

A Local Convergence Theory for the Stochastic Gradient Descent Method in Non-Convex Optimization with Non-Isolated Local Minima

Taehee Ko * ¹ and Xiantao Li ^{† 1}

¹Department of Mathematics, Pennsylvania State University, University Park PA 16802, USA

Abstract. Loss functions with non-isolated minima have emerged in several machine-learning problems, creating a gap between theoretical predictions and observations in practice. In this paper, we formulate a new type of local convexity condition that is suitable to describe the behavior of loss functions near non-isolated minima. We show that such a condition is general enough to encompass many existing conditions. In addition, we study the local convergence of the stochastic gradient descent (SGD) method under this mild condition by adopting the notion of stochastic stability. In the convergence analysis, we establish concentration inequalities for the iterates in SGD, which can be used to interpret the empirical observation from some practical training results.

Keywords:

Stochastic Gradient Descent,
Stochastic Stability,
Non-Convex Optimization,
Local Convergence,
Non-Isolated Minima.

Article Info.:

Volume: X
Number: X
Pages: 1 - 23
Date: /2023
doi.org/10.4208/jml.230106

Article History:

Received: 06/01/2023
Accepted: 03/05/2023

Communicated by:
Weinan E

1 Introduction

The stochastic gradient descent (SGD) and its variants have been predominantly applied in machine learning [1] due to the overall computational efficiency and robustness. In a typical training task, the objective function $f(\mathbf{x})$ is expressed as an empirical risk,

$$f(\mathbf{x}) = \mathbb{E}[f(\mathbf{x}; \omega)]. \quad (1.1)$$

Here, the randomness of ω stems from the sampling of training data set. The standard SGD updates the iterations by sampling the gradient $\nabla f(\mathbf{x}_n, \omega_n)$. With this unbiased estimator, we let $\xi_n := \nabla f(\mathbf{x}_n, \omega_n) - \nabla f(\mathbf{x}_n)$ be the noise induced by the sampling of the gradient. Then the iteration formula can be expressed as follows:

$$\text{SGD} : \mathbf{x}_{n+1} = \mathbf{x}_n - a_n(\nabla f(\mathbf{x}_n) + \xi_n), \quad (1.2)$$

where we have adopted the notation from [28], and for simplicity, we set the batch size to be 1. Here, a_n is known as the learning rate.

In the case when f is convex, convergence properties for SGD have been well established, examples include the stepsize policy [1], comparison to stochastic averaging methods [29], validation analysis [10], etc. On the other hand, a remarkable observation, as

*Corresponding author. tuk351@psu.edu.

[†]xiantao.li@psu.edu.

demonstrated in many recent studies, is that the stochastic gradient algorithms still perform well in non-convex optimization, e.g. the training of neural networks [4, 12, 32], and such success has intrigued many theoretical developments to explore convergence properties for non-convex optimization problems. It is important though, to point out that most of these results are obtained under some global assumptions, such as the global Lipschitz constant [14, 25, 27] or global Hölder constant [22] for the gradient or the globally bounded variance of the noise [18].

Despite these important theoretical results, training tasks with actual non-convex loss functions have exhibited many issues that are still difficult to interpret based on these analyses. For instance, the performance of training neural networks can be very sensitive to the initialization [8, 11, 24]. Intuitively, upon convergence, the behavior of SGD is largely determined by local properties. However, many of the existing results on non-convex optimizations are established under global assumptions, which may not hold for many practical training tasks, e.g., the lack of a global Lipschitz constant [16, 38], and globally bounded variance of the noise [13]. One recent development toward addressing this issue is the work [28] which proved the convergence of SGD under local assumptions on the initialization, the gradient, and the noise, but in the case of isolated local minima.

In addition to the aforementioned issues, another interesting issue in non-convex optimization is the presence of non-isolated minima, as observed in [3, 6, 9, 35]. Despite the clear intuition from Fig. 1.1, formulating local convexity conditions around such a set to precisely describe the behavior of the loss function in (1.1) is a non-trivial task. More importantly, it creates an important gap between practical optimization and theoretical analysis. This important scenario has received little attention until the recent analysis under some local convexity assumptions [7, 33]. Meanwhile, questions still remain as to whether more general characterizations exist for a wider variety of loss functions with non-isolated local minima, and more importantly, whether the convergence SGD can still be established for these general optimization problems. From a practical viewpoint, such convergence analysis will help interpret many observations from training algorithms, e.g., the roles of hyperparameters in stochastic optimization.

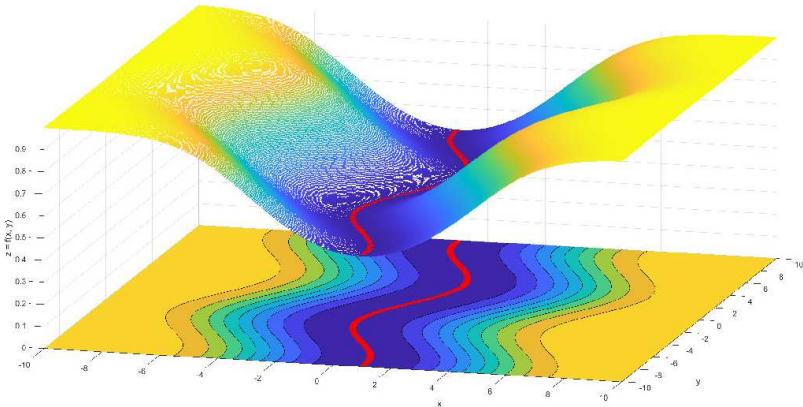


Figure 1.1: The landscape of a non-convex function with non-isolated global minima (red).

The main objective of this paper is two-fold: First, we introduce locally-defined mathematical conditions for the loss function with non-isolated minima. We establish relationships among such conditions, including those proposed in prior works [7, 15, 18]. Second, we prove that non-isolated minima still have a finite probability of attracting SGD iterates (1.2). Specifically, we establish concentration properties of SGD iterations, which lead to convergence to \mathcal{X} in probability and complexity bounds. These results are obtained only with local conditions, which sharpens results from those analyses obtained under global conditions. In the case where \mathcal{X} is compact, under more relaxed assumptions than those in [7], we prove a more general concentration inequality of SGD. Furthermore, even under the same assumption as in [7], our concentration inequalities are improved.

Specifically, we summarize our contributions as follows:

- We propose several conditions that describe loss functions with non-isolated minima, including a locally-defined restrict secant inequality (LRSI), the weak quasar-convexity (WQC) the non-negative support (NNS) conditions. Their mathematical expressions, as well as those from the existing literature, are summarized in Table 1.1.
- We propose a new approach to analyze SGD iterations. The approach combines the Lyapunov function approach [20] based on the optional stopping theorem [39], and yields nearly sharp theoretical results under only locally defined conditions. With the help of these results, we explicitly formulate the probability of convergence when the iteration from SGD starts near a set of non-isolated minima. This provides theoretical support for several empirical observations from the training of neural networks.
- We show the concentration properties of SGD that are also applicable to the regime of small batch sizes for a large class of landscapes arising in non-convex optimization. As an example, the result qualitatively elucidates the slowdown of SGD in a highly flat landscape near global minima. Additionally, we found that convergence properties of SGD near global minima are still very similar to those from strongly convex functions when the local landscape of non-isolated global minima satisfies the LRSI condition.

Table 1.1: The concentration results by previous works and this work under local conditions: the Local Strong Convexity (LSC), a locally-defined Polyak-Łojasiewicz condition (PL*), the Hession of the constant positive rank (HCPRC) on a Compact submanifold of global minima, a locally-defined restrict secant inequality (LRSI), the weak quasar-convexity (WQC) as an extension of the quasar convexity (QC) in [15] and the non-negative support (NNS). In these conditions, x^* refers to a local minimum, and x_p denotes a metric projection of x .

Condition	Definition	Concentration Inequality
LSC ([28])	$(\nabla f(x), x - x^*) \geq \mu \ x - x^*\ ^2$	$\mathbb{P}\{\ x - x^*\ ^2 \geq \epsilon\} \leq \mathcal{O}(1/(\epsilon n^\beta))$
PL* ([26])	$\ \nabla f(x)\ ^2 \geq \mu(f(x) - f^*)$	$\mathbb{P}\{f(x) - f^* \geq \epsilon\} \leq \mathcal{O}(\rho^n / \epsilon)$
HCPRC ([7])	$\text{rank}(\nabla^2 f(x)) = d - \delta$	$\mathbb{P}\{f(x) - f^* \geq \epsilon\} \leq \mathcal{O}(1/(\epsilon^2 n^\beta) + n^{1-\beta})$
LRSI (this work)	$(\nabla f(x), x - x_p) \geq \mu \ x - x_p\ ^2$	$\mathbb{P}\{f(x) - f^* \geq \epsilon\} \leq \mathcal{O}(1/(\epsilon n^\beta))$
WQC (this work)	$(\nabla f(x), x - x_p) \geq \zeta(f(x) - f^*)$	$\mathbb{P}\{f(x) - f^* \geq \epsilon\} \leq \mathcal{O}(1/(\epsilon n^{1-\beta}))$
NNS (this work)	$(\nabla f(x), x - x_p) \geq h(x) \geq 0$	$\mathbb{P}\{h(x) \geq \epsilon\} \leq \mathcal{O}(1/(\epsilon n^{1-\beta}))$

2 Preliminaries

2.1 Notation

Throughout the paper, $K \subset \mathbb{R}^d$ denotes a compact set unless it is stated otherwise. d represents the dimension of the parameter space. We use $\|\cdot\|$ for the ℓ_2 norm in \mathbb{R}^d . $[N]$ is the set of integers $[N] := \{1, 2, \dots, N\}$. For the complexity estimates, we define big- \mathcal{O} notation as follows: $a_n = \mathcal{O}(b_n)$ if $\limsup_n |a_n/b_n| < \infty$.

Next, we introduce the notions of the stable path and the stochastic stability for discrete stochastic processes.

Definition 2.1 (Stable Path). *With an initial iterate $x_1 \in K$, a realization $\{x_n\}_{n=2}^\infty$ from a stochastic algorithm is called a stable path if*

$$x_n \in K \quad \text{for all } n \geq 2. \quad (2.1)$$

Since the update rule for SGD in (1.2) uses the previous information, the iteration forms a discrete stochastic process with a filtration $\{\mathcal{F}_n\}_{n=1}^\infty$. Formally, the stochastic stability is defined by the measure of the set of stable paths with respect to the σ -algebra $\mathcal{F}_\infty := \sigma(\cup_n \mathcal{F}_n)$ as follows.

Definition 2.2 (Stochastic Stability). *With an initial guess $x_1 \in K$, an iteration from a stochastic algorithm is said to be stable with probability at least $1 - \eta$, if the following inequality is satisfied:*

$$\mathbb{P}\{x_n \in K \text{ for all } n \geq 2 | x_1\} \geq 1 - \eta. \quad (2.2)$$

The notion of stochastic stability can be found in [21], p 31].

2.2 Assumptions

Our regularity assumption on the loss function f is that it has a locally Lipschitz gradient. For example, a large class of neural networks satisfy this condition.

Assumption 2.1. *The gradient of f in the iteration (1.2) is locally Lipschitz continuous for any compact set K , i.e., there exists a constant $L_K > 0$ such that,*

$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L_K \|\mathbf{x} - \mathbf{y}\| \quad \text{for all } \mathbf{x}, \mathbf{y} \in K.$$

We make an assumption on the noise in SGD (1.2) as follows.

Assumption 2.2. *The noise satisfies that*

(i) $\mathbb{E}[\xi(\mathbf{x}, \omega) | \mathbf{x}] = 0$ for any $\mathbf{x} \in \mathbb{R}^d$, (Unbiased Stochastic Gradient).

(ii) *There exists some $\sigma_K > 0$ for any compact set K such that*

$$\sup_{\mathbf{x} \in K} \mathbb{E} \left[\|\xi(\mathbf{x}, \omega)\|^2 | \mathbf{x} \right] \leq \sigma_K.$$

The second part implies that the variance of the noise is locally bounded. This assumption is equivalent to the local conditions in [7].

Next, we make the standard assumption on the learning rates [1, 34], as follows.

Assumption 2.3. *The learning rates satisfy that*

$$a_n > 0, \quad \sum_{n=1}^{\infty} a_n = \infty, \quad \sum_{n=1}^{\infty} a_n^2 < \infty. \quad (2.3)$$

3 Stochastic gradient descent with local convexity conditions

In this section, we start by introducing a condition that is mild, but still sufficient to describe non-convex loss function near local minima. We present the relationships between the new condition and other existing characterizations of such loss functions. Then, under such conditions, we prove the concentration inequalities of the SGD method (1.2).

3.1 Local convexity conditions and their relations

Before we present the analysis of the SGD iterations (1.2), we first examine a number of important local convexity conditions. An important emphasis will be placed on their relations, since results that are established under one condition can be directly extended to situations where a weaker condition holds.

We first introduce the notion of a metric projection, which will be used to pose a new local convexity condition.

Definition 3.1. *For a compact set \mathcal{X} , the set-valued mapping $\Pi_{\mathcal{X}} : \mathbb{R}^d \rightarrow 2^{\mathcal{X}}$ is defined by*

$$\Pi_{\mathcal{X}}(\mathbf{x}) = \{\mathbf{z} \in \mathcal{X} : \|\mathbf{x} - \mathbf{z}\| = \text{dist}(\mathbf{x}, \mathcal{X})\}.$$

This is known as the metric projection [37] with respect to the usual Euclidean distance in our case.

We will denote by \mathcal{X} , a compact subset of global minima with or without connectivity. In the following assumption, we denote by $N_r(\mathcal{X})$ the closure of a r -neighborhood of \mathcal{X} .

Assumption 3.1. *There exists a non-negative function $h : N_r(\mathcal{X}) \rightarrow \mathbb{R}$ such that*

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p) \geq h(\mathbf{x}) \quad (3.1)$$

for any $\mathbf{x} \in N_r(\mathcal{X})$ and any projection $\mathbf{x}_p \in \Pi_{\mathcal{X}}(\mathbf{x})$. Due to the non-negative function h , we refer to this condition as non-negative support (NNS) condition.

As a pictorial illustration, Fig. 3.1 shows the examples that fulfill Assumption 3.1.

As a starting point to analyze optimization algorithms for non-convex functions, numerous conditions have been previously proposed to describe a non-convex loss function near local minima: [the restricted secant inequality (RSI) [41]], [the Polyak-Łojasiewicz (PL), the Quadratic Growth (QG) [18]], [the Łojasiewicz [5]], [the Kurdyka-Łojasiewicz (KL) [36]], [the quasar convexity (QC) [15]] and [the star convexity (*C) [31]]. Some of

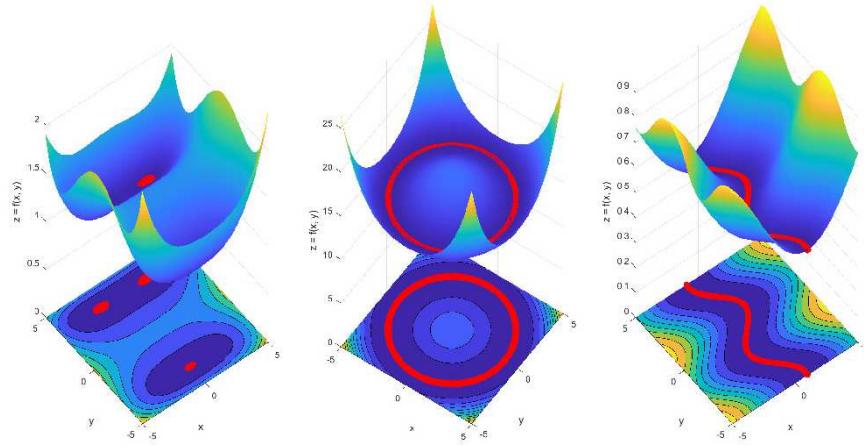


Figure 3.1: Examples of non-convex functions with isolated minima (left), or non-isolated minima (middle and right). The red regions indicate sets of local minima. The functions from all three figures satisfy the NNS condition (3.1).

these conditions were initially proposed to hold globally, but they can be easily relaxed to local conditions, in that they are assumed to hold in the neighborhood $N_r(\mathcal{X})$. Namely, for any $\mathbf{x} \in N_r(\mathcal{X})$,

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p) \geq \mu \|\mathbf{x} - \mathbf{x}_p\|^2 \quad (\text{LRSI}), \quad (3.2a)$$

$$\|\nabla f(\mathbf{x})\|^2 \geq \mu(f(\mathbf{x}) - f^*) \quad (\text{PL}^*), \quad (3.2b)$$

$$\|\nabla f(\mathbf{x})\|^\beta \geq \mu(f(\mathbf{x}) - f^*) \quad (\text{L}^*), \quad \beta \in (1, 2], \quad (3.2c)$$

$$\|\nabla f(\mathbf{x})\| \geq \varphi(f(\mathbf{x})) \quad (\text{KL}^*), \quad \varphi : \mathbb{R}_+ \mapsto \mathbb{R}_+ \text{ non-decreasing}, \quad (3.2d)$$

$$f(\mathbf{x}) - f^* \geq \mu \|\mathbf{x} - \mathbf{x}_p\|^2 \quad (\text{QG}^*), \quad (3.2e)$$

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^*) \geq f(\mathbf{x}) - f^* \quad (*\text{C}), \quad (3.2f)$$

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^*) \geq \zeta(f(\mathbf{x}) - f^*) \quad (\text{QC}), \quad (3.2g)$$

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p) \geq \zeta(f(\mathbf{x}) - f^*) \quad (\text{WQC}). \quad (3.2h)$$

Here $\mu > 0$, $0 < \zeta \leq 1$ and \mathbf{x}_p is a projection of \mathbf{x} onto \mathcal{X} .

Another interesting characterization is in terms of the local geometry [7]. Let $\mathcal{M} := \{\mathbf{x} : f(\mathbf{x}) = f^*\}$ be the set of all global minima and assume that there exists an open set $U \subset \mathbb{R}^d$ with some $\mathfrak{d} \in \{0, 1, \dots, d-1\}$ such that $\mathcal{M} \cap U$ is a non-empty \mathfrak{d} -dimensional C^2 -submanifold of \mathbb{R}^d with $f \in C^3(U)$ and

$$\text{rank}(\nabla^2 f(\mathbf{x})) = d - \mathfrak{d} \quad \text{for each } \mathbf{x} \in \mathcal{M} \cap U. \quad (3.3)$$

In particular, they also considered the case of such a submanifold being compact and without boundary [7, Section 6]. We will refer to this condition as the Hessian of the constant positive rank on a Compact submanifold (HCPRC) condition.

We now discuss the relations between various convexity conditions.

Remark 3.1. LRSI and WQC are special cases of the NNS condition since we can choose the following support functions respectively,

$$h(\mathbf{x}) = \mu \|\mathbf{x} - \mathbf{x}_p\|^2 \quad (\text{for LRSI}), \quad (3.4a)$$

$$h(\mathbf{x}) = \zeta(f(\mathbf{x}) - f^*) \quad (\text{for WQC}). \quad (3.4b)$$

Therefore, Assumption 3.1 encompasses very general local landscapes for $N_r(\mathcal{X})$ due to the flexibility of choosing $h(\mathbf{x})$. One can make an interpretation that $h(\mathbf{x})$ estimates how close the loss value $f(\mathbf{x})$ is to the minimum f^* by the mean value theorem,

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p) \approx f(\mathbf{x}) - f^*.$$

Now we turn to the relationship between NNS in (3.1) and PL* or QG* in (3.2). We first introduce a useful proposition of a projection map.

Proposition 3.1. *For a compact set $\mathcal{X} \subset \mathbb{R}^d$ and any $\mathbf{x} \in \mathbb{R}^d$, let $\mathbf{x}_p \in \Pi_{\mathcal{X}}(\mathbf{x})$. If $\mathbf{x} \neq \mathbf{x}_p$, then $\Pi_{\mathcal{X}}(\mathbf{y}) = \{\mathbf{x}_p\}$ for any $\mathbf{y} \in \{\mathbf{z} : \mathbf{z} = (1-t)\mathbf{x}_p + t\mathbf{x}, t \in (0, 1)\}$. That is, the projection of any \mathbf{y} between \mathbf{x} and \mathbf{x}_p onto \mathcal{X} is unique and equal to \mathbf{x}_p .*

In the following example, f represents functions whose local landscape in the vicinity of \mathcal{X} is highly flat, in which case (3.2c) (L*), (3.2d) (KL*), (3.2g) (WQC) and (3.2h) (NNS) conditions hold, but the conditions (3.2b) (PL*) and (3.2e) (QG*) fail.

Example 3.1 (A flat basin of attraction is neither PL* or QG* but is KL* and WQC). Let f be a non-convex, continuously differentiable function and satisfy that $\inf_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) = 0$. Let \mathcal{X} be a compact set. Assume that for a closed r -neighborhood, $N_r(\mathcal{X})$, and for some $C > 0$ and some $q > 2$,

$$f(\mathbf{x}) = C\|\mathbf{x} - \mathbf{x}_p\|^q \quad \text{for any } \mathbf{x} \in N_r(\mathcal{X}),$$

where \mathbf{x}_p is a projection from \mathbf{x} to \mathcal{X} . Due to the projection and $f \geq 0$, \mathcal{X} consists of global minima of $f(\mathbf{x})$. By Proposition 3.1, it follows that for any $\mathbf{x} \in N_r(\mathcal{X})$ and any $t \in [0, r]$,

$$f(\mathbf{x}_p + t\mathbf{v}) = Ct^q,$$

where \mathbf{v} is the unit direction from \mathbf{x}_p to \mathbf{x} . Then, by Lemma A.1, the function $f(\mathbf{x}_p + t\mathbf{v})$ in t is ζ -quasar-convex with some $\zeta > 0$, since it is continuously differentiable in $[0, r]$ and $f'(t) = 0$ only if $t = 0$. Finally, note that the form of such a function remains the same as is Ct^q regardless of $\mathbf{x}_p \in \mathcal{X}$. This implies the NNS condition with $h(\mathbf{x}) = \zeta f(\mathbf{x})$. As a result, $f(\mathbf{x})$ is WQC (3.2) as well as NNS in (3.1). Furthermore, a simple calculation yields that

$$\|\nabla f(\mathbf{x})\| = |(\nabla f(\mathbf{x}), \mathbf{v})| = Cq t_0^{q-1}, \quad (3.5)$$

where t_0 is defined by $\mathbf{x} = \mathbf{x}_p + t_0\mathbf{v}$. The first equality holds since f is defined by the distance $\|\mathbf{x} - \mathbf{x}_p\|$ and decreases the fastest towards \mathbf{x}_p at the point \mathbf{x} , which implies that \mathbf{v} and $\nabla f(\mathbf{x})$ are parallel. We note that by setting $\varphi(x) = Cqx^{1-1/q}$ for $q > 2$, the identity (3.5) can be rewritten as

$$\|\nabla f(\mathbf{x})\| = C^{\frac{1}{q}} q \varphi(f(\mathbf{x})).$$

Since φ is an increasing function, KL* (3.2) is satisfied and so is L*.

However, we note that $\|\nabla f(\mathbf{x})\|^2$ is of order t_0^{2q-2} , while $f(\mathbf{x})$ is of order t_0^q . Since $q > 2$, we have

$$\lim_{t \rightarrow 0} \frac{\|\nabla f(\mathbf{x}_p + t\mathbf{v})\|^2}{f(\mathbf{x}_p + t\mathbf{v})} = 0. \quad (3.6)$$

Thus, PL^* (3.2) does not hold even near \mathcal{X} . Similarly, QG^* (3.2) does not hold either, since

$$\lim_{t \rightarrow 0} \frac{f(\mathbf{x}_p + t\mathbf{v})}{t^2} = 0.$$

Now we summarize the relationships among the convexity conditions in the following diagram:

$$\begin{array}{ccc}
 (\text{HCPRC}) & \xrightarrow{(1)} & (\text{LRSI}) \xrightarrow{(2)} (\text{NNS}) \\
 & \Downarrow_{(3)} & \Uparrow_{(4)} \\
 (\text{QG}^*) \text{ or } (\text{PL}^*) & & (\text{WQC}) \\
 & \Downarrow_{(7)} & \Uparrow_{(5)} \\
 (\text{L}^*) & & (\text{QC}) \\
 & \Downarrow_{(6)} & \Uparrow_{(8)} \\
 (\text{KL}^*) & & (*\text{C})
 \end{array} \quad (3.7)$$

Figure 3.2: The inclusion relations among the convexity conditions. The proof of the implication of one condition to another: (1) [7, Lemma 14]; (2) Assumption 3.1 with $h(\mathbf{x}) = \mu\|\mathbf{x} - \mathbf{x}_p\|^2$; (3) [18, Appendix A], (Global conditions are given but the same technique also applies; (4) Assumption 3.1 with $h(\mathbf{x}) = \zeta(f(\mathbf{x}) - f^*)$; (5) WQC (3.2h) with $\mathcal{X} = \{\mathbf{x}^*\} \rightarrow \text{QC}$ (3.2g); (6) [15, Observation 3]; (7) L^* is a generalization of PL^* [5]; (8) KL^* is a generalization of L^* [36].

Note that (3) is the implication from LRSI to PL^* . The relationship between QG^* and LRSI and that between LRSI and WQC are open.

3.2 Main results on the convergence of SGD

The following theorem shows how likely the iterations from SGD will stay close to \mathcal{X} and eventually converge to \mathcal{X} if its local landscape satisfies Assumption 3.1.

Theorem 3.1 (Stochastic Stability and Probabilistic Convergence). *Suppose that there exists a closed r -neighborhood of a compact subset of global minima \mathcal{X} such that $N_r(\mathcal{X})$ satisfies Assumption 3.1. Under Assumptions 2.1 to 2.3, the following statements hold for the SGD iterations (1.2):*

(i) *For any initial $\mathbf{x}_1 \in N_r(\mathcal{X})$, the N iterates remain in $N_r(\mathcal{X})$ with a positive probability,*

$$\mathbb{P}\{\mathbf{x}_n \in N_r(\mathcal{X}) \text{ for each } n \in [N] | \mathbf{x}_1\} \geq 1 - C_N, \quad (3.8)$$

where

$$\begin{aligned} C_N &:= \frac{b_N}{r^2} \left(\text{dist}(\mathbf{x}_1, \mathcal{X})^2 + \frac{\sigma_r}{I} \sum_{n=1}^{N-1} \frac{a_n^2}{b_{n+1}} \right), \\ b_n &:= \prod_{j=1}^{n-1} \left(1 + L_r^2 a_j^2 \right), \quad b_1 := 1, \end{aligned} \quad (3.9)$$

L_r is the Lipschitz constant for $N_r(\mathcal{X})$ and I is the batch-size.

(ii) If $h(\mathbf{x}) > 0$ for each $\mathbf{x} \in N_r(\mathcal{X}) \setminus \mathcal{X}$ in Assumption 3.1, then

$$\mathbb{P} \left\{ \lim_{n \rightarrow \infty} \mathbf{x}_n = \mathbf{x}, f(\mathbf{x}) = f^* \mid \{\mathbf{x}_n\}_{n=1}^{\infty} \subset N_r(\mathcal{X}) \right\} = 1. \quad (3.10)$$

Consequently, these statements yield

$$\mathbb{P} \left\{ \lim_{n \rightarrow \infty} \mathbf{x}_n = \mathbf{x}, f(\mathbf{x}) = f^* \mid \mathbf{x}_1 \right\} \geq 1 - C_{\infty}.$$

Remark 3.2. The inequality explicitly shows the role of the hyperparameters (e.g. $\mathbf{x}_1, a_n, I, \sigma_r, L_r, r$) (3.8) in the stochastic stability of SGD in non-convex optimization. A qualitative interpretation is that if one starts with a good initialization and the initial guess \mathbf{x}_1 lies in a large nearly convex region (i.e., $r \gg 1$) near a set of global optima \mathcal{X} , then the optimization algorithm has little chance of leaving the region.

Remark 3.3. In terms of stochastic stability, we observe that C_{∞} in the definition (3.9) can be finite even for some small batch size, e.g. $I = 1$, as long as a_n is sufficiently small, since the learning rate a_n satisfies $\sum_n a_n^2 < \infty$ in Assumption 2.3. Thus, our result (3.8) theoretically ensures the regime of small constant batch size for SGD [12] with general learning rate a_n , which is in contrast to the similar result by [7] that requires a policy of increasing batch-size I as a polynomial order of the number of iterations N to keep the stability of the iteration.

Theorem 3.1 also allows us to examine the stochastic stability of SGD for both decreasing and constant learning rates, as follows.

Corollary 3.2 (Upper Bounds of Learning Rates for Non-Convex Loss). Under the assumptions in Theorem 3.1, the following statements hold for iterations of SGD:

(i) If the learning rate is decreasing, $a_n = a/n^{\beta}$ with $\beta \in (1/2, 1]$ and $n \geq 1$, the stochastic stability that $1 - C_{\infty} > 0$ in Theorems 3.1 is satisfied for any damping parameter $a > 0$ such that

$$e^{\frac{2\beta L_r}{2\beta-1} a^2} \left(\text{dist}(\mathbf{x}_1, \mathcal{X})^2 + \frac{2\beta \sigma_r a^2}{(2\beta-1)I} \right) < r^2. \quad (3.11)$$

(ii) If the learning rate is a constant, $a_n = a$ for $n \in [N]$, the stochastic stability that $1 - C_N > 0$ is satisfied if the learning rate a and the number of iteration N fulfill that

$$\left(1 + L_r^2 a^2 \right)^{N-1} \left(\text{dist}(\mathbf{x}_1, \mathcal{X})^2 + \frac{\sigma_r}{IL_r^2} \left(1 - \left(\frac{1}{1 + L_r^2 a^2} \right)^{N-1} \right) \right) < r^2. \quad (3.12)$$

Remark 3.4. In previous works, upper bounds for learning rates in the context of convex optimization have been studied for the optimization error in [1, 30]. As discussed in both works, upper bounds of learning rates are related to the optimal error $f(\mathbf{x}_n) - f^*$ and the condition number μ/L , where μ is the parameter in terms of the strong convexity and L is the global Lipschitz constant for the gradient. Corollary 3.2 suggests another type of upper bound pertaining to the stochastic stability in non-convex optimization.

In addition, Corollary 3.2 implies that for a large number of iterations ($N \gg 1$), using decreasing learning rates requires less effort than constant learning rates in terms of hyperparameters ($\mathbf{x}_1, L_r, \sigma_r, a$, etc.). Specifically, while the parameter a in the second term of (3.11) directly reduces the variance σ_r , such reduction in (3.12) occurs when a careful balance is struck between the number of iterations N and the batch-size I given a constant learning rate a .

Remark 3.5. When an empirical risk satisfies the NNS condition (3.1) for a compact set \mathcal{X} (not necessarily a set of local minima), and that $h(\mathbf{x}) > 0$ holds in (3.1) whenever $\mathbf{x} \in N_r(\mathcal{X}) \setminus \mathcal{X}$, the mean value theorem guarantees that \mathcal{X} is a setwise strict local minimum as defined in [23, Definition 3.4]. If the empirical risk is constructed under mild assumptions suggested in [23, Assumption 1], according to [23, Theorem 3.8], no suboptimal strict local minimum exists in \mathcal{X} .

Now, we consider the application of the SGD to the empirical risk with an initial guess in $N_r(\mathcal{X})$. Similar to the analysis in (4.15), we can achieve stochastic stability as the first result in Theorem 3.1 by giving an upper bound for the term $a_n^2 \|\nabla f(\mathbf{x}_n)\|^2$ in (4.15), for instance,

$$\left(\sup_{\mathbf{x} \in N_r(\mathcal{X})} \|\nabla f(\mathbf{x})\|^2 \right) a_n^2. \quad (3.13)$$

Note that this upper bound is still on the order of a_n^2 . Thus, a similar trick in our proof of Theorem 3.1 yields the stochastic stability of the SGD. By Lemma 4.3, the iterate converges to some stationary point of the empirical risk with a positive probability. On the other hand, [23, Theorem 3.8] guarantees that the stationary point is not a suboptimal local minimum.

The main idea for proving Theorem 3.1 will directly yield a concentration inequality of SGD to achieve $h(\mathbf{x}) \leq \epsilon$ as defined in Assumption 3.1. We provide more precise concentration properties of SGD for the non-isolated global minima, \mathcal{X} , under local conditions as follows.

Theorem 3.3 (Concentration Inequalities of SGD for Non-Convex Loss). *With the same assumptions and the notation in Theorem 3.1, SGD (1.2) with initialization $\mathbf{x}_1 \in N_r(\mathcal{X})$ satisfies that for any tolerance $\epsilon > 0$,*

$$\mathbb{P} \left\{ \min_{1 \leq n \leq N} h(\mathbf{x}_n) > \epsilon \mid \mathbf{x}_1 \right\} \leq \frac{1}{\epsilon} \cdot \frac{\text{dist}(\mathbf{x}_1, \mathcal{X})^2 + \sigma_r / I \sum_{n=1}^N a_n^2 / b_{n+1} + C_{N+1}}{2 \sum_{n=1}^N a_n / b_{n+1}}. \quad (3.14)$$

Let $a' > 0$ satisfy the inequality (3.11) for the case of decreasing learning rates. Then, for the learning rate $a_n = a/n^\beta$ with $\beta \in (1/2, 1)$,

(i) If f satisfies the WQC condition (3.2h) with $a \in (0, a')$, then

$$\mathbb{P} \left\{ \min_{1 \leq n \leq N} f(\mathbf{x}_n) - f^* > \epsilon | \mathbf{x}_1 \right\} \leq \mathcal{O} \left(\frac{1}{\epsilon N^{1-\beta}} \right) + C_{N+1}. \quad (3.15)$$

(ii) If f satisfies the HCPRC condition (3.3) with $a \in (0, \min\{1/2C, L_r/c^2, a'\})$ where C is defined in (4.11) and c in (4.23), then

$$\mathbb{P} \{ f(\mathbf{x}_N) - f^* > \epsilon | \mathbf{x}_1 \} \leq \mathcal{O} \left(\frac{1}{\epsilon N^\beta} \right) + C_{N+1}. \quad (3.16)$$

More generally, the same concentration result (3.16) is obtained if f satisfies the LRSI condition (3.2a).

Remark 3.6. The first result of Theorem 3.3 can quantitatively measure the slowdown of the convergence of SGD within a highly flat basin of attraction, which depends on the loss function f , the initialization \mathbf{x}_1 , the learning rate a_n , etc. As discussed after Assumption 3.1, the function h can be an implicit estimate for the optimization error. For example, in the case of a large flat landscape around global minima, h can be approximately equal to $\mathcal{O}(\|\mathbf{x} - \mathbf{x}_p\|^{q(\mathbf{x})})$ for some degree function $q(\mathbf{x}) \gg 1$ depending on \mathbf{x} . In this case, to reach a tolerance $\epsilon > 0$, SGD needs a very large number of iterations, since the right-hand side of (3.14) is of order $\mathcal{O}(1/\epsilon^q N^{1-\beta})$ for some large degree $q \gg 1$. For a general WQC (3.2h) loss function, we can make the same interpretation with (3.15).

On the other hand, if the loss function f starts with good initialization near which a submanifold of global minima fulfills the HCPRC condition (3.3), then SGD shows a better convergence rate as shown in (3.16) with an order of $\mathcal{O}(1/\epsilon N^\beta)$. As β approaches 1, the rate of convergence becomes sublinear, $\mathcal{O}(1/\epsilon N)$ with the first term (The second term accounts for the stochastic stability (3.8)). If we focus on the first term in (3.16), then the result (3.16) generalizes the convergence result for strongly convex functions with isolated minima [17, 29]. More importantly, our result (3.16) improves the bound in [7] whose concentration inequality involves the term $\mathcal{O}(N^{1-\beta})$ under equivalent assumptions.

4 Convergence analysis

In this section, we give the proof of the main result. First, we briefly review the notation for stopped stochastic processes introduced in [20, Section 4.5].

4.1 Stopped stochastic processes

Let \mathbf{x}_1 be the initialization and $\{(\mathbf{x}_n, \mathcal{F}_n)\}_{n=1}^\infty$ be the iteration from SGD or a stochastic process with a filtration. Let $V(\cdot), k(\cdot)$ be real-valued and non-negative functions on \mathbb{R}^d . Especially, $V(\cdot)$ will represent a Lyapunov function. In addition, we will denote a perturbed Lyapunov function by $V_n(\mathbf{x}_n)$ and a non-negative function scaled with learning rates by $k_n(\mathbf{x}_n)$. More importantly, any function or stochastic process with the tilde superscripted

will stand for a modified function in conjunction with a stopped process, respectively, which depends on a stopping time, specifically

$$\tilde{x}_n := \begin{cases} x_n, & n \leq \tau, \\ x_\tau, & n > \tau, \end{cases} \quad \tilde{V}(\tilde{x}_n) := \begin{cases} V_n(x_n), & n \leq \tau, \\ V_\tau(x_\tau), & n > \tau, \end{cases} \quad \tilde{k}(x) := \begin{cases} k(x), & x \in K, \\ 0, & \text{otherwise} \end{cases} \quad (4.1)$$

with the stopping time $\tau := \{n \geq 1 : x_n \notin K\}$ for a given set K . Throughout the convergence analysis, we will consider K as some set of global minima and \mathbb{I}_A as the indicator function which values one on the event A and zero otherwise.

4.2 Outline of the proof

For the proof of Theorem 3.1, we first construct a recursive inequality in terms of the distance between x_n and \mathcal{X} , which essentially yields a supermartingale property as in Lemma 4.1,

$$\mathbb{E} \left[\text{dist}(x_{n+1}, \mathcal{X})^2 | x_n \in N_r(\mathcal{X}) \right] \lesssim \mathbb{E} \left[\text{dist}(x_n, \mathcal{X})^2 \right] - a_n h(x_n), \quad (4.2)$$

where h is a non-negative function as in Assumption 3.1 and a_n is the learning rate. This property brings the problem into the framework of stochastic stability [20]. In particular, a Lyapunov function V and a non-negative function k can be defined from $\text{dist}(x_n, \mathcal{X})^2$ and $h(x_n)$, respectively. Our next step is to use Lemma 4.1 to estimate the probability of divergence from $N_r(\mathcal{X})$, i.e., the probability,

$$\mathbb{P} \{ \text{dist}(x_n, \mathcal{X}) > r | x_1 \}. \quad (4.3)$$

This gives the first result in Theorem 3.1.

For the second result, we observe that any stable path $\{x_n\}$ converges by Lemma 4.4. In such an event, $k(x_n)$ converges and it remains to show that $k(x_n) \approx h(x_n)$ will converge to 0 if $h(x) > 0$ is assumed in the region $N_r(\mathcal{X}) \setminus \mathcal{X}$. We prove it by contradiction that $k(x_n)$ converges to a positive random variable with positive probability. On the other hand, telescoping the above inequality yields an inequality of the form

$$\begin{aligned} \infty &= \left(\sum_{n=N}^{\infty} a_n \right) \delta \mathbb{P} \left\{ k(x_n) > \delta \text{ for all } n \geq N, \{x_n\}_{n=2}^{\infty} \subset B_r(x^*) | x_1 \right\} \\ &\lesssim \text{dist}(x_1, \mathcal{X})^2 < \infty. \end{aligned} \quad (4.4)$$

The left-hand side is infinite by Lemma 4.2 and Assumption 2.3, while the right-hand side is certainly finite.

For Theorem 3.3, we use the telescoping trick above and obtain an inequality in expectation as follows:

$$\mathbb{E} \left[\min_{1 \leq n \leq N} h(x_n) \mathbb{I}_{E_N} | x_1 \right] \lesssim \frac{\text{dist}(x_1, \mathcal{X})^2}{\sum_{n=1}^N a_n}. \quad (4.5)$$

Note that E_N represents the event in which the iterations remain near \mathcal{X} up to the N -th step. By applying Markov's inequality, we arrive at the concentration inequality,

$$\mathbb{P}\left\{\min_{1 \leq n \leq N} h(\mathbf{x}_n) \leq \epsilon \text{ or } E_N \text{ does not occur} \mid \mathbf{x}_1\right\} \gtrsim 1 - \frac{\text{dist}(\mathbf{x}_1, \mathcal{X})^2}{\epsilon \sum_{n=1}^N a_n}.$$

However, by Theorem 3.1, we know an upper bound for the probability of the event that the iterations diverge. Combining these results, we find the estimate,

$$\mathbb{P}\left\{\min_{1 \leq n \leq N} h(\mathbf{x}_n) \leq \epsilon \mid \mathbf{x}_1\right\} \lesssim \frac{\text{dist}(\mathbf{x}_1, \mathcal{X})^2}{\epsilon \sum_{n=1}^N a_n} + C_{N+1}.$$

Thanks to the fact that the NNS condition (3.2h) is weaker than the WQC condition (3.2g), we set $h(\mathbf{x}) = \zeta(f(\mathbf{x}) - f^*)$ and obtain the inequality by using the integral test for $a_n = a/n^\beta$. For the case of the HCPR condition (3.3), it suffices to show the same result under the LRSI condition (3.2), since it is weaker by Lemma 4.6. Observing that PL^* is weaker than LRSI in (3.7), Lemma 4.4 gives the recursive inequality of a form

$$\mathbb{E}[V_{n+1} \mid \mathcal{F}_n] \lesssim (1 - a_n)V_n + a_n^2 \quad (4.6)$$

and by using Lemma 4.5, we obtain the result under the HCPRC condition.

4.3 Some useful lemmas and theorems

In this section, we provide some prior lemmas and theorems. Some of them are restated so that they can be directly used for the proofs of the main results in the next section.

Lemma 4.1 ([20, Theorem 5.1]). *Let $\{\mathbf{x}_n\}$ be a Markov chain on \mathbb{R}^d . Let $V(\cdot)$ be a non-negative real-valued function on \mathbb{R}^d and for a given $r > 0$, define the set $N_r := \{\mathbf{x} : V(\mathbf{x}) \leq r\}$. Suppose that for any $\mathbf{x} \in N_r$ and each $n \geq 1$,*

$$\mathbb{E}[V(\mathbf{x}_{n+1}) \mid \mathbf{x}_n] - V(\mathbf{x}) \leq -k(\mathbf{x}), \quad (4.7)$$

where $k(\mathbf{x}) \geq 0$ and is continuous on N_r . Then, for any $\mathbf{x} \in N_r$,

$$\mathbb{P}\left\{\sup_{v+1 \leq n < \infty} V(\mathbf{x}_n) > r \mid \mathbf{x}_v = \mathbf{x}\right\} \leq \frac{V(\mathbf{x})}{r}. \quad (4.8)$$

Lemma 4.2 ([19, Lemma E.1]). *Let $\{k_n\}$ be a non-negative sequence of random variables. If $\mathbb{P}\{\liminf_n k_n > \delta\} > 0$ for some $\delta > 0$, then there exists a natural number $N \geq 1$ such that*

$$\mathbb{P}\left\{k_n > \delta \text{ for all } n \geq N \mid \liminf_n k_n > \delta\right\} > 0. \quad (4.9)$$

Lemma 4.3 ([33, Theorem 1]). *Under Assumptions 2.1-2.3, for any $\mathbf{x}_1 \in K$, the iteration from SGD (1.2) satisfies the property*

$$\mathbb{P}\left\{\lim_{n \rightarrow \infty} \mathbf{x}_n = \mathbf{x}, \nabla f(\mathbf{x}) = 0 \mid \{\mathbf{x}_n\}_{n=2}^\infty \subset K, \mathbf{x}_1\right\} = 1.$$

That is, the iteration almost surely converges to a critical point in the event that it stays in K .

Lemma 4.4 ([33, Lemma 2 when ∇f is locally Lipschitz]). Denote by $E_n := \{\mathbf{x}_i \in N_r(\mathcal{X}), \forall i \in [n]\}$, the event that iterations are stable up to the n -th step. Under Assumptions 2.1-2.3, for $\delta > 0$, SGD (1.2) has the property that for any $n \geq 1$,

$$\begin{aligned} & \mathbb{E} [(f(\mathbf{x}_{n+1}) - f^*) \mathbb{1}_{E_{n+1}} | \mathcal{F}_n] \\ & \leq \left(f(\mathbf{x}_n) - f^* - a_n(1 - Ca_n) \|\nabla f(\mathbf{x}_n)\|^2 + \sigma_{r+\delta} Ca_n^2 \right) \mathbb{1}_{E_n}. \end{aligned} \quad (4.10)$$

Specifically, the constant is given by,

$$C = \frac{L_{r+\delta}}{2} + \frac{\sup_{\mathbf{x} \in N_r(\mathcal{X})} \|\nabla f(\mathbf{x})\|}{\delta}. \quad (4.11)$$

$\sigma_{r+\delta}$ is the constant associated with $N_{r+\delta}(\mathcal{X})$ by Assumption 2.2.

Lemma 4.5 ([2, Lemma 4]). Let $\{e_n\}$ be a non-negative sequence such that

$$e_{n+1} \leq \left(1 - \frac{C}{n^\beta}\right) e_n + \frac{C'}{n^{\beta'}}, \quad (4.12)$$

where $\beta \in (0, 1)$, $\beta < \beta'$ and $C, C' > 0$. Then, there exist constants $C'' > 0$ and $n_0 \in \mathbb{N}$ such that

$$e_n \leq \frac{C''}{n^{\beta'-\beta}} \quad \text{for each } n \geq n_0. \quad (4.13)$$

Lemma 4.6 ([7, Lemma 14 in the case of a compact submanifold]). If f satisfies the HCPRC condition (3.3), for any $\mathbf{x}_0 \in \mathcal{M} \cap U$, then the compact submanifold $\mathcal{X} := \mathcal{M} \cap U$ has the following property that for some $r > 0$ and $c > 0$:

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p) \geq c \|\mathbf{x} - \mathbf{x}_p\|^2 \quad (4.14)$$

for any $\mathbf{x} \in N_r(\mathcal{X})$.

4.4 Proofs of main results

Proof of Theorem 3.1. First, we derive a recursive inequality for the distance between \mathbf{x}_n from SGD (1.2) and \mathbf{x}^* as follows:

$$\begin{aligned} & \mathbb{E} [\text{dist}(\mathbf{x}_{n+1}, \mathcal{X})^2 | \mathbf{x}_n] \leq \mathbb{E} [\|\mathbf{x}_{n+1} - (\mathbf{x}_n)_p\|^2 | \mathbf{x}_n] \\ & \leq \text{dist}(\mathbf{x}_n, \mathcal{X})^2 - 2a_n(\nabla f(\mathbf{x}_n), \mathbf{x}_n - (\mathbf{x}_n)_p) + a_n^2 \|\nabla f(\mathbf{x}_n)\|^2 + \sigma_r a_n^2 \\ & \leq (1 + L_r^2 a_n^2) \text{dist}(\mathbf{x}_n, \mathcal{X})^2 - 2a_n(\nabla f(\mathbf{x}_n), \mathbf{x}_n - (\mathbf{x}_n)_p) + \sigma_r a_n^2 \end{aligned} \quad (4.15)$$

for any $\mathbf{x}_n \in N_r(\mathcal{X})$. Here, L_r is the local Lipschitz constant from Assumption 2.1 and σ_r can be chosen by Assumption 2.2. The first inequality holds by the definition of the distance to a set from a point. The second inequality is a direct calculation using (1.2) and Assumption 2.2. The conditional expectation of the cross term with $\nabla f(\mathbf{x}_n)$ and ξ_n vanishes by Assumption 2.2. In the last inequality, we used the Lipschitz condition in Assumption 2.1.

To proceed, let us define the following:

$$b_n := \prod_{j=1}^{n-1} (1 + L_r^2 a_n^2), \quad b_1 := 1, \quad (4.16a)$$

$$V_n(\mathbf{x}_n) := \frac{\text{dist}(\mathbf{x}_n, \mathcal{X})^2}{b_n} + \sigma_K \sum_{j=n}^{N-1} \frac{a_j^2}{b_{j+1}}, \quad (4.16b)$$

$$k_n(\mathbf{x}_n) := \frac{a_n}{b_{n+1}} k(\mathbf{x}_n), \quad k(\mathbf{x}_n) := 2(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p). \quad (4.16c)$$

Based on the stochastic stability analysis [20, Theorem 5.1], by setting the stopping time $\tau := \{n \geq 2 : \mathbf{x}_n \notin N_r(\mathcal{X})\}$ with $K = N_r(\mathcal{X})$ in the inequality (4.1), we can modify the inequality (4.15) as follows:

$$\mathbb{E}[\tilde{V}_{n+1}(\tilde{\mathbf{x}}_{n+1}) | \mathcal{F}_n] \leq \tilde{V}_n(\tilde{\mathbf{x}}_n) - \tilde{k}_n(\tilde{\mathbf{x}}_n) \leq \tilde{V}_n(\tilde{\mathbf{x}}_n), \quad (4.17)$$

which implies that $\{\tilde{V}_n(\tilde{\mathbf{x}}_n)\}$ is a non-negative supermartingale. Furthermore, by applying Markov's inequality to this supermartingale, we find that for any $\mathbf{x}_1 \in N_r(\mathcal{X})$,

$$\mathbb{P}\left\{\sup_{N \geq n \geq 2} \tilde{V}_n(\tilde{\mathbf{x}}_n) > \frac{r^2}{b_N} \mid \mathbf{x}_1\right\} \leq \frac{b_N}{r^2} (V_1(\mathbf{x}_1)).$$

On the other hand, we observe that

$$\begin{aligned} & \mathbb{P}\left\{\sup_{N \geq n \geq 2} \text{dist}(\mathbf{x}_n, \mathcal{X}) > r \mid \mathbf{x}_1\right\} \\ & \leq \mathbb{P}\left\{\sup_{N \geq n \geq 2} \text{dist}(\tilde{\mathbf{x}}_n, \mathcal{X}) > r \mid \mathbf{x}_1\right\} \\ & \leq \mathbb{P}\left\{\sup_{N \geq n \geq 2} \tilde{V}_n(\tilde{\mathbf{x}}_n) > \frac{r^2}{b_N} \mid \mathbf{x}_1\right\}. \end{aligned}$$

The first inequality follows by the definition of the stopping time τ in the above. The second inequality can be deduced from the definition (4.16). Therefore, the first statement follows by combining these two inequalities. Moreover, Assumption 2.3 guarantees the limiting case $N = \infty$.

For the second statement, we telescope the inequality (4.17) and recall the definition $k_n(\mathbf{x}_n)$ in (4.16)

$$\sum_{n=1}^{\infty} \frac{a_n}{b_{n+1}} \mathbb{E}[\tilde{k}(\tilde{\mathbf{x}}_n) | \mathbf{x}_1] \leq V_1(\mathbf{x}_1) < \infty. \quad (4.18)$$

Note that conditioned on $\mathbf{x}_1 \in N_r(\mathcal{X})$, the event that $\{\mathbf{x}_n\}_{n=2}^{\infty} \subset N_r(\mathcal{X})$ occurs with some positive probability thanks to the first statement. In this event, $\{\mathbf{x}_n\}$ converges in $N_r(\mathcal{X})$ with probability 1 by Lemma 4.3. This guarantees $f(\mathbf{x}_n)$ converges to some random variable by the continuity of f , which is less than or equal to the global minimum f^* . In fact, we show that $f(\mathbf{x}_n)$ converges to f^* almost surely in this event.

Suppose on the contrary that $k(\mathbf{x}_n)$ in the definition (4.16) converges to a positive random variable with some positive probability. By the continuity of a measure, there exists a $\delta > 0$ such that

$$\mathbb{P} \left\{ \lim_{n \rightarrow \infty} k(\mathbf{x}_n) > \delta \mid \{\mathbf{x}_n\}_{n=2}^{\infty} \subset N_r(\mathcal{X}), \mathbf{x}_1 \right\} > 0.$$

By Lemma 4.2, there exists a $N \in \mathbb{N}$ satisfying

$$\mathbb{P} \left\{ k(\mathbf{x}_n) > \frac{\delta}{2} \text{ for all } n \geq N \mid \lim_{n \rightarrow \infty} k(\mathbf{x}_n) > \delta, \{\mathbf{x}_n\}_{n=2}^{\infty} \subset N_r(\mathcal{X}), \mathbf{x}_1 \right\} > 0.$$

These two inequalities imply that with some positive probability the event $k(\mathbf{x}_n) > \delta/2$ for all $n \geq N$ occurs when $\{\mathbf{x}_n\}_{n=2}^{\infty} \subset B_r(\mathbf{x}^*)$ given $\mathbf{x}_1 \in B_r(\mathbf{x}^*)$. However, by the definition of conditional probability, we have

$$\begin{aligned} & \mathbb{P} \left\{ k(\mathbf{x}_n) > \frac{\delta}{2} \text{ for all } n \geq N, \{\mathbf{x}_n\}_{n=2}^{\infty} \subset B_r(\mathbf{x}^*) \mid \mathbf{x}_1 \right\} \\ &= \mathbb{P} \left\{ k(\mathbf{x}_n) > \frac{\delta}{2} \text{ for all } n \geq N \mid \{\mathbf{x}_n\}_{n=2}^{\infty} \subset B_r(\mathbf{x}^*), \mathbf{x}_1 \right\} \\ & \quad \times \mathbb{P} \left\{ \{\mathbf{x}_n\}_{n=2}^{\infty} \subset B_r(\mathbf{x}^*) \mid \mathbf{x}_1 \right\} > 0 \end{aligned}$$

by the above inequality and the previous result of stability. Furthermore, by using the inequality (4.18), Assumption 2.2 and the Markov's inequality, one has,

$$\begin{aligned} \infty &= \left(\sum_{n=N}^{\infty} \frac{a_n}{b_{n+1}} \right) \frac{\delta}{2} \mathbb{P} \left\{ k(\mathbf{x}_n) > \frac{\delta}{2} \text{ for all } n \geq N, \{\mathbf{x}_n\}_{n=2}^{\infty} \subset B_r(\mathbf{x}^*) \mid \mathbf{x}_1 \right\} \\ &\leq V_1(\mathbf{x}_1) < \infty, \end{aligned}$$

which is a contradiction. Thus, in the event that $\{\mathbf{x}_n\}_{n=2}^{\infty} \subset N_r(\mathcal{X})$, $\lim_{n \rightarrow \infty} k(\mathbf{x}_n) = k(\mathbf{x}) = 0$ with probability 1. Finally, if $\mathbf{x} \in N_r(\mathcal{X}) - \mathcal{X}$, then the assumption that $h(\mathbf{x}) > 0$ for all $\mathbf{x} \in N(\mathcal{X}) - \mathcal{X}$ in the statement implies that

$$0 = \frac{k(\mathbf{x})}{2} = (\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p) \geq h(\mathbf{x}) > 0,$$

which is a contradiction. Therefore, the limit \mathbf{x} must lie in \mathcal{X} and $f(\mathbf{x}) = f^*$. \square

Proof of Corollary 3.2. By the integral test, we have

$$\sum_{n=1}^{\infty} \frac{1}{n^{2\beta}} \leq 1 + \int_1^{\infty} \frac{dx}{x^{\beta}} = \frac{2\beta}{2\beta - 1}.$$

With this, we use the inequality (3.11) and trace back to the condition of stability

$$r^2 > e^{\frac{2\beta L_r}{2\beta - 1} a^2} \left(\text{dist}(\mathbf{x}_1, \mathcal{X})^2 + \frac{2\beta \sigma_r}{2\beta - 1} a^2 \right)$$

$$\begin{aligned}
&\geq e^{L_r \sum_{n=1}^{\infty} a_n^2} \left(\text{dist}(\mathbf{x}_1, \mathcal{X})^2 + \sigma_r \sum_{n=1}^{\infty} a_n^2 \right) \\
&\geq b_{\infty} \left(\text{dist}(\mathbf{x}_1, \mathcal{X})^2 + \sigma_r \sum_{n=1}^{\infty} \frac{a_n^2}{b_{n+1}} \right) \implies 1 - C_{\infty} > 0,
\end{aligned}$$

where C_{∞} is defined in Theorem 3.1. In the third inequality, we used the well-known inequality that $\prod_n (1 + c_n) \leq e^{\sum_n c_n}$ for any real-valued sequence $\{c_n\}$ with $c_n > -1$.

In the case of a constant learning rate, we use the well-known formula

$$1 + x + \cdots + x^n = \frac{1 - x^n}{1 - x} \quad (4.19)$$

and a direct calculation shows the result (3.12). \square

Proof of Theorem 3.3. From the inequality (4.15) and its counterpart (4.17) with the stopped process, a telescoping trick can be used

$$\sum_{n=1}^N \mathbb{E} [k_n(\mathbf{x}_n) \mathbb{I}_{E_N} | \mathbf{x}_1] \leq \sum_{n=1}^N \mathbb{E} [\tilde{k}_n(\tilde{\mathbf{x}}_n) | \mathbf{x}_1] \leq V_1(\mathbf{x}_1). \quad (4.20)$$

By noting that $k_n(\mathbf{x}_n) \geq 2a_n/b_{n+1}h(\mathbf{x}_n)$ in (4.16), we have from above inequality

$$\mathbb{E} \left[\min_{1 \leq n \leq N} h(\mathbf{x}_n) \mathbb{I}_{E_N} | \mathbf{x}_1 \right] \leq \frac{V_1(\mathbf{x}_1)}{2 \sum_{n=1}^N a_n/b_{n+1}}, \quad (4.21)$$

where $V_1(\mathbf{x}_1)$ is defined in (4.16). This proves the first result.

For the rest of the statements in the theorem, we use the stability result in Theorem 3.1 as well as Markov's inequality. For simplicity, we keep using the above notations. By applying Markov's inequality to the inequality (4.20), we have

$$\mathbb{P} \left\{ \min_{1 \leq n \leq N} h(\mathbf{x}_n) > \epsilon \text{ and } E_N \text{ occurs } | \mathbf{x}_1 \right\} \leq \frac{V_1(\mathbf{x}_1)}{2\epsilon \sum_{n=1}^N a_n/b_{n+1}},$$

or equivalently,

$$\mathbb{P} \left\{ \min_{1 \leq n \leq N} h(\mathbf{x}_n) \leq \epsilon \text{ or } E_N \text{ does not occur } | \mathbf{x}_1 \right\} \geq 1 - \frac{V_1(\mathbf{x}_1)}{2\epsilon \sum_{n=1}^N a_n/b_{n+1}}.$$

However, according to the result in Theorem 3.1, we see that

$$\mathbb{P} \{ E_N \text{ does not occur } | \mathbf{x}_1 \} \leq C_{N+1},$$

which leads to

$$\mathbb{P} \left\{ \min_{1 \leq n \leq N} h(\mathbf{x}_n) \leq \epsilon | \mathbf{x}_1 \right\} \geq 1 - \frac{V_1(\mathbf{x}_1)}{2\epsilon \sum_{n=1}^N a_n/b_{n+1}} - C_{N+1}.$$

Thus, if f is WQC and $a_n = a/n^\beta$, then we have

$$\begin{aligned} & \mathbb{P}\left\{\min_{1 \leq n \leq N} f(\mathbf{x}_n) - f^* > \epsilon \mid \mathbf{x}_1\right\} \\ & \leq \frac{V_1(\mathbf{x}_1)}{2\zeta\epsilon \sum_{n=1}^N a_n/b_{n+1}} + C_{N+1} = \mathcal{O}\left(\frac{1}{\epsilon N^{1-\beta}}\right) + C_{N+1} \end{aligned} \quad (4.22)$$

by the integral test.

Secondly, we suppose that f is HCPRC (3.3). By Lemma 4.6, f is LRSI, that is, there exists a compact set of global minima \mathcal{X} , $r > 0$ and $c > 0$ such that $\mathbf{x}_0 \in \mathcal{X}$ and

$$(\nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}_p) \geq c\|\mathbf{x} - \mathbf{x}_p\|^2 \quad (4.23)$$

for any $\mathbf{x} \in N_r(\mathcal{X})$. This implies that

$$\|\nabla f(\mathbf{x})\| \geq c\|\mathbf{x} - \mathbf{x}_p\|$$

by the Cauchy-Schwarz inequality. By Lipschitz continuity in Assumption 2.1, we have

$$f(\mathbf{x}) \leq f^* + \frac{L_r}{2}\|\mathbf{x} - \mathbf{x}_p\|^2,$$

since $\nabla f(\mathbf{x}_p) = 0$. With these results, we achieve that for any $\mathbf{x} \in N_r(\mathcal{X})$,

$$f(\mathbf{x}) - f^* \leq \frac{L_r}{2}\|\mathbf{x} - \mathbf{x}_p\|^2 \leq \frac{L_r}{2c^2}\|\nabla f(\mathbf{x})\|^2,$$

which is the PL* condition (3.2b). Then, for learning rate $a_n \in (0, \min\{1/2C, L_r/c^2\})$ (C defined in (4.11) and c in (4.23)), by Lemma 4.4, we have

$$\mathbb{E}[(f(\mathbf{x}_{n+1}) - f^*)\mathbb{I}_{E_{n+1}} \mid \mathcal{F}_n] \leq \left(1 - \frac{c^2 a_n}{L_r}\right)(f(\mathbf{x}_n) - f^*)\mathbb{I}_{E_n} + \frac{\sigma_{r+\delta} C}{I} a_n^2, \quad (4.24)$$

$$\mathbb{E}[V_{n+1} \mid \mathbf{x}_1] \leq \left(1 - \frac{c^2 a_n}{L_r}\right)\mathbb{E}[V_n \mid \mathbf{x}_1] + \frac{\sigma_{r+\delta} C}{I} a_n^2 \quad (4.25)$$

by letting $V_n := (f(\mathbf{x}_n) - f^*)\mathbb{I}_{E_n}$ and taking expectation up to \mathbf{x}_1 .

In particular, by Lemma 4.5, if $a_n = a/n^\beta$ with $\beta \in (0, 1)$ and $a \in (0, \min\{1/2C, L_r/c^2\})$, there exist $C'' > 0$ and $N_0 \in \mathbb{N}$ such that

$$\mathbb{E}[V_N \mid \mathbf{x}_1] \leq \frac{C''}{N^\beta} \quad \text{for any } N \geq N_0$$

and by using Markov's inequality,

$$\mathbb{P}\{f(\mathbf{x}_N) - f^* > \epsilon \text{ and } E_N \text{ occurs} \mid \mathbf{x}_1\} \leq \frac{C''}{\epsilon N^\beta}.$$

As we did above, we achieve that for any $N \geq N_0$,

$$\mathbb{P}\{f(\mathbf{x}_N) - f^* \leq \epsilon \mid \mathbf{x}_1\} \geq 1 - \frac{C''}{\epsilon N^\beta} - C_{N+1}.$$

Finally, suppose that f satisfies the LRSI condition (3.2a). Then, based on the proof from the inequality (4.23) to the above concentration inequality, we can obtain the same concentration result for the case of the LRSI condition (3.2a). \square

5 Numerical results

Here we provide an example of a loss function, for which we verify the NNS condition (3.1) numerically. We use the pre-trained model, ResNet-56_cifar10¹, whose parameter dimension is about 0.85M, and obtain the density of states (DoS, or the spectral density) of the Hessian of the loss function at the local minimum x_* of the pre-trained model using the PyHessian [40]. We observe that most of the eigenvalues of the Hessian are close to 0 as shown in Fig. 5.1, indicating that in many directions, the landscape in the vicinity of x_* is flat. Following the procedure in [40], we pick a random direction v and verify that such a flat region near x_* exists as in Fig. 5.1. Within the flat region, we consider a sufficiently small segment \mathcal{X} centered at x_* , which can be viewed as a set of non-isolated local minima approximately.

We check whether the set \mathcal{X} fulfills the NNS condition (3.1) by sampling points z on \mathcal{X} and monitoring the sign of the quantity $(\nabla f(x), x - z)$ as in (3.1) for $x = z + tw$ with 50 sampled random directions w . The quantity $(\nabla f(x), x - z)$ is computed using a central difference method for loss function values. In the numerical experiments, with results shown in Fig. 5.2, we sample points uniformly from \mathcal{X} , which are of the form $x_* + av$ for $a = -5, -4, \dots, 5$ in the direction v used in Fig. 5.1. Each panel in Fig. 5.2 corresponds to points uniformly sampled from \mathcal{X} . We clearly see that most of those quantities $(\nabla f(x), x - z)$ are positive in the landscapes perturbed along 50 randomly sampled directions.

With the same procedure described above, we made similar observations in independent simulations for the given direction v with different sampled directions w defined at the points uniformly sampled from \mathcal{X} . In addition, we obtain similar results by following the same procedure with different random directions v .

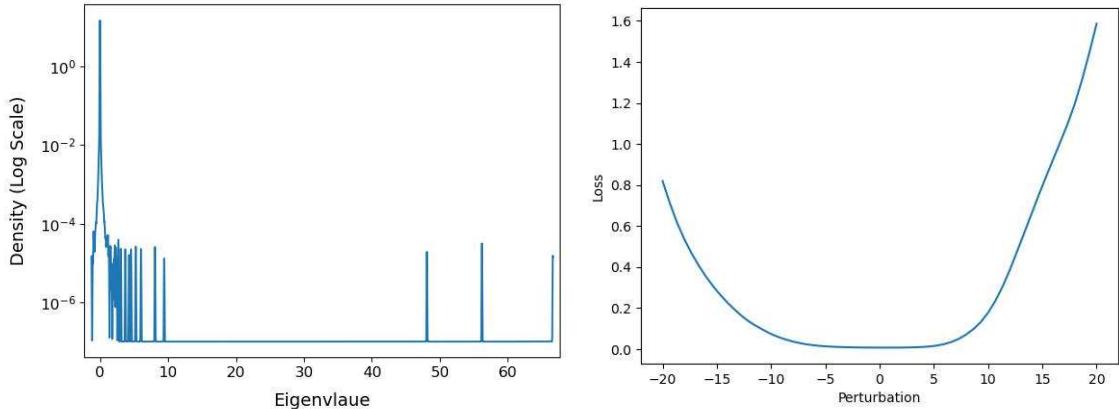


Figure 5.1: Numerical verification of the NNS condition (3.1). The left panel shows the density of states (DoS) of the Hessian of a loss function $f(x)$ defined by the pre-trained model, ResNet56_cifar10, at the local minimum x_* . Most of the eigenvalues are concentrated at 0 where the DoS is an order of 1, while at other positive eigenvalues, the DoS is below an order of 10^{-4} . In the right panel, a flat landscape of the loss function is observed in a direction v . The graph is obtained from the expression $f(x_* + av)$ with $\|v\| = 1$ and $a \in [-20, 20]$, where x_* denotes the local minimum of the pre-trained model.

¹The model is available at <https://pypi.org/project/pytorchcv/>

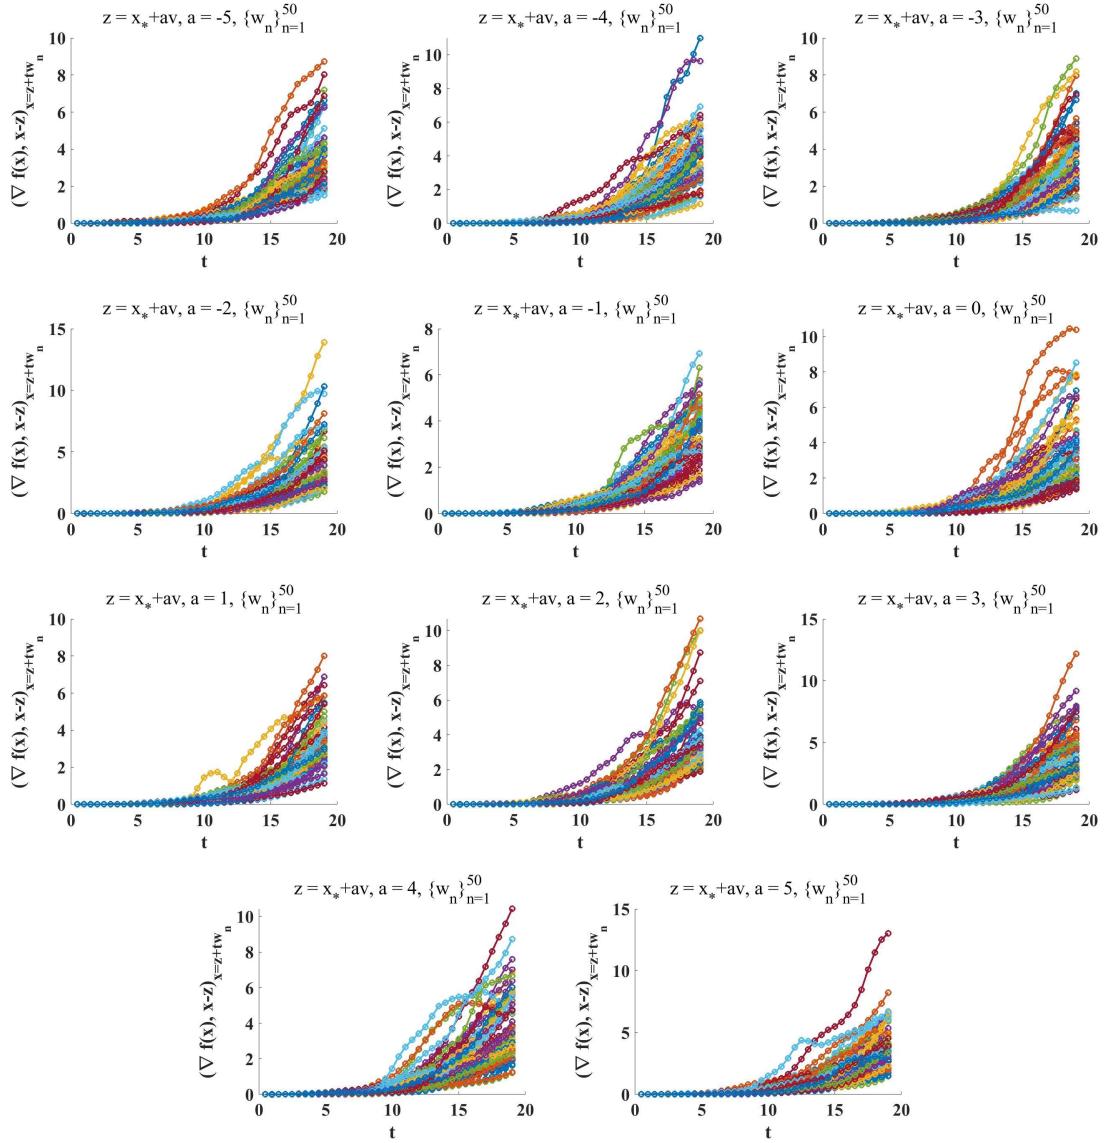


Figure 5.2: Numerical verification of the NNS condition (3.1). At each of the 11 points uniformly sampled from the flat region ($a \in [-5, 5]$) in Fig. 5.1, 50 random directions w are sampled, and the quantities $(\nabla f(x), x - z)$, where $x = z + tw$ for $t \in (0, 20)$, are plotted.

With the same procedure described above, we made similar observations in independent simulations for the given direction v with different sampled directions w defined at the points uniformly sampled from \mathcal{X} . In addition, we obtain similar results by following the same procedure with different random directions v . These numerical observations confirm that the loss function associated with the pre-trained model, ResNet56_cifar 10 satisfies the NNS condition (3.1).

6 Summary and discussions

This paper has focused on local convergence in the context of non-convex optimization, with the stochastic gradient descent (SGD), especially for loss functions with non-isolated minima. We have proved the convergence and concentration inequalities in terms of hyperparameters. The technical results rely on the stochastic stability analysis, the optional stopping theorem for discrete stochastic processes, and the geometric characterizations of non-isolated minima. An extension of our analysis to other variants of the stochastic gradient method is likely, and it is expected to provide sufficient conditions on the learning rates and the local Lipschitz constant to guarantee the convergence with high probability.

Appendix A

Proof of Proposition 3.1. To start, we show that $\Pi_{\mathcal{X}}(\mathbf{y})$ contains \mathbf{x}_p . Note that for any $\mathbf{z} \in \mathcal{X}$, $\|\mathbf{x} - \mathbf{x}_p\| \leq \|\mathbf{x} - \mathbf{z}\|$ by the definition of the projection. From this, the collinear relation between \mathbf{x} , \mathbf{y} and \mathbf{x}_p implies that

$$\|\mathbf{y} - \mathbf{x}_p\| + \|\mathbf{x} - \mathbf{y}\| = \|\mathbf{x} - \mathbf{x}_p\| \leq \|\mathbf{x} - \mathbf{z}\|.$$

Furthermore, by the triangle inequality,

$$\|\mathbf{y} - \mathbf{x}_p\| \leq \|\mathbf{x} - \mathbf{z}\| - \|\mathbf{x} - \mathbf{y}\| \leq \|\mathbf{y} - \mathbf{z}\|.$$

That is, $\Pi_{\mathcal{X}}(\mathbf{y})$ at least contains \mathbf{x}_p .

Now, we prove that this set is indeed a singleton. Suppose on the contrary that there is $\mathbf{y}_p \in \Pi_{\mathcal{X}}(\mathbf{y})$ and $\mathbf{y}_p \neq \mathbf{x}_p$. We show that \mathbf{y}_p does not lie on the line through \mathbf{x}_p and \mathbf{x} . First of all, \mathbf{y}_p cannot lie on $(\mathbf{x}_p, \mathbf{x})$. Otherwise, \mathbf{x}_p is not projection of \mathbf{x} . The other possibility is that $\mathbf{y}_p = \mathbf{x}_p + t(\mathbf{x} - \mathbf{x}_p)$ for some $t > 1$. However, we can see that

$$\|\mathbf{y} - \mathbf{x}_p\| < \|\mathbf{x} - \mathbf{x}_p\| \leq \|\mathbf{x} - \mathbf{y}_p\| < \|\mathbf{x} - \mathbf{y}_p\| + \|\mathbf{x} - \mathbf{y}\| = \|\mathbf{y} - \mathbf{y}_p\|.$$

In the second inequality, we recall the fact that \mathbf{x}_p is a projection of \mathbf{x} onto \mathcal{X} . This inequality is not true as opposed to the hypothesis that both sides must be equal as $\text{dist}(\mathbf{y}, \mathcal{X})$.

As a result, we can assume that \mathbf{y}_p is not on the line passing through \mathbf{x} , \mathbf{y} and \mathbf{x}_p . However, this results in the strict triangle inequality,

$$\|\mathbf{x} - \mathbf{y}_p\| < \|\mathbf{x} - \mathbf{y}\| + \|\mathbf{y} - \mathbf{y}_p\|.$$

This inequality, together with $\|\mathbf{y} - \mathbf{y}_p\| = \|\mathbf{y} - \mathbf{x}_p\| = \text{dist}(\mathbf{y}, \mathcal{X})$, leads to a contradiction to $\mathbf{x}_p \in \Pi_{\mathcal{X}}(\mathbf{x})$, i.e.,

$$\|\mathbf{x} - \mathbf{y}_p\| < \|\mathbf{x} - \mathbf{x}_p\|.$$

This completes the proof. □

Lemma A.1 ([15, Observation 1]). *Let $a < b$ and f be a real-valued continuously differentiable function on $[a, b]$. Then, f is unimodal, that is, $f'(c) \neq 0$ for all $c \in [a, b]$ such that $c \notin$*

$\operatorname{argmin}_{x \in [a,b]} f(x)$ if and only if f satisfies the QC for some $\zeta \in (0, 1]$ and some minimum $x^* \in [a, b]$,

$$f'(x)(x - x^*) \geq \zeta(f(x) - f^*).$$

Acknowledgments

We would like to acknowledge support for this project from the National Science Foundation (NSF Grant DMS-1953120).

References

- [1] L. Bottou, F. E. Curtis, and J. Nocedal, Optimization methods for large-scale machine learning, *SIAM Rev.*, **60**(2):223–311, 2018.
- [2] K. L. Chung, On a stochastic approximation method, *Ann. Math. Stat.*, 463–483, 1954.
- [3] Y. Cooper, Global minima of overparameterized neural networks, *SIAM J. Math. Data Sci.*, **3**(2):676–691, 2021.
- [4] A. Cutkosky and F. Orabona, Momentum-based variance reduction in non-convex SGD, *Adv. Neural. Inf. Process. Syst.*, **32**, 2019.
- [5] S. Dereich and S. Kassing, Convergence of stochastic gradient descent schemes for Łojasiewicz-landscapes, *arXiv:2102.09385*, 2021.
- [6] F. Draxler, K. Veschnini, M. Salmhofer, and F. Hamprecht, Essentially no barriers in neural network energy landscape. In: *International Conference on Machine Learning*, PMLR, 1309–1318, 2018.
- [7] B. Fehrman, B. Gess, and A. Jentzen, Convergence rates for the stochastic gradient descent method for non-convex objective functions, *J. Mach. Learn. Res.*, **21**:136, 2020.
- [8] S. Fort and A. Scherlis, The goldilocks zone: Towards better understanding of neural network loss landscapes. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, 3574–3581, 2019.
- [9] T. Garipov, P. Izmailov, D. Podoprikhin, D. P. Vetrov, and A. G. Wilson, Loss surfaces, mode connectivity, and fast ensembling of DNNs, *Adv. Neural. Inf. Process. Syst.*, **31**, 2018.
- [10] S. Ghadimi and G. Lan, Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization I: A generic algorithmic framework, *SIAM J. Optim.*, **22**(4):1469–1492, 2012.
- [11] X. Glorot and Y. Bengio, Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, JMLR Workshop and Conference Proceedings, 249–256, 2010.
- [12] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, Accurate, large minibatch SGD: Training imagenet in 1 hour, *arXiv:1706.02677*, 2017.
- [13] M. Gurbuzbalaban, U. Simsekli, and L. Zhu, The heavy-tail phenomenon in SGD. In: *International Conference on Machine Learning*, PMLR, 3964–3975, 2021.
- [14] F. Haddadpour, M. M. Kamani, M. Mahdavi, and V. Cadambe, Local SGD with periodic averaging: Tighter analysis and adaptive synchronization, *Adv. Neural. Inf. Process. Syst.*, **32**, 2019.
- [15] O. Hinder, A. Sidford, and N. Sohoni, Near-optimal methods for minimizing star-convex functions and beyond. In: *Conference on Learning Theory*, PMLR, 1894–1938, 2020.
- [16] L. Hodgkinson, U. Simsekli, R. Khanna, and M. W. Mahoney, Generalization properties of stochastic optimizers via trajectory analysis, *arXiv:2108.00781*, 2021.
- [17] P. Jain, D. Nagaraj, and P. Netrapalli, Making the last iterate of SGD information theoretically optimal. In: *Proceedings of the Annual Conference On Learning Theory (COLT)*, 1752–1755, 2019.
- [18] H. Karimi, J. Nutini, and M. Schmidt, Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, 795–811, 2016.

- [19] T. Ko and X. Li, Stochastic algorithms for self-consistent calculations of electronic structures, *Math. Comp.*, 2023.
- [20] H. Kushner and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, Vol. 35, Springer Science & Business Media, 2003.
- [21] H. J. Kushner, Stochastic stability and control, Technical Report, Brown Univ Providence RI, 1967.
- [22] Y. Lei, T. Hu, G. Li, and K. Tang, Stochastic gradient descent for nonconvex learning without bounded gradient assumptions, *IEEE Trans. Neural Netw. Learn. Syst.*, **31**(10):4394–4400, 2019.
- [23] D. Li, T. Ding, and R. Sun, On the benefit of width for neural networks: Disappearance of basins, *SIAM J. Optim.*, **32**(3):1728–1758, 2022.
- [24] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, Visualizing the loss landscape of neural nets, *Adv. Neural. Inf. Process. Syst.*, **31**, 2018.
- [25] X. Li and F. Orabona, On the convergence of stochastic gradient descent with adaptive stepsizes. In: *The 22nd International Conference on Artificial Intelligence and Statistics*, PMLR, 983–992, 2019.
- [26] C. Liu, L. Zhu, and M. Belkin, Loss landscapes and optimization in over-parameterized non-linear systems and neural networks, *Appl. Comput. Harmon. Anal.*, 2022.
- [27] L. Madden, E. Dall’Anese, and S. Becker, High probability convergence bounds for stochastic gradient descent assuming the Polyak-Łojasiewicz inequality, *arXiv:2006.05610*, 2020.
- [28] P. Mertikopoulos, N. Hallak, A. Kavis, and V. Cevher, On the almost sure convergence of stochastic gradient descent in non-convex problems, *Adv. Neural. Inf. Process. Syst.*, **33**:1117–1128, 2020.
- [29] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, Robust stochastic approximation approach to stochastic programming, *SIAM J. Optim.*, **19**(4):1574–1609, 2009.
- [30] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*, Vol. 87, Springer Science & Business Media, 2003.
- [31] Y. Nesterov and B. T. Polyak, Cubic regularization of Newton method and its global performance, *Math. Program.*, **108**(1):177–205, 2006.
- [32] B. Neyshabur, R. R. Salakhutdinov, and N. Srebro, Path-SGD: Path-normalized optimization in deep neural networks, *Adv. Neural. Inf. Process. Syst.*, **28**, 2015.
- [33] V. Patel, B. Tian, and S. Zhang, Global convergence and stability of stochastic gradient descent, *arXiv:2110.01663*, 2021.
- [34] H. Robbins and S. Monro, A stochastic approximation method, *Ann. Math. Stat.*, 400–407, 1951.
- [35] L. Sagun, U. Evci, V. U. Guney, Y. Dauphin, and L. Bottou, Empirical analysis of the Hessian of over-parametrized neural networks, *arXiv:1706.04454*, 2017.
- [36] K. Scaman, C. Malherbe, and L. Dos Santos, Convergence rates of non-convex stochastic gradient descent under a generic Łojasiewicz condition and local smoothness. In: *International Conference on Machine Learning*, PMLR, 19310–19327, 2022.
- [37] A. Shapiro, Existence and differentiability of metric projections in Hilbert spaces, *SIAM J. Optim.*, **4**(1): 130–141, 1994.
- [38] R. Sun, Optimization for deep learning: Theory and algorithms, *arXiv:1912.08957*, 2019.
- [39] D. Williams, *Probability with Martingales*, Cambridge University Press, 1991.
- [40] Z. Yao, A. Gholami, K. Keutzer, and M. W. Mahoney, PyHessian: Neural networks through the lens of the Hessian. In: *2020 IEEE International Conference on Big Data (Big data)*, IEEE, 581–590, 2020.
- [41] H. Zhang and W. Yin, Gradient methods for convex minimization: Better rates under weaker conditions, *arXiv:1303.4645*, 2013.