



# Stopping criteria for, and strong convergence of, stochastic gradient descent on Bottou-Curtis-Nocedal functions

Vivak Patel<sup>1</sup> 

Received: 1 April 2020 / Accepted: 19 September 2021 / Published online: 29 October 2021

© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2021

## Abstract

Stopping criteria for Stochastic Gradient Descent (SGD) methods play important roles from enabling adaptive step size schemes to providing rigor for downstream analyses such as asymptotic inference. Unfortunately, current stopping criteria for SGD methods are often heuristics that rely on asymptotic normality results or convergence to stationary distributions, which may fail to exist for nonconvex functions and, thereby, limit the applicability of such stopping criteria. To address this issue, in this work, we rigorously develop two stopping criteria for SGD that can be applied to a broad class of nonconvex functions, which we term Bottou-Curtis-Nocedal functions. Moreover, as a prerequisite for developing these stopping criteria, we prove that the gradient function evaluated at SGD's iterates converges strongly to zero for Bottou-Curtis-Nocedal functions, which addresses an open question in the SGD literature. As a result of our work, our rigorously developed stopping criteria can be used to develop new adaptive step size schemes or bolster other downstream analyses for nonconvex functions.

**Keywords** Stochastic gradient descent · Nonconvex · Stopping criteria · Strong convergence

**Mathematics Subject Classification** 65K05 · 68Q25 · 90C06 · 90C30 · 68T05

## 1 Introduction

In data-driven and simulation-based disciplines, the optimization problem

$$\min_{\theta} F(\theta) \tag{1}$$

---

This work is supported by the Wisconsin Alumni Research Foundation.

---

✉ Vivak Patel  
vivak.patel@wisc.edu

<sup>1</sup> Department of Statistics, University of Wisconsin, Madison, USA

is frequently solved, where  $F(\theta) = \mathbb{E}[f(\theta, X)]$ ;  $f : \mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}$ ;  $X$  is a random variable with arbitrary support; and  $\mathbb{E}$  is the expectation operator. Depending on  $X$ , the optimization problem's objective function and its gradient may be impractical or impossible to evaluate directly [2]. Fortunately, when certain regularity conditions hold on  $f$  and  $X$  (e.g., [4,49]), the optimization problem's structure is exploited to generate solvers that use the gradient of  $f$  with respect to the argument  $\theta$  for independent copies of  $X$  [7,41]. Moreover, when the gradient of  $f$  is significantly cheaper to compute than the gradient of  $F$ , the optimization problem can be efficiently solved using these so-called Stochastic Gradient Descent (SGD) methods [2].

Despite the potential efficiency of SGD methods, they are fickle algorithms as they are highly sensitive to the choice of step sizes [33,36]; and, consequently, SGD robustification using adaptive stepping procedures has been an active area of research since the inception of stochastic approximation procedures [11,12,14,24,32,43,44].<sup>1</sup> Critically, SGD methods with adaptive stepping procedures lay out *criteria* for when to *stop* the current step size schedule, and then provide a procedure for how to modify it. Thus, SGD's robustification is inherently tied to stopping criteria (see §17 of [37]).

While their utility in the adaptive step size selection for SGD methods cannot be overstated, stopping criteria play other, equally important roles. In the deterministic setting, stopping criteria determine whether asymptotic inference tools can be applied (e.g., Theorem 5.23 in [45]), serve as a proxy for generalization quality [38],<sup>2</sup> and indicate when to terminate an algorithm. Even in the stochastic setting, these applications of stopping criteria are essential and have warranted investigation.

Historically and even recently, stopping criteria—either studied independently or within the context of adaptive stepsize rules—have focused almost entirely on either determining when a confidence region of the iterates contains a strict minimizer [35,43,44,46,51], or have focused on determining when the iterates satisfy a notion of statistical stationarity [5,32,53]. Unfortunately, such stopping criteria are often limited on modern problems. For example, stopping criteria that rely on finding confidence regions that contain a strict minimizer are often limited to strongly convex problems,<sup>3</sup> and are not applicable more broadly, say, to the increasingly common class of overparametrized, nonconvex optimization problems found in machine learning [1]. Moreover, such confidence-region stopping criteria have had limited rigorous development except in special cases [35,37,46].

Thus, stopping criteria that rely on statistical stationarity would seem preferable, yet these criteria face two difficulties. First, such stopping criteria are generally only employable with sufficiently small step sizes (e.g., [53]), where the degree of smallness has yet to be clearly specified. Additionally, such stopping criteria are rather brittle for general nonconvex problems for two reasons. First, such stopping criteria may fail to exist as a stationary distribution need not exist. As an example the stopping criterion in [53] depends on the iterates achieving a stationary distribution, which cannot occur, say, for the simple function  $F(\theta) = \mathbb{E}[(X - \theta_1\theta_2 X)^2]$ , which is a toy example

<sup>1</sup> For a review of recent approaches to adaptive step size procedures, see [8].

<sup>2</sup> There has been some recent work that disagrees with whether early termination leads to better generalization. See [1] and related work. However, even in this case, one needs to know that a minimizer is achieved.

<sup>3</sup> We can include those problems that are nonconvex, yet are locally strongly convex around minimizers.

of overparametrized models that are increasingly common in machine learning. Second, such stopping criteria make use of only a single iterate sequence (i.e., a single “Markov chain”) from which to determine statistical stationarity, which is highly unreliable based on more rigorous stationarity detection methods used in Markov Chain Monte Carlo methods [42]. Unfortunately, just as for confidence-region stopping criteria, these statistical stationarity stopping criteria for stochastic optimization have had limited rigorous development, excepting special cases [5,32,37].

To begin addressing the aforementioned limitations, we expand the scope of stopping criteria for stochastic optimization by rigorously developing analogues to gradient-based stopping criteria used in deterministic optimization. By taking this gradient based perspective to stopping criteria,

1. we are able to supply stopping criteria that are applicable to any differentiable nonconvex problems with unbiased stochastic gradients, which renders our criteria applicable to a broader set of problems in comparison to confidence-region and statistical stationarity stopping criteria; and,
2. by focusing on a rigorous understanding of our stopping criteria, we demonstrate that the gradient of  $F$  evaluated at the iterates converges to zero with probability one, which improves the analysis of stochastic gradient methods for a broad class of nonconvex problems and noise models.<sup>4</sup>

As a result, our stopping criteria lay a rigorous foundation for developing adaptive step size algorithms, and our stopping criteria can be leveraged for other downstream analyses such as developing statistical uncertainty sets.

The remainder of this paper is organized as follows. In Sect. 2, we provide an overview of important notation. In Sect. 3, we detail the challenges of rigorously developing stopping criteria, and how our work contributes to the analysis of stochastic gradient descent on nonconvex functions. In Sect. 4, we review, catalog, and organize common conditions on  $F$  that specify its nonconvexity; and, in Sect. 4.4, we specify a general class of nonconvex functions which we refer to as Bottou-Curtis-Nocedal functions [4]. In Sect. 5, we precisely specify the stochastic gradient descent iterates (Sect. 5.1), and prove that the gradient evaluated at the iterates converge strongly to zero. In Sect. 6, we introduce and analyze our two stopping criteria. In Sect. 7, we examine the behavior of our stopping criteria on a real data analysis task. In Sect. 8, we conclude this work.

## 2 Notation

Below is a list of notation that is used throughout the text. The notation is organized into several groups: notation related to numerical analysis; notation related to probability; notation related the stochastic gradient method and stopping criteria; notation related to conditions placed on the objective function.

<sup>4</sup> Including a number of reports that came out after a preprint of this work was made public. See [18,23,31].

*Numerical Analysis*

$F$  Objective function that maps  $\mathbb{R}^p \rightarrow \mathbb{R}$ . Also referred to as “Deterministic Objective.”

$\dot{F}$  Gradient function of the deterministic objective that maps  $\mathbb{R}^p \rightarrow \mathbb{R}^p$ . Also referred to as “Deterministic Gradient.”

$f$  A function that maps  $\mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}$ . Used to construct the deterministic objective, and is referred to as the “Stochastic Objective.”

$\dot{f}$  Gradient function of the stochastic objective with respect to the first argument, and maps  $\mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}^p$ . Also referred to as the “Stochastic Gradient.”

$\theta$  The parameter to be optimized in the objective function. Takes value in  $\mathbb{R}^p$ .

$\|\cdot\|_2$  If the argument is a vector, this refers to the Euclidean norm of the vector. If the argument is a matrix, this refers to the Operator 2-norm.

$\lambda_{\max}$  The largest eigenvalue corresponding to a given symmetric, positive definite matrix argument.

$\lambda_{\min}$  The smallest eigenvalue corresponding to the given symmetric, positive definite matrix argument.

$\kappa$  The 2-norm condition number of the given matrix.

*Probability*

$\mathbb{P}$  The probability operator with respect to the underlying Probability Space.

$\mathbb{E}$  The expectation operator corresponding to the underlying Probability Space.

$X, X_k$  Independent, identically distributed random variables that map from the underlying Probability Space to  $\mathbb{R}^d$ .

$\tau, T_j$  Stopping times defined with respect to an appropriate filtration.

$\mathcal{F}_k$   $\sigma$ -algebras.  $\{\mathcal{F}_k\}$  denotes the filtration generated by  $\{X_k\}$ .

$\Delta$  Symmetric set difference operator.

$i.o.$  Infinitely often.

$\mathbf{1}$  The indicator function over a specified event.

*Stochastic Gradient Method and Stopping Criteria* $\beta_k$  $M_k$  $K$  $\psi_k$  $P_k$  $Z_k$  $\mathcal{G}_k$  $S$  $\mathbf{P}$  $\mathbf{SC}$  $F_{l,b}$  $C$  $\pi$  $\mathbf{IS}$  $\mathbf{DO}$  $\mathbf{SO}$  $\mathbf{DG}$  $\mathbf{SG}$  $\mathbf{NM}$  $\mathbf{T}$  $\mathbf{BCN}$ 

The iterates generated by Stochastic Gradient Descent (SGD).

The matrix-valued learning rate used to generate the SGD iterates.

A deterministic integer defined in Lemma 2

A resumed sequence of iterates of SGD defined with respect to a given finite stopping time.

The resumed matrix-valued learning rate of SGD defined with respect to a given finite stopping time.

A resumed sequence of random variables,  $\{X_k\}$ , defined with respect to a given finite stopping time.A resumed  $\sigma$ -algebra defined with respect to a stopped  $\sigma$ -algebra.A constant used to denote an upper bound for the sum of  $\lambda_{\max}(M_k)^2$ .

A designation referring to the four properties (P1, P2, P3, P4) required of SGD.

A designation referring to a stopping criterion. The abbreviation refers to “Stopping Criterion.”

The lower bound value for the objective function.

Constants used in specifying assumptions about the optimization problem.

Constants used in specifying assumptions about the tail probabilities.

A designation referring to assumptions about the domain of the objective function. The abbreviation stands for “Iterate Space.”

A designation referring to assumptions about the deterministic objective function. The abbreviation stands for “Deterministic Objective.”

A designation referring to assumptions about the stochastic objective function. The abbreviation stands for “Stochastic Objective.”

A designation referring to assumptions about the deterministic gradient function. The abbreviation stands for “Deterministic Gradient.”

A designation referring to assumptions about the stochastic gradient function. The abbreviation stands for “Stochastic Gradient.”

A designation referring to assumptions about the variability of the stochastic gradient function. The abbreviation refers to “Noise Model.”

A designation referring to assumptions about the tail behavior of the stochastic gradient function. The abbreviation stands for “Tail.”

A reference to a set of assumptions that define a class of functions which we refer to as Bottou-Curtis-Nocedal functions. See Sect. 4.4.

### 3 Challenges of rigorous stopping criteria

Fundamentally, stopping criteria are posing a zero-one question—that is, either some event does not happen (i.e., zero) or the event does happen (i.e., one). In a deterministic setting, useful stopping criteria are typically asked about observable quantities, such as the value of the objective function or the norm of the gradient function. Additionally, in the deterministic setting, rigorous stopping criteria theory requires ensuring that the stopping criterion will be triggered, which, in turn, relies on proving that the underlying optimization procedure generates iterates that converge to a stationary point or, more broadly, that induce a reduction in the objective.

In the stochastic setting, these useful deterministic stopping criteria cannot be directly applied as the quantities on which they are based (e.g., the objective function or gradient function) cannot be feasibly evaluated. Thus, in the stochastic setting, such deterministic stopping criteria can be adapted by either estimating the deterministic objective or deterministic gradient function (e.g., **SC-1**), or estimating the zero-one question posed by the deterministic stopping criteria (e.g., **SC-2**).

Unfortunately, these estimated stopping criteria inherently require a more complex analysis. We now discuss each of these complexities, and, when appropriate, mention how we contribute to the broader literature through our analysis.

#### 3.1 Triggering the underlying deterministic stopping criterion

The usual, deterministic gradient-based stopping criterion requires evaluating whether the norm of the gradient is less than some threshold. In the deterministic context, such a stopping criterion is guaranteed to be triggered if we can prove that the gradient function evaluated at the iterates is converging to zero. In the stochastic setting, any estimated stopping criterion for this underlying deterministic criterion requires establishing the same result with the additional high standard of strong convergence (i.e., almost sure convergence, with probability one convergence).

To illustrate why we need strong convergence, consider stopping a sequence of  $\{0, 1\}$ -valued, independent random variables when we observe a sequence of, say, twenty, consecutive zeros with (possibly) random evaluation points,  $\{T_j\} \subset \mathbb{N}$ . Specifically, let  $\{X_k\}$  be independent random variables such that  $\mathbb{P}[X_k = 0] = 1 - k^{-\alpha}$  for some  $\alpha \in (0, 1]$ . Then,  $\{X_k\}$  is converging to zero in probability, yet, by the second Borel-Cantelli lemma,  $\mathbb{P}[X_k = 1, \text{ i.o.}] = 1$ , where *i.o.* means infinitely often. As a result, it is entirely possible that a one appears in  $\{X_{T_j}, X_{T_j+1}, \dots, X_{T_j+19}\}$  for all  $j$ ; that is, because strong convergence does not hold, a poor choice of  $\{T_j\}$  would prevent us from ever observing a consecutive sequence of twenty zeros. Thus, stopping criteria are predicated on establishing strong convergence, rendering convergence in probability insufficient.

When  $F$  is convex, strong convergence can be readily established along with other convergence rate results [2,4,7], but such results have not translated to useful stopping criteria with limited exceptions (e.g., [5,32,35,37,46]). When  $F$  is nonconvex, convergence in probability is already sufficiently challenging, even under the myriad of notions of nonconvexity that we catalog in Sect. 4. In fact, even for the popular Bottou-

Curtis-Nocedal (**BCN**) functions [4]—see Sect. 4.4—convergence in probability has yet to be established. Despite this, there are results under more stringent conditions, notably Corollary 4.12 of [4], Theorem 2(c) of [26], and Theorem 1 of [27]. This last result is a rather important contribution, but the additional condition assumed excludes the simple linear regression problem and, consequently, BCN functions in general. In summary, strong convergence for general, BCN nonconvex functions has yet to be established.

In order to address this gap, our first contribution is to prove the strong convergence (i.e., convergence with probability one) of  $\hat{F}$  evaluated at the iterates of SGD to zero for **BCN** functions (see Corollary 1), which generalizes and strengthens the preceding results. Our proof employs several strategies that either refine current techniques or are atypical in the stochastic optimization literature; our general strategy and how it is distinct from [26] and [27] is discussed at the beginning of Sect. 5.2. Owing to this contribution, we can guarantee that the underlying deterministic stopping criterion will be triggered.

### 3.2 Triggering the estimated stopping criterion

Unfortunately, even if the underlying deterministic stopping criterion is triggered, the estimated stopping criterion is not necessarily subject to such a guarantee. As an example, consider **SC-2** with  $\epsilon \in (0, 1)$  (see Sect. 6) applied to  $f(\theta) = \theta X$ , where  $X$  is a Rademacher random variable. In this example, the stopping criterion would never be triggered. However, there are situations in which the stopping criterion would offer substantial computational and information theoretic-benefits (see **T1** in Sect. 6).

As this example demonstrates, whether an estimated stopping criterion is triggered is *not* subsumed by whether the underlying deterministic stopping criterion is triggered, and the choice of estimated stopping criterion will have implications for performance even for the same underlying deterministic stopping criterion. Thus, studying whether the estimated stopping criterion is triggered adds a layer of complexity that does not exist for the deterministic context. To our knowledge, such an investigation has not been undertaken in the literature, and will be a central issue that we will explore.

### 3.3 False negative and false positive control

Owing to the fact that estimated stopping criteria depend on the values of random quantities, an estimated stopping criterion may

1. fail to stop the algorithm even when the deterministic gradient is sufficiently small (i.e., a false negative);<sup>5</sup> or
2. incorrectly stop the algorithm even though the norm of the iterate's deterministic gradient is too large (i.e., a false positive).

<sup>5</sup> The notions of “sufficiently small” or “too large” are dependent on the application, just as they are in deterministic optimization.

Ideally, a stopping criterion is specified in order to control the probability of both types of falsehoods from occurring. Unfortunately, to our knowledge, there are no stopping criteria that have been specified that have rigorously addressed either of these concerns. Therefore, in this work, we will derive false negative probability controls. While our techniques can be applied to derive coarse false positive probability controls, we will not state formal results in this direction: in our view, a false positive is often deleterious as it would terminate the algorithm too early for any utility to be derived from the terminal iterate; thus, coarse false positive controls are rather useless for most practical situations (and, for those where it is allowable, they can be readily derived using our techniques). For meaningful false positive control, we would need assumptions on the cumulative distribution function of  $\|\dot{f}(\theta, X)\|_2$ , which is an atypical consideration in the stochastic optimization literature. Thus, we will leave it to future work.

## 4 Bottou-Curtis-Nocedal and other conditions on nonconvexity

For the specification of nonconvex functions, there are many conditions in the literature and many of them are intimately related. Our goal here is to define these conditions as they appear in the literature, and demonstrate how they are related. At the end, we state the Bottou-Curtis-Nocedal (BCN) conditions, which was popularized in [4] and which is the main set of conditions under consideration in this work.

Before enumerating these distinct conditions, we note that there is a core assumption that  $F$  is bounded from below, both  $\dot{F}$  and  $\dot{f}$  exist, and that  $\dot{f}$  is unbiased. While it is possible to allow for bias in  $\dot{f}$  (see Sect. 4 of [4]), this is only useful when we are using a scalar step size; if we consider a matrix-based step size, as we will do here, then the bias in  $\dot{f}$  cannot be allowed without extra assumptions. Hence, we will take  $\dot{f}$  to be unbiased and we simply point out that the biased case with scalar step sizes is an easier analysis than what we will present below.

### 4.1 Control on the iterate space

There are three common ways in which the space of iterates of  $\theta$  is controlled. The first condition assumes

**(IS-1)** The argument,  $\theta$ , is restricted to a compact, convex subset of  $\mathbb{R}^p$ , denoted by  $\Theta$  [25].

Condition **IS-1** is quite convenient in the analysis of SGD methods because it prevents the iterates from diverging out of a compact set—if an iterate is pushed outside of the convex, compact set then it is simply projected back into the set. As a result of this projection operation, the analysis of SGD methods where the iterates are

restricted to a convex, compact is a special case of when the situation where  $\Theta$  is not required to be bounded [2].

A less restrictive approach assumes

**(IS-2)** The argument,  $\theta$ , is restricted to a closed, convex subset of  $\mathbb{R}^p$ , denoted by  $\Theta$  [17,47].

Condition **IS-2** certainly contains Condition **IS-1**, and thus is more general. However, just as for Condition **IS-1**, Condition **IS-2** is imposed by projecting the iterates into  $\Theta$ . Again, as a result of this projection operation, the analysis of SGD methods where the iterates are restricted to a convex, closed set is a special case of the analysis in which  $\Theta$  is the whole space  $\mathbb{R}^p$ .

Thus, the final way in which the space of iterates of  $\theta$  is controlled is to allow  $\theta$  to be unrestricted. For uniformity, we define this as

**(IS-3)** The argument,  $\theta$ , is unrestricted. That is  $\theta$  can take any value in  $\mathbb{R}^p$ .

## 4.2 Control on the objective

There are several common ways in which the stochastic objective (SO) function,  $f$ , is directly controlled. The first one assumes

**(SO-1)** There exists a  $C > 0$  such that  

$$\mathbb{P}[\forall \theta, |f(\theta, X)| < C] = 1.$$

Condition **SO-1** is assumed in [19,54]. In fact, both of these works require more stringent controls on  $f$  and its derivatives. However, in [19], these added controls allow for the approximation of SGD methods by stochastic differential equations, which are then leveraged to characterize the escape times of an SGD method from a saddle point. It is worth noting that such a characterization supplies a more complete analysis than the more recent results of [13,21], which only supply conditions for finding an  $\epsilon$ -approximate stationary point with a nearly positive definite Hessian, but do not guarantee that such a point is stable (i.e., that the SGD method will not escape this point as well).

A natural relaxation of Condition **SO-1** is to assume that  $F(\theta)$  is bounded from above and below. While bounding  $F(\theta)$  from below is necessary when  $\theta$  is not

restricted to a compact subset, there were no works in our review that directly assumed that  $F(\theta)$  is bounded from above. The closest condition in this regard come from [15], which assumes

**(SO-2)** There exists a  $C > 0$  such that for all  $\theta \in \mathbb{R}^p$ ,

$$\mathbb{E}[|f(\theta, X)|^2] \leq C.$$

Note, by Jensen's inequality, Condition **SO-2** implies

$$|F(\theta)|^2 = |\mathbb{E}[f(\theta, X)]|^2 \leq \mathbb{E}[|f(\theta, X)|^2] \leq C; \quad (2)$$

that is, Condition **SO-2** implies  $F$  is bounded.

An alternative restriction to bounding the stochastic objective is to impose Lipschitz continuity, which is often more appropriate for a number of data-driven optimization problems. One specific form of using such continuity is assumed in [34,39], and specified by

**(SO-3)** There exists a  $C > 0$  such that

$$\mathbb{P}[\forall \theta_1, \theta_2, |f(\theta_1, X) - f(\theta_2, X)| \leq C \|\theta_1 - \theta_2\|_2] = 1.$$

Unfortunately, Condition **SO-3** excludes our litmus problem (i.e., the standard linear regression problem). Interestingly, Condition **SO-3** is assumed locally in [15] along with twice differentiability and constraints on the expected Hessian of  $F$  in order to establish local rates of convergence for nonconvex functions.

The natural relaxations of Conditions **SO-1** and **SO-3** to the deterministic objective (DO) function,  $F$ , are

**(DO-1)** There exists a  $C > 0$  such that  $|F(\theta)| \leq C$  for all  $\theta$ .

and

**(DO-2)** There exists a  $C > 0$  such that for all  $\theta_1$  and  $\theta_2$ ,

$$|F(\theta_1) - F(\theta_2)| \leq C \|\theta_1 - \theta_2\|_2.$$

By Jensen's inequality, it follows that Condition **SO-1** implies Condition **DO-1**, and that Condition **SO-3** implies Condition **DO-2**. Note, Condition **DO-1** is used in [19] to establish an SDE approximation to the SGD iterates, and establish rather fine-resolutions results about the escape times of the SGD from saddle points. Condition **DO-2** was used in the early nonconvex results of [17,39], and was used recently by [34] to establish that SGD finds  $\epsilon$ -approximate second-order stationary points with high probability. Importantly, Condition **DO-2** is essential in Theorem 1 of [27] (i.e., the aforementioned convergence with probability one result): it is used three times in the proof of the result and done in such a way that it cannot be relaxed.

### 4.3 Control on the gradient

There are many more conditions that are placed on the stochastic gradients (SG),  $\dot{f}$ , than the stochastic objective (SO),  $f$ , owing to the centrality of  $\dot{f}$  in SGD methods. The first set of conditions will be analogous to those conditions on  $f$ . The most restrictive condition, integral to the results in [3,19,29,39], is

$$\text{(SG-1)} \quad \text{There exists a } C > 0 \text{ such that} \\ \mathbb{P}[\forall \theta, \|\dot{f}(\theta, X)\|_2 \leq C] = 1.$$

A more applicable relative of this condition and an analogue of Condition **SO-3**, used in [1,3,13,20,21,28,29,34,40,47,54], is

$$\text{(SG-2)} \quad \text{There exists a } C > 0 \text{ such that} \\ \mathbb{P}[\forall \theta_1, \theta_2, \|\dot{f}(\theta_1, X) - \dot{f}(\theta_2, X)\|_2 \leq C \|\theta_1 - \theta_2\|_2] = 1.$$

Condition **SG-2** is also used locally in the results of [15]. More importantly, Condition **SG-2** is integral to the previously mentioned result, Theorem 2(c) of [26], which established convergence in probability. While we will not directly compare Condition **SG-2** to the BCN conditions, we will compare weaker conditions implied by Condition **SG-2** to the BCN conditions. In order to do so, we will first need to discuss common conditions placed on the noise model (NM).

The first common restriction on the noise model (NM), used in [6,27], is specified by

$$\text{(NM-1)} \quad \text{There exists a } C \geq 0 \text{ such that} \\ \mathbb{P}[\forall \theta, \|\dot{f}(\theta, X) - \dot{F}(\theta)\|_2 \leq C] = 1.$$

The second and less restrictive noise model condition, used in [13, 17, 21, 25, 28, 34, 52], is specified by

$$\text{(NM-2)} \quad \text{There exists a } C \geq 0 \text{ such that, } \forall \theta, \quad \mathbb{E} \left[ \left\| \dot{f}(\theta, X) - \dot{F}(\theta) \right\|_2^2 \right] \leq C.$$

As we will see subsequently, we have the following noise model condition, which is motivated by Condition **SG-2**.

$$\text{(NM-3)} \quad \text{There exists a } C_1, C_2 \geq 0 \text{ such that, } \forall \theta, \quad \mathbb{E} \left[ \left\| \dot{f}(\theta, X) - \dot{F}(\theta) \right\|_2^2 \right] \leq C_1 + C_2 F(\theta).$$

The following lemma relates Condition **SG-2** to Condition **NM-3**. Note, in the following lemma, the parameter,  $L$ , is taken to be 0 in [26].

**Lemma 1** *Suppose there exists  $L \in \mathbb{R}$  such that  $\mathbb{P}[\forall \theta, f(\theta, X) \geq L] = 1$  and suppose Condition **SG-2** holds. Then, for all  $\theta_1, \theta_2$ ,*

$$\left\| \dot{F}(\theta_1) - \dot{F}(\theta_2) \right\|_2 \leq C \left\| \theta_1 - \theta_2 \right\|_2, \quad (3)$$

and for some  $C_1, C_2 \geq 0$ , **NM-3** holds.

**Proof** The first result follows by an application of Jensen's inequality. For the second result, for any  $\theta$ , let

$$\tilde{\theta} = \theta - \frac{1}{C} \dot{f}(\theta, X). \quad (4)$$

By Taylor's theorem and Condition **SG-2**,

$$f(\tilde{\theta}, X) \leq f(\theta, X) + \dot{f}(\theta, X)'(\tilde{\theta} - \theta) + \frac{C}{2} \left\| \tilde{\theta} - \theta \right\|_2^2 \quad (5)$$

$$\leq f(\theta, X) - \frac{1}{C} \left\| \dot{f}(\theta, X) \right\|_2^2 + \frac{1}{2C} \left\| \dot{f}(\theta, X) \right\|_2^2 \quad (6)$$

$$\leq f(\theta, X) - \frac{1}{C} \left\| \dot{f}(\theta, X) \right\|_2^2. \quad (7)$$

Therefore, since  $L \leq f(\tilde{\theta}, X)$  with probability one,

$$\left\| \dot{f}(\theta, X) \right\|_2^2 \leq C[f(\theta, X) - f(\tilde{\theta}, X)] \leq C[f(\theta, X) - L]. \quad (8)$$

Moreover, this inequality with Jensen's inequality implies that

$$\|\dot{F}(\theta)\|_2^2 \leq \mathbb{E} \left[ \|\dot{f}(\theta, X)\|_2^2 \right] \leq C[F(\theta) - L] \quad (9)$$

Therefore,

$$\mathbb{E} \left[ \|\dot{f}(\theta, X) - \dot{F}(\theta)\|_2^2 \right] \leq 2\mathbb{E} \left[ \|\dot{f}(\theta, X)\|_2^2 \right] + 2 \|\dot{F}(\theta)\|_2^2 \leq 4C(F(\theta) - L). \quad (10)$$

If  $L$  is positive, then we can choose  $C_1 = 0$  and  $C_2 = 4C$ . If  $L$  is negative, then we can choose  $C_1 = -4CL$  and  $C_2 = 4C$ .  $\square$

While a variant of Condition **NM-3** was recently considered (see the discussion after **NM-4**), a direct application of **NM-3** and the Polyak-Łojasiewicz (PL) condition [22] provide a straightforward proof of convergence (in probability) of the objective function evaluated at the iterates to the optimal value.

The final noise model condition is assumed in the BCN conditions [4], and is specified by

(NM-4) There exists a  $C_1, C_2 \geq 0$  such that,  $\forall \theta \geq 0$ ,

$$\mathbb{E} \left[ \|\dot{f}(\theta, X) - \dot{F}(\theta)\|_2^2 \right] \leq C_1 + C_2 \|\dot{F}(\theta)\|_2^2.$$

In general, if  $\dot{F}$  is Lipschitz continuous, then **NM-4** implies **NM-3** by a simple analogue of Lemma 1. Given that  $\dot{F}$  is often taken to be Lipschitz continuous in the stochastic optimization literature, **NM-3** is effectively more general. Some recent works combining the two have also appeared under the names of expected smoothness [23] and [18], and show weak forms of convergence under this broader assumption. In fact, to strengthen their convergence results, the PL condition is employed [18,23], which renders **NM-3** and **NM-4** equivalent. Specifically, if  $\exists \mu > 0$  and an optimal objective function value  $F^*$  such that

$$\|\dot{F}(\theta)\|_2^2 \geq \mu[F(\theta) - F^*], \quad (11)$$

then

$$\begin{aligned} C_1 + C_2 F(\theta) &= (C_1 + C_2 F^*) + C_2 [F(\theta) - F^*] \\ &\leq (C_1 + C_2 F^*) + \frac{C_2}{\mu} \|\dot{F}(\theta)\|_2^2. \end{aligned} \quad (12)$$

We note that the PL condition, if assumed globally as in [18,23], is rather a strong assumption to make as it all stationary points to have value  $F^*$ , which is usually taken to be the global minimum.

The final two conditions are the less restrictive implications that follow from Conditions **SG-1** and **SG-2** about the deterministic gradient (DG),  $\dot{F}$ . The first condition, found in [6,16,20,48], is

**(DG-1)** There exists a  $C > 0$  such that, for all  $\theta$ ,

$$\|\dot{F}(\theta)\|_2 \leq C.$$

The second condition, found in [25,27,52,55], is

**(DG-2)** There exists a  $C > 0$  such that for all  $\theta_1, \theta_2$ ,

$$\|\dot{F}(\theta_1) - \dot{F}(\theta_2)\|_2 \leq C \|\theta_1 - \theta_2\|_2.$$

#### 4.4 Bottou-Curtis-Noceodal functions

The Bottou-Curtis-Noceodal (BCN) conditions, as popularized in [4], is a very general set of conditions as it—arguably—takes the least restrictive conditions from those discussed above. That is, the BCN conditions are defined as follows.

The BCN conditions are

**(BCN)**

1.  $F$  is bounded from below.
2.  $\dot{F}$  is Lipschitz continuous (DG-2).
3. For all  $\theta \in \mathbb{R}^p$ ,  $\dot{F}(\theta) = \mathbb{E}[\dot{f}(\theta, X)]$ .
4. The stochastic gradients satisfy NM-4.

**Definition 1** A function that satisfies the BCN conditions is referred to as a BCN function.

**Remark 1** While we have not explicitly done so, we can allow for biased stochastic gradients if we use scalar step sizes.

## 5 Strong global convergence

Here, we establish one of our key results: the deterministic gradient evaluated at the iterates generated by stochastic gradient descent (SGD) converges to zero with probability one for Bottou-Curtis-Noceodal (BCN) nonconvex functions. In order to do this, we will first need to specify the precise nature of the SGD iterates that we will consider. Then, we will establish global convergence of SGD with probability one for BCN functions, which includes a broad class of convex and nonconvex functions.

## 5.1 Stochastic gradient descent method

Here, we will consider SGD methods with matrix-valued learning rates. In doing so, we will require the particular **BCN** conditions described earlier; however, we emphasize that if the step size is scalar and the stochastic gradients are allowed to be bias, then the iterates can be analyzed with the same arguments below and with less difficulty.

Let  $\{X_k : k \in \mathbb{N}\}$  be independent and identically distributed random variables that have the same distribution as  $X$ . Let  $\beta_0$  be either a fixed or random quantity in  $\mathbb{R}^p$ . Let  $\mathcal{F}_0 = \sigma(\beta_0)$  and  $\mathcal{F}_k = \sigma(\beta_0, X_1, \dots, X_k)$  denote the corresponding elements of the usual filtration. Define the Stochastic Gradient Descent iterates  $\{\beta_k : k \in \mathbb{N}\}$  recursively by

$$\beta_{k+1} = \beta_k - M_k \dot{f}(\beta_k, X_{k+1}), \quad (13)$$

where  $\{M_k\}$  are matrices that satisfy:

**(P1)** The matrices  $\{M_k\}$  are symmetric and positive definite.

There exists an  $S > 0$  such that

$$(P2) \quad \sum_{k=0}^{\infty} \lambda_{\max}(M_k)^2 < S,$$

where  $\lambda_{\max}(\cdot)$  denotes the largest eigenvalue of the given matrix.

The sum,

$$(P3) \quad \sum_{k=0}^{\infty} \lambda_{\min}(M_k),$$

diverges, where  $\lambda_{\min}(\cdot)$  is the smallest eigenvalue of the given matrix.

The condition numbers of  $\{M_k\}$  with respect to operator

$$(P4) \quad 2 - \text{norm}, \{\kappa(M_k)\}, \text{ satisfy} \\ \lim_{k \rightarrow \infty} \lambda_{\max}(M_k) \kappa(M_k) = 0.$$

There are several remarks worth making at this point. First, if  $M_k$  are scalar multiples of the identity, we would recognize properties **P1** to **P3** as the Robbins-Monro conditions [41], and **P4** would be implied by **P2**. Second, **P4** ensures that the condition number of  $M_k$  cannot grow too rapidly such that subsequences of  $\{M_k\}$  may hinder convergence (see Lemma 2). Importantly, **P4** can be relaxed so that the limit is a positive number that would depend on certain constants governing the deterministic gradient behavior and the noise model (see (23)).

Finally, we make a brief note about an important property of SGD that follows from the independence of  $\{X_k\}$ , which we will formalize and make use of later. Specifically,  $\{\beta_k\}$  enjoy an analogue of the strong Markov property; that is, for any finite, stopping time  $\tau$ , the iterates  $\{\beta_{\tau+k}\}$  are independent of  $\mathcal{F}_\tau$  given  $\beta_\tau$  and  $M_\tau$ . Moreover, since properties **P1** to **P4** still hold for  $\{M_{\tau+k}\}$  (see Lemma 3), any property that holds for  $\{\beta_{\tau+k}\}$  (given  $\beta_\tau$  and  $M_\tau$ ) must also hold for  $\{\beta_k\}$ .

## 5.2 Strong global convergence

With the formulation of SGD in hand and the nature of the BCN functions specified, we are now ready to prove the strong global convergence of the iterates; that is, we will prove that

$$\mathbb{P} \left[ \lim_{k \rightarrow \infty} \|\dot{F}(\beta_k)\|_2 = 0 \right] = 1. \quad (14)$$

The proof of this result will proceed by two steps. First, we will establish that, for any  $\delta > 0$ ,

$$\mathbb{P} [\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}] = 1, \quad (15)$$

where *i.o.* means infinitely often (see Theorem 1). Unfortunately, this alone will not imply convergence. Therefore, we will use this result to prove that, for any  $\delta > 0$ ,

$$\mathbb{P} [\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}] = 0, \quad (16)$$

which implies that the deterministic gradient of the iterates converge to zero with probability one (see Theorem 2).

Before detailing the results, we would like to point out the general strategy that we use, and how it is similar to, or distinct from, previous efforts. First, we use coupling to relate the iterate sequence and a related sequence, for which we then establish an analogue of the strong Markov property mentioned previously. With these two pieces, a refinement of Zoutendijk's global convergence strategy [56], and an induction argument, we are able to prove (15). Then, we leverage (15), the inclusion-exclusion principle, and a conditional version of the Borel-Cantelli lemma to conclude (16).

As mentioned, our approach adapts Zoutendijk's global convergence strategy [56], which has been done previously in the stochastic optimization literature [39, 40]. However, our approach refines this argument by establishing an analogue of the strong Markov property through coupling, which we have not observed in any of the stochastic optimization literature. This allows us to state a much stronger result than what has previously been established.

Another point of departure is that our approach avoids restating the behavior of the iterates as a martingale, which is the primary strategy when  $F$  is convex (e.g., see [2]). Our approach also avoids restating any evaluation of the iterates with respect to the objective or the gradient as a martingale, which is the strategy that is used in Theorem 2(c) of [26], and which seems to require much stronger conditions than the general BCN conditions. In fact, the use of martingales only appears in the proof of the conditional Borel-Cantelli lemma, which we do not derive but rather cite from source material.

Finally, our proof does not require explicitly bounding the behavior of the differences between the norms of the deterministic gradients of sequential iterates by the steps,  $\{M_k\}$ , to prove the deterministic gradients of the iterates converge to zero with probability one. This type of argumentation is essential to the proof of [27]<sup>6</sup> and in [26].<sup>7</sup>

**Remark 2** We also note that all of the inequalities and equalities below hold with probability one, even if this is not explicitly stated.

**Remark 3** We also point out that the probabilities and expectations below should be conditional on  $\mathcal{F}_0$ . However, to avoid the additional cumbersome notation, we will not explicitly state this.

### 5.2.1 Strong markov property and coupling

Our first task will be to set the stage to the analogue of strong Markov property that will be relevant in analyzing the SGD method. Let  $\tau$  be a finite stopping time with respect to  $\{\mathcal{F}_k\}$ ; that is,  $\mathbb{P}[\tau < \infty] = 1$ . Moreover, as all of our arguments will be asymptotic in this section, we will assume also that  $\mathbb{P}[\tau \geq K] = 1$  where  $K \in \mathbb{N}$  is defined in the following lemma to satisfy the following property, which will be used later.

**Lemma 2** Suppose  $\{M_k\}$  satisfy **P1** to **P4**, and recall that  $C$  is the Lipschitz constant in **DG-2** and  $C_2$  is the scaling constant in **NM-4**. There exists a  $K \in \mathbb{N}$  such that, for all  $k \geq K$  and for any eigenvalue,  $\lambda$  of  $M_k$ ,

$$\frac{1}{2}\lambda_{\min}(M_k) \leq \lambda - \frac{C}{2}\lambda^2 - \frac{CC_2}{2}\lambda_{\max}(M_k)^2. \quad (17)$$

**Proof** We begin by showing that right hand side of (17) is minimized when  $\lambda = \lambda_{\min}(M_k)$  for all  $k$  sufficiently large and then establish (17) for this choice of  $\lambda$ . We first establish

$$\lambda_{\max}(M_k) + \lambda_{\min}(M_k) \leq \frac{2}{C}. \quad (18)$$

By **P2**,  $\lambda_{\max}(M_k) \rightarrow 0$ . Therefore, there exists a  $k_0 \in \mathbb{N}$  such that for all  $k \geq k_0$ ,  $\lambda_{\max}(M_k) \leq 1/C$ . Therefore, (18) follows for all  $k \geq k_0$ . Moreover, for every  $k \geq k_0$ ,

<sup>6</sup> This bound is required to apply Lemma 1 of [27]. See the second display equations of Page 6 in [27].

<sup>7</sup> This type of bound is established in (15) and the subsequent display equation in [26]. The argument then essentially reestablishes Lemma 1 of [27].

we see that (18) holds for any other eigenvalue of  $M_k$  in place of  $\lambda_{\max}(M_k)$ . Letting  $\lambda$  denote any eigenvalue of  $M_k$  and noting that  $\lambda - \lambda_{\min}(M_k) \geq 0$ ,

$$(\lambda - \lambda_{\min}(M_k))(\lambda + \lambda_{\min}(M_k)) \leq \frac{2}{C}(\lambda - \lambda_{\min}(M_k)), \quad (19)$$

which can be rearranged to conclude that, for any  $k \geq k_0$  and for any eigenvalue  $\lambda$  of  $M_k$ ,

$$\lambda_{\min}(M_k) - \frac{C}{2}\lambda_{\min}(M_k)^2 \leq \lambda - \frac{C}{2}\lambda^2. \quad (20)$$

It follows that the right hand side of (17) is minimized by  $\lambda_{\min}(M_k)$ .

We now establish (17) with  $\lambda = \lambda_{\min}(M_k)$  by contradiction. Suppose that there exists a subsequence  $\{k_j\}$  for which (17) does not hold; that is, by rearranging (17),

$$C\lambda_{\min}(M_{k_j})^2 + CC_2\lambda_{\max}(M_{k_j})^2 > \lambda_{\min}(M_{k_j}). \quad (21)$$

Taking an upper bound on the left hand side,

$$2C(C_2 + 1)\lambda_{\max}(M_{k_j})^2 > \lambda_{\min}(M_{k_j}). \quad (22)$$

Hence,

$$\lambda_{\max}(M_{k_j})\kappa(M_{k_j}) > \frac{1}{2C(C_2 + 1)}, \quad (23)$$

which contradicts **P4** for  $j$  sufficiently large. Hence, there is a  $K \geq k_0$  such that, for all  $k \geq K$ , (17) holds.  $\square$

Now, using  $\tau$ , we will define a sequence of iterates  $\{\psi_k\}$  that we will eventually couple with  $\{\beta_k\}$ . To define  $\{\psi_k\}$ , let

1.  $Z_k := X_{\tau+1+k}$  for all  $k \in \mathbb{N}$ .
2.  $\psi_0 := \beta_{\tau+1}$ .
3. For all  $k \in \{0\} \cup \mathbb{N}$ ,  $P_k := M_{\tau+1+k}$  and

$$\psi_{k+1} := \psi_k - P_k \dot{f}(\psi_k, Z_{k+1}) \mathbf{1} \left[ \|\dot{F}(\psi_k)\|_2^2 > \delta \right], \quad (24)$$

where  $\mathbf{1}[\cdot]$  is the indicator function for the given event.

There are two properties of these quantities that are worth noting:  $\{Z_k\}$  are independent and identically distributed;  $\{P_k\}$  inherit the properties of  $\{M_k\}$  with probability one. The former is verified by Theorem 4.1.3 of [10], which states that  $\{Z_k\}$  are mutually independent and independent of  $\mathcal{F}_{\tau+1}$ , and have the same distribution as  $X_1$ . The latter is verified by the following lemma.

**Lemma 3** *With probability one,  $\{P_k\}$  satisfy **P1** to **P4**. Moreover, with probability one, for all  $k \geq 0$  and any eigenvalue  $\lambda$  of  $P_k$ ,*

$$\frac{1}{2}\lambda_{\min}(P_k) \leq \lambda - \frac{C}{2}\lambda^2 - \frac{CC_2}{2}\lambda_{\max}(P_k)^2. \quad (25)$$

**Proof** The result follows from a standard divide and conquer argument. For **P1**,

$$\mathbb{P}[P_k = P'_k, P_k \succ 0] = \mathbb{P}[M'_{\tau+1+k} = M_{\tau+1+k}, M_{\tau+1+k} \succ 0] \quad (26)$$

$$= \sum_{j=0}^{\infty} \mathbb{P}[M'_{j+1+k} = M_{j+1+k}, M_{j+1+k} \succ 0 \mid \tau = j] \mathbb{P}[\tau = j] \quad (27)$$

$$= \sum_{j=0}^{\infty} \underbrace{\mathbb{P}[M'_{j+1+k} = M_{j+1+k}, M_{j+1+k} \succ 0]}_{=1 \text{ by P1 on } \{M_k\}} \mathbb{P}[\tau = j] \quad (28)$$

$$= \sum_{j=0}^{\infty} \mathbb{P}[\tau = j] \quad (29)$$

$$= \mathbb{P}[\tau < \infty] = 1. \quad (30)$$

Similarly, for **P2**,

$$\mathbb{P}\left[\sum_{k=0}^{\infty} \lambda_{\max}(P_k)^2 < S\right] \quad (31)$$

$$= \sum_{j=0}^{\infty} \underbrace{\mathbb{P}\left[\sum_{k=0}^{\infty} \lambda_{\max}(M_{j+1+k})^2 < S \mid \tau = j\right]}_{=1 \text{ by P2 on } \{M_k\}} \mathbb{P}[\tau = j] \quad (32)$$

$$= \mathbb{P}[\tau < \infty] = 1. \quad (33)$$

The analogous arguments will show that the remaining conclusions of the lemma hold.  $\square$

From these properties, we see that if (24) did not have the indicator term, then the only difference between  $\{\psi_k\}$  and  $\{\beta_k\}$  is the initialization—the fact that  $\{P_k\}$  and  $\{M_k\}$  are distinct is of little importance for our purposes as long as **P1** to **P4** are satisfied. Thus, we see that  $\{\beta_k\}$  exhibit an analogue of the strong Markov property.

Now, to couple these two iterate sequences, let  $\mathcal{G}_0 = \mathcal{F}_{\tau+1}$  and  $\mathcal{G}_k = \sigma(\mathcal{F}_{\tau+1}, Z_1, \dots, Z_k)$ , and, for  $\delta > 0$ , define  $\tau_\delta$  to be a stopping time with respect to  $\{\mathcal{G}_k\}$  such that

$$\tau_\delta = \min\left\{k \geq 0 : \|\dot{F}(\psi_k)\|_2^2 \leq \delta\right\}. \quad (34)$$

Then, on the event  $\{k \geq \tau_\delta\}$ ,  $\psi_{k+1} = \psi_k$ . Moreover, on the event  $\{k < \tau_\delta\}$ ,

$$\begin{aligned} \psi_{k+1} &= \psi_k - P_k \dot{f}(\psi_k, Z_{k+1}) \\ &= \beta_{\tau+1+k} - M_{\tau+1+k} \dot{f}(\beta_{\tau+1+k}, X_{\tau+2+k}) = \beta_{\tau+2+k} \end{aligned} \quad (35)$$

follows by induction. Therefore, for all  $0 \leq k \leq \tau_\delta$ ,  $\psi_k = \beta_{\tau+1+k}$ ; that is, the sequences are coupled in this interval. Owing to this coupling, we see that  $\tau_\delta + 1$  is the number of iterates for  $\beta_k$  to be within a “ $\delta$ -region of zero gradient” after iterate  $\tau$ .

We now apply Zoutendijk's global convergence approach to conclude that  $\tau_\delta$  is finite with probability one.

### 5.2.2 Zoutendijk's global convergence approach

We now apply Zoutendijk's global convergence approach [56] to  $\{\psi_k\}$  to conclude that  $\mathbb{P}[\tau_\delta < \infty] = 1$ . First, by the fundamental theorem of calculus and **DG-2** (recall, with constant  $C$ ),

$$F(\psi_{k+1}) \leq F(\psi_k) + \dot{F}(\psi_k)'(\psi_{k+1} - \psi_k) + \frac{C}{2} \|\psi_{k+1} - \psi_k\|_2^2 \quad (36)$$

$$\begin{aligned} &= F(\psi_k) - \dot{F}(\psi_k)'P_k \dot{f}(\psi_k, Z_{k+1}) \mathbf{1} \left[ \|\dot{F}(\psi_k)\|_2^2 > \delta \right] \\ &\quad + \frac{C}{2} \|P_k \dot{f}(\psi_k, Z_{k+1})\|_2^2 \mathbf{1} \left[ \|\dot{F}(\psi_k)\|_2^2 > \delta \right]. \end{aligned} \quad (37)$$

We now take the conditional expectation of the resulting inequality with respect to  $\mathcal{G}_k$ . Note, since  $\psi_k$ ,  $P_k$  are measurable with respect to  $\mathcal{G}_k$  and  $Z_{k+1}$  is independent of  $\mathcal{G}_k$ , then  $\mathbb{E}[F(\psi_k)|\mathcal{G}_k] = F(\psi_k)$  and

$$\mathbb{E} \left[ \dot{F}(\psi_k)'P_k \dot{f}(\psi_k, Z_{k+1}) \mid \mathcal{G}_k \right] = \dot{F}(\psi_k)'P_k \dot{F}(\psi_k). \quad (38)$$

For the third term in (37), we will need to make use of **NM-4** (with parameters  $C_1, C_2 \geq 0$ ).

$$\begin{aligned} &\mathbb{E} \left[ \|P_k [\dot{f}(\psi_k, Z_{k+1}) - \dot{F}(\psi_k) + \dot{F}(\psi_k)]\|_2^2 \mid \mathcal{G}_k \right] \\ &= \mathbb{E} \left[ \|P_k [\dot{f}(\psi_k, Z_{k+1}) - \dot{F}(\psi_k)]\|_2^2 \mid \mathcal{G}_k \right] + \dot{F}(\psi_k)'P_k^2 \dot{F}(\psi_k) \end{aligned} \quad (39)$$

$$\leq C_1 \lambda_{\max}(P_k)^2 + C_2 \lambda_{\max}(P_k)^2 \|\dot{F}(\psi_k)\|_2^2 + \dot{F}(\psi_k)'P_k^2 \dot{F}(\psi_k) \quad (40)$$

Putting the calculation for these three terms together in (37), we conclude

$$\begin{aligned} \mathbb{E} \left[ F(\psi_{k+1}) \mid \mathcal{G}_k \right] &\leq F(\psi_k) + \frac{CC_1}{2} \lambda_{\max}(P_k)^2 - \mathbf{1} \left[ \|\dot{F}(\psi_k)\|_2^2 > \delta \right] \\ &\quad \times \left[ \dot{F}(\psi_k)'P_k \dot{F}(\psi_k) - \frac{C}{2} \dot{F}(\psi_k)'P_k^2 \dot{F}(\psi_k) - \frac{CC_2}{2} \lambda_{\max}(P_k)^2 \|\dot{F}(\psi_k)\|_2^2 \right]. \end{aligned} \quad (41)$$

We can find an upper bound for (41) by finding a lower bound for the third term in the right hand side of the inequality. In particular, we will lower bound

$$\min_{v \in \mathbb{S}^{p-1}} v' P_k v - \frac{C}{2} v' P_k^2 v - \frac{CC_2}{2} \lambda_{\max}(P_k)^2 \|v\|_2^2, \quad (42)$$

where  $\mathbb{S}^{p-1}$  is the unit sphere in  $\mathbb{R}^p$ .<sup>8</sup>

<sup>8</sup> We could drop the last term in the optimization problem as it is a constant.

Using the Schur decomposition of  $P_k$ ,<sup>9</sup> we can transform (42) into the equivalent problem

$$\min_{v \in \mathbb{S}^{p-1}} \sum_{i=1}^p \left[ \lambda_i - \frac{C}{2} \lambda_i^2 - \frac{CC_2}{2} \lambda_1^2 \right] v_i^2, \quad (43)$$

where  $\lambda_{\max}(P_k) = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p = \lambda_{\min}(P_k)$ ; and  $v_i$  are the components of  $v$ . Applying Lemma 3, the solution to (43) is lower bounded by  $\lambda_{\min}(P_k)/2$ .

Plugging this lower bound into (41), we have that

$$\begin{aligned} \mathbb{E} [F(\psi_{k+1}) | \mathcal{G}_k] &\leq F(\psi_k) + \frac{CC_1}{2} \lambda_{\max}(P_k)^2 \\ &\quad - \frac{1}{2} \lambda_{\min}(P_k) \|\dot{F}(\psi_k)\|_2^2 \mathbf{1} [\|\dot{F}(\psi_k)\|_2^2 > \delta]. \end{aligned} \quad (44)$$

Rearranging and applying the condition in the indicator,

$$\begin{aligned} \frac{\delta}{2} \lambda_{\min}(P_k) \mathbf{1} [\|\dot{F}(\psi_k)\|_2^2 > \delta] \\ &\leq F(\psi_k) - \mathbb{E} [F(\psi_{k+1}) | \mathcal{G}_k] + \frac{CC_1}{2} \lambda_{\max}(P_k)^2. \end{aligned} \quad (45)$$

Now, recall that  $P_k$  are measurable with respect to  $\mathcal{F}_{\tau+1}$  and recall that  $\mathcal{F}_{\tau+1} \subset \mathcal{G}_k$  for all  $k$ . Therefore,

$$\begin{aligned} \frac{\delta}{2} \lambda_{\min}(P_k) \mathbb{P} [\|\dot{F}(\psi_k)\|_2^2 > \delta | \mathcal{F}_{\tau+1}] \\ &\leq \mathbb{E} [F(\psi_k) - F(\psi_{k+1}) | \mathcal{F}_{\tau+1}] + \frac{CC_1}{2} \lambda_{\max}(P_k)^2. \end{aligned} \quad (46)$$

Moreover, by (a) summing both sides from  $k = 0$  to  $n \in \mathbb{N}$ , (b) recalling that  $F(\psi_0)$  is finite with probability one given  $\mathcal{F}_{\tau+1}$ , and (c) applying **P2** from Lemma 3, we conclude

$$\begin{aligned} \frac{\delta}{2} \sum_{k=0}^n \lambda_{\min}(P_k) \mathbb{P} [\|\dot{F}(\psi_k)\|_2^2 > \delta | \mathcal{F}_{\tau+1}] \\ &\leq F(\psi_0) - \mathbb{E} [F(\psi_{n+1}) | \mathcal{F}_{\tau+1}] + \frac{SCC_1}{2}. \end{aligned} \quad (47)$$

Recall that, we have assumed that  $F(\theta)$  is bounded from below by some constant  $F_{l.b.}$  as a core assumption. Using this, we see that

$$\frac{\delta}{2} \sum_{k=0}^n \lambda_{\min}(P_k) \mathbb{P} [\|\dot{F}(\psi_k)\|_2^2 > \delta | \mathcal{F}_{\tau+1}] \leq F(\psi_0) - F_{l.b.} + \frac{SCC_1}{2}. \quad (48)$$

<sup>9</sup> Since  $P_k$  is random, its Schur decomposition is random.

Moreover, note that, for all  $k \geq 0$ ,

$$\mathbb{P} \left[ \bigcap_{j=0}^{\infty} \left\{ \|\dot{F}(\psi_j)\|_2^2 > \delta \right\} \middle| \mathcal{F}_{\tau+1} \right] \leq \mathbb{P} \left[ \|\dot{F}(\psi_k)\|_2^2 > \delta \middle| \mathcal{F}_{\tau+1} \right]. \quad (49)$$

Therefore, for arbitrary  $n$ ,

$$\mathbb{P} \left[ \bigcap_{j=0}^{\infty} \left\{ \|\dot{F}(\psi_j)\|_2^2 > \delta \right\} \middle| \mathcal{F}_{\tau+1} \right] \leq \frac{F(\psi_0) - F_{l.b.} + SCC_1/2}{(\delta/2) \sum_{k=0}^n \lambda_{\min}(P_k)} \quad (50)$$

By **P3** from Lemma 3, the right hand side of this inequality can be made arbitrarily small, which implies that the conditional probability on the left hand side is zero.

That is,

$$0 = \mathbb{P} \left[ \bigcap_{j=0}^{\infty} \left\{ \|\dot{F}(\psi_j)\|_2^2 > \delta \right\} \middle| \mathcal{F}_{\tau+1} \right] = \mathbb{P} \left[ \tau_{\delta} = \infty \middle| \mathcal{F}_{\tau+1} \right]. \quad (51)$$

In other words, we have concluded that  $\mathbb{P} \left[ \tau_{\delta} < \infty \middle| \mathcal{F}_{\tau+1} \right] = 1$  with probability one for any finite stopping time  $\tau$ . With this result, our last step is to use induction.

For now, we will say that the iterates are in a  $\delta$ -region of zero gradient if the squared-norm of the gradient of the iterate is no greater than  $\delta$ . If we let  $\tau = -1$ , then  $\tau_{\delta}$  is the first time the iterates enter a  $\delta$ -region of zero gradient. Let  $T_1(\delta) = \tau_{\delta}$  when  $\tau = -1$ . Then, from the above argument, we have shown that  $T_1(\delta)$  is a finite stopping time. Now, define  $T_j(\delta)$  to be the  $j^{th}$  time that the iterates enter a  $\delta$ -region of zero gradient. Suppose that  $T_j(\delta)$  is finite. Then, define  $\tau = T_j(\delta)$ . Then,  $\tau_{\delta}$  for this  $\tau$  is the next time that the iterates enter a  $\delta$ -region of zero gradient. That is,  $T_{j+1}(\delta) = \tau_{\delta} + T_j(\delta)$ . Since we have assumed that  $\tau = T_j(\delta)$  is finite, we conclude that  $\tau_{\delta}$  is finite, which implies that  $T_{j+1}(\delta)$  is finite. Therefore, by induction we have proven the following result.

**Theorem 1** *Let  $F$  be a Bottou-Curtis-Nocedal function (**BCN**) and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1), then, for any  $\delta > 0$ ,*

$$\mathbb{P} \left[ \|\dot{F}(\beta_k)\|_2^2 \leq \delta, \text{ i.o.} \middle| \mathcal{F}_0 \right] = 1 \text{ with probability 1,} \quad (52)$$

where  $\mathcal{F}_0 = \sigma(\beta_0)$ .

### 5.2.3 Inclusion-Exclusion and markov's inequality

Our next step is to prove that

$$\mathbb{P} \left[ \|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.} \middle| \mathcal{F}_0 \right] = 0 \text{ with probability 1.} \quad (53)$$

Again, we will temporarily drop the conditioning on  $\mathcal{F}_0$  for simplicity of the notation. By Theorem 1 and the inclusion-exclusion principle,

$$1 = \mathbb{P}[\{\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}\} \cup \{\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}\}] \quad (54)$$

$$\begin{aligned} &= \mathbb{P}[\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}] + \mathbb{P}[\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}] \\ &\quad - \mathbb{P}[\{\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}\} \cap \{\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}\}]. \end{aligned} \quad (55)$$

Applying Theorem 1 again, we conclude that

$$\begin{aligned} &\mathbb{P}[\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}] \\ &= \mathbb{P}[\{\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}\} \cap \{\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}\}]. \end{aligned} \quad (56)$$

We will now show that the probability of the right hand side is zero. Note, for any outcome

$$\omega \in \{\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}\} \cap \{\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}\}, \quad (57)$$

there must be an infinite subsequence of  $\mathbb{N}$  such that  $\beta_k$  is in a  $\delta^2$ -region of zero gradient and then  $\beta_{k+1}$  exits this  $\delta^2$ -region of zero gradient. Suppose this were not true. Then, there are two cases. In the first case,  $\beta_k$  enters a  $\delta^2$ -region of zero gradient and then never leaves, in which case

$$\omega \notin \{\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}\}. \quad (58)$$

In the second case, we have that  $\beta_k$  exits a  $\delta^2$ -region of zero gradient, and never enters again, which implies

$$\omega \notin \{\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}\}. \quad (59)$$

In both cases, we have a contradiction. Therefore, using just one of the cases, we conclude that

$$\begin{aligned} &\{\|\dot{F}(\beta_k)\|_2 \leq \delta, \text{ i.o.}\} \cap \{\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.}\} \\ &\subset \{\|\dot{F}(\beta_k)\|_2 \leq \delta, \|\dot{F}(\beta_{k+1})\|_2 > \delta, \text{ i.o.}\}. \end{aligned} \quad (60)$$

We can write this latter event as

$$\begin{aligned} &\{\|\dot{F}(\beta_k)\|_2 \leq \delta, \|\dot{F}(\beta_{k+1})\|_2 > \delta, \text{ i.o.}\} \\ &= \{\dot{F}(\beta_{k+1})\|_2 \mathbf{1}[\|\dot{F}(\beta_k)\|_2 \leq \delta] > \delta, \text{ i.o.}\}. \end{aligned} \quad (61)$$

We will now show that this ultimate event occurs with probability zero using Markov's inequality and the Borel-Cantelli lemma. Let  $\epsilon > 0$  and recall that  $C > 0$  is the parameter in **DG-2**. Then,

$$\mathbb{P}[\|\dot{F}(\beta_{k+1})\|_2 \mathbf{1}[\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq \delta + C\epsilon \mid \mathcal{F}_k] \quad (62)$$

$$\begin{aligned} &\leq \mathbb{P}\left[\left(\|\dot{F}(\beta_{k+1}) - \dot{F}(\beta_k)\|_2 + \|\dot{F}(\beta_k)\|_2\right) \mathbf{1}[\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq \delta + C\epsilon \mid \mathcal{F}_k\right] \\ &\geq \delta + C\epsilon \mid \mathcal{F}_k \end{aligned} \quad (63)$$

$$\leq \mathbb{P} [\|\dot{F}(\beta_{k+1}) - \dot{F}(\beta_k)\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] + \delta \geq \delta + C\epsilon | \mathcal{F}_k] \quad (64)$$

$$\leq \mathbb{P} [C \|\beta_{k+1} - \beta_k\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq C\epsilon | \mathcal{F}_k] \quad (65)$$

$$\leq \mathbb{P} [\|M_k \dot{f}(\beta_k, X_{k+1})\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq \epsilon | \mathcal{F}_k] \quad (66)$$

Applying Markov's inequality to the last conditional probability and using **NM-4**,

$$\mathbb{P} [\|\dot{F}(\beta_{k+1})\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq \delta + C\epsilon | \mathcal{F}_k] \quad (67)$$

$$\leq \frac{\lambda_{\max}(M_k)^2}{\epsilon^2} [C_1 + (C_2 + 1) \|\dot{F}(\beta_k)\|_2^2] \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] \quad (68)$$

$$\leq \frac{\lambda_{\max}(M_k)^2}{\epsilon^2} [C_1 + (C_2 + 1)\delta^2] \quad (69)$$

By **P2**, the sum of the right hand side is bounded with probability one. Therefore, by the conditional second Borel-Cantelli lemma (Theorem 5.3.2 of [10]),

$$\mathbb{P} [\|\dot{F}(\beta_{k+1})\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq \delta + C\epsilon, \text{ i.o.}] = 0. \quad (70)$$

Since  $\epsilon > 0$  is arbitrary, then this conclusion holds for each element in the sequence  $\{\epsilon_m\}$  where  $\epsilon_m \downarrow 0$ . Therefore,

$$\mathbb{P} [\|\dot{F}(\beta_{k+1})\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] > \delta, \text{ i.o.}] \quad (71)$$

$$= \mathbb{P} \left[ \bigcup_{m \in \mathbb{N}} \{\|\dot{F}(\beta_{k+1})\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq \delta + C\epsilon_m, \text{ i.o.}\} \right] \quad (72)$$

$$\leq \sum_{m=1}^{\infty} \mathbb{P} [\|\dot{F}(\beta_{k+1})\|_2 \mathbf{1} [\|\dot{F}(\beta_k)\|_2 \leq \delta] \geq \delta + C\epsilon_m, \text{ i.o.}] \quad (73)$$

$$= 0. \quad (74)$$

Therefore, by using this result with (56), (60) and (61), we conclude the following result.

**Theorem 2** *Let  $F$  be a Bottou-Curtis-Nocedal function (**BCN**) and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1), then, for any  $\delta > 0$ ,*

$$\mathbb{P} [\|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.} | \mathcal{F}_0] = 0 \text{ with probability 1,} \quad (75)$$

where  $\mathcal{F}_0 = \sigma(\beta_0)$ .

Theorem 2 supplies the following corollary.

**Corollary 1** *Let  $F$  be a Bottou-Curtis-Nocedal function (**BCN**) and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1), then*

$$\mathbb{P} \left[ \lim_{k \rightarrow \infty} \|\dot{F}(\beta_k)\|_2 = 0 \middle| \mathcal{F}_0 \right] = 1, \text{ with probability 1,} \quad (76)$$

where  $\mathcal{F}_0 = \sigma(\beta_0)$ .

**Proof** For any  $\delta > 0$ , by Theorem 2,

$$1 = \mathbb{P} \left[ \left\{ \|\dot{F}(\beta_k)\|_2 > \delta, \text{ i.o.} \right\}^c \middle| \mathcal{F}_0 \right] \quad (77)$$

$$= \mathbb{P} \left[ \limsup_{k \rightarrow \infty} \|\dot{F}(\beta_k)\|_2 \leq \delta \middle| \mathcal{F}_0 \right]. \quad (78)$$

Since  $\delta > 0$  is arbitrary, the preceding result applies to each element in the sequence  $\{\delta_m\}$  where  $\delta_m \downarrow 0$ . Since the countable intersection of probability one events has probability one,

$$1 = \mathbb{P} \left[ \bigcap_{m \in \mathbb{N}} \left\{ \limsup_{k \rightarrow \infty} \|\dot{F}(\beta_k)\|_2 \leq \delta_m \right\} \middle| \mathcal{F}_0 \right] \quad (79)$$

$$= \mathbb{P} \left[ \limsup_{k \rightarrow \infty} \|\dot{F}(\beta_k)\|_2 = 0 \middle| \mathcal{F}_0 \right], \quad (80)$$

which is the desired result.  $\square$

## 6 Stopping criteria

As discussed in Sect. 3, the key challenges related to rigorous estimated stopping criteria are: establishing that the underlying deterministic stopping criterion will be triggered; establishing that the estimated stopping criterion will be triggered; establishing false positive control; and establishing false negative control.

We have addressed the first challenge in Corollary 1 when the underlying deterministic stopping criterion relies on the gradient function becoming sufficiently small in norm. Therefore, we now address whether the estimated stopping criterion will be triggered, and the issue of false negative control. We again underscore that while the probability of false positives can be controlled using the techniques discussed below, such bounds would be rather coarse and would not be useful in the typical optimization context where a false positive would be particularly insidious.

To achieve this, we organize the remainder of this section as follows. In Sect. 6.1, we state our two estimated stopping criteria (**SC-1** and **SC-2**) that correspond to the aforementioned gradient-based deterministic stopping criterion. In Sect. 6.2, we superficially examine the consequences of the **BCN** conditions for our stopping criteria, which motivate two important specializations that we will consider in addition to the general **BCN** conditions. In Sect. 6.3, we establish that **SC-1** will be triggered with probability one and establish control over its false negative probability. In Sect. 6.4, we establish that **SC-2** will be triggered with probability one and establish control over its false negative probability.

## 6.1 Overview of estimated stopping criteria

Now, we specify the estimated stopping criteria for the underlying deterministic criterion. First, we recall that the underlying deterministic stopping criterion is when the norm of the gradient drops below a given value. Second, before stating these two criteria, we note that a stopping criterion *should not* be evaluated at each iteration in the stochastic setting, but rather ought to be evaluated at select, possibly random, iterations. While we will not specify the selection of these iterations in this work, we will allow for this generality by considering the situation in which the given estimated stopping criterion is evaluated at a strictly increasing sequence of finite stopping times,  $\{T_j : j \in \mathbb{N}\}$ , with respect to  $\{\mathcal{F}_k\}$ .

The first estimated stopping criterion directly attempts to estimate the gradient function using independent samples, while the second estimated stopping criterion attempts to estimate the zero-one question posed by the deterministic stopping criterion using independent samples.

Let  $\epsilon > 0$ . Let  $\{N_j\}$  be  $\mathbb{N}$ -valued random variables such that  $N_j$  is measurable with respect to  $\mathcal{F}_{T_j}$ . Moreover, for each  $j$ , let  $\{Z_{ij} : i = 1, \dots, N_j\}$  be copies of  $X$  that are independent of each other and  $\{\mathcal{F}_k\}$ . Then, the SGD iterates are stopped at iterate  $T_J$ , where

$$J = \min \left\{ j \geq 1 : \frac{1}{N_j} \left\| \sum_{i=1}^{N_j} \dot{f}(\beta_{T_j}, Z_{ij}) \right\|_2 \leq \epsilon \right\}.$$

Let  $\epsilon > 0$ . Let  $\{N_j\}$  be  $\mathbb{N}$ -valued random variables such that  $N_j$  is measurable with respect to  $\mathcal{F}_{T_j}$ . Let  $\bar{\delta} \in (0, 1)$  and let  $\{\delta_j\}$  be  $(0, \bar{\delta})$ -valued random variables such that  $\delta_j$  is measurable with respect to  $\mathcal{F}_{T_j}$ . Moreover, for each  $j$ , let  $\{Z_{ij} : i = 1, \dots, N_j\}$  be copies of  $X$  that are independent of each other and  $\{\mathcal{F}_k\}$ . Then, the SGD iterates are stopped at iterate  $T_J$ , where

$$J = \min \left\{ j \geq 1 : \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \left\| \dot{f}(\beta_{T_j}, Z_{ij}) \right\|_2 \leq \epsilon \right] \geq \delta_j \right\}.$$

**Remark 4** We can readily imagine more efficient, dependent variants of these two stopping criteria that make use of a recent set of historical values of the stochastic gradients used to update  $\{\beta_k\}$ . Indeed, such stopping criteria would be the most desirable and require an even more complex analysis, which we will leave to future work.

We now address one final technical point. Given that our estimated stopping criteria are only evaluated at iterations  $\{T_j : j \in \mathbb{N}\}$ , we need to ensure that the underlying deterministic stopping criterion would be triggered at these iterations in finite time with probability one. Thus, our first step is to specialize Corollary 1 to these iterations as follows.

**Corollary 2** *Let  $F$  be a Bottou-Curtis-Nocedal nonconvex function (Sect. 4.4) and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1). If  $\{T_j\}$  are positive-valued, strictly increasing, finite stopping times with respect to  $\{\mathcal{F}_k\}$ , then*

$$\mathbb{P} \left[ \lim_{j \rightarrow \infty} \|\dot{F}(\beta_{T_j})\|_2 = 0 \mid \mathcal{F}_0 \right] = 1, \text{ with probability 1,} \quad (81)$$

where  $\mathcal{F}_0 = \sigma(\beta_0)$ . Therefore, for any  $\epsilon > 0$ , there exist a stopping time  $J^*$  that is finite with probability one such that  $T_{J^*}$  is finite with probability one, and

$$\mathbb{P} \left[ \|\dot{F}(\beta_{T_{J^*}})\|_2 \leq \epsilon \mid \mathcal{F}_0 \right] = 1, \text{ with probability 1.} \quad (82)$$

**Proof** Note,  $\{T_j\}$  are strictly increasing, which implies  $T_j \geq j - 1$  for all  $j$ . Therefore, for any  $j \in \mathbb{N}$  and  $\epsilon > 0$ ,

$$\bigcap_{m=j}^{\infty} \{\|\dot{F}(\beta_{T_m})\|_2 \leq \epsilon\} \quad (83)$$

$$= \bigcap_{m=j}^{\infty} \left( \bigcup_{k=m-1}^{\infty} \{\|\dot{F}(\beta_k)\|_2 \leq \epsilon\} \cap \{T_m = k\} \right) \quad (84)$$

$$\supset \bigcap_{m=j}^{\infty} \left( \bigcup_{k=m-1}^{\infty} \left\{ \sup_{k \geq j-1} \|\dot{F}(\beta_k)\|_2 \leq \epsilon \right\} \cap \{T_m = k\} \right) \quad (85)$$

$$= \bigcap_{m=j}^{\infty} \left[ \left\{ \sup_{k \geq j-1} \|\dot{F}(\beta_k)\|_2 \leq \epsilon \right\} \cap \left( \bigcup_{k=m-1}^{\infty} \{T_m = k\} \right) \right] \quad (86)$$

$$= \bigcap_{m=j}^{\infty} \left\{ \sup_{k \geq j-1} \|\dot{F}(\beta_k)\|_2 \leq \epsilon \right\} \cap \{T_m < \infty\} \quad (87)$$

$$= \left\{ \sup_{k \geq j-1} \|\dot{F}(\beta_k)\|_2 \leq \epsilon \right\} \cap \left( \bigcap_{m=j}^{\infty} \{T_m < \infty\} \right). \quad (88)$$

Since  $T_m$  are assumed to be finite stopping times with probability one, we conclude

$$\mathbb{P} \left[ \sup_{T_m \geq j-1} \|\dot{F}(\beta_{T_m})\|_2 \leq \epsilon \mid \mathcal{F}_0 \right] \geq \mathbb{P} \left[ \sup_{k \geq j-1} \|\dot{F}(\beta_k)\|_2 \leq \epsilon \mid \mathcal{F}_0 \right]. \quad (89)$$

By Corollary 1, the limit of the left hand side converges to 1 as  $j \rightarrow \infty$ . Since  $\epsilon > 0$  is arbitrary, the first part of the result follows.

For the second part, recall that  $T_j$  are finite with probability one for all  $j$ . Therefore,

$$\begin{aligned}\mathbb{P}[T_{J^*} < \infty | \mathcal{F}_0] &= \sum_{j=0}^{\infty} \mathbb{P}[T_j < \infty | J^* = j, \mathcal{F}_0] \mathbb{P}[J^* = j | \mathcal{F}_0] \\ &= \sum_{j=0}^{\infty} \mathbb{P}[J^* = j | \mathcal{F}_0] = \mathbb{P}[J^* < \infty | \mathcal{F}_0].\end{aligned}\tag{90}$$

Hence, it is enough to prove that  $\mathbb{P}[J^* < \infty | \mathcal{F}_0] = 1$ . Note,

$$\mathbb{P}[J^* < \infty | \mathcal{F}_0] = \mathbb{P}\left[\bigcup_j \left\{ \|\dot{F}(\beta_{T_j})\|_2 \leq \epsilon \right\} \middle| \mathcal{F}_0\right]\tag{91}$$

$$\geq \mathbb{P}\left[\limsup_{j \rightarrow \infty} \|\dot{F}(\beta_{T_j})\|_2 \leq \epsilon \middle| \mathcal{F}_0\right],\tag{92}$$

where we have already shown that the right hand side has probability one.  $\square$

## 6.2 Consequences of the BCN conditions and specializations

For a **BCN** function, we will readily be able to establish that **SC-1** will be triggered with probability one. We will refer to the case where we consider a function satisfying **BCN** as Scenario (a). Unfortunately, we will have some difficulty with this general scenario for **SC-2** if the stochastic gradients are concentrated away from their mean, which happens for the **BCN** function induced by  $f(\theta, X) = \theta X$ , where  $X$  is a Rademacher random variable.

For this reason, we will consider two specializations of **BCN**. The first specialization is the case when  $C_1 = 0$  in the definition of **NM-4**. This specialization corresponds to the important case of over-parametrized models, and occurs when the minimizer of  $F$  is also a minimizer of all  $f$  with probability one [1]. This first specialization will be referred to as Scenario (b).

For the final specialization of **BCN**, which we refer to as Scenario (c), we consider the following additional condition.

**(T1)** Let  $\pi_1, \pi_2 \in (0, 1)$  and  $\pi_3 \geq 1$ . For  $\|\dot{F}(\theta)\|_2 \leq \pi_1$  and for any  $t \geq \pi_3 \|\dot{F}(\theta)\|_2$ ,

$$\mathbb{P}[\|\dot{f}(\theta, X)\|_2 \geq t] \leq \left(\frac{\pi_3 \|\dot{F}(\theta)\|_2}{t}\right)^{\pi_2}$$

The role of **T1** is to ensure that there is a concentration of the norm of the stochastic gradient function near small values when the deterministic gradient is sufficiently small, which would avoid the issues that occur with  $f(\theta, X) = \theta X$ , where  $X$  is a Rademacher random variable. Moreover, we note that the general **BCN** conditions do not imply **T1**, as illustrated by  $f(\theta, X) = \theta X$ . Conversely, **T1** does not imply **BCN**, as a random variable that satisfies **T1** is not even guaranteed to have a mean, leave alone a second moment (i.e., **NM-4**).

To summarize, Scenario (a) refers to objective functions satisfying the general **BCN** conditions; Scenario (b) refers to objective functions satisfying the **BCN** conditions with  $C_1 = 0$  in **NM-4**; and Scenario (c) refers to objective functions satisfying **BCN** and **T1**.

### 6.3 Stopping criterion by gradient estimation

We now establish that **SC-1** will be triggered with probability one and we derive bounds on the probability of a false negative. We begin with some general consequences of any effort to estimate the mean for the three aforementioned scenarios.

**Lemma 4** *Let  $F$  be a **BCN** function (Sect. 4.4). For  $N \in \mathbb{N}$ , let  $\{Z_1, \dots, Z_N\}$  be independent copies of  $X$ . Then for  $\epsilon > 0$ ,*

$$\mathbb{P}\left[\frac{1}{N} \left\| \sum_{i=1}^N \dot{f}(\theta, Z_i) \right\|_2 \leq \epsilon\right] \geq 1 - \frac{C_1 + (C_2 + N) \|\dot{F}(\theta)\|_2^2}{N\epsilon^2}. \quad (93)$$

Moreover, if **T1** holds and  $\|\dot{F}(\theta)\|_2 \leq \pi_1$ , then, for  $\epsilon \geq \pi_3 \|\dot{F}(\theta)\|_2$ ,

$$\mathbb{P}\left[\frac{1}{N} \left\| \sum_{i=1}^N \dot{f}(\theta, Z_i) \right\|_2 \leq \epsilon\right] \geq 1 - N \left(\frac{\pi_3 \|\dot{F}(\theta)\|_2}{\epsilon}\right)^{\pi_2}. \quad (94)$$

**Proof** In both cases, we will find the upper bound for the complement, from which a lower bound for the stated event is readily derived. For the first case, by Markov's inequality and **NM-4**,

$$\mathbb{P}\left[\frac{1}{N} \left\| \sum_{i=1}^N \dot{f}(\theta, Z_i) \right\|_2 > \epsilon\right] \leq \frac{\mathbb{E}\left[\left\| \frac{1}{N} \sum_{i=1}^N \dot{f}(\theta, Z_i) \right\|_2^2\right]}{\epsilon^2} \quad (95)$$

$$\leq \frac{C_1 + (C_2 + N) \|\dot{F}(\theta)\|_2^2}{N\epsilon^2}. \quad (96)$$

For the second case, if  $\epsilon \geq \pi_3 \|\dot{F}(\theta)\|_2$ , then

$$\mathbb{P}\left[\frac{1}{N} \left\| \sum_{i=1}^N \dot{f}(\theta, Z_i) \right\|_2 > \epsilon\right] \leq N\mathbb{P}\left[\|\dot{f}(\theta, Z_1)\|_2 > \epsilon\right], \quad (97)$$

to which we apply **T1**.  $\square$

The first part of Lemma 4 covers Scenarios (a) and (b), while the second part covers Scenario (c). We discuss each of these scenarios individually and in order. With regard to Scenario (a), Lemma 4 shows that the main hurdle in establishing a meaningful bound on the probability of a false negative is in choosing  $N\epsilon^2$  to be sufficiently large to account for the constant  $C_1 > 0$ . Thus, for applications where achieving a small bound on the deterministic gradient is necessary, Lemma 4 implies that  $N$  would need to be larger than the reciprocal squared of this small bound. On the other hand, for applications where achieving a bound on the order of the noise level is acceptable (e.g.,  $\epsilon \approx \sqrt{C_1}$ ), then moderate values of  $N$  are acceptable. The following result formalizes these concepts.

**Proposition 1** *Let  $F$  be a **BCN** function (Sect. 4.4) and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1). Let  $\epsilon > 0$ . If*

$$\mathbb{P} \left[ \liminf_{j \rightarrow \infty} N_j > \frac{C_1}{\epsilon^2} \middle| \mathcal{F}_0 \right] = 1 \text{ with probability one,} \quad (98)$$

then **SC-1** will be triggered in finite time with probability one. Moreover, for any  $\rho \in (0, 1)$  and  $\gamma > 1$ , if

$$N_j > \frac{\gamma C_1 + C_2(\epsilon\rho)^2}{(1 - \rho^2)\epsilon^2}, \quad (99)$$

then

$$\mathbb{P} \left[ \frac{1}{N_j} \left\| \sum_{i=1}^{N_j} \dot{f}(\beta_{T_j}, Z_{ij}) \right\|_2 > \epsilon, \quad \left\| \dot{F}(\beta_{T_j}) \right\|_2 \leq \frac{\rho\epsilon}{\sqrt{\gamma}} \middle| \mathcal{F}_{T_j} \right] \leq \frac{1}{\gamma}. \quad (100)$$

**Proof** When  $C_1 = 0$ , the proof of the result uses analogous reasoning to the case when  $C_1 > 0$ ; therefore, we will take  $C_1 > 0$ . By (98), there exist  $\delta > 1$  and a finite stopping time,  $\bar{J}$ , such that for all  $j \geq \bar{J}$ , with probability one,

$$N_j \geq \frac{\delta C_1}{\epsilon^2}. \quad (101)$$

Moreover, by Corollary 2, for  $\delta' \in (1, \delta)$ , there exists a finite stopping time  $J^*$  such that for all  $j \geq J^*$ , with probability one,

$$\frac{(\delta' - 1)C_1}{C_2 + \delta C_1/\epsilon^2} > \left\| \dot{F}(\beta_{T_j}) \right\|_2^2. \quad (102)$$

Therefore, for all  $j \geq \max\{J^*, \bar{J}\}$ , with probability one,

$$\frac{C_1 + (C_2 + N_j) \left\| \dot{F}(\beta_{T_j}) \right\|_2^2}{N_j \epsilon^2} < \frac{C_1 + (C_2 + \delta C_1/\epsilon^2) \left\| \dot{F}(\beta_{T_j}) \right\|_2^2}{\delta C_1} \quad (103)$$

$$< \frac{C_1 + (\delta' - 1)C_1}{\delta C_1}. \quad (104)$$

Therefore, for all  $j \geq \max\{J^*, \bar{J}\}$ , Lemma 4 implies

$$\mathbb{P} \left[ \frac{1}{N_j} \left\| \sum_{i=1}^{N_j} \dot{f}(\beta_{T_j}, Z_{ij}) \right\|_2 \leq \epsilon \middle| \mathcal{F}_{T_j} \right] \geq 1 - \frac{\delta'}{\delta} > 0. \quad (105)$$

Given that  $\{Z_{ij}\}$  are independent over  $i$  and  $j$  and recall that  $J$  is defined in **SC-1**, the probability that  $J > j$  is controlled by a geometric distribution for  $j \geq \max\{J^*, \bar{J}\}$ . Therefore, we conclude that  $J$  is finite with probability one: that is, **SC-1** is triggered in finite time with probability one.

For the second part of the result, note

$$\begin{aligned} & \mathbb{P} \left[ \frac{1}{N_j} \left\| \sum_{i=1}^{N_j} \dot{f}(\beta_{T_j}, Z_{ij}) \right\|_2 > \epsilon, \left\| \dot{F}(\beta_{T_j}) \right\|_2 \leq \frac{\rho\epsilon}{\sqrt{\gamma}} \middle| \mathcal{F}_{T_j} \right] \\ &= \mathbb{P} \left[ \frac{1}{N_j} \left\| \sum_{i=1}^{N_j} \dot{f}(\beta_{T_j}, Z_{ij}) \right\|_2 > \epsilon \middle| \mathcal{F}_{T_j} \right] \mathbf{1} \left[ \left\| \dot{F}(\beta_{T_j}) \right\|_2 \leq \frac{\rho\epsilon}{\sqrt{\gamma}} \right]. \end{aligned} \quad (106)$$

Applying Lemma 4 with  $\left\| \dot{F}(\beta_{T_j}) \right\|_2 \leq \rho\epsilon/\sqrt{\gamma}$  and the lower bound on  $N_j$  implies the bound on the false negative probability.  $\square$

Clearly, Proposition 1 applies to Scenario (b) which, we recall, is the special case of Scenario (a) in which  $C_1 = 0$ . Here, we see that even with  $N_j = 1$  for all  $j$ , Proposition 1 implies that **SC-1** would be triggered in finite time (though the false positive control would likely be quite poor). Similarly, Proposition 1 allows for controlling the false negative probability at  $1/\gamma$  so long as

$$N_j > \frac{C_2 \rho^2}{1 - \rho^2}, \quad \rho \in (0, 1). \quad (107)$$

With regard to Scenario (c), Lemma 4 demonstrates that having more samples does not help in the case where **T1** holds; in fact, it actually makes the situation worse. This is expected as such a fat-tailed distribution has a population mean of infinity, which, intuitively, would make a sample mean entirely useless. Thus, considering **T1** for **SC-1** does not provide any additional value, and we will not pursue it further.

To summarize, under the general **BCN** conditions and Scenario (b), Proposition 1 shows that **SC-1** will be triggered in finite time for  $N_j$  sufficiently large, and provides control over the false negative probability.

## 6.4 Stopping criterion by majority vote

We now establish that **SC-2** will be triggered with probability one and we derive bounds on the probability of a false negative. Just as we did in Sect. 6.3, we begin with a general result that will guide our specific analysis of **SC-2**.

**Lemma 5** *Let  $F$  be a Bottou-Curtis-Nocedal nonconvex function (Sect. 4.4). For  $N \in \mathbb{N}$ , let  $\{Z_1, \dots, Z_N\}$  be independent copies of  $X$ . Let  $\epsilon > 0$  and  $\delta \in (0, 1)$  and define*

$$\Delta = \delta - \mathbb{P}[\|\dot{f}(\theta, X)\|_2 \leq \epsilon]. \quad (108)$$

When  $\Delta < 0$ ,

$$\mathbb{P}\left[\frac{1}{N} \sum_{i=1}^N \mathbf{1}[\|\dot{f}(\theta, Z_i)\|_2 \leq \epsilon] \geq \delta\right] \geq 1 - \exp(-2N\Delta^2). \quad (109)$$

**Proof** The proof leverages McDiarmid's inequality (see Sect. 3 of [30]). Let the range of  $X$  be denoted by  $\mathcal{X}$ . Let  $z_1, \dots, z_N \in \mathcal{X}$  and define

$$h(z_1, \dots, z_N) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[\|\dot{f}(\theta, z_i)\|_2 \leq \epsilon]. \quad (110)$$

Then, for any  $j \in \{1, \dots, N\}$  and  $z_1, \dots, z_N, z'_j \in \mathcal{X}$ ,

$$|h(z_1, \dots, z_N) - h(z_1, \dots, z'_j, \dots, z_N)| \leq \frac{1}{N}. \quad (111)$$

Since

$$\mathbb{E}[h(Z_1, \dots, Z_N)] = \mathbb{P}[\|\dot{f}(\theta, Z_1)\|_2 \leq \epsilon], \quad (112)$$

McDiarmid's inequality implies

$$\begin{aligned} \mathbb{P}[h(Z_1, \dots, Z_N) < \delta] \\ &= \mathbb{P}[h(Z_1, \dots, Z_N) - \mathbb{E}[h(Z_1, \dots, Z_N)] < \Delta] \end{aligned} \quad (113)$$

$$\leq \exp(-2N\Delta^2). \quad (114)$$

By computing the complement, the result follows.  $\square$

If we compare Lemma 5 to Lemma 4, we get two very different controls on the triggering probability: in the former, we have an exponential control in  $N$ , while in the latter we have sublinear control in  $N$ . As a result, increases in  $N$  offer a dramatically greater benefit for **SC-2** over **SC-1**, and this phenomenon underlies the famous theorem of machine learning referred to as “classification is easier than regression” [9, Theorem 6.5]. However, this benefit comes at a slight loss in generality: Lemma 5 requires that  $\mathbb{P}[\|\dot{f}(\theta, X)\|_2 \leq \epsilon]$  exceeds  $\delta$ , which may not hold for arbitrary values

of  $\epsilon > 0$ , or for general **BCN** functions such as the one induced by  $f(\theta, X) = \theta X$  with  $X$  as a Rademacher variable.

If we can allow for  $\epsilon > 0$  to be large (e.g.,  $\epsilon \gg \sqrt{C_1}$ ), then we can account for general **BCN** functions. However, if we cannot, then we must restrict our analysis to the situation where  $\dot{F}(\theta, X)$  has some concentration near  $\dot{F}(\theta)$ , which is precisely induced in Scenarios (b) and (c). Therefore, we will restrict our attention to Scenarios (b) and (c), and we note that if  $\epsilon$  is allowed to be large than Scenario (a) can be accounted for analogously to Scenario (b). We begin with proving that **SC-2** will be triggered in finite time under Scenarios (b) and (c).

**Proposition 2** *Let  $F$  be a **BCN** function (Sect. 4.4) with either  $C_1 = 0$  in **NM-4** or **T1**; and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1). Let  $\epsilon > 0$ . Then **SC-2** is triggered in finite time with probability one.*

**Proof** Let  $\delta' \in (\bar{\delta}, 1)$ , where  $\bar{\delta}$  is defined in **SC-2**. Now, by Corollary 2, there exists a finite random variable  $J^*$  defined for each scenario such that for all  $j \geq J^*$ :

1. under Scenario (b),  $\|\dot{F}(\beta_{T_j})\|_2^2 < (1 - \delta')\epsilon^2/(C_2 + 1)$ ; or
2. under Scenario (c),  $\|\dot{F}(\beta_{T_j})\|_2^2 < \epsilon(1 - \delta')^{1/\pi_2}/\pi_3$ .

Then, by Lemma 4 (with  $N = 1$ ), under both scenarios, for  $j \geq J^*$ , with probability one,

$$\mathbb{P}\left[\|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \mid \mathcal{F}_{T_j}\right] > \delta'. \quad (115)$$

Therefore,

$$\delta_j - \mathbb{P}\left[\|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \mid \mathcal{F}_{T_j}\right] < \bar{\delta} - \delta' < 0. \quad (116)$$

Apply Lemma 5, we conclude, for  $j \geq J^*$ , with probability one,

$$\mathbb{P}\left[\frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1}\left[\|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon\right] \geq \delta_j \mid \mathcal{F}_{T_j}\right] > 1 - \exp(-2(\delta' - \bar{\delta})^2). \quad (117)$$

Hence, the distribution of  $J$ , as defied in **SC-2**, for  $j \geq J^*$  is controlled by a geometric distribution. We conclude  $J$  is finite with probability one; that is, **SC-2** is triggered in finite time with probability one.  $\square$

Our next step is to establish control over the probability of a false negative. Just as we did in Proposition 1, we will state these controls in terms of some parameters  $\rho$  and  $\gamma$ ; however, unlike in Proposition 1, we will need more control over the values of  $\rho$ . Specifically,

$$\rho \in \left(0, \sqrt{\frac{1 - \bar{\delta}}{C_2 + 1}}\right) \quad \text{or} \quad \rho \in \left(0, (1 - \bar{\delta})^{1/\pi_2}\right), \quad (118)$$

for Scenarios (b) or (c), respectively. The following result establishes control over the probability of a false negative for Scenario (b).

**Proposition 3** Let  $F$  be a **BCN** function (Sect. 4.4) with  $C_1 = 0$  in **NM-4**; and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1). Let  $\epsilon > 0$ . Suppose  $\rho$  satisfies (118) and  $\gamma > 1$ . If

$$N_j > \frac{\log(\gamma)}{2(1 - \bar{\delta} - \rho^2/\gamma)^2}, \quad (119)$$

then

$$\mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j, \quad \|\dot{F}(\beta_{T_j})\|_2 \leq \frac{\rho\epsilon}{\sqrt{\gamma}} \middle| \mathcal{F}_{T_j} \right] \leq \frac{1}{\gamma}. \quad (120)$$

**Proof** When  $\|\dot{F}(\beta_{T_j})\|_2 \leq \rho\epsilon/(\pi_3\gamma)$ , by (118) and Lemma 4,

$$\mathbb{P} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \middle| \mathcal{F}_{T_j} \right] \geq 1 - \rho^2/\gamma > \bar{\delta}. \quad (121)$$

Therefore,

$$\delta_j - \mathbb{P} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \middle| \mathcal{F}_{T_j} \right] \leq \bar{\delta} + \rho^2/\gamma - 1 < 0. \quad (122)$$

Hence, for the given choice of  $N_j$ ,

$$\exp \left\{ -2N_j \left( \delta_j - \mathbb{P} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \middle| \mathcal{F}_{T_j} \right] \right)^2 \right\} \quad (123)$$

$$\leq \exp \left\{ -2N_j \left( \bar{\delta} + \rho^2/\gamma - 1 \right)^2 \right\} \leq \frac{1}{\gamma}. \quad (124)$$

Noting that

$$\begin{aligned} & \mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j, \quad \|\dot{F}(\beta_{T_j})\|_2 \leq \frac{\rho\epsilon}{\sqrt{\gamma}} \middle| \mathcal{F}_{T_j} \right] \\ &= \mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j \middle| \mathcal{F}_{T_j} \right] \mathbf{1} \left[ \|\dot{F}(\beta_{T_j})\|_2 \leq \frac{\rho\epsilon}{\sqrt{\gamma}} \right], \end{aligned} \quad (125)$$

then, applying Lemma 5 when the norm of  $\dot{F}(\beta_{T_j})$  satisfies the hypothesized upper bound and  $N_j$  satisfies the hypothesized lower bound,

$$\mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j \middle| \mathcal{F}_{T_j} \right] \leq \frac{1}{\gamma}. \quad (126)$$

□

We can derive a similar result for Scenario (c).

**Proposition 4** Let  $F$  be a **BCN** function (Sect. 4.4) satisfying **TI**; and let  $\{\beta_k\}$  be the iterates generated by Stochastic Gradient Descent satisfying **P1** to **P4** (Sect. 5.1). Let  $\epsilon > 0$ . Suppose  $\rho$  satisfies (118) and  $\gamma > 1$ . If

$$N_j > \frac{\log(\gamma)}{2(1 - \bar{\delta} - (\rho/\gamma)^{\pi_2})^2}, \quad (127)$$

then

$$\mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j, \quad \|\dot{F}(\beta_{T_j})\|_2 \leq \frac{\rho\epsilon}{\pi_3\gamma} \middle| \mathcal{F}_{T_j} \right] \leq \frac{1}{\gamma}. \quad (128)$$

**Proof** When  $\|\dot{F}(\beta_{T_j})\|_2 \leq \rho\epsilon/(\pi_3\gamma)$ , by (118) and Lemma 4,

$$\mathbb{P} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \middle| \mathcal{F}_{T_j} \right] \geq 1 - (\rho/\gamma)^{\pi_2} > \bar{\delta}. \quad (129)$$

Therefore,

$$\delta_j - \mathbb{P} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \middle| \mathcal{F}_{T_j} \right] \leq \bar{\delta} + (\rho/\gamma)^{\pi_2} - 1 < 0. \quad (130)$$

Hence, for the given choice of  $N_j$ ,

$$\exp \left\{ -2N_j \left( \delta_j - \mathbb{P} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \middle| \mathcal{F}_{T_j} \right] \right)^2 \right\} \quad (131)$$

$$\leq \exp \left\{ -2N_j (\bar{\delta} + (\rho/\gamma)^{\pi_2} - 1)^2 \right\} \leq \frac{1}{\gamma}. \quad (132)$$

Noting that

$$\begin{aligned} & \mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j, \quad \|\dot{F}(\beta_{T_j})\|_2 \leq \frac{\rho\epsilon}{\pi_3\gamma} \middle| \mathcal{F}_{T_j} \right] \\ &= \mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j \middle| \mathcal{F}_{T_j} \right] \mathbf{1} \left[ \|\dot{F}(\beta_{T_j})\|_2 \leq \frac{\rho\epsilon}{\pi_3\gamma} \right], \end{aligned} \quad (133)$$

then, applying Lemma 5 when the norm of  $\dot{F}(\beta_{T_j})$  satisfies the hypothesized upper bound and  $N_j$  satisfies the hypothesized lower bound,

$$\mathbb{P} \left[ \frac{1}{N_j} \sum_{i=1}^{N_j} \mathbf{1} \left[ \|\dot{f}(\beta_{T_j}, Z_{ij})\|_2 \leq \epsilon \right] < \delta_j \middle| \mathcal{F}_{T_j} \right] \leq \frac{1}{\gamma}. \quad (134)$$

□

**Remark 5** The general case—Scenario (a)—can be addressed using these same techniques as long as  $\epsilon$  is allowed to be sufficiently large so that we can provide a nontrivial upper bound on  $\mathbb{P}[\|\dot{f}(\theta, X)\|_2 > \epsilon]$ .

To summarize, we have shown that under Scenarios (b) and (c), **SC-2** will be triggered in finite time with probability one and the false negative probability can be controlled when  $N_j$  is sufficiently large and  $\|\dot{F}(\beta_{T_j})\|_2$  is sufficiently small.

## 7 Numerical experiment

We will discuss, in order, the background of the experiment, the experimental setup, the results of the experiment, and the interpretation of the results.

### 7.1 Background

For the setting of this experiment, we train a neural network to classify greyscale images of ten different types of clothing [50]. We make use of 60,000 examples in this training, and the objective function is defined as the average sparse categorical cross entropy over the labels of the sixty thousand examples and the predicted labels as produced by a neural network that: vectorizes the image; passes it through three dense rectified linear unit layers (each with output dimension of 128); passes the result through a dense linear layer with output dimension 10; and chooses the index with the maximum value to generate a prediction. The resulting parameter that is being optimized is of dimension 151, 306.

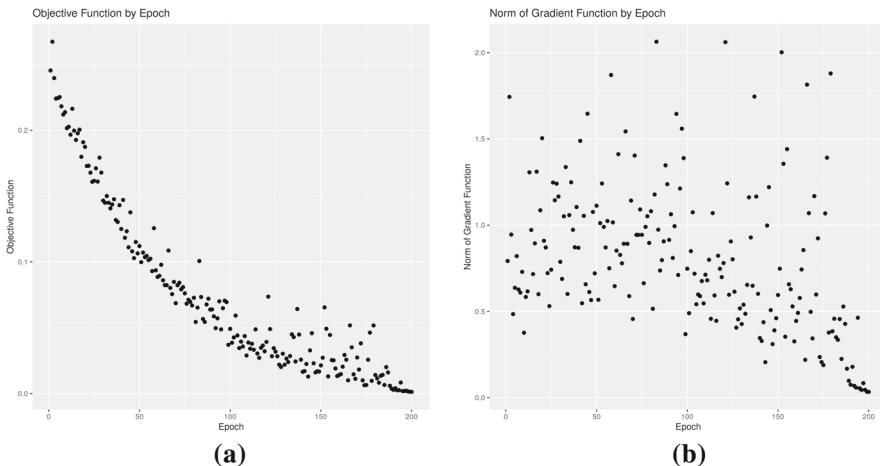
The objective function is optimized using batch stochastic gradient descent with a scalar learning rate

$$M_k = 0.2 \frac{375000}{k + 375000} I, \quad (135)$$

and is terminated after two hundred epochs (i.e., total passes through the data set). Note, at the beginning of the optimization, the learning rate is 0.2 and by the end it is 0.1, where  $k$  represents the total number of batches processed in the 200 epochs. At the end of each epoch, the parameter value is stored, the value of the objective function is calculated, and the norm of the gradient function is calculated. The value of the objective function and gradient function at the end of each epoch is plotted in Fig. 1.

### 7.2 Experimental setup

As an overview, at each of the recorded parameter values, the two stopping criteria are evaluated multiple times and the triggering rate (i.e., the number of times the stopping criterion is triggered divided by the total number of attempts) is recorded. Specifically, at each of the recorded parameter values—which occurs at the end of each epoch—,



**Fig. 1** The value of the objective function (left) and norm of the gradient function (right) at the end of each of the two hundred epochs

**SC-1** and **SC-2** are evaluated one hundred times each using independent samples from the underlying data with the following specifications:

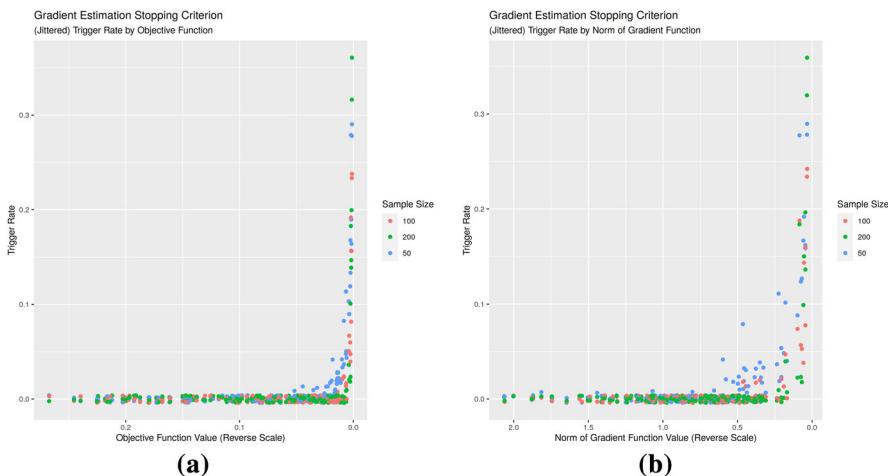
1. the threshold  $\epsilon = 0.07$  (i.e., about twice the norm of the gradient at the final epoch);
2. the sample size  $N_j$  takes value in  $\{50, 100, 200\}$ ; and,
3. for **SC-2**,  $\delta_j$  are kept constant and take value in  $\{0.5, 0.6, 0.7, 0.8\}$ .

In total, **SC-1** is evaluated on 600 combinations (200 epochs, three distinct values of  $N_j$ ), and **SC-2** is evaluated on 2,400 combinations (200 epochs, three distinct values of  $N_j$ , four distinct values for  $\delta_j$ ). Furthermore, each combination is evaluated independently 100 times. At the end of each 100 independent evaluations the triggering rate is calculated for each combination, which is the number of evaluated stopping criteria that are triggered divided by the total number of evaluated stopping criteria (i.e., 100) for the given combination.

### 7.3 Results

Figure 2 shows the estimated triggering rate of **SC-1** as a function of the objective and as a function of the norm of the gradient. Note, Fig. 2's x-axes are reversed so as to roughly align with the increase in epochs; moreover, the points are jittered to allow for the zero values to be observable. Figure 2 indicates that, as the objective function decreases (i.e., as we are presumably approach a solution) and as the norm of the gradient function decreases (i.e., as we are presumably approaching a stationary point), **SC-1** is more likely to be triggered. Moreover, Fig. 2 shows that the triggering rate is most variable for the case where  $N_j = 50$  and least variable for the case of  $N_j = 200$ .

Figure 3 shows the estimated triggering rate of **SC-2** as a function of the objective and as a function of the norm of the gradient. Note, Fig. 2's x-axes are reversed so



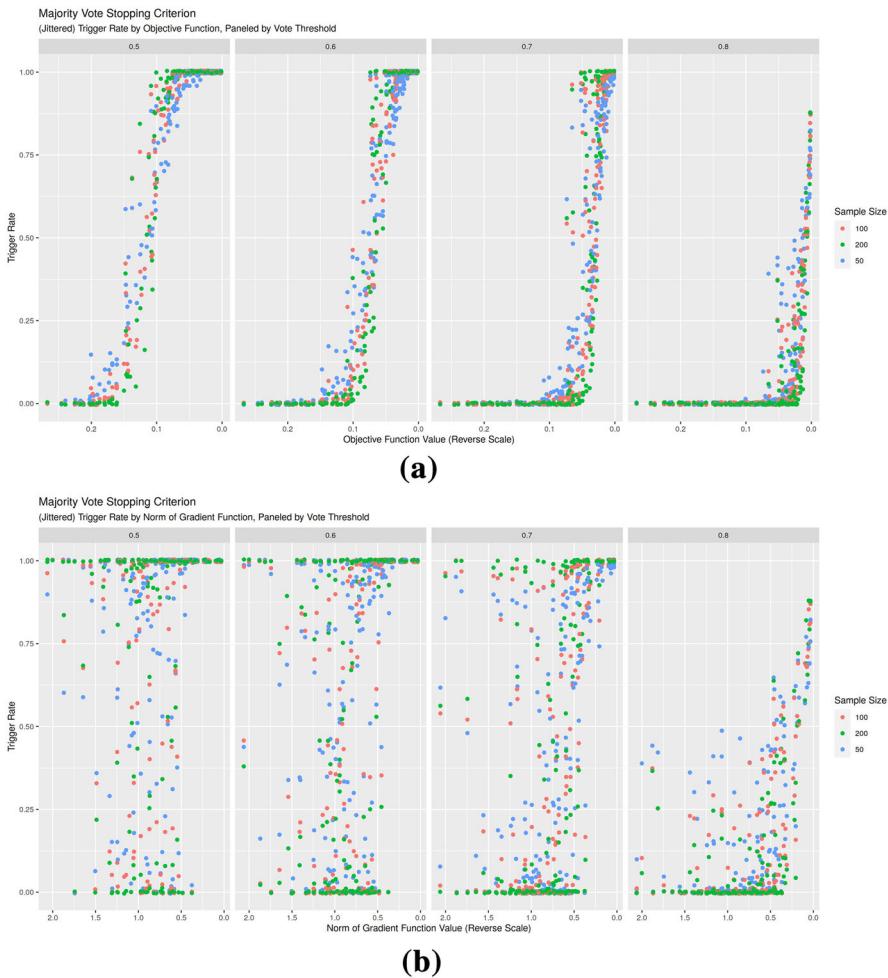
**Fig. 2** The triggering rate of **SC-1** by the value of the objective function (left) and norm of the gradient function (right). Note, the x-axis scales are reversed so as to roughly align with the increase in epochs

as to roughly align with the increase in epochs; the points are jittered to allow for the zero values to be observable; and the panels each correspond to the different values of  $\delta_j$  (i.e., the vote threshold). Figure 3 indicates that, as the objective function decreases in value, **SC-2** is more likely to be triggered—for the smaller vote thresholds, **SC-2** is likely to be triggered at less optimal values, while, for larger vote thresholds, **SC-2** is likely to be triggered close to the optimal value. On the other hand, Fig. 3 shows that, for all of the vote threshold values except for  $\delta_j = 0.8$ , **SC-2** has no discernible pattern for when it is triggered. When  $\delta_j = 0.8$ , Fig. 3 shows that **SC-2** is likely to be triggered only for small values of the gradient with the most variability occurring when  $N_j = 50$  and least variability occurring when  $N_j = 200$ .

## 7.4 Discussion

As we would expect from the theory presented, **SC-1** correlates well with the decaying norm of the gradient and the probability of a false negative decays with increasing  $N_j$ . Moreover, **SC-1** is rather conservative: the probability of a false negative, even when  $N_j = 200$  and the norm of the gradient is at its smallest, is 0.65. Of course,  $N_j = 200$  is less than 0.4% of the total set of examples, and using more examples will better control the false negative probability, as we would expect.

Whereas **SC-1** behaved as we would expect, **SC-2** is surprisingly more unpredictable. In particular, when the vote threshold was less than 0.8, **SC-2** was triggered for the larger values of the norm of the gradient. Unfortunately, for this particular problem, the distribution of the stochastic gradients are highly concentrated near zero with a handful that are far from zero. Therefore, when the vote threshold is low, this concentration of the stochastic gradients is likely to produce a false positive. However, when the vote threshold is increased, we see much more desirable behavior of **SC-2**: it is being triggered only for smaller values of the norm of the gradient, and the false



**Fig. 3** The triggering rate of **SC-2** by the value of the objective function (left) and norm of the gradient function (right). Note, the x-axis scales are reversed so as to roughly align with the increase in epochs

negative rate is closer to about 0.2. Again, as expected from the theory, **SC-2** is much more efficient with the observations in comparison to **SC-1**, but we must be careful and ensure that the value of  $\delta_j$  is sufficiently large to prevent false positive signals.

## 8 Conclusion

In this work, our goal was to lay a rigorous foundation for stopping criteria for stochastic gradient descent (SGD) as applied to Bottou-Curtis-Nocedal (BCN) functions, which includes a broad class of convex and nonconvex functions. We started by developing a strong global convergence result for SGD on BCN functions, which generalizes

previous results on the convergence of SGD on nonconvex functions. Then, we presented two stopping criteria and rigorously analyzed them.

This work has raised several questions that we enumerate below, and which we hope to address in future work.

1. Given strong global convergence, what is the local rate of convergence to a stationary point? Are we even guaranteed to find a stationary point? Is this stationary point guaranteed to be a minimum? As evidenced by the many works cited, these issues are of great importance and can be answered more completely now that we have established strong global convergence.
2. Can something be said about dependent versions of the estimated stopping criteria? In some sense, dependent stopping criteria are ideal as they are the least wasteful stopping criteria. However, to develop such results, we need a maximal inequality over the norms of the iterates.
3. For all of the stopping criteria, what are reasonable choices of  $\{T_j\}$  and  $\{N_j\}$ ?
4. What are reasonable conditions to place on the lower tail probabilities (analogous to **T1**) and what are their implications for controlling the false positive probability of the stopping criteria studied in this work?
5. Finally, is there a context in which **NM-3** is more appropriate than **NM-4**, and can the preceding results be developed in this context as well?

**Acknowledgements** We thank the reviewers for their detailed feedback, which has greatly improved the quality of this work.

## References

1. Bassily, R., Belkin, M., Ma, S.: On exponential convergence of sgd in non-convex over-parametrized learning. arXiv preprint [arXiv:1811.02564](https://arxiv.org/abs/1811.02564) (2018)
2. Bertsekas, D.P.: Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. *Optimization for Machine Learning* **2010**(1–38), 3 (2011)
3. Bi, J., Gunn, S.R.: A stochastic gradient method with biased estimation for faster nonconvex optimization. In: Pacific Rim International Conference on Artificial Intelligence, pp. 337–349. Springer (2019)
4. Bottou, L., Curtis, F.E., Nocedal, J.: Optimization methods for large-scale machine learning. *SIAM Rev.* **60**(2), 223–311 (2018)
5. Chee, J., Toulis, P.: Convergence diagnostics for stochastic gradient descent with constant learning rate. In: International Conference on Artificial Intelligence and Statistics, pp. 1476–1485 (2018)
6. Chen, X., Liu, S., Sun, R., Hong, M.: On the convergence of a class of adam-type algorithms for non-convex optimization. arXiv preprint [arXiv:1808.02941](https://arxiv.org/abs/1808.02941) (2018)
7. Chung, K.L., et al.: On a stochastic approximation method. *Ann. Math. Stat.* **25**(3), 463–483 (1954)
8. Curtis, F.E., Scheinberg, K.: Adaptive stochastic optimization. arXiv preprint [arXiv:2001.06699](https://arxiv.org/abs/2001.06699) (2020)
9. Devroye, L., Györfi, L., Lugosi, G.: A probabilistic theory of pattern recognition, vol. 31. Springer Science & Business Media, Berlin (2013)
10. Durrett, R.: Probability: theory and examples, 4th edn. Cambridge University Press, Cambridge (2010)
11. Ermoliev, Y.: Stochastic quasigradient methods and their application to system optimization. *Stoch.: Int. J. Probab. Stoch. Process.* **9**(1–2), 1–36 (1983)
12. Fabian, V.: Stochastic approximation of minima with improved asymptotic speed. *The Annals of Mathematical Statistics* pp. 191–200 (1967)
13. Fang, C., Lin, Z., Zhang, T.: Sharp analysis for nonconvex sgd escaping from saddle points. arXiv preprint [arXiv:1902.00247](https://arxiv.org/abs/1902.00247) (2019)

14. Farrell, R.: Bounded length confidence intervals for the zero of a regression function. *The Annals of Mathematical Statistics* pp. 237–247 (1962)
15. Fehrman, B., Gess, B., Jentzen, A.: Convergence rates for the stochastic gradient descent method for non-convex objective functions. arXiv preprint [arXiv:1904.01517](https://arxiv.org/abs/1904.01517) (2019)
16. Foster, D.J., Sekhari, A., Sridharan, K.: Uniform convergence of gradients for non-convex learning and optimization. In: *Advances in Neural Information Processing Systems*, pp. 8745–8756 (2018)
17. Ghadimi, S., Lan, G., Zhang, H.: Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Math. Program.* **155**(1–2), 267–305 (2016)
18. Gower, R.M., Sebbouh, O., Loizou, N.: Sgd for structured nonconvex functions: Learning rates, mini-batching and interpolation. arXiv preprint [arXiv:2006.10311](https://arxiv.org/abs/2006.10311) (2020)
19. Hu, W., Li, C.J., Li, L., Liu, J.G.: On the diffusion approximation of nonconvex stochastic gradient descent. arXiv preprint [arXiv:1705.07562](https://arxiv.org/abs/1705.07562) (2017)
20. Huang, F., Chen, S.: Linear convergence of accelerated stochastic gradient descent for nonconvex nonsmooth optimization. arXiv preprint [arXiv:1704.07953](https://arxiv.org/abs/1704.07953) (2017)
21. Jin, C., Netrapalli, P., Ge, R., Kakade, S.M., Jordan, M.I.: On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. arXiv preprint [arXiv:1902.04811](https://arxiv.org/abs/1902.04811) (2019)
22. Karimi, H., Nutini, J., Schmidt, M.: Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer (2016)
23. Khaled, A., Richtárik, P.: Better theory for sgd in the nonconvex world. arXiv preprint [arXiv:2002.03329](https://arxiv.org/abs/2002.03329) (2020)
24. Kiefer, J., Wolfowitz, J., et al.: Stochastic estimation of the maximum of a regression function. *Ann. Math. Stat.* **23**(3), 462–466 (1952)
25. Lei, J., Shanbhag, U.V.: A randomized block proximal variable sample-size stochastic gradient method for composite nonconvex stochastic optimization. arXiv preprint [arXiv:1808.02543](https://arxiv.org/abs/1808.02543) (2018)
26. Lei, Y., Hu, T., Li, G., Tang, K.: Stochastic gradient descent for nonconvex learning without bounded gradient assumptions. *IEEE Transactions on Neural Networks and Learning Systems* (2019)
27. Li, X., Orabona, F.: On the convergence of stochastic gradient descent with adaptive stepsizes. arXiv preprint [arXiv:1805.08114](https://arxiv.org/abs/1805.08114) (2018)
28. Li, Z., Li, J.: A simple proximal stochastic gradient method for nonsmooth nonconvex optimization. In: *Advances in Neural Information Processing Systems*, pp. 5564–5574 (2018)
29. Ma, Y., Klabjan, D.: Convergence analysis of batch normalization for deep neural nets. CoRR, [arXiv:1705.08011](https://arxiv.org/abs/1705.08011) 2 (2017)
30. McDiarmid, C.: Concentration. In: *Probabilistic methods for algorithmic discrete mathematics*, pp. 195–248. Springer (1998)
31. Mertikopoulos, P., Hallak, N., Kavis, A., Cevher, V.: On the almost sure convergence of stochastic gradient descent in non-convex problems. arXiv preprint [arXiv:2006.11144](https://arxiv.org/abs/2006.11144) (2020)
32. Mirozahmedov, F., Uryasev, S.: Adaptive stepsize regulation for stochastic optimization algorithm. *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki* **23**(6), 1314–1325 (1983)
33. Nemirovski, A., Juditsky, A., Lan, G., Shapiro, A.: Robust stochastic approximation approach to stochastic programming. *SIAM J. Optim.* **19**(4), 1574–1609 (2009)
34. Park, S., Jung, S.H., Pardalos, P.M.: Combining stochastic adaptive cubic regularization with negative curvature for nonconvex optimization. arXiv preprint [arXiv:1906.11417](https://arxiv.org/abs/1906.11417) (2019)
35. Patel, V.: Kalman-based stochastic gradient method with stop condition and insensitivity to conditioning. *SIAM J. Optim.* **26**(4), 2620–2648 (2016)
36. Patel, V.: The impact of local geometry and batch size on the convergence and divergence of stochastic gradient descent. arXiv preprint [arXiv:1709.04718](https://arxiv.org/abs/1709.04718) 9 (2017)
37. Pflug, G.C.: Stepsize rules, stopping times and their implementation in stochastic quasi-gradient algorithms. numerical techniques for stochastic optimization pp. 353–372 (1988)
38. Prechelt, L.: Early stopping-but when? In: *Neural Networks: Tricks of the trade*, pp. 55–69. Springer (1998)
39. Reddi, S.J., Hefny, A., Sra, S., Poczos, B., Smola, A.: Stochastic variance reduction for nonconvex optimization. In: *International conference on machine learning*, pp. 314–323 (2016)
40. Reddi, S.J., Sra, S., Poczos, B., Smola, A.J.: Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. In: *Advances in Neural Information Processing Systems*, pp. 1145–1153 (2016)

41. Robbins, H., Monro, S.: A stochastic approximation method. *The annals of mathematical statistics* pp. 400–407 (1951)
42. Roy, V.: Convergence diagnostics for markov chain monte carlo. *Annual Rev. Stat. Its Appl.* **7**, 387–412 (2020)
43. Sielken, R.L.: Stopping times for stochastic approximation procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **26**(1), 67–75 (1973)
44. Stroup, D.F., Braun, H.I.: On a new stopping rule for stochastic approximation. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **60**(4), 535–554 (1982)
45. Van der Vaart, A.W.: *Asymptotic statistics*, vol. 3. Cambridge University Press, Cambridge (2000)
46. Wada, T., Itani, T., Fujisaki, Y.: A stopping rule for linear stochastic approximation. In: 49th IEEE Conference on Decision and Control (CDC), pp. 4171–4176. IEEE (2010)
47. Wang, X., Wang, X., Yuan, Y.X.: Stochastic proximal quasi-newton methods for non-convex composite optimization. *Optim. Methods Softw.* **34**(5), 922–948 (2019)
48. Ward, R., Wu, X., Bottou, L.: Adagrad stepsizes: Sharp convergence over nonconvex landscapes, from any initialization. *arXiv preprint arXiv:1806.01811* (2018)
49. Wu, L.: Mixed effects models for complex data. Chapman and Hall/CRC, Florida (2009)
50. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms (2017)
51. Yin, G.: A stopping rule for the robbins-monro method. *J. Optim. Theory Appl.* **67**(1), 151–173 (1990)
52. Yu, H., Jin, R.: On the computation and communication complexity of parallel sgd with dynamic batch sizes for stochastic non-convex optimization. *arXiv preprint arXiv:1905.04346* (2019)
53. Zhang, P., Lang, H., Liu, Q., Xiao, L.: Statistical adaptive stochastic gradient methods. *arXiv preprint arXiv:2002.10597* (2020)
54. Zhou, Y.: Nonconvex optimization in machine learning: Convergence, landscape, and generalization. Ph.D. thesis, The Ohio State University (2018)
55. Zou, F., Shen, L., Jie, Z., Zhang, W., Liu, W.: A sufficient condition for convergences of adam and rmsprop. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 11127–11135 (2019)
56. Zoutendijk, G.: Nonlinear programming, computational methods. *Integer and nonlinear programming* pp. 37–86 (1970)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.