

INDUSTRIAL AND SYSTEMS ENGINEERING



Almost-sure convergence of iterates and multipliers in
stochastic sequential quadratic optimization

FRANK E. CURTIS, XIN JIANG, AND QI WANG

Department of Industrial and Systems Engineering, Lehigh University

COR@L Technical Report 23T-019



Almost-sure convergence of iterates and multipliers in stochastic sequential quadratic optimization

FRANK E. CURTIS^{*1}, XIN JIANG^{†1}, AND QI WANG^{‡1}

¹Department of Industrial and Systems Engineering, Lehigh University

October 6, 2023

Abstract

Stochastic sequential quadratic optimization (SQP) methods for solving continuous optimization problems with nonlinear equality constraints have attracted attention recently, such as for solving large-scale data-fitting problems subject to nonconvex constraints. However, for a recently proposed subclass of such methods that is built on the popular stochastic-gradient methodology from the unconstrained setting, convergence guarantees have been limited to the asymptotic convergence of the expected value of a stationarity measure to zero. This is in contrast to the unconstrained setting in which almost-sure convergence guarantees (of the gradient of the objective to zero) can be proved for stochastic-gradient-based methods. In this paper, new almost-sure convergence guarantees for the primal iterates, Lagrange multipliers, and stationarity measures generated by a stochastic SQP algorithm in this subclass of methods are proved. It is shown that the error in the Lagrange multipliers can be bounded by the distance of the primal iterate to a primal stationary point plus the error in the latest stochastic gradient estimate. It is further shown that, subject to certain assumptions, this latter error can be made to vanish by employing a running average of the Lagrange multipliers that are computed during the run of the algorithm. The results of numerical experiments are provided to demonstrate the proved theoretical guarantees.

1 Introduction

In this paper, we study convergence guarantees that can be offered for a stochastic algorithm for solving continuous optimization problems with nonlinear equality constraints. Such problems arise in a variety of important areas throughout science and engineering, including optimal control, PDE-constrained optimization, network optimization, and resource allocation [6, 8, 23, 34], and have recently arisen in new and interesting modern application areas such as constrained deep neural network training (e.g., physics-informed learning [15, 22] where one can impose hard constraints rather than merely define the loss function to minimize residual errors [3]). In certain instances of such problems, the objective function can be defined as an expectation over a random variable argument. In this context, the useful features of the algorithm that we study are that it is applicable when only (unbiased) stochastic estimates of the gradient of the objective function are tractable to obtain during the optimization, while at the same time the algorithm can exploit exact constraint function and derivative values that are tractable to obtain.

Sequential quadratic optimization (commonly known as SQP) methods are a popular and powerful class of derivative-based algorithms for solving continuous constrained optimization problems. SQP methods that are stochastic in nature (due to their use of stochastic gradient estimates in place of true gradients of the

*E-mail: frank.e.curtis@lehigh.edu

†E-mail: xjiang@lehigh.edu

‡E-mail: qi420@lehigh.edu

objective) have been investigated recently for solving problems of the aforementioned type, namely, ones for which the objective is defined by an expectation. For one subclass of stochastic SQP methods developed in [3, 2, 16, 17] that is based on the classical stochastic-gradient methodology (known more broadly as stochastic approximation) [35], convergence guarantees have been limited to asymptotic convergence of the expected value of a stationarity measure to zero; see, e.g., [3, Corollary 3.14]. In this paper, we present new analyses of situations in which, for an instance of a method in this subclass, almost-sure convergence guarantees of the primal iterates, Lagrange multipliers, and a stationarity measure can be guaranteed. (For a straightforward analysis, we focus on a simplified variant of the algorithm from [3].) This brings the analysis of this subclass of methods in line with analyses that show almost-sure convergence of stochastic-gradient-based methods for the unconstrained setting; see §1.2. Other stochastic SQP-based methods for solving continuous optimization problems with nonlinear equality constraints have been proposed recently (see, e.g., [1, 4, 5, 19, 27, 26, 30, 33]), and for these methods different types of convergence guarantees have been proved due to the fact that they impose different (often stronger) requirements on the stochastic gradient estimates. Therefore, our work in this paper is distinct from analyses for these other methods, although we contend that the tools employed in this paper might be applicable when analyzing other related algorithms as well. One article that presents results comparable to ours is [28], although in that article the algorithm and assumptions are different than our setting.

Our first contribution is an analysis of situations in which almost-sure convergence of the primal iterates of a stochastic SQP method can be guaranteed. Our main assumption for this analysis can be viewed as a generalization of the Polyak-Łojasiewicz (PL) condition used in unconstrained optimization.

Our second contribution is an analysis of convergence of the sequence of Lagrange multipliers generated by a stochastic SQP method. The convergence behavior of the Lagrange multipliers is of interest for several reasons. For one thing, Lagrange multipliers are used for common certificates of stationarity. They can also play an essential role in sensitivity analysis and active-set identification when solving inequality-constrained problems. Since various algorithms for solving inequality-constrained problems employ algorithms for solving equality-constrained subproblems, it is important to provide convergence guarantees for the Lagrange multipliers when solving equality-constrained problems. In our analysis, we first show that the expected error in the Lagrange multiplier computed in any iteration of our stochastic SQP method of interest can be bounded by the distance of the primal iterate to a primal stationary point plus a term related to the error in the latest stochastic gradient estimate. Such a result is natural, and as in other settings of statistical estimation it suggests that better Lagrange multiplier estimates can be obtained through averaging the multipliers obtained in each iteration. We consider such an approach as well; in particular, we show conditions under which averaging can cause the expected error to vanish asymptotically. (Our approach to averaging is related to the ergodic convergence analysis of optimization algorithms; see, e.g., [32, 12, 21]. It is worthwhile to emphasize that it is different from dual averaging ideas that have been developed for acceleration [29].)

Under our combined assumptions, the stochastic SQP method that we analyze possesses almost-sure convergence guarantees of the primal iterates, Lagrange multipliers, and a stationarity measure while only employing stochastic gradient estimates. To illustrate the practical relevance of our theoretical contributions, we provide the results of numerical experiments. For example, we demonstrate situations when averaging of the Lagrange multipliers results in more accurate stationarity measures, which, as mentioned, is useful in practice for recognizing (approximate) stationarity and other reasons.

1.1 Notation

We use \mathbb{R} to denote the set of real numbers and $\mathbb{R}_{\geq r}$ (resp., $\mathbb{R}_{> r}$) to denote the set of real numbers greater than or equal to (resp., greater than) $r \in \mathbb{R}$. We use \mathbb{R}^n to denote the set of n -dimensional real vectors and $\mathbb{R}^{m \times n}$ to denote the set of m -by- n -dimensional real matrices. We use \mathbb{S}^n to denote the set of symmetric matrices in $\mathbb{R}^{n \times n}$, $\mathbb{S}_{\geq 0}^n$ to denote the set of positive semidefinite matrices in \mathbb{S}^n , and $\mathbb{S}_{> 0}^n$ to denote the set of positive definite matrices in $\mathbb{S}_{\geq 0}^n$. We use \mathbb{N} to denote the positive integers, and for any $k \in \mathbb{N}$ define $[k] := \{1, \dots, k\}$.

Given a real matrix $A \in \mathbb{R}^{n \times m}$, we use $\sigma_{\min}(A)$ to denote the smallest singular value of A and $\sigma_{\max}(A) = \|A\|_2$ to denote the largest singular value of A , i.e., the spectral norm of A . Given such a matrix A with full

column rank, we denote by $A^\dagger = (A^T A)^{-1} A^T$ the Moore–Penrose pseudoinverse of A . For future reference, we state the following lemma pertaining to pseudoinverses.

Lemma 1.1 (see[37, Theorem 4.1]). *If $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{n \times m}$ have full column rank, then $\|A^\dagger - B^\dagger\|_2 \leq \|A^\dagger\|_2 \|B^\dagger\|_2 \|A - B\|_2$.*

For a real-valued sequence $\{x_k\}$ (of numbers, vectors, or matrices), we use $\{x_k\} \subset \mathbb{R}^n$ to indicate that $x_k \in \mathbb{R}^n$ for all $k \in \mathbb{N}$. Similarly, for a sequence of random variables $\{X_k\}$, we use $\{X_k\} \subset \mathbb{R}^n$ to indicate that $X_k \in \mathbb{R}^n$ for all $k \in \mathbb{N}$, which in turn means that, for all $k \in \mathbb{N}$, a realization/outcome of X_k is an element of \mathbb{R}^n . Generally, we use a capital letter to denote a random variable and the corresponding lower-case letter to denote a realization of the random variable. For example, a stochastic objective gradient estimator at iteration $k \in \mathbb{N}$ is denoted as G_k , of which a realization is written as g_k .

Let $\{V_k\} \subset \mathbb{R}^{n \times m}$ be a sequence of random variables and V be a random variable all defined with respect to a probability space $(\Omega, \mathcal{F}, \mathbb{P})$; i.e., an realization $\omega \in \Omega$ defines a realization $V_k(\omega)$ for any $k \in \mathbb{N}$ and a realization $V(\omega)$. The sequence $\{V_k\}$ converges in distribution to V if and only if the cumulative distribution functions (CDFs) of the elements of $\{V_k\}$ converge pointwise to the CDF of V as $k \rightarrow \infty$; e.g., for such $\{V_k\} \subset \mathbb{R}^n$ and a matrix $\Sigma \in \mathbb{S}^n$, we write $\{V_k\} \xrightarrow{d} \mathcal{N}(0, \Sigma)$ to indicate that $\{V_k\}$ converges in distribution to a multivariate normal random vector with mean zero and covariance matrix Σ . The sequence $\{V_k\}$ converges in probability to V as $k \rightarrow \infty$, which we denote by $\{V_k\} \xrightarrow{p} V$, if and only if for any $\epsilon \in \mathbb{R}_{>0}$ one finds

$$\lim_{k \rightarrow \infty} \mathbb{P}[\|V_k - V\|_2 > \epsilon] = 0.$$

Finally, $\{V_k\}$ converges almost-surely to V as $k \rightarrow \infty$, which we denote by $\{V_k\} \xrightarrow{a.s.} V$, if and only if there exists $\Omega_0 \subset \Omega$ with $\mathbb{P}(\Omega_0) = 0$ such that

$$\lim_{k \rightarrow \infty} V_k(\omega) = V(\omega) \quad \text{for all } \omega \in \Omega \setminus \Omega_0.$$

We use $\mathbf{1}_A$ to denote an indicator random variable for the event A , which takes the value 1 if event A occurs and takes the value 0 otherwise.

1.2 Mathematical Background

In the classical article by Robbins and Monro [35], it is shown that a straightforward approach of stochastic approximation for solving an equation with a unique root x_* can lead to convergence in probability of the iterate sequence. Specifically, under certain basic assumptions, as long as the prescribed sequence of *step sizes* $\{\alpha_k\}$ employed by the algorithm is unsummable, but square summable (e.g., $\alpha_k = 1/k$ for all $k \in \mathbb{N}$), it can be shown that the generated sequence of solution estimates $\{X_k\}$ satisfies

$$\lim_{k \rightarrow \infty} \mathbb{E}[(X_k - x_*)^2] = 0, \tag{1}$$

which in turn implies that $\{X_k\} \xrightarrow{p} x_*$. Cast into the context of minimizing a smooth, potentially nonconvex objective $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with a stochastic-gradient method, these same principles can be used to prove (see, e.g., [9])

$$\lim_{k \rightarrow \infty} \mathbb{E}[\|\nabla f(X_k)\|_2^2] = 0.$$

In a later article by Robbins and Siegmund [36], certain conditions are shown to offer an almost-sure convergence guarantee for a sequence generated by stochastic approximation. In particular, the article first proves a general theorem, which we state in a slightly simplified form for our purposes.

Lemma 1.2 (see[36, Theorem 1]). *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $\{\mathcal{F}_k\}$ with $\mathcal{F}_k \subseteq \mathcal{F}_{k+1}$ for all $k \in \mathbb{N}$ be a sequence of sub- σ -algebras of \mathcal{F} . Let $\{R_k\}$, $\{P_k\}$, and $\{Q_k\}$ be sequences of nonnegative random variables such that, for all $k \in \mathbb{N}$, the random variables R_k , P_k , and Q_k are \mathcal{F}_k -measurable. If $\sum_{k=1}^{\infty} Q_k < \infty$ and, for all $k \in \mathbb{N}$, one has*

$$\mathbb{E}[R_{k+1} | \mathcal{F}_k] \leq R_k - P_k + Q_k,$$

then, almost-surely, $\sum_{k=1}^{\infty} P_k < \infty$ and $\lim_{k \rightarrow \infty} R_k$ exists and is finite.

Applied to the context of stochastic approximation for solving an equation, it is shown in [36] that under certain assumptions (which we omit for brevity) the conclusion in (1) can be strengthened for the same algorithm to

$$\mathbb{P} \left[\lim_{k \rightarrow \infty} X_k = x_\star \right] = 1,$$

which is to say that $\{X_k\} \xrightarrow{a.s.} x_\star$. One can derive a similar such result in the context of minimizing a smooth, potentially nonconvex objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with a stochastic-gradient method. For example, under a related set of conditions, it has been shown by Bertsekas and Tsitsiklis [7] that a stochastic-gradient method can be guaranteed to yield

$$\mathbb{P} \left[\lim_{k \rightarrow \infty} \nabla f(X_k) = 0 \right] = 1, \quad \text{i.e., } \{\nabla f(X_k)\} \xrightarrow{a.s.} 0. \quad (2)$$

Our main contributions in this paper are also almost-sure convergence guarantees, but for a stochastic SQP method in the context of nonlinear-equality-constrained optimization. In particular, we show almost-sure convergence guarantees for the primal iterates, Lagrange multipliers, and stationarity measures for a simplified variant of the algorithm from [3]. Some of our almost-sure convergence guarantees for the Lagrange multipliers computed by our method of interest pertains to an averaged sequence. For our analysis of this sequence, we require two key results that are known from the literature. The first result that we need is the central limit theorem (CLT) for a multidimensional martingale difference triangular array stated as Lemma 1.3.

Lemma 1.3 (Multidimensional martingale central limit theorem). *Let $\{(\xi_{k,i}, \mathcal{F}_{k,i})\}_{k \in \mathbb{N}, i \in [k]}$ be an n -dimensional martingale difference triangular array, i.e., with an initial generating σ -algebra $\mathcal{F}_{k,1}$ for all $k \in \mathbb{N}$, one has*

(i) $\mathcal{F}_{k,i} = \sigma(\xi_{k,1}, \dots, \xi_{k,i-1})$ for all $k \in \mathbb{N}$ and $i \in \{2, \dots, k\}$ and

(ii) $\mathbb{E}[\xi_{k,i} | \mathcal{F}_{k,i}] = 0$ for all $k \in \mathbb{N}$ and $i \in [k]$.

If the array has the properties that

$$\xi_{k,i} \text{ is square-integrable, i.e., } \mathbb{E}[\|\xi_{k,i}\|_2^2] < \infty \text{ for all } (k,i) \in \mathbb{N} \times [k], \quad (3a)$$

$$\left\{ \sum_{i=1}^k \mathbb{E}[\|\xi_{k,i}\|_2^2 \mathbf{1}_{\{\|\xi_{k,i}\|_2 > \delta\}} | \mathcal{F}_{k,i}] \right\} \xrightarrow{p} 0 \text{ for all } \delta \in \mathbb{R}_{>0}, \text{ and} \quad (3b)$$

$$\left\{ \sum_{i=1}^k \mathbb{E}[\xi_{k,i} \xi_{k,i}^T | \mathcal{F}_{k,i}] \right\} \xrightarrow{p} \Sigma \text{ for some } \Sigma \in \mathbb{S}^n, \quad (3c)$$

$$\text{then } \left\{ \sum_{i=1}^k \xi_{k,i} \right\} \xrightarrow{d} \mathcal{N}(0, \Sigma).$$

The multidimensional martingale CLT in Lemma 1.3 can be derived by applying the one-dimensional martingale CLT (see, e.g., [25, Theorem 2.3] and [20, Corollary 3.1]) to $\xi_{k,i}^a := a^T \xi_{k,i}$ for arbitrary $a \in \mathbb{R}^n$ since $\{\sum_{i=1}^k \xi_{k,i}^a\} \xrightarrow{d} \mathcal{N}(0, a^T \Sigma a)$ for any given $a \in \mathbb{R}^n$ implies that $\{\sum_{i=1}^k \xi_{k,i}\} \xrightarrow{d} \mathcal{N}(0, \Sigma)$; see, e.g., [18, Exercise 3.10.8]. A similar result is used in [24]. We refer to (3b) as Lindeberg's condition, as is common in the literature.

The second key result that we need is the following, which we refer to as a moment convergence result. It follows, e.g., by [14, Theorem 4.5.2].

Lemma 1.4 (Moment convergence). *Let $\{X_k\} \subset \mathbb{R}^n$ be a sequence of random vectors such that $\{X_k\} \xrightarrow{d} X$ and $\sup_{k \in \mathbb{N}} \mathbb{E}[\|X_k\|_2^\Theta] < \infty$ for some $\Theta \in \mathbb{R}_{>0}$. Then, for all $\theta \in (0, \Theta)$, one finds that $\lim_{k \rightarrow \infty} \mathbb{E}[\|X_k\|_2^\theta] = \mathbb{E}[\|X\|_2^\theta]$.*

1.3 Outline

In Section 2, we present formally the continuous optimization problem of interest and the stochastic SQP algorithm for solving it that we analyze in the remainder of the paper. We also provide preliminary assumptions that we make about the problem and the algorithm, and state basic properties of the algorithm that transfer from the prior work in [3]. In Section 3, we prove convergence results for the primal iterates generated by the algorithm. In Section 4, we prove convergence results for Lagrange multiplier sequences that are generated by the algorithm. The results of numerical experiments are provided in Section 5 and concluding remarks are offered in Section 6.

2 Problem and Algorithm Descriptions

The algorithm that we study is designed to solve

$$\min_{x \in \mathbb{R}^n} f(x) \text{ subject to (s.t.) } c(x) = 0 \text{ with } f(x) = \mathbb{E}_\iota[F(x, \iota)], \quad (4)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are continuously differentiable, ι is a random variable with associated probability space $(\Omega_\iota, \mathcal{F}_\iota, \mathbb{P}_\iota)$, $F : \mathbb{R}^n \times \Omega_\iota \rightarrow \mathbb{R}$, and $\mathbb{E}_\iota[\cdot]$ denotes expectation taken with respect to \mathbb{P}_ι . Given an initial point $x_1 \in \mathbb{R}^n$, any run of the algorithm generates a sequence of iterates $\{x_k\} \subset \mathbb{R}^n$, i.e., a realization of a stochastic process $\{X_k\} \subset \mathbb{R}^n$. We make the following assumption about problem (4) and the generated sequence of iterates.

Assumption 2.1. *There exists open convex $\mathcal{X} \subseteq \mathbb{R}^n$ containing $\{X_k\} \subset \mathbb{R}^n$ such that the following hold. The objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable and bounded below over \mathcal{X} and its gradient function $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is bounded and Lipschitz continuous over \mathcal{X} . The constraint function $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$ (where $m \leq n$) is Lipschitz continuous, continuously differentiable, and bounded over \mathcal{X} , its Jacobian function $\nabla c^T : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$ is Lipschitz continuous over \mathcal{X} , and $\nabla c(x)^T$ has full row rank for all $x \in \mathcal{X}$ with singular values that are bounded below uniformly by a positive constant over \mathcal{X} .*

It would be possible to loosen Assumption 2.1 to require only that \mathcal{X} contains $\{X_k\}$ almost surely. However, since this would require repeated references to probability-one events throughout our analysis without strengthening our conclusions substantially, we do not bother with this level of generality.

Under Assumption 2.1, there exists a tuple of constants, which we denote by $(\kappa_{\mathcal{X}}, f_{\text{inf}}, \kappa_{\nabla f}, \kappa_c, \kappa_{\nabla c}, r) \in \mathbb{R}_{>0} \times \mathbb{R} \times \mathbb{R}_{>0} \times \mathbb{R}_{>0} \times \mathbb{R}_{>0} \times \mathbb{R}_{>0}$, such that

$$\begin{aligned} \|x\|_2 &\leq \kappa_{\mathcal{X}}, \quad f(x) \geq f_{\text{inf}}, \quad \|\nabla f(x)\|_2 \leq \kappa_{\nabla f}, \\ \|c(x)\|_2 &\leq \kappa_c, \quad \|\nabla c(x)\|_2 \leq \kappa_{\nabla c}, \quad \text{and } \sigma_{\min}(\nabla c(x)) \geq r \end{aligned} \quad (5)$$

for all $x \in \mathcal{X}$, and there exists $(L_{\nabla f}, L_c, \Gamma) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ such that

$$\begin{aligned} \|\nabla f(x) - \nabla f(\bar{x})\|_2 &\leq L_{\nabla f} \|x - \bar{x}\|_2, \quad \|c(x) - c(\bar{x})\|_2 \leq L_c \|x - \bar{x}\|_2, \\ \text{and } \|\nabla c(x) - \nabla c(\bar{x})\|_2 &\leq \Gamma \|x - \bar{x}\|_2 \end{aligned} \quad (6)$$

for all $(x, \bar{x}) \in \mathcal{X} \times \mathcal{X}$. The Lagrangian function $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ corresponding to problem (4) is defined by $\mathcal{L}(x, y) = f(x) + c(x)^T y$, and the first-order stationarity conditions for (4) are given by

$$0 = \begin{bmatrix} \nabla_x \mathcal{L}(x, y) \\ \nabla_y \mathcal{L}(x, y) \end{bmatrix} = \begin{bmatrix} \nabla f(x) + \nabla c(x)y \\ c(x) \end{bmatrix}. \quad (7)$$

We refer to any $(x, y) \in \mathbb{R}^n \times \mathbb{R}^m$ satisfying (7) as a stationary point for problem (4). In addition, we refer to any $x \in \mathbb{R}^n$ such that there exists $y \in \mathbb{R}^m$ with (x, y) satisfying (7) as a primal stationary point for (4).

Let us now describe the algorithm whose convergence properties are the subject of our study. We state the algorithm in terms of a realization of the quantities it generates. At an iterate $x_k \in \mathbb{R}^n$, the algorithm

generates a stochastic gradient estimate $g_k \approx \nabla f(x_k) \in \mathbb{R}^n$, and makes use of $h_k \in \mathbb{S}^n$ (see upcoming Assumptions 2.3 and 2.4 about these quantities). Given g_k and h_k , a direction $d_k \in \mathbb{R}^n$ is computed by solving

$$\min_{d \in \mathbb{R}^n} g_k^T d + \frac{1}{2} d^T h_k d \text{ s.t. } c(x_k) + \nabla c(x_k)^T d = 0. \quad (8)$$

Under upcoming Assumption 2.4 (which includes that $d^T h_k d > 0$ for all $d \in \text{Null}(\nabla c(x_k)^T)$), the solution d_k of (8) as well as a Lagrange multiplier for the constraints, call it $y_k \in \mathbb{R}^m$, can be obtained by solving

$$\begin{bmatrix} h_k & \nabla c(x_k) \\ \nabla c(x_k)^T & 0 \end{bmatrix} \begin{bmatrix} d_k \\ y_k \end{bmatrix} = - \begin{bmatrix} g_k \\ c_k \end{bmatrix}. \quad (9)$$

Upon computation of the search direction in iteration $k \in \mathbb{N}$, the algorithm selects a step size $\alpha_k \in (0, 1]$. Specifically, with a *merit parameter* $\tau \in \mathbb{R}_{>0}$, *ratio parameter* $\xi \in \mathbb{R}_{>0}$ (see upcoming Assumption 2.2), and $(L_{\nabla f}, \Gamma)$ from (6), the algorithm that we analyze selects the step size for all $k \in \mathbb{N}$ as

$$\alpha_k \leftarrow \frac{\beta_k \tau \xi}{\tau L_{\nabla f} + \Gamma} \text{ for some } \beta_k \in (0, 1], \quad (10)$$

where $\{\beta_k\}$ is unsummable, but square-summable. This choice of step sizes means that the algorithm that we analyze is a simplified variant of the algorithm from [3]; see [3] and below for further discussion about the merit parameter, ratio parameter, and the particular formula for the step size stated in (10). (We conjecture that our ultimate conclusions also hold for the original variant of the algorithm from [3] in the event considered in Section 3.2.1 of that paper, but for our aims in this paper we merely consider the simplified variant presented here in order to allow our analysis not to be obscured by auxiliary details.) One aspect that makes our variant a simplified one is that we assume that the merit and ratio parameters are initialized to values that are sufficiently small such that the algorithm does not need to update them adaptively. This means that, for our analysis, we can consider the merit function (with *fixed* merit parameter τ) $\phi_\tau : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by $\phi_\tau(x) = \tau f(x) + \|c(x)\|_1$. (Any convex norm can be used for the constraint violation. We use the ℓ_1 -norm for consistency with [3].) A local approximation of ϕ_τ at $x \in \mathbb{R}^n$ can be defined through a function $q_\tau : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{S}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$; specifically, with $g \in \mathbb{R}^n$ and $h \in \mathbb{S}^n$, a local approximation of ϕ_τ at x as a function of d is given by

$$q_\tau(x, g, h, d) = \tau(f(x) + g^T d + \frac{1}{2} \max\{d^T h d, 0\}) + \|c(x) + \nabla c(x)^T d\|_1. \quad (11)$$

The reduction in $q_\tau(x, g, h, \cdot)$ offered by d with $c(x) + \nabla c(x)^T d = 0$ can be defined through $\Delta q_\tau : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^{n \times n} \times \mathbb{R}^n \rightarrow \mathbb{R}$ defined by

$$\begin{aligned} \Delta q_\tau(x, g, h, d) &= q_\tau(x, g, h, 0) - q_\tau(x, g, h, d) \\ &= -\tau(g^T d + \frac{1}{2} \max\{d^T h d, 0\}) + \|c(x)\|_1. \end{aligned} \quad (12)$$

Consistent with [3, Section 3.2.1], we make the following assumption.

Assumption 2.2. *The parameters $\xi \in \mathbb{R}_{>0}$ and $\tau \in \mathbb{R}_{>0}$ are chosen such that, for some $\nu \in (0, 1)$ and all $k \in \mathbb{N}$ in any run, the following hold.*

(i) $\xi \leq \xi_k^{\text{trial}}$, where one defines

$$\xi_k^{\text{trial}} := \begin{cases} \infty & \text{if } d_k = 0 \\ \frac{\Delta q_\tau(x_k, g_k, h_k, d_k)}{\tau \|d_k\|_2^2} & \text{otherwise;} \end{cases}$$

(ii) $\tau \leq \tau_k^{\text{trial, true}}$, where, with d_k^{true} being the solution of (8) if g_k were replaced by $\nabla f(x_k)$ and $\rho_k := \nabla f(x_k)^T d_k^{\text{true}} + \max\{(d_k^{\text{true}})^T h_k d_k^{\text{true}}, 0\}$, one defines

$$\tau_k^{\text{trial, true}} := \begin{cases} \infty & \text{if } \rho_k \leq 0 \\ \frac{(1-\nu)\|c(x_k)\|_1}{\rho_k} & \text{otherwise.} \end{cases}$$

The definition of $\tau_k^{\text{trial, true}}$ and Assumption 2.2 ensure for all $k \in \mathbb{N}$ that

$$\Delta q_\tau(x_k, \nabla f(x_k), h_k, d_k^{\text{true}}) \geq \frac{1}{2}\tau \max\{(d_k^{\text{true}})^T h_k d_k^{\text{true}}, 0\} + \nu \|c_k\|_1. \quad (13)$$

Overall, the step-size choice in (10) along with Assumption 2.2 results in a simplified variant of the adaptive strategy presented in [3], and one finds from [3] that this setting ensures that the step from x_k to x_{k+1} yields a sufficient-decrease-type property that is relevant for our analysis; see Lemma 2.1 below.

The stochastic SQP method that we study is stated in Algorithm 1.

Algorithm 1 Stochastic SQP

Require: $x_1 \in \mathbb{R}^n$; $(L_{\nabla f}, \Gamma) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ (see (6)); $(\tau, \xi) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ satisfying Assumption 2.2; $\{\beta_k\} \subset (0, 1]$ that is unsummable, but square-summable.

- 1: **for** $k \in \mathbb{N}$ **do**
 - 2: Compute g_k (see Assumption 2.3).
 - 3: Compute h_k (see Assumption 2.4).
 - 4: Compute (d_k, y_k) by solving (8).
 - 5: Set α_k by (10).
 - 6: Set $x_{k+1} \leftarrow x_k + \alpha_k d_k$.
 - 7: **end for**
-

All that remains before presenting our analysis are to articulate some final assumptions and state a few key relations from [3]. Considering the outcomes of entire runs of Algorithm 1, henceforth we consider the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $\Omega := \prod_{k=1}^{\infty} \Omega_k$. In this manner, each realization of a run of the algorithm can be associated with $\omega \in \Omega$, an infinite-dimensional tuple whose k th element is $\omega_k \in \Omega_k$ which determines the stochastic gradient estimate. The stochastic process defined by Algorithm 1 can thus be expressed as

$$\{(X_k(\omega), G_k(\omega), H_k(\omega), D_k(\omega), D_k^{\text{true}}(\omega), Y_k(\omega), Y_k^{\text{true}}(\omega))\},$$

where, for all $k \in \mathbb{N}$, the random variables are the iterate $X_k(\omega)$, stochastic gradient estimator $G_k(\omega)$, symmetric matrix $H_k(\omega)$, search direction $D_k(\omega)$, true search direction $D_k^{\text{true}}(\omega)$, Lagrange multiplier $Y_k(\omega)$, and true Lagrange multiplier $Y_k^{\text{true}}(\omega)$. (As in Assumption 2.2, the true search direction and Lagrange multiplier are defined by the solution of (9) if the stochastic gradient were replaced by the true gradient.) For instance, given ω that specifies a particular outcome of a run of the algorithm, the quantity $X_k(\omega) \in \mathbb{R}^n$ is the k th iterate. Given initial conditions (including that $X_1(\omega) = x_1 \in \mathbb{R}^n$ for all $\omega \in \Omega$), let \mathcal{F}_1 denote the σ -algebra corresponding to the initial conditions and, for all $k \in \mathbb{N} \setminus \{1\}$, let \mathcal{F}_k denote the σ -algebra defined by the initial conditions and the random variables $\{G_1, \dots, G_{k-1}\}$. We assume the following.

Assumption 2.3. *For all $k \in \mathbb{N}$, the stochastic gradient estimator satisfies $\mathbb{E}[G_k | \mathcal{F}_k] = \nabla f(X_k)$. In addition, there exists $\sigma \in \mathbb{R}_{>0}$ such that, for all $k \in \mathbb{N}$, it holds that $\mathbb{E}[\|G_k - \nabla f(X_k)\|_2^2 | \mathcal{F}_k] \leq \sigma^2$.*

Assumption 2.4. *For all $k \in \mathbb{N}$, the matrix $H_k \in \mathbb{S}^n$ is \mathcal{F}_k -measurable and bounded in ℓ_2 -norm by $\kappa_H \in \mathbb{R}_{>0}$. In addition, there exists $\zeta \in (0, \kappa_H]$ such that, for all $k \in \mathbb{N}$, $u^T H_k u \geq \zeta \|u\|_2^2$ for all $u \in \text{Null}(\nabla c(X_k)^T)$.*

Under Assumptions 2.3 and 2.4, it follows that $\{\mathcal{F}_k\}$ is a filtration for the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. In particular, the initial conditions and a realization of $\{G_1, \dots, G_{k-1}\}$ determine the realizations of $\{(X_j, H_j, D_j^{\text{true}}, Y_j^{\text{true}})\}_{j=1}^k$ and $\{(D_j, Y_j)\}_{j=1}^{k-1}$, i.e., for all $k \in \mathbb{N}$, one has that $(X_k, H_k, D_k^{\text{true}}, Y_k^{\text{true}})$ is \mathcal{F}_k -measurable while (G_k, D_k, Y_k) is \mathcal{F}_{k+1} -measurable. Examples of symmetric matrices satisfying Assumption 2.4 for each $k \in \mathbb{N}$ are $H_k = I$ or H_k being (an approximation of) the Hessian of the Lagrangian at (X_k, Y_{k-1}) as long as safeguards are included to ensure that H_k is sufficiently positive definite in the null space of $\nabla c(X_k)^T$, as required by Assumption 2.4.

The following lemma characterizes decreases of the merit function ϕ_τ and the boundedness of Δq_τ ; it is adapted from Lemmas 3.7, 3.8, and 3.12 in [3].

Lemma 2.1. For all $k \in \mathbb{N}$, it follows that $\mathbb{E}[D_k | \mathcal{F}_k] = D_k^{\text{true}}$,

$$\begin{aligned} & \phi_\tau(X_k + \alpha_k D_k) - \phi_\tau(X_k) \\ & \leq -\alpha_k \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) \\ & \quad + \frac{1}{2} \alpha_k \beta_k \Delta q_\tau(X_k, G_k, H_k, D_k) + \alpha_k \tau \nabla f(X_k)^T (D_k - D_k^{\text{true}}), \end{aligned} \quad (14)$$

and

$$\mathbb{E}[\Delta q_\tau(X_k, G_k, H_k, D_k) | \mathcal{F}_k] \leq \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) + \tau \zeta^{-1} \sigma^2. \quad (15)$$

3 Convergence of primal iterates

One can derive from the analysis in [3] that, under Assumptions 2.1, 2.2, 2.3, and 2.4, Algorithm 1 yields (recall the stationarity conditions (7))

$$\liminf_{k \rightarrow \infty} \mathbb{E}[\|\nabla f(X_k) + \nabla c(X_k) Y_k^{\text{true}}\|_2^2 + \|c(X_k)\|_2] = 0; \quad (16)$$

see [3, Corollary 3.14] for further details. This lower limit in (16) does not imply that the primal iterate sequence converges in any particular sense to a primal stationary point. In this section, we establish conditions under which almost-sure convergence of the primal iterate sequence can be guaranteed, which when coupled with the results of the next section (in particular, Corollary 4.3) provides a rich picture of the convergence behavior of Algorithm 1.

Before presenting the results of this section, let us state a common decomposition of the search direction that is used in our analysis. Observe by the Fundamental Theorem of Linear Algebra that D_k can be decomposed as $D_k = U_k + V_k$, where $U_k \in \text{Null}(\nabla c(X_k)^T)$ and $V_k \in \text{Range}(\nabla c(X_k))$. Under Assumption 2.1, we can now introduce $Z_k \in \mathbb{R}^{n \times (n-m)}$ as a matrix whose columns form a basis for $\text{Null}(\nabla c(X_k)^T)$. It follows that $U_k = Z_k W_k$ for some $W_k \in \mathbb{R}^{n-m}$. Now, from (9), it follows that

$$-c(X_k) = \nabla c(X_k)^T D_k = \nabla c(X_k)^T (Z_k W_k + V_k) = \nabla c(X_k)^T V_k.$$

Since $\nabla c(X_k)^T$ has full row rank under Assumption 2.1, the unique $V_k \in \text{Range}(\nabla c(X_k))$ that solves this linear system is given by

$$V_k = -\nabla c(X_k) (\nabla c(X_k)^T \nabla c(X_k))^{-1} c(X_k) = -(\nabla c(X_k)^\dagger)^T c(X_k). \quad (17)$$

Hence, from (9) and (17), it follows under Assumption 2.4 that

$$\begin{aligned} & -G_k = H_k(Z_k W_k + V_k) + \nabla c(X_k) Y_k \\ \implies & \quad W_k = -(Z_k^T H_k Z_k)^{-1} Z_k^T (G_k + H_k V_k) \\ \implies & \quad Z_k W_k = -Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T (G_k + H_k V_k) \\ & \quad = -Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T (G_k - H_k (\nabla c(X_k)^\dagger)^T c(X_k)). \end{aligned} \quad (18)$$

Our main result of this section is Theorem 3.1 on page 11. Leading up to that result, we require two technical lemmas that are presented next. The first, Lemma 3.1 below, establishes a result about the asymptotic behavior of the sequence of merit function values, namely, $\{\phi_\tau(X_k)\}$, and about the sequence of reductions in the model of the merit function corresponding to the true objective gradients and corresponding search directions. Note that the proof of the lemma uses the fact that the sequence $\{\beta_k\}$ employed in Algorithm 1 is unsummable (i.e., $\sum_{k=1}^{\infty} \beta_k = \infty$), but square-summable (i.e., $\sum_{k=1}^{\infty} \beta_k^2 < \infty$).

Lemma 3.1. Suppose that Assumptions 2.1, 2.2, 2.3, and 2.4 hold. Then,

$$\begin{aligned} & \lim_{k \rightarrow \infty} \{\phi_\tau(X_k)\} \text{ exists and is finite almost surely, and} \\ & \liminf_{k \rightarrow \infty} \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) = 0 \text{ almost surely.} \end{aligned}$$

Proof. For all $k \in \mathbb{N}$, define $R_k := \phi_\tau(X_k) - \tau f_{\text{inf}}$, and observe from Assumption 2.1 that $R_k \geq 0$ for all $k \in \mathbb{N}$. Also, for arbitrary $k \in \mathbb{N}$, one finds

$$\begin{aligned}
& \mathbb{E}[R_{k+1} | \mathcal{F}_k] \\
&= \mathbb{E}[\phi_\tau(X_k + \alpha_k D_k) | \mathcal{F}_k] - \tau f_{\text{inf}} \\
&\leq \phi_\tau(X_k) - \tau f_{\text{inf}} - \alpha_k \mathbb{E}[\Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) | \mathcal{F}_k] \\
&\quad + \frac{1}{2} \alpha_k \beta_k \mathbb{E}[\Delta q_\tau(X_k, G_k, H_k, D_k) | \mathcal{F}_k] \\
&= R_k - \alpha_k \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) + \frac{1}{2} \alpha_k \beta_k \mathbb{E}[\Delta q_\tau(X_k, G_k, H_k, D_k) | \mathcal{F}_k] \\
&\leq R_k - \alpha_k \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) \\
&\quad + \frac{1}{2} \alpha_k \beta_k (\Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) + \tau \zeta^{-1} \sigma^2) \\
&= R_k - \alpha_k (1 - \frac{1}{2} \beta_k) \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) + \frac{1}{2} \alpha_k \beta_k \tau \zeta^{-1} \sigma^2 \\
&\leq R_k - \alpha_k \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) + \frac{1}{2} \alpha_k \beta_k \tau \zeta^{-1} \sigma^2 \\
&= R_k - \frac{\beta_k \tau \xi}{\tau L_{\nabla f} + \Gamma} \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) + \frac{\beta_k^2 \tau^2 \xi \sigma^2}{2\zeta(\tau L_{\nabla f} + \Gamma)}, \tag{19}
\end{aligned}$$

where the first inequality follows from Lemma 2.1, the second follows from (15), and the third follows from (13) and the fact that $\{\beta_k\} \subset (0, 1]$. Now define

$$P_k := \frac{\beta_k \tau \xi}{\tau L_{\nabla f} + \Gamma} \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) \quad \text{and} \quad Q_k := \frac{\beta_k^2 \tau^2 \xi \sigma^2}{2\zeta(\tau L_{\nabla f} + \Gamma)}$$

and observe that $P_k \geq 0$ for all $k \in \mathbb{N}$ follows by (13), $Q_k \geq 0$ for all $k \in \mathbb{N}$ follows since $\{\beta_k\} \subset (0, 1]$, and $\sum_{k=1}^{\infty} Q_k < \infty$ since $\sum_{k=1}^{\infty} \beta_k^2 < \infty$. Therefore, with (19), applying Lemma 1.2 shows that, almost surely,

$$\sum_{k=1}^{\infty} P_k = \sum_{k=1}^{\infty} \alpha_k \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) < \infty$$

and $\lim_{k \rightarrow \infty} R_k = \lim_{k \rightarrow \infty} (\phi_\tau(X_k) - f_{\text{inf}})$ exists and is finite. The latter fact leads directly to the first desired conclusion, whereas the former fact along with $\sum_{k=1}^{\infty} \beta_k = \infty$ (so $\sum_{k=1}^{\infty} \alpha_k = \infty$) yields the second. \square

The components of our second technical lemma can be derived from various results from [3], namely, Lemmas 2.10, 2.11, 2.12, and 3.4 in [3]. Nonetheless, we provide a proof of the lemma for the sake of completeness.

Lemma 3.2. *Suppose that Assumptions 2.1, 2.2, 2.3, and 2.4 hold. Then, with respect to the decomposition $D_k^{\text{true}} = U_k^{\text{true}} + V_k$ for all $k \in \mathbb{N}$, where $U_k^{\text{true}} \in \text{Null}(\nabla c(X_k)^T)$ and $V_k \in \text{Range}(\nabla c(X_k))$, and with*

$$\Psi_k := \begin{cases} \|U_k^{\text{true}}\|_2^2 + \|c(X_k)\|_2 & \text{if } \|U_k^{\text{true}}\|_2^2 \geq \kappa_{uv} \|V_k\|_2^2 \\ \|c(X_k)\|_2 & \text{otherwise.} \end{cases}$$

the following statements hold.

(a) There exists $\kappa_{uv} \in \mathbb{R}_{>0}$ such that, for all $k \in \mathbb{N}$,

$$\|U_k^{\text{true}}\|_2^2 \geq \kappa_{uv} \|V_k\|_2^2 \implies (D_k^{\text{true}})^T H_k D_k^{\text{true}} \geq \frac{1}{2} \zeta \|U_k^{\text{true}}\|_2^2.$$

(b) There exists $\kappa_\Psi \in \mathbb{R}_{>0}$ such that, for all $k \in \mathbb{N}$,

$$\|D_k^{\text{true}}\|_2^2 + \|c(X_k)\|_2 \leq (\kappa_\Psi + 1) \Psi_k.$$

(c) There exists $\kappa_q \in \mathbb{R}_{>0}$ such that, for all $k \in \mathbb{N}$,

$$\Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) \geq \kappa_q \Psi_k.$$

Proof. We prove each part in turn.

(a) For arbitrary $k \in \mathbb{N}$ and $\kappa \in \mathbb{R}_{>0}$, one finds under Assumption 2.4 that $\|U_k^{\text{true}}\|_2^2 \geq \kappa \|V_k\|_2^2$ implies

$$\begin{aligned} (D_k^{\text{true}})^T H_k D_k^{\text{true}} &= (U_k^{\text{true}})^T H_k U_k^{\text{true}} + 2(U_k^{\text{true}})^T H_k V_k + V_k^T H_k V_k \\ &\geq \zeta \|U_k^{\text{true}}\|_2^2 - 2\|U_k^{\text{true}}\|_2 \|H_k\|_2 \|V_k\|_2 - \|H_k\|_2 \|V_k\|_2^2 \\ &\geq (\zeta - \frac{2\kappa_H}{\sqrt{\kappa}} - \frac{\kappa_H}{\kappa}) \|U_k^{\text{true}}\|_2^2. \end{aligned}$$

Thus, the desired result holds for any $\kappa_{uv} \in \mathbb{R}_{>0}$ such that $\frac{2\kappa_H}{\sqrt{\kappa_{uv}}} + \frac{\kappa_H}{\kappa_{uv}} \leq \frac{\zeta}{2}$.

(b) If $\|U_k^{\text{true}}\|_2^2 \geq \kappa_{uv} \|V_k\|_2^2$, it follows that

$$\begin{aligned} \|D_k^{\text{true}}\|_2^2 &= \|U_k^{\text{true}}\|_2^2 + \|V_k\|_2^2 \leq (1 + \kappa_{uv}^{-1}) \|U_k^{\text{true}}\|_2^2 \\ &\leq (1 + \kappa_{uv}^{-1}) (\|U_k^{\text{true}}\|_2^2 + \|c(X_k)\|_2). \end{aligned}$$

Otherwise, with (17) and Assumption 2.1, it follows that

$$\begin{aligned} \|D_k^{\text{true}}\|_2^2 &= \|U_k^{\text{true}}\|_2^2 + \|V_k\|_2^2 < (\kappa_{uv} + 1) \|V_k\|_2^2 \\ &= (\kappa_{uv} + 1) \|(\nabla c(X_k)^\dagger)^T c(X_k)\|_2^2 \\ &\leq \kappa_c (\kappa_{uv} + 1) r^{-2} \|c(X_k)\|_2. \end{aligned}$$

Hence, $\|D_k^{\text{true}}\|_2^2 \leq \kappa_\Psi \Psi_k$ holds with $\kappa_\Psi := \max\{1 + \kappa_{uv}^{-1}, \kappa_c (\kappa_{uv} + 1) r^{-2}\}$. The desired conclusion then follows since, by definition, $\Psi_k \geq \|c(X_k)\|_2$.

(c) By (13) and part (a), it follows that $\|U_k^{\text{true}}\|_2^2 \geq \kappa_{uv} \|V_k\|_2^2$ implies

$$\Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) \geq \frac{1}{4} \tau \zeta \|U_k^{\text{true}}\|_2^2 + \nu \|c(X_k)\|_1,$$

and otherwise one still finds $\Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) \geq \nu \|c(X_k)\|_1$. Hence, since $\|\cdot\|_2 \leq \|\cdot\|_1$, the result follows with $\kappa_q := \min\{\frac{1}{4} \tau \zeta, \nu\}$. \square

We are now prepared to prove our main result of this section, which establishes conditions under which $\{\phi_\tau(X_k)\}$ converges almost-surely to a minimizer of the merit function at which the constraints of problem (4) are satisfied and that $\{X_k\}$ converges almost-surely to a primal stationary point (in fact, a local minimizer). After proving the result, we provide some additional commentary on the inequality (20) that is required for the theorem.

Theorem 3.1. *Suppose that Assumptions 2.1, 2.2, 2.3, and 2.4 hold. In addition, suppose that there exists $x_\star \in \mathcal{X}$ with $c(x_\star) = 0$, $\mu \in \mathbb{R}_{>1}$, and $\epsilon \in \mathbb{R}_{>0}$ such that for all $x \in \mathcal{X}_{\epsilon, x_\star} := \{x \in \mathcal{X} : \|x - x_\star\|_2 \leq \epsilon\}$ one finds*

$$\phi_\tau(x) - \phi_\tau(x_\star) \begin{cases} = 0 & \text{if } x = x_\star \\ \in (0, \mu(\tau \|Z(x)^T \nabla f(x)\|_2^2 + \|c(x)\|_2)] & \text{otherwise,} \end{cases} \quad (20)$$

where for all $x \in \mathcal{X}_{\epsilon, x_\star}$ one defines $Z(x) \in \mathbb{R}^{n \times (n-m)}$ as some orthonormal matrix whose columns form a basis for the null space of $\nabla c(x)^T$. Then, if $\limsup_{k \rightarrow \infty} \{\|X_k - x_\star\|_2\} \leq \epsilon$ almost surely, it follows that

$$\begin{aligned} &\{\phi_\tau(X_k)\} \xrightarrow{a.s.} \phi_\tau(x_\star), \quad \{X_k\} \xrightarrow{a.s.} x_\star, \\ \text{and} \quad &\left\{ \begin{bmatrix} \nabla f(X_k) + \nabla c(X_k) Y_k^{\text{true}} \\ c(X_k) \end{bmatrix} \right\} \xrightarrow{a.s.} 0. \end{aligned}$$

Proof. As before, for all $k \in \mathbb{N}$, let $Z_k \in \mathbb{R}^{n \times (n-m)}$ denote an orthonormal matrix whose columns form a basis for $\text{Null}(\nabla c(X_k))$, where if $X_k \in \mathcal{X}_{\epsilon, x_\star}$ then $Z_k = Z(X_k)$ is one such that (20) holds at $x = X_k$. For arbitrary $k \in \mathbb{N}$, it follows from (9) that the reduced gradient satisfies

$$Z_k^T \nabla f(X_k) = -Z_k^T H_k U_k^{\text{true}} - Z_k^T H_k V_k,$$

from which it follows with (17), (18), and Assumptions 2.1 and 2.4 that

$$\begin{aligned} \|Z_k^T \nabla f(X_k)\|_2^2 &= \|Z_k^T H_k U_k^{\text{true}} + Z_k^T H_k V_k\|_2^2 \\ &\leq 2(\|Z_k^T H_k U_k^{\text{true}}\|_2^2 + \|Z_k^T H_k V_k\|_2^2) \\ &\leq 2\kappa_H^2 (\|U_k^{\text{true}}\|_2^2 + \|V_k\|_2^2) \\ &\leq 2\kappa_H^2 (\|U_k^{\text{true}}\|_2^2 + \|(\nabla c(X_k)^\dagger)^T\|_2^2 \|c(X_k)\|_2^2) \\ &\leq 2\kappa_H^2 (\|U_k^{\text{true}}\|_2^2 + \kappa_c r^{-2} \|c(X_k)\|_2). \end{aligned}$$

Combining this with (20) and Lemma 3.2, one finds that $X_k \in \mathcal{X}_{\epsilon, x_\star}$ implies

$$\begin{aligned} &\phi_\tau(X_k) - \phi_\tau(x_\star) \\ &\leq \mu(2\tau\kappa_H^2 \|U_k^{\text{true}}\|_2^2 + (2\tau\kappa_c\kappa_H^2 r^{-2} + 1)\|c(X_k)\|_2) \\ &\leq \mu \max\{2\tau\kappa_H^2, 2\tau\kappa_c\kappa_H^2 r^{-2} + 1\} (\|D_k^{\text{true}}\|_2^2 + \|c(X_k)\|_2) \\ &\leq \mu \max\{2\tau\kappa_H^2, 2\tau\kappa_c\kappa_H^2 r^{-2} + 1\} (\kappa_\Psi + 1) \Psi_k \\ &\leq \mu \max\{2\tau\kappa_H^2, 2\tau\kappa_c\kappa_H^2 r^{-2} + 1\} (\kappa_\Psi + 1) \kappa_q^{-1} \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}). \end{aligned}$$

Since by Lemma 3.1 one has $\liminf_{k \rightarrow \infty} \Delta q_\tau(X_k, \nabla f(X_k), H_k, D_k^{\text{true}}) = 0$ almost surely, it follows from above and the conditions of the theorem that the limit $\lim_{k \rightarrow \infty} \phi_\tau(X_k)$ (which exists almost surely by Lemma 3.1) must be $\phi_\tau(x_\star)$, as desired. The remaining conclusions follow from $\{\phi_\tau(X_k)\} \xrightarrow{a.s.} \phi_\tau(x_\star)$ and the facts that, by (20), the point x_\star is the unique point in $\mathcal{X}_{\epsilon, x_\star}$ with merit function value equal to $\phi_\tau(x_\star)$ and that is a primal stationary point. \square

Observe that (20) can be viewed as a generalization of the well-known Polyak–Lojasiewicz condition from the unconstrained continuous optimization literature [31]. Indeed, if the constraints are affine and one considers x such that $x - x_\star \in \text{Null}(\nabla c(x_\star)^T)$, then (20) says that the squared ℓ_2 -norm of the reduced gradient $Z(x)^T \nabla f(x)$ is at least proportional to $\phi_\tau(x) - \phi_\tau(x_\star)$. On the other hand, if the objective function remains constant along displacements in $\text{Range}(\nabla c(x))$, then (20) says that the norm of the constraint violation is at least proportional to this difference in merit function values. More generally, such as for nonlinear f and c , the condition (20) is a generalization of these special cases that says that the combination of the squared ℓ_2 -norm of the reduced gradient and constraint violation is at least proportional to the optimality gap in the merit function in a neighborhood of the point x_\star .

4 Convergence of Lagrange multipliers

In this section, we study convergence properties of the sequence $\{Y_k\}$ generated by Algorithm 1. Let us begin by expressing the solution component Y_k of (9) for arbitrary $k \in \mathbb{N}$ using the step decomposition stated in (17) and (18). One finds by substituting (17) and (18) back into (9) gives

$$\begin{aligned} \nabla c(X_k) Y_k &= -(H_k(Z_k W_k + V_k) + G_k) \\ &= (I - H_k Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T) (H_k (\nabla c(X_k)^\dagger)^T c(X_k) - G_k), \end{aligned}$$

from which it follows under Assumption 2.1 that

$$Y_k = \nabla c(X_k)^\dagger (I - H_k Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T) (H_k (\nabla c(X_k)^\dagger)^T c(X_k) - G_k). \quad (21)$$

For notational convenience, let us now define the matrix

$$M_k := \nabla c(X_k)^\dagger (I - H_k Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T) \in \mathbb{R}^{m \times n}, \quad (22)$$

which can be viewed as the product of a pseudoinverse and a projection matrix, so that we may succinctly write from (21) that

$$Y_k = M_k (H_k (\nabla c(X_k)^\dagger)^T c(X_k) - G_k). \quad (23)$$

(We remark in passing that an alternative to our subsequent discussions and analysis could be considered where, for all $k \in \mathbb{N}$, the vector Y_k is not defined through (9), but rather is set as a so-called least squares multiplier, i.e., as $Y_k = \arg \min_{Y \in \mathbb{R}^m} \|G_k + \nabla c(X_k) Y\|_2^2$. For one thing, this choice removes the dependence of Y_k on H_k , which might have certain advantages. However, for our purposes, we focus on multipliers being computed through (9) since this is a popular approach in practice and does not require additional computation.)

Our goal in our analysis of properties of Lagrange multiplier estimators is to prove convergence of such estimators when the primal iterate approaches a primal stationary point for (4). After all, it is only when the primal iterate lies in such a neighborhood (or at least a neighborhood of the feasible region) that the Lagrange multiplier has meaning in terms of certifying stationarity. Consequently, a main focus in our results is how Y_k may be viewed through (23) as a function of the primal iterate X_k . For our analysis here, we make the following assumption, where given $x \in \mathbb{R}^n$ and $\epsilon \in \mathbb{R}_{>0}$ we define (consistent with Theorem 3.1) the neighborhood $\mathcal{X}_{\epsilon, x} := \{\bar{x} \in \mathbb{R}^n : \|\bar{x} - x\|_2 \leq \epsilon\}$.

Assumption 4.1. *Given $x_* \in \mathcal{X}$ as a primal stationary point for problem (4), there exist $\epsilon \in \mathbb{R}_{>0}$, $\mathcal{H} : \mathbb{R}^n \rightarrow \mathbb{S}^n$, $L_{\mathcal{H}} \in \mathbb{R}_{>0}$, $\mathcal{M} : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$, and $L_{\mathcal{M}} \in \mathbb{R}_{>0}$ such that the following hold.*

- (i) $H_k = \mathcal{H}(X_k)$ whenever $X_k \in \mathcal{X}_{\epsilon, x_*}$;
- (ii) $\|\mathcal{H}(x) - \mathcal{H}(\bar{x})\|_2 \leq L_{\mathcal{H}} \|x - \bar{x}\|_2$ for all $(x, \bar{x}) \in \mathcal{X}_{\epsilon, x_*} \times \mathcal{X}_{\epsilon, x_*}$;
- (iii) $M_k = \mathcal{M}(X_k)$ whenever $X_k \in \mathcal{X}_{\epsilon, x_*}$; and
- (iv) $\|\mathcal{M}(x) - \mathcal{M}(\bar{x})\|_2 \leq L_{\mathcal{M}} \|x - \bar{x}\|_2$ for all $(x, \bar{x}) \in \mathcal{X}_{\epsilon, x_*} \times \mathcal{X}_{\epsilon, x_*}$.

A few important observations and justifications are in order.

- Assumption 4.1 states that, in a neighborhood of a primal stationary point x_* , the algorithm sets the matrix H_k through a Lipschitz continuous function of X_k . This clearly holds if the algorithm chooses $H_k = H$ for all $k \in \mathbb{N}$ for some prescribed $H \in \mathbb{S}_{>0}^n$ that guarantees that Assumption 2.4 holds for all $k \in \mathbb{N}$. Alternatively, the requirements on $\{H_k\}$ in Assumption 4.1 can be satisfied if second-order derivatives of f are Lipschitz continuous and the algorithm chooses H_k as (an approximation of) the Hessian of the objective at X_k , or even as (an approximation of) the Hessian of the Lagrangian at (X_k, \bar{Y}_k) if the second-order derivatives of the components of c are also Lipschitz continuous and care is taken in the selection of \bar{Y}_k . In practice, one might consider $\bar{Y}_k = Y_{k-1}$ for all $k \in \mathbb{N}$, but as one can see, this choice is influenced by noise in the stochastic gradient estimators, which can cause $\{H_k\}$ to violate Assumption 2.4. An alternative choice that would not violate the assumption is to choose \bar{Y}_k as a prescribed vector or at least one that remains fixed for all sufficiently large $k \in \mathbb{N}$. In any case, one finds that there exist reasonable choices for the algorithm to choose $\{H_k\}$ such that Assumption 4.1 holds. (We remark in passing that deterministic Newton-based methods for solving constrained optimization problems can possess local convergence guarantees of the Lagrange multipliers under relatively strong assumptions when H_k is the Hessian of the Lagrangian at (X_k, Y_{k-1}) , and this might also be achievable for a stochastic algorithm with highly accurate gradient estimates. However, since our algorithm of interest operates in a highly stochastic regime in which only Assumption 2.3 is assumed to hold, allowing unadulterated use of the Hessian of the Lagrangian at (X_k, Y_{k-1}) is not reasonable.)

- Assumption 4.1 also states that, in a neighborhood of a primal stationary point x_* , the matrix M_k is defined through a Lipschitz continuous function of X_k . Given the assumptions about \mathcal{H} , Assumption 2.1 (which states that the constraint Jacobians have full row rank and the constraint Jacobian function is Lipschitz continuous over the set \mathcal{X} containing the iterates), and Lemma 1.1, it might seem at first glance that this assumption about \mathcal{M} is a straightforward consequence of these other assumptions. However, one needs to be careful. To understand a potential issue when trying to draw such a conclusion, recall that a product of functions that are Lipschitz continuous on a bounded set is itself Lipschitz continuous on the set. Hence, one might attempt to justify our Lipschitz-continuity assumption about \mathcal{M} by assuming that, for all $k \in \mathbb{N}$, the null space basis Z_k can be viewed as being generated by a Lipschitz continuous function of X_k . Note, however, that it has been shown in [11] that, even when $\nabla c(\cdot)^T$ is continuous with values that have full row rank for all $x \in \mathbb{R}^n$, there does not necessarily exist continuous $\mathcal{Z} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times (n-m)}$ such that for all $x \in \mathbb{R}^n$ the columns of $\mathcal{Z}(x)$ form a basis for $\text{Null}(\nabla c(x)^T)$. That being said, as also shown in [11], it is possible to employ procedures such that the reduced Hessian (approximation), namely, $Z_k^T H_k Z_k$, does not depend on the choice of Z_k , and, given a point $\bar{x} \in \mathbb{R}^n$ at which $\nabla c(\bar{x})^T$ has full row rank, there exists a neighborhood of \bar{x} over which one can define a “null space basis function” that is continuous over the neighborhood. See also [10], which discusses two procedures that ensure that, in the neighborhood of a given point, a null space basis can be defined by a Lipschitz continuous null space basis function. Overall, due to these observations, we contend that assuming Lipschitz continuity of \mathcal{M} —in a sufficiently small neighborhood of x_* , as is stated in Assumption 4.1—is justified for our analysis.

We are now prepared to prove that, near a primary stationary point, the expected error in the Lagrange multiplier estimator is bounded by the distance of the primal iterate to the primal stationary point plus an error due to the stochastic gradient estimator. The following theorem may be viewed as the main result of this section since the subsequent results follow from it.

Theorem 4.1. *Suppose that Assumptions 2.1, 2.2, 2.3, 2.4, and 4.1 hold, (x_*, y_*) is a stationary point for problem (4), and $\epsilon \in \mathbb{R}_{>0}$ is defined as in Assumption 4.1. Then, for any $k \in \mathbb{N}$, one finds $\|X_k - x_*\|_2 \leq \epsilon$ implies*

$$\|Y_k - y_*\|_2 \leq \kappa_y \|X_k - x_*\|_2 + r^{-1} \|\nabla f(X_k) - G_k\|_2, \quad (24)$$

where $\kappa_y := \kappa_H L_c r^{-2} + L_{\nabla f} r^{-1} + \kappa_{\nabla f} L_{\mathcal{M}}$.

Proof. Under the conditions of the theorem, $\|X_k - x_*\|_2 \leq \epsilon$ and (23) imply that $Y_k = M_k(H_k(\nabla c(X_k)^\dagger)^T c(X_k) - G_k)$, where $M_k = \mathcal{M}(X_k)$ and $H_k = \mathcal{H}(X_k)$, while one can similarly derive that $y_* = -M_* \nabla f(x_*)$, where $M_* = \mathcal{M}(x_*)$. Hence, the difference $Y_k - y_*$ can be decomposed into three parts as follows:

$$\begin{aligned} Y_k - y_* &= M_k H_k (\nabla c(X_k)^\dagger)^T c(X_k) \\ &\quad + M_k (\nabla f(x_*) - G_k) + (M_* - M_k) \nabla f(x_*). \end{aligned} \quad (25)$$

Consequently, to bound $\|Y_k - y_*\|_2$ when $\|X_k - x_*\|_2 \leq \epsilon$, one can employ the triangle inequality and bound the norms of the three terms on the right-hand side of (25) separately. First, one finds that

$$\begin{aligned} &\|M_k H_k (\nabla c(X_k)^\dagger)^T c(X_k)\|_2 \\ &\leq \|\nabla c(X_k)^\dagger\|_2 \|I - H_k Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T\|_2 \|H_k\|_2 \|(\nabla c(X_k)^\dagger)^T\|_2 \|c(X_k)\|_2 \\ &\leq \kappa_H r^{-2} \|c(X_k) - c(x_*)\|_2 \leq \kappa_H L_c r^{-2} \|X_k - x_*\|_2, \end{aligned} \quad (26)$$

where the first inequality follows from (22) and properties of norms; the second follows from $c(x_*) = 0$, Assumptions 2.1 and 2.4, and the fact that $P_k := I - H_k Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T$ is a projection matrix, so $\|P_k\|_2 \leq 1$; and the last follows from Lipschitz continuity of c (see (6)). Second,

$$\begin{aligned} &\|M_k (\nabla f(x_*) - G_k)\|_2 \\ &\leq \|\nabla c(X_k)^\dagger\|_2 \|I - H_k Z_k (Z_k^T H_k Z_k)^{-1} Z_k^T\|_2 \|\nabla f(x_*) - G_k\|_2 \end{aligned}$$

$$\begin{aligned}
&\leq r^{-1}\|\nabla f(x_\star) - G_k\|_2 \\
&\leq r^{-1}(\|\nabla f(x_\star) - \nabla f(X_k)\|_2 + \|\nabla f(X_k) - G_k\|_2) \\
&\leq L_{\nabla f}r^{-1}\|X_k - x_\star\|_2 + r^{-1}\|\nabla f(X_k) - G_k\|_2,
\end{aligned} \tag{27}$$

where the first inequality follows from (22) and properties of norms, the second follows from Assumption 2.1 and $\|P_k\|_2 \leq 1$ (see above), the third follows from $\nabla f(x_\star) - G_k = \nabla f(x_\star) - \nabla f(X_k) + \nabla f(X_k) - G_k$ and the triangle inequality, and the last follows from Lipschitz continuity of ∇f (see (6)). Third,

$$\|(M_\star - M_k)\nabla f(x_\star)\|_2 \leq \|M_\star - M_k\|_2\|\nabla f(x_\star)\|_2 \leq \kappa_{\nabla f}L_{\mathcal{M}}\|X_k - x_\star\|_2 \tag{28}$$

follows from properties of norms and Assumptions 2.1 and 4.1. Substituting (26), (27), and (28) into (25) yields the desired conclusion. \square

Theorem 4.1 shows that whenever X_k is sufficiently close to a primal stationary point x_\star , the Lagrange multiplier error $\|Y_k - y_\star\|_2$ is bounded by a constant times the primal iterate error $\|X_k - x_\star\|_2$ plus a constant times the error in the stochastic gradient estimator. The latter term is inevitable since, no matter the proximity of X_k to x_\star , the Lagrange multiplier estimator Y_k is influenced by G_k . To emphasize this point, we present the following corollary showing that, under the same conditions as Theorem 4.1, the Lagrange multiplier estimator Y_k^{true} —which appears in the convergence guarantee [3, Corollary 3.14], even though it is a quantity that is not actually computed in the algorithm—satisfies the same bound as Y_k , except without the additional noise.

Corollary 4.1. *Suppose that Assumptions 2.1, 2.2, 2.3, 2.4, and 4.1 hold, (x_\star, y_\star) is a stationary point for problem (4), and $\epsilon \in \mathbb{R}_{>0}$ is defined as in Assumption 4.1. Then, for any $k \in \mathbb{N}$, one finds $\|X_k - x_\star\|_2 \leq \epsilon$ implies*

$$\|Y_k^{\text{true}} - y_\star\|_2 \leq \kappa_y\|X_k - x_\star\|_2, \tag{29}$$

where $\kappa_y \in \mathbb{R}_{>0}$ is defined as in Theorem 4.1.

Proof. Following the same steps as in the proof of Theorem 4.1, the difference $Y_k^{\text{true}} - y_\star$ can now be decomposed as in (25) with $G_k = \nabla f(X_k)$. Therefore, in place of (27), one instead obtains the bound

$$\begin{aligned}
&\|\nabla c(X_k)^\dagger(I - H_k Z_k(Z_k^T H_k Z_k)^{-1}Z_k^T)(\nabla f(x_\star) - \nabla f(X_k))\|_2 \\
&\leq r^{-1}\|\nabla f(x_\star) - \nabla f(X_k)\|_2 \leq L_{\nabla f}r^{-1}\|X_k - x_\star\|_2,
\end{aligned}$$

which when combined with (26), (28), and (25) gives the desired result. \square

The previous theorem and corollary lead to the following corollary, which considers the special case that there exists an iteration number beyond which the primal iterates remain within a neighborhood of a particular primal stationary point almost surely. In such a case, the expected error in the Lagrange multiplier is bounded by a constant times the distance from the primal iterate to the primal stationary point plus a term that is proportional to the bound on the noise of the stochastic gradient estimators. The corollary also adds that the “true” multiplier does not have this additional error term. (We remind the reader that, in Section 3, we have shown conditions under which one can conclude that $\{X_k\} \xrightarrow{a.s.} x_\star$. While this convergence guarantee does not directly imply that $\bar{k} \in \mathbb{N}$ exists as in the following corollary, it motivates consideration of this special case as one of practical relevance.)

Corollary 4.2. *Suppose that Assumptions 2.1, 2.2, 2.3, 2.4, and 4.1 hold, (x_\star, y_\star) is a stationary point for problem (4), $\epsilon \in \mathbb{R}_{>0}$ is defined as in Assumption 4.1, and for some $\bar{k} \in \mathbb{N}$ one has almost surely that*

$$\|X_k - x_\star\|_2 \leq \epsilon \text{ for all } k \in \mathbb{N} \text{ with } k \geq \bar{k}.$$

Then, for any $k \in \mathbb{N}$ with $k \geq \bar{k}$, one finds that

$$\begin{aligned}
&\mathbb{E}[\|Y_k - y_\star\|_2 | \mathcal{F}_k] \leq \kappa_y\|X_k - x_\star\|_2 + r^{-1}\sigma \\
&\text{and } \mathbb{E}[\|Y_k^{\text{true}} - y_\star\|_2 | \mathcal{F}_k] \leq \kappa_y\|X_k - x_\star\|_2,
\end{aligned}$$

where $\kappa_y \in \mathbb{R}_{>0}$ is defined as in Theorem 4.1.

Proof. Assumption 2.3 and Jensen's inequality imply for all $k \in \mathbb{N}$ that

$$\mathbb{E}[\|G_k - \nabla f(X_k)\|_2 | \mathcal{F}_k]^2 \leq \mathbb{E}[\|G_k - \nabla f(X_k)\|_2^2 | \mathcal{F}_k] \leq \sigma^2.$$

The desired conclusions follow from this fact, the conditions of the corollary, and the conclusions of Theorem 4.1 and Corollary 4.1. \square

At this point, we have determined useful properties of two multiplier estimator sequences, $\{Y_k\}$ and $\{Y_k^{\text{true}}\}$. In the neighborhood of a stationary point (x_*, y_*) , each element Y_k^{true} of the latter sequence has error that is proportional to the distance from X_k to the primal stationary point x_* , which is as much as can be expected. However, this estimator is not realized by the algorithm as it requires explicit knowledge of the true gradient $\nabla f(X_k)$, which is not realized by the algorithm. Each element Y_k of the former sequence has the same error plus an additional error due to the stochastic gradient estimator. One can reduce this error by computing a more accurate gradient estimate, say, using a large mini-batch of gradient estimates. However, to avoid the need of computing a highly accurate gradient estimate in any given iteration, one might ask whether it is possible to obtain an estimator with lower error merely under Assumption 2.3? Using classical ideas from statistical estimation, specifically mean estimation, a natural idea to consider is whether a more accurate estimator can be defined by an average, say, $Y_k^{\text{avg}} := \frac{1}{k} \sum_{i=1}^k Y_k$. After all, if the stochastic process $\{X_k\}$ converges in some sense to x_* , then with Y_k defined through a continuous function of X_k in a neighborhood of x_* for all $k \in \mathbb{N}$, one might expect $\{Y_k^{\text{avg}}\}$ to converge in the same sense to y_* .

The main challenge in the analysis of $\{Y_k^{\text{avg}}\}$ is that, in contrast to the classical setting of mean estimation using independent and identically distributed random variables, the multipliers $\{Y_k\}$ are computed at points $\{X_k\}$ at which the distributions of $\{G_k\}$ (or $\{\nabla f(X_k)\}$ for that matter) are neither independent nor identically distributed. Hence, our analysis of the averaged sequence requires some additional assumptions that are introduced in the following lemma. We prove the lemma and a following theorem, then discuss and provide justifications for these required assumptions in further detail.

Lemma 4.1. *Suppose that Assumptions 2.1, 2.2, 2.3, 2.4, and 4.1 hold, and with M_k defined in (22) and $\Delta_k := \nabla f(X_k) - G_k$ for all $k \in \mathbb{N}$ one finds*

$$\frac{1}{k} \mathbb{E}[\|M_i \Delta_i\|_2^2] < \infty \text{ for all } (k, i) \in \mathbb{N} \times [k], \quad (30a)$$

$$\left\{ \frac{1}{k} \sum_{i=1}^k \mathbb{E} \left[\|M_i \Delta_i\|_2^2 \mathbf{1}_{\left\{ \frac{\|M_i \Delta_i\|_2}{\sqrt{k}} > \delta \right\}} \middle| \mathcal{F}_i \right] \right\} \xrightarrow{p} 0 \text{ for all } \delta \in \mathbb{R}_{>0}, \quad (30b)$$

$$\left\{ \frac{1}{k} \sum_{i=1}^k \mathbb{E}[M_i \Delta_i \Delta_i^T M_i^T | \mathcal{F}_i] \right\} \xrightarrow{p} \Sigma \text{ for some } \Sigma \in \mathbb{S}^n, \text{ and} \quad (30c)$$

$$\sup_{k \in \mathbb{N}} \mathbb{E} \left[\left\| \sum_{i=1}^k \frac{1}{\sqrt{k}} M_i \Delta_i \right\|_2^\Theta \right] < \infty \text{ for some } \Theta \in \mathbb{R}_{>1}. \quad (30d)$$

Then, it follows that

$$\lim_{k \rightarrow \infty} \frac{1}{k} \mathbb{E} \left[\left\| \sum_{i=1}^k M_i \Delta_i \right\|_2 \right] = 0. \quad (31)$$

Proof. For $(k, i) \in \mathbb{N} \times [k]$, let $\mathcal{F}_{k,i} := \mathcal{F}_i$ and $\xi_{k,i} := \frac{1}{\sqrt{k}} M_i \Delta_i$. Under Assumption 2.3, it follows that $\mathbb{E}[\xi_{k,i} | \mathcal{F}_{k,i}] = 0$ for all $(k, i) \in \mathbb{N} \times [k]$ and that $\{(\xi_{k,i}, \mathcal{F}_{k,i})\}_{k \in \mathbb{N}, i \in [k]}$ is an n -dimensional martingale difference triangular array (see Lemma 1.3). Employing Lemma 1.3, one finds that

$$\left\{ \sum_{i=1}^k \xi_{k,i} \right\} = \left\{ \sum_{i=1}^k \frac{1}{\sqrt{k}} M_i \Delta_i \right\} \xrightarrow{d} \mathcal{N}(0, \Sigma).$$

This fact, Lemma 1.4, and (30d) together yield

$$\lim_{k \rightarrow \infty} \mathbb{E} \left[\left\| \sum_{i=1}^k \frac{1}{\sqrt{k}} M_i \Delta_i \right\|_2 \right] = \tilde{\sigma} < \infty,$$

where $\tilde{\sigma} := \mathbb{E}[\|X\|_2]$ for $X \sim \mathcal{N}(0, \Sigma)$. Therefore,

$$\lim_{k \rightarrow \infty} \frac{1}{k} \mathbb{E} \left[\left\| \sum_{i=1}^k M_i \Delta_i \right\|_2 \right] = \left(\lim_{k \rightarrow \infty} \frac{1}{\sqrt{k}} \right) \left(\lim_{k \rightarrow \infty} \mathbb{E} \left[\left\| \frac{1}{\sqrt{k}} \sum_{i=1}^k M_i \Delta_i \right\|_2 \right] \right) = 0,$$

which gives the desired result (31). \square

The previous lemma allows us to prove the following theorem. We remark upfront that the iteration index $\bar{k} \in \mathbb{N}$ that is used in the definition of $\{Y_k^{\text{avg}}\}_{k=\bar{k}}^\infty$ in the theorem is not necessarily known in advance by the algorithm. In particular, it is not necessarily known by the algorithm when the primal iterate sequence may be guaranteed to have entered a neighborhood of a primary stationary point such that the conditions of Assumption 4.1 hold with a certain $\epsilon \in \mathbb{R}_{>0}$. A stronger assumption would be to assume that $\bar{k} = 1$, i.e., that the algorithm is initialized at a point within a sufficiently small neighborhood of x_* in which it remains almost surely. However, we state the result for general \bar{k} for the sake of generality, and since in practice it would be reasonable to consider a scheme that only computes an averaged Lagrange multiplier estimate over those recent multipliers such that the corresponding primal iterates are within a prescribed neighborhood of the current iterate. We consider such a strategy in our experiments in Section 5.

Theorem 4.2. *Suppose that Assumptions 2.1, 2.2, 2.3, 2.4, and 4.1 hold, (x_*, y_*) is a stationary point for problem (4), $\epsilon \in \mathbb{R}_{>0}$ is defined as in Assumption 4.1, for some $\bar{k} \in \mathbb{N}$ one has almost surely that*

$$\|X_k - x_*\|_2 \leq \epsilon \text{ for all } k \in \mathbb{N} \text{ with } k \geq \bar{k},$$

$\{M_k\}$ and $\{\Delta_k\}$ are defined as in Lemma 4.1, and (30) holds. Then, with $Y_k^{\text{avg}} := \frac{1}{k-\bar{k}+1} \sum_{i=\bar{k}}^k Y_i$ for all $k \in \mathbb{N}$ with $k \geq \bar{k}$, one finds that

$$\limsup_{k \geq \bar{k}, k \rightarrow \infty} \mathbb{E}[\|Y_k^{\text{avg}} - y_*\|_2] \leq \limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{\kappa_y}{k-\bar{k}+1} \sum_{i=\bar{k}}^k \mathbb{E}[\|X_i - x_*\|_2].$$

Proof. As in the proof of Theorem 4.1, one finds for any $k \in \mathbb{N}$ with $k \geq \bar{k}$ that

$$\begin{aligned} \|Y_k^{\text{avg}} - y_*\|_2 &= \frac{1}{k-\bar{k}+1} \left\| \sum_{i=\bar{k}}^k (Y_i - y_*) \right\|_2 \\ &\leq \frac{\kappa_H L_c r^{-2} + \kappa_{\nabla f} L_{\mathcal{M}}}{k-\bar{k}+1} \sum_{i=\bar{k}}^k \|X_i - x_*\|_2 \\ &\quad + \frac{1}{k-\bar{k}+1} \left\| \sum_{i=\bar{k}}^k M_i (\nabla f(x_*) - \nabla f(X_k) + \nabla f(X_k) - G_k) \right\|_2 \\ &\leq \frac{\kappa_y}{k-\bar{k}+1} \sum_{i=\bar{k}}^k \|X_i - x_*\|_2 + \frac{1}{k-\bar{k}+1} \left\| \sum_{i=\bar{k}}^k M_i \Delta_i \right\|_2, \end{aligned} \tag{32}$$

Thus, under the conditions of the theorem, it follows with Lemma 4.1 that

$$\limsup_{k \geq \bar{k}, k \rightarrow \infty} \mathbb{E}[\|Y_k^{\text{avg}} - y_*\|_2]$$

$$\begin{aligned}
&\leq \limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{\kappa_y}{k - \bar{k} + 1} \sum_{i=\bar{k}}^k \mathbb{E}[\|X_i - x_\star\|_2] + \limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{1}{k - \bar{k} + 1} \mathbb{E} \left[\left\| \sum_{i=\bar{k}}^k M_i \Delta_i \right\|_2 \right] \\
&= \limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{\kappa_y}{k - \bar{k} + 1} \sum_{i=\bar{k}}^k \mathbb{E}[\|X_i - x_\star\|_2],
\end{aligned}$$

as desired. \square

The conditions in (30) in Lemma 4.1 are admittedly nontrivial, but upon closer inspection it is not surprising that conditions akin to these would be needed in order to draw any useful conclusions about an averaged Lagrange multiplier sequence. First, consider conditions (30a) and (30d), which essentially require that the expected value of the norm of $M_k \Delta_k$ is bounded uniformly over $k \in \mathbb{N}$. Indeed, observe that if one were instead to assume that

$$\mathbb{E}[\|M_k \Delta_k\|_2^2] < \infty \text{ for all } k \in \mathbb{N}, \quad (33)$$

then (30a) follows directly while one also finds

$$\sup_{k \in \mathbb{N}} \mathbb{E} \left[\left\| \sum_{i=1}^k \frac{1}{\sqrt{k}} M_i \Delta_i \right\|_2^2 \right] \leq \sup_{k \in \mathbb{N}} \frac{1}{k} \sum_{i=1}^k \mathbb{E} [\|M_i \Delta_i\|_2^2] < \infty,$$

which yields (30d) for $\Theta = 2$. It is not surprising that this expected value would be required to be bounded uniformly since the elements of $\{M_k \Delta_k\}$ arise as the error terms in the proof of Theorem 4.2. Second, beyond merely being bounded, which would not be sufficient on its own, conditions (30b) and (30c) require convergence properties of $\{M_k \Delta_k\}$ that match those needed for the multidimensional martingale central limit theorem, which we have stated as Lemma 1.3. Notice that if one considers the conditions of Assumption 4.1, which have that the elements of $\{M_k\}$ are defined by a continuous (in fact, Lipschitz continuous) mapping \mathcal{M} of the elements of $\{X_k\}$, then (30b) and (30c) essentially require convergence in probability of averages of inner and outer products of the elements of $\{\mathcal{M}(X_k)(\nabla f(X_k) - G_k)\}$, which is a reasonable assumption to consider when, say, $\{X_k\} \xrightarrow{a.s.} x_\star$.

We close this section with the following corollary to Theorem 4.2.

Corollary 4.3. *Suppose that the conditions of Theorem 4.2 hold and that the iterate sequence converges almost surely to x_\star , i.e., $\{X_k\} \xrightarrow{a.s.} x_\star$. Then,*

$$\{Y_k^{\text{true}}\} \xrightarrow{a.s.} y_\star \text{ and } \{Y_k^{\text{avg}}\}_{k=\bar{k}}^\infty \xrightarrow{a.s.} y_\star.$$

Proof. The almost-sure convergence of $\{X_k\}$ to x_\star implies that

$$\mathbb{P} \left[\omega \in \Omega : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \|X_k(\omega) - x_\star\|_2 > 0 \right] = 0.$$

In addition, from (29), one finds for all $\omega \in \Omega$ that

$$\limsup_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{true}}(\omega) - y_\star\|_2 \leq \limsup_{k \geq \bar{k}, k \rightarrow \infty} \kappa_y \|X_k(\omega) - x_\star\|_2.$$

These facts combined with nonnegativity of norms yields

$$\begin{aligned}
&\mathbb{P} \left[\omega \in \Omega : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{true}}(\omega) - y_\star\|_2 > 0 \right] \\
&\leq \mathbb{P} \left[\omega \in \Omega : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \|X_k(\omega) - x_\star\|_2 > 0 \right] = 0,
\end{aligned}$$

which again with nonnegativity of norms further implies

$$\begin{aligned} & \mathbb{P} \left[\omega \in \Omega : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{true}}(\omega) - y_\star\|_2 = 0 \right] = 1 \\ \implies & \mathbb{P} \left[\omega \in \Omega : \lim_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{true}}(\omega) - y_\star\|_2 = 0 \right] = 1. \end{aligned}$$

Consequently, $\{Y_k^{\text{true}}\}_{k=\bar{k}}^\infty \xrightarrow{a.s.} y_\star$, so $\{Y_k^{\text{true}}\} \xrightarrow{a.s.} y_\star$, as desired.

Let us now consider the averaged sequence. Recall that under the conditions of the corollary one has for all $\omega \in \Omega$ that

$$\begin{aligned} & \sup_{k \geq \bar{k}, k \in \mathbb{N}} \frac{1}{\sqrt{k-\bar{k}+1}} \left\| \sum_{i=\bar{k}}^k M_i(\omega) \Delta_i(\omega) \right\|_2 < \infty \\ \implies & \limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{1}{k-\bar{k}+1} \left\| \sum_{i=\bar{k}}^k M_i(\omega) \Delta_i(\omega) \right\|_2 = 0. \end{aligned}$$

Thus, one obtains along with (30d) that

$$\begin{aligned} & \mathbb{P} \left[\omega \in \Omega : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{1}{k-\bar{k}+1} \left\| \sum_{i=\bar{k}}^k M_i(\omega) \Delta_i(\omega) \right\|_2 = 0 \right] \\ \geq & \mathbb{P} \left[\omega \in \Omega : \sup_{k \geq \bar{k}, k \in \mathbb{N}} \frac{1}{\sqrt{k-\bar{k}+1}} \left\| \sum_{i=\bar{k}}^k M_i(\omega) \Delta_i(\omega) \right\|_2 < \infty \right] = 1. \end{aligned}$$

Thus, there exists a null (possibly empty) set $\Omega_0 \subset \Omega$ such that

$$\limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{1}{k-\bar{k}+1} \left\| \sum_{i=\bar{k}}^k M_i(\omega) \Delta_i(\omega) \right\|_2 = 0 \text{ for all } \omega \in \Omega \setminus \Omega_0. \quad (34)$$

On the other hand, $\{X_k\} \xrightarrow{a.s.} x_\star$ implies $\{\|X_k - x_\star\|_2\} \xrightarrow{a.s.} 0$. Thus, there exists a null (possibly empty) set $\Omega_1 \subset \Omega$ such that

$$\limsup_{k \geq \bar{k}, k \rightarrow \infty} \frac{1}{k-\bar{k}+1} \sum_{i=\bar{k}}^k \|X_i(\omega) - x_\star\|_2 = 0 \text{ for all } \omega \in \Omega \setminus \Omega_1. \quad (35)$$

Then, recall from (32) that for all $\omega \in \Omega$ one finds

$$\begin{aligned} & \limsup_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{avg}}(\omega) - y_\star\|_2 \\ \leq & \limsup_{k \geq \bar{k}, k \rightarrow \infty} \left(\frac{\kappa_y}{k-\bar{k}+1} \sum_{i=\bar{k}}^k \|X_i(\omega) - x_\star\|_2 + \frac{1}{k-\bar{k}+1} \left\| \sum_{i=\bar{k}}^k M_i(\omega) \Delta_i(\omega) \right\|_2 \right), \end{aligned}$$

which along with the fact that Ω_0 and Ω_1 are null sets, (34), and (35) implies

$$\mathbb{P} \left[\omega \in \Omega : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{avg}}(\omega) - y_\star\|_2 = 0 \right]$$

$$\begin{aligned}
&= \mathbb{P} \left[\omega \in \Omega \setminus (\Omega_0 \cup \Omega_1) : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{avg}}(\omega) - y_\star\|_2 = 0 \right] \\
&\geq \mathbb{P} \left[\omega \in \Omega \setminus (\Omega_0 \cup \Omega_1) : \limsup_{k \geq \bar{k}, k \rightarrow \infty} \left(\frac{\kappa_y}{k - \bar{k} + 1} \sum_{i=\bar{k}}^k \|X_i(\omega) - x_\star\|_2 \right. \right. \\
&\quad \left. \left. + \frac{1}{k - \bar{k} + 1} \left\| \sum_{i=\bar{k}}^k M_i(\omega) \Delta_i(\omega) \right\|_2 \right) = 0 \right] = 1,
\end{aligned}$$

Combined with nonnegativity of norms, it follows that

$$\mathbb{P} \left[\omega \in \Omega : \lim_{k \geq \bar{k}, k \rightarrow \infty} \|Y_k^{\text{avg}}(\omega) - y_\star\|_2 = 0 \right] = 1,$$

which further implies that $\{Y_k^{\text{avg}}\}_{k=\bar{k}}^\infty \xrightarrow{a.s.} y_\star$, as desired. \square

5 Numerical demonstrations

In this section, we provide the results of a small set of numerical experiments to demonstrate the results of our theoretical analysis. Specifically, we show when solving each instance in a set of test problems that the primal iterates generated by an implementation of Algorithm 1 approach a primal stationary point (in fact, a global minimizer), and as this occurs one finds that the true Lagrange multiplier sequence converges and averaged Lagrange multiplier sequences converge despite the fact that the sequence of Lagrange multipliers itself does not converge (due to gradient errors). This behavior of the algorithm is representative of behavior that we have witnessed when solving other problems (not shown due to space considerations) as well. In terms of averaged Lagrange multipliers, we consider averages over all iterations as well as averages determined through practical strategies that attempt to ignore multipliers obtained at primal iterates that are far from a primal stationary point (or at least the current iterate generated by the algorithm).

For our numerical demonstrations, we consider constrained logistic regression; see also [2]. In particular, we consider problem instances of the form

$$\min_{x \in \mathbb{R}^n} \frac{1}{N} \sum_{i=1}^N \log(1 + e^{-\gamma_i d_i^T x}) \quad \text{s.t.} \quad Ax = b, \quad \|x\|_2^2 = 1,$$

where $D = [d_1 \ \dots \ d_N] \in \mathbb{R}^{n \times N}$ is a feature matrix, $\gamma \in \mathbb{R}^N$ is a label vector, $A \in \mathbb{R}^{m \times n}$, and $b \in \mathbb{R}^m$. This problem is nonconvex due to the nonlinear constraint $\|x\|_2^2 = 1$, although we confirmed in all runs that the primal iterates approach the globally optimal solution; see below. For the feature matrices and label vectors, we considered four datasets from LIBSVM [13], namely, **a9a** ($(n, N) = (123, 32561)$), **australian** ($(n, N) = (14, 690)$), **ionosphere** ($(n, N) = (34, 351)$), and **splice** ($(n, N) = (60, 1000)$). These datasets were selected as ones that led to relatively large errors in the stochastic gradient estimates. (See below for further information about how the gradient estimates were computed in our experiments.) For each problem instance, we choose $m = 10$ and randomly generated the initial point x_1 and constraint data (A, b) , keeping these quantities fixed for all runs for a given problem instance. Specifically, each entry of these quantities was drawn from a standard normal distribution and it was confirmed in each case that A had full row rank.

For the implementation of Algorithm 1, we used the Matlab implementation for [3] known as **StochasticSQP**.¹ For simplicity, we set h_k as the identity matrix I for all iterations in all runs of the algorithm. We solved each linear system (9) with Matlab's built-in `mldivide(\)` function, which in particular meant that—as would be recommended in general for implementations of the algorithm—all Lagrange multipliers were computed directly from this linear system, not by computing a null space basis Z_k . We chose $\beta_k = \mathcal{O}(\frac{1}{k})$ with $\beta_1 = 1$.

¹<https://github.com/frankecurtis/StochasticSQP>




Figure 1: Distances of primal iterates and Lagrange multiplier estimates to solution values over runs of Algorithm 1 to solve constrained logistic regression problems defined using four LIBSVM datasets.

For each problem instance, to obtain (x_*, y_*) , we initially ran `StochasticSQP` with true gradients in all iterations until a stationary point was found to high accuracy, then verified that at these points both first- and second-order stationary conditions held to high accuracy, thus verifying (approximate) global optimality of the final iterate. Then, for each instance, we ran `StochasticSQP` with mini-batch gradient estimates with mini-batch sizes of 16 to generate $\{(x_k, y_k, y_k^{\text{true}})\}$. (For each run, the sequence $\{\beta_k\}$ was chosen as an unsummable, yet square-summable sequence that was tuned in a manner such that the primal iterates tended toward the globally optimal primal solution.) We verified that in all runs the adaptive parameter rules implemented in `StochasticSQP` kept the merit parameter sequence $\{\tau_k\}$ constant at the initial value of 0.1 and kept the ratio parameter sequence $\{\xi_k\}$ constant at the initial value of 1.0; hence, the runs respected the simplified variant of the algorithm from [3] that we have analyzed as Algorithm 1.

The results for each problem instance are shown in Figure 1. Each plot in the figure shows the results of a single run with a budget of $k = 10^5$ iterations, although we verified that the results were qualitatively consistent in other runs (with different random seeds) as well. In each plot, the figure shows the sequences $\{\|x_k - x_*\|_2\}$, $\{\|y_k - y_*\|_2\}$, and $\{\|y_k^{\text{true}} - y_*\|_2\}$, where the quantities x_k , y_k , and y_k^{true} are defined as in the prior sections. The figure also shows the sequences $\{\|y_k^{\text{avg}} - y_*\|_2\}$ and $\{\|y_k^{\text{avg}, \epsilon} - y_*\|_2\}$ for a few values of $\epsilon \in \mathbb{R}_{>0}$, where y_k^{avg} is the averaged Lagrange multiplier over all iterations (from the initial point) whereas $y_k^{\text{avg}, \epsilon}$ is the averaged Lagrange multiplier over iterations $\{k', \dots, k\}$, where for each $k \in \mathbb{N}$ the index k' is the smallest value in $[k]$ such that $\|x_j - x_k\|_2 \leq \epsilon$ for all $j \in \{x_{k'}, \dots, x_k\}$. In this manner, for relatively large values of ϵ , the average is likely to be taken over more iterations (thus reducing noise), whereas for relatively small values of ϵ , the average is taken only over iterations in which the primal point is relatively close to the current primal iterate (to improve accuracy of the estimate as x_k nears x_*).

Figure 1 demonstrates the results of our theoretical analysis. In each run, within the iterations performed before the computational budget was exceeded, the distance of the primal iterate to the primal solution typically reduces monotonically (due to the tuned step-size sequence, as mentioned). Thus, the run provides an instance in which the primal iterates appear to converge despite the use of stochastic gradient estimates. At the same time, the distance of the Lagrange multiplier estimate to the optimal Lagrange multiplier oscillates between relatively large values due to the errors in the stochastic gradient estimates, whereas the

distances of true and averaged Lagrange multiplier estimates reach much smaller values. Since the true Lagrange multiplier estimates are unobtainable in practice, one finds that averaged Lagrange multipliers may be considered to obtain better estimates than the $\{y_k\}$ sequence itself. In some cases, averaging over all iterations may be sufficient, but if desired one may consider choosing $\epsilon \in \mathbb{R}_{>0}$ (tuning the value, if computationally feasible) and considering a sequence such as $\{y_k^{\text{avg}_\epsilon}\}$ instead.

6 Conclusion

We have presented new convergence analyses for a stochastic SQP method built on the stochastic gradient methodology. In particular, for a simplified variant of the algorithm from [3], we have proved almost-sure convergence guarantees for the primal iterates, Lagrange multiplier estimates, and stationarity measures. It has been shown that the error in the Lagrange multipliers is bounded by the distance of the primal iterate to a primal stationary point plus the error due to stochastic gradient estimate. Furthermore, under modest assumptions, the latter error can be shown to vanish by incorporating a running average of the Lagrange multipliers. The results of numerical experiments demonstrate our proved theoretical guarantees.

Acknowledgements

This work is supported by the U.S. National Science Foundation (NSF) award IIS-1704458 and U.S. Office of Naval Research award N00014-21-1-2532.

References

- [1] Albert S. Berahas, Raghu Bollapragada, and Baoyu Zhou. An adaptive sampling sequential quadratic programming method for equality constrained stochastic optimization. *arXiv e-prints*, arXiv:2206.00712v2, 2023.
- [2] Albert S. Berahas, Frank E. Curtis, Michael J. O’Neill, and Daniel P. Robinson. A stochastic sequential quadratic optimization algorithm for nonlinear equality constrained optimization with rank-deficient Jacobians. *arXiv e-prints*, arXiv:2106.13015v2, 2023.
- [3] Albert S. Berahas, Frank E. Curtis, Daniel P. Robinson, and Baoyu Zhou. Sequential quadratic optimization for nonlinear equality constrained stochastic optimization. *SIAM Journal on Optimization*, 31(2):1352–1379, 2021.
- [4] Albert S. Berahas, Jiahao Shi, Zihong Yi, and Baoyu Zhou. Accelerating stochastic sequential quadratic programming for equality constrained optimization using predictive variance reduction. *Computational Optimization and Applications*, 86(1):79–116, 2023.
- [5] Albert S. Berahas, Miaolan Xie, and Baoyu Zhou. A sequential quadratic programming method with high probability complexity bounds for nonlinear equality constrained stochastic optimization. *arXiv e-prints*, arXiv:2301.00477, 2023.
- [6] Dimitri P. Bertsekas. *Network Optimization: Continuous and Discrete Methods*. Optimization and Neural Computation Series. Athena Scientific, 1998.
- [7] Dimitri P. Bertsekas and John N. Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.
- [8] John T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. Advances in Design and Control. Society for Industrial and Applied Mathematics, 2nd edition, 2010.

- [9] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- [10] Richard H. Byrd and Jorge Nocedal. An analysis of reduced Hessian methods for constrained optimization. *Mathematical Programming*, 49:285–323, 1990.
- [11] Richard H. Byrd and Robert B. Schnabel. Continuity of the null space basis and constrained optimization. *Mathematical Programming*, 35:32–41, 1986.
- [12] Antonin Chambolle and Thomas Pock. On the ergodic convergence rates of a first-order primal–dual algorithm. *Mathematical Programming*, 159(1-2):253–287, 2016.
- [13] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):1–27, 2011.
- [14] Kai Lai Chung. *A Course in Probability Theory, Revised Edition.*, volume 3rd ed. Academic Press, 2001.
- [15] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(3):88, 2022.
- [16] Frank E. Curtis, Michael J. O’Neill, and Daniel P. Robinson. Worst-case complexity of an SQP method for nonlinear equality constrained stochastic optimization. *Mathematical Programming*, 2023. <https://doi.org/10.1007/s10107-023-01981-1>.
- [17] Frank E. Curtis, Daniel P. Robinson, and Baoyu Zhou. Sequential quadratic optimization for stochastic optimization with deterministic nonlinear inequality and equality constraints. *arXiv e-prints*, arXiv:2302.14790, 2023.
- [18] Richard Durrett. *Probability: Theory and Examples*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 5th edition, 2019.
- [19] Yuchen Fang, Sen Na, Michael W. Mahoney, and Mladen Kolar. Fully stochastic trust-region sequential quadratic programming for equality-constrained optimization problems. *arXiv e-prints*, arXiv:2211.15943, 2022.
- [20] Peter Hall and Christopher C. Heyde. *Martingale Limit Theory and Its Application*. Academic Press, 2014.
- [21] Xin Jiang and Lieven Vandenbergh. Bregman three-operator splitting methods. *Journal of Optimization Theory and Applications*, 196(3):936–972, 2023.
- [22] George E. Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- [23] F.-S. Kupfer and Ekkehard W. Sachs. Numerical solution of a nonlinear parabolic control problem by a reduced SQP method. *Computational Optimization and Applications*, 1(1):113–135, 1992.
- [24] Tiejun Li, Tiannan Xiao, and Guoguo Yang. Revisiting the central limit theorems for the SGD-type methods. *arXiv e-prints*, arXiv:2207.11755, 2023.
- [25] Don L. McLeish. Dependent central limit theorems and invariance principles. *The Annals of Probability*, 2(4), 1974.
- [26] Sen Na, Mihai Anitescu, and Mladen Kolar. An adaptive stochastic sequential quadratic programming with differentiable exact augmented Lagrangians. *Mathematical Programming*, 199(1-2):721–791, 2023.

- [27] Sen Na, Mihai Anitescu, and Mladen Kolar. Inequality constrained stochastic nonlinear optimization via active-set sequential quadratic programming. *Mathematical Programming*, 2023. <https://doi.org/10.1007/s10107-023-01935-7>.
- [28] Sen Na and Michael W. Mahoney. Asymptotic convergence rate and statistical inference for stochastic sequential quadratic programming. *arXiv e-prints*, arXiv:2205.13687, 2022.
- [29] Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- [30] Figen Oztoprak, Richard Byrd, and Jorge Nocedal. Constrained optimization in the presence of noise. *arXiv e-prints*, arXiv:2110.04355, 2021.
- [31] B. T. Polyak. Gradient methods for minimization of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(3):643–653, 1963.
- [32] Boris T. Polyak and Anatoli B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [33] Songqiang Qiu and Vyacheslav Kungurtsev. A sequential quadratic programming method for optimization with stochastic objective functions, deterministic inequality constraints and robust subproblems. *arXiv e-prints*, arXiv:2302.07947, 2023.
- [34] Tyrone Rees, H. Sue Dollar, and Andrew J. Wathen. Optimal solvers for PDE-constrained optimization. *SIAM Journal on Scientific Computing*, 32(1):271–298, 2010.
- [35] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- [36] Herbert Robbins and David Siegmund. A convergence theorem for non-negative almost supermartingales and some applications. In *Optimizing Methods in Statistics*, pages 233–257. Elsevier, 1971.
- [37] Per-Åke Wedin. Perturbation theory for pseudo-inverses. *BIT Numerical Mathematics*, 13(2):217–232, 1973.