
Fast Mixing of Stochastic Gradient Descent with Normalization and Weight Decay

Zhiyuan Li

Princeton University

zhiyuanli@cs.princeton.edu

Tianhao Wang

Yale University

tianhao.wang@yale.edu

Dingli Yu

Princeton University

dingliy@cs.princeton.edu

Abstract

We prove the Fast Equilibrium Conjecture proposed by Li et al. [1], *i.e.*, stochastic gradient descent (SGD) on a scale-invariant loss (*e.g.*, using networks with various normalization schemes) with learning rate η and weight decay factor λ mixes in function space in $\tilde{O}(1/(\eta\lambda))$ steps, under two standard assumptions: (1) the noise covariance matrix is non-degenerate and (2) the minimizers of the loss form a connected, compact and analytic manifold. The analysis uses the framework of Li et al. [2] and shows that for every $T > 0$, the iterates of SGD with learning rate η and weight decay factor λ on the scale-invariant loss converge in distribution in $\ln(1 + T\lambda/\eta)/(4\eta\lambda)$ iterations as $\eta\lambda \rightarrow 0$ while satisfying $\eta \leq O(\lambda) \leq O(1)$. Moreover, the evolution of the limiting distribution can be described by a stochastic differential equation that mixes to the same equilibrium distribution for every initialization around the manifold of minimizers as $T \rightarrow \infty$.

1 Introduction

Generalization in modern deep learning has significantly deviated from classical learning theory due to the vast overparametrization in deep neural networks and is underlain by the implicit bias of training algorithms [3]. Instead of decreasing the training objective as fast as possible, the training algorithm and its hyperparameters are often tuned for good implicit bias, *i.e.*, the ability to pick empirical minimizers with good generalization among various different minimizers. Sometimes good implicit bias occurs at the cost of less efficient optimization, including the usage of large learning rates (LR) [4] or small batch size [5, 6]. Thus the training objective alone is not an effective measure of the entire training progress. In other words, behind the minimization of the training objective, there potentially exists some *hidden progress*, and the evolution of the model therein plays a crucial role in the implicit bias.

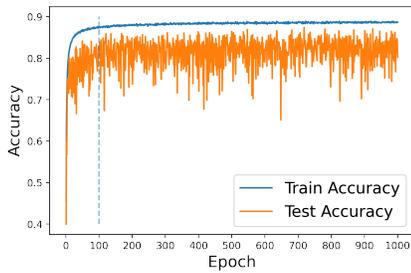
The current paper aims to provide a better theoretical understanding of such hidden progress for neural networks equipped with normalization layers (*e.g.*, BatchNorm [7], LayerNorm [8], and others [9–13]) trained by Stochastic Gradient Descent (SGD) with Weight Decay (WD), dubbed SGD+WD. For learning rate (LR) η and WD factor λ , we formulate SGD+WD as

$$x_{\eta,\lambda}(k+1) = (1 - \eta\lambda)x_{\eta,\lambda}(k) - \eta\nabla L_{\xi_k}(x_{\eta,\lambda}(k)) \quad (1)$$

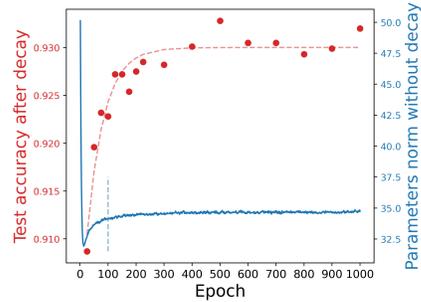
where $x_{\eta,\lambda}(k) \in \mathbb{R}^D$ is the parameter after k iterations, and L_{ξ_k} is the loss over the ξ_k -th sample with each ξ_k being sampled independently and uniformly randomly across all training data. In particular, we are interested in explaining the following phenomenon:

Longer training with SGD+WD after LR decay improves final test accuracy of normalized networks.

We demonstrate such phenomenon in Figure 1, where test accuracy after LR decay keeps improving when training accuracy plateaus. In an extreme case, Li et al. [1] empirically showed that the test



(a) Train and test accuracy for CIFAR-10 training with $\eta = 0.8$, $\lambda = 5 \cdot 10^{-4}$.



(b) Test accuracy after LR decay and the total norm of parameters before LR decay.

Figure 1: The train and test accuracy plateaus after parameter norm convergence within 100 epochs, but the generalization of SGD iterate after LR decay keeps improving. Figure 1a shows the train and test accuracy of scale invariant PreResNet trained by SGD+WD on CIFAR-10 with standard data augmentation. Each red dot in Figure 1b represents the test accuracy of model which decays LR to 10^{-3} at the corresponding epoch. The test accuracy is evaluated until achieving full training accuracy after LR decay.

accuracy of ResNet can still improve after maintaining nearly full training accuracy for thousands of epochs when trained by SGD+WD on CIFAR10. Such phenomenon is also demonstrated for a standard decoder-only Transformer trained by Adam on small arithmetic datasets and is named ‘grokking’ by Power et al. [14], where the validation accuracy can increase from random guess to full accuracy long after the almost perfect fitting of the training data.

Based on theoretical derivations, Li et al. [1] further proposed the *Fast Equilibrium Conjecture* (Conjecture 1.1), which informally says that for the normalized model trained by SGD+WD, such hidden progress happens in $\tilde{O}(1/(\eta\lambda))$ steps and the model converges to an equilibrium, and since then, further training can no longer improve the final test accuracy. A recent line of works [15, 16, 2] show that gradient noise in stochastic gradient can cause a higher order regularizing effect and improve generalization even when training loss is close to 0. In particular, Li et al. [2] proposed a mathematical framework for characterizing the implicit bias of SGD in the time scale of $O(1/\eta^2)$. Under such a time scale, the hidden progress of SGD is shown to be described by a Stochastic Differential Equation (SDE) termed the limiting diffusion, which can then be used to rigorously prove its generalization benefit in some cases (see Section 6 in [2]).

However, the $O(1/\eta^2)$ rate given by Li et al. [2] is not applicable to networks with normalization layers and WD, because the assumptions on the loss landscape made in [2] that minimizers of training loss connect as a manifold fail to hold for networks with normalization layers and WD, or more broadly, for all *scale invariant* loss (see Definition 2.2) with ℓ_2 regularization. Here a loss L is scale invariant means that $L(Cx) = L(x)$ for any $C > 0$ and parameter $x \neq 0$, which is a consequence of normalization layers. The assumption of manifold of minimizers fails because any ℓ_2 -regularized scale invariant loss has no local minimizer, not to mention the manifold of minimizers. To see this, simply note that for every x where scale invariant loss is well-defined, *i.e.*, $x \neq 0$, reducing its norm while keeping the direction of x strictly decreases ℓ_2 regularized scale invariant loss. Moreover, the loss landscape becomes unboundedly sharp around the origin. These drastic changes to loss landscape induced by normalization layers could lead to bizarre training dynamics beyond the scope of standard optimization viewpoint, *e.g.*, deep neural networks with normalization can even be trained with an exponentially increasing LR schedule [17].

1.1 Our Results

In this paper, we show that for networks with normalization trained by SGD+WD the $\tilde{O}(1/(\eta\lambda))$ rate is indeed the correct time scale for the aforementioned hidden progress and deliver a partial proof to the Fast Equilibrium Conjecture proposed by Li et al. [1]. The key observation here is that we need to rescale the SGD+WD dynamics both in time and parameter norm by leveraging the scale invariance of loss, so that the framework in [2] can again be applied to achieve an SDE-based characterization for the hidden progress. Our rescalings are motivated by the analysis in [1] for the parameter norm convergence which happens in $\tilde{O}(1/(\lambda\eta))$ steps.

Before stating the main theorem, we will first introduce some notations and restate the Fast Equilibrium Conjecture. Let $F_z(x)$ be the output of a scale invariant neural network with parameter x on data z , *i.e.*, $F_z(x) = F_z(Cx)$ for any parameter x , input data z and constant $C > 0$. In other words, the output of the network only depends on the direction of parameter, shorthand as $\bar{x} := x/\|x\|_2$.

Let Ξ be the total number of training data and $L(x) = \frac{1}{\Xi} \sum_{\xi=1}^{\Xi} L_{\xi}(x)$ be the empirical training loss. Denote $\sigma(x) = (\sigma_1(x), \dots, \sigma_{\Xi}(x))$ where each $\sigma_{\xi}(x) := (\nabla L_{\xi}(x) - \nabla L(x))/\sqrt{\Xi}$. Then we can rewrite SGD+WD (Equation (1)) as

$$x_{\eta,\lambda}(k+1) = (1 - \eta\lambda)x_{\eta,\lambda}(k) - \eta(\nabla L(x_{\eta,\lambda}(k)) + \sqrt{\Xi} \cdot \sigma_{\xi_k}(x_{\eta,\lambda}(k))). \quad (2)$$

Let $\{W(t)\}_{t \geq 0}$ be a Ξ -dimensional Brownian motion. As a common approach to analyzing SGD, the canonical SDE approximation of SGD+WD (Equation (2)) is

$$dX_{\eta,\lambda}(t) = -\eta\nabla L(X_{\eta,\lambda}(t))dt - \eta\lambda X_{\eta,\lambda}(t)dt + \eta\sigma(X_{\eta,\lambda}(t))dW(t) \quad (3)$$

The Fast Equilibrium Conjecture is stated below. The convergence rate is much faster than the $e^{-\Theta(1/\eta)}$ global mixing time of Langevin dynamics [18] and thus the conjecture gets its name.

Conjecture 1.1 (Fast Equilibrium Conjecture, Li et al. [1]). Suppose $X_{\eta,\lambda}(t)$ is a solution of (3), then for any input z , $F_z(X_{\eta,\lambda}(t))$ converges to the same equilibrium distribution independent of the initial parameter x_{init} in $\tilde{O}(1/(\eta\lambda))$ time.

We note that the above conjecture is implied by the convergence of the distribution of the parameter direction $\bar{X}_{\eta,\lambda}(t)$. Next, the main theorem of this paper is stated informally below.

Theorem 1.2 (Informal version of Theorem 5.5). *Suppose Γ is a connected manifold consisting only of local minimizers of L . Under some regularity assumptions, there is an open neighborhood U of Γ , such that for any initialization $x_{\text{init}} \in U$ and $T > 0$, as $\eta\lambda \rightarrow 0$ with $\eta \leq O(\lambda) \leq O(1)$, both $\bar{x}_{\eta,\lambda}\left(\left\lfloor \frac{\ln(\frac{2\lambda}{\eta}(e^{2T}-1)+1)}{4\eta\lambda} \right\rfloor\right)$ and $\bar{X}_{\eta,\lambda}\left(\frac{\ln(\frac{2\lambda}{\eta}(e^{2T}-1)+1)}{4\eta\lambda}\right)$ converge in distribution to the same distribution denoted by $\mu_{T,x_{\text{init}}}$. Moreover, as $T \rightarrow \infty$, $\mu_{T,x_{\text{init}}}$ weakly converges to the same equilibrium distribution for every $x_{\text{init}} \in U$.*

The **main contributions** of this paper are summarized as follows:

1. We give a SDE-based characterization (Theorem 4.4) for the limiting dynamics of SGD+WD for a scale invariant stochastic loss in the limit of $\eta\lambda \rightarrow 0$ with $\eta = O(\lambda)$ and $\lambda = O(1)$. By introducing a novel time-rescaling tailored to the scale invariant loss and weight decay, our analysis adapts the framework proposed by Li et al. [2].
2. We show that SGD without WD for a stochastic scale invariant loss has the same limiting dynamics as that of SGD+WD, but is exponentially slower (see Theorem 3.2). This is consistent to the empirical observation that turning on WD for SGD for scale invariant loss helps generalization [19, 1].
3. Under the assumption of all minimizers forming a manifold and noise being non-degenerate in the tangent space of the manifold, we show that from any initialization, the limiting dynamics of SGD+WD converges to a unique stationary distribution (see Theorem 5.4 and Theorem 5.5). This delivers a partial proof to the Fast Equilibrium Conjecture in [1].
4. Though our convergence result is asymptotic, we verify in simplified settings that the phenomena predicted by our theory happens with LR η and WD factor λ of practical scale (see Section 6 for details of experiments). We also show empirically that the mixing process exists in practical settings, and is beneficial for generalization.

2 Preliminary

Notations. We denote by \mathbb{N} the set of all nonnegative integers and \mathbb{R}_+ the set of all nonnegative real numbers. For any $k \in \mathbb{N}$, we denote by \mathcal{C}^k the set of all k times continuously differentiable functions. For any vector $u \in \mathbb{R}^D$, we denote its i -th coordinate by u_i . For any mapping $F : \mathbb{R}^D \rightarrow \mathbb{R}^D$, we denote the *Jacobian* of F at x by $\partial F(x) \in \mathbb{R}^{D \times D}$ where the (i, j) -th entry is $\partial_j F_i(x)$. We also use $\partial F(x)[u]$ and $\partial^2 F(x)[u, v]$ to denote the first and second order directional derivatives of F at x along the derivation of u (and v). With a slight abuse of notation, we view $\partial^2 F$ as a linear mapping on $\mathbb{R}^{D \times D}$ such that $\partial^2 F(x)[A] = \sum_{i,j=1}^D \partial^2 F(x)[e_i, e_j] A_{ij}$, for any $A \in \mathbb{R}^{D \times D}$. For any submanifold $\Gamma \subset \mathbb{R}^D$ and $x \in \Gamma$, we denote by $T_x(\Gamma)$ the tangent space of Γ at x . We denote by $\mathbb{1}_{\xi} \in \mathbb{R}^{\Xi}$ the one-hot vector where the ξ -th coordinate is 1, and $\mathbb{1}$ denotes the all 1 vector. We say $K \subset \mathbb{R}^D$ is a *cone* if and only if $0 \notin K$ and $\forall \alpha > 0, \alpha K \subseteq K$.

Recall that the training loss $L : \mathbb{R}^D \rightarrow \mathbb{R}$ is defined as $L = \frac{1}{\Xi} \sum_{i=1}^{\Xi} L_i$ where L_i is the loss on the i -th sample. With vast overparametrization in modern machine learning models, multiple minimizers can exist and form a manifold [20, 21]. Thus following Fehrman et al. [22], Li et al. [2], Arora et al. [23], we make the assumption below throughout this paper.

Assumption 2.1. Each loss function $L_i : \mathbb{R}^D \rightarrow \mathbb{R}$ is a \mathcal{C}^4 function. Γ is a $(D - M)$ -dimensional \mathcal{C}^2 -submanifold of \mathbb{R}^D for some integer $M \in [0, D - 1]$, where each $x \in \Gamma$ is a local minimizer of L and $\text{rank}(\nabla^2 L(x)) = M$ for all $x \in \Gamma$.

Note that $\nabla^2 L(x)$ must have zero eigenvalues in the tangent space of Γ at x , we are indeed assuming the Hessian $\nabla^2 L$ attains the maximal rank everywhere on the manifold Γ .

In this paper, we are interested in the behavior of SGD+WD with each L_i being scale invariant, or equivalently, 0-homogenous. We note that the level sets of scale invariant functions are always cones, which will be used frequently in our analysis. To this end, we make Assumption 2.3.

Definition 2.2 (Homogeneous Functions). We say a function $f : \mathbb{R}^D \setminus \{0\} \rightarrow \mathbb{R}^m$ is a k -homogeneous for some $k \in \mathbb{R}$ if and only if for all $x \in \mathbb{R}^D \setminus \{0\}$ and $\alpha > 0$, $f(\alpha x) = \alpha^k f(x)$. Specifically, we say a function f is *scale invariant* if and only if it is 0-homogeneous.

Assumption 2.3. L_i is scale invariant for each $1 \leq i \leq \Xi$ and Γ is a cone.

Below are two useful properties of homogeneous functions, whose proofs follow from directly applying chain rules.

Lemma 2.4. For any $l \in \mathbb{N}$ and k -homogeneous function f , $\nabla^l f$ is $(k - l)$ -homogeneous.

Lemma 2.5 (Euler's Theorem for Homogeneous Functions). For any real-valued k -homogeneous function f , $\langle x, \nabla f(x) \rangle = kf(x)$. Specifically, if f is scale invariant, $\langle x, \nabla f(x) \rangle \equiv 0$.

Recall that $\sigma_i(x) = \frac{1}{\sqrt{\Xi}}(\nabla L_i(x) - \nabla L(x))$, so the noise function σ is (-1) -homogeneous and thus the noise covariance $\Sigma(x) = \sigma(x)\sigma(x)^\top$ is (-2) -homogeneous.

Next, the following notion of limiting map of gradient flow plays a key role in our analysis.

Definition 2.6. For any $x \neq 0$, we define the *gradient flow* governed by $-\nabla L$ as the unique solution of $\phi(x, t) := x - \int_0^t \nabla L(\phi(x, s)) ds$ for $t \geq 0$ and denote its associated limiting map by $\Phi(x) = \lim_{t \rightarrow \infty} \phi(x, t)$ whenever the limit exists.

Throughout the paper, we use U to denote the attraction set of Γ under gradient flow, that is, $U = \{x \in \mathbb{R}^D \mid \Phi(x) \text{ is well-defined and } \Phi(x) \in \Gamma\}$. By Lemma B.15 of Arora et al. [23], U is open and Φ is \mathcal{C}^2 in U with $\nabla^2 \Phi$ being locally lipschitz.

Lemma 2.7. Under Assumption 2.3, U is a cone and Φ is 1-homogeneous in U .

Li et al. [2] established several important properties of Φ by relating the derivatives of Φ to those of L , and in particular, we recall the following characterization of $\partial \Phi$.

Lemma 2.8 (Lemma 4.3, [2]). For any $x \in \Gamma$, $\partial \Phi(x) \in \mathbb{R}^{D \times D}$ is the orthogonal projection matrix onto the tangent space $T_x(\Gamma)$. As a consequence, for any $x \in \Gamma$, Assumption 2.3 $\implies x \in T_x(\Gamma) \implies \partial \Phi(x)x = x$.

2.1 Limiting Diffusion on The Manifold of Local Minimizers

We recap the notion of *Katzenberger processes* proposed by Li et al. [2] and the characterization of the corresponding limiting diffusion based on Katzenberger's theorems [24]. In this subsection we only assume Assumption 2.1, but not Assumption 2.3.

Definition 2.9 (Uniform metric). The *uniform metric* between two functions $f, g : [0, \infty) \rightarrow \mathbb{R}^D$ is defined to be $d_U(f, g) = \sum_{T=1}^{\infty} 2^{-T} \min\{1, \sup_{t \in [0, T]} \|f(t) - g(t)\|_2\}$.

For each $n \in \mathbb{N}$, let $A_n : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and $B_n : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be two non-decreasing functions with $A_n(0) = B_n(0) = 0$, and $\{Z_n(t)\}_{t \geq 0}$ be a \mathbb{R}^Ξ -valued stochastic process. In our context of SGD+WD, given loss function $L : \mathbb{R}^D \rightarrow \mathbb{R}$, noise function $\sigma : \mathbb{R}^D \rightarrow \mathbb{R}^{D \times \Xi}$ and initialization $x_{\text{init}} \in U$, we call the following stochastic process (4) a *Katzenberger process*

$$X_n(t) = x_{\text{init}} + \int_0^t \sigma(X_n(s)) dZ_n(s) + \int_0^t X_n(t) dB_n(s) - \int_0^t \nabla L(X_n(s)) dA_n(s) \quad (4)$$

if as $n \rightarrow \infty$ the following conditions are satisfied:

1. A_n increases infinitely fast, i.e., $\forall \epsilon > 0, \inf_{t \geq 0} (A_n(t + \epsilon) - A_n(t)) \rightarrow \infty$;
2. $B_n(t)$ converges to $c \cdot t$ in uniform metric for some constant c .
3. Z_n converges in distribution to $(I_{\Xi} - \frac{1}{\Xi} \mathbf{1} \mathbf{1}^\top)W$ in uniform metric where W is a Ξ -dimension standard Brownian motion;

Theorem 2.10 (Adapted from Theorem 4.6 in Li et al. [2]). *Given a Katzenberger process $\{X_n(\cdot)\}_{n \in \mathbb{N}}$, if SDE (5) has a global solution Y in U with $Y(0) = \Phi(x_{\text{init}})$, then for any $t > 0$, $X_n(t)$ converges in distribution to $Y(t)$ as $n \rightarrow \infty$.*

$$Y(t) = \Phi(x_{\text{init}}) + \int_0^t c \partial \Phi(Y(s)) Y(s) ds + \int_0^t \partial \Phi(Y(s)) \sigma(Y(s)) dW(s) + \int_0^t \frac{1}{2} \partial^2 \Phi(Y(s)) [\Sigma(Y(s))] ds. \quad (5)$$

We note that the global solution always exists if the manifold Γ is compact. Our notion of Katzenberger process and theorem statement is slightly more general than those in Li et al. [2] to handle the weight decay. However, our formulation is still under the original framework of Katzenberger [24] and the proof in Li et al. [2] can be easily adapted to Theorem 2.10.

3 Warm-up: Simultaneous Limit Case

As a warm-up, we first consider the setting where $\eta, \lambda \rightarrow 0$ simultaneously with $\frac{\lambda}{\eta} \equiv C$ for some constant $C \geq 0$. In this special regime, we do not need to use the scale invariance property of the loss, and we can directly apply Theorem 2.10 to obtain the limiting diffusion of SGD+WD. Nonetheless, we will see the benefit of weight decay as a source of acceleration. While for the general case of $\eta\lambda \rightarrow 0$ that will be considered in Section 4, we need to carefully design a time rescaling by calibrating with the dynamics of parameter magnitude, so that under the new scaling the dynamics can still be understood as a Katzenberger process.

Now, recall the SGD+WD updates in Equation (2), and let us fix $x_{\eta, \lambda}(0) = x_{\text{init}}$ for some $x_{\text{init}} \in U$. Define $\check{X}_{\eta, \lambda}(t) = x_{\eta, \lambda}(\lfloor t/\eta^2 \rfloor)$, which is roughly equivalent to SDE (3) with $1/\eta^2$ times acceleration, and we can rewrite the discrete-time update of $x_{\eta, \lambda}$ as

$$\check{X}_{\eta, \lambda}(t) = x_{\text{init}} + \int_0^t \sigma(\check{X}_{\eta, \lambda}(s)) dZ_{\eta, \lambda}(s) + \int_0^t \check{X}_{\eta, \lambda}(s) dB_{\eta, \lambda}(s) - \int_0^t \nabla L(\check{X}_{\eta, \lambda}(s)) dA_{\eta, \lambda}(s) \quad (6)$$

where $A_{\eta, \lambda}, B_{\eta, \lambda}$ and $Z_{\eta, \lambda}$ are defined by

$$A_{\eta, \lambda}(t) = \eta \lfloor t/\eta^2 \rfloor, \quad B_{\eta, \lambda}(t) = \lambda \eta \lfloor t/\eta^2 \rfloor, \quad Z_{\eta, \lambda}(t) = \eta \sum_{k=1}^{\lfloor t/\eta^2 \rfloor} \sqrt{\Xi} \left(\mathbf{1}_{\xi_k} - \frac{1}{\Xi} \mathbf{1} \right). \quad (7)$$

Note that $A_{\eta, \lambda}(t)$ is roughly t/η which becomes very large for small η , thus the negative gradient part will drive $\check{X}_{\eta, \lambda}(t)$ rapidly towards the manifold Γ and force $\check{X}_{\eta, \lambda}(t)$ to stay close to Γ after that. On the other hand, as $\eta \rightarrow 0$, $B_{\eta, \lambda}(t)$ will converge to Ct and $Z_{\eta, \lambda}$ will weakly converge to a Brownian motion, and these terms make up the slow dynamics of SGD. More precisely, we have the following lemma summarizing the properties of these integrators, which shows Equation (6) is a valid Katzenberger process.

Lemma 3.1. *Let $A_{\eta, \lambda}, B_{\eta, \lambda}$ and $Z_{\eta, \lambda}$ be as defined in Equation (7). Then as $\eta, \lambda \rightarrow 0$ with $\frac{\lambda}{\eta} \equiv C$, it holds that (1) $A_{\eta, \lambda}$ increases infinitely fast, (2) $B_{\eta, \lambda}(t)$ converges to Ct in uniform metric and (3) $Z_{\eta, \lambda}$ converges in distribution to $(I_{\Xi} - \frac{1}{\Xi} \mathbf{1} \mathbf{1}^\top)^{1/2} W$ in uniform metric where $\{W(t)\}_{t \geq 0}$ is the Ξ -dimensional standard Brownian motion.*

Therefore, a direct application of Theorem 2.10 yields the limiting diffusion in this case.

Theorem 3.2. *Under Assumption 2.1, let $x_{\eta, \lambda}(0) \equiv x_{\text{init}} \in U, \forall \eta, \lambda > 0$ in SGD+WD (2). Consider*

$$dY_C(t) = -C \partial \Phi(Y_C) Y_C dt + \frac{1}{2} \partial^2 \Phi(Y_C) [\Sigma(Y_C)] dt + \partial \Phi(Y_C) \sigma(Y_C) dW(t) \quad (8)$$

where $\{W(t)\}_{t \geq 0}$ is the standard Brownian motion in \mathbb{R}^Ξ . Suppose SDE (8) has a global solution Y_C in U for some $C \geq 0$ with $Y_C(0) = \Phi(x_{\text{init}})$, then $x_{\eta, \lambda}(\lfloor t/\eta^2 \rfloor)$ converges in distribution to

$Y_C(t)$ as $\lambda, \eta \rightarrow 0$ with $\frac{\lambda}{\eta} \equiv C$. Also, under Assumption 2.3, SDE (8) with any $C' \geq 0$ has a global solution $Y_{C'}$ in U with $Y_{C'}(0) = \Phi(x_{\text{init}})$. Moreover, $Y_{C'}(t) \stackrel{d}{=} Y_0(\frac{e^{4C't}-1}{4C'})e^{-C't}$.

Remark 3.3. Theorem 3.2 shows that when there is no WD, the limiting diffusion is still the same as that with WD, but exponentially slower than that with WD in the regime of LR η and WD factor λ going to zero with a fixed ratio.

4 Limiting Diffusion for The General Case

In the previous section, we showed that the limiting diffusion exists when η and λ go to zero with a fixed ratio. However, the situation is more complicated in the general case where we drop the assumption of η/λ being fixed. Below we first explain the challenges in analysis and our solution for this general regime. Then we present the continuous-time analysis for SDE and the discrete-time analysis for SGD+WD. Our analysis applies for all the cases when $\eta\lambda \rightarrow 0$ with $\eta = O(\lambda) = O(1)$.

Challenges for the General Case. A concrete example for the challenge is when $\eta \rightarrow 0$ and λ be fixed as a constant, which is also the most natural and practical setting. We quickly find ourselves in a dilemma if we still want to apply Katzenberger’s theorem [24], or its simplified version Theorem 2.10: If we view WD or ℓ_2 regularization as the ‘fast’ part of the dynamics, that is, a part of the loss function, then there is no minimizer for the ℓ_2 regularized scale-invariant loss and thus it doesn’t satisfy the condition of Katzenberger’s theorem; if we view WD as some ‘slow’ dynamics and formulate it as $\frac{\lambda}{\eta} X_{\eta,\lambda} dB_{\eta,\lambda}(t)$ in Equation (6), then unlike the simultaneous limit case, $\frac{\lambda}{\eta}$ doesn’t necessarily have a limit, and thus the condition of Katzenberger’s theorem is again not met.

The above dilemma reflects two different roles of WD in early and late phase of training: in the early phase, when the norm is large, WD is more like a part of the loss function that executes the ℓ_2 regularization. In contrast, in the late phase of SGD training, especially when the norm square of parameters has stabilized to some value, *i.e.*, $\|x_{\eta,\lambda}(t)\|_2^2 \propto \sqrt{\frac{\eta}{\lambda}}$ (e.g. Figure 1b), WD should be viewed as the ‘slow’ dynamics and we can apply the analysis in the simultaneous limit case. This is because by the scale-invariance of loss L , Equation (2) can be rewritten as the following form, Equation (9), with $\tilde{\eta} = \sqrt{\eta\lambda}$, $\tilde{\lambda} = \sqrt{\eta\lambda}$, $\tilde{x}_{\tilde{\eta},\tilde{\lambda}} = (\frac{\lambda}{\eta})^{1/4} x_{\eta,\lambda}$.

$$\tilde{x}_{\tilde{\eta},\tilde{\lambda}}(k+1) = (1 - \tilde{\eta}\tilde{\lambda})\tilde{x}_{\tilde{\eta},\tilde{\lambda}}(k) - \tilde{\eta}(\nabla L(x_{\tilde{\eta},\tilde{\lambda}}(k)) + \sqrt{\Xi}\sigma_{\xi_k}(x_{\tilde{\eta},\tilde{\lambda}}(k))) \quad (9)$$

With such a rescaling, we successfully make the norm of parameters in constant scale, that is, $\|\tilde{x}_{\tilde{\eta},\tilde{\lambda}}\|_2^2 = \sqrt{\frac{\lambda}{\eta}}\|x_{\eta,\lambda}\|_2^2 = \Theta(1)$ and thus we can apply Katzenberger’s theorem. Note that we cannot do this in the early phase of SGD+WD as we start from a fixed initialization and such a rescaling will change the magnitude of the initialization.

Our Strategy for Analysis. To overcome the above dilemma, our core strategy is to introduce a novel combination of parameter rescaling $R_{\eta,\lambda}$ (Equation (11)) and time rescaling $\tau_{\eta,\lambda}$ (Equation (12)) which smoothly interpolates the early and late regime. Because the rescalings are adaptive to the norm of the parameter along the training trajectory, they allow us to apply Katzenberger’s theorem on the rescaled dynamics $\frac{X_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(t))}{R_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(t))}$. (See formal statements in Theorem 4.1). Compared with the ordinary SGD without WD studied in Li et al. [2] where the time rescaling is set to be a fixed acceleration by $1/\eta^2$ times, here the design of the time rescaling is more complicated.

Since the norm has no effect on the loss value but only affects the speed, we need to consider the dynamics of the parameter direction. To do so, we need to normalize the iterates properly. However, when the trace of the noise covariance is not constant, in general it is hard to find a close-form solution for $\|X_{\eta,\lambda}(t)\|_2$, but we can approximate it using the former special case. In specific, recall the canonical SDE approximation in (3). Li et al. [1] proved that the dynamics of $\|X_{\eta,\lambda}(t)\|_2^2$ is

$$d\|X_{\eta,\lambda}(t)\|_2^2 = -2\eta\lambda\|X_{\eta,\lambda}(t)\|_2^2 + \eta^2 \text{tr}(\Sigma(X_{\eta,\lambda}(t)))dt.$$

Suppose $\text{tr}(\Sigma(x)) \equiv \hat{\sigma}^2/\|x\|_2^2$ for some $\sigma > 0$, then the above further simplifies into $d\|X_{\eta,\lambda}(t)\|_2^2 = -2\eta\lambda\|X_{\eta,\lambda}(t)\|_2^2 dt + \frac{\eta^2 \hat{\sigma}^2}{\|X_{\eta,\lambda}(t)\|_2^2} dt$, which admits a closed-form solution:

$$\|X_{\eta,\lambda}(t)\|_2^4 = \frac{\eta\hat{\sigma}^2}{2\lambda} + e^{-4\lambda\eta t} \left(\|X_{\eta,\lambda}(0)\|_2^4 - \frac{\eta\hat{\sigma}^2}{2\lambda} \right). \quad (10)$$

This implies that the norm of the weights at the equilibrium is of order $(\eta/\lambda)^{1/4}$. Moreover, Equation (10) reflects the scaling of the norm of the iterates in terms of η and λ as we will see later.

4.1 Continuous-time Analysis for SDE

We first consider the continuous-time case of the SDE approximation (3). The main result is summarized in Theorem 4.1, which shows that the limiting diffusion exists for SDE with a suitable non-linear rescaling.

As mentioned in the previous discussion, we consider a scaling function $R_{\eta,\lambda}(t)$ inspired by the norm dynamics for the special case in Equation (10):

$$R_{\eta,\lambda}(t) = \left(\frac{\eta}{2\lambda} + e^{-4\eta\lambda t} \left(1 - \frac{\eta}{2\lambda} \right) \right)^{1/4}. \quad (11)$$

Next, to rescale the time, we define $\tau_{\eta,\lambda} : [0, \infty) \rightarrow [0, \infty)$ by

$$\tau_{\eta,\lambda}(t) = \int_{s=0}^t \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds = \frac{1}{2} \ln \left(1 + (e^{4\eta\lambda t} - 1) \frac{\eta}{2\lambda} \right) \quad (12)$$

where the second equality follows from a direct calculation. It is easy to show that $\tau_{\eta,\lambda}^{-1}(T) = \frac{\ln(\frac{2\lambda}{\eta}(e^{2T}-1)+1)}{4\eta\lambda}$. Then we have the following theorem (see Appendix E for the proof), which says the rescaled version of SDE approximation Equation (3) admits the following limiting diffusion Equation (13), where $\{W(t)\}_{t \geq 0}$ is the standard Brownian motion in \mathbb{R}^{Ξ} .

$$dY(t) = -\frac{1}{2}Y(t)dt + \frac{1}{2}\partial^2\Phi(Y(t))[\Sigma(Y(t))]dt + \partial\Phi(Y(t))\sigma(Y(t))dW(t) \quad (13)$$

Theorem 4.1. *Under Assumption 2.1 and 2.3, let $X_{\eta,\lambda}(0) \equiv x_{\text{init}} \in U$ for all $\eta, \lambda > 0$ in SDE (3). Let $R_{\eta,\lambda}(t)$ and $\tau_{\eta,\lambda}(t)$ be defined in (11) and (12). If SDE (13) has a global solution Y in U with $Y(0) = \Phi(x_{\text{init}})$, then $\frac{X_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(T))}{R_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(T))}$ converges in distribution to $Y(T)$ as $\eta\lambda \rightarrow 0$ with $\eta < 2\lambda < c$ for some constant c .*

Remark 4.2. *The additional constraint $\eta < 2\lambda$ when $\eta\lambda \rightarrow 0$ can be relaxed to $\eta < 2C\lambda$ for any constant $C > 0$. It suffices to note that for SDE (3), $X_{\eta,\lambda}$ with $X_{\eta,\lambda}(0) = x_{\text{init}}$ and $X_{C^{-1/2}\eta, C^{1/2}\lambda}$ with $X_{C^{-1/2}\eta, C^{1/2}\lambda}(0) = C^{-1/4}x_{\text{init}}$ have the same trajectories, up to a rescaling of $C^{-1/4}$. This is equivalent to replace $\eta/(2\lambda)$ by $\eta/(2C\lambda)$ in (11).*

4.2 Discrete-time Analysis for SGD+WD

Now, we proceed to analyze SGD+WD by mimicking the continuous-time behavior. Specifically, we view $R_{\eta,\lambda}(k)$ as an approximation of the norm of $x_{\eta,\lambda}(k)$, and consider the rescaled version of Equation (2) denoted by $\hat{x}_{\eta,\lambda}(k) = x_{\eta,\lambda}(k)/R_{\eta,\lambda}(k)$. Next, we introduce the time rescaling through the $\tau_{\eta,\lambda}(\cdot)$ defined in Equation (12) and denote $\tilde{t} = \tau_{\eta,\lambda}(t)$, so $t = \tau_{\eta,\lambda}^{-1}(\tilde{t})$. Now define $\tilde{X}_{\eta,\lambda}(\tilde{t}) := \hat{x}_{\eta,\lambda}(\lfloor \tilde{t} \rfloor)$, and it can be shown that (see Appendix E for the derivation)

$$\begin{aligned} \tilde{X}_{\eta,\lambda}(\tilde{t}) &= \tilde{X}_{\eta,\lambda}(\tilde{t}) + \int_{\tilde{s}=0}^{\tilde{t}} -\nabla L(\tilde{X}_{\eta,\lambda}(\tilde{s}))dA_{\eta,\lambda}(\tilde{s}) - \int_{\tilde{s}=0}^{\tilde{t}} \tilde{X}_{\eta,\lambda}(\tilde{s})dB_{\eta,\lambda}(\tilde{s}) \\ &\quad - \int_{\tilde{s}=0}^{\tilde{t}} \sigma(\tilde{X}_{\eta,\lambda}(\tilde{s}))dZ_{\eta,\lambda}(\tilde{s}) \end{aligned} \quad (14)$$

where $A_{\eta,\lambda}$, $B_{\eta,\lambda}$ and $Z_{\eta,\lambda}$ are defined by

$$A_{\eta,\lambda}(\tilde{t}) = \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(i+1)}, \quad (15)$$

$$B_{\eta,\lambda}(\tilde{t}) = \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \eta\lambda - (1 - \eta\lambda) \left(\frac{R_{\eta,\lambda}(i)}{R_{\eta,\lambda}(i+1)} - 1 \right), \quad (16)$$

$$Z_{\eta,\lambda}(\tilde{t}) = \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta\sqrt{\Xi}}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(i+1)} \left(\mathbb{1}_{\xi_i} - \frac{1}{\Xi} \mathbb{1} \right). \quad (17)$$

The convergence of $A_{\eta,\lambda}$, $B_{\eta,\lambda}$ and $Z_{\eta,\lambda}$ are summarized in the following lemma.

Lemma 4.3. Let $A_{\eta,\lambda}, B_{\eta,\lambda}$ and $Z_{\eta,\lambda}$ be as defined in Equation (15), (16) and (17) respectively. Then as $\eta\lambda \rightarrow 0$ with $\eta < 2\lambda < c$ for some constant c , it holds that (1) $A_{\eta,\lambda}$ increases infinitely fast, (2) $B_{\eta,\lambda}(t)$ converges to $\frac{t}{2}$ in uniform metric, and (3) $Z_{\eta,\lambda}$ converges in distribution to $(I_{\Xi} - \frac{1}{\Xi}\mathbf{1}\mathbf{1}^\top)^{1/2}W$ in uniform metric where $\{W(t)\}_{t \geq 0}$ is the Ξ -dimensional standard Brownian motion.

Therefore, $\tilde{X}_{\eta,\lambda}(\tilde{t})$ is also a valid Katzenberger process. Applying Theorem 2.10 yields:

Theorem 4.4. Under Assumption 2.1 and 2.3, let $x_{\eta,\lambda}(0) \equiv x_{\text{init}} \in U$ for all $\eta, \lambda > 0$ in SGD+WD (2). Let $R_{\eta,\lambda}(t)$ and $\tau_{\eta,\lambda}(t)$ be defined in (11) and (12). If SDE (13) has a global solution Y in U with $Y(0) = \Phi(x_{\text{init}})$, then for any $T > 0$, $\frac{x_{\eta,\lambda}(\lfloor \tau_{\eta,\lambda}^{-1}(T) \rfloor)}{R_{\eta,\lambda}(\lfloor \tau_{\eta,\lambda}^{-1}(T) \rfloor)}$ converges in distribution to $Y(T)$ as $\eta\lambda \rightarrow 0$ with $\eta < 2\lambda < c$ some constant c .

5 Mixing to Equilibrium

Now we proceed to study the ergodicity of the limiting diffusion (13). Omitted proofs of this section are delayed to Appendix F.

Due to the nature of the scale invariant of the loss L , we only care about the direction of $Y(t)$, i.e., $Y(t)/\|Y(t)\|_2$. To study the ergodicity of the normalized diffusion process, we need some additional assumptions on Γ and the noise covariance. For any $r > 0$, define $\Gamma_r := \Gamma \cap \{x \in \mathbb{R}^D : \|x\|_2 = r\}$. We assume that Γ_1 is compact manifold to ensure the existence of stationary distribution of the limiting diffusion process. We also need to assume that Γ_1 are connected (so is Γ) for the uniqueness of the stationary distribution.

Assumption 5.1. Γ_1 is compact and connected.

We further assume that the noise is non-degenerate on the manifold of local minimizers, so that as a Markov chain the limiting diffusion is irreducible.

Assumption 5.2 (Controllability). For each $x \in \Gamma$, $\text{span}(\{\partial\Phi(x)\sigma_i(x)\}_{i=1}^{\Xi}) = T_x(\Gamma_{\|x\|_2})$.

Assumption 5.3. $\text{tr}(\Sigma(\cdot))$ is an analytic function on $\mathbb{R}^D \setminus \{0\}$ and Γ is an analytic manifold.

Now we are ready to state our main result in this section, which is Theorem 5.4. It is proved in two cases respectively in Appendices F.3 and F.4, depending on whether the trace of gradient covariance $\text{tr}(\Sigma)$ is constant on Γ_1 or not. If it is, then the diffusion process essentially is on a $(D - M - 1)$ -dimensional manifold (after suitable rescaling), Γ_1 , just as in the analysis by Wang and Wang [25]. Otherwise, the situation becomes more complicated and the diffusion process is on Γ , a $(D - M)$ -dimensional manifold, in which case we will need the analyticity assumption (Assumption 5.3).

Theorem 5.4. Under Assumption 2.1, 2.3, 5.1, 5.2 and 5.3, starting from any initialization $Y(0) \in \Gamma$, the distribution of $\bar{Y}(t)$ converges to a unique stationary distribution π on Γ_1 in total variation.

Our main result on the fast mixing of SGD+WD and its SDE approximation follows from a direct combination of the convergence of the SGD+WD iterates proved in Theorem 4.4 and Theorem 5.4. Here note that as Γ_1 is compact, convergence in total variation implies convergence in distribution.

Theorem 5.5 (Fast Mixing of SGD+WD). Under Assumption 2.1, 2.3, 5.1, 5.2 and 5.3, let $x_{\eta,\lambda}(0) \equiv X_{\eta,\lambda}(0) \equiv x_{\text{init}} \in U$ for all $\eta, \lambda > 0$ for SGD+WD (2) and SDE approximation (3). For any $T > 0$, as $\eta\lambda \rightarrow 0$ with $\eta = O(\lambda)$ and $\lambda = O(1)$, both $\bar{x}_{\eta,\lambda}(\lfloor \frac{\ln(\frac{2\lambda}{\eta}(e^{2T}-1)+1)}{4\eta\lambda} \rfloor)$ and $\bar{X}_{\eta,\lambda}(\frac{\ln(\frac{2\lambda}{\eta}(e^{2T}-1)+1)}{4\eta\lambda})$ converge in distribution to the same distribution, denoted by $\mu_{T,x_{\text{init}}}$. Moreover, for every $x_{\text{init}} \in U$, $\mu_{T,x_{\text{init}}}$ weakly converges to the same equilibrium distribution π supported on Γ_1 as $T \rightarrow \infty$.

6 Experiments

In this section, we first empirically verify our theory for the time scaling of the dynamics in a simple setting where our theory applies. We then show that the diffusion process exists during the training of PreResNet on CIFAR-10, and it has implicit bias towards better generalization.

6.1 Verification of Time Scaling

Setting and Theoretical Prediction. We train the following normalized linear model by ℓ_2 regression: $F_z(x) = \langle \frac{x}{\|x\|_2}, z \rangle$, where $x \in \mathbb{R}^D$ is the model parameter, and $z \in \mathbb{R}^D$ is the input.

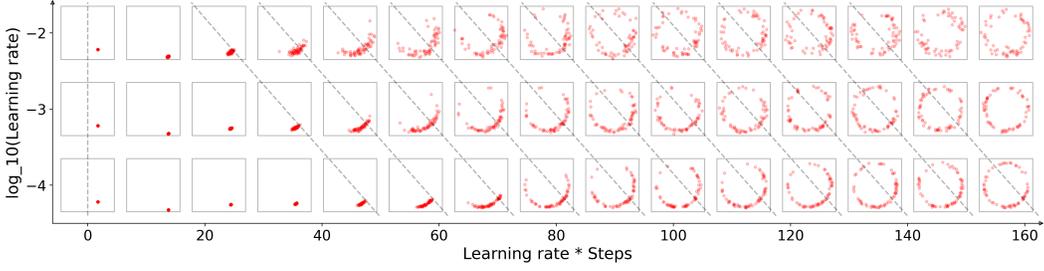


Figure 2: Scatter plots for projections of model parameters into the affine subspace containing the 1D zero-loss unit norm manifold. Models are trained with LR $\eta \in \{10^{-2}, 10^{-3}, 10^{-4}\}$ and WD factor $\lambda = 0.05$. Each small box in the figure contains 60 models that are trained with the same LR and the same number of steps. The x axis indicates the product of the LR and the number of steps of each box; the y axis indicates $\log_{10}(\text{LR})$ of each box. The dashed lines represent the time scaling $\frac{1}{2} \ln \left(1 + (e^{4\eta\lambda t} - 1) \frac{\eta}{2\lambda} \right) = T$ for different T 's (from 0 to 15.58), where t is the number of steps in SGD+WD. The dynamics are consistent with time scaling suggested by our theory (Theorem 4.1).

Let $\{z_i\}_{i=1}^N$ be the input samples, and $y_i = F_{z_i}(x^*)$ be the target label for each z_i , for some x^* . The training loss is $L(x) = \frac{1}{2N} \sum_{i=1}^N (F_{z_i}(x) - F_{z_i}(x^*))^2$. We set $N = D - 2$, so the solution space $S = \{w : \langle w - \frac{x^*}{\|x^*\|_2}, z_i \rangle = 0, \forall i \in [N]\}$ is a 2-dimensional linear space. The manifold of unit-norm global minimizers Γ_1 is then equal to $S \cap \{x \mid \|x\|_2 = 1\}$. We generate $\{z_i\}_{i=1}^N$ randomly in a way that almost surely x^* is not contained in the linear span of $\{z_i\}_{i=1}^N$. This implies that $M = \text{rank}(\nabla^2 L(x)) = D - 2$ on Γ_1 and that Γ_1 is a 1-dimensional manifold (a circle), thus Assumption 2.1 and 5.1 hold. During training, we set the loss at step t as $L^{(t)}(x) = L(x) + \langle \frac{x}{\|x\|_2}, \epsilon_t \rangle$ where $\epsilon_t \stackrel{iid}{\sim} N(0, \hat{\sigma}^2 I_D)$ for some $\hat{\sigma}$. Then the SGD+WD update rule is

$$x_{\eta,\lambda}(k+1) = (1 - \eta\lambda)x_{\eta,\lambda}(k) - \eta \left(\nabla L(x_{\eta,\lambda}(k)) + \left(I_D - \frac{xx^\top}{\|x\|_2^2} \right) \cdot \frac{\epsilon_t}{\|x\|_2} \right). \quad (18)$$

Let $\sigma(x) = \frac{\hat{\sigma}}{\|x\|_2} \left(I_D - \frac{xx^\top}{\|x\|_2^2} \right)$, and the canonical SDE approximation is

$$dX_{\eta,\lambda}(t) = -\eta \nabla L(X_{\eta,\lambda}(t)) dt - \eta \lambda X_{\eta,\lambda}(t) dt + \eta \sigma(X_{\eta,\lambda}(t)) dW(t). \quad (19)$$

For SDE approximation, we set $\Xi = D$, and thus Theorem 4.1 applies to Equation (19), suggesting that the correct time scale (number of steps for SGD+WD) for the limiting dynamics is $\tau_{\eta,\lambda}^{-1}(T) = \frac{\ln((2\lambda/\eta)(e^{2T} - 1) + 1)}{4\eta\lambda}$ for each $T \geq 0$. Furthermore, Assumption 5.3 holds because every term in this example is analytic. Assumption 5.2 holds because for any vector $v \in T_x(\Gamma_{\|x\|_2})$, we have $\langle x, v \rangle = 0$, and hence $\partial\Phi(x)\sigma(x)v = \partial\Phi(x)\frac{\hat{\sigma}}{\|x\|_2}v = \frac{\hat{\sigma}}{\|x\|_2}v$. Therefore, our main theorem (Theorem 5.5) predicts that the limiting dynamics mix in $\frac{\ln(\lambda/\eta) + O(1)}{4\eta\lambda}$ steps.

Remark 6.1. For ease of demonstration, our main theorem (Theorem 5.5) is proved for SGD+WD with finitely many samples and thus does not directly apply to Equation (18). However, our analysis can be extended to the case of Gaussian noise in an straightforward way and the claim in Theorem 5.5 indeed holds for Equation (18).

Experimental Results. In our experiments, we choose $D = 10$, $\bar{\sigma} = 0.3$, the WD factor $\lambda = 0.05$, and LR $\eta \in \{10^{-2}, 10^{-3}, 10^{-4}\}$. In Figure 2, we plot the projections of $\frac{x}{\|x\|_2}$ on the solution space for 60 different runs with identical initialization for each η . For each run, the only differences are the LR η and/or the noise ξ_t . The dashed lines in the figure indicates our time scaling, i.e., $\frac{1}{2} \ln \left(1 + (e^{4\eta\lambda t} - 1) \frac{\eta}{2\lambda} \right) = T$. Figure 2 shows that time scaling of $O((\ln 2\lambda/\eta + T)/(\eta\lambda))$ fits the dynamics better, compared to $O(T/(\eta\lambda))$.

6.2 Limiting Diffusion on CIFAR-10

Beyond the toy example, we further study the limiting diffusion of PreResNet on CIFAR-10 [26]. We train a 32-layer PreResNet [27] with initial LR $\eta = 0.8$ and WD factor $\lambda = 5 \cdot 10^{-4}$. Unlike the normalized linear model, it is hard to visualize the model projection of PreResNet on the manifold. Instead, we choose the test accuracy of $\Phi(x_t)$ as a test function. In particular, we decay the LR to 10^{-3} to approximate gradient flow at different time t , and record the test accuracy after training 1000 more epochs. The results are shown in Figure 1 and Figure 3.

We observe that without LR decay, the train accuracy, test accuracy and parameter norm converge quickly after training 100 epochs; the test accuracy of $\Phi(x_t)$ converges much slower. It suggests that there exists a mixing process after reaching the manifold. Moreover, we observe the test accuracy of $\Phi(x_t)$ after convergence is significantly higher than $\Phi(x_t)$ at 100th epoch. It indicates that the mixing process is beneficial for generalization.

Our time scaling (12) suggests that the optimal step of LR decay grows no faster than $\tilde{\Omega}(1/\eta\lambda)$ as $\eta \rightarrow 0$. Unfortunately, (12) alone is not sufficient for deciding the optimal step for decaying LR as the mixing time T for the continuous dynamics is unknown. A potential usage of time scaling (12) is to first estimate T via another SGD run with a larger LR, which we leave for future work.

7 Related Work

Normalization and Scale Invariance. Previous works have analyzed the benefits of normalization layers from different viewpoints [28, 13, 29–50]. As noted before, normalization layers induce the scale invariance. It has been shown that scale invariance enables robust and efficient training of SGD+WD [51]. Scale invariance also brings about the interesting equivalence between the effect of WD and LR schedules [52, 19, 17]. Moreover, for SGD+WD with LR η and WD factor λ , the parameter norm will converge to $(\lambda/\eta)^{1/4}$ [1, 53, 54], and this induces the intrinsic LR which is equal to $\eta\lambda$ [1]. These observations are crucial to our derivations in the current paper.

Fast Equilibrium Conjecture. Recently, Wang and Wang [25] proposed a spherical SDE model to approximate SGD+WD with constant LR. Using a novel adaption of Simon’s theory, they justified the Fast Equilibrium Conjecture by showing that SGD+WD dynamics consists of three stages: descent ($O(1/\sqrt{\lambda\eta})$ time), diffusion ($O(1/(\lambda\eta))$ time) and tunneling ($O(e^{C/(\lambda\eta)})$ time). However, their analysis relies on the strong assumption of the minimizers of L being isolated, which is against the empirical evidence that the level sets of deep learning loss functions are connected [55]. As a result, the diffusion phase shall bring no generalization benefit and cannot explain the improvement of final generalization if staying at training loss plateau for a longer time (see Figure 1). We allow the local minimizers to form a connected manifold, which can be viewed a generalization of the Morse function assumption [25], as an isolated minimizer is just a manifold of dimension 0.

Another common weakness of existing analyses in [1, 25] is that they only work for the SDE approximation (3), and do not apply to the actual discrete-time dynamics (2). In contrast, our results can handle both the continuous and discrete time dynamics under more reasonable assumptions.

SDE Approximation. Continuous-time tools such as SDE have been a popular lens for studying optimization algorithms including SGD [56–61]. Many interesting properties of SGD have been discovered through this approach [1, 62–65].

Slow Dynamics of SGD Around Zero Loss Manifold. Recent works [66, 16, 2] show that under the assumption that the minimizers locally connect as a manifold, SGD with label noise with small learning rate will move around the manifold after convergence, towards the direction of smaller trace of Hessian, at a very slow rate of $O(\eta^2)$ per step. Arora et al. [23] show that such slow dynamics on manifold can happen without stochastic gradient noise, if the update rule is non-smooth around the manifold of minimizers and GD enters Edge of Stability regime ([67]). Concretely, they show that normalized GD implicitly penalizes the largest eigenvalue of the Hessian at the rate of $O(\eta^2)$. Additional related works are deferred to Appendix A.

8 Conclusion and Future Work

We provide an SDE-based characterization for the limiting dynamics of SGD+WD for a scale invariant loss as $\eta\lambda \rightarrow 0$ with $\eta = O(\lambda)$ and $\lambda = O(1)$. Under some technical assumptions, we further show that the limiting diffusion converges to a unique stationary distribution. It leaves as future work to relax the technical assumptions. Another interesting and important direction for future work is to understand and characterize the benefit on generalization induced by the limiting diffusion.

Acknowledgements

This work is funded by NSF, ONR, Simons Foundation, DARPA and SRC. ZL is also supported by Microsoft Research PhD Fellowship.

References

- [1] Zhiyuan Li, Kaifeng Lyu, and Sanjeev Arora. Reconciling modern deep learning with traditional optimization analyses: The intrinsic learning rate. *Advances in Neural Information Processing Systems*, 33, 2020.
- [2] Zhiyuan Li, Tianhao Wang, and Sanjeev Arora. What happens after sgd reaches zero loss?—a mathematical framework. In *International Conference on Learning Representations*, 2021.
- [3] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy8gdB9xx>.
- [4] Yuanzhi Li, Colin Wei, and Tengyu Ma. Towards explaining the regularization effect of initial large learning rate in training neural networks. *arXiv preprint arXiv:1907.04595*, 2019.
- [5] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- [6] Dominic Masters and Carlo Luschi. Revisiting small batch training for deep neural networks. *arXiv preprint arXiv:1804.07612*, 2018.
- [7] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [8] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [9] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [10] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29:901–909, 2016.
- [11] Siyuan Qiao, Huiyu Wang, Chenxi Liu, Wei Shen, and Alan Yuille. Micro-batch training with batch-channel normalization and weight standardization. *arXiv preprint arXiv:1903.10520*, 2019.
- [12] Lei Huang, Xianglong Liu, Yang Liu, Bo Lang, and Dacheng Tao. Centered weight normalization in accelerating training of deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2803–2811, 2017.
- [13] Andrew Brock, Soham De, and Samuel L Smith. Characterizing signal propagation to close the performance gap in unnormalized resnets. In *International Conference on Learning Representations*, 2021.
- [14] Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.
- [15] Guy Blanc, Neha Gupta, Gregory Valiant, and Paul Valiant. Implicit regularization for deep neural networks driven by an ornstein-uhlenbeck like process. In *Conference on learning theory*, pages 483–513. PMLR, 2020.
- [16] Alex Damian, Tengyu Ma, and Jason Lee. Label noise sgd provably prefers flat global minimizers. *arXiv preprint arXiv:2106.06530*, 2021.
- [17] Zhiyuan Li and Sanjeev Arora. An exponential learning rate schedule for deep learning. In *International Conference on Learning Representations*, 2019.
- [18] Bin Shi, Weijie J Su, and Michael I Jordan. On learning rates and schrödinger operators. *arXiv preprint arXiv:2004.06977*, 2020.

- [19] Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization. *arXiv preprint arXiv:1810.12281*, 2018.
- [20] Yaim Cooper. The loss landscape of overparameterized neural networks. *arXiv preprint arXiv:1804.10200*, 2018.
- [21] Y Cooper. The critical locus of overparameterized neural networks. *arXiv preprint arXiv:2005.04210*, 2020.
- [22] Benjamin Fehrman, Benjamin Gess, and Arnulf Jentzen. Convergence rates for the stochastic gradient descent method for non-convex objective functions. *Journal of Machine Learning Research*, 21, 2020.
- [23] Sanjeev Arora, Zhiyuan Li, and Abhishek Panigrahi. Understanding gradient descent on the edge of stability in deep learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 948–1024. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/arora22a.html>.
- [24] Gary Shon Katzenberger. Solutions of a stochastic differential equation forced onto a manifold by a large drift. *The Annals of Probability*, pages 1587–1628, 1991.
- [25] Yi Wang and Zhiren Wang. Three-stage evolution and fast equilibrium for sgd with non-degenerate critical points. In *International Conference on Machine Learning*, pages 23092–23113. PMLR, 2022.
- [26] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [28] Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the generalization benefit of normalization layers: Sharpness reduction. *arXiv preprint arXiv:2206.07085*, 2022.
- [29] David Balduzzi, Marcus Frean, Lennox Leary, JP Lewis, Kurt Wan-Duo Ma, and Brian McWilliams. The shattered gradients problem: If resnets are the answer, then what is the question? In *International Conference on Machine Learning*, pages 342–350. PMLR, 2017.
- [30] Soham De and Sam Smith. Batch normalization biases residual blocks towards the identity function in deep networks. *Advances in Neural Information Processing Systems*, 33:19964–19975, 2020.
- [31] Hadi Daneshmand, Jonas Kohler, Francis Bach, Thomas Hofmann, and Aurelien Lucchi. Batch normalization provably avoids ranks collapse for randomly initialised deep networks. *Advances in Neural Information Processing Systems*, 33:18387–18398, 2020.
- [32] Hadi Daneshmand, Amir Joudaki, and Francis Bach. Batch normalization orthogonalizes representations in deep random networks. *Advances in Neural Information Processing Systems*, 34:4896–4906, 2021.
- [33] Ekdeep S Lubana, Robert Dick, and Hidenori Tanaka. Beyond batchnorm: towards a unified understanding of normalization in deep learning. *Advances in Neural Information Processing Systems*, 34:4778–4791, 2021.
- [34] Antoine Labatie, Dominic Masters, Zach Eaton-Rosen, and Carlo Luschi. Proxy-normalizing activations to match batch normalization while removing batch dependence. *Advances in Neural Information Processing Systems*, 34:16990–17006, 2021.
- [35] Ping Luo, Xinjiang Wang, Wenqi Shao, and Zhanglin Peng. Towards understanding regularization in batch normalization. In *International Conference on Learning Representations*, 2019.

- [36] Mattias Teye, Hossein Azizpour, and Kevin Smith. Bayesian uncertainty estimation for batch normalized deep networks. In *International Conference on Machine Learning*, pages 4907–4916. PMLR, 2018.
- [37] Alexander Shekhovtsov and Boris Flach. Stochastic normalizations as bayesian learning. In *Asian Conference on Computer Vision*, pages 463–479. Springer, 2018.
- [38] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? *Advances in neural information processing systems*, 31, 2018.
- [39] Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normalization. *Advances in neural information processing systems*, 31, 2018.
- [40] Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. In *International Conference on Machine Learning*, pages 2232–2241. PMLR, 2019.
- [41] Ryo Karakida, Shotaro Akaho, and Shun-ichi Amari. The normalization method for alleviating pathological sharpness in wide neural networks. *Advances in neural information processing systems*, 32, 2019.
- [42] Zinan Lin, Vyas Sekar, and Giulia Fanti. Why spectral normalization stabilizes gans: Analysis and improvements. *Advances in neural information processing systems*, 34:9625–9638, 2021.
- [43] Susanna Lange, Kyle Helfrich, and Qiang Ye. Batch normalization preconditioning for neural network training. *Journal of Machine Learning Research*, 23(72):1–41, 2022.
- [44] Sanjeev Arora, Zhiyuan Li, and Kaifeng Lyu. Theoretical analysis of auto rate-tuning by batch normalization. In *International Conference on Learning Representations*, 2019.
- [45] Yongqiang Cai, Qianxiao Li, and Zuowei Shen. A quantitative analysis of the effect of batch normalization on gradient descent. In *International Conference on Machine Learning*, pages 882–890. PMLR, 2019.
- [46] Jonas Kohler, Hadi Daneshmand, Aurelien Lucchi, Thomas Hofmann, Ming Zhou, and Klaus Neymeyr. Exponential convergence rates for batch normalization: The power of length-direction decoupling in non-convex optimization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 806–815. PMLR, 2019.
- [47] Xiaoxia Wu, Edgar Dobriban, Tongzheng Ren, Shanshan Wu, Zhiyuan Li, Suriya Gunasekar, Rachel Ward, and Qiang Liu. Implicit regularization and convergence for weight normalization. *Advances in Neural Information Processing Systems*, 33:2835–2847, 2020.
- [48] Chao Ma and Lexing Ying. A riemannian mean field formulation for two-layer neural networks with batch normalization. *Research in the Mathematical Sciences*, 9(3):1–15, 2022.
- [49] Yonatan Dukler, Quanquan Gu, and Guido Montúfar. Optimization theory for relu neural networks trained with normalization layers. In *International conference on machine learning*, pages 2751–2760. PMLR, 2020.
- [50] Shengchao Liu, Dimitris Papailiopoulos, and Dimitris Achlioptas. Bad global minima exist and sgd can reach them. *Advances in Neural Information Processing Systems*, 33:8543–8552, 2020.
- [51] Zhiyuan Li, Srinadh Bhojanapalli, Manzil Zaheer, Sashank Reddi, and Sanjiv Kumar. Robust training of neural networks using scale invariant architectures. In *International Conference on Machine Learning*, pages 12656–12684. PMLR, 2022.
- [52] Elad Hoffer, Ron Banner, Itay Golan, and Daniel Soudry. Norm matters: efficient and accurate normalization schemes in deep networks. *Advances in Neural Information Processing Systems*, 31, 2018.
- [53] Twan Van Laarhoven. L2 regularization versus batch and weight normalization. *arXiv preprint arXiv:1706.05350*, 2017.

- [54] Ruosi Wan, Zhanxing Zhu, Xiangyu Zhang, and Jian Sun. Spherical motion dynamics: Learning dynamics of neural network with normalization, weight decay, and sgd. *arXiv preprint arXiv:2006.08419*, 2020.
- [55] Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *International conference on machine learning*, pages 1309–1318. PMLR, 2018.
- [56] Stephan Mandt, Matthew D Hoffman, and David M Blei. Stochastic gradient descent as approximate bayesian inference. *arXiv preprint arXiv:1704.04289*, 2017.
- [57] Qianxiao Li, Cheng Tai, and E Weinan. Stochastic modified equations and adaptive stochastic gradient algorithms. In *International Conference on Machine Learning*, pages 2101–2110. PMLR, 2017.
- [58] Justin Sirignano and Konstantinos Spiliopoulos. Stochastic gradient descent in continuous time. *SIAM Journal on Financial Mathematics*, 8(1):933–961, 2017.
- [59] Qianxiao Li, Cheng Tai, and E Weinan. Stochastic modified equations and dynamics of stochastic gradient algorithms i: Mathematical foundations. *The Journal of Machine Learning Research*, 20(1):1474–1520, 2019.
- [60] Xiang Cheng, Dong Yin, Peter Bartlett, and Michael Jordan. Stochastic gradient and langevin processes. In *International Conference on Machine Learning*, pages 1810–1819. PMLR, 2020.
- [61] Justin Sirignano and Konstantinos Spiliopoulos. Stochastic gradient descent in continuous time: A central limit theorem. *Stochastic Systems*, 10(2):124–151, 2020.
- [62] Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. *arXiv preprint arXiv:1803.00195*, 2018.
- [63] Zeke Xie, Issei Sato, and Masashi Sugiyama. A diffusion theory for deep learning dynamics: Stochastic gradient descent escapes from sharp minima exponentially fast. *arXiv preprint arXiv:2002.03495*, 2020.
- [64] Scott Pesme, Loucas Pillaud-Vivien, and Nicolas Flammarion. Implicit bias of sgd for diagonal linear networks: a provable benefit of stochasticity. *arXiv preprint arXiv:2106.09524*, 2021.
- [65] Stephan Wojtowysch. Stochastic gradient descent with noise of machine learning type. part ii: Continuous time analysis. *arXiv preprint arXiv:2106.02588*, 2021.
- [66] Guy Blanc, Neha Gupta, Gregory Valiant, and Paul Valiant. Implicit regularization for deep neural networks driven by an ornstein-uhlenbeck like process. *arXiv preprint arXiv:1904.09080*, 2019.
- [67] Jeremy M. Cohen, Simran Kaur, Yanzhi Li, J. Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability, 2021.
- [68] Chao Ma, Lei Wu, and Lexing Ying. The multiscale structure of neural network loss functions: The effect on optimization and origin. *arXiv preprint arXiv:2204.11326*, 2022.
- [69] Kaiyue Wen, Tengyu Ma, and Zhiyuan Li. How does sharpness-aware minimization minimize sharpness? *arXiv preprint arXiv:2211.05729*, 2022.
- [70] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.
- [71] Xinran Gu, Kaifeng Lyu, Longbo Huang, and Sanjeev Arora. Why (and when) does local sgd generalize better than sgd? In *OPT 2022: Optimization for Machine Learning (NeurIPS 2022 Workshop)*.
- [72] Hong Liu, Sang Michael Xie, Zhiyuan Li, and Tengyu Ma. Same pre-training loss, better downstream: Implicit bias matters for language models. *arXiv preprint arXiv:2210.14199*, 2022.

- [73] Wolfgang Kliemann. Recurrence and invariant measures for degenerate diffusions. *The annals of probability*, pages 690–707, 1987.
- [74] Kanji Ichihara and Hiroshi Kunita. A classification of the second order degenerate elliptic operators and its probabilistic characterization. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 30(3):235–254, 1974.
- [75] Martin Hairer. Ergodic theory for stochastic pdes. *preprint*, 2008.
- [76] Giuseppe Da Prato and Jerzy Zabczyk. *Ergodicity for infinite dimensional systems*, volume 229. Cambridge university press, 1996.
- [77] Andrei Agrachev, Davide Barilari, and Ugo Boscain. *A comprehensive introduction to sub-Riemannian geometry*, volume 181. Cambridge University Press, 2019.
- [78] Jean-Michel Coron. *Control and nonlinearity*. Number 136. American Mathematical Soc., 2007.
- [79] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [80] Stewart N Ethier and Thomas G Kurtz. *Markov processes: characterization and convergence*, volume 282. John Wiley & Sons, 2009.
- [81] Ward Whitt. Proofs of the martingale fclt. *Probability Surveys*, 4:268–302, 2007.
- [82] David Pollard. *Convergence of stochastic processes*. Springer Science & Business Media, 2012.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? **[Yes]** See Section X.
- Did you include the license to the code and datasets? **[No]** The code and the data are proprietary.
- Did you include the license to the code and datasets? **[N/A]**

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

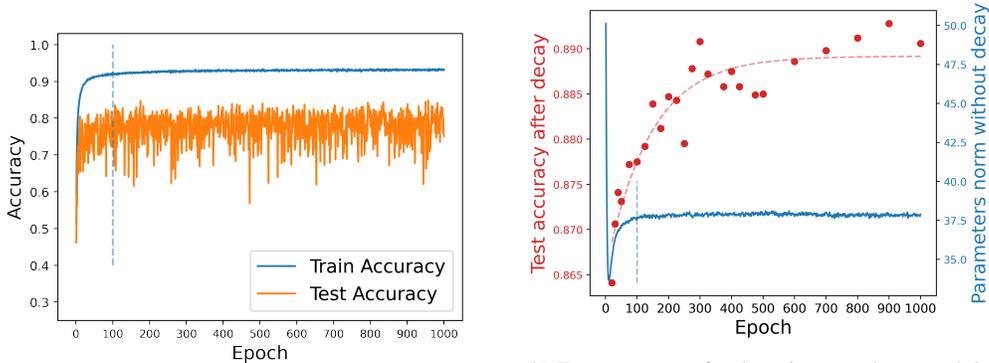
1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]**
 - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]** This paper studies the theoretical properties of stochastic gradient descent with weight decay.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]**
 - (b) Did you include complete proofs of all theoretical results? **[Yes]**
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[No]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]**
 - (b) Did you mention the license of the assets? **[N/A]**
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[No]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

A Additional Related Work

Slow Dynamics of SGD Around Zero Loss Manifold. Ma et al. [68] argues that such flatness driven phenomenon can also be caused by a multi-scale loss landscape. The recent paper Lyu et al. [28] is probably the most related work to us, which shows that GD with weight decay on a scale invariant loss function implicitly penalize the spherical sharpness, *i.e.*, the largest eigenvalue of the Hessian evaluated at the normalized parameter, under the same set of assumptions as ours. The timescale for their sharpness-reduction phenomenon is also $\Theta(1/\eta\lambda)$ steps. The main difference is that the our slow dynamics (limiting diffusion) is caused by the gradient noise while the slow dynamics in [28] is caused by Edge of Stability, thus does not have mixing property on manifold of minimizers. Wen et al. [69] shows that the slow dynamics of sharpness-aware minimization (SAM, proposed by [70]) penalizes some notion of sharpness depending on the batch size. Gu et al. [71] studies the slow dynamics of local SGD around manifold and use it to explain the generalization benefit of local SGD over SGD. Liu et al. [72] studies the slow dynamics of SGD on language models using the results from [2] and empirically observed that different pretraining procedures can result in same pretraining loss but with different downstream performance.

B Additional experimental results

Implementation: We use the implementation of 32-layer-PreResNet from <https://github.com/bearpaw/pytorch-classification>, with slight changes (following Appendix C of [17]) to ensure the model is scale invariant.



(a) Train and test accuracy for CIFAR-10 training with $\eta = 0.8$, $\lambda = 0.0005$.

(b) Test accuracy after learning rate decay and the total norm of parameters before learning rate decay.

Figure 3: PreResNet trained without data augmentation on CIFAR-10. The train and test accuracy plateaus after parameter norm convergence within 100 epochs, but the generalization of SGD iterate after LR decay keeps improves. Figure 3a shows the train and test accuracy of scale invariant PreResNet trained with standard data augmentation on CIFAR-10. Each red dot in Figure 3b represents the test accuracy of model which decays LR to 10^{-3} at the corresponding epoch. The test accuracy is evaluated at full training accuracy after LR decay.

C Preliminaries

We refer the reader to Appendix A in Li et al. [2] for preliminaries on stochastic processes that are useful for our results.

C.1 Omitted proofs for properties induced by scale-invariance

Proof of Lemma 2.7. By the definition of U , for any $x \in U$, we have $\Phi(x) = \lim_{t \rightarrow \infty} \phi(x, t) \in \Gamma$. Define $x(t) = \phi(x, t)$, and for any $\alpha > 0$, let $\tilde{x}(t) = \alpha \cdot x(t/\alpha^2)$, then we have $\tilde{x}(0) = \alpha x$ and $\frac{d\tilde{x}(t)}{dt} = -\alpha^{-1} \nabla L(x(t/\alpha^2)) = -\nabla L(\tilde{x}(t))$ where the second equality follows from the scale invariance of L . Therefore, by the definition of $\phi(x, t)$, we see that $\phi(\alpha x, t) = \tilde{x}(t) = \alpha \cdot x(t/\alpha^2) =$

$\alpha \cdot \phi(x, t/\alpha^2)$. Let $t \rightarrow \infty$, and we get $\Phi(\alpha x) = \alpha \cdot \Phi(x)$. Since Γ itself is a cone, it follows that $\Phi(\alpha x) \in \Gamma$, and thus $\alpha x \in U$. Hence, we conclude that U is also a cone. \square

D Omitted Proofs for the simultaneous limit case in Section 3

Now we present the proofs for the results in Section 3 where $\lambda, \eta \rightarrow 0$ simultaneously, that is, $\frac{\eta}{\lambda} \equiv C$ where C is some positive constant. This case can be regarded as a direct application of the framework in Li et al. [2].

Proof of Lemma 3.1. The proof can be directly adapted from that of Lemma 4.2 in Li et al. [2]. \square

Proof of Theorem 3.2. The first claim is a direct consequence of Lemma 3.1 and Theorem 2.10.

For the second claim, we first show that if Y_0 exists, then for any $C > 0$, Y_C also exists and they are equal in distribution after the desired rescaling. By Lemma G.1, there is some Brownian motion $W'(t)$ such that $W\left(\frac{e^{4Ct}-1}{4C}\right) = \int_{\tau=0}^t e^{2C\tau} dW'(\tau)$, $\forall t \geq 0$ almost surely. For convenience, we denote $Y_0\left(\frac{e^{4Ct}-1}{4C}\right)e^{-Ct}$ by $\tilde{Y}_C(t)$. Thus we have

$$\begin{aligned} & d\tilde{Y}_C(t) + C\tilde{Y}_C(t)dt \\ &= e^{-Ct} \left(\frac{1}{2} \partial^2 \Phi(\tilde{Y}_C e^{Ct}) [\Sigma(\tilde{Y}_C e^{Ct})] d\left(\frac{e^{4Ct}-1}{4C}\right) - \partial \Phi(\tilde{Y}_C e^{Ct}) \sigma(\tilde{Y}_C e^{Ct}) dW\left(\frac{e^{4Ct}-1}{4C}\right) \right) \\ &= e^{3Ct} \frac{1}{2} \partial^2 \Phi(\tilde{Y}_C e^{Ct}) [\Sigma(\tilde{Y}_C e^{Ct})] dt - e^{Ct} \partial \Phi(\tilde{Y}_C e^{Ct}) \sigma(\tilde{Y}_C e^{Ct}) dW'(t) \\ &= \frac{1}{2} \partial^2 \Phi(\tilde{Y}_C) [\Sigma(\tilde{Y}_C)] dt - \partial \Phi(\tilde{Y}_C) \sigma(\tilde{Y}_C) dW'(t), \end{aligned}$$

where the last equality uses the fact that $\partial^k \Phi$ is $(1-k)$ -homogeneous, σ is (-1) -homogeneous and Σ is (-2) -homogeneous by Lemma 2.4. Thus we construct a global solution for (8) with hyperparameter C . Since U is a cone and $Y_0(t) \in U$ for all $t \geq 0$, it also holds that $\tilde{Y}_C(t) \in U$.

By applying the above argument in the other direction, we can show that if Y_C exists then Y_0 exists and they are equivalent up to some rescaling. This finishes the proof. \square

E Omitted proofs for general case in Section 4

Here we provide the proofs for the limiting diffusion for continuous-time SDE and discrete SGD+WD respectively, in a general case of $\eta\lambda \rightarrow 0$ where $\eta \leq 2\lambda \leq c$ and c is any constant. Note the analysis here can be easily generalized to the case $\eta\lambda \rightarrow 0$ where $\eta = \mathcal{O}(\lambda) = \mathcal{O}(1)$, that is, exists constants C_1, C_2 , such that $\eta \leq C_1\lambda \leq C_2$. Such a limit is more interesting than the simultaneous limit as the it is more common to tune LR without fixing the ratio between WD and LR.

Below we will show how these rescalings work first in the continuous case in Appendix E.1, and then generalize the same proof idea to the discrete case, SGD+WD, in Appendix E.2.

E.1 Proof for the limiting diffusion of SDE

Proof of Theorem 4.1. Recall the scaling function $R_{\eta,\lambda}(t)$ defined in Equation (11):

$$R_{\eta,\lambda}(t) = \left(\frac{\eta}{2\lambda} + e^{-4\eta\lambda t} \left(1 - \frac{\eta}{2\lambda} \right) \right)^{1/4}.$$

which is exactly Equation (10) except $\|X(0)\|_2$ is replaced with 1. A simple calculation gives

$$d \ln R_{\eta,\lambda}(t) = -\frac{\eta\lambda e^{-4\eta\lambda t} (1 - \eta/(2\lambda))}{\eta/(2\lambda) + e^{-4\eta\lambda t} (1 - \eta/(2\lambda))}. \quad (20)$$

Now, define $\widehat{X}_{\eta,\lambda}(t) = X_{\eta,\lambda}(t)/R_{\eta,\lambda}(t)$, a normalized version of $X_{\eta,\lambda}(t)$. Its dynamics is

$$\begin{aligned} d\widehat{X}_{\eta,\lambda}(t) &= -\frac{\eta}{R_{\eta,\lambda}(t)}\nabla L(X_{\eta,\lambda}(t))dt - \frac{\eta}{R_{\eta,\lambda}(t)}\sigma(X_{\eta,\lambda}(t))dW(t) \\ &\quad - \frac{\eta\lambda X_{\eta,\lambda}(t)}{R_{\eta,\lambda}(t)}dt - \frac{X_{\eta,\lambda}(t)}{R_{\eta,\lambda}(t)}d\ln R_{\eta,\lambda}(t) \\ &= -\frac{\eta}{R_{\eta,\lambda}(t)^2}\nabla L(\widehat{X}_{\eta,\lambda}(t))dt - \frac{\eta}{R_{\eta,\lambda}(t)^2}\sigma(\widehat{X}_{\eta,\lambda}(t))dW(t) \\ &\quad - \widehat{X}_{\eta,\lambda}(t)(\eta\lambda dt + d\ln R_{\eta,\lambda}(t)). \end{aligned} \quad (21)$$

Next, to rescale the time, recall the function $\tau_{\eta,\lambda}$ defined in Equation (12):

$$\tau_{\eta,\lambda}(t) = \int_{s=0}^t \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds,$$

and consider $\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t)) = \widehat{X}_{\eta,\lambda}(t)$. Note that $dR_{\eta,\lambda}(t)^4 = -4\eta\lambda R_{\eta,\lambda}(t)^4 + 2\eta^2 dt$, so we have

$$d\ln R_{\eta,\lambda}(t)^4 = \frac{dR_{\eta,\lambda}(t)^4}{R_{\eta,\lambda}(t)^4} = -4\eta\lambda dt + 2\frac{\eta^2}{R_{\eta,\lambda}(t)^4} dt, \quad (22)$$

which implies that

$$\tau_{\eta,\lambda}(t) = \frac{1}{2} \int_{s=0}^t 4\eta\lambda ds + d\ln R_{\eta,\lambda}(s)^4 = 2\eta\lambda t + 2\ln R_{\eta,\lambda}(t) = \frac{1}{2} \ln \left(1 + (e^{4\eta\lambda t} - 1) \frac{\eta}{2\lambda} \right). \quad (23)$$

Based on the above closed-form expression of $\tau_{\eta,\lambda}(t)$, it is straightforward to verify that

$$\tau_{\eta,\lambda}^{-1}(T) = \frac{\ln \left(\frac{2\lambda}{\eta} (e^{2T} - 1) + 1 \right)}{4\eta\lambda}. \quad (24)$$

Applying Lemma G.1, we have

$$\begin{aligned} d\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t)) &= -\frac{R_{\eta,\lambda}(t)^2}{\eta}\nabla L(\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t)))d\tau_{\eta,\lambda}(t) - \sigma(\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t)))dW(\tau_{\eta,\lambda}(t)) \\ &\quad - \widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t))(\eta\lambda dt + d\ln R_{\eta,\lambda}(t)). \end{aligned} \quad (25)$$

By the property of $R_{\eta,\lambda}(t)$ in Equation (20) and the definition of $\tau_{\eta,\lambda}(t)$ in Equation (12), we have

$$\eta\lambda dt + d\ln R_{\eta,\lambda}(t) = \frac{\eta^2/2}{\eta/(2\lambda) + e^{-4\eta\lambda t}(1 - \eta/(2\lambda))} dt = \frac{\eta^2}{2R_{\eta,\lambda}(t)^4} dt = \frac{1}{2} d\tau_{\eta,\lambda}(t)$$

where the second equality follows from the definition of $R_{\eta,\lambda}(t)$ in Equation (11). Then combining the above identity and Equation (25) yields

$$\begin{aligned} d\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t)) &= -\frac{R_{\eta,\lambda}(t)^2}{\eta}\nabla L(\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t)))d\tau_{\eta,\lambda}(t) - \frac{\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t))}{2}d\tau_{\eta,\lambda}(t) \\ &\quad - \sigma(\widetilde{X}_{\eta,\lambda}(\tau_{\eta,\lambda}(t)))dW(\tau_{\eta,\lambda}(t)). \end{aligned}$$

Note that $\tau_{\eta,\lambda} : [0, \infty) \rightarrow [0, \infty)$ is a bijection, so we can rewrite the above in $\tilde{t} := \tau_{\eta,\lambda}(t)$ to get

$$d\widetilde{X}_{\eta,\lambda}(\tilde{t}) = -\frac{R(\tau_{\eta,\lambda}^{-1}(\tilde{t}))^2}{\eta}\nabla L(\widetilde{X}_{\eta,\lambda}(\tilde{t}))d\tilde{t} + \sigma(\widetilde{X}_{\eta,\lambda}(\tilde{t}))dW(\tilde{t}) - \frac{\widetilde{X}_{\eta,\lambda}(\tilde{t})}{2}d\tilde{t} \quad (26)$$

where by writing $\tau_{\eta,\lambda}^{-1}(\cdot)$ we mean the functional inverse of $\tau_{\eta,\lambda}(\cdot)$ defined in Equation (12).

By assumption $\eta < 2\lambda$, can see from Equation (11) that $R_{\eta,\lambda}(t)^2$ is decreasing on $[0, \infty)$ and is bounded from below by $\sqrt{\eta/(2\lambda)}$. Therefore, for any $t \geq 0$, $R_{\eta,\lambda}(t)^2/\eta \geq 1/\sqrt{2\lambda\eta}$, so we can again apply Theorem 2.10 to get the desired limiting diffusion. This finishes the proof. \square

E.2 Proof for the limiting diffusion of SGD+WD

Recall the update of SGD+WD in Equation (2) and the scaling function $R_{\eta,\lambda}$ in Equation (11), and the normalized version $x_{\eta,\lambda}(k)/R_{\eta,\lambda}(k)$ satisfies the following:

$$\begin{aligned} \frac{x_{\eta,\lambda}(k+1)}{R_{\eta,\lambda}(k+1)} &= (1-\eta\lambda) \frac{x_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k+1)} - \frac{\eta}{R_{\eta,\lambda}(k+1)} (\nabla L(x_{\eta,\lambda}(k)) + \sqrt{\Xi} \sigma_{\xi_k}(x_{\eta,\lambda}(k))) \\ &= \frac{x_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k)} - \frac{x_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k)} \cdot \frac{\eta\lambda R_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k+1)} + \left(\frac{R_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k+1)} - 1 \right) \frac{x_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k)} \\ &\quad - \frac{\eta}{R_{\eta,\lambda}(k)R_{\eta,\lambda}(k+1)} \left(\nabla L \left(\frac{x_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k)} \right) + \sqrt{\Xi} \sigma_{\xi_k} \left(\frac{x_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k)} \right) \right). \end{aligned}$$

Then $\hat{x}_{\eta,\lambda}(k) = x_{\eta,\lambda}(k)/R_{\eta,\lambda}(k)$ satisfies

$$\begin{aligned} \hat{x}_{\eta,\lambda}(k+1) &= \hat{x}_{\eta,\lambda}(k) - \frac{\eta\lambda R_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k+1)} \hat{x}_{\eta,\lambda}(k) + \left(\frac{R_{\eta,\lambda}(k)}{R_{\eta,\lambda}(k+1)} - 1 \right) \hat{x}_{\eta,\lambda}(k) \\ &\quad - \frac{\eta}{R_{\eta,\lambda}(k)R_{\eta,\lambda}(k+1)} (\nabla L(\hat{x}_{\eta,\lambda}(k)) + \sqrt{\Xi} \sigma_{\xi_k}(\hat{x}_{\eta,\lambda}(k))) \end{aligned}$$

Define $\hat{X}_{\eta,\lambda}(t) = \hat{x}_{\eta,\lambda}(\lfloor t \rfloor)$, and rewriting the above equation yields

$$\begin{aligned} \hat{X}_{\eta,\lambda}(t) &= \int_{s=0}^t \left[-\eta\lambda \hat{X}_{\eta,\lambda}(s) + (1-\eta\lambda) \left(\frac{R_{\eta,\lambda}(s)}{R_{\eta,\lambda}(s+1)} - 1 \right) \hat{X}_{\eta,\lambda}(s) \right] d[s] \\ &\quad - \int_{s=0}^t \frac{\eta}{R_{\eta,\lambda}(s)R_{\eta,\lambda}(s+1)} \nabla L(\hat{X}_{\eta,\lambda}(s)) d[s] \\ &\quad - \int_{s=0}^t \frac{\eta\sqrt{\Xi}}{R_{\eta,\lambda}(s)R_{\eta,\lambda}(s+1)} \sigma(\hat{X}_{\eta,\lambda}(s)) d \sum_{i=1}^{\lfloor s \rfloor} \left(\mathbf{1}_{\xi_i} - \frac{1}{\Xi} \mathbf{1} \right). \end{aligned}$$

Recall that $\tilde{X}_{\eta,\lambda}(\tilde{t}) = \hat{x}_{\eta,\lambda}(\lfloor \tilde{t} \rfloor) = \hat{X}_{\eta,\lambda}(t)$, thus we obtain the expression of $\tilde{X}_{\eta,\lambda}(\tilde{t})$ in (14).

For clarity, we break Lemma 4.3 into the following series of lemmas and prove them respectively.

Lemma E.1. *Let $\{A_{\eta,\lambda}\}_{\eta,\lambda>0}$ be defined as in Equation (15). Then $A_{\eta,\lambda}$ increases infinitely fast as $\eta\lambda \rightarrow 0$ with $\eta < 2\lambda < c$ for some constant c , that is, for any $\epsilon > 0$,*

$$\liminf_{\eta \rightarrow 0} \inf_{\tilde{t} \geq 0} (A_{\eta,\lambda}(\tilde{t} + \epsilon) - A_{\eta,\lambda}(\tilde{t})) = \infty.$$

Proof of Lemma E.1. Fix any $\epsilon > 0$ and $\tilde{t} \geq 0$, and let η be sufficiently small. By definition, we have

$$\begin{aligned} A_{\eta,\lambda}(\tilde{t} + \epsilon) - A_{\eta,\lambda}(\tilde{t}) &= \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon) \rfloor} \frac{\eta}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(i+1)} - \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(i+1)} \\ &= \sum_{i=\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor + 1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon) \rfloor} \frac{\eta}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(i+1)} \\ &\geq \sum_{i=\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon) \rfloor} \frac{\eta}{R_{\eta,\lambda}(i)^2} \geq \int_{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon) \rfloor - 1} \frac{\eta}{R_{\eta,\lambda}(s)^2} ds \end{aligned}$$

where the two inequalities follow from the fact that $R_{\eta,\lambda}(t)$ is monotonically decreasing. Moreover, since $R_{\eta,\lambda}(t) \in ((\frac{\eta}{2\lambda})^{1/4}, 1]$, it holds that $\eta/R_{\eta,\lambda}(t)^2 \in [\eta, \sqrt{2\eta\lambda})$, for all $t \geq 0$. We then have

$$\begin{aligned} A_{\eta,\lambda}(\tilde{t} + \epsilon) - A_{\eta,\lambda}(\tilde{t}) &\geq \int_{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon) \rfloor - 1} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} \cdot \frac{R_{\eta,\lambda}(s)^2}{\eta} ds \\ &\geq \frac{1}{\sqrt{2\eta\lambda}} \int_{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon) \rfloor - 1} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds \\ &\geq \frac{1}{\sqrt{2\eta\lambda}} \int_{\tau_{\eta,\lambda}^{-1}(\tilde{t})}^{\tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon)} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds - 3\sqrt{2\eta\lambda}. \end{aligned}$$

Recall the definition of $\tau_{\eta,\lambda}(t)$ in Equation (12), we further have

$$A_{\eta,\lambda}(\tilde{t} + \epsilon) - A_{\eta,\lambda}(\tilde{t}) \geq \frac{\tau_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(\tilde{t} + \epsilon)) - \tau_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(\tilde{t}))}{\sqrt{2\eta\lambda}} - 3\sqrt{2\eta\lambda} = \frac{\epsilon}{\sqrt{2\eta\lambda}} - 3\sqrt{2\eta\lambda},$$

which holds for all $\tilde{t} \geq 0$ simultaneously. Hence, we obtain the desired result:

$$\liminf_{\eta \rightarrow 0} \liminf_{\tilde{t} \geq 0} A_{\eta,\lambda}(\tilde{t} + \epsilon) - A_{\eta,\lambda}(\tilde{t}) \geq \lim_{\eta \rightarrow 0} \frac{\epsilon}{\sqrt{2\eta\lambda}} - 3\sqrt{2\eta\lambda} = \infty.$$

□

Lemma E.2. Let $\{B_{\eta,\lambda}\}_{\eta>0}$ be defined as in Equation (16). Then as $\eta\lambda \rightarrow 0$ with $\eta < 2\lambda < c$ for some constant c , $B_{\eta,\lambda}(\tilde{t})$ converges to $\tilde{t}/2$ in uniform metric.

Proof of Lemma E.2. Since $\tilde{t} = \tau_{\eta,\lambda}(t) = \int_0^t \tau'_{\eta,\lambda}(s) ds$, the idea is to show that the above sum can be seen as an approximation of the Riemann sum, which then yields the approximation of the integral. To see this, recall the definition of $R_{\eta,\lambda}(t)$ in Equation (11), and note that there exists some $t_i \in [i, i+1]$ such that

$$\frac{R_{\eta,\lambda}(i)}{R_{\eta,\lambda}(i+1)} - 1 = \frac{-R'_{\eta,\lambda}(t_i)}{R_{\eta,\lambda}(i+1)} = \frac{\eta\lambda e^{-4\eta\lambda t_i} (1 - \frac{\eta}{2\lambda})}{R_{\eta,\lambda}(i+1) R_{\eta,\lambda}(t_i)^3}. \quad (27)$$

Then since $R_{\eta,\lambda}(t)$ is positive and monotonically decreasing, it follows that

$$\begin{aligned} B_{\eta,\lambda}(\tilde{t}) &= \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \eta\lambda - (1 - \eta\lambda) \frac{\eta\lambda e^{-4\eta\lambda t_i} (1 - \frac{\eta}{2\lambda})}{R_{\eta,\lambda}(i+1) R_{\eta,\lambda}(t_i)^3} \\ &\leq \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \eta\lambda - (1 - \eta\lambda) \frac{\eta\lambda e^{-4\eta\lambda t_i} (1 - \frac{\eta}{2\lambda})}{R_{\eta,\lambda}(t_i)^4} \\ &= \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2/2 + \eta^2\lambda^2 e^{-4\eta\lambda t_i} (1 - \frac{\eta}{2\lambda})}{R_{\eta,\lambda}(t_i)^4} \\ &\leq \frac{1}{2} \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4} + \left(1 - \frac{\eta}{2\lambda}\right) \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2\lambda^2 e^{-4\eta\lambda i}}{R_{\eta,\lambda}(i+1)^4} \end{aligned}$$

where the last inequality again follows from the monotonicity of $R(t)$.

Note that when $i \leq \frac{1}{8\eta\lambda} \ln \frac{1}{\eta}$, we have $R_{\eta,\lambda}(i)^4 \geq \sqrt{\eta}(1 - \frac{\eta}{2\lambda})$, and when $i > \frac{1}{8\eta\lambda} \ln \frac{1}{\eta}$, we have $e^{-4\eta\lambda i} \leq \sqrt{\eta}$. Thus we further have

$$\begin{aligned}
B_{\eta,\lambda}(\tilde{t}) &\leq \frac{1}{2} \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4} + \left(1 - \frac{\eta}{2\lambda}\right) \sum_{i=1}^{\lfloor \frac{1}{8\eta\lambda} \ln \frac{1}{\eta} \rfloor - 1} \frac{\eta^2 \lambda^2 e^{-4\eta\lambda i}}{R_{\eta,\lambda}(i+1)^4} \\
&\quad + \left(1 - \frac{\eta}{2\lambda}\right) \sum_{i=\lfloor \frac{1}{8\eta\lambda} \ln \frac{1}{\eta} \rfloor}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2 \lambda^2 e^{-4\eta\lambda i}}{R_{\eta,\lambda}(i+1)^4} \\
&\leq \frac{1}{2} \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4} + \sum_{i=1}^{\lfloor \frac{1}{8\eta\lambda} \ln \frac{1}{\eta} \rfloor - 1} \eta^{3/2} \lambda^2 e^{-4\eta\lambda i} \\
&\quad + \sqrt{\eta} \left(\lambda^2 - \frac{\eta\lambda}{2} \right) \sum_{i=\lfloor \frac{1}{8\eta\lambda} \ln \frac{1}{\eta} \rfloor}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4} \\
&\leq \left(\frac{1}{2} + \sqrt{\eta} \lambda^2 \right) \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4} + \frac{1}{8} \sqrt{\eta} \lambda \ln \frac{1}{\eta}.
\end{aligned}$$

Then since $R_{\eta,\lambda}(t)$ is monotonically decreasing, it follows that

$$\begin{aligned}
B_{\eta,\lambda}(\tilde{t}) &\leq \left(\frac{1}{2} + \sqrt{\eta} \lambda^2 \right) \int_3^{\tau_{\eta,\lambda}^{-1}(\tilde{t})+2} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds + \frac{1}{8} \sqrt{\eta} \lambda \ln \frac{1}{\eta} \\
&\leq \left(\frac{1}{2} + \sqrt{\eta} \lambda^2 \right) \tau_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(\tilde{t})) + 2\eta\lambda + 4\eta^{3/2} \lambda^3 + \frac{1}{8} \sqrt{\eta} \lambda \ln \frac{1}{\eta} \\
&= \left(\frac{1}{2} + \sqrt{\eta} \lambda^2 \right) \tilde{t} + 2\eta\lambda + 4\eta^{3/2} \lambda^3 + \frac{1}{8} \sqrt{\eta} \lambda \ln \frac{1}{\eta}
\end{aligned} \tag{28}$$

where the second inequality follows from the fact that $\frac{\eta^2}{R_{\eta,\lambda}(t)^4} \in [\eta^2, 2\eta\lambda]$ for all $t \geq 0$.

We can similarly establish a lower bound for $B_{\eta,\lambda}(\tilde{t})$. It follows from Equation (27) that

$$\begin{aligned}
B_{\eta,\lambda}(\tilde{t}) &\geq \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \left(\eta\lambda - (1 - \eta\lambda) \frac{\eta\lambda e^{-4\eta\lambda t_i} (1 - \frac{\eta}{2\lambda})}{R_{\eta,\lambda}(i+1)^4} \right) \\
&= \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2/2 + [\eta\lambda e^{-4\eta\lambda(i+1)} - (1 - \eta\lambda)\eta\lambda e^{-4\eta\lambda t_i}] (1 - \frac{\eta}{2\lambda})}{R_{\eta,\lambda}(i+1)^4} \\
&\geq \frac{1}{2} \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4} + \left(1 - \frac{\eta}{2\lambda}\right) \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta\lambda(e^{-4\eta\lambda} - 1)e^{-4\eta\lambda i}}{R_{\eta,\lambda}(i+1)^4} \\
&\geq \frac{1}{2} \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4} - \left(4 - \frac{2\eta}{\lambda}\right) \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2 \lambda^2 e^{-4\eta\lambda i}}{R_{\eta,\lambda}(i+1)^4}
\end{aligned}$$

where the last inequality is due to the fact that $1 - e^{-x} \leq x$ for $x \geq 0$. Then applying the similar argument as we have used for the upper bound, we get

$$B_{\eta,\lambda}(\tilde{t}) \geq \left(\frac{1}{2} - 4\sqrt{\eta} \lambda^2 \right) \tilde{t} - C\sqrt{\eta} \lambda \ln \frac{1}{\eta} \tag{29}$$

for some universal constant $C > 0$.

Now, combining the upper bound in Equation (28) and the lower bound in Equation (29), we obtain

$$\left| B_{\eta,\lambda}(\tilde{t}) - \frac{\tilde{t}}{2} \right| \leq 4\sqrt{\eta} \lambda^2 \tilde{t} + C\sqrt{\eta} \lambda \ln \frac{1}{\eta}$$

for all $\tilde{t} \geq 0$, where $C > 0$ is some universal constant. Recall the definition of the uniform metric, and we then have

$$\begin{aligned} d_U(B_{\eta,\lambda}(\tilde{t}), \tilde{t}/2) &\leq \sum_{T=1}^{\infty} 2^{-T} \min \left\{ 1, \sup_{i \in [0, T)} \left| B_{\eta,\lambda}(\tilde{t}) - \frac{\tilde{t}}{2} \right| \right\} \\ &\leq \sum_{T=1}^{\infty} 2^{-T} \min \left\{ 1, 4\sqrt{\eta}\lambda^2 T + C\sqrt{\eta} \ln \frac{1}{\eta} \right\} \end{aligned}$$

where the right-hand-side converges to 0 by a standard ϵ - δ argument as $\eta \rightarrow 0$. This completes the proof. \square

Lemma E.3. *Let $R_{\eta,\lambda}(t)$ be defined in Equation (11), and $\tau_{\eta,\lambda}(t)$ in Equation (12). Then it holds that*

$$\lim_{\eta \rightarrow 0, \eta\lambda \rightarrow 0} \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i)^2 R_{\eta,\lambda}(i+1)^2} = \tilde{t}$$

uniformly for all $\tilde{t} \geq 0$.

Proof of Lemma E.3. First, since $R_{\eta,\lambda}(t)$ is positive and monotonically decreasing, we have

$$\sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i)^4} \leq \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i)^2 R_{\eta,\lambda}(i+1)^2} \leq \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i+1)^4}$$

which further implies that

$$\int_1^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds \leq \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i)^2 R_{\eta,\lambda}(i+1)^2} \leq \int_2^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor + 2} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds.$$

Then since $\eta^2/R_{\eta,\lambda}(t)^4 \leq \max\{2\eta\lambda, \eta^2\}$ for all $t \geq 0$, we have

$$\begin{aligned} \int_0^{\tau_{\eta,\lambda}^{-1}(\tilde{t})} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds - 2 \max\{2\eta\lambda, \eta^2\} &\leq \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i)^2 R_{\eta,\lambda}(i+1)^2} \\ &\leq \int_0^{\tau_{\eta,\lambda}^{-1}(\tilde{t})} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds + 2 \max\{2\eta\lambda, \eta^2\}. \end{aligned}$$

Recall the definition of $\tau_{\eta,\lambda}(t)$ in Equation (12), so $\int_0^{\tau_{\eta,\lambda}^{-1}(\tilde{t})} \frac{\eta^2}{R_{\eta,\lambda}(s)^4} ds = \tau_{\eta,\lambda}(\tau_{\eta,\lambda}^{-1}(\tilde{t})) = \tilde{t}$. Therefore, we get

$$\left| \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i)^2 R_{\eta,\lambda}(i+1)^2} - \tilde{t} \right| \leq 2 \max\{2\eta\lambda, \eta^2\}$$

for all $\tilde{t} \geq 0$. This implies the uniform convergence for all $\tilde{t} \geq 0$ and thus completes the proof. \square

Lemma E.4. *Let $\{Z_{\eta,\lambda}\}_{\eta>0}$ be defined as in Equation (17). Then as $\eta\lambda \rightarrow 0$ with $\eta < 2\lambda < c$ for some constant c , $Z_{\eta,\lambda}$ weakly converges to $(I_{\Xi} - \frac{1}{\Xi} \mathbb{1} \mathbb{1}^{\top})^{1/2} W$ in the uniform metric, where W is a Ξ -dimensional standard Brownian motion.*

Proof of Lemma E.4. For any $\eta > 0$, note that $Z_{\eta,\lambda}$ is a stochastic process with independent increments, so $Z_{\eta,\lambda}$ is indeed a Ξ -dimensional martingale. Thus we can apply the multidimensional martingale functional central limit theorem (Theorem G.2) to get the desired convergence.

To do so, we need to verify two facts: 1) the expected value of the maximum jump in $Z_{\eta,\lambda}$ is asymptotically negligible; 2) for each fixed $\tilde{t} \geq 0$, $\lim_{\eta \rightarrow 0} \mathbb{E}[Z_{\eta,\lambda}(\tilde{t}) Z_{\eta,\lambda}(\tilde{t})^{\top}] = \tilde{t} (I_{\Xi} - \frac{1}{\Xi} \mathbb{1} \mathbb{1}^{\top})$. These two facts correspond to conditions (55) and (56) respectively.

For the first fact, note that the norm of the i -th jump in $Z_{\eta,\lambda}$ is bounded by $\frac{\eta\sqrt{\Xi}}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(i+1)}\|\mathbf{1}_{\xi_i} - \frac{1}{\Xi}\mathbf{1}\|_2$. Since $R_{\eta,\lambda}(t) \geq \min\{(\frac{\eta}{2\lambda})^{1/4}, 1\}$ for all $t \geq 0$, we see that $\frac{\eta\sqrt{\Xi}}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(i+1)}\|\mathbf{1}_{\xi_i} - \frac{1}{\Xi}\mathbf{1}\|_2 \leq (1+\sqrt{\Xi})\max\{\sqrt{2\eta\lambda}, \eta\}$, which implies that any jump in $Z_{\eta,\lambda}$ is asymptotically negligible as $\eta \rightarrow 0$ and $\eta\lambda \rightarrow 0$.

For the second fact, fixing any $\tilde{t} \geq 0$, we have

$$\begin{aligned}\mathbb{E}[Z_{\eta,\lambda}(\tilde{t})Z_{\eta,\lambda}(\tilde{t})^\top] &= \sum_{i,j=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2\Xi}{R_{\eta,\lambda}(i)R_{\eta,\lambda}(j)R_{\eta,\lambda}(i+1)R_{\eta,\lambda}(j+1)} \mathbb{E}\left[\left(\mathbf{1}_{\xi_i} - \frac{1}{\Xi}\mathbf{1}\right)\left(\mathbf{1}_{\xi_j} - \frac{1}{\Xi}\mathbf{1}\right)^\top\right] \\ &= \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2\Xi}{R_{\eta,\lambda}(i)^2R_{\eta,\lambda}(i+1)^2} \mathbb{E}\left[\left(\mathbf{1}_{\xi_i} - \frac{1}{\Xi}\mathbf{1}\right)\left(\mathbf{1}_{\xi_i} - \frac{1}{\Xi}\mathbf{1}\right)^\top\right] \\ &= \sum_{i=1}^{\lfloor \tau_{\eta,\lambda}^{-1}(\tilde{t}) \rfloor} \frac{\eta^2}{R_{\eta,\lambda}(i)^2R_{\eta,\lambda}(i+1)^2} \left(I_\Xi - \frac{1}{\Xi}\mathbf{1}\mathbf{1}^\top\right)\end{aligned}$$

where the second equality follows from the independence between ξ_i and ξ_j for $i \neq j$. Now, applying Lemma E.3 yields $\lim_{\eta \rightarrow 0} \mathbb{E}[Z_{\eta,\lambda}(\tilde{t})Z_{\eta,\lambda}(\tilde{t})^\top] = \tilde{t}(I_\Xi - \frac{1}{\Xi}\mathbf{1}\mathbf{1}^\top)$.

Combining the above two facts, applying Theorem G.2 and ?? completes the proof. \square

F Omitted Proofs in in Section 5

Here we provide the proof of the fast mixing of SGD+WD. By Theorem 4.4, we know that the iterates of SGD+WD track the solution to the limiting diffusion in Equation (13) at time T after roughly $\frac{1}{4\eta\lambda} \ln(\frac{2\lambda}{\eta}(e^{2T}-1)+1)$ steps. Here T is the time index for the limiting diffusion, which can be shown to mix to the equilibrium and the mixing speed is independent of η or λ since Equation (13) does not rely on these parameters. Therefore, it suffices to establish the ergodicity of the limiting diffusion, and the mixing of SGD+WD immediately follows. Then the total number of steps of SGD+WD is on the order of $O(\frac{1}{\eta\lambda} \ln \frac{\lambda}{\eta})$, yielding the fast mixing of SGD+WD.

F.1 Preliminary on the ergodic theory for SDEs

We first briefly review some preliminaries on the ergodic theory for diffusion processes. Our main references are from [73–75].

Let Γ be a M -dimensional smooth submanifold of \mathbb{R}^D . Recall that in the main context we specify Γ to be the manifold of some scale invariant loss, here we consider the general case with slight abuse of notation. Consider the following diffusion process on Γ :

$$dX(t) = \hat{f}_0(X(t))dt + \sum_{i=1}^{\Xi} f_i(X(t))dW_i(t) \quad (30)$$

where $\hat{f}_0, f_1, \dots, f_\Xi$ are smooth vector fields on Γ and $W(t) = (W_1(t), \dots, W_\Xi(t))^\top$ is a Ξ -dimensional Brownian motion. Writing the above diffusion process in the Stratonovich form, it is equivalent to

$$dX(t) = \underbrace{\left(\hat{f}_0(X(t)) - \frac{1}{2} \sum_{i=1}^{\Xi} \partial f_i(X(t))f_i(X(t))\right)}_{f_0(X(t))} dt + \sum_{i=1}^{\Xi} f_i(X(t)) \circ dW_i(t). \quad (31)$$

We denote the associated Markov transition kernel of $\{X(t)\}_{t \geq 0}$ by $\{\mathcal{P}_t\}_{t \geq 0}$, where each $\mathcal{P}_t : \Gamma \times \mathcal{B}(\Gamma) \rightarrow \mathbb{R}_+$ and $\mathcal{P}_t(x, S)$ is equal to the probability that starting from x the process is in the set $S \subseteq \Gamma$ at time t . Each \mathcal{P}_t is a Markov operator on the space of probability measures on Γ , where for any probability measure μ on Γ , $\mathcal{P}_t\mu$ is defined by $(\mathcal{P}_t\mu)(S) := \int_{x \in \mathcal{X}} \mathcal{P}_t(x, S)\mu(dx)$. For

convenience, we denote $f_0 := \hat{f}_0 - \frac{1}{2} \sum_{i=1}^{\Xi} \partial f_i f_i$. $\{\mathcal{P}_t\}_{t \geq 0}$ is also known as a *Markov semigroup*, and the corresponding generator is

$$\mathcal{L} := f_0 + \frac{1}{2} \sum_{i=1}^{\Xi} f_i^2.$$

To study the ergodicity of $\{\mathcal{P}_t\}_{t \geq 0}$, the approach used in [74, 73] associates the diffusion process to a deterministic control system

$$\frac{d\psi(t)}{dt} = f_0(\psi(t)) + \sum_{i=1}^{\Xi} f_i(\psi(t))u_i(t) \quad (32)$$

where $u : \mathbb{R}_+ \rightarrow \mathbb{R}^{\Xi}$ is any piecewise continuous function, which is called the control function. The controllability of the control system (32) can be characterized using the property of Lie algebra generated by $\{f_i\}_{i=0}^{\Xi}$, which is closely related to the ergodic properties of $\{\mathcal{P}_t\}_{t \geq 0}$.

Next, we introduce some useful notions and results corresponding to the control system (32). Denote by $\psi(t, x, u)$ the solution to (32) at time t with initialization x under the control u . Define the positive orbit of $x \in \Gamma$ at time t as $\mathcal{O}^+(x, t) = \{y \in \Gamma : \text{there exists } u \in \mathcal{U} \text{ such that } y = \psi(t, x, u)\}$, and $\mathcal{O}^+(x) = \cup_{t \geq 0} \mathcal{O}^+(x, t)$. Let $S \subseteq \Gamma$ satisfy that for any $x, y \in S$, it holds that $y \in \mathcal{O}^+(x)$. For any such S , there exists a unique maximal set $R \supseteq S$ with this property, and such a R is called a *control set* for the control system (32).

Definition F.1 (Invariant control set). A set $S \subseteq \Gamma$ is called an *invariant set* for the control system (32) if $\mathcal{O}^+(x) \subseteq S$ for all $x \in S$. A control set $R \subseteq \Gamma$ is called an *invariant control set* for the control system (32) if for any $x \in R$, $\mathcal{O}^+(x) = R$.

A probability measure μ on Γ is called an *invariant probability measure* for the diffusion process $X(t)$ if it satisfies that $\mu(S) = \int_{x \in \Gamma} \mathcal{P}_t(x, S) \mu(dx)$ for any Borel set $S \subseteq \Gamma$. Moreover, if μ cannot be decomposed into the sum of two different invariant measures, then we say μ is an *extremal invariant probability measure*. The following lemma characterize the relationship between extremal invariant probability measures and invariant control sets.

Lemma F.2 (Lemma 4.1, Kliemann [73]). *Let μ be an extremal invariant probability measure for $\{X(t)\}_{t \geq 0}$. Then the support of μ , denoted by $\text{supp } \mu$, is an invariant control set, and μ is the unique invariant probability measure on $\text{supp } \mu$.*

We can view each \mathcal{P}_t as a mapping from the space of functions on Γ to itself, by generalizing the fact that for any probability measure μ on Γ , $\mathcal{P}_t \mu$ is still a probability measure on Γ . Then we say the Markov semigroup $\{\mathcal{P}_t\}_{t \geq 0}$ satisfies the *Feller property* if each \mathcal{P}_t maps the space of continuous bounded functions into itself. The following theorem guarantees the existence of invariant probability measures for any Feller process.

Theorem F.3 (Krylov-Bogolioubov Theorem, [76]). *Let $\{\mathcal{P}_t\}_{t \geq 0}$ be a Feller Markov semigroup over Γ . Assume that there exists a probability measure μ_0 on Γ such that the sequence $\{\mathcal{P}_t \mu_0\}_{t \geq 0}$ is tight, then there exists at least one invariant probability measure for $\{\mathcal{P}_t\}_{t \geq 0}$.*

However, the Feller property does not guarantee the ergodicity of a Markov semigroup, as there might be many invariant probability measures having disjoint supports. To this end, we need a stronger regularity condition known as the *strong Feller property*. Specifically, the Markov semigroup $\{\mathcal{P}_t\}_{t \geq 0}$ is strong Feller if each \mathcal{P}_t maps all bounded Borel functions to bounded continuous functions. The well known Hörmander's theorem provides a sufficient condition for the strong Feller property, which depends on the Lie algebra generated by $\{f_i\}_{i=0}^{\Xi}$.

For any vector fields f, g on Γ , the Lie bracket between f and g is defined by

$$[f, g](x) = \partial g(x)f(x) - \partial f(x)g(x).$$

The Lie algebra generated by $\{f_i\}_{i=0}^{\Xi}$, denoted by $\mathcal{LA}(f_0, f_1, \dots, f_{\Xi})$, is the smallest vector space containing f_0, f_1, \dots, f_{Ξ} that is further closed under the Lie bracket operation. The following lemma characterizes when the Markov semigroup $\{\mathcal{P}_t\}_{t \geq 0}$ is strong Feller.

Lemma F.4 (Lemma 5.1, Ichihara and Kunita [74]). *For the Markov semigroup $\{\mathcal{P}_t\}_{t \geq 0}$ associated with the diffusion process (30), regard $f_0 + \frac{\partial}{\partial t}, f_1, \dots, f_\Xi$ as vector fields on the product manifold $\mathbb{R}_+ \times \Gamma$, where \mathbb{R}_+ is the time set and $\frac{\partial}{\partial t}$ is the shift invariant vector field in the time variable. Suppose $\dim \mathcal{L}\mathcal{A}(f_0 + \frac{\partial}{\partial t}, f_1, \dots, f_t)(t, x) = \dim(\Gamma) + 1$ for all $(t, x) \in \mathbb{R}_+ \times \Gamma$. Then $\{\mathcal{P}_t\}_{t \geq 0}$ is a strong Feller Markov semigroup.*

To determine the dimension of $\mathcal{L}\mathcal{A}(f_0 + \frac{\partial}{\partial t}, f_1, \dots, f_t)$, we can apply Lemma F.5.

Lemma F.5 (Lemma 2.1, Ichihara and Kunita [74]). *Assume that $\dim \mathcal{L}\mathcal{A}(f_0, f_1, \dots, f_\Xi) = \dim(\Gamma)$. Then for any $x \in \Gamma$, $\dim \mathcal{L}\mathcal{A}(f_0 + \frac{\partial}{\partial t}, f_1, \dots, f_\Xi)(x)$ is $\dim(\Gamma)$ or $\dim(\Gamma) - 1$. Furthermore, the following three conditions are equivalent:*

- (a) $\dim \mathcal{L}\mathcal{A}(f_0 + \frac{\partial}{\partial t}, f_1, \dots, f_\Xi)(x) = \dim(\Gamma)$.
- (b) $\dim \mathcal{L}\mathcal{A}_0(x) = \dim(\Gamma) - 1$.
- (c) $f_0(x) \notin \mathcal{L}\mathcal{A}_0(x)$.

Here $\mathcal{L}\mathcal{A}_0$ is defined as $\mathcal{L}\mathcal{A}_0 = \{\sum_{i=1}^k \lambda_i f_i + g : k \in [\Xi], \lambda_i \in \mathbb{R}, \forall i \in [k], g \in \mathcal{L}\mathcal{A}'\}$, where $\mathcal{L}\mathcal{A}'$ is the set of all linear sums of Lie brackets between any two iterative Lie brackets of $\{f_0, f_1, \dots, f_\Xi\}$.

The strong Feller property further implies that the transition probability admits a continuous density.

Theorem F.6 (Theorem, Ichihara and Kunita [74]). *Under the setting of Lemma F.4, there exists a C^∞ function $p_t(x, y)$ on $(0, \infty) \times \Gamma \times \Gamma$ such that $p_t(x, y)dy = \mathcal{P}_t(x, dy)$.*

Finally, for a strong Feller $\{\mathcal{P}_t\}_{t \geq 0}$, its ergodic property is given in the following proposition.

Proposition F.7 (Proposition 5.1, Ichihara and Kunita [74]). *Under the setting of Lemma F.4, further suppose the Markov semigroup $\{\mathcal{P}_t\}_{t \geq 0}$ has an extremal invariant probability measure μ . Then $\lim_{t \rightarrow \infty} \|\mathcal{P}_t(x, \cdot) - \mu\|_{TV} = 0$ for any $x \in \text{supp } \mu$.*

With the above tools in hand, we are ready to show the mixing of the limiting diffusion on the manifold of local minimizers. The proofs in the following subsections are organized as follows:

1. We first prove Lemma F.9 in Appendix F.2, which describes the dynamics of the parameter norm of the limiting diffusion.
2. Then in Appendix F.3 we prove the mixing of the limiting diffusion in the special case when the trace of the noise covariance is constant on Γ_1 .
3. Finally, we provide the proof for the mixing in the general case in Appendix F.4.

We note that our proofs only requires a strictly weaker assumption than Assumption 5.2, which is stated below (Assumption F.8).

Assumption F.8. $\dim \mathcal{L}\mathcal{A}(f_1, \dots, f_\Xi) = \dim(\Gamma) - 1 = D - M - 1$.

F.2 Norm Dynamics of The Limiting Diffusion

First, we characterize the norm of the limiting dynamics as follows.

Lemma F.9. *For SDE (13), it holds that $d\|Y(t)\|_2^2 = -\|Y(t)\|_2^2 dt + \text{tr}(\Sigma(Y(t)))dt$.*

We need the following intermediate lemma.

Lemma F.10. *Under Assumption 2.3, for any $x \in U$ and $Q \in \mathbb{R}^{D \times D}$, it holds that*

$$\langle \Phi(x), \partial^2 \Phi(x)[Q] \rangle = \text{tr}(Q) - \text{tr}(\partial \Phi(x) Q \partial \Phi(x)^\top).$$

Proof of Lemma F.10. Note that for any $x \in U$, since L is scale invariant $\langle x, \nabla L(x) \rangle = 0$, we have $\|\Phi(x)\|_2^2 = \|x\|_2^2$, and thus $\partial^2 \|\Phi(x)\|_2^2 = 2I_D$, that is,

$$\partial^2 \|\Phi(x)\|_2^2[Q] = 2I_D[Q] = 2 \text{tr}(Q). \quad (33)$$

for any $Q \in \mathbb{R}^{D \times D}$. On the other hand, we have

$$\partial_j \|\Phi(x)\|_2^2 = 2 \sum_{i=1}^D \partial_j \Phi_i(x) \Phi_i(x)$$

for any $j \in [D]$, and then

$$\partial_{jk} \|\Phi(x)\|_2^2 = 2 \sum_{i=1}^D \partial_{jk} \Phi_i(x) \Phi_i(x) + 2 \sum_{i=1}^D \partial_j \Phi_i(x) \partial_k \Phi_i(x)$$

for any $j, k \in [D]$. Therefore, for any $Q \in \mathbb{R}^{D \times D}$, we have

$$\begin{aligned} \partial^2 \|\Phi(x)\|_2^2 [Q] &= \sum_{j,k=1}^D \partial_{jk} \|\Phi(x)\|_2^2 Q_{jk} \\ &= 2 \sum_{i,j,k=1}^D \partial_{jk} \Phi_i(x) \Phi_i(x) Q_{jk} + 2 \sum_{i,j,k=1}^D \partial_j \Phi_i(x) \partial_k \Phi_i(x) Q_{jk} \\ &= 2 \langle \Phi(x), \partial^2 \Phi(x) [Q] \rangle + 2 \text{tr}(\partial \Phi(x) Q \partial \Phi(x)^\top). \end{aligned} \quad (34)$$

Finally, combining (33) and (34) yields the desired result. \square

Proof of Lemma F.9. Applying Ito's lemma, we have

$$\begin{aligned} d\|Y(t)\|_2^2 &= -\langle Y(t), \partial \Phi(Y(t)) Y(t) \rangle dt + \langle Y(t), \partial^2 \Phi(Y(t)) [\Sigma(Y(t))] \rangle dt \\ &\quad + 2 \langle Y(t), \partial \Phi(Y(t)) \sigma(Y(t)) dW(t) \rangle + \text{tr} [\partial \Phi(Y(t)) \Sigma(Y(t)) \partial \Phi(Y(t))^\top] dt. \end{aligned}$$

Note that since $\partial \Phi(Y(t)) Y(t) = Y(t)$, we have

$$\begin{aligned} \langle Y(t), \partial \Phi(Y(t)) \sigma(Y(t)) dW(t) \rangle &= \langle \partial \Phi(Y(t)) Y(t), \sigma(Y(t)) dW(t) \rangle \\ &= \langle Y(t), \sigma(Y(t)) dW(t) \rangle = 0. \end{aligned}$$

Since $Y(t)$ stays on Γ , we always have $\Phi(Y(t)) = Y(t)$, then it follows from Lemma F.10 that $\langle Y(t), \partial^2 \Phi(Y(t)) [\Sigma(Y(t))] \rangle = \text{tr}(\Sigma(Y(t))) - \text{tr}(\partial \Phi(Y(t)) \Sigma(Y(t)) \partial \Phi(Y(t))^\top)$. Combining the above yields the desired formula of the dynamics of $\|Y(t)\|_2^2$. \square

Next, we divide the proof of Theorem 5.4 into two parts stratified by whether the trace of the noise covariance is constant on Γ_1 or not.

F.3 Proof for the Mixing When Trace is a Constant

From now on, we specify Γ to be the manifold of local minimizers of the loss L , as given in Assumption 2.1. Recall that $\Gamma_1 = \{x \in \Gamma \mid \|x\|_2 = 1\}$ is the restriction of Γ on the unit sphere. In this subsection, we prove the mixing result when the $\text{tr}(\Sigma)$ is a constant on Γ_1 , that is, under Assumption F.11.

Assumption F.11. For any $x \in \Gamma_1$, $\text{tr}(\Sigma(x)) = 1$.¹

Note that, by Lemma F.9, Assumption F.11 implies that

$$d\|Y(t)\|_2^2 = \left(-\|Y(t)\|_2^2 + \frac{1}{\|Y(t)\|_2^2} \right) dt,$$

which is deterministic and has a closed-form solution:

$$\|Y(t)\|_2^2 = \sqrt{1 + e^{-2t} (\|Y(0)\|_2^4 - 1)}. \quad (35)$$

Therefore, regardless of the initialization, $\|Y(t)\|_2$ will always (monotonically) converge to 1 as $t \rightarrow \infty$. In this case, we only need to consider the mixing of the direction of $Y(t)$, i.e. $Y(t)/\|Y(t)\|_2$.

¹Here 1 can be replaced by any positive constant and the proof still works.

Lemma F.12. Under Assumption F.11, for SDE (13), define $\bar{Y}(t) = Y(t)/\|Y(t)\|_2$. Then

$$d\bar{Y}(t) = \frac{\partial^2 \Phi(\bar{Y}(t))[\Sigma(\bar{Y}(t))] - \bar{Y}(t)}{2 + 2e^{-2t}(\|Y(0)\|_2^4 - 1)} dt + \frac{\partial \Phi(\bar{Y}(t))\sigma(\bar{Y}(t))}{\sqrt{1 + e^{-2t}(\|Y(0)\|_2^4 - 1)}} dW(t). \quad (36)$$

Proof of Lemma F.12. By Lemma F.9, we have

$$\begin{aligned} d\bar{Y}(t) &= \frac{1}{\|Y(t)\|_2} dY(t) - \frac{Y(t)}{2\|Y(t)\|_2^3} d\|Y(t)\|_2^2 \\ &= -\frac{\partial \Phi(Y(t))Y(t)}{2\|Y(t)\|_2} dt + \frac{\partial^2 \Phi(Y(t))[\Sigma(Y(t))]}{2\|Y(t)\|_2} dt + \frac{\partial \Phi(Y(t))\sigma(Y(t))}{\|Y(t)\|_2} dW(t) \\ &\quad + \frac{Y(t)}{2\|Y(t)\|_2} dt - \frac{\text{tr}(\Sigma(Y(t)))Y(t)}{2\|Y(t)\|_2^3} dt \\ &= \frac{\partial^2 \Phi(\bar{Y}(t))[\Sigma(\bar{Y}(t))] - \bar{Y}(t)}{2\|Y(t)\|_2^4} dt + \frac{\partial \Phi(\bar{Y}(t))\sigma(\bar{Y}(t))}{\|Y(t)\|_2^2} dW(t) \end{aligned} \quad (37)$$

where the third equality follows from Lemma 2.8 and the assumption that $\text{tr}(\Sigma(x)) \equiv 1/\|x\|_2^2$. Finally plugging-in the expression for $\|Y(t)\|_2^2$ in Equation (35) finishes the proof. \square

Now according to Lemma F.12, $\{\bar{Y}(t)\}_{t \geq 0}$ is a diffusion process on Γ_1 . We denote the associated class of Markov transition kernel by $\{\tilde{\mathcal{P}}_t\}_{t \geq 0}$. To align with the notation in Equation (30) and (31), we define

$$\hat{f}_0(x) = \frac{1}{2}(\partial^2 \Phi(x)[\Sigma(x)] - x) \quad \text{and} \quad f_i(x) = \partial \Phi(x)\sigma_i(x) \quad \text{for each } i \in [\Xi].$$

We rescale the time by defining $\tau(t) = \int_0^t \frac{1}{1 + e^{-2s}(\|Y(0)\|_2^4 - 1)} ds$. Let $\tilde{Y}(\tau(t)) \equiv \bar{Y}(t)$, identifying $\tilde{t} \equiv \tau(t)$, then by Lemma F.12, \tilde{Y} admits

$$\begin{aligned} d\tilde{Y}(\tilde{t}) &= \frac{1}{2} \frac{\partial^2 \Phi(\bar{Y}(t))[\Sigma(\bar{Y}(t))] - \bar{Y}(t)}{1 + e^{-2t}(\|Y(0)\|_2^4 - 1)} dt + \frac{\partial \Phi(\bar{Y}(t))\sigma(\bar{Y}(t))}{\sqrt{1 + e^{-2t}(\|Y(0)\|_2^4 - 1)}} dW(t) \\ &= \frac{1}{2} (\partial^2 \Phi(\bar{Y}(t))[\Sigma(\bar{Y}(t))] - \bar{Y}(t)) d\tau(t) + \partial \Phi(\bar{Y}(t))\sigma(\bar{Y}(t)) dW(\tau(t)) \\ &= \hat{f}_0(\tilde{Y}(\tilde{t})) d\tilde{t} + \sum_{i=1}^{\Xi} f_i(\tilde{Y}(\tilde{t})) dW_i(\tilde{t}) \end{aligned}$$

where the second equality follows from Lemma G.1 and the definition of $\tau(t)$. For \tilde{Y} , we denote the associated Markov semigroup by $\{\tilde{\mathcal{P}}_t\}_{t \geq 0}$. Since there is a bijection between t and \tilde{t} and $\lim_{t \rightarrow \infty} \tilde{t}/t = 1$, the mixing of $\tilde{\mathcal{P}}_t$ implies that of $\bar{\mathcal{P}}_t$.

Next, as introduced in the previous subsection, writing the dynamics of \tilde{Y} as a Stratonovich SDE, we get

$$\begin{aligned} d\tilde{Y}(\tilde{t}) &= \hat{f}_0(\tilde{Y}(\tilde{t})) d\tilde{t} - \frac{1}{2} \sum_{i=1}^{\Xi} \partial f_i(\tilde{Y}(\tilde{t})) f_i(\tilde{Y}(\tilde{t})) d\tilde{t} + \sum_{i=1}^{\Xi} f_i(\tilde{Y}(\tilde{t})) \circ dW_i(\tilde{t}) \\ &= f_0(\tilde{Y}(\tilde{t})) d\tilde{t} + \sum_{i=1}^{\Xi} f_i(\tilde{Y}(\tilde{t})) \circ dW_i(\tilde{t}) \end{aligned}$$

whose associated deterministic control system is given by

$$\frac{d\psi(t)}{dt} = f_0(\psi(t)) + \sum_{i=1}^{\Xi} f_i(\psi(t)) u_i(t) \quad (38)$$

where u is any piecewise continuous \mathbb{R}^{Ξ} -valued functions.

The following lemma establishes the key property of the control system in Equation (38).

Lemma F.13. *Under Assumption 5.1 and F.8, Γ_1 itself is the unique invariant control set contained in Γ_1 for the control system (38).*

Proof of Lemma F.13. It suffices to show that for any $x, y \in \Gamma_1, y \in \overline{O^+(x)}$. By Assumption F.8, $\{f_i\}_{i=1}^\Xi$ satisfies the Lie algebra rank condition, so the following driftless control system

$$\frac{d\hat{\psi}(t)}{dt} = \sum_{i=1}^{\Xi} f_i(\hat{\psi}(t))u_i(t) \quad (39)$$

is globally controllable by Rashevski-Chow theorem (see, e.g., [77, 78]), which means that there exists some $t_{x,y} > 0$ and a control $u : \mathbb{R}_+ \rightarrow \mathbb{R}^\Xi$ such that $y = \hat{\psi}(t_{x,y}, x; u)$, where $\hat{\psi}(t_{x,y}, x; u)$ is the solution of Equation (39) at time $t_{x,y}$ with initialization x under the control u . Next, we can use the global controllability of the driftless control system (39) to show that $y \in \overline{O^+(x)}$ for the original control system (38).

For any $\delta > 0$, define a control $u_\delta : \mathbb{R}_+ \rightarrow \mathbb{R}^\Xi$ as $u_\delta(t) = u(t/\delta)$, then we have $y = \hat{\psi}(t_{x,y}, x; u) = \hat{\psi}(t_{x,y}\delta, x; u_\delta)$. Now using u_δ for the original control system (38), it follows that $\lim_{\delta \rightarrow 0} \psi(t_{x,y}\delta, x; u_\delta) = y$ as the drift is dominated by the controlled terms. Thus we conclude that $y \in \overline{O^+(x)}$, and this completes the proof. \square

Now we are ready to prove Theorem 5.4 in the special case of constant trace.

Proof of Theorem 5.4 when trace is constant. Since $\tilde{\mathcal{P}}_{\tilde{t}} = \bar{\mathcal{P}}_{\tilde{t}}$ and $\lim_{\tilde{t} \rightarrow \infty} \tilde{t}/t = 1$ by L'Hospital's rule, it suffices to show the ergodicity of $\{\tilde{\mathcal{P}}_{\tilde{t}}\}_{\tilde{t} \geq 0}$. Let $\mathcal{L}\mathcal{A}(f_1, \dots, f_\Xi)$ be the Lie algebra generated by $\{f_1, \dots, f_\Xi\}$. Assumption F.8 (implied by Assumption 5.2) assumes that $\dim \mathcal{L}\mathcal{A}(f_1, \dots, f_\Xi) = D - M - 1$, which is equal to the dimension of Γ_1 . Therefore, by Lemma F.5, it follows that $\dim \mathcal{L}\mathcal{A}(f_0 + \frac{\partial}{\partial \tilde{t}}, f_1, \dots, f_\Xi) = D - M = \dim(\Gamma_1) + 1$. Then applying Lemma F.4 yields that $\tilde{\mathcal{P}}_{\tilde{t}}$ satisfies the strong Feller property. Since Γ_1 is compact by Assumption 5.1, there exists at least one invariant measure by Theorem F.3, which further implies that there exists at least one extremal invariant measure by Krein-Milman theorem. Moreover, combining Lemma F.13 and Lemma F.2, we see that $\{\tilde{\mathcal{P}}_{\tilde{t}}\}_{\tilde{t} \geq 0}$ has a unique invariant measure, which we denote by π , and $\text{supp } \pi = \Gamma_1$. Finally, by Proposition F.7, $\tilde{\mathcal{P}}_{\tilde{t}}$ is ergodic in the sense that

$$\lim_{\tilde{t} \rightarrow \infty} \|\tilde{\mathcal{P}}_{\tilde{t}}(x, \cdot) - \pi(\cdot)\|_{\text{TV}} = 0, \quad \forall x \in \Gamma_1.$$

where $\|\cdot\|_{\text{TV}}$ denotes the total variation distance. This completes the proof. \square

F.4 Proof for the Mixing When Trace is not Constant

Recall the limiting diffusion given in Equation (13):

$$dY(t) = -\frac{1}{2}Y(t)dt + \frac{1}{2}\partial^2\Phi(Y(t))[\Sigma(Y(t))]dt + \partial\Phi(Y(t))\sigma(Y(t))dW(t).$$

We continue to use the notation $\hat{f}_0(x) = \frac{1}{2}(\partial^2\Phi(x)[\Sigma(x)] - x)$, $f_i(x) = \partial\Phi(x)\sigma_i(x)$ for each $i \in [\Xi]$, and $f_0(x) = \hat{f}_0(x) - \frac{1}{2}\sum_{i=1}^{\Xi}\partial f_i(x)f_i(x)$. Writing the dynamics of Y in (13) as a Stratonovich SDE, we again get

$$dY(t) = f_0(Y(t))dt + \sum_{i=1}^{\Xi} f_i(Y(t)) \circ dW_i(t).$$

Similarly we define its associated deterministic control system:

$$\frac{d\psi(t)}{dt} = f_0(\psi(t)) + \sum_{i=1}^{\Xi} f_i(\psi(t))u_i(t) \quad (40)$$

Since $\text{tr}(\Sigma(x))$ is continuous and Γ_1 is compact, we know that both $\beta_{\max} := \max_{x \in \Gamma_1} \text{tr}(\Sigma(x))$ and $\beta_{\min} := \min_{x \in \Gamma_1} \text{tr}(\Sigma(x))$ can be attained in Γ_1 . Note that $\beta_{\min} > 0$ by Assumption F.8.

Recall the dynamics of $\|Y(t)\|_2^2$ given by Lemma F.9, which implies that if $\|Y(t)\|_2 \geq \beta_{\max}^{1/4}$ then $\frac{d\|Y(t)\|_2^2}{dt} \leq 0$ and that if $\|Y(t)\|_2 \leq \beta_{\min}^{1/4}$ then $\frac{d\|Y(t)\|_2^2}{dt} \geq 0$. Similarly for the control system (40), we have

$$\frac{d\|\psi(t)\|_2^2}{dt} = 2\left\langle \psi(t), \frac{d\psi(t)}{dt} \right\rangle = 2\langle \psi(t), f_0(\psi(t)) \rangle + 2 \sum_{i=1}^{\Xi} u_i(t) \langle \psi(t), f_i(\psi(t)) \rangle$$

where the second term is equal to 0 as each $f_i(\psi(t)) = \partial\Phi(\psi(t))\sigma_i(\psi(t))$ is in $T_{\psi(t)}(\Gamma_{\|\psi(t)\|_2})$. To further simplify the first term, we need the following lemma.

Lemma F.14. *Under Assumption 2.3, for any $x \in \Gamma$, it holds that $\text{tr}(\Sigma(x)) = 2\langle x, f_0(x) + \frac{x}{2} \rangle$.*

Proof of Lemma F.14. Applying Lemma F.10, we have

$$\begin{aligned} 2\left\langle x, f_0(x) + \frac{x}{2} \right\rangle &= \langle x, \partial^2\Phi(x)[\Sigma(x)] \rangle - \sum_{i=1}^{\Xi} \langle x, \partial f_i(x) f_i(x) \rangle \\ &= \text{tr}(\Sigma(x)) - \text{tr}(\partial\Phi(x)\Sigma(x)\partial\Phi(x)^\top) - \sum_{i=1}^{\Xi} \langle x, \partial f_i(x) f_i(x) \rangle. \end{aligned}$$

Recall that $\Sigma(x) = \sigma(x)\sigma(x)^\top$ and $\sigma(x) = (\sigma_1(x), \dots, \sigma_{\Xi}(x))$. Moreover, for all $i \in [\Xi]$, $\langle x, f_i(x) \rangle = \langle x, \partial\Phi(x)\sigma_i(x) \rangle = 0$, which by differentiating with x further implies that $\partial f_i(x)^\top x = -f_i(x)$. Therefore, it follows that

$$\begin{aligned} 2\left\langle x, f_0(x) + \frac{x}{2} \right\rangle &= \text{tr}(\Sigma(x)) - \sum_{i=1}^{\Xi} (\text{tr}(\partial\Phi(x)\sigma_i(x)\sigma_i(x)^\top\partial\Phi(x)^\top) - \langle \partial\Phi(x)\sigma_i(x), \partial\Phi(x)\sigma_i(x) \rangle) \\ &= \text{tr}(\Sigma(x)). \end{aligned}$$

This completes the proof. \square

Hence, similar to $\|Y(t)\|_2^2$, the dynamics of $\|\psi(t)\|_2^2$ for the control system (40) can also be simplified into

$$\frac{d\|\psi(t)\|_2^2}{dt} = -\|\psi(t)\|_2^2 + \text{tr}(\Sigma(\psi(t))).$$

Now, define

$$\Gamma_{a,b} := \{x \in \Gamma \mid a \leq \|x\|_2 \leq b\},$$

and the above calculation implies that $\Gamma_{a,b}$ is an invariant set for the control system (40) if $a \leq \beta_{\min}^{1/4}$ and $b \geq \beta_{\max}^{1/4}$. This motivates us to define

$$\Gamma_* := \{x \in \Gamma \mid \beta_{\min} \leq \|x\|_2^4 \leq \beta_{\max}\}, \quad (41)$$

which can be shown to be the unique invariant control set in Γ for the control system (40).

Lemma F.15. *Under Assumption 5.1 and F.8, Γ_* is the unique invariant control set contained in Γ for the control system (40).*

Proof of Lemma F.15. As discussed above, Γ_* itself is an invariant set, so it suffices to show that for every $x \in \Gamma$, $\Gamma_* \subseteq \overline{\mathcal{O}^+(x)}$. Specifically, below we will show that for any $y \in \Gamma_*$ and $x \in \Gamma$, $y \in \overline{\mathcal{O}^+(x)}$.

By the same argument as in the proof of Lemma F.13, Assumption 5.1 and F.8 imply that for any $x \in \Gamma$, we have $\Gamma_{\|x\|_2} \subset \overline{\mathcal{O}^+(x)}$. Then, it suffices to show there exist $u \in \Gamma_{\|x\|_2}$ and $v \in \Gamma_{\|y\|_2}$ such that $v \in \overline{\mathcal{O}^+(u)}$, as it would follow that $y \in \overline{\mathcal{O}^+(v)} \subseteq \overline{\mathcal{O}^+(u)} \subseteq \overline{\mathcal{O}^+(x)}$.

Without loss of generality, we assume that $\|x\|_2 > \|y\|_2$. Our strategy is to pick an arbitrary $z_* \in \text{argmin}_{z \in \Gamma_1} \text{tr}(\Sigma(z))$ and let $v = \|y\|_2 z_*$ and $u = \|x\|_2 z_*$. Note that for any $r \in (\|y\|_2, \|x\|_2]$, we have

$$-\|rz_*\|_2^2 + \text{tr}(\Sigma(rz_*)) = -r^2 + \frac{1}{r^2} \text{tr}(\Sigma(z_*)) < \frac{1}{r^2} (-\|y\|_2^4 + \text{tr}(\Sigma(z_*))) \leq 0 \quad (42)$$

where the last inequality is due to the definition of Γ_* . Thus the idea is to use the controlled terms in (40), i.e., $\sum_{i=1}^{\Xi} f_i(\psi(t))u_i(t)$, to cancel the part of $f_0(\psi(t))$ in the tangent space of $\Gamma_{\|\psi(t)\|_2}$, so that the remaining dynamics becomes

$$\frac{d\tilde{\psi}(t)}{dt} = -\frac{1}{2}\tilde{\psi}(t) + \frac{1}{2}\left\langle \frac{\tilde{\psi}(t)}{\|\tilde{\psi}(t)\|_2}, f_0(\tilde{\psi}(t)) \right\rangle \frac{\tilde{\psi}(t)}{\|\tilde{\psi}(t)\|_2} \quad (43)$$

under which v is in the positive orbit of u by the above calculation.

To do so, for any $r \in (\|y\|_2, \|x\|_2]$, consider the part of $f_0(rz_*)$ in the tangent space of Γ_r , which is given by

$$\begin{aligned} (I - z_* z_*^\top) f_0(rz_*) &= \frac{1}{2}(I - z_* z_*^\top) \left(\partial^2 \Phi(rz_*) [\Sigma(rz_*)] - \sum_{i=1}^{\Xi} \partial(\partial \Phi \sigma_i)(rz_*) \partial \Phi(rz_*) \sigma_i(rz_*) \right) \\ &= \frac{1}{2r^3}(I - z_* z_*^\top) \left(\partial^2 \Phi(z_*) [\Sigma(z_*)] - \sum_{i=1}^{\Xi} \partial(\partial \Phi \sigma_i)(z_*) \partial \Phi(z_*) \sigma_i(z_*) \right) \\ &= \frac{1}{2r^3}(I - z_* z_*^\top) f_0(z_*) \end{aligned} \quad (44)$$

where the second equality follows from the fact that $\partial \Phi(x)$ is 0-homogeneous, $\partial^2 \Phi(x)$ and each $\sigma_i(x)$ are (-1) -homogeneous, and $\Sigma(x)$ is (-2) -homogeneous. By Assumption F.8, there exists some $\lambda \in \mathbb{R}^{\Xi}$ such that

$$\frac{1}{2}(I - z_* z_*^\top) f_0(z_*) = \sum_{i=1}^{\Xi} \lambda_i \partial \Phi(z_*) \sigma_i(z_*). \quad (45)$$

For the ordinary differential equation (10) with initialization $\tilde{\psi}(0) = \|x\|_2 z_*$, denote $r(t) = \|\tilde{\psi}(t)\|_2$ which is continuously decreasing in time t . Then for the control system (40) initialized at $\psi(0) = \|x\|_2 z_*$, we choose the control function as $u_i(t) = \frac{\lambda_i}{r(t)^3}$ for all $i \in [\Xi]$, it follows from (44) and (45) that

$$\begin{aligned} \frac{d\psi(t)}{dt} &= f_0(\psi(t)) + \sum_{i=1}^{\Xi} \partial \Phi(\psi(t)) \sigma_i(\psi(t)) u_i(t) \\ &= f_0(\psi(t)) + \frac{\lambda_i}{r(t)^3} \sum_{i=1}^{\Xi} \partial \Phi(\psi(t)) \sigma_i(\psi(t)) \\ &= -\frac{1}{2}\psi(t) + \frac{1}{2}\left\langle \frac{\psi(t)}{\|\psi(t)\|_2}, f_0(\psi(t)) \right\rangle \frac{\psi(t)}{\|\psi(t)\|_2} \end{aligned}$$

whose solution is given by $\psi(t) = r(t)z_*$. By (42), it holds that either $r(t) = \|y\|_2$ for some $t > 0$ or $\lim_{t \rightarrow \infty} r(t) = \|y\|_2$. Therefore, we conclude that $\|y\|_2 z_* \in \mathcal{O}^+(\|x\|_2 z_*)$. This finishes the proof. \square

Next, to show that the limiting diffusion is still a strong Feller process, we need to verify the condition for Lemma F.4.

Lemma F.16. *Let f and g be two vector fields on Γ . Suppose $\langle f(x), x \rangle = 0$ for every $x \in \Gamma$, then it holds that $\langle x, [f, g](x) \rangle = \langle \nabla \langle x, g(x) \rangle, f(x) \rangle$. In particular, $\langle x, [x, f](x) \rangle = 0$.*

Proof of Lemma F.16. By definition, we have

$$\langle x, [f, g](x) \rangle = \langle x, \partial g(x) f(x) - \partial f(x) g(x) \rangle \quad (46)$$

and

$$\langle \nabla \langle x, g(x) \rangle, f(x) \rangle = \langle g(x) + \partial g(x)^\top x, f(x) \rangle. \quad (47)$$

Since $\langle f(x), x \rangle \equiv 0$, differentiating both sides with x yields

$$f(x) + \partial f(x)^\top x = 0 \quad (48)$$

Combining (47) and (48), we obtain

$$\langle \nabla \langle x, g(x) \rangle, f(x) \rangle = \langle g(x), -\partial f(x)^\top x \rangle + \langle \partial g(x)^\top x, f(x) \rangle. \quad (49)$$

Finally, comparing (46) with (49), we conclude that $\langle x, [f, g](x) \rangle = \langle \nabla \langle x, g(x) \rangle, f(x) \rangle$. Moreover, in the special case of $g(x) \equiv x$, it follows that

$$\langle x, [x, f](x) \rangle = -\langle x, [f, x](x) \rangle = -\langle \nabla \langle x, x \rangle, f(x) \rangle = -2\langle x, f(x) \rangle = 0$$

where the first equality follows from the anti-symmetry of Lie bracket. \square

Definition F.17 (Analytic Function). Let $V \subseteq \mathbb{R}^D$ be any open set, a function $f : V \rightarrow \mathbb{R}$ is *analytic* if for each $x \in V$, there is a neighborhood V_x of x such that the Taylor series of f expanded at x converges to f everywhere in V_x . A vector-valued function is analytic if each of its coordinate functions is analytic.

Lemma F.18. For any positive integer k , define

$$A^k := \{[[[f_{i_1}, f_{i_2}], f_{i_3}], \dots, f_{i_k}] \mid i_1 \in [\Xi], i_2, \dots, i_k \in \{0\} \cup [\Xi]\}$$

and $A = \bigcup_{k=1}^{\infty} A_k$. Then under Assumption F.8, 2.3 and 5.3, it holds that for all $x \in \Gamma$,

$$\text{rank}(\text{span}(\{f(x) \mid f \in A\})) = D - M.$$

Proof of Lemma F.18. By Assumption F.8, it suffices to show for each $x^* \in \Gamma_1$, that there exists some $f \in A$ such that $\langle x^*, f(x^*) \rangle \neq 0$. Now that $\text{tr}(\Sigma(x))$ is a non-constant analytic function, we claim that for every $x^* \in \Gamma_1$, there must exist a relatively open neighborhood $V \subseteq \Gamma_1$ of x^* , some $k^* \geq 1$, some elements of A , $\{h_i\}_{i=1}^N$, and some linear combination denoted by $f = \sum_{i=1}^N \alpha_i h_i$, such that

$$f^{(k^*)}(\text{tr}(\Sigma))(x) = \underbrace{(f \circ \dots \circ f)}_{k^* \text{ times}}(\text{tr}(\Sigma))(x) \neq 0.$$

Here for a smooth vector field f on Γ_1 and a smooth function g on Γ_1 , we define $f(g)(x) \triangleq \frac{\partial g(x+tf(x))}{\partial t} \Big|_{t=0} = \langle \nabla g(x), f(x) \rangle$ for $x \in \mathbb{R}^D$. Moreover, we define $f^{(k)}$ as the composition of f for k times.

To prove the above claim, we first note that since Γ_1 is an analytic manifold, there is an analytic diffeomorphism $\psi : V \rightarrow \mathbb{R}^{D-M}$ such that the push-forward mapping of h_i, f and $\text{tr}(\Sigma)$ are all analytic, which are defined as $\psi_* h_i(y) := \frac{\partial \psi(x)}{\partial x} \Big|_{x=\psi^{-1}(y)} h_i(\psi^{-1}(y))$, $\psi_* f(y) := f(\psi^{-1}(y))$ and $\psi_* \text{tr}(\Sigma)(y) := \text{tr}(\Sigma(\psi^{-1}(y)))$ respectively. For convenience, we denote $\psi_* \text{tr}(\Sigma)$ by g . Since g is analytic and non-constant, there exists $k^* \geq 1$ and $v \in \mathbb{R}^{D-M}$ such that $\nabla^{k^*} g(x^*)[v, \dots, v] \neq 0$, and $\nabla^k g(\psi(x^*)) = 0$ for all $k \leq k^* - 1$. By Assumption F.8, we know there are some elements of A , $\{h_i\}_{i=1}^N$, and some linear combination denoted by $f = \sum_{i=1}^N \alpha_i h_i$, such that $\psi_* f(\psi(x^*)) = v$. Therefore,

$$f^{(k^*)}(\text{tr}(\Sigma))(x^*) = (\psi_* f)^{(k^*)}(g)(\psi(x^*)) = \nabla^{k^*} g(\psi(x^*)) [v, \dots, v] \neq 0,$$

where the second equality we use the property that $\nabla^k g(\psi(x^*)) = 0$ for all $k \leq k^* - 1$.

Now we continue the proof of Lemma F.18. By Lemma F.14, we know that $\text{tr}(\Sigma(x)) = 2\langle x, f_0(x) + \frac{x}{2} \rangle$. Since $\langle f(x), x \rangle = 0$ for every $x \in \Gamma$, we can apply Lemma F.16 iteratively to get

$$\begin{aligned} \left\langle x, \left[\left[\left[\left[f_0(x) + \frac{x}{2}, f \right], f \right], \dots \right], f \right](x) \right\rangle &= \left\langle \nabla \left\langle \dots \left\langle \nabla \left\langle \nabla \left\langle x, f_0(x) + \frac{x}{2} \right\rangle, f \right\rangle, f \right\rangle \dots, f \right\rangle, f \right\rangle(x) \\ &= \frac{1}{2} \left\langle \nabla \left\langle \dots \left\langle \nabla \left\langle \nabla \text{tr}(\Sigma), f \right\rangle, f \right\rangle \dots, f \right\rangle, f \right\rangle(x) \\ &= \underbrace{(f \circ \dots \circ f)}_{k^* \text{ times}}(\text{tr}(\Sigma))(x) \neq 0 \end{aligned}$$

Applying the second claim of Lemma F.16 iteratively, we have that $\langle x, \left[\left[\left[\left[\frac{x}{2}, f \right], f \right], \dots \right], f \right](x) \rangle = 0$. This completes the proof. \square

Proof of Theorem 5.4 when trace is not constant. Let $\{\mathcal{P}_t\}_{t \geq 0}$ be the Markov semigroup associated with the limiting diffusion. Recall the unique invariant control set $\Gamma_* = \{x \in \Gamma \mid \beta_{\min} \leq \|x\|_2^4 \leq \beta_{\max}\}$ defined in (41). As discussed above, the dynamics of $\|Y(t)\|_2^2$ satisfies that $\frac{d\|Y(t)\|_2^2}{dt} \leq 0$ when $\|Y(t)\|_2 \geq \beta_{\max}^{1/4}$ and $\frac{d\|Y(t)\|_2^2}{dt} \geq 0$ when $\|Y(t)\|_2 \leq \beta_{\min}^{1/4}$. Thus, without loss of generality, we can assume that the initialization x_{init} is in a compact manifold $\widehat{\Gamma} = \Gamma_{a,b}$ where $0 < a \leq \beta_{\min}^{1/4}$ and $\beta_{\max}^{1/4} \leq b < \infty$.

Based on Lemma F.15 and F.18, following the same argument as that in the proof for case of constant trace, we get that $\{\mathcal{P}_t\}_{t \geq 0}$ is strong Feller, has a unique invariant measure π_* and is ergodic in the sense that

$$\lim_{t \rightarrow \infty} \|\mathcal{P}_t(x, \cdot) - \pi_*(\cdot)\|_{\text{TV}} = 0, \quad \forall x \in \text{supp}(\pi_*) \quad (50)$$

where $\text{supp}(\pi_*) = \Gamma_*$. It remains to generalize the above convergence guarantee to all $x \in \widehat{\Gamma}$.

Note that $\mathcal{P}_t(x, \Gamma_*) = 1$ for all $t \geq 0$ and $x \in \Gamma_*$ as Γ_* is the unique invariant control set. Then fixing any $s_* > 0$ and $\epsilon > 0$, by the strong Feller property and Theorem F.6, there exists some $a_{s_*, \epsilon} \in (0, \beta_{\min}^{1/4})$ and $b_{s_*, \epsilon} \in (\beta_{\max}^{1/4}, \infty)$ such that $\mathcal{P}_{s_*}(x, \Gamma_*) \geq 1 - \epsilon$ for all $x \in \widetilde{\Gamma} := \Gamma_{a_{s_*, \epsilon}, b_{s_*, \epsilon}}$. Further note that for any $t > s_*$ and $x \in \widetilde{\Gamma}$,

$$\begin{aligned} \mathcal{P}_t(x, \Gamma_*) &= \mathbb{P}(Y(t) \in \Gamma_* \mid Y(0) = x) \\ &= \mathbb{P}(Y(t) \in \Gamma_* \mid Y(s_*) \in \Gamma_*, Y(0) = x) \cdot \mathbb{P}(Y(s_*) \in \Gamma_* \mid Y(0) = x) \\ &\quad + \mathbb{P}(Y(t) \in \Gamma_* \mid Y(s_*) \notin \Gamma_*, Y(0) = x) \cdot \mathbb{P}(Y(s_*) \notin \Gamma_* \mid Y(0) = x) \\ &\geq \mathbb{P}(Y(t) \in \Gamma_* \mid Y(s_*) \in \Gamma_*) \cdot \mathcal{P}_{s_*}(x, \Gamma_*) = 1 - \epsilon \end{aligned} \quad (51)$$

where for the inequality we use the Markov property. Moreover, by Lemma F.9, we have $\frac{d\|Y(t)\|_2^2}{dt} < -c_{s_*, \epsilon} < 0$ when $\|Y(t)\|_2 \leq a_{s_*, \epsilon}$ or $\|Y(t)\|_2 \geq b_{s_*, \epsilon}$, for some constant $c_{s_*, \epsilon} > 0$. This implies that for any initialization $Y(0) \in \widehat{\Gamma}$, $Y(t)$ will reach $\widetilde{\Gamma}$ after an initial burn-in period of length at most $s_0 = \max\{\frac{a_{s_*, \epsilon} - a}{c_{s_*, \epsilon}}, \frac{b - b_{s_*, \epsilon}}{c_{s_*, \epsilon}}\}$. Note that $\widetilde{\Gamma}$ is also an invariant set.

Next, for any $x \in \widehat{\Gamma}$ and $t \geq s_* + s_0$, we can bound the TV distance between $\mathcal{P}_t(x, \cdot)$ and π_* as follows:

$$\begin{aligned} \|\mathcal{P}_t(x, \cdot) - \pi_*(\cdot)\|_{\text{TV}} &= \frac{1}{2} \int_{\widehat{\Gamma}} |\mathcal{P}_t(x, dy) - \pi_*(dy)| \\ &= \frac{1}{2} \int_{y \in \widehat{\Gamma}} \left| \int_{z \in \widehat{\Gamma}} (\mathcal{P}_{t-s_0}(z, dy) - \pi_*(dy)) \mathcal{P}_{s_0}(x, dz) \right| \\ &\leq \frac{1}{2} \int_{y \in \widehat{\Gamma}} \int_{z \in \widetilde{\Gamma}} \left| \int_{w \in \widetilde{\Gamma}} (\mathcal{P}_{t-s_*-s_0}(w, dy) - \pi_*(dy)) \mathcal{P}_{s_*}(z, dw) \right| \mathcal{P}_{s_0}(x, dz) \\ &\leq \frac{1}{2} \int_{y \in \widehat{\Gamma}} \int_{z \in \widetilde{\Gamma}} \int_{w \in \Gamma_*} |\mathcal{P}_{t-s_*-s_0}(w, dy) - \pi_*(dy)| \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \\ &\quad + \frac{1}{2} \int_{y \in \widehat{\Gamma}} \int_{z \in \widetilde{\Gamma}} \int_{w \in \widetilde{\Gamma} \setminus \Gamma_*} (\mathcal{P}_{t-s_*-s_0}(w, dy) + \pi_*(dy)) \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \end{aligned}$$

where for the first equality we apply $\mathcal{P}_{s_0}(x, \tilde{\Gamma}) = 1$, and the second equality is due to the fact that $\tilde{\Gamma}$ is an invariant set. Applying Fubini's theorem, we further have

$$\begin{aligned}
\|\mathcal{P}_t(x, \cdot) - \pi_*(\cdot)\|_{\text{TV}} &\leq \frac{1}{2} \int_{z \in \tilde{\Gamma}} \int_{w \in \Gamma_*} \int_{y \in \hat{\Gamma}} |\mathcal{P}_{t-s_*-s_0}(w, dy) - \pi_*(dy)| \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \\
&\quad + \frac{1}{2} \int_{z \in \tilde{\Gamma}} \int_{w \in \tilde{\Gamma} \setminus \Gamma_*} \int_{y \in \hat{\Gamma}} (\mathcal{P}_{t-s_*-s_0}(w, dy) + \pi_*(dy)) \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \\
&= \frac{1}{2} \int_{z \in \tilde{\Gamma}, w \in \Gamma_*} \|\mathcal{P}_{t-s_*-s_0}(w, \cdot) - \pi_*(\cdot)\|_{\text{TV}} \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \\
&\quad + \int_{z \in \tilde{\Gamma}} \int_{w \in \tilde{\Gamma} \setminus \Gamma_*} \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \\
&= \frac{1}{2} \int_{z \in \tilde{\Gamma}, w \in \Gamma_*} \|\mathcal{P}_{t-s_*-s_0}(w, \cdot) - \pi_*(\cdot)\|_{\text{TV}} \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \\
&\quad + \int_{z \in \tilde{\Gamma}} (1 - \mathcal{P}_{s_*}(z, \Gamma_*)) \mathcal{P}_{s_0}(x, dz) \\
&\leq \frac{1}{2} \int_{z \in \tilde{\Gamma}, w \in \Gamma_*} \|\mathcal{P}_{t-s_*-s_0}(w, \cdot) - \pi_*(\cdot)\|_{\text{TV}} \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) + \epsilon
\end{aligned} \tag{52}$$

where the second inequality is because $\mathcal{P}_{s_*}(z, \Gamma_*) \geq 1 - \epsilon$ for all $z \in \tilde{\Gamma}$ by the definition of $\tilde{\Gamma}$. For any fixed $w \in \Gamma_*$, by (50) we have $\lim_{t \rightarrow \infty} \|\mathcal{P}_{t-s_*-s_0}(w, \cdot) - \pi_*(\cdot)\|_{\text{TV}} = 0$. Then since the TV distance is always bounded by 1, it follows from Dominated Convergence Theorem that there exists some $T > 0$ such that

$$\frac{1}{2} \int_{z \in \tilde{\Gamma}, w \in \Gamma_*} \|\mathcal{P}_{t-s_*-s_0}(w, \cdot) - \pi_*(\cdot)\|_{\text{TV}} \mathcal{P}_{s_*}(z, dw) \mathcal{P}_{s_0}(x, dz) \leq \epsilon, \quad \forall t \geq T. \tag{53}$$

Then combining (52) and (53), we get

$$\|\mathcal{P}_t(x, \cdot) - \pi(\cdot)\|_{\text{TV}} \leq 2\epsilon, \quad \forall t \geq T.$$

Since ϵ is arbitrary, we conclude that

$$\lim_{t \rightarrow \infty} \|\mathcal{P}_t(x, \cdot) - \pi(\cdot)\|_{\text{TV}} = 0, \quad \forall x \in \hat{\Gamma}, \tag{54}$$

so the distribution of $Y(t)$ converges in TV distance to a unique stationary distribution π_* on Γ .

Finally, by data processing inequality, the distribution of $\bar{Y}(t) = \frac{Y(t)}{\|Y(t)\|_2}$ also converges in TV distance to the distribution π on Γ_1 which is uniquely induced by π_* . This completes the proof. \square

G Auxiliary Results

The following lemma is a classic result. See, e.g., Corollary 8.5.5 in Oksendal [79].

Lemma G.1 (Time-changed Brownian motion). *Let $g : [0, \infty) \rightarrow [0, \infty)$ be a differentiable function with nonnegative first-order derivative. Then it holds that*

$$W(g(t)) \stackrel{d}{=} \int_{s=0}^t \sqrt{g'(t)} dW(s).$$

The multidimensional martingale functional central limit theorem can be found in, e.g., Theorem 7.1 in Ethier and Kurtz [80] or Theorem 2.1 in Whitt [81]. We use a version of Theorem 2.1 in Whitt [81] given in Theorem G.2.

Theorem G.2 (Multidimensional martingale FCLT, [81]). *For $n \geq 1$, let $M_n = (M_{n,1}, \dots, M_{n,k})$ be a martingale satisfying $M_n(0) = (0, \dots, 0)$. Let $C \in \mathbb{R}^{k \times k}$ be a semi-positive definite matrix. Suppose the following holds:*

- The expected value of the maximum jump in M_n is asymptotically negligible, i.e., for each $T > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{E}[J(M_n, T)] = 0 \quad (55)$$

where $J(x, T) = \sup\{|x(t) - x(t-)| : 0 < t \leq T\}$ for any function $x : [0, \infty) \rightarrow \mathbb{R}$.

- For each pair $(i, j) \in [k] \times [k]$ and each $t > 0$,

$$[M_{n,i}, M_{n,j}](t) \Rightarrow C_{ij}t \text{ as } n \rightarrow \infty. \quad (56)$$

Then $M_n \Rightarrow \sqrt{C}W$ in Skorokhod metric where W is the k -dimensional standard Brownian motion.

Lemma G.3 (Problem 7, Section 5, Pollard [82]). *If $X_n \Rightarrow X$ in the Skorokhod metric, and X has sample paths in $\mathcal{C}_{\mathbb{R}^D}[0, \infty)$, then $X_n \Rightarrow X$ in the uniform metric.*