
PAGE: A Simple and Optimal Probabilistic Gradient Estimator for Nonconvex Optimization

Zhize Li¹ Hongyan Bao¹ Xiangliang Zhang¹ Peter Richtárik¹

Abstract

In this paper, we propose a novel stochastic gradient estimator—Probabilistic Gradient Estimator (PAGE)—for nonconvex optimization. PAGE is easy to implement as it is designed via a small adjustment to vanilla SGD: in each iteration, PAGE uses the vanilla minibatch SGD update with probability p_t or reuses the previous gradient with a small adjustment, at a much lower computational cost, with probability $1 - p_t$. We give a simple formula for the optimal choice of p_t . Moreover, we prove the first tight lower bound $\Omega(n + \frac{\sqrt{n}}{\epsilon^2})$ for nonconvex finite-sum problems, which also leads to a tight lower bound $\Omega(b + \frac{\sqrt{b}}{\epsilon^2})$ for nonconvex online problems, where $b := \min\{\frac{\sigma^2}{\epsilon^2}, n\}$. Then, we show that PAGE obtains the optimal convergence results $O(n + \frac{\sqrt{n}}{\epsilon^2})$ (finite-sum) and $O(b + \frac{\sqrt{b}}{\epsilon^2})$ (online) matching our lower bounds for both nonconvex finite-sum and online problems. Besides, we also show that for nonconvex functions satisfying the Polyak-Łojasiewicz (PL) condition, PAGE can automatically switch to a faster linear convergence rate $O(\cdot \log \frac{1}{\epsilon})$. Finally, we conduct several deep learning experiments (e.g., LeNet, VGG, ResNet) on real datasets in PyTorch showing that PAGE not only converges much faster than SGD in training but also achieves the higher test accuracy, validating the optimal theoretical results and confirming the practical superiority of PAGE.

1. Introduction

Nonconvex optimization is ubiquitous across many domains of machine learning, including robust regression, low rank matrix recovery, sparse recovery and supervised learning

¹King Abdullah University of Science and Technology, Thuwal, Kingdom of Saudi Arabia. Correspondence to: Zhize Li <zhize.li@kaust.edu.sa>.

(Jain & Kar, 2017). Driven by the applied success of deep neural networks (LeCun et al., 2015), and the critical place nonconvex optimization plays in training them, research in nonconvex optimization has been undergoing a renaissance (Ghadimi & Lan, 2013; Ghadimi et al., 2016; Zhou et al., 2018; Fang et al., 2018; Li, 2019; Li & Richtárik, 2020).

1.1. The problem

Motivated by this development, we consider the general optimization problem

$$\min_{x \in \mathbb{R}^d} f(x), \quad (1)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a differentiable and possibly nonconvex function. We are interested in functions having the *finite-sum* form

$$f(x) := \frac{1}{n} \sum_{i=1}^n f_i(x), \quad (2)$$

where the functions f_i are also differentiable and possibly nonconvex. Form (2) captures the standard empirical risk minimization problems in machine learning (Shalev-Shwartz & Ben-David, 2014). Moreover, if the number of data samples n is very large or even infinite, e.g., in the online/streaming case, then $f(x)$ usually is modeled via the *online* form

$$f(x) := \mathbb{E}_{\zeta \sim \mathcal{D}}[F(x, \zeta)], \quad (3)$$

which we also consider in this work. For notational convenience, we adopt the notation of the finite-sum form (2) in the descriptions and algorithms in the rest of this paper. However, our results apply to the online form (3) as well by letting $f_i(x) := F(x, \zeta_i)$ and treating n as a very large value or even infinite.

1.2. Gradient complexity

To measure the efficiency of algorithms for solving the nonconvex optimization problem (1), it is standard to bound the number of stochastic gradient computations needed to find a solution of suitable characteristics. In this paper we

use the standard term *gradient complexity* to describe such bounds. In particular, our goal will be to find a (possibly random) point $\hat{x} \in \mathbb{R}^d$ such that $\mathbb{E}[\|\nabla f(\hat{x})\|] \leq \epsilon$, where the expectation is with respect to the randomness inherent in the algorithm. We use the term ϵ -approximate solution to refer to such a point \hat{x} .

Two of the most classical gradient complexity results for solving problem (1) are those for gradient descent (GD) and stochastic gradient descent (SGD). In particular, the gradient complexity of GD is $O(n/\epsilon^2)$ in this nonconvex regime, and assuming that the stochastic gradient satisfies a (uniform) bounded variance assumption (Assumption 1), the gradient complexity of SGD is $O(1/\epsilon^4)$. Note that although SGD has a worse dependence on ϵ , it typically only needs to compute a constant minibatch of stochastic gradients in each iteration instead of the full batch (i.e., n stochastic gradients) used in GD. Hence, SGD is better than GD if the number of data samples n is very large or the error tolerance ϵ is not very small.

There has been extensive research in designing gradient-type methods with an improved dependence on n and/or ϵ (Nesterov, 2004; Nemirovski et al., 2009; Ghadimi & Lan, 2013; Ghadimi et al., 2016). In particular, the SVRG method of Johnson & Zhang (2013), the SAGA method of Defazio et al. (2014) and the SARAH method of Nguyen et al. (2017) are representatives of what is by now a large class of *variance-reduced* methods, which have played a particularly important role in this effort. However, the analyses in these papers focused on the convex regime. Furthermore, several accelerated (momentum) methods have been designed as well (Nesterov, 1983; Lan & Zhou, 2015; Lin et al., 2015; Lan & Zhou, 2018; Allen-Zhu, 2017; Li & Li, 2020; Lan et al., 2019; Li et al., 2020b; Li, 2021), with or without variance reduction. There are also some lower bounds given by (Lan & Zhou, 2015; Woodworth & Srebro, 2016; Xie et al., 2019).

Coming back to problem (1) in the nonconvex regime studied in this paper, interesting recent development starts with the work of Reddi et al. (2016), and Allen-Zhu & Hazan (2016), who have concurrently shown that if f has the finite-sum form (2), a suitably designed minibatch version of SVRG enjoys the gradient complexity $O(n + n^{2/3}/\epsilon^2)$, which is an improvement on the $O(n/\epsilon^2)$ gradient complexity of GD. Subsequently, other variants of SVRG were shown to possess the same improved rate, including those developed by (Lei et al., 2017; Li & Li, 2018; Ge et al., 2019; Horváth & Richtárik, 2019; Qian et al., 2019). More recently, Fang et al. (2018) proposed the SPIDER method, and Zhou et al. (2018) proposed the SNVRG method, both of them further improve the gradient complexity to $O(n + \sqrt{n}/\epsilon^2)$. Further variants of the SARAH method (e.g., Wang et al., 2018; Li, 2019; Pham et al., 2019; Li et al., 2020a; Horváth et al.,

2020; Li & Richtárik, 2021) which also achieve the same $O(n + \sqrt{n}/\epsilon^2)$ gradient complexity have been developed. Also there are some lower bounds given by (Fang et al., 2018; Zhou & Gu, 2019; Arjevani et al., 2019). See Table 1 for an overview of results.

2. Our Contributions

As we show in through this work, despite enormous effort by the community to design efficient methods for solving (1) in the nonconvex regime, there is still a considerable gap in our understanding. First, while optimal methods for (1) in the finite-sum regime exist (e.g., SPIDER (Fang et al., 2018), SpiderBoost (Wang et al., 2018), SARAH (Pham et al., 2019), SSRGD (Li, 2019)), the known lower bound $\Omega(\sqrt{n}/\epsilon^2)$ (Fang et al., 2018) used to establish their optimality works only for $n \leq O(1/\epsilon^4)$, i.e., in the small data regime (see Table 1). Moreover, these methods are unnecessarily complicated, often with a double loop structure, and reliance on several hyperparameters. Besides, there is also no tight lower bound to show the optimality of optimal methods in the online regime.

In this paper, we resolve the above issues by designing a simple Probabilistic Gradient Estimator (PAGE) described in Algorithm 1 for achieving optimal convergence results in nonconvex optimization. Moreover, PAGE is very simple and easy to implement. In each iteration, PAGE uses minibatch SGD update with probability p_t , or reuses the previous gradient with a small adjustment (at a low computational cost) with probability $1 - p_t$ (see Line 4 of Algorithm 1). We would like to highlight the following results:

- We provide tight lower bounds to close the gap for both nonconvex finite-sum problem (2) and online problem (3) (see Theorem 2 and Corollary 5). Our lower bounds are based and inspired by recent work (Fang et al., 2018; Arjevani et al., 2019). Then we show the optimality of PAGE by proving that PAGE achieves the optimal convergence results matching our lower bounds for both nonconvex finite-sum problem (2) and online problem (3) (see Corollaries 2 and 4). See Table 1 for a detailed comparison.

- Moreover, we show that PAGE can automatically switch to a faster *linear convergence* $O(\cdot \log \frac{1}{\epsilon})$ by exploiting the local structure of the objective function, via the PL condition (Assumption 3), although the objective function f is globally nonconvex. See the middle and the last row of Table 1 (highlighted with green color). For example, PAGE automatically switches from the sublinear rate $O(n + \sqrt{n}/\epsilon^2)$ to the faster linear rate $O((n + \sqrt{n}\kappa) \log \frac{1}{\epsilon})$ for nonconvex finite-sum problem (2).

- PAGE is simple and easy to implement via a small adjustment to vanilla minibatch SGD, and takes a lower computational cost than SGD (i.e., $p_t = 1$ in Algorithm 1) since

Table 1. Stochastic gradient complexity for finding an ϵ -approximate solution $\mathbb{E}[\|\nabla f(\hat{x})\|] \leq \epsilon$ for nonconvex problems

Problem	Assumption	Algorithm or Lower Bound	Gradient complexity
Finite-sum (2)	Asp. 2	GD (Nesterov, 2004)	$O(\frac{n}{\epsilon^2})$
Finite-sum (2)	Asp. 2	SVRG (Allen-Zhu & Hazan, 2016; Reddi et al., 2016) SCSG (Lei et al., 2017), SVRG+ (Li & Li, 2018)	$O(n + \frac{n^{2/3}}{\epsilon^2})$
Finite-sum (2)	Asp. 2	SNVRG (Zhou et al., 2018), Geom-SARAH (Horváth et al., 2020)	$\tilde{O}\left(n + \frac{\sqrt{n}}{\epsilon^2}\right)$
Finite-sum (2)	Asp. 2	SPIDER (Fang et al., 2018), SpiderBoost (Wang et al., 2018), SARAH (Pham et al., 2019), SSRGD (Li, 2019)	$O(n + \frac{\sqrt{n}}{\epsilon^2})$
Finite-sum (2)	Asp. 2	PAGE (this paper)	$O(n + \frac{\sqrt{n}}{\epsilon^2})$
Finite-sum (2)	Asp. 2	Lower bound (Fang et al., 2018)	$\Omega(\frac{\sqrt{n}}{\epsilon^2})$ if $n \leq O(\frac{1}{\epsilon^4})$
Finite-sum (2)	Asp. 2	Lower bound (this paper)	$\Omega(n + \frac{\sqrt{n}}{\epsilon^2})$
Finite-sum (2)	Asp. 2 and 3 (PL setting)	PAGE (this paper)	$O\left((n + \sqrt{n\kappa}) \log \frac{1}{\epsilon}\right)^a$
Online (3) ^b	Asp. 1 and 2	SGD (Ghadimi et al., 2016; Khaled & Richtárik, 2020; Li & Richtárik, 2020)	$O(\frac{\sigma^2}{\epsilon^4})$
Online (3)	Asp. 1 and 2	SCSG (Lei et al., 2017), SVRG+ (Li & Li, 2018)	$O(b + \frac{b^{2/3}}{\epsilon^2})$
Online (3)	Asp. 1 and 2	SNVRG (Zhou et al., 2018), Geom-SARAH (Horváth et al., 2020)	$\tilde{O}\left(b + \frac{\sqrt{b}}{\epsilon^2}\right)$
Online (3)	Asp. 1 and 2	SPIDER (Fang et al., 2018), SpiderBoost (Wang et al., 2018), SARAH (Pham et al., 2019), SSRGD (Li, 2019)	$O(b + \frac{\sqrt{b}}{\epsilon^2})$
Online (3)	Asp. 1 and 2	PAGE (this paper)	$O(b + \frac{\sqrt{b}}{\epsilon^2})^c$
Online (3)	Asp. 1 and 2	Lower bound (this paper)	$\Omega(b + \frac{\sqrt{b}}{\epsilon^2})$
Online (3)	Asp. 1, 2 and 3 (PL setting)	PAGE (this paper)	$O\left((b + \sqrt{b\kappa}) \log \frac{1}{\epsilon}\right)$

^aNote that PAGE can switch to a faster *linear convergence* $O(\cdot \log \frac{1}{\epsilon})$ instead of sublinear rate $O(\cdot \frac{1}{\epsilon^2})$ by exploiting the local structure of the objective function via the PL condition (Assumption 3).

^bNote that we refer the online problem (3) as the finite-sum problem (2) with large or infinite n as discussed in the introduction Section 1.1. In this online case, the full gradient may not be available (e.g., if n is infinite), thus the bounded variance of stochastic gradient Assumption 1 is needed in this case.

^cIn the online case, $b := \min\{\frac{\sigma^2}{\epsilon^2}, n\}$, and σ is defined in Assumption 1. If n is very large, i.e., $b := \min\{\frac{\sigma^2}{\epsilon^2}, n\} = \frac{\sigma^2}{\epsilon^2}$, then $O(b + \frac{\sqrt{b}}{\epsilon^2}) = O(\frac{\sigma^2}{\epsilon^2} + \frac{\sigma}{\epsilon^3})$ is better than the rate $O(\frac{\sigma^2}{\epsilon^4})$ of SGD by a factor of $\frac{1}{\epsilon^2}$ or $\frac{\sigma}{\epsilon}$.

$b' < b$. We conduct several deep learning experiments (e.g., LeNet, VGG, ResNet) on real datasets in PyTorch showing that PAGE indeed not only converges much faster than SGD in training but also achieves higher test accuracy. This validates our theoretical results and confirms the practical superiority of PAGE.

2.1. The PAGE gradient estimator

In this section, we describe PAGE, an SGD variant employing a new, simple and optimal gradient estimator (see Algorithm 1). In particular, PAGE was inspired by algorithmic design elements coming from methods such as SARAH (Nguyen et al., 2017), SPIDER (Fang et al., 2018), SSRGD (Li, 2019) (usage of a recursive estimator), and L-SVRG (Kovalev et al., 2020) and SAGD (Bibi et al., 2018) (probabilistic switching between two estimators to avoid a double loop structure).

bilistic switching between two estimators to avoid a double loop structure).

In each iteration, the gradient estimator g^{t+1} of PAGE is defined in Line 4 of Algorithm 1, which indicates that PAGE uses the vanilla minibatch SGD update with probability p_t , and reuses the previous gradient g^t with a small adjustment (which lowers the computational cost since $b' \ll b$) with probability $1 - p_t$. In particular, the $p_t \equiv 1$ case reduces to vanilla minibatch SGD, and to GD if we further set the minibatch size to $b = n$. We give a simple formula for the optimal choice of p_t , i.e., $p_t \equiv \frac{b'}{b+b'}$ is enough for PAGE to obtain the optimal convergence rates. More details can be found in the convergence results of Section 4.

Note that PAGE with constant probability $p_t \equiv p$ can be reduced to an equivalent form of the double loop algorithm

Algorithm 1 Probabilistic Gradient Estimator (PAGE)

Input: initial point x^0 , stepsize η , minibatch size b , $b' < b$, probability $\{p_t\} \in (0, 1]$
 1: $g^0 = \frac{1}{b} \sum_{i \in I} \nabla f_i(x^0)$ // I denotes random minibatch samples with $|I| = b$
 2: **for** $t = 0, 1, 2, \dots$ **do**
 3: $x^{t+1} = x^t - \eta g^t$
 4: $g^{t+1} = \begin{cases} \frac{1}{b} \sum_{i \in I} \nabla f_i(x^{t+1}) & \text{with probability } p_t \\ g^t + \frac{1}{b'} \sum_{i \in I'} (\nabla f_i(x^{t+1}) - \nabla f_i(x^t)) & \text{with probability } 1 - p_t \end{cases}$
 5: **end for**
Output: \hat{x}_T chosen uniformly from $\{x^t\}_{t \in [T]}$

with geometric distribution Geom-SARAH (Horváth et al., 2020), but our single-loop PAGE is more flexible and also leads to simpler and better analysis. Similar to L-SVRG (Kovalev et al., 2020) which switches between GD and SVRG probabilistically, L2S (Li et al., 2020a) switches between GD and SARAH and uses a fixed probability p (i.e., equivalent to Geom-SARAH (Horváth et al., 2020)). However, PAGE is more general which switches between minibatch SGD and minibatch SARAH and also allows a flexible probability p_t . More importantly, the minibatch SGD update instead of GD can allow PAGE to solve both nonconvex finite-sum and online problems, while L2S (Li et al., 2020a) can only deal with the finite-sum case. Besides, our convergence analysis of PAGE is simple and clean, which is totally different from L2S (Li et al., 2020a). Concretely, our analysis of PAGE directly shows the decrease for each iteration (see (19) or (22)), i.e., *truly loopless analysis*. However, L2S (Li et al., 2020a) still uses a *double loop analysis* where they transform the probabilistic switch steps to an equivalent double loop structure and upper bound the variance term by considering all inner loop iterations together not just one iteration as ours (see Lemma 5 of L2S vs. our Lemma 3).

3. Notation and Assumptions

Let $[n]$ denote the set $\{1, 2, \dots, n\}$ and $\|\cdot\|$ denote the Euclidean norm for a vector and the spectral norm for a matrix. Let $\langle u, v \rangle$ denote the inner product of two vectors u and v . We use $O(\cdot)$ and $\Omega(\cdot)$ to hide the absolute constant, and $\tilde{O}(\cdot)$ to hide the logarithmic factor. We will write $\Delta_0 := f(x^0) - f^*$ and $f^* := \min_{x \in \mathbb{R}^d} f(x)$.

In order to prove convergence results, one usually needs the following standard assumptions depending on the setting (see e.g., Ghadimi et al., 2016; Lei et al., 2017; Li & Li, 2018; Allen-Zhu, 2018; Zhou et al., 2018; Fang et al., 2018).

Assumption 1 (Bounded variance) *The stochastic gradient has bounded variance if $\exists \sigma > 0$, such that*

$$\mathbb{E}_i[\|\nabla f_i(x) - \nabla f(x)\|^2] \leq \sigma^2, \quad \forall x \in \mathbb{R}^d. \quad (4)$$

Assumption 2 (Average L -smoothness) *A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is average L -smooth if $\exists L > 0$,*

$\mathbb{E}_i[\|\nabla f_i(x) - \nabla f_i(y)\|^2] \leq L^2 \|x - y\|^2, \quad \forall x, y \in \mathbb{R}^d. \quad (5)$

Moreover, we also prove faster linear convergence rates for nonconvex functions under the Polyak-Łojasiewicz (PL) condition (Polyak, 1963).

Assumption 3 (PL condition) *A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies PL condition¹ if $\exists \mu > 0$, such that*

$$\|\nabla f(x)\|^2 \geq 2\mu(f(x) - f^*), \quad \forall x \in \mathbb{R}^d. \quad (6)$$

4. General Convergence Results

In this section, we present two main convergence theorems for PAGE (Algorithm 1): i) for nonconvex finite-sum problem (2) (Section 4.1), and ii) for nonconvex online problem (3) (Section 4.2). Subsequently, we formulate several corollaries which lead to the optimal convergence results. Finally, we provide tight lower bounds for both types of nonconvex problems to close the gap and validate the optimality of PAGE. See Table 1 for an overview.

4.1. Convergence for nonconvex finite-sum problems

In this section, we focus on the nonconvex finite-sum problems defined via (2). In this case, we do not need the bounded variance assumption (Assumption 1).

Theorem 1 (Nonconvex finite-sum problem (2))

Suppose that Assumption 2 holds. Choose the stepsize $\eta \leq \frac{1}{L(1 + \sqrt{\frac{1-p}{pb'}})}$, minibatch size $b = n$, secondary minibatch size $b' < b$, and probability $p_t \equiv p \in (0, 1]$. Then the number of iterations performed by PAGE sufficient for finding an ϵ -approximate solution (i.e., $\mathbb{E}[\|\nabla f(\hat{x}_T)\|] \leq \epsilon$) of nonconvex finite-sum problem (2) can be bounded by

$$T = \frac{2\Delta_0 L}{\epsilon^2} \left(1 + \sqrt{\frac{1-p}{pb'}}\right).$$

¹It is worth noting that the PL condition does not imply convexity of f . For example, $f(x) = x^2 + 3\sin^2 x$ is a nonconvex function but it satisfies PL condition with $\mu = 1/32$.

Moreover, according to the gradient estimator of PAGE (Line 4 of Algorithm 1), we know that it uses $pb + (1 - p)b'$ stochastic gradients for each iteration on expectation. Thus, the number of stochastic gradient computations (i.e., gradient complexity) is

$$\begin{aligned} \#\text{grad} &= b + T(pb + (1 - p)b') \\ &= b + \frac{2\Delta_0 L}{\epsilon^2} \left(1 + \sqrt{\frac{1-p}{pb'}}\right) (pb + (1 - p)b'). \end{aligned}$$

Note that the first b in $\#\text{grad}$ is due to the computation of g^0 (see Line 1 in Algorithm 1).

As we mentioned before, if we choose $p_t \equiv 1$ and $b = n$ (see Line 4 of Algorithm 1), PAGE reduces to the vanilla GD method. We now show that our main theorem indeed recovers the convergence result of GD.

Corollary 1 (We recover GD by letting $p_t \equiv 1$) Suppose that Assumption 2 holds. Choose the stepsize $\eta \leq \frac{1}{L}$, minibatch size $b = n$ and probability $p_t \equiv 1$. Then PAGE reduces to GD, and the number of iterations performed by PAGE to find an ϵ -approximate solution of the nonconvex finite-sum problem (2) can be bounded by $T = \frac{2\Delta_0 L}{\epsilon^2}$. Moreover, the number of stochastic gradient computations (i.e., gradient complexity) is

$$\#\text{grad} = n + \frac{2\Delta_0 L n}{\epsilon^2} = O\left(\frac{n}{\epsilon^2}\right).$$

Next, we provide a parameter setting that leads to the optimal convergence result for nonconvex finite-sum problem (2), which corresponds to the 6th row of Table 1. Note that a fixed p_t is enough for PAGE to obtain the optimal convergence result although people can choose different p_t in practice.

Corollary 2 (Optimal result for problem (2)) Suppose that Assumption 2 holds. Choose the stepsize $\eta \leq \frac{1}{L(1+\sqrt{b/b'})}$, minibatch size $b = n$, secondary minibatch size $b' \leq \sqrt{b}$ and probability $p_t \equiv \frac{b'}{b+b'}$. Then the number of iterations performed by PAGE to find an ϵ -approximate solution of the nonconvex finite-sum problem (2) can be bounded by $T = \frac{2\Delta_0 L}{\epsilon^2} \left(1 + \frac{1}{\sqrt{b/b'}}\right)$. Moreover, the number of stochastic gradient computations (i.e., gradient complexity) is

$$\#\text{grad} \leq n + \frac{8\Delta_0 L \sqrt{n}}{\epsilon^2} = O\left(n + \frac{\sqrt{n}}{\epsilon^2}\right).$$

Finally, we establish a lower bound matching the above upper bound, which shows that the convergence result obtained by PAGE in Corollary 2 is indeed optimal. This lower bound corresponds to the 8th row of Table 1.

Theorem 2 (Lower bound) For any $L > 0$, $\Delta_0 > 0$ and $n > 0$, there exists a large enough dimension d and a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying Assumption 2 in the finite-sum case such that any linear-span first-order algorithm needs $\Omega\left(n + \frac{\Delta_0 L \sqrt{n}}{\epsilon^2}\right)$ stochastic gradient computations in order to finding an ϵ -approximate solution, i.e., a point \hat{x} such that $\mathbb{E}\|\nabla f(\hat{x})\| \leq \epsilon$.

4.2. Convergence for nonconvex online problems

In this section, we focus on the nonconvex online problems, i.e., (3). Recall that we refer this online problem (3) as the finite-sum problem (2) with large or infinite n . Also, we need the bounded variance assumption (Assumption 1) in this online case. Similarly, we first present the main theorem in this online case and then provide corollaries with the optimal convergence results. Finally, we provide tight lower bound for validating the optimality of PAGE.

Theorem 3 (Nonconvex online problem (3)) Suppose that Assumptions 1 and 2 hold. Choose the stepsize $\eta \leq \frac{1}{L(1+\sqrt{\frac{1-p}{pb'}})}$, minibatch size $b = \min\{\lceil \frac{2\sigma^2}{\epsilon^2} \rceil, n\}$, secondary minibatch size $b' < b$ and probability $p_t \equiv p \in (0, 1]$. Then the number of iterations performed by PAGE to find an ϵ -approximate solution ($\mathbb{E}\|\nabla f(\hat{x}_T)\| \leq \epsilon$) of nonconvex online problem (3) can be bounded by

$$T = \frac{4\Delta_0 L}{\epsilon^2} \left(1 + \sqrt{\frac{1-p}{pb'}}\right) + \frac{1}{p}.$$

Moreover, the number of stochastic gradient computations (gradient complexity) $\#\text{grad} = b + T(pb + (1 - p)b')$ is

$$2b + \frac{(1-p)b'}{p} + \frac{4\Delta_0 L}{\epsilon^2} \left(1 + \sqrt{\frac{1-p}{pb'}}\right) (pb + (1 - p)b').$$

Similarly, if we choose $p_t \equiv 1$ (see Line 4 of Algorithm 1), the PAGE method reduces to the vanilla minibatch SGD method. Here we theoretically show that our main theorem with $p_t \equiv 1$ can recover the convergence result of SGD in the following Corollary 3.

Corollary 3 (We recover SGD by letting $p_t \equiv 1$) Suppose that Assumptions 1 and 2 hold. Let stepsize $\eta \leq \frac{1}{L}$, minibatch size $b = \lceil \frac{2\sigma^2}{\epsilon^2} \rceil$ and probability $p_t \equiv 1$, then the number of iterations performed by PAGE to find an ϵ -approximate solution of nonconvex online problem (3) can be bounded by $T = \frac{4\Delta_0 L}{\epsilon^2} + 1$. Moreover, the number of stochastic gradient computations (gradient complexity) is

$$\#\text{grad} = \frac{4\sigma^2}{\epsilon^2} + \frac{8\Delta_0 L \sigma^2}{\epsilon^4} = O\left(\frac{\sigma^2}{\epsilon^4}\right).$$

Now, we provide a parameter setting that leads to the optimal convergence result of our main theorem for nonconvex

online problem (3), which corresponds to the 14th row of Table 1. Similarly, a fixed p_t is enough for PAGE to obtain the optimal convergence result in this online case.

Corollary 4 (Optimal result for problem (3)) *Suppose that Assumptions 1 and 2 hold. Choose the stepsize $\eta \leq \frac{1}{L(1+\sqrt{b}/b')}$, minibatch size $b = \min\{\lceil \frac{2\sigma^2}{\epsilon^2} \rceil, n\}$, secondary minibatch size $b' \leq \sqrt{b}$ and probability $p_t \equiv \frac{b'}{b+b'}$. Then the number of iterations performed by PAGE sufficient to find an ϵ -approximate solution of nonconvex online problem (3) can be bounded by $T = \frac{4\Delta_0 L}{\epsilon^2} (1 + \frac{\sqrt{b}}{b'}) + \frac{b+b'}{b'}$. Moreover, the number of stochastic gradient computations (i.e., gradient complexity) is*

$$\#\text{grad} \leq 3b + \frac{16\Delta_0 L \sqrt{b}}{\epsilon^2} = O\left(b + \frac{\sqrt{b}}{\epsilon^2}\right).$$

Before we provide our lower bound, we first recall the lower bound established by Arjevani et al. (2019).

Theorem 4 (Arjevani et al., 2019) *For any $L > 0$, $\Delta_0 > 0$ and $\sigma^2 > 0$, there exists a large enough dimension d and function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying Assumptions 1 and 2 in the online case (here n is infinite) such that any linear-span first-order algorithm needs $\Omega(\frac{\sigma^2}{\epsilon^2} + \frac{\Delta_0 L \sigma}{\epsilon^3})$ stochastic gradient computations in order to find an ϵ -approximate solution, i.e., a point \hat{x} such that $\mathbb{E}\|\nabla f(\hat{x})\| \leq \epsilon$.*

Now, we provide a lower bound corollary which directly follows from the lower bound Theorem 4 given by Arjevani et al. (2019) and our Theorem 2. It indicates that the convergence result obtained by PAGE in Corollary 4 is indeed optimal.

Corollary 5 (Lower bound) *For any $L > 0$, $\Delta_0 > 0$, $\sigma^2 > 0$ and $n > 0$, there exists a large enough dimension d and a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying Assumptions 1 and 2 in the online case (here n may be finite) such that any linear-span first-order algorithm needs $\Omega(b + \frac{\Delta_0 L \sqrt{b}}{\epsilon^2})$, where $b = \min\{\frac{\sigma^2}{\epsilon^2}, n\}$, stochastic gradient computations for finding an ϵ -approximate solution, i.e., a point \hat{x} such that $\mathbb{E}\|\nabla f(\hat{x})\| \leq \epsilon$.*

5. Better Convergence under PL Condition

In this section, we show that better convergence can be achieved if the loss function f satisfies the PL condition (Assumption 3). Note that under the PL condition, one can obtain a faster linear convergence $O(\cdot \log \frac{1}{\epsilon})$ (see Corollary 6) rather than the sublinear convergence $O(\cdot \frac{1}{\epsilon^2})$ (see Corollary 2). In many cases, although the loss function f is globally nonconvex, some local regions (e.g., large gradient regions) may satisfy the PL condition. We prove that PAGE can automatically switch to the faster convergence rate in these regions where f satisfies PL condition locally.

As in Section 4, here we also establish two main theorems and deduce corollaries for both finite-sum and online regimes. The convergence results are also listed in Table 1 (i.e., the middle row and last row).

Theorem 5 (Nonconvex finite-sum problem (2) under PL)

Suppose that Assumptions 2 and 3 hold. Choose the stepsize $\eta \leq \min\{\frac{1}{L(1+\sqrt{\frac{1-p}{pb'}})}, \frac{p}{2\mu}\}$, minibatch size $b = n$, secondary minibatch size $b' < b$, and probability $p_t \equiv p \in (0, 1]$. Then the number of iterations performed by PAGE sufficient for finding an ϵ -solution ($\mathbb{E}[f(x^T) - f^] \leq \epsilon$) of nonconvex finite-sum problem (2) can be bounded by*

$$T = \left(\left(1 + \sqrt{\frac{1-p}{pb'}} \right) \kappa + \frac{2}{p} \right) \log \frac{\Delta_0}{\epsilon}.$$

Moreover, the number of stochastic gradient computations (i.e., gradient complexity) $\#\text{grad} = b + T(pb + (1-p)b')$ is

$$b + (pb + (1-p)b') \left(\left(1 + \sqrt{\frac{1-p}{pb'}} \right) \kappa + \frac{2}{p} \right) \log \frac{\Delta_0}{\epsilon}.$$

Corollary 6

Suppose that Assumptions 2 and 3 hold. Let stepsize $\eta \leq \min\{\frac{1}{L(1+\sqrt{b}/b')}, \frac{b'}{2\mu(b+b')}\}$, minibatch size $b = n$, secondary minibatch size $b' \leq \sqrt{b}$, and probability $p_t \equiv \frac{b'}{b+b'}$. Then the number of iterations performed by PAGE to find an ϵ -solution of nonconvex finite-sum problem (2) can be bounded by $T = \left((1 + \frac{\sqrt{b}}{b'})\kappa + \frac{2(b+b')}{b'} \right) \log \frac{\Delta_0}{\epsilon}$. Moreover, the number of stochastic gradient computations (gradient complexity) is

$$\#\text{grad} = O\left((n + \sqrt{n}\kappa) \log \frac{1}{\epsilon} \right).$$

Remark: Note that Corollary 6 uses exactly the same parameter setting as in Corollary 2 in the large condition number case (i.e., $\kappa := \frac{L}{\mu} \geq 2\sqrt{n}$, then the stepsize turns to $\eta \leq \frac{1}{L(1+\sqrt{b}/b')}$). Thus, PAGE can automatically switch to this faster linear convergence rate $O(\cdot \log \frac{1}{\epsilon})$ instead of the sublinear convergence $O(\cdot \frac{1}{\epsilon^2})$ in Corollary 2 in some regions where f satisfies the PL condition locally.

Theorem 6 (Nonconvex online problem (3) under PL)

Suppose that Assumptions 1, 2 and 3 hold. Choose the stepsize $\eta \leq \min\{\frac{1}{L(1+\sqrt{\frac{1-p}{pb'}})}, \frac{p}{2\mu}\}$, minibatch size $b = \min\{\lceil \frac{2\sigma^2}{\mu\epsilon} \rceil, n\}$, secondary minibatch size $b' < b$, and probability $p_t \equiv p \in (0, 1]$. Then the number of iterations performed by PAGE sufficient for finding an ϵ -solution ($\mathbb{E}[f(x^T) - f^] \leq \epsilon$) of nonconvex finite-sum problem (2) can be bounded by*

$$T = \left(\left(1 + \sqrt{\frac{1-p}{pb'}} \right) \kappa + \frac{2}{p} \right) \log \frac{2\Delta_0}{\epsilon}.$$

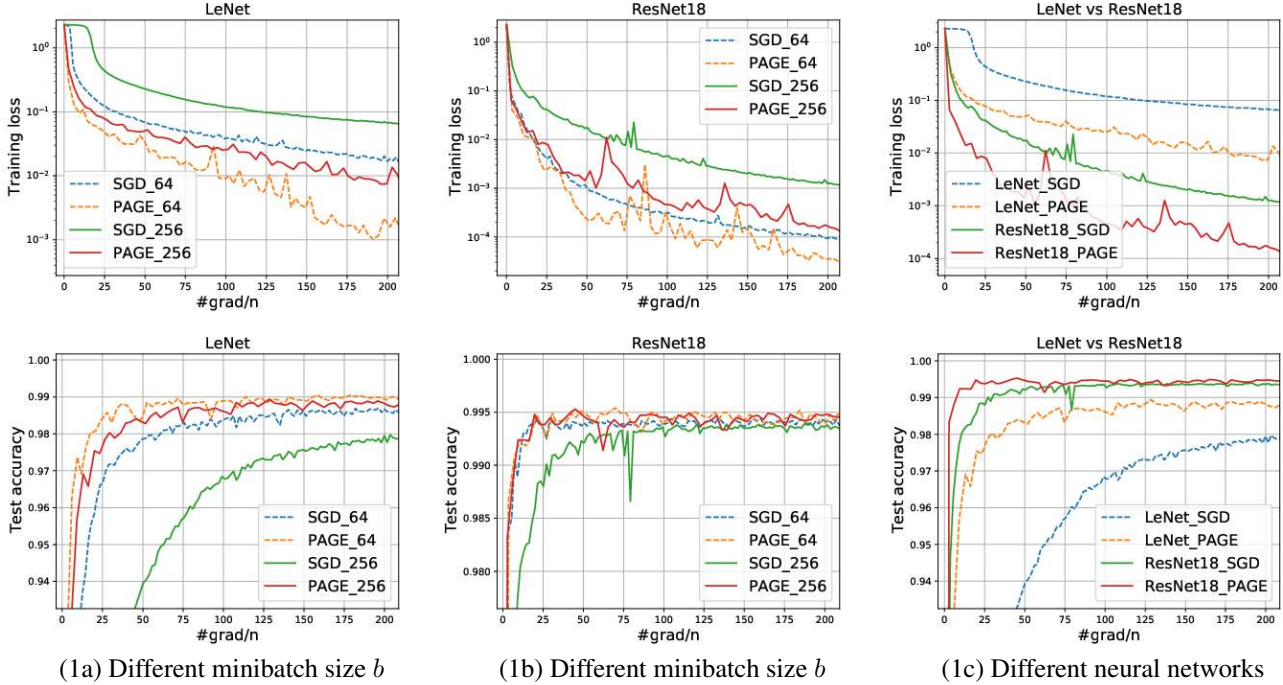


Figure 1. LeNet and ResNet18 on MNIST dataset

Moreover, the number of stochastic gradient computations (i.e., gradient complexity) $\#\text{grad} = b + T(pb + (1-p)b')$ is

$$b + (pb + (1-p)b') \left(\left(1 + \sqrt{\frac{1-p}{pb'}}\right) \kappa + \frac{2}{p} \right) \log \frac{2\Delta_0}{\epsilon}.$$

Corollary 7 Suppose that Assumptions 1, 2 and 3 hold. Choose the stepsize $\eta \leq \min\{\frac{1}{L(1+\sqrt{b/b'})}, \frac{b'}{2\mu(b+b')}\}$, minibatch size $b = \min\{\lceil \frac{2\sigma^2}{\mu\epsilon} \rceil, n\}$, secondary minibatch $b' \leq \sqrt{b}$ and probability $p_t \equiv \frac{b'}{b+b'}$. Then the number of iterations performed by PAGE to find an ϵ -solution of nonconvex online problem (3) can be bounded by $T = \left((1 + \frac{\sqrt{b}}{b'})\kappa + \frac{2(b+b')}{b'} \right) \log \frac{2\Delta_0}{\epsilon}$. Moreover, the number of stochastic gradient computations (gradient complexity) is

$$\#\text{grad} = O\left((b + \sqrt{b}\kappa) \log \frac{1}{\epsilon} \right).$$

6. Experiments

In this section, we conduct several deep learning experiments for multi-class image classification. Concretely, we compare our PAGE algorithm with vanilla SGD by running standard LeNet (LeCun et al., 1998), VGG (Simonyan & Zisserman, 2014) and ResNet (He et al., 2016) models on MNIST (LeCun et al., 1998) and CIFAR-10 (Krizhevsky, 2009) datasets. We implement the algorithms in PyTorch (Paszke et al., 2019) and run the experiments on several NVIDIA Tesla V100 GPUs.

According to the update form in PAGE (see Line 4 of Algorithm 1), PAGE enjoys a lower computational cost than vanilla minibatch SGD (i.e., $p_t \equiv 1$ in PAGE) since $b' < b$. Thus, in the experiments we want to show how the performance of PAGE compares with vanilla minibatch SGD under different minibatch sizes b (i.e., $b = 64, 256, 512$). Note that we do not tune the parameters for PAGE, i.e., we set $b' = \sqrt{b}$ and $p_t \equiv \frac{b'}{b+b'} = \frac{\sqrt{b}}{b+\sqrt{b}}$ according to our theoretical results (see e.g., Corollary 2 and 4). For the stepsize/learning rate η , we choose the same one for both PAGE and minibatch SGD according to the theoretical results.

Concretely, in Figure 1, we choose standard minibatch $b = 64$ and $b = 256$ for both PAGE and vanilla minibatch SGD for MNIST experiments. In Figure 2, we choose $b = 256$ and $b = 512$ for CIFAR-10 experiments. The first row of Figures 1 and 2 denotes the training loss with respect to the gradient computations, and the second row denotes the test accuracy with respect to the gradient computations. Both Figures 1 and 2 demonstrate that PAGE not only converges much faster than SGD in training but also achieves higher test accuracy (which is typically very important in practice, e.g., lead to a better model). Moreover, the performance gap between PAGE and SGD is larger when the minibatch size b is larger (i.e, gap between solid lines in Figures 1a, 1b, 2a, 2b), which is consistent with the update form of PAGE, i.e, it reuses the previous gradient with a small adjustment (lower computational cost $b' = \sqrt{b}$ instead of b) with probability $1 - p_t$. The experimental results validate our theoretical

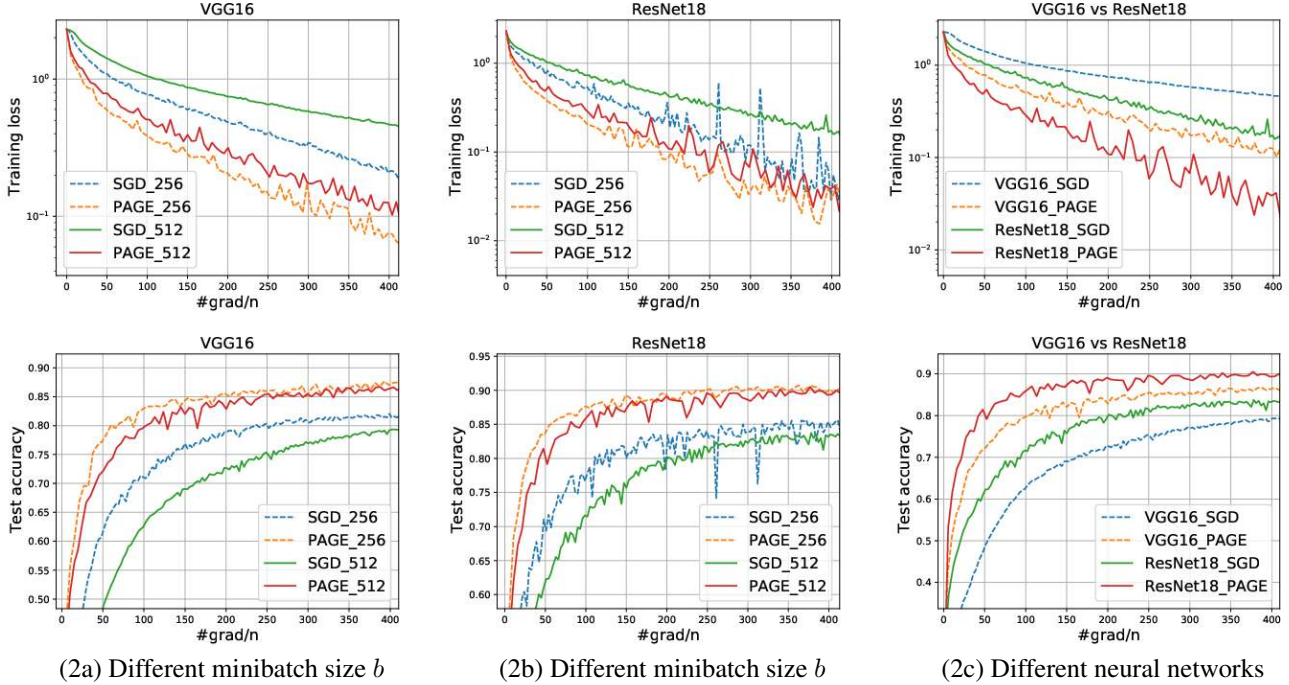


Figure 2. VGG16 and ResNet18 on CIFAR-10 dataset

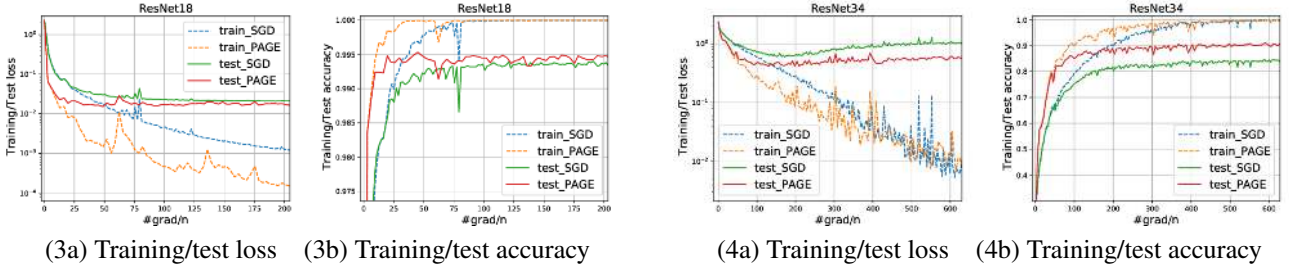


Figure 3. ResNet18 on MNIST dataset

Figure 4. ResNet34 on CIFAR-10 dataset

results and confirm the practical superiority of PAGE.

In the following, we conduct extra experiments for comparing the training loss and test loss (Figure 3a, 4a), and training accuracy and test accuracy (Figure 3b, 4b) between PAGE and SGD. Note that Figure 3 (i.e., 3a, 3b) uses MNIST dataset and Figure 4 (i.e., 4a, 4b) uses CIFAR-10 dataset. Figures (3a) and (4a) also demonstrate that PAGE converges much faster than SGD both in training loss and test loss. Moreover, Figures (3b) and (4b) demonstrate that PAGE achieves the higher test accuracy than SGD and converges faster in training accuracy. Thus, our PAGE is not only converging faster than SGD in training but also achieves the higher test accuracy (which is typically very important in practice, e.g., lead to a better model). Again, the experimental results validate our theoretical results and confirm the practical superiority of PAGE.

7. Conclusion

In this paper, we propose a simple and optimal PAGE algorithm for both nonconvex finite-sum and online optimization. We prove tight lower bounds and show that PAGE achieves the optimal convergence results matching our lower bounds for nonconvex functions satisfying the PL condition, PAGE can automatically switch to a faster linear convergence rate. Besides, PAGE is easy to implement and we conduct several deep learning experiments (e.g., LeNet, VGG, ResNet) in PyTorch which confirm the practical superiority of PAGE. More importantly, the novel convergence analysis of PAGE is very simple and clean. Thus PAGE and its analysis can be easily adopted and generalized to other works. In fact, it already leads to some further breakthroughs in communication-efficient distributed learning (e.g., Gorbunov et al., 2021; Richtárik et al., 2021).

Acknowledgements

Zhize Li and Peter Richtárik were supported by KAUST Baseline Research Fund. Hongyan Bao and Xiangliang Zhang were supported by KAUST Competitive Research Grant URF/1/3756-01-01.

References

- Allen-Zhu, Z. Katyusha: the first direct acceleration of stochastic gradient methods. In *ACM Symposium on Theory of Computing*, pp. 1200–1205. ACM, 2017.
- Allen-Zhu, Z. Natasha 2: Faster non-convex optimization than SGD. In *Advances in Neural Information Processing Systems*, pp. 2680–2691, 2018.
- Allen-Zhu, Z. and Hazan, E. Variance reduction for faster non-convex optimization. In *International Conference on Machine Learning*, pp. 699–707, 2016.
- Arjevani, Y., Carmon, Y., Duchi, J. C., Foster, D. J., Srebro, N., and Woodworth, B. Lower bounds for non-convex stochastic optimization. *arXiv preprint arXiv:1912.02365*, 2019.
- Bibi, A., Sailanbayev, A., Ghanem, B., Gower, R. M., and Richtárik, P. Improving SAGA via a probabilistic interpolation with gradient descent. *arXiv preprint arXiv:1806.05633*, 2018.
- Defazio, A., Bach, F., and Lacoste-Julien, S. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing Systems*, pp. 1646–1654, 2014.
- Fang, C., Li, C. J., Lin, Z., and Zhang, T. SPIDER: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. In *Advances in Neural Information Processing Systems*, pp. 687–697, 2018.
- Ge, R., Li, Z., Wang, W., and Wang, X. Stabilized SVRG: Simple variance reduction for nonconvex optimization. In *Conference on Learning Theory*, pp. 1394–1448, 2019.
- Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Ghadimi, S., Lan, G., and Zhang, H. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2):267–305, 2016.
- Gorbunov, E., Burlachenko, K., Li, Z., and Richtárik, P. MARINA: Faster non-convex distributed learning with compression. In *International Conference on Machine Learning*, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Horváth, S. and Richtárik, P. Nonconvex variance reduced optimization with arbitrary sampling. In *International Conference on Machine Learning*, pp. 2781–2789, 2019.
- Horváth, S., Lei, L., Richtárik, P., and Jordan, M. I. Adaptivity of stochastic gradient methods for nonconvex optimization. *arXiv preprint arXiv:2002.05359*, 2020.
- Jain, P. and Kar, P. Non-convex optimization for machine learning. *Foundations and Trends in Machine Learning*, 10(3-4):142–336, 2017.
- Johnson, R. and Zhang, T. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pp. 315–323, 2013.
- Khaled, A. and Richtárik, P. Better theory for SGD in the nonconvex world. *arXiv preprint arXiv:2002.03329*, 2020.
- Kovalev, D., Horváth, S., and Richtárik, P. Don’t jump through hoops and remove those loops: SVRG and Katyusha are better without the outer loop. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, 2020.
- Krizhevsky, A. Learning multiple layers of features from tiny images. 2009.
- Lan, G. and Zhou, Y. An optimal randomized incremental gradient method. *arXiv preprint arXiv:1507.02000*, 2015.
- Lan, G. and Zhou, Y. Random gradient extrapolation for distributed and stochastic optimization. *SIAM Journal on Optimization*, 28(4):2753–2782, 2018.
- Lan, G., Li, Z., and Zhou, Y. A unified variance-reduced accelerated gradient method for convex optimization. In *Advances in Neural Information Processing Systems*, pp. 10462–10472, 2019.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *Nature*, 521:436–444, 2015.
- Lei, L., Ju, C., Chen, J., and Jordan, M. I. Non-convex finite-sum optimization via SCSG methods. In *Advances in Neural Information Processing Systems*, pp. 2345–2355, 2017.

- Li, B., Ma, M., and Giannakis, G. B. On the convergence of sarah and beyond. In *International Conference on Artificial Intelligence and Statistics*, pp. 223–233, 2020a.
- Li, Z. SSRGD: Simple stochastic recursive gradient descent for escaping saddle points. In *Advances in Neural Information Processing Systems*, pp. 1521–1531, 2019.
- Li, Z. ANITA: An optimal loopless accelerated variance-reduced gradient method. *arXiv preprint arXiv:2103.11333*, 2021.
- Li, Z. and Li, J. A simple proximal stochastic gradient method for nonsmooth nonconvex optimization. In *Advances in Neural Information Processing Systems*, pp. 5569–5579, 2018.
- Li, Z. and Li, J. A fast Anderson-Chebyshev acceleration for nonlinear optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 1047–1057, 2020.
- Li, Z. and Richtárik, P. A unified analysis of stochastic gradient methods for nonconvex federated optimization. *arXiv preprint arXiv:2006.07013*, 2020.
- Li, Z. and Richtárik, P. ZeroSARAH: Efficient nonconvex finite-sum optimization with zero full gradient computation. *arXiv preprint arXiv:2103.01447*, 2021.
- Li, Z., Kovalev, D., Qian, X., and Richtárik, P. Acceleration for compressed gradient descent in distributed and federated optimization. In *International Conference on Machine Learning*, 2020b.
- Lin, H., Mairal, J., and Harchaoui, Z. A universal catalyst for first-order optimization. In *Advances in neural information processing systems*, pp. 3384–3392, 2015.
- Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.
- Nesterov, Y. A method for unconstrained convex minimization problem with the rate of convergence $o(1/k^2)$. In *Doklady AN USSR*, volume 269, pp. 543–547, 1983.
- Nesterov, Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, 2004.
- Nguyen, L. M., Liu, J., Scheinberg, K., and Takáč, M. SARAH: A novel method for machine learning problems using stochastic recursive gradient. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2613–2621. JMLR. org, 2017.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pp. 8024–8035, 2019.
- Pham, N. H., Nguyen, L. M., Phan, D. T., and Tran-Dinh, Q. ProxSARAH: An efficient algorithmic framework for stochastic composite nonconvex optimization. *arXiv preprint arXiv:1902.05679*, 2019.
- Polyak, B. T. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963.
- Qian, X., Qu, Z., and Richtárik, P. L-SVRG and L-Katyusha with arbitrary sampling. *arXiv preprint arXiv:1906.01481*, 2019.
- Reddi, S. J., Hefny, A., Sra, S., Póczos, B., and Smola, A. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*, pp. 314–323, 2016.
- Richtárik, P., Sokolov, I., and Fatkhullin, I. EF21: A new, simpler, theoretically better, and practically faster error feedback. *arXiv preprint arXiv:2106.05203*, 2021.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: from theory to algorithms*. Cambridge University Press, 2014.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Wang, Z., Ji, K., Zhou, Y., Liang, Y., and Tarokh, V. Spiderboost: A class of faster variance-reduced algorithms for nonconvex optimization. *arXiv preprint arXiv:1810.10690*, 2018.
- Woodworth, B. E. and Srebro, N. Tight complexity bounds for optimizing composite objectives. In *Advances in neural information processing systems*, pp. 3639–3647, 2016.
- Xie, G., Luo, L., and Zhang, Z. A general analysis framework of lower complexity bounds for finite-sum optimization. *arXiv preprint arXiv:1908.08394*, 2019.
- Zhou, D. and Gu, Q. Lower bounds for smooth nonconvex finite-sum optimization. *arXiv preprint arXiv:1901.11224*, 2019.
- Zhou, D., Xu, P., and Gu, Q. Stochastic nested variance reduction for nonconvex optimization. In *Advances in Neural Information Processing Systems*, pp. 3925–3936, 2018.