal{W}$ and all $z \in \mathcal{Z}$. Then, we obtain*

$$I(W_T; S) \stackrel{(i)}{\leq} \sum_{t=0}^{T-1} \frac{d}{2} \log \left(1 + \frac{\beta_t \eta_t}{d} \text{Var}[\nabla f(W_t, B_t) | W_t] \right) \leq \sum_{t=0}^{T-1} \frac{d}{2} \log \left(1 + \frac{\beta_t \eta_t L^2}{d} \right), \quad (3)$$

where $\text{Var}[\nabla f(W, B) | W] := \mathbb{E} \mathbb{E}_B [\|\nabla_W f(W, B) - \mathbb{E}_B [\nabla_W f(W, B)]\|^2 | W]$ is the conditional variance.

Note that bound (i) can be obtained by the data-processing inequality [6].

The merit of such information-theoretic generalization bounds is that we can evaluate the bound value using the empirically estimated gradient variance per iteration. However, unfortunately, from Eq. (3), this bound is *time-dependent*; namely, the bound value can diverge unless the gradient variance or $\beta_t \eta_t$ approaches 0 as $T \rightarrow \infty$. This is due to the data-processing inequality when deriving upper bound (i) in Eq. (3). By the data-processing inequality, we obtain $I(W_T; S) \leq I(W^{(T)}; S)$, where

$W^{(T)} := (W_0, W_1, \dots, W_T)$ denotes the joint random variables appearing in all the iterations in the algorithm. Since $W^{(T)}$ is treated simultaneously, the bound is inevitably linear in T .

Another limitation of the information-theoretic approach appears in the setting where training losses (f) are also used for performance evaluation, which is often employed in sampling and non-convex optimization studies of SGLD [29, 41]. In this setting, the generalization error is defined as

$$\text{gen}(\mu, P_{W|S}; F) := \mathbb{E}_{S,W}[F_\mu(W) - F_S(W)], \quad (4)$$

where $F_\mu := \mathbb{E}_Z[f(w, Z)]$. We cannot conduct the information-theoretic analysis for Eq. (4) because the tail behavior of the distribution of the training loss is unclear.

Time-independent generalization bounds for Eq. (4). To solve the above problems, Farghly and Rebeschini [10] provided the generalization error bounds of Eq. (4) from the stability perspective under the following assumptions widely used in the non-convex optimization analysis of SGLD [29, 41, 19].

Assumption 2 (Smoothness). *For each $z \in \mathcal{Z}$, $f(\cdot, z)$ is differentiable and M -smooth. That is, there is a positive constant M for all $w, \bar{w} \in \mathcal{W}$ and all $z \in \mathcal{Z}$ such that*

$$\|\nabla f(w, z) - \nabla f(\bar{w}, z)\| \leq M\|w - \bar{w}\|.$$

Assumption 3 (Dissipativity [14]). *For each $z \in \mathcal{Z}$, $f(\cdot, z)$ is (m, b) -dissipative.¹ That is, there are positive constants m and b for all $w \in \mathcal{W}$ and $z \in \mathcal{Z}$ such that*

$$m\|w\|^2 - b \leq \nabla f(w, z) \cdot w.$$

The discussion regarding loss functions that satisfy Assumption 3 is presented in Appendix B.

Hereafter, we eliminate the time dependence of the step size and temperature by setting $\eta_t = \eta$ and $\beta_t = \beta$. With this notation, Farghly and Rebeschini [10] derived the following generalization bound.

Theorem 3 (Farghly and Rebeschini [10]). *Suppose that Assumptions 2 and 3 hold. Assume that the initial law of W_0 has a finite fourth moment σ . Then, if $\eta \leq 1/2m$, for any $T \in \mathbb{N}$, we have*

$$|\text{gen}(\mu, P_{W_T|S}; F)| < C_1 \left(\eta T \wedge \frac{n(C_2 + 1)}{n - k} \right) \left(\frac{k}{n\eta^{1/2}} + \eta^{1/2} \right), \quad (5)$$

where $(x \wedge y) = \min\{x, y\}$, and C_1 and C_2 are the positive constant terms w.r.t. $\{M, m, b, d, \beta, \sigma\}$ and $\{M, m, b, d, \beta\}$, respectively.

Farghly and Rebeschini [10] utilized the Wasserstein stability on the basis of the contraction property of Langevin diffusion under reflection coupling. The important technique to derive the above bound is that we only focus on W_T differently from $W^{(T)}$ of the information-theoretic approach when deriving the contraction property. In this way, the resulting bounds do not suffer from divergence as $T \rightarrow \infty$; however, it still has a problem. That is, Eq. (5) depends on the factor $\eta^{-1/2}$, which implies that it becomes vacuous or even diverges with decreasing $\eta (= \eta_T)$ as $T \rightarrow \infty$.

In this paper, we propose new generalization bounds to address the drawbacks of the information-theoretic and stability-based approaches. Specifically, the proposed bounds are established on the basis of the two expected generalization errors outlined in Eqs. (1) and (4), which remain time-independent and do not diverge as the step size decreases.

3 Time-independent generalization error bound for SGLD

Here, we explain our time-independent bound of $\text{gen}(\mu, P_{W|S}; L)$ for SGLD. We first introduce the main result (Section 3.1) and then summarize its proof outline (Sections 3.2 and 3.3). Finally, in Section 3.4, we provide a detailed discussion on our bound with concrete examples.

¹ This assumption holds not only for (strongly) convex losses but also for many practically used non-convex loss functions [24]. For example, it applies to non-convex loss functions with l_2 constraints and likelihood functions that satisfy Poincaré inequality [2, 32].

3.1 Main result

Our key idea is to derive the generalized error bound using the FP equation. To use the FP equation, we impose the following regularity condition for P_{W_0} .

Assumption 4 (Regularity of the initial distribution). *The initial distribution of W_0 : P_{W_0} is a Gaussian distribution² with a finite variance $s^2 > 0$, which is independent of η and T .*

Our analysis is also grounded in the time evolution of the FP equation using the logarithmic Sobolev inequality (LSI) [2] associated with π described as follows. We state that π satisfies the LSI with constant c_{LS} , if for any $\rho \ll \pi$, the following relation holds:

$$\text{KL}(\rho|\pi) \leq c_{\text{LS}} \mathbb{E} \|\nabla \log \rho - \nabla \log \pi\|^2.$$

Raginsky et al. [29] showed the existence of c_{LS} under Assumptions 2, 3, and $\beta \geq 2/m$. Note that c_{LS} is expressed by the problem-dependent constant (see Appendix F.1 for details).

We now introduce our generalization error bound.

Theorem 4. *Suppose that Assumptions 1, 2, 3, and 4 are satisfied. Then, for any $\beta \geq 2/m$ and $\eta \in (0, 1 \wedge \frac{m}{5M^2} \wedge 4\beta c_{\text{LS}})$ and any $T \in \mathbb{N}$, we have*

$$|\text{gen}(\mu, P_{W_T|S}; L)| \leq \sqrt{\frac{2c_1\sigma_g^2}{n} \left(1 \wedge \frac{\eta T}{4\beta c_{\text{LS}}}\right) (V_{\nabla} + c_2)}, \quad (6)$$

where c_1 , c_2 , and V_{∇} are the positive constant terms w.r.t. $\{M, m, b, d, \beta, s^2\}$.

The above theorem shows $|\text{gen}(\mu, P_{W_T|S}; L)| = \mathcal{O}(\sqrt{(\eta T \wedge 1)/n})$, which implies time independence since it does not diverge even if $T \rightarrow \infty$ and thus converges as $n \rightarrow \infty$.

In Eq. (6), the term V_{∇} corresponds to *stability*, which is expressed as the upper bound of the difference of the expected conditional gradients with respect to changes in training datasets at each iteration. This shows a certain similarity to existing information-theoretic generalization bounds, such as Theorem 2, expressed by the variance of gradients with respect to the training datasets. This similarity is discussed in detail in Section 3.4. Additionally, detailed information on the explicit expression of c_1 , c_2 , and V_{∇} can be found in Appendix F.

3.2 Proof outline of Theorem 4

In this section, we present how to derive our bound in Theorem 4. Our aim here is to share the ideas behind our analysis and an outline of the proof, providing the detailed proof in Appendix F.

We adopt the information-theoretic approach and focus on the MI in Eq. (2). By using the Jensen inequality, we have the following upper bound of the MI:

$$I(W_T; S) \leq \mathbb{E}_{S, S'} \text{KL}(P_{W_T|S} | P_{W_T|S'}), \quad (7)$$

where S and S' are random variables drawn independently from μ^n , and $\text{KL}(P_{W_T|S} | P_{W_T|S'})$ is the KL divergence from $P_{W_T|S'}$ to $P_{W_T|S}$. Note that this KL divergence indicates the stability of the learned parameter from two datasets, S and S' . We also note that $P_{W_T|S'}$ can be regarded as the data-dependent prior. Thus, this KL divergence is tighter than that of the data-independent prior, which is often used in the probably approximately correct (PAC)-Bayes bound³.

The key idea is to analyze *the time evolution of the KL divergence*, which is summarized in the following lemma:

Lemma 1. *Suppose that the same assumptions in Theorem 4 hold. Then, for any $t \in \mathbb{N}$, we have*

$$\text{KL}(P_{W_t|S} | P_{W_t|S'}) \leq e^{-\frac{\eta}{4\beta c_{\text{LS}}}} \text{KL}(P_{W_{t-1}|S} | P_{W_{t-1}|S'}) + \eta V_{\Delta} + c_3 \eta, \quad (8)$$

where V_{Δ} and c_3 is the constant term w.r.t. $\{M, m, b, d, \beta, s^2\}$.

² The Gaussian assumption can be relaxed, e.g., to a Gaussian mixture, in the theorems and corollaries shown in this paper. The detailed discussions are provided in Appendix F.3. ³ We can confirm this from the fact that $\mathbb{E}_{S, S'} \text{KL}(P_{W_T|S} | P_{W_T|S'}) = \mathbb{E}_S \text{KL}(P_{W_T|S} | P_{W_T}) - \mathbb{E}_{S'} \text{KL}(P_{W_T|S'} | P_{W_T})$, where P_{W_T} is a data-independent prior distribution.

We will discuss the details of Lemma 1 in Section 3.3. By recursively applying Eq. (8) from $t = 0$ to T , we obtain

$$\text{KL}(P_{W_T|S}|P_{W_T|S'}) \leq \frac{1 - e^{-\frac{\eta T}{4\beta c_{\text{LS}}}}}{1 - e^{-\frac{\eta}{4\beta c_{\text{LS}}}}} \eta (V_{\nabla} + c_3) \stackrel{(i)}{\leq} 4\beta c_{\text{LS}} \left(1 \wedge \frac{\eta T}{4\beta c_{\text{LS}}}\right) \frac{1}{1 - \frac{\eta}{4\beta c_{\text{LS}}}} (V_{\nabla} + c_3),$$

which is based on the fact that $\text{KL}(P_{W_0|S}|P_{W_0|S'}) = 0$. Note that bound (i) can be obtained from $e^{-\frac{\eta}{4\beta c_{\text{LS}}}} < 1 - \frac{\eta}{4\beta c_{\text{LS}}} + \frac{\eta^2}{16\beta^2 c_{\text{LS}}^2}$ when $\frac{\eta}{4\beta c_{\text{LS}}} \leq 1$, and $e^{-\frac{\eta T}{4\beta c_{\text{LS}}}} \geq 1 - \frac{\eta T}{4\beta c_{\text{LS}}}$.

3.3 Proof outline of Lemma 1 under the continuous Langevin diffusion

Here, we organize our ideas for the proof of Lemma 1 that are important in the derivation of Theorem 4. For simplicity, we now provide an intuitive explanation and an outline of the proof under the continuous Langevin diffusion setting. Note that the results of Theorem 4 and Lemma 1 are based on the SGLD setting, and their proofs are shown in Appendix F.

The Langevin diffusion is defined as

$$dW_t = -\nabla F(W_t, S)dt + \sqrt{2\beta^{-1}}dH_t, \quad (9)$$

where dH_t is the standard Brownian motion in \mathbb{R}^d . Note that, in this context, t expresses the continuous time and the *full-batch* gradient $\nabla F(W_t, S)$ is used. The stationary distribution of Eq. (9) is given as the Gibbs distribution $\pi(dw) \propto \exp(-\beta F(w, S))$.

With some abuse of notation, let us denote $P_{W_t|S}$ as the conditional distribution obtained using Eq. (9) and express its density as ρ_t . Then, the FP equation [2] for Eq. (9) can be obtained as

$$\frac{\partial \rho_t(w, t)}{\partial t} = \nabla \cdot \left(\frac{1}{\beta} \nabla \rho_t + \rho_t \nabla F(w, S) \right). \quad (10)$$

Similarly, we can define the Langevin diffusion when we use dataset S' and the conditional distribution using that diffusion as $P_{W_t|S'}$ with the density γ_t , obtaining the FP equation in the form of ρ_t replaced by γ_t in Eq. (10).

Now we analyze the time evolution of $\text{KL}(P_{W_t|S}|P_{W_t|S'}) = \text{KL}(\rho_t|\gamma_t)$ at time t , i.e., $\partial \text{KL}(\rho_t|\gamma_t)/\partial t$. By utilizing the FP equations of ρ_t and γ_t and the Cauchy–Schwartz inequality, we obtain the following upper bound:

$$\frac{\partial \text{KL}(\rho_t|\gamma_t)}{\partial t} \leq -\frac{1}{2\beta} \mathbb{E} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2 + \frac{\beta}{2} \mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2. \quad (11)$$

The second term on the right-hand side of Eq. (11) represents the stability of the gradient with respect to the randomness of the training dataset $S, S' \sim \mu^n$, which leads to V_{∇} in Lemma 1 under the SGLD setting. Hereafter, we define $\mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2$ as \tilde{V}_{∇_t} .

By introducing $\nabla \log \pi(w)$ into $\mathbb{E} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2$ in Eq. (11), we obtain

$$\begin{aligned} \frac{\partial \text{KL}(\rho_t|\gamma_t)}{\partial t} &\leq -\frac{1}{4\beta} \mathbb{E} \|\nabla \log \rho_t - \nabla \log \pi\|^2 + \frac{1}{2\beta} \mathbb{E} \|\nabla \log \pi\|^2 + \frac{1}{\beta} \mathbb{E} \nabla \log \rho_t \cdot \nabla \log \gamma_t + \frac{\beta}{2} \tilde{V}_{\nabla_t} \\ &\leq -\frac{1}{4\beta c_{\text{LS}}} \text{KL}(\rho_t|\pi) + \frac{1}{2\beta} \Omega(\rho_t, \gamma_t, \pi) + \frac{\beta}{2} \tilde{V}_{\nabla_t} \\ &\leq -\frac{1}{4\beta c_{\text{LS}}} \left(\text{KL}(\rho_t|\gamma_t) + \mathbb{E} \log \frac{\gamma_t}{\pi} \right) + \frac{1}{2\beta} \Omega(\rho_t, \gamma_t, \pi) + \frac{\beta}{2} \tilde{V}_{\nabla_t}, \end{aligned} \quad (12)$$

where the first inequality is from the fact that $-x^2 \leq -\|x - y\|^2/2 + y^2$ for $x, y \in \mathbb{R}^d$ and the second one is from the LSI. We introduced $\Omega(\rho_t, \gamma_t, \pi) := \mathbb{E}_{\rho_t} \|\nabla \log \pi\|^2 + 2\mathbb{E}_{\rho_t} \nabla \log \rho_t \cdot \nabla \log \gamma_t$ to simplify the notation.

By integrating $e^{\frac{t}{4\beta c_{\text{LS}}}} \frac{\partial \text{KL}(\rho_t|\gamma_t)}{\partial t}$ in Eq. (12) over $t \in [0, \eta]$ and rearranging it, we obtain

$$\begin{aligned} \text{KL}(\rho_{\eta}|\gamma_{\eta}) &\leq e^{-\frac{\eta}{4\beta c_{\text{LS}}}} \text{KL}(\rho_0|\gamma_0) \\ &\quad + \int_0^{\eta} e^{-\frac{(\eta-t)}{4\beta c_{\text{LS}}}} \left(\frac{\beta}{2} \tilde{V}_{\nabla_t} - \frac{1}{4\beta c_{\text{LS}}} \mathbb{E} \log \frac{\gamma_t}{\pi} + \frac{1}{2\beta} \Omega(\rho_t, \gamma_t, \pi) \right) dt. \end{aligned} \quad (13)$$

In Appendix F, we show that the terms related to π in Eq. (13) can be bounded by using the techniques of Raginsky et al. [29] and Vempala and Wibisono [32].

We next derive an upper bound for the following terms in Eq. (13): $\mathbb{E}[\log \frac{\gamma_t}{\pi}]$ and $\Omega(\rho_t, \gamma_t, \pi)$ by using the *parametrix method* for the FP equation [12, 27], which allows us to expand the FP equation's solution via the heat kernel. On the basis of this expansion, we can upper bound Eq. (13) as

$$\int_0^\eta e^{-\frac{(\eta-t)}{4\beta c_{LS}}} \left(-\frac{1}{4\beta c_{LS}} \mathbb{E} \log \frac{\gamma_t}{\pi} + \frac{1}{2\beta} \Omega(\rho_t, \gamma_t, \pi) \right) dt \leq \mathcal{O}(\eta). \quad (14)$$

By combining Eq. (14) with Eq. (13), we obtain the continuous version of Lemma 1.

The same procedure can be used for the SGLD setup. The difference from the continuous Langevin diffusion case is that the discretization errors and the effects of using a stochastic gradient are taken into account, resulting in the appearance of an additional constant⁴ in the above bounds (see Appendix F for details).

3.4 Additional discussion on our bound in terms of stability

We conclude this section by presenting further discussion on our bound in terms of stability with a concrete example.

As shown in Eq. (7), the information-theoretic generalization bound is closely related to the stability in KL divergence under the different training datasets. However, our bound in Theorem 4 incorporates the constant term c_2 , which is irrelevant to stability, alongside the stability term V_∇ . If we can avoid the occurrence of c_2 , the resulting upper bound of $\text{KL}(P_{W_t|S}|P_{W_t|S'})$ would be dominantly expressed by V_∇ , and as a result, we may obtain a bound where the relationship between generalization and stability is more directly represented.

The problematic constant term c_2 arises from $c_3\eta$ in Lemma 1 analyzing the time evolution of stability in KL divergence. Specifically, the term $c_3\eta$ is the byproduct of treating the general dissipative function using LSI. Actually, it is possible to avoid the problematic constant term $c_3\eta$ and derive bounds that are evaluated solely on the basis of stability-related metrics in specific examples, such as strongly convex or bounded (non-convex) losses with l_2 -regularization. For simplicity, we show this fact using the following theorem under the Langevin diffusion (LD) setting, where the probability induced by Eq. (9) is expressed as $P_{W_T|S}$.

Theorem 5. *Suppose that Assumptions 1 and 2 are satisfied and that $F(w, z)$ is R -strongly convex ($0 < R < \infty$). Then, for any $T \in \mathbb{R}_+$, we have*

$$\frac{\partial \text{KL}(\rho_t|\gamma_t)}{\partial t} \leq -\frac{R}{4} \text{KL}(\rho_t|\gamma_t) + \frac{\beta}{2} \mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2, \quad (15)$$

and

$$|\text{gen}(\mu, P_{W_T|S}; L)| \leq \sqrt{\frac{2\beta\sigma_g^2}{n} \int_0^T e^{-\frac{(T-t)R}{4}} \mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2 dt}. \quad (16)$$

A similar bound in Eq. (16) (with R replaced by $\lambda/e^{8\beta C}$) can be obtained for bounded non-convex losses with l_2 -regularization, where $F(w, z) = F_0(w, z) + \frac{\lambda}{2} \|w\|^2$ ($0 < \lambda < \infty$) and $F_0(w, z)$ is C -bounded ($0 \leq C < \infty$). The full proof is summarized in Appendix F.4.

When comparing with Lemma 1, we can see that, in Eq. (15), stability-unrelated constants do not appear in the time evolution of KL divergence at each time step. Therefore, the resulting generalization bound is also independent of such constants. Furthermore, when compared with Theorem 2, which adds up the stability terms at all time steps, our bound is dominated by the stability terms near the final time step, as those at earlier time steps decrease geometrically by $e^{-\frac{R}{4}}$. This indicates that the stability around the initial time steps is of lesser importance in evaluating the final generalization performance.

Note that our bounds are closely related to the bound indicated in Proposition 9 of Mou et al. [23], which was also derived by focusing on stability. The bound of Mou et al. [23] primarily assesses

⁴ This constant is evaluable (see Vempala and Wibisono [32] or Kinoshita and Suzuki [19]).

generalization errors focusing on the *gradient norm* near the conclusion of training. In contrast, our bounds evaluate it through the norm of *differences in gradients*, emphasizing the state in the proximity of training completion. In other words, our bound allows for the evaluation of generalization errors using a stability measure that is more closely related to generalization performance than the gradient norm. This benefit originates from our approach, which tracks the time evolution of MI-related stability in Eq. (15) on the basis of information-theoretic generalization bounds, in contrast to the PAC-Bayes bounds derived from the direct analysis of stability measures as in Mou et al. [23].

4 Generalization analysis for SGLD directly using a training loss

In this section, we consider the setting that the generalization performance is measured by a training loss f directly as in Eq. (4). We show that this is possible by demonstrating that loss functions of SGLD are sub-exponential under smooth and dissipative assumptions (Section 4.1). On the basis of this fact, we obtain for the first time an information-theoretic generalization bound of SGLD that is similar to Theorem 4. Finally, combining these results with existing optimization error bounds provides an excess risk bound with improved convergence (Section 4.2).

4.1 Smooth and dissipative loss function of SGLD is sub-exponential

To perform an information-theoretic analysis for SGLD, it is necessary to know the tail behavior of $f(W, Z)$. Our contribution here is showing that a loss function of SGLD under smooth and dissipative assumptions is sub-exponential.

Theorem 6. *Suppose that Assumptions 2, 3 and 4 are satisfied. Let $P_{W_T} = \mathbb{E}_S[P_{W_T|S}]$ be the marginal distribution of the output obtained using the SGLD algorithm at the T -th iteration. Then, for any $\eta \in (0, 1 \wedge \frac{m}{5M^2})$ and $T \in \mathbb{N}$, $f(W_T, Z)$ is sub-exponential under the distribution $P_{W_T} \otimes \mu$. That is, there exist positive constants σ_e^2 and ν w.r.t. $\{m, \beta, M, b, d, s^2\}$ ⁵ such that*

$$\log \mathbb{E}_{W_T \otimes Z} \left[e^{\lambda(f(W_T, Z) - \mathbb{E}_{W_T \otimes Z}[f(W_T, Z)])} \right] \leq \frac{\sigma_e^2 \lambda^2}{2} \quad \text{for all } |\lambda| < \frac{1}{\nu}.$$

Proof sketch. First, note that under Assumptions 2 and 3, for any $z \in \mathcal{Z}$, we obtain

$$\frac{m}{3} \|w\|^2 - \frac{b}{2} \log 3 \leq f(w, z) \leq \frac{M}{2} \|w\|^2 + M \sqrt{\frac{b}{m}} \|w\| + A, \quad (17)$$

where A is a positive constant (see Lemma 8 in Appendix G.1 for its explicit form). By employing Lemma B.2 from Farghly and Rebeschini [10], we can show the following fact: for any $p \in \mathbb{N}$,

$$\mathbb{E} \|W_T\|_2^{2p} \leq \mathbb{E} \|W_0\|_2^{2p} + c(p), \quad (18)$$

where

$$c(p) := \frac{1}{m} \left(\frac{6}{m} \right)^{p-1} \left(1 + \frac{2^{2p} p (2p-1) d}{m\beta} \right) \left[\left(2b + 8 \frac{M^2}{m^2} b \right)^p + 1 + 2 \left(\frac{d}{\beta} \right)^{p-1} (2p-1)^p \right].$$

This implies that W_T is a sub-Gaussian random variable according to Proposition 2.5.2 in Vershynin [33]. To show the sub-exponential property, we directly upper-bound $\mathbb{E}_{W_T \otimes Z} [e^{\lambda(f(W_T, Z) - \mathbb{E}_{W_T \otimes Z}[f(W_T, Z)])}]$ by considering the Taylor expansion of the exponential moment and using Eqs. (17) and (18) (see Appendix G.2 for the complete proof). \square

Remark 1. *In previous information-theoretic analysis studies [28, 25, 35], it is often assumed that a loss function $l(w, Z)$ is sub-Gaussian under the distribution μ for all $w \in \mathcal{W}$. In contrast, Theorem 6 holds under the distribution $P_{W_T} \otimes \mu$, not conditioned on $w \in \mathcal{W}$.*

We can interpret the sub-exponential property of SGLD intuitively as follows. Under Assumptions 2 and 3, the loss function grows at most as a quadratic function shown in Eq. (17). The conditional distribution of the parameters follows the Gaussian distribution, and the square of the Gaussian random variable is known as the chi-square (χ^2) random variable [34]. According to these facts, we expect that the behavior of the loss function resembles that of the χ^2 -random variable; therefore, it is sub-exponential since the χ^2 -distribution is also sub-exponential [34]. Theorem 6 validates this intuition.

⁵ The explicit form of σ_e^2 and ν can be seen in Appendix G.2.

4.2 Generalization bounds for SGLD using the same loss for training and evaluation

On the basis of Theorem 6, we can derive the following information-theoretic generalization bound for SGLD even if a surrogate loss is not used. In contrast to Theorem 4, an assumption regarding the tail behavior of a loss function such as Assumption 1 is not necessary.

Corollary 1. *Suppose that Assumptions 2, 3, and 4 are satisfied. Then, for any $\beta \geq 2/m$, $\eta \in (0, 1 \wedge \frac{m}{5M^2} \wedge 4\beta c_{LS})$, and $T \in \mathbb{N}$, we obtain*

$$|\text{gen}(\mu, P_{W_T|S}; F)| \leq \Psi^{*-1} \left(\frac{c_1}{n} \left(1 \wedge \frac{\eta T}{4\beta c_{LS}} \right) (V_{\nabla} + c_2) \right),$$

where

$$\Psi^{*-1}(y) = \begin{cases} \sqrt{2\sigma_e^2 y} & \text{if } y \leq \frac{\sigma_e^2}{2\nu} \\ \nu y + \frac{\sigma_e^2}{2\nu} & \text{otherwise} \end{cases},$$

c_1 and c_2 are the same as in Theorem 4, and σ_e^2 and ν are the same as in Theorem 6.

Proof sketch. This is the direct consequence of the sub-exponential property from Theorem 6 and the upper bound of MI in Eq. (7) (see Appendix G.3 for the complete proof). \square

Remark 2. *Despite the assumptions of Corollary 1 being the same as those made by Farghly and Rebeschini [10] except for the initial distribution and step size, the resulting bound becomes 0 as $n \rightarrow \infty$ without being dependent on inverse stepsize.*

We conclude this section by introducing our excess risk bound. Let us define the excess risk as follows: $\text{Excess}(\mu, P_{W|S}) := \mathbb{E}_{W,S}[F_{\mu}(W) - F_{\mu}(w^*)]$, where $w^* = \text{argmin}_{w \in \mathcal{W}} F_{\mu}(w)$. Under this definition, we derive the following upper bound for the excess risk by utilizing Corollary 1.

Corollary 2. *Suppose that Assumptions 2, 3, and 4 are satisfied. Then, for any $\beta \geq 2/m$, $\eta \in (0, 1 \wedge \frac{m}{5M^2} \wedge 4\beta c_{LS})$, and $T \in \mathbb{N}$, we obtain*

$$\text{Excess}(\mu, P_{W_T|S}) = \mathcal{O} \left(\sqrt{\frac{(\eta T \wedge 1)}{n}} + e^{-\eta T/c_{LS}} + \sqrt{\eta} + c_{\text{err}} \right),$$

where c_{err} is the positive constant w.r.t. $\{M, m, b, d, \beta\}$ corresponding to the optimization error.

We show the complete proof in Appendix G.4. In contrast with the existing excess risk studies, our bound does not diverge with increasing t owing to the time-independent generalization bound in Corollary 1.

5 Related studies and discussion

In this section, we compare our generalization bounds with those in related studies. Table 1 shows the order of each bound value along with its assumptions for a loss function.

SGLD analysis with/without changing losses. The existing generalization error bounds in Table 1 are time-dependent; namely, we need to impose restrictive conditions for the step size η in terms of t to achieve a generalization bound that decays to zero with increasing sample size [29, 28, 25, 35] (see the right column in Table 1). Some important applications of SGLD do not satisfy these conditions. For instance, the short-run Markov chain Monte Carlo [26] method used in energy-based models [17] adopts SGLD with a *fixed* step size. Another example is the cyclic SGLD [42] used in deep learning, where the step size is *periodically increased or decreased* to facilitate escape from local optima.

Farghly and Rebeschini [10] first analyzed the generalization error of SGLD by using smoothness and dissipative assumptions, which are broadly used in sampling and non-convex optimization studies [29, 41, 4, 44]. Their bound is time-independent; the bound does not diverge with time and achieves the order $\mathcal{O}(n^{-1/2})$. However, the bound depends on the inverse of step size $\eta^{-1/2}$ owing to the reflection coupling [9], which results in the unnatural behavior of decreasing η with increasing t . Farghly and Rebeschini [10] also derived a bound that does not suffer from this problem by assuming the Lipschitz loss function with weight decay; however, these assumptions excessively restrict the

Table 1: Comparison of our bounds with those in existing studies. Our bounds are time-independent and bounded even if $\eta \rightarrow 0$. **(I)** denotes the information-theoretic approach and **(S)** denotes the stability analysis approach. The symbol * means that the sub-Gaussian assumption is unnecessary for our bounds when using the same loss for training and generalization performance evaluation. Namely, our bounds can be derived under more relaxed assumptions for a loss function in this case.

Study	Assumptions for a loss function	Expected generalization error bound
(S) Raginsky et al. [29] (Thm. 2.1.)	Dissipative, Smoothness	$\mathcal{O}(\eta t + e^{-\eta t/c} + 1/n)$
(S) Mou et al. [23] (Thm. 1.)	Bounded, Lipschitz	$\mathcal{O}(\sqrt{\eta t}/n)$
(S) Mou et al. [23] (Thm. 2.)	Lipschitz, Sub-Gaussian, (Weight decay) ⁶	$\mathcal{O}(\sqrt{\eta} \log(t+1)/n)$
(I) Pensia et al. [28] (Cor. 1.)	Lipschitz, Sub-Gaussian	$\mathcal{O}(\sqrt{\eta t}/n)$
(I) Negrea et al. [25] (Thm. 3.1.)	Sub-Gaussian	$\mathcal{O}(\sqrt{\eta t}/n)$
(S) Farghly and Rebeschini [10] (Thm. 3.1.)	Lipschitz, Smoothness, Weight decay	$\mathcal{O}((\eta t \wedge 1)(1/n + \sqrt{\eta}))$
(S) Farghly and Rebeschini [10] (Thm. 4.1.)	Dissipative, Smoothness	$\mathcal{O}((\eta t \wedge 1)(\sqrt{\eta^{-1}}/n + \sqrt{\eta}))$
(I) Wang et al. [35] (Thm. 1.)	Sub-Gaussian	$\mathcal{O}(\sqrt{\eta t}/n)$
(I) Ours (Thm. 4 and Cor. 1)	Dissipative, Smoothness, Sub-Gaussian*	$\mathcal{O}(\sqrt{(\eta t \wedge 1)/n})$

class of loss functions and algorithms. In contrast to these bounds, our bound is time-independent and does not require scaling η , t , and n to achieve $\mathcal{O}(n^{-1/2})$.

Time evolution analysis of MI via FP equation. The analysis of SGLD using the FP equation has been successfully used in the convergence analysis of SGLD [32, 19]. These studies present analyses of the discretization errors and convergence properties of the unadjusted Langevin algorithm, SGLD, and variance-reduction SGLD (SVRG-LD) [8], comparing them with the continuous Langevin dynamics through the FP equation.

In generalization error analysis, the FP equation is mainly used to analyze the time evolution of KL divergence appearing in a generalization bound on the basis of the stability approach. Li et al. [20] analyzed the time evolution of the KL divergence between the probability densities of the parameters obtained from two training datasets that differ by only one data point under the bounded loss assumption. Mou et al. [23] also studied the KL divergence and Hellinger divergence, and they derived a generalization error bound on the basis of the PAC-Bayes notion [22]. Our idea is similar to these: we analyze the time evolution of the KL divergence between the probability densities of the parameters obtained from two training datasets. The differences between our approach and other approaches are twofold. First, we do not assume weight decay or Lipschitz continuity but instead derive our analysis assuming smoothness and dissipativity. Second, the KL divergence we analyzed is tighter than that of the PAC-Bayes bound with data-independent prior dealt by Mou et al. [23].

6 Limitations and future work

In this paper, we provide a generalization analysis of SGLD, where Gaussian noise is a fundamental assumption for our theoretical results. Thus, it is difficult to extend our analysis to other noisy iterative algorithm variants with a different noise, such as differentially private SGD with Laplace or uniform noise [36]. Another limitation of this study is that we have estimated the sub-exponential parameter roughly with respect to the dimensions of the model parameters. Further investigation of the sub-exponentiality of smooth and dissipative losses, and improvement of the dependence on dimensionality, are crucial for enhancing the practicality of our generalization bounds. The sub-exponential property of a loss function is expected to be helpful in fields other than generalized error analysis. For example, this property opens up room for new theoretical analysis policies that employ useful concentration and transport inequalities [34] in the sampling and optimization context. We hope that the analysis presented in this paper goes beyond generalization analysis and provides valuable insights into understanding the characteristics of machine learning.

⁶ The order of the bound in Mou et al. [23] varies with the choice of regularization parameters and decay factors. In this paper, we adopt the order of this bound in Table 1 of Farghly and Rebeschini [10]. For a more comprehensive discussion, we refer to Section 5.2 of Mou et al. [23].

Acknowledgments and Disclosure of Funding

We sincerely appreciate the anonymous reviewers for their insightful feedback. FF was supported by JSPS KAKENHI Grant Number JP23K16948. FF was supported by JST, PRESTO Grant Number JPMJPR22C8, Japan. MF was supported by RIKEN Special Postdoctoral Researcher Program. MF was supported by JST, ACT-X Grant Number JPMJAX210K, Japan.

References

- [1] R. Amit, B. Epstein, S. Moran, and R. Meir. Integral probability metrics PAC-bayes bounds. In *Advances in Neural Information Processing Systems*, 2022.
- [2] D. Bakry, I. Gentil, and M. Ledoux. *Analysis and Geometry of Markov Diffusion Operators*, volume 348. Springer Science & Business Media, 2013.
- [3] Y. Bu, S. Zou, and V. V. Veeravalli. Tightening mutual information-based bounds on generalization error. *IEEE Journal on Selected Areas in Information Theory*, 1(1):121–130, 2020.
- [4] N. H. Chau, É. Moulines, M. Rásonyi, S. Sabanis, and Y. Zhang. On stochastic gradient Langevin dynamics with dependent data streams: The fully nonconvex case. *SIAM Journal on Mathematics of Data Science*, 3(3):959–986, 2021.
- [5] T. Chen, E. Fox, and C. Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pages 1683–1691, 2014.
- [6] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- [7] T. Deck and S. Kruse. Parabolic differential equations with unbounded coefficients – A generalization of the parametrix method. *Acta Applicandae Mathematica*, 74:71–91, 2002.
- [8] K. A. Dubey, S. J. Reddi, S. A. Williamson, B. Póczos, A. J. Smola, and E. P. Xing. Variance reduction in stochastic gradient Langevin dynamics. In *Advances in Neural Information Processing Systems*, volume 29, pages 1154–1162, 2016.
- [9] A. Eberle. Reflection couplings and contraction rates for diffusions. *Probability Theory and Related Fields*, 166:851–886, 2021.
- [10] T. Farghly and P. Rebeschini. Time-independent generalization bounds for SGLD in non-convex settings. In *Advances in Neural Information Processing Systems*, volume 34, pages 19836–19846, 2021.
- [11] V. Feldman, I. Mironov, K. Talwar, and A. Thakurta. Privacy amplification by iteration. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 521–532, 2018.
- [12] A. Friedman. *Partial Differential Equations of Parabolic Type*. Courier Dover Publications, 2008.
- [13] S. B. Gelfand and S. K. Mitter. Recursive stochastic algorithms for global optimization in \mathbb{R}^d . *SIAM Journal on Control and Optimization*, 29(5):999–1018, 1991.
- [14] J. K. Hale. *Asymptotic Behavior of Dissipative Systems*. Mathematical surveys and monographs. American Mathematical Society, 1988.
- [15] H. Harutyunyan, M. Raginsky, G. Ver Steeg, and A. Galstyan. Information-theoretic generalization bounds for black-box learning algorithms. In *Advances in Neural Information Processing Systems*, pages 24670–24682, 2021.
- [16] U. G. Haussmann and E. Pardoux. Time reversal of diffusions. *The Annals of Probability*, 14(4):1188–1205, 1986.
- [17] G. E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002.

- [18] Y. Jiang, B. Neyshabur, H. Mobahi, D. Krishnan, and S. Bengio. Fantastic generalization measures and where to find them. In *The Tenth International Conference on Learning Representations*, 2020.
- [19] Y. Kinoshita and T. Suzuki. Improved convergence rate of stochastic gradient Langevin dynamics with variance reduction and its application to optimization. In *Advances in Neural Information Processing Systems*, pages 19022–19034, 2022.
- [20] J. Li, X. Luo, and M. Qiao. On generalization error bounds of noisy gradient methods for non-convex learning. In *The Eighth International Conference on Learning Representations*, 2020.
- [21] R. Livni. Information theoretic lower bounds for information theoretic upper bounds. *arXiv preprint arXiv:2302.04925*, 2022.
- [22] D. A. McAllester. PAC-Bayesian stochastic model selection. *Machine Learning*, 51(1):5–21, 2003.
- [23] W. Mou, L. Wang, X. Zhai, and K. Zheng. Generalization bounds of SGLD for non-convex learning: Two theoretical viewpoints. In *Proceedings of the 31st Conference on Learning Theory*, volume 75, pages 605–638, 2018.
- [24] W. Mou, N. Flammarion, M. J. Wainwright, and P. L. Bartlett. Improved bounds for discretization of Langevin diffusions: Near-optimal rates without convexity. *Bernoulli*, 28(3):1577 – 1601, 2022.
- [25] J. Negrea, M. Haghifam, G. K. Dziugaite, A. Khisti, and D. M. Roy. Information-theoretic generalization bounds for SGLD via data-dependent estimates. In *Advances in Neural Information Processing Systems*, volume 32, pages 11015–11025, 2019.
- [26] E. Nijkamp, M. Hill, S.-C. Zhu, and Y. N. Wu. Learning non-convergent non-persistent short-run MCMC toward energy-based model. In *Advances in Neural Information Processing Systems*, volume 32, pages 5232–5242, 2019.
- [27] G. A. Pavliotis. *Stochastic processes and applications: Diffusion processes, the Fokker-Planck and Langevin equations*, volume 60. Springer, 2014.
- [28] A. Pensia, V. Jog, and P.-L. Loh. Generalization error bounds for noisy, iterative algorithms. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 546–550, 2018.
- [29] M. Raginsky, A. Rakhlin, and M. Telgarsky. Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis. In *Proceedings of the 30th Conference on Learning Theory*, volume 65, pages 1674–1703, 2017.
- [30] D. Russo and J. Zou. Controlling bias in adaptive data analysis using information theory. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51, pages 1232–1240, 2016.
- [31] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Stochastic convex optimization. In *Proceedings of the 22nd Conference on Learning Theory*, 2009.
- [32] S. Vempala and A. Wibisono. Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices. In *Advances in Neural Information Processing Systems*, volume 32, pages 8094–8106, 2019.
- [33] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [34] M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- [35] H. Wang, Y. Huang, R. Gao, and F. Calmon. Analyzing the generalization capability of SGLD using properties of Gaussian channels. In *Advances in Neural Information Processing Systems*, volume 34, pages 24222–24234, 2021.

- [36] H. Wang, R. Gao, and F. P. Calmon. Generalization bounds for noisy iterative algorithms using properties of additive noise channels. *Journal of Machine Learning Research*, 24(26):1–43, 2023.
- [37] Z. Wang and Y. Mao. On the generalization of models trained with SGD: Information-theoretic bounds and implications. In *The Tenth International Conference on Learning Representations*, 2022.
- [38] Z. Wang and Y. Mao. Tighter information-theoretic generalization bounds from supersamples. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pages 36111–36137, 2023.
- [39] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pages 681–688, 2011.
- [40] A. Xu and M. Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *Advances in Neural Information Processing Systems*, volume 30, pages 2524–2533, 2017.
- [41] P. Xu, J. Chen, D. Zou, and Q. Gu. Global convergence of Langevin dynamics based algorithms for nonconvex optimization. In *Advances in Neural Information Processing Systems*, volume 31, pages 3126–3137, 2018.
- [42] R. Zhang, C. Li, J. Zhang, C. Chen, and A. G. Wilson. Cyclical stochastic gradient MCMC for Bayesian deep learning. In *The Eighth International Conference on Learning Representations*, 2020.
- [43] Y. Zhang, P. Liang, and M. Charikar. A hitting time analysis of stochastic gradient Langevin dynamics. In *Proceedings of the 30th Conference on Learning Theory*, volume 65, pages 1980–2022, 2017.
- [44] Y. Zhang, Ö. D. Akyildiz, T. Damoulas, and S. Sabanis. Nonasymptotic estimates for stochastic gradient Langevin dynamics under local conditions in nonconvex optimization. *Applied Mathematics & Optimization*, 87(25), 2021.

A Notation used in the main paper

We summarize the notation we used in the main part of this paper.

Category	Symbol	Meaning
Scalars and constants	$n \in \mathbb{N}$	The sample size
	$w \in \mathbb{R}$	Model parameters (deterministic)
	$w^* \in \mathbb{R}$	$\operatorname{argmin}_{w \in \mathcal{W}} L_\mu(w)$ (deterministic)
	$W \in \mathbb{R}$	Model parameters (random variables)
	$t, T \in \mathbb{N}$	An iteration of the SGLD algorithm
	$W^{(T)}$	The joint random variables appearing in all the iterations until T
	$d \in \mathbb{R}$	The number of parameter dimensions
	$k \in \mathbb{N}$	The number of samples in a mini-batch B ($\leq n$)
	\mathbf{I}_d	Identity matrix with d rows and d columns
	$\eta (= \eta) \in \mathbb{R}$	The learning rate
	$\beta_t (= \beta) \in \mathbb{R}$	The inverse temperature
	$\xi_t \in \mathbb{R}$	Gaussian noise sampled from $\mathcal{N}(0, \mathbf{I}_d)$
	σ_g^2	A positive constant for sub-Gaussian random variables
	L	A positive constant in the Lipschitz continuous function
	M	A positive constant in the smoothness condition
	m, b	Positive constants in the dissipative condition
	s^2	A positive finite Gaussian variance for the initial parameter distribution P_{W_0}
	$\sigma \in \mathbb{R}$	A finite fourth moment of the initial parameters W_0
	$c_{LS} \in \mathbb{R}$	The logarithmic Sobolev constant
	$\tilde{V}_{\nabla_t} \in \mathbb{R}$	The expected value of $\mathbb{E} \ \nabla F(W_t, S) - \nabla F(W_t, S')\ ^2$
σ_e^2	A positive constant for sub-exponential random variables	
$c_{\text{err}} \in \mathbb{R}$	The constant w.r.t. $\{M, m, b, d, \beta\}$ corresponding to the optimization error	
Sets and sequences	\mathcal{Z}	The instance space
	\mathcal{W}	The parameter space
	\mathbb{R}, \mathbb{R}_+	The set of real numbers and that of positive real numbers
	\mathbb{N}	The set of natural numbers
	$[n] := \{1, \dots, n\}$	The set of all integers between 1 and n
	$S, S' := \{Z_i\}_{i=1}^n (\in \mathbb{R})$	The i.i.d. samples from μ^n
	$B \subset [n]$	A mini-batch set
Probability and information theory	$(B_i)_{i=1}^\infty$	An i.i.d. sequence of random variables specifying indexes
	$(\xi_t)_{t=1}^\infty$	An i.i.d. sequence of Gaussian noise ξ_t
	μ	An unknown data generating distribution
	$P_{W S}$	A conditional distribution w.r.t. W given S via SGLD (or the continuous Langevin diffusion)
	$\mathcal{N}(\mathbf{m}, \Sigma)$	Gaussian distribution with mean $\mathbf{m} \in \mathbb{R}^d$ and covariance $\Sigma \in \mathbb{R}^{d \times d}$
	$P \otimes Q$	The product distribution
	$I(W; S)$	The mutual information between W and S
	\mathbb{E}_x	The expectation w.r.t. x
	\mathbb{E}	The expectation w.r.t. all randomness
	$\operatorname{Var}[\nabla f(W, B) W]$	The gradient variance w.r.t. B conditioned by W
$\operatorname{Var}[\nabla f(W, B)]$	The gradient variance w.r.t. B	
$\operatorname{KL}(P Q)$	The Kullback–Leibler divergence of P from Q	
Functions	ρ_t	The density of $P_{W_t S}$
	γ_t	The density of $P_{W_t S'}$
	π	The Gibbs distribution (stationary distribution of the continuous Langevin diffusion)
	$l : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$	an original loss function
	$f : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$	a surrogate loss function
	L_μ, F_μ	The population risk based on an original or a surrogate loss
	L_S, F_S	The empirical risk based on an original or a surrogate loss
	$F(w, B)$	The empirical risk with l or f on a mini-batch B
$\nabla F(w, B)$	The gradient of $F(w, B)$ w.r.t. w	
$\operatorname{gen}(\mu, P_{W S}; L), \operatorname{gen}(\mu, P_{W S}; F)$	The expected generalization error based on an original or a surrogate loss	
$\operatorname{Excess}(\mu, P_{W S})$	The excess risk defined as $\mathbb{E}_{W,S}[F_\mu(W) - F_\mu(w^*)]$	

B Additional information for dissipative losses

Here, we provide the additional information for losses with dissipativity in Assumption 3.

The dissipative assumption plays an essential role in guaranteeing the geometrical convergence of SGLD to the stationary distribution. We note that convergence to the stationary distribution is crucial for reducing training error since the stationary distribution in this context corresponds to the Gibbs posterior distribution of the given loss function. The dissipative assumption is widely used in the research on sampling or non-convex potential function optimization; thus, it is a fundamental property that enables optimization with SGLD rather than strong constraint conditions for generalization. As Mou et al. [24] discussed, the dissipative assumption is weaker than convexity and strong convexity.

Many non-convex losses commonly used in practice satisfy the dissipative property. First, all strongly convex and convex losses obviously satisfy dissipativity. The dissipative losses also include losses that are strongly convex or convex when sufficiently far from zero, that is, there exists $m, R > 0$ such that $\|x - y\|^2 \geq R, (x - y) \cdot (\nabla f(x, z) - \nabla f(y, z)) \geq m\|x - y\|^2$ or that $\|x - y\|^2 \geq R, (x - y) \cdot (\nabla f(x, z) - \nabla f(y, z)) \geq 0$ for all x and y (refer to Mou et al. [24]). This means that the dissipative losses include the non-convex losses that have a local optimum somewhat close to zero and losses whose tail behavior is similar to the strongly convex losses. A typical example of losses that satisfy dissipativity is non-convex losses with l_2 regularization that used for many machine learning models including deep learning models (see Mou et al. [23]). Because of its capability to handle many non-convex losses, the dissipative condition is often employed in the theoretical analysis of non-convex optimization, such as Raginsky et al. [29] and Xu et al. [41].

In the Bayesian context, for example, we often use the negative log-likelihood losses, which satisfy the dissipative property if the likelihood distribution satisfies the Poincaré inequality [2, 32]. Poincaré inequality is applicable to a wide range of practical likelihood distributions, such as log-concave distributions, distributions obtained via bounded perturbations of Poincaré-inequality-satisfying (PIS) distributions, distributions with Gaussian convolution added to bounded losses, distributions formed by Lipschitz continuous transformations of PIS distributions, and direct sums of PIS distributions [2]. Therefore, the dissipative assumption covers many useful Bayesian models, including Bayesian deep learning models [32, 24].

In essence, the dissipative assumption allows for the broad treatment of not only general non-convex optimization in (deep) machine learning but also non-convex losses used in Bayesian inference and Bayesian machine learning. On the other hand, it is essential to note that thick-tailed losses, such as long-tailed t -distributions or Cauchy distributions, cannot be handled as a dissipative loss [24].

C Difference from generalization error bounds based on uniform convergence

In this section, we discuss the difference between the generalization error bounds based on the information-theoretic (IT) approach and that on the basis of the uniform convergence (UC) notion.

The generalization error bound based on the UC notion guarantees that the generalization error of all hypotheses in the algorithm’s output space simultaneously vanishes as the size of the training data increases, ensuring the convergence of generalization error. Furthermore, this bound asserts that within the empirical risk minimization (ERM) principle, it suffices to output any hypothesis from the class that minimizes empirical risk, and by measuring model complexities such as VC-dimension or Rademacher complexity, one can evaluate generalization performance. In other words, the UC-based generalization bounds offer non-trivial guarantees only when the hypothesis class utilized by the algorithm, along with its complexity, is moderately constrained.

On the other hand, deep neural networks (DNNs) are included in vast hypothesis classes where the model complexity drastically increases with the size and depth of the network. When applied to such models, UC-based bounds turn into a vacuous metric due to the exceedingly large model complexity. Furthermore, bounds based on UC rely solely on the hypothesis space and are unable to leverage beneficial statistics obtained from algorithms or datasets, which sometimes results in an inability to capture the true essence of generalization performance. For instance, the gradient variance w.r.t. model parameters exhibit a strong correlation with the generalization performance of deep learning [18]; however, this correlation cannot be represented within UC-type bounds (see Amit et al. [1] for details). This observation leads to the recent interest in data and algorithmic-dependent generalization bounds, such as the PAC-Bayes and IT-based generalization error bounds.

The strength of the IT-based analysis lies in its capacity to directly incorporate the algorithm- and data-dependent statistics related to the generalization performance, such as the gradient variance instead of the model complexity, into the generalization error upper bounds. Especially, the gradient variance is empirically known to have a stronger correlation with the generalization performance of DNNs [18] in comparison to statistics appearing in uniform convergence analysis contexts (e.g., VC dimension, the number of parameters d , and the norm of parameters). Although the gradient variance implicitly depends on d , it is widely recognized that, in practice, the gradient variance becomes reasonably small as the training proceeds [18].

In our bounds, such as Theorem 4, the generalization error bound is expressed through a quantity that reflects the stability w.r.t. variations in the training data, which is closely related to the generalization properties [23, 25, 35]. While this quantity is expressed via the expectation of gradients and thus is implicitly dependent on d like the gradient variance, it is expected to decrease as training progresses and generalization performance is being enhanced [23, 25]⁷. Consequently, IT-based bounds offer a sensible generalization bound even in models with significantly high complexity, such as DNNs. This is why it has gained attention in the context of SGLD’s generalization analysis.

In short, the core aim of IT-based analysis is to offer practical bounds that effectively account for models with high complexity, like DNNs, by directly integrating empirically validated statistics associated with generalization obtained from datasets and an algorithm. Active discussions within

⁷ Note that some constants in our bounds explicitly depends on d under non-convex and dissipative losses. Removing this dependency is our significant future work. As can be seen in Appendix F.3, however, this dependence does not occur in the convex loss case.

the realm of IT-based analysis revolve around how to analyze the generalization performance of DNNs, which involve non-convex losses, to derive bounds that lead to an accurate understanding of generalization performance.

D Further discussion for convergence and dependence on dimensions

Here, we provide further discussions on IT-based bounds including ours from the perspectives of convergence and dependence on parameter dimension especially focusing on convex losses.

D.1 On the convergence of SGLD and our bounds in convex losses

As shown by Shalev-Shwartz et al. [31], there is a convex problem in which the unique solution of ERM fails to generalize. This shows that any optimization algorithm executed for an infinite algorithm iteration must fall into one of two categories: either it never converges to the minimum, or it fails to generalize. The SGLD algorithm leads to the former phenomenon because the obtained parameters via SGLD do not converge to the (local) minima under a fixed temperature parameter β for the Gaussian noise coefficient even after infinite iterations.

SGLD rather ensures convergence to a stationary distribution, known as the Gibbs posterior distribution $\pi(dw) \propto e^{-\beta F(w,S)}$, when β remains fixed over time steps. In essence, the trajectory of parameters via the SGLD algorithm gets closer to the minima and then explores its vicinity due to the addition of Gaussian noise to the gradient. Therefore, while convergence to a target distribution occurs, convergence to the minima itself is not achieved without controlling the noise via β .

Although SGLD does not converge to the minima, it boasts a distinct edge in its ability to explore parameters globally, even within non-convex problems, thanks to the Gaussian noise. This property enables the evaluation of how the obtained expected loss w.r.t. the stationary distribution deviates from that with the global minima. Specifically, we can evaluate this difference by factors that depend on parameter dimensions d , β , and the constants appearing in the assumptions for the potential function, such as dissipativity and smoothness, as elaborated in Appendix G.4. Furthermore, we can also derive the upper bounds for the population risk and excess risk both for convex and non-convex losses (see Raginsky et al. [29], Xu et al. [41]).

D.2 On dependence on parameter dimension of IT-based bounds

Recently, Livni [21] has shown that every algorithm that guarantees non-trivial population loss on convex problems, must carry dimension-dependent information on the sample. Together with our Theorem 4, this fact implies that, if the temperature β is dimension independent, then SGLD will not achieve non-trivial population loss on the (convex) construction in Livni [21]. Alternatively, one could choose dimension-dependent β in SGLD but then algorithmic-independent generalization bounds can be easily (and have been) obtained via standard uniform convergence argument.

Unfortunately, removing the dependence on the parameter dimension d is difficult or unavoidable even if our framework is utilized when analyzing the generalization error of discretized Langevin dynamics such as SGLD through the MI between the dataset and parameters. On the other hand, existing and our IT-based bounds such as Theorems 2 and 5 are expressed by the gradient variance or the stability of the expected gradient, which implicitly depend on d but could be smaller than it as training proceeds and the generalization performance is enhanced. We refer to Appendix C for an explanation of the advantages of this property in the IT-based bounds.

In order to theoretically mitigate this reliance on dimensionality, it could be imperative to explore an alternative approach to evaluating generalization that deviates from the MI between parameters and data, which forms the cornerstone of this paper. One possible avenue is, for instance, the utilization of conditional mutual information (CMI) involving super-samples (e.g., Wang and Mao [38]), as highlighted in Livni [21], as well as methods to quantify the MI between the learned hypothesis and dataset (hMI), instead of focusing on the parameters [15]. However, the drawback of these approaches is that it becomes challenging to explicitly incorporate statistics directly obtained from algorithms, such as the gradient variance, into the understanding and evaluation of generalization despite being analyses of algorithm-dependent generalization performance. Seeking IT-based bounds that not only represent algorithm- and data-dependent statistics related to generalization performance, such as

gradient variance but also theoretically eliminate dimension dependence constitutes a significant future work in the context of the IT-based generalization analysis field.

E Theoretical properties of SGLD and Langevin diffusion

Here, we show some theoretical properties of SGLD under Assumptions 2 and 3.

Lemma 2 (Adapted from Farghly and Rebeschini [10]). *Suppose that Assumptions 2 and 3 are satisfied. Then, for any $z \in \mathcal{Z}$, we have*

$$\|\nabla f(0, z)\| \leq M\sqrt{\frac{b}{m}}.$$

Proof. We straightforwardly obtain the above claim from Assumptions 2 and 3 with $w = 0$. \square

Lemma 3 (Modified version from Raginsky et al. [29]). *Suppose that Assumption 2 is satisfied. Then, for any $z \in \mathcal{Z}$ and all $w \in \mathcal{W}$, we have*

$$\|\nabla f(w, z)\| \leq M\|w\| + M\sqrt{\frac{b}{m}}.$$

Proof. Raginsky et al. [29] derived the upper bound of the gradient as $\|\nabla f(w, z)\| \leq M\|w\| + B$ by assuming the the following condition: $\|\nabla f(0, z)\| \leq B$ ($B > 0$). We replace the constant B by $M\sqrt{\frac{b}{m}}$ based on Lemma 2. \square

Lemma 4 (Modified version from Xu et al. [41]). *Suppose that Assumptions 2 and 3 are satisfied. Then, for any $z \in \mathcal{Z}$, we have*

$$\mathbb{E}\|\nabla F_S(w) - \nabla F(w, B)\|^2 \leq \frac{8(n-k)M^2(\|w\|^2 + \frac{k}{m})}{k(n-1)} := 8\delta M^2 \left(\|w\|^2 + \frac{k}{m} \right),$$

where $\delta := \frac{n-k}{k(n-1)} \in (0, 1]$.

Proof. Xu et al. [41] assumed that $\nabla F_S(w)$ is dissipative. In contrast, we posed the dissipative assumption on $\nabla f(w, z)$ for each z following Farghly and Rebeschini [10]. We then modified the upper bound of the stochastic gradient shown in Xu et al. [41]. \square

Lemma 5 (Modified version from Raginsky et al. [29] and Xu et al. [41]). *Suppose that Assumptions 2 and 3 are satisfied. Let $\eta \in (0, 1 \wedge \frac{m}{5M^2})$ be fixed. Then, for any $z \in \mathcal{Z}$ and any $t \in \mathbb{N}$, we have*

$$\mathbb{E}\|W_{t+1}\|^2 \leq (1 - 2\eta m + 10\eta^2 M^2)\mathbb{E}\|W_t\|^2 + 2\eta \left(b + 10\eta M^2 \frac{b}{m} + \frac{d}{\beta} \right),$$

and

$$\mathbb{E}\|W_{t+1}\|^2 \leq \begin{cases} 2\eta(b + 10\eta M^2 \frac{b}{m} + \frac{d}{\beta}) & (A_{\eta, m, M} \leq 0) \\ (1 - 2\eta m + 10\eta^2 M^2)^t \mathbb{E}\|W_0\|^2 + 2\frac{b + 10\eta M^2 \frac{b}{m} + \frac{d}{\beta}}{m - 5\eta M^2} & (0 \leq A_{\eta, m, M} \leq 1), \end{cases} \quad (19)$$

where $A_{\eta, m, M} := (1 - 2\eta m + 10\eta^2 M^2)$. Combining the inequalities in Eq. (19), we have

$$\begin{aligned} \mathbb{E}\|W_t\|^2 &\leq \mathbb{E}\|W_0\|^2 + 2 \left(1 \vee \frac{1}{m} \right) \left(b + 10\eta M^2 \frac{b}{m} + \frac{d}{\beta} \right) \\ &\leq s^2 + 2 \left(1 \vee \frac{1}{m} \right) \left(b + 10\eta M^2 \frac{b}{m} + \frac{d}{\beta} \right) =: C_0, \end{aligned}$$

where C_0 is independent of η and β and s^2 is the square moment of the initial distribution.

Proof. We slightly modified the coefficients of the upper bound of the l_2 norm of the parameter shown in Raginsky et al. [29] and Xu et al. [41] based on the upper bound of the stochastic gradient Lemma 4. \square

Combining Lemmas 3 and 5, we have the following upper bound for the stochastic gradient.

Lemma 6. *Suppose that Assumptions 2 and 3 are satisfied. Let $\eta \in (0, 1 \wedge \frac{m}{5M^2})$ be fixed. Then, for any $z \in \mathcal{Z}$ any $t \in \mathbb{N}$, we have*

$$\mathbb{E}\|\nabla F(W_t, B)\|^2 \leq \begin{cases} 2M^2\eta^2(b + 4\eta M^2 \frac{b}{m} + \frac{d}{\beta}) + M^2 \frac{b}{m} & (A_{\eta, m, M} \leq 0) \\ M^2(1 - 2\eta m + 10\eta^2 M^2)^t \mathbb{E}\|W_0\|^2 + 2M^2 \frac{b+4\eta M^2 \frac{b}{m} + \frac{d}{\beta}}{m - \eta M^2} + M^2 \frac{b}{m} & (0 \leq A_{\eta, m, M} \leq 1), \end{cases}$$

where $A_{\eta, m, M} := (1 - 2\eta m + 10\eta^2 M^2)$. Combining the above inequalities, we have

$$\mathbb{E}\|\nabla F(W_t, B)\|^2 \leq M^2 C_0 + M^2 \frac{b}{m}.$$

Proof. We obtain the result by combining Lemma 3 and 5 and using the Jensen inequality. \square

F Proofs of the generalization error bound with surrogate loss in Section 3

This section provides the complete proof of Theorem 4 restated as follows.

Theorem 4. *Suppose that Assumptions 1, 2, 3, and 4 are satisfied. Then, for any $\beta \geq 2/m$ and $\eta \in (0, 1 \wedge \frac{m}{5M^2} \wedge 4\beta c_{LS})$ and any $T \in \mathbb{N}$, we have*

$$|\text{gen}(\mu, P_{W_T|S}; L)| \leq \sqrt{\frac{2c_1\sigma_g^2}{n} \left(1 \wedge \frac{\eta T}{4\beta c_{LS}}\right)} (V_{\nabla} + c_2), \quad (6)$$

where c_1, c_2 , and V_{∇} are the positive constant terms w.r.t. $\{M, m, b, d, \beta, s^2\}$.

Our proof consists of the following three steps. We construct the FP equations for the density of the parameters under two different datasets and derive the time evolution of the KL divergence as the upper bound of the MI (Appendix F.1). We then analyze this time evolution by using the *parametrix method* (Appendix F.2) for solving the FP equation. The distinction between this section and Section 3.3 lies in the focus of SGLD, which employs stochastic gradients and random noise from $\mathcal{N}(0, \mathbf{I}_d)$, as opposed to the continuous Langevin diffusion that employs full-batch gradients and standard Brownian motion.

F.1 FP equation for SGLD and time evolution of the KL divergence

As the first step, we construct the two different FP equations for the parameter density.

We define a one-step SGLD at the initial step as follows:

$$dW_t = -\nabla F(W_0, B_0)dt + \sqrt{2\beta^{-1}}dH_t.$$

Note that, at time $t = \eta$,

$$W_{\eta} = W_0 - \eta \nabla F(W_0, B_0) + \sqrt{2\eta\beta^{-1}}H_{\eta},$$

is distributionally equivalent to

$$W_{\eta} = W_0 - \eta \nabla F(W_0, B_0) + \sqrt{2\eta\beta^{-1}}\xi,$$

where $\xi \sim N(0, \mathbf{I}_d)$.

The distribution ρ_t of W_t depends on random variables W_0 and B_0 . We thus denote the joint distribution of $\{W_0, W_t, B_0\}$ under a dataset S as $\rho_{0tB}(W_0, W_t, B_0)$, where ρ_0 is the distribution of W_0 and U is the *uniform* distribution of B_0 . Then, its conditional and marginal distribution is expressed as

$$\rho_{0tB}(W_0, W_t, B_0) = \rho_0(w_0)U(B_0)\rho_{t|0,B}(W_t|W_0, B_0) = \rho_{tB}(W_t, B)\rho_{0|t,B}(W_0|W_t, B_0).$$

Since we have introduced so many notations, for the sake of simplicity, we allow the abuse of notation and let $\rho_{t|0B}$ and ρ_t denote both the distribution and density.

On the basis of these facts, we can obtain the FP equation for $\rho_{t|0B}(w_t)$ as

$$\frac{\partial \rho_{t|0B}}{\partial t} = \nabla \cdot \left(\frac{1}{\beta} \nabla \rho_{t|0B} + \rho_{t|0B} \nabla F(w_0, B_0) \right), \quad (20)$$

and its marginal process as

$$\frac{\partial \rho_t}{\partial t} = \nabla \cdot \left(\frac{1}{\beta} \nabla \rho_t + \rho_t \mathbb{E}_{\rho_{0B|t}}[\nabla F(w_0, B_0)|w_t = w] \right), \quad (21)$$

which is derived in Vempala and Wibisono [32] and Kinoshita and Suzuki [19]. As shown above, the randomness associated with the dataset can be handled by simply taking the expectation for conditional gradients with respect to a uniform distribution U . To avoid cumbersome discussions, we omit descriptions related to the expectation with respect to B_0 from here on.

As with the first step, we can define a one-step SGLD with the joint density $\gamma_{0t}(w_0, w_t)$ and the conditional distribution $\gamma_0(w_0)\gamma_{t|0}(w_t|w_0)$ under a dataset $S' (\neq S)$, where $\gamma_0(w_0)\gamma_{t|0}(w_t|w_0)$ corresponds to the marginal distribution, i.e., $\gamma_0(w_0)\gamma_{t|0}(w_t|w_0) = \gamma_t(w_t)\gamma_{0|t}(w_0|w_t)$. We also can obtain the FP equation and its marginal process in the form of $\rho_{t|0}$ and ρ_t replaced by $\gamma_{t|0}$ and γ_t in Eqs. (20) and (21).

In Section 3.3, we obtain the upper bound of the MI as follows:

$$I(W_t; S) \leq \mathbb{E}_{S, S'} \text{KL}(P_{W_t|S} | P_{W_t|S'}) = \mathbb{E}_{S, S'} \text{KL}(\rho_t | \gamma_t).$$

By taking the derivation w.r.t. t , we have

$$\frac{\partial \text{KL}(\rho_t | \gamma_t)}{\partial t} = \int dw \left(\frac{\partial \rho_t}{\partial t} \log \frac{\rho_t}{\gamma_t} \right) - \int dw \left(\frac{\rho_t}{\gamma_t} \frac{\partial \gamma_t}{\partial t} \right). \quad (22)$$

The first and second terms can be expressed as

$$\int dw \left(\frac{\partial \rho_t}{\partial t} \log \frac{\rho_t}{\gamma_t} \right) = -\frac{1}{\beta} \int dw \nabla \log \rho_t \cdot \nabla \log \frac{\rho_t}{\gamma_t} - \int dw \rho_{0t} \nabla \log \frac{\rho_t}{\gamma_t} \cdot \nabla F(w_0, B_0),$$

and

$$\int dw \left(\frac{\rho_t}{\gamma_t} \frac{\partial \gamma_t}{\partial t} \right) = -\frac{1}{\beta} \int dw \nabla \frac{\rho_t}{\gamma_t} \cdot \nabla \gamma_t - \int dw \frac{\rho_t}{\gamma_t} \nabla \log \frac{\rho_t}{\gamma_t} \cdot \gamma_{0t} \nabla F(W'_0, B'_0),$$

where W'_0 follows γ_0 , which is the density of the initial distribution and B'_0 is the stochastic gradient based on S' .

According to these facts, Eq. (22) can be rewritten as

$$\begin{aligned} & \frac{\partial \text{KL}(\rho_t | \gamma_t)}{\partial t} \\ &= -\frac{1}{\beta} \mathbb{E}_{\rho_t} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2 \\ & \quad - \int \rho_{0t} dw \nabla \log \frac{\rho_t}{\gamma_t} \cdot \nabla F(W_0, B_0) + \int dw \frac{\rho_t}{\gamma_t} \nabla \log \frac{\rho_t}{\gamma_t} \cdot \gamma_{0t} \nabla F(W'_0, B'_0) \\ &= -\frac{1}{\beta} \mathbb{E}_{\rho_t} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2 - \int \rho_t dw \nabla \log \frac{\rho_t}{\gamma_t} \cdot \mathbb{E}_{\rho_{0|t}}[\nabla F(w_0, B_0)|W_t = w] \\ & \quad + \int dw \frac{\rho_t}{\gamma_t} \nabla \log \frac{\rho_t}{\gamma_t} \cdot \gamma_t \mathbb{E}_{\gamma_{0|t}}[\nabla F(w'_0, B'_0)|W_t = w] \\ &= -\frac{1}{\beta} \mathbb{E}_{\rho_t} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2 \\ & \quad - \int \rho_t dw \nabla \log \frac{\rho_t}{\gamma_t} \cdot (\mathbb{E}_{\rho_{0|t}}[\nabla F(W_0, B_0)|W_t = w] - \mathbb{E}_{\gamma_{0|t}}[\nabla F(W'_0, B'_0)|W_t = w]) \\ &\leq -\frac{1}{2\beta} \mathbb{E}_{\rho_t} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2 \\ & \quad + \frac{\beta}{2} \int \rho_t dw \|\mathbb{E}_{\rho_{0|t}}[\nabla F(W_0, B_0)|W_t = w] - \mathbb{E}_{\gamma_{0|t}}[\nabla F(W'_0, B'_0)|W_t = w]\|^2, \end{aligned}$$

where the final inequality comes from the Cauchy–Schwartz inequality. We define $\tilde{V}_{\nabla_t} := \int \rho_t dw \|\mathbb{E}_{\rho_{0,t}}[\nabla F(W_0, B_0)|W_t = w] - \mathbb{E}_{\gamma_{0,t}}[\nabla F(W'_0, B'_t)|W_t = w]\|^2$ for simplicity. We evaluate this term in Appendix F.2.3.

In the same way as Section 3.3, we have the following inequality by introducing the logarithm of the stationary distribution $\nabla \log \pi(w)$ into $\mathbb{E} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2$ in the above:

$$\begin{aligned} & -\mathbb{E}_{\rho_t} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2 \\ &= -\mathbb{E}_{\rho_t} \|\nabla \log \rho_t\|^2 - \mathbb{E}_{\rho_t} \|\nabla \log \gamma_t\|^2 + 2\mathbb{E}_{\rho_t} \nabla \log \rho_t \cdot \nabla \log \gamma_t \\ &\leq -\frac{1}{2}\mathbb{E}_{\rho_t} \|\nabla \log \rho_t - \nabla \log \pi\|^2 + \mathbb{E}_{\rho_t} \|\nabla \log \pi\|^2 + 2\mathbb{E}_{\rho_t} \nabla \log \rho_t \cdot \nabla \log \gamma_t, \end{aligned}$$

where we used $-x^2 \leq -\|x - y\|^2/2 + y^2$ for all $x, y \in \mathbb{R}^d$. To simplify the notation, we express the second and third terms as $\Omega(\rho_t, \gamma_t, \pi) := \mathbb{E}_{\rho_t} \|\nabla \log \pi\|^2 + 2\mathbb{E}_{\rho_t} \nabla \log \rho_t \cdot \nabla \log \gamma_t$. From the above fact, we obtain

$$\begin{aligned} \frac{\partial \text{KL}(\rho_t|\gamma_t)}{\partial t} &\leq -\frac{1}{4\beta} \mathbb{E}_{\rho_t} \|\nabla \log \rho_t - \nabla \log \pi\|^2 + \frac{1}{2\beta} \Omega(\rho_t, \gamma_t, \pi) + \frac{\beta}{2} \tilde{V}_{\nabla_t} \\ &\leq -\frac{1}{4\beta c_{\text{LS}}} \text{KL}(\rho_t|\pi) + \frac{1}{2\beta} \Omega(\rho_t, \gamma_t, \pi) + \frac{\beta}{2} \tilde{V}_{\nabla_t} \\ &\leq -\frac{1}{4\beta c_{\text{LS}}} \left(\text{KL}(\rho_t|\gamma_t) + \mathbb{E}_{\rho_t} \log \frac{\gamma_t}{\pi} \right) + \frac{1}{2\beta} \Omega(\rho_t, \gamma_t, \pi) + \frac{\beta}{2} \tilde{V}_{\nabla_t}, \quad (23) \end{aligned}$$

where the second inequality is from the LSI [2]. In the above, the LSI constant c_{LS} is defined by Bakry et al. [2] as follows:

$$\begin{aligned} c_{\text{LS}} &\leq \lambda_l := 2D_1 + 2\rho_0^{-1}(D_2 + 2), \\ \rho_0^{-1} &\leq \frac{2C(d + b\beta)}{m\beta} \exp\left(\frac{2}{m}(M + B)(b\beta + d) + \beta(A + B)\right) + \frac{1}{m\beta(d + b\beta)}, \end{aligned}$$

where $D_1 = \frac{2m^2 + 8M^2}{\beta m^2 M}$, $D_2 \leq \frac{6M(d + \beta)}{m}$, and C is the universal constant (see also Appendices B and E in Raginsky et al. [29]).

Multiplying $e^{\frac{t}{4\beta c_{\text{LS}}}}$ for both hands in Eq. (23) yields

$$e^{\frac{t}{4\beta c_{\text{LS}}}} \frac{\partial \text{KL}(\rho_t|\gamma_t)}{\partial t} \leq -e^{\frac{t}{4\beta c_{\text{LS}}}} \frac{1}{4\beta c_{\text{LS}}} \mathbb{E}_{\rho_t} \log \frac{\gamma_t}{\pi} + \frac{1}{2\beta} e^{\frac{t}{4\beta c_{\text{LS}}}} \Omega(\rho_t, \gamma_t, \pi) + \frac{\beta}{2} \tilde{V}_{\nabla_t}.$$

By integrating time $t = 0 \rightarrow \eta$, we obtain

$$\begin{aligned} e^{\frac{\eta}{4\beta c_{\text{LS}}}} \text{KL}(\rho_\eta|\gamma_\eta) &\leq \text{KL}(\rho_0|\gamma_0) - \int_0^\eta dt e^{\frac{t}{4\beta c_{\text{LS}}}} \frac{1}{4\beta c_{\text{LS}}} \mathbb{E}_{\rho_t} \log \frac{\gamma_t}{\pi} \\ &\quad + \int_0^\eta dt \frac{1}{2\beta} e^{\frac{t}{4\beta c_{\text{LS}}}} \Omega(\rho_t, \gamma_t, \pi) + \frac{\beta}{2} \int_0^\eta dt e^{\frac{t}{4\beta c_{\text{LS}}}} \tilde{V}_{\nabla_t}. \quad (24) \end{aligned}$$

We later evaluate the second and third terms of Eq. (24) in Appendix F.2. We then evaluate the fourth term \tilde{V}_{∇_t} of Eq. (24) as $\tilde{V}_{\nabla_t} \leq D_1$, where D_1 is problem dependent constant that can be independent of η , in Appendix F.2.3.

F.2 The solution for the FP equation via the parametrix method

In this section, we evaluate the second and third terms in Eq. (24) by utilizing the *parametrix method* for the FP equation [12, 27].

F.2.1 Consequences of the parametrix method for the FP equation

We summarize two essential consequences of the parametrix method used in Appendix F.2.2.

Solution expansion of the FP equation. The first consequence of this method is as follows. Given the following FP equation,

$$\nabla \cdot \left(\frac{1}{\beta} \nabla \rho_t(w) + \rho_t(w) b(w, t) \right) - \frac{\partial \rho_t(w)}{\partial t} = 0,$$

and initial condition $\rho_{t=0} = \rho_0$, the solution can be expanded as

$$\rho_t(w) = \mathbb{E}_{\xi \sim \rho_0} Z(w, t; \xi, 0) + \mathbb{E}_{\xi \sim \rho_0} \int_0^t \int_{\mathbb{R}^d} dz Z(w, t; z, \tau) \Phi(z, \tau; \xi, 0), \quad (25)$$

where

$$Z(x, t; z, \tau) := \frac{1}{\left(\frac{4\pi}{\beta}(t-\tau)\right)^{d/2}} e^{-\frac{\beta\|x-z\|^2}{4(t-\tau)}},$$

and

$$\Phi(z, \tau; \xi, t) := \sum_{n=1}^{\infty} L^n Z(z, \tau; \xi, t). \quad (26)$$

In Eq. (26), $L^n Z(z, \tau; \xi, t)$ is defined through

$$L^{n+1} Z(x, \tau; \xi, t) := \int_t^\tau \int_{\mathbb{R}^d} ds dy (LZ(x, \tau; y, s))(L^n Z(y, s; \xi, t)),$$

where

$$LZ(x, \tau; \xi, t) := b(x, \tau) \cdot \nabla_x Z(x, \tau; \xi, t),$$

and thus $L^1 Z(x, \tau; \xi, t) = LZ(x, \tau; \xi, t)$.

The above expansion requires the convergence of Eq. (26). Fortunately, this condition holds for the Langevin diffusion, for example, because $\rho_0(w)$ and $b(w, t) = \nabla F(w, S)$ satisfies the following two assumptions for the initial state and $b(w, t)$ from Lemma 3 and Assumption 4: (i) there exist positive constants a and b such that $\rho_0(w) \leq ae^{b\|w\|^2} < \infty$ for all $w \in \mathcal{W}$, and (ii) there exist some positive constants a' and b' such that $\|b(w, t)\| < a'\|w\| + b'$ for all $w \in \mathcal{W}$.

Parametrix solution is twice differentiable. Another important consequence is that the parametrix solution $\rho_t(w)$ is twice differentiable with respect to w . Under the initial distribution $\mathcal{N}(0, s^2 \mathbf{I}_d)$ with Assumption 4, we can obtain the following facts according to Friedman [12] and Pavliotis [27]:

$$\rho_t(w) \leq \frac{1}{(2\pi(s^2 + \frac{2t}{\beta}))^{d/2}} e^{-\frac{\|w\|^2}{2(s^2 + \frac{2t}{\beta})}} + \left(s^2 + \frac{2t}{\beta}\right)^{1/2} \frac{C_0}{(2\pi(s^2 + \frac{2t}{\beta}))^{d/2}} e^{-\frac{\|w\|^2}{2(s^2 + \frac{2t}{\beta})}}, \quad (27)$$

$$\sum_{i=1}^d \left| \frac{\partial \rho_t(w)}{\partial w_i} \right| \leq \frac{C_1}{(2\pi(s^2 + \frac{2t}{\beta}))^{(d+1)/2}} e^{-\frac{\|w\|^2}{2(s^2 + \frac{2t}{\beta})}} + \frac{C_2}{(2\pi(s^2 + \frac{2t}{\beta}))^{d/2}} e^{-\frac{\|w\|^2}{2(s^2 + \frac{2t}{\beta})}}, \quad (28)$$

and

$$\begin{aligned} \sum_{i,j=1}^d \left| \frac{\partial^2 \rho_t(w)}{\partial w_i \partial w_j} \right| &\leq \frac{C_3}{(2\pi(s^2 + \frac{2t}{\beta}))^{(d+2)/2}} e^{-\frac{\|w\|^2}{2(s^2 + \frac{2t}{\beta})}} \\ &+ \left(s^2 + \frac{2t}{\beta}\right)^{-1/2} \frac{C_4}{(2\pi(s^2 + \frac{2t}{\beta}))^{d/2}} e^{-\frac{\|w\|^2}{2(s^2 + \frac{2t}{\beta})}}, \end{aligned} \quad (29)$$

where $\{C_0, C_1, C_2, C_3, C_4\}$ are positive constants w.r.t. $\{m, M, \beta, d, b, s^2\}$.

Eqs. (27) and (28) can be derived by following the proof of Theorem 11 in Friedman [12]. The statement of this theorem is about the transition kernel; therefore, it corresponds to the case where the expectation with respect to $\xi \sim \rho_0$ is excluded from Eq. (25). In light of this fact, we take the convolution by the initial distribution $\xi \sim \rho_0$ for the beginning part of the proof of Theorem 11.

This approach is equivalent to taking convolutions in the overall discussion of Section 4 described by Friedman [12]. After convolution, following the proof of Theorem 11 leads to Eqs. (27) and (28). Theorem 11 does not yield results related to second-order differentials; however, we can derive Eq. (29) by combining Lemma 3 from Friedman [12] into the proof of Theorem 11 and employing a similar way as described above. While it is assumed that the coefficients of the FP equation are bounded in Friedman [12], we can relax this assumption to unbounded coefficients as shown in subsequent work such as Deck and Kruse [7].

F.2.2 Applying the parametrix method for SGLD's FP equations

Now, we get back to the SGLD setting. First of all, it should be mentioned that we do not lose generality by focusing solely on the initial iteration, i.e., $t = 0 \rightarrow \eta$. This reason is as follows.

For the initial iteration ($t = 0 \rightarrow \eta$), we can see that $b(x, t) = \mathbb{E}_{\rho_{0|t}}[\nabla F(W_0, B_0)|W_t = w]$ from Eq. (21). As we explained in Appendix F.2.1, the condition of the expansion is satisfied under Lemma 3 and Assumption 4. Thus, the solution and its differentiation can be obtained via the parametrix solution, expressed as Eqs. (27), (28), and (29) with the constants $\{C_0, C_1, C_2, C_3, C_4\}$ that depend on the problem except η . When considering the second iteration ($t = \eta \rightarrow 2\eta$), the initial distribution is expressed as ρ_η . The concern here is whether the solution of the FP for SGLD satisfies the conditions of the parametrix method in this case. Fortunately, these conditions are also satisfied in the second iteration. The initial condition of the expansion is satisfied from Eq. (27), and the condition $b(x, t) = \mathbb{E}_{\rho_{\eta|t}}[\nabla F(W_\eta, B_1)|W_t = w]$ also satisfies the condition of the FP expansion from Lemma 3. We thus have the same form of the solution in Eqs. (27), (28), and (29) at time $t = \eta \rightarrow 2\eta$. In the same way, the solution at $t \in (s\eta, (s+1)\eta]$ for $s \in \mathbb{N}$ can be expanded as the same parametrix expansion.

Bounding $\Omega(\rho_t, \gamma_t, \pi)$ (related to the third term in Eq. (24)). We can decompose $\Omega(\rho_t, \gamma_t, \pi)$ as

$$\mathbb{E}_{\rho_t} \|\nabla \log \pi\|^2 + 2\mathbb{E}_{\rho_t} \nabla \log \rho_t \cdot \nabla \log \gamma_t = \mathbb{E}_{\rho_t} \|\nabla \log \pi\|^2 + 2 \int dw \nabla \rho_t \cdot \nabla \log \gamma_t.$$

To derive the upper bound of the right-hand side, we focus on the following facts:

$$\int dw \nabla \rho_t \cdot \nabla \log \gamma_t = \int dw \sum_{i=1}^d \frac{\partial \rho_t}{\partial w_i} \frac{\partial \log \gamma_t}{\partial w_i}. \quad (30)$$

For the i -th dimension, we have

$$\int dw \frac{\partial \rho_t}{\partial w_i} \frac{\partial \log \gamma_t}{\partial w_i} = - \int dw \frac{\partial^2 \rho_t}{\partial w_i^2} \log \gamma_t \leq \int dw \left| \frac{\partial^2 \rho_t}{\partial w_i^2} \right| |\log \gamma_t|, \quad (31)$$

where we used the integration by parts from the fact that $\frac{\partial \rho_t}{\partial w_i} \rightarrow 0$ as $\|w\| \rightarrow 0$ according to the expansion in Eq. (28). Details of this argument can be found in Mou et al. [24]. By using the result in Eq. (27) for ρ_t and γ_t , we have

$$\int dw \left| \frac{\partial^2 \rho_t}{\partial w_i^2} \right| |\log \gamma_t| \leq \frac{C'_1}{(2\pi(s^2 + \frac{2t}{\beta}))^{1/2}} + C'_2, \quad (32)$$

where the Gaussian integral is used for $\frac{\partial^2 \rho_t}{\partial w_i^2}$ in Eq. (29) and ρ_t in Eq. (27) is replaced to γ_t . We note that C'_1 and C'_2 only depend on $\{m, M, \beta, d, b, s^2\}$.

Substituting Eqs. (31) and (32) into Eq. (30), we obtain

$$\int dw \nabla \rho_t \cdot \nabla \log \gamma_t \leq \frac{dC'_1}{(2\pi(s^2 + \frac{2t}{\beta}))^{1/2}} + dC'_2. \quad (33)$$

From Lemma 6, we have

$$\begin{aligned} \mathbb{E}_{\rho_t} \|\nabla \log \pi\|^2 &\leq \beta^2 \mathbb{E}_{\rho_t} \|\nabla F_S(w)\|^2 \\ &\leq \beta^2 M^2 \left(s^2 + 2 \left(1 \vee \frac{1}{m} \right) \left(b + 10M^2 \frac{b}{m} + \frac{d}{\beta} \right) \right) + \beta^2 M^2 \frac{b}{m}. \end{aligned} \quad (34)$$

Eqs. (33) and (34) leads to $\Omega(\rho_t, \gamma_t, \pi) \leq D_2$, where

$$D_2 := \beta^2 M^2 \left(s^2 + 2 \left(1 \vee \frac{1}{m} \right) \left(b + 10M^2 \frac{b}{m} + \frac{d}{\beta} \right) + \frac{b}{m} \right) + \frac{dC'_1}{(2\pi(s^2 + \frac{2t}{\beta}))^{1/2}} + dC'_2. \quad (35)$$

Bounding $-\mathbb{E}_{\rho_t} \log \frac{\gamma_t}{\pi}$ (related to the second term in Eq. (24)). By using the Kolmogorov solution of the FP equation [2], we have

$$\begin{aligned}\gamma_t(w) &= \mathbb{E}_{W_T}[\gamma_0(W_T)|W_0 = w] \\ &= \mathbb{E}_{W_T} \left[\frac{1}{(2\pi s^2)^{d/2}} e^{-\frac{\|W_T\|^2}{2s^2}} \Big| W_0 = w \right] \geq \frac{1}{(2\pi s^2)^{d/2}} e^{-\frac{\mathbb{E}_{W_T}[\|W_T\|^2|W_0=w]}{2s^2}},\end{aligned}\quad (36)$$

where the last inequality comes from Jensen's inequality. The above inequality gives us

$$\begin{aligned}-\mathbb{E}_{\tilde{W}_T} \log \gamma_t(\tilde{W}_T) &= -\mathbb{E}_{\tilde{W}_T} \log \mathbb{E}_{W_T}[\gamma_0(W_T)|W_0 = \tilde{W}_T] \\ &\leq -\mathbb{E}_{\tilde{W}_T} \mathbb{E}_{W_T}[\log \gamma_0(W_T)|W_0 = \tilde{W}_T] \\ &\leq \frac{d}{2} \log(2\pi s^2) + \frac{1}{2s^2} \mathbb{E}_{\tilde{W}_T} \mathbb{E}_{W_T}[\|W_T\|^2|W_0 = \tilde{W}_T],\end{aligned}$$

where \tilde{W}_T is the independent copy of W_T and the first and second inequalities are obtained from Jensen's inequality and Eq. (36), respectively. By using Lemma 5 twice, we obtain

$$\begin{aligned}-\mathbb{E}_{\tilde{W}_T} \log \gamma_t(\tilde{W}_T) &\leq \frac{d}{2} \log(2\pi s^2) + \frac{1}{2s^2} \left(\mathbb{E}_{\tilde{W}_T} \|\tilde{W}_T\|^2 + 2 \left(1 \vee \frac{1}{m}\right) \left(b + 10M^2 \frac{b}{m} + \frac{d}{\beta}\right) \right) \\ &\leq \underbrace{\frac{d}{2} \log(2\pi s^2) + \frac{1}{2s^2} \left(s^2 + 4 \left(1 \vee \frac{1}{m}\right) \left(b + 10M^2 \frac{b}{m} + \frac{d}{\beta}\right) \right)}_{=: B_1}.\end{aligned}$$

In addition, from Lemma 8, we have

$$\begin{aligned}\mathbb{E}_{\rho_t} \log \pi &= \beta M \mathbb{E}_{\rho_t} \|W\|^2 + \frac{\beta b}{2m} + A \\ &\leq \underbrace{\beta M \left(s^2 + 2 \left(1 \vee \frac{1}{m}\right) \left(b + 10M^2 \frac{b}{m} + \frac{d}{\beta}\right) \right)}_{=: B_2} + \frac{\beta b}{2m} + A.\end{aligned}$$

Thus, we have the upper bound of $-\mathbb{E}_{\rho_t} \log \frac{\gamma_t}{\pi}$ as

$$-\mathbb{E}_{\rho_t} \log \frac{\gamma_t}{\pi} \leq B_1 + B_2 =: D_3, \quad (37)$$

where D_3 is the positive constant only depends on $\{m, M, \beta, d, b, s^2\}$.

F.2.3 Bounding the stability term and finalizing proof

Finally, we show the upper bound of the stability term expressed as \tilde{V}_{∇_t} in Eq. (24). Similarly to Appendix F.2.2, we focus on the initial iteration $t = 0 \rightarrow \eta$. From the definition of \tilde{V}_{∇_t} , we have

$$\begin{aligned}\tilde{V}_{\nabla_t} &= \int \rho_t dw \|\mathbb{E}_{\rho_{0|t}}[\nabla F(W_0, B_0)|W_t = w] - \mathbb{E}_{\gamma_{0|t}}[\nabla F(W'_0, B'_t)|W_t = w]\|^2 \\ &\leq 2 \int \rho_t dw \|\mathbb{E}_{\rho_{0|t}}[\nabla F(W_0, B_0)|W_t = w]\|^2 + 2 \int \rho_t dw \|\mathbb{E}_{\gamma_{0|t}}[\nabla F(W'_0, B'_t)|W_t = w]\|^2.\end{aligned}$$

The first term of the above can be rewritten as

$$\int \rho_t dw \|\mathbb{E}_{\rho_{0|t}}[\nabla F(W_0, B_0)|W_t = w]\|^2 \leq \mathbb{E}_{\rho_0} \|\nabla F(W_0, B_0)\|^2,$$

by using Jensen's inequality for the conditional distribution. Since $\mathbb{E}_{\rho_0} \|\nabla F(W_0, B_0)\|^2$ can be bounded by using Eq. (34), we have

$$\begin{aligned}&\int \rho_t dw \|\mathbb{E}_{\rho_{0|t}}[\nabla F(W_0, B_0)|W_t = w]\|^2 \\ &\leq M^2 \left(s^2 + 2 \left(1 \vee \frac{1}{m}\right) \left(b + 10M^2 \frac{b}{m} + \frac{d}{\beta}\right) \right) + M^2 \frac{b}{m} =: D_4.\end{aligned}\quad (38)$$

Next, we derive the upper bound of

$$\int \rho_t dw \|\mathbb{E}_{\gamma_{0|t}}[\nabla F(W'_0, B'_t)|W_t = w]\|^2 = \mathbb{E}_{\rho_t} \mathbb{E}_{\gamma_{0|t}}[\|\nabla F(W'_0, B'_t)|W_t = w]\|^2.$$

From Lemma 3, we have $\|\nabla F(W'_0, B'_t)\|^2 \leq 2M^2\|W'_0\|^2 + 2M^2\frac{b}{m}$. Thus, we need to evaluate $\mathbb{E}_{\rho_t} \mathbb{E}_{\gamma_{0|t}}[\|W'_0\|^2|W'_t = w]$; however, it is difficult to analyze this expectation because the densities ρ and γ at time t and 0 are different.

Fortunately, we can circumvent this difficulty by using the reverse process formulae shown in Haussmann and Pardoux [16]. According to the fact that the conditional expectation $\mathbb{E}_{\gamma_{0|t}}[\cdot]$ implies the reverse process of Eq. (20). This formulae gives us the following reverse process for time s ($0 \leq s \leq t$):

$$d\tilde{W}_s = [\nabla F(\tilde{W}_{s=t}, B'_{s=t}) + 2\beta^{-1}\nabla \log \gamma_{t-s}]dt + \sqrt{2\beta^{-1}}dH_s, \quad \tilde{W}_0 \sim \gamma_t. \quad (39)$$

In the above, $B'_{s=t}$ implies B'_0 in the original forward process and thus a mini-batch sample is fixed. We obtain the relationship $\tilde{\gamma}_s = \gamma_{t-s}$, where $\tilde{\gamma}_s$ is the distribution of \tilde{W}_s . This relationship reflects the inverse process of γ_t , and we also have $\tilde{\gamma}_t = \gamma_0$ and $\tilde{\gamma}_0 = \gamma_t$. We can analyze Eq. (39) by using the parametrix method [7]. We refer to Remark 3 for the explanation that Eq. (39) satisfies the assumptions of the parametrix method [7].

Let us express $p_\gamma(y, s|x, s')$ as the transition kernel of Eq. (39). For simplicity, we express the conditional distribution given $\tilde{W}_0 \sim \tilde{\gamma}_0$ as $\tilde{\gamma}_{s=t|s=0}$, which corresponds to the above transition kernel: $\tilde{\gamma}_{s=t|s=0}(y) = p_\gamma(y, s|x = w, s' = 0)$. By fixing \tilde{W}_0 as w , we obtain $\mathbb{E}_{\gamma_{0|t}}[\|W'_0\|^2|W'_t = w] = \mathbb{E}_{\tilde{\gamma}_{s=t|s=0}}[\|\tilde{W}_t\|^2|\tilde{W}_0 = w]$. By analyzing the reverse process of γ_t , we can evaluate the second term in the upper bound of \tilde{V}_{∇_t} , i.e., $\int \rho_t dw \|\mathbb{E}_{\gamma_{0|t}}[\nabla F(W'_0, B'_t)|W_t = w]\|^2$.

We consider approximating $\tilde{\gamma}_{s=t|s=0}$ by the parametrix method to derive the upper bound of $\mathbb{E}_{\tilde{\gamma}_{s=t|s=0}}[\|\tilde{W}_t\|^2|\tilde{W}_0 = w]$. By using the upper bound of the transition kernel provided by the parametrix method in Deck and Kruse [7], we obtain

$$p_\gamma(y, s|x, s') \leq K_1(s - s')^{-d/2} e^{-K_2 \frac{\|y-x\|^2}{s-s'}},$$

where K_1 and K_2 are positive and problem-dependent constants and do not depend on $s - s'$. From the above inequality, by setting $x = w$ and $s - s' = t$, we have

$$\mathbb{E}_{\gamma_{0|t}}[\|W_0\|^2|W_t = w] \leq \tilde{C}_0(\|w\|^2 + \tilde{C}_1 t^2),$$

where \tilde{C}_0 and \tilde{C}_1 are positive and problem-dependent constants. Thus, we have

$$\begin{aligned} & \mathbb{E}_{\rho_t} \mathbb{E}_{\gamma_{0|t}}[\|\nabla F(W'_0, B'_t)|W_t = w]\|^2 \\ & \leq 2M^2 \tilde{C}_0 \left(\tilde{C}_1 t^2 + s^2 + 2 \left(1 \vee \frac{1}{m} \right) \left(b + 10M^2 \frac{b}{m} + \frac{d}{\beta} \right) \right) + 2M^2 \frac{b}{m} =: D_5. \end{aligned} \quad (40)$$

In the above, $\tilde{C}_1 t^2$ is negligibly much smaller than the other terms within $0 \leq t \leq \eta$.

From Eq. (38) and Eq. (40), we obtain

$$\tilde{V}_{\nabla_t} \leq 2(D_4 + D_5) := D_1. \quad (41)$$

We conclude this section by finalizing the proof of Theorem 4. By combining Eqs. (35), (37), and (41) with Eq. (24) and taking the expectation with respect to all of the randomness, we obtain

$$\begin{aligned} \mathbb{E}_{S, S'} \text{KL}(\rho_\eta | \gamma_\eta) & \leq e^{\frac{-\eta}{4\beta c_{\text{LS}}}} \mathbb{E}_{S, S'} \text{KL}(\rho_0 | \gamma_0) + (1 - e^{\frac{-\eta}{4\beta c_{\text{LS}}}}) D_2 \\ & \quad + 2c_{\text{LS}}(1 - e^{\frac{-\eta}{4\beta c_{\text{LS}}}}) D_3 + 2\beta^2 c_{\text{LS}}(1 - e^{\frac{-\eta}{4\beta c_{\text{LS}}}}) D_1, \end{aligned}$$

where we used the following fact:

$$\int_0^\eta dt e^{\frac{t}{4\beta c_{\text{LS}}}} = 4\beta c_{\text{LS}}(e^{\frac{\eta}{4\beta c_{\text{LS}}}} - 1).$$

Since $e^{\frac{-\eta}{4\beta c_{LS}}} \geq 1 - \frac{\eta}{4\beta c_{LS}}$ from the assumption, we have

$$\mathbb{E}_{S,S'} \text{KL}(\rho_\eta | \gamma_\eta) \leq e^{\frac{-\eta}{4\beta c_{LS}}} \mathbb{E}_{S,S'} \text{KL}(\rho_0 | \gamma_0) + \frac{\eta}{4\beta c_{LS}} D_2 + \frac{\eta}{2\beta} D_3 + \frac{\eta\beta}{2} D_1.$$

This concludes the proof.

Remark 3. We show that Eq. (39) satisfies the assumption of the parametrix method [7]. First, Deck and Kruse [7] assumes the strong regularity condition for the diffusion coefficient, which is satisfied because the diffusion coefficient in our setting is a constant. Next, we confirm the assumptions that the drift coefficient $b(w, s) := \nabla F(\tilde{W}_{s=t}, B'_{s=t}) + 2\beta^{-1} \nabla \log \gamma_{t-s}$ must satisfy. Specifically, the following two assumptions for $b(w, s)$ must be satisfied: (i) the locally Hölder continuous condition on some bounded subset in \mathbb{R}^d and (ii) the global growing condition, that is, $\|b(w, s)\| \leq c_0(\|x\| + 1)$ with some positive constant c_0 . Fortunately, for ∇F in $b(w, s)$, the assumption (i) is satisfied by Assumption 2, and the assumption (ii) holds from Lemma 3. Furthermore, $\nabla \log \gamma_{t-s}$ in $b(w, s)$ also satisfies the assumption (ii) from Lemma E.1 in Mou et al. [24]. According to the fact that $\nabla \gamma_{t-s}$ satisfies the Hölder continuous as shown in [12], we can see that $\frac{1}{\gamma_{t-s}}$ is bounded by considering the bounded set in \mathbb{R}^d . This means that $\nabla \log \gamma_{t-s} = \frac{1}{\gamma_{t-s}} \nabla \gamma_{t-s}$ in $b(w, s)$ is Hölder continuous and satisfies the assumption (i).

F.3 On relaxing Gaussian condition in Assumption 4

The Gaussian initial distribution assumption for W_0 could be relaxed. Let us consider the case when the initial distribution P_{W_0} is a mixture of Gaussian distribution, where each component of P_{W_0} satisfies Assumption 4. In our original proof, the Gaussian assumption is used when deriving the upper bound of the finite second moment at the initial state, and when analytically marginalizing out the initial state of the transition kernel given by the fundamental solution of the parametrix method. Even when using the mixture of Gaussian distribution as the initial distribution, it is possible to satisfy these conditions. The finite second-moment condition can easily be satisfied and the integration of the transition kernel can be executed by focusing on each component of the mixture distribution. Thus, by repeating the similar derivation in Appendices F.1 and F.2, we get the similar upper bound of $\mathbb{E}_{S,S'} \text{KL}(\rho_\eta | \gamma_\eta)$ even when the initial distribution is the Gaussian mixture distribution.

F.4 Proof of Theorem 5

We first show the proof of Theorem 5.

Theorem 5. Suppose that Assumptions 1 and 2 are satisfied and that $F(w, z)$ is R -strongly convex ($0 < R < \infty$). Then, for any $T \in \mathbb{R}_+$, we have

$$\frac{\partial \text{KL}(\rho_t | \gamma_t)}{\partial t} \leq -\frac{R}{4} \text{KL}(\rho_t | \gamma_t) + \frac{\beta}{2} \mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2, \quad (15)$$

and

$$|\text{gen}(\mu, P_{W_T|S}; L)| \leq \sqrt{\frac{2\beta\sigma_g^2}{n} \int_0^T e^{-\frac{(T-t)R}{4}} \mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2 dt}. \quad (16)$$

Proof. Since $f(w, z)$ is R -strongly convex function for any z , we have

$$\begin{aligned} \frac{\partial \text{KL}(\rho_t | \gamma_t)}{\partial t} &\leq -\frac{1}{2\beta} \mathbb{E} \|\nabla \log \rho_t - \nabla \log \gamma_t\|^2 + \frac{\beta}{2} \mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2 \\ &\leq -\frac{R}{4} \text{KL}(\rho_t | \gamma_t) + \frac{\beta}{2} \mathbb{E} \|\nabla F(W_t, S) - \nabla F(W_t, S')\|^2, \end{aligned}$$

where we utilized the local LSI in Theorem 5.5.2 of Bakry et al. [2]. Since the stationary distribution is $\pi \propto \exp(-\beta F(x))$, From Theorem 5.5.2 of Bakry et al. [2], γ_t satisfies the LSI with the LSI constant $2/(\beta R)$. By integrating $e^{\frac{tR}{4}} \frac{\partial \text{KL}(\rho_t | \gamma_t)}{\partial t}$ over $t \in [0, T]$ and rearranging the above, we obtain the upper bound of $I(W_T; S)$. This concludes the proof. \square

We can obtain the similar result for bounded non-convex losses with l_2 -regularization $F(w, z) = F_0(w, z) + \frac{\lambda}{2}\|w\|^2$ ($0 < \lambda < \infty$), where $F_0(w, z)$ is C -bounded ($0 \leq C < \infty$) with the initial distribution $\pi_0 \propto e^{-\frac{\beta\lambda}{2}\|w\|^2}$. From Lemma 34 in Li et al. [20], γ_t satisfies the LSI with the constant $\frac{\lambda}{e^{8\beta C}}$. Then, following the way in the proof of Theorem 5, we obtain the bound with replacing R of Eq. (16) to $\frac{\lambda}{e^{8\beta C}}$.

G Proofs of generalization analyses directly using a training loss

In this section, we provide our proof for our generalization bounds in the case when the same loss is used for training and the generalization performance evaluation (Corollaries 1 and 2). The key to deriving these bounds is showing that f in SGLD is sub-exponential under Assumptions 2, 3 (Theorem 6). Therefore, we explain how to obtain this result in Appendices G.1 and G.2 before introducing the details of proofs for our bounds in Appendices G.3 and G.4.

G.1 Preparation for the proof of sub-exponential property

We introduce some auxiliary lemmas that assure the existence of bounded local minima. These are used later for showing the sub-exponential property of a loss function in SGLD.

Lemma 7. *Suppose that Assumptions 2 and 3 are satisfied. Then, for each $z \in \mathcal{Z}$, there exists a positive constant A such that*

$$|f(0, z)| \leq A.$$

Proof. Denote \tilde{w}_z^* as a global minima of $f(\cdot, z)$ for each $z \in \mathcal{Z}$. By using Taylor's theorem around \tilde{w}_z^* , for $t \in (0, 1]$, we obtain the following equation with a parameter $\tilde{w}_z = t\tilde{w}_z^*$:

$$f(0, z) = f(\tilde{w}_z^*, z) + \nabla f(\tilde{w}_z^*, z) \cdot \tilde{w}_z^* + \frac{1}{2}\tilde{w}_z^* \cdot \nabla^2 f(\tilde{w}_z, z) \cdot \tilde{w}_z^*.$$

According to Assumption 2 and the fact that \tilde{w}_z^* is the global minima (i.e., $\nabla f(\tilde{w}_z^*, z) = 0$), we obtain

$$f(0, z) \leq f(\tilde{w}_z^*, z) + \frac{1}{2}M\|\tilde{w}_z^*\|^2. \quad (42)$$

Farghly and Rebeschini [10] has shown that all the local minima \tilde{w}_z^* are inside the ball in the Euclidean space. That is, for each $z \in \mathcal{Z}$, all \tilde{w}_z^* are located in $\overline{B(0, r)}$ with $r = \sqrt{b/m}$, where $B(x, r)$ ($r > 0$) is the ball in the Euclidean space defined as $B(x, r) := \{x \in \mathbb{R}^d : \|x - y\| < r\}$ and $\overline{B(x, r)}$ is the closure of $B(x, r)$. From this fact, we obtain $\|\tilde{w}_z^*\|^2 \leq b/m$ and Eq. (42) can be upper bounded as

$$f(0, z) \leq f(\tilde{w}_z^*, z) + \frac{Mb}{2m}.$$

Next, we show that, for each $z \in \mathcal{Z}$, the global minima $f(\tilde{w}_z^*, z)$ is bounded uniformly. Since $f(\tilde{w}_z, z)$ is continuous with respect to w for each $z \in \mathcal{Z}$ under Assumption 2, it is continuous in C when considering the closed set $C := \overline{B(0, r)}$ with $r = \sqrt{b/m}$. From the property of the continuous function in the closed set, the maximum and minimum value of $f(\tilde{w}_z, z)$ is always bounded, i.e., we have $f(\tilde{w}_z^*, z) < \infty$ for each $z \in \mathcal{Z}$. By considering the largest global minimum and denote it as \tilde{A} , we obtain $f(\tilde{w}_z^*, z) \leq \tilde{A}$ and thus

$$f(0, z) \leq \tilde{A} + \frac{Mb}{2m}.$$

This concludes the proof. \square

Under Lemma 7, we can modify the upper and lower bound for $f(w, z)$ in Raginsky et al. [29] as follows.

Lemma 8 (Modified version from Raginsky et al. [29]). *Suppose that Assumptions 2 and 3 are satisfied. Then, for any $z \in \mathcal{Z}$, we have*

$$\frac{m}{3}\|w\|^2 - \frac{b}{2}\log 3 \leq f(w, z) \leq \frac{M}{2}\|w\|^2 + M\sqrt{\frac{b}{m}}\|w\| + A.$$

Proof. Raginsky et al. [29] assumed that for any $z \in \mathcal{Z}$, there exists constant A such that $\|f(0, z)\| \leq A$. Instead, we show the existence of such A by Lemma 7. On the basis of this fact, we obtain the claim in the same way with Raginsky et al. [29]. \square

Finally, we provide the following useful lemma from Farghly and Rebeschini [10].

Lemma 9 (Lemma B.2 in Farghly and Rebeschini [10]). *Suppose that Assumptions 2, 3 and 4 are satisfied. Then, for any $\eta \in (0, 1 \wedge \frac{m}{5M^2})$, $T \in \mathbb{N}$ and $p \in \mathbb{N}$, we have*

$$\mathbb{E}\|W_T\|_2^{2p} \leq \mathbb{E}\|W_0\|_2^{2p} + c(p), \quad (43)$$

where

$$c(p) := \frac{1}{m} \left(\frac{6}{m}\right)^{p-1} \left(1 + \frac{2^{2p} p (2p-1) d}{m\beta}\right) \left[\left(2b + 8\frac{M^2}{m^2}b\right)^p + 1 + 2\left(\frac{d}{\beta}\right)^{p-1} (2p-1)^p \right].$$

If, for any $p \geq 1$, the L_p norm of a random variable X is upper-bounded as $(\mathbb{E}[X^p])^{1/p} \leq C\sqrt{p}$ with some positive constant C , then X is a sub-Gaussian random variable (see Proposition 2.5.2 in Vershynin [33]). According to the fact that $p^2 \leq e^{2p}$, it is evident that W_T is a sub-Gaussian random variable from Eq. (43).

G.2 Sub-exponential property for a loss function

We now provide the complete proof of Theorem 6. We first show the fact that a loss function in SGLD has the sub-exponential property (Appendix G.2) and explain how to evaluate the constants in the sub-exponential condition for deriving our generalization bounds (Appendix G.3).

Recall that the statement of Theorem 6 is as follows.

Theorem 6. *Suppose that Assumptions 2, 3 and 4 are satisfied. Let $P_{W_T} = \mathbb{E}_S[P_{W_T|S}]$ be the marginal distribution of the output obtained using the SGLD algorithm at the T -th iteration. Then, for any $\eta \in (0, 1 \wedge \frac{m}{5M^2})$ and $T \in \mathbb{N}$, $f(W_T, Z)$ is sub-exponential under the distribution $P_{W_T} \otimes \mu$. That is, there exist positive constants σ_e^2 and ν w.r.t. $\{m, \beta, M, b, d, s^2\}$ such that*

$$\log \mathbb{E}_{W_T \otimes Z} \left[e^{\lambda(f(W_T, Z) - \mathbb{E}_{W_T \otimes Z}[f(W_T, Z)])} \right] \leq \frac{\sigma_e^2 \lambda^2}{2} \quad \text{for all } |\lambda| < \frac{1}{\nu}.$$

Proof. Now, we proceed to the proof of sub-exponential property. To show the sub-exponential property for $f(w, z)$, it is sufficient to show that there exists a positive number c_0 such that $\mathbb{E}e^{\lambda f(w, z)} < \infty$ for all $|\lambda| \leq c_0$ (see Theorem 2.13 in Wainwright [34]). To show this, we follow the proof of Proposition 2.7.1 in Vershynin [33], which uses the Taylor expansion of the exponential moment. By considering the Taylor expansion, we have

$$\begin{aligned} \mathbb{E}e^{\lambda(f(W_T, Z) - \mathbb{E}f(W_T, Z))} &\leq 1 + \mathbb{E} \sum_{p=2}^{\infty} \frac{\lambda^p (f(W_T, Z) - \mathbb{E}f(W_T, Z))^p}{p!} \\ &\leq 1 + \mathbb{E} \sum_{p=2}^{\infty} \frac{(\lambda e)^p (f(W_T, Z) - \mathbb{E}f(W_T, Z))^p}{p^p} \end{aligned}$$

where we used $p! \geq (p/e)^p$, which is obtained by the Stirling's approximation. Later, we restrict the λ such that this series converges, and thus, we can swap the sum and expectation. From the fact that $(x+y)^p \leq 2^{p-1}x^p + 2^{p-1}y^p$ for $x, y \geq 0$, we obtain

$$\mathbb{E}(f(W_T, Z) - \mathbb{E}f(W_T, Z))^p \leq 2^{p-1} \mathbb{E}[|f(W_T, Z)|^p] + 2^{p-1} (|\mathbb{E}f(W_T, Z)|)^p.$$

Given the marginal distribution of the parameters obtained by the T -th iterate of the SGLD algorithm, i.e., $W_T \sim p(W_T)$, we have the following fact by using the result of Lemma 8 and the Cauchy-Schwartz inequality:

$$\frac{m}{3} \|W_T\|^2 - \frac{b}{2} \log 3 \leq f(W_T, Z) \leq \frac{M}{2} \|W_T\|^2 + M \sqrt{\frac{b}{m}} \|W_T\| + A \leq M \|W_T\|^2 + \frac{b}{2m} + A. \quad (44)$$

By using Eq. (44) and the inequality $(x + y)^p \leq 2^{p-1}x^p + 2^{p-1}y^p$ for $x, y \geq 0$, we have

$$\begin{aligned} \mathbb{E}f(W_T, Z)^p &\leq \mathbb{E}|f(W_T, Z)|^p \\ &\leq \left(2^{p-1}M^p\mathbb{E}\|W_T\|_2^{2p} + 2^{p-1}\left(\frac{b}{2m} + A\right)^p\right) \vee \left(2^{p-1}\left(\frac{m}{3}\right)^p\mathbb{E}\|W_T\|_2^{2p} + 2^{p-1}\left(\frac{b}{2}\log 3\right)^p\right). \end{aligned} \quad (45)$$

The term $\mathbb{E}\|W_T\|_2^{2p}$ in Eq. (45) can further be upper bounded by utilizing Eq. (43) in Lemma 9. As discussed immediately below Lemma 9, W_T is a sub-Gaussian random variable. Thus, for any $p \in \mathbb{N}$, $(\mathbb{E}[W_T^p])^{1/p} \leq C\sqrt{p}$ holds with some positive constant C . Combining Eq. (45) and the above fact leads to

$$\mathbb{E}f(W_T, Z)^p \leq C_0^p + C_1^p p^p,$$

where C_0 and C_1 are positive constants that only depend on s^2, m, M, b, d, A , and β . For the latter purpose, we introduce C_5 as

$$\mathbb{E}f(W_T, Z)^p \leq C_5 p^p,$$

where C_5 only depends on s^2, m, M, b, d, A and β . Then, we have

$$\mathbb{E}e^{\lambda(f(W_T, Z) - \mathbb{E}f(W_T, Z))} \leq 1 + \sum_{p=2}^{\infty} \frac{(\lambda e)^p C_5 p^p}{p^p} = 1 + \sum_{p=2}^{\infty} (\lambda e C_5)^p = 1 + \frac{(\lambda e C_5)^2}{1 - \lambda e C_5},$$

where $\lambda e C_5 < 1$. Moreover, by setting $\lambda e C_5 < 1/2$, we have

$$\mathbb{E}e^{\lambda(f(W_T, Z) - \mathbb{E}f(W_T, Z))} \leq 1 + 2\lambda^2 e^2 C_5^2 \leq e^{2\lambda^2 e^2 C_5^2}.$$

From the above, we can see that $f(W_T, Z)$ is a sub-exponential function with the following constants: $\sigma_e^2 := 4e^2 C_5^2$ and $\nu := \frac{1}{2eC_5}$ where C_5 only depends on s^2, m, M, b, d and A . \square

G.3 Proof of generalization error bound directly using a training loss

Here, we provide the complete proof of Corollary 1.

Corollary 1. *Suppose that Assumptions 2, 3, and 4 are satisfied. Then, for any $\beta \geq 2/m$, $\eta \in (0, 1 \wedge \frac{m}{5M^2} \wedge 4\beta_{\text{CLS}})$, and $T \in \mathbb{N}$, we obtain*

$$|\text{gen}(\mu, P_{W_T|S}; F)| \leq \Psi^{*-1} \left(\frac{c_1}{n} \left(1 \wedge \frac{\eta T}{4\beta_{\text{CLS}}} \right) (V_{\nabla} + c_2) \right),$$

where

$$\Psi^{*-1}(y) = \begin{cases} \sqrt{2\sigma_e^2 y} & \text{if } y \leq \frac{\sigma_e^2}{2\nu} \\ \nu y + \frac{\sigma_e^2}{2\nu} & \text{otherwise} \end{cases},$$

c_1 and c_2 are the same as in Theorem 4, and σ_e^2 and ν are the same as in Theorem 6.

Proof. We use the following theorem in Bu et al. [3] to derive the generalization error for sub-exponential losses.

Theorem 7 (Bu et al. [3]). *Suppose that there exist positive constants σ_e^2 and ν such that*

$$\log \mathbb{E}_{W_T \otimes Z} \left[e^{\lambda(f(W_T, Z) - \mathbb{E}_{W_T \otimes Z}[f(W_T, Z)])} \right] \leq \frac{\sigma_e^2 \lambda^2}{2} \quad \text{for all } |\lambda| < \frac{1}{\nu}.$$

Then, we have

$$|\text{gen}(\mu, P_{W_T|S}; F)| \leq \Psi^{*-1} \left(\frac{I(W_T; S)}{n} \right),$$

where

$$\Psi^{*-1}(y) = \begin{cases} \sqrt{2\sigma_e^2 y} & \text{if } y \leq \frac{\sigma_e^2}{2\nu} \\ \nu y + \frac{\sigma_e^2}{2\nu} & \text{otherwise.} \end{cases}$$

Substituting the constants of the sub-exponential property shown in Theorem 6 and the upper bound of $I(W_T; S)$ in Theorem 4 into the above completes the proof. \square

G.4 Proof of an excess risk

We rewrite our corollary as follows.

Corollary 2. *Suppose that Assumptions 2, 3, and 4 are satisfied. Then, for any $\beta \geq 2/m$, $\eta \in (0, 1 \wedge \frac{m}{5M^2} \wedge 4\beta c_{LS})$, and $T \in \mathbb{N}$, we obtain*

$$\text{Excess}(\mu, P_{W_T|S}) = \mathcal{O}\left(\sqrt{\frac{(\eta T \wedge 1)}{n}} + e^{-\eta T/c_{LS}} + \sqrt{\eta} + c_{\text{err}}\right),$$

where c_{err} is the positive constant w.r.t. $\{M, m, b, d, \beta\}$ corresponding to the optimization error.

Proof. We can decompose the excess risk at T as

$$\begin{aligned} \text{Excess}(\mu, P_{W_T|S}) &= \mathbb{E}_{W_T, S}[F_\mu(W_T) - F_S(W_T) + F_S(W_T) - F_\mu(w^*)] \\ &= \text{gen}(\mu, P_{W_T|S}; F) + \mathbb{E}_{W_T, S}[F_S(W_T) - F_\mu(w^*)], \end{aligned}$$

where the last term is called the *optimization error*. The optimization error can be bounded as

$$\begin{aligned} \mathbb{E}_{W_T, S}[F_S(W_T) - F_\mu(w^*)] &= \mathbb{E}_{W_T, S}[F_S(W_T) - \min_w F_S(w) + \min_w F_S(w) - F_S(w^*)] \\ &\leq \mathbb{E}_{W_T, S}[F_S(W_T) - \min_w F_S(w)], \end{aligned}$$

where the above inequality comes from the fact that $\mathbb{E}_{W_T, S}[\min_w F_S(w) - F_S(w^*)] \leq 0$. Let us denote ϵ_{opt} as $\mathbb{E}_{W_T, S}[F_S(W_T) - \min_w F_S(w)]$. Then, we can express the upper bound of the excess risk as follows:

$$\text{Excess}(\mu, P_{W_T|S}) \leq |\text{gen}(\mu, P_{W_T|S}; F)| + \epsilon_{\text{opt}}.$$

We first bound the ϵ_{opt} term. Using the Gibbs distribution $\pi(dw) \propto \exp(-\beta F(w, S))$ and the triangle inequality, we obtain

$$\epsilon_{\text{opt}} \leq |\mathbb{E}_{W_T, S} F_S(W_T) - \mathbb{E}_{\pi, S} F_S(W)| + |\mathbb{E}_{\pi, S} F_S(W) - \min_w F_S(w)|, \quad (46)$$

where we express the expectation under the joint distribution $\mu^N \otimes \pi$ as $\mathbb{E}_{\pi, S}$. The first term on the right-hand side of Eq. (46) is the *convergence error* of the SGLD algorithm, which can be seen as $\mathcal{O}(e^{-k\eta/\beta c_{LS}} + \sqrt{\eta})$. From Lemma 6 in Raginsky et al. [29], we have

$$|\mathbb{E}_{W_T, S} F_S(W_T) - \mathbb{E}_{\pi, S} F_S(W_T)| \leq \left(M\sigma + M\sqrt{\frac{b}{m}}\right) \mathbb{E}_S W_2(P_{W_T|S}, \pi),$$

where $\sigma^2 := \mathbb{E}_{P_{W_T|S}} \|W_T\|^2 \vee \mathbb{E}_\pi \|W\|^2$ and W_2 is the 2-Wasserstein distance. By using the T_2 inequality, we obtain $W_2(P_{W_T|S}, \pi) \leq \sqrt{c_{LS} \text{KL}(P_{W_T|S} \| \pi)}$. From Theorem 1 in Vempala and Wibisono [32], we further obtain $\text{KL}(P_{W_T|S} \| \pi) \leq \mathcal{O}(e^{-2k\eta/\beta c_{LS}} + \eta)$. Combining these results leads to $|\mathbb{E}_{W_T, S} F_S(W_T) - \mathbb{E}_{\pi, S} F_S(W_T)| = \mathcal{O}(e^{-k\eta/\beta c_{LS}} + \sqrt{\eta})$.

The second term $|\mathbb{E}_{\pi, S} F_S(W_T) - \min_w F_S(w)|$ corresponds to the *minimization error*, which can be upper-bounded by $c_{\text{err}} := \frac{d}{2\beta} \log\left(\frac{\epsilon M}{m} \left(\frac{b\beta}{d} + 1\right)\right)$ according to Proposition 11 in Raginsky et al. [29]. This completes the proof. \square