

STOCHASTIC QUASI-NEWTON METHODS FOR NONCONVEX STOCHASTIC OPTIMIZATION*

XIAO WANG[†], SHIQIAN MA[‡], DONALD GOLDFARB[§], AND WEI LIU[¶]

Abstract. In this paper we study stochastic quasi-Newton methods for nonconvex stochastic optimization, where we assume that noisy information about the gradients of the objective function is available via a stochastic first-order oracle (\mathcal{SFO}). We propose a general framework for such methods, for which we prove almost sure convergence to stationary points and analyze its worst-case iteration complexity. When a randomly chosen iterate is returned as the output of such an algorithm, we prove that in the worst case, the \mathcal{SFO} -calls complexity is $O(\epsilon^{-2})$ to ensure that the expectation of the squared norm of the gradient is smaller than the given accuracy tolerance ϵ . We also propose a specific algorithm, namely, a stochastic damped limited-memory BFGS (SdLBFGS) method, that falls under the proposed framework. Moreover, we incorporate the stochastic variance reduced gradient variance reduction technique into the proposed SdLBFGS method and analyze its \mathcal{SFO} -calls complexity. Numerical results on a nonconvex binary classification problem using a support vector machine and a multiclass classification problem using neural networks are reported.

Key words. nonconvex stochastic optimization, stochastic approximation, quasi-Newton method, damped limited-memory BFGS method, variance reduction

AMS subject classifications. 90C15, 90C30, 62L20, 90C60

DOI. 10.1137/15M1053141

1. Introduction. In this paper, we consider the following stochastic optimization problem:

$$(1.1) \quad \min_{x \in \mathbb{R}^n} f(x) = \mathbb{E}[F(x, \xi)],$$

where $F : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$ is continuously differentiable and possibly nonconvex, $\xi \in \mathbb{R}^d$ denotes a random variable with distribution function P , and $\mathbb{E}[\cdot]$ denotes the expectation taken with respect to ξ . In many cases the function $F(\cdot, \xi)$ is not given explicitly and/or the distribution function P is unknown, or the function values and gradients of f cannot be easily obtained and only noisy information about the gradient of f is available. In this paper we assume that noisy gradients of f can be obtained via calls to a *stochastic first-order oracle* (\mathcal{SFO}). Problem (1.1) arises in many applications in statistics and machine learning [36, 52], mixed logit modeling problems in economics

*Received by the editors December 16, 2015; accepted for publication (in revised form) January 24, 2017; published electronically May 18, 2017.

<http://www.siam.org/journals/siopt/27-2/M105314.html>

Funding: The research of the first author was supported in part by UCAS President Grant Y35101AY00 and NSFC grant 11301505. The research of the second author was supported in part by the Hong Kong Research Grants Council General Research Fund (grant 14205314). The research of the third author was supported in part by NSF grant CCF-1527809.

[†]School of Mathematical Sciences, University of Chinese Academy of Sciences, and Key Laboratory of Big Data Mining and Knowledge Management, Chinese Academy of Sciences, Beijing, China (wangxiao@ucas.ac.cn).

[‡]Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, Shatin, N. T., Hong Kong, China (sqma@se.cuhk.edu.hk).

[§]Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027 (goldfarb@columbia.edu).

[¶]Tencent AI Lab, Shenzhen, China (wliu@ee.columbia.edu).

and transportation [7, 4, 26], as well as many other areas. A special case of (1.1) that arises frequently in machine learning is the empirical risk minimization problem

$$(1.2) \quad \min_{x \in \mathbb{R}^n} f(x) = \frac{1}{T} \sum_{i=1}^T f_i(x),$$

where $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is the loss function that corresponds to the i th data sample, and T denotes the number of data samples which is assumed to be extremely large.

The idea of employing stochastic approximation (SA) to solve stochastic programming problems can be traced back to the seminal work of Robbins and Monro [47]. The classical SA method, also referred to as stochastic gradient descent (SGD), mimics the steepest gradient descent method, i.e., it updates iterate x_k via $x_{k+1} = x_k - \alpha_k g_k$, where the stochastic gradient g_k is an unbiased estimate of the gradient $\nabla f(x_k)$ of f at x_k , and α_k is the step size. The SA method has been studied extensively in [12, 17, 19, 44, 45, 49, 50], where the main focus has been the convergence of SA in different settings. Recently, there has been a lot of interest in analyzing the worst-case complexity of SA methods, stimulated by the complexity theory developed by Nesterov for first-order methods for solving convex optimization problems [42, 43]. Nemirovski et al. [41] proposed a mirror descent SA method for solving the convex stochastic programming problem $x^* := \operatorname{argmin}\{f(x) \mid x \in X\}$, where f is nonsmooth and convex and X is a convex set, and proved that for any given $\epsilon > 0$, the method needs $O(\epsilon^{-2})$ iterations to obtain an \bar{x} such that $\mathbb{E}[f(\bar{x}) - f(x^*)] \leq \epsilon$. Other SA methods with provable complexities for solving convex stochastic optimization problems have also been studied (e.g., see [20, 28, 29, 30, 31, 32, 3, 14, 2, 53, 56]).

Recently there has been a lot of interest in SA methods for stochastic optimization problem (1.1) in which f is a nonconvex function. In [6], an SA method to minimize a general cost function was proposed by Bottou and proved to be convergent to stationary points. Ghadimi and Lan [21] proposed a randomized stochastic gradient (RSG) method that returns an iterate from a randomly chosen iteration as an approximate solution. It is shown in [21] that to return a solution \bar{x} such that $\mathbb{E}[\|\nabla f(\bar{x})\|^2] \leq \epsilon$, where $\|\cdot\|$ denotes the Euclidean norm, the total number of \mathcal{SFO} -calls needed by RSG is $O(\epsilon^{-2})$. Ghadimi and Lan [22] also studied an accelerated SA method for solving (1.1) based on Nesterov's accelerated gradient method [42, 43], which improved the \mathcal{SFO} -call complexity for convex cases from $O(\epsilon^{-2})$ to $O(\epsilon^{-4/3})$. In [23], Ghadimi, Lan, and Zhang proposed a mini-batch SA method for solving problems in which the objective function is a composition of a nonconvex smooth function f and a convex nonsmooth function and analyzed its worst-case \mathcal{SFO} -call complexity. In [13], a method that incorporates a block-coordinate decomposition scheme into stochastic mirror descent methodology was proposed by Dang and Lan for a nonconvex stochastic optimization problem $x^* = \operatorname{argmin}\{f(x) : x \in X\}$ in which the convex set X has a block structure. More recently, Wang, Ma, and Yuan [55] proposed a penalty method for nonconvex stochastic optimization problems with nonconvex constraints and analyzed its \mathcal{SFO} -call complexity.

In this paper, we study stochastic quasi-Newton (SQN) methods for solving the nonconvex stochastic optimization problem (1.1). In the deterministic optimization setting, quasi-Newton methods are more robust and achieve higher accuracy than gradient methods, because they use approximate second-order derivative information. Quasi-Newton methods usually employ the following updates for solving (1.1):

$$(1.3) \quad x_{k+1} = x_k - \alpha_k B_k^{-1} \nabla f(x_k) \quad \text{or} \quad x_{k+1} = x_k - \alpha_k H_k \nabla f(x_k),$$

where B_k is an approximation to the Hessian matrix $\nabla^2 f(x_k)$ at x_k , or H_k is an approximation to $[\nabla^2 f(x_k)]^{-1}$. The most widely used quasi-Newton method, the BFGS method [8, 18, 24, 54], updates B_k via

$$(1.4) \quad B_k = B_{k-1} + \frac{y_{k-1}y_{k-1}^\top}{s_{k-1}^\top y_{k-1}} - \frac{B_{k-1}s_{k-1}s_{k-1}^\top B_{k-1}}{s_{k-1}^\top B_{k-1}s_{k-1}},$$

where $s_{k-1} := x_k - x_{k-1}$ and $y_{k-1} := \nabla f(x_k) - \nabla f(x_{k-1})$. By using the Sherman–Morrison–Woodbury formula, it is easy to derive that the equivalent update to $H_k = B_k^{-1}$ is

$$(1.5) \quad H_k = (I - \rho_{k-1}s_{k-1}y_{k-1}^\top)H_{k-1}(I - \rho_{k-1}y_{k-1}s_{k-1}^\top) + \rho_{k-1}s_{k-1}s_{k-1}^\top,$$

where $\rho_{k-1} := 1/(s_{k-1}^\top y_{k-1})$. For stochastic optimization, there has been some work in designing SQN methods that update the iterates via (1.3) using the stochastic gradient g_k in place of $\nabla f(x_k)$. Specific examples include the following. The adaptive subgradient (AdaGrad) method proposed by Duchi, Hazan, and Singer [15], which takes B_k to be a diagonal matrix that estimates the diagonal of the square root of the uncentered covariance matrix of the gradients, has been proven to be quite efficient in practice. In [5], Bordes, Bottou, and Gallinari studied SGD with a diagonal rescaling matrix based on the secant condition associated with quasi-Newton methods. Roux and Fitzgibbon [48] discussed the necessity of including both Hessian and covariance matrix information in a stochastic Newton type method. Byrd et al. [9] proposed a quasi-Newton method that uses the sample average approximation approach to estimate Hessian-vector multiplications. In [10], Byrd et al. proposed a stochastic limited-memory BFGS (L-BFGS) [34] method based on SA and proved its convergence for strongly convex problems. Stochastic BFGS and L-BFGS methods were also studied for online convex optimization by Schraudolph, Yu, and Günter in [51]. For strongly convex problems, Mokhtari and Ribeiro proposed a regularized stochastic BFGS method and analyzed its convergence in [38] and studied an online L-BFGS method in [39]. Recently, Moritz, Nishihara, and Jordan [40] proposed a linearly convergent method that integrates the L-BFGS method in [10] with the variance reduction technique (stochastic variance reduced gradient (SVRG)) proposed by Johnson and Zhang in [27] to alleviate the effect of noisy gradients. A related method that incorporates SVRG into a quasi-Newton method was studied by Lucchi, McWilliams, and Hofmann in [35]. In [25], Gower, Goldfarb, and Richtárik proposed a variance reduced block L-BFGS method that converges linearly for convex functions. It should be noted that all of the above SQN methods are designed for solving convex or even strongly convex problems.

Challenges. The key challenge in designing SQN methods for nonconvex problem lies in the difficulty in preserving the positive-definiteness of B_k (and H_k), due to the nonconvexity of the problem and the presence of noise in estimating the gradient. It is known that the BFGS update (1.4) preserves the positive-definiteness of B_k as long as the curvature condition

$$(1.6) \quad s_{k-1}^\top y_{k-1} > 0$$

holds, which can be guaranteed for strongly convex problems. For the nonconvex problems, the curvature condition (1.6) can be satisfied by performing a line search. However, doing this is no longer feasible for (1.1) in the stochastic setting, because exact function values and gradient information are not available. As a result, an

important issue in designing SQN methods for nonconvex problems is how to preserve the positive-definiteness of B_k (or H_k) without line search.

Our contributions. Our contributions (and where they appear) in this paper are as follows:

1. We propose a general framework for SQN methods for solving the nonconvex stochastic optimization problem (1.1) and prove its almost sure convergence to a stationary point when the step size α_k is diminishing. We also prove that the number of iterations N needed to obtain $\frac{1}{N} \sum_{k=1}^N \mathbb{E}[\|\nabla f(x_k)\|^2] \leq \epsilon$ is $N = O(\epsilon^{-\frac{1}{1-\beta}})$, for α_k chosen proportional to $k^{-\beta}$, where $\beta \in (0.5, 1)$ is a constant. (See section 2.)
2. When a randomly chosen iterate x_R is returned as the output of SQN, we prove that the worst-case \mathcal{SFO} -calls complexity needed to guarantee $\mathbb{E}[\|\nabla f(x_R)\|^2] \leq \epsilon$ is $O(\epsilon^{-2})$. (See section 2.2.)
3. We propose a stochastic damped L-BFGS (SdLBFGS) method that fits into the proposed framework. This method adaptively generates a positive definite matrix H_k that approximates the inverse Hessian matrix at the current iterate x_k . Convergence and complexity results for this method are provided. Moreover, our method does not generate H_k explicitly, and only its multiplication with vectors is computed directly. (See section 3.)
4. Motivated by the recent advance of SVRG for nonconvex minimization [46, 1], we propose a variance reduced variant of SdLBFGS and analyze its \mathcal{SFO} -calls complexity. (See section 4.)

2. A general framework for SQN methods for nonconvex optimization.

In this section, we study SQN methods for the (possibly nonconvex) stochastic optimization problem (1.1). We assume that an \mathcal{SFO} outputs a stochastic gradient $g(x, \xi)$ of f for a given x , where ξ is a random variable whose distribution is supported on $\Xi \subseteq \mathbb{R}^d$. Here we assume that Ξ does not depend on x .

We now give some assumptions that are required throughout this paper.

- AS.1. $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable. $f(x)$ is lower bounded by a real number f^{low} for any $x \in \mathbb{R}^n$. ∇f is globally Lipschitz continuous with Lipschitz constant L ; namely for any $x, y \in \mathbb{R}^n$,

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|.$$

- AS.2. For any iteration k , we have

$$(2.1) \quad \text{a) } \mathbb{E}_{\xi_k} [g(x_k, \xi_k)] = \nabla f(x_k),$$

$$(2.2) \quad \text{b) } \mathbb{E}_{\xi_k} [\|g(x_k, \xi_k) - \nabla f(x_k)\|^2] \leq \sigma^2,$$

where $\sigma > 0$ is the noise level of the gradient estimation, and ξ_k , $k = 1, 2, \dots$, are independent samples, and for a given k the random variable ξ_k is independent of $\{x_j\}_{j=1}^k$.

Remark 2.1. Note that the stochastic BFGS methods studied in [38, 10, 39] require that the noisy gradient is bounded, i.e.,

$$(2.3) \quad \mathbb{E}_{\xi_k} [\|g(x_k, \xi_k)\|^2] \leq M_g,$$

where $M_g > 0$ is a constant. Our assumption (2.2) is weaker than (2.3).

Analogous to deterministic quasi-Newton methods, our SQN method takes steps

$$(2.4) \quad x_{k+1} = x_k - \alpha_k H_k g_k,$$

where g_k is defined as a mini-batch estimate of the gradient,

$$(2.5) \quad g_k = \frac{1}{m_k} \sum_{i=1}^{m_k} g(x_k, \xi_{k,i}),$$

and $\xi_{k,i}$ denotes the random variable generated by the i th sampling in the k th iteration. From AS.2 we can see that g_k has the following properties:

$$(2.6) \quad \mathbb{E}[g_k | x_k] = \nabla f(x_k), \quad \mathbb{E}[\|g_k - \nabla f(x_k)\|^2 | x_k] \leq \frac{\sigma^2}{m_k}.$$

AS.3. There exist two positive constants $\underline{\kappa}, \bar{\kappa}$ such that

$$\underline{\kappa}I \preceq H_k \preceq \bar{\kappa}I \quad \forall k,$$

where the notation $A \succeq B$ with $A, B \in \mathbb{R}^{n \times n}$ means that $A - B$ is positive semidefinite.

We denote by $\xi_k = (\xi_{k,1}, \dots, \xi_{k,m_k})$ the random samplings in the k th iteration, and we denote by $\xi_{[k]} := (\xi_1, \dots, \xi_k)$ the random samplings in the first k iterations. Since H_k is generated iteratively based on historical gradient information by a random process, we make the following assumption on $H_k (k \geq 2)$ to control the randomness (note that H_1 is given in the initialization step).

AS.4. For any $k \geq 2$, the random variable H_k depends only on $\xi_{[k-1]}$. It then follows directly from AS.4 and (2.6) that

$$(2.7) \quad \mathbb{E}[H_k g_k | \xi_{[k-1]}] = H_k \nabla f(x_k),$$

where the expectation is taken with respect to ξ_k generated in the computation of g_k .

We will not specify how to compute H_k until section 3, where a specific updating scheme for H_k satisfying both assumptions AS.3 and AS.4 will be proposed.

We now present our SQN method for solving (1.1) as Algorithm 2.1.

Algorithm 2.1. SQN method for nonconvex stochastic optimization.

Input: Given $x_1 \in \mathbb{R}^n$, a positive definite matrix $H_1 \in \mathbb{R}^{n \times n}$, batch sizes $\{m_k\}_{k \geq 1}$, and step sizes $\{\alpha_k\}_{k \geq 1}$

- 1: **for** $k = 1, 2, \dots$ **do**
 - 2: Calculate $g_k = \frac{1}{m_k} \sum_{i=1}^{m_k} g(x_k, \xi_{k,i})$.
 - 3: Generate a positive definite Hessian inverse approximation H_k .
 - 4: Calculate $x_{k+1} = x_k - \alpha_k H_k g_k$.
 - 5: **end for**
-

2.1. Convergence and complexity of SQN with diminishing step size.

In this subsection, we analyze the convergence and complexity of SQN under the condition that the step size α_k in (2.4) is diminishing. Specifically, in this subsection we assume α_k satisfies the following condition:

$$(2.8) \quad \sum_{k=1}^{+\infty} \alpha_k = +\infty, \quad \sum_{k=1}^{+\infty} \alpha_k^2 < +\infty,$$

which is a standard assumption in SA algorithms (see, e.g., [10, 39, 41]). One very simple choice of α_k that satisfies (2.8) is $\alpha_k = O(1/k)$.

The following lemma shows that a descent property in terms of the expected objective value holds for SQN. Our analysis is similar to analyses that have been used in [6, 38].

LEMMA 2.2. *Suppose that $\{x_k\}$ is generated by SQN and assumptions AS.1–AS.4 hold. Further assume that (2.8) holds, and $\alpha_k \leq \frac{\kappa}{L\bar{\kappa}^2}$ for all k . (Note that this can be satisfied if α_k is nonincreasing and the initial step size $\alpha_1 \leq \frac{\kappa}{L\bar{\kappa}^2}$.) Then the following inequality holds:*

$$(2.9) \quad \mathbb{E}[f(x_{k+1})|x_k] \leq f(x_k) - \frac{1}{2}\alpha_k \underline{\kappa} \|\nabla f(x_k)\|^2 + \frac{L\sigma^2 \bar{\kappa}^2}{2m_k} \alpha_k^2 \quad \forall k \geq 1,$$

where the conditional expectation is taken with respect to ξ_k .

Proof. Define $\delta_k = g_k - \nabla f(x_k)$. From (2.4), and assumptions AS.1 and AS.3, we have

$$\begin{aligned} f(x_{k+1}) &\leq f(x_k) + \langle \nabla f(x_k), x_{k+1} - x_k \rangle + \frac{L}{2} \|x_{k+1} - x_k\|^2 \\ &= f(x_k) - \alpha_k \langle \nabla f(x_k), H_k g_k \rangle + \frac{L}{2} \alpha_k^2 \|H_k g_k\|^2 \\ (2.10) \quad &\leq f(x_k) - \alpha_k \langle \nabla f(x_k), H_k \nabla f(x_k) \rangle - \alpha_k \langle \nabla f(x_k), H_k \delta_k \rangle + \frac{L}{2} \alpha_k^2 \bar{\kappa}^2 \|g_k\|^2. \end{aligned}$$

Taking expectation with respect to ξ_k on both sides of (2.10) conditioned on x_k , we obtain

$$(2.11) \quad \mathbb{E}[f(x_{k+1})|x_k] \leq f(x_k) - \alpha_k \langle \nabla f(x_k), H_k \nabla f(x_k) \rangle + \frac{L}{2} \alpha_k^2 \bar{\kappa}^2 \mathbb{E}[\|g_k\|^2|x_k],$$

where we used (2.7) and the fact that $\mathbb{E}[\delta_k|x_k] = 0$. From (2.6) and $\mathbb{E}[\delta_k|x_k] = 0$, it follows that

$$\begin{aligned} \mathbb{E}[\|g_k\|^2|x_k] &= \mathbb{E}[\|g_k - \nabla f(x_k) + \nabla f(x_k)\|^2|x_k] \\ &= \mathbb{E}[\|\nabla f(x_k)\|^2|x_k] + \mathbb{E}[\|g_k - \nabla f(x_k)\|^2|x_k] + 2\mathbb{E}[\langle \delta_k, \nabla f(x_k) \rangle|x_k] \\ (2.12) \quad &= \|\nabla f(x_k)\|^2 + \mathbb{E}[\|g_k - \nabla f(x_k)\|^2|x_k] \leq \|\nabla f(x_k)\|^2 + \sigma^2/m_k, \end{aligned}$$

which together with (2.11) and AS.3 yields that

$$(2.13) \quad \mathbb{E}[f(x_{k+1})|x_k] \leq f(x_k) - \left(\alpha_k \underline{\kappa} - \frac{L}{2} \alpha_k^2 \bar{\kappa}^2 \right) \|\nabla f(x_k)\|^2 + \frac{L\sigma^2 \bar{\kappa}^2}{2m_k} \alpha_k^2.$$

Then (2.13) combined with the assumption $\alpha_k \leq \frac{\kappa}{L\bar{\kappa}^2}$ implies (2.9). \square

Before proceeding further, we introduce the definition of a *supermartingale* (see [16] for more details).

DEFINITION 2.3. *Let $\{\mathcal{F}_k\}$ be an increasing sequence of σ -algebras. If $\{X_k\}$ is a stochastic process satisfying (i) $\mathbb{E}[|X_k|] < \infty$, (ii) $X_k \in \mathcal{F}_k$ for all k , and (iii) $\mathbb{E}[X_{k+1}|\mathcal{F}_k] \leq X_k$ for all k , then $\{X_k\}$ is called a supermartingale.*

PROPOSITION 2.4 (see, e.g., Theorem 5.2.9 in [16]). *If $\{X_k\}$ is a nonnegative supermartingale, then $\lim_{k \rightarrow \infty} X_k \rightarrow X$ almost surely and $\mathbb{E}[X] \leq \mathbb{E}[X_0]$.*

We are now ready to give convergence results for SQN (Algorithm 2.1).

THEOREM 2.5. Suppose that assumptions AS.1–AS.4 hold for $\{x_k\}$ generated by SQN with batch size $m_k = m$ for all k . If the step size α_k satisfies (2.8) and $\alpha_k \leq \frac{\kappa}{L\bar{\kappa}^2}$ for all k , then

$$(2.14) \quad \liminf_{k \rightarrow \infty} \|\nabla f(x_k)\| = 0 \quad \text{with probability 1.}$$

Moreover, there exists a positive constant M_f such that

$$(2.15) \quad \mathbb{E}[f(x_k)] \leq M_f \quad \forall k.$$

Proof. Define $\beta_k := \frac{\alpha_k \kappa}{2} \|\nabla f(x_k)\|^2$ and $\gamma_k := f(x_k) + \frac{L\sigma^2 \bar{\kappa}^2}{2m} \sum_{i=k}^{\infty} \alpha_i^2$. Let \mathcal{F}_k be the σ -algebra measuring β_k , γ_k , and x_k . From (2.9) we know that for any k ,

$$(2.16) \quad \begin{aligned} \mathbb{E}[\gamma_{k+1} | \mathcal{F}_k] &= \mathbb{E}[f(x_{k+1}) | \mathcal{F}_k] + \frac{L\sigma^2 \bar{\kappa}^2}{2m} \sum_{i=k+1}^{\infty} \alpha_i^2 \\ &\leq f(x_k) - \frac{\alpha_k \kappa}{2} \|\nabla f(x_k)\|^2 + \frac{L\sigma^2 \bar{\kappa}^2}{2m} \sum_{i=k}^{\infty} \alpha_i^2 = \gamma_k - \beta_k, \end{aligned}$$

which implies that $\mathbb{E}[\gamma_{k+1} - f^{low} | \mathcal{F}_k] \leq \gamma_k - f^{low} - \beta_k$. Since $\beta_k \geq 0$, we have $0 \leq \mathbb{E}[\gamma_k - f^{low}] \leq \gamma_1 - f^{low} < +\infty$, which implies (2.15). According to Definition 2.3, $\{\gamma_k - f^{low}\}$ is a supermartingale. Therefore, Proposition 2.4 shows that there exists a γ such that $\lim_{k \rightarrow \infty} \gamma_k = \gamma$ with probability 1, and $\mathbb{E}[\gamma] \leq \mathbb{E}[\gamma_1]$. Note that from (2.16) we have $\mathbb{E}[\beta_k] \leq \mathbb{E}[\gamma_k] - \mathbb{E}[\gamma_{k+1}]$. Thus,

$$\mathbb{E} \left[\sum_{k=1}^{\infty} \beta_k \right] \leq \sum_{k=1}^{\infty} (\mathbb{E}[\gamma_k] - \mathbb{E}[\gamma_{k+1}]) < +\infty,$$

which further yields that

$$(2.17) \quad \sum_{k=1}^{\infty} \beta_k = \frac{\kappa}{2} \sum_{k=1}^{\infty} \alpha_k \|\nabla f(x_k)\|^2 < +\infty \quad \text{with probability 1.}$$

Since $\sum_{k=1}^{\infty} \alpha_k = +\infty$, it follows that (2.14) holds. \square

Under the assumption (2.3) used in [38, 10, 39], we now prove a stronger convergence result showing that any limit point of $\{x_k\}$ generated by SQN is a stationary point of (1.1) with probability 1.

THEOREM 2.6. Assume the same assumptions hold as in Theorem 2.5 and that (2.3) holds. Then

$$(2.18) \quad \lim_{k \rightarrow \infty} \|\nabla f(x_k)\| = 0 \quad \text{with probability 1.}$$

Proof. For any given $\epsilon > 0$, according to (2.14), there exist infinitely many iterates x_k such that $\|\nabla f(x_k)\| < \epsilon$. Then if (2.18) does not hold, there must exist two infinite sequences of indices $\{m_i\}$, $\{n_i\}$ with $n_i > m_i$, such that for $i = 1, 2, \dots$,

$$(2.19) \quad \|\nabla f(x_{m_i})\| \geq 2\epsilon, \quad \|\nabla f(x_{n_i})\| < \epsilon, \quad \|\nabla f(x_k)\| \geq \epsilon, \quad k = m_i + 1, \dots, n_i - 1.$$

Then from (2.17) it follows that

$$+\infty > \sum_{k=1}^{+\infty} \alpha_k \|\nabla f(x_k)\|^2 \geq \sum_{i=1}^{+\infty} \sum_{k=m_i}^{n_i-1} \alpha_k \|\nabla f(x_k)\|^2 \geq \epsilon^2 \sum_{i=1}^{+\infty} \sum_{k=m_i}^{n_i-1} \alpha_k \quad \text{with probability 1,}$$

which implies that

$$(2.20) \quad \sum_{k=m_i}^{n_i-1} \alpha_k \rightarrow 0 \text{ with probability 1, as } i \rightarrow +\infty.$$

According to (2.12), we have that

$$(2.21) \quad \begin{aligned} \mathbb{E}[\|x_{k+1} - x_k\| | x_k] &= \alpha_k \mathbb{E}[\|H_k g_k\| | x_k] \leq \alpha_k \bar{\kappa} \mathbb{E}[\|g_k\| | x_k] \\ &\leq \alpha_k \bar{\kappa} (\mathbb{E}[\|g_k\|^2 | x_k])^{\frac{1}{2}} \leq \alpha_k \bar{\kappa} (M_g/m)^{\frac{1}{2}}, \end{aligned}$$

where the last inequality is due to (2.3) and the convexity of $\|\cdot\|^2$. Then it follows from (2.21) that

$$\mathbb{E}[\|x_{n_i} - x_{m_i}\|] \leq \bar{\kappa} (M_g/m)^{\frac{1}{2}} \sum_{k=m_i}^{n_i-1} \alpha_k,$$

which together with (2.20) implies that $\|x_{n_i} - x_{m_i}\| \rightarrow 0$ with probability 1, as $i \rightarrow +\infty$. Hence, from the Lipschitz continuity of ∇f , it follows that $\|\nabla f(x_{n_i}) - \nabla f(x_{m_i})\| \rightarrow 0$ with probability 1 as $i \rightarrow +\infty$. However, this contradicts (2.19). Therefore, the assumption that (2.18) does not hold is not true. \square

Remark 2.7. Note that our result in Theorem 2.6 is stronger than the ones given in existing works such as [38] and [39]. Moreover, although Bottou [6] also proves that the SA method for nonconvex stochastic optimization with diminishing step size is almost surely convergent to a stationary point, our analysis requires weaker assumptions. For example, [6] assumes that the objective function is three times continuously differentiable, while our analysis does not require this. Furthermore, we are able to analyze the iteration complexity of SQN, for a specifically chosen step size α_k (see Theorem 2.8 below), which is not provided in [6].

We now analyze the iteration complexity of SQN.

THEOREM 2.8. *Suppose that assumptions AS.1–AS.4 hold for $\{x_k\}$ generated by SQN with batch size $m_k = m$ for all k . We also assume that α_k is specifically chosen as*

$$(2.22) \quad \alpha_k = \frac{\underline{\kappa}}{L\bar{\kappa}^2} k^{-\beta}$$

with $\beta \in (0.5, 1)$. Note that this choice satisfies (2.8) and $\alpha_k \leq \frac{\underline{\kappa}}{L\bar{\kappa}^2}$ for all k . Then

$$(2.23) \quad \frac{1}{N} \sum_{k=1}^N \mathbb{E}[\|\nabla f(x_k)\|^2] \leq \frac{2L(M_f - f^{\text{low}})\bar{\kappa}^2}{\underline{\kappa}^2} N^{\beta-1} + \frac{\sigma^2}{(1-\beta)m} (N^{-\beta} - N^{-1}),$$

where N denotes the iteration number. Moreover, for a given $\epsilon \in (0, 1)$, to guarantee that $\frac{1}{N} \sum_{k=1}^N \mathbb{E}[\|\nabla f(x_k)\|^2] < \epsilon$, the number of iterations N needed is at most $O(\epsilon^{-\frac{1}{1-\beta}})$.

Proof. Taking expectation on both sides of (2.9) and summing over $k = 1, \dots, N$ yields

$$\begin{aligned} \frac{1}{2} \frac{\underline{\kappa}}{L} \sum_{k=1}^N \mathbb{E}[\|\nabla f(x_k)\|^2] &\leq \sum_{k=1}^N \frac{1}{\alpha_k} (\mathbb{E}[f(x_k)] - \mathbb{E}[f(x_{k+1})]) + \frac{L\sigma^2\bar{\kappa}^2}{2m} \sum_{k=1}^N \alpha_k \\ &= \frac{1}{\alpha_1} f(x_1) + \sum_{k=2}^N \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}} \right) \mathbb{E}[f(x_k)] \end{aligned}$$

$$\begin{aligned}
& -\frac{\mathbb{E}[f(x_{N+1})]}{\alpha_N} + \frac{L\sigma^2\bar{\kappa}^2}{2m} \sum_{k=1}^N \alpha_k \\
& \leq \frac{M_f}{\alpha_1} + M_f \sum_{k=2}^N \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}} \right) - \frac{f^{low}}{\alpha_N} + \frac{L\sigma^2\bar{\kappa}^2}{2m} \sum_{k=1}^N \alpha_k \\
& = \frac{M_f - f^{low}}{\alpha_N} + \frac{L\sigma^2\bar{\kappa}^2}{2m} \sum_{k=1}^N \alpha_k \\
& \leq \frac{L(M_f - f^{low})\bar{\kappa}^2}{\underline{\kappa}} N^\beta + \frac{\sigma^2 \underline{\kappa}}{2(1-\beta)m} (N^{1-\beta} - 1),
\end{aligned}$$

which results in (2.23), where the second inequality is due to (2.15) and the last inequality is due to (2.22). Then for a given $\epsilon > 0$, to guarantee that $\frac{1}{N} \sum_{k=1}^N \mathbb{E}[\|\nabla f(x_k)\|^2] \leq \epsilon$, it suffices to require that

$$\frac{2L(M_f - f^{low})\bar{\kappa}^2}{\underline{\kappa}^2} N^{\beta-1} + \frac{\sigma^2}{(1-\beta)m} (N^{-\beta} - N^{-1}) < \epsilon.$$

Since $\beta \in (0.5, 1)$, it follows that the number of iterations N needed is at most $O(\epsilon^{-\frac{1}{1-\beta}})$. \square

Remark 2.9. Note that Theorem 2.8 also provides iteration complexity analysis for the classic SGD method, which can be regarded as a special case of SQN with $H_k = I$. To the best of our knowledge, our complexity result in Theorem 2.8 is new for both SGD and SQN methods.

2.2. Complexity of SQN with random output and constant step size.

We analyze the \mathcal{SFO} -calls complexity of SQN when the output is randomly chosen from $\{x_i\}_{i=1}^N$, where N is the maximum iteration number. Our results in this subsection are motivated by the RSG method proposed by Ghadimi and Lan [21]. RSG runs SGD for R iterations, where R is a randomly chosen integer from $\{1, \dots, N\}$ with a specifically defined probability mass function P_R . In [21] it is proved that under certain conditions on the step size and P_R , $O(1/\epsilon^2)$ \mathcal{SFO} -calls are needed by SGD to guarantee $\mathbb{E}[\|\nabla f(x_R)\|^2] \leq \epsilon$. We show below that under similar conditions, the same complexity holds for our SQN.

THEOREM 2.10. *Suppose that assumptions AS.1–AS.4 hold and that α_k in SQN (Algorithm 2.1) is chosen such that $0 < \alpha_k \leq 2\underline{\kappa}/(L\bar{\kappa}^2)$ for all k with $\alpha_k < 2\underline{\kappa}/(L\bar{\kappa}^2)$ for at least one k . Moreover, for a given integer N , let R be a random variable with the probability mass function*

$$(2.24) \quad P_R(k) := \text{Prob}\{R = k\} = \frac{\alpha_k \underline{\kappa} - \alpha_k^2 L \bar{\kappa}^2 / 2}{\sum_{k=1}^N (\alpha_k \underline{\kappa} - \alpha_k^2 L \bar{\kappa}^2 / 2)}, \quad k = 1, \dots, N.$$

Then we have

$$(2.25) \quad \mathbb{E}[\|\nabla f(x_R)\|^2] \leq \frac{D_f + (\sigma^2 L \bar{\kappa}^2) / 2 \sum_{k=1}^N (\alpha_k^2 / m_k)}{\sum_{k=1}^N (\alpha_k \underline{\kappa} - \alpha_k^2 L \bar{\kappa}^2 / 2)},$$

where $D_f := f(x_1) - f^{low}$ and the expectation is taken with respect to R and $\xi_{[N]}$. Moreover, if we choose $\alpha_k = \underline{\kappa}/(L\bar{\kappa}^2)$ and $m_k = m$ for all $k = 1, \dots, N$, then (2.25) reduces to

$$(2.26) \quad \mathbb{E}[\|\nabla f(x_R)\|^2] \leq \frac{2L\bar{\kappa}^2 D_f}{N\underline{\kappa}^2} + \frac{\sigma^2}{m}.$$

Proof. From (2.10) it follows that

$$\begin{aligned} f(x_{k+1}) &\leq f(x_k) - \alpha_k \langle \nabla f(x_k), H_k \nabla f(x_k) \rangle - \alpha_k \langle \nabla f(x_k), H_k \delta_k \rangle \\ &\quad + \frac{L}{2} \alpha_k^2 \bar{\kappa}^2 [\|\nabla f(x_k)\|^2 + 2\langle \nabla f(x_k), \delta_k \rangle + \|\delta_k\|^2] \\ &\leq f(x_k) - \left(\alpha_k \underline{\kappa} - \frac{1}{2} \alpha_k^2 L \bar{\kappa}^2 \right) \|\nabla f(x_k)\|^2 + \frac{1}{2} \alpha_k^2 L \bar{\kappa}^2 \|\delta_k\|^2 + \alpha_k^2 L \bar{\kappa}^2 \langle \nabla f(x_k), \delta_k \rangle \\ &\quad - \alpha_k \langle \nabla f(x_k), H_k \delta_k \rangle, \end{aligned}$$

where $\delta_k = g_k - \nabla f(x_k)$. Now summing $k = 1, \dots, N$ and noticing that $\alpha_k \leq 2\underline{\kappa}/(L\bar{\kappa}^2)$ yields

$$\begin{aligned} &\sum_{k=1}^N \left(\alpha_k \underline{\kappa} - \frac{L\bar{\kappa}^2}{2} \alpha_k^2 \right) \|\nabla f(x_k)\|^2 \\ &\leq f(x_1) - f^{\text{low}} + \frac{L\bar{\kappa}^2}{2} \sum_{k=1}^N \alpha_k^2 \|\delta_k\|^2 + \sum_{k=1}^N (\alpha_k^2 L \bar{\kappa}^2 \langle \nabla f(x_k), \delta_k \rangle - \alpha_k \langle \nabla f(x_k), H_k \delta_k \rangle). \end{aligned} \quad (2.27)$$

By AS.2 and AS.4 we have that

$$\mathbb{E}_{\xi_k} [\langle \nabla f(x_k), \delta_k \rangle | \xi_{[k-1]}] = 0 \quad \text{and} \quad \mathbb{E}_{\xi_k} [\langle \nabla f(x_k), H_k \delta_k \rangle | \xi_{[k-1]}] = 0.$$

Moreover, from (2.6) it follows that $\mathbb{E}_{\xi_k} [\|\delta_k\|^2 | \xi_{[k-1]}] \leq \sigma^2/m_k$. Therefore, taking the expectation on both sides of (2.27) with respect to $\xi_{[N]}$ yields

$$(2.28) \quad \sum_{k=1}^N (\alpha_k \underline{\kappa} - \alpha_k^2 L \bar{\kappa}^2 / 2) \mathbb{E}_{\xi_{[N]}} [\|\nabla f(x_k)\|^2] \leq f(x_1) - f^{\text{low}} + \frac{L\bar{\kappa}^2 \sigma^2}{2} \sum_{k=1}^N \frac{\alpha_k^2}{m_k}.$$

It follows from the definition of P_R in (2.24) that

$$(2.29) \quad \mathbb{E} [\|\nabla f(x_R)\|^2] = \mathbb{E}_{R, \xi_{[N]}} [\|\nabla f(x_R)\|^2] = \frac{\sum_{k=1}^N (\alpha_k \underline{\kappa} - \alpha_k^2 L \bar{\kappa}^2 / 2) \mathbb{E}_{\xi_{[N]}} [\|\nabla f(x_k)\|^2]}{\sum_{k=1}^N (\alpha_k \underline{\kappa} - \alpha_k^2 L \bar{\kappa}^2 / 2)},$$

which together with (2.28) implies (2.25). \square

Remark 2.11. Note that in Theorem 2.10, α_k 's are not required to be diminishing, and they can be constant as long as they are upper bounded by $2\underline{\kappa}/(L\bar{\kappa}^2)$.

We now show that the \mathcal{SFO} complexity of SQN with random output and constant step size is $O(\epsilon^{-2})$.

COROLLARY 2.12. *Assume the conditions in Theorem 2.10 hold, and $\alpha_k = \underline{\kappa}/(L\bar{\kappa}^2)$ and $m_k = m$ for all $k = 1, \dots, N$. Let \bar{N} be the total number of \mathcal{SFO} -calls needed to calculate stochastic gradients g_k in SQN (Algorithm 2.1). For a given accuracy tolerance $\epsilon > 0$, we assume that*

$$(2.30) \quad \bar{N} \geq \max \left\{ \frac{C_1^2}{\epsilon^2} + \frac{4C_2}{\epsilon}, \frac{\sigma^2}{L^2 \bar{D}} \right\}, \quad \text{where} \quad C_1 = \frac{4\sigma \bar{\kappa}^2 D_f}{\underline{\kappa}^2 \sqrt{\bar{D}}} + \sigma L \sqrt{\bar{D}}, \quad C_2 = \frac{4L\bar{\kappa}^2 D_f}{\underline{\kappa}^2},$$

where \tilde{D} is a problem-independent positive constant. Moreover, we assume that the batch size satisfies

$$(2.31) \quad m_k = m := \left\lceil \min \left\{ \bar{N}, \max \left\{ 1, \frac{\sigma}{L} \sqrt{\frac{\bar{N}}{\tilde{D}}} \right\} \right\} \right\rceil.$$

Then $\mathbb{E}[\|\nabla f(x_R)\|^2] \leq \epsilon$, where the expectation is taken with respect to R and $\xi_{[N]}$.

Proof. Note that the number of iterations of SQN is at most $N = \lceil \bar{N}/m \rceil$. Obviously, $N \geq \bar{N}/(2m)$. From (2.26) we have that

$$(2.32) \quad \begin{aligned} \mathbb{E}[\|\nabla f(x_R)\|^2] &\leq \frac{2L\bar{\kappa}^2 D_f}{N\bar{\kappa}^2} + \frac{\sigma^2}{m} \leq \frac{4L\bar{\kappa}^2 D_f}{\bar{N}\bar{\kappa}^2} m + \frac{\sigma^2}{m} \\ &\leq \frac{4L\bar{\kappa}^2 D_f}{\bar{N}\bar{\kappa}^2} \left(1 + \frac{\sigma}{L} \sqrt{\frac{\bar{N}}{\tilde{D}}} \right) + \max \left\{ \frac{\sigma^2}{\bar{N}}, \frac{\sigma L \sqrt{\tilde{D}}}{\sqrt{\bar{N}}} \right\}. \end{aligned}$$

Equation (2.30) indicates that

$$\sqrt{\bar{N}} \geq \frac{\sqrt{C_1^2 + 4\epsilon C_2}}{\epsilon} \geq \frac{\sqrt{C_1^2 + 4\epsilon C_2} + C_1}{2\epsilon}.$$

Equation (2.30) also implies that $\sigma^2/\bar{N} \leq \sigma L \sqrt{\tilde{D}}/\sqrt{\bar{N}}$. Then (2.32) yields that

$$\mathbb{E}[\|\nabla f(x_R)\|^2] \leq \frac{4L\bar{\kappa}^2 D_f}{\bar{N}\bar{\kappa}^2} \left(1 + \frac{\sigma}{L} \sqrt{\frac{\bar{N}}{\tilde{D}}} \right) + \frac{\sigma L \sqrt{\tilde{D}}}{\sqrt{\bar{N}}} = \frac{C_1}{\sqrt{\bar{N}}} + \frac{C_2}{\bar{N}} \leq \epsilon,$$

which completes the proof. \square

Remark 2.13. In Corollary 2.12 we did not consider the \mathcal{SFO} -calls that are involved in updating H_k in line 3 of SQN. In the next section, we consider a specific updating scheme to generate H_k and analyze the total \mathcal{SFO} -calls complexity of SQN including the generation of the H_k .

3. Stochastic damped L-BFGS method. In this section, we propose a specific way, namely a damped L-BFGS method (SdLBFGS), to generate H_k in SQN (Algorithm 2.1) that satisfies assumptions AS.3 and AS.4. We also provide an efficient way to compute $H_k g_k$ without generating H_k explicitly.

Before doing this, we first describe a stochastic damped BFGS method as follows. We generate an auxiliary stochastic gradient at x_k using the samplings from the $(k-1)$ -st iteration:

$$\bar{g}_k := \frac{1}{m_{k-1}} \sum_{i=1}^{m_{k-1}} g(x_k, \xi_{k-1,i}).$$

Note that we assume that our \mathcal{SFO} can separate two arguments x_k and ξ_k in the stochastic gradient $g(x_k, \xi_{k-1})$ and generate an output $g(x_k; \xi_{k-1,i})$. The stochastic gradient difference is defined as

$$(3.1) \quad y_{k-1} := \bar{g}_k - g_{k-1} = \frac{\sum_{i=1}^{m_{k-1}} g(x_k, \xi_{k-1,i}) - g(x_{k-1}, \xi_{k-1,i})}{m_{k-1}}.$$

The iterate difference is still defined as $s_{k-1} = x_k - x_{k-1}$. We then define

$$(3.2) \quad \bar{y}_{k-1} = \hat{\theta}_{k-1} y_{k-1} + (1 - \hat{\theta}_{k-1}) B_{k-1} s_{k-1},$$

where

$$(3.3) \quad \hat{\theta}_{k-1} = \begin{cases} \frac{0.75 s_{k-1}^\top B_{k-1} s_{k-1}}{s_{k-1}^\top B_{k-1} s_{k-1} - s_{k-1}^\top y_{k-1}} & \text{if } s_{k-1}^\top y_{k-1} < 0.25 s_{k-1}^\top B_{k-1} s_{k-1}, \\ 1 & \text{otherwise.} \end{cases}$$

Note that if $B_{k-1} \succ 0$, then $0 < \hat{\theta}_{k-1} \leq 1$. Our stochastic damped BFGS approach updates B_{k-1} as

$$(3.4) \quad B_k = B_{k-1} + \frac{\bar{y}_{k-1} \bar{y}_{k-1}^\top}{s_{k-1}^\top \bar{y}_{k-1}} - \frac{B_{k-1} s_{k-1} s_{k-1}^\top B_{k-1}}{s_{k-1}^\top B_{k-1} s_{k-1}}.$$

According to the Sherman–Morrison–Woodbury formula, this corresponds to updating $H_k = B_k^{-1}$ as

$$(3.5) \quad H_k = (I - \rho_{k-1} s_{k-1} \bar{y}_{k-1}^\top) H_{k-1} (I - \rho_{k-1} \bar{y}_{k-1} s_{k-1}^\top) + \rho_{k-1} s_{k-1} s_{k-1}^\top,$$

where $\rho_{k-1} = (s_{k-1}^\top \bar{y}_{k-1})^{-1}$. The following lemma shows that the damped BFGS updates (3.4) and (3.5) preserve the positive definiteness of B_k and H_k .

LEMMA 3.1. *For \bar{y}_{k-1} defined in (3.2), $s_{k-1}^\top \bar{y}_{k-1} \geq 0.25 s_{k-1}^\top B_{k-1} s_{k-1}$. Moreover, if $B_{k-1} = H_{k-1}^{-1} \succ 0$, then B_k and H_k generated by the damped BFGS updates (3.4) and (3.5) are both positive definite.*

Proof. From (3.2) and (3.3) we have that

$$\begin{aligned} s_{k-1}^\top \bar{y}_{k-1} &= \hat{\theta}_{k-1} (s_{k-1}^\top y_{k-1} - s_{k-1}^\top B_{k-1} s_{k-1}) + s_{k-1}^\top B_{k-1} s_{k-1} \\ &= \begin{cases} 0.25 s_{k-1}^\top B_{k-1} s_{k-1} & \text{if } s_{k-1}^\top y_{k-1} < 0.25 s_{k-1}^\top B_{k-1} s_{k-1}, \\ s_{k-1}^\top y_{k-1} & \text{otherwise,} \end{cases} \end{aligned}$$

which implies $s_{k-1}^\top \bar{y}_{k-1} \geq 0.25 s_{k-1}^\top B_{k-1} s_{k-1}$. Therefore, if $B_{k-1} \succ 0$, it follows that $\rho_{k-1} > 0$. As a result, for H_k defined in (3.5) and any nonzero vector $z \in \mathbb{R}^n$, we have

$$z^\top H_k z = z^\top (I - \rho_{k-1} s_{k-1} \bar{y}_{k-1}^\top) H_{k-1} (I - \rho_{k-1} \bar{y}_{k-1} s_{k-1}^\top) z + \rho_{k-1} (s_{k-1}^\top z)^2 > 0,$$

given that $H_{k-1} \succ 0$. Therefore, both H_k and B_k defined in (3.5) and (3.4) are positive definite. \square

Computing H_k by the stochastic damped BFGS update (3.5) and computing the step direction $H_k g_k$ requires $O(n^2)$ multiplications. This is costly if n is large. The L-BFGS method originally proposed by Liu and Nocedal [34] can be adopted here to reduce this computational cost. The L-BFGS method can be described as follows for deterministic optimization problems. Given an initial estimate $H_{k,0} \in \mathbb{R}^{n \times n}$ of the inverse Hessian at the current iterate x_k and two sequences $\{s_j\}$, $\{y_j\}$, $j = k-p, \dots, k-1$, where p is the memory size, the L-BFGS method updates $H_{k,i}$ recursively as

$$(3.6) \quad H_{k,i} = (I - \rho_j s_j y_j^\top) H_{k,i-1} (I - \rho_j y_j s_j^\top) + \rho_j s_j s_j^\top, \quad j = k - (p - i + 1); i = 1, \dots, p,$$

where $\rho_j = (s_j^\top y_j)^{-1}$. The output $H_{k,p}$ is then used as the estimate of the inverse Hessian at x_k to compute the search direction at the k th iteration. It can be shown that if the sequence of pairs $\{s_j, y_j\}$ satisfies the curvature condition $s_j^\top y_j > 0$, $j = k-1, \dots, k-p$, then $H_{k,p}$ is positive definite provided that $H_{k,0}$ is positive definite. Recently, stochastic L-BFGS methods have been proposed for solving strongly convex problems in [10, 39, 40, 25]. However, the theoretical convergence analyses in these papers do not apply to nonconvex problems. We now show how to design an SdLBFGS formula for nonconvex problems.

Suppose that in the past iterations the algorithm generated s_j and \bar{y}_j that satisfy

$$s_j^\top \bar{y}_j \geq 0.25 s_j^\top H_{j+1,0}^{-1} s_j, \quad j = k-p, \dots, k-2.$$

Then at the current iterate, we compute $s_{k-1} = x_k - x_{k-1}$ and y_{k-1} by (3.1). Since $s_{k-1}^\top y_{k-1}$ may not be positive, motivated by the stochastic damped BFGS update (3.2)–(3.5), we define a new vector $\{\bar{y}_{k-1}\}$ as

$$(3.7) \quad \bar{y}_{k-1} = \theta_{k-1} y_{k-1} + (1 - \theta_{k-1}) H_{k,0}^{-1} s_{k-1},$$

where

$$(3.8) \quad \theta_{k-1} = \begin{cases} \frac{0.75 s_{k-1}^\top H_{k,0}^{-1} s_{k-1}}{s_{k-1}^\top H_{k,0}^{-1} s_{k-1} - s_{k-1}^\top y_{k-1}} & \text{if } s_{k-1}^\top y_{k-1} < 0.25 s_{k-1}^\top H_{k,0}^{-1} s_{k-1}, \\ 1 & \text{otherwise.} \end{cases}$$

Similar to Lemma 3.1, we can prove that

$$s_{k-1}^\top \bar{y}_{k-1} \geq 0.25 s_{k-1}^\top H_{k,0}^{-1} s_{k-1}.$$

Using s_j and \bar{y}_j , $j = k-p, \dots, k-1$, we define the SdLBFGS formula as

$$(3.9) \quad H_{k,i} = (I - \rho_j s_j \bar{y}_j^\top) H_{k,i-1} (I - \rho_j \bar{y}_j s_j^\top) + \rho_j s_j s_j^\top, \quad j = k - (p - i + 1); i = 1, \dots, p,$$

where $\rho_j = (s_j^\top \bar{y}_j)^{-1}$. As in the analysis in Lemma 3.1, by induction we can show that $H_{k,i} \succ 0$, $i = 1, \dots, p$. Note that when $k < p$, we use s_j and \bar{y}_j , $j = 1, \dots, k$, to execute the SdLBFGS update.

We next discuss the choice of $H_{k,0}$. A popular choice in the standard L-BFGS method is to set $H_{k,0} = \frac{s_{k-1}^\top y_{k-1}}{y_{k-1}^\top y_{k-1}} I$. Since $s_{k-1}^\top y_{k-1}$ may not be positive for nonconvex problems, we set

$$(3.10) \quad H_{k,0} = \gamma_k^{-1} I, \quad \text{where } \gamma_k = \max \left\{ \frac{y_{k-1}^\top y_{k-1}}{s_{k-1}^\top y_{k-1}}, \delta \right\} \geq \delta,$$

where $\delta > 0$ is a given constant.

To prove that $H_k = H_{k,p}$ generated by (3.9)–(3.10) satisfies assumptions AS.3 and AS.4, we need to make the following assumption.

AS.5. The function $F(x, \xi)$ is twice continuously differentiable with respect to x .

The stochastic gradient $g(x, \xi)$ is computed as $g(x, \xi) = \nabla_x F(x, \xi)$, and there exists a positive constant κ such that $\|\nabla_{xx}^2 F(x, \xi)\| \leq \kappa$ for any x, ξ .

Note that AS.5 is equivalent to requiring that $-\kappa I \preceq \nabla_{xx}^2 F(x, \xi) \preceq \kappa I$, rather than the strong convexity assumption $0 \prec \underline{\kappa} I \preceq \nabla_{xx}^2 F(x, \xi) \preceq \kappa I$ required in [10, 39]. The following lemma shows that the eigenvalues of H_k are bounded below away from zero under assumption AS.5.

LEMMA 3.2. Suppose that AS.5 holds. Given $H_{k,0}$ defined in (3.10), suppose that $H_k = H_{k,p}$ is updated through the SdLBFGS formula (3.9). Then all the eigenvalues of H_k satisfy

$$(3.11) \quad \lambda(H_k) \geq \left(\frac{4p\kappa^2}{\delta} + (4p+1)(\kappa + \delta) \right)^{-1}.$$

Proof. According to Lemma 3.1, $H_{k,i} \succ 0$, $i = 1, \dots, p$. To prove that the eigenvalues of H_k are bounded below away from zero, it suffices to prove that the eigenvalues of $B_k = H_k^{-1}$ are bounded from above. From the damped L-BFGS formula (3.9), $B_k = B_{k,p}$ can be computed recursively as

$$B_{k,i} = B_{k,i-1} + \frac{\bar{y}_j \bar{y}_j^\top}{s_j^\top \bar{y}_j} - \frac{B_{k,i-1} s_j s_j^\top B_{k,i-1}}{s_j^\top B_{k,i-1} s_j}, \quad j = k - (p - i + 1); i = 1, \dots, p,$$

starting from $B_{k,0} = H_{k,0}^{-1} = \gamma_k I$. Since $B_{k,0} \succ 0$, Lemma 3.1 indicates that $B_{k,i} \succ 0$ for $i = 1, \dots, p$. Moreover, the following inequalities hold:

$$(3.12) \quad \begin{aligned} \|B_{k,i}\| &\leq \left\| B_{k,i-1} - \frac{B_{k,i-1} s_j s_j^\top B_{k,i-1}}{s_j^\top B_{k,i-1} s_j} \right\| + \left\| \frac{\bar{y}_j \bar{y}_j^\top}{s_j^\top \bar{y}_j} \right\| \\ &\leq \|B_{k,i-1}\| + \left\| \frac{\bar{y}_j \bar{y}_j^\top}{s_j^\top \bar{y}_j} \right\| = \|B_{k,i-1}\| + \frac{\bar{y}_j^\top \bar{y}_j}{s_j^\top \bar{y}_j}. \end{aligned}$$

From the definition of \bar{y}_j in (3.7) and the facts that $s_j^\top \bar{y}_j \geq 0.25 s_j^\top B_{j+1,0} s_j$ and $B_{j+1,0} = \gamma_{j+1} I$ from (3.10), we have that for any $j = k - 1, \dots, k - p$

$$(3.13) \quad \frac{\bar{y}_j^\top \bar{y}_j}{s_j^\top \bar{y}_j} \leq 4 \frac{\|\theta_j y_j + (1 - \theta_j) B_{j+1,0} s_j\|^2}{s_j^\top B_{j+1,0} s_j} = 4\theta_j^2 \frac{y_j^\top y_j}{\gamma_{j+1} s_j^\top s_j} + 8\theta_j(1 - \theta_j) \frac{y_j^\top s_j}{s_j^\top s_j} + 4(1 - \theta_j)^2 \gamma_{j+1}.$$

Note that from (3.1) we have

$$y_j = \frac{\sum_{l=1}^{m_j} g(x_{j+1}, \xi_{j,l}) - g(x_j, \xi_{j,l})}{m_j} = \frac{1}{m_j} \left(\sum_{l=1}^{m_j} \bar{\nabla}_{xx}^2 F(x_j, \xi_{j,l}, s_j) \right) s_j,$$

where $\bar{\nabla}_{xx}^2 F(x_j, \xi_{j,l}, s_j) = \int_0^1 \nabla_{xx}^2 F(x_j + ts_j, \xi_{j,l}) dt$, because $g(x_{j+1}, \xi_{j,l}) - g(x_j, \xi_{j,l}) = \int_0^1 \frac{dg}{dt}(x_j + ts_j, \xi_{j,l}) dt = \int_0^1 \nabla_{xx}^2 F(x_j + ts_j, \xi_{j,l}) s_j dt$. Therefore, for any $j = k - 1, \dots, k - p$, from (3.13), and the facts that $0 < \theta_j \leq 1$ and $\delta \leq \gamma_{j+1} \leq \kappa + \delta$ and the assumption AS.5 it follows that

$$(3.14) \quad \begin{aligned} \frac{\bar{y}_j^\top \bar{y}_j}{s_j^\top \bar{y}_j} &\leq \frac{4\theta_j^2 \kappa^2}{\gamma_{j+1}} + 8\theta_j(1 - \theta_j)\kappa + 4(1 - \theta_j)^2 \gamma_{j+1} \\ &\leq \frac{4\theta_j^2 \kappa^2}{\delta} + 4[(1 - \theta_j^2)\kappa + (1 - \theta_j)^2 \delta] \leq \frac{4\kappa^2}{\delta} + 4(\kappa + \delta). \end{aligned}$$

Combining (3.12) and (8) yields

$$\|B_{k,i}\| \leq \|B_{k,i-1}\| + 4 \left(\frac{\kappa^2}{\delta} + \kappa + \delta \right).$$

By induction, we have that

$$\|B_k\| = \|B_{k,p}\| \leq \|B_{k,0}\| + 4p \left(\frac{\kappa^2}{\delta} + \kappa + \delta \right) \leq \frac{4p\kappa^2}{\delta} + (4p+1)(\kappa + \delta),$$

which implies (3.11). \square

We now prove that H_k is uniformly bounded above.

LEMMA 3.3. *Suppose that the assumption AS.5 holds. Given $H_{k,0}$ defined in (3.10), suppose that $H_k = H_{k,p}$ is updated through the SdLBFGS formula (3.9). Then H_k satisfies*

$$(3.15) \quad \lambda_{\max}(H_k) = \|H_k\| \leq \left(\frac{\alpha^{2p} - 1}{\alpha^2 - 1} \right) \frac{4}{\delta} + \frac{\alpha^{2p}}{\delta},$$

where $\alpha = (4\kappa + 5\delta)/\delta$, and $\lambda_{\max}(H_k)$ and $\|H_k\|$ denote, respectively, the maximum eigenvalue and operator norm $\|\cdot\|$ of H_k .

Proof. For notational simplicity, let $H = H_{k,i-1}$, $H^+ = H_{k,i}$, $s = s_j$, $\bar{y} = \bar{y}_j$, $\rho = (s_j^\top \bar{y}_j)^{-1} = (s^\top \bar{y})^{-1}$. Now (3.5) can be written as

$$H^+ = H - \rho(H\bar{y}s^\top + s\bar{y}^\top H) + \rho ss^\top + \rho^2(\bar{y}^\top H\bar{y})ss^\top.$$

Using the facts that $\|uv^\top\| = \|u\| \cdot \|v\|$ for any vectors u and v , $\rho s^\top s = \rho\|s\|^2 = \frac{s^\top s}{s^\top \bar{y}} \leq \frac{4}{\delta}$, and $\frac{\|\bar{y}\|^2}{s^\top \bar{y}} \leq 4(\frac{\kappa^2}{\delta} + \kappa + \delta) < \frac{4}{\delta}(\kappa + \delta)^2$, which follows from (8), we have that

$$\|H^+\| \leq \|H\| + \frac{2\|H\| \cdot \|\bar{y}\| \cdot \|s\|}{s^\top \bar{y}} + \frac{s^\top s}{s^\top \bar{y}} + \frac{s^\top s}{s^\top \bar{y}} \cdot \frac{\|H\| \cdot \|\bar{y}\|^2}{s^\top \bar{y}}.$$

Noting that $\frac{\|\bar{y}\| \|s\|}{s^\top \bar{y}} = [\frac{\|\bar{y}\|^2}{s^\top \bar{y}} \cdot \frac{\|s\|^2}{s^\top \bar{y}}]^{1/2}$, it follows that

$$\|H^+\| \leq \left(1 + 2 \cdot \frac{4}{\delta}(\kappa + \delta) + \left(\frac{4}{\delta}(\kappa + \delta) \right)^2 \right) \|H\| + \frac{4}{\delta} = (1 + (4\kappa + 4\delta)/\delta)^2 \|H\| + \frac{4}{\delta}.$$

Hence, by induction we obtain (3.15). \square

Lemmas 3.2 and 3.3 indicate that H_k generated by (3.7)–(3.9) satisfies assumption AS.3. Moreover, since y_{k-1} defined in (3.1) does not depend on random samplings in the k th iteration, it follows that H_k depends only on $\xi_{[k-1]}$ and assumption AS.4 is satisfied.

To analyze the cost of computing the step direction $H_k g_k$, note that from (3.9), H_k can be represented as

$$(3.16) \quad H_{k,i} = (I - \rho_j s_j \bar{y}_j^\top) H_{k,i-1} (I - \rho_j \bar{y}_j s_j^\top) + \rho_j s_j s_j^\top, \quad j = k - (p - i + 1); i = 1, \dots, p,$$

which is the same as the classical L-BFGS formula in (3.6), except that y_j is replaced by \bar{y}_j . Hence, we can compute the step direction by the two-loop recursion, implemented in the following procedure.

We now analyze the computational cost of Procedure 3.1. In step 2, the computation of γ_k involves $y_{k-1}^\top y_{k-1}$ and $s_{k-1}^\top y_{k-1}$, which take $2n$ multiplications. In step 3, from the definition of \bar{y}_k in (3.7), since $s_{k-1}^\top y_{k-1}$ has been obtained in a previous

Procedure 3.1. Step computation using SdLBFGS.

Input: Let x_k be a current iterate. Given the stochastic gradient g_{k-1} at iterate x_{k-1} , the random variable ξ_{k-1} , the batch size m_{k-1} , s_j , \bar{y}_j and ρ_j , $j = k-p, \dots, k-2$, and $u_0 = g_k$.

Output: $H_k g_k = v_p$.

- 1: Set $s_{k-1} = x_k - x_{k-1}$ and calculate y_{k-1} through (3.1)
- 2: Calculate γ_k through (3.10)
- 3: Calculate \bar{y}_{k-1} through (3.7) and $\rho_{k-1} = (s_{k-1}^\top \bar{y}_{k-1})^{-1}$
- 4: **for** $i = 0, \dots, \min\{p, k-1\} - 1$ **do**
- 5: Calculate $\mu_i = \rho_{k-i-1} u_i^\top s_{k-i-1}$
- 6: Calculate $u_{i+1} = u_i - \mu_i \bar{y}_{k-i-1}$
- 7: **end for**
- 8: Calculate $v_0 = \gamma_k^{-1} u_p$
- 9: **for** $i = 0, \dots, \min\{p, k-1\} - 1$ **do**
- 10: Calculate $\nu_i = \rho_{k-p+i} v_i^\top \bar{y}_{k-p+i}$
- 11: Calculate $v_{i+1} = v_i + (\mu_{p-i-1} - \nu_i) s_{k-p+i}$
- 12: **end for**

step, one only needs to compute $s_{k-1}^\top s_{k-1}$ and some scalar-vector products, thus the computation of \bar{y}_{k-1} takes $3n$ multiplications. Due to the fact that

$$\rho_{k-1}^{-1} = s_{k-1}^\top \bar{y}_{k-1} = \begin{cases} 0.25\gamma_k s_{k-1}^\top s_{k-1} & \text{if } s_{k-1}^\top y_{k-1} < 0.25\gamma_k s_{k-1}^\top s_{k-1}, \\ s_{k-1}^\top y_{k-1} & \text{otherwise,} \end{cases}$$

all involved computations have been done for ρ_{k-1} . Furthermore, the first loop steps 4–7 involve p scalar-vector multiplications and p vector inner products. So do the second loop steps 9–12. Including the product $\gamma_k^{-1} u_p$, the whole procedure takes $(4p+6)n$ multiplications.

Notice that in step 1 of Procedure 3.1, the computation of y_{k-1} involves the evaluation of $\sum_{i=1}^{m_{k-1}} g(x_k, \xi_{k-1,i})$, which requires m_{k-1} \mathcal{SFO} -calls. As a result, when Procedure 3.1 is plugged into SQN (Algorithm 2.1), the total number of \mathcal{SFO} -calls needed in the k th iteration becomes $m_k + m_{k-1}$. This leads to the following overall \mathcal{SFO} -calls complexity result for our SdLBFGS method.

THEOREM 3.4. *Suppose that AS.1, AS.2, and AS.5 hold. Let N_{sfo} denote the total number of \mathcal{SFO} -calls in SQN (Algorithm 2.1) in which Procedure 3.1 is used to compute $H_k g_k$. Under the same conditions as in Corollary 2.12, to achieve $\mathbb{E}[\|\nabla f(x_R)\|^2] \leq \epsilon$, $N_{sfo} \leq 2\bar{N}$, where \bar{N} satisfies (2.30), i.e., is $O(\epsilon^{-2})$.*

4. SdLBFGS with a variance reduction technique. Motivated by the recent advance of SVRG for nonconvex minimization proposed in [46] and [1], we now present a variance reduced SdLBFGS method, which we call SdLBFGS-VR, for solving (1.2). Here, the mini-batch stochastic gradient is defined as $g(x) = \frac{1}{|\mathcal{K}|} \sum_{i \in \mathcal{K}} \nabla f_i(x)$, where the subsample set $\mathcal{K} \subseteq [T]$ is randomly chosen from $\{1, \dots, T\}$. SdLBFGS-VR allows a constant step size and thus can accelerate the convergence speed of SdLBFGS. SdLBFGS-VR is summarized in Algorithm 4.1.

We now analyze the \mathcal{SFO} -calls complexity of Algorithm 4.1. We first analyze the convergence rate of SdLBFGS-VR, essentially following [46].

Algorithm 4.1. SdLBFGS with variance reduction (SdLBFGS-VR).

Input: $\tilde{x}_0 \in \mathbb{R}^d, \{\alpha_t\}_{t=0}^{m-1}$
Output: Iterate x chosen uniformly random from $\{x_t^{k+1} : t = 0, \dots, q-1; k = 0, \dots, N-1\}$

```

1: for  $k = 0, \dots, N-1$  do
2:    $x_0^{k+1} = \tilde{x}_k$ 
3:   compute  $\nabla f(\tilde{x}_k)$ 
4:   for  $t = 0, \dots, q-1$  do
5:     Sample a set  $\mathcal{K} \subseteq [T]$  with  $|\mathcal{K}| = m$ 
6:     Compute  $g_t^{k+1} = \nabla f_{\mathcal{K}}(x_t^{k+1}) - \nabla f_{\mathcal{K}}(\tilde{x}_k) + \nabla f(\tilde{x}_k)$  where  $\nabla f_{\mathcal{K}}(x_t^{k+1}) = \frac{1}{m} \sum_{i \in \mathcal{K}} \nabla f_i(x_t^{k+1})$ 
7:     Compute  $d_t^{k+1} = -H_t^{k+1} g_t^{k+1}$  through Procedure 3.1
8:     Set  $x_{t+1}^{k+1} = x_t^{k+1} + \alpha_t^{k+1} d_t^{k+1}$ 
9:   end for
10:  Set  $\tilde{x}_{k+1} = x_q^{k+1}$ 
11: end for

```

LEMMA 4.1. Suppose assumptions AS.1, AS.2, and AS.5 hold. Set $c_t^{k+1} = c_{t+1}^{k+1}(1 + \alpha_t^{k+1}\beta_t + 2L^2(\alpha_t^{k+1})^2\bar{\kappa}^2/m) + (\alpha_t^{k+1})^2L^3\bar{\kappa}^2/m$. Then

$$(4.1) \quad \Gamma_t^{k+1} \mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] \leq R_t^{k+1} - R_{t+1}^{k+1},$$

where $R_t^{k+1} = \mathbb{E}[f(x_t^{k+1}) + c_t^{k+1}\|x_t^{k+1} - \tilde{x}_k\|^2]$ and $\Gamma_t^{k+1} = \alpha_t^{k+1}\underline{\kappa} - c_{t+1}^{k+1}\alpha_t^{k+1}\bar{\kappa}^2/\beta_t - (\alpha_t^{k+1})^2L\bar{\kappa}^2 - 2c_{t+1}^{k+1}(\alpha_t^{k+1})^2\bar{\kappa}^2$ for any $\beta_t > 0$.

Proof. It follows from AS.1 and the fact that $\underline{\kappa}I \preceq H_t^{k+1} \preceq \bar{\kappa}I$, for any $t = 0, \dots, q-1; k = 0, \dots, N-1$, which follows under assumption AS.5, from Lemmas 3.2 and 3.3, that

$$\begin{aligned}
 \mathbb{E}[f(x_{t+1}^{k+1})] &\leq \mathbb{E}\left[f(x_t^{k+1}) + \langle \nabla f(x_t^{k+1}), x_{t+1}^{k+1} - x_t^{k+1} \rangle + \frac{L}{2}\|x_{t+1}^{k+1} - x_t^{k+1}\|^2\right] \\
 &= \mathbb{E}\left[f(x_t^{k+1}) - \alpha_t^{k+1}\langle \nabla f(x_t^{k+1}), H_t^{k+1}g_t^{k+1} \rangle + \frac{(\alpha_t^{k+1})^2L}{2}\|H_t^{k+1}g_t^{k+1}\|^2\right] \\
 (4.2) \quad &\leq \mathbb{E}\left[f(x_t^{k+1}) - \alpha_t^{k+1}\underline{\kappa}\|\nabla f(x_t^{k+1})\|^2 + \frac{(\alpha_t^{k+1})^2L\bar{\kappa}^2}{2}\|g_t^{k+1}\|^2\right].
 \end{aligned}$$

Moreover, for any $\beta_t > 0$, since g_t^{k+1} is an unbiased estimate of $\nabla f(x_t^{k+1})$,

$$\begin{aligned}
 (4.3) \quad \mathbb{E}[\|x_{t+1}^{k+1} - \tilde{x}_k\|^2] &= \mathbb{E}[\|x_{t+1}^{k+1} - x_t^{k+1}\|^2 + \|x_t^{k+1} - \tilde{x}_k\|^2 + 2\langle x_{t+1}^{k+1} - x_t^{k+1}, x_t^{k+1} - \tilde{x}_k \rangle] \\
 &= \mathbb{E}[(\alpha_t^{k+1})^2\|H_t^{k+1}g_t^{k+1}\|^2 + \|x_t^{k+1} - \tilde{x}_k\|^2 \\
 &\quad - 2\alpha_t^{k+1}\langle H_t^{k+1}\nabla f(x_t^{k+1}), x_t^{k+1} - \tilde{x}_k \rangle] \\
 &\leq \mathbb{E}[(\alpha_t^{k+1})^2\bar{\kappa}^2\|g_t^{k+1}\|^2 + \|x_t^{k+1} - \tilde{x}_k\|^2 \\
 &\quad + \alpha_t^{k+1}\mathbb{E}[\beta_t^{-1}\|H_t^{k+1}\nabla f(x_t^{k+1})\|^2 + \beta_t\|x_t^{k+1} - \tilde{x}_k\|^2].
 \end{aligned}$$

Furthermore, the following inequality holds:

$$\begin{aligned}
\mathbb{E}[\|g_t^{k+1}\|^2] &= \mathbb{E}[\|\nabla f_{\mathcal{K}}(x_t^{k+1}) - \nabla f(x_t^{k+1}) + \nabla f(x_t^{k+1}) - \nabla f_{\mathcal{K}}(\tilde{x}_k) + \nabla f(\tilde{x}_k)\|^2] \\
&\leq 2\mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] + 2\mathbb{E}[\|\nabla f_{\mathcal{K}}(x_t^{k+1}) - \nabla f_{\mathcal{K}}(\tilde{x}_k) \\
&\quad - (\nabla f(x_t^{k+1}) - \nabla f(\tilde{x}_k))\|^2] \\
&= 2\mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] + 2\mathbb{E}[\|\nabla f_{\mathcal{K}}(x_t^{k+1}) - \nabla f_{\mathcal{K}}(\tilde{x}_k) \\
&\quad - \mathbb{E}[\nabla f_{\mathcal{K}}(x_t^{k+1}) - \nabla f_{\mathcal{K}}(\tilde{x}_k)]\|^2] \\
&\leq 2\mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] + 2\mathbb{E}[\|\nabla f_{\mathcal{K}}(x_t^{k+1}) - \nabla f_{\mathcal{K}}(\tilde{x}_k)\|^2] \\
&= 2\mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] + \frac{2}{m}\mathbb{E}[\|\nabla f_i(x_t^{k+1}) - \nabla f_i(\tilde{x}_k)\|^2] \\
(4.4) \quad &\leq 2\mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] + \frac{2L^2}{m}\mathbb{E}[\|x_t^{k+1} - \tilde{x}_k\|^2].
\end{aligned}$$

Combining (4.2), (4.3), and (4.4) yields that

$$\begin{aligned}
R_{t+1}^{k+1} &\leq \mathbb{E}[f(x_{t+1}^{k+1}) + c_{t+1}^{k+1}\|x_{t+1}^{k+1} - \tilde{x}_k\|^2] \\
&\leq \mathbb{E}\left[f(x_t^{k+1}) - \alpha_t^{k+1}\underline{\kappa}\|\nabla f(x_t^{k+1})\|^2 + \frac{(\alpha_t^{k+1})^2 L \bar{\kappa}^2}{2}\|g_t^{k+1}\|^2\right. \\
&\quad \left.+ c_{t+1}^{k+1}(\alpha_t^{k+1})^2 \bar{\kappa}^2 \mathbb{E}[\|g_t^{k+1}\|^2] + c_{t+1}^{k+1} \mathbb{E}[\|x_t^{k+1} - \tilde{x}_k\|^2]\right] \\
&\quad + \alpha_t^{k+1} c_{t+1}^{k+1} \mathbb{E}[\beta_t^{-1} \bar{\kappa}^2 \|\nabla f(x_t^{k+1})\|^2 + \beta_t \|x_t^{k+1} - \tilde{x}_k\|^2] \\
&= \mathbb{E}[f(x_t^{k+1})] - (\alpha_t^{k+1} \underline{\kappa} - c_{t+1}^{k+1} \alpha_t^{k+1} \bar{\kappa}^2 / \beta_t) \mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] \\
&\quad + ((\alpha_t^{k+1})^2 L \bar{\kappa}^2 / 2 + c_{t+1}^{k+1} (\alpha_t^{k+1})^2 \bar{\kappa}^2) \mathbb{E}[\|g_t^{k+1}\|^2] \\
&\quad + c_{t+1}^{k+1} (1 + \alpha_t^{k+1} \beta_t) \mathbb{E}[\|x_t^{k+1} - \tilde{x}_k\|^2] \\
&\leq \mathbb{E}[f(x_t^{k+1})] + (c_{t+1}^{k+1} 1 + c_{t+1}^{k+1} \alpha_t^{k+1} \beta_t + 2c_{t+1}^{k+1} L^2 (\alpha_t^{k+1})^2 \bar{\kappa}^2 / m \\
&\quad + (\alpha_t^{k+1})^2 L^3 \bar{\kappa}^2 / m) \mathbb{E}[\|x_t^{k+1} - \tilde{x}_k\|^2] \\
&\quad - (\alpha_t^{k+1} \underline{\kappa} - c_{t+1}^{k+1} \alpha_t^{k+1} \bar{\kappa}^2 / \beta_t - (\alpha_t^{k+1})^2 L \bar{\kappa}^2 - 2c_{t+1}^{k+1} (\alpha_t^{k+1})^2 \bar{\kappa}^2) \mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] \\
&= R_t^{k+1} - \Gamma_t^{k+1} \mathbb{E}[\|\nabla f(x_t^{k+1})\|^2],
\end{aligned}$$

which further implies (4.1). \square

THEOREM 4.2. *Suppose assumptions AS.1, AS.2, and AS.5 hold. Set $\beta_t = \beta = \frac{L\bar{\kappa}}{T^{1/3}}$, $c_q^{k+1} = c_q = 0$. Suppose that there exist two positive constants $\nu, \mu_0 \in (0, 1)$ such that*

$$(4.5) \quad \left(1 - \frac{\nu \bar{\kappa}}{\mu_0 \underline{\kappa}}\right) \underline{\kappa} \geq \mu_0 \bar{\kappa} (e - 1) + \mu_0 m \bar{\kappa} + 2\mu_0^2 \bar{\kappa} m (e - 1)$$

holds. Set $\alpha_t^{k+1} = \alpha = \frac{\mu_0 m}{L \bar{\kappa} T^{2/3}}$, and $q = \lfloor \frac{T}{3\mu_0 m} \rfloor$.

$$(4.6) \quad \mathbb{E}[\|\nabla f(x)\|^2] \leq \frac{T^{2/3} L [f(x_0) - f(x_*)]}{q N m \nu}.$$

Proof. Denote $\theta = \alpha \beta + 2\alpha^2 L^2 \bar{\kappa}^2 / m$. It then follows that $\theta = \mu_0 m / T + 2\mu_0^2 m / T^{4/3} \leq 3\mu_0 m / T$, and $(1 + \theta)^q \leq e$, where e is the Euler's number. Because $c_q = 0$, for any $k \geq 0$, we have

$$\begin{aligned}
c_0 := c_0^{k+1} &= \frac{\alpha^2 L^3 \bar{\kappa}^2}{m} \cdot \frac{(1+\theta)^q - 1}{\theta} = \frac{\mu_0^2 m L((1+\theta)^q - 1)}{T^{4/3} \theta} = \frac{\mu_0^2 L((1+\theta)^q - 1)}{(\mu_0 T^{1/3} + 2\mu_0^2)} \\
&\leq \frac{\mu_0 L((1+\theta)^q - 1)}{T^{1/3}} \leq \frac{\mu_0 L(e - 1)}{T^{1/3}}.
\end{aligned}$$

Therefore, it follows that

$$\min_t \Gamma_t^{k+1} \geq \frac{\nu m}{LT^{2/3}}.$$

As a result, we have

$$\sum_{t=0}^{q-1} \mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] \leq \frac{R_0^{k+1} - R_q^{k+1}}{\min_t \Gamma_t^{k+1}} = \frac{\mathbb{E}[f(\tilde{x}_k) - f(\tilde{x}_{k+1})]}{\min_t \Gamma_t^{k+1}},$$

which further yields that

$$\mathbb{E}[\|\nabla f(x)\|^2] = \frac{1}{qN} \sum_{k=0}^{N-1} \sum_{t=0}^{q-1} \mathbb{E}[\|\nabla f(x_t^{k+1})\|^2] \leq \frac{f(x_0) - f(x_*)}{qN \min_t \Gamma_t^{k+1}} \leq \frac{T^{2/3} L[f(x_0) - f(x_*)]}{qNm\nu}.$$

□

COROLLARY 4.3. *Under the same conditions as Theorem 4.2, to achieve $\mathbb{E}[\|\nabla f(x)\|^2] \leq \epsilon$, the total number of component gradient evaluations required in Algorithm 4.1 is $O(T^{2/3}/\epsilon)$.*

Proof. From Theorem 4.2, it follows that to obtain an ϵ -solution, the outer iteration number N of Algorithm 4.1 should be in the order of $O(\frac{T^{2/3}}{qm\epsilon}) = O(\frac{T^{-1/3}}{\epsilon})$, which is due to the fact that $qm = O(T)$. As a result, the total number of component gradient evaluations is $(T + qm)N$, which is $O(T^{2/3}/\epsilon)$. □

5. Numerical experiments. In this section, we empirically study the performance of the proposed SdLBFGS and SdLBFGS-VR methods. We compare SdLBFGS with SGD with g_k given by (2.5) using a diminishing step size $\alpha_k = \beta/k$ in both methods, for solving the following nonconvex support vector machine (SVM) problem with a sigmoid loss function, which has been considered in [21, 37]:

$$(5.1) \quad \min_{x \in \mathbb{R}^n} f(x) := \mathbb{E}_{u,v} [1 - \tanh(v\langle x, u \rangle)] + \lambda \|x\|_2^2,$$

where $\lambda > 0$ is a regularization parameter, $u \in \mathbb{R}^n$ denotes the feature vector, and $v \in \{-1, 1\}$ refers to the corresponding label. In our experiments, λ was chosen as 10^{-4} . We also compare SdLBFGS-VR with SVRG [46], using a constant step size in both methods, for solving

$$(5.2) \quad \min_{x \in \mathbb{R}^n} \frac{1}{T} \sum_{i=1}^T f_i(x) + \lambda \|x\|^2,$$

where $f_i(x) = 1 - \tanh(v_i \langle x, u_i \rangle)$, $i = 1, \dots, T$. In all of our tests, we used the same mini-batch size m at every iteration. All the algorithms were implemented in MATLAB R2013a on a PC with a 2.60 GHz Intel microprocessor and 8 GB of memory.

5.1. Numerical results for SdLBFGS on synthetic data. In this subsection, we report numerical results for SdLBFGS and SGD for solving (5.1) on synthetic data for problems of dimension $n = 500$. We set the initial point for both methods to




FIG. 1. Comparison on a synthetic dataset of SGD and SdLBFGS variants with different memory size p , with respect to the squared norm of the gradient and the percent of correctly classified data, respectively. A step size of $\alpha_k = 10/k$ and batch size of $m = 100$ were used for all SdLBFGS variants. Step sizes of $10/k$ and $20/k$ were used for SGD.

$x_1 = 5 * \bar{x}_1$, where \bar{x}_1 was drawn from the uniform distribution over $[0, 1]^n$. We generated training and testing points (u, v) in the following manner. We first generated a sparse vector u with 5% nonzero components following the uniform distribution on $[0, 1]^n$, and then set $v = \text{sign}(\langle \bar{x}, u \rangle)$ for some $\bar{x} \in \mathbb{R}^n$ drawn from the uniform distribution on $[-1, 1]^n$. Every time we computed a stochastic gradient of the sigmoid loss function in (5.1), we accessed m of these data points (u, v) without replacement. Drawing data points without replacement is a common practice in testing the performance of stochastic algorithms, although in this case the random samples are not necessarily independent.

In Figure 1 we compare the performance of SGD and SdLBFGS with various memory sizes p . The batch size was set to $m = 100$ and the step size to both $\alpha_k = 10/k$ and $\alpha_k = 20/k$ for SGD and $10/k$ for SdLBFGS. In the left figure we plot the squared norm of the gradient (SNG) versus the number of iterations, up to a total of 1000. The SNG was computed using $N = 5000$ randomly generated testing points $(u_i, v_i), i = 1, \dots, N$, as

$$(5.3) \quad \text{SNG} = \left\| \frac{1}{N} \sum_{i=1}^N \nabla_x F(\tilde{x}; u_i, v_i) + 2\lambda \tilde{x} \right\|^2,$$

where \tilde{x} is the output of the algorithm and $F(x; u, v) = 1 - \tanh(v\langle x, u \rangle)$. In the right figure we plot the percentage of correctly classified testing data. From Figure 1 we can see that increasing the memory size improves the performance of SdLBFGS. When the memory size $p = 0$, we have $H_k = \gamma_k^{-1}I$ and SdLBFGS reduces to SGD with an adaptive step size. SdLBFGS with memory size $p = 1$ and 3 oscillates quite a lot. The variants of SdLBFGS with memory size $p = 5, 10$, and 20 all oscillate

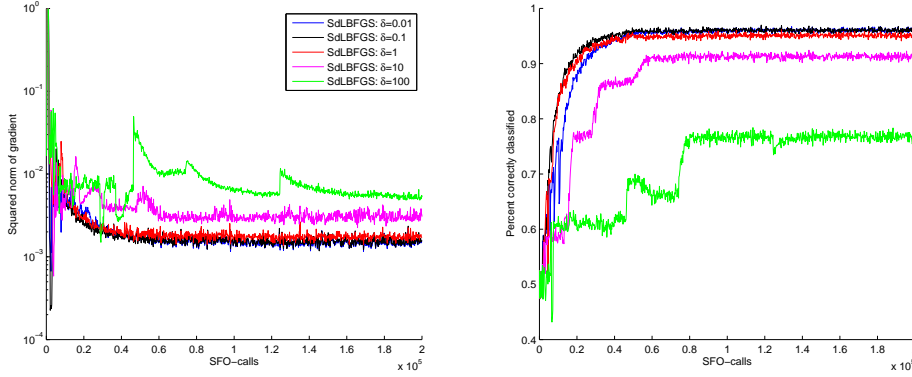


FIG. 2. Comparison on a synthetic dataset of SdLBFGS variants with different initial Hessian approximations, by varying the value of δ , with respect to the squared norm of the gradient and the percent of correctly classified data, respectively. A step size of $\alpha_k = 10/k$, memory size of $p = 20$, and batch size of $m = 100$ were used for all SdLBFGS variants.

less and perform quite similarly, and they all significantly outperform SdLBFGS with $p = 0$ and $p = 1$.

In Figure 2, we report the performance of SdLBFGS with different δ used in (3.10). From Figure 2 we see that SdLBFGS performs best with small δ such as $\delta = 0.01, 0.1$, and 1.

In Figure 3 we report the effect of the batch size m on the performance of SGD and SdLBFGS with memory size $p = 20$. For SdLBFGS, the left figure shows that $m = 500$ gives the best performance among the three choices 50, 100, and 500 tested, with respect to the total number of iterations taken. This is because a larger batch size leads to gradient estimation with lower variance. The right figure shows that if the total number of \mathcal{SFO} -calls is fixed, then because of the trade-off between the number of iterations and the batch size, i.e., because the number of iterations is proportional to the reciprocal of the batch size, the SdLBFGS variant corresponding to $m = 100$ slightly outperforms the $m = 500$ variant.

In Figure 4, we report the percentage of correctly classified data for 5000 randomly generated testing points. The results are consistent with the one shown in the left figure of Figure 3, i.e., the ones with a lower squared norm of the gradient give a higher percentage of correctly classified data.

Moreover, we also counted the number of steps taken by SdLBFGS in which $s_{k-1}^\top y_{k-1} < 0$. We set the total number of iterations to 1000 and tested the effect of the memory size and batch size of SdLBFGS on the number of such steps. For fixed batch size $m = 50$, the average numbers of such steps over 10 runs of SdLBFGS were respectively equal to (178, 50, 36, 42, 15) when the memory sizes were (1, 3, 5, 10, 20). For fixed memory size $p = 20$, the average numbers of such steps over 10 runs of SdLBFGS were respectively equal to (15, 6, 1) when the batch sizes were (50, 100, 500). Therefore, the number of such steps roughly decreases as the memory size p and the batch size m increase. This is to be expected because as p increases, there is less negative effect caused by “limited memory,” and as m increases, the gradient estimation has lower variance.




FIG. 3. Comparison of SGD and SdLBFGS variants with different batch size m on a synthetic dataset. The memory size of SdLBFGS was $p = 20$ and the step size of SdLBFGS was $\alpha_k = 10/k$, while the step size of SGD was $\alpha_k = 20/k$.




FIG. 4. Comparison of correct classification percentage on a synthetic dataset by SGD and SdLBFGS with different batch sizes. The memory size of SdLBFGS was $p = 20$ and the step size of SdLBFGS was $\alpha_k = 10/k$. The step size of SGD was $\alpha_k = 20/k$.

5.2. Numerical results for SdLBFGS on the RCV1 dataset. In this subsection, we compare SGD and SdLBFGS for solving (5.1) on a real dataset: RCV1 [33], which is a collection of newswire articles produced by Reuters in 1996–1997. In our tests, we used a subset¹ of RCV1 used in [11] that contains 9625 articles with 29992 distinct words. The articles are classified into four categories, C15, ECAT, GCAT, and MCAT, each with 2022, 2064, 2901, and 2638 articles, respectively. We consider the binary classification problem of predicting whether an article is in the second and fourth categories, i.e., the entry of each label vector is 1 if a given article appears in category MCAT or ECAT and is -1 otherwise. We used 60% of the articles (5776) as training data and the remaining 40% (3849) as testing data.

In Figure 5, we compare SdLBFGS with various memory sizes and SGD on the RCV1 dataset. For SGD and SdLBFGS, we used the step size $\alpha_k = 10/k$ and the batch size $m = 100$. We also used a second step size of $20/k$ for SGD. Note that the SNG computed via (5.3) uses $N = 3849$ testing data. The left figure shows that for the RCV1 dataset, increasing the memory size improves the performance of

¹Downloaded from <http://www.cad.zju.edu.cn/home/dengcai/Data/TextData.html>.




FIG. 5. Comparison of SdLBFGS variants with different memory sizes on the RCV1 dataset. The step size of SdLBFGS was $\alpha_k = 10/k$ and the batch size was $m = 100$.




FIG. 6. Comparison on a RCV dataset of SdLBFGS variants with different initial Hessian approximations, by varying the value of δ , with respect to the squared norm of the gradient and the percent of correctly classified data, respectively. A step size of $\alpha_k = 10/k$, memory size of $p = 10$, and batch size of $m = 100$ were used for all SdLBFGS variants.

SdLBFGS. The performance of SdLBFGS with memory sizes $p = 10, 20$ was similar, although for $p = 10$ it was slightly better. The right figure also shows that larger memory sizes can achieve higher correct classification percentages.

In Figure 6, we report the performance of SdLBFGS on RCV1 dataset with different δ used in (3.10). Similar to Figure 2, we see from Figure 6 that SdLBFGS works best for small δ such as $\delta = 0.01, 0.1$, and 1 .

Figure 7 compares SGD and SdLBFGS with different batch sizes. The step size of SGD and SdLBFGS was set to $\alpha_k = 20/k$ and $\alpha_k = 10/k$, respectively. The memory size of SdLBFGS was chosen as $p = 10$. We tested SGD with batch size $m = 1, 50, 100$, and SdLBFGS with batch size $m = 1, 50, 75, 100$. From Figure 7 we can see that SGD performs worse than SdLBFGS. For SdLBFGS, from Figure 7(a) we observe that larger batch sizes give better results in terms of SNG. If we fix the total number of SFO-calls to 2×10^4 , SdLBFGS with $m = 1$ performs the worst among the




FIG. 7. Comparison of SGD and SdLBFGS with different batch sizes on the RCV1 dataset. For SdLBFGS the step size was $\alpha_k = 10/k$ and the memory size was $p = 10$. For SGD the step size was $\alpha_k = 20/k$.




FIG. 8. Comparison of correct classification percentage by SGD and SdLBFGS with different batch sizes on the RCV1 dataset. For SdLBFGS the step size was $\alpha_k = 10/k$ and the memory size was $p = 10$. For SGD the step size was $\alpha_k = 20/k$.

different batch sizes and exhibits dramatic oscillation. The performance gets much better when the batch size becomes larger. In this set of tests, the performance with $m = 50, 75$ was slightly better than $m = 100$. One possible reason is that for the same number of SFO-calls, a smaller batch size leads to a larger number of iterations and thus gives better results.

In Figure 8, we report the percentage of correctly classified data points for both SGD and SdLBFGS with different batch sizes. These results are consistent with the ones in Figure 7. Roughly speaking, the algorithm that gives a lower SNG leads to a higher percentage of correctly classified data points.

We also counted the number of steps taken by SdLBFGS in which $s_{k-1}^\top y_{k-1} < 0$. We again set the total number of iterations of SdLBFGS to 1000. For fixed batch size $m = 50$, the average numbers of such steps over 10 runs of SdLBFGS were respectively equal to (3, 5, 8, 6, 1) when the memory sizes were (1, 3, 5, 10, 20). For fixed memory size $p = 10$, the average numbers of such steps over 10 runs of SdLBFGS




FIG. 9. Comparison of SdLBFGS-VR with different memory sizes p on RCV1 dataset. The step size is set as $\alpha = 0.1$. We set $q = 115$, and batch size $m = 50$, so that $T \approx qm$.




FIG. 10. Comparison of SdLBFGS-VR with different batch sizes on RCV1 dataset. We set the step size as $\alpha = 0.1$ and the memory size as $p = 20$.

were respectively equal to $(308, 6, 2)$ when the batch sizes were $(1, 50, 75)$. This is qualitatively similar to our observations in section 5.1, except that for the fixed batch size $m = 50$, fewer such steps were required by the SdLBFGS variants with memory sizes $p = 1$ and $p = 3$ compared with $p = 5$ and $p = 10$.

5.3. Numerical results for SdLBFGS-VR on the RCV1 dataset. In this subsection, we compare SdLBFGS-VR, SVRG [46], and SdLBFGS for solving (5.2) on the RCV1 dataset with $T = 5776$. We here follow the same strategy as suggested in [46] to initialize SdLBFGS-VR and SVRG. In particular, we run SGD with step size $1/t$ and batch size 20 for T iterations to get the initial point for SdLBFGS-VR and SVRG, where t denotes the iteration counter. In all the tests, we use a constant step size α for both methods. The comparison results are shown in Figures 9–12. In these figures, the “SFO-calls” in the x -axis include both the number of stochastic gradients and T gradient evaluations of the individual component functions when computing the full gradient $\nabla f(\tilde{x}_k)$ in each outer loop.




FIG. 11. Comparison of SdLBFGS-VR and SVRG on RCV1 dataset with different step sizes. The memory size of SdLBFGS-VR is set as $p = 20$. For both algorithms, we set the inner iteration number $q = 115$ and the batch size $m = 50$.




FIG. 12. Comparison of SdLBFGS-VR and SdLBFGS on RCV1 dataset. For both algorithms, the memory size is $p = 20$ and batch size is $m = 50$. For SdLBFGS-VR, the inner iteration number is $q = 115$.

Figure 9 compares the performance of SdLBFGS-VR with different memory size p . It shows that the L-BFGS improves performance, even when $p = 1$. Moreover, larger memory size usually provides better performance, but the difference is not very significant.

Figure 10 compares the performance of SdLBFGS-VR with different batch sizes m and shows that SdLBFGS-VR is not very sensitive to m . In these tests, we always set $q = \lfloor T/m \rfloor$.

The impact of step size on SdLBFGS-VR and SVRG is shown in Figure 11 for three step sizes: $\alpha = 0.1, 0.01$, and 0.001 . Clearly, for the same step size, SdLBFGS-VR gives better results than SVRG. From our numerical tests, we also observed that neither SdLBFGS-VR nor SVRG is stable when $\alpha \geq 1$.

In Figure 12, we report the performance of SdLBFGS-VR with different constant step sizes α for $\alpha = 0.1, 0.01, 0.001$ and SdLBFGS with different diminishing step sizes




FIG. 13. Comparison of SdLBFGS-VR with different memory sizes p on MNIST dataset. The step size is set as $\alpha = 0.1$. We set $q = 1200$ and batch size $m = 50$, so that $T = qm$.




FIG. 14. Comparison of SdLBFGS-VR with different batch sizes on MNIST dataset. We set the step size as $\alpha = 0.1$ and the memory size as $p = 20$.

β/k for $\beta = 10, 1, 0.1$, since SdLBFGS needs a diminishing step size to guarantee convergence. We see there that SdLBFGS-VR usually performs better than SdLBFGS. The performance of SdLBFGS with $\beta = 10$ is in fact already very good but still is inferior to SdLBFGS-VR. This indicates that the variance reduction technique is indeed helpful.

5.4. Numerical results for SdLBFGS-VR on the MNIST dataset. In this section, we report the numerical results of SdLBFGS-VR for solving a multiclass classification problem using neural networks on a standard testing dataset MNIST.² All the experimental settings are the same as in [46]. In particular, $T = 60000$. The numerical results are reported in Figures 13–16, and their purposes are the same as Figures 9–12. From these figures, we have similar observations as those from Figures 9–12. Note that in Figures 13–16, the “SFO-calls” in the x -axis again include both the number of stochastic gradients and T gradient evaluations of the individual component functions when computing the full gradient $\nabla f(\tilde{x}_k)$ in each outer loop.

6. Conclusions. In this paper we proposed a general framework for stochastic quasi-Newton methods for nonconvex stochastic optimization. Global convergence, iteration complexity, and SFO-calls complexity were analyzed under different conditions on the step size and the output of the algorithm. Specifically, a stochastic damped limited-memory BFGS method was proposed, which falls under the pro-

²<http://yann.lecun.com/exdb/mnist>.




FIG. 15. Comparison of SdLBFGS-VR and SVRG on MNIST dataset with different step sizes. The memory size of SdLBFGS-VR is set as $p = 20$. For both algorithms, we set the inner iteration number $q = 1200$ and the batch size $m = 50$.




FIG. 16. Comparison of SdLBFGS-VR and SdLBFGS on MNIST dataset. For both algorithms, the memory size is $p = 20$ and batch size is $m = 50$. For SdLBFGS-VR, the inner iteration number is $q = 1200$.

posed framework and does not generate H_k explicitly. The damping technique was used to preserve the positive definiteness of H_k , without requiring the original problem to be convex. A variance reduced stochastic L-BFGS method was also proposed for solving the empirical risk minimization problem. Encouraging numerical results were reported for solving nonconvex classification problems using SVM and neural networks.

Acknowledgments. The authors are grateful to two anonymous referees for their insightful comments and constructive suggestions that have improved the presentation of this paper greatly. The authors also thank Conghui Tan for helping conduct the numerical tests in section 5.4.

REFERENCES

- [1] Z. ALLEN-ZHU AND E. HAZAN, *Variance reduction for faster non-convex optimization*, in Proceedings of ICML, 2016.
- [2] F. BACH, *Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression*, J. Mach. Learn. Res., 15 (2014), pp. 595–627.
- [3] F. BACH AND E. MOULINES, *Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$* , in Neural Inf. Process. Ser., MIT Press, Cambridge, MA, 2013.
- [4] F. BASTIN, C. CIRILLO, AND P. L. TOINT, *Convergence theory for nonconvex stochastic programming with an application to mixed logit*, Math. Program., 108 (2006), pp. 207–234.

- [5] A. BORDES, L. BOTTOU, AND P. GALLINARI, *SGD-QN: Careful quasi-Newton stochastic gradient descent*, J. Mach. Learn. Res., 10 (2009), pp. 1737–1754.
- [6] L. BOTTOU, *Online algorithms and stochastic approximations*, in Online Learning and Neural Networks, D. Saad, ed., Cambridge University Press, Cambridge, UK, 1998.
- [7] D. BROWNSTONE, D. S. BUNCH, AND K. TRAIN, *Joint mixed logit models of stated and revealed preferences for alternative-fuel vehicles*, Transport. Res. B, 34 (2000), pp. 315–338.
- [8] C. G. BROYDEN, *The convergence of a class of double-rank minimization algorithms*, J. Inst. Math. Appl., 6 (1970), pp. 76–90.
- [9] R. H. BYRD, G. CHIN, W. NEVEITT, AND J. NOCEDAL, *On the use of stochastic Hessian information in optimization methods for machine learning*, SIAM J. Optim., 21 (2011), pp. 977–995.
- [10] R. H. BYRD, S. L. HANSEN, J. NOCEDAL, AND Y. SINGER, *A stochastic quasi-Newton method for large-scale optimization*, SIAM J. Optim., 26 (2016), pp. 1008–1031.
- [11] D. CAI AND X. HE, *Manifold adaptive experimental design for text categorization*, IEEE Trans. Knowledge Data Engineering, 4 (2012), pp. 707–719.
- [12] K. L. CHUNG, *On a stochastic approximation method*, Ann. Math. Statist., 25 (1954), pp. 463–483.
- [13] C. D. DANG AND G. LAN, *Stochastic block mirror descent methods for nonsmooth and stochastic optimization*, SIAM J. Optim., 25 (2015), pp. 856–881.
- [14] A. DEFAZIO, F. BACH, AND S. LACOSTE-JULIEN, *SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives*, in Neural Inf. Process. Ser. MIT Press, Cambridge, MA, 2014.
- [15] J. C. DUCHI, E. HAZAN, AND Y. SINGER, *Adaptive subgradient methods for online learning and stochastic optimization*, J. Mach. Learn. Res., 12 (2011), pp. 2121–2159.
- [16] R. DURRETT, *Probability: Theory and Examples*, Cambridge University Press, London, 2010.
- [17] Y. ERMOLIEV, *Stochastic quasigradient methods and their application to system optimization*, Stochastics, 9 (1983), pp. 1–36.
- [18] R. FLETCHER, *A new approach to variable metric algorithms*, Computer J., 13 (1970), pp. 317–322.
- [19] A. A. GAIVORONSKI, *Nonstationary stochastic programming problems*, Kibernetika, 4 (1978), pp. 89–92.
- [20] S. GHADIMI AND G. LAN, *Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization, I: A generic algorithmic framework*, SIAM J. Optim., 22 (2012), pp. 1469–1492.
- [21] S. GHADIMI AND G. LAN, *Stochastic first- and zeroth-order methods for nonconvex stochastic programming*, SIAM J. Optim., 15 (2013), pp. 2341–2368.
- [22] S. GHADIMI AND G. LAN, *Accelerated gradient methods for nonconvex nonlinear and stochastic programming*, Math. Program., 156 (2016), pp. 59–99.
- [23] S. GHADIMI, G. LAN, AND H. ZHANG, *Mini-batch stochastic approximation methods for non-convex stochastic composite optimization*, Math. Program., 155 (2016), pp. 267–305.
- [24] D. GOLDFARB, *A family of variable metric updates derived by variational means*, Math. Comp., 24 (1970), pp. 23–26.
- [25] R. M. GOWER, D. GOLDFARB, AND P. RICHTÁRIK, *Stochastic block BFGS: Squeezing more curvature out of data*, in Proceedings of ICML, 2016.
- [26] D. A. HENSHER AND W. H. GREENE, *The mixed logit model: The state of practice*, Transportation, 30 (2003), pp. 133–176.
- [27] R. JOHNSON AND T. ZHANG, *Accelerating stochastic gradient descent using predictive variance reduction*, Neural Inf. Process. Ser., MIT Press, Cambridge, MA, 2013.
- [28] A. JUDITSKY, A. NAZIN, A. B. TSYBAKOV, AND N. VAYATIS, *Recursive aggregation of estimators via the mirror descent algorithm with average*, Probl. Inf. Transm., 41 (2005), pp. 368–384.
- [29] A. JUDITSKY, P. RIGOLLET, AND A. B. TSYBAKOV, *Learning by mirror averaging*, Ann. Statist., 36 (2008), pp. 2183–2206.
- [30] G. LAN, *An optimal method for stochastic composite optimization*, Math. Program., 133 (2012), pp. 365–397.
- [31] G. LAN, A. S. NEMIROVSKI, AND A. SHAPIRO, *Validation analysis of mirror descent stochastic approximation method*, Math. Program., 134 (2012), pp. 425–458.
- [32] N. LE ROUX, M. SCHMIDT, AND F. BACH, *A stochastic gradient method with an exponential convergence rate for strongly-convex optimization with finite training sets*, in Neural Inf. Process. Ser., MIT Press, Cambridge, MA, 2012.
- [33] D. LEWIS, Y. YANG, T. ROSE, AND F. LI, *RCV1: A new benchmark collection for text categorization research*, J. Mach. Learn. Res., 5 (2004), pp. 361–397.
- [34] D. C. LIU AND J. NOCEDAL, *On the limited memory BFGS method for large scale optimization*, Math. Program. Ser. B, 45 (1989), pp. 503–528.

- [35] A. LUCCHI, B. MCWILLIAMS, AND T. HOFMANN, *A Variance Reduced Stochastic Newton Method*, preprint, arXiv:1503.08316, 2015.
- [36] J. MAIRAL, F. BACH, J. PONCE, AND G. SAPIRO, *Online dictionary learning for sparse coding*, in Proceedings of ICML, 2009.
- [37] L. MASON, J. BAXTER, P. BARTLETT, AND M. FREAN, *Boosting algorithms as gradient descent in function space*, in Neural Inf. Process. Ser. 12, MIT Press, Cambridge, MA, 1999, pp. 512–518.
- [38] A. MOKHTARI AND A. RIBEIRO, *RES: Regularized stochastic BFGS algorithm*, IEEE Trans. Signal Process., 62 (2014), pp. 6089–6104.
- [39] A. MOKHTARI AND A. RIBEIRO, *Global convergence of online limited memory BFGS*, J. Mach. Learn. Res., 16 (2015), pp. 3151–3181.
- [40] P. MORITZ, R. NISHIHARA, AND M. I. JORDAN, *A linearly-convergent stochastic L-BFGS algorithm*, in Proceedings of AISTATS, 2016, pp. 249–258.
- [41] A. S. NEMIROVSKI, A. JUDITSKY, G. LAN, AND A. SHAPIRO, *Robust stochastic approximation approach to stochastic programming*, SIAM J. Optim., 19 (2009), pp. 1574–1609.
- [42] Y. E. NESTEROV, *A method for unconstrained convex minimization problem with the rate of convergence $\mathcal{O}(1/k^2)$* , Dokl. Akad. Nauk SSSR, 269 (1983), pp. 543–547.
- [43] Y. E. NESTEROV, *Introductory lectures on convex optimization: A basic course*, Appl. Optim., Kluwer Academic Publishers, Boston, MA, 2004.
- [44] B. T. POLYAK, *New stochastic approximation type procedures*, Automat. i Telemekh., 7 (1990), pp. 98–107.
- [45] B. T. POLYAK AND A. B. JUDITSKY, *Acceleration of stochastic approximation by averaging*, SIAM J. Control Optim., 30 (1992), pp. 838–855.
- [46] S. J. REDDI, A. HEFNY, S. SRA, B. PÓCZÓS, AND A. SMOLA, *Stochastic variance reduction for nonconvex optimization*, in Proceedings of ICML, 2016.
- [47] H. ROBBINS AND S. MONRO, *A stochastic approximatin method*, Ann. Math. Statist., 22 (1951), pp. 400–407.
- [48] N. L. ROUX AND A.W. FITZGIBBON, *A fast natural Newton method*, in Proceedings of ICML, 2010, pp. 623–630.
- [49] A. RUSZCZYNSKI AND W. SYSKI, *A method of aggregate stochastic subgradients with on-line stepsize rules for convex stochastic programming problems*, Math. Prog. Stud., 28 (1986), pp. 113–131.
- [50] J. SACKS, *Asymptotic distribution of stochastic approximation*, Ann. Math. Statist., 29 (1958), pp. 373–409.
- [51] N. N. SCHRAUDOLPH, J. YU, AND S. GÜNTHER, *A stochastic quasi-Newton method for online convex optimization*, in Proceedings of AISTATS, 2007, pp. 436–443.
- [52] S. SHALEV-SHWARTZ AND S. BEN-DAVID, *Understanding Machine Learning: From Theory to Algorithms*, Cambridge University Press, Cambridge, UK, 2014.
- [53] O. SHAMIR AND T. ZHANG, *Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes*, in Proceedings of ICML, 2013.
- [54] D. F. SHANNO, *Conditioning of quasi-Newton methods for function minimization*, Math. Comp., 24 (1970), pp. 647–656.
- [55] X. WANG, S. MA, AND Y. YUAN, *Penalty methods with stochastic approximation for stochastic nonlinear programming*, Math. Comp., 86 (2017), pp. 1793–1820.
- [56] L. XIAO AND T. ZHANG, *A proximal stochastic gradient method with progressive variance reduction*, SIAM J. Optim., 24 (2014), pp. 2057–2075.