

Modified Line Search Sequential Quadratic Methods for Equality-Constrained Optimization with Unified Global and Local Convergence Guarantees

Albert S. Berahas^{†*} Raghu Bollapragada[‡] Jiahao Shi[†]

Abstract

In this paper, we propose a method that has foundations in the line search sequential quadratic programming paradigm for solving general nonlinear equality constrained optimization problems. The method employs a carefully designed modified line search strategy that utilizes second-order information of both the objective and constraint functions, as required, to mitigate the Maratos effect. Contrary to classical line search sequential quadratic programming methods, our proposed method is endowed with global convergence and local superlinear convergence guarantees. Moreover, we extend the method and analysis to the setting in which the constraint functions are deterministic but the objective function is stochastic or can be represented as a finite-sum. We also design and implement a practical inexact matrix-free variant of the method. Finally, numerical results illustrate the efficiency and efficacy of the method.

1 Introduction

In this paper, we propose a sequential quadratic programming (SQP) algorithm for solving equality constrained optimization problems of the form

$$\min_{x \in \mathbb{R}^n} f(x) \quad \text{s.t. } c(x) = 0, \quad (1.1)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are twice continuously differentiable nonlinear functions. Problems of this structure arise in a plethora of science and engineering applications; see [27, 60] and the references therein. Our proposed algorithm has foundations in the line search SQP paradigm, utilizes second-order information of the objective and constraint functions, and employs a carefully designed modified line search strategy. The method is

[†]Department of Industrial and Operations Engineering, University of Michigan. (aberahas@umich.edu, jiahaos@umich.edu)

[‡]Operations Research Program, University of Texas at Austin. (raghu.bollapragada@utexas.edu)

*Corresponding author.

endowed with both global convergence and local fast (superlinear) convergence guarantees. This is contrary to classical line search (and other) SQP methods that assume proximity of the initial iterate to the solution and a unit step size to prove analogous results; see e.g., [41, 53, 57].

Moreover, we extend the modified line search SQP methodology to the equality constrained stochastic setting, i.e.,

$$\min_{x \in \mathbb{R}^n} f(x) = \mathbb{E}[F(x, \omega)] \quad \text{s.t. } c(x) = 0, \quad (1.2)$$

where ω is a random variable with probability space (Ω, \mathcal{F}, P) that represents the stochasticity in the problem, $F : \mathbb{R}^n \times \Omega \rightarrow \mathbb{R}$, and $\mathbb{E}[\cdot]$ is the expectation taken with respect to the distribution of ω . In this problem setting, it is assumed that one has access to exact values of the constraint function and its derivatives, but only stochastic estimates of the objective function and its derivatives. Such problems arise in numerous application areas including machine learning [45, 63], optimal power flow [72, 76], statistical estimation [18, 36], portfolio optimization [69, 78], PDE-constrained optimization [9, 61], and simulation-based optimization [1, 33]. As a special case, and since the distribution of ω is often unknown in practice, we also consider the empirical risk variant [70] of (1.2), i.e.,

$$\min_{x \in \mathbb{R}^n} f(x) = \frac{1}{N} \sum_{i=1}^N f_i(x) = \frac{1}{N} \sum_{i=1}^N F(x, \omega_i) \quad \text{s.t. } c(x) = 0, \quad (1.3)$$

where N denotes the number of samples, and $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ (for $i \in \{1, 2, \dots, N\}$) denote the component functions that are twice continuously differentiable, ω_i is a random realization generated from the distribution P , and $F : \mathbb{R}^n \times \Omega \rightarrow \mathbb{R}$. The proposed approaches for solving (1.2) and (1.3) employ inexact objective function information, and as such the strategies need to be modified appropriately. Under reasonable assumptions on the stochasticity, we show local fast (superlinear) convergence rates analogous to those in the deterministic setting for these problem classes.

1.1 Background and Related Work

As a disclaimer, we note that this is not an extensive review of methods for solving problems of the form (1.1), (1.2) and (1.3). Rather, it is a concise summary of the works most closely related to our proposed approach. We begin with a detour to the unconstrained (deterministic and stochastic) settings and the fast local convergence of Newton’s method, and then discuss SQP methods for equality constrained deterministic and stochastic optimization.

In the unconstrained setting, it is well established that Newton’s method, under reasonable assumptions, is globally convergent and has fast local convergence [51, 56]. This requires a globalization mechanism (e.g., line search) that controls the length of the step when the iterates are far from stationarity, and results in an adaptive algorithm with

eventual fast local quadratic convergence. In the unconstrained stochastic setting, several second-order methods have been proposed; see e.g., [2, 10, 29, 55, 62]. We focus on [10, 62]. In [10], a subsampled Newton method is analyzed, and under appropriate sampling (gradient and Hessian) conditions, global linear and local superlinear convergence guarantees are established in expectation. In [62], the same method is explored and high probability convergence guarantees are derived. We note that for both approaches (and all other stochastic Newton methods listed above) the fast local results are derived under a proximity to the optimal solution assumption and under the assumption that the unit step size is employed. Our goal is to derive analogous results for the equality constrained setting without the last two assumptions.

Classical deterministic second-order SQP methods have proven robust and efficient in practice across many applications [38, 65], and are endowed with local superlinear convergence guarantees; see e.g., [41, 53, 57, 75]. To establish such fast local guarantees, the results critically rely on a proximity assumption between the starting point and the optimal solution and on the assumption that the unit step size is employed. The former assumption depends on problem specific, often unknown, quantities, and employing the unit step size far from the solution can lead to undesirable behavior. To alleviate these issues and to globalize SQP methods, several techniques have been proposed; see e.g., [6, 14, 42, 58, 71]. That said, unlike the unconstrained setting for which (as mentioned above) globalized versions of Newton’s method have provable “two-phase” behavior (global convergence and local fast convergence) [56, Theorem 6.2.31], in the equality constrained setting such results cannot be proven due to the Maratos effect [51, Section 15.5]. In fact, there exist simple counterexamples for which the conditions required for local fast convergence are never guaranteed to be satisfied for classical globalized SQP methods [44, Chapter III.5]. Several remedies have been proposed to mitigate these issues; see e.g., [16, 34, 64] and Section 2 for more details. On a high level, these approaches either change the algorithm, modify the SQP subproblem, allow for nonmonotonicity or restrict the problem class considered. In this work, we propose a second-order SQP method with a modified line search condition that allows for both global and fast local convergence guarantees without the unit step size and proximity assumptions.

Several methods have been proposed for solving constrained stochastic optimization problems of the form (1.2) and (1.3); see e.g., [3, 5–8, 24–26, 30, 48, 49, 52]. While stochastic penalty methods have been proposed for this problem class, we focus here on SQP methods due to their demonstrated empirical superiority in both the deterministic [40, 42, 58] and stochastic [6, 48] settings. The proposed SQP-type methods are adaptive (line search, trust region or other), and utilize, insofar as possible, well-established techniques from the deterministic setting. That said, with the exception of [48, 50], all methods are essentially first-order methods. We should note that the presentation of most of the algorithms is general and allows for second-order information to be utilized, however, the analyses (even in [48, 50]) do not explicitly account for second-order information and no local fast results are established, and numerical experiments consider only first-order variants. Our goal is

to develop a practical stochastic SQP method endowed with local superlinear convergence rates.

1.2 Summary of Contributions

- We develop, analyze, and implement a second-order line search SQP method with a novel modified line search condition. The method is designed to solve problems of the form (1.1), and is extended to tackle stochastic problems. Our algorithm leverages the global convergence properties of classical line search SQP methods, and employs a carefully modified line search condition, that incorporates curvature information of both objective and constraint functions, as appropriate and required to mitigate the Maratos effect. The condition that triggers the modified line search is adaptive and does not require knowledge of unknown quantities.
- To set the stage for our main theoretical results, we derive local convergence guarantees for the SQP method under the proximity and unit step size assumptions. We include this section for completeness and to compare with existing results in the literature, but most importantly to discuss the sampling requirements. Specifically, we utilize adaptive sampling strategies within the SQP methodology (applicable to both problems (1.2) and (1.3)) and, under reasonable assumptions, show that locally the primal-dual iterates converge to a first-order stationary point at a superlinear rate. These results are natural extensions of the local convergence results of existing deterministic SQP methodologies.
- Our main theoretical contributions are as follows. For problems (1.1) and (1.3), we prove that the iterates generated by the proposed modified line search SQP method, from any starting point, converge to a first-order stationary point (global convergence) and, moreover, eventually, the unit step size is always accepted and the algorithm achieves a superlinear convergence rate asymptotically (local convergence). Our methodology is simple and adaptive, makes adjustments as needed and detected over the course of the optimization, and has provable fast local convergence guarantees, without requiring knowledge of problem specific parameters or the proximity and unit step size assumptions.
- To tackle large-scale problems, we propose a practical inexact matrix-free variant of the modified line search SQP method that employs the minimum residual (MINRES) method [21] as a subroutine to compute a step. The method reduces the computational and storage costs and allows for inexact subproblem solutions. Under reasonable conditions and a sufficient number of MINRES iterations, the method enjoys an asymptotic linear rate of convergence with constant one half.
- Finally, we empirically demonstrate that our proposed method is competitive with a classical SQP method and other adaptations on the CUTEst test problems [11]. In the

stochastic setting, we compare the performance of our proposed method to a stochastic SQP methodology with an adaptive step size selection scheme [26] on constrained binary classification logistic regression problems using the LIBSVM datasets [17]. We illustrate the efficiency and robustness of the practical variant.

1.3 Paper Organization

We conclude this section by introducing some notation. In Section 2, we present a discrepancy in global and local convergence guarantees of line search SQP methods due to the Maratos effect, and examine various remedies. General assumptions and preliminaries are given in Section 3. Our proposed modified line search SQP method is presented in Section 4. In Section 5, we present local convergence guarantees. The main theoretical results of the proposed methodology are given in Section 6. In Section 7, we introduce and analyze a practical inexact variant for the large scale setting. Numerical results (deterministic and stochastic settings) are presented in Section 8, and concluding remarks are given in Section 9.

1.4 Notation

Let \mathbb{R} denote the set of real numbers, \mathbb{R}^n denote the set of n -dimensional real vectors, $\mathbb{R}^{m \times n}$ denote the set of m -by- n -dimensional real matrices, and $\mathbb{N} := \{0, 1, 2, \dots\}$ denote the set of natural numbers. Unless specified, all norms $\|\cdot\|$ are 2-norms. Let $f_k := f(x_k)$, $g_k := \nabla f(x_k)$, and $H_k := \nabla^2 f(x_k)$ denote the value, gradient and Hessian of the objective function, respectively, evaluated at x_k . Let $c_k := c(x_k)$, and $J_k := \nabla c(x_k)^T$ denote the Jacobian matrix, evaluated at x_k . Let $y \in \mathbb{R}^m$ denote the dual variables, and $w = [x^T \ y^T]^T \in \mathbb{R}^{n+m}$ a pair of primal-dual variables. Let $\mathcal{L}(x, y) := f(x) + y^T c(x)$ denote the Lagrangian function and $W(x, y) \in \mathbb{R}^{n \times n}$ denote the Hessian of the Lagrangian function, where $W_k := W(x_k, y_k) := H(x_k) + \sum_{i=1}^m y_{k,i} \nabla^2 c_i(x_k)$. Throughout the paper we use overbars to denote approximations. We use the sets $S_k^f, S_k^g, S_k^H \subseteq \{\omega_1, \omega_2, \dots\}$ to denote the samples used to compute function, gradient, and Hessian estimates, respectively,

$$\bar{f}_k = \frac{1}{|S_k^f|} \sum_{i \in S_k^f} f_i(x_k), \quad \bar{g}_k = \frac{1}{|S_k^g|} \sum_{i \in S_k^g} \nabla f_i(x_k), \quad \bar{H}_k = \frac{1}{|S_k^H|} \sum_{i \in S_k^H} \nabla^2 f_i(x_k), \quad (1.4)$$

where $f_i(x_k) = F(x_k, \omega_i)$, $\nabla f_i(x_k) = \nabla F(x_k, \omega_i)$, and $\nabla^2 f_i(x_k) = \nabla^2 F(x_k, \omega_i)$. These samples (ω_i) are drawn at random from the distribution P or uniformly from the set $\{\omega_1, \omega_2, \dots, \omega_N\}$ in the context of problems (1.2) or (1.3), respectively.

2 Analysis of Deterministic SQP Methods

In this section, we present the line search SQP methodology and its associated global and local convergence guarantees. The goal is to highlight an existing gap in the analysis of line

search (and other) SQP methods for solving deterministic equality constrained problems of the form (1.1) and to discuss several remedies. By gap in this setting we refer to the need for proximity and unit step size assumptions or modifications to classical and practical methodologies in order to prove strong global and fast local convergence guarantees. We make this explicit below.

We begin by introducing the foundation of the classical SQP methodology¹ for equality constrained optimization [51, Chapter 18]. SQP methods are iterative and solve a sequence of subproblems at each iteration in order to compute a step and update the iterate. Specifically, at each iteration $k \in \mathbb{N}$,

$$x_{k+1} \leftarrow x_k + \alpha_k d_k, \quad (2.1)$$

where $\alpha_k \in \mathbb{R}_{>0}$ is the step size and $d_k \in \mathbb{R}^n$ is the search direction defined as the (possibly inexact) solution of the following SQP subproblem,

$$\min_{d \in \mathbb{R}^n} f(x_k) + \nabla f(x_k)^T d + \frac{1}{2} d^T W_k d \quad \text{s.t.} \quad c(x_k) + \nabla c(x_k)^T d = 0, \quad (2.2)$$

where $W_k \in \mathbb{R}^{n \times n}$ is a matrix that captures curvature information of the objective and constraint functions. To the best of our knowledge, the first SQP method was proposed in [75], with $W_k = \nabla^2 \mathcal{L}(x_k, y_k)$, where $\mathcal{L}(x_k, y_k) = f(x_k) + y_k^T c(x_k)$ is the Lagrangian function and $y_k \in \mathbb{R}^m$ are the Lagrange multipliers. This method, as well as variants for which W_k is some economical and sufficiently accurate approximation of the Hessian of the Lagrangian, e.g., [41, 53, 57], are endowed with local superlinear convergence guarantees under the assumptions that the unit step size is employed, i.e., $\alpha_k = 1$ in (2.1), and that the starting point (x_0, y_0) is in the neighborhood of a first-order stationary point (x^*, y^*) , i.e., points for which

$$\nabla_x \mathcal{L}(x^*, y^*) = 0 \quad \text{and} \quad c(x^*) = 0. \quad (2.3)$$

Unfortunately, the neighborhood is defined by problem-specific unknown parameters, and employing a unit step size far from the solution may lead to undesirable behavior.

To overcome the two limitations and ensure global convergence, i.e., convergence to a first-order stationary point from any initial starting point, two main algorithmic paradigms have prevailed, line search [42, 58] and trust region [14, 71]. While these approaches ensure convergence from any starting point, local fast convergence can only be proven under the same proximity and unit step size conditions mentioned above. Consider a detour to the unconstrained setting. The iterate update form in this setting, (2.1)–(2.2), results in Newton steps. If the classical Armijo sufficient decrease line search condition with the

¹By *classical SQP methodology* we refer to line search SQP methods; e.g., [58], compute a step by solving (1.1), that utilize the ℓ_1 -merit function, update the merit parameter via dual variables estimates (or other mechanisms), and that impose a sufficient decrease condition. See e.g., [51, Chapter 18, Algorithm 18.3]

initial trial step size set to unity is employed, it is well established that under reasonable assumptions after sufficiently many iterations, the line search condition will be satisfied with the unit step size [56]. This is known as the “two-phase” property of Newton’s method. Perhaps surprisingly, this property does not extend to the general nonlinear equality constrained setting, and as such, there exists a gap between the globalizing phase and the fast local phase for classical SQP methods. That is, for some problems, the unit step size ($\alpha_k = 1$, in (2.1)) does not satisfy the sufficient decrease condition imposed by classical line search SQP methods irrespective of the proximity of the primal-dual iterate to a first-order stationary point. An example of such a problem is given in [44],

$$\min z_1^2 + z_2^2 \quad \text{s.t.} \quad (z_1 + 1)^2 + z_2^2 - 4 = 0, \quad (2.4)$$

for $(z_1, z_2) \in \{(z_1 + 1)^2 + z_2^2 = 4, z_1 \in (-1, 1)\}$. An intuitive explanation for the discrepancy is that unlike the unconstrained setting it is not always possible to simultaneously reduce both the constraint violation and the objective function value along a given search direction, regardless of its magnitude. As a result, there are problems for which even steps that make good progress towards a solution and keep the constraint violation reasonably small can be rejected by the line search SQP algorithm. This unfavorable phenomenon is known as the Maratos effect [51, Section 15.5].

To mitigate this issue, a sufficient condition is derived in [44], related to the curvature in the objective and constraint functions and a measure of progress, that allows for classical line search SQP methods to have provable “two-phase” guarantees. Rather than a remedy, the condition restricts the general problem class and highlights the gap in the global/local analysis of classical SQP methods. Our modified line search strategy draws inspiration from this condition. Below, we discuss some algorithmic approaches that mitigate the issue and ensure, after a sufficiently large number of iterations, the unit step length is accepted. This, in conjunction with the established global convergence, yields the asymptotic superlinear convergence rate.

Modified Merit Function. One remedy is to modify the classical line search SQP method that utilizes the ℓ_1 merit function and instead use the augmented Lagrangian function as the merit function [64]. By modifying the merit function in this way, the constrained problem can be formulated as an unconstrained problem with $n + m$ variables and the analysis for the line search Newton’s method, e.g., [56, Theorem 6.2.31], can be generalized to this setting. Several works have used similar strategies, e.g., Fletcher’s exact penalty function [59] and the (augmented) Lagrangian function [16, 37]. Modifying the merit function is a reasonable approach to fill the gap between local and global convergence analysis. However, the issue of these merit functions is their sensitivity to the merit parameter updating scheme.

Modified Subproblem and Update Rule. One can also modify the constraint approximation in the SQP subproblem (2.2) as a correction; see e.g., [22, 34, 46]. Specifi-


cally, after computing the solution to (2.2), an auxiliary subproblem is solved for which the affine approximation to the constraints is replaced by the quadratic approximation, $c_i(x_k) + \nabla c_i(x_k)^T d + \frac{1}{2} d^T \nabla^2 c_i(x_k) d = 0$, for $i = \{1, \dots, m\}$. The algorithm then searches along an arc defined by a combination of the solution of the original and auxiliary subproblems. Both line search [22, 34, 46] and trust region [31, 35] SQP variants have been developed that utilize this strategy. While this reformulation guarantees that after sufficiently many iterations the unit step size is accepted, this does not come for free as the cost of solving the auxiliary subproblem (even approximately) can be higher than that of solving (2.2).

Watchdog and Other Methods. The watchdog method [16], a nonmonotone method that do not require or enforce sufficient decrease at every iteration and allows for occasional aggressive (unit step size) steps, presents another alternative. A drawback of this method is the possibly expensive (in terms of function evaluations) restart procedure that is required when aggressive steps fail. Another drawback is the sensitivity to algorithmic parameters such as number of consecutive aggressive steps. Other approaches have either used nonmonotone line search, trust-region or filter techniques to address the Maratos effect [19, 20, 39, 68].

Our approach and remedy, presented in Section 4, is algorithmic. We maintain the same subproblem, merit function and merit parameter update, and carefully modify the sufficient decrease condition as required, and derive asymptotic local superlinear and global convergence guarantees, i.e., we design and analyze a line search SQP method for general equality constraints endowed with the “two-phase” behavior of Newton’s method. In contrast to preceding approaches, our approach strives to preserve the classical line search SQP methodology to the extent possible and only deviates from the classical method when the Maratos effect is detected. In such cases, a modified line search condition is employed to produce more suitable step sizes and guide the iterates away from undesired regions, while maintaining global convergence guarantees and adding fast local convergence.

We conclude this section with some empirical results on two handpicked examples that showcase the differences between the methods. We compare the classical line search SQP method (**SQP-L1**) [42], the **2nd-corr**, **Watchdog**, **SQP-AugLag** alternatives (discussed above), and our proposed method (**Our method**). Figure 1a illustrates the performance of the methods on the counter-example (2.4), starting from $(z_1, z_2) = (\sqrt{2} - 1, \sqrt{2})$, a point that is unfavorable for the **SQP-L1** method. As expected, **SQP-L1** produces small step sizes from this starting point, and takes several iterations to reach a point where the unit step size is eventually accepted. All other remedies, including our proposed method, despite traversing different trajectories, accept the unit step size from the second iteration onwards and converge much faster. Figure 1b illustrates the performance of the methods on a constrained version of the Rosenbrock problem ($c(z) = (z_1 + 2)^2 + (z_2 - 1)^2 - 9$), with starting point $(z_1, z_2) = (-1.1, 1)$. Notably, the **SQP-L1** and **SQP-AugLag** struggle on

this problem, and fail to make progress. On the other hand, **2nd-corr**, **Watchdog** and **Our method** accept the unit step size for the majority of the iterations, traverse different trajectories and exhibit fast convergence. The **2nd-corr** method performs particularly well on this problem, this is not surprising given the constraint is quadratic. Overall, our approach appears to be competitive with contemporary approaches.



(a) **Example 1:** Counter-example (2.4), starting point $(\sqrt{2} - 1, \sqrt{2})$.

(b) **Example 2:** $\min_{z_1, z_2} (1 - z_1)^2 + 100(z_2 - z_1^2)^2$ s.t. $(z_1 + 2)^2 + (z_2 - 1)^2 = 9$, starting point $(-1.1, 1)$.

Figure 1: Comparison of classical line search SQP method (**SQP-L1**), remedies to address Maratos effect (**2nd-corr**, **SQP-AugLag**, **Watchdog**) and our proposed SQP method (**Our method**). Trajectory plots (contours in grayscale, constraints in orange), step size, feasibility and stationarity. (Trajectories outside region excluded.)

3 General Assumptions and Preliminaries

In this section, we first present the general assumptions that are applicable to all the problems considered in this paper and then state and discuss some problem-specific assumptions. Additional technical assumptions are stated, as required, in Sections 5 and 6. Following each set of assumptions, we introduce lemmas that are used in Sections 5 and 6.

We make the following assumption about problems (1.1), (1.2), and (1.3).

Assumption 3.1. *Let $\mathcal{X} \subseteq \mathbb{R}^n$ be an open convex set containing the iterates $\{x_k\}$ and first trial points $\{x_k + d_k\}$. The objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is twice continuously differentiable and bounded below over \mathcal{X} with $|f(x)| \leq \kappa_f$, the gradients $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ are Lipschitz continuous with constant L_1 and bounded over \mathcal{X} with $\|g(x)\|_2 \leq \kappa_g$, and the Hessians $H : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ are Lipschitz continuous with constant L_2 and bounded over \mathcal{X} with $\|H(x)\|_2 \leq \kappa_H$. The constraint function $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$ (where $m \leq n$) is twice continuously differentiable and bounded over \mathcal{X} with $\|c(x)\|_2 \leq \kappa_c$ and the associated Jacobian matrix $\nabla c^T : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$ is Lipschitz continuous with constant Γ_1 and bounded over \mathcal{X} . The*

Hessian of $c(\cdot)$ is Lipschitz continuous with constant Γ_2 . Finally, we assume the singular values of ∇c^T are bounded away from zero over \mathcal{X} such that $\|(J_k J_k^T)^{-1} J_k\|_2 \leq \kappa_{J^\dagger}$.

We make a few remarks about Assumption 3.1. Most of the statements in Assumption 3.1 constitute standard smoothness assumptions for deterministic equality constrained nonlinear optimization problems of the form (1.1); see e.g., [14, 73]. With regards to (1.2) and (1.3), smoothness assumptions similar to those in Assumption 3.1 are common in the unconstrained stochastic setting; see e.g., [12, 66, 67]. When considering problems (1.2) and (1.3), the iterates generated by the proposed algorithms are stochastic, and admittedly, it is not ideal to assume that the gradient and Hessian of the objective function are bounded over the set of iterates and trial iterates. That said, we believe this assumption is reasonable due to the fact that any well designed algorithm should be driven towards the deterministic feasible region. Moreover, we note that similar assumptions are common in the constrained stochastic optimization literature; see e.g., [6, Assumption 3.2], [48, Assumption 1]. Finally, the assumption on the boundedness of the singular values of the Jacobian of the constraints is satisfied by the common linear independence constraint qualification (LICQ) condition [51].

The next assumption pertains to the Lipschitz continuity of the gradient and Hessian of the Lagrangian function near a first-order primal-dual stationary point.

Assumption 3.2. *For some neighborhood of $w^* = (x^{*T}, y^{*T})^T$ with radius $r \in \mathbb{R}_{>0}$, where x^*, y^* is a first-order primal-dual stationary point, i.e., points satisfying (2.3), the gradient and Hessian of the Lagrangian function, $\mathcal{L}(x, y) = f(x) + y^T c(x)$, are Lipschitz continuous with constants $L_{\nabla \mathcal{L}}$ and L_W , respectively.*

Assumption 3.2 is a consequence of Assumption 3.1, and can be regarded as an analogue to the Lipschitz continuity of the Hessian of the objective function assumption in the context of unconstrained optimization and Newton’s method [51, Theorem 3.5]. Assumptions 3.1 and 3.2 are required for all problems (1.1), (1.2) and (1.3).

We require an additional assumption common in the stochastic optimization literature in order to establish theoretical results for the fully stochastic regime and problem (1.2); see e.g., [12, 67]. Specifically, unbiased gradient estimates with bounded variance, and Hessians estimates with bounded variance.

Assumption 3.3. *There exists $\sigma_g, \sigma_H \in \mathbb{R}_{>0}$ such that for all $x \in \mathcal{X}$ and $i = 1, 2, \dots$, $\mathbb{E}[\nabla f_i(x)] = \nabla f(x)$, $\text{Var}(\nabla f_i(x)) \leq \sigma_g^2$ and $\|\mathbb{E}[(\nabla^2 f_i(x) - \nabla^2 f(x))^2]\|_2 \leq \sigma_H^2$.*

Assumption 3.3 pertains only to the derivatives of the objective function of (1.2) and not the function values. To this end, we make the following disclaimer. In Section 5, we present local convergence results (under proximity and unit step size assumptions) for all problem settings, however, our main theoretical results (Section 6) are only for problems (1.1) and (1.3). That said, at the expense of additional complexity in the algorithm and analysis, one could extend the results to the fully stochastic setting (1.2) by utilizing approaches for

adaptive step sizes developed for the unconstrained stochastic setting [4, 15, 54]. This is not in the scope of this paper.

In the context of (1.3), we replace Assumption 3.3 with the following assumption.

Assumption 3.4. *There exists $\kappa_0, \mu_0, \kappa_1, \mu_1, \kappa_2, \mu_2 \in \mathbb{R}_{>0}$ such that for all $x \in \mathcal{X}$ and $i = 1, \dots, N$, $|f_i(x)| \leq \kappa_0 + \mu_0|f(x)|$, $\|\nabla f_i(x)\|_2 \leq \kappa_1 + \mu_1\|\nabla f(x)\|_2$, and $\|\nabla^2 f_i(x)\|_2 \leq \kappa_2 + \mu_2\|\nabla^2 f(x)\|_2$.*

The next lemma characterizes the errors in the approximations.

Lemma 3.5. *Suppose Assumptions 3.1 and 3.4 hold. Then, for $k \in \mathbb{N}$,*

$$|f_k - \bar{f}_k| \leq 2 \left(\frac{N - |S_k^f|}{N} \right) (\kappa_0 + \mu_0 \kappa_f) := \epsilon_{f_k}, \quad (3.1)$$

$$\|g_k - \bar{g}_k\|_2 \leq 2 \left(\frac{N - |S_k^g|}{N} \right) (\kappa_1 + \mu_1 \kappa_g) := \epsilon_{g_k}, \quad (3.2)$$

$$\|H_k - \bar{H}_k\|_2 \leq 2 \left(\frac{N - |S_k^H|}{N} \right) (\kappa_2 + \mu_2 \kappa_H) := \epsilon_{H_k}. \quad (3.3)$$

Proof. For brevity, we defer the proof to Appendix A. □

Lemma 3.5 shows an inverse relationship between the approximation errors and the sample sizes. In this setting, the sample sizes are capped at N , and when $|S_k^f| = |S_k^g| = |S_k^H| = N$, it follows that $\epsilon_{f_k} = \epsilon_{g_k} = \epsilon_{H_k} = 0$.

4 A Modified Line Search SQP Method

In this section, we present in detail our proposed second-order modified line search SQP method. The goal is to develop a practical algorithm that exhibits the “two-phase” behavior of Newton’s method. That is, global convergence to first-order critical points from any starting point, and fast superlinear local convergence after sufficiently many iterations without imposing proximity or unit step size conditions. For simplicity, we present a deterministic variant of the method for solving (1.1), and discuss the necessary algorithmic changes required in order to tackle problems (1.2) and (1.3).

Our proposed modified line search SQP method is iterative (2.1). The search direction is computed by solving an SQP subproblem (2.2), and then a step size is selected guided by a merit function and a sufficient decrease condition. At every iteration $k \in \mathbb{N}$, a primal-dual search direction $(d_k, \delta_k) \in \mathbb{R}^n \times \mathbb{R}^m$ is computed by solving (possibly inexactly) the subproblem given in (2.2). Under Assumption 3.1, this is equivalent to solving the following linear system of equations

$$M_k \begin{bmatrix} d_k \\ \delta_k \end{bmatrix} = \begin{bmatrix} W_k & J_k^T \\ J_k & 0 \end{bmatrix} \begin{bmatrix} d_k \\ \delta_k \end{bmatrix} = - \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix}, \quad (4.1)$$

where $W_k \in \mathbb{R}^{n \times n}$ is defined as the Hessian of the Lagrangian. We make the following standard assumptions about the matrices M_k and W_k that appear in (4.1).

Assumption 4.1. *The sequence $\{W_k\}$ is bounded in norm by $\kappa_W \in \mathbb{R}_{>0}$ over $x \in \mathcal{X}$. In addition, there exists a constant $\zeta \in \mathbb{R}_{>0}$ such that, for all $k \in \mathbb{N}$, $u^T W_k u \geq \zeta \|u\|_2^2$ for all $u \in \mathbb{R}^n$ such that $J_k u = 0$.*

Assumption 4.2. *There are $\Lambda, \kappa_M \in \mathbb{R}_{>0}$ such that $\|M_k^{-1}\| \leq \Lambda$, $\|M_k\| \leq \kappa_M$.*

Under Assumptions 3.1 and 4.1, the linear system (4.1) has a unique solution. Moreover, in the deterministic setting and under these assumptions, if $d_k = 0$, then (x_k, y_k) satisfy (2.3). Assumption 4.1 is common in the equality constrained optimization literature [53, 57]. The existence of $\Lambda \in \mathbb{R}_{>0}$ in Assumption 4.2 is guaranteed by Assumptions 3.1 (singular values of Jacobian bounded from zero) and 4.1; see [77, Lemma 16.1]. The existence of $\kappa_M \in \mathbb{R}_{>0}$ is guaranteed since both W_k and J_k are bounded.

Given a search direction $d_k \in \mathbb{R}^n$, the method proceeds to select a positive step size to update the iterate (2.1). This selection is guided by a merit function. As is common with popular line search SQP methods, we consider the ℓ_1 -merit function

$$\phi(x, \tau) := \tau f(x) + \|c(x)\|_1, \quad (4.2)$$

where $\phi : \mathbb{R}^n \times \mathbb{R}_{>0} \rightarrow \mathbb{R}$ and $\tau \in \mathbb{R}_{>0}$ is the merit parameter that is chosen adaptively as the optimization progresses and whose goal is to balance the two possibly competing goals of minimizing the objective function and satisfying the constraints. We make use of a local model of the merit function $l : \mathbb{R}^n \times \mathbb{R}_{>0} \times \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, defined as

$$l(x, \tau, \nabla f(x), d) := \tau(f(x) + \nabla f(x)^T d) + \|c(x) + \nabla c(x)^T d\|_1. \quad (4.3)$$

In particular, a critical quantity in our algorithm is the reduction (across a step for which $c(x) + \nabla c(x)^T d = 0$) in the local model of the merit function, $\Delta l : \mathbb{R}^n \times \mathbb{R}_{>0} \times \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, defined as

$$\Delta l(x, \tau, \nabla f(x), d) := l(x, \tau, \nabla f(x), 0) - l(x, \tau, \nabla f(x), d) = -\tau \nabla f(x)^T d + \|c(x)\|_1. \quad (4.4)$$

After computing a step and before updating the iterate (taking the step) the merit parameter τ_k is updated. We follow standard approaches [6, 13]. For some $\sigma, \epsilon_\tau \in (0, 1)$, the merit parameter is updated as follows,

$$\tau_k^{trial} \leftarrow \begin{cases} \infty & \text{if } g_k^T d_k + \max\{d_k^T W_k d_k, 0\} \leq 0 \\ \frac{(1-\sigma)\|c_k\|_1}{g_k^T d_k + \max\{d_k^T W_k d_k, 0\}} & \text{otherwise,} \end{cases} \quad (4.5)$$

followed by

$$\tau_k \leftarrow \begin{cases} \tau_{k-1} & \text{if } \bar{\tau}_{k-1} \leq \tau_k^{trial} \\ (1 - \epsilon_\tau) \tau_k^{trial} & \text{otherwise.} \end{cases} \quad (4.6)$$

In both cases, $\tau_k \leq \tau_k^{trial}$ and $\Delta l(x_k, \tau_k, g_k, d_k) \geq \tau_k \max\{d_k^T W_k d_k, 0\} + \sigma \|c_k\|_1$. This inequality is encouraging as it guarantees $\Delta l(x_k, \tau_k, g_k, d_k)$ is always nonnegative, and is zero only when $\|c_k\|_1 = \|d_k\| = 0$, points that satisfy first-order stationarity (2.3).

Our proposed algorithm computes a step size by ensuring sufficient decrease on the merit function akin to that used in classical line search SQP algorithms, e.g., [42],

$$\phi(x_k + \alpha_k d_k, \tau_k) \leq \phi(x_k, \tau_k) - \eta \alpha_k \Delta l(x_k, \tau_k, g_k, d_k), \quad (4.7)$$

where $\eta \in (0, 1)$, until a certain condition is satisfied. The distinctive feature of our approach is a carefully crafted modified line search condition that is employed as required over the course of the optimization. The intuition of the condition and modification is to detect and correct for the Maratos effect.

As revealed by classical SQP analysis, under the proximity (iterates to the optimal solution) and unit step size assumptions, the method enjoys a superlinear rate of convergence; see e.g., [41, 53, 57]. In Section 5, we prove a slightly different result and show that there exists a threshold γ_d such that if we run the classical line search SQP method until $\|d_j\| \leq \gamma_d$ for some iteration $j \in \mathbb{N}$ and simply set $\alpha_k = 1$ for $k \geq j$, superlinear convergence is retained. However, γ_d depends on unknown parameters (see Lemma 5.9), and it is therefore impractical to implement such rule. We propose two modifications to this naive and impractical approach. First, we consider an adaptive threshold parameter (γ_k instead of γ_d), that is updated (reduced), as necessary, over the course of the optimization. Second, since it is not known if γ_k is sufficiently small, and employing a unit step size for all iterates for which $\|d_k\| \leq \gamma_k$ may be too aggressive, we propose a carefully crafted sufficient decrease condition that incorporates second-order information of the objective and constraint functions, i.e.,

$$\begin{aligned} \phi(x_k + \alpha_k d_k, \tau_k) &\leq \phi(x_k, \tau_k) - \eta \alpha_k \Delta l(x_k, \tau_k, g_k, d_k) \\ &\quad + \frac{1}{2} \alpha_k^2 \tau_k d_k^T H_k d_k + \frac{1}{2} \alpha_k^2 \sum_{i=1}^m |d_k^T \nabla^2 c_i(x) d_k|. \end{aligned} \quad (4.8)$$

The final component of the algorithm is the fact that (4.8) is satisfied with $\alpha_k = 1$ when $\|d_k\|$ is sufficiently small (Lemma 6.4). As such, our proposed adaptive mechanism reduces γ_k when (4.8) is triggered to ensure that the threshold parameter is eventually, and as needed, sufficiently small. We note that this novel step size condition only needs to be triggered when, for small search directions, the classical line search fails to accept the unit step size. Overall, our proposed approach only deviates from the classical line search approach when the occurrence of the Maratos effect is detected. The full algorithm is given in Algorithm 1.

Remark 4.3. *We make a few remarks about Algorithm 1.*

- **Adaptive Condition (Line 5) and Behavior of γ_k :** *The purpose of the condition in Line 5 is to ensure that $\|d_k\|$ is sufficiently small before any modifications are employed.*

Algorithm 1 Modified Line Search SQP

Require: $x_0 \in \mathbb{R}^n$; $y_0 \in \mathbb{R}^m$; $\tau_{-1} \in \mathbb{R}_{>0}$; $\gamma_0 \in \mathbb{R}_{>0}$; $\nu_\alpha, \nu_\gamma, \epsilon_\tau, \sigma \in (0, 1)$.

```

1: for all  $k \in \mathbb{N}$  do
2:   Compute  $f_k, g_k, H_k$ .
3:   Compute  $(d_k, \delta_k)$  as the solution of (4.10).
4:   Set  $\tau_k^{trial}$  and  $\tau_k$  by (4.5) and (4.6).
5:   if  $\|d_k\| > \gamma_k$  then
6:     until (4.7) is achieved, set  $\alpha_k = \nu_\alpha \alpha_k$ .
7:   else
8:     if (4.7) is satisfied with  $\alpha_k = 1$  then set  $\alpha_k = 1$ .
9:     else
10:      until (4.8) is achieved, set  $\alpha_k = \nu_\alpha \alpha_k$ .
11:      Set  $\gamma_{k+1} = \nu_\gamma \gamma_k$ .
12:     end if
13:   end if
14:   Set  $x_{k+1} \leftarrow x_k + \alpha_k d_k$  and  $y_{k+1} \leftarrow y_k + \delta_k$ .
15: end for

```

The else condition is guaranteed to eventually be satisfied (see Lemma 6.3) since γ_k remains fixed while $\|d_k\| > \gamma_k$. When $\|d_k\| \leq \gamma_k$, the method attempts to detect if the Maratos effect is present by checking whether $\alpha_k = 1$ satisfies (4.7). If $\alpha_k = 1$ satisfies (4.7), the classical line search condition is good enough and the step size is set to one. Otherwise, we either have Marato's effect or we are not yet sufficiently close to the optimal solution (γ_k not sufficiently small). In this case, we consider the modified line search condition (4.8), and reduce γ_k . The sequence γ_k , which is non-increasing, is fundamental to our algorithm. This sequence reduces by a factor of ν_γ upon each decrement. Importantly, while γ_k experiences a decrease at a linear rate at most, we prove that $\|d_k\|$ undergoes a decrease at a superlinear rate after a sufficient number of iterations (see Lemma 6.3), leading to the consistent engagement of Line 7. If $\alpha_k = 1$ does not satisfy (4.7) infinitely often, the algorithm ends up with $\gamma_k \rightarrow 0$.

- **Marato's Condition:** In [44, Proposition 8] it is mentioned that if

$$\tau_k d_k^T H_k d_k + \sum_{i=1}^m |d_k^T \nabla^2 c_i(x) d_k| \leq 2\sigma_m \eta \Delta l(x_k, \tau_k, g_k, d_k) \quad (4.9)$$

is satisfied for some $\sigma_m \in \mathbb{R}_{>0}$ when $\|d_k\|$ is sufficiently small, then the Maratos effect is not present. Thus, one option is to only modify the algorithm when (4.9) is not satisfied. We deviate from this strategy in two ways. First, we directly use an adaptive condition on the search direction for simplicity. Second, we only apply the modification

when $\alpha_k = 1$ does not satisfy condition (4.7) to reduce the frequency of employment of the modified line search condition.

- **Comparison to Classical Sufficient Decrease Condition:** Compared to classical sufficient decrease (on the merit function) conditions employed by SQP methods (4.7), condition (4.8) has two additional terms related to second-order information of the objective and constraint functions. As we show (see Lemma 6.4), $\alpha_k = 1$ satisfies (4.8) when $\|d_k\|$ is sufficiently small. Hence, we reduce γ_k by a factor of ν_γ to ensure asymptotic unit step size acceptance.
- **Comparison to Existing Remedies:** Similar to the Watchdog method [16], our algorithm adopts a nonmonotone line search condition when $\tau_k d_k^T H_k d_k + \sum_{i=1}^m |d_k^T \nabla^2 c_i(x) d_k| \gg 0$. That said, the nonmonotonicity is controlled in our approach, and the step size is selected by condition (4.8), and thus is more conservative than the unit step size that is employed by the Watchdog method [16].
- **Comparison to Adaptive Step Size Scheme:** Alternative step size schemes that do not require function evaluations have been proposed for the constrained stochastic setting [6, 7]. Such schemes are particularly effective in the stochastic setting, however, we refrain from utilizing such approaches and retain a traditional line search approach for two main reasons. First, for these adaptive schemes, asymptotic unit step sizes can only be guaranteed for well-conditioned problems. Second, the necessity of the Lipschitz constants (or an estimate) in these algorithm poses challenges in the implementation since these constants are usually unknown.
- **Parameters:** Algorithm 1 has several parameters. Most importantly, γ_0 , the initial value of the threshold parameter. We recommend and set $\gamma_0 \leftarrow \bar{\gamma} \|d_0\|$ for some $\bar{\gamma} \in (0, 1)$ to make sure that Line 5 is triggered in the first iteration. The other parameters (τ_{-1} , η and ν_α) are universal to all methods considered, with the additional parameter ν_γ , the reduction factor in the threshold parameter.

4.1 The Stochastic Setting

The development of adaptive line search SQP methods for the stochastic and finite sum settings is not as straightforward. Several modifications are necessary to ensure convergence (global and fast local) as well as practicality. We focus on the finite sum setting (1.3) in this section. In this context, we follow the same algorithmic structure as Algorithm 1 but replace the exact function (and derivatives) with their stochastic counterparts.

At every iteration $k \in \mathbb{N}$, we first select subsets of all the data $S_k^f, S_k^g, S_k^H \subseteq \{\omega_1, \omega_2, \dots, \omega_N\}$ (specific rules given in Section 6), compute the estimates of function value, gradient, and Hessian via (1.4) as $\bar{f}_k, \bar{g}_k, \bar{H}_k$ (replace f_k, g_k, H_k Line 2 of Algorithm 1), and set $\bar{W}_k = \bar{H}_k + \sum_{i=1}^m y_{k,i} \nabla^2 c_i(x_k)$. Then, we obtain the search direction $(\bar{d}_k, \bar{\delta}_k) \in \mathbb{R}^n \times \mathbb{R}^m$

by solving a linear system of the form

$$\bar{M}_k \begin{bmatrix} \bar{d}_k \\ \bar{\delta}_k \end{bmatrix} = \begin{bmatrix} \bar{W}_k & J_k^T \\ J_k & 0 \end{bmatrix} \begin{bmatrix} \bar{d}_k \\ \bar{\delta}_k \end{bmatrix} = - \begin{bmatrix} \bar{g}_k + J_k^T \bar{y}_k \\ c_k \end{bmatrix}. \quad (4.10)$$

To ensure the uniqueness of solution to (4.10) and other theoretical properties (similar to those discussed above for the deterministic), we make the following assumption.

Assumption 4.4. *Assumptions 4.1 and 4.2 hold (with the same constants) for the matrices \bar{W}_k and \bar{M}_k .*

Further, the update rule for the merit parameter ($\bar{\tau}_k$) in the finite sum setting is similar to (4.5)–(4.6), with approximations replacing all deterministic quantities. The classical and modified line search conditions (4.7) and (4.8) are relaxed as follows

$$\bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_k) \leq \bar{\phi}(x_k, \bar{\tau}_k) - \eta \bar{\alpha}_k \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) + \epsilon_{A_k}, \quad (4.11)$$

$$\begin{aligned} \bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_k) &\leq \bar{\phi}(x_k, \bar{\tau}_k) - \eta \bar{\alpha}_k \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) \\ &\quad + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{d}_k^T H_k \bar{d}_k + \frac{1}{2} \bar{\alpha}_k^2 \sum_{i=1}^m |\bar{d}_k^T \nabla^2 c_i(x) \bar{d}_k| + \epsilon_{A_k}, \end{aligned} \quad (4.12)$$

respectively, where $\bar{\phi}(x_k, \bar{\tau}_k) = \bar{\tau}_k \bar{f}_k + \|c_k\|_1$ and $\epsilon_{A_k} \in \mathbb{R}_{\geq 0}$ is a relaxation term added to account for the noise. The relaxation term is of paramount importance to both the theoretical analysis as well as the practical implementation. This term is proportional to the accuracy of the function and gradient approximations. The explicit expression for ϵ_{A_k} is given in Section 6 (Lemma 6.3). Finally, the iterates are updated via

$$x_{k+1} = x_k + \bar{\alpha}_k \bar{d}_k, \quad \text{and} \quad y_{k+1} = y_k + \bar{\delta}_k. \quad (4.13)$$

5 Preliminary Local Convergence Guarantees

Before we present the theoretical convergence guarantees for Algorithm 1 (Section 6), in this section, we set the stage by formalizing the superlinear convergence results for a general SQP method under the proximity and unit step size assumptions. We consider both the deterministic and stochastic settings, derive neighborhood conditions, and provide fast local convergence guarantees. In contrast to the conventional neighborhood proximity condition, which assumes $\|w_0 - w^*\|$ is sufficiently small [41, 53, 57], we propose an alternative explicit condition based on the norm of the search direction $\|d_k\|$ (or $\|\bar{d}_k\|$). This modification is motivated by the fact that it is computable within an algorithm. We show that under reasonable assumptions, when $\|d_k\|$ (or $\|\bar{d}_k\|$) is below a certain threshold (depends on unknown problem-specific parameters) the conventional neighborhood condition is satisfied. In Section 6, we analyze the adaptive algorithm presented in Section 4 that does not require knowledge of the unknown parameters.

We make the following assumption throughout this section.

Assumption 5.1. *There exists $\kappa_d, \mu_M \in \mathbb{R}_{>0}$, such that for all $k \in \{j \in \mathbb{N} \mid \|d_j\| \leq \kappa_d\}$, it follows that $\frac{\mu_M}{2} \|w_k - w^*\| \leq \left\| \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\|$, where d_k is the solution of (4.1) and $w_k = (x_k^T, y_k^T)^T$ is generated by Algorithm 1.*

Remark 5.2. *When (x_k, y_k) satisfies $\|\nabla \mathcal{L}(x_k, y_k)\| \leq \Lambda^{-1} \kappa_d$, it follows that $\|d_k\| \leq \Lambda \|\nabla \mathcal{L}(x_k, y_k)\| \leq \kappa_d$. Thus, Assumption 5.1 is satisfied if $\mathcal{L}(x, y)$ is $\frac{\mu_M}{2}$ -strongly convex with respect to w . Consequently, in this scenario, the iterates converge to a second-order stationary point. We note that this assumption is not needed for all iterates; instead, it is required only after running the algorithm for a sufficiently large number of iterations, such that the iterates enter a locally strongly convex regime.*

5.1 Deterministic Problem (1.1)

We present local analysis for the SQP method for (1.1) first.

Theorem 5.3. *Suppose Assumptions 3.1, 3.2, 4.1, 4.2, and 5.1 hold. Let $w_k = [x_k^T, y_k^T]^T$, $x_{k+1} = x_k + d_k$ and $y_{k+1} = y_k + \delta_k$, where (d_k, δ_k) are computed via (4.1). Then, for $k \in \mathbb{N}$, $\|w_{k+1} - w^*\| \leq \frac{\Lambda L_W}{2} \|w_k - w^*\|^2$. If the starting point satisfies $\|d_0\| \leq \min \left\{ \frac{\mu_M}{2\kappa_M(1+2\kappa_{j^+}\kappa_W)\Lambda L_W}, \kappa_d \right\}$, then, $\|w_k - w^*\| \rightarrow 0$ at a Q -quadratic rate.*

Proof. Since a more general result is proven in Theorem 5.6, we defer the proof of this theorem to Appendix B. \square

Note that similar local quadratic convergence results for the deterministic SQP method are established in [23, Theorem 15.2.1] and [51, 75].

5.2 Stochastic Problem (1.2)

At each iteration $k \in \mathbb{N}$, given $S_k^g, S_k^H \subseteq \{1, 2, \dots\}$, we compute $(\bar{d}_k, \bar{\delta}_k)$ via (4.10) and update the iterates via (4.13) with $\bar{\alpha}_k = 1$. Since the step size is predefined (unit), function estimates are not required. Lemma 5.4 ensures that when $\|\bar{d}_k\|$ is sufficiently small, the neighborhood condition is satisfied, and is important in establishing the superlinear convergence result.

Lemma 5.4. *Suppose Assumptions 3.1, 3.2, 3.4, 4.4, and 5.1 hold. For any $\kappa_w \in \mathbb{R}_{>0}$ and $k \in \mathbb{N}$, suppose that $\|g_k - \bar{g}_k\| \leq \min \left\{ \frac{\mu_M \kappa_w}{4}, \frac{\kappa_d}{2\Lambda} \right\}$. When $\|\bar{d}_k\| \leq \min \left\{ \frac{\mu_M \kappa_w}{4\kappa_M(1+2\kappa_{j^+}\kappa_W)}, \frac{\kappa_d}{2} \right\}$, it follows that $\|w_k - w^*\| \leq \kappa_w$.*

Proof. We use the orthogonal decomposition of the search direction $\bar{d}_k \in \mathbb{R}^n$, i.e.,

$$\bar{d}_k = \bar{u}_k + v_k, \quad \bar{u}_k \in \text{Null}(J_k), \quad v_k \in \text{Range}(J_k^T) \quad \text{for all } k \in \mathbb{N}. \quad (5.1)$$

Let Z_k be an orthogonal basis for the null space of J_k which under Assumption 3.1 is a matrix in $\mathbb{R}^{n \times (n-m)}$. By (4.10), it follows that

$$v_k = -J_k^T (J_k J_k^T)^{-1} c_k, \quad \bar{u}_k = -Z_k (Z_k^T \bar{W}_k Z_k)^{-1} Z_k^T (\bar{g}_k + J_k^T y_k + \bar{W}_k v_k), \quad (5.2)$$

$$\bar{\delta}_k = -(J_k J_k^T)^{-1} J_k (\bar{g}_k + J_k^T y_k + \bar{W}_k \bar{d}_k). \quad (5.3)$$

By (5.1), (5.2), (5.3), the Cauchy-Schwartz inequality, and Assumption 3.1,

$$\begin{aligned} \|\bar{\delta}_k\| &\leq \|(J_k J_k^T)^{-1} J_k\| \|\bar{g}_k + J_k^T y_k + \bar{W}_k v_k + \bar{W}_k \bar{u}_k\| \\ &\leq \kappa_{J^\dagger} (\|\bar{g}_k + J_k^T y_k + \bar{W}_k v_k\| + \|\bar{W}_k \bar{u}_k\|) \\ &\leq \kappa_{J^\dagger} (\|Z_k (Z_k^T \bar{W}_k Z_k)^{-1} Z_k^T \bar{u}_k\| + \|\bar{W}_k \bar{u}_k\|) \\ &\leq 2\kappa_{J^\dagger} \kappa_W \|\bar{d}_k\|. \end{aligned} \quad (5.4)$$

When $\|\bar{d}_k\| \leq \frac{\kappa_d}{2}$, the condition in Assumption 5.1 is satisfied since it follows by Assumptions 4.4 that

$$\|d_k\| \leq \|\bar{d}_k\| + \|d_k - \bar{d}_k\| \leq \|\bar{d}_k\| + \bar{M}_k^{-1} \begin{bmatrix} g_k - \bar{g}_k \\ 0 \end{bmatrix} \leq \|\bar{d}_k\| + \Lambda \epsilon_{g_k} \leq \kappa_d.$$

By Assumptions 5.1 and the gradient error condition,

$$\begin{aligned} \|w_k - w^*\| &\leq 2\mu_M^{-1} \left\| \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\| \\ &\leq 2\mu_M^{-1} \left\| \begin{bmatrix} \bar{g}_k + J_k^T y_k \\ c_k \end{bmatrix} + \begin{bmatrix} g_k - \bar{g}_k \\ 0 \end{bmatrix} \right\| \\ &\leq 2\mu_M^{-1} \left\| \bar{M}_k \bar{M}_k^{-1} \begin{bmatrix} \bar{g}_k + J_k^T y_k \\ c_k \end{bmatrix} \right\| + 2\mu_M^{-1} \epsilon_{g_k} \\ &\leq 2\mu_M^{-1} \kappa_M \left\| \begin{bmatrix} \bar{d}_k \\ \bar{\delta}_k \end{bmatrix} \right\| + \frac{\kappa_w}{2} \\ &\leq 2\mu_M^{-1} \kappa_M (1 + 2\kappa_{J^\dagger} \kappa_W) \|\bar{d}_k\| + \frac{\kappa_w}{2} \leq \kappa_w. \end{aligned}$$

□

Our goal is to prove that in expectation $\|w_k - w^*\| \rightarrow 0$ at a superlinear rate. If there is a deterministic uniform bound for $\|w_k - w^*\|$, the result follows directly. However, in the general stochastic setting, this assumption does not hold. Instead, we introduce an assumption on the second moment of the distance of the iterates to w^* . This assumption is also made in the stochastic unconstrained setting [10].

Assumption 5.5. *There exists a constant $\iota \in \mathbb{R}_{>0}$ such that for all $k \in \mathbb{N}$, $\mathbb{E}[\|w_k - w^*\|^2] \leq \iota (\mathbb{E}[\|w_k - w^*\|])^2$.*

The local superlinear convergence guarantees, under the proximity and unit step size assumptions, of an adaptive sampling SQP method for (1.2) is given below.

Theorem 5.6. *Suppose Assumptions 3.1, 3.2, 3.3, 4.4, 5.1, 5.5 hold. Further, suppose that for all $k \in \mathbb{N}$:*

$$(i) \quad \|g_k - \bar{g}_k\| \leq \min \left\{ \frac{\mu_M}{12\iota L_W \Lambda}, \frac{\kappa_d}{2\Lambda} \right\};$$

$$(ii) \quad |S_k^g| \geq |S_0^g| \beta_k^k, \text{ with } |S_0^g| \geq (6\sigma_g \iota L_W \Lambda^2)^2, \beta_k > \beta_{k-1}, \beta_k \rightarrow \infty, \text{ and } \beta_1 > 1;$$

$$(iii) \quad |S_k^H| > |S_{k-1}^H|, \text{ with } \lim_{k \rightarrow \infty} |S_k^H| = \infty, \text{ and } |S_0^H| \geq (4\sigma_H \Lambda)^2.$$

Then, if the starting point satisfies $\|\bar{d}_0\| \leq \min \left\{ \frac{\mu_M}{12\iota \Lambda L_W \kappa_M (1+2\kappa_j + \kappa_W)}, \frac{\kappa_d}{2} \right\}$, it follows that $\mathbb{E}[\|w_k - w^*\|] \rightarrow 0$ at an R -superlinear rate, i.e., there exists a positive sequence $\{\xi_k\}$ such that $\mathbb{E}[\|w_k - w^*\|] \leq \xi_k$ and $\frac{\xi_{k+1}}{\xi_k} \rightarrow 0$.

Proof. It follows by Lemma 5.4 and condition (i) that $\|w_0 - w^*\| \leq \frac{1}{3\iota L_W \Lambda}$. The remainder of the proof is similar to [10, Theorem 2.4] if we view the Lagrangian function as the objective function. Thus, for brevity we omit the rest of the proof. \square

Remark 5.7. *Theorem 5.6 is a stochastic analogue of Theorem 5.3. The sample growth rates for the gradient and Hessian estimates match those derived for local superlinear guarantees in the unconstrained setting [10], i.e., the gradient sample size is required to grow at a rate faster than geometric and the Hessian sample size is required to grow monotonically. Condition (i) is required due to the usage of the $\|\bar{d}_k\|$ as neighborhood proximity condition. We acknowledge that this condition is relatively strong and is one of the main reasons we focus on problem (1.3) in Section 6.*

5.3 Finite-sum Problem (1.3)

In this subsection, we present local analysis for the SQP method (iterates updated via (4.13)) for solving (1.3). We start with a preliminary lemma that bounds the optimality gap between subsequent iterates.

Lemma 5.8. *Suppose that Assumptions 3.1, 3.2, 3.4, and 4.4 hold. Then, for $k \in \mathbb{N}$,*

$$\|w_{k+1} - w^*\| \leq \frac{\Lambda L_W}{2} \|w_k - w^*\|^2 + \Lambda \epsilon_{H_k} \|w_k - w^*\| + \Lambda \epsilon_{g_k}.$$

Proof. The proof is similar to [10, Lemma 2.2] if we remove the expectation, view the Lagrangian function as the objective function, and view the primal-dual variables as the decision variables. Thus, for brevity we omit the proof. \square

We now present the local fast convergence result (analogous to Theorem 5.6).

Theorem 5.9. *Suppose Assumptions 3.1, 3.2, 3.4, 4.4, and 5.1 hold. Further, suppose that for all $k \in \mathbb{N}$:*

- (i) $|S_k^g| \geq \left(1 - \frac{\min\{\mu_M \Lambda, 6\Lambda L_W \kappa_d, 2\} \beta_k^{k/2}}{24\Lambda^2 L_W (\kappa_1 + \mu_1 \kappa_g)}\right) N$ with $\beta_k < \beta_{k-1}$, $\beta_k \rightarrow 0$, $\beta_0 \leq 1$;
- (ii) $|S_k^H| > |S_{k-1}^H|$, $\lim_{k \rightarrow \infty} |S_k^H| = N$, and $|S_0^H| \geq \max\left\{\frac{8\Lambda(\kappa_2 + \mu_2 \kappa_H) - 1}{8\Lambda(\kappa_2 + \mu_2 \kappa_H)} N, 1\right\}$.

If the starting point satisfies

$$\|\bar{d}_0\| \leq \min\left\{\frac{\mu_M}{12\Lambda L_W \kappa_M (1 + 2\kappa_{j^*} \kappa_W)}, \frac{\kappa_d}{2}\right\}, \quad (5.5)$$

then, $\|w_k - w^*\| \rightarrow 0$ at an R -superlinear rate, i.e., there is a positive sequence $\{\xi_k\}$ such that $\|w_k - w^*\| \leq \xi_k$ and $\frac{\xi_{k+1}}{\xi_k} \rightarrow 0$. In addition, for any $k \in \mathbb{N}$, $\|w_k - w^*\| \leq \frac{1}{3\Lambda L_W}$.

Proof. For brevity, we defer the proof to Appendix C. □

Remark 5.10. *The requirements for the rate of increase in the samples sizes of the gradient and Hessian estimates are similar to those in Theorem 5.6 with the exception that the sample sizes are capped at N . If $|S_k^H| = |S_k^g| = N$ for all $k \in \mathbb{N}$, then this theorem recovers Theorem 5.3.*

6 Convergence Analysis of Modified Line Search SQP Method

In this section, we show that the proposed algorithm (Algorithm 1) has an asymptotic superlinear convergence rate without the proximity and unit step size assumptions. Due to the challenges that arise with stochastic problems (1.2) (discussed in Sections 4 and 5.2), we focus on solving finite-sum problems (1.3). As a special case, we recover results for the deterministic problem (1.1). Specifically, we show that after sufficiently many iterations, the unit step size is always accepted, and that the iterates approach a first-order critical point (2.3). The need for the proximity assumption is circumvented by the adaptivity of the parameter γ_k (proximity check) in our proposed algorithm.

We start by presenting two fundamental results that are used in the analysis of SQP methods. The first lemma pertains to the relationship between the reduction in the model of the merit function and the search direction.

Lemma 6.1. *Suppose Assumptions 3.1 and 4.4 hold. There exists a constant $\kappa_l \in \mathbb{R}_{>0}$ such that for all $k \in \mathbb{N}$, $\Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) \geq \kappa_l \bar{\tau}_k \|\bar{d}_k\|^2$.*

Proof. The result is independent of the step size scheme. Similar results are proven in [6, Lemma 3.4] and, thus, for brevity we omit the proof. □

The next lemma shows that the merit parameter sequence is bounded below.

Lemma 6.2. *Suppose Assumptions 3.1, 3.4, and 4.4 hold. There exists $\bar{\tau}_{\min} \in \mathbb{R}_{>0}$ such that $\bar{\tau}_k \geq \bar{\tau}_{\min}$ for all $k \in \mathbb{N}$.*

Proof. By Lemma 3.5, $\|\bar{g}_k - g_k\| \leq 2 \left(\frac{N-1}{N}\right) (\kappa_1 + \mu_1 \kappa_g) \leq 2(\kappa_1 + \mu_1 \kappa_g)$. The desired result then follows by [6, Proposition 3.18]. \square

It is well established that, under reasonable assumptions, the SQP method with the classical line search condition (4.7) is globally convergent [42, 58]. The next lemma implicitly proves such a result for our proposed method for the finite-sum setting (1.3). Specifically, there does not exist a constant $\gamma_{\tilde{k}} \in \mathbb{R}_{>0}$ such that $\|\bar{d}_k\| > \gamma_{\tilde{k}}$ for all $k > \tilde{k}$, which implicitly proves global convergence. The result ensures that Line 7 of Algorithm 1 is triggered infinitely often.

Lemma 6.3. *Suppose Assumptions 3.1, 3.4, and 4.4 hold. Let the relaxation parameter in (4.11) be defined as*

$$\epsilon_{A_k} = \bar{\tau}_k \left(\kappa_{J^\dagger \kappa_c} + \zeta^{-1}(\kappa_g + \epsilon_{g_k} + \kappa_W \kappa_{J^\dagger \kappa_c}) \right) \epsilon_{g_k} + 2\bar{\tau}_k \epsilon_{f_k}, \quad (6.1)$$

where ϵ_{f_k} and ϵ_{g_k} are given in Lemma 3.5. Further, suppose that the sampling conditions in line 2 of Algorithm 1 satisfy:

$$(i) \sum_{k=0}^{\infty} \frac{N - |S_k^f|}{N} < \infty; \quad (ii) \sum_{k=0}^{\infty} \frac{N - |S_k^g|}{N} < \infty.$$

For any iteration $\tilde{k} \in \mathbb{N}$, if $\|\bar{d}_{\tilde{k}}\| > \gamma_{\tilde{k}}$, then there must exist $k > \tilde{k}$ such that $\|\bar{d}_k\| \leq \gamma_{\tilde{k}}$.

Proof. We derive the desired result by contradiction. Suppose that $\|\bar{d}_k\| > \gamma_{\tilde{k}}$ for all $k > \tilde{k}$, the algorithm will always select a step size using the classical line search condition (4.11) for all $k > \tilde{k}$. By Assumption 3.1, (5.1), (5.2), and (6.1), and the fact that $\bar{\alpha}_k \leq 1$,

$$\begin{aligned} & \phi(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_k) - \phi(x_k, \bar{\tau}_k) \\ & \leq \bar{\alpha}_k \bar{\tau}_k g_k^T \bar{d}_k - \bar{\alpha}_k \|c_k\|_1 + \frac{1}{2} (\bar{\tau}_k L_1 + \Gamma_1) \bar{\alpha}_k^2 \|\bar{d}_k\|^2 \\ & = -\bar{\alpha}_k \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) + \bar{\alpha}_k \bar{\tau}_k \bar{d}_k^T (g_k - \bar{g}_k) + \frac{1}{2} (\bar{\tau}_k L_1 + \Gamma_1) \bar{\alpha}_k^2 \|\bar{d}_k\|^2 \\ & \leq -\bar{\alpha}_k \left(1 - \frac{\bar{\tau}_k L_1 + \Gamma_1}{2\bar{\tau}_k \kappa_l} \bar{\alpha}_k \right) \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) + \bar{\alpha}_k \bar{\tau}_k \bar{d}_k^T (g_k - \bar{g}_k) \\ & \leq -\bar{\alpha}_k \left(1 - \frac{\bar{\tau}_k L_1 + \Gamma_1}{2\bar{\tau}_k \kappa_l} \bar{\alpha}_k \right) \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) + \epsilon_{A_k} - 2\bar{\tau}_k \epsilon_{f_k}. \end{aligned}$$

When $\bar{\alpha}_k \leq \min \left\{ 1, \frac{2\bar{\tau}_k \kappa_l (1-\eta)}{\bar{\tau}_k L_1 + \Gamma_1} \right\}$, by Lemma 3.5, condition (4.11) is satisfied. Let $\tilde{\alpha}_{\min} := \left\{ 1, \nu_\alpha \frac{2\bar{\tau}_{\min} \kappa_l (1-\eta)}{\bar{\tau}_{\min} L_1 + \Gamma_1} \right\}$. Since $\frac{2\bar{\tau}_k \kappa_l (1-\eta)}{\bar{\tau}_k L_1 + \Gamma_1} \geq \frac{2\bar{\tau}_{\min} \kappa_l (1-\eta)}{\bar{\tau}_{\min} L_1 + \Gamma_1}$, (4.11) is satisfied with $\bar{\alpha}_k = \tilde{\alpha}_{\min}$. Hence, for all $k \geq \tilde{k}$,

$$\bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_{k+1}) - \bar{\phi}(x_k, \bar{\tau}_k)$$

$$\begin{aligned}
&= \bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_{k+1}) - \bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_k) + \bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_k) - \bar{\phi}(x_k, \bar{\tau}_k) \\
&\leq (\bar{\tau}_{k+1} - \bar{\tau}_k) f_{\text{inf}} - \eta \tilde{\alpha}_{\min} \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) + \epsilon_{A_k}.
\end{aligned}$$

For any $K > \tilde{k}$, summing over $k = \tilde{k}, \dots, K$, by Lemma 6.2

$$\begin{aligned}
&\sum_{k=\tilde{k}}^K \bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_{k+1}) - \bar{\phi}(x_k, \bar{\tau}_k) \\
&\leq -\eta \tilde{\alpha}_{\min} \sum_{k=\tilde{k}}^K \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) + \sum_{k=\tilde{k}}^K (\bar{\tau}_{k+1} - \bar{\tau}_k) f_{\text{inf}} + \sum_{k=\tilde{k}}^K \epsilon_{A_k} \quad (6.2) \\
&\leq -\eta \tilde{\alpha}_{\min} \kappa_l \bar{\tau}_{\min} \sum_{k=\tilde{k}}^K \|\bar{d}_k\|^2 + \bar{\tau}_{-1} |f_{\text{inf}}| + \sum_{k=\tilde{k}}^K \epsilon_{A_k}.
\end{aligned}$$

By Assumption 3.1 the merit function is bounded below, and by conditions (i) and (ii) the relaxation term ϵ_{A_k} is summable. Thus, by (6.2), $\sum_{k=\tilde{k}}^K \|\bar{d}_k\|^2 < \infty$, which leads to a contradiction, and the desired result follows. \square

The sampling condition (i) (and (ii) with $|S_k^g|$) can be replaced by a sufficient condition following d'Alembert's ratio test [74], $\lim_{k \rightarrow \infty} (N - |S_{k+1}^f|) / (N - |S_k^f|) < 1$. Another surrogate condition is to ensure that $|S_k^f| = N$ after sufficiently many iterations. Lemma 6.3 shows that Algorithm 1 is guaranteed to eventually trigger the *else* condition (Line 7). Since (4.11) is not guaranteed to be satisfied with the unit step size (Line 8 of Algorithm 1), we need to examine the step sizes returned by the modified line search condition (4.12). Specifically, we show that there exists a lower bound for $\bar{\alpha}_k$ returned by (4.12), and when $\|\bar{d}_k\|$ is sufficiently small, $\bar{\alpha}_k = 1$ satisfies (4.12).

Lemma 6.4. *Suppose the assumptions and conditions in Lemma 6.3 hold. Additionally, let $|S_k^H| \geq \left(1 - \frac{\bar{\epsilon}_H}{2(\kappa_2 + \mu_2 \kappa_H)}\right) N$ where $\bar{\epsilon}_H \in (0, 2(1 - \eta)\kappa_l)$. For all $k \in \mathbb{N}$, (4.12) is satisfied with $\bar{\alpha}_k \geq \min\left\{1, \nu_\alpha \sqrt{\frac{6\bar{\tau}_{\min} \kappa_l (1 - \eta) - 3\bar{\tau}_{\min} \bar{\epsilon}_H}{(\bar{\tau}_{\min} L_2 + \Gamma_2) \gamma_0}}\right\} := \bar{\alpha}_{\min}$. Moreover, if $\|\bar{d}_k\| \leq \frac{6\bar{\tau}_{\min} \kappa_l (1 - \eta) - 3\bar{\tau}_{\min} \bar{\epsilon}_H}{\bar{\tau}_{\min} L_2 + \Gamma_2}$, (4.12) is satisfied with $\bar{\alpha}_k = 1$.*

Proof. When $|S_k^H| \geq \left(1 - \frac{\bar{\epsilon}_H}{2(\kappa_2 + \mu_2 \kappa_H)}\right) N$, by Lemma 3.5, it follows that $\epsilon_{H_k} = 2 \frac{N - |S_k^H|}{N} (\kappa_2 + \mu_2 \kappa_H) \leq \bar{\epsilon}_H < 2(1 - \eta)\kappa_l$, for all $k \in \mathbb{N}$. Thus, it follows that

$$\begin{aligned}
&\bar{\phi}(x_k + \bar{\alpha}_k \bar{d}_k, \bar{\tau}_k) - \bar{\phi}(x_k, \bar{\tau}_k) \\
&\leq \bar{\alpha}_k \bar{\tau}_k g_k^T \bar{d}_k - \bar{\alpha}_k \|c_k\|_1 + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{d}_k^T \nabla^2 f(x_k) \bar{d}_k \\
&\quad + \frac{1}{2} \bar{\alpha}_k^2 \sum_{i=1}^m |d_k^T \nabla^2 c_i(x) \bar{d}_k| + \frac{1}{6} \bar{\alpha}_k^3 (\bar{\tau}_k L_2 + \Gamma_2) \|\bar{d}_k\|^3 + 2\bar{\tau}_k \epsilon_{f_k}
\end{aligned}$$

$$\begin{aligned}
&= \bar{\alpha}_k (\bar{\tau}_k \bar{g}_k^T \bar{d}_k - \|c_k\|_1) + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{d}_k^T \bar{H}_k \bar{d}_k + \frac{1}{2} \bar{\alpha}_k^2 \sum_{i=1}^m |\bar{d}_k^T \nabla^2 c_i(x) \bar{d}_k| + \bar{\alpha}_k \bar{\tau}_k (g_k - \bar{g}_k)^T \bar{d}_k \\
&\quad + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{d}_k^T (\nabla^2 f(x_k) - \bar{H}_k) \bar{d}_k + \frac{1}{6} \bar{\alpha}_k^3 (\bar{\tau}_k L_2 + \Gamma_2) \|\bar{d}_k\|^3 + 2\bar{\tau}_k \epsilon_{f_k} \\
&\leq -\bar{\alpha}_k \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{\epsilon}_H \|\bar{d}_k\|^2 + \frac{1}{6} \bar{\alpha}_k^3 (\bar{\tau}_k L_2 + \Gamma_2) \|\bar{d}_k\|^3 \\
&\quad + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{d}_k^T \bar{H}_k \bar{d}_k + \frac{1}{2} \bar{\alpha}_k^2 \sum_{i=1}^m |\bar{d}_k^T \nabla^2 c_i(x) \bar{d}_k| + \epsilon_{A_k} \\
&\leq -\bar{\alpha}_k \left(1 - \frac{\bar{\epsilon}_H}{2\kappa_l} \bar{\alpha}_k - \frac{\bar{\tau}_k L_2 + \Gamma_2}{6\bar{\tau}_k \kappa_l} \bar{\alpha}_k^2 \|\bar{d}_k\| \right) \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) \\
&\quad + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{d}_k^T \bar{H}_k \bar{d}_k + \frac{1}{2} \bar{\alpha}_k^2 \sum_{i=1}^m |\bar{d}_k^T \nabla^2 c_i(x) \bar{d}_k| + \epsilon_{A_k} \\
&\leq -\bar{\alpha}_k \left(1 - \frac{\bar{\epsilon}_H}{2\kappa_l} - \frac{\bar{\tau}_k L_2 + \Gamma_2}{6\bar{\tau}_k \kappa_l} \bar{\alpha}_k^2 \gamma_0 \right) \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) \\
&\quad + \frac{1}{2} \bar{\alpha}_k^2 \bar{\tau}_k \bar{d}_k^T \bar{H}_k \bar{d}_k + \frac{1}{2} \bar{\alpha}_k^2 \sum_{i=1}^m |\bar{d}_k^T \nabla^2 c_i(x) \bar{d}_k| + \epsilon_{A_k},
\end{aligned}$$

where the first inequality follows due to Assumption 3.1 and Lemma 3.5, the second inequality follows by Lemma 3.5 and the definitions of $\Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k)$ and ϵ_{A_k} , the third inequality follows by Lemma 6.1, and the last inequality follows by the fact that $\bar{\alpha}_k \leq 1$ and $\|\bar{d}_k\| \leq \gamma_k \leq \gamma_0$. When $\bar{\alpha}_k \leq \min \left\{ 1, \sqrt{\frac{6\bar{\tau}_k \kappa_l (1-\eta) - 3\bar{\tau}_k \bar{\epsilon}_H}{(\bar{\tau}_k L_2 + \Gamma_2) \gamma_0}} \right\}$, the relaxed line search condition (4.12) is satisfied. Since $\sqrt{\frac{6\bar{\tau}_k \kappa_l (1-\eta) - 3\bar{\tau}_k \bar{\epsilon}_H}{(\bar{\tau}_k L_2 + \Gamma_2) \gamma_0}} \geq \sqrt{\frac{6\bar{\tau}_{\min} \kappa_l (1-\eta) - 3\bar{\tau}_{\min} \bar{\epsilon}_H}{(\bar{\tau}_{\min} L_2 + \Gamma_2) \gamma_0}}$, the desired conclusion holds with the line search procedure $\bar{\alpha}_k = \nu_\alpha \bar{\alpha}_k$.

Plugging $\bar{\alpha}_k = 1$ into the third inequality above,

$$\begin{aligned}
\bar{\phi}(x_k + \bar{d}_k, \bar{\tau}_k) - \bar{\phi}(x_k, \bar{\tau}_k) &\leq - \left(1 - \frac{\bar{\epsilon}_H}{2\kappa_l} - \frac{\bar{\tau}_k L_2 + \Gamma_2}{6\bar{\tau}_k \kappa_l} \|\bar{d}_k\| \right) \Delta l(x_k, \bar{\tau}_k, \bar{g}_k, \bar{d}_k) \\
&\quad + \frac{1}{2} \bar{\tau}_k \bar{d}_k^T \bar{H}_k \bar{d}_k + \frac{1}{2} \sum_{i=1}^m |\bar{d}_k^T \nabla^2 c_i(x) \bar{d}_k| + \epsilon_{A_k}. \tag{6.3}
\end{aligned}$$

When $\|\bar{d}_k\| \leq \frac{6\bar{\tau}_{\min} \kappa_l (1-\eta) - 3\bar{\tau}_{\min} \bar{\epsilon}_H}{\bar{\tau}_{\min} L_2 + \Gamma_2}$, it follows that $1 - \frac{\bar{\epsilon}_H}{2\kappa_l} - \frac{\bar{\tau}_k L_2 + \Gamma_2}{6\bar{\tau}_k \kappa_l} \|\bar{d}_k\| \geq \eta$. Thus, the modified line search condition (4.12) is satisfied. \square

We now present our main convergence result for Algorithm 1. We establish asymptotic local superlinear convergence results under appropriate sampling conditions.

Theorem 6.5. *Let $\{w_k\}_{k \in \mathbb{N}}$ be the primal-dual iterates generated by Algorithm 1. Suppose Assumptions 3.1, 3.2, 3.4, 4.4, and 5.1 hold. Further, suppose that the following conditions are satisfied for all $k \in \mathbb{N}$:*

$$(i) \sum_{k=0}^{\infty} \frac{N - |S_k^f|}{N} < \infty; \quad (ii) \sum_{k=0}^{\infty} \frac{N - |S_k^g|}{N} < \infty;$$

(iii) $|S_k^H| \geq \left(1 - \frac{\bar{\epsilon}_H}{2(\kappa_2 + \mu_2 \kappa_H)}\right) N$, where $\bar{\epsilon}_H < 2(1 - \eta)\kappa_l$.

Moreover, suppose that for sufficiently large $\hat{k} \in \mathbb{N}$, the following sampling conditions are satisfied for all $k \geq \hat{k}$, where $t = k - \hat{k}$:

(iv) $|S_k^g| \geq \left(1 - \frac{\kappa_w \rho_0 \beta_t^{t/2}}{4\Lambda(\kappa_1 + \mu_1 \kappa_g)}\right) N$ for some $\kappa_w \in \mathbb{R}_{>0}$, with $\beta_t < \beta_{t-1}$, $\beta_t \rightarrow 0$, $\beta_1 \leq \rho_0^4$

and $\rho_0 := \min \left\{ \nu_\gamma, \frac{\nu_\gamma^2 \mu_M}{4\kappa_M(1+2\kappa_{j^\dagger} \kappa_W)} \max\{2\Lambda L_{\nabla \mathcal{L}}, \nu_\gamma\}, \frac{\mu_M \Lambda}{2} \right\}$;

(v) $|S_{k+1}^H| > |S_k^H|$, $\lim_{k \rightarrow \infty} |S_k^H| = N$, and $|S_{\hat{k}}^H| \geq \max \left\{ \frac{8\Lambda(\kappa_2 + \mu_2 \kappa_H) - \rho_0}{8\Lambda(\kappa_2 + \mu_2 \kappa_H)} N, 1 \right\}$.

Then, $\|w_k - w^*\| \rightarrow 0$ at an R -superlinear rate.

Proof. Conditions (iv) and (v) imply, $\epsilon_{g_k} = 2 \frac{N - |S_k^g|}{N} (\kappa_1 + \mu_1 \kappa_g) \leq \frac{\kappa_w \rho_0 \beta_t^{t/2}}{2\Lambda}$ and $\epsilon_{H_{\hat{k}}} = 2 \frac{N - |S_{\hat{k}}^H|}{N} (\kappa_2 + \mu_2 \kappa_H) \leq \frac{\rho_0}{4\Lambda}$. It follows by Lemma 6.3 that $\|\bar{d}_k\| \leq \gamma_k$ infinitely often. We consider two cases:

- (1) Line 9 of Algorithm 1 is triggered infinitely often;
- (2) Line 9 of Algorithm 1 is not triggered infinitely often.

Let us first consider case (1). Define

$$\kappa_w := \min \left\{ \frac{\nu_\gamma}{3\Lambda L_W}, \frac{\kappa_d}{\rho_0}, \frac{\nu_\gamma^2 \mu_M}{8\kappa_M(1+2\kappa_{j^\dagger} \kappa_W) \Lambda^2 L_{\nabla \mathcal{L}} L_W}, \frac{\nu_\gamma^2 \min \left\{ \frac{\kappa_d}{2}, \frac{6\bar{\tau}_{\min} \kappa_l (1-\eta) - 3\bar{\tau}_{\min} \bar{\epsilon}_H}{\bar{\tau}_{\min} L_2 + \Gamma_2} \right\}}{\max\{2\Lambda L_{\nabla \mathcal{L}}, \nu_\gamma\} \rho_0} \right\},$$

$$\gamma_d := \min \left\{ \frac{\mu_M \kappa_w}{4\kappa_M(1+2\kappa_{j^\dagger} \kappa_W)}, \frac{\kappa_d}{2}, \frac{6\bar{\tau}_{\min} \kappa_l (1-\eta) - 3\bar{\tau}_{\min} \bar{\epsilon}_H}{\bar{\tau}_{\min} L_2 + \Gamma_2} \right\}.$$

Algorithm 1 sets $\gamma_{k+1} = \nu_\gamma \gamma_k$ when it enters Line 9. Since $\gamma_{k+1} = \nu_\gamma \gamma_k$ happens infinitely often, there must exist some $\tilde{k} > 0$ such that the algorithm sets $\gamma_{\tilde{k}+1} = \nu_\gamma \gamma_{\tilde{k}} \in [\nu_\gamma \gamma_d, \gamma_d]$. Suppose that $\|\bar{d}_{\tilde{k}+1}\| > \gamma_{\tilde{k}+1}$, it follows by Lemma 6.3 that there must exist $\hat{k} > \tilde{k} + 1$ such that $\|\bar{d}_{\hat{k}}\| \leq \gamma_{\tilde{k}+1}$ for the first time, γ_k is therefore not updated and $\gamma_{\hat{k}} = \gamma_{\tilde{k}+1} \in [\nu_\gamma \gamma_d, \gamma_d]$. If $\|\bar{d}_{\hat{k}+1}\| \leq \gamma_{\hat{k}+1}$, let $\hat{k} = \tilde{k} + 1$ and $t = k - \hat{k}$. We define sequences $\{\xi_t\}_{t=0}^\infty$ and $\{\rho_t\}_{t=1}^\infty$ as follows:

$$\xi_{t+1} = \max \left\{ \xi_t \rho_t, \beta_{t+1}^{(t+1)/4} \right\}, \quad \xi_0 = 1, \quad \rho_t = \frac{\Lambda L_W \kappa_w \xi_t}{2} + \Lambda \epsilon_{H_{\hat{k}}} \frac{N - |S_k^H|}{N - |S_{\hat{k}}^H|} + \frac{\rho_0 \beta_t^{t/4}}{2}.$$

We prove the following results for all $k \geq \hat{k}$ ($t = k - \hat{k}$, $t \geq 0$):

- (a) **(Unit step size):** $\bar{\alpha}_k = 1$;
- (b) **(Superlinear convergence):** $\|w_k - w^*\| \leq \kappa_w \xi_t$;

(c) **(Controlled sequences)**: $\xi_t \leq 1, \rho_t < 1$;

(d) ($\|\bar{d}_{k+1}\| \leq \gamma_{k+1}$): $2\Lambda L_{\nabla\mathcal{L}}\kappa_w\xi_{t+1} \leq \gamma_{\hat{k}}\nu_\gamma^{t+1}$.

Note that (a), (b), (c) are required in Theorem 5.9, we also need (d) to enforce the algorithm to always enter line 7. Otherwise, our algorithm may consider to employing the line search condition (4.11) and the resulting step size could be smaller than 1.

Let us first give two general results. Assume that (a) and (b) hold for iteration k , and consider iteration $k+1$. By Lemma 5.8, conditions (iv) and (v), it follows that

$$\begin{aligned} \|w_{k+1} - w^*\| &\leq \frac{\Lambda L_W}{2} \|w_k - w^*\|^2 + \Lambda \epsilon_{H_k} \|w_k - w^*\| + \frac{\kappa_w \rho_0 \beta_t^{t/2}}{2} \\ &\leq \frac{\Lambda L_W}{2} \kappa_w^2 \xi_t^2 + \Lambda \epsilon_{H_k} \frac{N - |S_k^H|}{N - |S_k^H|} \kappa_w \xi_t + \frac{\kappa_w \rho_0 \beta_t^{t/2}}{2} \\ &\leq \kappa_w \xi_t \left(\frac{\Lambda L_W \kappa_w \xi_t}{2} + \Lambda \epsilon_{H_k} \frac{N - |S_k^H|}{N - |S_k^H|} + \frac{\rho_0 \beta_t^{t/4}}{2} \right). \end{aligned} \quad (6.4)$$

Moreover, when $\|w_k - w^*\| \leq \kappa_w$, by (4.10), Assumptions 3.2 and condition (iv) that

$$\begin{aligned} \|\bar{d}_k\| &\leq \|\bar{M}_k^{-1}\| \left\| \begin{bmatrix} \bar{g}_k + J_k^T y_k \\ c_k \end{bmatrix} \right\| \leq \|\bar{M}_k^{-1}\| \left\| \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} + \begin{bmatrix} \bar{g}_k - g_k \\ 0 \end{bmatrix} \right\| \\ &\leq \Lambda L_{\nabla\mathcal{L}} \|w_k - w^*\| + \frac{\kappa_w \rho_0 \beta_t^{t/2}}{2}. \end{aligned} \quad (6.5)$$

We then use induction to prove (a)–(d). For the base case $k = \hat{k}$, since $\|\bar{d}_{\hat{k}}\| \leq \gamma_d$, (a) is satisfied by Lemma 6.4. Since $\rho_0 \leq \frac{\mu_M \Lambda}{2}$ and $\kappa_w \leq \frac{\kappa_d}{\rho_0}$, we have $\epsilon_{g_{\hat{k}}} \leq \frac{\kappa_w \rho_0}{2\Lambda} \leq \min\left\{\frac{\mu_M \kappa_w}{4}, \frac{\kappa_d}{2\Lambda}\right\}$. It then follows by Lemma 5.4 that $\|w_{\hat{k}} - w^*\| \leq \kappa_w$ by the definitions of γ_d and κ_w , (b) is then satisfied. Since $\nu_\gamma < 1$, (c) holds trivially by the definitions of ξ_0 and ρ_0 . By (6.4), $\kappa_w \leq \frac{\nu_\gamma^2 \mu_M}{8\kappa_M(1+2\kappa_{j^*}\kappa_W)\Lambda^2 L_{\nabla\mathcal{L}} L_W}$, $\xi_0 = 1$, the definition of ρ_0 , and conditions (iv) and (v) it follows that

$$\|w_{\hat{k}+1} - w^*\| \leq \kappa_w \xi_0 \left(\frac{\Lambda L_W \kappa_w}{2} + \Lambda \epsilon_{H_{\hat{k}}} + \frac{\rho_0}{2} \right) \leq \kappa_w \rho_0.$$

Since $\beta_1^{1/4} \leq \rho_0$, we have $\xi_1 = \rho_0$, and it follows from the definition of ρ_0 that

$$\max\{2\Lambda L_{\nabla\mathcal{L}}\kappa_w\rho_0, \nu_\gamma\kappa_w\rho_0\} \leq \frac{\nu_\gamma^2 \mu_M \kappa_w}{4\kappa_M(1+2\kappa_{j^*}\kappa_W)}. \quad (6.6)$$

Furthermore, since $\kappa_w \leq \frac{\nu_\gamma^2}{\max\{2\Lambda L_{\nabla\mathcal{L}}, \nu_\gamma\}\rho_0} \min\left\{\frac{\kappa_d}{2}, \frac{6\bar{\tau}_{\min}\kappa_l(1-\eta)-3\bar{\tau}_{\min}\bar{\epsilon}_H}{\bar{\tau}_{\min}L_2+\Gamma_2}\right\}$,

$$\max\{2\Lambda L_{\nabla\mathcal{L}}\kappa_w\rho_0, \nu_\gamma\kappa_w\rho_0\} \leq \nu_\gamma^2 \min\left\{\frac{\kappa_d}{2}, \frac{6\bar{\tau}_{\min}\kappa_l(1-\eta)-3\bar{\tau}_{\min}\bar{\epsilon}_H}{\bar{\tau}_{\min}L_2+\Gamma_2}\right\}. \quad (6.7)$$

It then follows from (6.5), (6.6), and (6.7) that

$$\|\bar{d}_{\hat{k}+1}\| \leq \Lambda L_{\nabla \mathcal{L}} \kappa_w \xi_1 + \frac{\kappa_w \rho_0 \nu_\gamma}{2} \leq \nu_\gamma^2 \gamma_d \leq \nu_\gamma \gamma_{\hat{k}} \leq \gamma_{\hat{k}+1},$$

and (d) is satisfied for the base case.

Let us then assume that (a) and (b) hold for iteration k . At iteration $k+1$, (a) is satisfied since $\|\bar{d}_{k+1}\| \leq \gamma_{k+1} \leq \gamma_d$. (b) is satisfied since $\|w_{k+1} - w^*\| \leq \kappa_w \xi_t \rho_t \leq \kappa_w \xi_{t+1}$. As for (c), it follows that

$$\xi_{t+1} \leq \max \left\{ \rho_t, \beta_1^{(t+1)/4} \right\} < 1, \quad \rho_{t+1} \leq \frac{\Lambda L_W \kappa_w}{2} + \Lambda \epsilon_{H_{\hat{k}}} + \frac{\rho_0}{2} = \rho_0 \leq \nu_\gamma < 1.$$

Since (a) and (b) hold for iteration $k+1$, we can apply (6.4) again at iteration $k+1$ and obtain $\|w_{k+2} - w^*\| \leq \kappa_w \xi_{t+1} \rho_{t+1}$ by the definition of ρ_{t+1} . It then follows from (6.5), (6.6), and (6.7) that

$$\begin{aligned} \|\bar{d}_{k+2}\| &\leq \Lambda L_{\nabla \mathcal{L}} \|w_{k+2} - w^*\| + \frac{\kappa_w \rho_0 \beta_1^{(t+2)/2}}{2} \leq \Lambda L_{\nabla \mathcal{L}} \kappa_w \xi_{t+1} \rho_{t+1} + \frac{\kappa_w \rho_0 \beta_1^{(t+2)/4}}{2} \\ &\leq \frac{\gamma_{\hat{k}} \nu_\gamma^{t+1} \rho_{t+1}}{2} + \frac{\kappa_w \rho_0 \nu_\gamma^{t+2}}{2} \leq \frac{\gamma_{\hat{k}} \nu_\gamma^{t+2}}{2} + \frac{\nu_\gamma \gamma_d \nu_\gamma^{t+2}}{2} = \gamma_{\hat{k}} \nu_\gamma^{t+2}. \end{aligned}$$

(d) is therefore satisfied. We then conclude that $\{\xi_t\} \rightarrow 0$ by the fact that $\rho_t \leq \nu_\gamma$ and $\beta_t \rightarrow 0$. In addition, (iv) and (v) implies that $\{\rho_t\} \rightarrow 0$. We can conclude that

$$\lim_{t \rightarrow \infty} \frac{\xi_{t+1}}{\xi_t} = \lim_{t \rightarrow \infty} \max \left\{ \rho_t, \frac{\beta_{t+1}^{(t+1)/4}}{\xi_t} \right\} \leq \lim_{t \rightarrow \infty} \max \left\{ \rho_t, \left(\frac{\beta_{t+1}}{\beta_t} \right)^{t/4} \beta_{t+1}^{1/4} \right\} = 0.$$

Hence, $\|w_k - w^*\| \rightarrow 0$ at an R-superlinear rate.

As for case when Line 9 of Algorithm 1 is not triggered infinitely often, Algorithm 1 either enters Line 6 or Line 8 after some iteration k_c . In this scenario, line search condition (4.11) is always considered and γ_k remains fixed at some $\bar{\gamma}$ after iteration k_c . It then follows by the proof of Lemma 6.3 that there exists $k_d \in \mathbb{R}_{>0}$ such that $\|\bar{d}_k\| \leq \min\{\bar{\gamma}, \bar{\gamma}_d\}$ for all $k \geq k_d$ ($\bar{\gamma}_d$ is defined in Theorem 5.9), which implies Algorithm 1 enters Line 8 for all $k \geq \max\{k_c, k_d\}$. Let $k_s = \max\{k_c, k_d, \hat{k}\}$, excluding first k_s iterations, and assuming (without loss of generality) that w_{k_s} as the starting point, it follows by the definition of κ_w and ρ_0 that $\kappa_w \rho_0 \leq \frac{1}{3\Lambda L_W} \min \left\{ \frac{\mu_M \Delta}{2}, 3\Lambda L_W \kappa_d, 1 \right\}$. In addition, $\beta_k < \beta_{k-1}$ and $\rho_0 < 1$ hold, and as a result, conditions (i) and (ii) in Theorem 5.9 are satisfied. Thus, considering w_{k_s} as the starting point and applying the result from Theorem 5.9 completes the proof. \square

Remark 6.6. *Theorem 6.5 provides an asymptotic superlinear convergence rate with respect to $\|w_k - w^*\|$. It unifies the global convergence and the local superlinear convergence results of classical SQP methods. The sampling conditions (i), (ii) are required for the global convergence result. Condition (iii) is needed for proving that the modified line search condition (4.12) produces unit step size when $\|\bar{d}_k\|$ is sufficiently small, and conditions*

(iv) and (v) are required for showing the superlinear convergence rate of $\|w_k - w^*\|$ and need to be satisfied after \hat{k} iterations. The iteration \hat{k} determines when $\|\bar{d}_{\hat{k}}\|$ is sufficiently small (similar to the conditions required in Theorem 5.9). Although \hat{k} depends on unknown problem-specific constants, the initial conditions of (iv) and (v) are guaranteed to be satisfied after a sufficient number of iterations as long as (ii) and $\lim_{k \rightarrow \infty} |S_k^H| = N$ hold.

7 Practical Inexact SQP

In Algorithm 1, the exact solution of the linear system (4.10) is required to compute a step. This poses challenges in large-scale (n and m) settings. In this section, we present an inexact matrix-free variant of Algorithm 1 that utilizes the Minimum Residual (MINRES) method [21] to solve (4.10) at every iteration. The rationale behind choosing MINRES over other widely used approaches lies in the facts that the matrix \bar{M}_k is not assumed to be positive definite, that the method can be implemented matrix-free, and due to the robust theoretical guarantees and empirical performance [21, 32, 43].

Let $(\bar{d}_k^{(t)}, \bar{\delta}_k^{(t)})$ denote the approximate solution obtained at the t -th iteration of the MINRES method (at the k -th iteration of Algorithm 1) such that

$$\begin{bmatrix} \rho_k^{(t)} \\ r_k^{(t)} \end{bmatrix} = \begin{bmatrix} \bar{W}_k & J_k^T \\ J_k & 0 \end{bmatrix} \begin{bmatrix} \bar{d}_k^{(t)} \\ \bar{\delta}_k^{(t)} \end{bmatrix} + \begin{bmatrix} \bar{g}_k + J_k^T y_k \\ c_k \end{bmatrix}, \quad (7.1)$$

where $\rho_k^{(t)} \in \mathbb{R}^n$ and $r_k^{(t)} \in \mathbb{R}^m$ denote the vectors of residuals. In the practical version of our algorithm, the residual vectors are nonzero and we utilize the inexact solutions $\bar{d}_k^{(t)}, \bar{\delta}_k^{(t)}$ in lieu of the exact solutions $\bar{d}_k, \bar{\delta}_k$ of subproblem (4.10) in Algorithm 1. Restrictions are imposed on the residual vectors which in turn provide lower bounds on the number of MINRES iterations.

For ease of presentation and analysis, we assume the true gradient of the objective function is available and used ($\bar{g}_k = g_k = \nabla f_k$) and only approximate the Hessian. One can consider the setting with inexact gradients at the cost of weaker theoretical results and more complicated analysis. For brevity we do not provide a comprehensive convergence analysis of the practical version of Algorithm 1. Instead, we explore the relationship between the accuracy in (7.1), the Hessian sample size, and the rate of convergence of the inexact method. Similar to Section 5 we make the proximity and unit step size assumptions.

The following lemma characterizes the accuracy of the approximate solutions obtained with respect to the number of MINRES iterations.

Lemma 7.1. *Suppose Assumptions 3.1 and 4.4 hold. There exists $\theta \in (0, 1)$ such that for all $k, t \in \mathbb{N}$, $\left\| \begin{bmatrix} \rho_k^{(t)} \\ r_k^{(t)} \end{bmatrix} \right\| \leq 2\theta^t \left\| \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\|$.*

Proof. The proof can be found in Appendix D. \square

We make two comments about the above result. First, the constant θ has dependence on the condition number of the matrix \bar{M}_k (4.10). Second, if $t = m + n$ MINRES iterations are performed, then the exact solution of the linear system is guaranteed to be obtained under the given assumptions [21].

The next lemma is an inexact (in terms of the solution of the linear system) counterpart of Lemma 5.8.

Lemma 7.2. *Suppose Assumptions 3.1, 3.2, and 4.4 hold. For all $k \in \mathbb{N}$, let $|S_k^g| = N$ and $|S_k^H| \leq N$. Then, when $\|w_k - w^*\| \leq r$ for some $r \in \mathbb{R}_{>0}$ (defined in Assumption 3.2),*

$$\|w_{k+1} - w^*\| \leq \frac{\Lambda L_W}{2} \|w_k - w^*\|^2 + 2\Lambda \left(\frac{(N - |S_k^H|)}{N} (\kappa_2 + \mu_2 \kappa_H) + \theta^t L_{\nabla \mathcal{L}} \right) \|w_k - w^*\|.$$

Proof. The proof is identical to [10, Lemma 3.1] if we remove the expectation, view the Lagrangian function as the objective function, and view the primal-dual variables as the decision variables. Thus, for brevity we omit the proof. \square

Under the proximity and unit step size assumptions, we prove a linear rate of convergence for the proposed practical inexact matrix-free SQP method.

Theorem 7.3. *Suppose Assumptions 3.1, 3.2, and 4.4 hold. For $k \in \mathbb{N}$, let $|S_k^g| = N$ and $|S_k^H| \geq \max \left\{ \frac{12\Lambda(\kappa_2 + \mu_2 \kappa_H) - 1}{12\Lambda(\kappa_2 + \mu_2 \kappa_H)} N, 1 \right\}$. If $t \geq \frac{\log(12\Lambda L_{\nabla \mathcal{L}})}{\log(\frac{1}{\theta})}$ and $\|\bar{d}_0\| \leq \bar{\gamma}_d$ ($\bar{\gamma}_d$ defined in Theorem 5.9), then $\|w_{k+1} - w^*\| \leq \frac{1}{2} \|w_k - w^*\|$.*

Proof. By the conditions given in the theorem, it follows that $\frac{(N - |S_k^H|)}{N} (\kappa_2 + \mu_2 \kappa_H) \leq \frac{1}{12}$ and $\Lambda L_{\nabla \mathcal{L}} \theta^t \leq \frac{1}{12}$. Moreover, as proved in Theorem 5.9, $\|w_0 - w^*\| \leq \frac{1}{3\Lambda L_W}$. The remainder of the proof is similar to [10, Theorem 3.2], with $r = \frac{1}{3\Lambda L_W}$ in Lemma 7.2. \square

Remark 7.4. *Theorem 7.3 shows that under the proximity and unit step size assumptions, if a sufficient number of MINRES iterations are employed, then the primal and dual iterates converge at a linear rate. The linear rate of convergence is a fixed constant and the number of required MINRES iterations has only a logarithmic dependence on the condition number. In the inexact gradient setting, one can prove a similar result if the gradient approximation is sufficiently small as compared to the proximity measure. We note that the analysis presented in this section can be integrated into the adaptive algorithm of Section 6. Finally, in the stochastic setting (1.2), one requires that the gradient approximation error is sufficiently small and Assumption 5.5 to derive a linear rate of convergence in expectation.*

8 Numerical Experiments

In this section, we demonstrate the empirical performance of our proposed algorithm in both the deterministic and stochastic settings using a MATLAB implementation. We

first compare the performance of Algorithm 1 to the classical line search SQP method and other adaptations (discussed in Section 2) on a subset of the nonlinear equality constrained problems from the CUTEst collection [11]. In the stochastic setting, we compare with the adaptive step size SQP method proposed in [6] on equality constrained logistic regression problems. Finally, we compare the exact and inexact variants of Algorithm 1.

8.1 Deterministic Setting

The goal of this section is to demonstrate the robustness and efficiency of our proposed modified line search SQP method in the deterministic setting. We compare our proposed algorithm to the classical line search SQP method [42] and other adaptations (discussed in Section 2) on the CUTEst collection of test problems. Among 123 equality constrained test problems, we select 41 problems with the following characteristics: (1) f is not a constant function; (2) $n + m \leq 1000$; (3) the LICQ holds at all iterates in all runs of all algorithms (4 LICQ failures); and, (4) at least one method was able to solve the problem (12 unsolved).

We consider 5 methods in this subsection: **SQP-L1** [42], **2nd-corr** [34], **Watchdog** [16], **SQP-AugLag** [64] and **Our method** (Algorithm 1). Common parameters, used across methods, were set as: $\tau_{-1} = 1$, $\nu_\alpha = 0.5$, $\eta = 10^{-4}$ as in [6]. Method-specific parameters were set as:

- **Our Method** (Algorithm 1): $\gamma_0 = 0.999\|d_0\|$, $\nu_\gamma = 0.7$;
- **Watchdog** [16]: 5 iterations for the relaxed step before a restart, as recommended in [51, Section 15.6];
- **SQP-AugLag** [64]: $r_{-1} = 10^6$ (tuned over $r_{-1} = \{10, 10^2, \dots, 10^8\}$).

Although there are two additional hyperparameters γ_0 and ν_γ in Algorithm 1, the performance of the method is robust with respect to these hyperparameters, and the choices stated above work well across most problems tested. For **SQP-AugLag**, we found that the performance is sensitive to the initial penalty parameter r_{-1} , and tuned this value. A run terminated with a message of success if for $k \leq 100$,

$$\|g_k + J_k^T y_k\|_\infty \leq 10^{-6} \max\{1, \|g_0 + J_0^T y_0\|_\infty\} \quad \text{and} \quad \|c_k\|_\infty \leq 10^{-6} \max\{1, \|c_0\|_\infty\}.$$

Otherwise, the run was considered a failure.

Figure 2 presents the performance of the methods in terms of iterations and function evaluations using Dolan-Moré performance profiles [28]. As is clear, given the budget of iterations, our proposed method is robust and efficient, and, in fact, slightly outperforms the other methods on this test set. Moreover, we investigate the 7 instances for which the performance of the **SQP-L1** method differs from that of our proposed line search method indicating that the modified line search condition was triggered (i.e., Line 9, Algorithm 1). For these problems, the number of iterations for which the modified line search is employed




Figure 2: Dolan-Moré performance profiles [28] comparing methods on CUTEst collection of test problems in terms of iterations (left) and function evaluations (right).

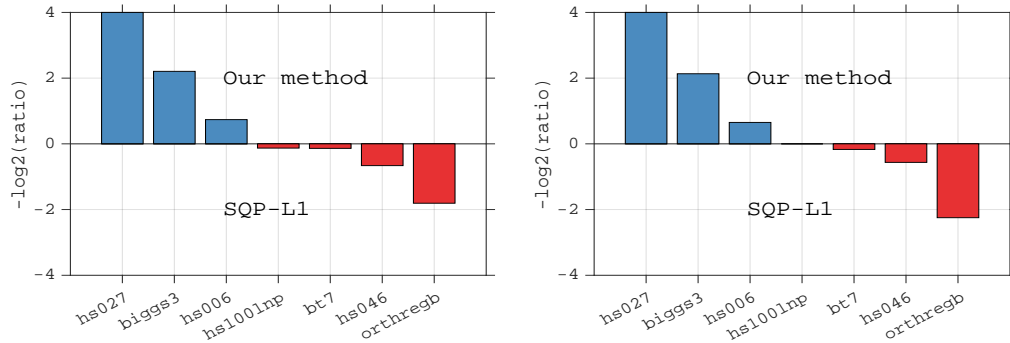


Figure 3: Morales performance profiles [47] comparing the classical line search SQP method (**SQP-L1**) and our proposed method on seven test problems in terms of iterations (left) and function evaluations (right).

is not insignificant; see Table 1. We present a comparison of the two methods in the form of Morales profiles [47] (Figure 3). While the overall performance is comparable, the results suggest that when our proposed method outperforms **SQP-L1** method, it does so by a significant margin.

Table 1: Percentage of iterations for which modified line search condition is employed.

Problem	hs027	biggs3	hs006	hs100lnp	bt7	hs046	orthregb
Percentage	8%	23%	33%	18%	14%	16%	71%

8.2 Stochastic Setting

In this section, we consider the following constrained binary classification problem,

$$\min_{x \in \mathbb{R}^n} f(x) = \frac{1}{N} \sum_{i=1}^N \log \left(1 + e^{-y_i (X_i x)} \right) \quad \text{s.t.} \quad A_1 x = a_1, \quad x^T A_2 x = a_2, \quad (8.1)$$

where $X \in \mathbb{R}^{N \times n}$ is the data matrix (feature data for N data points), $X_i \in \mathbb{R}^{1 \times n}$ is the i th row of X , $y \in \{-1, 1\}^N$ are the labels (for each data point), and $A_1 \in \mathbb{R}^{m \times n}$ (full row rank), $a_1 \in \mathbb{R}^m$, $A_2 \in \mathbb{R}^{n \times n}$ (positive definite) and $a_2 \in \mathbb{R}$. In these experiments, we consider datasets from the LIBSVM collection [17]. For brevity, we present results only on two datasets: **australian**: $n = 14$, $N = 621$; **mushroom**: $n = 112$, $N = 5500$. With regards to the constraints, $m = 5$, the matrix and vector A_1 and a_1 were generated from a normal distribution, the matrix A_2 was a positive definite matrix with eigenvalues evenly distributed between 1 and 10, and $a_2 = 5$.

In this set of experiments, the methods have access to exact function and gradient information of the objective and constraint functions, but inexact Hessian information of the objective function. The goal of these experiments is to investigate the effect of the Hessian approximation, the step size scheme and search direction (linear system solution) quality. We consider two batch sizes, $b = 5\%N$ (small batch) and $b = 50\%N$ (large batch) for the Hessian, and also the exact Hessian setting ($b = 100\%N$, full batch). We compare the adaptive step size method proposed in [6] and Algorithm 1. For each step size scheme we consider first-order ($W_k = I$) and second-order (W_k approximation of the Hessian of the Lagrangian) variants. We note that the first-order variants of the methods are essentially deterministic algorithms. A budget of 50 iterations ($50N$ gradient evaluations), $100N$ function evaluations, and $50N$ Hessian evaluations was used for all methods. We use the same parameter values for Algorithm 1 as described in Section 8.1. For the adaptive step size method, we use the default values given in [6, Section 4.2]. We should note that this adaptive step size scheme requires at least an estimate of the Lipschitz constants of the objective and constraint gradients. For all problems and algorithms, the initial primal iterate (x_0) was set to a normal random vector scaled to have norm 0.1, and the multipliers were initialized as $y_0 = \arg \min_{y \in \mathbb{R}^m} \|g_0 + J_0^T y\|_2^2$. Finally, for the small and large batch instances, we ran 10 replications with different random seeds for each problem, dataset, and algorithm, and report the average performance.

Figure 4 presents the evolution of the feasibility and stationarity errors and the step size with respect to iterations and number of Hessian evaluations. Overall, the second-order variants of the methods (even small batch) converge faster than the first-order variants in terms of iterations. In the small batch setting (5%), Algorithm 1 yields smaller step sizes as compared to the adaptive step size scheme counterpart. This observation corroborates Lemma 6.4, which posits that a small Hessian approximation error is requisite. When more accurate Hessian approximations are employed, our results suggest that in terms of

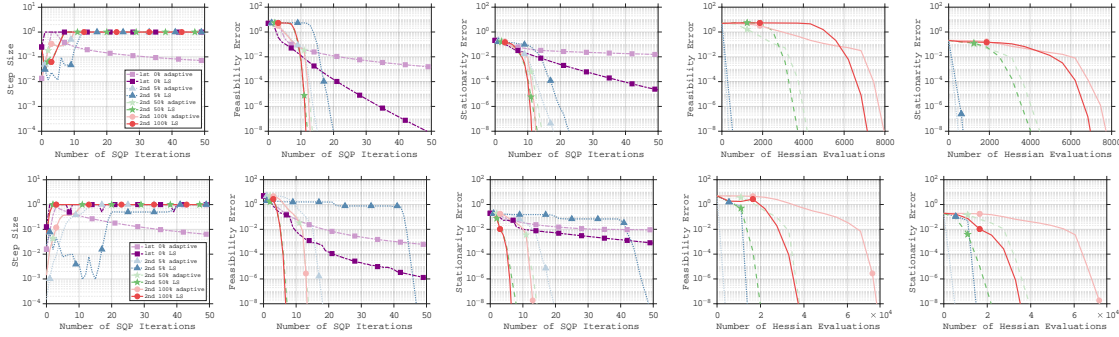


Figure 4: Comparison of first-order (**1st**) and second-order (**2nd**) SQP methods with adaptive step size (**adaptive**) and line search (**LS**) schemes on the **australian** (top) and **mushroom** (below) datasets [17]. True function and gradient information and inexact Hessian information (batch sizes **5%**, **50%**, **100%** of total sample size).

iterations the line search variant (Algorithm 1) outperforms the adaptive step size variant. This, of course, comes at the cost of function evaluations that are not required by the adaptive step size method. Finally, as expected, the small batch variants of the two methods outperform the large and full batch variants when the metric of comparison is the number of Hessian evaluations.

Next, we compare exact and inexact (solution of the linear system (4.1)) variants of Algorithm 1 in Figure 5. In these experiments, we utilize the MINRES method [21] to solve the linear system (4.1) to a prescribed accuracy. Specifically, the MINRES method is terminated when the residual vectors $\rho_k^{(t)}$ and $r_k^{(t)}$ satisfy

$$\left\| \begin{bmatrix} \rho_k^{(t)} \\ r_k^{(t)} \end{bmatrix} \right\|_{\infty} \leq \max \left\{ \kappa \left\| \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\|_{\infty}, 10^{-12} \right\} \quad (8.2)$$

where $\kappa = 10^{-12}$ (exact) and $\kappa = 10^{-1}$ (inexact) is the accuracy parameter. In addition to the budgets given earlier in this subsection, a budget of $(m+n)N$ total MINRES iterations was imposed. We compare exact and inexact variants of Algorithm 1 with $b = 5\%N$ (small batch), $b = 50\%N$ (large batch) and $b = 100\%N$ (full batch), and an adaptive Hessian batch size scheme, $b_k = \min \{ \lfloor (1 - 0.95^{(k+2)/2})N \rfloor, N \}$, which guarantees that the initial Hessian estimate uses 5% of the total samples. In Figure 5, we observe that while exact methods generally yield more precise solutions in terms of iterations, our proposed inexact approach is competitive and requires fewer MINRES iterations to achieve convergence. Moreover, the adaptive methods achieve a good balance across all metrics making them a viable option in practice.

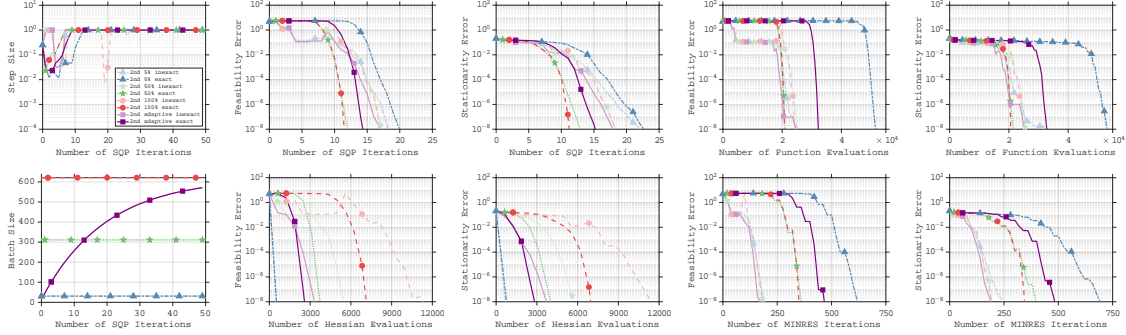


Figure 5: Comparison of **exact** and **inexact** second-order (**2nd**) line search SQP methods with adaptive step size (**adaptive**) and line search (**LS**) schemes on the **australian** dataset [17]. True function and gradient information and inexact Hessian information (batch sizes **5%**, **50%**, **100%** of total sample size, and **adaptive**).

9 Final Remarks

This paper introduces a novel second-order line search sequential quadratic programming (SQP) method designed for solving general nonlinear equality constrained optimization problems in both deterministic and stochastic settings. Our proposed methodology not only capitalizes on the global convergence properties of classical line search SQP methods, but also alleviates the adverse Maratos effect by selectively employing a carefully designed modified line search, and is endowed with local superlinear convergence guarantees without the need for a priori knowledge of problem-specific constants. For large-scale problems, we have developed an inexact matrix-free variant, which integrates the MINRES method, offering a compromise between efficiency and accuracy. Numerical experiments conducted on the CUTEst test set and logistic regression problems demonstrate the efficacy, efficiency and robustness of the proposed method. Finally, we highlight the adaptivity of our algorithm that makes our approach not only theoretically appealing but also practical.

References

- [1] Sigrún Andradóttir. Simulation optimization. *Handbook of simulation*, pages 307–334, 1998.
- [2] Albert S Berahas, Raghu Bollapragada, and Jorge Nocedal. An investigation of newton-sketch and subsampled newton methods. *Optimization Methods and Software*, 35(4):661–680, 2020.

- [3] Albert S Berahas, Raghu Bollapragada, and Baoyu Zhou. An adaptive sampling sequential quadratic programming method for equality constrained stochastic optimization. *arXiv preprint arXiv:2206.00712*, 2022.
- [4] Albert S Berahas, Liyuan Cao, and Katya Scheinberg. Global convergence rate analysis of a generic line search algorithm with noise. *SIAM Journal on Optimization*, 31(2):1489–1518, 2021.
- [5] 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. *Mathematics of Operations Research*, 2023.
- [6] Albert S Berahas, Frank E Curtis, Daniel Robinson, and Baoyu Zhou. Sequential Quadratic Optimization for Nonlinear Equality Constrained Stochastic Optimization. *SIAM Journal on Optimization*, 31(2):1352–1379, 2021.
- [7] 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.
- [8] 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 preprint arXiv:2301.00477*, 2023.
- [9] Lorenz T Biegler, Omar Ghattas, Matthias Heinkenschloss, and Bart van Bloemen Waanders. Large-scale pde-constrained optimization: an introduction. In *Large-Scale PDE-Constrained Optimization*, pages 3–13. Springer, 2003.
- [10] Raghu Bollapragada, Richard H Byrd, and Jorge Nocedal. Exact and inexact subsampled newton methods for optimization. *IMA Journal of Numerical Analysis*, 39(2):545–578, 2019.
- [11] Ingrid Bongartz, Andrew R Conn, Nick Gould, and Ph L Toint. CUTE: Constrained and unconstrained testing environment. *ACM Transactions on Mathematical Software*, 21(1):123–160, 1995.
- [12] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM review*, 60(2):223–311, 2018.
- [13] Richard H Byrd, Frank E Curtis, and Jorge Nocedal. An inexact SQP method for equality constrained optimization. *SIAM Journal on Optimization*, 19(1):351–369, 2008.

- [14] Richard H Byrd, Robert B Schnabel, and Gerald A Shultz. A trust region algorithm for nonlinearly constrained optimization. *SIAM Journal on Numerical Analysis*, 24(5):1152–1170, 1987.
- [15] Coralia Cartis and Katya Scheinberg. Global convergence rate analysis of unconstrained optimization methods based on probabilistic models. *Mathematical Prog.*, 169:337–375, 2018.
- [16] RM Chamberlain, MJD Powell, C Lemarechal, and HC Pedersen. The watchdog technique for forcing convergence in algorithms for constrained optimization. *Algorithms for constrained minimization of smooth nonlinear functions*, pages 1–17, 1982.
- [17] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, 2011.
- [18] Nilanjan Chatterjee, Yi-Hau Chen, Paige Maas, and Raymond J Carroll. Constrained maximum likelihood estimation for model calibration using summary-level information from external big data sources. *Journal of the American Statistical Association*, 111(513):107–117, 2016.
- [19] Zhongwen Chen, Yu-Hong Dai, and Jiangyan Liu. A penalty-free method with super-linear convergence for equality constrained optimization. *Computational Optimization and Applications*, 76:801–833, 2020.
- [20] Zhongwen Chen, Yu-Hong Dai, and Tauyou Zhang. A line search penalty-free sqp method for equality-constrained optimization without maratos effect. *IMA Journal of Numerical Analysis*, 42(4):3771–3802, 2022.
- [21] Sou-Cheng T Choi, Christopher C Paige, and Michael A Saunders. Minres-qlp: A krylov subspace method for indefinite or singular symmetric systems. *SIAM Journal on Scientific Computing*, 33(4):1810–1836, 2011.
- [22] Thomas F Coleman and Andrew R Conn. Nonlinear programming via an exact penalty function: asymptotic analysis. *Mathematical programming*, 24(1):123–136, 1982.
- [23] Andrew R Conn, Nicholas IM Gould, and Philippe L Toint. *Trust region methods*. SIAM, 2000.
- [24] 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*, pages 1–53, 2023.
- [25] Frank E Curtis, Daniel P Robinson, and Baoyu Zhou. Inexact sequential quadratic optimization for minimizing a stochastic objective function subject to deterministic nonlinear equality constraints. *arXiv preprint arXiv:2107.03512*, 2021.

- [26] Frank E Curtis, Daniel P Robinson, and Baoyu Zhou. Sequential quadratic optimization for stochastic optimization with deterministic nonlinear inequality and equality constraints. *arXiv preprint arXiv:2302.14790*, 2023.
- [27] Kalyanmoy Deb. *Optimization for engineering design: Algorithms and examples*. Learning Ltd., 2012.
- [28] Elizabeth D Dolan and Jorge J Moré. Benchmarking optimization software with performance profiles. *Mathematical programming*, 91:201–213, 2002.
- [29] Murat A Erdogdu and Andrea Montanari. Convergence rates of sub-sampled newton methods. *Advances in Neural Information Processing Systems*, 28, 2015.
- [30] Yuchen Fang, Sen Na, Michael W Mahoney, and Mladen Kolar. Fully stochastic trust-region sequential quadratic programming for equality-constrained optimization problems. *arXiv preprint arXiv:2211.15943*, 2022.
- [31] Roger Fletcher. Second order corrections for non-differentiable optimization. In *Numerical Analysis: Proceedings of the 9th Biennial Conference Held at Dundee, Scotland, June 23–26, 1981*, pages 85–114. Springer, 2006.
- [32] David Chin-Lung Fong and Michael Saunders. Cg versus minres: An empirical comparison. *Sultan Qaboos University Journal for Science [SQUJS]*, 17(1):44–62, 2012.
- [33] Michael C Fu. Optimization via simulation: A review. *Annals of OR*, 53:199–247, 1994.
- [34] Masao Fukushima. A successive quadratic programming algorithm with global and superlinear convergence properties. *Mathematical Programming*, 35(3):253–264, 1986.
- [35] Daniel Gabay. Reduced quasi-newton methods with feasibility improvement for nonlinearly constrained optimization. *Algorithms for Constrained Minimization of Smooth Nonlinear Functions*, pages 18–44, 1982.
- [36] Charles J Geyer. Constrained maximum likelihood exemplified by isotonic convex logistic regression. *Journal of the American Statistical Association*, 86(415):717–724, 1991.
- [37] Philip E Gill, Vyacheslav Kungurtsev, and Daniel P Robinson. A stabilized sqp method: superlinear convergence. *Mathematical Programming*, 163(1-2):369–410, 2017.
- [38] Philip E Gill, Walter Murray, and Michael A Saunders. Snopt: An sqp algorithm for large-scale constrained optimization. *SIAM review*, 47(1):99–131, 2005.

- [39] Chao Gu and Detong Zhu. A secant algorithm with line search filter method for nonlinear optimization. *Applied mathematical modelling*, 35(2):879–894, 2011.
- [40] S P Han and Olvi L Mangasarian. Exact penalty functions in nonlinear programming. *Mathematical programming*, 17:251–269, 1979.
- [41] Shih-Ping Han. Superlinearly convergent variable metric algorithms for general nonlinear programming problems. *Mathematical Programming*, 11(1):263–282, 1976.
- [42] Shih-Ping Han. A globally convergent method for nonlinear programming. *Journal of optimization theory and applications*, 22(3):297–309, 1977.
- [43] Yang Liu and Fred Roosta. A newton-mr algorithm with complexity guarantees for nonconvex smooth unconstrained optimization. *arXiv preprint arXiv:2208.07095*, 2022.
- [44] Nicholas Maratos. Exact penalty function algorithms for finite dimensional and control optimization problems. 1978.
- [45] Pablo Márquez-Neila, Mathieu Salzmann, and Pascal Fua. Imposing hard constraints on deep networks: Promises and limitations. *arXiv: 1706.02025*, 2017.
- [46] David Q Mayne and Elijah Polak. *A superlinearly convergent algorithm for constrained optimization problems*. Springer, 1982.
- [47] José Luis Morales. A numerical study of limited memory BFGS methods. *Applied Mathematics Letters*, 15(4):481–487, 2002.
- [48] Sen Na, Mihai Anitescu, and Mladen Kolar. An adaptive stochastic sequential quadratic programming with differentiable exact augmented lagrangians. *Math. Prog.*, 199(1-2):721–791, 2023.
- [49] Sen Na, Mihai Anitescu, and Mladen Kolar. Inequality constrained stochastic nonlinear optimization via active-set sequential quadratic programming. *Mathematical Prog.*, pages 1–75, 2023.
- [50] Sen Na and Michael W Mahoney. Asymptotic convergence rate and statistical inference for stochastic sequential quadratic programming. *arXiv preprint arXiv:2205.13687*, 2022.
- [51] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer-Verlag New York, New York, 2006.
- [52] Figen Oztoprak, Richard Byrd, and Jorge Nocedal. Constrained optimization in the presence of noise. *SIAM Journal on Optimization*, 33(3):2118–2136, 2023.

- [53] UM Garcia Palomares and Olvi Mangasarian. Superlinearly convergent quasi-newton algorithms for nonlinearly constrained optimization problems. *Mathematical Prog.*, 11(1):1–13, 1976.
- [54] Courtney Paquette and Katya Scheinberg. A stochastic line search method with expected complexity analysis. *SIAM Journal on Optimization*, 30(1):349–376, 2020.
- [55] Mert Pilanci and Martin J Wainwright. Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence. *SIAM Journal on Optimization*, 27(1):205–245, 2017.
- [56] Elijah Polak. *Computational methods in optimization: a unified approach*, volume 77. Academic press, 1971.
- [57] Michael JD Powell. The convergence of variable metric methods for nonlinearly constrained optimization calculations. In *Nonlinear programming 3*, pages 27–63. Elsevier, 1978.
- [58] Michael JD Powell. A fast algorithm for nonlinearly constrained optimization calculations. In *Numerical Analysis: Proceedings of the Biennial Conference Held at Dundee, June 28–July 1, 1977*, pages 144–157. Springer, 2006.
- [59] Michael JD Powell and Y Yuan. A recursive quadratic programming algorithm that uses differentiable exact penalty functions. *Mathematical programming*, 35:265–278, 1986.
- [60] Singiresu S Rao. *Engineering optimization: theory and practice*. John Wiley & Sons, 2019.
- [61] 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.
- [62] Farbod Roosta-Khorasani and Michael W Mahoney. Sub-sampled newton methods. *Mathematical Programming*, 174(1):293–326, 2019.
- [63] Soumava Kumar Roy, Zakaria Mhammedi, and Mehrtash Harandi. Geometry aware constrained optimization techniques for deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4469, 2018.
- [64] Klaus Schittkowski. The nonlinear programming method of wilson, han, and powell with an augmented lagrangian type line search function: Part 1: Convergence analysis. *Numerische Mathematik*, 38:83–114, 1982.
- [65] Klaus Schittkowski. *Nonlinear programming codes: information, tests, performance*, volume 183. Springer Science & Business Media, 2012.

- [66] Jiahao Shi and James C Spall. Sqp-based projection spsa algorithm for stochastic optimization with inequality constraints. In *2021 American control conference (ACC)*, pages 1244–1249. IEEE, 2021.
- [67] Suvrit Sra, Sebastian Nowozin, and Stephen J Wright. *Optimization for machine learning*. Mit Press, 2012.
- [68] Stefan Ulbrich. On the superlinear local convergence of a filter-sqp method. *Mathematical Programming*, 100:217–245, 2004.
- [69] Stanislav Uryasev and Panos M Pardalos. *Stochastic optimization: algorithms and applications*, volume 54. Springer Science & Business Media, 2013.
- [70] Vladimir Vapnik. Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4, 1991.
- [71] Avi Vardi. A trust region algorithm for equality constrained minimization: convergence properties and implementation. *SIAM Journal on Numerical Analysis*, 22(3):575–591, 1985.
- [72] Maria Vrakopoulou, Johanna L Mathieu, and Göran Andersson. Stochastic optimal power flow with uncertain reserves from demand response. In *2014 47th Hawaii International Conference on System Sciences*, pages 2353–2362. IEEE, 2014.
- [73] Andreas Wächter and Lorenz T Biegler. Line search filter methods for nonlinear programming: Motivation and global convergence. *SIAM Journal on Optimization*, 16(1):1–31, 2005.
- [74] Edmund Taylor Whittaker and George Neville Watson. *A course of modern analysis: an introduction to the general theory of infinite processes and of analytic functions; with an account of the principal transcendental functions*. University press, 1920.
- [75] Robert B Wilson. A simplicial algorithm for concave programming. *Ph. D. Dissertation, Graduate School of Business Administration*, 1963.
- [76] Allen J Wood, Bruce F Wollenberg, and Gerald B Sheblé. *Power generation, operation, and control*. John Wiley & Sons, New Jersey, USA, 2013.
- [77] Stephen Wright, Jorge Nocedal, et al. Numerical optimization. *Springer Science*, 35(67-68):7, 1999.
- [78] William T Ziemba and Raymond G Vickson. *Stochastic optimization models in finance*. Aca. Press, 2014.

A Appendix: Proof of Lemma 3.5

Proof. It follows by Assumptions 3.1 and 3.4, the triangular inequality that

$$\begin{aligned} \|g_k - \bar{g}_k\| &= \left\| \left(\frac{N - |S_k^g|}{N |S_k^g|} \right) \sum_{i \in S_k^g} \nabla f_i(x_k) - \frac{1}{N} \sum_{i \in [N] \setminus S_k^g} \nabla f_i(x_k) \right\| \\ &\leq \left(\frac{N - |S_k^g|}{N |S_k^g|} \right) \sum_{i \in S_k^g} \|\nabla f_i(x_k)\| + \frac{1}{N} \sum_{i \in [N] \setminus S_k^g} \|\nabla f_i(x_k)\| \\ &\leq 2 \left(\frac{N - |S_k^g|}{N} \right) (\kappa_1 + \beta_1 \|\nabla f(x)\|) \leq 2 \left(\frac{N - |S_k^g|}{N} \right) (\kappa_1 + \beta_1 \kappa_g). \end{aligned}$$

Note that the inequalities still hold for any types of norm. Similarly,

$$\begin{aligned} \|H_k - \bar{H}_k\| &\leq 2 \left(\frac{N - |S_k^H|}{N} \right) (\kappa_2 + \beta_2 \|\nabla^2 f(x)\|) \leq 2 \left(\frac{N - |S_k^H|}{N} \right) (\kappa_2 + \beta_2 \kappa_H). \\ |f_k - \bar{f}_k| &\leq 2 \left(\frac{N - |S_k^f|}{N} \right) (\kappa_0 + \beta_0 |f(x)|) \leq 2 \left(\frac{N - |S_k^f|}{N} \right) (\kappa_0 + \beta_0 \kappa_f). \end{aligned}$$

□

B Appendix: Proof of Theorem 5.3

Proof. Replacing stochastic quantities $\bar{d}_k, \bar{\delta}_k$ with deterministic quantities d_k, δ_k , it follows by (5.4) that $\|\delta_k\| \leq 2\kappa_{J^\dagger} \kappa_W \|d_k\|$. When $\|d_k\| \leq \kappa_d$, it then follows by assumptions 5.1 that

$$\begin{aligned} \|w_k - w^*\| &\leq 2\mu_M^{-1} \left\| \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\| \leq 2\mu_M^{-1} \left\| M_k M_k^{-1} \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\| \\ &\leq 2\mu_M^{-1} \kappa_M \left\| \begin{bmatrix} d_k \\ \delta_k \end{bmatrix} \right\| \leq 2\mu_M^{-1} \kappa_M (1 + 2\kappa_{J^\dagger} \kappa_W) \|d_k\|. \end{aligned} \quad (\text{B.1})$$

If $\|d_0\| \leq \frac{1}{2\mu_M^{-1} \kappa_M (1 + 2\kappa_{J^\dagger} \kappa_W) \Lambda L_W}$, we have $\|w_0 - w^*\| \leq \frac{1}{\Lambda L_W}$ by (B.1). It then follows by Cauchy-Schwartz inequality and assumption 3.2 that when $r = \frac{1}{\Lambda L_W}$,

$$\begin{aligned} \|w_{k+1} - w^*\| &= \left\| w_k - w^* - M_k^{-1} \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\| \\ &\leq \|M_k^{-1}\| \left\| M_k (w_k - w^*) - \begin{bmatrix} g_k + J_k^T y_k \\ c_k \end{bmatrix} \right\| \leq \frac{\Lambda L_W}{2} \|w_k - w^*\|^2. \end{aligned} \quad (\text{B.2})$$

We then use induction to prove $\|w_k - w^*\| \leq \frac{1}{\Lambda L_W}$. The base case holds trivially, suppose that the statement holds for iteration k , then it follows from (B.2) that

$$\|w_{k+1} - w^*\| \leq \frac{\Lambda L_W}{2} \|w_k - w^*\|^2 \leq \frac{\Lambda L_W}{2} \|w_k - w^*\| \|w_k - w^*\| \leq \frac{1}{2} \|w_k - w^*\|. \quad (\text{B.3})$$

Hence, $\|w_k - w^*\| \rightarrow 0$ at a Q-quadratic rate. \square

C Appendix: Proof of Theorem 5.9

Proof. We use induction to show that for all k ,

$$\|w_k - w^*\| \leq \frac{1}{3\Lambda L_W} \xi_k. \quad (\text{C.1})$$

By condition (i), the initial gradient approximation satisfies $\|g_0 - \bar{g}_0\| \leq \min \left\{ \frac{\mu_M}{12\Lambda L_W}, \frac{\kappa_d}{2\Lambda} \right\}$. The base case, $k = 0$, is therefore satisfied. Moreover, by Lemma 5.4, it follows that $\|w_0 - w^*\| \leq \frac{1}{3\Lambda L_W}$. Define sequences $\{\xi_k\}_{k=0}^\infty$ and $\{\rho_k\}_{k=0}^\infty$ as follows:

$$\xi_{k+1} = \max \left\{ \xi_k \rho_k, \beta_{k+1}^{(k+1)/4} \right\}, \quad \xi_0 = 1, \quad \rho_k = \frac{\xi_k}{6} + \frac{1}{4} \frac{N - |S_k^H|}{N - |S_0^H|} + \frac{\beta_k^{k/4}}{2}. \quad (\text{C.2})$$

Assume (C.1) is true for iteration k , and consider iteration $k + 1$,

$$\begin{aligned} & \|w_{k+1} - w^*\| \\ & \leq \frac{\Lambda L_W}{2} \|w_k - w^*\|^2 + \frac{2(\kappa_2 + \mu_2 \kappa_H) \Lambda (N - |S_k^H|)}{N} \|w_k - w^*\| + \frac{2(\kappa_1 + \mu_1 \kappa_g) \beta_k^{k/2}}{12\Lambda L_W (\kappa_1 + \mu_1 \kappa_g)} \\ & \leq \frac{\xi_k}{3\Lambda L_W} \left(\frac{\xi_k}{6} \right) + \frac{\xi_k}{3\Lambda L_W} \left(\frac{1}{4} \frac{N - |S_k^H|}{N - |S_0^H|} \right) + \frac{1}{3\Lambda L_W} \left(\frac{\beta_k^{k/2}}{2} \right) \\ & = \frac{\xi_k}{3\Lambda L_W} \left(\frac{\xi_k}{6} + \frac{1}{4} \frac{N - |S_k^H|}{N - |S_0^H|} + \frac{\beta_k^{k/2}}{2\xi_k} \right) \\ & \leq \frac{\xi_k}{3\Lambda L_W} \left(\frac{\xi_k}{6} + \frac{1}{4} \frac{N - |S_k^H|}{N - |S_0^H|} + \frac{\beta_k^{k/4}}{2} \right) = \frac{\xi_k}{3\Lambda L_W} \rho_k \leq \frac{\xi_{k+1}}{3\Lambda L_W}. \end{aligned}$$

(C.1) is therefore satisfied on iteration $k + 1$. We then use induction to verify that $\xi_k < 1$ for $k \geq 1$. Considering the base case where $k = 1$, we find that

$$\rho_0 = \frac{\xi_0}{6} + \frac{1}{4} + \frac{1}{2} = \frac{11}{12} < 1, \quad \text{and} \quad \xi_1 = \max \left\{ \xi_0 \rho_0, \beta_1^{1/4} \right\} = \max \left\{ \rho_0, \beta_1^{1/4} \right\} < 1.$$

Let us assume that $\xi_k < 1$ for some $k > 1$. By the fact that $\{|S_k^H|\}$ is increasing and $\{\beta_k\}$ is decreasing, we have

$$\xi_{k+1} \leq \max \left\{ \rho_k, \beta_1^{(k+1)/4} \right\} < 1, \quad \rho_k = \frac{\xi_k}{6} + \frac{1}{4} \frac{N - |S_k^H|}{N - |S_0^H|} + \frac{\beta_k^{k/4}}{2} \leq \frac{1}{6} + \frac{1}{4} + \frac{1}{2} = \frac{11}{12} < 1.$$

This substantiates that $\rho_k \leq 11/12$, which implies $\{\xi_k\} \rightarrow 0$ along with the fact that $\beta_{k+1} < \beta_k \leq 1$. Consequently, it follows by (C.2) that $\{\rho_k\} \rightarrow 0$. Based on these findings, we deduce that

$$\lim_{k \rightarrow \infty} \frac{\xi_{k+1}}{\xi_k} = \lim_{k \rightarrow \infty} \max \left\{ \rho_k, \frac{\beta_{k+1}^{(k+1)/4}}{\xi_k} \right\} \leq \lim_{k \rightarrow \infty} \max \left\{ \rho_k, \left(\frac{\beta_{k+1}}{\beta_k} \right)^{k/4} \beta_{k+1}^{1/4} \right\} = 0.$$

□

D Appendix: Proof of Lemma 7.1

Proof. Since \bar{M}_k is nonsingular by assumption 3.1, 4.4, and [51, Lemma 16.1]. Let $\Theta(\bar{M}_k)$ denote the set of eigenvalues of \bar{M}_k and $\mathcal{E}_\lambda(\bar{M}_k)$ denote the eigenspace corresponding to an eigenvalue $\lambda \in \Theta(\bar{M}_k)$. Suppose that $\Theta_{+, \max}, \Theta_{+, \min}$ are the maximum and the minimum of positive element in $\Theta(\bar{M}_k) \setminus \Theta_\perp(\bar{M}_k)$, $\Theta_{-, \max}, \Theta_{-, \min}$ are the negative element in maximum and the minimum in $\Theta(\bar{M}_k) \setminus \Theta_\perp(\bar{M}_k)$. Let $\kappa_i^+ := \frac{\Theta_{+, \max}}{\Theta_{+, \min}}$, $\kappa_i^- := \frac{\Theta_{-, \min}}{\Theta_{-, \max}}$, and $\theta = \max \left\{ \frac{\sqrt{\kappa_i^+} - 1}{\sqrt{\kappa_i^+} + 1}, \frac{\sqrt{\kappa_i^-} - 1}{\sqrt{\kappa_i^-} + 1} \right\}$. The proof then follows directly from [43, Lemma 4] by the fact that \bar{M}_k does not have zero eigenvalues and taking $i = \psi_+$, $j = \psi_+ + \psi_0 + 1$ in its context. □