Global Convergence and Stability of Stochastic Gradient Descent

Vivak Patel
Department of Statistics
University of Wisconsin – Madison
Madison, WI 53706
vivak.patel@wisc.edu

Shushu Zhang
Department of Statistics
University of Michigan – Ann Arbor
shushuz@umich.edu

Bowen Tian
Department of Statistics
The Ohio State University
tian.837@buckeyemail.osu.edu

Abstract

In machine learning, stochastic gradient descent (SGD) is widely deployed to train models using highly non-convex objectives with equally complex noise models. Unfortunately, SGD theory often makes restrictive assumptions that fail to capture the non-convexity of real problems, and almost entirely ignore the complex noise models that exist in practice. In this work, we demonstrate the restrictiveness of these assumptions using three canonical models in machine learning. Then, we develop novel theory to address this shortcoming in two ways. First, we establish that SGD's iterates will either globally converge to a stationary point or diverge under nearly arbitrary nonconvexity and noise models. Under a slightly more restrictive assumption on the joint behavior of the non-convexity and noise model that generalizes current assumptions in the literature, we show that the objective function cannot diverge, even if the iterates diverge. As a consequence of our results, SGD can be applied to a greater range of stochastic optimization problems with confidence about its global convergence behavior and stability.

1 Introduction

Stochastic Gradient Descent (SGD) and its variants are dominant algorithms for solving stochastic optimization problems arising in machine learning, and have expanded their reach to more complex problems from estimating Gaussian Processes [Chen et al., 2020], covariance estimation in stochastic filters [Kim et al., 2021], and systems identification [Hardt et al., 2016, Zhang and Patel, 2020]. Accordingly, understanding the behavior of SGD and its variants has been crucial to their reliable application in machine learning and beyond. As a result, the theory of these methods has greatly advanced, most notably for SGD as it is the basis for, and simplest of, these methods. Indeed, SGD has been analyzed from many perspectives: global convergence analysis [Lei et al., 2019, Gower et al., 2020, Khaled and Richtárik, 2020, Mertikopoulos et al., 2020, Patel, 2021], local convergence analysis [Mertikopoulos et al., 2020], greedy and global complexity analysis [Gower et al., 2020, Khaled and Richtárik, 2020], asymptotic weak convergence [Wang et al., 2021], and saddle point analysis [Fang et al., 2019, Mertikopoulos et al., 2020, Jin et al., 2021].

While all of these perspectives add new dimensions to our understanding of SGD, the global convergence analysis of SGD is the foundation as it dictates whether local analyses, complexity analyses or saddle point analyses are even warranted. As surveyed in Patel [2021], these current global

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

convergence analyses of SGD make a wide variety of assumptions, most commonly: (1) the objective function is bounded from below, (2) the gradient function is globally Lipschitz continuous, (3) the stochastic gradients are unbiased, and (4) the variance of the stochastic gradients are bounded. While the first and third assumption are generally reasonable,[1] the second and fourth assumptions *and their more recent generalizations* are not usually applicable to machine learning problems as we now demonstrate through three simple examples. Note, in the first two examples, we make use of a penalty function, which can be removed without impacting the result.

Example 1: Feed Forward Neural Network. Consider the example (Y, Z) where Y is a binary label and Z is a feature vector. We will attempt to predict Y from Z using a simple multi-layer feed forward network as shown in Fig. 1. The next result states that for a simple distribution over the example space and for a simple, archetype network, the gradient function is not globally Lipschitz continuous, nor does it satisfy the (possibly) more general (L_0, L_1)-smooth assumption [Zhang et al., 2019, Definition 1, Assumption 3]. Moreover, the variance of the stochastic gradients is unbounded. See Appendix A.2 for a proof.

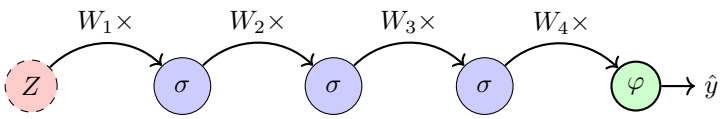

Figure 1: A diagram of a simple feed forward network for binary classification.

Proposition 1. *Consider the feed forward network in Fig. 1 with σ linear and φ sigmoid trained with a binary cross-entropy loss with a Ridge penalty. There exists a finite, discrete distribution for (Y, Z) such that the risk function's gradient is not globally Lipschitz continuous, nor does it satisfy the (L_0, L_1)-smooth assumption. Moreover, the variance of the stochastic gradients is not bounded.*

Example 2: Recurrent Neural Network. Consider the example (Y, Z_0, Z_1, Z_2, Z_3) where Y is a binary label and $\{Z_0, Z_1, Z_2, Z_3\}$ are sequentially observed. We will attempt to predict Y from Z using a simple recurrent network as shown in Fig. 2. The next result states that for a simple distribution over the example space and for a simple, archetype network, the training function violates the aforementioned assumptions. See Appendix A.3 for a proof.

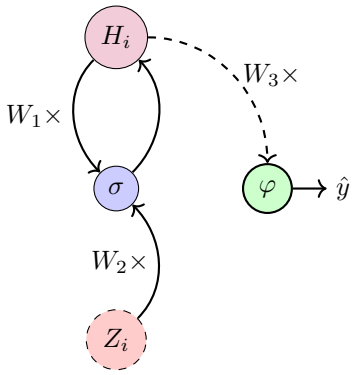

Figure 2: A diagram of a recurrent neural network for a binary classification task.

Proposition 2. *Consider the recurrent network in Fig. 2 with σ linear and φ sigmoid trained with a binary cross-entropy loss with a Ridge penalty. There exists a finite, discrete distribution for (Y, Z_0, Z_1, Z_2, Z_3) such that the risk function's gradient is not globally Lipschitz continuous, nor does it satisfy the (L_0, L_1)-smooth assumption. Moreover, the variance of the stochastic gradients is not bounded.*

Example 3: Poisson Regression. Consider fitting a Poisson regression model by the standard maximum likelihood framework using independent copies of the example (Y, Z), where Y is a count

[1] See Bottou et al. [2018] for a simple relaxation of unbiased stochastic gradients.

response variable and Z is a predictor. The next result states that for a very nice Poisson regression problem, the stochastic gradients violate the bounded variance assumption, its generalization [Bottou et al., 2018, Assumption 4.3c], and, in turn, its generalization, expected smoothness [Khaled and Richtárik, 2020, Assumption 2]. See Appendix A.4 for a proof.

Proposition 3. *Let Y and Z be independent Poisson random variables with mean one. Consider estimating a Poisson regression model of Y as a function of Z. Then, the gradient function is not globally Lipschitz continuous, nor does it satisfy the (L_0, L_1)-smooth assumption. Moreover, the variance of the stochastic gradients is not bounded, does not satisfy [Bottou et al., 2018, Assumption 4.3c], nor does it satisfy [Khaled and Richtárik, 2020, Assumption 2].*

As these three examples show, global convergence analyses that make use of the aforementioned assumptions do not apply to these canonical examples of machine learning problems. In fact, to our knowledge and as summarized in Table 1, *there are no global convergence analyses of SGD* that apply to these examples.

The Problem:
As a result, we do not have guarantees about how SGD behaves on these simple machine learning problems, which calls into question what SGD and its variants are doing on more complicated machine learning models.

Contributions. To address this problem,

1. We relax the global Lipschitz continuous gradient assumption, the bounded variance assumption, and their aforementioned generalizations to assumptions that are applicable to the examples above. Specifically, we relax the global Lipschitz continuity assumption and the (L_0, L_1)-smooth assumption to local α-Hölder continuity of the gradient for $\alpha \in (0, 1]$ (see Assumption 2), which is even a relaxation even for deterministic gradient algorithms (c.f. Nocedal and Wright [2006],Theorems 3.2, 3.8, 4.5, 4.6). For the α in the local Hölder assumption, we also relax the bounded variance assumption to only require that the $(1 + \alpha)$-moment of the stochastic gradient is bounded by an *arbitrary* upper semi-continuous function (see Assumption 4). Our assumption allows stochastic gradients whose noise may not have a variance. Moreover, our assumption generalizes the noise assumption of Bottou et al. [2018], the expected smoothness assumption of Gower et al. [2020] and Khaled and Richtárik [2020], and the noise assumption of Asi and Duchi [2019]. We also point out that we do not require the coercivity or asymptotic flatness assumptions that are commonly considered in the analysis of SGD (e.g., Mertikopoulos et al. [2020], Assumptions 2 and 3).
2. Owing to the relaxation in the assumptions, we cannot apply the standard analysis as the local Hölder constant and the iterate difference are conditionally dependent random variables (see the discussion after Lemma 1). As a result, by generalizing our previous techniques in Patel [2021], Patel and Zhang [2021] to the α-Hölder continuous setting, we innovate a new analysis strategy (see Section 4.1) to prove that, with probability one, either SGD's iterates will converge to a stationary point *or* they diverge (see Theorem 2). Importantly, our new analysis strategy can be broadly applied even to deterministic algorithms to relax the assumptions found in the literature.
3. The divergence component of our Theorem 2 is somewhat disconcerting as we cannot say exactly what happens if the iterates diverge. To ameliorate this concern, we add an additional assumption (see Assumption 5) and introduce another analysis strategy (see Section 4.2) to prove, even if the iterates are diverging, the objective function converges to a finite random variable with probability one (see Theorem 3).

Our results *do not* supply a rate of convergence as this is impossible for the broad class of nonconvex functions and the generality of the noise models studied in this work [Wolpert and Macready, 1997]; in other words, we can always construct a nonconvex objective function, a noise process, and choose an initialization such that any rate of convergence statement will be violated. Indeed, we find it remarkable that it is even possible to provide a global convergence statement for such a broad class of nonconvex functions and general noise models.

[2] In Majewski et al. [2018], this is implied by their third assumption.

Table 1: A survey of influential, recent global analyses of SGD and their dependence on the two assumptions that are either individually or both violated by the simple examples discussed in Section 1.

Assumption	**Works Depending on the Assumption**
Global Lipschitz or Hölder Continuity of Gradient	Reddi et al. [2016a], Ma and Klabjan [2017], Zhou et al. [2018], Bassily et al. [2018], Lei et al. [2019], Li and Orabona [2019], Gower et al. [2020], Khaled and Richtárik [2020], Mertikopoulos et al. [2020], Patel [2021], Jin et al. [2021], Wang et al. [2021].
Bounded Variance of Stochastic Gradients	Reddi et al. [2016b], Ma and Klabjan [2017], Majewski et al. [2018], Hu et al. [2019], Bi and Gunn [2019], Zou et al. [2019], Mertikopoulos et al. [2020].[2]

Limitations. We make note of two important limitations in our work. First, we do not consider the important case of nonsmoothness in this work as we require that the gradients of the stochastic optimization function are continuous. However, we note that if the results of Bianchi et al. [2022] are broadly applicable, then SGD never observes a point of nonsmoothness and our results would then be applicable. Second, we do not have a simple interpretation of Assumption 5—though it seems to have a close relative in another analysis (see Wang and Wu [2020])—, nor have we been able to construct a relevant counterexample that can illuminate the limitations of this assumption.

2 Stochastic Optimization

We consider solving the optimization problem

$$\min_{\theta \in \mathbb{R}^p} \{F(\theta) := \mathbb{E}\left[f(\theta, X)\right]\}, \tag{1}$$

where F maps $\mathbb{R}^p$ into $\mathbb{R}$; f maps $\mathbb{R}^p$ and the co-domain of the random variable X into $\mathbb{R}$; and $\mathbb{E}$ is the expectation operator. As we require gradients, we take F and f to differentiable in θ, and denote its derivatives with respect to θ by $\dot{F}(\theta)$ and $\dot{f}(\theta, X)$. With this notation, we make the following general assumptions about the deterministic portion of the objective function.

Assumption 1. There exists $F_{l.b.} \in \mathbb{R}$ such that $\forall \theta \in \mathbb{R}^p$, $F_{l.b.} \leq F(\theta)$.

Assumption 2. There exists $\alpha \in (0, 1]$ such that $\dot{F}(\theta)$ is locally α-Hölder continuous.

Remark 1. For $\dot{F}$ to be locally α-Hölder continuous for some $\alpha \in (0, 1]$, for every compact set $K \subset \mathbb{R}^p$ there exists a constant $L \geq 0$ such that for every $\theta, \psi \in K$,

$$\left\| \dot{F}(\theta) - \dot{F}(\psi) \right\|_2 \leq L \left\| \theta - \psi \right\|_2^{\alpha}. \tag{2}$$

Remark 2. As an example, an empirical risk minimization problem for a deep neural network with twice continuously differentiable activation functions with a twice continuously differentiable loss function will readily satisfy the above conditions.

Assumptions 1 and 2 would even be considered mild in the context of non-convex deterministic optimization, in which it is also common to assume that the objective function has well-behaved level sets [e.g., Nocedal and Wright, 2006, Theorems 3.2, 3.8, 4.5, 4.6]. Importantly, Assumption 2 relaxes the common restrictive assumption of globally Hölder continuous gradient functions that is common in other analyses (see Table 1).

Our final step is to make some assumptions about the stochastic portion of the objective function. The first assumption requires the stochastic gradients to be unbiased, which can readily be relaxed [Bottou et al., 2018]. The second assumption allows for a generic noise model for an α-Hölder continuous gradient function, and even allows for the second moment to not exist when $\alpha < 1$ [c.f. Wang et al., 2021, which requires a decomposition of the noise term that we do not require].

Assumption 3. For all $\theta \in \mathbb{R}^p$, $\dot{F}(\theta) = \mathbb{E}[\dot{f}(\theta, X)]$.

Assumption 4. Let $\alpha \in (0, 1]$ be as in Assumption 2. There exists an upper semi-continuous function $G(\theta)$ such that $\mathbb{E}[\|\dot{f}(\theta, X)\|_2^{1+\alpha}] \leq G(\theta)$.

Remark 3. For $G(\theta)$ to be upper semi-continuous, then for all $g > 0$, $\{\theta \in \mathbb{R}^p : G(\theta) < g\}$ are open in $\mathbb{R}^p$.

We will show that Assumptions 1 to 4 are sufficient for a global convergence result (see Theorem 2).

Remark 4. As shown in §A, our examples from §1 satisfy Assumptions 1 to 4.

Remark 5. It is entirely possible that $\mathbb{E}[\|\dot{f}(\theta, X)\|_2^{1+\alpha}]$ is (at least) upper semi-continuous, and to set $G(\theta)$ equal to this function. In the case that $\mathbb{E}[\|\dot{f}(\theta, X)\|_2^{1+\alpha}]$ is not upper semi-continuous, it is possible to specify $G(\theta)$ as the upper envelope of $\mathbb{E}[\|\dot{f}(\theta, X)\|_2^{1+\alpha}]$ (i.e., its limit supremum function). However, it is unlikely that $\mathbb{E}[\|\dot{f}(\theta, X)\|_2^{1+\alpha}]$ nor its upper envelope are easy to specify explicitly, and it is more likely to be able to find an upper bound.

Remark 6. We use Assumption 4 to conclude that on any compact set, the $1 + \alpha$ moment of the stochastic gradient is bounded. Of course, we can assume this directly (i.e., on any compact set, the $1 + \alpha$ moment is bounded), which, at first glance, appears to be a relaxation. However, if we assume that on any compact set, the $1 + \alpha$ moment is bounded, we can use this to construct a $G(\theta)$ that is upper semi-continuous. Thus, the two assumptions are equivalent.

Remark 7. A simple example that shows the utility of Assumption 4 is to optimize $\mathbb{E}[\theta^X]$ where X is an exponential random variable with parameter 1 and $\theta \in [1, u]$ where $u < \exp(1)$. First, it is easy to confirm that the objective function is differentiable and its derivative is globally Lipschitz continuous. Moreover, given that we are on a bounded interval, we conclude that the derivative is globally α-Hölder continuous for any $\alpha \in (0, 1]$; therefore, we are free to choose the α as we see fit. Now, when $u < \exp(1/2)$, we have that second moment of the stochastic gradient function exists. However, when $\exp(1/2) < u < \exp(1)$, only smaller moments of the stochastic gradient will exist. Specifically, only for $u < \exp(1/(1 + \epsilon))$ with $\epsilon \in (0, 1)$ will the $1 + \epsilon$ moment of the stochastic gradient will exist. Thus, depending on the size of our interval, we may not have the existence of the second moment, and, consequently, we may not have the existence of the variance.

In order to show that the objective function cannot diverge (i.e., to prove stability), we will need an additional assumption. This assumption will relate the gradient function, noise model and variation on the local Hölder constant. To begin, we define the variation on the local Hölder constant. Let $\alpha \in (0, 1]$ be as in Assumption 2 and $\epsilon > 0$ be arbitrary, and define

$$\mathcal{L}_\epsilon(\theta) = \begin{cases} \sup_\varphi \left\{ \frac{\left\|\dot{F}(\varphi) - \dot{F}(\theta)\right\|_2}{\|\varphi - \theta\|_2^\alpha} : \|\varphi - \theta\|_2 \leq (G(\theta) \vee \epsilon)^{\frac{1}{1+\alpha}} \right\} & \text{if this quantity is nonzero} \\ \epsilon & \text{otherwise,} \end{cases} \tag{3}$$

where $\vee$ indicates the maximum between two quantities. Note, the choice of ϵ is irrelevant, and they can be distinct for the two cases in the definition of $\mathcal{L}_\epsilon$, but we fix them to be the same for simplicity. Importantly, the quantity, $\mathcal{L}_\epsilon$, is defined at every parameter θ under Assumption 2.

With this quantity, we can state a nonintuitive, technical assumption that is needed to prove stability.

Assumption 5. There exists $C_1, C_2, C_3 \geq 0$ such that, $\forall \theta \in \mathbb{R}^p$,

$$\mathcal{L}_\epsilon(\theta)G(\theta) + \alpha \left(\frac{\left\|\dot{F}(\theta)\right\|_2^{1+\alpha}}{\mathcal{L}_\epsilon(\theta)} \right)^{1/\alpha} \leq C_1 + C_2(F(\theta) - F_{l.b.}) + C_3 \left\|\dot{F}(\theta)\right\|_2^2. \tag{4}$$

Assumption 5 generalizes Assumption 4.3(c) of Bottou et al. [2018], which is satisfied for a large swath of statistical models. Moreover, Assumption 5 generalizes the notion of expected smoothness [see Khaled and Richtárik, 2020, for a history of the assumption], which expanded the optimization problems covered by the theory of Bottou et al. [2018]. Note, Assumption 5 is about the asymptotic properties of the stochastic optimization problem as the left hand side of the inequality in Assumption 5 can be bounded inside of any compact set. Thus, Assumption 5 covers a variety of asymptotic behaviors, such as $\exp(\|\theta\|_2^2)$, $\exp(\|\theta\|_2)$, $\|\theta\|_2^r$ for $r \in \mathbb{R}$, $\log(\|\theta\|_2 + 1)$, and $\log(\log(\|\theta\|_2 + 1) + 1)$. Therefore, Assumption 5 holds for functions with a variety of different asymptotic behaviors.

We will show that Assumptions 1 to 5 are sufficient for a stability result (see Theorem 3).

Now that we have specified the nature of the stochastic optimization problem, we turn our attention to the algorithm used to solve the problem, namely, stochastic gradient descent (SGD).

3 Stochastic Gradient Descent

SGD starts with an arbitrary initial value, $\theta_0 \in \mathbb{R}^p$, and generates a sequence of iterates $\{\theta_k : k \in \mathbb{N}\}$ according to the rule

$$\theta_{k+1} = \theta_k - M_k \dot{f}(\theta_k, X_{k+1}), \tag{5}$$

where $\{M_k : k+1 \in \mathbb{N}\} \subset \mathbb{R}^{p \times p}$; and $\{X_k : k \in \mathbb{N}\}$ are independent and identically distributed copies of X. Importantly, $\{M_k\}$ cannot be arbitrary, and the following properties specify a generalization of the Robbins and Monro [1951] conditions for matrix-valued learning rates [c.f. Patel, 2021].

The first condition requires a positive learning rate, and imposes symmetry to ensure the existence of real eigenvalues.

Property 1. $\{M_k : k+1 \in \mathbb{N}\}$ are symmetric, positive definite matrices.

The next two properties are a natural generalization of the Robbins-Monro conditions. Let $\alpha \in (0, 1]$ be as in Assumption 2.

Property 2. Let $\lambda_{\max}(\cdot)$ denote the largest eigenvalue of a symmetric, positive definite matrix. Then, $\sum_{k=0}^{\infty} \lambda_{\max}(M_k)^{1+\alpha} =: S < \infty$.

Property 3. Let $\lambda_{\min}(\cdot)$ denote the smallest eigenvalue of a symmetric, positive definite matrix. Then, $\sum_{k=0}^{\infty} \lambda_{\min}(M_k) = \infty$.

> We will show that Properties 1 to 3 are sufficient for a global convergence result (see Theorem 2).

The final property ensures the stability of the condition number of $\{M_k\}$. Note, this property is readily satisfied for scalar learning rates satisfying the Robbins-Monro conditions.

Property 4. Let $\kappa(\cdot)$ denote the ratio of the largest and smallest eigenvalues of a symmetric, positive definite matrix. Then, $\lim_{k \to \infty} \lambda_{\max}(M_k)^{\alpha} \kappa(M_k) = 0$.

> We will show that Properties 1 to 4 are sufficient for stability (see Theorem 3).

4 Global Convergence & Stability

With the stochastic optimization problem and with stochastic gradient descent (SGD) specified, we now turn our attention to what happens when SGD is applied to a stochastic optimization problem. The key step in the analysis of SGD on any objective function is to establish a bound between the optimality gap at θ_{k+1} with that of θ_k. This step is achieved by using the local Hölder continuity of the gradient function and the fundamental theorem of calculus. Using Assumption 2, we first specify the local Hölder constant.

Definition 1. For any $\theta, \varphi \in \mathbb{R}^p$, define

$$L(\theta, \varphi) = \sup_{\psi} \left\{ \frac{\left\| \dot{F}(\psi) - \dot{F}(\theta) \right\|_2}{\|\psi - \theta\|_2^{\alpha}} : \psi \in \overline{B(\theta, \|\varphi - \theta\|_2)} \right\}, \tag{6}$$

where $B(\theta, r)$ is an open ball around θ of radius $r > 0$, and $\overline{B(\theta, r)}$ is its closure. Moreover, for any $R \geq 0$, let L_R be the supremum of $L(\theta, \varphi)$ for any distinct $\theta, \varphi \in \overline{B(0, R)}$.

Remark 8. Note, when the gradient is locally Hölder continuous, L_R is finite for any $R \geq 0$.

With this definition, we can now relate the optimality gap of θ_{k+1} with that of θ_k by using the following result and proved in Appendix B.

Lemma 1. *Suppose Assumptions 1 and 2 hold. Then, for any $\theta, \varphi \in \mathbb{R}^p$,*

$$F(\varphi) - F_{l.b.} \leq F(\theta) - F_{l.b.} + \dot{F}(\theta)'(\varphi - \theta) + \frac{L(\theta, \varphi)}{1+\alpha} \|\varphi - \theta\|_2^{1+\alpha}. \tag{7}$$

Now, if we simply set $\varphi = \theta_{k+1}$ and $\theta = \theta_k$ in Lemma 1 and try to take expectations to manage the randomness of the stochastic gradient, we will run into the problem that $L(\theta_k, \theta_{k+1})$ and $\|\theta_{k+1} - \theta_k\|_2$ are potentially dependent,[3] and we cannot compute its expectation. In previous work, this technical challenge is waived away by using a global Hölder constant to upper bound $L(\theta_k, \theta_{k+1})$, which is unrealistic even for simple problems (see Section 1).

To address this technical challenge, we innovate two new strategies for handling the dependence between $L(\theta_k, \theta_{k+1})$ and $\|\theta_{k+1} - \theta_k\|_2$. In both strategies, we follow the same general approach:

1. We begin by restricting our analysis to specific events, which will allow us to decouple $L(\theta_k, \theta_{k+1})$ and $\|\theta_{k+1} - \theta_k\|_2$.
2. With these two quantities decoupled, we will develop a recurrence relationship between the optimality gap at θ_{k+1} and that of θ_k.
3. We apply this recurrence relationship with refinements of standard arguments or new ones to derive the desired property about the objective function.
4. Finally, we state the generality of the specific events on which we have studied SGD's iterates.

Thus, it follows, we will define two distinct series of events for the two strategies. The first strategy, which we refer to as the pseudo-global strategy, will provide the global convergence analysis. The second strategy, which we refer to as the local strategy, will provide the stability result.

4.1 Pseudo-Global Strategy and Global Convergence Analysis

For the first strategy, which supplies the global convergence result, we study SGD on the events

$$\mathcal{B}_k(R) := \bigcap_{j=0}^{k} \left\{ \|\theta_j\|_2 \leq R \right\}, \; k+1 \in \mathbb{N}, \tag{8}$$

for every $R \geq 0$. We now try to control the optimality gap at iteration $k+1$ with that of iteration k, which will result in two cases.

1. (Case 1) $\mathcal{B}_{k+1}(R)$ holds. We can bound $L(\theta_k, \theta_{k+1})$ by L_R, and $G(\theta)$ is also bounded in the ball of radius R about the origin (which follows from G being upper semi-continuous in Assumption 4). As a result, we could then proceed with the analysis in a manner that is similar to having a global Hölder constant.
2. (Case 2) $\|\theta_{k+1}\|_2 > R$ and $\mathcal{B}_k(R)$ holds. In this case, controlling $L(\theta_k, \theta_{k+1})$ is very challenging and, to our knowledge, was not solved before our work.

Our approach for controlling the optimality gap in both cases is supplied in the next lemma, whose proof is in Appendix C.

Lemma 2. *Let $\{M_k\}$ satisfy Property 1. Suppose Assumptions 1 to 4 hold. Let $\{\theta_k\}$ satisfy (5). Then, $\forall R \geq 0$,*

$$\begin{aligned} \mathbb{E}\left[[F(\theta_{k+1}) - F_{l.b.}] \mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] | \, \mathcal{F}_k \right] &\leq [F(\theta_k) - F_{l.b.}] \mathbf{1}\left[\mathcal{B}_k(R)\right] \\ &- \lambda_{\min}(M_k) \left\| \dot{F}(\theta_k) \right\|_2^2 \mathbf{1}\left[\mathcal{B}_k(R)\right] + \frac{L_{R+1} + \partial F_R}{1+\alpha} \lambda_{\max}(M_k)^{1+\alpha} G_R, \end{aligned} \tag{9}$$

where $G_R = \sup_{\theta \in \overline{B(0,R)}} G(\theta) < \infty$ with $G(\theta)$; and $\partial F_R = \sup_{\theta \in \overline{B(0,R)}} \|\dot{F}(\theta)\|_2 (1+\alpha) < \infty$.

With this recursion and standard martingale results [Robbins and Siegmund, 1971, Neveu and Speed, 1975, Exercise II.4], the limit of $[F(\theta_k) - F_{l.b.}]\mathbf{1}\left[\mathcal{B}_k(R)\right]$ exists with probability one and is finite for

[3] While it is possible that these two terms are independent, we would require a lot more information to determine this and it would likely be on an iterate-by-iterate basis for the general class of problems considered in this work. Thus, in this general setting, we cannot assume independence and need to default to treating these terms as dependent.

every $R \geq 0$. As a result, the limit of $F(\theta_k) - F_{l.b.}$ exists and is finite on the event $\{\sup_k \|\theta_k\|_2 < \infty\}$ (see Corollary 1).

We can also use Lemma 2 to make a statement about the gradient. Specifically, we can show that the limit infimum of $\mathbb{E}[\|\dot{F}(\theta_k)\|_2^2 \mathbf{1}[\mathcal{B}_k(R)]]$ must be zero, which is now a standard argument that mimics Zoutendijk's theorem [Nocedal and Wright, 2006, Theorem 3.2]. By Markov's inequality, this result implies that $\|\dot{F}(\theta_k)\|_2 \mathbf{1}[\mathcal{B}_k(R)]$ gets arbitrarily close to 0 infinitely often (see Lemma 9). To show convergence to zero, however, is not standard. Several strategies have been developed, namely those of Li and Orabona [2019], Lei et al. [2019], Mertikopoulos et al. [2020], Patel [2021], Patel and Zhang [2021]. Unfortunately, the approaches of Li and Orabona [2019], Lei et al. [2019] rely intimately on the existence of a global Hölder constant, while that of Mertikopoulos et al. [2020] requires even more restrictive assumptions. Fortunately, the approach of Patel [2021], Patel and Zhang [2021] can be improved and generalized to the current context (see Lemma 10). Thus, we show that $\lim_{k\to\infty} \|\dot{F}(\theta_k)\|_2 = 0$ on $\{\sup_k \|\theta_k\|_2 < \infty\}$ (see Corollary 2).

Our final step is to clarify the role of $\{\sup_k \|\theta_k\|_2 < \infty\}$ in the asymptotics of SGD's iterates. At first glance, this event seems to imply that the iterates converge to a point. However, owing to the general nature of the noise, it is also possible, say, that the iterates approach a limit cycle or oscillate between points with the same norm. Even beyond this event, the generality of the noise model may allow for substantial excursions between $F_{l.b.}$ and infinity (c.f., a simple random walk, which has a limit supremum of infinity and a limit infinimum of negative infinity). Thankfully, we can prove that either the iterates converge to a point or they must diverge—a result that we refer to as the Capture Theorem (see Appendix C).

Theorem 1 (Capture Theorem)**.** *Let $\{\theta_k\}$ be defined as in (5), and let $\{M_k\}$ satisfy Properties 1 and 2. If Assumption 4 holds, then either $\{\lim_{k\to\infty} \theta_k \text{ exists}\}$ or $\{\liminf_{k\to\infty} \|\theta_k\|_2 = \infty\}$ must occur.*

By putting together the above arguments and results, we can conclude that either SGD's iterates diverge or SGD's iterates converge to a stationary point.

Theorem 2 (Global Convergence)**.** *Let θ_0 be arbitrary, and let $\{\theta_k : k \in \mathbb{N}\}$ be defined according to (5) with $\{M_k : k+1\}$ satisfying Properties 1 to 3. Suppose Assumptions 1 to 4 hold. Let $\mathcal{A}_1 = \{\liminf_{k\to\infty} \|\theta_k\|_2 = \infty\}$ and $\mathcal{A}_2 = \{\lim_{k\to\infty} \theta_k \text{ exists}\}$. Then, the following statements hold.*

1. *$\mathbb{P}[\mathcal{A}_1] + \mathbb{P}[\mathcal{A}_2] = 1$.*
2. *On $\mathcal{A}_2$, there exists a finite random variable, $F_{\lim}$, such that $\lim_{k\to\infty} F(\theta_k) = F_{\lim}$ and $\lim_{k\to\infty} \dot{F}(\theta_k) = 0$ with probability one.*

Proof. By Theorem 1, we have that $\mathbb{P}[\mathcal{A}_1] + \mathbb{P}[\mathcal{A}_2] = 1$. Then, on $\mathcal{A}_2$, Corollaries 1 and 2 imply that $F(\theta_k) \to F_{\lim}$, which is finite, and $\dot{F}(\theta_k) \to 0$. □

We pause to stress to a fact about Theorem 2: it is nonobvious. To be specific, under such general nonconvexity and noise, we should anticipate any number of asymptotic behaviors for the iterates: convergence to a stationary point, convergence to a nonstationary point, being trapped in a cycle, convergence to a limit cycle, and divergence to infinity. However and very surprisingly, we are able to show that only two possible outcomes can occur: convergence to a stationary point or divergence. Indeed, in previous results [e.g., Mertikopoulos et al., 2020, Patel, 2021], only less specific determinations could be made under much more limited settings.

4.2 Local Strategy and Stability Analysis

While Theorem 2 provides a complete global convergence result, it allows for the possibility of diverging iterates. The possibility of divergent iterates raises the spectre of whether the objective function can also diverge along this sequence. That is, there is a possibility that SGD may be unstable, which would be highly unexpected and undesirable, especially when the objective function is coercive (e.g., has an ℓ^1 penalty on the parameters). To formalize this concept, we define a relevant notion of stability.

Definition 2. Stochastic Gradient Descent is stable if

$$\mathbb{P}\left[\limsup_{k\to\infty} F(\theta_k) = \infty\right] = 0, \tag{10}$$

where $\{\theta_k\}$ satisfy (5).

We now state the stopping times that we will use to decouple the relationship between $L(\theta_k, \theta_{k+1})$ and $\|\theta_{k+1} - \theta_k\|_2$. For every $j+1 \in \mathbb{N}$, define

$$\tau_j = \min\left\{ k : \begin{array}{c} F(\theta_{k+1}) - F_{l.b.} > F(\theta_k) - F_{l.b.} + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) \\ + \dfrac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha} \|\theta_{k+1} - \theta_k\|_2^{1+\alpha}, \text{ and } k \geq j \end{array} \right\}. \tag{11}$$

Now, we will use (11) to establish the stability of the objective function. Just as we did with $\mathcal{B}_k(R)$, we will derive a recursion on the optimality gap over the events $\{\{\tau_j > k\} : k+1 \in \mathbb{N}\}$. Of course, just as before, the main challenge in deriving a recursive formula is to address $\{\tau_j = k\}$. Our solution is supplied in the following lemma, whose proof is in Appendix D.

Lemma 3. *Let $\{M_k\}$ satisfy Property 1. Suppose Assumptions 1 to 4 hold. Let $\{\theta_k\}$ satisfy (5). Then, for any $j+1 \in \mathbb{N}$ and $k > j$,*

$$\begin{aligned} \mathbb{E}\left[(F(\theta_{k+1}) - F_{l.b.}) \boldsymbol{1}\left[\tau_j > k\right] | \mathcal{F}_k \right] &\leq \left(F(\theta_k) - F_{l.b.} - \dot{F}(\theta_k)' M_k \dot{F}(\theta_k) \right) \boldsymbol{1}\left[\tau_j > k-1\right] \\ &+ \frac{\lambda_{\max}(M_k)^{1+\alpha}}{1+\alpha} \left[\mathcal{L}_\epsilon(\theta_k) G(\theta_k) + \alpha \left[\frac{\left\| \dot{F}(\theta_k) \right\|_2^{1+\alpha}}{\mathcal{L}_\epsilon(\theta_k)} \right]^{1/\alpha} \right] \boldsymbol{1}\left[\tau_j > k-1\right]. \end{aligned} \tag{12}$$

From Lemma 3, there is a clear motivation for Assumption 5. Indeed, if we apply Assumption 5, Lemma 3 produces the following simple recursive relationship.

Lemma 4. *If Assumptions 1 to 5, and Properties 1 and 4 hold, and $\{\theta_k\}$ satisfy (5), then there exists a $K \in \mathbb{N}$ such that for any $j+1 \in \mathbb{N}$ and any $k \geq \min\{K, j+1\}$,*

$$\begin{aligned} &\mathbb{E}\left[(F(\theta_{k+1}) - F_{l.b.}) \boldsymbol{1}\left[\tau_j > k\right] | \mathcal{F}_k \right] \\ &\leq \left(1 + \lambda_{\max}(M_k)^{1+\alpha} \frac{C_2}{1+\alpha} \right) (F(\theta_k) - F_{l.b.}) \boldsymbol{1}\left[\tau_j > k-1\right] \\ &- \frac{1}{2} \lambda_{\min}(M_k) \left\| \dot{F}(\theta_k) \right\|_2^2 \boldsymbol{1}\left[\tau_j > k-1\right] + \lambda_{\max}(M_k)^{1+\alpha} \frac{C_1}{1+\alpha}. \end{aligned} \tag{13}$$

Just as in the pseudo-global strategy, Lemma 4 can be combined with standard martingale results [Robbins and Siegmund, 1971, Neveu and Speed, 1975, Exercise II.4] to conclude that the limit of $F(\theta_k)$ exists and is finite on the event $\cup_{j=0}^\infty \{\tau_j = \infty\}$ (see Corollary 3). Also as in the pseudo-global strategy, by improving on the arguments in Patel [2021], Patel and Zhang [2021], we show that $\liminf_k \dot{F}(\theta_k) = 0$ on the event $\cup_{j=0}^\infty \{\tau_j = \infty\}$ (see Lemma 14).

Finally, we show that $\cup_{j=0}^\infty \{\tau_j = \infty\}$ is a probability one event (see Theorem 5). This statement should not come as a surprise on the event $\{\lim_k \theta_k \text{ exists}\}$, but is slightly surprising that it must also hold on $\{\lim_k \|\theta_k\|_2 = \infty\}$. By combining these results, we can conclude as follows.

Theorem 3 (Stability)**.** *Let θ_0 be arbitrary, and let $\{\theta_k : k \in \mathbb{N}\}$ be defined according to (5) with $\{M_k : k+1\}$ satisfying Properties 1 to 4. Suppose Assumptions 1 to 5 hold. Then,*

1. *There exists a finite random variable, $F_{\lim}$, such that $\lim_{k\to\infty} F(\theta_k) = F_{\lim}$ with probability one;*
2. $\liminf_{k\to\infty} \dot{F}(\theta_k) = 0$ *with probability one.*

Proof. Using Corollary 3, we conclude that $\exists F_{\lim}$ that is finite such that $\lim_k F(\theta_k) = F_{\lim}$ on $\cup_{j=0}^\infty \{\tau_j = \infty\}$. Using Lemma 14, we conclude that $\liminf_k \dot{F}(\theta_k) = 0$ on $\cup_{j=0}^\infty \{\tau_j = \infty\}$. Finally, we apply Theorem 5 to conclude that $\mathbb{P}[\cup_{j=0}^\infty \{\tau_j = \infty\}] = 1$. □

We would like to demonstrate a simple example of how we would use Theorem 3. Consider applying SGD to linear regression as specified in Appendix A.1. For this example, it is straightforward to verify that the assumptions of Theorem 3 are satisfied. Therefore, if we are to apply SGD to linear regression, we know that with probability one, $F(\theta_k) \to F_{\lim}$ which is finite. Since $F(\theta) \to \infty$ as $\theta \to \infty$, we know that $\{\theta_k\}$ cannot diverge. Hence, $\{\theta_k\}$ must remain finite with probability one. By Theorem 2, $\{\theta_k\}$ must converge to a stationary point. Since this stationary point is unique in our specific example of linear regression, we know that SGD must converge to the global minimizer of the linear regression problem. Note, we can follow this outline to draw similar conclusions in more complex situations.

5 Conclusion

In this work, we studied the global convergence analysis of Stochastic Gradient Descent with diminishing step size. We began our discussion by producing three simple problems for which the common assumptions (i.e., global Hölder continuity, bounded variance) and their generalizations in the SGD literature are violated. Indeed, to our knowledge, there does not exist theory that covers these problems. For example, prior to our work, it was unknown what SGD with arbitrary initialization and diminishing step sizes will do on simple neural network problems, which raised the question of what SGD is doing on more complicated learning problems.

Motivated by our example problems, we considered a more general set of assumptions (see Assumptions 2 and 4). Given the generality of our assumptions, we developed a new analysis technique that is of interest beyond this work, and we proved that SGD's iterates either converge to a stationary point or diverge. Thus, we now know how SGD with arbitrary initialization and diminishing step sizes will behave on a much larger class of learning problems.

We note that we do not provide rate of convergence results mainly because it is *impossible* for the broad class of functions admitted by our assumptions [Wolpert and Macready, 1997]. We stress that global rates of convergence (e.g., complexity statements) results that exist do not apply to the two simple neural network problems that we supplied at the beginning of this work.

We also studied what happens when SGD's iterates diverge. To this end, we required an additional assumption under which we developed another novel analysis technique and showed that, regardless of SGD's iterates' behavior, the objective function will converge to a finite random variable with probability one. Unfortunately, we make an assumption that we were not able to interpret, but we will leave this to future work.

References

H. Asi and J. C. Duchi. Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019.

R. Bassily, M. Belkin, and S. Ma. On exponential convergence of sgd in non-convex over-parametrized learning. *arXiv preprint arXiv:1811.02564*, 2018.

J. Bi and S. R. Gunn. A stochastic gradient method with biased estimation for faster nonconvex optimization. In *Pacific Rim International Conference on Artificial Intelligence*, pages 337–349. Springer, 2019.

P. Bianchi, W. Hachem, and S. Schechtman. Convergence of constant step stochastic gradient descent for non-smooth non-convex functions. *Set-Valued and Variational Analysis*, pages 1–31, 2022.

L. Bottou, F. E. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.

H. Chen, L. Zheng, R. Al Kontar, and G. Raskutti. Stochastic gradient descent in correlated settings: A study on gaussian processes. In *NeurIPS*, 2020.

C. Fang, Z. Lin, and T. Zhang. Sharp analysis for nonconvex sgd escaping from saddle points. In *Conference on Learning Theory*, pages 1192–1234. PMLR, 2019.

R. M. Gower, O. Sebbouh, and N. Loizou. Sgd for structured nonconvex functions: Learning rates, minibatching and interpolation. *arXiv preprint arXiv:2006.10311*, 2020.

M. Hardt, T. Ma, and B. Recht. Gradient descent learns linear dynamical systems. *arXiv preprint arXiv:1609.05191*, 2016.

W. Hu, C. J. Li, L. Li, and J.-G. Liu. On the diffusion approximation of nonconvex stochastic gradient descent. *Annals of Mathematical Sciences and Applications*, 4(1), 2019.

C. Jin, P. Netrapalli, R. Ge, S. M. Kakade, and M. I. Jordan. On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. *Journal of the ACM (JACM)*, 68(2):1–29, 2021.

A. Khaled and P. Richtárik. Better theory for sgd in the nonconvex world. *arXiv preprint arXiv:2002.03329*, 2020.

H.-S. Kim, L. Zhang, A. Bienkowski, and K. R. Pattipati. Multi-pass sequential mini-batch stochastic gradient descent algorithms for noise covariance estimation in adaptive kalman filtering. *IEEE Access*, 9:99220–99234, 2021.

Y. Lei, T. Hu, G. Li, and K. Tang. Stochastic gradient descent for nonconvex learning without bounded gradient assumptions. *IEEE transactions on neural networks and learning systems*, 31(10):4394–4400, 2019.

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*, pages 983–992. PMLR, 2019.

Y. Ma and D. Klabjan. Convergence analysis of batch normalization for deep neural nets. *arXiv preprint arXiv:1705.08011*, 2017.

S. Majewski, B. Miasojedow, and E. Moulines. Analysis of nonsmooth stochastic approximation: the differential inclusion approach. *arXiv preprint arXiv:1805.01916*, 2018.

P. Mertikopoulos, N. Hallak, A. Kavis, and V. Cevher. On the almost sure convergence of stochastic gradient descent in non-convex problems. *arXiv preprint arXiv:2006.11144*, 2020.

J. Neveu and T. Speed. *Discrete-parameter martingales*, volume 10. North-Holland Amsterdam, 1975.

J. Nocedal and S. Wright. *Numerical optimization*. Springer Science & Business Media, 2006.

V. Patel. Stopping criteria for, and strong convergence of, stochastic gradient descent on bottou-curtis-nocedal functions. *Mathematical Programming*, pages 1–42, 2021.

V. Patel and S. Zhang. Stochastic gradient descent on nonconvex functions with general noise models. *arXiv preprint arXiv:2104.00423*, 2021.

S. Reddi, S. Sra, B. Poczos, and A. J. Smola. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. *Advances in neural information processing systems*, 29:1145–1153, 2016a.

S. J. Reddi, A. Hefny, S. Sra, B. Poczos, and A. Smola. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*, pages 314–323. PMLR, 2016b.

H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

H. Robbins and D. Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics*, pages 233–257. Elsevier, 1971.

H. Wang, M. Gürbüzbalaban, L. Zhu, U. Şimşekli, and M. A. Erdogdu. Convergence rates of stochastic gradient descent under infinite noise variance. *arXiv preprint arXiv:2102.10346*, 2021.

Y. Wang and S. Wu. Asymptotic analysis via stochastic differential equations of gradient descent algorithms in statistical and computational paradigms. *Journal of machine learning research*, 21, 2020.

D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1):67–82, 1997.

J. Zhang, T. He, S. Sra, and A. Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2019.

S. Zhang and V. Patel. Stochastic approximation for high-frequency observations in data assimilation. *arXiv preprint arXiv:2011.02672*, 2020.

D. Zhou, P. Xu, and Q. Gu. Stochastic nested variance reduction for nonconvex optimization. *arXiv preprint arXiv:1806.07811*, 2018.

F. Zou, L. Shen, Z. Jie, W. Zhang, and W. Liu. A sufficient condition for convergences of adam and rmsprop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11127–11135, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Each claim in the abstract corresponds to a statement in the introduction, which then references the corresponding result in the paper.
   (b) Did you describe the limitations of your work? [Yes] See the end of Section 1
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] SGD is already a widely existing and deployed method.
   (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] Each result references the precise assumptions that are used and these are clearly stated in Section 2.
   (b) Did you include complete proofs of all theoretical results? [Yes] Outlines are provided in the main document, while complete proofs of each result are provided in the appendix/supplement.
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)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (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)? [N/A]
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? [N/A]
   (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? [N/A]
   (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 Details for Counter Examples

A.1 Simple Linear Regression

We begin our exploration of assumptions with a rather simple problem. Let $Z \in \mathbb{R}$ be a random variable such that $\mathbb{E}\left[Z^2\right] = 1$ and $\mathbb{E}\left[Z^4\right] = 2$. Moreover, let ϵ be an independent random variable with mean zero and variance 1. Finally, let $\theta^* \in \mathbb{R}$ and define $Y = Z\theta^* + \epsilon$. Consider the estimation problem of minimizing $F(\theta)$ where

$$F(\theta) = \frac{1}{2}\mathbb{E}\left[(Z\theta - Y)^2\right] = \frac{1}{2}(\theta - \theta^*)^2 + \frac{1}{2}. \tag{14}$$

Letting $X = (Y, Z)$, let $f(\theta, X) = 0.5(Z\theta - Y)^2$. Now, the variance of $\dot{f}(\theta, X)$ is

$$\begin{aligned}
&\mathbb{E}\left[((Z^2 - 1)(\theta - \theta^*) - Z\epsilon)^2\right] \\
&= \mathbb{E}\left[(Z^2 - 1)^2\right](\theta - \theta^*)^2 - 2(\theta - \theta^*)\mathbb{E}\left[(Z^2 - 1)Z\epsilon\right] + \mathbb{E}\left[Z^2\epsilon^2\right] && (15)\\
&= (\theta - \theta^*)^2 + 1. && (16)
\end{aligned}$$

Clearly, the variance scales with the error in the parameter, which violates the common bounded noise model assumption. In particular, as $|\theta| \to \infty$, the variance diverges.

On the other hand, the simple linear regression problem does satisfy our assumptions. In particular,

1. Assumptions 1 and 3 are easily verified.
2. Given that $\dot{F}$ is globally Lipschitz continuous, it is locally Lipschitz continuous. Therefore, Assumption 2 is satisfied.
3. From the variance calculation of $\dot{f}(\theta, X)$, we conclude
$$\mathbb{E}\left[\dot{f}(\theta, X)^2\right] = 2(\theta - \theta^*)^2 + 1, \tag{17}$$
which is a continuous function. Hence, Assumption 4 is satisfied.

A.2 Feed Forward Network for Binary Classification

We now prove Proposition 1. Consider the binary classification problem with label Y and feature Z where $(Y, Z) = (0, 0)$ with probability $1/2$ and $(Y, Z) = (1, 1)$ with probability $1/2$. We solve this classification problem using the network shown in Fig. 1 with σ linear and φ sigmoid. We will train this model using the binary cross entropy loss function. Letting $X = (Y, Z)$ and $\theta = (W_1, W_2, W_3, W_4)$,

$$f(\theta, X) = -Y\log(\hat{y}) - (1 - Y)\log(1 - \hat{y}) + \frac{1}{2}\sum_{i=1}^{4} W_i^2, \tag{18}$$

and

$$\hat{y} = \begin{cases} \frac{1}{2} & Z = 0 \\ \frac{1}{1+\exp(-W_4W_3W_2W_1)} & Z = 1. \end{cases} \tag{19}$$

From this, we compute,

$$F(\theta) = \frac{1}{2}\log(2) + \frac{1}{2}\log[1 + \exp(-W_4W_3W_2W_1)] + \frac{1}{2}\sum_{i=1}^{4} W_i^2. \tag{20}$$

Moreover,

$$\dot{f}(\theta, X) = \begin{cases} \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix} & (Y, Z) = (0, 0) \\ \frac{-1}{1+\exp(W_4W_3W_2W_1)} \begin{bmatrix} W_4W_3W_2 \\ W_4W_3W_1 \\ W_4W_2W_1 \\ W_3W_2W_1 \end{bmatrix} + \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix} & (Y, Z) = (1, 1), \end{cases} \tag{21}$$

and, consequently,

$$\dot{F}(\theta) = \frac{-1/2}{1+\exp(W_4W_3W_2W_1)} \begin{bmatrix} W_4W_3W_2 \\ W_4W_3W_1 \\ W_4W_2W_1 \\ W_3W_2W_1 \end{bmatrix} + \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix}. \tag{22}$$

Finally, letting $\ddot{F}(\theta) = \nabla^2 F(\psi)|_{\psi=\theta}$,

$$\begin{aligned} \ddot{F}(\theta) = & \frac{-0.5}{1+\exp(W_4W_3W_2W_1)} \begin{bmatrix} 0 & W_4W_3 & W_4W_2 & W_3W_2 \\ W_4W_3 & 0 & W_4W_1 & W_3W_1 \\ W_4W_2 & W_4W_1 & 0 & W_2W_1 \\ W_3W_2 & W_3W_1 & W_2W_1 & 0 \end{bmatrix} \\ & + \frac{0.5\exp(W_4W_3W_2W_1)}{[1+\exp(W_4W_3W_2W_1)]^2} \begin{bmatrix} W_4W_3W_2 \\ W_4W_3W_1 \\ W_4W_2W_1 \\ W_3W_2W_1 \end{bmatrix} \begin{bmatrix} W_4W_3W_2 \\ W_4W_3W_1 \\ W_4W_2W_1 \\ W_3W_2W_1 \end{bmatrix}' + I_4, \end{aligned} \tag{23}$$

where I_4 is the 4×4 identity matrix.

We first establish that $\dot{F}(\theta)$ is not globally Lipschitz continuous. With $\theta = (1, -1, W_3, W_3)$ and $\phi = (1, -1, W_3, 0)$, it is enough to find a lower bound for the first component of $\dot{F}(\theta) - \dot{F}(\phi)$, denoted by $\dot{F}_1(\theta) - \dot{F}_1(\phi)$. To this end,

$$|\dot{F}_1(\theta) - \dot{F}_1(\phi)| = \frac{0.5W_3^2}{1+\exp(-W_3^2)} \geq \frac{1}{4}|W_3 - 0|^2. \tag{24}$$

Thus, $\dot{F}$ is not globally Lipschitz.

We now establish that F does not satisfy (L_0, L_1)-smoothness. That is, we show that there is no $L_0, L_1 \geq 0$ such that $\|\ddot{F}(\theta)\| \leq L_0\|\dot{F}(\theta)\| + L_1$, where the norms can be chosen arbitrarily owing to the equivalence of norms in finite-dimensional vector spaces. To see this, note that the Frobenius norm of $\ddot{F}(\theta)$ is lower bounded by the absolute value of the $[1, 1]$ entry. Using notation,

$$\frac{0.5\exp(W_4W_3W_2W_1)}{[1+\exp(W_4W_3W_2W_1)]^2}(W_4W_3W_2)^2 + 1 \leq \left\|\ddot{F}(\theta)\right\|_F. \tag{25}$$

Let $\theta = (0, W_4, W_4, W_4)$, then the lower bound is

$$\frac{1}{8}W_4^6 \leq \left\|\ddot{F}(\theta)\right\|_F. \tag{26}$$

Notice, for this same choice of θ, the l^1 norm of the gradient is bounded above by

$$\left\|\dot{F}(\theta)\right\|_1 \leq \frac{1}{4}|W_4|^3 + 3|W_4|. \tag{27}$$

For any choice of $L_0, L_1 > 0$, we conclude that there is a W_4 sufficiently large such that, for this parametrization of θ,

$$L_0\left\|\dot{F}(\theta)\right\| + L_1 \leq L_0[\frac{1}{4}|W_4|^3 + 3|W_4|] + L_1 < \frac{1}{8}W_4^6 \leq \left\|\ddot{F}(\theta)\right\|_F. \tag{28}$$

Thus, we see that no L_0 nor L_1 can exist that will satisfy the (L_0, L_1)-smooth assumption for all choices of θ.

To show that the variance is not bounded, we study the variance of the first component of $\dot{f}(\theta, X)$ which we denote by $\dot{f}_1(\theta, X)$. By direct calculation,

$$\mathbb{E}\left[(\dot{f}_1(\theta, X) - \dot{F}_1(\theta))^2\right] = \frac{1}{4}\frac{W_4^2W_3^2W_2^2}{[1+\exp(W_4W_3W_2W_1)]^2}. \tag{29}$$

We again consider $\theta = (1, -1, W_3, W_3)$, then the variance at this value of θ is

$$\frac{1}{4}\frac{W_3^4}{[1+\exp(-W_3^2)]^2} \geq \frac{1}{16}W_3^4. \tag{30}$$

Therefore, as $W_3 \to \infty$, the variance goes to infinity. That is, the variance of the stochastic gradients is unbounded.

On the other hand, the problem does satisfy our assumptions. In particular,

1. Assumptions 1 and 3 are easily verified.
2. Given that $\dot{F}$ is continuously differentiable, then compactness and continuity of the derivative of $\dot{F}$ imply that it is locally Lipschitz continuous. Therefore, Assumption 2 is satisfied.
3. Given the computation of the variance for the first component, we have $\mathbb{E}[\dot{f}_1(\theta, X)^2]$ is

$$\frac{1}{4}\frac{W_4^2 W_3^2 W_2^2}{[1+\exp(W_4 W_3 W_2 W_1)]^2} + \dot{F}_1(\theta)^2, \tag{31}$$

which is a continuous function. By repeating this argument for each component, we conclude that Assumption 4 is satisfied.

A.3 Recurrent Neural Network for Binary Classification

Consider observing one of two sequences $(1, 0, 0, 0)$ or $(0, 0, 0, 0)$ with equal probabilities, and suppose that each sequence corresponds to the label 1 or 0, respectively. Now consider Fig. 2 to be a 1-dimensional linear recurrent neural network which reads each element of the sequence and uses a logistic output layer to predict either a label of one or zero. If we fix $H_0 = 0$ and $W_3 = 1$, then the model predicts the probability of a 1 label as

$$\hat{y}(Z_0, Z_1, Z_2, Z_3) = \frac{\exp(W_1^3 W_2 Z_0)}{1+\exp(W_1^3 W_2 Z_0)}. \tag{32}$$

If we use the binary cross entropy loss with ℓ^2 regularization, and let $X = (Y, Z_0, Z_1, Z_2, Z_3)$ and $\theta = (W_1, W_2)$ then

$$f(\theta, X) = -Y \log \hat{y}(Z_0, Z_1, Z_2, Z_3) - (1-Y)\log[1-\hat{y}(Z_0, Z_1, Z_2, Z_3)] + \frac{1}{2}(W_1^2 + W_2^2) \tag{33}$$

$$= -Y\left[W_1^3 W_2 Z_0 - \log(1+\exp(W_1^3 W_2 Z_0))\right] + (1-Y)\log(1+\exp(W_1^3 W_2 Z_0)) + \frac{1}{2}(W_1^2 + W_2^2) \tag{34}$$

$$= -W_1^3 W_2 Z_0 Y + \log(1+\exp(W_1^3 W_2 Z_0)) + \frac{1}{2}(W_1^2 + W_2^2), \tag{35}$$

and

$$\dot{f}(\theta, X) = \begin{bmatrix} -3W_1^2 W_2 Z_0 Y + \frac{3W_1^2 W_2 Z_0 \exp(W_1^3 W_2 Z_0)}{1+\exp(W_1^3 W_2 Z_0)} + W_1 \\ -W_1^3 Z_0 Y + \frac{W_1^3 Z_0 \exp(W_1^3 W_2 Z_0)}{1+\exp(W_1^3 W_2 Z_0)} + W_2 \end{bmatrix} \tag{36}$$

Taking the expectations, we compute

$$F(\theta) = \frac{1}{2}\left[\log(2) + \log(1+\exp(W_1^3 W_2)) - W_1^3 W_2 + W_1^2 + W_2^2\right], \tag{37}$$

and

$$\dot{F}(\theta) = \begin{bmatrix} \frac{-3W_1^2 W_2}{2}\frac{1}{1+\exp(W_1^3 W_2)} + W_1 \\ \frac{-W_1^3}{2}\frac{1}{1+\exp(W_1^3 W_2)} + W_2 \end{bmatrix}. \tag{38}$$

Taking another derivative and letting $\ddot{F}(\theta) = \nabla^2 F(\psi)|_{\psi=\theta}$,

$$\ddot{F}(\theta) = \begin{bmatrix} \frac{9W_1^4 W_2^2 \exp(W_1^3 W_2)}{2(1+\exp(W_1^3 W_2))^2} - \frac{3W_1 W_2}{1+\exp(W_1^3 W_2)} + 1 & \frac{3W_1^5 W_2 \exp(W_1^3 W_2)}{2(1+\exp(W_1^3 W_2))^2} - \frac{3W_1^2}{2}\frac{1}{1+\exp(W_1^3 W_2)} \\ \frac{3W_1^5 W_2 \exp(W_1^3 W_2)}{2(1+\exp(W_1^3 W_2))^2} - \frac{3W_1^2}{2}\frac{1}{1+\exp(W_1^3 W_2)} & \frac{W_1^6 \exp(W_1^3 W_2)}{2(1+\exp(W_1^3 W_2))^2} + 1 \end{bmatrix}. \tag{39}$$

We first establish that $\dot{F}$ is not globally Lipschitz continuous. Notice, if we set $W_2 = 1$, then the first and second component of $\dot{F}(\theta)$ are proportional to $-W_1^2$ and $-W_1^3$ respectively, which are not globally Lipschitz continuous functions.

We now show that F also does not satisfy (L_0, L_1)-smoothness. Notice that, using the bottom right entry of $\ddot{F}(\theta)$,

$$\frac{W_1^6 \exp(W_1^3 W_2)}{2(1+\exp(W_1^3 W_2))^2} < \left\|\ddot{F}(\theta)\right\|_F, \tag{40}$$

and

$$\left\|\dot{F}(\theta)\right\|_1 \leq \frac{3W_1^2|W_2| + |W_1|^3}{2[1+\exp(W_1^3 W_2)]} + |W_1| + |W_2|. \tag{41}$$

If we choose $W_2 = 0$, then, for any $L_0, L_1 \geq 0$ there exists a $|W_1|$ sufficiently large such that

$$\frac{W_1^6}{8} < \left\|\ddot{F}(\theta)\right\|_F \not\leq L_0 \left\|\dot{F}(\theta)\right\|_1 + L_1 \leq L_0 \left(\frac{|W_1|^3}{4} + |W_1|\right) + L_1. \tag{42}$$

Hence, $F(\theta)$ is not (L_0, L_1)-smooth.

Moreover, computing the trace of the variance of $\dot{f}(\theta, X)$, we recover

$$\mathbb{E}\left[\left\|\dot{f}(\theta, X) - \dot{F}(\theta)\right\|_2^2\right] = \left(\frac{3W_1^2 W_2}{2[1+\exp(W_1^3 W_2)]}\right)^2 + \left(\frac{W_1^3}{2[1+\exp(W_1^3 W_2)]}\right)^2, \tag{43}$$

which does not satisfy a bounded variance assumption (choose $W_2 = 0$ and let $W_1 \to \infty$). Thus, any work that makes either a global Lipschitz bound on the gradient or a global noise model bound fails to apply to this simple recurrent neural network training problem.

On the other hand, the problem does satisfy our assumptions. In particular,

1. Assumptions 1 and 3 are easily verified.
2. Given that $\dot{F}$ is continuously differentiable, then compactness and continuity of the derivative of $\dot{F}$ imply that it is locally Lipschitz continuous. Therefore, Assumption 2 is satisfied.
3. Moreover,

$$\mathbb{E}\left[\left\|\dot{f}(\theta, X)\right\|_2^2\right] = \left(\frac{3W_1^2 W_2}{2[1+\exp(W_1^3 W_2)]}\right)^2 + \left(\frac{W_1^3}{2[1+\exp(W_1^3 W_2)]}\right)^2 + \left\|\dot{F}(\theta)\right\|_2^2, \tag{44}$$

which is a continuous function. Hence, Assumption 4 is satisfied.

A.4 Poisson Regression

Here, we consider the task of estimating a Poisson regression model for data $X = (Y, Z)$ where Y is a count response variable and Z is the covariate. To make this problem simpler, we will assume that both Y and Z are independent Poisson random variables with mean 1, which implies that the parameter in the model, $\theta^* = 0$. If we use a likelihood framework, then, up to a constant depending on Y,

$$f(\theta, X) = -YZ\theta + \exp(\theta Z), \tag{45}$$

and

$$\dot{f}(\theta, X) = -YZ + Z\exp(\theta Z). \tag{46}$$

From this, we compute

$$F(\theta) = -\theta + \exp(\exp(\theta) - 1), \tag{47}$$

$$\dot{F}(\theta) = -1 + \exp(\exp(\theta) + \theta - 1), \tag{48}$$

and, letting $\nabla^2 F(\psi)|_{\psi=\theta} = \ddot{F}(\theta)$,

$$\ddot{F}(\theta) = (\exp(\theta) + 1)\exp(\exp(\theta) + \theta - 1). \tag{49}$$

We begin by showing that $\dot{F}(\theta)$ is not globally Lipschitz continuous. To do so, for any $\theta > 0$, note

$$|\dot{F}(\theta) - \dot{F}(0)| = \exp(\exp(\theta) + \theta - 1) - 1 > \exp(\theta) - 1 \geq \theta + \theta^2/2. \tag{50}$$

Thus, for any $L > 0$ there exists a $\theta > 0$ such that $|\dot{F}(\theta) - \dot{F}(0)| > L|\theta|$.

We now show that $F(\theta)$ does not satisfy the (L_0, L_1)-smooth assumption. Note, for $\theta \geq 0$,

$$\exp(\exp(\theta) + 2\theta - 1) < \ddot{F}(\theta), \tag{51}$$

and

$$\dot{F}(\theta) < \exp(\exp(\theta) + \theta - 1). \tag{52}$$

It follows that for any $L_0, L_1 > 0$, there exists a $\theta > 0$ such that $L_0|\dot{F}(\theta)| + L_1 < \ddot{F}(\theta)$.

For the noise, we compute the second moment of $\dot{f}(\theta, X)$. That is,

$$\begin{aligned}
\mathbb{E}\left[\dot{f}(\theta, X)^2\right] &= \mathbb{E}\left[Y^2 Z^2 - 2YZ^2\exp(\theta Z) + Z^2\exp(2\theta Z)\right] && (53)\\
&= 4 - 2\mathbb{E}\left[Z^2\exp(\theta Z)\right] + \mathbb{E}\left[Z^2\exp(2\theta Z)\right] && (54)\\
&= 4 - 2(\exp(\theta) + 1)\exp(\exp(\theta) + \theta - 1)\\
&\quad + (\exp(2\theta) + 1)\exp(\exp(2\theta) + 2\theta - 1). && (55)
\end{aligned}$$

It is clear from this calculation that the variance (computed by subtracting off $\dot{F}(\theta)^2$) will diverge as θ tends to infinity. To show that [Bottou et al., 2018, Assumption 4.3c] does not apply, it is enough to show that its generalization, [Khaled and Richtárik, 2020, Assumption 2] does not apply. To this end, we must show that there does not exists a $C_0, C_1, C_2 \geq 0$ such that, $\forall\theta$,

$$\mathbb{E}\left[\dot{f}(\theta, X)^2\right] \leq C_0 + C_1 F(\theta) + C_2|\dot{F}(\theta)|^2. \tag{56}$$

From our calculations, it is easy to verify that $F(\theta)$ and $\dot{F}(\theta)$ are dominated by $\exp(2\exp(\theta))$, and that the second moment of the stochastic gradient is bounded from below by $\exp(\exp(2\theta))$ for $\theta \geq \log(4)$. Hence, for any $C_0, C_1, C_2 \geq 0$, there exists θ sufficiently large such that

$$C_0 + C_1 F(\theta) + C_2|\dot{F}(\theta)|^2 \leq C_0 + (C_1 + C_2)\exp(2\exp(\theta)) < \exp(\exp(2\theta)) \leq \mathbb{E}\left[\dot{f}(\theta, X)^2\right]. \tag{57}$$

Thus, [Bottou et al., 2018, Assumption 4.3c] and [Khaled and Richtárik, 2020, Assumption 2] do not hold.

On the other hand, the problem does satisfy our assumptions. In particular,

1. Assumptions 1 and 3 are easily verified.
2. Given that $\ddot{F}$ is continuous, Assumption 2 is satisfied.
3. Moreover, we can use the calculated value $\mathbb{E}[\dot{f}(\theta, X)^2]$, which is a continuous function, as $G(\theta)$ to satisfy Assumption 4.

A.5 Noiseless Feed Forward Network for Binary Classification

Out of interest, we reconsider the second example but construct a different data distribution that produces noiseless stochastic gradient. Consider the binary classification problem with label Y and feature Z where $(Y, Z) = (0, -1)$ with probability $1/2$ and $(Y, Z) = (1, 1)$ with probability $1/2$. We solve this classification problem using the network shown in Fig. 1 with σ linear and φ sigmoid.

We will train this model using the binary cross entropy loss function. Letting $X = (Y, Z)$ and $\theta = (W_1, W_2, W_3, W_4)$,

$$f(\theta, X) = -Y \log(\hat{y}) - (1 - Y) \log(1 - \hat{y}) + \frac{1}{2} \sum_{i=1}^{4} W_i^2, \tag{58}$$

and

$$\hat{y} = \frac{1}{1 + \exp(-W_4 W_3 W_2 W_1 Z)}. \tag{59}$$

Moreover,

$$\dot{f}(\theta, X) = Z(\hat{y} - Y) \begin{bmatrix} W_4 W_3 W_2 \\ W_4 W_3 W_1 \\ W_4 W_2 W_1 \\ W_3 W_2 W_1 \end{bmatrix} + \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix}, \tag{60}$$

and, consequently,

$$\dot{F}(\theta) = \frac{-1}{1 + \exp(W_4 W_3 W_2 W_1)} \begin{bmatrix} W_4 W_3 W_2 \\ W_4 W_3 W_1 \\ W_4 W_2 W_1 \\ W_3 W_2 W_1 \end{bmatrix} + \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix}. \tag{61}$$

We first establish that $\dot{F}(\theta)$ is not globally Lipschitz continuous. With $\theta = (1, -1, W_3, W_3)$ and $\phi = (1, 0, 0, 0)$, it is enough to find a lower bound for the first component of $\dot{F}(\theta) - \dot{F}(\phi)$, denoted by $\dot{F}_1(\theta) - \dot{F}_1(\phi)$. To this end,

$$|\dot{F}_1(\theta) - \dot{F}_1(\phi)| = \frac{W_3^2}{1 + \exp(-W_3^2)} \geq \frac{1}{2} |W_3 - 0|^2. \tag{62}$$

Thus, $\dot{F}$ is not globally Lipschitz.

On the other hand, the problem does satisfy our assumptions. In particular,

1. Assumptions 1 and 3 are easily verified.
2. Given that $\dot{F}$ is continuously differentiable, then compactness and continuity of the derivative of $\dot{F}$ imply that it is locally Lipschitz continuous. Therefore, Assumption 2 is satisfied.
3. Moreover, $\dot{f}(\theta, Z) = \dot{F}(\theta)$—that is, there $\dot{f}(\theta, Z)$ has zero variance for the distribution that we have constructed. Therefore,

$$\mathbb{E}\left[\left\|\dot{f}(\theta, X)\right\|_2^2\right] = \left\|\dot{F}(\theta)\right\|_2^2, \tag{63}$$

which is a continuous function. Hence, Assumption 4 is satisfied.

B Technical Lemmas

Lemma 5 (Lemma 1)**.** *Suppose Assumptions 1 and 2 hold. Then, for any $\theta, \varphi \in \mathbb{R}^p$,*

$$F(\varphi) - F_{l.b.} \leq F(\theta) - F_{l.b.} + \dot{F}(\theta)'(\varphi - \theta) + \frac{L(\theta, \varphi)}{1 + \alpha} \|\varphi - \theta\|_2^{1+\alpha}. \tag{64}$$

Proof. By Taylor's theorem,

$$F(\varphi) - F_{l.b.} = F(\theta) - F_{l.b.} + \int_0^1 \dot{F}(\theta + t(\varphi - \theta))'(\varphi - \theta) dt. \tag{65}$$

Now, add and subtract $\dot{F}(\theta)$ to $\dot{F}(\theta + t(\varphi - \theta))$ in the integral, then apply Assumption 2. By Definition 1,

$$\left\|\dot{F}(\theta + t(\varphi - \theta)) - \dot{F}(\theta)\right\|_2 \leq L(\theta, \varphi) t^\alpha \|\theta - \varphi\|_2^\alpha. \tag{66}$$

We conclude,

$$F(\varphi) - F_{l.b.} \leq F(\theta) - F_{l.b.} + \dot{F}(\theta)'(\varphi - \theta) + L(\theta, \varphi) \|\varphi - \theta\|_2^{1+\alpha} \int_0^1 t^\alpha dt. \tag{67}$$

By computing the integral, the result follows. □

Lemma 6. *Suppose* $\{M_k : k+1 \in \mathbb{N}\}$ *satisfy Properties 1 and 4. Then* $\forall C > 0$, $\exists K \in \mathbb{N}$ *such that* $\forall k \geq K$,

$$\lambda_{\min}(M_k) - \frac{C}{2}\lambda_{\max}(M_k)^{1+\alpha} \geq \frac{1}{2}\lambda_{\min}(M_k). \tag{68}$$

Proof. Fix $C > 0$. Rearranging the conclusion, we see that it is equivalent to prove that $\exists K \in \mathbb{N}$ such that $\forall k \geq K$, $1/C \geq \lambda_{\max}(M_k)^\alpha \kappa(M_k)$. This follows from Property 4. □

Lemma 7. *For any* $\theta \in \mathbb{R}^p$, $v \in \mathbb{R}$, $L > 0$ *and* $\alpha \in (0, 1]$,

$$\frac{L}{1+\alpha} v^{1+\alpha} - \left\|\dot{F}(\theta)\right\|_2 v \geq -\frac{\alpha}{1+\alpha}\left[\frac{\left\|\dot{F}(\theta)\right\|_2^{1+\alpha}}{L}\right]^{1/\alpha}. \tag{69}$$

Proof. If we minimize the left hand side of the inequality, we see that a minimum value occurs when $v^\alpha = \left\|\dot{F}(\theta)\right\|_2 / L \geq 0$. Solving for v and plugging this back into the left hand side, we conclude that the inequality holds. □

C Global Convergence Analysis

We begin by first deriving a recursive relationship between the optimality gap at iteration $k+1$ and the optimality gap at iteration k on the events $\{\mathcal{B}_j(R)\}$ as defined in (8) for arbitrary $R \geq 0$. Using this result, we then provide an analysis of the convergence of the objective function. Then, we turn our attention to the gradient function. Note, $B(\theta, r)$ is the open ball around θ of radius r.

C.1 A Recursive Relationship

Lemma 8 (Lemma 2)**.** *Let* $\{M_k\}$ *satisfy Property 1. Suppose Assumptions 1 to 4 hold. Let* $\{\theta_k\}$ *satisfy (5). Then,* $\forall R \geq 0$,

$$\begin{aligned}\mathbb{E}\left[\left[F(\theta_{k+1}) - F_{l.b.}\right]\boldsymbol{1}\left[\mathcal{B}_{k+1}(R)\right] \middle| \mathcal{F}_k\right] &\leq \left[F(\theta_k) - F_{l.b.}\right]\boldsymbol{1}\left[\mathcal{B}_k(R)\right] \\ &- \lambda_{\min}(M_k)\left\|\dot{F}(\theta_k)\right\|_2^2 \boldsymbol{1}\left[\mathcal{B}_k(R)\right] + \frac{L_{R+1} + \partial F_R}{1+\alpha}\lambda_{\max}(M_k)^{1+\alpha} G_R,\end{aligned} \tag{70}$$

where $G_R = \sup_{\theta \in \overline{B(0,R)}} G(\theta) < \infty$ *with* $G(\theta)$*; and* $\partial F_R = \sup_{\theta \in \overline{B(0,R)}} \|\dot{F}(\theta)\|_2 (1+\alpha) < \infty$.

Proof. Fix $R \geq 0$. For any $k+1 \in \mathbb{N}$, the definition of local Hölder continuity implies that L_{R+1} is well defined (see Definition 1). Therefore, Lemma 1 implies

$$\begin{aligned}&\left[F(\theta_{k+1}) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_{k+1}(R+1)\right] \\ &\leq \left(\left[F(\theta_k) - F_{l.b.}\right] + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{L_{R+1}}{1+\alpha}\|\theta_{k+1} - \theta_k\|_2^{1+\alpha}\right)\mathbf{1}\left[\mathcal{B}_{k+1}(R+1)\right].\end{aligned} \tag{71}$$

Now, since $\overline{B(0,R)} \subset \overline{B(0,R+1)}$, it also holds true that

$$\begin{aligned}&\left[F(\theta_{k+1}) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] \\ &\leq \left(\left[F(\theta_k) - F_{l.b.}\right] + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{L_{R+1}}{1+\alpha}\|\theta_{k+1} - \theta_k\|_2^{1+\alpha}\right)\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right].\end{aligned} \tag{72}$$

Our goal now is to replace $\mathcal{B}_{k+1}(R)$ on the right hand side by $\mathcal{B}_k(R)$. However, there is a technical difficulty which we must address. First, it follows from the preceding inequality that

$$\begin{aligned}
&[F(\theta_{k+1}) - F_{l.b.}]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] \\
&\leq \left([F(\theta_k) - F_{l.b.}] + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{L_{R+1}}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right) \\
&\quad \times \left(\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] - \mathbf{1}\left[\mathcal{B}_k(R)\right]\right) \\
&\quad + \left([F(\theta_k) - F_{l.b.}] + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{L_{R+1}}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right)\mathbf{1}\left[\mathcal{B}_k(R)\right].
\end{aligned} \tag{73}$$

The first term on the right hand side of the inequality only contributes meaningfully if it is positive. Since $\mathbf{1}\left[\mathcal{B}_k(R)\right] \geq \mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]$, then two statements hold: (i) $\mathbf{1}\left[\mathcal{B}_k(R)\right]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] = \mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]$; and (ii) the first term of the right hand side of (73) is positive if and only if

$$\left([F(\theta_k) - F_{l.b.}] + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{L_{R+1}}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right)\mathbf{1}\left[\mathcal{B}_k(R)\right] < 0. \tag{74}$$

By the choice of L_{R+1}, Assumption 1 and Lemma 1 imply that if (74) occurs, then $\|\theta_{k+1}\|_2 > R + 1 \geq \|\theta_k\|_2 + 1$. By the reverse triangle inequality and (5), if (74) occurs, then $\|M_k\dot{f}(\theta_k, X_{k+1})\|_2 \geq 1$. Hence,

$$\begin{aligned}
&\left([F(\theta_k) - F_{l.b.}] + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{L_{R+1}}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right) \\
&\quad \times \left(\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] - \mathbf{1}\left[\mathcal{B}_k(R)\right]\right) \\
&\leq \left(-[F(\theta_k) - F_{l.b.}] - \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) - \frac{L_{R+1}}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right) \\
&\quad \times \left(\mathbf{1}\left[\mathcal{B}_k(R)\right] - \mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\right)\mathbf{1}\left[\mathcal{B}_k(R)\right]\mathbf{1}\left[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2 \geq 1\right].
\end{aligned} \tag{75}$$

We now compute another coarse upper bound for this inequality. Note, by Assumption 1 and Cauchy-Schwarz,

$$\begin{aligned}
&\left(-[F(\theta_k) - F_{l.b.}] - \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) - \frac{L_{R+1}}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right) \\
&\quad \times \left(\mathbf{1}\left[\mathcal{B}_k(R)\right] - \mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\right)\mathbf{1}\left[\mathcal{B}_k(R)\right]\mathbf{1}\left[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2 \geq 1\right] && (76) \\
&\leq \left\|\dot{F}(\theta_k)\right\|_2\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right]\mathbf{1}\left[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2 \geq 1\right] && (77) \\
&\leq \left\|\dot{F}(\theta_k)\right\|_2\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha}\mathbf{1}\left[\mathcal{B}_k(R)\right] && (78) \\
&\leq \frac{\partial F_R}{1+\alpha}\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha}\mathbf{1}\left[\mathcal{B}_k(R)\right], && (79)
\end{aligned}$$

where $\partial F_R = \sup_{\theta \in \overline{B(0,R)}} \|\dot{F}(\theta)\|_2(1+\alpha) < \infty$ given that $\|\dot{F}(\theta)\|_2$ is a continuous function of θ.

Applying this inequality to (73), we conclude

$$\begin{aligned}
&[F(\theta_{k+1}) - F_{l.b.}]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] \\
&\leq \left([F(\theta_k) - F_{l.b.}] - \dot{F}(\theta_k)'M_k\dot{f}(\theta_k, X_{k+1}) + \frac{L_{R+1} + \partial F_R}{1+\alpha}\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha}\right) \\
&\quad \times \mathbf{1}\left[\mathcal{B}_k(R)\right].
\end{aligned} \tag{80}$$

By Assumption 3,

$$\begin{aligned}
&\mathbb{E}\left[[F(\theta_{k+1}) - F_{l.b.}]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] \middle| \mathcal{F}_k\right] \\
&\leq \left([F(\theta_k) - F_{l.b.}] - \dot{F}(\theta_k)'M_k\dot{F}(\theta_k) + \frac{L_{R+1} + \partial F_R}{1+\alpha}\mathbb{E}\left[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha} \middle| \mathcal{F}_k\right]\right) \\
&\quad \times \mathbf{1}\left[\mathcal{B}_k(R)\right].
\end{aligned} \tag{81}$$

Using Property 1 and Assumption 4,

$$\mathbb{E}\left[\left[F(\theta_{k+1}) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\middle|\,\mathcal{F}_k\right]$$
$$\leq \left(\left[F(\theta_k) - F_{l.b.}\right] - \lambda_{\min}(M_k)\left\|\dot{F}(\theta_k)\right\|_2^2 + \frac{L_{R+1} + \partial F_R}{1+\alpha}\lambda_{\max}(M_k)^{1+\alpha}G(\theta_k)\right)\mathbf{1}\left[\mathcal{B}_k(R)\right]. \tag{82}$$

By Assumption 4, G is upper semicontinuous and $\overline{B(0,R)}$ is compact, which implies that G_R is well defined and finite. The result follows. □

C.2 Objective Function Analysis

Corollary 1. *Let $\{\theta_k\}$ be defined as in (5) satisfying Properties 1 and 2. Suppose Assumptions 1 to 4 hold. Then, there exists a finite random variable $F_{\lim}$ such that on the event $\{\sup_k \|\theta_k\|_2 < \infty\}$, $\lim_{k\to\infty} F(\theta_k) = F_{\lim}$ with probability one.*

Proof. By Lemma 2, for every $R \geq 0$,

$$\mathbb{E}\left[\left[F(\theta_{k+1}) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\middle|\,\mathcal{F}_k\right]$$
$$\leq \left[F(\theta_k) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_k(R)\right] + \frac{(L_{R+1} + \partial F_R)G_R}{1+\alpha}\lambda_{\max}(M_k)^{1+\alpha}. \tag{83}$$

By Neveu and Speed [1975, Exercise II.4] (cf. Robbins and Siegmund [1971]) and Property 2, $\lim_{k\to\infty}[F(\theta_k) - F_{l.b.}]\mathbf{1}\left[\mathcal{B}_k(R)\right]$ converges to a finite random variable with probability one. Since $R \geq 0$ is arbitrary, we conclude that there exists a finite random variable $F_{\lim}$ such that the set $\{\sup_k \|\theta_k\|_2 \leq R\}$ is a subset of $\{\lim_k F(\theta_k) = F_{\lim}\}$ up to a measure zero set. Since the countable union of measure zero sets has measure zero,

$$\left\{\sup_k \|\theta_k\|_2 < \infty\right\} = \bigcup_{R\in\mathbb{N}}\left\{\sup_k \|\theta_k\|_2 \leq R\right\} \subset \left\{\lim_{k\to\infty} F(\theta_k) = F_{\lim}\right\}, \tag{84}$$

up to a measure zero set. The result follows. □

C.3 Gradient Function Analysis

We now prove that the gradient norm evaluated at SGD's iterates must, repeatedly, get arbitrarily close to zero.

Lemma 9. *Let $\{\theta_k\}$ be defined as in (5) satisfying Properties 1 to 3. Suppose Assumptions 1 to 4 hold. Then, $\forall R \geq 0$ and for all $\delta > 0$,*

$$\mathbb{P}\left[\left\|\dot{F}(\theta_k)\right\|_2^2 \boldsymbol{1}\left[\mathcal{B}_k(R)\right] \leq \delta,\ i.o.\right] = 1. \tag{85}$$

Proof. By Lemma 2,

$$\lambda_{\min}(M_k)\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}\left[\mathcal{B}_k(R)\right]\right] \leq \mathbb{E}\left[\left[F(\theta_k) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_k(R)\right]\right]$$
$$- \mathbb{E}\left[\left[F(\theta_{k+1}) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\right] + \frac{(L_{R+1} + \partial F_R)G_R}{1+\alpha}\lambda_{\max}(M_k)^{1+\alpha}. \tag{86}$$

Taking the sum of this equation for all k from 0 to $j \in \mathbb{N}$, we have

$$\sum_{k=0}^{j}\lambda_{\min}(M_k)\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}\left[\mathcal{B}_k(R)\right]\right] \leq \left[F(\theta_0) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_0(R)\right]$$
$$- \mathbb{E}\left[\left[F(\theta_{j+1}) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_{j+1}(R)\right]\right] + \frac{(L_{R+1} + \partial F_R)G_R}{1+\alpha}\sum_{k=0}^{j}\lambda_{\max}(M_k)^{1+\alpha}. \tag{87}$$

By Assumption 1 and Property 2, the right hand side is bounded by

$$\left[F(\theta_0) - F_{l.b.}\right]\mathbf{1}\left[\mathcal{B}_0(R)\right] + \frac{(L_{R+1} + \partial F_R)G_R}{1+\alpha}S, \tag{88}$$

which is finite. Therefore, $\sum_{k=0}^{\infty} \lambda_{\min}(M_k)\mathbb{E}[\|\dot{F}(\theta_k)\|_2^2 \mathbf{1}[\mathcal{B}_k(R)]]$ is finite. Furthermore, by Property 3, $\liminf_k \mathbb{E}[\|\dot{F}(\theta_k)\|_2^2 \mathbf{1}[\mathcal{B}_k(R)]] = 0$.

Now, for any $\delta > 0$, Markov's inequality implies that for all $j+1 \in \mathbb{N}$, and for all $k \geq j$

$$\mathbb{P}\left[\bigcap_{k=j}^{\infty}\left\{\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}[\mathcal{B}_k(R)] > \delta\right\}\right] \leq \mathbb{P}\left[\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}[\mathcal{B}_k(R)] > \delta\right] \tag{89}$$

$$\leq \frac{1}{\delta}\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}[\mathcal{B}_k(R)]\right]. \tag{90}$$

Since the last inequality holds for every $k \geq j$, then, in particular, for all $j+1 \in \mathbb{N}$,

$$\mathbb{P}\left[\bigcap_{k=j}^{\infty}\left\{\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}[\mathcal{B}_k(R)] > \delta\right\}\right] \leq \frac{1}{\delta}\min_{j\leq k}\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}[\mathcal{B}_k(R)]\right], \tag{91}$$

where the right hand side is zero because $\liminf_k \mathbb{E}[\|\dot{F}(\theta_k)\|_2^2 \mathbf{1}[\mathcal{B}_k(R)]] = 0$.

As the countable union of measure zero sets has measure zero, we conclude that for all $\delta > 0$,

$$\mathbb{P}\left[\left\|\dot{F}(\theta_k)\right\|_2^2 \mathbf{1}[\mathcal{B}_k(R)] \leq \delta,\ i.o.\right] = 1. \tag{92}$$

□

Unfortunately, Lemma 9 does not guarantee that the gradient norm will be captured within a region of zero. In order to prove this, we first show that it is not possible (i.e., a zero probability event) for the limit supremum and limit infimum of the gradients to be distinct.

Lemma 10. *Let $\{\theta_k\}$ be defined as in (5) satisfying Properties 1 and 2. Suppose Assumptions 1 to 4 hold. Then, $\forall R \geq 0$ and for all $\delta > 0$,*

$$\mathbb{P}\left[\left\|\dot{F}(\theta_{k+1})\right\|_2 \mathbf{1}[\mathcal{B}_{k+1}(R)] > \delta, \left\|\dot{F}(\theta_k)\right\|_2 \mathbf{1}[\mathcal{B}_k(R)] \leq \delta,\ i.o.\right] = 0. \tag{93}$$

Proof. Let $\gamma > 0$. Let L_R be as in Definition 1, and G_R be as in Lemma 2. Then, for $\delta > 0$,

$$\mathbb{P}\left[\left\|\dot{F}(\theta_{k+1})\right\|_2 \mathbf{1}[\mathcal{B}_{k+1}(R)]\, \mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2 \mathbf{1}[\mathcal{B}_k(R)] \leq \delta\right] > \delta + L_R\gamma^{\alpha}\right] \tag{94}$$

$$= \mathbb{P}\Bigg[\left(\left\|\dot{F}(\theta_{k+1})\right\|_2 - \left\|\dot{F}(\theta_k)\right\|_2 + \left\|\dot{F}(\theta_k)\right\|_2\right) \mathbf{1}[\mathcal{B}_{k+1}(R)] \tag{95}$$

$$\times\, \mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2 \mathbf{1}[\mathcal{B}_k(R)] \leq \delta\right] > \delta + L_R\gamma^{\alpha}\Bigg]. \tag{96}$$

Using the reverse triangle inequality, $\|\dot{F}(\theta_{k+1})\|_2 - \|\dot{F}(\theta_k)\|_2 \leq \|\dot{F}(\theta_{k+1}) - \dot{F}(\theta_k)\|_2$. Now, making use of the restriction to $\mathcal{B}_{k+1}(R)$, $\|\dot{F}(\theta_{k+1}) - \dot{F}(\theta_k)\|_2 \leq L_R\|\theta_{k+1} - \theta_k\|_2^{\alpha}$. Moreover, on $\|\dot{F}(\theta_k)\|_2 \leq \delta$, $\|\dot{F}(\theta_k)\|_2 \mathbf{1}[\mathcal{B}_{k+1}(R)] \leq \delta$. Putting these two observations together,

$$\mathbb{P}\Bigg[\left(\left\|\dot{F}(\theta_{k+1})\right\|_2 - \left\|\dot{F}(\theta_k)\right\|_2 + \left\|\dot{F}(\theta_k)\right\|_2\right) \mathbf{1}[\mathcal{B}_{k+1}(R)] \tag{97}$$

$$\times\, \mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2 \mathbf{1}[\mathcal{B}_k(R)] \leq \delta\right] > \delta + L_R\gamma^{\alpha}\Bigg] \tag{98}$$

$$\leq \mathbb{P}\Bigg[L_R\,\|\theta_{k+1} - \theta_k\|_2^{\alpha}\, \mathbf{1}[\mathcal{B}_{k+1}(R)]\, \mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2 \mathbf{1}[\mathcal{B}_k(R)] \leq \delta\right] > L_R\gamma^{\alpha}\Bigg] \tag{99}$$

$$= \mathbb{P}\Bigg[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2 \mathbf{1}[\mathcal{B}_{k+1}(R)]\, \mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2 \mathbf{1}[\mathcal{B}_k(R)] \leq \delta\right] > \gamma\Bigg]. \tag{100}$$

Now, using $\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right]\right] \leq \mathbf{1}\left[\mathcal{B}_k(R)\right]$,

$$\mathbb{P}\Big[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] \leq \delta\right] > \gamma\Big] \tag{101}$$

$$\leq \mathbb{P}\left[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] > \gamma\right] \tag{102}$$

$$\leq \mathbb{P}\left[\left\|M_k\dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha}\mathbf{1}\left[\mathcal{B}_k(R)\right] > \gamma^{1+\alpha}\right] \tag{103}$$

$$\leq \frac{1}{\gamma^{1+\alpha}}\left\|M_k\right\|_2^{1+\alpha}\mathbb{E}\left[\mathbb{E}\left[\left\|\dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha}\middle|\mathcal{F}_k\right]\mathbf{1}\left[\mathcal{B}_k(R)\right]\right], \tag{104}$$

where the last inequality is a consequence of Markov's inequality, $\|M_k\dot{f}(\theta_k, X_{k+1})\|_2 \leq \|M_k\|_2\|\dot{f}(\theta_k, X_{k+1})\|_2$, and $\mathbf{1}\left[\mathcal{B}_k(R)\right]$ being measurable with respect to $\mathcal{F}_k$.

By Assumption 4, $\mathbb{E}\left[\|\dot{f}(\theta_k, X_{k+1})\|_2^{1+\alpha}\middle|\mathcal{F}_k\right] \leq G(\theta_k)$. Moreover, on $\mathcal{B}_k(R)$, $G(\theta_k) \leq G_R$. Using this in the expectation, we conclude

$$\mathbb{P}\left[\left\|\dot{F}(\theta_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right]\mathbf{1}\left[\left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] \leq \delta\right] > \delta + L_R\gamma^\alpha\right] \leq \frac{1}{\gamma^{1+\alpha}}\left\|M_k\right\|_2^{1+\alpha}G_R. \tag{105}$$

By Property 2, the sum of the last expression over all $k+1 \in \mathbb{N}$ is finite. By the Borel-Cantelli lemma, for all $R \geq 0$, $\delta > 0$ and $\gamma > 0$,

$$\mathbb{P}\left[\left\|\dot{F}(\theta_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > \delta + L_R\gamma^\alpha, \left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] \leq \delta,\ i.o.\right] = 0. \tag{106}$$

Since this holds for any $\gamma > 0$, it will hold for every value in a sequence $\gamma_n \downarrow 0$. Since the countable union of measure zero events has measure zero, for any $R \geq 0$ and $\delta > 0$,

$$\mathbb{P}\left[\left\{\left\|\dot{F}(\theta_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > \delta, \left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] \leq \delta,\ i.o.\right\} \cap \Omega_\delta^c\right] = 0, \tag{107}$$

where $\Omega_\delta = \{\limsup_k \|\dot{F}(\theta_{k+1})\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] = \delta\}$.

We now show that Ω_δ is a probability zero event. Notice, by Lemma 9 and the definition of Ω_δ,

$$\Omega_\delta \subset \left\{\left\|\dot{F}(\theta_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > \delta/2, \left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] \leq \delta/2,\ i.o.\right\} \cap \Omega_{\delta/2}^c, \tag{108}$$

up to a set of measure zero. By applying (107) with $\delta/2$, $\mathbb{P}[\Omega_\delta] = 0$. The conclusion of the result follows. □

We now put together Lemmas 9 and 10 to show that, on the event $\{\sup_k \|\theta_k\|_2 < \infty\}$, $\|\dot{F}(\theta_k)\|_2$ converges to 0 with probability one.

Corollary 2. *Let $\{\theta_k\}$ be defined as in (5) satisfying Properties 1 to 3. Suppose Assumptions 1 to 4 hold. Then, on the event $\{\sup_k \|\theta_k\|_2 < \infty\}$, $\lim_{k\to\infty} \|\dot{F}(\theta_k)\|_2 = 0$.*

Proof. For any $R \geq 0$ and $\delta > 0$, Lemma 9 implies

$$\begin{aligned}&\mathbb{P}\left[\left\|\dot{F}(\theta_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > \delta,\ i.o.\right]\\ &\quad= \mathbb{P}\left[\left\{\left\|\dot{F}(\theta_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > \delta,\ i.o.\right\} \cap \left\{\left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] \leq \delta,\ i.o.\right\}\right].\end{aligned} \tag{109}$$

We see that this latter event is exactly,

$$\mathbb{P}\left[\left\|\dot{F}(\theta_{k+1})\right\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > \delta, \left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\mathcal{B}_k(R)\right] \leq \delta,\ i.o.\right], \tag{110}$$

which, by Lemma 10, is zero with probability one. Therefore, $\mathbb{P}[\|\dot{F}(\theta_{k+1})\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > \delta,\ i.o.]$ is zero. Letting $\delta_n \downarrow 0$ and noting that the countable union of measure zero sets has measure zero, we conclude $\mathbb{P}[\|\dot{F}(\theta_{k+1})\|_2\mathbf{1}\left[\mathcal{B}_{k+1}(R)\right] > 0,\ i.o.] = 0$.

Therefore, for all $R \geq 0$, $\{\sup_k \|\theta_k\|_2 \leq R\} \subset \{\lim_{k\to\infty} \|\dot{F}(\theta_k)\|_2 = 0\}$ up to a measure zero set. Since $\{\sup_k \|\theta_k\|_2 < \infty\} = \cup_{R\in\mathbb{N}}\{\sup_k \|\theta_k\|_2 \leq R\}$, the result follows. □

C.4 Capture Theorem

The final step in our proof is to study the event $\{\sup_k \|\theta_k\| < \infty\}$.

Theorem 4 (Theorem 1)**.** *Let $\{\theta_k\}$ be defined as in (5), and let $\{M_k\}$ satisfy Properties 1 and 2. If Assumption 4 holds, then either $\{\lim_{k\to\infty} \theta_k \text{ exists}\}$ or $\{\liminf_{k\to\infty} \|\theta_k\|_2 = \infty\}$ must occur.*

Proof. Let $\bar{\theta} \in \mathbb{R}^p$. Fix $R \geq 0$ and let $\gamma > 0$. Then,

$$\begin{aligned}
&\mathbb{P}\left[\left\|\theta_{k+1} - \bar{\theta}\right\|_2 \geq R + \gamma, \left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \\
&\quad = \mathbb{P}\left[\left\|\theta_{k+1} - \bar{\theta}\right\|_2 \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \geq R + \gamma\right] && (111)\\
&\quad = \mathbb{P}\left[\left(\left\|\theta_{k+1} - \bar{\theta}\right\|_2 - \left\|\theta_k - \bar{\theta}\right\|_2 + \left\|\theta_k - \bar{\theta}\right\|_2\right) \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \geq R + \gamma\right]. && (112)
\end{aligned}$$

Now, $\|\theta_k - \bar{\theta}\|_2 \mathbf{1}\left[\|\theta_k - \bar{\theta}\|_2 \leq R\right] \leq R$. Therefore,

$$\begin{aligned}
&\mathbb{P}\left[\left(\left\|\theta_{k+1} - \bar{\theta}\right\|_2 - \left\|\theta_k - \bar{\theta}\right\|_2 + \left\|\theta_k - \bar{\theta}\right\|_2\right) \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \geq R + \gamma\right] && (113)\\
&\quad \leq \mathbb{P}\left[\left(\left\|\theta_{k+1} - \bar{\theta}\right\|_2 - \left\|\theta_k - \bar{\theta}\right\|_2\right) \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] + R \geq R + \gamma\right] && (114)\\
&\quad \leq \mathbb{P}\left[\|\theta_{k+1} - \theta_k\|_2 \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \geq \gamma\right], && (115)
\end{aligned}$$

where the last line follows by applying the reverse triangle inequality. By using (5) and Markov's inequality,

$$\begin{aligned}
&\mathbb{P}\left[\|\theta_{k+1} - \theta_k\|_2 \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \geq \gamma\right] && (116)\\
&\quad \leq \mathbb{P}\left[\left\|M_k \dot{f}(\theta_k, X_{k+1})\right\|_2 \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \geq \gamma\right] && (117)\\
&\quad \leq \frac{1}{\gamma^{1+\alpha}} \|M_k\|_2^{1+\alpha} \mathbb{E}\left[\mathbb{E}\left[\left\|\dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha} \middle| \mathcal{F}_k\right] \mathbf{1}\left[\left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right]\right]. && (118)
\end{aligned}$$

By applying Assumption 4, $\mathbb{E}\left[\|\dot{f}(\theta_k, X_{k+1})\|_2^{1+\alpha} \middle| \mathcal{F}_k\right] \leq G(\theta_k)$. Moreover, on $\|\theta_k - \bar{\theta}\|_2 \leq R$, $G(\theta_k) \leq \sup_{\theta:\|\theta\|_2 \leq R + \|\bar{\theta}\|_2} G(\theta) =: G_{R+\|\bar{\theta}\|_2} < \infty$ since G is upper semi-continuous. Combining these steps,

$$\mathbb{P}\left[\left\|\theta_{k+1} - \bar{\theta}\right\|_2 \geq R + \gamma, \left\|\theta_k - \bar{\theta}\right\|_2 \leq R\right] \leq \frac{1}{\gamma^{1+\alpha}} \|M_k\|_2^{1+\alpha} G_{R+\|\bar{\theta}\|_2}, \tag{119}$$

By Property 2, we see that the sum of the probabilities is finite. Together with the Borel-Cantelli lemma, $\forall R \geq 0$, $\forall \gamma > 0$, and for all $\bar{\theta} \in \mathbb{R}^p$,

$$\mathbb{P}\left[\|\theta_{k+1} - \bar{\theta}\|_2 \geq R + \gamma,\ \|\theta_k - \bar{\theta}\|_2 \leq R,\ i.o.\right] = 0. \tag{120}$$

Since $\gamma > 0$ is arbitrary, we can show that this statement holds for a countable sequence of $\gamma_n \downarrow 0$. Therefore, $\forall R \geq 0$ and all $\bar{\theta} \in \mathbb{R}^p$,

$$\mathbb{P}\left[\limsup_k \left\|\theta_k - \bar{\theta}\right\|_2 > R, \liminf_k \left\|\theta_k - \bar{\theta}_2\right\|_2 \leq R\right] = 0. \tag{121}$$

Since R is arbitrary, we conclude that for any ordering of positive rational numbers, $\{R_n\}$, $\mathbb{P}[\limsup_k \|\theta_{k+1} - \bar{\theta}\|_2 > R_n, \liminf_k \|\theta_k - \bar{\theta}\|_2 \leq R_n] = 0$ for every n. Again, the countable union of measure zero sets is measure zero. Hence, we conclude that $\mathbb{P}[\limsup_k \|\theta_k - \bar{\theta}\|_2 > \liminf_k \|\theta_k - \bar{\theta}\|_2] = 0$. Thus, either $\lim_k \|\theta_k - \bar{\theta}\|_2$ exists and is either infinite or finite.

Moreover, on the event that the limit is finite, since $\bar{\theta}$ is arbitrary, we can choose $p + 1$ distinct values of $\bar{\theta}$ which do not belong to a hyperplane of dimension smaller than p, and, by triangulation, the $\lim_k \theta_k$ converges to a fixed point (up to a set of measure zero). □

D Stability Analysis

We begin with a recursive relationship on the events $\{\tau_j > k\}$. We use this result to prove that the objective function converges to a finite limit on these events. Then, we use this result to conclude that the gradient function visits to a region of zero on the same event. Finally, we study this event to establish that the two statements above hold with probability one.

D.1 A Recursive Relationship

Lemma 11 (Lemma 3)**.** *Let $\{M_k\}$ satisfy Property 1. Suppose Assumptions 1 to 4 hold. Let $\{\theta_k\}$ satisfy (5). Then, for any $j+1 \in \mathbb{N}$ and $k > j$,*

$$\mathbb{E}\left[\left(F(\theta_{k+1}) - F_{l.b.}\right)\boldsymbol{1}\left[\tau_j > k\right] \middle| \mathcal{F}_k\right] \leq \left(F(\theta_k) - F_{l.b.} - \dot{F}(\theta_k)' M_k \dot{F}(\theta_k)\right)\boldsymbol{1}\left[\tau_j > k-1\right]$$
$$+ \frac{\lambda_{\max}(M_k)^{1+\alpha}}{1+\alpha}\left[\mathcal{L}_\epsilon(\theta_k)G(\theta_k) + \alpha\left[\frac{\left\|\dot{F}(\theta_k)\right\|_2^{1+\alpha}}{\mathcal{L}_\epsilon(\theta_k)}\right]^{1/\alpha}\right]\boldsymbol{1}\left[\tau_j > k-1\right]. \tag{122}$$

Proof. By the construction of τ_j, when $\tau_j > k$, then

$$F(\theta_{k+1}) - F_{l.b.} \leq F(\theta_k) - F_{l.b.} + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}. \tag{123}$$

Using this relationship and using $\mathbf{1}\left[\tau_j > k\right] = \mathbf{1}\left[\tau_j > k-1\right] - \mathbf{1}\left[\tau_j = k\right]$,

$$\begin{aligned}
&\mathbb{E}\left[\left\{F(\theta_{k+1}) - F_{l.b.}\right\}\mathbf{1}\left[\tau_j > k\right] \middle| \mathcal{F}_k\right] \\
&\leq \mathbb{E}\left[\left\{F(\theta_k) - F_{l.b.} + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right\}\mathbf{1}\left[\tau_j > k-1\right] \middle| \mathcal{F}_k\right] \\
&- \mathbb{E}\left[\left\{F(\theta_k) - F_{l.b.} + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right\}\mathbf{1}\left[\tau_j = k\right] \middle| \mathcal{F}_k\right]
\end{aligned} \tag{124}$$

For the first term on the right hand side, we can apply Assumptions 3 and 4, Property 1, and (5) to calculate

$$\begin{aligned}
&\mathbb{E}\left[\left\{F(\theta_k) - F_{l.b.} + \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) + \frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha}\right\}\mathbf{1}\left[\tau_j > k-1\right] \middle| \mathcal{F}_k\right] \\
&\leq \left\{F(\theta_k) - F_{l.b.} - \dot{F}(\theta_k)' M_k \dot{F}(\theta_k) + \frac{\lambda_{\max}(M_k)^{1+\alpha}}{1+\alpha}\mathcal{L}_\epsilon(\theta_k)G(\theta_k)\right\}\mathbf{1}\left[\tau_j > k-1\right].
\end{aligned} \tag{125}$$

For the second term on the right hand side of (124), we require two facts. The first fact is $\mathbf{1}\left[\tau_j = k\right] \leq \mathbf{1}\left[\tau_j > k-1\right]$ which implies $\mathbf{1}\left[\tau_j = k\right] = \mathbf{1}\left[\tau_j = k\right]\mathbf{1}\left[\tau_j > k-1\right]$. For the second fact, the Cauchy-Schwarz inequality and Lemma 7 imply

$$\begin{aligned}
&- F(\theta_k)' M_k \dot{f}(\theta_k, X_{k+1}) + \frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|M_k f(\theta_k, X_{k+1})\right\|_2^{1+\alpha} \\
&\geq -\left\|F(\theta_k)\right\|_2 \left\|M_k \dot{f}(\theta_k, X_{k+1})\right\|_2 + \frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|M_k f(\theta_k, X_{k+1})\right\|_2^{1+\alpha} && (126) \\
&\geq -\frac{\alpha}{1+\alpha}\left[\frac{\left\|\dot{F}(\theta_k)\right\|_2^{1+\alpha}}{\mathcal{L}_\epsilon(\theta_k)}\right]^{1/\alpha}. && (127)
\end{aligned}$$

Hence, using (5),

$$\begin{aligned}
&- \left[F(\theta_k) - F_{l.b.}\right] - \dot{F}(\theta_k)'(\theta_{k+1} - \theta_k) - \frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|\theta_{k+1} - \theta_k\right\|_2^{1+\alpha} \\
&\leq -\left[F(\theta_k) - F_{l.b.}\right] + \frac{\alpha}{1+\alpha}\left[\frac{\left\|\dot{F}(\theta_k)\right\|_2^{1+\alpha}}{\mathcal{L}_\epsilon(\theta_k)}\right]^{1/\alpha}.
\end{aligned} \tag{128}$$

Putting together these two preceding facts together,

$$
\begin{aligned}
&-\mathbb{E}\left[\left\{F(\theta_k)-F_{l.b.}+\dot{F}(\theta_k)'(\theta_{k+1}-\theta_k)+\frac{\mathcal{L}_\epsilon(\theta_k)}{1+\alpha}\left\|\theta_{k+1}-\theta_k\right\|_2^{1+\alpha}\right\}\mathbf{1}\left[\tau_j=k\right]\middle|\mathcal{F}_k\right] \\
&\leq\left\{-\left[F(\theta_k)-F_{l.b.}\right]+\frac{\alpha}{1+\alpha}\left[\frac{\left\|\dot{F}(\theta_k)\right\|_2^{1+\alpha}}{\mathcal{L}_\epsilon(\theta_k)}\right]^{1/\alpha}\right\}\mathbb{P}\left[\tau_j=k|\,\mathcal{F}_k\right]\mathbf{1}\left[\tau_j>k-1\right] \qquad (129)\\
&\leq\frac{\alpha\lambda_{\max}(M_k)^{1+\alpha}}{1+\alpha}\left[\frac{\left\|\dot{F}(\theta_k)\right\|_2^{1+\alpha}}{\mathcal{L}_\epsilon(\theta_k)}\right]^{1/\alpha}\mathbf{1}\left[\tau_j>k-1\right], \qquad (130)
\end{aligned}
$$

where we bound $\mathbb{P}[\tau_j=k|\mathcal{F}_k]$ using Theorem 5. By applying the bounds on the first term, (125), and second term, (130), to (124), the result follows. □

By applying Assumption 5 to Lemma 3, we have the following simplified form.

Lemma 12 (Lemma 4)**.** *If Assumptions 1 to 5, and Properties 1 and 4 hold, and $\{\theta_k\}$ satisfy (5), then there exists a $K\in\mathbb{N}$ such that for any $j+1\in\mathbb{N}$ and any $k\geq\min\{K,j+1\}$,*

$$
\begin{aligned}
&\mathbb{E}\left[(F(\theta_{k+1})-F_{l.b.})\boldsymbol{1}\left[\tau_j>k\right]|\,\mathcal{F}_k\right] \\
&\leq\left(1+\lambda_{\max}(M_k)^{1+\alpha}\frac{C_2}{1+\alpha}\right)(F(\theta_k)-F_{l.b.})\boldsymbol{1}\left[\tau_j>k-1\right] \\
&-\frac{1}{2}\lambda_{\min}(M_k)\left\|\dot{F}(\theta_k)\right\|_2^2\boldsymbol{1}\left[\tau_j>k-1\right]+\lambda_{\max}(M_k)^{1+\alpha}\frac{C_1}{1+\alpha}.
\end{aligned} \qquad (131)
$$

Proof. The result follows by first using Assumption 5 in Lemma 3. Then, collecting similar terms, we apply Lemma 6 to find K. □

D.2 Objective Function Analysis

With this recursive formula, we now have the first result.

Corollary 3. *If Assumptions 1 to 5 and Properties 1, 2 and 4 hold, and $\{\theta_k\}$ satisfy (5), then $\lim_{k\to\infty}F(\theta_k)$ exists and is finite on $\cup_{j=0}^{\infty}\{\tau_j=\infty\}$.*

Proof. By Lemma 4 and Robbins and Siegmund [1971], Neveu and Speed [1975, Exercise II.4], the limit as k goes to infinity of $(F(\theta_k)-F_{l.b.})\mathbf{1}\left[\tau_j>k-1\right]$ exists with probability one and is integrable. Therefore, on the event $\{\tau_j=\infty\}$, the limit of $F(\theta_k)-F_{l.b.}$ exists and is integrable. As a result, the limit of $F(\theta_k)-F_{l.b.}$ exists and is finite on $\cup_{j=0}^{\infty}\{\tau_j=\infty\}$. □

Additionally, we can state the following useful result.

Lemma 13. *If Assumptions 1 to 5, and Properties 1, 2 and 4 hold, and $\{\theta_k\}$ satisfy (5), then $\exists K\in\mathbb{N}$ such that for any $j>K$, $\exists N_j>0$ for which*

$$
\sup_{k>j}\mathbb{E}\left[(F(\theta_k)-F_{l.b.})\boldsymbol{1}\left[\tau_j>k-1\right]\right]\leq N_j. \qquad (132)
$$

Proof. In Lemma 4, we upper bound the right hand side by removing the negative term, and, by Property 2, we add $C_1(1+\alpha)^{-1}\sum_{\ell=k+1}^{\infty}\lambda_{\max}(M_\ell)^{1+\alpha}$ to both side. Then, for all $k\geq j$,

$$
\begin{aligned}
&\mathbb{E}\left[(F(\theta_{k+1})-F_{l.b.})\mathbf{1}\left[\tau_j>k\right]\right]+\frac{C_1}{1+\alpha}\sum_{\ell=k+1}^{\infty}\lambda_{\max}(M_\ell)^{1+\alpha} \\
&\leq\left(1+\lambda_{\max}(M_k)^{1+\alpha}\frac{C_2}{1+\alpha}\right)\mathbb{E}\left[(F(\theta_k)-F_{l.b.})\mathbf{1}\left[\tau_j>k-1\right]\right]+\frac{C_1}{1+\alpha}\sum_{\ell=k}^{\infty}\lambda_{\max}(M_\ell)^{1+\alpha}.
\end{aligned} \qquad (133)
$$

Using $1 + C_2(1+\alpha)^{-1}\lambda_{\max}(M_k)^{1+\alpha} \leq \exp(C_2(1+\alpha)^{-1}\lambda_{\max}(M_k)^{1+\alpha})$, it follows

$$\begin{aligned}
&\mathbb{E}\left[(F(\theta_{k+1}) - F_{l.b.})\mathbf{1}\left[\tau_j > k\right]\right] + \frac{C_1}{1+\alpha}\sum_{\ell=k+1}^{\infty}\lambda_{\max}(M_\ell)^{1+\alpha} \\
&\leq \exp\left(\frac{C_2}{1+\alpha}\lambda_{\max}(M_k)^{1+\alpha}\right)\left[\mathbb{E}\left[(F(\theta_k) - F_{l.b.})\mathbf{1}\left[\tau_j > k-1\right]\right] + \frac{C_1}{1+\alpha}\sum_{\ell=k}^{\infty}\lambda_{\max}(M_\ell)^{1+\alpha}\right].
\end{aligned} \tag{134}$$

Hence,

$$\begin{aligned}
&\mathbb{E}\left[(F(\theta_{k+1}) - F_{l.b.})\mathbf{1}\left[\tau_j > k\right]\right] + \frac{C_1}{1+\alpha}\sum_{\ell=k+1}^{\infty}\lambda_{\max}(M_\ell)^{1+\alpha} \\
&\leq \exp\left(\frac{C_2}{1+\alpha}\sum_{\ell=j}^{k}\lambda_{\max}(M_\ell)^{1+\alpha}\right)\left[\mathbb{E}\left[(F(\theta_j) - F_{l.b.})\right] + \frac{C_1}{1+\alpha}\sum_{\ell=j}^{\infty}\lambda_{\max}(M_\ell)^{1+\alpha}\right],
\end{aligned} \tag{135}$$

where we have used $\mathbf{1}\left[\tau_j > j-1\right] = 1$. By Property 2, the summation in the exponent is finite, which implies the result. □

D.3 Gradient Function Analysis

Lemma 14. *If Assumptions 1 to 5, and Properties 1 to 4 hold, and $\{\theta_k\}$ satisfy (5), then, for any $\delta > 0$,*

$$\mathbb{P}\left[\left\|\dot{F}(\theta_k)\right\|_2 \boldsymbol{1}\left[\tau_j > k-1\right] \leq \delta \text{ i.o.}\right] = 1. \tag{136}$$

Proof. By Lemma 4,

$$\begin{aligned}
&\frac{1}{2}\lambda_{\min}(M_k)\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2\mathbf{1}\left[\tau_j > k-1\right]\right] \leq \mathbb{E}\left[(F(\theta_k) - F_{l.b.})\mathbf{1}\left[\tau_j > k-1\right]\right] \\
&\quad - \mathbb{E}\left[(F(\theta_{k+1}) - F_{l.b.})\mathbf{1}\left[\tau_j > k\right]\right] + \frac{C_2}{1+\alpha}\lambda_{\max}(M_k)^{1+\alpha}\mathbb{E}\left[(F(\theta_k) - F_{l.b.})\mathbf{1}\left[\tau_j > k-1\right]\right] \\
&\quad + \frac{C_1}{1+\alpha}\lambda_{\max}(M_k)^{1+\alpha}.
\end{aligned} \tag{137}$$

By applying Lemma 13,

$$\begin{aligned}
&\frac{1}{2}\lambda_{\min}(M_k)\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2\mathbf{1}\left[\tau_j > k-1\right]\right] \leq \mathbb{E}\left[(F(\theta_k) - F_{l.b.})\mathbf{1}\left[\tau_j > k-1\right]\right] \\
&\quad - \mathbb{E}\left[(F(\theta_{k+1}) - F_{l.b.})\mathbf{1}\left[\tau_j > k\right]\right] + \lambda_{\max}(M_k)^{1+\alpha}\left(\frac{C_2N_j + C_1}{1+\alpha}\right).
\end{aligned} \tag{138}$$

By summing and using Assumption 1,

$$\begin{aligned}
&\frac{1}{2}\sum_{k=j}^{\infty}\lambda_{\min}(M_k)\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2\mathbf{1}\left[\tau_j > k-1\right]\right] \\
&\quad \leq \mathbb{E}\left[F(\theta_j) - F_{l.b.}\right] + \frac{C_2N_j + C_1}{1+\alpha}\sum_{k=j}^{\infty}\lambda_{\max}(M_k)^{1+\alpha}.
\end{aligned} \tag{139}$$

By Property 2, the right hand side is bounded. Now, by Property 3,

$$\liminf_{k\to\infty}\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2\mathbf{1}\left[\tau_j > k-1\right]\right] = 0. \tag{140}$$

Using Markov's inequality, for any $\ell \in \mathbb{N}$ and any $\delta > 0$,

$$\mathbb{P}\left[\bigcap_{k=\ell}^{\infty}\left\{\left\|\dot{F}(\theta_k)\right\|_2\mathbf{1}\left[\tau_j > k-1\right] > \delta\right\}\right] \leq \frac{1}{\delta^2}\min_{k\geq\ell}\mathbb{E}\left[\left\|\dot{F}(\theta_k)\right\|_2^2\mathbf{1}\left[\tau_j > k-1\right]\right] = 0. \tag{141}$$

As the countable union of sets of measure zero have measure zero, the result follows. □

D.4 Stopping Time Analysis

WE compute the probability of $\{\tau_j = k\}$.

Theorem 5. *Let $\{\tau_j : j+1 \in \mathbb{N}\}$ be defined as in (11). If Assumptions 1, 2 and 4 and Property 1 hold, and $\{\theta_k\}$ satisfy (5), then, for any $j+1 \in \mathbb{N}$ and any $k+1 \in \mathbb{N}$,*

$$\mathbb{P}\left[\tau_j = k \middle| \mathcal{F}_k\right] \leq \begin{cases} 0 & k \leq j, \\ \lambda_{\max}(M_k)^{1+\alpha} & k > j. \end{cases} \tag{142}$$

Moreover, if Property 2 also holds, then $\mathbb{P}\left[\cup_{j=0}^{\infty}\{\tau_j = \infty\}\right] = 1$.

Proof. The case of $k \leq j$ is trivial. So consider only $k > j$. By the construction of $L(\cdot,\cdot)$ and $\mathcal{L}_\epsilon(\cdot)$, $\omega \in \{L(\theta_k, \theta_{k+1}) > \mathcal{L}_\epsilon(\theta_k)\}$ implies $\omega \in \{\|\theta_{k+1} - \theta_k\|_2 > (G(\theta_k) \vee \epsilon)^{\frac{1}{1+\alpha}}\}$. Using (5), Markov's inequality, Property 1, we conclude

$$\mathbb{P}\left[\tau_j = k \middle| \mathcal{F}_k\right] \leq \mathbb{P}\left[\left\|M_k \dot{f}(\theta_k, X_{k+1})\right\|_2^{1+\alpha} > G(\theta_k) \vee \epsilon \middle| \mathcal{F}_k\right] \tag{143}$$

$$\leq \frac{\lambda_{\max}(M_k)^{1+\alpha}\mathbb{E}\left[\left\|\dot{f}(\theta_k, X_{k+1}\right\|_2^{1+\alpha} \middle| \mathcal{F}_k\right]}{G(\theta_k) \vee \epsilon}. \tag{144}$$

Applying Assumption 4 supplies the bound on $\mathbb{P}\left[\tau_j = k \middle| \mathcal{F}_k\right]$. For the second part, note

$$\mathbb{P}\left[\tau_j = \infty\right] \geq 1 - \mathbb{P}\left[\tau_j < \infty\right] \geq 1 - \sum_{k=j+1}^{\infty} \lambda_{\max}(M_k)^{1+\alpha}. \tag{145}$$

Therefore,

$$\mathbb{P}\left[\bigcup_{j=0}^{\infty}\{\tau_j = \infty\}\right] = \lim_{j\to\infty} \mathbb{P}\left[\tau_j = \infty\right]. \tag{146}$$

Since $\lim_j \mathbb{P}[\tau_j = \infty] \geq 1 - \lim_j \sum_{k=j+1}^{\infty} \lambda_{\max}(M_k)^{1+\alpha}$, applying Property 2 supplies the final result. □