
A Unified Convergence Theorem for Stochastic Optimization Methods

Xiao Li

School of Data Science (SDS)
Shenzhen Institute of Artificial Intelligence
and Robotics for Society (AIRS)
The Chinese University of Hong Kong, Shenzhen
Shenzhen, China
lixiao@cuhk.edu.cn

Andre Milzarek

School of Data Science (SDS)
Shenzhen Research Institute of Big Data (SRIBD)
The Chinese University of Hong Kong, Shenzhen
Shenzhen, China
andremilzarek@cuhk.edu.cn

Abstract

In this work, we provide a fundamental unified convergence theorem used for deriving expected and almost sure convergence results for a series of stochastic optimization methods. Our unified theorem only requires to verify several representative conditions and is not tailored to any specific algorithm. As a direct application, we recover expected and almost sure convergence results of the stochastic gradient method (SGD) and random reshuffling (RR) under more general settings. Moreover, we establish new expected and almost sure convergence results for the stochastic proximal gradient method (prox-SGD) and stochastic model-based methods for nonsmooth nonconvex optimization problems. These applications reveal that our unified theorem provides a plugin-type convergence analysis and strong convergence guarantees for a wide class of stochastic optimization methods.

1 Introduction

Stochastic optimization methods are widely used to solve stochastic optimization problems and empirical risk minimization, serving as one of the foundations of machine learning. Among the many different stochastic methods, the most classic one is the stochastic gradient method (SGD), which dates back to Robbins and Monro [36]. If the problem at hand has a finite-sum structure, then another popular stochastic method is random reshuffling (RR) [20]. When the objective function has a composite form or is weakly convex (nonsmooth and nonconvex), then the stochastic proximal gradient method (prox-SGD) and stochastic model-based algorithms are the most typical approaches [18, 11]. Apart from the mentioned stochastic methods, there are many others like SGD with momentum, Adam, stochastic higher order methods, etc. In this work, our goal is to establish and understand fundamental *convergence* properties of these stochastic optimization methods via a novel unified convergence framework.

Motivations. Suppose we apply SGD to minimize a smooth nonconvex function f . SGD generates a sequence of iterates $\{\mathbf{x}^k\}_{k \geq 0}$, which is a stochastic process due to the randomness of the algorithm and the utilized stochastic oracles. The most commonly seen ‘convergence result’ for SGD is the

expected iteration complexity, which typically takes the form [17]

$$\min_{k=0,\dots,T} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|^2] \leq \mathcal{O}\left(\frac{1}{\sqrt{T+1}}\right) \quad \text{or} \quad \mathbb{E}[\|\nabla f(\mathbf{x}^{\bar{k}})\|^2] \leq \mathcal{O}\left(\frac{1}{\sqrt{T+1}}\right), \quad (1)$$

where T denotes the total number of iterations and \bar{k} is an index sampled uniformly at random from $\{0, \dots, T\}$. Note that we ignored some higher-order convergence terms and constants to ease the presentation. Complexity results are integral to understand core properties and progress of the algorithm during the first T iterations, while the asymptotic convergence behavior plays an equally important role as it characterizes whether an algorithm can eventually approach an exact stationary point or not. We refer to [Appendix H](#) for additional motivational background for studying asymptotic convergence properties of stochastic optimization methods. Here, an *expected convergence result*, associated with the nonconvex minimization problem $\min_{\mathbf{x}} f(\mathbf{x})$, has the form

$$\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|] = 0. \quad (2)$$

Intuitively, it should be possible to derive expected convergence from the expected iteration complexity (1) by letting $T \rightarrow \infty$. However, this is not the case as the ‘min’ operator and the sampled \bar{k} are not well defined or become meaningless when T goes to ∞ .

The above results are stated in expectation and describe the behavior of the algorithm by averaging infinitely many runs. Though this is an important convergence measure, in practical situations the algorithm is often only run once and the last iterate is returned as a solution. This observation motivates and necessitates *almost sure convergence results*, which establish convergence with probability 1 for a single run of the stochastic method:

$$\lim_{k \rightarrow \infty} \|\nabla f(\mathbf{x}^k)\| = 0 \quad \text{almost surely.} \quad (3)$$

Backgrounds. Expected and almost sure convergence results have been extensively studied for convex optimization; see, e.g., [10, 34, 42, 46, 5, 41]. Almost sure convergence of SGD for minimizing a smooth nonconvex function f was provided in the seminal work [3] using very standard assumptions, i.e., Lipschitz continuous ∇f and bounded variance. Under the same conditions, the same almost sure convergence of SGD was established in [33] based on a much simpler argument than that of [3]. A weaker ‘lim inf’-type almost sure convergence result for SGD with AdaGrad step sizes was shown in [26]. Recently, the work [28] derives almost sure convergence of SGD under the assumptions that f and ∇f are Lipschitz continuous, f is coercive, f is not asymptotically flat, and the v -th moment of the stochastic error is bounded with $v \geq 2$. This result relies on stronger assumptions than the base results in [3]. Nonetheless, it allows more aggressive diminishing step sizes if $v > 2$. Apart from standard SGD, almost sure convergence of different respective variants for min-max problems was discussed in [22]. In terms of expected convergence, the work [6] showed $\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|] = 0$ under the additional assumptions that f is twice continuously differentiable and the multiplication of the Hessian and gradient $\nabla^2 f(\mathbf{x})\nabla f(\mathbf{x})$ is Lipschitz continuous.

Though the convergence of SGD is well-understood and a classical topic, asymptotic convergence results of the type (2) and (3) often require a careful and separate analysis for other stochastic optimization methods — especially when the objective function is simultaneously nonsmooth and nonconvex. In fact and as outlined, a direct transition from the more common complexity results (1) to the full convergence results (2) and (3) is often not possible without further investigation.

Main contributions. We provide a fundamental *unified convergence theorem* (see [Theorem 2.1](#)) for deriving both expected and almost sure convergence of stochastic optimization methods. Our theorem is not tailored to any specific algorithm, instead it incorporates several abstract conditions that suit a vast and general class of problem structures and algorithms. The proof of this theorem is elementary.

We then apply our novel theoretical framework to several classical stochastic optimization methods to recover existing and to establish new convergence results. Specifically, we recover expected and almost sure convergence results for SGD and RR. Though these results are largely known in the literature, we derive unified and slightly stronger results under a general ABC condition [24, 23] rather than the standard bounded variance assumption. We also remove the stringent assumption used in [6] to show (2) for SGD. As a core application of our framework, we derive expected and almost sure convergence results for prox-SGD in the nonconvex setting and under the more general ABC condition and for stochastic model-based methods under very standard assumptions. In particular, we show that the iterates $\{\mathbf{x}^k\}_{k \geq 0}$ generated by prox-SGD and other stochastic model-based methods

will approach the set of stationary points almost surely and in an expectation sense. These results are *new* to our knowledge (see also [Subsection 3.5](#) for further discussion).

The above applications illustrate the general plugin-type purpose of our unified convergence analysis framework. Based on the given recursion and certain properties of the algorithmic update, we can derive broad convergence results by utilizing our theorem, which can significantly simplify the convergence analysis of stochastic optimization methods; see [Subsection 2.1](#) for a summary.

2 A unified convergence theorem

Throughout this work, let $(\Omega, \mathcal{F}, \{\mathcal{F}_k\}_{k \geq 0}, \mathbb{P})$ be a filtered probability space and let us assume that the sequence of iterates $\{\mathbf{x}^k\}_{k \geq 0}$ is adapted to the filtration $\{\mathcal{F}_k\}_{k \geq 0}$, i.e., each of the random vectors $\mathbf{x}^k : \Omega \rightarrow \mathbb{R}^n$ is \mathcal{F}_k -measurable.

In this section, we present a unified convergence theorem for the sequence $\{\mathbf{x}^k\}_{k \geq 0}$ based on an abstract convergence measure Φ . To make the abstract convergence theorem more accessible, the readers may momentarily regard Φ and $\{\mu_k\}_{k \geq 0}$ as ∇f and the sequence related to the step sizes, respectively. We then present the main steps for showing the convergence of a stochastic optimization method by following a step-by-step verification of the conditions in our unified convergence theorem.

Theorem 2.1. *Let the mapping $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and the sequences $\{\mathbf{x}^k\}_{k \geq 0} \subseteq \mathbb{R}^n$ and $\{\mu_k\}_{k \geq 0} \subseteq \mathbb{R}_{++}$ be given. Consider the following conditions:*

(P.1) *The function Φ is L_Φ -Lipschitz continuous for some $L_\Phi > 0$, i.e., we have $\|\Phi(\mathbf{x}) - \Phi(\mathbf{y})\| \leq L_\Phi \|\mathbf{x} - \mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.*

(P.2) *There exists a constant $a > 0$ such that $\sum_{k=0}^{\infty} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] < \infty$.*

The following statements are valid:

(i) *Let the conditions (P.1)–(P.2) be satisfied and suppose further that*

(P.3) *There exist constants $A, B, b \geq 0$ and $p_1, p_2, q > 0$ such that*

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^q] \leq A\mu_k^{p_1} + B\mu_k^{p_2} \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^b].$$

(P.4) *The sequence $\{\mu_k\}_{k \geq 0}$ and the parameters a, b, q, p_1, p_2 satisfy*

$$\{\mu_k\}_{k \geq 0} \text{ is bounded, } \sum_{k=0}^{\infty} \mu_k = \infty, \text{ and } a, q \geq 1, a \geq b, p_1, p_2 \geq q.$$

Then, it holds that $\lim_{k \rightarrow \infty} \mathbb{E}[\|\Phi(\mathbf{x}^k)\|] = 0$.

(ii) *Let the properties (P.1)–(P.2) hold and assume further that*

(P.3') *There exist constants $A, b \geq 0, p_1, p_2, q > 0$ and random vectors $\mathbf{A}_k, \mathbf{B}_k : \Omega \rightarrow \mathbb{R}^n$ such that*

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \mu_k^{p_1} \mathbf{A}_k + \mu_k^{p_2} \mathbf{B}_k$$

and for all k , $\mathbf{A}_k, \mathbf{B}_k$ are \mathcal{F}_{k+1} -measurable and we have $\mathbb{E}[\mathbf{A}_k \mid \mathcal{F}_k] = 0$ almost surely, $\mathbb{E}[\|\mathbf{A}_k\|^q] \leq A$, and $\limsup_{k \rightarrow \infty} \|\mathbf{B}_k\|^q / (1 + \|\Phi(\mathbf{x}^k)\|^b) < \infty$ almost surely.

(P.4') *The sequence $\{\mu_k\}_{k \geq 0}$ and the parameters a, b, q, p_1, p_2 satisfy $\mu_k \rightarrow 0$,*

$$\sum_{k=0}^{\infty} \mu_k = \infty, \quad \sum_{k=0}^{\infty} \mu_k^{2p_1} < \infty, \quad \text{and } q \geq 2, \quad qa \geq b, \quad p_1 > \frac{1}{2}, \quad p_2 \geq 1.$$

Then, it holds that $\lim_{k \rightarrow \infty} \|\Phi(\mathbf{x}^k)\| = 0$ almost surely.

The proof of [Theorem 2.1](#) is elementary. We provide the core ideas here and defer its proof to [Appendix A](#). Item (i) is proved by contradiction. An easy first result is $\liminf_{k \rightarrow \infty} \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] = 0$. We proceed and assume that $\{\mathbb{E}[\|\Phi(\mathbf{x}^k)\|]\}_{k \geq 0}$ does not converge to zero. Then, for some $\delta > 0$, we can construct two subsequences $\{\ell_t\}_{t \geq 0}$ and $\{u_t\}_{t \geq 0}$ such that $\ell_t < u_t$ and $\mathbb{E}[\|\Phi(\mathbf{x}^{\ell_t})\|] \geq 2\delta$, $\mathbb{E}[\|\Phi(\mathbf{x}^{u_t})\|^a] \leq \delta^a$, and $\mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] > \delta^a$ for all $\ell_t < k < u_t$. Based on this construction, the conditions in the theorem, and a set of inequalities, we will eventually reach a contradiction. We notice that the Lipschitz continuity of Φ plays a prominent role when establishing this contradiction. Our overall proof strategy is inspired by the analysis of classical trust region-type methods, see, e.g., [\[9, Theorem 6.4.6\]](#). Let us also mention that a different strategy for the fully deterministic setting and

scalar case $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}$ was provided in [8]. For item (ii), we first control the stochastic behavior of the error terms \mathbf{A}_k by martingale convergence theory. We can then conduct sample-based arguments to derive the final result, which is essentially deterministic and hence, follows similar arguments to that of item (i).

The major application areas of our unified convergence framework comprise stochastic optimization methods that have non-vanishing stochastic errors or that utilize diminishing step sizes. In the next subsection, we state the main steps for showing convergence of stochastic optimization methods. This also clarifies the abstract conditions listed in the theorem.

2.1 The steps for showing convergence of stochastic optimization methods

In order to apply the unified convergence theorem, we have to verify the conditions stated in the theorem, resulting in three main phases below.

Phase I: Verifying (P.1)–(P.2). Conditions (P.1)–(P.2) are used for both the expected and the almost sure convergence results. Condition (P.1) is a problem property and is very standard. We present the final convergence results in terms of the abstract measure Φ . This measure can be regarded as $f - f^*$ in convex optimization, ∇f in smooth nonconvex optimization, the gradient of the Moreau envelope in weakly convex optimization, etc. In all the situations, assuming Lipschitz continuity of the convergence measure Φ is standard and is arguably a minimal assumption in order to obtain iteration complexity and/or convergence results.

Condition (P.2) is typically a result of the algorithmic property or complexity analysis. To verify this condition, one first establishes the recursion of the stochastic method, which almost always has the form

$$\mathbb{E}[\mathbf{y}_{k+1} \mid \mathcal{F}_k] \leq (1 + \beta_k) \mathbf{y}_k - \mu_k \|\Phi(\mathbf{x}^k)\|^a + \zeta_k.$$

Here, \mathbf{y}_k is a suitable Lyapunov function measuring the (approximate) descent property of the stochastic method, ζ_k represents the error term satisfying $\sum_{k=0}^{\infty} \zeta_k < \infty$, β_k is often related to the step sizes and satisfies $\sum_{k=0}^{\infty} \beta_k < \infty$. Then, applying the supermartingale convergence theorem (see Theorem B.1), we obtain $\sum_{k=0}^{\infty} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] < \infty$, i.e., condition (P.2).

Since condition (P.2) is typically a consequence of the underlying algorithmic recursion, one can also derive the standard finite-time complexity bound (1) in terms of the measure $\mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a]$ based on it. Hence, non-asymptotic complexity results are also included implicitly in our framework as a special case. To be more specific, (P.2) implies $\sum_{k=0}^T \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] \leq M$ for some constant $M > 0$ and some total number of iterations T . This then yields $\min_{0 \leq k \leq T} \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] \leq M / \sum_{k=0}^T \mu_k$. Note that the sequence $\{\mu_k\}_{k \geq 0}$ is often related to the step sizes. Thus, choosing the step sizes properly results in the standard finite-time complexity result.

Phase II: Verifying (P.3)–(P.4) for showing expected convergence. Condition (P.3) requires an upper bound on the step length of the update in terms of expectation, including upper bounds for the search direction and the stochastic error of the algorithm. It is often related to certain bounded variance-type assumptions for analyzing stochastic methods. For instance, (P.3) is satisfied under the standard bounded variance assumption for SGD, the more general ABC assumption for SGD, the bounded stochastic subgradients assumption, etc. Condition (P.4) is a standard diminishing step sizes condition used in stochastic optimization.

Then, one can apply item (i) of Theorem 2.1 to obtain $\mathbb{E}[\|\Phi(\mathbf{x}^k)\|] \rightarrow 0$.

Phase III: Verifying (P.3')–(P.4') for showing almost sure convergence. Condition (P.3') is parallel to (P.3). It decomposes the update into a martingale term \mathbf{A}_k and a bounded error term \mathbf{B}_k . We will see later that this condition holds true for many stochastic methods. Though this condition requires the update to have a certain decomposable form, it indeed can be verified by bounding the step length of the update in conditional expectation, which is similar to (P.3). Hence, (P.3') can be interpreted as a conditional version of (P.3). To see this, we can construct

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \underbrace{\mu_k \cdot \frac{1}{\mu_k} (\mathbf{x}^{k+1} - \mathbf{x}^k - \mathbb{E}[\mathbf{x}^{k+1} - \mathbf{x}^k \mid \mathcal{F}_k])}_{\mathbf{A}_k} + \underbrace{\mu_k \cdot \frac{1}{\mu_k} \mathbb{E}[\mathbf{x}^{k+1} - \mathbf{x}^k \mid \mathcal{F}_k]}_{\mathbf{B}_k}. \quad (4)$$

By Jensen's inequality, we then have $\mathbb{E}[\mathbf{A}_k \mid \mathcal{F}_k] = 0$,

$$\mathbb{E}[\|\mathbf{A}_k\|^q] \leq 2^q \mu_k^{-q} \cdot \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^q], \quad \text{and} \quad \|\mathbf{B}_k\|^q \leq \mu_k^{-q} \cdot \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^q \mid \mathcal{F}_k].$$

Thus, once it is possible to derive $\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^q \mid \mathcal{F}_k] = \mathcal{O}(\mu_k^q)$ in an almost sure sense, condition (P.3') is verified with $p_1 = p_2 = 1$. Condition (P.4') is parallel to (P.4) and is standard in stochastic optimization. Application of item (ii) of [Theorem 2.1](#) then yields $\|\Phi(\mathbf{x}^k)\| \rightarrow 0$ almost surely.

In the next section, we will illustrate how to show convergence for a set of classic stochastic methods by following the above three steps.

3 Applications to stochastic optimization methods

3.1 Convergence results of SGD

We consider the standard SGD method for solving the smooth optimization problem $\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x})$, where the iteration of SGD is given by

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \alpha_k \mathbf{g}^k. \quad (5)$$

Here, \mathbf{g}^k denotes a stochastic approximation of the gradient $\nabla f(\mathbf{x}^k)$. We assume that each stochastic gradient \mathbf{g}^k is \mathcal{F}_{k+1} -measurable and that the generated stochastic process $\{\mathbf{x}^k\}_{k \geq 0}$ is adapted to the filtration $\{\mathcal{F}_k\}_{k \geq 0}$. We consider the following standard assumptions:

- (A.1) The mapping $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz continuous on \mathbb{R}^n with modulus $L > 0$.
- (A.2) The objective function f is bounded from below on \mathbb{R}^n , i.e., there is \bar{f} such that $f(\mathbf{x}) \geq \bar{f}$ for all $\mathbf{x} \in \mathbb{R}^n$.
- (A.3) Each oracle \mathbf{g}^k defines an unbiased estimator of $\nabla f(\mathbf{x}^k)$, i.e., it holds that $\mathbb{E}[\mathbf{g}^k \mid \mathcal{F}_k] = \nabla f(\mathbf{x}^k)$ almost surely, and there exist $C, D \geq 0$ such that

$$\mathbb{E}[\|\mathbf{g}^k - \nabla f(\mathbf{x}^k)\|^2 \mid \mathcal{F}_k] \leq C[f(\mathbf{x}^k) - \bar{f}] + D \quad \text{almost surely} \quad \forall k \in \mathbb{N}.$$

- (A.4) The step sizes $\{\alpha_k\}_{k \geq 0}$ satisfy $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\sum_{k=0}^{\infty} \alpha_k^2 < \infty$.

We now derive the convergence of SGD below by setting $\Phi \equiv \nabla f$ and $\mu_k \equiv \alpha_k$.

Phase I: Verifying (P.1)–(P.2). (A.1) verifies condition (P.1) with $L_\Phi \equiv L$. We now check (P.2). Using (A.2), (A.3), and a standard analysis for SGD gives the following recursion (see [Appendix C.1](#) for the full derivation):

$$\mathbb{E}[f(\mathbf{x}^{k+1}) - \bar{f} \mid \mathcal{F}_k] \leq \left(1 + \frac{LC\alpha_k^2}{2}\right) [f(\mathbf{x}^k) - \bar{f}] - \alpha_k \left(1 - \frac{L\alpha_k}{2}\right) \|\nabla f(\mathbf{x}^k)\|^2 + \frac{LD\alpha_k^2}{2}. \quad (6)$$

Taking total expectation, using (A.4), and applying the supermartingale convergence theorem ([Theorem B.1](#)) gives $\sum_{k=0}^{\infty} \alpha_k \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|^2] < \infty$. Furthermore, the sequence $\{\mathbb{E}[f(\mathbf{x}^k)]\}_{k \geq 0}$ converges to some finite value. This verifies (P.2) with $a = 2$.

Phase II: Verifying (P.3)–(P.4) for showing expected convergence. For (P.3), we have by (5) and (A.3) that

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2] \leq \alpha_k^2 \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|^2] + C\alpha_k^2 \mathbb{E}[f(\mathbf{x}^k) - \bar{f}] + D\alpha_k^2.$$

Due to the convergence of $\{\mathbb{E}[f(\mathbf{x}^k)]\}_{k \geq 0}$, there exists F such that $\mathbb{E}[f(\mathbf{x}^k) - \bar{f}] \leq F$ for all k . Thus, condition (P.3) holds with $q = 2$, $A = CF + D$, $p_1 = 2$, $B = 1$, $p_2 = 2$, and $b = 2$. Condition (P.4) is verified by (A.4) and the previous parameters choices. Therefore, we can apply [Theorem 2.1](#) to deduce $\mathbb{E}[\|\nabla f(\mathbf{x}^k)\|] \rightarrow 0$.

Phase III: Verifying (P.3')–(P.4') for showing almost sure convergence. For (P.3'), it follows from the update (5) that

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \alpha_k (\mathbf{g}^k - \nabla f(\mathbf{x}^k)) - \alpha_k \nabla f(\mathbf{x}^k).$$

We have $p_1 = 1$, $\mathbf{A}_k = \mathbf{g}^k - \nabla f(\mathbf{x}^k)$, $p_2 = 1$, and $\mathbf{B}_k = \nabla f(\mathbf{x}^k)$. Using (A.2), (A.3), $\mathbb{E}[f(\mathbf{x}^k) - \bar{f}] \leq F$, and choosing any $q = b > 0$ establishes (P.3'). As before, condition (P.4') follows from (A.4) and the previous parameters choices. Applying [Theorem 2.1](#) yields $\|\nabla f(\mathbf{x}^k)\| \rightarrow 0$ almost surely.

Finally, we summarize the above results in the following corollary.

Corollary 3.1. *Let us consider SGD (5) for smooth nonconvex optimization problems under (A.1)–(A.4). Then, we have $\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|] = 0$ and $\lim_{k \rightarrow \infty} \|\nabla f(\mathbf{x}^k)\| = 0$ almost surely.*

3.2 Convergence results of random reshuffling

We now consider random reshuffling (RR) applied to problems with a finite sum structure

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}, i),$$

where each component function $f(\cdot, i) : \mathbb{R}^n \rightarrow \mathbb{R}$ is supposed to be smooth. At iteration k , RR first generates a random permutation σ^{k+1} of the index set $\{1, \dots, N\}$. It then updates \mathbf{x}^k to \mathbf{x}^{k+1} through N consecutive gradient descent-type steps by accessing and using the component gradients $\{\nabla f(\cdot, \sigma_1^{k+1}), \dots, \nabla f(\cdot, \sigma_N^{k+1})\}$ sequentially. Specifically, one update-loop (epoch) of RR is given by

$$\tilde{\mathbf{x}}_0^k = \mathbf{x}^k, \quad \tilde{\mathbf{x}}_i^k = \tilde{\mathbf{x}}_{i-1}^k - \alpha_k \nabla f(\tilde{\mathbf{x}}_{i-1}^k, \sigma_i^{k+1}), \quad i = 1, \dots, N, \quad \mathbf{x}^{k+1} = \tilde{\mathbf{x}}_N^k. \quad (7)$$

After one such loop, the step size α_k and the permutation σ^{k+1} is updated accordingly; cf. [20, 30, 32]. We make the following standard assumptions:

- (B.1) For all $i \in \{1, \dots, N\}$, $f(\cdot, i)$ is bounded from below by some \bar{f} and the gradient $\nabla f(\cdot, i)$ is Lipschitz continuous on \mathbb{R}^n with modulus $L > 0$.
- (B.2) The step sizes $\{\alpha_k\}_{k \geq 0}$ satisfy $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\sum_{k=0}^{\infty} \alpha_k^3 < \infty$.

A detailed derivation of the steps shown in Subsection 2.1 for RR is deferred to Appendix D.2. Based on the discussion in Appendix D.2 and on Theorem 2.1, we obtain the following results for RR.

Corollary 3.2. *We consider RR (7) for smooth nonconvex optimization problems under (B.1)–(B.2). Then it holds that $\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|] = 0$ and $\lim_{k \rightarrow \infty} \|\nabla f(\mathbf{x}^k)\| = 0$ almost surely.*

3.3 Convergence of the proximal stochastic gradient method

We consider the composite-type optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} \psi(\mathbf{x}) := f(\mathbf{x}) + \varphi(\mathbf{x}) \quad (8)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a continuously differentiable function and $\varphi : \mathbb{R}^n \rightarrow (-\infty, \infty]$ is τ -weakly convex (see Appendix E.1), proper, and lower semicontinuous. In this section, we want to apply our unified framework to study the convergence behavior of the well-known proximal stochastic gradient method (prox-SGD):

$$\mathbf{x}^{k+1} = \text{prox}_{\alpha_k \varphi}(\mathbf{x}^k - \alpha_k \mathbf{g}^k), \quad (9)$$

where $\mathbf{g}^k \approx \nabla f(\mathbf{x}^k)$ is a stochastic approximation of $\nabla f(\mathbf{x}^k)$, $\{\alpha_k\}_{k \geq 0} \subseteq \mathbb{R}_+$ is a suitable step size sequence, and $\text{prox}_{\alpha_k \varphi} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $\text{prox}_{\alpha_k \varphi}(\mathbf{x}) := \text{argmin}_{\mathbf{y} \in \mathbb{R}^n} \varphi(\mathbf{y}) + \frac{1}{2\alpha_k} \|\mathbf{x} - \mathbf{y}\|^2$ is the well-known proximity operator of φ .

3.3.1 Assumptions and preparations

We first recall several useful concepts from nonsmooth and variational analysis. For a function $h : \mathbb{R}^n \rightarrow (-\infty, \infty]$, the Fréchet (or regular) subdifferential of h at the point \mathbf{x} is given by

$$\partial h(\mathbf{x}) := \{\mathbf{g} \in \mathbb{R}^n : h(\mathbf{y}) \geq h(\mathbf{x}) + \langle \mathbf{g}, \mathbf{y} - \mathbf{x} \rangle + o(\|\mathbf{y} - \mathbf{x}\|) \text{ as } \mathbf{y} \rightarrow \mathbf{x}\},$$

see, e.g., [39, Chapter 8]. If h is convex, then the Fréchet subdifferential coincides with the standard (convex) subdifferential. It is well-known that the associated first-order optimality condition for the composite problem (8) — $0 \in \partial \psi(\mathbf{x}) = \nabla f(\mathbf{x}) + \partial \varphi(\mathbf{x})$ — can be represented as a nonsmooth equation, [39, 21],

$$F_{\text{nat}}^\alpha(\mathbf{x}) := \mathbf{x} - \text{prox}_{\alpha \varphi}(\mathbf{x} - \alpha \nabla f(\mathbf{x})) = 0, \quad \alpha \in (0, \tau^{-1}),$$

where F_{nat}^α denotes the so-called *natural residual*. The natural residual F_{nat}^α is a common stationarity measure for the nonsmooth problem (8) and widely used in the analysis of proximal methods.

We will make the following assumptions on f , φ , and the stochastic oracles $\{\mathbf{g}^k\}_{k \geq 0}$:

- (C.1) The function f is bounded from below on \mathbb{R}^n , i.e., there is \bar{f} such that $f(\mathbf{x}) \geq \bar{f}$ for all $\mathbf{x} \in \mathbb{R}^n$, and the gradient mapping ∇f is Lipschitz continuous (on \mathbb{R}^n) with modulus $L > 0$.
- (C.2) The function φ is τ -weakly convex, proper, lower semicontinuous, and bounded from below on $\text{dom } \varphi$, i.e., we have $\varphi(\mathbf{x}) \geq \bar{\varphi}$ for all $\mathbf{x} \in \text{dom } \varphi$.

(C.3) There exists $L_\varphi > 0$ such that $\varphi(\mathbf{x}) - \varphi(\mathbf{y}) \leq L_\varphi \|\mathbf{x} - \mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \text{dom } \varphi$.

(C.4) Each \mathbf{g}^k defines an unbiased estimator of $\nabla f(\mathbf{x}^k)$, i.e., we have $\mathbb{E}[\mathbf{g}^k \mid \mathcal{F}_k] = \nabla f(\mathbf{x}^k)$ almost surely, and there exist $C, D \geq 0$ such that

$$\mathbb{E}[\|\mathbf{g}^k - \nabla f(\mathbf{x}^k)\|^2 \mid \mathcal{F}_k] \leq C[f(\mathbf{x}^k) - \bar{f}] + D \quad \text{almost surely} \quad \forall k \in \mathbb{N}.$$

(C.5) The step sizes $\{\alpha_k\}_{k \geq 0}$ satisfy $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\sum_{k=0}^{\infty} \alpha_k^2 < \infty$.

Here, we again assume that the generated stochastic processes $\{\mathbf{x}^k\}_{k \geq 0}$ is adapted to the filtration $\{\mathcal{F}_k\}_{k \geq 0}$. The assumptions (C.1), (C.2), (C.4), and (C.5) are fairly standard and broadly applicable. In particular, (C.1), (C.4), and (C.5) coincide with the conditions (A.1)–(A.4) used in the analysis of SGD. We continue with several remarks concerning condition (C.3).

Remark 3.3. Assumption (C.3) requires the mapping φ to be Lipschitz continuous on its effective domain $\text{dom } \varphi$. This condition holds in many important applications, e.g., when φ is chosen as a norm or indicator function. Nonconvex examples satisfying (C.2) and (C.3) include, e.g., the minimax concave penalty (MCP) function [45], the smoothly clipped absolute deviation (SCAD) [15], or the student-t loss function. We refer to [4] and Appendix E.2 for further discussion.

3.3.2 Convergence results of prox-SGD

We now analyze the convergence of the random process $\{\mathbf{x}^k\}_{k \geq 0}$ generated by the stochastic algorithmic scheme (9). As pioneered in [11], we will use the Moreau envelope $\text{env}_{\theta\psi}$,

$$\text{env}_{\theta\psi} : \mathbb{R}^n \rightarrow \mathbb{R}, \quad \text{env}_{\theta\psi}(\mathbf{x}) := \min_{\mathbf{y} \in \mathbb{R}^n} \psi(\mathbf{y}) + \frac{1}{2\theta} \|\mathbf{x} - \mathbf{y}\|^2, \quad (10)$$

as a smooth Lyapunov function to study the descent properties and convergence of prox-SGD.

We first note that the conditions (C.1) and (C.2) imply θ^{-1} -weak convexity of ψ for every $\theta \in (0, (L + \tau)^{-1}]$. In this case, the Moreau envelope $\text{env}_{\theta\psi}$ is a well-defined and continuously differentiable function with gradient $\nabla \text{env}_{\theta\psi}(\mathbf{x}) = \frac{1}{\theta}(\mathbf{x} - \text{prox}_{\theta\psi}(\mathbf{x}))$; see, e.g., [38, Theorem 31.5].

As shown in [13, 11], the norm of the Moreau envelope — $\|\nabla \text{env}_{\theta\psi}(\mathbf{x})\|$ — defines an alternative stationarity measure for problem (8) that is equivalent to the natural residual if θ is chosen sufficiently small. A more explicit derivation of this connection is provided in Lemma E.1.

Next, we establish convergence of prox-SGD by setting $\Phi \equiv \nabla \text{env}_{\theta\psi}$ and $\mu_k \equiv \alpha_k$. Our analysis is based on the following two estimates which are verified in Appendix E.4 and Appendix E.5.

Lemma 3.4. *Let $\{\mathbf{x}^k\}_{k \geq 0}$ be generated by prox-SGD and let the assumptions (C.1)–(C.4) be satisfied. Then, for $\theta \in (0, [3L + \tau]^{-1})$ and all k with $\alpha_k \leq \min\{\frac{1}{2\tau}, \frac{1}{2(\theta^{-1} - [L + \tau])}\}$, it holds that*

$$\begin{aligned} \mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi} \mid \mathcal{F}_k] &\leq (1 + 4C\theta^{-1}\alpha_k^2) \cdot [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] \\ &\quad - L\theta\alpha_k \|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2 + 2\alpha_k^2 (CL_\varphi^2 + D\theta^{-1}), \end{aligned} \quad (11)$$

almost surely, where $\bar{\psi} := \bar{f} + \bar{\varphi}$.

Lemma 3.5. *Let $\{\mathbf{x}^k\}_{k \geq 0}$ be generated by prox-SGD and suppose that the assumptions (C.1)–(C.4) hold. Then, for $\theta \in (0, [\frac{4}{3}L + \tau]^{-1})$ and all k with $\alpha_k \leq \frac{1}{2\tau}$, we have almost surely*

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \leq 8(2L + C)\alpha_k^2 \cdot [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + 4((2L + C)\theta + 1)L_\varphi^2 + D)\alpha_k^2. \quad (12)$$

Phase I: Verifying (P.1)–(P.2). In [21, Corollary 3.4], it is shown that the gradient of the Moreau envelope is Lipschitz continuous with modulus $L_e := \max\{\theta^{-1}, (1 - [L + \tau]\theta)^{-1}[L + \tau]\}$ for all $\theta \in (0, [L + \tau]^{-1})$. Thus, condition (P.1) is satisfied.

Furthermore, due to $\alpha_k \rightarrow 0$ and choosing $\theta \in (0, [3L + \tau]^{-1})$, the estimate (11) in Lemma 3.4 holds for all k sufficiently large. Consequently, due to $\text{env}_{\theta\psi}(\mathbf{x}) \geq \psi(\text{prox}_{\theta\psi}(\mathbf{x})) \geq \bar{\psi}$ and (C.5), Theorem B.1 is applicable and upon taking total expectation, $\{\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^k)]\}_{k \geq 0}$ converges to some $E \in \mathbb{R}$. In addition, the sequence $\{\text{env}_{\theta\psi}(\mathbf{x}^k)\}_{k \geq 0}$ converges almost surely to some random variable e^* and we have $\sum_{k=0}^{\infty} \alpha_k \mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] < \infty$. This verifies condition (P.2) with $a = 2$.

Phase II: Verifying (P.3)–(P.4) for showing convergence in expectation. Assumptions (C.1)–(C.5) and Lemma 3.5 allow us to establish the required bound stated in (P.3). Specifically, taking total

expectation in (12), we have

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2] \leq 8(2L + C)\alpha_k^2 \cdot \mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + 4(((2L + C)\theta + 1)L_\varphi^2 + D)\alpha_k^2$$

for all k sufficiently large. Due to $\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^k)] \rightarrow \mathbf{E}$, there exists F such that $\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] \leq F$ for all k . Hence, (P.3) holds with $q = 2$, $A = 8(2L + C)F + 4(((2L + C)\theta + 1)L_\varphi^2 + D)$, $p_1 = 2$, and $B = 0$. The property (P.4) easily follows from (C.5) and the parameter choices. Consequently, using Theorem 2.1, we can infer $\mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|] \rightarrow 0$.

Phase III: Verifying (P.3')–(P.4') for showing almost sure convergence. We follow the construction in (4) and set $\mathbf{A}_k = \alpha_k^{-1}(\mathbf{x}^{k+1} - \mathbf{x}^k - \mathbb{E}[\mathbf{x}^{k+1} - \mathbf{x}^k \mid \mathcal{F}_k])$, $\mathbf{B}_k = \alpha_k^{-1}\mathbb{E}[\mathbf{x}^{k+1} - \mathbf{x}^k \mid \mathcal{F}_k]$, and $p_1, p_2 = 1$. Clearly, we have $\mathbb{E}[\mathbf{A}_k \mid \mathcal{F}_k] = 0$ and based on the previous results in **Phase II**, we can show $\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2] = \mathcal{O}(\alpha_k^2)$ which establishes boundedness of $\{\mathbb{E}[\|\mathbf{A}_k\|^2]\}_{k \geq 0}$. Similarly, for \mathbf{B}_k and by Lemma 3.5 and Jensen's inequality, we obtain

$$\|\mathbf{B}_k\|^2 \leq \alpha_k^{-2}\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \leq 8(2L + C) \cdot [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + \mathcal{O}(1).$$

Due to $\text{env}_{\theta\psi}(\mathbf{x}^k) \rightarrow e^*$ almost surely, this shows $\limsup_{k \rightarrow \infty} \|\mathbf{B}_k\|^2 < \infty$ almost surely. Hence, all requirements in (P.3') are satisfied with $q = 2$ and $b = 0$. Moreover, it is easy to see that property (P.4') also holds in this case. Overall, Theorem 2.1 implies $\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\| \rightarrow 0$ almost surely.

As mentioned, it is possible to express the obtained convergence results in terms of the natural residual $F_{\text{nat}} = F_{\text{nat}}^1$, see, e.g., Lemma E.1. We summarize our observations in the following corollary.

Corollary 3.6. *Let us consider prox-SGD (9) for the composite problem (8) under (C.1)–(C.5). Then, we have $\lim_{k \rightarrow \infty} \mathbb{E}[\|F_{\text{nat}}(\mathbf{x}^k)\|] = 0$ and $\lim_{k \rightarrow \infty} \|F_{\text{nat}}(\mathbf{x}^k)\| = 0$ almost surely.*

Remark 3.7. As a byproduct, Lemma 3.4 also leads to an expected iteration complexity result of prox-SGD by using the ABC condition (C.4) rather than the standard bounded variance assumption. This is a nontrivial extension of [11, Corollary 3.6]. We provide a full derivation in Appendix E.6.

3.4 Convergence of stochastic model-based methods

In this section, we consider the convergence of stochastic model-based methods for nonsmooth weakly convex optimization problems

$$\min_{\mathbf{x} \in \mathbb{R}^n} \psi(\mathbf{x}) := f(\mathbf{x}) + \varphi(\mathbf{x}) = \mathbb{E}_{\xi \sim D}[f(\mathbf{x}, \xi)] + \varphi(\mathbf{x}), \quad (13)$$

where both f and φ are assumed to be (nonsmooth) weakly convex functions and ψ is lower bounded, i.e., $\psi(\mathbf{x}) \geq \bar{\psi}$ for all $\mathbf{x} \in \text{dom } \varphi$. Classical stochastic optimization methods — including proximal stochastic subgradient, stochastic proximal point, and stochastic prox-linear methods — for solving (13) are unified by the stochastic model-based methods (SMM) [14, 11]:

$$\mathbf{x}^{k+1} = \text{argmin}_{\mathbf{x} \in \mathbb{R}^n} f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) + \varphi(\mathbf{x}) + \frac{1}{2\alpha_k} \|\mathbf{x} - \mathbf{x}^k\|^2, \quad (14)$$

where $f_{\mathbf{x}^k}(\mathbf{x}, \xi^k)$ is a stochastic approximation of f around \mathbf{x}^k using the sample ξ^k ; see Appendix F.1 for descriptions of three major types of SMM. Setting $\mathcal{F}_k := \sigma(\xi^0, \dots, \xi^{k-1})$, it is easy to see that $\{\mathbf{x}^k\}_{k \geq 0}$ is adapted to $\{\mathcal{F}_k\}_{k \geq 0}$. We analyze convergence of SMM under the following assumptions.

(D.1) The stochastic approximation function $f_{\mathbf{x}}$ satisfies a one-sided accuracy property, i.e., we have $\mathbb{E}_\xi[f_{\mathbf{x}}(\mathbf{x}, \xi)] = f(\mathbf{x})$ for all $\mathbf{x} \in U$ and

$$\mathbb{E}_\xi[f_{\mathbf{x}}(\mathbf{y}, \xi) - f(\mathbf{y})] \leq \frac{\tau}{2} \|\mathbf{x} - \mathbf{y}\|^2 \quad \forall \mathbf{x}, \mathbf{y} \in U,$$

where U is an open convex set containing $\text{dom } \varphi$.

(D.2) The function $\mathbf{y} \mapsto f_{\mathbf{x}}(\mathbf{y}, \xi) + \varphi(\mathbf{y})$ is η -weakly convex for all $\mathbf{x} \in U$ and almost every ξ .

(D.3) There exists $L > 0$ such that the stochastic approximation function $f_{\mathbf{x}}$ satisfies

$$f_{\mathbf{x}}(\mathbf{x}, \xi) - f_{\mathbf{x}}(\mathbf{y}, \xi) \leq L\|\mathbf{x} - \mathbf{y}\| \quad \forall \mathbf{x}, \mathbf{y} \in U, \quad \text{and almost every } \xi.$$

(D.4) The function φ is L_φ -Lipschitz continuous.

(D.5) The step sizes $\{\alpha_k\}_{k \geq 0}$ satisfy $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\sum_{k=0}^{\infty} \alpha_k^2 < \infty$.

Assumptions (D.1), (D.2), (D.3) are standard for analyzing SMM and identical to that of [11]. (D.5) is convention for stochastic methods. Assumption (D.4) mimics (C.3); see Remark 3.3 for discussions.

We now derive the convergence of SMM below by setting $\Phi \equiv \nabla \text{env}_{\theta\psi}$ and $\mu_k \equiv \alpha_k$. Our derivation is based on the following two estimates, in which the proof of Lemma 3.9 is given in Appendix F.2.

Lemma 3.8 (Theorem 4.3 of [11]). *Let $\theta \in (0, (\tau + \eta)^{-1})$ and $\alpha_k < \theta$ be given. Then, we have*

$$\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] \leq \text{env}_{\theta\psi}(\mathbf{x}^k) - \frac{(1 - [\tau + \eta]\theta)\alpha_k}{2(1 - \eta\alpha_k)} \|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2 + \frac{2L^2\alpha_k^2}{(1 - \eta\alpha_k)(\theta - \alpha_k)}.$$

Lemma 3.9. *For all k with $\alpha_k \leq 1/(2\eta)$, it holds that*

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \leq (16(L + L_\varphi)^2 + 8L^2)\alpha_k^2.$$

Phase I: Verifying (P.1)–(P.2). As before, [21, Corollary 3.4] implies that the mapping $\nabla \text{env}_{\theta\psi}$ is Lipschitz continuous for all $\theta \in (0, (\tau + \eta)^{-1})$. Hence, condition (P.1) is satisfied. Using $\alpha_k \rightarrow 0$, we can apply Theorem B.1 to the recursion obtained in Lemma 3.8 for all k sufficiently large and it follows $\sum_{k=0}^{\infty} \alpha_k \mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] < \infty$. Thus, condition (P.2) holds with $a = 2$.

Phase II: Verifying (P.3)–(P.4) for showing convergence in expectation. Taking total expectation in Lemma 3.9 verifies condition (P.3) with $q = 2$, $A = (16(L + L_\varphi)^2 + 8L^2)$, $p_1 = 2$, $B = 0$. Moreover, condition (P.4) is true by assumption (D.5) and the previous parameters choices. Thus, applying Theorem 2.1 gives $\mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|] \rightarrow 0$.

Phase III: Verifying (P.3')–(P.4') for showing almost sure convergence. As in (4), we can set $\mathbf{A}_k = \alpha_k^{-1}(\mathbf{x}^{k+1} - \mathbf{x}^k - \mathbb{E}[\mathbf{x}^{k+1} - \mathbf{x}^k \mid \mathcal{F}_k])$, $\mathbf{B}_k = \alpha_k^{-1}\mathbb{E}[\mathbf{x}^{k+1} - \mathbf{x}^k \mid \mathcal{F}_k]$. Applying Lemma 3.9 and utilizing Jensen's inequality, we have $\mathbb{E}[\mathbf{A}_k \mid \mathcal{F}_k] = 0$, $\mathbb{E}[\|\mathbf{A}_k\|^2] \leq (4/\alpha_k^2)\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2] \leq 4(16(L + L_\varphi)^2 + 8L^2)$ and $\|\mathbf{B}_k\|^2 \leq 16(L + L_\varphi)^2 + 8L^2$. Thus, condition (P.3') is satisfied with $p_1 = p_2 = 1$, $q = 2$. Assumption (D.5), together with the previous parameter choices verifies condition (P.4') and hence, applying Theorem 2.1 yields $\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\| \rightarrow 0$ almost surely.

Summarizing this discussion, we obtain the following convergence results for SMM.

Corollary 3.10. *We consider the family of stochastic model-based methods (14) for the optimization problem (13) under assumptions (D.1)–(D.5). Let $\{\mathbf{x}^k\}_{k \geq 0}$ be a generated sequence. Then, we have $\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|] = 0$ and $\lim_{k \rightarrow \infty} \|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\| = 0$ almost surely.*

Remark 3.11. The results presented in Corollary 3.10 also hold under certain extended settings. In fact, we can replace (D.3) by a slightly more general Lipschitz continuity assumption on f . Moreover, it is possible to establish convergence in the case where f is not Lipschitz continuous but has Lipschitz continuous gradient, which is particularly useful when we apply stochastic proximal point method for smooth f . A more detailed derivation and discussion of such extensions is deferred to Appendix F.3.

3.5 Related work and discussion

SGD and RR. The literature for SGD is extremely rich and several connected and recent works have been discussed in Section 1. Our result in Corollary 3.1 unifies many of the existing convergence analyses of SGD and is based on the general ABC condition (A.3) (see [23, 24, 19] for comparison) rather than on the standard bounded variance assumption. Our expected convergence result generalizes the one in [6] using much weaker assumptions. Our results for RR are in line with the recent theoretical observations in [30, 32, 25]. In particular, Corollary 3.2 recovers the almost sure convergence result shown in [25], while the expected convergence result appears to be new.

Prox-SGD and SMM. The work [11] established one of the first complexity results for prox-SGD using the Moreau envelope. Under a bounded variance assumption ($C = 0$ in condition (C.4)) and for general nonconvex and smooth f , the authors showed $\mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] = \mathcal{O}((T + 1)^{-1/2})$, where \mathbf{x}^k is sampled uniformly from the past $T + 1$ iterates $\mathbf{x}^0, \dots, \mathbf{x}^T$. As mentioned, this result cannot be easily extended to the asymptotic convergence results discussed in this paper. Earlier studies of prox-SGD for nonconvex f and $C = 0$ include [18] where convergence of prox-SGD is established if the variance parameter $D = D_k \rightarrow 0$ vanishes as $k \rightarrow \infty$. This can be achieved by progressively increasing the size of the selected mini-batches or via variance reduction techniques as in prox-SVRG and prox-SAGA, see [35]. The question whether prox-SGD can converge and whether the accumulation points of the iterates $\{\mathbf{x}^k\}_{k \geq 0}$ correspond to stationary points was only addressed recently in [27]. The authors use a differential inclusion approach to study convergence of prox-SGD. However, additional compact constraints $\mathbf{x} \in \mathcal{X}$ have to be introduced in the model (8) to guarantee sure boundedness of $\{\mathbf{x}^k\}_{k \geq 0}$ and applicability of the differential inclusion techniques. Lipschitz

continuity of φ also appears as an essential requirement in [27, Theorem 5.4]. The analyses in [14, 12] establish asymptotic convergence guarantees for SMM. However, both works require a priori (almost) sure boundedness of $\{\mathbf{x}^k\}_{k \geq 0}$ and a density / Sard-type condition in order to show convergence. We refer to [16] for an extension of the results in [27, 12] to prox-SGD in Hilbert spaces. By contrast, our convergence framework allows to complement these differential inclusion-based results and — for the first time — fully removes any stringent boundedness assumption on $\{\mathbf{x}^k\}_{k \geq 0}$. Instead, our analysis relies on more transparent assumptions that are verifiable and common in stochastic optimization and machine learning. In summary, we are now able to claim: *prox-SGD and SMM converge under standard stochastic conditions if φ is Lipschitz continuous*. In the easier convex case, analogous results have been obtained, e.g., in [18, 1, 40].

We provide an overview of several related and representative results in Table 1 in Appendix G.

4 Conclusion

In this work, we provided a novel convergence framework that allows to derive expected and almost sure convergence results for a vast class of stochastic optimization methods under state-of-the-art assumptions and in a unified way. We specified the steps on how to utilize our theorem in order to establish convergence results for a given stochastic algorithm. As concrete examples, we applied our theorem to derive asymptotic convergence guarantees for SGD, RR, prox-SGD, and SMM. To our surprise, some of the obtained results appear to be new and provide new insights into the convergence behavior of some well-known and standard stochastic methodologies. These applications revealed that our unified theorem can serve as a plugin-type tool with the potential to facilitate the convergence analysis of a wide class of stochastic optimization methods.

Finally, it is important to investigate in which situations our convergence results in terms of the stationarity measure Φ can be strengthened — say to almost sure convergence guarantees for the iterates $\{\mathbf{x}^k\}_{k \geq 0}$. We plan to consider such a possible extension in future work.

Acknowledgments and Disclosure of Funding

The authors would like to thank the Area Chair and anonymous reviewers for their detailed and constructive comments, which have helped greatly to improve the quality and presentation of the manuscript. In addition, we would like to thank Michael Ulbrich for valuable feedback and comments on an earlier version of this work.

X. Li was partially supported by the National Natural Science Foundation of China (NSFC) under Grant No. 12201534 and 72150002, by the Shenzhen Science and Technology Program under Grant No. RCBS20210609103708017, and by the Shenzhen Institute of Artificial Intelligence and Robotics for Society (AIRS) under Grant No. AC01202101108. A. Milzarek was partly supported by the National Natural Science Foundation of China (NSFC) – Foreign Young Scholar Research Fund Project (RFIS) under Grant No. 12150410304 and by the Fundamental Research Fund – Shenzhen Research Institute of Big Data (SRIBD) Startup Fund JCYJ-AM20190601.

References

- [1] Yves F. Atchadé, Gersende Fort, and Eric Moulines. On perturbed proximal gradient algorithms. *J. Mach. Learn. Res.*, 18(10):1–33, 2017.
- [2] Amir Beck. *First-order methods in optimization*. SIAM, Philadelphia, PA, 2017.
- [3] Dimitri P. Bertsekas and John N. Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM J. Optim.*, 10(3):627–642, 2000.
- [4] Axel Böhm and Stephen J. Wright. Variable smoothing for weakly convex composite functions. *J. Optim. Theory Appl.*, 188(3):628–649, 2021.
- [5] Léon Bottou. Stochastic learning. In *Summer School on Machine Learning*, pages 146–168. Springer, 2003.
- [6] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.

- [7] Leo Breiman. *Probability*, volume 7 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992. Corrected reprint of the 1968 original.
- [8] Guy Cohen and Daoli Zhu. Decomposition and coordination methods in large scale optimization problems: The nondifferentiable case and the use of augmented lagrangians. *Adv. in Large Scale Systems*, 1:203–266, 1984.
- [9] Andrew R. Conn, Nicholas I. M. Gould, and Philippe L. Toint. *Trust-region methods*. MPS/SIAM Series on Optim. SIAM; MPS, Philadelphia, PA, 2000.
- [10] J-C Culioli and Guy Cohen. Decomposition/coordination algorithms in stochastic optimization. *SIAM J. Control Optim.*, 28(6):1372–1403, 1990.
- [11] Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM J. Optim.*, 29(1):207–239, 2019.
- [12] Damek Davis, Dmitriy Drusvyatskiy, Sham Kakade, and Jason D. Lee. Stochastic subgradient method converges on tame functions. *Found. Comput. Math.*, 20(1):119–154, 2020.
- [13] Dmitriy Drusvyatskiy and Adrian S. Lewis. Error bounds, quadratic growth, and linear convergence of proximal methods. *Math. Oper. Res.*, 43(3):919–948, 2018.
- [14] John C. Duchi and Feng Ruan. Stochastic methods for composite and weakly convex optimization problems. *SIAM J. Optim.*, 28(4):3229–3259, 2018.
- [15] Jianqing Fan. Comments on “wavelets in statistics: A review” by A. Antoniadis. *Journal of the Italian Statistical Society*, 6(2):131, 1997.
- [16] Caroline Geiersbach and Teresa Scarinci. Stochastic proximal gradient methods for nonconvex problems in Hilbert spaces. *Comput. Optim. Appl.*, 78(3):705–740, 2021.
- [17] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM J. Optim.*, 23(4):2341–2368, 2013.
- [18] Saeed Ghadimi, Guanghui Lan, and Hongchao Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Math. Program.*, 155(1-2, Ser. A):267–305, 2016.
- [19] Robert Gower, Othmane Sebbouh, and Nicolas Loizou. SGD for structured nonconvex functions: Learning rates, minibatching and interpolation. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130, pages 1315–1323. PMLR, 13–15 Apr 2021.
- [20] Mert Gürbüzbalaban, Asu Ozdaglar, and P. A. Parrilo. Why random reshuffling beats stochastic gradient descent. *Math. Program.*, 186(1-2):49–84, 2021.
- [21] Tim Hoheisel, Maxime Laborde, and Adam Oberman. A regularization interpretation of the proximal point method for weakly convex functions. *J. Dyn. Games*, 7(1):79–96, 2020.
- [22] Ya-Ping Hsieh, Panayotis Mertikopoulos, and Volkan Cevher. The limits of min-max optimization algorithms: Convergence to spurious non-critical sets. In *International Conference on Machine Learning*, pages 4337–4348. PMLR, 2021.
- [23] Ahmed Khaled and Peter Richtárik. Better theory for SGD in the nonconvex world. arXiv preprint, arXiv:2002.03329, 2020.
- [24] Yunwen Lei, Ting Hu, Guiying Li, and Ke Tang. Stochastic gradient descent for nonconvex learning without bounded gradient assumptions. *IEEE Transactions on Neural Networks and Learning Systems*, 31(10):4394–4400, 2019.
- [25] Xiao Li, Andre Milzarek, and Junwen Qiu. Convergence of random reshuffling under the Kurdyka-Łojasiewicz inequality. arXiv preprint, arXiv:2110.04926, 2021.

- [26] Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 983–992. PMLR, 2019.
- [27] Szymon Majewski, Blazej Miasojedow, and Eric Moulines. Analysis of nonsmooth stochastic approximation: the differential inclusion approach. arXiv preprint, arXiv:1805.01916v1, 2018.
- [28] Panayotis Mertikopoulos, Nadav Hallak, Ali Kavis, and Volkan Cevher. On the almost sure convergence of stochastic gradient descent in non-convex problems. *Advances in Neural Information Processing Systems*, 33:1117–1128, 2020.
- [29] Andre Milzarek, Fabian Schaipp, and Michael Ulbrich. A semismooth Newton stochastic proximal point algorithm with variance reduction. arXiv preprint, arXiv:2204.00406v1, 2022.
- [30] Konstantin Mishchenko, Ahmed Khaled, and Peter Richtárik. Random reshuffling: Simple analysis with vast improvements. *Advances in Neural Information Processing Systems*, 33:17309–17320, 2020.
- [31] Yu. Nesterov. Gradient methods for minimizing composite functions. *Math. Program.*, 140(1, Ser. B):125–161, 2013.
- [32] Lam M Nguyen, Quoc Tran-Dinh, Dzung T Phan, Phuong Ha Nguyen, and Marten van Dijk. A unified convergence analysis for shuffling-type gradient methods. *J. Mach. Learn.*, 22(207):1–44, 2021.
- [33] Francesco Orabona. Almost sure convergence of SGD on smooth non-convex functions. <https://parameterfree.com/2020/10/05/almost-sure-convergence-of-sgd-on-smooth-non-convex-functions/>, 2020.
- [34] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM J. Control Optim.*, 30(4):838–855, 1992.
- [35] Sashank J. Reddi, Suvrit Sra, Barnabas Póczos, and Alexander J Smola. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. In *Advances in Neural Information Processing Systems 29*, pages 1145–1153. Curran Associates, Inc., 2016.
- [36] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- [37] Herbert Robbins and David Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing Methods in Statistics*, pages 233–257. Elsevier, 1971.
- [38] R. Tyrrell Rockafellar. *Convex analysis*. Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1997. Reprint of the 1970 original, Princeton Paperbacks.
- [39] R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational analysis*, volume 317 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1998.
- [40] Lorenzo Rosasco, Silvia Villa, and B. Công Vũ. Convergence of stochastic proximal gradient algorithm. *Appl. Math. Optim.*, 82(3):891–917, 2020.
- [41] Othmane Sebbouh, Robert M Gower, and Aaron Defazio. Almost sure convergence rates for stochastic gradient descent and stochastic heavy ball. In *Conference on Learning Theory*, pages 3935–3971. PMLR, 2021.
- [42] Ohad Shamir and Tong Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning*, pages 71–79. PMLR, 2013.
- [43] Daniel W. Stroock. *Probability theory*. Cambridge University Press, Cambridge, second edition, 2011. An analytic view.

- [44] Jean-Philippe Vial. Strong and weak convexity of sets and functions. *Math. Oper. Res.*, 8(2):231–259, 1983.
- [45] Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, 38(2):894–942, 2010.
- [46] Tong Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the Twenty-First International Conference on Machine Learning*, page 116, 2004.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] The limitations are written in equivalent forms as future works in the conclusion section; see Section 4.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] We conduct theoretical investigation about the fundamental stochastic optimization methods, which will not bring any negative societal impact.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] We put them in the appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Contents

1	Introduction	1
2	A unified convergence theorem	3
2.1	The steps for showing convergence of stochastic optimization methods	4
3	Applications to stochastic optimization methods	5
3.1	Convergence results of SGD	5
3.2	Convergence results of random reshuffling	6
3.3	Convergence of the proximal stochastic gradient method	6
3.3.1	Assumptions and preparations	6
3.3.2	Convergence results of prox-SGD	7
3.4	Convergence of stochastic model-based methods	8
3.5	Related work and discussion	9
4	Conclusion	10
A	Proof of Theorem 2.1	15
B	The supermartingale convergence theorem	17
C	Proofs for Subsection 3.1	17
C.1	Derivation of (6)	17
D	Results and proofs for Subsection 3.2	17
D.1	A bound on the step length	18
D.2	Proof of Corollary 3.2	18
E	Results and proofs for Subsection 3.3	19
E.1	Weakly convex functions	19
E.2	Examples of weakly convex and Lipschitz continuous regularizers	19
E.3	Equivalent stationarity measures	20
E.4	Proof of Lemma 3.4	20
E.5	Proof of Lemma 3.5	22
E.6	Expected iteration complexity of prox-SGD	23
F	Results and proofs for Subsection 3.4	23
F.1	Three stochastic model-based methods	23
F.2	Proof of Lemma 3.9	24
F.3	Convergence of SMM: An extended setting	24
G	Comparison: related literature	27
H	Non-asymptotic complexity vs. asymptotic convergence	28

A Proof of Theorem 2.1

Part I: Proof of item (i). First, combining $\sum_{k=0}^{\infty} \mu_k = \infty$ in (P.4) and (P.2), we can infer $\liminf_{k \rightarrow \infty} \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] = 0$. Let us assume that $\{\mathbb{E}[\|\Phi(\mathbf{x}^k)\|]\}_{k \geq 0}$ does not converge to zero. Then, there exist $\delta > 0$ and two subsequences $\{\ell_t\}_{t \geq 0}$ and $\{u_t\}_{t \geq 0}$ such that $\ell_t < u_t \leq \ell_{t+1}$ and

$$\mathbb{E}[\|\Phi(\mathbf{x}^{\ell_t})\|] \geq 2\delta, \quad \mathbb{E}[\|\Phi(\mathbf{x}^{u_t})\|^a] \leq \delta^a, \quad \text{and} \quad \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] > \delta^a \quad (15)$$

for all $\ell_t < k < u_t$ and $t \in \mathbb{N}$. Combining (P.2) and (15) and applying Jensen's inequality ($a \geq 1$), this yields

$$\infty > \sum_{t=0}^{\infty} \sum_{k=\ell_t}^{u_t-1} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] \geq \sum_{t=0}^{\infty} \left[\sum_{k=\ell_t+1}^{u_t-1} \mu_k \delta^a + \mu_{\ell_t} (\mathbb{E}[\|\Phi(\mathbf{x}^{\ell_t})\|])^a \right] \geq \delta^a \sum_{t=0}^{\infty} \sum_{k=\ell_t}^{u_t-1} \mu_k,$$

which immediately implies $\beta_t := \sum_{k=\ell_t}^{u_t-1} \mu_k \rightarrow 0$. Since $\{\mu_k\}_{k \geq 0}$ is bounded, there exists $\bar{\mu}$ such that $\mu_k \leq \bar{\mu}$ for all k . For any $p \geq 1$, this further implies $\mu_k^p \leq \bar{\mu}^{p-1} \mu_k$ for all k . Using Hölder's and Jensen's inequality, $q \geq 1$, and (P.3), we have

$$\begin{aligned} \mathbb{E}[\|\mathbf{x}^{u_t} - \mathbf{x}^{\ell_t}\|] &\leq \sum_{k=\ell_t}^{u_t-1} \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|] \leq \left[\sum_{k=\ell_t}^{u_t-1} \mu_k \right]^{1-\frac{1}{q}} \left[\sum_{k=\ell_t}^{u_t-1} \mu_k^{-(q-1)} \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^q] \right]^{\frac{1}{q}} \\ &\leq \beta_t^{1-\frac{1}{q}} \left[\sum_{k=\ell_t}^{u_t-1} A \mu_k^{p_1-(q-1)} + B \mu_k^{p_2-(q-1)} \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^b] \right]^{\frac{1}{q}} \\ &\leq \beta_t^{1-\frac{1}{q}} \left[\sum_{k=\ell_t}^{u_t-1} A \bar{\mu}^{p_1-q} \mu_k + B \bar{\mu}^{p_2-q} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^b] \right]^{\frac{1}{q}} \\ &= \beta_t^{1-\frac{1}{q}} \left[A \bar{\mu}^{p_1-q} \beta_t + B \bar{\mu}^{p_2-q} \sum_{k=\ell_t}^{u_t-1} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^b] \right]^{\frac{1}{q}}, \end{aligned} \quad (16)$$

where we have utilized $p_1, p_2 \geq q$ in the third line. We first consider the case $a > b$. We can apply Jensen's and Hölder's inequality to obtain

$$\begin{aligned} \sum_{k=\ell_t}^{u_t-1} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^b] &\leq \sum_{k=\ell_t}^{u_t-1} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a]^{\frac{b}{a}} \\ &\leq \left[\sum_{k=\ell_t}^{u_t-1} \mu_k \right]^{\frac{a-b}{a}} \left[\sum_{k=\ell_t}^{u_t-1} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] \right]^{\frac{b}{a}} \\ &= \beta_t^{\frac{a-b}{a}} \left[\sum_{k=\ell_t}^{u_t-1} \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] \right]^{\frac{b}{a}}. \end{aligned} \quad (17)$$

Plugging (17) into (16) and invoking (P.2) and $\beta_t \rightarrow 0$, we have $\mathbb{E}[\|\mathbf{x}^{u_t} - \mathbf{x}^{\ell_t}\|] \rightarrow 0$ as $t \rightarrow \infty$. We now consider the case $a = b$. Let us introduce $E_k := \sum_{i=0}^{k-1} \mu_i \mathbb{E}[\|\Phi(\mathbf{x}^i)\|^a]$. By (P.2), the sequence $\{E_k\}_{k \geq 1}$ converges and hence, we have $E_{u_t} - E_{\ell_t} \rightarrow 0$ as $t \rightarrow \infty$. It then follows from (16) that

$$\mathbb{E}[\|\mathbf{x}^{u_t} - \mathbf{x}^{\ell_t}\|] \leq \beta_t^{1-\frac{1}{q}} [A\beta_t + B(E_{u_t} - E_{\ell_t})]^{\frac{1}{q}} \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

Together, this establishes $\mathbb{E}[\|\mathbf{x}^{u_t} - \mathbf{x}^{\ell_t}\|] \rightarrow 0$ as $t \rightarrow \infty$ in both cases. However, by the Lipschitz continuity of Φ in (P.1), $a \geq 1$ in (P.4), and the construction (15), we have

$$\delta \leq \mathbb{E}[\|\Phi(\mathbf{x}^{\ell_t})\|] - \mathbb{E}[\|\Phi(\mathbf{x}^{u_t})\|] \leq \mathbb{E}[\|\Phi(\mathbf{x}^{u_t}) - \Phi(\mathbf{x}^{\ell_t})\|] \leq L_{\Phi} \mathbb{E}[\|\mathbf{x}^{u_t} - \mathbf{x}^{\ell_t}\|]. \quad (18)$$

By letting $t \rightarrow \infty$ in (18), we get a contradiction. This concludes the proof of item (i).

Part II: Proof of item (ii). In order to control the stochastic behavior of the error terms \mathbf{A}_k and to establish the almost sure convergence of the sequence $\{\|\Phi(\mathbf{x}^k)\|\}_{k \geq 0}$, we will utilize several results from martingale theory.

Definition A.1 (Martingale). *Let $(\Omega, \mathcal{U}, \mathbb{P})$ be a probability space and $\{\mathcal{U}_k\}_{k \geq 0}$ a family of increasing sub- σ -fields of \mathcal{U} . A random process $\{M_k\}_{k \geq 0}$ defined on this probability space is said to be a*

martingale with respect to the family $\{\mathcal{U}_k\}_{k \geq 0}$ (or an $\{\mathcal{U}_k\}_{k \geq 0}$ -martingale) if each \mathbf{M}_k is integrable and \mathcal{U}_k -measurable and we have $\mathbb{E}[\mathbf{M}_{k+1} | \mathcal{U}_k] = \mathbf{M}_k$ a.s. for all k .

Next, we state a standard convergence theorem for vector-valued martingales, see, e.g., [43, Theorem 5.2.22 and Section 5.3] or [7, Theorem 5.14].

Theorem A.2 (Martingale Convergence Theorem). *Let $\{\mathbf{M}_k\}_{k \geq 0}$ be a given vector-valued $\{\mathcal{U}_k\}_{k \geq 0}$ -martingale as specified in Definition A.1. If $\sup_k \mathbb{E}[\|\mathbf{M}_k\|] < \infty$, then $\{\mathbf{M}_k\}_{k \geq 0}$ converges almost surely to an integrable random vector \mathbf{M} .*

Step (a): Analysis of the error terms $\{\mathbf{A}_k\}_{k \geq 0}$. Let us define $\mathbf{M}_k := \sum_{i=0}^{k-1} \mu_i^{p_1} \mathbf{A}_i$. Then, it holds that

$$\mathbb{E}[\mathbf{M}_{k+1} | \mathcal{F}_k] = \sum_{i=0}^{k-1} \mu_i^{p_1} \mathbf{A}_i + \mu_k^{p_1} \mathbb{E}[\mathbf{A}_k | \mathcal{F}_k] = \mathbf{M}_k,$$

i.e., $\{\mathbf{M}_k\}_{k \geq 0}$ defines a martingale adapted to the filtration $\{\mathcal{F}_k\}_{k \geq 0}$. In addition, inductively, we obtain

$$\begin{aligned} \mathbb{E}[\|\mathbf{M}_k\|^2] &= \mathbb{E}[\|\mathbf{M}_{k-1}\|^2] + 2\mu_{k-1}^{p_1} \mathbb{E}[\langle \mathbf{M}_{k-1}, \mathbf{A}_{k-1} \rangle] + \mu_{k-1}^{2p_1} \mathbb{E}[\|\mathbf{A}_{k-1}\|^2] \\ &= \dots = \sum_{i=0}^{k-1} \mu_i^{2p_1} \mathbb{E}[\|\mathbf{A}_i\|^2] \leq \sum_{i=0}^{k-1} \mu_i^{2p_1} \mathbb{E}[\|\mathbf{A}_i\|^q]^{\frac{2}{q}} \leq A^{\frac{2}{q}} \sum_{i=0}^{k-1} \mu_i^{2p_1}, \end{aligned}$$

where we utilized Jensen's inequality, the concavity of the mapping $x \mapsto x^{\frac{2}{q}}$, and $\mathbb{E}[\|\mathbf{A}_i\|^q] \leq A$ in (P.3'). Hence, by (P.4') and Jensen's inequality, we can infer $\sup_k \mathbb{E}[\|\mathbf{M}_k\|] \leq A^{\frac{1}{q}} [\sum_{i=0}^{\infty} \mu_i^{2p_1}]^{\frac{1}{2}} < \infty$. Theorem A.2 then implies that $\{\mathbf{M}_k\}_{k \geq 0}$ converges almost surely to some integrable random vector \mathbf{M} .

Next, we establish almost sure convergence of $\{\|\Phi(\mathbf{x}^k)\|\}_{k \geq 0}$. Our derivation generally mimics the proof of item (i), but uses sample-based arguments.

Step (b): Almost sure convergence. First, applying the monotone convergence theorem to (P.2) gives $\mathbb{E}[\sum_{k=0}^{\infty} \mu_k \|\Phi(\mathbf{x}^k)\|^a] = \lim_{T \rightarrow \infty} \sum_{k=0}^T \mu_k \mathbb{E}[\|\Phi(\mathbf{x}^k)\|^a] < \infty$, which immediately implies $\sum_{k=0}^{\infty} \mu_k \|\Phi(\mathbf{x}^k)\|^a < \infty$ almost surely. This, together with $\sum_{k=0}^{\infty} \mu_k = \infty$, yields $\liminf_{k \rightarrow \infty} \|\Phi(\mathbf{x}^k)\| = 0$ almost surely. We now consider an arbitrary sample $\omega \in \mathcal{M}$ where

$$\begin{aligned} \mathcal{M} := \left\{ \omega : \sum_{k=0}^{\infty} \mu_k \|\Phi(\mathbf{x}^k(\omega))\|^a < \infty, \quad \lim_{k \rightarrow \infty} \mathbf{M}_k(\omega) = \mathbf{M}(\omega), \right. \\ \left. \text{and } \limsup_{k \rightarrow \infty} \frac{\|\mathbf{B}_k(\omega)\|^q}{1 + \|\Phi(\mathbf{x}^k(\omega))\|^b} < \infty \right\}. \end{aligned}$$

Our preceding discussion implies that the event \mathcal{M} occurs with probability 1 and it holds that $\liminf_{k \rightarrow \infty} \|\Phi(\mathbf{x}^k(\omega))\| = 0$. Let us assume that $\{\|\Phi(\mathbf{x}^k(\omega))\|\}_{k \geq 0}$ does not converge to zero. Then, there exist $\delta > 0$ and two subsequences $\{\ell_t\}_{t \geq 0}$ and $\{u_t\}_{t \geq 0}$ such that $\ell_t < u_t \leq \ell_{t+1}$ and

$$\|\Phi(\mathbf{x}^{\ell_t}(\omega))\| \geq 2\delta, \quad \|\Phi(\mathbf{x}^{u_t}(\omega))\| \leq \delta, \quad \text{and} \quad \|\Phi(\mathbf{x}^k(\omega))\| > \delta \quad (19)$$

for all $\ell_t < k < u_t$ and $t \in \mathbb{N}$. (Notice that in contrast to the proof of item (i), the sequences $\{\ell_t\}_{t \geq 0}$, $\{u_t\}_{t \geq 0}$, and δ will now generally depend on the selected sample ω). Due to $\omega \in \mathcal{M}$, this yields

$$\infty > \sum_{t=0}^{\infty} \sum_{k=\ell_t}^{u_t-1} \mu_k \|\Phi(\mathbf{x}^k(\omega))\|^a \geq \delta^a \sum_{t=0}^{\infty} \sum_{k=\ell_t}^{u_t-1} \mu_k,$$

which again implies $\beta_t := \sum_{k=\ell_t}^{u_t-1} \mu_k \rightarrow 0$. Furthermore, there exist $K \in \mathbb{N}$ and $B \in \mathbb{R}$ such that $\|\mathbf{B}_k(\omega)\| \leq B(1 + \|\Phi(\mathbf{x}^k(\omega))\|^b)^{1/q}$ for all $k \geq K$. We first consider the case $qa > b$. It follows from (P.3') that

$$\begin{aligned} \|\mathbf{x}^{u_t}(\omega) - \mathbf{x}^{\ell_t}(\omega)\| &\leq B \sum_{k=\ell_t}^{u_t-1} \mu_k^{p_2} (1 + \|\Phi(\mathbf{x}^k(\omega))\|^b)^{\frac{1}{q}} + \left\| \sum_{k=\ell_t}^{u_t-1} \mu_k^{p_1} \mathbf{A}_k(\omega) \right\| \\ &\leq B \sum_{k=\ell_t}^{u_t-1} \mu_k + B \sum_{k=\ell_t}^{u_t-1} \mu_k \|\Phi(\mathbf{x}^k(\omega))\|^{\frac{b}{q}} + \left\| \sum_{k=\ell_t}^{u_t-1} \mu_k^{p_1} \mathbf{A}_k(\omega) \right\| \\ &\leq B\beta_t + B \left(\sum_{k=\ell_t}^{u_t-1} \mu_k \right)^{1 - \frac{b}{qa}} \left(\sum_{k=\ell_t}^{u_t-1} \mu_k \|\Phi(\mathbf{x}^k(\omega))\|^a \right)^{\frac{b}{qa}} \\ &\quad + \|\mathbf{M}_{u_t}(\omega) - \mathbf{M}_{\ell_t}(\omega)\| \end{aligned}$$

$$\leq B\beta_t + B\beta_t^{1-\frac{b}{qa}} \left(\sum_{k=0}^{\infty} \mu_k \|\Phi(\mathbf{x}^k(\omega))\|^a \right)^{\frac{b}{qa}} + \|\mathbf{M}_{u_t}(\omega) - \mathbf{M}_{\ell_t}(\omega)\|$$

for all t sufficiently large, where we used the subadditivity of $x \mapsto x^{\frac{1}{q}}$, $p_2 \geq 1$, and $\mu_k \rightarrow 0$ in the second inequality, and Hölder's inequality in the third inequality. In the case $qa = b$, let us introduce $\mathbf{E}_k(\omega) := \sum_{i=0}^{k-1} \mu_k \|\Phi(\mathbf{x}^i(\omega))\|^a$. Then, it follows

$$\|\mathbf{x}^{u_t}(\omega) - \mathbf{x}^{\ell_t}(\omega)\| \leq B\beta_t + B[\mathbf{E}_{u_t}(\omega) - \mathbf{E}_{\ell_t}(\omega)] + \|\mathbf{M}_{u_t}(\omega) - \mathbf{M}_{\ell_t}(\omega)\|$$

for all t sufficiently large. Due to $\beta_t \rightarrow 0$, $\mathbf{E}_k(\omega) \rightarrow \mathbf{E}(\omega) := \sum_{k=0}^{\infty} \mu_k \|\Phi(\mathbf{x}^k(\omega))\|^a$, and $\mathbf{M}_k(\omega) \rightarrow \mathbf{M}(\omega)$, we can infer $\|\mathbf{x}^{u_t}(\omega) - \mathbf{x}^{\ell_t}(\omega)\| \rightarrow 0$ in both cases. As before, the Lipschitz continuity of Φ in (P.1) yields the contradiction

$$\delta \leq \|\Phi(\mathbf{x}^{\ell_t}(\omega))\| - \|\Phi(\mathbf{x}^{u_t}(\omega))\| \leq \|\Phi(\mathbf{x}^{u_t}(\omega)) - \Phi(\mathbf{x}^{\ell_t}(\omega))\| \leq L_{\Phi} \|\mathbf{x}^{u_t}(\omega) - \mathbf{x}^{\ell_t}(\omega)\| \rightarrow 0.$$

Thus, for all $\omega \in \mathcal{M}$, we have $\lim_{k \rightarrow \infty} \|\Phi(\mathbf{x}^k(\omega))\| = 0$. Since the event \mathcal{M} occurs with probability 1, this concludes the proof of item (ii).

B The supermartingale convergence theorem

The following well-known and celebrated convergence theorem for supermartingale-type stochastic processes is due to Robbins and Siegmund [37].

Theorem B.1 (Supermartingale Convergence Theorem). *Let $\{y_k\}_{k \geq 0}$, $\{p_k\}_{k \geq 0}$, and $\{q_k\}_{k \geq 0}$ be sequences of nonnegative integrable random variables adapted to a filtration $\{\mathcal{U}_k\}_{k \geq 0}$. Furthermore, let $\{\beta_k\}_{k \geq 0} \subseteq \mathbb{R}_+$ be given with $\sum_{k=0}^{\infty} \beta_k < \infty$ and assume that we have*

$$\mathbb{E}[y_{k+1} \mid \mathcal{U}_k] \leq (1 + \beta_k)y_k - p_k + q_k$$

for all k and $\sum_{k=0}^{\infty} q_k < \infty$ almost surely. Then, it holds that

- (a) *If $\sum_{k=0}^{\infty} \mathbb{E}[q_k] < \infty$, then the sequence $\{\mathbb{E}[y_k]\}_{k \geq 0}$ converges to a finite number \bar{y} and we have $\sum_{k=0}^{\infty} \mathbb{E}[p_k] < \infty$.*
- (b) *$\{y_k\}_{k \geq 0}$ almost surely converges to a nonnegative finite random variable y and it follows $\sum_{k=0}^{\infty} p_k < \infty$ almost surely.*

C Proofs for Subsection 3.1

C.1 Derivation of (6)

Using the Lipschitz continuity of ∇f and the descent lemma, we obtain

$$\begin{aligned} f(\mathbf{x}^{k+1}) &\leq f(\mathbf{x}^k) + \langle \nabla f(\mathbf{x}^k), \mathbf{x}^{k+1} - \mathbf{x}^k \rangle + \frac{L}{2} \|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \\ &= f(\mathbf{x}^k) + \alpha_k \langle \nabla f(\mathbf{x}^k), \nabla f(\mathbf{x}^k) - \mathbf{g}^k \rangle - \alpha_k \|\nabla f(\mathbf{x}^k)\|^2 + \frac{L\alpha_k^2}{2} \|\mathbf{g}^k\|^2 \\ &= f(\mathbf{x}^k) - \alpha_k \left(1 - \frac{L\alpha_k}{2} \right) \|\nabla f(\mathbf{x}^k)\|^2 + \alpha_k (1 - L\alpha_k) \langle \nabla f(\mathbf{x}^k), \nabla f(\mathbf{x}^k) - \mathbf{g}^k \rangle \\ &\quad + \frac{L\alpha_k^2}{2} \|\nabla f(\mathbf{x}^k) - \mathbf{g}^k\|^2. \end{aligned}$$

Taking conditional expectation gives (6).

D Results and proofs for Subsection 3.2

We now provide more detailed algorithmic procedure and motivations for RR. First, let us define the set of all possible permutations of $\{1, 2, \dots, N\}$ as

$$\Lambda := \{\sigma : \sigma \text{ is a permutation of } \{1, 2, \dots, N\}\}.$$

At each iteration k , a permutation σ^{k+1} is generated according to an i.i.d. uniform distribution over Λ . Then, RR updates \mathbf{x}^k to \mathbf{x}^{k+1} through N consecutive gradient descent-type steps by using the

components $\{f(\cdot, \sigma_1^{k+1}), \dots, f(\cdot, \sigma_N^{k+1})\}$ sequentially, where σ_i^{k+1} represents the i -th element of σ^{k+1} . In each step, only one component $f(\cdot, \sigma_i^{k+1})$ is selected for updating. To be more specific, this method starts with $\tilde{\mathbf{x}}_0^k = \mathbf{x}^k$ and then uses $f(\cdot, \sigma_i^{k+1})$ to update $\tilde{\mathbf{x}}_i^k$ as

$$\tilde{\mathbf{x}}_i^k = \tilde{\mathbf{x}}_{i-1}^k - \alpha_k \nabla f(\tilde{\mathbf{x}}_{i-1}^k, \sigma_i^{k+1})$$

for $i = 1, 2, \dots, N$, resulting in $\mathbf{x}^{k+1} = \tilde{\mathbf{x}}_N^k$.

RR is used in a vast variety of engineering fields. Most notably, RR is extensively applied in practice for training deep neural networks; see, e.g., [20, 30, 32] and the references therein.

D.1 A bound on the step length

We need the following lemma to for our later analysis.

Lemma D.1. *We have the following estimate:*

$$\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^q \leq \bar{A}G(\mathbf{x}^0)^{\frac{q}{2}}\alpha_k^q \quad (20)$$

for all k sufficiently large and some positive constants \bar{A} and $G(\mathbf{x}^0)^{\frac{q}{2}}$.

Proof. Using $|\sum_{i=1}^N a_i|^q \leq N^{q-1} \sum_{i=1}^N |a_i|^q$, (7), and $\nabla f(\mathbf{x}^k) = \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{x}^k, \sigma_i^{k+1})$, we have

$$\begin{aligned} \|\mathbf{x}^{k+1} - \mathbf{x}^k\|^q &= \alpha_k^q \left\| \sum_{i=1}^N \nabla f(\tilde{\mathbf{x}}_{i-1}^k, \sigma_i^{k+1}) \right\|^q \\ &\leq 2^{q-1} N^q \alpha_k^q \|\nabla f(\mathbf{x}^k)\|^q + 2^{q-1} \alpha_k^q \left\| \sum_{i=1}^N [\nabla f(\tilde{\mathbf{x}}_{i-1}^k, \sigma_i^{k+1}) - \nabla f(\mathbf{x}^k, \sigma_i^{k+1})] \right\|^q \\ &\leq 2^{q-1} N^q \alpha_k^q \|\nabla f(\mathbf{x}^k)\|^q + (2N)^{q-1} L^q \alpha_k^q \cdot \sum_{i=1}^N \|\tilde{\mathbf{x}}_{i-1}^k - \mathbf{x}^k\|^q. \end{aligned}$$

Setting $V_k := \sum_{i=1}^N \|\tilde{\mathbf{x}}_{i-1}^k - \mathbf{x}^k\|^q$, we recursively obtain

$$\begin{aligned} V_k &= \alpha_k^q \sum_{i=2}^N \left\| \sum_{j=1}^{i-1} \nabla f(\tilde{\mathbf{x}}_{j-1}^k, \sigma_j^{k+1}) \right\|^q \\ &\leq 2^{q-1} \alpha_k^q \sum_{i=2}^N \left[\left\| \sum_{j=1}^{i-1} [\nabla f(\tilde{\mathbf{x}}_{j-1}^k, \sigma_j^{k+1}) - \nabla f(\mathbf{x}^k, \sigma_j^{k+1})] \right\|^q + \left\| \sum_{j=1}^{i-1} \nabla f(\mathbf{x}^k, \sigma_j^{k+1}) \right\|^q \right] \\ &\leq 2^{q-1} \alpha_k^q \sum_{i=2}^N \left[L^q (i-1)^{q-1} \sum_{j=1}^{i-1} \|\tilde{\mathbf{x}}_{j-1}^k - \mathbf{x}^k\|^q + \left[(i-1) \sum_{j=1}^{i-1} \|\nabla f(\mathbf{x}^k, \sigma_j^{k+1})\|^2 \right]^{\frac{q}{2}} \right] \\ &\leq 2^{q-1} L^q \alpha_k^q \left(\sum_{i=1}^{N-1} i^{q-1} \right) V_k + 2^{\frac{3q}{2}-1} L^{\frac{q}{2}} \alpha_k^q \sum_{i=2}^N (i-1)^{\frac{q}{2}} \left[\sum_{j=1}^{i-1} (f(\mathbf{x}^k, \sigma_j^{k+1}) - \bar{f}) \right]^{\frac{q}{2}} \\ &\leq 2^{q-1} L^q \left[\frac{N^q}{q} \right] \alpha_k^q V_k + 2^{\frac{3q-2}{2}} (LN)^{\frac{q}{2}} \left[\frac{2N^{\frac{q}{2}+1}}{q+2} \right] \alpha_k^q (f(\mathbf{x}^k) - \bar{f})^{\frac{q}{2}} \end{aligned}$$

where we applied the estimate $\|\nabla f(\mathbf{x}, i)\|^2 \leq 2L(f(\mathbf{x}, i) - \bar{f})$ (see also (24) for comparison). Clearly, this establishes $V_k = \mathcal{O}(\alpha_k^q (f(\mathbf{x}^k) - \bar{f})^{\frac{q}{2}})$. Furthermore, following the proof of [25, Lemma 3.2], we have $f(\mathbf{x}^k) - \bar{f} \leq G(\mathbf{x}^0)$ where $G(\mathbf{x}^0) = (f(\mathbf{x}^0) - \bar{f}) \exp(2L^3 N^3 \sum_{j=0}^{\infty} \alpha_j^3)$. Hence, using $\|\nabla f(\mathbf{x})\|^2 \leq 2L(f(\mathbf{x}) - \bar{f})$ and (B.2), there exists a constant \bar{A} such that (20) holds for all sufficiently large k . \square

D.2 Proof of Corollary 3.2

We derive the convergence of RR below by setting $\Phi \equiv \nabla f$ and $\mu_k \equiv \alpha_k$.

Phase I: Verifying (P.1)–(P.2). (B.1) verifies condition (P.1) with $L_\Phi = L$. Towards verifying (P.2), note that under the above assumptions for RR, [25, Lemma 3.1] establishes the recursion:

$$f(\mathbf{x}^{k+1}) - \bar{f} \leq (1 + 2L^3 N^3 \alpha_k^3) [f(\mathbf{x}^k) - \bar{f}] - \frac{N\alpha_k}{2} \|\nabla f(\mathbf{x}^k)\|^2 - \frac{1 - LN\alpha_k}{2N\alpha_k} \|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \quad (21)$$

for all k as long as $\alpha_k < 1/(\sqrt{2}LN)$. (This always holds for large enough k as $\alpha_k \rightarrow 0$). Taking total expectation and applying [Theorem B.1](#) provides $\mathbb{E}[f(\mathbf{x}^k) - \bar{f}] \rightarrow F$ and $\sum_{k=0}^{\infty} \alpha_k \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|^2] < \infty$. This verifies condition [\(P.2\)](#) with $a = 2$.

Phase II: Verifying [\(P.3\)](#)–[\(P.4\)](#) for showing expected convergence. We can infer from [Lemma D.1](#) that [\(P.3\)](#) holds with arbitrary $q \geq 1$ and $q \in \mathbb{N}$, $A = \bar{A}G(\mathbf{x}^0)^{\frac{q}{2}}$, $p_1 = q$, and $B = 0$. [\(P.4\)](#) can be easily verified by these parameter choices and [\(B.2\)](#). Thus, applying [Theorem 2.1](#) gives $\mathbb{E}[\|\nabla f(\mathbf{x}^k)\|] \rightarrow 0$.

Phase III: Verifying [\(P.3'\)](#)–[\(P.4'\)](#) for showing almost sure convergence. According to the update [\(7\)](#), we can let $A_k \equiv 0$, $p_2 = 1$, $B_k = \sum_{i=1}^N \nabla f(\tilde{\mathbf{x}}_{i-1}^k, \sigma_i^{k+1})$, $q = 2$, and $b = 2$, in [\(P.3'\)](#). We have

$$\limsup_{k \rightarrow \infty} \frac{\|B_k\|^2}{1 + \|\nabla f(\mathbf{x}^k)\|^2} \leq \limsup_{k \rightarrow \infty} \frac{2\|B_k - N\nabla f(\mathbf{x}^k)\|^2 + 2N^2\|\nabla f(\mathbf{x}^k)\|^2}{1 + \|\nabla f(\mathbf{x}^k)\|^2} < \infty,$$

since $\|B_k - N\nabla f(\mathbf{x}^k)\|^2 \leq \mathcal{O}(\alpha_k^2)$ — as shown in [Appendix D.1](#) — which converges to 0 as $k \rightarrow \infty$. Condition [\(P.3'\)](#) is verified. Assumption [\(B.2\)](#) and the previous parameter choices verify condition [\(P.4'\)](#). Hence, we can apply [Theorem 2.1](#) to derive $\|\nabla f(\mathbf{x}^k)\| \rightarrow 0$ almost surely.

Summarizing the above results yields [Corollary 3.2](#).

E Results and proofs for [Subsection 3.3](#)

E.1 Weakly convex functions

A function $h : \mathbb{R}^n \rightarrow (-\infty, \infty]$ is said to be τ -weakly convex if

$$h + \frac{\tau}{2} \|\cdot\|^2$$

is convex (on its effective domain $\text{dom } h = \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) < \infty\}$). The class of weakly convex functions covers many important nonsmooth nonconvex problems. For instance, any function that has the composite form

$$h(\mathbf{x}) = r(c(\mathbf{x}))$$

with r being convex Lipschitz continuous and the Jacobian of c being Lipschitz continuous is weakly convex. We refer to [\[11\]](#) for more discussions.

If h is τ -weakly convex, proper, and lower semicontinuous, then the proximity operator $\text{prox}_{\alpha h} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, given by

$$\text{prox}_{\alpha h}(\mathbf{x}) := \underset{\mathbf{y} \in \mathbb{R}^n}{\text{argmin}} h(\mathbf{y}) + \frac{1}{2\alpha} \|\mathbf{x} - \mathbf{y}\|^2,$$

is a well-defined function for all $\alpha \in (0, \tau^{-1})$ and Lipschitz continuous with constant $(1 - \alpha\tau)^{-1}$, see, e.g., [\[39, Proposition 12.19\]](#) or [\[21, Proposition 3.3\]](#). The τ -weak convexity is equivalent to

$$h(\mathbf{y}) \geq h(\mathbf{x}) + \langle \mathbf{s}, \mathbf{y} - \mathbf{x} \rangle - \frac{\tau}{2} \|\mathbf{x} - \mathbf{y}\|^2, \quad \forall \mathbf{x}, \mathbf{y}, \quad \forall \mathbf{s} \in \partial h(\mathbf{x}) \quad (22)$$

see, e.g., [\[44, 11\]](#).

E.2 Examples of weakly convex and Lipschitz continuous regularizers

In this section, we discuss several common nonconvex regularizers that can be shown to be weakly convex and Lipschitz continuous. The functions presented here have been mentioned in [Remark 3.3](#).

The minimax concave penalty (MCP), introduced in [\[45\]](#), is the parametrized function $\varphi_{\lambda, \theta} : \mathbb{R} \rightarrow \mathbb{R}_+$ given by

$$\varphi_{\lambda, \theta}(x) := \begin{cases} \lambda|x| - \frac{x^2}{2\theta} & \text{if } |x| \leq \theta\lambda, \\ \frac{\theta\lambda^2}{2} & \text{otherwise,} \end{cases}$$

where $\lambda, \theta > 0$ are two positive parameters. This function is θ^{-1} -weakly convex and smooth for $x \neq 0$. Discussing the subdifferential $\partial\varphi_{\lambda, \theta}(x)$ and using [\[39, Theorem 9.13\]](#), it can be shown that $\varphi_{\lambda, \theta}$ is Lipschitz continuous on \mathbb{R} with modulus λ .

The smoothly clipped absolute deviation (SCAD) [15] is defined by

$$\varphi_{\lambda,\theta}(x) := \begin{cases} \lambda|x| & \text{if } |x| \leq \lambda, \\ \frac{-x^2+2\theta\lambda|x|-\lambda^2}{2(\theta-1)} & \text{if } \lambda < |x| \leq \theta\lambda, \\ \frac{(\theta+1)\lambda^2}{2} & \text{if } |x| > \theta\lambda, \end{cases}$$

where $\lambda > 0$ and $\theta > 2$ are given parameters. The SCAD function is $(\theta - 1)^{-1}$ -weakly convex and Lipschitz continuous with modulus λ .

The student-t loss function is given by $\varphi_\theta(x) := \frac{\theta^2}{2} \log(1 + \theta^{-2}x^2)$ for some $\theta \neq 0$. The first- and second-order derivative of φ_θ can be calculated as follows:

$$\varphi'_\theta(x) = \frac{\theta^2 x}{\theta^2 + x^2} \quad \text{and} \quad \varphi''_\theta(x) = \frac{\theta^2(\theta^2 - x^2)}{(\theta^2 + x^2)^2}.$$

Some simple computations then show that φ_θ is $\frac{1}{8}$ -weakly convex and Lipschitz continuous with modulus $|\theta|/2$. Additional examples can be found in [4].

E.3 Equivalent stationarity measures

We now first show that the two stationarity measures $\mathbf{x} \mapsto \|F_{\text{nat}}^1(\mathbf{x})\|$ and $\mathbf{x} \mapsto \|\nabla \text{env}_{\theta\psi}(\mathbf{x})\|$ are equivalent. This can be used to verify [Corollary 3.6](#).

Lemma E.1. *Suppose that the conditions (C.1) and (C.2) are satisfied and let $\theta \in (0, [3L + \tau]^{-1})$ and $\mathbf{x} \in \mathbb{R}^n$ be given. Then, we have*

$$(1 - [3L + \tau]\theta)\gamma\theta^{-2}\|F_{\text{nat}}(\mathbf{x})\| \leq \|\nabla \text{env}_{\theta\psi}(\mathbf{x})\| \leq (1 + [L - \tau]\theta)(\gamma + \tau)\theta^{-2}\|F_{\text{nat}}(\mathbf{x})\|,$$

where $\gamma = \theta/(1 - [L + \tau]\theta)$ and $F_{\text{nat}} := F_{\text{nat}}^1$.

Proof. For every fixed $\mathbf{x} \in \mathbb{R}^n$, let us define $\psi_{\mathbf{x}}(\mathbf{y}) := \psi(\mathbf{y}) + \frac{L+\tau}{2}\|\mathbf{x} - \mathbf{y}\|^2$ and $\varphi_{\mathbf{x}}(\mathbf{y}) := \varphi(\mathbf{y}) + \frac{\tau}{2}\|\mathbf{x} - \mathbf{y}\|^2$. Then, setting $\gamma = \theta/(1 - [L + \tau]\theta)$, we may write

$$\text{prox}_{\theta\psi}(\mathbf{x}) = \text{prox}_{\gamma\varphi_{\mathbf{x}}}(\mathbf{x}) \quad \text{and} \quad \nabla \text{env}_{\theta\psi}(\mathbf{x}) = \theta^{-1}\gamma\nabla \text{env}_{\gamma\varphi_{\mathbf{x}}}(\mathbf{x}).$$

Applying [13, Theorem 3.5] with $G \equiv \partial\varphi_{\mathbf{x}}$, $\Phi \equiv \partial\psi_{\mathbf{x}}$, $F \equiv \nabla f + L(\cdot - \mathbf{x})$, $t \equiv \gamma$, and $\beta \equiv 2L$, we can establish the following estimates for all $\mathbf{x} \in \mathbb{R}^n$:

$$(1 - 2L\gamma)\theta^{-1}\|\mathbf{x} - \text{prox}_{\gamma\varphi_{\mathbf{x}}}(\mathbf{x} - \gamma F(\mathbf{x}))\| \leq \|\nabla \text{env}_{\theta\psi}(\mathbf{x})\|$$

and $\|\nabla \text{env}_{\theta\psi}(\mathbf{x})\| \leq (1 + 2L\gamma)\theta^{-1}\|\mathbf{x} - \text{prox}_{\gamma\varphi_{\mathbf{x}}}(\mathbf{x} - \gamma F(\mathbf{x}))\|$. Moreover, it holds that

$$\begin{aligned} \text{prox}_{\gamma\varphi_{\mathbf{x}}}(\mathbf{x} - \gamma F(\mathbf{x})) &= \text{prox}_{\gamma\varphi_{\mathbf{x}}}(\mathbf{x} - \gamma\nabla f(\mathbf{x})) \\ &= \underset{\mathbf{y} \in \mathbb{R}^n}{\text{argmin}} \varphi(\mathbf{y}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{1}{2} \left[\frac{1}{\gamma} + \tau \right] \|\mathbf{y} - \mathbf{x}\|^2 \\ &= \text{prox}_{\frac{\gamma}{\gamma+\tau}\varphi}(\mathbf{x} - \frac{\gamma}{\gamma+\tau}\nabla f(\mathbf{x})) \end{aligned}$$

In addition, as shown in [31, Lemma 2], the functions $\delta \mapsto \|F_{\text{nat}}^{1/\delta}(\mathbf{x})\|$ and $\delta \mapsto \|F_{\text{nat}}^{1/\delta}(\mathbf{x})\|/\delta$ are decreasing and increasing in δ , respectively. Hence, we can infer

$$\min\{1, \lambda_2/\lambda_1\}\|F_{\text{nat}}^{\lambda_2}(\mathbf{x})\| \leq \|F_{\text{nat}}^{\lambda_1}(\mathbf{x})\| \leq \max\{1, \lambda_2/\lambda_1\}\|F_{\text{nat}}^{\lambda_2}(\mathbf{x})\| \quad (23)$$

for all $\lambda_1, \lambda_2 > 0$ and $\mathbf{x} \in \mathbb{R}^n$. Choosing $\lambda_1 = \gamma/(\gamma + \tau)$ and $\lambda_2 = 1$ and noticing $1 - 2L\gamma = (1 - [3L + \tau]\theta)\theta^{-1} > 0$, we can conclude the proof of [Lemma E.1](#). \square

E.4 Proof of Lemma 3.4

Let us again note that the τ -weak convexity of φ — as stated in (C.2) — implies Lipschitz continuity of the proximity operator $\text{prox}_{\theta\varphi}$ with modulus $(1 - \theta\tau)^{-1}$. In the following results, we first bound the term $\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 = \|\mathbf{x}^{k+1} - \text{prox}_{\theta\psi}(\mathbf{x}^k)\|^2$.

Lemma E.2. *Suppose that (C.1) and (C.2) are satisfied and let $\theta \in (0, [L + \tau]^{-1})$, $\alpha > 0$, $\mathbf{x} \in \mathbb{R}^n$, and $\bar{\mathbf{x}} = \text{prox}_{\theta\psi}(\mathbf{x})$ be given. Then, it holds that*

$$\bar{\mathbf{x}} = \text{prox}_{\alpha\varphi}(\bar{\mathbf{x}} - \alpha\nabla f(\bar{\mathbf{x}}) - \alpha\theta^{-1}[\bar{\mathbf{x}} - \mathbf{x}]).$$

Proof. This is Lemma 3.2 in [11]. □

Our second result is analogous to Lemma 3.3 in [11].

Lemma E.3. *Suppose that the conditions (C.1), (C.2), and (C.4) are satisfied and let $\theta \in (0, [L + \tau]^{-1})$ be given. Defining $\bar{\mathbf{x}}^k = \text{prox}_{\theta\psi}(\mathbf{x}^k)$, $k \in \mathbb{N}$, it holds that*

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \mid \mathcal{F}_k] \leq (1 - \alpha_k\theta_k)^2\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \frac{\alpha_k^2\sigma_k^2}{(1 - \alpha_k\tau)^2} \quad \text{almost surely} \quad \forall k \in \mathbb{N},$$

where $\theta_k := (\theta^{-1} - [L + \tau])/(1 - \alpha_k\tau)$ and $\sigma_k^2 := \mathbb{E}[\|\mathbf{g}^k - \nabla f(\mathbf{x}^k)\|^2 \mid \mathcal{F}_k]$.

Proof. Invoking Lemma E.2 and the Lipschitz continuity of $\text{prox}_{\alpha_k\varphi}$ and ∇f , it follows

$$\begin{aligned} & (1 - \alpha_k\tau)^2\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \\ &= (1 - \alpha_k\tau)^2\|\text{prox}_{\alpha_k\varphi}(\mathbf{x}^k - \alpha_k\mathbf{g}^k) - \text{prox}_{\alpha_k\varphi}(\bar{\mathbf{x}}^k - \alpha_k\nabla f(\bar{\mathbf{x}}^k) - \alpha_k\theta^{-1}[\bar{\mathbf{x}}^k - \mathbf{x}^k])\|^2 \\ &\leq \|(1 - \alpha_k\theta^{-1})[\bar{\mathbf{x}}^k - \mathbf{x}^k] + \alpha_k[\mathbf{g}^k - \nabla f(\bar{\mathbf{x}}^k)]\|^2 \\ &= (1 - \alpha_k\theta^{-1})^2\|\bar{\mathbf{x}}^k - \mathbf{x}^k\|^2 + 2\alpha_k(1 - \alpha_k\theta^{-1})\langle \bar{\mathbf{x}}^k - \mathbf{x}^k, \mathbf{g}^k - \nabla f(\bar{\mathbf{x}}^k) \rangle \\ &\quad + \alpha_k^2[\|\mathbf{g}^k - \nabla f(\mathbf{x}^k)\|^2 + 2\langle \mathbf{g}^k - \nabla f(\mathbf{x}^k), \nabla f(\mathbf{x}^k) - \nabla f(\bar{\mathbf{x}}^k) \rangle + \|\nabla f(\mathbf{x}^k) - \nabla f(\bar{\mathbf{x}}^k)\|^2] \\ &\leq (1 - \alpha_k[\theta^{-1} - L])^2\|\bar{\mathbf{x}}^k - \mathbf{x}^k\|^2 + 2\alpha_k(1 - \alpha_k\theta^{-1})\langle \bar{\mathbf{x}}^k - \mathbf{x}^k, \mathbf{g}^k - \nabla f(\mathbf{x}^k) \rangle \\ &\quad + 2\alpha_k^2\langle \mathbf{g}^k - \nabla f(\mathbf{x}^k), \nabla f(\mathbf{x}^k) - \nabla f(\bar{\mathbf{x}}^k) \rangle + \alpha_k^2\|\mathbf{g}^k - \nabla f(\mathbf{x}^k)\|^2. \end{aligned}$$

Taking conditional expectation, using (C.4) and $\mathbf{x}^k, \bar{\mathbf{x}}^k \in \mathcal{F}_k$, we obtain

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \mid \mathcal{F}_k] \leq (1 - \alpha_k\theta_k)^2\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \frac{\alpha_k^2\sigma_k^2}{(1 - \alpha_k\tau)^2}$$

almost surely, where $\theta_k = (\theta^{-1} - [L + \tau])/(1 - \alpha_k\tau)$. □

By (C.4), the stochastic error term σ_k^2 is bounded by $C[f(\mathbf{x}^k) - \bar{f}] + D$ almost surely. In order to proceed, we need to link the function values “ $f(\mathbf{x}^k) - \bar{f}$ ” and “ $\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}$ ” where $\bar{\psi} = \bar{f} + \bar{\varphi}$. The following lemma precisely establishes such a connection under assumption (C.3).

Lemma E.4. *Suppose that the conditions (C.1)–(C.3) are satisfied and let $\theta \in (0, [\frac{4}{3}L + \tau]^{-1})$ be given. Then, it holds that*

$$f(\mathbf{x}) - \bar{f} \leq 2[\text{env}_{\theta\psi}(\mathbf{x}) - \bar{\psi}] + L_\varphi^2\theta \quad \forall \mathbf{x} \in \text{dom } \varphi.$$

Proof. Notice that the Lipschitz continuity of ∇f and assumption (C.2) imply $f(\mathbf{x} - L^{-1}\nabla f(\mathbf{x})) - f(\mathbf{x}) \leq -\frac{1}{2L}\|\nabla f(\mathbf{x})\|^2$ and

$$\|\nabla f(\mathbf{x})\|^2 \leq 2L[f(\mathbf{x}) - \bar{f}] = 2L[\psi(\mathbf{x}) - \varphi(\mathbf{x}) - \bar{f}] \leq 2L[\psi(\mathbf{x}) - \bar{\psi}] \quad (24)$$

for all $\mathbf{x} \in \text{dom } \varphi$. Using (C.3), the Lipschitz continuity of ∇f , Young’s inequality (twice), and $\theta < \frac{3}{4}L^{-1}$ and setting $\bar{\mathbf{x}} = \text{prox}_{\theta\psi}(\mathbf{x}) \in \text{dom } \varphi$, this yields

$$\begin{aligned} \psi(\mathbf{x}) - \bar{\psi} &\leq \text{env}_{\theta\psi}(\mathbf{x}) - \bar{\psi} + \langle \nabla f(\bar{\mathbf{x}}), \mathbf{x} - \bar{\mathbf{x}} \rangle + \frac{1}{2}[L - \theta^{-1}]\|\mathbf{x} - \bar{\mathbf{x}}\|^2 + \varphi(\mathbf{x}) - \varphi(\bar{\mathbf{x}}) \\ &\leq \text{env}_{\theta\psi}(\mathbf{x}) - \bar{\psi} + \frac{1}{2L}\|\nabla f(\bar{\mathbf{x}})\|^2 + \left[L - \frac{1}{2\theta}\right]\|\mathbf{x} - \bar{\mathbf{x}}\|^2 + \frac{1}{4\theta}\|\mathbf{x} - \bar{\mathbf{x}}\|^2 + L_\varphi^2\theta \\ &\leq \text{env}_{\theta\psi}(\mathbf{x}) - \bar{\psi} + [\psi(\bar{\mathbf{x}}) - \bar{\psi}] + \frac{1}{2\theta}\|\mathbf{x} - \bar{\mathbf{x}}\|^2 + L_\varphi^2\theta = 2[\text{env}_{\theta\psi}(\mathbf{x}) - \bar{\psi}] + L_\varphi^2\theta. \end{aligned}$$

This finishes the proof of Lemma E.4. □

We now verify [Lemma 3.4](#). Choosing $\theta \in (0, [3L + \tau]^{-1})$ ensures that [Lemma E.1](#) and [Lemma E.4](#) are applicable. Let $k \in \mathbb{N}$ be given with $\alpha_k \leq \min\{\frac{1}{2\tau}, \frac{1}{2(\theta^{-1} - [L + \tau])}\}$. This implies $1 - \alpha_k\tau > \frac{1}{2}$,

$$2L \leq \theta^{-1} - [L + \tau] \leq \theta_k \leq 2(\theta^{-1} - [L + \tau]),$$

and $2 - \theta_k\alpha_k \geq 1$. Utilizing the definition of the Moreau envelope, we can further infer

$$\begin{aligned} \text{env}_{\theta\psi}(\mathbf{x}^{k+1}) &= \min_{\mathbf{y} \in \mathbb{R}^n} \psi(\mathbf{y}) + \frac{1}{2\theta} \|\mathbf{y} - \mathbf{x}^{k+1}\|^2 \\ &\leq \text{env}_{\theta\psi}(\mathbf{x}^k) + \frac{1}{2\theta} [\|\bar{\mathbf{x}}^k - \mathbf{x}^{k+1}\|^2 - \|\bar{\mathbf{x}}^k - \mathbf{x}^k\|^2]. \end{aligned} \quad (25)$$

Taking conditional expectation in (25) and applying [Lemma E.3](#), (C.4), and [Lemma E.4](#), we obtain

$$\begin{aligned} \mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi} \mid \mathcal{F}_k] &\leq [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + \frac{1}{2\theta} [\mathbb{E}[\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \mid \mathcal{F}_k] - \|\bar{\mathbf{x}}^k - \mathbf{x}^k\|^2], \\ &\leq [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] - \frac{2 - \theta_k\alpha_k}{2\theta} \theta_k\alpha_k \|\bar{\mathbf{x}}^k - \mathbf{x}^k\|^2 + \frac{\alpha_k^2 \sigma_k^2}{2\theta(1 - \alpha_k\tau)^2} \\ &\leq \left[1 + \frac{4C\alpha_k^2}{\theta}\right] [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] - L\theta\alpha_k \|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2 + 2\alpha_k^2 \left[CL_\varphi^2 + \frac{D}{\theta}\right], \end{aligned}$$

almost surely, where we used $\mathbf{x}^k - \bar{\mathbf{x}}^k = \theta \nabla \text{env}_{\theta\psi}(\mathbf{x}^k)$. This finishes the proof of [Lemma 3.4](#).

E.5 Proof of [Lemma 3.5](#)

We first establish an upper bound for the natural residual that is analogous to [[12](#), [Lemma A.1](#)].

Lemma E.5. *Suppose that the conditions (C.2) and (C.3) are satisfied. Let $\mathbf{x} \in \text{dom } \varphi$, $\mathbf{v} \in \mathbb{R}^n$, and $\alpha \in (0, \frac{1}{2}\tau^{-1})$ be given and set $F_v^\alpha(\mathbf{x}) := \mathbf{x} - \text{prox}_{\alpha\varphi}(\mathbf{x} - \alpha\mathbf{v})$. Then, it holds that*

$$\|F_v^\alpha(\mathbf{x})\|^2 \leq 4L_\varphi^2\alpha^2 + 4\|\mathbf{v}\|^2\alpha^2. \quad (26)$$

Proof. Let us first note that if h is τ -weakly convex, proper, and lower semicontinuous, then the proximity operator $\text{prox}_{\alpha h}$ can be equivalently characterized via the optimality condition:

$$\mathbf{p} = \text{prox}_{\alpha h}(\mathbf{x}) \iff \mathbf{x} - \mathbf{p} \in \alpha \partial h(\mathbf{p}) \quad (27)$$

for all $\alpha \in (0, \tau^{-1})$. Consequently, since φ is τ -weakly convex, we have $\alpha^{-1}(F_v^\alpha(\mathbf{x}) - \alpha\mathbf{v}) \in \partial\varphi(\text{prox}_{\alpha\varphi}(\mathbf{x} - \alpha\mathbf{v}))$. Using (22), this means

$$\varphi(\mathbf{x}) - \varphi(\text{prox}_{\alpha\varphi}(\mathbf{x} - \alpha\mathbf{v})) \geq \frac{1}{\alpha} \langle F_v^\alpha(\mathbf{x}) - \alpha\mathbf{v}, F_v^\alpha(\mathbf{x}) \rangle - \frac{\tau}{2} \|F_v^\alpha(\mathbf{x})\|^2.$$

In addition, due to $\mathbf{x}, \text{prox}_{\alpha\varphi}(\mathbf{x} - \alpha\mathbf{v}) \in \text{dom } \varphi$ and using Young's inequality, we have

$$\varphi(\mathbf{x}) - \varphi(\text{prox}_{\alpha\varphi}(\mathbf{x} - \alpha\mathbf{v})) \leq L_\varphi \|F_v^\alpha(\mathbf{x})\| \leq L_\varphi^2\alpha + \frac{1}{4\alpha} \|F_v^\alpha(\mathbf{x})\|^2.$$

Using Young's inequality once more — $\langle \mathbf{v}, F_v^\alpha(\mathbf{x}) \rangle \leq \alpha\|\mathbf{v}\|^2 + \frac{1}{4\alpha} \|F_v^\alpha(\mathbf{x})\|^2$ — it follows $\frac{1}{2}[\frac{1}{\alpha} - \tau] \|F_v^\alpha(\mathbf{x})\|^2 \leq L_\varphi^2\alpha + \alpha\|\mathbf{v}\|^2$. The choice of α then readily implies (26). \square

Based on [Lemma E.5](#), we now verify [Lemma 3.5](#). Reusing the notation in [Lemma E.5](#), we first observe

$$\mathbf{x}^{k+1} = \text{prox}_{\alpha_k\varphi}(\mathbf{x}^k - \alpha_k\mathbf{g}^k) = \mathbf{x}^k - F_{\mathbf{g}^k}^{\alpha_k}(\mathbf{x}^k).$$

Hence, by (24), [Lemma E.4](#), and (C.4), it follows

$$\begin{aligned} \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] &= \mathbb{E}[\|F_{\mathbf{g}^k}^{\alpha_k}(\mathbf{x}^k)\|^2 \mid \mathcal{F}_k] \leq 4L_\varphi^2\alpha_k^2 + 4\alpha_k^2\mathbb{E}[\|\mathbf{g}^k\|^2 \mid \mathcal{F}_k] \\ &= 4L_\varphi^2\alpha_k^2 + 4\alpha_k^2\|\nabla f(\mathbf{x}^k)\|^2 + 4\alpha_k^2\mathbb{E}[\|\nabla f(\mathbf{x}^k) - \mathbf{g}^k\|^2 \mid \mathcal{F}_k] \\ &\leq 4(2L + C)\alpha_k^2 \cdot [f(\mathbf{x}^k) - \bar{f}] + 4(L_\varphi^2 + D)\alpha_k^2 \\ &\leq 8(2L + C)\alpha_k^2 \cdot [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + 4(((2L + C)\theta + 1)L_\varphi^2 + D)\alpha_k^2. \end{aligned}$$

This is exactly (12) in [Lemma 3.5](#).

E.6 Expected iteration complexity of prox-SGD

As a byproduct of our analysis, we provide the expected iteration complexity of prox-SGD below.

Corollary E.6. *Let $\{\mathbf{x}^k\}_{k \geq 0}$ be generated by prox-SGD and let the assumptions (C.1)–(C.5) be satisfied. Then, for $\theta \in (0, [3\mathbf{L} + \tau]^{-1})$ and all k with $\alpha_k \leq \min\{\frac{1}{2\tau}, \frac{1}{2(\theta^{-1} - \lfloor \mathbf{L} + \tau \rfloor)}\}$, it holds that*

$$\min_{k=0, \dots, T} \mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] \leq \frac{\text{env}_{\theta\psi}(\mathbf{x}^0) - \bar{\psi} + \mathbf{K} \sum_{k=0}^T \alpha_k^2}{\mathbf{L}\theta \sum_{k=0}^T \alpha_k}.$$

Here, $\mathbf{K} > 0$ is defined in the proof. Consequently, if $\alpha_k = \frac{c}{(k+1)\log(k+2)}$ with some proper $c > 0$, then we have

$$\min_{k=0, \dots, T} \mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] \leq \mathcal{O}\left(\frac{\log(T+2)}{\sqrt{T+1}}\right).$$

Proof. Taking total expectation in Lemma 3.4 gives

$$\begin{aligned} \mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi}] &\leq (1 + 4\mathbf{C}\theta^{-1}\alpha_k^2)\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] \\ &\quad - \mathbf{L}\theta\alpha_k\mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] + 2\alpha_k^2(\mathbf{C}\mathbf{L}_\varphi^2 + \mathbf{D}\theta^{-1}). \end{aligned} \quad (28)$$

Then, by unrolling this recursion and setting $c_1 := 4\mathbf{C}\theta^{-1}$ and $c_2 := 2(\mathbf{C}\mathbf{L}_\varphi^2 + \mathbf{D}\theta^{-1})$, we have

$$\begin{aligned} \mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi}] &\leq \left\{ \prod_{j=0}^k (1 + c_1\alpha_j^2) \right\} \underbrace{\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^0) - \bar{\psi}]}_{E_1} \\ &\quad + c_2 \sum_{j=0}^{k-1} \left\{ \prod_{i=j+1}^k (1 + c_1\alpha_i^2) \right\} \alpha_j^2 + c_2\alpha_k^2. \end{aligned}$$

Next, using $\log(1+x) \leq x$ for all $x \geq 0$, we obtain the estimate

$$\prod_{i=j}^k (1 + c_1\alpha_i^2) = \exp\left(\sum_{i=j}^k \log(1 + c_1\alpha_i^2)\right) \leq \exp\left(\sum_{i=0}^\infty c_1\alpha_i^2\right) =: c_3$$

for all $j = 0, \dots, k$ and $k \in \mathbb{N}$. Thus, it follows $\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi}] \leq c_3 E_1 + c_2 c_3 \sum_{j=0}^k \alpha_j^2$ and we can infer $\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] \leq \mathbf{G}$ for all k and some constant $\mathbf{G} > 0$. Invoking this bound in (28) yields

$$\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi}] \leq \mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] - \mathbf{L}\theta\alpha_k\mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] + \alpha_k^2 \underbrace{(c_1\mathbf{G} + c_2)}_{\mathbf{K}}.$$

Finally, summing the above recursion from $k = 0$ to $k = T$ and rearranging the terms, we have

$$\min_{k=0, \dots, T} \mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|^2] \leq \frac{\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^0) - \bar{\psi}] + \mathbf{K} \sum_{k=0}^T \alpha_k^2}{\mathbf{L}\theta \sum_{k=0}^T \alpha_k}.$$

□

F Results and proofs for Subsection 3.4

F.1 Three stochastic model-based methods

Depending on the choice of the model function $f_{\mathbf{x}^k}(\cdot, \xi^k)$, we recover different stochastic optimization methods. Specifically, we can consider the following examples:

$$\begin{cases} f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) = f(\mathbf{x}^k, \xi^k) + \langle \mathbf{s}(\mathbf{x}^k, \xi^k), \mathbf{x} - \mathbf{x}^k \rangle & \text{stochastic subgradient model,} \\ f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) = f(\mathbf{x}, \xi^k) & \text{stochastic proximal point model,} \end{cases}$$

where $\mathbf{s}(\mathbf{x}^k, \xi^k) \in \partial f(\mathbf{x}^k, \xi^k)$ is a selected subgradient. In addition, if the weakly convex function f has the composite form $f(\mathbf{x}, \xi) = h(c(\mathbf{x}, \xi), \xi)$, where $h(\cdot, \xi)$ is convex and Lipschitz continuous, $c(\cdot, \xi)$ is smooth and has Lipschitz continuous Jacobian, we can further cover the prox-linear model:

$$f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) = h(c(\mathbf{x}^k, \xi^k) + \nabla c(\mathbf{x}^k, \xi^k)(\mathbf{x} - \mathbf{x}^k), \xi^k), \quad \text{stochastic prox-linear model.}$$

F.2 Proof of Lemma 3.9

Under our assumptions, the estimate [11, (4.8)] is valid. Furthermore, in [11, Lemma 4.1], it is shown that the assumptions (D.1) and (D.3) imply Lipschitz continuity of the mapping f (with constant L). Hence, setting $\mathbf{x} = \mathbf{x}^k$ in [11, (4.8)], we have

$$\begin{aligned} \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] &\leq -\frac{2\alpha_k}{1-\eta\alpha_k} \mathbb{E}[\psi(\mathbf{x}^{k+1}) - \psi(\mathbf{x}^k) \mid \mathcal{F}_k] + \frac{2L^2\alpha_k^2}{1-\eta\alpha_k} \\ &\leq \frac{2(L+L_\varphi)\alpha_k}{1-\eta\alpha_k} \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\| \mid \mathcal{F}_k] + \frac{2L^2\alpha_k^2}{1-\eta\alpha_k} \\ &\leq \frac{1}{2} \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] + \frac{2(L+L_\varphi)^2\alpha_k^2}{(1-\eta\alpha_k)^2} + \frac{2L^2\alpha_k^2}{1-\eta\alpha_k}, \end{aligned}$$

where we have used the Lipschitz continuity of ψ in the second inequality, while the last inequality is due to Young's and Jensen inequalities. Rearranging terms gives

$$\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \leq \left(\frac{4(L+L_\varphi)^2}{(1-\eta\alpha_k)^2} + \frac{4L^2}{1-\eta\alpha_k} \right) \alpha_k^2.$$

Thus, for all k with $\alpha_k \leq 1/(2\eta)$, we obtain $\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \leq (16(L+L_\varphi)^2 + 8L^2)\alpha_k^2$.

F.3 Convergence of SMM: An extended setting

In the following, we show that the analysis and results presented in the previous section and in subsection 3.4 can be further strengthened and generalized. In particular, it is possible to work with assumptions that are more aligned with the conditions (C.1)–(C.4) for prox-SGD. We consider the assumptions:

- (E.1) The stochastic model function $f_{\mathbf{x}}$ satisfies a one-sided accuracy property, i.e., we have $\mathbb{E}[f_{\mathbf{x}^k}(\mathbf{x}^k, \xi^k) \mid \mathcal{F}_k] = f(\mathbf{x}^k)$ for all k and
$$\mathbb{E}[f_{\mathbf{x}^k}(\mathbf{y}, \xi^k) - f(\mathbf{y}) \mid \mathcal{F}_k] \leq \frac{\tau}{2} \|\mathbf{x}^k - \mathbf{y}\|^2 \quad \forall \mathbf{y} \in \text{dom } \varphi, \quad \forall k, \quad \text{almost surely.}$$
- (E.2) The function $\mathbf{y} \mapsto f_{\mathbf{x}^k}(\mathbf{y}, \xi^k)$ is η -weakly convex for all k .
- (E.3) There exists a sequence $\{\mathbf{g}^k\}_{k \geq 0}$ and constants $C, D \geq 0$ such that $\mathbf{g}^k \in \partial f_{\mathbf{x}^k}(\mathbf{x}^k, \xi^k)$ and $\mathbb{E}[\|\mathbf{g}^k\|^2 \mid \mathcal{F}_k] \leq C[f(\mathbf{x}^k) - \bar{f}] + D$ almost surely for all k .
- (E.4) The function φ is ρ -weakly convex, proper, lower semicontinuous, and L_φ -Lipschitz continuous on $\text{dom } \varphi$. In addition, φ is bounded from below on $\text{dom } \varphi$, i.e., we have $\varphi(\mathbf{x}) \geq \bar{\varphi}$ for all $\mathbf{x} \in \text{dom } \varphi$.
- (E.5) The step sizes $\{\alpha_k\}_{k \geq 0}$ satisfy $\sum_{k=0}^{\infty} \alpha_k = \infty$ and $\sum_{k=0}^{\infty} \alpha_k^2 < \infty$.

Concerning the regularity of f , we will work with *one* of the following scenarios:

- (F.1) The mapping f is bounded from below and L_f -Lipschitz continuous on $\text{dom } \varphi$, i.e., there exists \bar{f} such that $f(\mathbf{x}) \geq \bar{f}$ for all $\mathbf{x} \in \text{dom } \varphi$ and we have $|f(\mathbf{x}) - f(\mathbf{y})| \leq L_f \|\mathbf{x} - \mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \text{dom } \varphi$.
- (F.2) The function f is bounded from below on \mathbb{R}^n , i.e., there is \bar{f} such that $f(\mathbf{x}) \geq \bar{f}$ for all $\mathbf{x} \in \mathbb{R}^n$, and the gradient mapping ∇f is Lipschitz continuous (on \mathbb{R}^n) with modulus $L > 0$.

In contrast to subsection 3.4 and [11], we will not require explicit Lipschitz continuity of the model function $f_{\mathbf{x}}$. Instead — as we will verify now — convergence of SMM can be established by only assuming Lipschitz continuity or Lipschitz smoothness of the mapping f . As mentioned in Appendix F.2, the conditions (D.1) and (D.3) already imply Lipschitz continuity of f (see again [11, Lemma 4.1]). Hence, assuming (F.1) can be more general. Moreover, the alternative assumption (F.2) allows to cover the case where f is not Lipschitz continuous but sufficiently smooth. This situation appears more frequently in stochastic proximal point methods where the model function is chosen as $f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) = f(\mathbf{x}, \xi^k)$. Assumption (E.2) is parallel to assumption (D.2). Assumption (E.3) is a mild variance-type condition that is similar to (C.4) and which can be weaker than (D.3). We now establish the core estimates provided in Lemma 3.8 and Lemma 3.9 under the more general assumptions (E.1)–(E.5) and (F.1) or (F.2). Our analysis is inspired by the proof of [11, Lemma 4.2] but is closer to [29, Appendix C].

Lipschitz continuous f. We first investigate core properties of SMM under (F.1). Let us define $\Psi_k(\mathbf{x}) := f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) + \varphi(\mathbf{x})$ and $F_k(\mathbf{x}) := f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) + \frac{\eta}{2}\|\mathbf{x} - \mathbf{x}^k\|^2$. Then, by assumption, the mapping $\mathbf{x} \mapsto \Psi_k(\mathbf{x}) + \frac{1}{2\alpha_k}\|\mathbf{x} - \mathbf{x}^k\|^2$ is strongly convex with parameter $\alpha_k^{-1} - \zeta$, where $\zeta = \eta + \rho$. Specifically, if $\alpha_k < \zeta^{-1}$ and by (14), we have

$$\Psi_k(\mathbf{x}) + \frac{1}{2\alpha_k}\|\mathbf{x} - \mathbf{x}^k\|^2 \geq \Psi_k(\mathbf{x}^{k+1}) + \frac{1}{2\alpha_k}\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 + \frac{1}{2}[\alpha_k^{-1} - \zeta]\|\mathbf{x} - \mathbf{x}^{k+1}\|^2 \quad (29)$$

for all $\mathbf{x} \in \text{dom } \varphi$ and due to the convexity of F_k and the Lipschitz continuity of f , it holds that

$$\begin{aligned} \Psi_k(\mathbf{x}) - \Psi_k(\mathbf{x}^{k+1}) &= [\psi(\mathbf{x}) - \psi(\mathbf{x}^{k+1})] + [f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) - f(\mathbf{x})] + [f(\mathbf{x}^{k+1}) - f(\mathbf{x}^k)] \\ &\quad + [f(\mathbf{x}^k) - f_{\mathbf{x}^k}(\mathbf{x}^k, \xi^k)] + [F_k(\mathbf{x}^k) - F_k(\mathbf{x}^{k+1})] + \frac{\eta}{2}\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \\ &\leq [\psi(\mathbf{x}) - \psi(\mathbf{x}^{k+1})] + [f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) - f(\mathbf{x})] + [f(\mathbf{x}^k) - f_{\mathbf{x}^k}(\mathbf{x}^k, \xi^k)] \\ &\quad + \langle \mathbf{g}^k, \mathbf{x}^k - \mathbf{x}^{k+1} \rangle + \mathsf{L}_f\|\mathbf{x}^{k+1} - \mathbf{x}^k\| + \frac{\eta}{2}\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2, \end{aligned} \quad (30)$$

where we used $\mathbf{g}^k \in \partial f_{\mathbf{x}^k}(\mathbf{x}^k, \xi^k) = \partial F_k(\mathbf{x}^k)$. Upon taking conditional expectation and using the Cauchy-Schwarz and Young's inequality and (E.1), we obtain

$$\begin{aligned} \mathbb{E}[\Psi_k(\mathbf{x}) - \Psi_k(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] &\leq \mathbb{E}[\psi(\mathbf{x}) - \psi(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] + 2\alpha_k\mathbb{E}[\|\mathbf{g}^k\|^2 \mid \mathcal{F}_k] + 2\mathsf{L}_f^2\alpha_k \\ &\quad + \frac{\tau}{2}\|\mathbf{x}^k - \mathbf{x}\|^2 + \frac{1}{2}\left[\eta + \frac{1}{2\alpha_k}\right]\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k]. \end{aligned}$$

Rearranging the terms in (29), this yields

$$\begin{aligned} &\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}\|^2 \mid \mathcal{F}_k] \\ &\leq \frac{1 + \alpha_k\tau}{1 - \alpha_k\zeta}\|\mathbf{x}^k - \mathbf{x}\|^2 + \frac{2\alpha_k}{1 - \alpha_k\zeta}\mathbb{E}[\psi(\mathbf{x}) - \psi(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] \\ &\quad + \frac{4\alpha_k^2}{1 - \alpha_k\zeta}\mathbb{E}[\|\mathbf{g}^k\|^2 \mid \mathcal{F}_k] + \frac{4\mathsf{L}_f^2\alpha_k^2}{1 - \alpha_k\zeta} - \frac{1 - 2\alpha_k\eta}{2(1 - \alpha_k\zeta)}\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k]. \end{aligned} \quad (31)$$

Let us now define $\bar{\mathbf{x}}^k := \text{prox}_{\theta\psi}(\mathbf{x}^k)$ for $\theta \in (0, \zeta^{-1})$. Then, due to $\mathbf{x}^k, \bar{\mathbf{x}}^k \in \text{dom } \varphi$ and the Lipschitz continuity of f and φ and applying Young's inequality, it holds that

$$\begin{aligned} f(\mathbf{x}^k) - \bar{f} &\leq \psi(\mathbf{x}^k) - \bar{\psi} \leq \text{env}_{\theta\psi}(\mathbf{x}^k) + (\mathsf{L}_f + \mathsf{L}_\varphi)\|\mathbf{x}^k - \bar{\mathbf{x}}^k\| - \frac{1}{2\theta}\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 - \bar{\psi} \\ &\leq \text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi} + 0.5\theta(\mathsf{L}_f + \mathsf{L}_\varphi)^2. \end{aligned} \quad (32)$$

In addition, setting $\mathbf{x} = \bar{\mathbf{x}}^k$ in (31), using

$$\psi(\bar{\mathbf{x}}^k) = \text{env}_{\theta\psi}(\mathbf{x}^k) - \frac{1}{2\theta}\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 \leq \psi(\mathbf{x}^{k+1}) + \frac{1}{2\theta}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 - \|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2]$$

and combining (E.3) and (32), we obtain

$$\begin{aligned} \mathbb{E}[\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \mid \mathcal{F}_k] &\leq \left[1 - \frac{(\theta^{-1} - \zeta - \tau)\alpha_k}{1 - \alpha_k\zeta}\right]\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \frac{4\mathsf{C}\alpha_k^2}{1 - \alpha_k\zeta}[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] \\ &\quad + \frac{\mathsf{G}\alpha_k^2}{1 - \alpha_k\zeta} - \frac{1 - 2[\eta + \theta^{-1}]\alpha_k}{2(1 - \alpha_k\zeta)}\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k], \end{aligned} \quad (33)$$

where $\mathsf{G} := 2\mathsf{C}\theta(\mathsf{L}_f + \mathsf{L}_\varphi)^2 + 4(\mathsf{L}_f^2 + \mathsf{D})$. Hence, using the definition of the Moreau envelope and (33), for $\theta \in (0, (\zeta + \tau)^{-1})$ and all k with $\alpha_k \leq \min\{\frac{1}{2\zeta}, \frac{\theta}{2(1+\theta\eta)}\}$, it follows

$$\begin{aligned} &\mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi} \mid \mathcal{F}_k] \\ &\leq [\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + \frac{1}{2\theta}[\mathbb{E}[\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \mid \mathcal{F}_k] - \|\bar{\mathbf{x}}^k - \mathbf{x}^k\|^2], \\ &\leq \left[1 + \frac{4\mathsf{C}\alpha_k^2}{\theta}\right][\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] - \frac{1 - \theta(\zeta + \tau)}{2}\alpha_k\|\nabla\text{env}_{\theta\psi}(\mathbf{x}^k)\|^2 + \frac{\mathsf{G}\alpha_k^2}{\theta}. \end{aligned} \quad (34)$$

This gives an estimate similar to Lemma 3.8.

Next, let us consider an iteration k with $\alpha_k \leq \min\{\frac{1}{2\zeta}, \frac{1}{2\eta}\}$. Using (31) with $\mathbf{x} = \mathbf{x}^k$ and Young's inequality, we obtain

$$\begin{aligned} & \frac{1}{2}\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \\ & \leq 2\alpha_k\mathbb{E}[\psi(\mathbf{x}^k) - \psi(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] + 4\alpha_k^2\mathbb{E}[\|\mathbf{g}^k\|^2 \mid \mathcal{F}_k] + 4\mathbf{L}_f^2\alpha_k^2 \\ & \leq 2(\mathbf{L}_f + \mathbf{L}_\varphi)\alpha_k\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\| \mid \mathcal{F}_k] + 4\mathbf{C}\alpha_k^2[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + \mathbf{G}\alpha_k^2 \\ & \leq \frac{1}{4}\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] + 4\mathbf{C}\alpha_k^2[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + [\mathbf{G} + 4(\mathbf{L}_f + \mathbf{L}_\varphi)^2]\alpha_k^2. \end{aligned}$$

Rearranging the terms yields $\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \leq 16\mathbf{C}\alpha_k^2[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + 4[\mathbf{G} + 4(\mathbf{L}_f + \mathbf{L}_\varphi)^2]\alpha_k^2$, which is similar to the estimates in Lemma 3.9 (indeed, it is more similar to Lemma 3.5).

With these estimates and by following exactly the derivations in Subsection 3.3, we can show $\mathbb{E}[\|\nabla\text{env}_{\theta\psi}(\mathbf{x}^k)\|] \rightarrow 0$ and $\|\nabla\text{env}_{\theta\psi}(\mathbf{x}^k)\| \rightarrow 0$ almost surely.

Lipschitz smooth f . We continue our discussion under (F.2). Using the Lipschitz continuity of ∇f , the estimate (30) changes to

$$\begin{aligned} \Psi_k(\mathbf{x}) - \Psi_k(\mathbf{x}^{k+1}) & \leq [\psi(\mathbf{x}) - \psi(\mathbf{x}^{k+1})] + [f_{\mathbf{x}^k}(\mathbf{x}, \xi^k) - f(\mathbf{x})] + [f(\mathbf{x}^k) - f_{\mathbf{x}^k}(\mathbf{x}^k, \xi^k)] \\ & \quad + \langle \mathbf{g}^k - \nabla f(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^{k+1} \rangle + \frac{\mathbf{L} + \eta}{2}\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2. \end{aligned}$$

Taking conditional expectation and applying Young's inequality and (E.1), this yields

$$\begin{aligned} \mathbb{E}[\Psi_k(\mathbf{x}) - \Psi_k(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] & \leq \mathbb{E}[\psi(\mathbf{x}) - \psi(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] + 2\alpha_k\mathbb{E}[\|\mathbf{g}^k\|^2 \mid \mathcal{F}_k] + 2\alpha_k\|\nabla f(\mathbf{x}^k)\|^2 \\ & \quad + \frac{\tau}{2}\|\mathbf{x}^k - \mathbf{x}\|^2 + \frac{1}{2}\left[\mathbf{L} + \eta + \frac{1}{2\alpha_k}\right]\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k]. \end{aligned}$$

By Lemma E.4 and (24), we can further infer $f(\mathbf{x}) - \bar{f} \leq 2[\text{env}_{\theta\psi}(\mathbf{x}) - \bar{\psi}] + \mathbf{L}_\varphi^2\theta$ and $\|\nabla f(\mathbf{x})\|^2 \leq 2\mathbf{L}[f(\mathbf{x}) - \bar{f}]$ for all $\mathbf{x} \in \text{dom } \varphi$. (Notice that the conditions (C.1)–(C.3) and (E.4) and (F.2) are identical). Hence, similar to (31), it follows

$$\begin{aligned} & \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}\|^2 \mid \mathcal{F}_k] \\ & \leq \frac{1 + \alpha_k\tau}{1 - \alpha_k\zeta}\|\mathbf{x}^k - \mathbf{x}\|^2 + \frac{2\alpha_k}{1 - \alpha_k\zeta}\mathbb{E}[\psi(\mathbf{x}) - \psi(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] + \frac{\tilde{\mathbf{G}}\alpha_k^2}{1 - \alpha_k\zeta} \\ & \quad + \frac{8(\mathbf{C} + 2\mathbf{L})\alpha_k^2}{1 - \alpha_k\zeta}[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] - \frac{1 - 2\alpha_k(\mathbf{L} + \eta)}{2(1 - \alpha_k\zeta)}\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k], \quad (35) \end{aligned}$$

where $\tilde{\mathbf{G}} = 4((\mathbf{C} + 2\mathbf{L})\mathbf{L}_\varphi^2 + \mathbf{D})$ and $\zeta = \eta + \rho$. At this point, we can fully mimic our earlier calculations. In particular, similar to (34), for $\theta \in (0, (\zeta + \tau)^{-1})$ and all k with $\alpha_k \leq \min\{\frac{1}{2\zeta}, \frac{\theta}{2(1+\theta(\mathbf{L}+\eta))}\}$, we obtain

$$\begin{aligned} \mathbb{E}[\text{env}_{\theta\psi}(\mathbf{x}^{k+1}) - \bar{\psi} \mid \mathcal{F}_k] & \leq \left[1 + \frac{8(\mathbf{C} + 2\mathbf{L})\alpha_k^2}{\theta}\right][\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] \\ & \quad - \frac{1 - \theta(\zeta + \tau)}{2}\alpha_k\|\nabla\text{env}_{\theta\psi}(\mathbf{x}^k)\|^2 + \frac{\tilde{\mathbf{G}}\alpha_k^2}{\theta}. \end{aligned}$$

Setting $\mathbf{x} = \mathbf{x}^k$ in (35) and using the Lipschitz continuity of φ and ∇f and Young's inequality, it holds that

$$\begin{aligned} & \frac{1}{2}\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \leq 2\alpha_k\mathbb{E}[\psi(\mathbf{x}^k) - \psi(\mathbf{x}^{k+1}) \mid \mathcal{F}_k] \\ & \quad + 8(\mathbf{C} + 2\mathbf{L})\alpha_k^2[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + \tilde{\mathbf{G}}\alpha_k^2 \\ & \leq 2\mathbf{L}_\varphi\alpha_k\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\| \mid \mathcal{F}_k] + \mathbf{L}\alpha_k\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] \\ & \quad - 2\alpha_k\mathbb{E}[\langle \nabla f(\mathbf{x}^k), \mathbf{x}^{k+1} - \mathbf{x}^k \rangle \mid \mathcal{F}_k] + 8(\mathbf{C} + 2\mathbf{L})\alpha_k^2[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + \tilde{\mathbf{G}}\alpha_k^2 \\ & \leq \frac{\mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k]}{4} + 16\alpha_k^2\|\nabla f(\mathbf{x}^k)\|^2 + 8(\mathbf{C} + 2\mathbf{L})\alpha_k^2[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + [\tilde{\mathbf{G}} + 16\mathbf{L}_\varphi^2]\alpha_k^2 \end{aligned}$$

$$\leq \frac{1}{4} \mathbb{E}[\|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2 \mid \mathcal{F}_k] + 8(C + 10L)\alpha_k^2[\text{env}_{\theta\psi}(\mathbf{x}^k) - \bar{\psi}] + [\tilde{G} + 16(1 + 2L\theta)L_\varphi^2]\alpha_k^2$$

for all k with $\alpha_k \leq \min\{\frac{1}{2\zeta}, \frac{1}{8L}, \frac{1}{2(L+\eta)}\}$. As before, this establishes variants of [Lemma 3.8](#) and [Lemma 3.9](#) (more precisely, [Lemma 3.5](#)) and allows to follow the derivations in [Subsection 3.3](#) to establish convergence results.

Finally, we summarize all the above observations in the following corollary.

Corollary F.1. *Let us consider the family of stochastic model-based methods (14) for the problem (13) under assumptions (E.1)–(E.5) and (F.1) or (F.2). Then, for all $\theta \in (0, (\eta + \rho + \tau)^{-1})$, we have $\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\|] = 0$ and $\lim_{k \rightarrow \infty} \|\nabla \text{env}_{\theta\psi}(\mathbf{x}^k)\| = 0$ almost surely.*

G Comparison: related literature

Table 1: Summary and comparison of related and representative literature.

SGD	Assumptions	Convergence	
		in expectation	almost surely
[3]	(A.1), (A.2), (A.4), bounded variance ((A.3) with $C = 0$)	✗	✓
[28]	(A.1), (A.4) f is coercive (\implies (A.2)) f is Lipschitz $\liminf_{\ x\ \rightarrow \infty} \ \nabla f(x)\ > 0$ bounded variance ((A.3) with $C = 0$)	✗	✓
[6]	(A.1), (A.2), (A.4) f is twice differentiable $x \mapsto \nabla^2 f(x)\nabla f(x)$ is Lipschitz bounded variance ((A.3) with $C = 0$)	✓	✗
This work	(A.1)–(A.4)	✓	✓
RR	Assumptions	Convergence	
		in expectation	almost surely
[25]	(B.1)–(B.2)	✗	✓
This work	(B.1)–(B.2)	✓	✓
prox-SGD	Assumptions	Convergence	
		in expectation	almost surely
[27]	(C.1), (C.3), (C.5) $\{\mathbf{x}^k\}_{k \geq 0}$ is surely bounded almost surely bounded noise	✗	✓
This work	(C.1)–(C.5)	✓	✓
SMM	Assumptions	Convergence	
		in expectation	almost surely
[14]	(D.3), (D.5) surely one-sided accuracy (\implies (D.1)) φ is convex (\implies (D.2)) $\{\mathbf{x}^k\}_{k \geq 0}$ is surely bounded (compact constraint) density / Sard-type condition	✗	✓
This work	(D.1)–(D.5)	✓	✓

H Non-asymptotic complexity vs. asymptotic convergence

In this subsection, we provide several additional arguments and illustrations that can help to explain and illuminate the potential differences between typical finite-step complexity rates and asymptotic convergence results — as obtained in [Theorem 2.1](#). To this end, we consider the standard optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}), \quad (36)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a given, smooth, and nonconvex function. As motivated in the introduction, complexity bounds for the nonconvex problem (36) typically take the form

$$\min_{k=0, \dots, T} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|^2] = \mathcal{O}((T+1)^{-\frac{1}{2}}) \quad \text{or} \quad \mathbb{E}[\|\nabla f(\mathbf{x}^{\bar{k}})\|^2] \leq \mathcal{O}((T+1)^{-\frac{1}{2}}). \quad (37)$$

Here, T denotes the total number of iterations, the index \bar{k} is sampled uniformly at random from $\{0, \dots, T\}$, and the iterates $\{\mathbf{x}^k\}_{k \geq 0}$ are assumed to be generated by the stochastic gradient descent method, see [\[17, 6, 24, 23\]](#). Similar (deterministic) complexity results are also available for the basic gradient descent method, [\[2\]](#), and many other related algorithmic schemes, [\[18, 11, 30, 32\]](#). In particular, for the gradient descent method, the complexity bounds (37) can be strengthened to

$$\min_{k=0, \dots, T} \|\nabla f(\mathbf{x}^k)\|^2 = \mathcal{O}((T+1)^{-1}), \quad (38)$$

see, e.g., [\[2, Theorem 10.15\]](#).

While the complexity bounds shown in (37) (and (38)) allow to capture and characterize the overall trend of the minimization procedure during the first T iterations, they cannot fully justify a common practice in stochastic optimization: the last iterate \mathbf{x}^T is returned as final output of the algorithm. In fact, even as T increases, the term $\mathbb{E}[\|\nabla f(\mathbf{x}^T)\|]$ can be arbitrarily large whereas the complexity measure $\min_{k=0, \dots, T} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|^2]$ decreases at its respective rate. The asymptotic convergence results

$$\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla f(\mathbf{x}^k)\|] = 0 \quad \text{or} \quad \lim_{k \rightarrow \infty} \|\nabla f(\mathbf{x}^k)\| = 0 \quad \text{almost surely,}$$

can provide additional information: as $\{\mathbb{E}[\|\nabla f(\mathbf{x}^k)\|]\}_{k \geq 0}$ converges to zero, the term $\mathbb{E}[\|\nabla f(\mathbf{x}^T)\|]$ will stay small (below any predefined threshold) for all T sufficiently large. In tandem with the complexity bounds (37), this supports common output strategies that return the last iterate \mathbf{x}^T — at least for large T . Hence, both non-asymptotic and asymptotic convergence analyses are important and informative — especially in the nonconvex and stochastic setting — and the combination of non-asymptotic and asymptotic convergence guarantees can paint a more complete picture of the convergence behavior of stochastic optimization methods.

We continue with a specific example that can illustrate the mentioned discrepancies and differences between non-asymptotic complexity and asymptotic convergence results. In particular, we construct a nonconvex function $f : \mathbb{R} \rightarrow \mathbb{R}$, a corresponding step size sequence $\{\alpha_k\}_{k \geq 0}$, and an initial point $\mathbf{x}^0 \in \mathbb{R}$, for which the standard gradient descent method generates a sequence of iterates $\{\mathbf{x}^k\}_{k \geq 0}$ that satisfies the non-asymptotic complexity bound (38), but we can *not* observe asymptotic convergence $f'(\mathbf{x}^k) \rightarrow 0$.

Let $0 \leq \nu < \kappa$ be given parameters. We consider the functions

$$c(x) := \begin{cases} e^{-\frac{1}{x}} & \text{if } x > 0, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad \bar{c}_{\kappa, \nu}(x) := \frac{c(\kappa^2 - x)}{c(\kappa^2 - x) + c(x - \nu^2)}.$$

The mappings c and $\bar{c}_{\kappa, \nu}$ are obviously C^∞ and we have $\bar{c}_{\kappa, \nu}(x) = 0$ for all $x \geq \kappa^2$ and $\bar{c}_{\kappa, \nu}(x) = 1$ for all $x \leq \nu^2$. Moreover, it holds that

$$c'(x) = \begin{cases} \frac{1}{x^2} e^{-\frac{1}{x}} & \text{if } x > 0, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad \bar{c}'_{\kappa, \nu}(x) = -\frac{c'(\kappa^2 - x)c(x - \nu^2) + c(\kappa^2 - x)c'(x - \nu^2)}{(c(\kappa^2 - x) + c(x - \nu^2))^2}.$$

We now define $\kappa_j := \frac{1}{4j(2j+1)}$, $\nu_j := \frac{1}{8j(2j+1)}$, $g(x) := x - \frac{1}{2}x^2$, $\gamma_j(x) := g(x)\bar{c}_{\kappa_j, \nu_j}((x - \frac{1}{2j})^2)$, and

$$f(x) := \begin{cases} h(x) + \sum_{k=1}^{\infty} \gamma_k(x) & \text{if } x \neq 0, \\ 0 & \text{if } x = 0, \end{cases} \quad \text{and} \quad h(x) := \begin{cases} \frac{1}{2}x^2 & \text{if } x \geq 0, \\ 8x^2(8x^2 - 1) & \text{if } x < 0. \end{cases}$$




Figure 1: Plot of f . The function f is depicted using a blue color. The dotted orange lines correspond to the functions $x \mapsto x$ and $x \mapsto \frac{1}{2}x^2$, respectively. The read points in subfigure (a) show the iterates \mathbf{x}^k , $k \in \mathbb{N}$. Subfigure (b) shows a logarithmic plot of f on $[0.01, 1]$.

An exemplary plot of the function f is shown in Figure 1. The function f is continuous on \mathbb{R} and smooth for all $x \neq 0$. We now want to run gradient descent on f with initial point $\mathbf{x}^0 = 1$ and diminishing step sizes

$$\alpha_k := \frac{1}{k+2}, \text{ if } k \text{ is even} \quad \text{and} \quad \alpha_k := \frac{1}{(k+1)(k+2)}, \text{ if } k \text{ is odd.}$$

Then, it follows

$$\mathbf{x}^k = \frac{1}{k+1} \quad \text{and} \quad f'(\mathbf{x}^k) = \begin{cases} 1 & \text{if } k \text{ is even,} \\ \frac{1}{k+1} & \text{if } k \text{ is odd.} \end{cases} \quad (39)$$

We now verify this statement by induction. We first notice that the functions $\{\gamma_k\}_{k \geq 0}$ and derivatives $\{\gamma'_k\}_{k \geq 0}$ have disjoint supports $[\frac{1}{2k} - \frac{1}{4k(2k+1)}, \frac{1}{2k} + \frac{1}{4k(2k+1)}] = [\frac{1}{2}(\frac{1}{2k} + \frac{1}{2k+1}), \frac{1}{2}(\frac{3}{2k} - \frac{1}{2k+1})]$ with center point $\frac{1}{2k}$, $k \in \mathbb{N}$. More specifically, we have

$$\gamma_k(x) = \begin{cases} 0 & \text{if } |x - \frac{1}{2k}| \geq \frac{1}{4k(2k+1)}, \\ g(x) & \text{if } |x - \frac{1}{2k}| \leq \frac{1}{8k(2k+1)}, \end{cases} \quad \text{and} \quad \gamma'_k(x) = \begin{cases} 0 & \text{if } |x - \frac{1}{2k}| \geq \frac{1}{4k(2k+1)}, \\ g'(x) & \text{if } |x - \frac{1}{2k}| \leq \frac{1}{8k(2k+1)}. \end{cases}$$

Hence, (39) clearly holds for the base case $k = 0$. Let us now assume that the induction hypothesis (39) is true for some k and let $k+1 = 2j$, $j \in \mathbb{N}$, be an even number. Then, we obtain

$$\begin{aligned} \mathbf{x}^{k+1} &= \mathbf{x}^k - \alpha_k f'(\mathbf{x}^k) \\ &= \frac{1}{k+1} - \frac{1}{(k+1)(k+2)} \cdot \left[\frac{1}{k+1} + \gamma'_j\left(\frac{1}{k+1}\right) \right] \\ &= \frac{1}{k+1} - \frac{1}{(k+1)(k+2)} \cdot 1 = \frac{1}{k+2}. \end{aligned}$$

Similarly, if $k+1 = 2j+1$, $j \in \mathbb{N}$, is odd, we then have $\frac{1}{2j} - \frac{1}{k+1} = \frac{1}{2j(2j+1)} > \frac{1}{4j(2j+1)}$ and $\frac{1}{k+1} - \frac{1}{2j+2} = \frac{1}{(2j+1)(2j+2)} > \frac{1}{(2j+2)(2j+3)}$. This yields

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \alpha_k f'(\mathbf{x}^k) = \frac{1}{k+1} - \frac{1}{k+2} \cdot \frac{1}{k+1} = \frac{1}{k+2}.$$

This finishes the proof of (39). We further notice

$$\sum_{k=0}^{\infty} \alpha_k \geq \sum_{k=0}^{\infty} \alpha_{2k} = \sum_{k=0}^{\infty} \frac{1}{2(k+1)} = \infty \quad \text{and} \quad \sum_{k=0}^{\infty} \alpha_k^2 < \infty.$$

Consequently, the step sizes $\{\alpha_k\}_k$ satisfy all standard requirements. In addition, we have

$$\sum_{k=0}^{\infty} \alpha_k |f'(\mathbf{x}^k)|^2 \leq \sum_{k=0}^{\infty} \frac{1}{(k+1)(k+2)} < \infty \quad \text{and} \quad \min_{k=0, \dots, T} |f'(\mathbf{x}^k)|^2 \leq \frac{1}{T^2}.$$

Thus, the complexity results in (38) obviously hold, but the gradient values $f'(\mathbf{x}^k)$ do not converge to zero. In particular, for even T , the last iterate \mathbf{x}^T satisfies $f'(\mathbf{x}^T) = 1$ which obstructs interpretability of \mathbf{x}^T and of the complexity bounds $\min_{k=0, \dots, T} |f'(\mathbf{x}^k)|^2 \leq \varepsilon$ or $|f'(\mathbf{x}^{\bar{k}})|^2 \leq \varepsilon$. Notice that the mapping f is not Lipschitz smooth around $x = 0$ and hence, the convergence results in [Theorem 2.1](#) are not applicable here. Of course, these observations have even higher significance in the stochastic setting, when evaluation of the bounds (37) is generally restrictive or not possible within the algorithmic procedure.

We conclude and summarize our discussion with a comment by Francesco Orabona on non-asymptotic and asymptotic convergence analyses for SGD ([\[33\]](#), blog post: “*Almost sure convergence of SGD on smooth non-convex functions*”, section 5 and 6, Oct. 05, 2020):

“Note that the 20-30 years ago there were many papers studying the asymptotic convergence of SGD and its variants in various settings. Then, the taste of the community changed moving from asymptotic convergence to finite-time rates. As it often happens when a new trend takes over the previous one, new generations tend to be oblivious to the old results and proof techniques. The common motivation to ignore these past results is that the finite-time analysis is superior to the asymptotic one, but this is clearly false (ask a statistician!). It should be instead clear to anyone that both analyses have pros and cons.”