

Retrospective Approximation Sequential Quadratic Programming for Stochastic Optimization with General Deterministic Nonlinear Constraints

Albert S. Berahas* Raghu Bollapragada^{†‡} Shagun Gupta[†]

May 25, 2025

Abstract

In this paper, we propose a framework based on the Retrospective Approximation (**RA**) paradigm to solve optimization problems with a stochastic objective function and general nonlinear deterministic constraints. This framework sequentially constructs increasingly accurate approximations of the true problems which are solved to a specified accuracy via a deterministic solver, thereby decoupling the uncertainty from the optimization. Such frameworks retain the advantages of deterministic optimization methods, such as fast convergence, while achieving the optimal performance of stochastic methods without the need to redesign algorithmic components. For problems with general nonlinear equality constraints, we present a framework that can employ any deterministic solver and analyze its theoretical work complexity. We then present an instance of the framework that employs a deterministic Sequential Quadratic Programming (**SQP**) method and that achieves optimal complexity in terms of gradient evaluations and linear system solves for this class of problems. For problems with general nonlinear constraints, we present an **RA**-based algorithm that employs an **SQP** method with robust subproblems. Finally, we demonstrate the empirical performance of the proposed framework on multi-class logistic regression problems and benchmark instances from the CUTEst test set, comparing its results to established methods from the literature.

1 Introduction

In this paper, we propose a framework for solving optimization problems with a stochastic objective function and general nonlinear deterministic constraints of the form,

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

*Department of Industrial and Operations Engineering, University of Michigan. (albertberahas@gmail.com)

[†]Operations Research and Industrial Engineering Program, University of Texas at Austin. (raghu.bollapragada@utexas.edu, shagungupta@utexas.edu)

[‡]Corresponding author.

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $c_E : \mathbb{R}^n \rightarrow \mathbb{R}^{m_E}$ and $c_I : \mathbb{R}^n \rightarrow \mathbb{R}^{m_I}$ are continuously differentiable functions. We consider two forms for the objective function,

$$f(x) = \frac{1}{|\mathcal{S}|} \sum_{\xi \in \mathcal{S}} F(x, \xi) \quad (1.2), \quad \text{and} \quad f(x) = \mathbb{E}[F(x, \xi)], \quad (1.3)$$

where the objective function (1.2) results in a deterministic finite-sum problem over a finite dataset $\mathcal{S} = \{\xi_1, \xi_2, \dots\}$ with $F : \mathbb{R}^n \times \mathcal{S} \rightarrow \mathbb{R}$, and the objective function (1.3) results in an expectation minimization problem over the random variable ξ with the associated probability space $(\Xi, \Omega, \mathcal{P})$, $F : \mathbb{R}^n \times \Xi \rightarrow \mathbb{R}$ and $\mathbb{E}[\cdot]$ denotes the expectation with respect to \mathcal{P} . Such problems arise in various applications including machine learning [35, 36], computational physics for climate modeling [32, 63], portfolio optimization [49], multistage optimization [58], and optimal power flow [59].

Methods for solving (1.1) with a deterministic objective function have been extensively studied; see e.g., [4, 42]. A key challenge in solving problem (1.1) is balancing the two possibly competing goals of minimizing the objective function and satisfying the constraints at each step of the method. Recently, stochastic methods have been developed for minimizing (1.2) and (1.3), although many of these methods are limited to problems with only equality constraints. These include penalty-based techniques [12, 39, 55, 56, 62], projection methods [27, 33, 34], Sequential Quadratic Programming (SQP) [1, 2, 19–21, 25, 37, 38, 45, 54], and Interior Point methods [15–17]. These methods retain the search direction computation mechanism of the deterministic approaches by using an approximation of the objective function gradient. However, these approaches redesign several components of the deterministic algorithms to make them suitable for stochastic settings and, as a result, lose some of the associated benefits such as automatic step size selection, and result in slower convergence.

The Retrospective Approximation (RA) framework is an alternate approach for solving stochastic optimization problems [13, 47]. The framework iteratively constructs deterministic subsampled approximations of increasing accuracy and solves each one to a certain specified accuracy. As a result, the approach decouples the uncertainty from the optimization process, enabling the use of deterministic methods to solve stochastic problems. This framework has been successfully employed in various areas, including stochastic root finding [14, 46], two-stage stochastic optimization [48], simulation optimization [22, 31], and optimal design problems [50, 51, 57, 61]. More recently, it has also been proposed for unconstrained stochastic optimization [40, 41]. As compared to deterministic methods, these methods are more efficient due to their use of less expensive subsampled function (and gradient) approximations in the initial stages, and as compared to fully stochastic methods, these methods achieve the optimal iteration and work complexity, can attain high accuracy solutions and can take advantage of developments in deterministic methods.

In this work, we propose frameworks based on the RA paradigm to solve stochastic optimization problems with general deterministic nonlinear constraints, and employ the SQP method, a state-of-the-art technique for optimization problems with general nonlinear constraints [42], as our deterministic optimization method. We propose: (1) a framework for problems with only equality constraints that can incorporate any deterministic method, and an instance that employs a line search SQP method [30, 52]; and, (2) an RA based algorithm for general nonlinear constraints that utilizes the robust-SQP method [6]. We analyze the theoretical convergence of the two frameworks and evaluate their empirical performance on regularized logistic regression multi-class classification problems as well as the S2MPJ CUTEst problem set [28], comparing them with established methods from the literature.

1.1 Constrained Stochastic Optimization

In this section, we present the fundamentals of constrained optimization for smooth problems, review the literature on constrained stochastic optimization, introduce the SQP method, and outline our contributions relative to existing work. We define the Lagrangian for (1.1) with $\lambda_E \in \mathbb{R}^{m_E}$ and $\lambda_I \in \mathbb{R}^{m_I}$ as $\mathcal{L}(x, \lambda_E, \lambda_I) = f(x) + c_E(x)^T \lambda_E + c_I(x)^T \lambda_I$. The Karush-Kuhn-Tucker (KKT) conditions state that x^* is a first-order stationary point if there exist λ_E^* and λ_I^* such that,

$$\nabla_x \mathcal{L}(x^*, \lambda_E^*, \lambda_I^*) = 0, \quad c_E(x^*) = 0, \quad c_I(x^*) \leq 0, \quad c_I(x^*) \odot \lambda_I^* = 0, \quad \lambda_I^* \geq 0, \quad (1.4)$$

where $c_I(x^*) \odot \lambda_I^*$ denotes the element-wise product of the two vectors.

While many powerful methods for obtaining a (or an approximate) KKT point in the deterministic optimization setting have existed for decades, see e.g., [42, Chapters 12-19], counterparts for constrained problems with stochastic objective functions are significantly fewer. Penalty-based methods [12, 39, 55, 56, 62] were among the first methods extended to this problem class due to their intuitive design and ease of implementation. However, the empirical performance of penalty-based methods is sensitive to the choice of penalty function and parameter. Projection methods [27, 33, 34] have also been proposed for this problem class for instances where the projection mapping onto the feasible region is easily computable. Moreover, stochastic interior points methods [15–17, 23] have also been recently developed. While these methods have sound theoretical results and show promising empirical performance, they are primarily applicable to bound constraints and require strong assumptions to handle the general nonlinear constraint setting. Finally, SQP methods have been proposed and are discussed in detail below.

SQP methods [42, Chapter 18] are iterative methods that update the current estimate of the solution by computing a search direction that minimizes a quadratic model of the objective function subject to a linear approximation of the constraints, i.e., at each iterate x ,

$$\begin{aligned} \min_{d \in \mathbb{R}^n} \quad & \frac{1}{2} d^T H d + \nabla f(x)^T d \\ \text{s.t.} \quad & c_E(x) + \nabla c_E(x)^T d = 0 \\ & c_I(x) + \nabla c_I(x)^T d \leq 0, \end{aligned} \quad (1.5)$$

where $d \in \mathbb{R}^n$ is the search direction and $H \in \mathbb{R}^{n \times n}$ is an approximation of the Hessian of the Lagrangian. Several methods have been developed for constrained stochastic optimization settings that incorporate an approximation of the gradient of the objective function in (1.5). Early methods were primarily designed for problems with only equality constraints [2, 3, 19, 45]. Building upon these methods, subsequent works in the equality-constrained setting leveraged features from deterministic SQP methods such as inexact solutions to (1.5) [1, 20], adaptive test to control the gradient approximation accuracy [1, 37], and line search and trust-region variants [3, 25]. All aforementioned methods use, modify and adapt mechanisms from the deterministic SQP methodology. That said, to overcome the challenges that arise due to the stochastic nature of (1.1), most advanced features of deterministic SQP methods such as quasi-Newton Hessian approximations, inexact search direction solutions, and line search mechanisms are omitted or used only in solitude, or strong assumptions on the behavior of the algorithm or problem are imposed. By incorporating a deterministic SQP solver within the RA framework, one can use all the advanced SQP features in harmony when solving constrained stochastic optimization problems.

When solving equality-constrained stochastic problems using SQP, computing a search direction via (1.5) reduces to solving a linear system of equations. As a result, the two main computational

costs of **SQP** methods in this setting are the number of gradient evaluations and the number of linear systems to be solved. Previously proposed methods such as [2, 3, 19, 45] achieve the optimal complexity for gradient evaluations $\mathcal{O}(\epsilon^{-4})$ but exhibit high complexity in terms of the number of linear system solves $\mathcal{O}(\epsilon^{-4})$. In contrast, the method in [1] achieves the optimal complexity in terms of the number of linear system solves $\mathcal{O}(\epsilon^{-2})$ but incurs a higher complexity in terms of gradient evaluations $\mathcal{O}(\epsilon^{-2(1+\nu)})$, where $\nu > 1$. Our proposed **RA**-based **SQP** method for stochastic optimization with deterministic equality constraints achieves optimal complexity in terms of both costs.

The general equality and inequality constrained setting presents additional challenges in the design of **SQP** methods primarily, due to the difficulty in solving the **SQP** subproblem (1.5) while ensuring feasibility. Few **SQP**-based methods have been developed to handle both constraint types with a stochastic objective function. These include an active-set method [38], a decomposition approach [21], and a two-step method [54]. The integration of deterministic **SQP** methods within the **RA** framework allows for the incorporation of general nonlinear constraints. However, the choice of method and other aspects of **RA** must be carefully considered to ensure strong empirical performance in the presence of inequality constraints, as will be discussed in Section 3.

1.2 Retrospective Approximation (RA)

In this section, we describe the **RA** framework and our approach for applying it to constrained stochastic optimization. The framework consists of nested outer and inner loops. In each outer iteration, a deterministic subsampled problem is constructed as a sample average approximation of the objective function in (1.1),

$$\begin{aligned} \min_{x \in \mathbb{R}^n} F_S(x) \quad s.t. \quad c_E(x) = 0, \quad c_I(x) \leq 0, \\ \text{where} \quad F_S(x) = \frac{1}{|S|} \sum_{\xi \in S} F(x, \xi) \quad \text{and} \quad g_S(x) = \frac{1}{|S|} \sum_{\xi \in S} \nabla F(x, \xi), \end{aligned} \tag{1.6}$$

$S \subseteq \mathcal{S}$ for (1.2) and S is a set of independent and identically distributed (i.i.d.) samples drawn from the probability space $(\Xi, \Omega, \mathcal{P})$ for (1.3). The deterministic subsampled problem (1.6) is then solved to a specified accuracy in the inner loop. The **RA** approach assumes greater control over the samples compared to other works in the constrained stochastic optimization literature that are in the fully stochastic regime and use only noisy gradient approximations [2, 17, 45]. The approach requires objective function (and gradient) evaluations using the same sample set across iterates and the ability to control the sample size (accuracy), which is feasible in most practical applications discussed earlier. Once the inner loop terminates, a new batch of samples is obtained, and the formulated subsampled problem is warm-started at the output of the previous outer iteration. The framework is detailed in Framework 1.1 and the following remark.

Framework 1.1 Retrospective Approximation for Constrained Stochastic Optimization

Inputs: Initial iterate $x_{0,0}$, batch size sequence $\{|S_k|\}$, termination test sequence $\{\mathcal{T}_k\}$.

```
1: for  $k = 0, 1, 2, \dots$  do
2:   Construct subsampled problem (1.6) from sample set  $S_k$ 
3:   for  $j = 0, 1, 2, \dots$  do
4:     if  $\mathcal{T}_k$  is satisfied;  $N_k = j$ , break
5:     Update  $x_{k,j+1}$ 
6:   end for
7:   Set  $x_{k+1,0} = x_{k,N_k}$ 
8: end for
```

Remark 1.1. We make the following remarks about Framework 1.1.

1. **Indexing:** For quantities with double indexing, e.g., $x_{k,j}$, the first index represents the outer iteration and the second denotes the inner iteration. Quantities with a single index correspond solely to an outer iteration.
2. **Outer Loop:** One obtains a sample set S_k of pre-specified batch size and formulates the deterministic subsampled problem in Line 2, independent of the current iterate $x_{k,0}$. The subsampled problem is solved using a deterministic solver initialized at $x_{k,0}$ in the inner loop.
3. **Inner Loop (Lines 4-5):** The inner loop terminates when the specified accuracy for the subsampled problem is achieved, as tested in Line 4 using the termination criterion \mathcal{T}_k for the outer iteration k . If the criterion is not met, a new iterate $x_{k,j+1}$ is obtained in Line 5. Since the goal is to solve a deterministic subsampled problem to a prescribed accuracy, one can employ deterministic optimization methods to update the iterate within the inner loop.

Framework 1.1 requires the specification of three key components: (1) the batch size sequence $\{|S_k|\}$, (2) the termination test $\{\mathcal{T}_k\}$, and (3) the update mechanism in the inner loop. The sequence $\{|S_k|\}$ is an increasing sequence. In previous work [21, 40, 41, 57], this sequence has been predefined. In this work, an adaptive strategy tailored to constrained optimization settings is employed to determine the batch size at the start of each outer iteration. In earlier works, the termination test \mathcal{T}_k either used a predefined sequence or adaptively determined N_k , the number of inner loop iterations, at the start of each outer iteration [50, 51, 57, 61]. Our goal is to adaptively set the required solution accuracy for each subsampled problem, similar to [40, 41]. The main challenge lies in determining how to measure solution accuracy. For unconstrained smooth optimization, the subsampled gradient norm is the natural choice, as done in [40, 41]. However, for constrained optimization, this estimation is more complex, with various options depending on the type of constraints (equality or general) and the update method in the inner loop. In this work, we examine these options across different scenarios. Due to the adaptive nature of the batch size and termination criterion, we are also able to provide detailed convergence rates and complexity results for the number of gradient evaluations and inner loop iterations, unlike previous works [50, 57, 61]. Lastly, we use deterministic SQP methods in the inner loop to update the iterate, unlike earlier work that employed projection-based methods to implicitly ensure feasibility [50, 51, 57, 61], and solve problems with general nonlinear constraints.

1.3 Contributions

We propose frameworks for solving optimization problems with a stochastic objective function and deterministic general nonlinear constraints (1.1) based on RA. Our contributions can be summarized as follows.

1. We propose a framework for solving stochastic optimization problems with general nonlinear equality constraints that can incorporate any deterministic solver (see Framework 2.1). Within this framework, we introduce adaptive conditions for selecting the batch size sequence and the termination test sequence, which lead to fast linear convergence of the outer iterates (Theorem 2.5). Furthermore, assuming that the deterministic solver exhibits sublinear convergence, we establish complexity guarantees for the finite-sum objective (1.2) that match those of deterministic methods applied to the full objective (Theorem 2.8). For the expectation-based problem (1.3), the framework achieves optimal complexity with respect to the number of gradient evaluations, $\mathcal{O}(\epsilon^{-4})$, and the number of inner loop iterations, $\mathcal{O}(\epsilon^{-2})$ (Theorem 2.8).
2. We propose a variant of the framework for equality constraints that uses the deterministic SQP method (see Algorithm 2.2) in the inner loop for problems with equality constraints. This variant introduces alternate customized conditions for the batch size sequence and termination criteria that leverage information from the SQP method to enhance empirical performance. The algorithm achieves optimal theoretical complexity results for number of gradient evaluations, $\mathcal{O}(\epsilon^{-4})$, and number of SQP linear system solves, $\mathcal{O}(\epsilon^{-2})$ (Theorem 2.10), and also allows for the addition of advanced techniques from deterministic methods to enhance empirical performance in stochastic settings. These include employing inexact SQP subproblem solutions and the use of quasi-Newton Hessian approximations.
3. We propose an RA based algorithm for stochastic problems with general (equality and inequality) nonlinear constraints (see Algorithm 3.1), an area with relatively limited literature [21, 38, 54]. Our proposed algorithm employs deterministic robust-SQP subproblems [6] within the inner loop, and we present two specific instances of such subproblems. Both lead to similarly fast linear convergence in the outer iterations (Theorem 3.5), but differ in practical performance due to variations in computational cost and subproblem accuracy.
4. We illustrate the empirical performance of the proposed methods (Algorithm 2.2 and Algorithm 3.1) on regularized logistic regression multi-class classification problems and the S2MPJ CUTEst problem set [28]. We show the performance benefits of utilizing deterministic solvers and the use of techniques such as quasi-Newton Hessian approximations and inexact subproblem solves in harmony with the stochasticity. We also compare with algorithms from the literature.

1.4 Notation

Let \mathbb{R} denote the set of real numbers, \mathbb{R}^n denote the set of n dimensional real vectors, and $\mathbb{R}^{n \times m}$ denote the set of n by m real matrices. Unless specified otherwise, $\|\cdot\|$ represents the Euclidean norm of a vector and the Frobenius norm of a matrix, $|\cdot|$ denotes the absolute value of a real number and the cardinality of a set, and $[\cdot]_+$ denotes the positive elements of a vector. The ceiling function is denoted as $\lceil \cdot \rceil$. The element wise product of $a, b \in \mathbb{R}^n$ is denoted as $a \odot b$. The vector of all ones of dimension n is denoted as e_n , and I_n denotes the identity matrix of dimension n . Expectation and variance with respect to the distribution \mathcal{P} are denoted as $\mathbb{E}[\cdot]$ and $\text{Var}(\cdot)$, respectively.

1.5 Paper Organization

The paper is organized as follows. In Section 2, we formalize the framework for equality constrained problems. We describe the framework for general deterministic solvers in Subsection 2.1 and provide theoretical convergence and complexity results in Subsection 2.2. We then describe the SQP variant of the framework for equality constrained problems in Subsection 2.3 and the corresponding theoretical analysis in Subsection 2.4. In Section 3, we present an RA based method for problems with general nonlinear constraints using the robust-SQP method, describe the algorithm in Subsection 3.1, and present the theoretical analysis in Subsection 3.2. In Section 4, we illustrate the empirical performance of the proposed algorithms on multi-class classification logistic regression and the S2MPJ CUTEst problem set, and compare with methods from the literature. Finally, we conclude with final remarks in Section 5.

2 Equality Constrained Problems

In this section, we consider problem (1.1) with only equality constraints, i.e.,

$$\min_{x \in \mathbb{R}^n} f(x) \quad \text{s.t.} \quad c(x) = 0, \quad (2.1)$$

where the constraint function c_E is denoted by $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$. We describe the proposed RA based framework in Subsection 2.1 for general deterministic solvers and provide the corresponding convergence and complexity results in Subsection 2.2. We then focus on using SQP methods [1, 2, 8] in the inner loop in Subsection 2.3 with the corresponding theoretical analysis in Subsection 2.4.

We introduce some additional notation for this section. The Lagrangian of problem (2.1) is $\mathcal{L}(x, \lambda) = f(x) + \lambda^T c(x)$ where $\lambda \in \mathbb{R}^m$ is the vector of dual variables. The KKT conditions (1.4) in the equality constraint setting reduce to

$$T(x, \lambda) = \begin{bmatrix} \nabla_x \mathcal{L}(x, \lambda) \\ c(x) \end{bmatrix} = \begin{bmatrix} \nabla f(x) + \nabla c(x) \lambda \\ c(x) \end{bmatrix} = 0,$$

where x^* is a first-order stationary point if there exists $\lambda^* \in \mathbb{R}^m$ such that $T(x^*, \lambda^*) = 0$. The corresponding quantities for subsampled problem (1.6) with sample set S are denoted as

$$\mathcal{L}_S(x, \lambda) = F_S(x) + \lambda^T c(x) \quad \text{and} \quad T_S(x, \lambda) = \begin{bmatrix} \nabla_x \mathcal{L}_S(x, \lambda) \\ c(x) \end{bmatrix}.$$

We denote the iterate (primal variable) and the dual variable at outer iteration k and inner iteration j as $x_{k,j}$ and $\lambda_{k,j}$, respectively. For convenience, we use the shorthand notation $c_{k,j} = c(x_{k,j})$ and $J_{k,j} = \nabla c(x_{k,j})^T$.

2.1 Retrospective Approximation General Framework

In this section, we describe the proposed RA based framework with nested inner and outer loops (Framework 1.1) to solve stochastic problems with deterministic equality constraints. In outer iteration k , we obtain a batch of samples S_k , independent of the current iterate, and define the subsampled problem (1.6). The subsampled problem is then solved up to a specified accuracy in the inner loop. For now, we assume that the batch size sequence is pre-specified and later introduce an adaptive strategy to select the batch size.

The primal iterate in the inner loop is initialized as the output of the previous outer iteration, i.e., $x_{k+1,0} = x_{k,N_k}$ where N_k is the number of inner iterations at outer iteration k . However, the dual variable iterate is initialized as

$$\lambda_{k,0} = \begin{cases} \lambda_{k-1,N_{k-1}} & \text{if } \text{Carry-Over}, \\ -(J_{k,0}J_{k,0}^T)^{-1}J_{k,0}g_{S_k}(x_{k,0}) & \text{if } \text{Reinitialize}. \end{cases} \quad (2.2)$$

The **Carry-Over** option reuses the dual variable from the previous outer iteration, requiring no additional computation. In contrast, the **Reinitialize** option solves the least squares problem $\lambda_{k,0} = \operatorname{argmin}_{\lambda \in \mathbb{R}^m} \|T_{S_k}(x_{k,0}, \lambda)\|^2$, which minimizes the KKT error for the subsampled problem at $x_{k,0}$ to obtain a dual variable estimate for the inner loop. The two options represent extreme choices for initialization. While the **Reinitialize** option provides a better dual variable for the subsampled problem, improving the initialization for the inner loop, it also incurs additional computational cost. Therefore, the choice between the two should be made based on the specific problem requirements. Alternatively, an inexact solution to the least squares problem is acceptable if it performs at least as well as the **Carry-Over** option, i.e., $\|T_{S_k}(x_{k,0}, \lambda_{k,0})\| \leq \|T_{S_k}(x_{k,0}, \lambda_{k-1,N_{k-1}})\|$.

In outer iteration k , the inner loop terminates when the subsampled problem is solved to a specified accuracy, as determined by termination criterion $\{\mathcal{T}_k\}$ in Framework 1.1 (Line 4). For the case of equality constraints, it is defined as

$$\mathcal{T}_k : \|T_{S_k}(x_{k,j}, \lambda_{k,j})\| \leq \gamma_k \|T_{S_k}(x_{k,0}, \lambda_{k,0})\| + \epsilon_k, \quad (2.3)$$

where $\gamma_k \in [0, 1)$ and $\epsilon_k \geq 0$ are user-defined parameters. This criterion measures solution accuracy in terms of the KKT error, and ensures a reduction in the KKT error of the subsampled problem relative to its initial value while allowing for an additional tolerance ϵ_k . When the initial KKT error is high, the relative nature of (2.3) prevents large deviations from the current solution. A high initial KKT error may indicate that the current iterate and dual variable are either poor solutions to the true problem, or they are good solutions but the chosen sample set is dominated by outliers. In either case, one would like to proceed with caution as in the former, subsequent outer iterations should reaffirm the scenario, leading to a better solution, while in the latter, one would avoid biasing the solution to outlier samples. Conversely, the term ϵ_k prevents unnecessary precision when the initial KKT error is low, which can occur when the solutions to successive subsampled problems are close to each other. In such cases, solving the subsampled problem to high accuracy may bias the iterates towards the solution of the subsampled problem rather than the solution of the true problem.

Any deterministic solver that updates the iterate and dual variable to satisfy \mathcal{T}_k within a finite number of iterations can be used in the inner loop. The complete framework is presented in Framework 2.1 and the remark that follows.

Framework 2.1 Equality Constrained Retrospective Approximation

Inputs: Sequences $\{|S_k|\}$, $\{\gamma_k\} < 1$ and $\{\epsilon_k\} \geq 0$; initial iterate $x_{0,0}$ and dual variable $\lambda_{-1,0}$.

```
1: for  $k = 0, 1, 2, \dots$  do
2:   Construct subsampled problem (1.6) from sample set  $S_k$  and set  $\lambda_{k,0}$  via (2.2)
3:   for  $j = 0, 1, 2, \dots$  do
4:     if (2.3) is satisfied;  $N_k = j$ , break
5:     Update  $x_{k,j+1}$  and  $\lambda_{k,j+1}$ 
6:   end for
7:   Set  $x_{k+1,0} = x_{k,N_k}$ 
8: end for
```

Remark 2.1. We make the following remarks about Framework 2.1, building upon Framework 1.1.

1. **Outer Loop (Line 2):** We obtain a sample set S_k and use it to form the subsampled problem (1.6). The batch size is a user-specified sequence and an adaptive strategy is described later. Next, we initialize the dual variable according to (2.2). The update can follow either *Carry-Over* or *Reinitialize* options, or any dual variable estimate that obtains a KKT error less than $\|T_{S_k}(x_{k,0}, \lambda_{k-1,N_{k-1}})\|$ for the subsampled problem at $x_{k,0}$. To ensure the algorithm is well defined at $k = 0$, we define $N_{-1} = 0$.
2. **Inner Loop (Lines 4-5):** We update the iterate and dual variable until the termination criterion (2.3) is satisfied. Any deterministic solver that can achieve this in a finite number of steps can be used in the inner loop.

In Framework 2.1, one should choose a deterministic solver suitable for the subsampled problems. While SQP methods work well for general equality constraints, one might prefer projection-based methods when the constraints are simple and projections are inexpensive. If the selected algorithm does not update the dual variable, a least squares problem, similar to *Reinitialize* in (2.2), can be solved within the inner loop to update it. Alternatively, the termination criterion \mathcal{T}_k may use solver-specific optimality measures instead of the KKT error. We discuss such criterion when discussing the SQP method in Subsection 2.3.

2.2 General RA Framework Convergence Analysis

In this section, we present theoretical results for Framework 2.1. We first establish the basic conditions required for convergence with any deterministic solver. Then, we propose an adaptive sampling strategy for selecting the batch size and provide complexity results in terms of the number of inner loop iterations and gradient evaluations.

To begin, we define \mathcal{F}_k as the σ -algebra corresponding to the initial conditions, i.e., $\{x_{0,0}, \lambda_{-1,0}\}$ and sample sets selected up to outer iteration k , i.e., $\{S_0, S_1, \dots, S_{k-1}\}$. This results in the nested structure $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}_k$, with $\mathcal{F}_0 = \{x_{0,0}, \lambda_{-1,0}\}$. Given a sample set S_k , the next outer iterate $x_{k+1,0} = x_{k,N_k}$ is obtained via a deterministic process starting from $x_{k,0}$. Thus, the iterate $x_{k,0}$ is specified under \mathcal{F}_k while the sample set S_k is independent of \mathcal{F}_k , and we only need to define the filtration over the outer iterations. The total expectation is then defined as the expectation conditioned on the initial conditions, i.e., $\mathbb{E}[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_0]$. We now state some basic assumptions for optimization of smooth functions over equality constraints.

Assumption 2.1. Let $\chi \subset \mathbb{R}^n$ be an open convex set containing the sequence of iterates generated by any run of Framework 2.1. The objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable and bounded below over χ , and its gradient function $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz continuous over χ . The constraint function $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuously differentiable and the Jacobian function, $J = \nabla c^T : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$ is Lipschitz continuous over χ . The Jacobian $J(x)$ over $x \in \chi$ has full row rank, i.e., the linear independence constraint qualification (LICQ) holds.

Assumption 2.2. For any run of Framework 2.1, the inner loop terminates finitely satisfying (2.3), i.e., $N_k < \infty$ for all $k \geq 0$.

Assumption 2.3. Let $\{x_{k,N_k}\}$ be the sequence of outer iterates generated by any run of Framework 2.1. There exists $\kappa_g \geq 0$ such that:

1. For the finite-sum problem (1.2): $\|\nabla f(x_{k,N_k})\|^2 \leq \kappa_g^2$ for all $k \geq 0$.
2. For the expectation problem (1.3): $\mathbb{E} [\|\nabla f(x_{k,N_k})\|^2 | \mathcal{F}_k] \leq \tilde{\kappa}_g^2$ for all $k \geq 0$.

Assumption 2.1 is a standard assumption in the equality constrained optimization literature, see e.g., [8, 42], that ensures the existence of a first-order stationary point under the LICQ. Assumption 2.2 guarantees that the deterministic solver used in the inner loop is well-posed and will terminate finitely. Assumption 2.3 imposes bounds on the gradient of the true objective function over the iterates generated by Algorithm 2.1. This assumption is common in constrained optimization for both deterministic [8, 42, 52] and stochastic [2, 19, 45] settings, where the gradient of the true problem is typically assumed to be deterministically bounded over the set of iterates. Thus, Assumption 2.3 does not pose strict restrictions on our analysis.

We now present standard assumptions on the error in the subsampled objective function gradients, similar to [1, 5, 26].

Assumption 2.4. For all $x \in \chi$, the individual component gradients are bounded relative to the gradient of the true objective functions:

1. For the finite-sum problem (1.2): There exist constants $\omega_1, \omega_2 \geq 0$ such that,

$$\|\nabla F(x, \xi)\|^2 \leq \omega_1^2 \|\nabla f(x)\|^2 + \omega_2^2 \quad \text{for all } \xi \in \mathcal{S}. \quad (2.4)$$

2. For the expectation problem (1.3): There exist constants $\tilde{\omega}_1, \tilde{\omega}_2 \geq 0$ such that,

$$\mathbb{E} [\|\nabla f(x) - \nabla F(x, \xi)\|^2 | x] \leq \tilde{\omega}_1^2 \|\nabla f(x)\|^2 + \tilde{\omega}_2^2. \quad (2.5)$$

For the finite-sum problem (1.2), under Assumption 2.4 by [26, Section 3.1], the error in the subsampled gradient over a sample set $S \subseteq \mathcal{S}$ can be bounded as,

$$\|\nabla f(x) - g_S(x)\|^2 \leq 4 \left(1 - \frac{|S|}{|\mathcal{S}|}\right)^2 (\omega_1^2 \|\nabla f(x)\|^2 + \omega_2^2) \quad \text{for all } x \in \chi. \quad (2.6)$$

For the expectation problem (1.3), the error in the subsampled gradient over a set of i.i.d. samples S from the distribution \mathcal{P} , independent of x , can be bounded as,

$$\mathbb{E} [\|\nabla f(x) - g_S(x)\|^2 | x] = \frac{1}{|S|^2} \sum_{\xi \in S} \mathbb{E} [\|\nabla f(x) - \nabla F(x, \xi)\|^2 | x] \leq \frac{\tilde{\omega}_1^2 \|\nabla f(x)\|^2 + \tilde{\omega}_2^2}{|S|}. \quad (2.7)$$

We define an additional gradient error metric for the expectation problem (1.3) over a set of samples S , given by

$$G_S = \max_{x \in \mathcal{X}} \frac{\|\nabla f(x) - g_S(x)\|}{\epsilon_G + \|\nabla f(x)\|}, \quad (2.8)$$

where $\epsilon_G > 0$, which represents the maximum error in the subsampled gradient relative to the true gradient, as defined in [40, 41] for analyzing the RA framework for unconstrained optimization.

To analyze Framework 2.1, we first establish a recursive bound on the KKT error of the true problem across outer iterations.

Lemma 2.2. *Suppose Assumptions 2.1 and 2.2 hold. Then, for all $k \geq 0$, the outer iterates generated by Framework 2.1 satisfy*

$$\begin{aligned} \|T(x_{k,N_k}, \lambda_{k,N_k})\| &\leq \gamma_k \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \epsilon_k \\ &\quad + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma_k \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|. \end{aligned} \quad (2.9)$$

For the expectation problem (1.3), for all $k \geq 0$

$$\begin{aligned} \mathbb{E}[\|T(x_{k,N_k}, \lambda_{k,N_k})\| | \mathcal{F}_k] &\leq \gamma_k \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \mathbb{E}[\epsilon_k | \mathcal{F}_k] \\ &\quad + \mathbb{E}[\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\ &\quad + \gamma_k \mathbb{E}[\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k]. \end{aligned} \quad (2.10)$$

Proof. The KKT error for the true problem at the end of outer iteration $k \geq 0$ can be bounded as,

$$\begin{aligned} \|T(x_{k,N_k}, \lambda_{k,N_k})\| &\leq \|T_{S_k}(x_{k,N_k}, \lambda_{k,N_k})\| + \|T(x_{k,N_k}, \lambda_{k,N_k}) - T_{S_k}(x_{k,N_k}, \lambda_{k,N_k})\| \\ &= \|T_{S_k}(x_{k,N_k}, \lambda_{k,N_k})\| + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| \\ &\leq \gamma_k \|T_{S_k}(x_{k,0}, \lambda_{k,0})\| + \epsilon_k + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| \\ &\leq \gamma_k \|T_{S_k}(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \epsilon_k + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| \\ &\leq \gamma_k \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \epsilon_k + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| \\ &\quad + \gamma_k \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}}) - T_{S_k}(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| \\ &= \gamma_k \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \epsilon_k \\ &\quad + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma_k \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|, \end{aligned}$$

where the second inequality follows from (2.3), the third inequality follows from $x_{k-1,N_{k-1}} = x_{k,0}$ and the update rule (2.2) for $\lambda_{k,0}$, thus completing the proof for (2.9). Taking the conditional expectation of (2.9) given \mathcal{F}_k yields (2.10), since $(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})$ are known under \mathcal{F}_k . \square

We now establish the convergence of Framework 2.1.

Theorem 2.3. *Suppose Assumptions 2.1, 2.2, 2.3 and 2.4 hold for any run of Framework 2.1.*

1. *For the finite-sum problem (1.2): If the sample set sequence is chosen such that $\{S_k\} \rightarrow \mathcal{S}$ and the termination criterion parameters as $0 \leq \{\gamma_k\} \leq \gamma < 1$ and $\epsilon_k \rightarrow 0$, then, $\|T(x_{k,N_k}, \lambda_{k,N_k})\| \rightarrow 0$ and $\{x_{k,N_k}\}$ converges to a first-order stationary point of (1.1).*
2. *For the expectation problem (1.3): If the batch size sequence is chosen such that $\{|S_k|\} \rightarrow \infty$, the termination criterion parameters as $0 \leq \{\gamma_k\} \leq \tilde{\gamma} < 1$ and $\mathbb{E}[\epsilon_k | \mathcal{F}_k] \rightarrow 0$, and given that $\mathbb{E}[G_{S_k}^2 | \mathcal{F}_k] \rightarrow 0$ as $\{|S_k|\} \rightarrow \infty$, then, $\mathbb{E}[\|T(x_{k,N_k}, \lambda_{k,N_k})\|] \rightarrow 0$ and $\{x_{k,N_k}\}$ converges to a first-order stationary point of (1.1) in expectation.*

Proof. For the finite-sum problem (1.2), unrolling (2.9) from Lemma 2.2 yields for all $k \geq 0$,

$$\begin{aligned}
\|T(x_{k,N_k}, \lambda_{k,N_k})\| &\leq \left\{ \prod_{h=0}^k \gamma_h \right\} \|T(x_{0,0}, \lambda_{-1,0})\| + \epsilon_k \\
&\quad + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma_k \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} [\epsilon_i + \|\nabla f(x_{i,N_i}) - g_{S_i}(x_{i,N_i})\|] \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \gamma_i \|\nabla f(x_{i,0}) - g_{S_i}(x_{i,0})\| \\
&\leq \left\{ \prod_{h=0}^k \gamma_h \right\} \|T(x_{0,0}, \lambda_{-1,0})\| + \epsilon_k \\
&\quad + 2 \left(\frac{|S| - |S_k|}{|S|} \right) (\omega_1 (\|\nabla f(x_{k,N_k})\| + \gamma_k \|\nabla f(x_{k,0})\|) + \omega_2 (1 + \gamma_k)) \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \left[\epsilon_i + \left(\frac{|S| - |S_i|}{|S|} \right) (\omega_1 \|\nabla f(x_{i,N_i})\| + \omega_2) \right] \\
&\quad + 2 \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \gamma_i \left(\frac{|S| - |S_i|}{|S|} \right) [\omega_1 \|\nabla f(x_{i,0})\| + \omega_2] \\
&\leq \left\{ \prod_{h=0}^k \gamma_h \right\} \|T(x_{0,0}, \lambda_{-1,0})\| + \epsilon_k + 2 \left(\frac{|S| - |S_k|}{|S|} \right) (1 + \gamma_k) (\omega_1 \kappa_g + \omega_2) \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \left[\epsilon_i + 2 \left(\frac{|S| - |S_i|}{|S|} \right) (1 + \gamma_i) (\omega_1 \kappa_g + \omega_2) \right],
\end{aligned} \tag{2.11}$$

where the second inequality follows from (2.6) and third inequality follows by Assumption 2.3. The first term in the above inequality converges to zero as $\left\{ \prod_{h=0}^k \gamma_h \right\} \leq \gamma^{k+1}$ where $\gamma < 1$. The second and third terms converge to zero under the specifications $\epsilon_k \rightarrow 0$ and $\{|S_k|\} \rightarrow |S|$, respectively. The final term is a sum of the form $\sum_{i=0}^{k-1} a_i \left\{ \prod_{h=i+1}^k \gamma_h \right\} \leq \sum_{i=0}^{k-1} a_i \gamma^{k-i}$ which converges to zero by Lemma A.1, with $b_i = a_i = \epsilon_i + 2 \left(\frac{|S| - |S_i|}{|S|} \right) (1 + \gamma_i) (\omega_1 \kappa_g + \omega_2) \rightarrow 0$. Thus, all terms in the KKT error bound converge to zero completing the proof.

For the expectation problem (1.3), taking the total expectation of (2.11) and applying the tower

property, i.e., $\mathbb{E}[\epsilon_k] = \mathbb{E}[\mathbb{E}[\epsilon_k|\mathcal{F}_k]]$, for all relevant quantities yields for all $k \geq 0$,

$$\begin{aligned}
\mathbb{E}[\|T(x_{k,N_k}, \lambda_{k,N_k})\|] &\leq \left\{ \prod_{h=0}^k \gamma_h \right\} \|T(x_{0,0}, \lambda_{-1,0})\| + \mathbb{E}[\mathbb{E}[\epsilon_k|\mathcal{F}_k]] \\
&\quad + \mathbb{E}[\mathbb{E}[\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\||\mathcal{F}_k]] \\
&\quad + \gamma_k \mathbb{E}[\mathbb{E}[\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\||\mathcal{F}_k]] \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \mathbb{E}[\mathbb{E}[\epsilon_i|\mathcal{F}_i] + \mathbb{E}[\|\nabla f(x_{i,N_i}) - g_{S_i}(x_{i,N_i})\||\mathcal{F}_i]] \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \gamma_i \mathbb{E}[\mathbb{E}[\|\nabla f(x_{i,0}) - g_{S_i}(x_{i,0})\||\mathcal{F}_i]].
\end{aligned} \tag{2.12}$$

The result contains two types of errors for the subsampled gradients. The first is $\mathbb{E}[\|\nabla f(x_{i,0}) - g_{S_i}(x_{i,0})\||\mathcal{F}_i]$, where S_i is independent of \mathcal{F}_i (and thus $x_{i,0}$), and (2.7) holds. The second is $\mathbb{E}[\|\nabla f(x_{i,N_i}) - g_{S_i}(x_{i,N_i})\||\mathcal{F}_i]$ where x_{i,N_i} is dependent on S_i . Here, we utilize the metric defined in (2.8) as,

$$\begin{aligned}
\mathbb{E}[\|\nabla f(x_{i,N_i}) - g_{S_i}(x_{i,N_i})\||\mathcal{F}_i] &\leq \mathbb{E}[G_{S_i}(\epsilon_G + \|\nabla f(x_{i,N_i})\|)|\mathcal{F}_i] \\
&= \epsilon_G \mathbb{E}[G_{S_i}|\mathcal{F}_i] + \mathbb{E}[G_{S_i} \|\nabla f(x_{i,N_i})\||\mathcal{F}_i] \\
&\leq \epsilon_G \mathbb{E}[G_{S_i}|\mathcal{F}_i] + \sqrt{\mathbb{E}[G_{S_i}^2|\mathcal{F}_i] \mathbb{E}[\|\nabla f(x_{i,N_i})\|^2|\mathcal{F}_i]},
\end{aligned} \tag{2.13}$$

where the last inequality follows from the expectation version of the Cauchy–Schwarz inequality. Substituting (2.13) and (2.7) into (2.12), along with the bound on the gradient from Assumption 2.3 yields,

$$\begin{aligned}
\mathbb{E}[\|T(x_{k,N_k}, \lambda_{k,N_k})\|] &\leq \|T(x_{0,0}, \lambda_{-1,0})\| \left\{ \prod_{h=0}^k \gamma_h \right\} \\
&\quad + \mathbb{E} \left[\mathbb{E}[\epsilon_k|\mathcal{F}_k] + \epsilon_G \mathbb{E}[G_{S_k}|\mathcal{F}_k] + \tilde{\kappa}_g \sqrt{\mathbb{E}[G_{S_k}^2|\mathcal{F}_k]} + \gamma_k \frac{\tilde{\omega}_1 \tilde{\kappa}_g + \tilde{\omega}_2}{\sqrt{|S_k|}} \right] \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \mathbb{E} \left[\mathbb{E}[\epsilon_i|\mathcal{F}_i] + \epsilon_G \mathbb{E}[G_{S_i}|\mathcal{F}_i] + \tilde{\kappa}_g \sqrt{\mathbb{E}[G_{S_i}^2|\mathcal{F}_i]} \right] \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \gamma_i \mathbb{E} \left[\frac{\tilde{\omega}_1 \tilde{\kappa}_g + \tilde{\omega}_2}{\sqrt{|S_i|}} \right].
\end{aligned}$$

In the above bound, the first term is the same as the first term for the finite-sum problem (above). Moreover, all terms within the second term expectation converge to zero under the specifications $\mathbb{E}[\epsilon_k|\mathcal{F}_k] \rightarrow 0$ and $\{|S_k|\} \rightarrow \infty$, and the assumption $\mathbb{E}[G_{S_k}^2|\mathcal{F}_k] \rightarrow 0$. The last two terms together are a series of the form $\sum_{i=0}^{k-1} \mathbb{E}[a_i] \left\{ \prod_{h=i+1}^k \gamma_h \right\} \leq \sum_{i=0}^{k-1} \mathbb{E}[a_i] \tilde{\gamma}^{k-i}$ which converges to zero by Lemma A.1, with $b_i = \mathbb{E}[a_i] \rightarrow 0$ as $a_i \rightarrow 0$, thus completing the proof. \square

Theorem 2.3 establishes the convergence of Framework 2.1 for stochastic problems with deterministic equality constraints, using any deterministic solver in the inner loop. While most of the

requirements are standard assumptions, for the expectation problem (1.3), an additional assumption that $\mathbb{E}[G_{S_k}^2 | \mathcal{F}_k] \rightarrow 0$ is needed. This ensures that the maximum deviation of the subsampled gradient from the true problem gradient decreases with the sample size when normalized by the true problem gradient, thereby enforcing the uniform convergence of the error in subsampled gradients. This condition is not strict and is satisfied by many error classes due to the relative nature of the error, as shown in [40, 41] for unconstrained stochastic optimization.

We now introduce a strategy to adaptively select the batch size $|S_k|$ in Framework 2.1. This strategy is a modification of the well-known norm test [5, 7, 10, 43].

Condition 2.1. *At the beginning of outer iteration $k \geq 0$ in Framework 2.1, $|S_k|$ is selected such that:*

1. *For the finite-sum problem (1.2): With constants $\theta, \hat{\theta}, a \geq 0$ and $\beta \in (0, 1)$,*

$$\begin{aligned} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 &\leq \theta^2 \|T(x_{k,0}, \lambda_{k-1, N_{k-1}})\|^2 + a^2 \beta^{2k}, \\ \|\nabla f(x_{k, N_k}) - g_{S_k}(x_{k, N_k})\|^2 &\leq \hat{\theta}^2 (\theta^2 \|T(x_{k,0}, \lambda_{k-1, N_{k-1}})\|^2 + a^2 \beta^{2k}). \end{aligned}$$

2. *For the expectation problem (1.3): With constants $\tilde{\theta}, \tilde{a} \geq 0$ and $\tilde{\beta} \in (0, 1)$,*

$$\mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 | \mathcal{F}_k] \leq \tilde{\theta}^2 \|T(x_{k,0}, \lambda_{k-1, N_{k-1}})\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}.$$

Remark 2.4. *We make the following remarks regarding Condition 2.1.*

1. **Batch size:** Under Assumption 2.4, Condition 2.1 can always be satisfied with a sufficiently large batch size. If θ (or $\hat{\theta}$) = 0, the condition is equivalent to increasing the batch size at a geometric rate. When θ (or $\hat{\theta}$) > 0, the condition takes the KKT error in the true problem at the current iterate and dual variable into account when selecting the batch size. We later provide sequences of sufficiently large batch sizes that satisfy Condition 2.1.
2. **True problem estimates:** The right-hand side of the inequalities in Condition 2.1 require access to true problem quantities which cannot be computed. However, we utilize these conditions to understand the permissible errors in the algorithm that guarantee convergence and develop a practical strategy using sampled estimates, detailed in Section 4. The condition can also be utilized with any other measure of optimality for the problem instead of the KKT error, as will be discussed in the context of solver-specific termination criteria in Subsection 2.4.
3. **Finite-sum condition:** For the finite-sum problem (1.2), Condition 2.1 bounds the gradient error at the start and end of the inner loop. While this is a stronger condition than usually considered for the norm test, due to the finite-sum nature of the problem, if one controls the gradient error at a particular iterate, one can ensure the gradient error is bounded over the space of iterates. Thus, with a large enough constant $\hat{\theta}$, the condition can be satisfied by selecting a batch size with small enough gradient error at the start of the inner loop.

We proceed to analyze Framework 2.1 with batch sizes selected to satisfy Condition 2.1. First, we introduce additional assumptions to characterize the variance of the expectation problem (1.3).

Assumption 2.5. *We make the following assumptions regarding the expectation problem (1.3).*

1. **CLT Scaling:** The gradient error metric (2.8) follows CLT-scaling, i.e., there exists $\kappa_G \geq 0$ such that for any set of i.i.d. samples S from \mathcal{P} , $\mathbb{E}[G_S^2] \leq \frac{\kappa_G^2}{|S|}$.

2. *Variance Lower Bound: The variance of the subsampled gradients is lower bounded, i.e., there exists $\kappa_\sigma > 0$ such that $\text{Var}(\nabla F(x_{k,0}, \xi) | \mathcal{F}_k) \geq \kappa_G^2 \kappa_\sigma^2$ for all $k \geq 0$.*

In Assumption 2.5, the CLT scaling assumption states that the maximum error in the subsampled gradient, when normalized by the true gradient norm, scales inversely with the sample size. This assumption, also used in [40, 41] for unconstrained stochastic optimization, parallels conditions in stochastic gradient methods but applies over the entire space of iterates rather than at each individual iterate, as required in the RA framework.

The variance lower bound in Assumption 2.5 is counterintuitive to the conditions prescribed in most works for stochastic optimization where the variance is assumed to be upper bounded [1, 21, 43]. Such an assumption becomes essential in our work to ensure that the adaptive sampling strategy of satisfying Condition 2.1 results in an increasing batch size sequence, as opposed to using a deterministically prescribed sequence in [21, 40, 41, 57]. While the assumption does not pose any reasonable barriers in practice, it is also not required when one uses a prescribed sequence for the batch sizes, as will also be demonstrated. When $\kappa_G = 0$, the lower bound does not exist but increasing the batch size is unnecessary since the deviation of the subsampled gradient from the true problem gradient is zero from CLT scaling. When $\kappa_G > 0$, the assumption can be relaxed to requiring non-zero variance infinitely often. This ensures that the batch size increases across multiple outer iterations when Condition 2.1 is satisfied for the expectation problem. If the variance is zero at outer iterates infinitely often, increasing the batch size is unnecessary, as the error in the subsampled gradient is zero infinitely often. Thus, the variance lower bound assumption excludes such degenerate cases from the theoretical analysis. In practice, since $|S_k| \geq |S_{k-1}|$ is enforced, the batch size remains unchanged when zero variance is encountered. As better solutions are found, the batch size eventually increases, preventing any violations of this assumption.

We now establish the complexity guarantees for Framework 2.1 to obtain an $\epsilon > 0$ stationary point, defined as $\|T(x_{k,N_k}, \lambda_{k,N_k})\| \leq \epsilon$ for the finite-sum problem (1.2) and in expectation as $\mathbb{E}[\|T(x_{k,N_k}, \lambda_{k,N_k})\|] \leq \epsilon$ for the expectation problem (1.3), with batch sizes selected to satisfy Condition 2.1. We first look at this complexity across outer iterations.

Theorem 2.5. *Suppose Assumptions 2.1, 2.2, 2.3 and 2.4 hold and the batch size sequence $\{|S_k|\}$ is chosen to satisfy Condition 2.1 in Framework 2.1.*

1. *For the finite-sum problem (1.2): For all $k \geq 0$, if the termination criterion parameters are chosen such that $0 \leq \gamma_k \leq \gamma < 1$ and $\epsilon_k = \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k$ with $\omega, \hat{\omega} \geq 0$, and Condition 2.1 parameters are chosen such that $a_1 = \left[\gamma + \theta(\omega + \hat{\theta} + \gamma) \right] < 1$, then, the KKT error converges at a linear rate across outer iterations, i.e.,*

$$\|T(x_{k,N_k}, \lambda_{k,N_k})\| \leq \max\{a_1 + \nu, \beta\}^{k+1} \max\{\|T(x_{0,0}, \lambda_{-1,0})\|, \frac{a_2}{\nu}\},$$

where $a_2 = a(\omega + \hat{\theta} + \gamma) + \hat{\omega}$ and $\nu > 0$ such that $a_1 + \nu < 1$.

2. *For the expectation problem (1.3): For all $k \geq 0$, if Assumption 2.5 is satisfied, the termination criterion parameters are chosen such that $0 \leq \gamma_k \leq \tilde{\gamma} < 1$ and $\epsilon_k = \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}, \xi) | \mathcal{F}_k)}{|S_k|}}$ where $\tilde{\omega} \geq 0$, and Condition 2.1 parameters are chosen such that $\tilde{a}_1 = \left[\tilde{\gamma} + \tilde{\theta} \left(\tilde{\omega} + \frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} + \tilde{\gamma} \right) \right] < 1$, then, the expected KKT error converges at a linear rate across outer iterations, i.e.,*

$$\mathbb{E}[\|T(x_{k,N_k}, \lambda_{k,N_k})\|] \leq \max\{\tilde{a}_1 + \tilde{\nu}, \tilde{\beta}\}^{k+1} \max\{\|T(x_{0,0}, \lambda_{-1,0})\|, \frac{\tilde{a}_2}{\tilde{\nu}}\},$$

where $\tilde{a}_2 = \tilde{a} \left(\tilde{\omega} + \frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} + \tilde{\gamma} \right)$ and $\tilde{\nu} > 0$ such that $\tilde{a}_1 + \tilde{\nu} < 1$.

Proof. For the finite-sum problem (1.2), under the stated conditions, the error bound (2.9) from Lemma 2.2 can be further simplified as,

$$\begin{aligned} \|T(x_{k,N_k}, \lambda_{k,N_k})\| &\leq \gamma \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k \\ &\quad + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| \\ &\leq \gamma \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \omega (\theta \|T(x_{k,0}, \lambda_{k-1,N_{k-1}})\| + a \beta^k) + \hat{\omega} \beta^k \\ &\quad + \hat{\theta} (\theta \|T(x_{k,0}, \lambda_{k-1,N_{k-1}})\| + a \beta^k) + \gamma (\theta \|T(x_{k,0}, \lambda_{k-1,N_{k-1}})\| + a \beta^k) \\ &= a_1 \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + a_2 \beta^k, \end{aligned}$$

where the second inequality follows from Condition 2.1 and the rest follows from the defined constants with $x_{k,0} = x_{k-1,N_{k-1}}$. Using the above bound, applying Lemma A.2 with $Z_k = \|T(x_{k,N_k}, \lambda_{k,N_k})\|$, $\rho_1 = a_1$, $\rho_2 = \beta$ and $b = a_2$ completes the proof.

For the expectation problem (1.3), under the stated conditions, the error bound (2.10) from Lemma 2.2 can be further simplified as,

$$\begin{aligned} \mathbb{E} [\|T(x_{k,N_k}, \lambda_{k,N_k})\| | \mathcal{F}_k] &\leq \tilde{\gamma} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}, \xi) | \mathcal{F}_k)}{|S_k|}} \\ &\quad + \mathbb{E} [\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\ &\quad + \tilde{\gamma} \mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\ &\leq \tilde{\gamma} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}, \xi) | \mathcal{F}_k)}{|S_k|}} \\ &\quad + \kappa_G \frac{(\epsilon_G + \kappa_g)}{\sqrt{|S_k|}} + \tilde{\gamma} \mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\ &\leq \tilde{\gamma} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| \\ &\quad + \tilde{\omega} \left[\tilde{\theta} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \tilde{a} \tilde{\beta}^k \right] \\ &\quad + \frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} \left[\tilde{\theta} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \tilde{a} \tilde{\beta}^k \right] \\ &\quad + \tilde{\gamma} \left[\tilde{\theta} \|T(x_{k,0}, \lambda_{k-1,N_{k-1}})\| + \tilde{a} \tilde{\beta}^k \right] \\ &= \tilde{a}_1 \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \tilde{a}_2 \tilde{\beta}^k, \end{aligned}$$

where the second inequality follows from (2.13) and Assumption 2.5, the third inequality follows by substituting $\kappa_G \leq \frac{\sqrt{\text{Var}(\nabla F(x_{k,0}, \xi) | \mathcal{F}_k)}}{\kappa_\sigma}$ and Condition 2.1 and the last equality substitutes the defined constants with $x_{k,0} = x_{k-1,N_{k-1}}$. Taking the total expectation of the above bound yields,

$$\mathbb{E} [\|T(x_{k,N_k}, \lambda_{k,N_k})\|] \leq \tilde{a}_1 \mathbb{E} [\|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\|] + \tilde{a}_2 \tilde{\beta}^k,$$

where applying Lemma A.2 with $Z_k = \mathbb{E} [\|T(x_{k,N_k}, \lambda_{k,N_k})\| | \mathcal{F}_k]$, $\rho_1 = \tilde{a}_1$, $\rho_2 = \tilde{\beta}$ and $b = \tilde{a}_2$ completes the proof. \square

Theorem 2.5 establishes linear convergence across outer iterations for Framework 2.1 when the batch sizes are chosen to satisfy Condition 2.1. While the additional tolerance term ϵ_k in termination criterion (2.3) is controlled by the variance for the expectation problem (1.3), an extra

relaxation term $\hat{\omega}\beta^k$ is introduced in ϵ_k for the finite-sum problem (1.2). In the finite-sum setting, Condition 2.1 evaluates the gradient error only over the selected sample set S_k , which may yield a low error at the current iterate but fail to represent the full problem, potentially leading to over-solving the subsampled problem. To avoid this, a geometrically decreasing error term is added to ϵ_k . In contrast, for the expectation problem, Condition 2.1 considers the expected error over the entire probability space \mathcal{P} , yielding a more representative measure of gradient error and eliminating the need for such additional error terms.

The conditions proposed in Theorem 2.5 show an interesting trade-off between the parameters of Condition 2.1 and termination criterion (2.3). If the subsampled problems are solved to high accuracy by setting small values of γ , ω and $\hat{\omega}$ (or $\tilde{\gamma}$ and $\tilde{\omega}$), a larger θ (or $\tilde{\theta}$) is permitted. This suggests that achieving high accuracy in solving subproblems reduces the necessity for an aggressive increase in batch size, as the high accuracy inherently introduces significant bias in the solutions. Conversely, if γ , ω and $\hat{\omega}$ (or $\tilde{\gamma}$ and $\tilde{\omega}$) are set to larger values, θ (or $\tilde{\theta}$) needs to be decreased for an aggressive increase in batch sizes to achieve progress every outer iteration. Setting θ (or $\tilde{\theta}$) to zero effectively enforces a geometric increase in batch size under Condition 2.1. This eliminates the need for the parameter conditions in Theorem 2.5 and the variance lower bound in Assumption 2.5. The results for this simplified case are presented in the following corollary.

Corollary 2.6. *Suppose the conditions of Theorem 2.5 hold.*

1. *For the finite-sum problem (1.2): For all $k \geq 0$, if the batch size is chosen as $|S_k| = \lceil (1 - \beta^k)|S| \rceil$ with $\beta \in (0, 1)$ and the termination criterion parameters are chosen as $0 \leq \gamma_k \leq \gamma < 1$ and $\epsilon_k = \omega \left(1 - \frac{|S_k|}{|S|}\right)$ with $\omega \geq 0$, then, the KKT error converges at a linear rate across outer iterations as expressed in Theorem 2.5 with $a_1 = \gamma$ and $a_2 = \omega + 2(1 + \gamma)(\omega_1\kappa_g + \omega_2)$.*
2. *For the expectation problem (1.3): For all $k \geq 0$, if Assumption 2.5 is satisfied, the batch size is chosen as $|S_{k+1}| = \left\lceil \frac{|S_k|}{\tilde{\beta}^2} \right\rceil$ with $\tilde{\beta} \in (0, 1)$ and the termination criterion parameters are chosen as $0 \leq \{\gamma_k\} \leq \tilde{\gamma} < 1$ and $\epsilon_k = \frac{\tilde{\omega}}{\sqrt{|S_k|}}$ where $\tilde{\omega} \geq 0$, then, the expected KKT error converges at a linear rate across outer iterations as expressed in Theorem 2.5 with $\tilde{a}_1 = \tilde{\gamma}$ and $\tilde{a}_2 = \frac{\tilde{\omega} + \kappa_G(\epsilon_G + \tilde{\kappa}_g) + \tilde{\gamma}(\tilde{\omega}_1\tilde{\kappa}_g + \tilde{\omega}_2)}{\sqrt{|S_0|}}$, where $|S_0|$ is the initial user-specified batch size.*

Proof. The proof follows the same procedure as Theorem 2.5, but with more pessimistic bounds on the gradient errors, using the explicitly specified form for the batch size. For the finite-sum problem (1.2), under the chosen parameters, the error bound (2.9) from Lemma 2.2 can be simplified as,

$$\begin{aligned}
\|T(x_{k,N_k}, \lambda_{k,N_k})\| &\leq \gamma \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \omega \left(1 - \frac{|S_k|}{|S|}\right) \\
&\quad + \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| \\
&\leq \gamma \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \omega \left(\frac{|S| - (1 - \beta^k)|S|}{|S|}\right) \\
&\quad + 2 \left(\frac{|S| - (1 - \beta^k)|S|}{|S|}\right) [\omega_1(\|\nabla f(x_{k,N_k})\| + \gamma \|\nabla f(x_{k,0})\|) + \omega_2(1 + \gamma)] \\
&\leq \gamma \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + (\omega + 2(1 + \gamma)(\omega_1\kappa_g + \omega_2))\beta^k \\
&= \gamma \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + a_2\beta^k,
\end{aligned}$$

where the second inequality follows from (2.6) and the chosen batch size sequence, the third inequality follows from Assumption 2.3 and the final equality substitutes the defined constants. Using

the above bound, applying Lemma A.2 with $Z_k = \|T(x_{k,N_k}, \lambda_{k,N_k})\|$, $\rho_1 = \gamma$, $\rho_2 = \beta$ and $b = a_2$ completes the proof.

For the expectation problem (1.3), under the stated conditions, the error bound (2.10) from Lemma 2.2 can be further simplified as,

$$\begin{aligned}
\mathbb{E} [\|T(x_{k,N_k}, \lambda_{k,N_k})\| | \mathcal{F}_k] &\leq \tilde{\gamma} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \frac{\tilde{\omega}}{\sqrt{|S_k|}} \\
&\quad + \mathbb{E} [\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\
&\quad + \tilde{\gamma} \mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\
&\leq \tilde{\gamma} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| + \frac{\tilde{\omega}}{\sqrt{|S_k|}} \\
&\quad + \kappa_G \frac{(\epsilon_G + \tilde{\kappa}_g)}{\sqrt{|S_k|}} + \tilde{\gamma} \frac{\tilde{\omega}_1 \|\nabla f(x_{k,0})\| + \tilde{\omega}_2}{\sqrt{|S_k|}} \\
&\leq \tilde{\gamma} \|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\| \\
&\quad + \frac{\tilde{\beta}^k}{\sqrt{|S_0|}} [\tilde{\omega} + \kappa_G(\epsilon_G + \tilde{\kappa}_g) + \tilde{\gamma}(\tilde{\omega}_1 \|\nabla f(x_{k,0})\| + \tilde{\omega}_2)]
\end{aligned}$$

where the second inequality follows from (2.7) and (2.13) for the gradient error with bias and Assumption 2.4 for the unbiased gradient error, and the third inequality follows from the chosen batch size sequence. Taking total expectation of the above bound, with $\mathbb{E}[\|\nabla f(x_{k,0})\|] = \mathbb{E}[\mathbb{E}[\|\nabla f(x_{k-1,N_{k-1}})\| | \mathcal{F}_{k-1}]] \leq \tilde{\kappa}_g$ by Assumption 2.3 and substituting the defined constants,

$$\mathbb{E} [\|T(x_{k,N_k}, \lambda_{k,N_k})\|] \leq \tilde{\gamma} \mathbb{E} [\|T(x_{k-1,N_{k-1}}, \lambda_{k-1,N_{k-1}})\|] + \tilde{a}_2 \tilde{\beta}^k,$$

where applying Lemma A.2 with $Z_k = \mathbb{E} [\|T(x_{k,N_k}, \lambda_{k,N_k})\| | \mathcal{F}_k]$, $\rho_1 = \tilde{\gamma}$, $\rho_2 = \tilde{\beta}$ and $b = \tilde{a}_2$ completes the proof. \square

Corollary 2.6 establishes linear convergence for Framework 2.1 under a pre-specified geometrically increasing batch size sequence, a simplified alternative to Condition 2.1. Since the batch sizes are deterministically chosen, the tolerance sequence ϵ_k is controlled via the batch size, and the conditions and results in Theorem 2.5 are simplified, particularly for the finite-sum problem (1.2). Alternatively, one can select sufficiently large batch sizes when θ (or $\tilde{\theta}$) is non-zero to ensure that Condition 2.1 holds. These batch sizes are conservative and derived under a stronger bound on the true problem objective function gradient than Assumption 2.3 for the expectation problem (1.3). Since a deterministic bound on the true problem gradient is a standard assumption in both deterministic [8, 42, 52] and stochastic [2, 19, 45] constrained optimization as previously discussed, the assumption is reasonable.

Assumption 2.6. Let $\{x_{k,N_k}\}$ be the sequence of outer iterates generated by any run of Framework 2.1. There exists $\kappa_g \geq 0$ such that $\|\nabla f(x_{k,N_k})\| \leq \kappa_g$ for all $k \geq 0$.

The next lemma provides sufficiently large batch sizes that ensure Condition 2.1 holds.

Lemma 2.7. Suppose Assumptions 2.1, 2.2, 2.4 and 2.6 hold in Framework 2.1.

1. For the finite-sum problem (1.2): For $k \geq 0$, Condition 2.1 is satisfied if

$$|S_k| \geq |\mathcal{S}| \left(1 - \sqrt{\frac{\theta^2 \|T(x_{k,0}, \lambda_{k-1,N_{k-1}})\|^2 + a^2 \beta^{2k}}{4(\omega_1^2 \kappa_g^2 + \omega_2^2)}} \right) \quad \text{with } \hat{\theta} = 1.$$

2. For the expectation problem (1.3): For $k \geq 0$, Condition 2.1 is satisfied if

$$|S_k| \geq \frac{\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2}{\tilde{\theta}^2 \|T(x_{k,0}, \lambda_{k-1, N_{k-1}})\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}}.$$

Proof. For the finite-sum problem (1.2),

$$\begin{aligned} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 &\leq 4 \left(1 - \frac{|S_k|}{|S|}\right)^2 (\omega_1^2 \|\nabla f(x_{k,0})\|^2 + \omega_2^2) \\ &\leq 4 \left(\sqrt{\frac{\theta^2 \|T(x_{k,0}, \lambda_{k-1, N_{k-1}})\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}}{4(\omega_1^2 \kappa_g^2 + \omega_2^2)}} \right)^2 (\omega_1^2 \kappa_g^2 + \omega_2^2), \end{aligned}$$

where the first inequality follows from (2.6) and the second inequality follows from Assumption 2.6 and the lower bound on $|S_k|$, thereby satisfying Condition 2.1. The gradient error condition on x_{k, N_k} with $\hat{\theta} = 1$ can be verified in the same way. For the expectation problem (1.3),

$$\begin{aligned} \mathbb{E}[\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 | \mathcal{F}_k] &\leq \frac{\tilde{\omega}_1^2 \|\nabla f(x_{k,0})\|^2 + \tilde{\omega}_2^2}{|S_k|} \\ &\leq (\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2) \frac{\tilde{\theta}^2 \|T(x_{k,0}, \lambda_{k-1, N_{k-1}})\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}}{\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2}, \end{aligned}$$

where the first inequality follows from (2.7) and the second inequality follows from Assumption 2.6 and the suggested lower bound on $|S_k|$, thereby satisfying Condition 2.1. \square

We have analyzed the iterates generated by Framework 2.1 under Condition 2.1 for sampling across outer iterations. Theorem 2.5 established a linear convergence rate across outer iterations, implying that an ϵ -accurate solution can be achieved in K_ϵ outer iterations, where $K_\epsilon = \mathcal{O}(\log \frac{1}{\epsilon})$. While this result provides an upper bound on the number of outer iterations, the dominant computational cost in Framework 2.1 arises from the inner loop—specifically, from solving deterministic subproblems (e.g., the SQP subproblem (1.5)) and computing the subsampled gradients.

We now analyze the total work complexity of Framework 2.1 to achieve an ϵ -accurate solution in terms of two metrics: (1) inner iteration complexity, which is the total number of inner iterations performed $\left(\sum_{k=0}^{K_\epsilon} N_k\right)$, and (2) gradient complexity, which is the total number of subsampled gradients computed $\left(\sum_{k=0}^{K_\epsilon} |S_k| N_k\right)$, assuming one gradient computation per inner iteration. To derive these results, we need to characterize the performance of the deterministic solver used in the inner loop beyond Assumption 2.2. We focus on a specific sublinear rate of convergence for the deterministic solver, most commonly achieved for deterministic problems with a nonconvex objective function and nonconvex equality constraints [1, 19].

Assumption 2.7. *The inner loop in Framework 2.1 exhibits sublinear convergence, i.e., for all $k \geq 0$ an ϵ -accurate solution to the subsampled problem, $\|T_{S_k}(x_{k, N_k}, \lambda_{k, N_k})\| \leq \epsilon$, is achieved in $N_k = \mathcal{O}(\epsilon^{-2})$ iterations.*

Given Assumption 2.7, we analyze the total work complexity of Framework 2.1.

Theorem 2.8. *Suppose Assumptions 2.1, 2.4, 2.6 and 2.7 hold and the batch size sequence $\{|S_k|\}$ is chosen to satisfy Condition 2.1 in Framework 2.1.*

1. For the finite-sum problem (1.2): For all $k \geq 0$, if termination criterion parameters are chosen as $0 \leq \gamma_k \leq \gamma < 1$ and $\epsilon_k = \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k$ with $\omega, \hat{\omega} > 0$, and Condition 2.1 parameters are chosen such that $\left[\gamma + \theta(\omega + \hat{\theta} + \gamma) \right] < \beta$, one achieves an $\epsilon > 0$ accurate solution in $\mathcal{O}(\epsilon^{-2})$ inner iterations and $\mathcal{O}(|S|\epsilon^{-2})$ gradient evaluations.
2. For the expectation problem (1.3): For all $k \geq 0$, if Assumption 2.5 holds, the termination criterion parameters are chosen as $0 \leq \gamma_k \leq \tilde{\gamma} < 1$, $\epsilon_k = \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}, \xi)|\mathcal{F}_k)}{|S_k|}}$ where $\tilde{\omega} > 0$, and Condition 2.1 parameters are chosen such that $\left[\tilde{\gamma} + \tilde{\theta} \left(\tilde{\omega} + \frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} \right) + \tilde{\gamma} \right] < \tilde{\beta}$ and $\tilde{a} = \sqrt{\frac{\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2}{|S_0|}}$, where $|S_0|$ is the initial batch size, one achieves an $\epsilon > 0$ accurate solution in expectation in $\mathcal{O}(\epsilon^{-2})$ inner iterations and $\mathcal{O}(\epsilon^{-4})$ gradient evaluations.

Proof. For the finite-sum problem (1.2), the number of outer iteration to reach an ϵ accurate solution is $K_\epsilon = \mathcal{O}\left(\frac{1}{\log(1/\beta)} \log\left(\frac{1}{\epsilon}\right)\right)$ under the stated conditions from Theorem 2.5. Thus, under Assumption 2.7, the total number of inner iterations can be bounded with a constant $C > 0$ as,

$$\sum_{k=0}^{K_\epsilon} N_k \leq \sum_{k=0}^{K_\epsilon} \frac{C}{(\gamma_k \|T_{S_k}(x_{k,0}, \lambda_{k,0})\| + \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k)^2} \leq \sum_{k=0}^{K_\epsilon} \frac{C}{\hat{\omega}^2 \beta^{2k}} = \frac{C}{\hat{\omega}^2} \frac{(1/\beta)^{2(K_\epsilon+1)} - 1}{(1/\beta)^2 - 1},$$

which is of the order $\left(\frac{1}{\beta}\right)^{2K_\epsilon} = \mathcal{O}(\epsilon^{-2})$. As the batch size is bounded by $|S_k| \leq |S|$, the total number of gradient evaluations can be bounded as $\sum_{k=0}^{K_\epsilon} |S_k| N_k \leq |S| \sum_{k=0}^{K_\epsilon} N_k = \mathcal{O}(|S|\epsilon^{-2})$.

For the expectation problem (1.3), the number of outer iterations to achieve an ϵ accurate solution in expectation is $K_\epsilon = \mathcal{O}\left(\frac{1}{\log(1/\tilde{\beta})} \log\left(\frac{1}{\epsilon}\right)\right)$ under the stated conditions from Theorem 2.5. The expected number of inner iterations in outer iteration $k \geq 0$ can be bounded with a constant $C > 0$ as,

$$\mathbb{E}[N_k | \mathcal{F}_k] \leq \mathbb{E} \left[\frac{C}{\left(\gamma_k \|T_{S_k}(x_{k,0}, \lambda_{k,0})\| + \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}, \xi)|\mathcal{F}_k)}{|S_k|}} \right)^2} \middle| \mathcal{F}_k \right] \leq \frac{C|S_k|}{\tilde{\omega}^2 \text{Var}(\nabla F(x_{k,0}, \xi)|\mathcal{F}_k)},$$

where the final bound is a deterministic quantity given \mathcal{F}_k . Let $|S_k|$ be selected such that Condition 2.1 holds with equality, i.e.,

$$|S_k| = \left\lceil \frac{\text{Var}(\nabla F(x_{k,0}, \xi)|\mathcal{F}_k)}{\tilde{\theta}^2 \|T(x_{k,0}, \lambda_{k-1}, N_{k-1})\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}} \right\rceil \leq \frac{\text{Var}(\nabla F(x_{k,0}, \xi)|\mathcal{F}_k)}{\tilde{\theta}^2 \|T(x_{k,0}, \lambda_{k-1}, N_{k-1})\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}} + 1.$$

With this selection, the bound on the expected number of inner iterations can be further refined,

$$\mathbb{E}[N_k | \mathcal{F}_k] \leq \frac{C}{\tilde{\omega}^2 (\tilde{\theta}^2 \|T(x_{k,0}, \lambda_{k-1}, N_{k-1})\|^2 + \tilde{a}^2 \tilde{\beta}^{2k})} + \frac{C}{\tilde{\omega}^2 \text{Var}(\nabla F(x_{k,0}, \xi)|\mathcal{F}_k)} \leq \frac{C}{\tilde{\omega}^2 \tilde{a}^2 \tilde{\beta}^{2k}} + \frac{C}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2},$$

where the final bound follows from Assumption 2.5. Thus, the expected total number of inner iterations can be bounded as,

$$\sum_{k=0}^{K_\epsilon} \mathbb{E}[\mathbb{E}[N_k | \mathcal{F}_k]] \leq \sum_{k=0}^{K_\epsilon} \frac{C}{\tilde{\omega}^2 \tilde{a}^2 \tilde{\beta}^{2k}} + \frac{C}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2} = \frac{C}{\tilde{\omega}^2 \tilde{a}^2} \frac{(1/\tilde{\beta})^{2(K_\epsilon+1)} - 1}{(1/\tilde{\beta})^2 - 1} + \frac{C(K_\epsilon+1)}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2},$$

which is of the order $\left(\frac{1}{\beta}\right)^{2K_\epsilon} = \mathcal{O}(\epsilon^{-2})$. The expected number of gradient evaluations in outer iteration $k \geq 0$ can be bounded as

$$\mathbb{E}[|S_k|N_k|\mathcal{F}_k] \leq \frac{C|S_k|}{\tilde{\omega}^2 \tilde{a}^2 \beta^{2k}} + \frac{C|S_k|}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2},$$

which is a deterministic quantity given \mathcal{F}_k . Under the stated parameter selection, if one constructs a batch size sequence $|\tilde{S}_k| = \left\lceil \frac{|S_0|}{\beta^{2k}} \right\rceil$, where $|S_0|$ is the initial batch size, $\mathbb{E}[\|g_{\tilde{S}_k}(x_{k,0}) - \nabla f(x_{k,0})\|^2|\mathcal{F}_k] \leq \tilde{\beta}^{2k} \frac{\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2}{|S_0|} = \tilde{a}^2 \tilde{\beta}^{2k}$ based on Lemma 2.7. Thus, one can satisfy a stronger gradient error condition with $|\tilde{S}_k|$ than Condition 2.1, implying $|S_k| \leq |\tilde{S}_k|$. The expected total number of gradient evaluations can then be bounded as,

$$\begin{aligned} \sum_{k=0}^{K_\epsilon} \mathbb{E}[\mathbb{E}[|S_k|N_k|\mathcal{F}_k]] &\leq \sum_{k=0}^{K_\epsilon} \frac{C|S_k|}{\tilde{\omega}^2 \tilde{a}^2 \beta^{2k}} + \frac{C|S_k|}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2} \leq \sum_{k=0}^{K_\epsilon} \frac{C|\tilde{S}_k|}{\tilde{\omega}^2 \tilde{a}^2 \beta^{2k}} + \frac{C|\tilde{S}_k|}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2} \\ &= \sum_{k=0}^{K_\epsilon} \frac{C|S_0|}{\tilde{\omega}^2 \tilde{a}^2 \beta^{4k}} + \frac{C|S_0|}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2 \beta^{2k}} \\ &= \frac{C|S_0|}{\tilde{\omega}^2 \tilde{a}^2} \frac{(1/\tilde{\beta})^{4(K_\epsilon+1)} - 1}{(1/\tilde{\beta})^2 - 1} + \frac{C|S_0|}{\tilde{\omega}^2 \kappa_G^2 \kappa_\sigma^2} \frac{(1/\tilde{\beta})^{2(K_\epsilon+1)} - 1}{(1/\tilde{\beta})^2 - 1}, \end{aligned}$$

which is of the order $\left(\frac{1}{\beta}\right)^{4K_\epsilon} = \mathcal{O}(\epsilon^{-4})$. □

The results of Theorem 2.8 hold for any deterministic solver satisfying Assumption 2.7, which assumes sublinear convergence in the inner loop. Improved complexity results are possible if faster convergence rates are achieved within the inner loop. For the finite-sum problem (1.2), the derived complexity matches that of a deterministic solver applied directly to the full problem using the true gradient. The current analysis is somewhat pessimistic in terms of gradient evaluations in bounding $|S_k| \leq |\mathcal{S}|$. A more refined argument, similar to that for the expectation problem with $a = 2\sqrt{\omega_1^2 \kappa_g^2 + \omega_2^2}$ in Condition 2.1, yields

$$\sum_{k=0}^{K_\epsilon} N_k |S_k| \leq \sum_{k=0}^{K_\epsilon} \frac{C}{\tilde{\omega}^2 \beta^{2k}} (1 - \beta^k) |\mathcal{S}| = |\mathcal{S}| \frac{C}{\tilde{\omega}^2} \sum_{k=0}^{K_\epsilon} \left(\frac{1}{\beta^{2k}} - \frac{1}{\beta^k} \right)$$

which is of the same order, since $|\mathcal{S}|(\epsilon^{-2} - \epsilon^{-1}) = \mathcal{O}(|\mathcal{S}|\epsilon^{-2})$. Thus, we present the simpler analysis. For the expectation problem, the gradient evaluation complexity of $\mathcal{O}(\epsilon^{-4})$ aligns with the optimal complexity bounds established for stochastic SQP methods [3, 19, 45]. The gradient and inner iteration complexities are independent of the choice of inner loop solver under Assumption 2.7, assuming one subsampled gradient computation per inner iteration. However, the actual computational cost also depends on the subproblem being solved in each inner iteration by the deterministic solver. For instance, SQP-based methods solve a linear system of equations for general nonlinear equality constraints, while projection-based methods compute a projection onto the feasible region. The relative cost of these subproblems depends on the subsampled problem and is an important consideration when selecting the deterministic solver.

2.3 Retrospective Approximation SQP Framework

In this section, we present an instance of the RA-based framework (Framework 2.1) that employs a deterministic SQP solver in the inner loop to solve stochastic problems with deterministic equality constraints. We describe the general framework, illustrate how different options can be implemented within the SQP method, and propose alternative termination criteria to (2.3) that leverage quantities from the SQP method to improve empirical performance.

In outer iteration k , a subsampled problem (1.6) is constructed using a sample set S_k , and the dual variable is initialized according to (2.2) for the inner loop. Within each inner iteration j of outer iteration k , the termination criterion (2.3) is tested. If the criterion is not satisfied, the iterate and dual variable are updated using an SQP method as follows, based on the approaches in [2, 8, 9, 42]. Given an approximation of the Hessian of the Lagrangian $H_{k,j}$ for the subsampled problem, the method computes a search direction via the subproblem given in (1.5). Under the LICQ and assuming the matrix $H_{k,j}$ is positive definite in the null space of $J_{k,j}$, the solution to the SQP subproblem (1.5) can be obtained by solving the linear system

$$\begin{bmatrix} H_{k,j} & J_{k,j}^T \\ J_{k,j} & 0 \end{bmatrix} \begin{bmatrix} d_{k,j} \\ \delta_{k,j} \end{bmatrix} = -T_{S_k}(x_{k,j}, \lambda_{k,j}) + \begin{bmatrix} \rho_{k,j} \\ r_{k,j} \end{bmatrix}, \quad (2.14)$$

where $(d_{k,j}, \delta_{k,j})$ is the search direction, and $(\rho_{k,j}, r_{k,j})$ are vectors that represent the residual in the linear system. The linear system can be solved either exactly, i.e., $\rho_{k,j} = 0, r_{k,j} = 0$, or inexactly, where $\rho_{k,j}$ and $r_{k,j}$ satisfy certain conditions to ensure $(d_{k,j}, \delta_{k,j})$ is a productive search direction [1, 8, 20].

To balance the two possibly competing goals of minimizing the subsampled objective function and minimizing the constraint violation when updating the iterate, the l_1 merit function $\phi_{S_k} : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$ defined over the set of samples S_k as,

$$\phi_{S_k}(x, \tau) = \tau F_{S_k}(x) + \|c(x)\|_1$$

is employed, where $\tau > 0$ is the merit parameter updated as necessary within the inner loop. A local linear model of the merit function $l_{S_k} : \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$ defined as,

$$l_{S_k}(x_{k,j}, \tau_{k,j}, d_{k,j}) = \tau_{k,j} (F_{S_k}(x_{k,j}) + g_{S_k}(x_{k,j})^T d_{k,j}) + \|c_{k,j} + J_{k,j} d_{k,j}\|_1.$$

and the reduction in this local model for the search direction $d_{k,j}$ is used to guide the merit parameter selection, defined as $\Delta l_{S_k} : \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$,

$$\begin{aligned} \Delta l_{S_k}(x_{k,j}, \tau_{k,j}, d_{k,j}) &= l_{S_k}(x_{k,j}, \tau_{k,j}, 0) - l_{S_k}(x_{k,j}, \tau_{k,j}, d_{k,j}) \\ &= -\tau_{k,j} g_{S_k}(x_{k,j})^T d_{k,j} + \|c_{k,j}\|_1 - \|c_{k,j} + J_{k,j} d_{k,j}\|_1 \\ &= -\tau_{k,j} g_{S_k}(x_{k,j})^T d_{k,j} + \|c_{k,j}\|_1 - \|r_{k,j}\|_1. \end{aligned} \quad (2.15)$$

The merit parameter $\tau_{k,j}$ is updated to ensure the search direction is a descent direction for the merit function. With user-defined parameters $\epsilon_\sigma, \epsilon_\tau, \epsilon_d \in (0, 1)$, a trial value $\tau_{k,j}^{trial}$ is obtained as,

$$\tau_{k,j}^{trial} = \begin{cases} \infty & \text{if } g_{S_k}(x_{k,j})^T d_{k,j} + \max\{d_{k,j}^T H_{k,j} d_{k,j}, \epsilon_d \|d_{k,j}\|^2\} \leq 0, \\ \frac{(1-\epsilon_\sigma)(\|c_{k,j}\|_1 - \|r_{k,j}\|_1)}{g_{S_k}(x_{k,j})^T d_{k,j} + \max\{d_{k,j}^T H_{k,j} d_{k,j}, \epsilon_d \|d_{k,j}\|^2\}} & \text{otherwise,} \end{cases}$$

and then the merit parameter is updated via,

$$\tau_{k,j} = \begin{cases} \tau_{k,j-1} & \text{if } \tau_{k,j-1} \leq \tau_{k,j}^{trial}, \\ (1 - \epsilon_\tau) \tau_{k,j}^{trial} & \text{otherwise,} \end{cases} \quad (2.16)$$

ensuring a non-decreasing positive sequence $\{\tau_{k,j}\}$ within the inner loop of each outer iteration k . At the beginning of each outer iteration, the merit parameter is reinitialized ($\{\tau_{k,-1}\} = \bar{\tau} > 0$) rather than carried over from previous iterations. While the subsampled objective function can drastically change across outer iterations, carrying over the merit parameter would not affect convergence. However, carrying over the merit parameter might initialize it at a smaller value than required, in turn significantly slowing down the empirical performance of the SQP method in the inner loop. Finally, the step size $\alpha_{k,j}$ is obtained to satisfy the Armijo sufficient decrease condition on the merit function, i.e.,

$$\phi_{S_k}(x_{k,j} + \alpha_{k,j} d_{k,j}, \tau_{k,j}) \leq \phi_{S_k}(x_{k,j}, \tau_{k,j}) - \eta \alpha_{k,j} \Delta l_{S_k}(x_{k,j}, \tau_{k,j}, d_{k,j}), \quad (2.17)$$

where $\eta \in (0, 1)$ is a user-defined parameter, and the iterate and dual variable are updated via $(x_{k,j+1}, \lambda_{k,j+1}) = (x_{k,j}, \lambda_{k,j}) + \alpha_{k,j}(d_{k,j}, \delta_{k,j})$.

The SQP linear system (2.14) can be solved inexactly using a numerical solver while still ensuring that $(d_{k,j}, \delta_{k,j})$ is a productive direction. We adopt the inexactness conditions as proposed in [1, 8]. Under these conditions, a solution to the SQP subproblem $(d_{k,j}, \delta_{k,j})$ is accepted if it satisfies either inexactness conditions I, which ensures sufficient decrease in the merit function model,

$$\begin{aligned} \Delta l_{S_k}(x_{k,j}, \tau_{k,j-1}, d_{k,j}) &\geq \epsilon_\sigma (1 - \epsilon_{feas}) \max\{\|c_{k,j}\|_1, \|r_{k,j}\| - \|c_{k,j}\|_1\} \\ &\quad + \epsilon_\sigma (1 - \epsilon_{feas}) \tau_{k,j-1} \max\{d_{k,j}^T H_{k,j} d_{k,j}, \epsilon_d \|d_{k,j}\|^2\}, \\ \left\| \begin{bmatrix} \rho_{k,j} \\ r_{k,j} \end{bmatrix} \right\| &\leq \kappa_T \min\{\|T_{S_k}(x_{k,j}, \lambda_{k,j})\|, \|d_{k,j}\|\}, \\ \text{and } \|\rho_{k,j}\| &\leq \kappa' \max\{\|J_{k,j}\|, \|g_{S_k}(x_{k,j})\|\}, \end{aligned} \quad (2.18)$$

where $\epsilon_\sigma, \epsilon_d \in (0, 1)$, $\kappa_T \in (0, \frac{1}{2})$ and $\kappa' > 0$ are user-defined parameters, or inexactness conditions II, which guarantees sufficient decrease in the linear model of the constraints,

$$\|r_{k,j}\| \leq \epsilon_{feas} \|c_{k,j}\| \quad \text{and} \quad \|\rho_{k,j}\| \leq \epsilon_{opt} \|c_{k,j}\|, \quad (2.19)$$

where $\epsilon_{feas}, \epsilon_{opt} \in (0, \frac{1}{2})$ are user-defined parameters. If inexactness condition I is satisfied, the merit parameter remains unchanged, i.e., $\tau_{k,j} = \tau_{k,j-1}$. Compared to the conditions proposed in [8, 9], inexactness condition I is a stronger condition due to the introduction of the norm of the search direction $\|d_{k,j}\|$ in the second inequality in (2.18). This extra condition is not required when using termination criterion (2.3) in the inner loop, but is introduced to ensure convergence when inexact solutions to the SQP linear system (2.14) are used along with the alternate termination criteria for the inner loop described next.

We summarize the SQP-based variant of Framework 2.1 in Algorithm 2.2 and the following remark.

Remark 2.9. *Algorithm 2.2 follows the framework of Framework 2.1, incorporating an SQP method within the inner loop. We make the following remarks regarding the inner loop.*

Algorithm 2.2 Equality Constrained RA-SQP

Inputs: Sequences $\{|S_k|\}$, $\{\gamma_k\} < 1$ and $\{\epsilon_k\} \geq 0$; initial iterate $x_{0,0}$ and dual variable $\lambda_{-1,0}$; initial merit parameter sequence $\{\tau_{k,-1}\} = \bar{\tau} > 0$; other parameters $\epsilon_\sigma, \epsilon_\tau, \epsilon_d, \epsilon_\alpha \in (0, 1)$.

```
1: for  $k = 0, 1, 2, \dots$  do
2:   Construct subsampled problem (1.6) and initialize  $\lambda_{k,0}$  according to (2.2)
3:   for  $j = 0, 1, 2, \dots$  do
4:     if (2.3) is satisfied;  $N_k = j$ , break
5:     Obtain approximation of Hessian of Lagrangian  $H_{k,j}$ 
6:     Solve (2.14) exactly or inexactly and obtain  $(d_{k,j}, \delta_{k,j})$ 
7:     Update merit parameter via (2.16)
8:     Set  $\alpha_{k,j} = 1$ 
9:     while (2.17) is not satisfied do
10:      Update  $\alpha_{k,j} = \epsilon_\alpha \alpha_{k,j}$ 
11:    end while
12:    Set  $(x_{k,j+1}, \lambda_{k,j+1}) = (x_{k,j}, \lambda_{k,j}) + \alpha_{k,j}(d_{k,j}, \delta_{k,j})$ 
13:  end for
14:  Set  $x_{k+1,0} = x_{k,N_k}$ 
15: end for
```

1. **Lagrangian Hessian approximation (Line 5):** While the identity matrix is a common choice for $H_{k,j}$ in constrained stochastic optimization, the deterministic nature of the subsampled problem allows for the use of more sophisticated approximations. In particular, one may employ the subsampled Hessian of the Lagrangian or a quasi-Newton approximation. When quasi-Newton methods, such as BFGS or L-BFGS, are used, the Hessian approximation is updated at each inner iteration and initialized using the final approximation from the previous outer iteration, i.e., in outer iteration k , we set $H_{k,0} = H_{k-1,N_{k-1}}$.
2. **SQP linear system (Line 6):** The SQP subproblem (2.14) can be solved either exactly or inexactly. If a numerical solver is used, one can apply the inexactness conditions (2.18) and (2.19) based on [1, 8] to ensure productive search directions.
3. **Merit parameter and step size (Lines 7-12):** The merit parameter update follows from [1, 2]. Since the objective function can change drastically across outer iterations, the merit parameter is reinitialized at the start of each outer iteration to a user-specified constant $\bar{\tau}$. The step size is computed via a backtracking Armijo line search procedure [42, 53] on the merit function (2.17).

For SQP methods, we propose two alternative termination criteria that leverage quantities from the SQP method to assess the progress of the algorithm on the subsampled problem. The first termination criterion is based on the norm of the search direction,

$$\mathcal{T}_k : \|d_{k,N_k}\| \leq \gamma_k \|d_{k,0}\| + \epsilon_k, \quad (2.20)$$

where the search direction $d_{k,j}$ serves as a measure of optimality. The second termination criterion is based on the decrease in the merit function model,

$$\mathcal{T}_k : \Delta l_{S_k}(x_{k,N_k}, \tau_{k,N_k}, d_{k,N_k}) \leq \gamma_k \min\{\Delta l_{S_k}(x_{k,0}, \tau_{k,0}, d_{k,0}), \kappa_d \|d_{k,0}\|^2\} + \epsilon_k, \quad (2.21)$$

where $\kappa_d > 0$. While the decrease in the merit function model is a bounded quantity, the bound can be quite large, necessitating the inclusion of the search direction norm for proper scaling. In practice, κ_d is chosen to be quite large, and the search direction norm term on the right-hand side of (2.21) is typically inactive.

The termination criterion (2.3), based on the KKT error of the subsampled problem, while theoretically sound presents challenges in practice. Specifically, the two components of the KKT error, $\nabla_x \mathcal{L}_{S_k}(x_{k,j}, \lambda_{k,j})$ and $c_{k,j}$ can differ significantly in magnitude, potentially causing the algorithm to focus disproportionately on either reducing the objective function or satisfying the constraints, rather than a balanced approach. Termination criterion (2.20) avoids these scaling issues as it uses the norm of the search direction instead of combining multiple quantities. Termination criterion (2.21) on the other hand explicitly balances the objective function and constraint violation using the merit parameter within the merit function model.

Unlike (2.3), criteria (2.20) and (2.21) require the SQP linear system (2.14) solution for evaluation. Thus, they perform one extra solve of the SQP subproblem beyond the number of inner iterations executed during each outer iteration. When employing criteria (2.20) or (2.21), we use the **Carry-Over** option in (2.2) for dual variable initialization and require the Hessian approximation of the Lagrangian to be carried over between outer iterations, i.e., $H_{k,0} = H_{k-1,N_{k-1}}$. This Hessian requirement is naturally satisfied when using either a constant Hessian approximation or a quasi-Newton method. Since $x_{k,0} = x_{k-1,N_{k-1}}$ and $J_{k,0} = J_{k-1,N_{k-1}}$, the Hessian approximation remains positive definite over the null space of $J_{k,0}$. Thus, if this restriction poses challenges, e.g., when using subsampled Hessians, one can either use a different Hessian approximation for the first inner iteration or instead apply the termination criterion (2.3). If additional assumptions regarding the error in the subsampled Hessians are made, the requirement to carry over the Hessian can be dropped. However, for brevity, we maintain this requirement in this work. Termination criteria (2.20) and (2.21) can be used when the SQP linear system is solved exactly or inexactly. While we limit our analysis and empirical exploration to these termination criterion and find that (2.21) performs best empirically, one can consider other criteria based on the deterministic method employed in the inner loop.

2.4 Retrospective Approximation SQP Convergence Analysis

In this section, we present the convergence properties of Algorithm 2.2 for problems with only equality constraints. We first establish the required conditions under which Algorithm 2.2 adheres to the theoretical results developed for Framework 2.1 in Subsection 2.2. We then present results for the alternate termination tests proposed in Subsection 2.3 for the SQP method. We begin by stating the assumptions required to ensure the well-posedness of the algorithm.

Assumption 2.8. *Let $\chi \subset \mathbb{R}^n$ be a closed bounded convex set containing the sequence of iterates generated by any run of Algorithm 2.2. The objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable and its gradient function $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz continuous over χ . The subsampled objective function $F_{S_k} : \mathbb{R}^n \rightarrow \mathbb{R}$ for all $k \geq 0$ is continuously differentiable with Lipschitz continuous gradients over χ . The constraint function $c : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuously differentiable and the Jacobian function, $J = \nabla c^T : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$ is Lipschitz continuous over χ . The Jacobian $J(x)$ over $x \in \chi$ has full row rank (LICQ). The sequence of symmetric Hessian approximation matrices $\{H_{k,j}\}$ is bounded in norm by $\kappa_H > 0$ such that $\|H_{k,j}\| \leq \kappa_H$. In addition, there exists a constant $\mu_H > 0$ such that the matrix $u^T H_{k,j} u \geq \mu_H \|u\|^2$ for all $u \in \mathbb{R}^n$, where $J_{k,j} u = 0$.*

Under Assumption 2.8, there exist $\kappa_J, \kappa_g, \kappa_{HJ} > 0$ such that $\{\|J_{k,j}\|\} \leq \kappa_J$, $\{\|\nabla f(x_{k,j})\|\} \leq \kappa_g$ and $\left\{\left\|\begin{bmatrix} H_{k,j} & J_{k,j}^T \\ J_{k,j} & 0 \end{bmatrix}^{-1}\right\|\right\} \leq \frac{1}{\kappa_{HJ}}$. Thus, Assumption 2.1 and 2.6 are satisfied under Assumption 2.8. Compared to [1–3, 8, 45], this is a stronger assumption for SQP based algorithms for stochastic problems with equality constraints, as it requires a closed bounded set of iterates and all subsampled problems to be smooth functions. These requirements stem from the fundamental structure of RA, which treats each subsampled problem as a deterministic problem, requiring each to satisfy the standard assumptions for deterministic SQP methods, where the objective functions must be smooth with bounded gradients and be bounded below over the set of iterates χ . In extreme cases, this assumption may not hold for every subsampled problem. As a result, in practice, an upper limit is imposed on the number of inner iterations per outer iteration.

We now formally state the well-posedness of Algorithm 2.2.

Theorem 2.10. *Suppose Assumptions 2.4, 2.5 and 2.8 hold. Then Algorithm 3.1 is well-posed and the inner loop converges at a sublinear rate satisfying Assumption 2.7, achieving the gradient and inner iteration complexity described in Theorem 2.8.*

Proof. The convergence of the SQP algorithm in the inner loop is established in [19] for exact subproblem solutions and in [8] for inexact subproblem solutions, under Assumption 2.8. These results apply because, over the bounded set of iterates χ , the subsampled problem $F_{S_k}(x)$ for all $k \geq 0$ is bounded below and has bounded gradients, thereby satisfying Assumption 2.2. Under Assumption 2.8 and by [19, Theorem 1], when exact search direction solutions are used, the number of inner loop iterations to achieve an ϵ accurate solution for the subsampled problem in outer iteration k is

$$N_k \leq \left\lceil \frac{\bar{\tau}(F_{S_k}(x_{k,0}) - \min_{x \in \chi} F_{S_k}(x)) + \|c_{k,0}\|_1}{\kappa_k} \right\rceil \frac{1}{\epsilon^2},$$

where κ_k is a finite constant based on $F_{S_k}(x)$ and the term within the right-hand side brackets is bounded as the set of iterates χ is a bounded set, thereby satisfying Assumption 2.7. This ensures that Algorithm 2.2 achieves the complexity results stated in Theorem 2.8. Similar results can be obtained when inexact search direction solutions are used, as shown in [1, 8, 9, 20]. \square

Theorem 2.10 establishes that Algorithm 2.2 achieves the optimal gradient ($\mathcal{O}(\epsilon^{-4})$) and SQP linear system subproblem complexity ($\mathcal{O}(\epsilon^{-2})$) for stochastic problems with equality constraints. We now analyze the convergence of Algorithm 2.2 under the alternate termination criteria proposed in Subsection 2.3. To establish the results in Subsection 2.2, we need to show an error bound analogous to Lemma 2.2 and adapt Condition 2.1 to align with the optimality measure used in the termination criteria.

We first consider the termination criterion (2.20), which is based on the norm of the search direction. We define $(\hat{d}_{k,j}, \hat{\delta}_{k,j})$ as the exact solutions to the SQP linear system (2.14), and $(d_{k,j}^{true}, \delta_{k,j}^{true})$ as the exact solutions to the SQP linear system (2.14) when the KKT error of the true problem is used instead of the KKT error of the subsampled problem. If $d_{k,j}^{true} = 0$, then $x_{k,j}$ is a first-order stationary point for the true problem (1.1).

Lemma 2.11. *Suppose Assumption 2.8 holds. Then for all $k \geq 0$, $j \geq 0$, if κ_T in (2.18), and ϵ_{feas} and ϵ_{opt} in (2.19) are chosen to be such that $\frac{\kappa_T}{\kappa_{HJ}} < \frac{1}{2}$ and $\frac{\kappa_J(\epsilon_{feas} + \epsilon_{opt})}{\kappa_{HJ}} < 1$, there exists*

$\eta_d \in (0, 1)$ such that the search direction generated by Algorithm 2.2 satisfies,

$$\|d_{k,j} - \hat{d}_{k,j}\| \leq \eta_d \|\hat{d}_{k,j}\|.$$

Proof. If the SQP linear system (2.14) is solved exactly, the relation is trivially satisfied with $\eta_d = 0$. When $d_{k,j}$ is an inexact solution, it satisfies either (2.18) or (2.19). By the definitions of $d_{k,j}$ and $\hat{d}_{k,j}$ from (2.14), we have $\begin{bmatrix} H_{k,j} & J_{k,j}^T \\ J_{k,j} & 0 \end{bmatrix} \begin{bmatrix} d_{k,j} - \hat{d}_{k,j} \\ \delta_{k,j} - \hat{\delta}_{k,j} \end{bmatrix} = \begin{bmatrix} \rho_{k,j} \\ r_{k,j} \end{bmatrix}$. Thus, by Assumption 2.8,

$$\|d_{k,j} - \hat{d}_{k,j}\| \leq \left\| \begin{bmatrix} d_{k,j} - \hat{d}_{k,j} \\ \delta_{k,j} - \hat{\delta}_{k,j} \end{bmatrix} \right\| = \left\| \begin{bmatrix} H_{k,j} & J_{k,j}^T \\ J_{k,j} & 0 \end{bmatrix}^{-1} \begin{bmatrix} \rho_{k,j} \\ r_{k,j} \end{bmatrix} \right\| \leq \kappa_{HJ}^{-1} \left\| \begin{bmatrix} \rho_{k,j} \\ r_{k,j} \end{bmatrix} \right\|.$$

If the inexactness conditions (2.18) are satisfied,

$$\|d_{k,j} - \hat{d}_{k,j}\| \leq \frac{\kappa_T}{\kappa_{HJ}} \min \{ \|T_{S_k}(x_{k,j}, \lambda_{k,j})\|, \|d_{k,j}\| \} \leq \frac{\kappa_T}{\kappa_{HJ}} \|d_{k,j}\| \leq \eta_d \|\hat{d}_{k,j}\|,$$

where $\eta_d = \frac{\frac{\kappa_T}{\kappa_{HJ}}}{1 - \frac{\kappa_T}{\kappa_{HJ}}} < 1$ as $\frac{\kappa_T}{\kappa_{HJ}} < \frac{1}{2}$. If the inexactness conditions (2.19) are satisfied,

$$\|d_{k,j} - \hat{d}_{k,j}\| \leq \frac{\|\rho_{k,j}\| + \|r_{k,j}\|}{\kappa_{HJ}} \leq \frac{\epsilon_{feas} + \epsilon_{opt}}{\kappa_{HJ}} \|c_{k,j}\| = \frac{\epsilon_{feas} + \epsilon_{opt}}{\kappa_{HJ}} \|J_{k,j} \hat{d}_{k,j}\| \leq \eta_d \|\hat{d}_{k,j}\|,$$

where $\eta_d = \frac{\kappa_J(\epsilon_{feas} + \epsilon_{opt})}{\kappa_{HJ}} < 1$ as specified. \square

We now establish an error bound similar to Lemma 2.2, under the termination criterion (2.20).

Lemma 2.12. *Suppose Assumption 2.8 holds. Then for all $k \geq 0$, the true search directions for the outer iterates generated by Algorithm 2.2 with termination criterion (2.20) satisfy*

$$\begin{aligned} \|d_{k,N_k}^{true}\| &\leq \gamma_k \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\epsilon_k}{1-\eta_d} + \frac{1+\eta_d}{1-\eta_d} \kappa_{HJ}^{-1} \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| \\ &\quad + \gamma_k \frac{1+\eta_d}{1-\eta_d} \kappa_{HJ}^{-1} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|. \end{aligned}$$

For the expectation problem (1.3), for all $k \geq 0$

$$\begin{aligned} \mathbb{E}[\|d_{k,N_k}^{true}\| | \mathcal{F}_k] &\leq \gamma_k \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\mathbb{E}[\epsilon_k | \mathcal{F}_k]}{1-\eta_d} + \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \mathbb{E}[\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\ &\quad + \gamma_k \frac{1+\eta_d}{1-\eta_d} \kappa_{HJ}^{-1} \mathbb{E}[\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k]. \end{aligned}$$

Proof. We start with a bound on the error in the search direction due to subsampled gradients. By the definition of $d_{k,j}^{true}$ and $\hat{d}_{k,j}$ from (2.14),

$$\begin{bmatrix} H_{k,j} & J_{k,j}^T \\ J_{k,j} & 0 \end{bmatrix} \begin{bmatrix} \hat{d}_{k,j} - d_{k,j}^{true} \\ \hat{\delta}_{k,j} - \delta_{k,j}^{true} \end{bmatrix} = - \begin{bmatrix} g_{S_k}(x_{k,j}) - \nabla f(x_{k,j}) \\ 0 \end{bmatrix}.$$

By Assumption 2.8,

$$\|\hat{d}_{k,j} - d_{k,j}^{true}\| \leq \left\| \begin{bmatrix} \hat{d}_{k,j} - d_{k,j}^{true} \\ \hat{\delta}_{k,j} - \delta_{k,j}^{true} \end{bmatrix} \right\| \leq \kappa_{HJ}^{-1} \|\nabla f(x_{k,j}) - g_{S_k}(x_{k,j})\|.$$

At the end of outer iteration k , the norm of the true search direction can be bounded as,

$$\begin{aligned}
\|d_{k,N_k}^{true}\| &\leq \|d_{k,N_k}\| + \|\hat{d}_{k,N_k} - d_{k,N_k}\| + \|d_{k,N_k}^{true} - \hat{d}_{k,N_k}\| \\
&\leq \gamma_k \|d_{k,0}\| + \epsilon_k + \eta_d \|\hat{d}_{k,N_k}\| + \|d_{k,N_k}^{true} - \hat{d}_{k,N_k}\| \\
&\leq \gamma_k \|\hat{d}_{k,0}\| + \gamma_k \|d_{k,0} - \hat{d}_{k,0}\| + \epsilon_k + \eta_d \|d_{k,N_k}^{true}\| + \eta_d \|\hat{d}_{k,N_k} - d_{k,N_k}^{true}\| + \|d_{k,N_k}^{true} - \hat{d}_{k,N_k}\| \\
&\leq \gamma_k (1 + \eta_d) \left(\|\hat{d}_{k,0} - d_{k,0}^{true}\| + \|d_{k,0}^{true}\| \right) + \epsilon_k + (1 + \eta_d) \|d_{k,N_k}^{true} - \hat{d}_{k,N_k}\| + \eta_d \|d_{k,N_k}^{true}\| \\
&= \gamma_k \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\epsilon_k}{1-\eta_d} + \frac{1+\eta_d}{1-\eta_d} \|d_{k,N_k}^{true} - \hat{d}_{k,N_k}\| + \gamma_k \frac{1+\eta_d}{1-\eta_d} \|\hat{d}_{k,0} - d_{k,0}^{true}\|,
\end{aligned}$$

where the second inequality follows from the termination criterion (2.20) and Lemma 2.11, the fourth inequality follows from Lemma 2.11 and the final equality follows from subtracting $\eta_d \|d_{k,N_k}^{true}\|$ from both sides and dividing by $1-\eta_d$, and from $d_{k,0}^{true} = d_{k-1,N_{k-1}}^{true}$ as $x_{k,0} = x_{k-1,N_{k-1}}$, $H_{k,0} = H_{k-1,N_{k-1}}$ and $\lambda_{k,0} = \lambda_{k-1,N_{k-1}}$. Employing the search direction error bound in the above inequality yields the desired deterministic bound. Taking the conditional expectation given \mathcal{F}_k yields the expectation bound as $d_{k-1,N_{k-1}}^{true}$ is defined under \mathcal{F}_k . \square

Using Lemma 2.12, we can establish results analogous to Theorem 2.3 by incorporating the new scaling factor $\kappa_{HJ}^{-1} \left(\frac{1+\eta_d}{1-\eta_d} \right)$ and requiring the termination criterion parameter to satisfy $\gamma_k < \frac{1-\eta_d}{1+\eta_d}$ for all $k \geq 0$. When exact solutions to the SQP linear system are used, $\eta_d = 0$, and it suffices to have $\{\gamma_k\} < 1$. We now modify Condition 2.1 to align the adaptive sampling strategy with termination criterion (2.20).

Condition 2.2. *At the beginning of outer iteration $k \geq 0$ in Algorithm 2.2 with termination criterion (2.20), the batch size $|S_k|$ is selected such that:*

1. *For the finite-sum problem (1.2): With constants $\theta, \hat{\theta}, a \geq 0$ and $\beta \in (0, 1)$,*

$$\begin{aligned}
\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 &\leq \theta^2 \|d_{k,0}^{true}\|^2 + a^2 \beta^{2k}, \\
\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\|^2 &\leq \hat{\theta}^2 (\theta^2 \|d_{k,0}^{true}\|^2 + a^2 \beta^{2k}).
\end{aligned}$$

2. *For the expectation problem (1.3): With constants $\tilde{\theta}, \tilde{a} \geq 0$ and $\tilde{\beta} \in (0, 1)$,*

$$\mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 | \mathcal{F}_k] \leq \tilde{\theta}^2 \|d_{k,0}^{true}\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}.$$

Following the approach of Theorem 2.5, we can establish linear convergence for the true search direction norm across outer iterations under Assumption 2.8, termination criterion (2.20) and Condition 2.2. The result is formalized in the following corollary.

Theorem 2.13. *Suppose Assumptions 2.4 and 2.8 hold and that the batch size sequence $\{|S_k|\}$ is chosen to satisfy Condition 2.2 in Algorithm 2.2 with termination criterion (2.20).*

1. *For the finite-sum problem (1.2): For all $k \geq 0$, if termination criterion parameters are chosen as $0 \leq \gamma_k \leq \gamma < \frac{1-\eta_d}{1+\eta_d}$, $\epsilon_k = \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k$ with $\omega, \hat{\omega} \geq 0$ and Condition 2.2 parameters are chosen such that $a_1 = \frac{1}{1-\eta_d} \left[\gamma(1+\eta_d) + \theta \left(\omega + (\gamma + \hat{\theta}) \frac{1+\eta_d}{\kappa_{HJ}} \right) \right] < 1$, then, the true search direction norm converges at a linear rate across outer iterations, i.e.,*

$$\|d_{k,N_k}^{true}\| \leq \max\{a_1 + \nu, \beta\}^{k+1} \max\{\|d_{0,0}^{true}\|, \frac{a_2}{\nu}\},$$

where $a_2 = \frac{a}{1-\eta_d} \left(\omega + (\gamma + \hat{\theta}) \frac{1+\eta_d}{\kappa_{HJ}} \right) + \frac{\hat{\omega}}{1-\eta_d}$ and $\nu > 0$ such that $a_1 + \nu < 1$.

2. For the expectation problem (1.3): For all $k \geq 0$, if Assumption 2.5 holds, the termination criterion parameters are chosen as $0 \leq \gamma_k \leq \tilde{\gamma} < \frac{1-\eta_d}{1+\eta_d}$, $\epsilon_k = \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0})|\mathcal{F}_k)}{|S_k|}}$ where $\tilde{\omega} \geq 0$ and Condition 2.2 parameters are chosen such that $\tilde{a}_1 = \frac{1}{1-\eta_d} \left[\tilde{\gamma}(1+\eta_d) + \tilde{\theta} \left(\tilde{\omega} + \frac{1+\eta_d}{\kappa_{HJ}} \left(\frac{\epsilon_G + \kappa_g}{\kappa_\sigma} + \tilde{\gamma} \right) \right) \right] < 1$, then, the expected true search direction norm converges at a linear rate across outer iterations, i.e.,

$$\mathbb{E} [\|d_{k,N_k}^{true}\|] \leq \max\{\tilde{a}_1 + \tilde{\nu}, \tilde{\beta}\}^{k+1} \max\{\|d_{0,0}^{true}\|, \frac{\tilde{a}_2}{\tilde{\nu}}\},$$

where $\tilde{a}_2 = \frac{\tilde{a}}{1-\eta_d} \left[\tilde{\omega} + \frac{1+\eta_d}{\kappa_{HJ}} \left(\frac{\epsilon_G + \kappa_g}{\kappa_\sigma} + \tilde{\gamma} \right) \right]$ and $\tilde{\nu} > 0$ such that $\tilde{a}_1 + \tilde{\nu} < 1$.

Proof. The proof follows the same procedure as the proof of Theorem 2.5, with adjustments to account for the changes in Lemma 2.12 relative to Lemma 2.2. Specifically, it incorporates the scaling factor $\kappa_{HJ}^{-1} \left(\frac{1+\eta_d}{1-\eta_d} \right)$ and the condition $\gamma_k < \frac{1-\eta_d}{1+\eta_d}$. The complete proof is presented in Appendix B. \square

Using the same approach, we can also establish a result analogous to Corollary 2.6 for the case of a geometrically increasing batch size sequence. This variant removes several of the requirements imposed in Theorem 2.13, while still guaranteeing convergence under milder conditions.

We now analyze the termination criterion (2.21), employing the merit function model decrease. We establish convergence under this criterion by showing that it is equivalent, up to constant factors, to the search direction norm-based termination criterion (2.20).

Lemma 2.14. *Suppose Assumption 2.8 holds. Then, there exists $\hat{\gamma} < 1$ such that termination criterion (2.21) with $\{\gamma_k\} < \hat{\gamma}$ satisfies termination criterion (2.20) for all $k \geq 0$ in Algorithm 2.2.*

Proof. Under termination criterion (2.21),

$$\begin{aligned} \|d_{k,N_k}\|^2 &\leq \kappa_{d,\Delta l} \Delta l_{S_k}(x_{k,N_k}, \tau_{k,N_k}, d_{k,N_k}) \\ &\leq \kappa_{d,\Delta l} \gamma_k \min\{\Delta l_{S_k}(x_{k,0}, \tau_{k,0}, d_{k,0}), \kappa_d \|d_{k,0}\|^2\} + \kappa_{d,\Delta l} \epsilon_k \\ &\leq \kappa_{d,\Delta l} \gamma_k \kappa_d \|d_{k,0}\|^2 + \kappa_{d,\Delta l} \epsilon_k, \end{aligned}$$

where the first inequality follows from [1, Lemma 4.11] with $\kappa_{d,\Delta l} > 0$, since, by Lemma 2.11 and Assumption 2.8, the error $\|d_{k,N_k} - \hat{d}_{k,N_k}\|$ is bounded and from [8, Lemma 4.7], the merit parameter is bounded below. Therefore, termination criterion (2.20) is satisfied with parameters $(\sqrt{\kappa_{d,\Delta l} \gamma_k \kappa_d}, \sqrt{\kappa_{d,\Delta l} \epsilon_k})$, which satisfy the conditions for convergence if γ_k is sufficiently small. \square

Thus, we have established the convergence of Algorithm 2.2 under termination criterion (2.21). Similarly, Condition 2.2 can be adapted for this termination criterion (2.21). We first define $\tau_{k,0}^{true}$ as the merit parameter obtained in the first inner iteration of outer iteration k if the true problem gradient and the true search direction $d_{k,0}^{true}$ are used for the update. We also define the true problem merit function model decrease as

$$\Delta l(x_{k,0}, \tau_{k,0}^{true}, d_{k,0}^{true}) = -\tau_{k,0}^{true} \nabla f(x_{k,0})^T d_{k,0}^{true} + \|c_{k,0}\|_1 - \|c_{k,0} + J_{k,0} d_{k,0}^{true}\|_1.$$

Next, we present the sampling condition to accompany termination criterion (2.21).

Condition 2.3. *At the beginning of outer iteration $k \geq 0$ in Algorithm 2.2 with termination criterion (2.21), the batch size $|S_k|$ is selected such that:*

1. For the finite-sum problem (1.2): With constants $\theta, \hat{\theta}, a \geq 0$ and $\beta \in (0, 1)$,

$$\begin{aligned}\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 &\leq \theta^2 \min\{\Delta l(x_{k,0}, \tau_{k,0}^{true}, d_{k,0}^{true}), \kappa_d \|d_{k,0}^{true}\|^2\} + a^2 \beta^{2k}, \\ \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\|^2 &\leq \hat{\theta}^2 (\theta^2 \min\{\Delta l(x_{k,0}, \tau_{k,0}^{true}, d_{k,0}^{true}), \kappa_d \|d_{k,0}^{true}\|^2\} + a^2 \beta^{2k}).\end{aligned}$$

2. For the expectation problem (1.3): With constants $\tilde{\theta}, \tilde{a} \geq 0$ and $\tilde{\beta} \in (0, 1)$,

$$\mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 | \mathcal{F}_k] \leq \tilde{\theta}^2 \min\{\Delta l(x_{k,0}, \tau_{k,0}^{true}, d_{k,0}^{true}), \kappa_d \|d_{k,0}^{true}\|^2\} + \tilde{a}^2 \tilde{\beta}^{2k}.$$

It is straightforward to observe that satisfying Condition 2.3 also implies Condition 2.2. If the parameter θ (and $\tilde{\theta}$) in Condition 2.3 is chosen to be sufficiently small so as to meet the requirements in Theorem 2.13, then linear convergence across outer iterations can be achieved while using the termination criterion (2.21), under Condition 2.3.

We established the convergence properties of Algorithm 2.2 under the alternate termination criterion (2.20) and (2.21) with inexact solutions to the SQP linear system (2.14) that satisfy either (2.18) or (2.19). When the SQP linear system is solved exactly, $\eta_d = 0$ in Lemma 2.11 reducing the conditions on parameters under the alternate termination criterion. The results presented in this section can be derived for termination criteria based on other optimality metrics, provided an error bound analogous to Lemma 2.2 can be established.

3 Inequality Constrained Problems

In this section, we present our proposed method for solving problem (1.1) with general nonlinear equality and inequality constraints. We begin by describing the challenges introduced by the presence of general nonlinear inequality constraints and highlighting the differences from the framework proposed in Section 2. We describe the proposed algorithm in Subsection 3.1 and provide theoretical guarantees in Subsection 3.2.

As mentioned in Section 1, x^* is a first-order stationary point for problem (1.1) with general nonlinear equality and inequality constraints if there exist $\lambda_E^* \in \mathbb{R}^{m_E}$ and $\lambda_I^* \in \mathbb{R}^{m_I}$ such that,

$$\nabla_x \mathcal{L}(x^*, \lambda_E^*, \lambda_I^*) = 0, \quad c_E(x^*) = 0, \quad c_I(x^*) \leq 0, \quad c_I(x^*) \odot \lambda_I^* = 0, \quad \lambda_I^* \geq 0.$$

While one could extend the general framework from Subsection 2.1 by incorporating the complementary slackness error from (1.4) into the termination criterion (2.3), instead we directly present a solver-specific termination criterion due to the various challenges associated with the general constrained setting. First, incorporating the complementary slackness error would add a third term in termination criterion (2.3), exacerbating the scaling issue discussed in Subsection 2.3. Second, evaluating the termination criterion would require good dual variable estimates for the subsampled problem, something not readily provided or updated by deterministic methods when dealing with general nonlinear constraints. Unlike Section 2, if the deterministic method in the inner loop does not provide good dual variable estimates, computing them requires solving a linear or quadratic program due to the presence of the complimentary slackness conditions in (1.4), which is significantly more expensive than solving the least squares problem in Section 2. Thus, to address these challenges, in order to establish theoretical results and achieve good empirical performance, we propose a framework tailored to a deterministic solver with a termination criterion that leverages

solver-specific quantities. Given our goal of handling general nonlinear constraints, we adopt an SQP-based deterministic algorithm for the inner loop.

An important consideration when SQP-based methods are used for general nonlinear constraints are infeasible stationary points. The convergence of SQP-based methods towards a feasible solution can be characterized as reaching a stationary point of an optimization problem that minimizes the constraint violation, i.e., $\min_{x \in \mathbb{R}^n} \frac{1}{2} \|c_E(x)\|^2 + \frac{1}{2} \|[c_I(x)]_+\|^2$. In the case of only equality constraints, LICQ conditions ensure that every stationary point to this problem is feasible. However, no such guarantees exist for general nonlinear constraints. If such a stationary point is reached and found to be infeasible, termination is necessary, returning an infeasible stationary point. The precise definition of this concept varies by algorithm and is provided later.

When solving a problem with general nonlinear constraints using SQP methods, one must solve subproblem (1.5) at each iteration. This subproblem is more computationally expensive than the SQP linear system (2.14) in the equality constraint only setting, and its feasibility is not guaranteed even when the original problem is feasible, as demonstrated in [42, Example 19.2]. Few works exist for the general constrained stochastic setting that handle this issue. Among these, [38] bypassed the issue of infeasibility by assuming that every SQP subproblem (1.5) is feasible, and, subsequent works have tackled this challenge using techniques such as step decomposition [21] and robust subproblems [54], which ensure subproblem feasibility while identifying infeasible stationary points. The latter two are techniques adapted from deterministic constrained optimization. Naturally, we wish to incorporate such SQP methods within the RA framework. We use SQP methods with robust subproblems [6, 44, 54], instead of a slack variable reformulation of the constraints, i.e., $s \in \mathbb{R}^{m_I}$, $c_I(x) + s = 0$ and $s \geq 0$, as is common, e.g., [21, 42, 52]. Within the RA framework, we observe that the slack variable reformulation incurs significant computational costs due to the need to update slack variables and track the active set, even after the inner loop achieves sufficient optimality for termination. In contrast, robust-SQP methods eliminate these inefficiencies while preserving strong convergence properties.

3.1 Algorithm Description

In this section, we describe the proposed method for solving problems with stochastic objective functions and general nonlinear constraints based on Framework 1.1 and the robust-SQP method. Specifically, we adopt the robust-SQP algorithm [6] that uses a two-step approach. The two-step approach first minimizes the linearized constraint violation and then proceeds to compute a search direction that minimizes a quadratic approximation of the objective function while maintaining the constraint violation achieved in the first step. This ensures the subproblems for the search direction calculation are feasible, contrary to subproblem (1.5), and are able to detect infeasible stationary points via the constraint minimization problem in the first step. We now detail the inner loop update mechanism and the termination criterion for the inner loop. We adopt notation similar to that in Section 2, where $x_{k,j}$ denotes the iterate at the outer iteration k and inner iteration j .

In outer iteration k , a set of samples S_k is obtained to form the subsampled problem (1.6). To update the iterate in inner iteration j , we first compute a search direction $p_{k,j}$ that minimizes the linearized constraint violation, based on a chosen measure of constraint violation. This violation can be measured using either the l_∞ or the l_1 norm, resulting in the following linear programs (l_∞

on the left and l_1 on the right),

$$\begin{aligned}
& \min_{\substack{y \in \mathbb{R}, \\ p_{k,j} \in \mathbb{R}^n}} y & \quad \min_{\substack{y_E \in \mathbb{R}^{m_E}, \\ y_I \in \mathbb{R}^{m_I}, \\ p_{k,j} \in \mathbb{R}^n}} e_{m_E}^T y_E + e_{m_I}^T y_I \\
& \text{s.t. } c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) p_{k,j} \geq -y e_{m_E} & \text{or} & \quad \text{s.t. } c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) p_{k,j} \geq -y_E \\
& c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) p_{k,j} \leq y e_{m_E} & & \quad c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) p_{k,j} \leq y_E \\
& c_I(x_{k,j}) + \nabla c_I^T(x_{k,j}) p_{k,j} \leq y e_{m_I} & & \quad c_I(x_{k,j}) + \nabla c_I^T(x_{k,j}) p_{k,j} \leq y_I \\
& y \geq 0, \|p_{k,j}\|_\infty \leq \sigma_p & & \quad y_E \geq 0, y_I \geq 0, \|p_{k,j}\|_1 \leq \sigma_p,
\end{aligned} \tag{3.1}$$

where y is the minimized linearized constraint violation when measured with the l_∞ norm, (y_E, y_I) represent the minimized linearized constraint violation components when measured with the l_1 norm, and σ_p is the bound on the norm of the search direction. The linear program corresponding to the l_∞ norm involves fewer variables and constraints compared to the one based on the l_1 norm. However, the l_1 norm provides a more precise measure of constraint violation which can potentially lead to improved performance. Although other norms can be used to measure constraint violation, they often lead to more challenging subproblems, for example, the Euclidean norm results in quadratically constrained quadratic programs (QCQPs). Thus, we restrict our focus to the l_∞ and l_1 norms in this work.

Next, a quadratic program is solved to compute the search direction $d_{k,j}$, which minimizes a quadratic approximation of the objective function at $x_{k,j}$, while preserving the constraint violation achieved by $p_{k,j}$. The resulting quadratic programs (l_∞ on the left and l_1 on the right) are

$$\begin{aligned}
& \min_{d_{k,j} \in \mathbb{R}^n} g_{S_k}(x_{k,j})^T d_{k,j} + \frac{1}{2} d_{k,j}^T H_{k,j} d_{k,j} & \quad \min_{d_{k,j} \in \mathbb{R}^n} g_{S_k}(x_{k,j})^T d_{k,j} + \frac{1}{2} d_{k,j}^T H_{k,j} d_{k,j} \\
& \text{s.t. } c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) d_{k,j} \geq -y e_{m_E} & \text{or} & \quad \text{s.t. } c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) d_{k,j} \geq -y_E \\
& c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) d_{k,j} \leq y e_{m_E} & & \quad c_E(x_{k,j}) + \nabla c_E^T(x_{k,j}) d_{k,j} \leq y_E \\
& c_I(x_{k,j}) + \nabla c_I^T(x_{k,j}) d_{k,j} \leq y e_{m_I} & & \quad c_I(x_{k,j}) + \nabla c_I^T(x_{k,j}) d_{k,j} \leq y_I \\
& \|d_{k,j}\|_\infty \leq \sigma_d & & \quad \|d_{k,j}\|_1 \leq \sigma_d,
\end{aligned} \tag{3.2}$$

where $y, (y_E, y_I)$ are the outputs from solving (3.1), $H_{k,j}$ is an approximation of the Hessian of the Lagrangian, and σ_d is the bound on the norm of the search direction. The bound constraint parameters σ_p and σ_d are user defined constants such that $0 < \sigma_p < \sigma_d < \infty$, ensuring that (3.2) remains feasible. These finite bound constraints on the search direction are essential for the identification of infeasible stationary points and to establish convergence theory with $\{\|d_{k,j}\|\} \rightarrow 0$, as demonstrated in [6, Example 2.1]. While σ_p and σ_d can be constants, in practice they are often set proportional to the constraint violation of the current iterate, following the approach in [18, 54]; further details are provided in Section 4.

A merit function is employed to balance the two possibly competing goals of decreasing the objective function and reducing constraint violation. Depending on the norm used to measure the constraint violation, the merit function $\phi_{S_k} : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$ is defined as (l_∞ on left and l_1 on right),

$$\phi_{S_k}(x, \tau) = \tau F_{S_k}(x) + \left\| \begin{array}{c} c_E(x) \\ [c_I(x)]_+ \end{array} \right\|_\infty \quad \text{or,} \quad \phi_{S_k}(x, \tau) = \tau F_{S_k}(x) + \left\| \begin{array}{c} c_E(x) \\ [c_I(x)]_+ \end{array} \right\|_1, \tag{3.3}$$

where $\tau > 0$ is the merit parameter. Note that in (3.3) the penalty parameter τ is on the objective function whereas in many robust-SQP methods [6, 44, 54] the penalty parameter ρ is on the constraint

violation term. The two forms are equivalent for appropriate choices of these parameters, and we prefer the former to maintain consistency with the merit function defined in Subsection 2.3 for the equality constrained setting. To update the merit parameter, we define the predicted reduction in a linear model of the merit function $\Delta l_{S_k} : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$ as,

$$\Delta l_{S_k}(x_{k,j}, d_{k,j}, \tau_{k,j}) = -\tau_{k,j} g_{S_k}(x_{k,j})^T d_{k,j} + \Delta_c, \quad (3.4)$$

where Δ_c denotes the improvement in the linearized constraint violation, defined as (l_∞ on the left and l_1 on the right),

$$\Delta_c = \left\| \begin{bmatrix} c_E(x_{k,j}) \\ [c_I(x_{k,j})]_+ \end{bmatrix} \right\|_\infty - y \quad \text{or,} \quad \Delta_c = \left\| \begin{bmatrix} c_E(x_{k,j}) \\ [c_I(x_{k,j})]_+ \end{bmatrix} \right\|_1 - (e_{m_E}^T y_E + e_{m_I}^T y_I). \quad (3.5)$$

The merit parameter is updated to ensure that the search direction $d_{k,j}$ is a descent direction for the merit function, similar to the approaches in [15, 54]. Given user-defined parameters $\epsilon_\sigma, \epsilon_\tau \in (0, 1)$, a trial value $\tau_{k,j}^{trial}$ is computed as,

$$\tau_{k,j}^{trial} = \begin{cases} \infty & \text{if } g_{S_k}(x_{k,j})^T d_{k,j} + d_{k,j}^T H_{k,j} d_{k,j} \leq 0, \\ \frac{(1-\epsilon_\sigma)\Delta_c}{g_{S_k}(x_{k,j})^T d_{k,j} + d_{k,j}^T H_{k,j} d_{k,j}} & \text{otherwise,} \end{cases}$$

and then the merit parameter is updated via

$$\tau_{k,j} = \begin{cases} \tau_{k,j-1} & \text{if } \tau_{k,j-1} \leq \tau_{k,j}^{trial}, \\ \min \left\{ (1-\epsilon_\tau)\tau_{k,j}, \tau_{k,j}^{trial} \right\} & \text{otherwise.} \end{cases} \quad (3.6)$$

This rule ensures that $\{\tau_{k,j}\}$ is a monotonically non-increasing sequence. The merit parameter is reinitialized to a user-defined value at the beginning of each inner loop. Finally, the step size $\alpha_{k,j}$ is selected to satisfy the sufficient decrease condition on the merit function, i.e.,

$$\phi_{S_k}(x_{k,j} + \alpha_{k,j} d_{k,j}, \tau_j) \leq \phi_{S_k}(x_{k,j}, \tau_{k,j}) + \eta \alpha_{k,j} \Delta l_{S_k}(x_{k,j}, d_{k,j}, \tau_{k,j}) \quad (3.7)$$

where $\eta \in (0, 1)$ is a user defined line search parameter.

Within an outer iteration k , under appropriate constraint qualifications, the search direction $d_{k,j}$ goes to zero, and the algorithm either approaches a KKT point for the subsampled problem or terminates at an infeasible stationary point [6, 54]. An infeasible stationary point is infeasible to the original problem, and a stationary point to the constraint minimization problem (l_∞ on the left and l_1 on the right)

$$\min_{x \in \mathbb{R}^n} \left\| \begin{bmatrix} c_E(x) \\ [c_I(x)]_+ \end{bmatrix} \right\|_\infty \quad \text{or,} \quad \min_{x \in \mathbb{R}^n} \left\| \begin{bmatrix} c_E(x) \\ [c_I(x)]_+ \end{bmatrix} \right\|_1. \quad (3.8)$$

If the search direction $p_{k,j}$ is zero while the constraints are violated, an infeasible stationary point has been reached, and both the inner and outer loops can be terminated, since the constraints in the subsampled problem are deterministic. Otherwise, in outer iteration k , the inner loop terminates when the termination criterion \mathcal{T}_k from Framework 1.1 (Line 4) is satisfied, defined as,

$$\mathcal{T}_k : \|d_{k,j}\| \leq \gamma_k \|d_{k,0}\| + \epsilon_k. \quad (3.9)$$

The algorithm is summarized in Algorithm 3.1 and the following remark.

Another possible measure of solution quality for the termination criterion is the merit function model reduction Δl_{S_k} . However, since it includes the positive part of the inequality constraints and the active set may change significantly within the inner loop due to possible drastic variations in the objective function across outer iterations, employing Δl_{S_k} leads to instability and thus it is not used. Moreover, one can use any other deterministic method in the inner loop but then the measure of optimality for termination must be appropriately chosen.

Algorithm 3.1 Inequality Constrained RA-SQP

Inputs: Sequences $\{|S_k|\}$, $\{\gamma_k\} < 1$ and $\{\epsilon_k\} \geq 0$; initial iterate $x_{0,0}$; initial merit parameter sequence $\{\tau_{k,-1}\} = \bar{\tau} > 0$; other parameters $\epsilon_\sigma, \epsilon_\tau, \epsilon_\alpha \in (0, 1)$, $\sigma_p, \sigma_d > 0$.

```

1: for  $k = 0, 1, 2, \dots$  do
2:   Construct subsampled problem (1.6) from sample set  $S_k$ 
3:   for  $j = 0, 1, 2, \dots$  do
4:     Solve (3.1) and obtain  $p_{k,j}$  and  $y$  (or  $y_E$  and  $y_I$ )
5:     if  $\|p_{k,j}\| = 0$  and  $y > 0$  (or either  $\|y_E\| > 0$  or  $\|y_I\| > 0$ )
6:       return  $x_{k,j}$  as an infeasible stationary point
7:     Obtain approximation of Hessian of Lagrangian  $H_{k,j}$ 
8:     Solve (3.2) and obtain  $d_{k,j}$ 
9:     if (3.9) is satisfied;  $N_k = j$ , break
10:    Update merit parameter via (3.6)
11:    Set  $\alpha_{k,j} = 1$ 
12:    while (3.7) is not satisfied do
13:      Update  $\alpha_{k,j} = \epsilon_\alpha \alpha_{k,j}$ 
14:    end while
15:    Set  $x_{k,j+1} = x_{k,j} + \alpha_{k,j} d_{k,j}$ 
16:  end for
17:  Set  $x_{k+1,0} = x_{k,N_k}$ 
18: end for

```

Remark 3.1. We make the following remarks about Algorithm 3.1.

1. **Outer Iteration (Line 2):** A set of samples S_k is obtained to construct the subsampled problem (1.6). The batch size is a prespecified user-defined sequence and an adaptive strategy is described later.
2. **SQP subproblem (Lines 4-8):** The search direction $d_{k,j}$ is computed via a two-step process. First, a search direction $p_{k,j}$ is computed via (3.1) that minimizes the linearized constraint violation. If $\|p_{k,j}\| = 0$ and $x_{k,j}$ is infeasible to the original problem, the algorithm terminates and returns $x_{k,j}$ as an infeasible stationary point. Then, the search direction $d_{k,j}$ is computed via (3.1) that minimizes a quadratic approximation of the objective function while achieving linearized constraint violation at least as good as the linearized constraint violation achieved by $p_{k,j}$. The formulations used to compute $p_{k,j}$ and $d_{k,j}$ depend on the choice of constraint violation measure (l_1 or l_∞ norm).
3. **Termination Test (Line 9):** The inner loop is terminated when the search direction norm $\|d_{k,j}\|$ falls below a certain tolerance (3.9). Thus, one extra SQP subproblem is solved in the

inner loop than the number of iterate updates, as $\|d_{k,j}\|$ is required to evaluate the termination criterion (3.9).

4. **Merit Parameter and Line Search (Lines 10-15):** The mechanism for updating the merit parameter and performing the line search to obtain the step size follows standard procedures from [6, 15, 42, 54]. Since the objective function can change drastically across outer iterations, similar to Algorithm 2.2, the merit parameter is reinitialized at the start of each outer iteration to the user-specified value $\bar{\tau}$ similar to Algorithm 2.2.

3.2 Convergence Analysis

In this section, we establish convergence guarantees for Algorithm 3.1 for problems with general nonlinear constraints (1.1). We first demonstrate the well-posedness of Algorithm 3.1 and establish conditions for convergence, and then present the adaptive strategy for batch size selection and complexity results.

To begin, we borrow a few definitions from Subsection 2.2. We use Assumption 2.4 and error metric G_S from (2.8) to characterize the error in the subsampled gradients, and define $\{\mathcal{F}_k\}$ as the same sequence of σ -algebras with $\mathcal{F}_0 = \{x_{0,0}\}$, i.e., without the dual variable. Next, we introduce the assumptions on (1.1). These assumptions are necessary to establish the well-posedness of Algorithm 3.1 that uses robust-SQP subproblems in the inner loop.

Assumption 3.1. Let $\chi \subset \mathbb{R}^n$ be a closed bounded convex set containing the sequence of iterates generated by any run of Algorithm 3.1. The objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable and its gradient function $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz continuous over χ . The subsampled objective function $F_{S_k} : \mathbb{R}^n \rightarrow \mathbb{R}$ for all $k \geq 0$ is continuously differentiable with Lipschitz continuous gradients over χ . The constraint functions $c_E : \mathbb{R}^n \rightarrow \mathbb{R}^{m_E}$ and $c_I : \mathbb{R}^n \rightarrow \mathbb{R}^{m_I}$ are continuously differentiable. The Jacobian functions, $\nabla c_E^T : \mathbb{R}^n \rightarrow \mathbb{R}^{m_E \times n}$ and $\nabla c_I^T : \mathbb{R}^n \rightarrow \mathbb{R}^{m_I \times n}$ are Lipschitz continuous over χ . The sequence of Hessian approximation matrices $\{H_{k,j}\} \in \mathbb{R}^{n \times n}$ are symmetric positive definite matrices with $\kappa_H \geq \mu_H > 0$, such that, $\mu_H I_n \preceq H_{k,j} \preceq \kappa_H I_n$.

While assuming a closed bounded set of iterates in Assumption 2.8 was a strong condition for problems with equality constraints, similar conditions (Assumption 3.1) are commonly required in both deterministic [6] and stochastic optimization [38, 54] with inequality constraints, thereby justifying their use in this scenario. However, we also need the subsampled problems to be smooth as it is required by the robust SQP method [54, Assumption 2.2] and the RA framework treats each subsampled problem as a deterministic problem, similar to Assumption 2.8. These conditions can be relaxed if permitted by the inner loop deterministic solver employed to establish convergence under conditions similar to Subsection 2.2. The conditions on the approximation of the Hessian of the Lagrangian can also be relaxed to positive definiteness over the null space of the active set. Under Assumption 3.1, there exists $\kappa_g \geq 0$ such that $\|\nabla f(x)\| \leq \kappa_g$ for all $x \in \chi$, and Assumption 2.6 is satisfied.

Assumption 3.2. The extended MFCQ hold at all points $x \in \chi$, i.e.,

1. the Jacobian for the equality constraints is full row rank, i.e., $\text{rank}(\nabla c_E(x)) = m_E$,
2. there exists $u \in \mathbb{R}^n$ such that $\nabla c_E(x)^T u = 0$ and $[\nabla c_I(x)^T]_i u > 0$ for all $i \in \{i : [c_I(x)]_i \geq 0\}$.

The constraint qualification in Assumption 3.2 is an extension of the MFCQ (Mangasarian-Fromovitz constraint qualification), also found in [6, 54], that incorporates infeasible points into the constraint qualification. These conditions are necessary to ensure the boundedness of the merit parameter ($\{\tau_{k,j}\} > 0$), thereby guaranteeing convergence to either a first-order stationary point or an infeasible stationary point. Under the stated assumptions, we now establish the well-posedness of Algorithm 3.1.

Theorem 3.2. *Suppose Assumptions 3.1 and 3.2 hold. Then, for each outer iteration $k \geq 0$ in Algorithm 3.1, $\{d_{k,j}\} \rightarrow 0$ and the inner loop terminates with either an infeasible stationary point or upon satisfying termination criterion (3.9).*

Proof. Based on [6, Theorem 6.1] (as restated in [54, Theorem 2.4]), under the stated assumptions, the robust SQP method as described in Subsection 3.1 in outer iteration k either terminates finitely with $d_{k,N_k} = 0$ or $\|d_{k,j}\| \rightarrow 0$ as $j \rightarrow \infty$, thus satisfying termination criterion (3.9) finitely. If $d_{k,N_k} = 0$ but the constraints are not satisfied, the algorithm terminates returning an infeasible stationary point based on [6, Lemma 2.2] and the deterministic nature of constraints. \square

At $x_{k,j}$, the solution to (3.1) is independent of the gradient approximation employed. Therefore, the feasible region of (3.2) is independent of the sample set given $x_{k,j}$. Let $d_{k,j}^{true}$ be the solution to the subproblem (3.2) where the true gradient of the objective function is used. If $d_{k,j}^{true} = 0$, $x_{k,j}$ is either a first-order stationary point or an infeasible stationary point by [6, Lemma 2.2]. We analyze the convergence of Algorithm 3.1 by analyzing $\|d_{k,j}^{true}\|$, similar to the analysis presented in Subsection 2.2. We first establish a bound on $\|d_{k,N_k}^{true}\|$ similar to Lemma 2.2.

Lemma 3.3. *Suppose Assumptions 3.1 and 3.2 hold. Then for all $k \geq 0$, the norm of the true search direction at outer iterates generated by Algorithm 3.1 satisfy*

$$\begin{aligned} \|d_{k,N_k}^{true}\| &\leq \gamma_k \|d_{k-1,N_{k-1}}^{true}\| + \epsilon_k \\ &\quad + \mu_H^{-1} \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma_k \mu_H^{-1} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|. \end{aligned} \quad (3.10)$$

For the expectation problem (1.3), for all $k \geq 0$

$$\begin{aligned} \mathbb{E}[\|d_{k,N_k}^{true}\| | \mathcal{F}_k] &\leq \gamma_k \|d_{k-1,N_{k-1}}^{true}\| + \mathbb{E}[\epsilon_k | \mathcal{F}_k] \\ &\quad + \mu_H^{-1} \mathbb{E}[\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\ &\quad + \gamma_k \mu_H^{-1} \mathbb{E}[\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k]. \end{aligned} \quad (3.11)$$

Proof. We first provide a bound on the error in the search direction based on the error in the subsampled gradient, following the procedure in [21, Lemma 4.16]. We then follow the procedure from Lemma 2.12 to establish the desired bounds.

At $x_{k,j}$, the feasible region to the subproblems for $d_{k,j}$ and $d_{k,j}^{true}$ are the same. The objective function of (3.2) can be rewritten as $(H_{k,j}^{-1} g_{S_k}(x_{k,j}) + d_{k,j})^T H_{k,j} (H_{k,j}^{-1} g_{S_k}(x_{k,j}) + d_{k,j})$ and as $(H_{k,j}^{-1} \nabla f(x_{k,j}) + d_{k,j}^{true})^T H_{k,j} (H_{k,j}^{-1} \nabla f(x_{k,j}) + d_{k,j}^{true})$ when the true problem gradient is employed. As the feasible region is convex for both l_1 and l_∞ norm based formulations in (3.2), from [4, Proposition 1.1.9],

$$\begin{aligned} (d_{k,j} - d_{k,j}^{true})^T H_{k,j} (H_{k,j}^{-1} \nabla f(x_{k,j}) - d_{k,j}^{true}) &\leq 0 \quad \text{and,} \\ (d_{k,j}^{true} - d_{k,j})^T H_{k,j} (H_{k,j}^{-1} g_{S_k}(x_{k,j}) - d_{k,j}) &\leq 0. \end{aligned}$$

Summing the two inequalities,

$$\begin{aligned}
0 &\leq (d_{k,j}^{true} - d_{k,j})^T H_{k,j} (H_{k,j}^{-1} \nabla f(x_{k,j}) - d_{k,j}^{true} - H_{k,j}^{-1} g_{S_k}(x_{k,j}) + d_{k,j}) \\
&= -(d_{k,j}^{true} - d_{k,j})^T H_{k,j} (d_{k,j}^{true} - d_{k,j}) \\
&\quad + (d_{k,j}^{true} - d_{k,j})^T H_{k,j} H_{k,j}^{-1} (\nabla f(x_{k,j}) - g_{S_k}(x_{k,j})) \\
&\leq -\mu_H \|d_{k,j}^{true} - d_{k,j}\|^2 + (d_{k,j}^{true} - d_{k,j})^T (\nabla f(x_{k,j}) - g_{S_k}(x_{k,j})) \\
&\leq -\mu_H \|d_{k,j}^{true} - d_{k,j}\|^2 + \|d_{k,j}^{true} - d_{k,j}\| \|\nabla f(x_{k,j}) - g_{S_k}(x_{k,j})\|,
\end{aligned}$$

where the second inequality follows from $\mu_H I_n \preceq H_{k,j}$ and the third inequality follows from Cauchy-Schwarz inequality. Using the above inequality, it follows that

$$\|d_{k,j}^{true} - d_{k,j}\| \leq \mu_H^{-1} \|\nabla f(x_{k,j}) - g_{S_k}(x_{k,j})\|.$$

The true search direction norm can be bounded as,

$$\begin{aligned}
\|d_{k,N_k}^{true}\| &\leq \|d_{k,N_k}\| + \|d_{k,N_k}^{true} - d_{k,N_k}\| \\
&\leq \gamma_k \|d_{k,0}\| + \epsilon_k + \|d_{k,N_k}^{true} - d_{k,N_k}\| \\
&\leq \gamma_k \|d_{k,0}^{true}\| + \epsilon_k + \|d_{k,N_k}^{true} - d_{k,N_k}\| + \gamma_k \|d_{k,0}^{true} - d_{k,0}\| \\
&= \gamma_k \|d_{k-1,N_{k-1}}^{true}\| + \epsilon_k + \|d_{k,N_k}^{true} - d_{k,N_k}\| + \gamma_k \|d_{k,0}^{true} - d_{k,0}\|,
\end{aligned}$$

where the second inequality follows from termination criterion (3.9) and the equality follows because $x_{k,0} = x_{k-1,N_{k-1}}$, $H_{k,0} = H_{k-1,N_{k-1}}$ resulting in the same subproblems for $d_{k,0}^{true}$ and $d_{k-1,N_{k-1}}^{true}$. Substituting the derived bound on the search direction error yields (3.10). Taking conditional expectation of (3.10) given \mathcal{F}_k yields (3.11) as $d_{k-1,N_{k-1}}^{true}$ is known under \mathcal{F}_k . \square

Using Lemma 3.3, we now establish the convergence of Algorithm 3.1, following the procedure in Theorem 2.3.

Theorem 3.4. *Suppose Assumptions 2.4, 3.1 and 3.2 hold for any run of Algorithm 3.1.*

1. *For the finite-sum problem (1.2): If the batch size sequence is chosen such that $\{S_k\} \rightarrow \mathcal{S}$ and the termination criterion parameters are chosen such that $\epsilon_k \rightarrow 0$ and $0 \leq \{\gamma_k\} \leq \gamma < 1$, then $\|d_{k,N_k}^{true}\| \rightarrow 0$.*
2. *For the expectation problem (1.3): If the batch size sequence is chosen such that $\{|S_k|\} \rightarrow \infty$ and the termination criterion parameters are chosen such that $\mathbb{E}[\epsilon_k | \mathcal{F}_k] \rightarrow 0$ and $0 \leq \{\gamma_k\} \leq \tilde{\gamma} < 1$, and given $\mathbb{E}[G_{S_k}^2 | \mathcal{F}_k] \rightarrow 0$ as $\{|S_k|\} \rightarrow \infty$, then $\mathbb{E}[\|d_{k,N_k}^{true}\|] \rightarrow 0$.*

Proof. The proof follows the same procedure as Theorem 2.3 and has been included here for completeness. For the finite-sum problem (1.2), unrolling (3.10) from Lemma 3.3 and substituting (2.6)

yields for all $k \geq 0$,

$$\begin{aligned}
\|d_{k,N_k}^{true}\| &\leq \left\{ \prod_{h=0}^k \gamma_h \right\} \|d_{0,0}^{true}\| + \epsilon_k \\
&\quad + 2\mu_H^{-1} \left(\frac{|S|-|S_k|}{|S|} \right) (\omega_1(\|\nabla f(x_{k,N_k})\| + \gamma_k \|\nabla f(x_{k,0})\|) + \omega_2(1 + \gamma_k)) \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \left[\epsilon_i + 2\mu_H^{-1} \left(\frac{|S|-|S_i|}{|S|} \right) (\omega_1 \|\nabla f(x_{i,N_i})\| + \omega_2) \right] \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \left[2\mu_H^{-1} \gamma_i \left(\frac{|S|-|S_i|}{|S|} \right) (\omega_1 \|\nabla f(x_{i,0})\| + \omega_2) \right] \\
&\leq \left\{ \prod_{h=0}^k \gamma_h \right\} \|d_{0,0}^{true}\| + \epsilon_k + 2\mu_H^{-1} \left(\frac{|S|-|S_k|}{|S|} \right) (\omega_1 \kappa_g + \omega_2) (1 + \gamma_k) \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \left[\epsilon_i + 2\mu_H^{-1} \left(\frac{|S|-|S_i|}{|S|} \right) (\omega_1 \kappa_g + \omega_2) (1 + \gamma_i) \right],
\end{aligned}$$

where the second inequality follows from Assumption 3.1. The first term in the above inequality converges to zero as $\left\{ \prod_{h=0}^k \gamma_h \right\} \leq \gamma^{k+1}$ where $\gamma < 1$. The second and third terms converge to zero based on the specifications $\epsilon_k \rightarrow 0$ and $\{|S_k|\} \rightarrow |S|$, respectively. The final term is a series of the form $\sum_{i=0}^{k-1} a_i \left\{ \prod_{h=i+1}^k \gamma_h \right\} \leq \sum_{i=0}^{k-1} a_i \gamma^{k-i}$ which converges to zero by Lemma A.1, with $b_i = a_i = \epsilon_i + 2\mu_H^{-1} \left(\frac{|S|-|S_i|}{|S|} \right) (1 + \gamma_i) (\omega_1 \kappa_g + \omega_2) \rightarrow 0$, thus completing the proof.

For the expectation problem (1.3), the expectation result (3.11) from Lemma 3.3 can be further refined following the procedure from Theorem 2.3. For all $k \geq 0$,

$$\begin{aligned}
\mathbb{E}[\|d_{k,N_k}^{true}\|] &\leq \left\{ \prod_{h=0}^k \gamma_h \right\} \|d_{k,N_k}^{true}\| + \mathbb{E} \left[\mathbb{E}[\epsilon_k | \mathcal{F}_k] + \epsilon_G \mu_H^{-1} \mathbb{E}[G_{S_k} | \mathcal{F}_k] + \kappa_g \mu_H^{-1} \sqrt{\mathbb{E}[G_{S_k}^2 | \mathcal{F}_k]} \right] \\
&\quad + \mathbb{E} \left[\gamma_k \mu_H^{-1} \frac{\tilde{\omega}_1 \kappa_g + \tilde{\omega}_2}{\sqrt{|S_k|}} \right] + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \mathbb{E}[\mathbb{E}[\epsilon_i | \mathcal{F}_i]] \\
&\quad + \sum_{i=0}^{k-1} \left\{ \prod_{h=i+1}^k \gamma_h \right\} \mathbb{E} \left[\epsilon_G \mu_H^{-1} \mathbb{E}[G_{S_i} | \mathcal{F}_i] + \kappa_g \mu_H^{-1} \sqrt{\mathbb{E}[G_{S_i}^2 | \mathcal{F}_i]} + \gamma_i \mu_H^{-1} \frac{\tilde{\omega}_1 \kappa_g + \tilde{\omega}_2}{\sqrt{|S_i|}} \right].
\end{aligned}$$

In the above bound, the first term is the same as it appeared for deterministic errors. All the terms within the second and third term expectation converge to zero based on the specifications $\mathbb{E}[\epsilon_k | \mathcal{F}_k] \rightarrow 0$, $\{|S_k|\} \rightarrow \infty$, and the assumption $\mathbb{E}[G_{S_k}^2 | \mathcal{F}_k] \rightarrow 0$. The final terms together are a series of the form $\sum_{i=0}^{k-1} \mathbb{E}[a_i] \left\{ \prod_{h=i+1}^k \gamma_h \right\} \leq \sum_{i=0}^{k-1} \mathbb{E}[a_i] \tilde{\gamma}^{k-i}$ which converges to zero by Lemma A.1, with $b_i = \mathbb{E}[a_i] \rightarrow 0$ as $a_i \rightarrow 0$, thus completing the proof. \square

Theorem 3.4 establishes the convergence of the norm of the true search direction to zero. As established in [6, 54], if the $\|d_{k,N_k}^{true}\| \rightarrow 0$, the limit point is a stationary point for the constraint violation minimization problem (3.8). If the limit point is feasible, it is a first-order stationary

point for (1.1). While the convergence of the true search direction norm holds deterministically for the finite-sum problem (1.2), it holds in expectation for the expectation problem (1.3). Thus, for the expectation problem (1.3), if there exists a subsequence of $\{d_{k,N_k}\}$ that converges to zero, the limit point is either an infeasible stationary point or a first-order stationary point. However, this scenario is not guaranteed from the convergence in expectation of the true search direction norm, similar to the result in [21]. Convergence can be established under stricter assumptions on the gradient errors, similar to the conditions in [54].

We now introduce the adaptive strategy for sample set size selection in Algorithm 3.1 as a modification of the well-known norm test, similar to Condition 2.1.

Condition 3.1. *At the beginning of outer iteration $k \geq 0$ in Algorithm 3.1, the batch size $|S_k|$ is selected such that:*

1. *For the finite-sum problem (1.2): With constants $\theta, \hat{\theta}, a \geq 0$ and $\beta \in (0, 1)$,*

$$\begin{aligned} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 &\leq \theta^2 \|d_{k,0}^{true}\|^2 + a^2 \beta^{2k}, \\ \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\|^2 &\leq \hat{\theta}^2 (\theta^2 \|d_{k,0}^{true}\|^2 + a^2 \beta^{2k}). \end{aligned}$$

2. *For the expectation problem (1.3): With constants $\tilde{\theta}, \tilde{a} \geq 0$ and $\tilde{\beta} \in (0, 1)$,*

$$\mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\|^2 | \mathcal{F}_k] \leq \tilde{\theta}^2 \|d_{k,0}^{true}\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}.$$

Condition 3.1 is analogous to Condition 2.1, and a discussion of its implications can be found in Remark 2.4. The practical aspects of its implementation are discussed later in Section 4. We now establish the complexity of Algorithm 3.1 under Condition 3.1 across outer iterations, assuming that Assumption 2.5 holds for the expectation problem (1.3), in a manner similar to Theorem 2.5.

Theorem 3.5. *Suppose Assumptions 2.4, 3.1 and 3.2 hold, and the batch size sequence $\{|S_k|\}$ is chosen to satisfy Condition 3.1 in Algorithm 3.1.*

1. *For the finite-sum problem (1.2): If termination criterion parameters are chosen as $0 \leq \{\gamma_k\} \leq \gamma < 1$, $\epsilon_k = \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k$ with $\omega, \hat{\omega} \geq 0$ and Condition 3.1 parameters are chosen such that $a_1 = \left[\gamma + \theta(\omega + \mu_H^{-1}(\hat{\theta} + \gamma)) \right] < 1$, the true search direction converges to zero at a linear rate across outer iterations as,*

$$\|d_{k,N_k}^{true}\| \leq \max\{a_1 + \nu, \beta\}^{k+1} \max\{\|d_{0,0}^{true}\|, \frac{a_2}{\nu}\},$$

where $a_2 = a(\omega + \mu_H^{-1}(\hat{\theta} + \gamma)) + \hat{\omega}$ and $\nu > 0$ such that $a_1 + \nu < 1$.

2. *For the expectation problem (1.3): If Assumption 2.5 is satisfied and the termination criterion parameters are chosen such that $0 \leq \{\gamma_k\} \leq \tilde{\gamma} < 1$, $\epsilon_k = \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}) | \mathcal{F}_k)}{|S_k|}}$ where $\tilde{\omega} \geq 0$ and Condition 3.1 parameters are chosen such that $\tilde{a}_1 = \left[\tilde{\gamma} + \tilde{\theta} \left(\tilde{\omega} + \mu_H^{-1} \left(\frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} + \tilde{\gamma} \right) \right) \right] < 1$, the true search direction converges to zero in expectation at a linear rate across outer iterations as,*

$$\mathbb{E} [\|d_{k,N_k}^{true}\|] \leq \max\{\tilde{a}_1 + \tilde{\nu}, \tilde{\beta}\}^{k+1} \max\{\|d_{0,0}^{true}\|, \frac{\tilde{a}_2}{\tilde{\nu}}\},$$

where $\tilde{a}_2 = \tilde{a} \left[\tilde{\omega} + \mu_H^{-1} \left(\frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} + \tilde{\gamma} \right) \right]$ and $\tilde{\nu} > 0$ such that $\tilde{a}_1 + \tilde{\nu} < 1$.

Proof. For the finite-sum problem (1.2), the error bound (3.10) can be further refined as,

$$\begin{aligned}
\|d_{k,N_k}^{true}\| &\leq \gamma \|d_{k-1,N_{k-1}}^{true}\| + \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k \\
&\quad + \mu_H^{-1} \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma \mu_H^{-1} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| \\
&\leq \gamma \|d_{k-1,N_{k-1}}^{true}\| + \omega \left(\theta \|d_{k-1,N_{k-1}}^{true}\| + a \beta^k \right) + \hat{\omega} \beta^k \\
&\quad + \mu_H^{-1} \hat{\theta} (\theta \|d_{k-1,N_{k-1}}^{true}\| + a \beta^k) + \mu_H^{-1} \gamma \left(\theta \|d_{k-1,N_{k-1}}^{true}\| + a \beta^k \right) \\
&= a_1 \|d_{k-1,N_{k-1}}^{true}\| + a_2 \beta^k,
\end{aligned}$$

where the second inequality follows from Condition 3.1 and the final equality follows from the defined constants. Using the above bound, applying Lemma A.2 with $Z_k = \|d_{k,N_k}^{true}\|$, $\rho_1 = a_1$, $\rho_2 = \beta$ and $b = a_2$ completes the proof.

For the expectation problem (1.3), the expectation error bound (3.11) can be further refined as,

$$\begin{aligned}
\mathbb{E} [\|d_{k,N_k}^{true}\| | \mathcal{F}_k] &\leq \tilde{\gamma} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}) | \mathcal{F}_k)}{|S_k|}} \\
&\quad + \mu_H^{-1} \mathbb{E} [\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\
&\quad + \tilde{\gamma} \mu_H^{-1} \mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\
&\leq \tilde{\gamma} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}) | \mathcal{F}_k)}{|S_k|}} \\
&\quad + \mu_H^{-1} \kappa_G \frac{(\epsilon_G + \kappa_g)}{\sqrt{|S_k|}} + \tilde{\gamma} \mu_H^{-1} \mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\
&\leq \tilde{\gamma} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{\omega} \left[\tilde{\theta} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a} \tilde{\beta}^k \right] \\
&\quad + \mu_H^{-1} \frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} \left[\tilde{\theta} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a} \tilde{\beta}^k \right] + \tilde{\gamma} \mu_H^{-1} \left[\tilde{\theta} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a} \tilde{\beta}^k \right] \\
&= \tilde{a}_1 \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a}_2 \tilde{\beta}^k,
\end{aligned}$$

where the second inequality follows from (2.13) and Assumption 2.5, the third inequality follows by substituting $\kappa_G \leq \frac{\sqrt{\text{Var}(\nabla F(x_{k,0}) | \mathcal{F}_k)}}{\kappa_\sigma}$ and Condition 3.1 and the last equality follows from the defined constants. Taking the total expectation of the bound yields,

$$\mathbb{E} [\|d_{k,N_k}^{true}\|] \leq \tilde{a}_1 \mathbb{E} [\|d_{k-1,N_{k-1}}^{true}\|] + \tilde{a}_2 \tilde{\beta}^k.$$

Applying Lemma A.2 with $Z_k = \mathbb{E} [\|d_{k,N_k}^{true}\|]$, $\rho_1 = \tilde{a}_1$, $\rho_2 = \tilde{\beta}$ and $b = \tilde{a}_2$ completes the proof. \square

Theorem 3.5 provides results analogous to those in Theorem 2.13, and the parameter trade-offs discussed in Theorem 2.5 remain applicable. Unlike Theorem 2.5, while Theorem 3.5 establishes the convergence of the norm of the true search direction at a linear rate, this does not necessarily result in linear convergence of the error in KKT conditions (1.4). Although $d_{k,N_k}^{true} = 0$ guarantees that a KKT point or an infeasible stationary point has been reached, we are not aware of any result that bounds the KKT error in terms of the true search direction norm, as is possible in the setting with only equality constraints.

We now present a corollary for the simpler case of geometrically increasing batch sizes, similar to Corollary 2.6. This results from setting θ (or $\tilde{\theta}$) to zero in Condition 3.1, removing all the parameter selection conditions in Theorem 3.5 and the variance lower bound in Assumption 2.5.

Corollary 3.6. *Suppose the conditions of Theorem 3.5 hold.*

1. *For the finite-sum problem (1.2): If the batch size is selected as $|S_k| = \lceil (1 - \beta^k)|S| \rceil$ with $\beta \in (0, 1)$ and termination criterion parameters are chosen as $0 \leq \{\gamma_k\} \leq \gamma < 1$, $\epsilon_k = \omega \left(1 - \frac{|S_k|}{|S|}\right)$ with $\omega \geq 0$, then the true search direction norm converges to zero at a linear rate across outer iterations, as expressed in Theorem 3.5 with $a_1 = \gamma$ and $a_2 = \omega + 2\mu_H^{-1}(1 + \gamma)(\omega_1\kappa_g + \omega_2)$.*
2. *For the expectation problem (1.3): If Assumption 2.5 is satisfied, the sample set size is chosen as $|S_{k+1}| = \left\lceil \frac{|S_k|}{\tilde{\beta}^2} \right\rceil$ with $\tilde{\beta} \in (0, 1)$ and the termination criterion parameters are chosen such that $0 \leq \{\gamma_k\} \leq \tilde{\gamma} < 1$, $\epsilon_k = \frac{\tilde{\omega}}{\sqrt{|S_k|}}$ where $\tilde{\omega} \geq 0$, then the true search direction norm converges to zero in expectation at a linear rate across outer iterations, as expressed in Theorem 3.5 with $\tilde{a}_1 = \tilde{\gamma}$ and $\tilde{a}_2 = \frac{\tilde{\omega} + \kappa_G \mu_H^{-1}(\epsilon_G + \kappa_g) + \tilde{\gamma} \mu_H^{-1}(\tilde{\omega}_1 \kappa_g + \tilde{\omega}_2)}{\sqrt{|S_0|}}$, where $|S_0|$ is the initial batch size.*

Proof. The proof follows the same procedure as Corollary 2.6 with adjustments to account for changes from Lemma 2.2 to Lemma 3.3 by incorporating the scaling factor μ_H^{-1} , as done in Theorem 3.5. The complete proof is presented in Appendix B. \square

Corollary 3.6 presents a geometrically increasing batch size scheme that satisfies Condition 3.1. One can also select batch sizes sufficiently large with θ (or $\hat{\theta}$) non-zero to ensure Condition 3.1 is satisfied, similar to Lemma 2.7.

Lemma 3.7. *Suppose Assumptions 2.4, 3.1 and 3.2 hold in Algorithm 3.1.*

1. *For the finite-sum problem (1.2): For $k \geq 0$, Condition 3.1 is satisfied if*

$$|S_k| \geq |S| \left(1 - \sqrt{\frac{\theta^2 \|d_{k,0}^{true}\|^2 + a^2 \beta^{2k}}{4(\omega_1^2 \kappa_g^2 + \omega_2^2)}} \right) \quad \text{with } \hat{\theta} = 1.$$

2. *For the expectation problem (1.3): For $k \geq 0$, Condition 3.1 is satisfied if*

$$|S_k| \geq \frac{\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2}{\tilde{\theta}^2 \|d_{k,0}^{true}\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}}.$$

Proof. The proof follows the same procedure as Lemma 2.7 with suitable adjustments for Condition 3.1. The complete proof can be found in Appendix B. \square

Lemma 3.7, while somewhat pessimistic, demonstrates that there exists a sequence of batch sizes $|S_k|$ that satisfies Condition 3.1. Although our analysis focuses on Framework 1.1 with a specific deterministic solver for inequality constraints, it can be readily extended to other deterministic solvers used within the inner loop, provided an appropriate termination criterion is defined. However, in contrast to the case with only equality constraints, the empirical performance in the presence of inequality constraints is significantly more sensitive to these choices.

4 Numerical Experiments

In this section, we present numerical results that illustrate the empirical performance of the proposed algorithms. All methods were implemented by the authors in Python¹. We begin by introducing the test problems, followed by a discussion of key implementation considerations for the adaptive sampling strategy. We report numerical results, first for problems with only equality constraints (Subsection 4.2), and then for problems involving general nonlinear constraints (Subsection 4.3). Implementation details for all methods are discussed in the respective subsections.

We evaluate the algorithms on two types of problems: regularized multi-class logistic regression tasks, and problems from the S2MPJ CUTEst problem set [28]. For the classification tasks, the objective function is defined using a logistic regression model with a cross-entropy loss over a set of classes \mathcal{K} , and is given by

$$f(x) = \frac{1}{|\mathcal{S}|} \sum_{(y,t) \in \mathcal{S}} \sum_{i \in \mathcal{K}} t^i \log \left(\frac{1}{1 + \exp(-y^T x^i)} \right), \quad (4.1)$$

where \mathcal{S} denotes the set of samples, and each sample $(y, t) \in \mathcal{S}$ consists of a feature vector $y \in \mathbb{R}^{n_f}$ (where n_f is the number of features including the bias term) and a label vector $t \in \{0, 1\}^{|\mathcal{K}|}$, with $t^i = 1$ if the sample belongs to class $i \in \mathcal{K}$, and $t^i = 0$ otherwise. Each $x^i \in \mathbb{R}^{n_f}$ represents the classifier corresponding to class $i \in \mathcal{K}$, and the variable $x \in \mathbb{R}^n$ is the concatenation of classifiers for all classes with $n = n_f |\mathcal{K}|$. We report results on the covtype dataset ($n_f = 55$, $|\mathcal{K}| = 7$, $|\mathcal{S}| = 581,012$) and the mnist dataset ($n_f = 781$, $|\mathcal{K}| = 10$, $|\mathcal{S}| = 60,000$) [11]. Constraints for these problems are described later.

The objective function for problems from the S2MPJ CUTEst set [28] are augmented as,

$$F(x, \xi) = f(x) + \xi \|x - x_{init} - e_n\|^2, \quad (4.2)$$

where $f(x)$ is the deterministic objective function of the true problem, x_{init} is the starting point and ξ is a uniformly distributed random variable in $[-0.1, 0.1]$ independent of x , ensuring $\mathbb{E}[F(x, \xi)] = f(x)$. This augmentation differs from those considered in [2, 3, 21, 45], as the noise in the objective function now depends on the current iterate and increases as the algorithm moves farther from the initial point. This results in an adversarial setting that demands more accurate function estimates to achieve good empirical performance. Additionally, the offset in $\|x - x_{init} - e_n\|^2$ prevents the algorithms from receiving perfect information at the initial point. The constraints remain the deterministic constraints specified in the problem set.

4.1 Adaptive Sampling Implementation

This section outlines the practical implementation of our adaptive sampling strategy for Algorithms 2.1, 2.2 and 3.1. In practice, the finite-sum problem (1.2) is treated as an expectation problem (1.3) within the adaptive sampling procedure. The adaptive sampling conditions at the start of outer iteration $k \geq 0$ can be expressed in general as $\mathbb{E}[\|g_{S_k}(x_{k,0}) - \nabla f(x_{k,0})\|^2] \leq \tilde{\theta}^2 Z_k^2 + \tilde{a}^2 \tilde{\beta}^{2k}$, where Z_k represents the true problem quantity in the adaptive sampling condition, determined

¹Our code will be made publicly available upon publication of the manuscript within the GitHub repository: <https://github.com/SANDOPT/Gradient-Tracking-Algorithmic-Framework.git>. All our experiments were conducted on a PowerEdge R760 server equipped with two Intel Xeon Gold 16-core processors and 256 GB of RDIMM memory

by the termination criterion \mathcal{T}_k used in the inner loop. Specifically, for Condition 2.1, $Z_k = \|T(x_{k,0}, \lambda_{k-1, N_{k-1}})\|$; for Condition 2.2, $Z_k = \|d_{k,0}^{true}\|$; for Condition 2.3, $Z_k^2 = \Delta l(x_{k,0}, \tau_{k,0}^{true}, d_{k,0}^{true})$; and for Condition 3.1, $Z_k = \|d_{k,0}^{true}\|$. The sampling condition can be rearranged as

$$\frac{\text{Var}(\nabla F(x_{k,0}, \xi)|x_{k,0})}{|S_k|} \leq \tilde{\theta}^2 Z_k^2 + \tilde{a}^2 \tilde{\beta}^{2k} \Rightarrow |S_k| \geq \frac{\text{Var}(\nabla F(x_{k,0}, \xi)|x_{k,0})}{\tilde{\theta}^2 Z_k^2 + \tilde{a}^2 \tilde{\beta}^{2k}},$$

where sample estimates are used to approximate the terms in the condition, similar to [1, 5, 7, 43].

Unlike earlier works, we cannot use estimates from the previous outer iteration to approximate the condition. This is because multiple inner iterations in outer iteration $k-1$ introduce a strong bias in the new starting point $x_{k,0}$ toward the sample set S_{k-1} . To mitigate this, we instead draw a new sample set \tilde{S}_{k-1} , independent of $x_{k,0}$, with $|\tilde{S}_{k-1}| = |S_{k-1}|$. We then use the sample variance over \tilde{S}_{k-1} and \tilde{Z}_k as an estimate for Z_k to approximate the condition as

$$|S_k| \geq \frac{\text{Var}_{\xi \in \tilde{S}_{k-1}}(\nabla F(x_{k,0}, \xi)|x_{k,0})}{\tilde{\theta}^2 \tilde{Z}_k^2 + \tilde{a}^2 \tilde{\beta}^{2k}}.$$

For Condition 2.1, $\tilde{Z}_k = \|T_{\tilde{S}_{k-1}}(x_{k,0}, \lambda_{k-1, N_{k-1}})\|$ can be readily computed as an estimate of Z_k . For Conditions 2.2, 2.3, and 3.1, we form the subsampled problem $F_{\tilde{S}_{k-1}}$ and perform a single inner loop iteration initialized at $x_{k,0}$ to compute the corresponding value of \tilde{Z}_k . No updates are made during this iteration. This adds the computational cost of one inner iteration, but incurs no additional gradient cost, as the computed gradients can be reused in the first inner iteration of outer iteration k . To enable this, the sample set S_k is constructed by augmenting \tilde{S}_{k-1} with additional samples, as needed, ensuring that $\tilde{S}_{k-1} \subseteq S_k$.

The geometric component in the adaptive sampling condition is intended to prevent large, erratic increases in batch sizes that may result from noisy sample estimates. In practice, we replace this component with a hard geometric increase limit, restricting batch size growth to at most a factor of $\tilde{\beta} > 1$ between successive outer iterations. Additionally, batch sizes are enforced to be non-decreasing and are bounded above by the total number of samples in the finite-sum problem (1.2). Therefore, the batch size at outer iteration k is selected as

$$|S_k| = \min \left\{ |S|, \tilde{\beta} |S_{k-1}|, \max \left\{ |S_{k-1}|, \left\lceil \frac{\text{Var}_{\xi \in \tilde{S}_{k-1}}(\nabla F(x_{k,0}, \xi)|x_{k,0})}{\tilde{\theta}^2 \tilde{Z}_k^2} \right\rceil \right\} \right\}.$$

We set $\tilde{\theta} = 0.5$ and $\tilde{\beta} = 5$ for the proposed algorithms and the batch size is initialized at 32. While prior works such as [1, 5, 43] often use larger values of $\tilde{\theta}$, we prefer a smaller value to allow for more significant adjustments to the batch size, given that multiple updates are performed between successive evaluations of the sampling condition.

4.2 Equality Constrained Problems

We first consider problems with only general nonlinear equality constraints. We evaluate the performance of five variants of Algorithm 2.2, each referred to as “RA-SQP” with additional descriptors, differing across three key characteristics. First, the termination criterion \mathcal{T}_k is either based on the norm of the search direction (2.20), labeled “ $\|d\|$ ” (with $\gamma_k = \gamma = 0.5$), or on the decrease in the merit function model (2.21), labeled “ Δl ” (with $\gamma_k = \gamma = 0.1$). Second, the Hessian of the Lagrangian is approximated either by the identity matrix (default) or using an L-BFGS update (with memory size as $\min\{n, 10\}$), labeled “L-BFGS”. Third, the linear system in the SQP subproblem

(2.14) is solved via MINRES either exactly (default option to a relative error tolerance of 10^{-6}) or approximately, using inexact solutions that satisfy (2.18) or (2.19) with $\kappa_T = 10^{-1}$, $\epsilon_{feas} = 10^{-4}$ and $\epsilon_{opt} = 10^{-4}$, as discussed in Subsection 2.3, and labeled “Inexact”. For the termination criterion based on $\|d\|$ (2.20), we only report results with the identity Hessian and exact linear system solutions, as the results for other configurations are not competitive. To avoid solving subsampled subproblems to unnecessary accuracy, the additional tolerance in the termination criterion is set to $\epsilon_k = 10^{-6}$. Furthermore, to prevent excessive computation on subproblems, we limit the number of inner iterations per outer iteration to a maximum of 500 for reasons discussed in Subsection 2.4.

We implement and compare against the stochastic SQP method [2], labeled “S-SQP”, and the adaptive sampling SQP method [1], labeled “AS-SQP”. We also include comparisons with the deterministic SQP method [2, 42], labeled “SQP”, applied to the finite-sum logistic regression problems. In all three baselines, “SQP”, “S-SQP”, and “AS-SQP”, we use the identity matrix as the Lagrangian Hessian approximation and solve the SQP linear system exactly. For “S-SQP”, we use a batch size of 1024, which is larger than commonly used, to ensure the method reaches solutions with accuracy comparable to the proposed RA-based algorithms. For “AS-SQP”, the batch size is initialized at 32 and increases according to the adaptive sampling conditions of the algorithm.

We employ the MINRES algorithm to solve the SQP linear system using its scipy implementation [60]. To evaluate algorithm performance, we track two metrics: the constraint violation $\|c(x)\|_\infty$ and the norm of the Lagrangian gradient $\|\nabla_x \mathcal{L}(x, \lambda^*)\|_\infty$, where $\lambda^* \in \mathbb{R}^m$ is the dual variable that minimizes the KKT error for the true problem at x . These metrics are plotted against two cost measures: the number of gradient evaluations and the total number of MINRES iterations used to solve the linear systems. For the RA-based algorithms, the MINRES count also includes the cost of solving the subproblem used to estimate quantities in the sampling condition.

4.2.1 Logistic Regression

We consider the logistic regression multi-class classification problem (4.1) with equality regularization constraints, i.e., $\|x^i\|^2 = 1$ for all $i \in \mathcal{K}$. In Figures 1 and 2, we present comparisons against existing methods on the covtype and mnist datasets, respectively. For a fair comparison, all algorithms use the identity matrix as the approximation for the Lagrangian Hessian. When comparing “RA-SQP $\|d\|$ ” and “RA-SQP Δl ” with algorithms from the literature, all of which use exact solutions to the SQP linear system, our algorithms achieve better performance in terms of the number of gradient evaluations and are comparable to deterministic SQP in terms of MINRES iterations. Furthermore, the flexibility of our framework is evident from the improved performance of “RA-SQP Δl Inexact”, which uses inexact solutions to the SQP linear system. This variant significantly reduces the number of MINRES iterations compared to “RA-SQP Δl ”, demonstrating how performance can be enhanced by modifying the underlying deterministic SQP method.

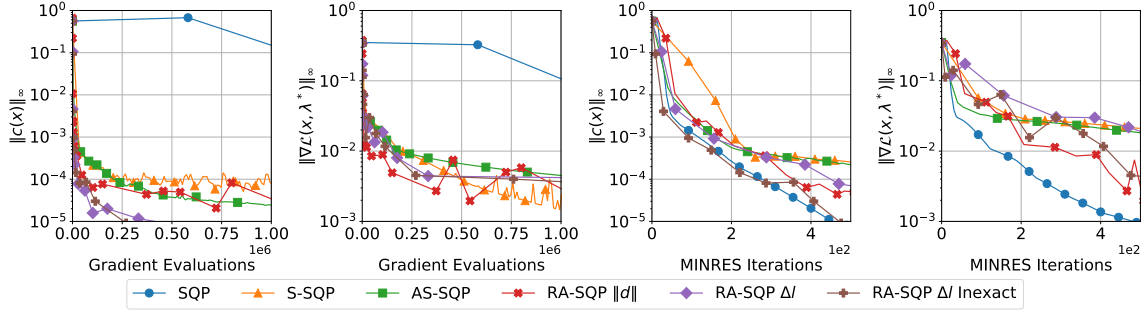


Figure 1: Constraint violation ($\|c(x)\|_\infty$) and Lagrangian gradient norm ($\|\nabla_x \mathcal{L}(x, \lambda^*)\|_\infty$) with optimized dual variable λ^* , with respect to number of gradient evaluations and number of MINRES iterations for stochastic SQP (“S-SQP” [2]), adaptive sampling SQP (“AS-SQP” [1]), deterministic SQP (“SQP” [2, 42]) and our proposed algorithms “RA-SQP $\|d\|$ ”, “RA-SQP $\Delta/$ ”, and “RA-SQP $\Delta/$ Inexact” over the multi-class logistic regression problem (4.1) with equality regularization constraints for the covtype dataset ($n_f = 55$, $|\mathcal{K}| = 7$, $|\mathcal{S}| = 581,012$, [11]).




Figure 2: Constraint violation ($\|c(x)\|_\infty$) and Lagrangian gradient norm ($\|\nabla_x \mathcal{L}(x, \lambda^*)\|_\infty$) with optimized dual variable λ^* , with respect to number of gradient evaluations and number of MINRES iterations for stochastic SQP (“S-SQP” [2]), adaptive sampling SQP (“AS-SQP” [1]), deterministic SQP (“SQP” [2, 42]) and our proposed algorithms “RA-SQP $\|d\|$ ”, “RA-SQP $\Delta/$ ”, and “RA-SQP $\Delta/$ Inexact” over the multi-class logistic regression problem (4.1) with equality regularization constraints for the mnist dataset ($n_f = 781$, $|\mathcal{K}| = 10$, $|\mathcal{S}| = 60,000$, [11]).

In Figure 3, we present results for our proposed algorithms on the mnist dataset, highlighting the flexibility of the RA framework and the performance gains achievable through enhancements such as the L-BFGS Hessian approximation and inexact solutions to the SQP linear system. Incorporating the L-BFGS approximation improves performance in terms of gradient evaluations compared to the identity Hessian counterparts. However, adopting the L-BFGS approximation results in harder to solve SQP linear systems for the higher quality search directions as depicted with the increase in MINRES iterations when calculating exact solutions. Adopting inexact solutions to the linear system with the L-BFGS approximation reduces this cost. These results demonstrate how techniques from deterministic SQP methods can be seamlessly integrated into the RA framework to enhance its effectiveness on stochastic optimization problems.




Figure 3: Constraint violation ($\|c(x)\|_\infty$) and Lagrangian gradient norm ($\|\nabla_x \mathcal{L}(x, \lambda^*)\|_\infty$) with optimized dual variable λ^* , with respect to number of gradient evaluations and number of MINRES iterations for “RA-SQP Δl ” with and without L-BFGS approximations and exact and inexact SQP linear system solutions over the multi-class logistic regression problem (4.1) with equality regularization constraints for the mnist dataset ($n_f = 781$, $|\mathcal{K}| = 10$, $|\mathcal{S}| = 60,000$, [11]).

In Figure 4, we present how the number of inner iterations (N_k) and the batch size ($|S_k|$) change across outer iterations and with respect to the number of gradient evaluations for the results presented in Figure 3. First, the number of inner iterations increases during the initial phase for all methods and subsequently stabilizes, never reaching the upper limit of 500 inner iterations. This behavior is desirable, as it aligns with the growth of the batch size. Second, while the rate of increase for the batch sizes is different for each method across outer iterations depending on the progress made, it is relatively consistent when measured against the number of gradient evaluations. This adaptive behavior is one of the main reasons for the good empirical performance of the framework with respect to the number of gradient evaluations. Third, the methods that use an L-BFGS approximation perform fewer inner iterations, as expected. Finally, when the L-BFGS approximation is used, the batch size increases slower than when the identity matrix is used across outer iterations when exact solutions to the SQP linear system are employed. Conversely, the increase is faster when inexact solutions are employed. As discussed in Subsection 2.2, “RA-SQP Δl L-BFGS” frequently solves the subsampled problems to a higher accuracy than required by the inner loop termination test \mathcal{T}_k , introducing additional bias in the solutions. In such cases, a slower increase in the batch size works better. Alternatively, a faster batch size increase can be employed by decreasing $\tilde{\theta}$ to take full advantage of this behavior. We note again that we do not tune the value of $\tilde{\theta}$ across methods to display the robustness of the framework.




Figure 4: Number of Inner Iterations (N_k) and Batch Size ($|S_k|$) with respect to outer iterations and number of gradient evaluations for “RA-SQP Δl ” with and without L-BFGS approximations and exact and inexact SQP linear system solutions over the multi-class logistic regression problem (4.1) with equality regularization constraints for the mnist dataset ($n_f = 781$, $|\mathcal{K}| = 10$, $|\mathcal{S}| = 60,000$, [11]).

4.2.2 CUTest Problems

In this section, we evaluate the empirical performance of Algorithm 2.2 on the CUTest problem set, using the augmented formulation in (4.2), and compare our results with methods from the literature. We focus on problems that have only equality constraints, satisfy the constraint qualifications (verified via solutions obtained using deterministic SQP), possess a non-constant objective, and have total dimension $n + m < 1000$. This filtering yields a benchmark suite of 88 problems. For each problem, we perform 10 runs with different random seeds, resulting in a total of 880 problem instances per algorithm. Each run is allotted a gradient evaluation budget of 10^6 .

In Figure 5, we present the results using Dolan-Moré performance profiles [24] with respect to the number of gradient evaluations and MINRES iterations, evaluating both feasibility and stationarity errors. The performance profiles are constructed as follows. For each cost metric, an algorithm is considered to have successfully solved a problem instance in a given seed run at iterate x_{out} , starting from initial point x_{init} , if it satisfies a prescribed tolerance $\epsilon_{tol} \in (0, 1)$. Specifically, the algorithm is deemed to have solved a given problem with respect to feasibility at x_{out} if $\|c(x_{out})\|_\infty \leq \epsilon_{tol} \max\{1, \|c(x_{init})\|_\infty\}$. The problem is further considered solved with respect to stationarity at x_{out} if it satisfies the feasibility criterion and, in addition, $\|\nabla_x \mathcal{L}(x_{out}, \lambda_{out}^*)\|_\infty \leq \epsilon_{tol} \max\{1, \|\nabla_x \mathcal{L}(x_{init}, \lambda_{init}^*)\|_\infty\}$, where λ_{out}^* and λ_{init}^* denote the dual variables that minimize the true problem’s KKT error at x_{out} and x_{init} respectively. Overall, across all tolerance levels, the RA-based algorithms demonstrate strong initial progress and exhibit greater robustness. The results also highlight the performance gains from incorporating the L-BFGS Hessian approximation, particularly at tighter tolerance levels where higher solution accuracy is required.

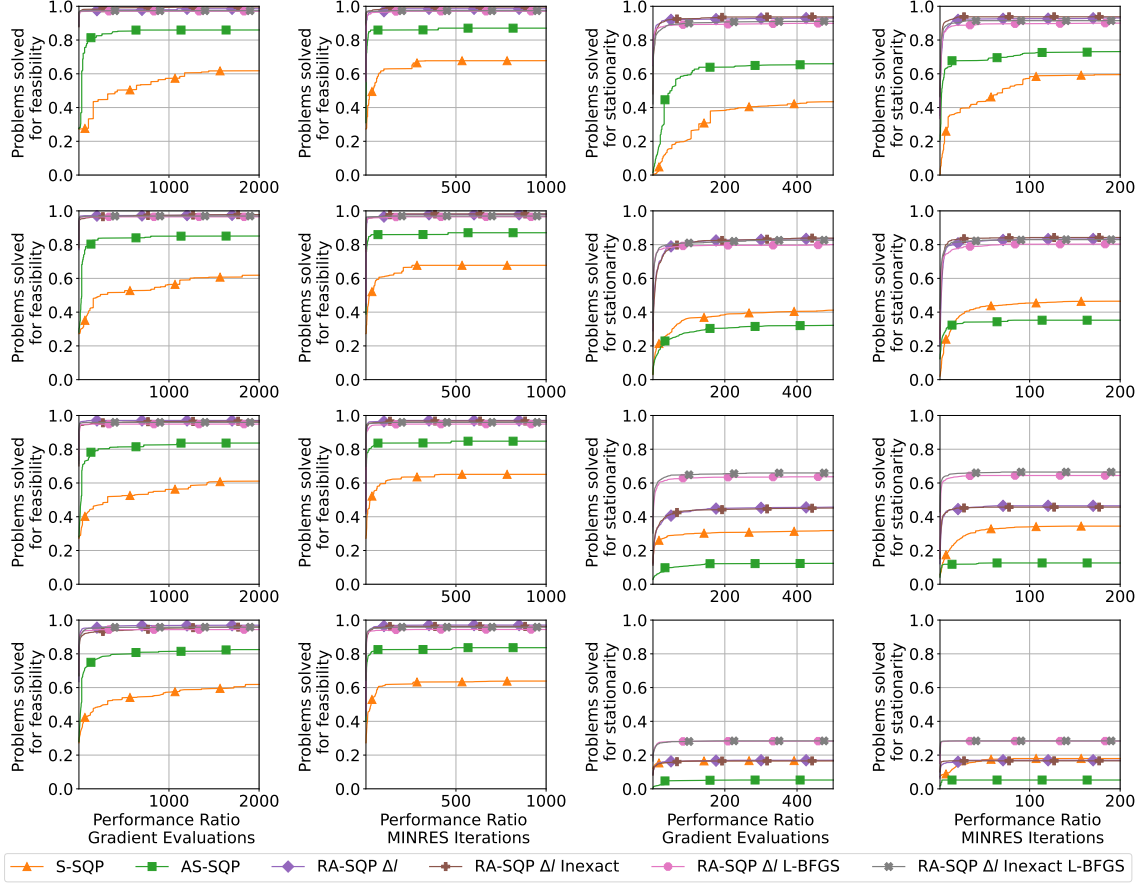


Figure 5: Performance profiles for feasibility and stationarity errors with respect to number of gradient evaluations and number of MINRES iterations for stochastic SQP (“S-SQP” [2]), adaptive sampling SQP (“AS-SQP” [1]) and our proposed algorithm “RA-SQP Δl ” with and without L-BFGS Hessian approximations and exact and inexact SQP linear system solutions over the CUTEst problem set [28] to accuracy levels $\epsilon_{tol} \in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$ that are decreasing going down the rows.

4.3 Inequality Constrained Problems

In this section, we illustrate the empirical performance of Algorithm 3.1 for solving problems with general nonlinear constraints. We consider regularized logistic regression multi-class classification problems as well as the CUTEst problem set, and compare our methods with algorithms from the literature. We also empirically evaluate the stability of the active set in the RA framework across outer iterations.

We implement and compare our algorithm with stochastic SQP [21] (labeled “S-SQP”) and stochastic robust SQP [54] (labeled “Robust-S-SQP”). We also include comparisons with deterministic robust SQP methods [6, 54], using the l_∞ norm (labeled “Robust-SQP- l_∞ ”) and the l_1 norm

(labeled “Robust-SQP- l_1 ”) for the finite-sum logistic regression problems. For fair comparison, all baseline algorithms use the identity matrix as the approximation to the Lagrangian Hessian. The batch size for both “S-SQP” and “Robust-SQP” is set to 1024, larger than typically used, to allow these methods to attain solution accuracies comparable to those achieved by the proposed RA-based algorithms.

We present two versions of Algorithm 3.1, each differing in the robust SQP method employed within the inner loop. The first uses the l_∞ norm and is labeled “RA-SQP- l_∞ ”, while the second employs the l_1 norm and is labeled “RA-SQP- l_1 ”. Both versions use the termination criterion (3.9), based on the search direction norm, with $\gamma_k = \gamma = 0.5$, and use the identity matrix to approximate the Lagrangian Hessian. Following the setup from the equality-constrained case, to prevent excessive computation on subsampled problems, we set $\epsilon_k = 10^{-6}$ and cap the number of inner iterations per outer iteration at 500 for reasons discussed in Subsection 3.2. The bound on the search direction norm in the SQP subproblems is set as $\sigma_d = 2\sigma_p$, where σ_p is proportional to the constraint violation, similar to the strategy in [54]. Specifically, for the l_∞ norm, we set $\sigma_p = 10 \left\| \begin{smallmatrix} c_E(x) \\ [c_I(x)]_+ \end{smallmatrix} \right\|_\infty$, while ensuring $\sigma_p \in [10^2, 10^4]$. For the l_1 norm, we set $\sigma_p = 10 \left\| \begin{smallmatrix} c_E(x) \\ [c_I(x)]_+ \end{smallmatrix} \right\|_1$, constrained to the range $[n \cdot 10^2, n \cdot 10^4]$ scaling the range by the dimension of the problem to account for the scaling difference in l_1 and l_∞ norms.

We employ GUROBI [29] to solve the linear and quadratic programs that arise in the SQP subproblems. The progress of the algorithms is evaluated using two primary metrics: the constraint violation, defined as $\|c(x)\|_\infty = \left\| \begin{smallmatrix} c_E(x) \\ [c_I(x)]_+ \end{smallmatrix} \right\|_\infty$, and the KKT residual, which is computed at the point x as the objective value of the linear program

$$\begin{aligned} KKT(x) = \min_{t \in \mathbb{R}, \lambda_E \in \mathbb{R}^{m_E}, \lambda_I \in \mathbb{R}^{m_I}} t \\ \text{s.t.} \quad & \|\nabla f(x) + \nabla c_E(x)\lambda_E + \nabla c_I(x)\lambda_I\|_\infty \leq t, \\ & \|\lambda_I \odot c_I(x)\|_\infty \leq t, \quad \lambda_I \geq 0, \end{aligned} \tag{4.3}$$

that measures violation of the general KKT conditions (1.4) using the best dual variables. These metrics are evaluated with respect to two computational costs: the number of gradient evaluations and the number of Barrier method iterations performed by GUROBI in solving the associated linear and quadratic programs. For the RA-based algorithms, the Barrier iterations include those required to solve subproblems for estimating quantities in the sampling condition.

4.3.1 Logistic Regression

We evaluate the empirical performance of Algorithm 3.1 on the logistic regression multi-class classification problem (4.1) with inequality regularization constraints, i.e., $\|x^i\|^2 \leq 1$ for all $i \in \mathcal{K}$. In Figures 6 and 7, we compare our algorithm with methods from the literature on the covtype and MNIST datasets, respectively. When comparing “RA-SQP- l_∞ ” and “RA-SQP- l_1 ” with existing methods, our algorithms outperform the others in terms of the number of gradient evaluations and achieve performance close to deterministic SQP methods with respect to Barrier iterations. Furthermore, we observe a difference in Barrier iteration performance between “RA-SQP- l_1 ” and “RA-SQP- l_∞ ”, with the l_∞ variant performing fewer Barrier iterations. This is likely due to the smaller size of the SQP subproblems when using the l_∞ norm, which reduces solver overhead without affecting performance with respect to gradient evaluations.

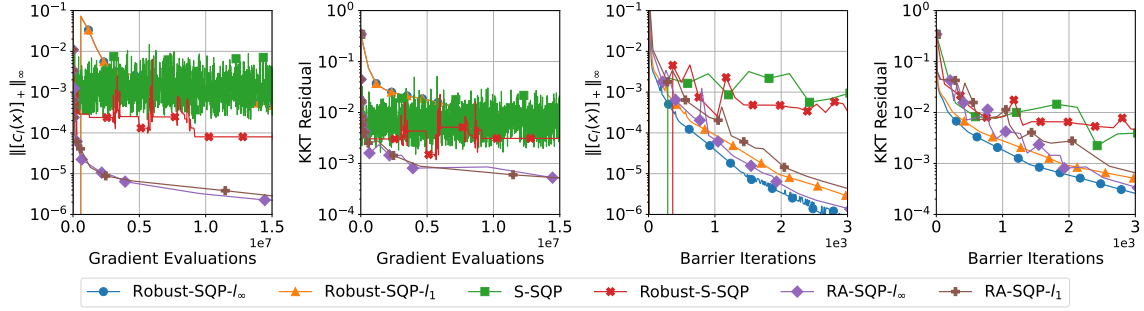


Figure 6: Constraint violation ($\|c_I(x)_+\|_\infty$) and KKT residual (4.3) with respect to number of gradient evaluations and number of Barrier iterations for stochastic SQP (“S-SQP” [21]), robust stochastic SQP (“Robust-S-SQP” [54]), deterministic robust SQP (“Robust-SQP- l_∞ ” and “Robust-SQP- l_1 ” [6, 54]) and our proposed algorithms “RA-SQP- l_∞ ” and “RA-SQP- l_1 ” over the multi-class logistic regression problem (4.1) with inequality regularization constraints for the covetype dataset ($n_f = 55$, $|\mathcal{K}| = 7$, $|\mathcal{S}| = 581,012$, [11]).

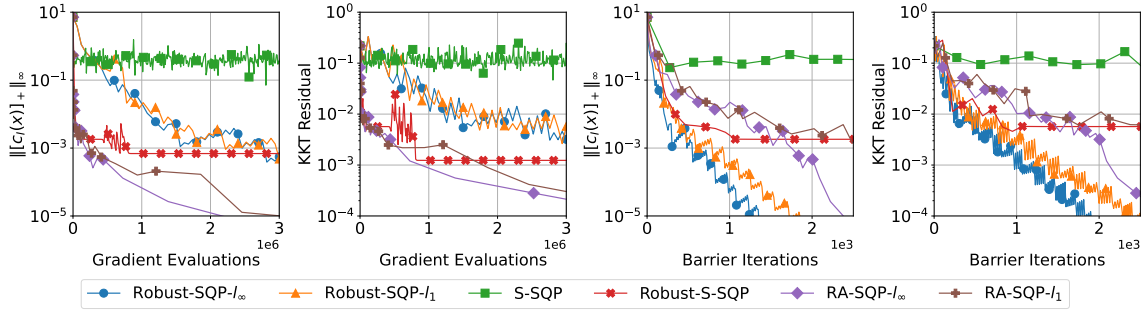


Figure 7: Constraint violation ($\|c_I(x)_+\|_\infty$) and KKT residual (4.3) with respect to number of gradient evaluations and number of Barrier iterations for stochastic SQP (“S-SQP” [21]), robust stochastic SQP (“Robust-S-SQP” [54]), deterministic robust SQP (“Robust-SQP- l_∞ ” and “Robust-SQP- l_1 ” [6, 54]) and our proposed algorithms “RA-SQP- l_∞ ” and “RA-SQP- l_1 ” over the multi-class logistic regression problem (4.1) with inequality regularization constraints for the mnist dataset ($n_f = 781$, $|\mathcal{K}| = 10$, $|\mathcal{S}| = 60,000$, [11]).

4.3.2 CUTEst Problems

In this section, we evaluate the empirical performance of Algorithm 3.1 on the CUTEst problem set in the form (4.2), and compare our algorithms with methods from the literature. We consider problems that contain general inequality constraints (excluding problems with only bound constraints), satisfy constraint qualifications (verified using a deterministic SQP solver), and have a non-constant objective function. Additionally, we restrict the problem size to $n + m_E + m_I < 2000$, resulting in a final set of 248 problems. For each problem, we perform 10 independent runs with different random seeds, totaling 2480 runs per algorithm. A gradient evaluation budget of 10^6 is allocated for each problem instance and seed run.

In Figure 8, we present results using Dolan-Moré performance profiles [24], evaluating both feasibility and stationarity errors with respect to the number of gradient evaluations and Barrier iterations. The profiles are constructed similarly to those in Fig. 5, as follows. For each cost metric, an algorithm is considered to have successfully solved a problem instance in a given seed run at iterate x_{out} , starting from initial iterate x_{init} , if it satisfies a prescribed tolerance $\epsilon_{tol} \in (0, 1)$. Specifically, the algorithm is deemed to have solved a problem with respect to feasibility if $\|c(x_{out})\|_{\infty} \leq \epsilon_{tol} \max\{1, \|c(x_{init})\|_{\infty}\}$. The algorithm is further considered to have solved the problem with respect to stationarity if it satisfies the feasibility criterion and $KKT(x_{out}) \leq \epsilon_{tol} \max\{1, KKT(x_{init})\}$, where the KKT residual is defined in (4.3).

Overall, across all tolerance levels, the RA-based algorithms demonstrate good performance in the early stages, are robust and achieve high accuracy. While we observe that “S-SQP” slightly outperforms the RA-based methods at the highest accuracy level with respect to stationarity, the RA-based algorithms remain efficient in reducing constraint violation.

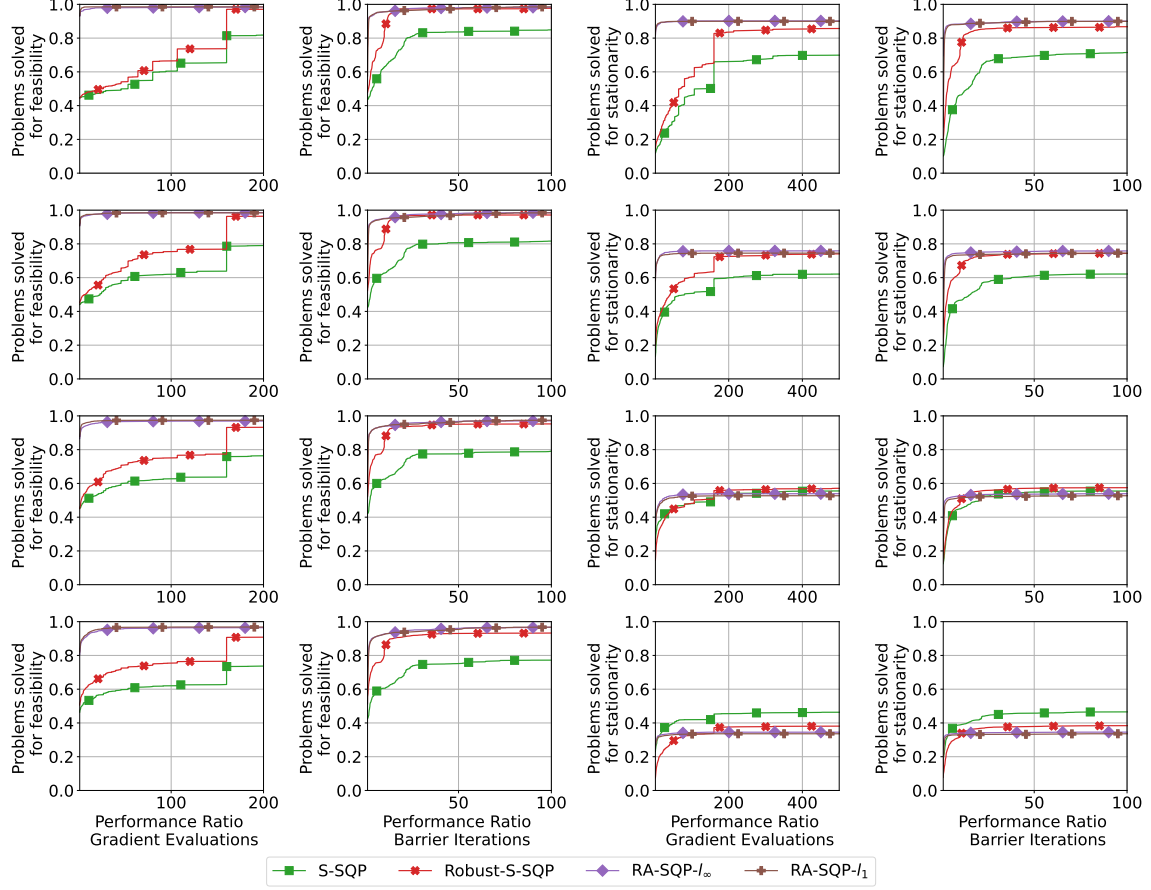


Figure 8: Performance profiles for feasibility and stationarity errors with respect to number of gradient evaluations and number of Barrier iterations for stochastic SQP (“S-SQP” [21]), robust stochastic SQP (“Robust-S-SQP” [54]) and our proposed algorithms “RA-SQP- l_∞ ” and “RA-SQP- l_1 ” over the CUTEst problem set [28] to accuracy levels $\epsilon_{tol} \in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$ decreasing going down the rows.

4.3.3 Active Set Stability

In this section, we empirically investigate the stability of the active set for problems with general nonlinear constraints across the outer iterations of Algorithm 3.1. Since the constraints are deterministic while the objective function varies between outer iterations due to subsampling, it is desirable for the algorithm to identify and maintain a stable active set after a few outer iterations, particularly once the batch size in the subsampled problem becomes sufficiently large. To study this behavior, we select problems from the CUTEst problem set [28] that feature a large number of inequality constraints (m_I), where the number of active inequality constraints at the optimal solution (m_I^*) constitutes only a small subset.

To measure this, we first define the active set of inequality constraints at a point x as $A(x) =$

$\{i : [c_I(x)]_i \geq -10^{-6}\}$, where the threshold 10^{-6} corresponds to the feasibility tolerance used by GUROBI when solving the SQP subproblems. Let A^* denote the active set at the optimal solution obtained using a deterministic solver. We employ two metrics to assess active set stability: “optimal set similarity”, defined as the Jaccard similarity between the current active set and the optimal active set, $\frac{|A(x) \cap A^*|}{|A(x) \cup A^*|}$, and the constraint violation. These metrics are evaluated over outer iterations for Algorithm 3.1 with the l_1 norm, and with respect to the number of gradient evaluations to facilitate comparison with “S-SQP” and “Robust-S-SQP” methods. While we have used a large batch size of 1024 for both “S-SQP” and “Robust-S-SQP” in our previous experiments to enable comparable accuracy, we also include results for both with a smaller batch size of 128 in this analysis, labeled “S-SQP (small)” and “Robust-S-SQP (small)”, respectively.

The results are summarized in Figure 9. Achieving the optimal active set is only meaningful if the corresponding solution is feasible. For the BATCH problem, both batch sizes of “S-SQP” fail to produce feasible solutions. While “RA-SQP- l_1 ”, “Robust-S-SQP” and “Robust-S-SQP (small)” converge to the optimal active set, the feasibility error remains too large to claim successful recovery. For the GPP problem, all algorithms eventually achieve feasibility. However, both “S-SQP” variants exhibit significant volatility in feasibility error and active set identification, especially the small batch size variant, and fail to stabilize on a consistent active set. Similar instability is observed in the behavior of “Robust-S-SQP”, with “Robust-S-SQP (small)” failing to identify the optimal active set. In contrast, “RA-SQP- l_1 ”, while slower, eventually recovers a feasible solution along with the optimal active set and maintains it consistently. For the OET1 problem, all algorithms rapidly achieve and maintain feasibility. However, with the exception of “RA-SQP- l_1 ”, all other methods exhibit erratic changes in their active sets across iterations or fail to identify the optimal active set. While this is an empirical discussion, the observations motivate further theoretical investigation into active set stability within the RA-based framework, potentially paving the way for an active set method tailored to inequality constrained problems.

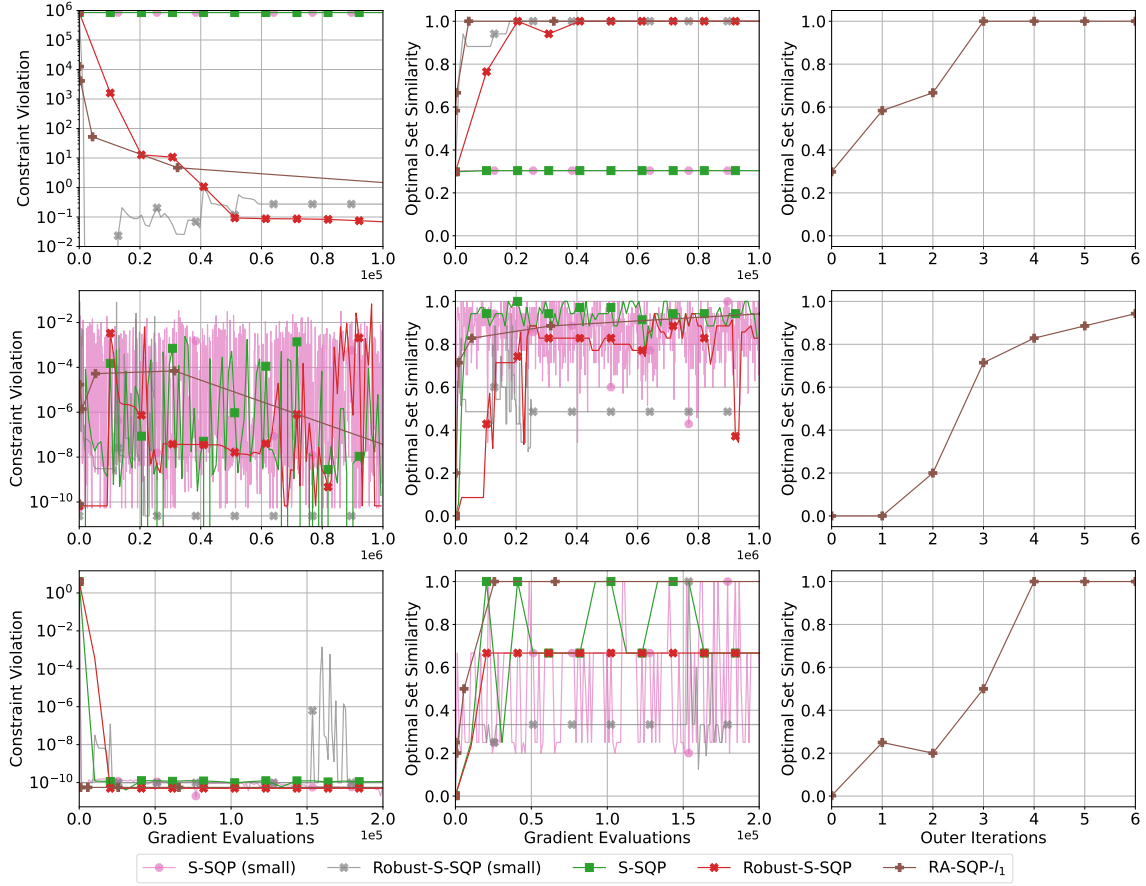


Figure 9: Constraint violation and optimal set similarity with respect to number of outer iterations and gradient evaluations for “Robust-S-SQP (small)”, “S-SQP (small)”, “S-SQP”, “Robust-S-SQP” and “RA-SQP- l_1 ” across selected problems from the CUTEst problem set [28]. Each row corresponds to a different problem, listed from top to bottom as follows: BATCH($m_I = 61$, $m_I^* = 17$), GPP($m_I = 198$, $m_I^* = 35$) and OET1($m_I = 1002$, $m_I^* = 3$).

5 Final Remarks

In this paper, we have proposed an algorithmic framework based on the Retrospective Approximation (RA) paradigm to solve stochastic optimization problems with deterministic constraints using deterministic optimization methods. The use of deterministic optimization methods has addressed challenges such as small step sizes and slow convergence that arise in recent stochastic approaches. They also have enabled the use of techniques from deterministic optimization, such as line search and quasi-Newton Hessian approximations, to improve empirical performance, techniques that are not readily applicable to the stochastic setting. We have presented a framework designed for problems with general nonlinear equality constraints that can leverage any deterministic solver for a stochastic problem. We have followed that with an instance of the framework that employs a de-

terministic line search Sequential Quadratic Programming (SQP) method. This instance is endowed with optimal complexity with respect to both gradient evaluations and SQP linear system solves for stochastic problems with equality constraints, and has allowed for the effective use of advanced techniques such as quasi-Newton Hessian approximations and inexact search direction computations, as demonstrated through our numerical experiments. For problems with general nonlinear equality and inequality constraints, we have proposed an RA-based algorithm that uses the deterministic robust-SQP method as the deterministic solver. The method ensures the consistency of the SQP subproblem in the presence of inequality constraints while identifying infeasible stationary points and is efficient as demonstrated in our numerical experiments. We present empirical results over multiple problems and compare with algorithms from literature to showcase the competitive performance of the proposed framework.

References

- [1] 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.
- [2] 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.
- [3] Albert S Berahas, Miaolan Xie, and Baoyu Zhou. A sequential quadratic programming method with high-probability complexity bounds for nonlinear equality-constrained stochastic optimization. *SIAM Journal on Optimization*, 35(1):240–269, 2025.
- [4] Dimitri Bertsekas. *Convex optimization theory*, volume 1. Athena Scientific, Belmont, Massachusetts, 2009.
- [5] Raghu Bollapragada, Richard Byrd, and Jorge Nocedal. Adaptive sampling strategies for stochastic optimization. *SIAM Journal on Optimization*, 28(4):3312–3343, 2018.
- [6] James V Burke and Shih-Ping Han. A robust sequential quadratic programming method. *Mathematical Programming*, 43(1):277–303, 1989.
- [7] Richard H Byrd, Gillian M Chin, Jorge Nocedal, and Yuchen Wu. Sample size selection in optimization methods for machine learning. *Mathematical programming*, 134(1):127–155, 2012.
- [8] 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.
- [9] Richard H Byrd, Frank E Curtis, and Jorge Nocedal. An inexact newton method for nonconvex equality constrained optimization. *Mathematical programming*, 122(2):273–299, 2010.
- [10] Richard G Carter. On the global convergence of trust region algorithms using inexact gradient information. *SIAM Journal on Numerical Analysis*, 28(1):251–265, 1991.
- [11] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.

- [12] Changan Chen, Frederick Tung, Naveen Vedula, and Greg Mori. Constraint-aware deep neural network compression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 400–415, 2018.
- [13] Huifen Chen and Bruce W Schmeiser. Retrospective approximation algorithms for stochastic root finding. In *Proceedings of Winter Simulation Conference*, pages 255–261. IEEE, 1994.
- [14] Huifen Chen and Bruce W Schmeiser. Stochastic root finding via retrospective approximation. *IEEE Transactions*, 33(3):259–275, 2001.
- [15] Frank E Curtis, Shima Dezfulian, and Andreas Wächter. An interior-point algorithm for continuous nonlinearly constrained optimization with noisy function and derivative evaluations. *arXiv preprint arXiv:2502.11302*, 2025.
- [16] Frank E Curtis, Xin Jiang, and Qi Wang. Single-loop deterministic and stochastic interior-point algorithms for nonlinearly constrained optimization. *arXiv preprint arXiv:2408.16186*, 2024.
- [17] Frank E Curtis, Vyacheslav Kungurtsev, Daniel P Robinson, and Qi Wang. A stochastic-gradient-based interior-point algorithm for solving smooth bound-constrained optimization problems. *arXiv preprint arXiv:2304.14907*, 2023.
- [18] Frank E Curtis, Jorge Nocedal, and Andreas Wächter. A matrix-free algorithm for equality constrained optimization problems with rank-deficient jacobians. *SIAM Journal on Optimization*, 20(3):1224–1249, 2010.
- [19] 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*, 205(1):431–483, 2024.
- [20] 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.
- [21] Frank E Curtis, Daniel P Robinson, and Baoyu Zhou. Sequential quadratic optimization for stochastic optimization with deterministic nonlinear inequality and equality constraints. *SIAM Journal on Optimization*, 34(4):3592–3622, 2024.
- [22] Geng Deng and Michael C Ferris. Variable-number sample-path optimization. *Mathematical Programming*, 117(1-2):81–109, 2009.
- [23] Shima Dezfulian and Andreas Wächter. On the convergence of interior-point methods for bound-constrained nonlinear optimization problems with noise. *arXiv preprint arXiv:2405.11400*, 2024.
- [24] Elizabeth D Dolan and Jorge J Moré. Benchmarking optimization software with performance profiles. *Mathematical programming*, 91:201–213, 2002.
- [25] Yuchen Fang, Sen Na, Michael W Mahoney, and Mladen Kolar. Fully stochastic trust-region sequential quadratic programming for equality-constrained optimization problems. *SIAM Journal on Optimization*, 34(2):2007–2037, 2024.

- [26] Michael P Friedlander and Mark Schmidt. Hybrid deterministic-stochastic methods for data fitting. *SIAM Journal on Scientific Computing*, 34(3):A1380–A1405, 2012.
- [27] Saeed Ghadimi, Guanghui Lan, and Hongchao Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2):267–305, 2016.
- [28] Serge Gratton and Philippe L Toint. S2mpj and cutest optimization problems for matlab, python and julia. *arXiv preprint arXiv:2407.07812*, 2024.
- [29] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024.
- [30] Shih-Ping Han. A globally convergent method for nonlinear programming. *Journal of optimization theory and applications*, 22(3):297–309, 1977.
- [31] Afrooz Jalilzadeh and Uday V Shanbhag. eg-vssa: An extragradient variable sample-size stochastic approximation scheme: Error analysis and complexity trade-offs. In *2016 Winter Simulation Conference (WSC)*, pages 690–701. IEEE, 2016.
- [32] James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, and Bryan Wilder. End-to-end constrained optimization learning: A survey. *arXiv preprint arXiv:2103.16378*, 2021.
- [33] Guanghui Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1-2):365–397, 2012.
- [34] Guanghui Lan. *First-order and stochastic optimization methods for machine learning*, volume 1. Springer, Atlanta, USA, 2020.
- [35] Tai Le Quy, Arjun Roy, Vasileios Iosifidis, Wenbin Zhang, and Eirini Ntoutsis. A survey on datasets for fairness-aware machine learning. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(3):e1452, 2022.
- [36] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35, 2021.
- [37] Sen Na, Mihai Anitescu, and Mladen Kolar. An adaptive stochastic sequential quadratic programming with differentiable exact augmented lagrangians. *Mathematical Programming*, 199(1-2):721–791, 2023.
- [38] Sen Na, Mihai Anitescu, and Mladen Kolar. Inequality constrained stochastic nonlinear optimization via active-set sequential quadratic programming. *Mathematical Programming*, 202(1):279–353, 2023.
- [39] Yatin Nandwani, Abhishek Pathak, and Parag Singla. A primal dual formulation for deep learning with constraints. *Advances in Neural Information Processing Systems*, 32, 2019.
- [40] David Newton. *Retrospective approximation for smooth stochastic optimization*. PhD thesis, Purdue University, 2023.

- [41] David Newton, Raghu Bollapragada, Raghu Pasupathy, and Nung Kwan Yip. A retrospective approximation approach for smooth stochastic optimization. *Mathematics of Operations Research*, 2024.
- [42] J Nocedal and SJ Wright. Numerical optimization 2nd edition springer. *New York*, 2006.
- [43] Thomas O’Leary-Roseberry and Raghu Bollapragada. Fast unconstrained optimization via hessian averaging and adaptive gradient sampling methods. *arXiv preprint arXiv:2408.07268*, 2024.
- [44] EO Omojokun. *Trust region algorithm for optimization with equalities and inequalities constraints*. PhD thesis, Ph. D Thesis, University of Colorado at Boulder, 1989.
- [45] Michael J O’Neill. A two stepsize sqp method for nonlinear equality constrained stochastic optimization. *arXiv preprint arXiv:2408.16656*, 2024.
- [46] Raghu Pasupathy. On choosing parameters in retrospective-approximation algorithms for stochastic root finding and simulation optimization. *Operations Research*, 58(4-part-1):889–901, 2010.
- [47] Raghu Pasupathy. An introspective on the retrospective-approximation paradigm. In *Proceedings of the 2011 Winter Simulation Conference (WSC)*, pages 412–421. IEEE, 2011.
- [48] Raghu Pasupathy and Yongjia Song. Adaptive sequential sample average approximation for solving two-stage stochastic linear programs. *SIAM Journal on Optimization*, 31(1):1017–1048, 2021.
- [49] Andre F Perold. Large-scale portfolio optimization. *Management science*, 30(10):1143–1160, 1984.
- [50] Chris Phelps, Johannes O Royset, and Qi Gong. Optimal control of uncertain systems using sample average approximations. *SIAM Journal on Control and Optimization*, 54(1):1–29, 2016.
- [51] E Polak and JO Royset. Efficient sample sizes in stochastic nonlinear programming. *Journal of Computational and Applied Mathematics*, 217(2):301–310, 2008.
- [52] 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.
- [53] Michael JD Powell and YJMP Yuan. A recursive quadratic programming algorithm that uses differentiable exact penalty functions. *Mathematical programming*, 35:265–278, 1986.
- [54] Songqiang Qiu and Vyacheslav Kungurtsev. A sequential quadratic programming method for optimization with stochastic objective functions, deterministic inequality constraints and robust subproblems. *arXiv preprint arXiv:2302.07947*, 2023.
- [55] Sathya N Ravi, Tuan Dinh, Vishnu Suresh Lokhande, and Vikas Singh. Explicitly imposing constraints in deep networks via conditional gradients gives improved generalization and faster convergence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4772–4779, 2019.

- [56] 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.
- [57] Johannes O Royset, Armen Der Kiureghian, and Elijah Polak. Optimal design with probabilistic objective and constraints. *Journal of Engineering Mechanics*, 132(1):107–118, 2006.
- [58] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, Philadelphia, PA, 2021.
- [59] Tyler Summers, Joseph Warrington, Manfred Morari, and John Lygeros. Stochastic optimal power flow based on conditional value at risk and distributional robustness. *International Journal of Electrical Power & Energy Systems*, 72:116–125, 2015.
- [60] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [61] David A Vondrak and R Kevin Wood. *Adaptive Selections of Sample Size and Solver Iterations in Stochastic Optimization with Application to Nonlinear Commodity Flow Problems*. PhD thesis, Naval Postgraduate School, 2009.
- [62] Xiao Wang, Shiqian Ma, and Ya-xiang Yuan. Penalty methods with stochastic approximation for stochastic nonlinear programming. *Mathematics of computation*, 86(306):1793–1820, 2017.
- [63] Wen Li Zhao, Pierre Gentine, Markus Reichstein, Yao Zhang, Sha Zhou, Yeqiang Wen, Changjie Lin, Xi Li, and Guo Yu Qiu. Physics-constrained machine learning of evapotranspiration. *Geophysical Research Letters*, 46(24):14496–14507, 2019.

A Technical Results

We present some technical results that have been used throughout the paper.

Lemma A.1. *Given a non-negative sequence $\{b_k\} \rightarrow 0$ and a constant $\gamma \in [0, 1)$,*

$$\lim_{K \rightarrow \infty} \sum_{k=0}^K b_k \gamma^{K-k} = 0.$$

Proof. Consider a finite index \bar{K} and a corresponding value $\bar{b} \geq 0$ such that $a_k \leq \bar{b}$ for all $k \geq \bar{K}$, which exist as $\{b_k\}$ is a non-negative sequence that converges to zero. For $K > \bar{K}$, the sum can be bounded as

$$\sum_{k=0}^K b_k \gamma^{K-k} = \sum_{k=0}^{\bar{K}} b_k \gamma^{K-k} + \sum_{k=\bar{K}+1}^K b_k \gamma^{K-k} \leq b_0 \sum_{k=0}^{\bar{K}} \gamma^{K-k} + \bar{b} \sum_{k=\bar{K}+1}^K \gamma^{K-k} \leq b_0 \frac{\gamma^{K-\bar{K}}}{1-\gamma} + \frac{\bar{b}\gamma}{1-\gamma}.$$

Given $\epsilon > 0$, if one selects \bar{K} large enough such that $\bar{b} \leq \frac{\epsilon(1-\gamma)}{2\gamma}$ and K large enough such that $b_0 \frac{\gamma^{K-\bar{K}}}{1-\gamma} \leq \frac{\epsilon}{2}$, then $\sum_{k=0}^K b_k \gamma^{K-k} \leq \epsilon$. As one can always select K and \bar{K} large enough to make the non-negative sum arbitrarily small, it converges to zero. \square

Lemma A.2. *Given a non-negative sequence $\{Z_k\}$ such that $Z_{k+1} \leq \rho_1 Z_k + b\rho_2^k$ where $\rho_1, \rho_2 \in [0, 1)$ and $0 \leq b < \infty$, the sequence $\{Z_k\} \rightarrow 0$ at a linear rate as,*

$$Z_k \leq \max\{\rho_1 + \nu, \rho_2\}^{k+1} \max\left\{Z_0, \frac{b}{\nu}\right\},$$

where $\nu > 0$ such that $\rho_1 + \nu < 1$.

Proof. The proof follows by induction. For Z_0 , the result trivially holds. Then, if the result holds for Z_k for $k \geq 0$,

$$\begin{aligned} Z_{k+1} &\leq \rho_1 Z_k + b\rho_2^k \leq \rho_1 \max\{\rho_1 + \nu, \rho_2\}^k \max\left\{Z_0, \frac{b}{\nu}\right\} + b\rho_2^k \\ &\leq \max\{\rho_1 + \nu, \rho_2\}^k \max\left\{Z_0, \frac{b}{\nu}\right\} (\rho_1 + \nu) \\ &\leq \max\{\rho_1 + \nu, \rho_2\}^{k+1} \max\left\{Z_0, \frac{b}{\nu}\right\}, \end{aligned}$$

thus completing the proof. \square

B Additional Proofs

In this section, we present certain proofs that have omitted from the paper for brevity.

Theorem 2.13. *Suppose Assumptions 2.4 and 2.8 hold and that the batch size sequence $\{|S_k|\}$ is chosen to satisfy Condition 2.2 in Algorithm 2.2 with termination criterion (2.20).*

1. *For the finite-sum problem (1.2): For all $k \geq 0$, if termination criterion parameters are chosen as $0 \leq \gamma_k \leq \gamma < \frac{1-\eta_d}{1+\eta_d}$, $\epsilon_k = \omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k$ with $\omega, \hat{\omega} \geq 0$ and Condition 2.2 parameters are chosen such that $a_1 = \frac{1}{1-\eta_d} \left[\gamma(1+\eta_d) + \theta \left(\omega + (\gamma + \hat{\theta}) \frac{1+\eta_d}{\kappa_{HJ}} \right) \right] < 1$, then, the true search direction norm converges at a linear rate across outer iterations, i.e.,*

$$\|d_{k,N_k}^{true}\| \leq \max\{a_1 + \nu, \beta\}^{k+1} \max\{\|d_{0,0}^{true}\|, \frac{a_2}{\nu}\},$$

where $a_2 = \frac{a}{1-\eta_d} \left(\omega + (\gamma + \hat{\theta}) \frac{1+\eta_d}{\kappa_{HJ}} \right) + \frac{\hat{\omega}}{1-\eta_d}$ and $\nu > 0$ such that $a_1 + \nu < 1$.

2. *For the expectation problem (1.3): For all $k \geq 0$, if Assumption 2.5 holds, the termination criterion parameters are chosen as $0 \leq \gamma_k \leq \tilde{\gamma} < \frac{1-\eta_d}{1+\eta_d}$, $\epsilon_k = \tilde{\omega} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0})|\mathcal{F}_k)}{|S_k|}}$ where $\tilde{\omega} \geq 0$ and Condition 2.2 parameters are chosen such that $\tilde{a}_1 = \frac{1}{1-\eta_d} \left[\tilde{\gamma}(1+\eta_d) + \tilde{\theta} \left(\tilde{\omega} + \frac{1+\eta_d}{\kappa_{HJ}} \left(\frac{\epsilon_G + \kappa_g}{\kappa_\sigma} + \tilde{\gamma} \right) \right) \right] < 1$, then, the expected true search direction norm converges at a linear rate across outer iterations, i.e.,*

$$\mathbb{E}[\|d_{k,N_k}^{true}\|] \leq \max\{\tilde{a}_1 + \tilde{\nu}, \tilde{\beta}\}^{k+1} \max\{\|d_{0,0}^{true}\|, \frac{\tilde{a}_2}{\tilde{\nu}}\},$$

where $\tilde{a}_2 = \frac{\tilde{a}}{1-\eta_d} \left[\tilde{\omega} + \frac{1+\eta_d}{\kappa_{HJ}} \left(\frac{\epsilon_G + \kappa_g}{\kappa_\sigma} + \tilde{\gamma} \right) \right]$ and $\tilde{\nu} > 0$ such that $\tilde{a}_1 + \tilde{\nu} < 1$.

Proof. The proof follows a similar procedure to Theorem 2.5. For the finite-sum problem (1.2), the result from Lemma 2.12 can be further refined as,

$$\begin{aligned}
\|d_{k,N_k}^{true}\| &\leq \gamma \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{1}{1-\eta_d} (\omega \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| + \hat{\omega} \beta^k) \\
&\quad + \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| \\
&\leq \gamma \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{1}{1-\eta_d} \left(\omega \left(\theta \|d_{k-1,N_{k-1}}^{true}\| + a \beta^k \right) + \hat{\omega} \beta^k \right) \\
&\quad + \kappa_{HJ}^{-1} \hat{\theta} \frac{1+\eta_d}{1-\eta_d} \left(\theta \|d_{k-1,N_{k-1}}^{true}\| + a \beta^k \right) + \gamma \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \left(\theta \|d_{k-1,N_{k-1}}^{true}\| + a \beta^k \right) \\
&= a_1 \|d_{k-1,N_{k-1}}^{true}\| + a_2 \beta^k,
\end{aligned}$$

where the second inequality follows from Condition 2.2 and the final equality follows from the defined constants. Using the above bound, applying Lemma A.2 with $Z_k = \|d_{k,N_k}^{true}\|$, $\rho_1 = a_1$, $\rho_2 = \beta$ and $b = a_2$ completes the proof.

For the expectation problem (1.3), the expectation result from (3.11) can be further refined as,

$$\begin{aligned}
\mathbb{E}[\|d_{k,N_k}^{true}\| | \mathcal{F}_k] &\leq \tilde{\gamma} \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\tilde{\omega}}{1-\eta_d} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}) | \mathcal{F}_k)}{|S_k|}} \\
&\quad + \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \mathbb{E}[\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\
&\quad + \tilde{\gamma} \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \mathbb{E}[\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\
&\leq \tilde{\gamma} \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\tilde{\omega}}{1-\eta_d} \sqrt{\frac{\text{Var}(\nabla F(x_{k,0}) | \mathcal{F}_k)}{|S_k|}} \\
&\quad + \kappa_{HJ}^{-1} \kappa_G \frac{1+\eta_d}{1-\eta_d} \frac{(\epsilon_G + \kappa_g)}{\sqrt{|S_k|}} + \tilde{\gamma} \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \mathbb{E}[\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\
&\leq \tilde{\gamma} \frac{1+\eta_d}{1-\eta_d} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\tilde{\omega}}{1-\eta_d} \left(\tilde{\theta} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a} \tilde{\beta}^k \right) \\
&\quad + \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \frac{(\epsilon_G + \kappa_g)}{\kappa_\sigma} \left(\tilde{\theta} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a} \tilde{\beta}^k \right) + \tilde{\gamma} \kappa_{HJ}^{-1} \frac{1+\eta_d}{1-\eta_d} \left(\tilde{\theta} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a} \tilde{\beta}^k \right) \\
&= \tilde{a}_1 \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a}_2 \tilde{\beta}^k,
\end{aligned}$$

where the second inequality follows from (2.13) and Assumption 2.5, the third inequality follows by substituting $\kappa_G \leq \frac{\sqrt{\text{Var}(\nabla F(x_{k,0}) | \mathcal{F}_k)}}{\kappa_\sigma}$ and Condition 2.2 and the last equality follows from the defined constants. Taking the total expectation of the bound yields,

$$\mathbb{E}[\|d_{k,N_k}^{true}\|] \leq \tilde{a}_1 \mathbb{E}[\|d_{k-1,N_{k-1}}^{true}\|] + \tilde{a}_2 \tilde{\beta}^k.$$

Applying Lemma A.2 with $Z_k = \mathbb{E}[\|d_{k,N_k}^{true}\|]$, $\rho_1 = \tilde{a}_1$, $\rho_2 = \tilde{\beta}$ and $b = \tilde{a}_2$ completes the proof. \square

Corollary 3.6. *Suppose the conditions of Theorem 3.5 hold.*

1. *For the finite-sum problem (1.2): If the batch size is selected as $|S_k| = \lceil (1 - \beta^k) |\mathcal{S}| \rceil$ with $\beta \in (0, 1)$ and termination criterion parameters are chosen as $0 \leq \{\gamma_k\} \leq \gamma < 1$, $\epsilon_k = \omega \left(1 - \frac{|S_k|}{|\mathcal{S}|}\right)$ with $\omega \geq 0$, then the true search direction norm converges to zero at a linear rate across outer iterations, as expressed in Theorem 3.5 with $a_1 = \gamma$ and $a_2 = \omega + 2\mu_H^{-1}(1 + \gamma)(\omega_1 \kappa_g + \omega_2)$.*

2. For the expectation problem (1.3): If Assumption 2.5 is satisfied, the sample set size is chosen as $|S_{k+1}| = \left\lceil \frac{|S_k|}{\tilde{\beta}^2} \right\rceil$ with $\tilde{\beta} \in (0, 1)$ and the termination criterion parameters are chosen such that $0 \leq \{\gamma_k\} \leq \tilde{\gamma} < 1$, $\epsilon_k = \frac{\tilde{\omega}}{\sqrt{|S_k|}}$ where $\tilde{\omega} \geq 0$, then the true search direction norm converges to zero in expectation at a linear rate across outer iterations, as expressed in Theorem 3.5 with $\tilde{a}_1 = \tilde{\gamma}$ and $\tilde{a}_2 = \frac{\tilde{\omega} + \kappa_G \mu_H^{-1}(\epsilon_G + \kappa_g) + \tilde{\gamma} \mu_H^{-1}(\tilde{\omega}_1 \kappa_g + \tilde{\omega}_2)}{\sqrt{|S_0|}}$, where $|S_0|$ is the initial batch size.

Proof. The proof follows from the same procedure as Corollary 2.6. For the finite-sum problem (1.2), substituting (2.6) and the defined parameters into error bound (3.10),

$$\begin{aligned}
\|d_{k,N_k}^{true}\| &\leq \gamma \|d_{k-1,N_{k-1}}^{true}\| + \omega \left(1 - \frac{|S_k|}{|S|}\right) \\
&\quad + \mu_H^{-1} \|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| + \gamma \mu_H^{-1} \|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| \\
&\leq \gamma \|d_{k-1,N_{k-1}}^{true}\| + \omega \left(\frac{|S| - (1 - \beta^k)|S|}{|S|}\right) \\
&\quad + 2\mu_H^{-1} \left(\frac{|S| - (1 - \beta^k)|S|}{|S|}\right) (\omega_1 \|\nabla f(x_{k,N_k})\| + \omega_2 + \gamma(\omega_1 \|\nabla f(x_{k,0})\| + \omega_2)) \\
&\leq \gamma \|d_{k-1,N_{k-1}}^{true}\| + (\omega + 2\mu_H^{-1}(1 + \gamma)(\omega_1 \kappa_g + \omega_2)) \beta^k \\
&= \gamma \|d_{k-1,N_{k-1}}^{true}\| + a_2 \beta^k,
\end{aligned}$$

where the third inequality follows from Assumption 3.1 and the final equality from the defined constants. Using the above bound, applying Lemma A.2 with $Z_k = \|d_{k,N_k}^{true}\|$, $\rho_1 = \gamma$, $\rho_2 = \beta$ and $b = a_2$ completes the proof.

For the expectation problem (1.3), the expectation error bound (3.11) can be further refined as,

$$\begin{aligned}
\mathbb{E} [\|d_{k,N_k}^{true}\| | \mathcal{F}_k] &\leq \tilde{\gamma} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\tilde{\omega}}{\sqrt{|S_k|}} + \mu_H^{-1} \mathbb{E} [\|\nabla f(x_{k,N_k}) - g_{S_k}(x_{k,N_k})\| | \mathcal{F}_k] \\
&\quad + \tilde{\gamma} \mu_H^{-1} \mathbb{E} [\|\nabla f(x_{k,0}) - g_{S_k}(x_{k,0})\| | \mathcal{F}_k] \\
&\leq \tilde{\gamma} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\tilde{\omega}}{\sqrt{|S_k|}} + \kappa_G \mu_H^{-1} \frac{(\epsilon_G + \kappa_g)}{\sqrt{|S_k|}} + \tilde{\gamma} \mu_H^{-1} \frac{\tilde{\omega}_1 \|\nabla f(x_{k,0})\| + \tilde{\omega}_2}{\sqrt{|S_k|}} \\
&\leq \tilde{\gamma} \|d_{k-1,N_{k-1}}^{true}\| + \frac{\tilde{\beta}^k}{\sqrt{|S_0|}} [\tilde{\omega} + \kappa_G \mu_H^{-1}(\epsilon_G + \kappa_g) + \tilde{\gamma} \mu_H^{-1}(\tilde{\omega}_1 \kappa_g + \tilde{\omega}_2)] \\
&= \tilde{\gamma} \|d_{k-1,N_{k-1}}^{true}\| + \tilde{a}_2 \tilde{\beta}^k,
\end{aligned}$$

where the second inequality follows from (2.13) with Assumption 2.5 and (2.7), the third inequality follows from the form of S_k and Assumption 3.1 and the equality follows from the defined constants. Taking the total expectation of the above bound yields,

$$\mathbb{E} [\|d_{k,N_k}^{true}\|] \leq \tilde{\gamma} \mathbb{E} [\|d_{k-1,N_{k-1}}^{true}\|] + \tilde{a}_2 \tilde{\beta}^k.$$

Applying Lemma A.2 with $Z_k = \mathbb{E} [\|d_{k,N_k}^{true}\|]$, $\rho_1 = \tilde{\gamma}$, $\rho_2 = \tilde{\beta}$ and $b = \tilde{a}_2$ completes the proof. \square

Lemma 3.7. Suppose Assumptions 2.4, 3.1 and 3.2 hold in Algorithm 3.1.

1. For the finite-sum problem (1.2): For $k \geq 0$, Condition 3.1 is satisfied if

$$|S_k| \geq |S| \left(1 - \sqrt{\frac{\theta^2 \|d_{k,0}^{true}\|^2 + a^2 \beta^{2k}}{4(\omega_1^2 \kappa_g^2 + \omega_2^2)}}\right) \quad \text{with } \hat{\theta} = 1.$$

2. For the expectation problem (1.3): For $k \geq 0$, Condition 3.1 is satisfied if

$$|S_k| \geq \frac{\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2}{\tilde{\theta}^2 \|d_{k,0}^{true}\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}}.$$

Proof. Similar to the procedure employed in Lemma 2.7, for the finite-sum problem (1.2), from Assumption 2.4

$$\begin{aligned} \|\nabla f(x_k, 0) - g_{S_k}(x_k, 0)\|^2 &\leq 4 \left(1 - \frac{|S_k|}{|S|}\right)^2 (\omega_1^2 \|\nabla f(x_{k,0})\|^2 + \omega_2) \\ &\leq 4 \left(\sqrt{\frac{\theta^2 \|d_{k,0}^{true}\|^2 + a^2 \beta^{2k}}{4(\omega_1^2 \kappa_g^2 + \omega_2^2)}} \right)^2 (\omega_1^2 \kappa_g^2 + \omega_2) \end{aligned}$$

where the first inequality follows from (2.6) and second inequality from Assumption 3.1 and the suggested lower bound on $|S_k|$, thereby satisfying Condition 3.1. The gradient error condition on x_{k,N_k} with $\hat{\theta} = 1$ can be verified in the same way.

For the expectation problem (1.3),

$$\begin{aligned} \mathbb{E}[\|\nabla f(x_k, 0) - g_{S_k}(x_k, 0)\|^2 | \mathcal{F}_k] &\leq \frac{\tilde{\omega}_1^2 \|\nabla f(x_{k,0})\|^2 + \tilde{\omega}_2^2}{|S_k|} \\ &\leq (\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2) \frac{\tilde{\theta}^2 \|d_{k,0}^{true}\|^2 + \tilde{a}^2 \tilde{\beta}^{2k}}{\tilde{\omega}_1^2 \kappa_g^2 + \tilde{\omega}_2^2} \end{aligned}$$

where the first inequality follows from (2.7) and second inequality from Assumption 3.1 and the suggested lower bound on $|S_k|$, thereby satisfying Condition 3.1, thus completing the proof. \square