

ISE



Industrial and
Systems Engineering

Active-Set Identification in Noisy and Stochastic Optimization

FRANK E. CURTIS, DANIEL P. ROBINSON, AND LARA ZEBIANE

Department of Industrial and Systems Engineering, Lehigh University

COR@L Technical Report 25T-015-R1



LEHIGH
UNIVERSITY.

COR@L
COMPUTATIONAL OPTIMIZATION
RESEARCH AT LEHIGH 

Active-Set Identification in Noisy and Stochastic Optimization

FRANK E. CURTIS^{*1}, DANIEL P. ROBINSON^{†1}, AND LARA ZEBIANE^{‡1}

¹Department of Industrial and Systems Engineering, Lehigh University

Original Publication: August 30, 2025

Last Revised: October 15, 2026

Abstract

Identifying active constraints from a point near a local minimizer is important both theoretically and practically in constrained continuous optimization as it can help identify optimal Lagrange multipliers and essentially reduce an inequality-constrained problem to an equality-constrained one. Traditional active-set identification guarantees have been proved under various types of constraint qualifications. Importantly, they also assume exact function and derivative values. This work extends such results to settings when both objective and constraint function and derivative values have deterministic or stochastic noise. Two strategies are analyzed. For each strategy, it is shown that, under constraint qualifications and other standard conditions, the active set of a local minimizer can be identified correctly when a point is close enough to the minimizer and the noise is sufficiently small. It is also shown that regularity of a minimizer is critical. Under strict complementarity, a fixed noise level suffices, whereas in the presence of a weakly active constraint the noise must vanish at least linearly in the distance to the minimizer in general, a rate that is also shown to be sufficient for correct identification. Guarantees are also stated for the use of active-set identification strategies within a stochastic algorithm, with explicit sample-size conditions. The theoretical findings are demonstrated with two simple illustrative examples and a more realistic constrained neural-network training task.

1 Introduction

A critical feature of an algorithm for solving constrained continuous optimization problems is its ability to identify the inequality constraints that are active at a local minimizer. Suppose that, with respect to a function $c : \mathbb{R}^n \rightarrow \mathbb{R}^q$, a problem defined with respect to a decision variable $x \in \mathbb{R}^n$ involves inequality constraints of the form $c(x) \leq 0$. The aim of active-set identification is for an algorithm to be able to determine, with respect to a local minimizer and/or KKT point $x_* \in \mathbb{R}^n$, the set $\{i \in \{1, \dots, q\} : c_i(x_*) = 0\}$ of inequality constraints that hold at equality—i.e., are *active*—at x_* , even if the algorithm only has access to a solution estimate that is within a local neighborhood of x_* . Active-set identification is critical in multiple ways. One way is in the context of algorithms that, in each iteration, compute a search direction by (a) estimating the active set, and then (b) computing the search direction with respect to the estimated active set (while ignoring all of the constraints that are estimated to be *inactive*) [19, 18]. Another manner in which active-set identification is critical is in the computation of Lagrange multipliers, which in turn can be useful algorithmically and/or for sensitivity analysis of the solution of interest.

Various strong active-set identification guarantees have been proved for settings in which a problem’s objective and constraint functions are continuously differentiable and an algorithm has access to exact function and derivative values; see, e.g., [12, 16, 32, 40, 46, 52, 62]. Typically, these guarantees require that a

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

[†]E-mail: daniel.p.robinson@lehigh.edu

[‡]E-mail: lara.zebiane@lehigh.edu

constraint qualification—such as the linear independence constraint qualification (LICQ) or Mangasarian-Fromovitz constraint qualification (MFCQ) [47]—and other standard conditions hold at the local minimizer or Karush-Kuhn-Tucker (KKT) point x_* , such as strict complementarity or a second-order sufficiency condition. Each such guarantee states that, if a point x is sufficiently close to x_* , then the active set at x_* can be identified correctly by solving a subproblem involving problem function and derivative values at x .

The contributions of this paper are to show that such guarantees can be extended to settings in which the problem function and derivative values contain errors. The design and analysis of optimization algorithms when function and derivative values are subject to errors are active areas of research, which motivates such extensions. Errors can arise in multiple settings. One setting of modern interest is that of *deterministic noise*, which we use to refer to situations in which the evaluation of a problem function or derivative always results in the same approximate value whenever the evaluation occurs multiple times at the same point. This may represent situations in which the function is defined by a simulation that is subject to computational noise. We refer the reader to the recent articles [1, 23, 29, 54, 60, 64] on the design and analysis of algorithms for solving unconstrained or constrained optimization problems subject to deterministic noise. Another setting of modern interest is that of *stochastic noise*. This refers to situations in which a problem function value is defined for any point in its domain by, e.g., a large finite sum and/or an expectation over a random variable argument. The noise arises from the fact that, in practice, it is intractable to evaluate the value or derivative of such a function exactly, and instead one obtains an approximate value defined by a random sample of inputs. There are a large number of articles on the design and analysis of algorithms for solving unconstrained optimization problems when evaluations are subject to stochastic noise; see, e.g., the references in [9, 28]. For recent articles on solving constrained problems of this type, see [3, 4, 24, 25, 26, 27, 33, 48, 49].

Despite these efforts on algorithm design and analysis, little has been done to extend active-set identification guarantees to settings with errors. (An exception is [49] wherein an active-set algorithm is devised for solving problems when objective function and derivative values are subject to stochastic noise while constraint function and derivative values are exact. However, in that setting, the active-set estimation is tied to a specific algorithm and the results do not allow noise in the constraint function or derivative values.) In this paper, we consider strategies for active-set identification that may be employed irrespective of any particular algorithm, and when function and derivative values may be subject to noise. Following primarily [52], we show that when a point x is sufficiently close to a local minimizer x_* and the noise is sufficiently small, the active set at x_* can be identified correctly through solving a few types of subproblems at x . In this respect, it is particularly comforting to confirm that state-of-the-art active-set identification techniques are robust enough to extend readily to settings involving function and derivative noise. We also state how our active-set identification guarantees can be employed for accurate active-set identification in the context of a stochastic algorithm for constrained optimization. Finally, we demonstrate our theoretical guarantees through a pair of illustrative numerical experiments, one involving a couple of small-dimensional examples where the results can be seen visually and another involving a constrained neural-network training problem.

Our contributions are best understood within the context of perturbation theory of nonlinear optimization; see, e.g., [8, 44, 34, 31, 41, 56, 58, 45]. That theory describes how the solutions and optimal Lagrange multipliers of a constrained optimization problem may change with perturbations of the problem data. Broadly speaking, perturbation theory has shown that, under regularity conditions, the active set of a perturbed solution coincides with that of the unperturbed one. Active-set identification in the presence of noise asks a related, but distinct question. Since a point x at which one may attempt to perform active-set identification might not be the solution of any perturbation of the original optimization problem of interest and the errors in the function and derivative values at x do not correspond to a perturbation of the problem in the classical sense, the errors that may be observed at two different points can in general be unrelated, so there might be no perturbed problem whose solution one could follow as x approaches a minimizer x_* . Rather, what one can do is consider a subproblem at a given point x with noisy data as a perturbation of the same type of subproblem at x_* with exact data, whose solution is known, and ask when the active-set estimate produced by the former equals the active-set at x_* . This is the approach taken in some of the active-set identification literature when employing exact data, and it is the approach that we take as well.

There are two features of our theoretical results for noisy settings that deserve emphasis, since they would

not be evident from a direct appeal to classical perturbation theory or prior results on active-set identification with exact values. First, our results quantify the level of noise that is admissible for correct identification. In particular, for one of our proposed strategies that relies on strict complementarity, we state an admissible noise level explicitly in terms of the conditioning of the KKT system at x_* (see Lemma 4.5). On the other hand, for another proposed approach, we show a rate at which it suffices for the noise to shrink as $x \rightarrow x_*$, a result that is unique to our noisy setting (see Corollary 4.1). Second, we show that the dependence on regularity is critical. In particular, we show that a fixed noise level suffices under strict complementarity, whereas in the presence of a weakly active constraint the noise must, in general, vanish at least linearly in the distance $\|x - x_*\|_2$, which is the rate that we show suffices (see Theorem 4.3 and Corollary 4.1). Overall, our analysis shows that the two strategies behave differently under noise, where the primary reason for the difference is not the strategies themselves, but rather possible degeneracy at x_* .

1.1 Notation

We use $\mathbb{R}$ to denote the set of real numbers and use $\mathbb{R}_{\geq r}$ to denote the set of real numbers greater than or equal to $r \in \mathbb{R}$. Real-valued n -vectors and m -by- n matrices are denoted by $\mathbb{R}^n$ and $\mathbb{R}^{m \times n}$, respectively. Given an integer $q \geq 1$, we define the corresponding set of integers $[q] := \{1, \dots, q\}$. Given $\mathcal{A} \subseteq [q]$, we use $\mathcal{A}^c \subseteq [q]$ to denote the complement of $\mathcal{A}$ with respect to the set $[q]$, namely, $\mathcal{A}^c := [q] \setminus \mathcal{A}$. We use $\mathbf{1}$ to denote the vector with all components equal to 1; in each instance, the length of the vector is determined by the context in which it appears. Similarly, we use I to denote the identity matrix whose size is determined by its context. Given any vector (function) v , its i th component is written as v_i . In addition, given such v , we define $[v]_+ := \max\{v, 0\}$ and $[v]_- := \max\{-v, 0\}$, where in each case the maximum is taken component-wise. Given a pair of vectors $(a, b) \in \mathbb{R}^n \times \mathbb{R}^n$, we use $a \perp b$ to indicate that $a_i b_i = 0$ for all $i \in [n]$. Given sets $\mathcal{A} \subseteq \mathbb{R}^n$ and $\mathcal{B} \subseteq \mathbb{R}^n$, we denote the ℓ_2 -norm-induced distance between them as

$$\text{dist}(\mathcal{A}, \mathcal{B}) := \inf_{(a, b) \in \mathcal{A} \times \mathcal{B}} \|a - b\|_2. \quad (1)$$

1.2 Organization

In §2, we state the form of the optimization problems that we consider throughout the paper, and provide background on optimality conditions and related concepts that are referenced. In §3, we discuss active-set identification guarantees in the context of having exact problem function and derivative values. This section includes one new contribution in the form of a modification of a result from [52] from the setting of employing a linear optimization problem (LP) with a trust-region constraint to one employing a quadratic optimization problem (QP). Our main contributions are contained in §4. In §4.1 and §4.2, we prove that a Lagrange multiplier-based strategy and a primal-step-based strategy, respectively, identify the active set correctly with explicit tolerances on the allowable noise. In §4.3, we show that the different noise requirements of the two strategies reflect the regularity, or lack thereof, of the optimization problem at the local minimizer of interest. In particular, we show that in the presence of a weakly active constraint at a minimizer, no rule can identify the active set under a fixed noise level, and that the linear rate at which our analysis requires the noise to vanish cannot be improved in general. In §4.4, we demonstrate the use of our results for a stochastic algorithm, with explicit sample-size conditions. We provide the results of illustrative numerical experiments in §5 and provide concluding remarks in §6.

2 Preliminaries

Given an objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, equality-constraint function $e : \mathbb{R}^n \rightarrow \mathbb{R}^p$, and inequality-constraint function $c : \mathbb{R}^n \rightarrow \mathbb{R}^q$, where all functions (i.e., f , e , and c) are continuously differentiable, consider the constrained continuous optimization problem

$$\min_{x \in \mathbb{R}^n} f(x) \text{ s.t. } e(x) = 0 \text{ and } c(x) \leq 0. \quad (2)$$

Let the Lagrangian $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}$ of (2) be defined by $\mathcal{L}(x, y, z) = f(x) + e(x)^T y + c(x)^T z$, where $y \in \mathbb{R}^p$ and $z \in \mathbb{R}^q$ are Lagrange multiplier variables. Under a constraint qualification, first-order necessary conditions for (2) are that at a local minimizer $x \in \mathbb{R}^n$ there exist $y \in \mathbb{R}^p$ and $z \in \mathbb{R}^q$ such that

$$\nabla f(x) + \nabla e(x)y + \nabla c(x)z = 0, \quad e(x) = 0, \quad \text{and} \quad 0 \leq z \perp c(x) \leq 0, \quad (3)$$

where the first two equations correspond to the conditions $\nabla_x \mathcal{L}(x, y, z) = 0$ and $\nabla_y \mathcal{L}(x, y, z) = 0$. We refer to any tuple of vectors $(x, y, z) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^q$ satisfying (3) as a KKT point. We refer to $x \in \mathbb{R}^n$ as a primal KKT point if and only if there exists $(y, z) \in \mathbb{R}^p \times \mathbb{R}^q$ such that (x, y, z) is a KKT point.

At a local minimizer $x \in \mathbb{R}^n$ of (2), a pair (y, z) such that (x, y, z) satisfies (3) is not necessarily unique. Therefore, for any $x \in \mathbb{R}^n$, let us define $\mathcal{D}(x) := \{(y, z) \in \mathbb{R}^p \times \mathbb{R}^q : (x, y, z) \text{ satisfies (3)}\}$. That is, for any $x \in \mathbb{R}^n$, the set $\mathcal{D}(x)$ (which may be empty) contains all Lagrange multiplier pairs that yield a KKT point along with x . At a local minimizer x of (2) the set $\mathcal{D}(x)$ is nonempty and bounded if and only if the MFCQ holds at x [36], and if the LICQ holds at a local minimizer x then $\mathcal{D}(x)$ is a singleton. Strict complementarity holds for a KKT point (x, y, z) if and only if $z_i > 0$ for all $i \in [q]$ such that $c_i(x) = 0$.

Let us indicate the set of active constraints at any primal point $x \in \mathbb{R}^n$ as the set of indices of the inequality constraints that are satisfied at equality at x , namely, $\mathcal{A}(x) := \{i \in [q] : c_i(x) = 0\}$. Let (x, y, z) be a KKT point. Of the constraints whose indices are contained in $\mathcal{A}(x)$, the set of *weakly* active constraints at x are those such that the corresponding Lagrange multiplier is equal to zero for all elements of $\mathcal{D}(x)$, i.e., $\mathcal{A}_0(x) := \{i \in \mathcal{A}(x) : z_i = 0 \text{ for all } (y, z) \in \mathcal{D}(x)\}$. Later, we reference a particular second-order optimality condition. In particular, let us define the critical cone at a local minimizer $x \in \mathbb{R}^n$ as

$$\begin{aligned} \mathcal{C}(x) := \{d \in \mathbb{R}^n : \nabla e_i(x)^T d = 0 \text{ for all } i \in [p], \quad \nabla c_i(x)^T d \leq 0 \text{ for all } i \in \mathcal{A}_0(x), \text{ and} \\ \nabla c_i(x)^T d = 0 \text{ for all } i \in \mathcal{A}(x) \setminus \mathcal{A}_0(x)\}. \end{aligned}$$

Then, the second-order condition that we reference requires that, at x , one has

$$d^T \nabla_{xx}^2 \mathcal{L}(x, y, z) d > 0 \quad \text{for all } d \in \mathcal{C}(x) \setminus \{0\} \quad \text{and all } (y, z) \in \mathcal{D}(x). \quad (4)$$

Condition (4) is the second-order condition employed, e.g., in [52, 63]. It is related to the standard second-order condition as follows. For a fixed $(y, z) \in \mathcal{D}(x)$, the standard critical cone is

$$\begin{aligned} \mathcal{C}_{(y,z)}(x) := \{d \in \mathbb{R}^n : \nabla e(x)^T d = 0, \quad \nabla c_i(x)^T d \leq 0 \text{ for all } i \in \mathcal{A}(x) \text{ with } z_i = 0, \text{ and} \\ \nabla c_i(x)^T d = 0 \text{ for all } i \in \mathcal{A}(x) \text{ with } z_i > 0\}, \end{aligned}$$

and in fact $\mathcal{C}_{(y,z)}(x) = \mathcal{C}(x)$ for every $(y, z) \in \mathcal{D}(x)$. Indeed, $\mathcal{C}(x) \subseteq \mathcal{C}_{(y,z)}(x)$ since $\mathcal{A}_0(x) \subseteq \{i \in \mathcal{A}(x) : z_i = 0\}$. Conversely, any $d \in \mathcal{C}_{(y,z)}(x)$ satisfies $\nabla f(x)^T d = -\sum_{i \in \mathcal{A}(x)} z_i \nabla c_i(x)^T d = 0$ by (3), so for any $(\bar{y}, \bar{z}) \in \mathcal{D}(x)$ one has $\sum_{i \in \mathcal{A}(x)} \bar{z}_i \nabla c_i(x)^T d = 0$ with every term nonpositive, which yields $\nabla c_i(x)^T d = 0$ for all $i \in \mathcal{A}(x) \setminus \mathcal{A}_0(x)$. Thus, (4) is the standard second-order sufficient condition [51] imposed for every multiplier in $\mathcal{D}(x)$, rather than for some multiplier, and so it is stronger in general. For example, for the problem to minimize $x_1 + \frac{1}{2}\|x\|_2^2$ subject to $-x_1 \leq 0$ and $-x_1 - x_2^2 \leq 0$, the origin is a strict global minimizer at which the MFCQ holds, $\mathcal{D}(0) = \{z \in \mathbb{R}_{\geq 0}^2 : z_1 + z_2 = 1\}$, $\mathcal{A}_0(0) = \emptyset$, and $\mathcal{C}(0) = \{d \in \mathbb{R}^2 : d_1 = 0\}$. Since the Hessian of the Lagrangian at $(0, z)$ is $\text{diag}(1, 1 - 2z_2)$, the standard condition holds with $z = (1, 0)$, whereas (4) fails for every $z \in \mathcal{D}(0)$ with $z_2 \geq \frac{1}{2}$. Condition (4) is implied by the strong second-order condition in [58] holding for every $(y, z) \in \mathcal{D}(x)$ since that requires positive curvature on a subspace containing $\mathcal{C}(x)$. Under the LICQ, the optimal multiplier set $\mathcal{D}(x)$ is a singleton and $\mathcal{A}_0(x) = \{i \in \mathcal{A}(x) : z_i = 0\}$, so condition (4) reduces to the standard second-order sufficient condition in this case.

We close this section with an observation and a lemma that will be referenced later on. First, recall that by Farkas' theorem, one has that for any $g \in \mathbb{R}^n$, $A \in \mathbb{R}^{n \times p}$, and $B \in \mathbb{R}^{n \times \ell}$ where $\ell \in [q]$, one and only one of the following two systems of equations and inequalities has a solution:

$$g + Ay + B\hat{z} = 0 \text{ and } \hat{z} \geq 0 \text{ for some } (y, \hat{z}) \in \mathbb{R}^p \times \mathbb{R}^\ell; \text{ or}$$

$$A^T d = 0, \quad B^T d \leq 0, \quad \text{and } g^T d < 0 \text{ for some } d \in \mathbb{R}^n.$$

Thus, at any $x \in \mathbb{R}^n$ that is feasible for (2) in the sense that $e(x) = 0$ and $c(x) \leq 0$, either $\mathcal{D}(x)$ is nonempty or for any trust-region radius $\Delta \in (0, \infty)$ the LP given by

$$\min_{d \in \mathbb{R}^n} \nabla f(x)^T d \quad \text{s.t.} \quad e(x) + \nabla e(x)^T d = 0, \quad c(x) + \nabla c(x)^T d \leq 0, \quad \text{and} \quad \|d\|_\infty \leq \Delta \quad (5)$$

has a negative optimal value. Similarly, by the same reasoning, at any such feasible x , either $\mathcal{D}(x)$ is nonempty or for any $\theta \in (0, \infty)$ the QP given by

$$\min_{d \in \mathbb{R}^n} \nabla f(x)^T d + \frac{1}{2} \theta \|d\|_2^2 \quad \text{s.t.} \quad e(x) + \nabla e(x)^T d = 0 \text{ and } c(x) + \nabla c(x)^T d \leq 0 \quad (6)$$

has a negative optimal value. In addition, recall the following lemma pertaining to solutions of perturbed linear systems; see, e.g., [10, Theorem 7.29] or [53, Theorem 2.1.2].

Lemma 2.1. *Suppose $K \in \mathbb{R}^{n \times n}$ is nonsingular and $b \in \mathbb{R}^n \setminus \{0\}$. Let $v \in \mathbb{R}^n$ be the unique solution of $Kv = b$ (meaning $v \neq 0$ since $b \neq 0$) and let $\tilde{v}$ satisfy $\tilde{K}\tilde{v} = \tilde{b}$, where $\tilde{K} = K + \delta K$ and $\tilde{b} = b + \delta b$. Then, for any vector-induced norm $\|\cdot\|$ and with $\text{cond}(K) := \|K\| \|K^{-1}\|$, one has that*

$$\|\delta K\| \|K^{-1}\| < 1 \quad \implies \quad \frac{\|v - \tilde{v}\|}{\|v\|} \leq \frac{\text{cond}(K)}{1 - \|\delta K\| \|K^{-1}\|} \left(\frac{\|\delta K\|}{\|K\|} + \frac{\|\delta b\|}{\|b\|} \right).$$

We also make use of the following absolute form of a perturbation bound.

Lemma 2.2. *Suppose $K \in \mathbb{R}^{n \times n}$ is nonsingular, and let $\tilde{K} = K + \delta K$, where for some vector-induced norm $\|\cdot\|$ one has $\|\delta K\| \|K^{-1}\| \leq \frac{1}{2}$. Then, $\tilde{K}$ is nonsingular with $\|\tilde{K}^{-1}\| \leq 2\|K^{-1}\|$, and if $Kv = b$ and $\tilde{K}\tilde{v} = \tilde{b}$ for some $(v, b, \tilde{v}, \tilde{b}) \in \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n$, then one has that*

$$\|\tilde{v} - v\| \leq 2\|K^{-1}\| (\|\tilde{b} - b\| + \|\delta K\| \|v\|).$$

Proof. Proof. Since $\|K^{-1}\delta K\| \leq \|\delta K\| \|K^{-1}\| \leq \frac{1}{2}$, by the properties of Neumann series $I + K^{-1}\delta K$ is nonsingular with an inverse having norm at most 2. Thus, $\tilde{K}^{-1} = (K(I + K^{-1}\delta K))^{-1} = (I + K^{-1}\delta K)^{-1} K^{-1} \implies \|\tilde{K}^{-1}\| \leq 2\|K^{-1}\|$. Since subtracting $\tilde{K}v = b + \delta Kv$ from $\tilde{K}\tilde{v} = \tilde{b}$ yields $\tilde{K}(\tilde{v} - v) = (\tilde{b} - b) - \delta Kv$, multiplying this equation by $\tilde{K}^{-1}$ and employing the bound above yields the desired conclusion. $\square$ $\square$

3 Exact Function and Derivative Values

The purpose of this section is to state three active-set identification results for the setting of having exact function and derivative values. These are the results upon which we build for our new theoretical guarantees for noisy settings.

Let $x_* \in \mathbb{R}^n$ be a local minimizer of problem (2). Each active-set identification result that we present is based on solving a subproblem that is designed to determine an estimate of the optimal active set with respect to x_* , namely, $\mathcal{A}(x_*)$, at some distinct point $x \in \mathbb{R}^n$. In particular, the results that we present state that, under certain conditions, the estimate of the optimal active set is correct as long as x is sufficiently close to x_* . The first subproblem that we introduce in this section, an LP defined with respect to problem function and derivative values at x , involves computing an estimate of an element of $\mathcal{D}(x_*)$ that can, in turn, be employed to estimate $\mathcal{A}(x_*)$. The second and third subproblems that we introduce, an LP and a QP also defined with respect to function and derivative values at x , each involve computing a step in the primal space that in turn can be employed to estimate $\mathcal{A}(x_*)$. Both of these LP techniques are based directly on ones established in [52]. The QP technique is similar to the latter LP technique from [52], except that we consider a QP subproblem. Our motivation for investigating this QP technique is that this allows active-set identification to be done through subproblems that are commonly solved for computing search directions within the context of recently designed algorithms for noisy and stochastic settings.

Let $x \in \mathbb{R}^n$ be given. A strategy for estimating the active set at a nearby local minimizer x_* involves first obtaining Lagrange multiplier values $(y, z) \in \mathbb{R}^p \times \mathbb{R}^q$, and then evaluating the KKT error given by the function $\psi : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^q$ defined as

$$\psi(x, y, z) = \|\nabla f(x) + \nabla e(x)y + \nabla c(x)z\|_1 + \|e(x)\|_1 + \|\min\{z, -c(x)\}\|_1, \quad (7)$$

where the minimum is taken component-wise; recall (3). The motivation for obtaining Lagrange multiplier values and evaluating the KKT error can be seen through the following from [52]; see also [32, 38, 63].

Theorem 3.1. [52, Theorem 2.1] *Suppose that, with respect to (2), the MFCQ and the second-order condition (4) hold at a local minimizer $x_* \in \mathbb{R}^n$. Then, there exist $\epsilon \in (0, 1]$ and $C \in \mathbb{R}_{>0}$ such that, for all $(x, y, z) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^q$ with $z \geq 0$ and $\text{dist}(\{(x, y, z)\}, \{x_*\} \times \mathcal{D}(x_*)) \leq \epsilon$, one has that $C^{-1}\psi(x, y, z) \leq \text{dist}(\{(x, y, z)\}, \{x_*\} \times \mathcal{D}(x_*)) \leq C\psi(x, y, z)$.*

Motivated by Theorem 3.1, a technique is proposed in [52] for estimating the active set from a primal point $x \in \mathbb{R}^n$. At x , the idea is to compute Lagrange multipliers by minimizing $\psi(x, y, z)$ with respect to (y, z) subject to $z \geq 0$, i.e., to evaluate $\omega : \mathbb{R}^n \rightarrow \mathbb{R}$ at x that is defined by

$$\omega(x) := \min_{(y, z) \in \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q} \psi(x, y, z). \quad (8)$$

This can be done by solving a linear program with equilibrium constraints (LPEC), which in turn can be solved as a mixed-integer linear program (MIP). However, since it can be computationally expensive to solve such a problem, a technique is proposed in [52] to solve an LP approximation to the LPEC, the solution of which can in turn be used to estimate the active set. Let us refer to this as the LP-LPEC approach. This is the approach that we extend to the noisy setting in the next section.

To state the LP-LPEC approach, let us first introduce $\kappa : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}$, $\rho : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q \rightarrow \mathbb{R}$, and $\bar{\rho} : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q \rightarrow \mathbb{R}$ defined respectively by

$$\kappa(x, y, z) := \|\nabla f(x) + \nabla e(x)y + \nabla c(x)z\|_1 + \|e(x)\|_1, \quad (9a)$$

$$\rho(x, y, z) := \kappa(x, y, z) + \sum_{\{i: c_i(x) < 0\}} -c_i(x)z_i + \sum_{\{i: c_i(x) \geq 0\}} c_i(x), \quad (9b)$$

$$\text{and } \bar{\rho}(x, y, z) := \kappa(x, y, z) + \sum_{\{i: c_i(x) < 0\}} (-c_i(x)z_i)^{1/2} + \sum_{\{i: c_i(x) \geq 0\}} c_i(x). \quad (9c)$$

At $x \in \mathbb{R}^n$, an approximate solution of the aforementioned LPEC can then be obtained by solving the following subproblem for the given constant $M \in \mathbb{R}_{>0}$:

$$\begin{aligned} & \min_{(y, z) \in \mathbb{R}^p \times \mathbb{R}^q} \rho(x, y, z) \quad \text{s.t.} \quad 0 \leq z \leq M\mathbf{1} \\ & \text{where } M := \max \left\{ \|c(x_*)\|_\infty, \max_{(y_*, z_*) \in \mathcal{D}(x_*)} \left\| \begin{bmatrix} y_* \\ z_* \end{bmatrix} \right\|_\infty \right\} + 1. \end{aligned} \quad (10)$$

The constant M is finite if and only if the MFCQ holds at x_* ; see [36]. A variant of (10) where the upper bound is dropped is considered in [52] and in our experiments it proved to be effective, although this variant is not covered by the theory in this paper. The following remark concerns a practical implementation of (10).

Remark 3.1. *Since the value of M in (10) is usually unknown in practice, one could set an initial estimate for M and double it every time a bound constraint involving this estimate is active at the computed solution of (10), or of its noisy counterpart (27) in §4.1 (similar to penalty-parameter updates for exact penalty methods [13, 21, 39]). If this strategy terminates, then none of these bound constraints is active at the final computed solution, so by convexity of the subproblem this solution also solves the subproblem with these bound constraints removed. Consequently, if the final estimate is at most the value of M defined in (10), then the final computed solution also solves the subproblem with that value of M , and our identification results apply*

to it. If instead the final estimate, call it M' , exceeds M , then for any $\bar{M} \geq M'$ the upcoming results of §4.1 still hold with M replaced by M' in (27), with thresholds that depend on $\bar{M}$ but not on M' , since the proof of Theorem 4.1 uses the bounds in (27) only to ensure that certain points with ℓ_∞ -norm at most $M - \frac{1}{2}$ are feasible and that Lemma 4.2 applies at the computed solution.

A solution of subproblem (10) can be obtained by solving the LP

$$\begin{aligned} \min_{(y,z,r,s) \in \mathbb{R}^p \times \mathbb{R}^q \times \mathbb{R}^n \times \mathbb{R}^n} \quad & \mathbf{1}^T r + \mathbf{1}^T s + \sum_{\{i: c_i(x) < 0\}} -c_i(x) z_i \\ \text{s.t.} \quad & \nabla f(x) + \nabla e(x)y + \nabla c(x)z = r - s, \quad r \geq 0, \quad s \geq 0, \quad \text{and } M\mathbf{1} \geq z \geq 0. \end{aligned} \quad (11)$$

Given x , let (y_x, z_x) be an optimal solution of (10). An active-set estimate can then be obtained through evaluation of the function $\bar{\rho}$; in particular, let the estimate be

$$\mathcal{A}_{\text{LP}}(y_x, z_x; x) := \{i \in [q] : c_i(x) \geq -(\beta \bar{\rho}(x, y_x, z_x))^\sigma\}, \quad (12)$$

where $\beta \in \mathbb{R}_{>0}$ and $\sigma \in (0, 1)$ are user-defined constants. Theorem 3.2 below, from [52], shows that this LP-LPEC approach can be employed under certain assumptions to provide an accurate active-set estimate from x near a local minimizer x_* . This theorem (and Theorem 4.1 in the noisy setting) holds for any $\beta \in \mathbb{R}_{>0}$ and $\sigma \in (0, 1)$, and the radius of the neighborhood on which identification is guaranteed depends on both of these parameters. The parameter σ , in particular, yields a tradeoff. Since $\bar{\rho}(x, y_x, z_x) \rightarrow 0$ as $x \rightarrow x_*$, a value of σ near 1 makes the threshold $(\beta \bar{\rho}(x, y_x, z_x))^\sigma$ small, and hence rejects inactive constraints on a larger neighborhood. On the other hand, a value of σ near 0 keeps the threshold large relative to this residual, which makes it more robust to errors in $\bar{\rho}(x, y_x, z_x)$. In §4, this robustness is what is needed to identify active constraints subject to noise. The parameter β should be scaled to the problem data; see §5.

Theorem 3.2. [52, Theorem 3.7] *Suppose that, with respect to (2), the MFCQ and the second-order condition (4) hold at a local minimizer $x_* \in \mathbb{R}^n$. Then, with $\epsilon \in (0, 1]$ defined as in Theorem 3.1 and any $\beta \in \mathbb{R}_{>0}$ and $\sigma \in (0, 1)$, there exists $\hat{\epsilon} \in (0, \epsilon]$ such that for all $x \in \mathbb{R}^n$ with $\|x - x_*\|_2 \leq \hat{\epsilon}$, the following hold.*

- (a) *There exists a solution (y_x, z_x) of (10) such that $\text{dist}(\{(y_x, z_x)\}, \mathcal{D}(x_*)) \leq \frac{1}{2}\epsilon$.*
- (b) *Any (y_x, z_x) solving (10) has $M^{-1}\rho(x, y_x, z_x) \leq \omega(x) \leq \bar{\rho}(x, y_x, z_x)$.*
- (c) *Any (y_x, z_x) solving (10) yields $\mathcal{A}_{\text{LP}}(y_x, z_x; x) = \mathcal{A}(x_*)$.*

This completes our discussion of the LP-LPEC approach that we extend in the next section.

Let us now introduce the other LP-based approach and a related QP-based approach that we extend to the noisy setting as well. If a local minimizer x_* is a primal KKT point and $x = x_*$, then both (5) and (6) are feasible with an optimal value of zero. In particular, $d_* = 0$ is an optimal solution of each, and it is the unique optimal solution of (6) by strong convexity. (For the LP (5), any feasible d in the critical cone with $\nabla f(x_*)^T d = 0$ is also optimal, so uniqueness cannot be expected.) At a point $x \neq x_*$, a negative optimal value is to be expected and is precisely what a primal search direction requires. The difficulty, however, is that the linearized constraints may be inconsistent, in which case (5) and (6) are infeasible. The standard remedy, which goes back to the $S\ell_1$ QP method of Fletcher [35] and the robust SQP methods of Burke and Han [11, 15], and which is used in certain solvers in so-called elastic mode [37, 20], is to penalize violation of the linearized constraints with an exact ℓ_1 -norm penalty. To that end, and in particular to compute a primal search direction that can also be used to determine an optimal active-set estimate at x , one can choose $(\Delta, \nu) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ and solve the nonsmooth subproblem

$$\min_{d \in \mathbb{R}^n} \quad \nabla f(x)^T d + \nu \|e(x) + \nabla e(x)^T d\|_1 + \nu \|[c(x) + \nabla c(x)^T d]_+\|_1 \quad \text{s.t.} \quad \|d\|_\infty \leq \Delta \quad (13)$$

or choose $(\theta, \nu) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ and solve the nonsmooth subproblem

$$\min_{d \in \mathbb{R}^n} \quad \nabla f(x)^T d + \frac{1}{2}\theta \|d\|_2^2 + \nu \|e(x) + \nabla e(x)^T d\|_1 + \nu \|[c(x) + \nabla c(x)^T d]_+\|_1. \quad (14)$$

Alternatively, one can solve the constrained LP subproblem

$$\begin{aligned}
& \min_{(d,r,s,t) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^p \times \mathbb{R}^q} \nabla f(x)^T d + \nu(\mathbf{1}^T r + \mathbf{1}^T s + \mathbf{1}^T t) \\
& \text{s.t. } e(x) + \nabla e(x)^T d = r - s, \quad c(x) + \nabla c(x)^T d \leq t, \\
& \quad r \geq 0, \quad s \geq 0, \quad t \geq 0, \quad \text{and } \|d\|_\infty \leq \Delta
\end{aligned} \tag{15}$$

that is equivalent to (13), or one can solve the constrained QP subproblem

$$\begin{aligned}
& \min_{(d,r,s,t) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^p \times \mathbb{R}^q} \nabla f(x)^T d + \nu(\mathbf{1}^T r + \mathbf{1}^T s + \mathbf{1}^T t) + \frac{1}{2}\theta\|d\|_2^2 \\
& \text{s.t. } e(x) + \nabla e(x)^T d = r - s, \quad c(x) + \nabla c(x)^T d \leq t, \\
& \quad r \geq 0, \quad s \geq 0, \quad \text{and } t \geq 0
\end{aligned} \tag{16}$$

that is equivalent to (14). These subproblems are feasible for any such (x, Δ, θ, ν) ; indeed, $d = 0$, $r = [e(x)]_+$, $s = [e(x)]_-$, and $t = [c(x)]_+$ form a feasible point for either of the subproblems. The equivalence of nonsmooth (13)–(14) with smooth constrained (15)–(16), and the form of the duals (17)–(18) below, are classical; see [35, 11], and see [21] for strategies for updating ν . Our interest in (14) specifically is that it is the subproblem solved in several recently proposed methods for noisy and stochastic constrained optimization [3, 4, 24, 25, 26, 27]. The dual of (15) can be written as the following (equivalent to a constrained LP)

$$\begin{aligned}
& \max_{(\alpha, \beta) \in \mathbb{R}^p \times \mathbb{R}^q} e(x)^T \alpha + c(x)^T \beta - \Delta \|\nabla f(x) + \nabla e(x)\alpha + \nabla c(x)\beta\|_1 \\
& \text{s.t. } -\nu\mathbf{1} \leq \alpha \leq \nu\mathbf{1} \text{ and } 0 \leq \beta \leq \nu\mathbf{1},
\end{aligned} \tag{17}$$

whereas the dual of (16) can be written as the QP

$$\begin{aligned}
& \max_{(\alpha, \beta) \in \mathbb{R}^p \times \mathbb{R}^q} e(x)^T \alpha + c(x)^T \beta - \frac{1}{2}\theta^{-1} \|\nabla f(x) + \nabla e(x)\alpha + \nabla c(x)\beta\|_2^2 \\
& \text{s.t. } -\nu\mathbf{1} \leq \alpha \leq \nu\mathbf{1} \text{ and } 0 \leq \beta \leq \nu\mathbf{1}.
\end{aligned} \tag{18}$$

These dual subproblems are also feasible for any (x, Δ, θ, ν) ; indeed, $(\alpha, \beta) = (0, 0)$ is feasible for both dual subproblems. Consequently, by weak duality, the optimal values of (15), (16), (17), and (18) are all finite, and in each case these values are attained by some primal solution (d, r, s, t) and some dual solution (α, β) , respectively; see, e.g., [5, Proposition 1.4.12].

Given any (x, Δ, ν) and a corresponding solution $(d_{x,\Delta,\nu}, r_{x,\Delta,\nu}, s_{x,\Delta,\nu}, t_{x,\Delta,\nu})$ of subproblem (15), an estimate of the optimal active set is given by

$$\mathcal{A}_{\text{LP}}(d_{x,\Delta,\nu}; x) := \{i \in [q] : c_i(x) + \nabla c_i(x)^T d_{x,\Delta,\nu} \geq 0\}. \tag{19}$$

Similarly, given any (x, θ, ν) and a corresponding solution $(d_{x,\theta,\nu}, r_{x,\theta,\nu}, s_{x,\theta,\nu}, t_{x,\theta,\nu})$ of subproblem (16), an estimate of the optimal active set is given by

$$\mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x) := \{i \in [q] : c_i(x) + \nabla c_i(x)^T d_{x,\theta,\nu} \geq 0\}. \tag{20}$$

Theorem 3.3 below shows that this QP-based approach is successful when x is sufficiently close to a local minimizer of (2) at which the LICQ and strict complementarity hold. The same kind of result can be proved for the LP-based approach; indeed, this is what has been proved in Lemma 3.1, Theorem 3.2, and Corollary 3.3 in [52]. (For the sake of brevity, we do not restate those results here, since they are identical to our theorem below, except for the use of an LP subproblem with a trust-region constraint instead of our QP.) Our proof of the following theorem follows closely the analysis in [52], although details are different due to our use of a QP subproblem; see Appendix A.

Theorem 3.3. *Suppose that, with respect to (2), the MFCQ holds at a local minimizer $x_* \in \mathbb{R}^n$.*

- (a) There exists $(\epsilon, \underline{\theta}) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ such that, for any $(x, \theta, \nu) \in \mathbb{R}^n \times \mathbb{R}_{>\underline{\theta}} \times \mathbb{R}_{>0}$ with $\|x - x_*\|_2 \leq \epsilon$, the solution $d_{x,\theta,\nu}$ of (14) yields $\mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x) \subseteq \mathcal{A}(x_*)$.
- (b) For any $\zeta \in (0, 1)$, there exists $(\epsilon, \underline{\theta}, \underline{\nu}) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0} \times \mathbb{R}_{>0}$ such that for any $(x, \theta, \nu) \in \mathbb{R}^n \times \mathbb{R}_{>\underline{\theta}} \times \mathbb{R}_{>\underline{\nu}}$ with $\|x - x_*\|_2 \leq \epsilon$ and $\theta \leq \|x - x_*\|_2^{-\zeta}$, the solution $d_{x,\theta,\nu}$ of (14) yields

$$-\nabla f(x_*) \in \text{range}[\nabla e(x_*)] + \text{pos}[\{\nabla c_i(x_*) : i \in \mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x)\}], \quad (21)$$

where $\text{range}[\cdot]$ denotes the range space of the matrix input and $\text{pos}[\cdot]$ denotes the set of nonnegative combinations of the input set of vectors.

- (c) If the LICQ holds at x_* , then for any $(\zeta, \epsilon, \underline{\theta}, \underline{\nu}, x, \theta, \nu)$ satisfying the conditions of part (b), the solution $d_{x,\theta,\nu}$ of (14) yields $\mathcal{A}(x_*) \setminus \mathcal{A}_0(x_*) \subseteq \mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x) \subseteq \mathcal{A}(x_*)$. If strict complementarity also holds at x_* , then $\mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x) = \mathcal{A}(x_*)$.

The proof of Theorem 3.3 reveals the influence of the parameters of the QP approach. Intuitively, the lower threshold $\underline{\theta}$ ensures that $d_{x,\theta,\nu}$ is sufficiently small in norm such that the step in (20) does not “go too far” to include indices of constraints that are inactive at x_* . The lower threshold $\underline{\nu}$, on the other hand, is a typical type of bound for exact penalization; it ensures that, when possible, the solution of (16) will have the appropriate auxiliary variable values equal to zero. Part (b) of Theorem 3.3 requires $\underline{\theta} < \theta \leq \|x - x_*\|_2^{-\zeta}$, which is a range that is nonempty for x sufficiently close to x_* . The lower bound ensures that the step is sufficiently small while the upper bound ensures that the regularization does not dominate the linear model.

4 Noisy Function and Derivative Values

We now assume that, at any $x \in \mathbb{R}^n$, one only has access to approximations of the values of f , e , and c and their first-order derivatives. More precisely, we assume that there exist positive error bounds $(\epsilon_f, \epsilon_e, \epsilon_c, \epsilon_{\nabla f}, \epsilon_{\nabla e}, \epsilon_{\nabla c})$ such that at all $x \in \mathbb{R}^n$ one can only compute $\tilde{f}(x) \in \mathbb{R}$, $\tilde{e}(x) \in \mathbb{R}^p$, $\tilde{c}(x) \in \mathbb{R}^q$, $\widetilde{\nabla f}(x) \in \mathbb{R}^n$, $\widetilde{\nabla e}(x) \in \mathbb{R}^{n \times p}$, and $\widetilde{\nabla c}(x) \in \mathbb{R}^{n \times q}$ satisfying

$$\begin{aligned} \|f(x) - \tilde{f}(x)\|_2 &\leq \epsilon_f, \quad \|e(x) - \tilde{e}(x)\|_2 \leq \epsilon_e, \quad \|c(x) - \tilde{c}(x)\|_2 \leq \epsilon_c, \\ \|\nabla f(x) - \widetilde{\nabla f}(x)\|_2 &\leq \epsilon_{\nabla f}, \quad \|\nabla e(x) - \widetilde{\nabla e}(x)\|_2 \leq \epsilon_{\nabla e}, \quad \|\nabla c(x) - \widetilde{\nabla c}(x)\|_2 \leq \epsilon_{\nabla c}. \end{aligned} \quad (22)$$

Here and throughout, vectors are measured in the ℓ_2 -norm and matrices in the induced (spectral) norm, so, e.g., $|c_i(x) - \tilde{c}_i(x)| \leq \epsilon_c$ for every $i \in [q]$. The results of §4.1 and §4.2 are statements about any realization of the approximate values satisfying (22). They hold regardless of whether the noise is deterministic or stochastic, and no property of the error distribution other than (22) is used. In §4.4, our probabilistic results are based on the probability with which (22) holds. For convenience, let us define the vector

$$\epsilon_{\text{err}} := [\epsilon_f, \epsilon_e, \epsilon_c, \epsilon_{\nabla f}, \epsilon_{\nabla e}, \epsilon_{\nabla c}]^T. \quad (23)$$

4.1 Identification through Lagrange Multipliers

We first describe a technique for identifying the active set by estimating Lagrange multipliers at a primal point $x \in \mathbb{R}^n$. This approach extends the LP-LPEC method described in §3, adapting it to noisy settings.

With the goal of again constructing a two-sided bound on $\omega(x)$, we follow (9) to define the noisy counterparts $\tilde{\kappa} : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}$, $\tilde{\rho} : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q \rightarrow \mathbb{R}$, and $\tilde{\bar{\rho}} : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q \rightarrow \mathbb{R}$ by

$$\tilde{\kappa}(x, y, z) := \|\widetilde{\nabla f}(x) + \widetilde{\nabla e}(x)y + \widetilde{\nabla c}(x)z\|_1 + \|\tilde{e}(x)\|_1, \quad (24a)$$

$$\tilde{\rho}(x, y, z) := \tilde{\kappa}(x, y, z) + \sum_{\{i: \tilde{c}_i(x) < 0\}} -\tilde{c}_i(x)z_i + \sum_{\{i: \tilde{c}_i(x) \geq 0\}} \tilde{c}_i(x), \quad (24b)$$

$$\text{and } \tilde{\rho}(x, y, z) := \tilde{\kappa}(x, y, z) + \sum_{\{i: \tilde{c}_i(x) < 0\}} (-\tilde{c}_i(x) z_i)^{1/2} + \sum_{\{i: \tilde{c}_i(x) \geq 0\}} \tilde{c}_i(x). \quad (24c)$$

The following relationship can be established for $\tilde{\rho}$ and $\tilde{\rho}$, drawing on [52, Lemma 3.5].

Lemma 4.1. *For any $(x, y, z) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q$, $\tilde{\rho}(x, y, z) \leq \tilde{\rho}(x, y, z) + q^{1/2} \tilde{\rho}(x, y, z)^{1/2}$.*

Proof. Proof. The proof follows the same logic as [52, Lemma 3.5]. $\square$ $\square$

A relationship between $\tilde{\rho}$, $\tilde{\rho}$, and ψ (recall (7)) is given in our next lemma. The existence of the error bounds in the lemma is crucial for our subsequent analysis for the noisy setting.

Lemma 4.2. *Let $\bar{M} \in \mathbb{R}_{\geq 1}$, define the error bounds*

$$\epsilon_\kappa := p^{1/2} \epsilon_e + n^{1/2} (\epsilon_{\nabla f} + p^{1/2} \bar{M} \epsilon_{\nabla e} + q^{1/2} \bar{M} \epsilon_{\nabla c}), \quad \epsilon_\rho := \epsilon_\kappa + q \bar{M} \epsilon_c, \quad \text{and} \quad \epsilon_\psi := \epsilon_\kappa + q \epsilon_c, \quad (25)$$

and let $\tilde{\psi}(x, y, z) := \tilde{\kappa}(x, y, z) + \|\min\{z, -\tilde{c}(x)\}\|_1$ denote the KKT error (7) evaluated with the approximations in place of the exact values. Then, $\epsilon_\rho \rightarrow 0$ and $\epsilon_\psi \rightarrow 0$ as $\epsilon_{\text{err}} \rightarrow 0$ (recall (23)). Furthermore, for all $(x, y, z) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q$ with $z \geq 0$ and $\|(c(x), y, z)\|_\infty \leq \bar{M}$ one has with these definitions that

$$|\rho(x, y, z) - \tilde{\rho}(x, y, z)| \leq \epsilon_\rho \quad \text{and} \quad |\psi(x, y, z) - \tilde{\psi}(x, y, z)| \leq \epsilon_\psi, \quad (26a)$$

$$\tilde{\psi}(x, y, z) \leq \tilde{\rho}(x, y, z), \quad (26b)$$

$$\text{which further imply } \bar{M}^{-1} (\tilde{\rho}(x, y, z) - \epsilon_\rho) \leq \psi(x, y, z) \leq \tilde{\rho}(x, y, z) + \epsilon_\psi. \quad (26c)$$

Moreover, if such a tuple (x, y, z) has $\rho(x, y, z) > 2\epsilon_\rho$, then the left-hand side of (26c) is strictly positive.

Proof. Proof. The fact that $\epsilon_\rho \rightarrow 0$ and $\epsilon_\psi \rightarrow 0$ as $\epsilon_{\text{err}} \rightarrow 0$ follows directly from (23) and (25). Now consider (x, y, z) as stated in the lemma and, for notational convenience, abbreviate $c_i := c_i(x)$ and $\tilde{c}_i := \tilde{c}_i(x)$ so that $|c_i - \tilde{c}_i| \leq \epsilon_c$ for all $i \in [q]$ by (22). For all $i \in [q]$, define the function $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ by $\phi_i(t) = -tz_i$ for all $t < 0$ and $\phi_i(t) = t$ for all $t \geq 0$ so that, by (9) and (24), one has that

$$\rho(x, y, z) = \kappa(x, y, z) + \sum_{i \in [q]} \phi_i(c_i) \quad \text{and} \quad \tilde{\rho}(x, y, z) = \tilde{\kappa}(x, y, z) + \sum_{i \in [q]} \phi_i(\tilde{c}_i).$$

The leading-term difference can be bounded since (22), $\|y\|_2 \leq p^{1/2} \bar{M}$, and $\|z\|_2 \leq q^{1/2} \bar{M}$ give

$$\begin{aligned} & |\kappa(x, y, z) - \tilde{\kappa}(x, y, z)| \\ & \leq \|e(x) - \tilde{e}(x)\|_1 + \|(\nabla f(x) - \tilde{\nabla} f(x)) + (\nabla e(x) - \tilde{\nabla} e(x))y + (\nabla c(x) - \tilde{\nabla} c(x))z\|_1 \\ & \leq p^{1/2} \epsilon_e + n^{1/2} (\epsilon_{\nabla f} + \epsilon_{\nabla e} \|y\|_2 + \epsilon_{\nabla c} \|z\|_2) \leq \epsilon_\kappa. \end{aligned}$$

Next, observe that ϕ_i is piecewise linear and continuous with slope either $-z_i$ or 1 at any point, hence it is Lipschitz with constant $\max\{1, z_i\} \leq \bar{M}$. Therefore $|\phi_i(c_i) - \phi_i(\tilde{c}_i)| \leq \bar{M} \epsilon_c$, and summing over $i \in [q]$ and combining with the bound above gives the first inequality in (26a). It remains to establish the second inequality in (26a) along with (26b), (26c), and the final claim. With $\tilde{\psi}$ as defined, for each $i \in [q]$ one has $|\min\{z_i, -\tilde{c}_i\}| = \tilde{c}_i$ if $\tilde{c}_i \geq 0$ and $|\min\{z_i, -\tilde{c}_i\}| \leq (-\tilde{c}_i z_i)^{1/2}$ if $\tilde{c}_i < 0$ since the minimum of two nonnegative numbers is at most their geometric mean. Hence $\tilde{\psi}(x, y, z) \leq \tilde{\rho}(x, y, z)$, which is (26b). Moreover, since $\psi(x, y, z) = \kappa(x, y, z) + \sum_{i \in [q]} |\min\{z_i, -c_i\}|$ and $t \mapsto |\min\{z_i, -t\}|$ is Lipschitz with constant 1 for each $i \in [q]$, one has $|\psi(x, y, z) - \tilde{\psi}(x, y, z)| \leq |\kappa(x, y, z) - \tilde{\kappa}(x, y, z)| + \sum_{i \in [q]} |c_i - \tilde{c}_i| \leq \epsilon_\kappa + q \epsilon_c = \epsilon_\psi$, which is the second inequality in (26a). Thus, with (26b), $\psi(x, y, z) \leq \tilde{\psi}(x, y, z) + \epsilon_\psi \leq \tilde{\rho}(x, y, z) + \epsilon_\psi$, which is the upper bound in (26c). Now, for the lower bound in (26c), one has $\psi(x, y, z) \geq \bar{M}^{-1} \rho(x, y, z)$ (see [52, Lemma 3.6]) since $\bar{M} \geq 1$, the terms of $\rho(x, y, z)$ and $\psi(x, y, z)$ with $c_i \geq 0$ coincide, and $c_i < 0$ implies $-c_i z_i \leq \max\{z_i, -c_i\} \min\{z_i, -c_i\} \leq \bar{M} \min\{z_i, -c_i\}$ by $\|(c(x), y, z)\|_\infty \leq \bar{M}$; combining with the first inequality in (26a) yields $\psi(x, y, z) \geq \bar{M}^{-1} (\tilde{\rho}(x, y, z) - \epsilon_\rho)$. Finally, if $\rho(x, y, z) > 2\epsilon_\rho$, then (26a) gives $\tilde{\rho}(x, y, z) - \epsilon_\rho \geq \rho(x, y, z) - 2\epsilon_\rho > 0$, so the left-hand side of (26c) is strictly positive. $\square$ $\square$

The LP-LPEC-based active-set identification strategy can now be established as follows. First, one computes a solution (y_x, z_x) of the optimization problem

$$\min_{(y,z) \in \mathbb{R}^p \times \mathbb{R}^q} \tilde{\rho}(x, y, z) \text{ s.t. } -M\mathbf{1} \leq y \leq M\mathbf{1}, \quad 0 \leq z \leq M\mathbf{1}, \quad (27)$$

where M is defined in (10). Concerning the challenge of not (typically) knowing the value of M in practice, see Remark 3.1. The solution of (27) can be obtained by solving the LP

$$\begin{aligned} \min_{(y,z,r,s) \in \mathbb{R}^p \times \mathbb{R}^q \times \mathbb{R}^n \times \mathbb{R}^n} \quad & \mathbf{1}^T r + \mathbf{1}^T s + \sum_{\{i: \tilde{c}_i(x) < 0\}} -\tilde{c}_i(x) z_i \\ \text{s.t.} \quad & \widetilde{\nabla} f(x) + \widetilde{\nabla} e(x) y + \widetilde{\nabla} c(x) z = r - s, \\ & r \geq 0, \quad s \geq 0, \quad -M\mathbf{1} \leq y \leq M\mathbf{1}, \quad 0 \leq z \leq M\mathbf{1}. \end{aligned}$$

The active-set estimate is then defined, similarly to (12), by

$$\tilde{\mathcal{A}}_{\text{LP}}(y_x, z_x; x) := \{i \in [q] : \tilde{c}_i(x) \geq -(\beta \tilde{\rho}(x, y_x, z_x))^\sigma\} \quad (28)$$

where $\beta \in \mathbb{R}_{>0}$ and $\sigma \in (0, 1)$ are user-defined constants.

In the following theorem, we demonstrate that this methodology yields a two-sided bound on the value of $\omega(x)$ and can accurately identify the optimal active set at a point near a local minimizer.

Theorem 4.1. *Suppose that, with respect to (2), the MFCQ and (4) hold at a local minimizer $x_* \in \mathbb{R}^n$. Let ϵ and C be defined as in Theorem 3.1, let $\epsilon' \in (0, \epsilon/2]$ be the radius as in [52, Theorem 3.4], let $\bar{M} \geq M$ with M defined in (10), and let ϵ_ρ and ϵ_ψ be defined by (25). Consider any $\beta \in \mathbb{R}_{>0}$ and $\sigma \in (0, 1)$. Then, there exist $\hat{\epsilon} \in (0, \epsilon']$ and $\bar{\epsilon} \in \mathbb{R}_{>0}$ such that for any $x \in \mathbb{R}^n$ with $\|x - x_*\|_2 \leq \hat{\epsilon}$ and any approximations satisfying (22) with $\|\epsilon_{\text{err}}\|_\infty \leq \bar{\epsilon}$, the following hold, where (y_x, z_x) denotes any solution of (27).*

- (a) *If $\rho(x, y_x, z_x) > 2\epsilon_\rho$, then $0 < \bar{M}^{-1}(\tilde{\rho}(x, y_x, z_x) - \epsilon_\rho) \leq \omega(x) \leq \tilde{\rho}(x, y_x, z_x) + \epsilon_\psi$.*
- (b) *If $\epsilon_\psi \leq (2C)^{-1}\|x - x_*\|_2$, then $\tilde{\mathcal{A}}_{\text{LP}}(y_x, z_x; x) = \mathcal{A}(x_*)$.*

Proof. Throughout the proof, let us restrict attention to $x \in \mathbb{R}^n$ satisfying $\|x - x_*\|_2 \leq \epsilon'$ and approximations satisfying (22). The (potentially smaller) radius $\hat{\epsilon} \in (0, \epsilon']$ and the (potentially tighter) bound on the approximations $\bar{\epsilon}$ that are required by the theorem are developed in the proof below.

Let us begin by observing that there exists a minimizer of $\psi(x, \cdot, \cdot)$ that is feasible for (27). Indeed, by [52, Theorem 3.4], there exists $(\hat{y}, \hat{z})$ attaining the minimum in (8), where in particular $\hat{z} \geq 0$ and $\omega(x) = \psi(x, \hat{y}, \hat{z})$, with $\text{dist}(\{(\hat{y}, \hat{z})\}, \mathcal{D}(x_*)) \leq \epsilon/2$. Since $\mathcal{D}(x_*)$ is nonempty and closed, this distance is attained, and since every element of $\mathcal{D}(x_*)$ has ℓ_∞ -norm at most $M - 1$ by (10), it follows with $\epsilon \leq 1$ that $\|(\hat{y}, \hat{z})\|_\infty \leq M - 1 + \epsilon/2 \leq M - \frac{1}{2}$. Thus, $(\hat{y}, \hat{z})$ is feasible for (27), as claimed.

Next, let us derive a lower bound for $\bar{\rho}(x, y_x, z_x)$ that does not depend on the noise, namely,

$$\bar{\rho}(x, y_x, z_x) \geq \psi(x, y_x, z_x) \geq \omega(x) = \psi(x, \hat{y}, \hat{z}) \geq C^{-1}\|x - x_*\|_2. \quad (29)$$

For the first inequality, observe that $\kappa(x, y, z)$ in (9) is the sum of the first two terms of $\psi(x, y, z)$ in (7), so all that remains is to consider the final terms in the definitions of $\bar{\rho}(x, y, z)$ and $\psi(x, y, z)$. If $c_i(x) \geq 0$, then $|\min\{z_i, -c_i(x)\}| = c_i(x)$ and the terms match. If $c_i(x) < 0$, then z_i and $-c_i(x)$ are both nonnegative, and since the minimum of two nonnegative numbers is at most their geometric mean, it follows that $|\min\{z_i, -c_i(x)\}| \leq (-c_i(x)z_i)^{1/2}$. Overall, the first inequality in (29) follows. The second inequality in (29) holds since (y_x, z_x) is feasible for (8), and the equality holds by the definition of $(\hat{y}, \hat{z})$. For the last inequality in (29), consider Theorem 3.1 at $(x, \hat{y}, \hat{z})$, the conditions of which hold since $\hat{z} \geq 0$ and

$$\text{dist}(\{(x, \hat{y}, \hat{z})\}, \{x_*\} \times \mathcal{D}(x_*)) \leq \|x - x_*\|_2 + \text{dist}(\{(\hat{y}, \hat{z})\}, \mathcal{D}(x_*)) \leq \epsilon' + \epsilon/2 \leq \epsilon.$$

Theorem 3.1 then gives $\|x - x_*\|_2 \leq \text{dist}(\{(x, \hat{y}, \hat{z})\}, \{x_*\} \times \mathcal{D}(x_*)) \leq C\psi(x, \hat{y}, \hat{z})$, where the first inequality holds since every element of $\{x_*\} \times \mathcal{D}(x_*)$ has x_* as its first block, which gives the final inequality in (29).

(As an aside, note that Theorem 3.1 cannot be applied directly at (x, y_x, z_x) , since nothing ensures that a solution of (27), which is computed from approximate values, lies within distance ϵ of $\{x_*\} \times \mathcal{D}(x_*)$. This is why the bound is obtained here through $(\hat{y}, \hat{z})$ attaining the minimum in (8).)

Our aim now is to introduce conditions that define the thresholds $\hat{\epsilon}$ and $\bar{\epsilon}$ such that the desired conclusions of the theorem hold. Toward this end, let $L_c \in \mathbb{R}_{>0}$ be a Lipschitz constant for c on $\{x \in \mathbb{R}^n : \|x - x_*\|_2 \leq \epsilon'\}$, let $m := \min_{i \notin \mathcal{A}(x_*)} |c_i(x_*)| > 0$ (with $m := \infty$ if $\mathcal{A}(x_*) = [q]$), and define

$$\tau := \min \left\{ 2, \frac{1}{4\beta} \left(\frac{m}{4} \right)^{1/\sigma} \right\}, \quad t_\tau := \min \left\{ \frac{\tau}{2}, \frac{\tau^2}{4q} \right\}, \quad \text{and} \quad \bar{\delta} := \left(\frac{(\beta/(2C))^\sigma}{2 \max\{L_c, (2Cq)^{-1}\}} \right)^{1/(1-\sigma)}.$$

Fix $(y_*, z_*) \in \mathcal{D}(x_*)$ and choose $\hat{\epsilon} \in (0, \min\{\epsilon', \bar{\delta}\}]$ and $\bar{\epsilon} \in \mathbb{R}_{>0}$ sufficiently small that, for all x with $\|x - x_*\|_2 \leq \hat{\epsilon}$ and all approximations with $\|\epsilon_{\text{err}}\|_\infty \leq \bar{\epsilon}$, one has (i) $\|c(x)\|_\infty \leq \bar{M}$; (ii) $c_i(x) \leq \frac{1}{2}c_i(x_*) < 0$ for all $i \notin \mathcal{A}(x_*)$; (iii) $\rho(x, y_*, z_*) \leq \frac{1}{2}t_\tau$ and $\epsilon_\rho \leq \frac{1}{4}t_\tau$; and (iv) $\epsilon_c \leq \frac{1}{4}m$. Conditions (i), (ii), and the first part of (iii) can be ensured by reducing $\hat{\epsilon}$ toward zero, as needed. In particular, with respect to (i) and (ii), this can be done by continuity of c , and with respect to the first part of (iii) it follows since $\rho(\cdot, y_*, z_*)$ is continuous (the functions ϕ_i for all $i \in [q]$ in the proof of Lemma 4.2 are continuous) and $\rho(x_*, y_*, z_*) = 0$ by $(y_*, z_*) \in \mathcal{D}(x_*)$ and (3). The second part of (iii) and (iv), on the other hand, can be ensured by reducing $\bar{\epsilon}$ as needed since ϵ_ρ is a linear function of ϵ_{err} according to (25) and $\epsilon_c \leq \|\epsilon_{\text{err}}\|_\infty$.

Let us now derive upper bounds that hold for all such x and all such approximations (satisfying (i)–(iv)). The point (y_*, z_*) is feasible for (27) since $\|(y_*, z_*)\|_\infty \leq M - 1$, and both (x, y_x, z_x) and (x, y_*, z_*) satisfy the conditions of Lemma 4.2 by (i) and the box constraints in (27). Hence, by the former inequality in (26a), the optimality of (y_x, z_x) with respect to (27), and (iii), it follows that

$$\rho(x, y_x, z_x) - \epsilon_\rho \leq \tilde{\rho}(x, y_x, z_x) \leq \tilde{\rho}(x, y_*, z_*) \leq \rho(x, y_*, z_*) + \epsilon_\rho \leq \frac{3}{4}t_\tau. \quad (30)$$

In particular, $\rho(x, y_x, z_x) \leq t_\tau \leq \frac{1}{2}\tau \leq 1$. (This fact is used later in Corollary 4.1.) Moreover, by Lemma 4.1 and since $t \mapsto t + (qt)^{1/2}$ is increasing, (30) gives $\tilde{\rho}(x, y_x, z_x) \leq \tilde{\rho}(x, y_x, z_x) + q^{1/2}\tilde{\rho}(x, y_x, z_x)^{1/2} \leq t_\tau + (qt_\tau)^{1/2} \leq \frac{1}{2}\tau + \frac{1}{2}\tau = \tau$, so that, by the definition of τ , one obtains

$$(\beta\tilde{\rho}(x, y_x, z_x))^\sigma \leq (\beta\tau)^\sigma \leq 4^{-\sigma}\frac{1}{4}m < \frac{1}{4}m. \quad (31)$$

Let us now prove part (a) of the theorem. By (i), the box constraints in (27), and $\|(\hat{y}, \hat{z})\|_\infty \leq M \leq \bar{M}$ (proved earlier), Lemma 4.2 applies at both (x, y_x, z_x) and $(x, \hat{y}, \hat{z})$. The upper bound in part (a) follows from (8) and (26c) at (x, y_x, z_x) , namely, $\omega(x) \leq \psi(x, y_x, z_x) \leq \tilde{\rho}(x, y_x, z_x) + \epsilon_\psi$. For the lower bound, (26c) at $(x, \hat{y}, \hat{z})$ and the optimality of (y_x, z_x) for (27), for which $(\hat{y}, \hat{z})$ is feasible, give

$$\omega(x) = \psi(x, \hat{y}, \hat{z}) \geq \bar{M}^{-1}(\tilde{\rho}(x, \hat{y}, \hat{z}) - \epsilon_\rho) \geq \bar{M}^{-1}(\tilde{\rho}(x, y_x, z_x) - \epsilon_\rho),$$

and the last quantity is positive by the conditions of part (a) and the final part of Lemma 4.2 at (x, y_x, z_x) .

Finally, let us prove part (b). Let $i \notin \mathcal{A}(x_*)$. By (22), (ii), (iv), and (31), it follows that

$$\tilde{c}_i(x) \leq c_i(x) + \epsilon_c \leq \frac{1}{2}c_i(x_*) + \epsilon_c \leq -\frac{1}{2}m + \frac{1}{4}m = -\frac{1}{4}m < -(\beta\tilde{\rho}(x, y_x, z_x))^\sigma,$$

so $i \notin \tilde{\mathcal{A}}_{\text{LP}}(y_x, z_x; x)$ by (28). Now let $i \in \mathcal{A}(x_*)$. By (26a)–(26b), (29), and $\epsilon_\psi \leq (2C)^{-1}\|x - x_*\|_2$,

$$\tilde{\rho}(x, y_x, z_x) \geq \tilde{\psi}(x, y_x, z_x) \geq \psi(x, y_x, z_x) - \epsilon_\psi \geq C^{-1}\|x - x_*\|_2 - \epsilon_\psi \geq (2C)^{-1}\|x - x_*\|_2.$$

Moreover, $\epsilon_c \leq q^{-1}\epsilon_\psi \leq (2Cq)^{-1}\|x - x_*\|_2$ and $|c_i(x)| = |c_i(x) - c_i(x_*)| \leq L_c\|x - x_*\|_2$. Since $\|x - x_*\|_2 \leq \hat{\epsilon} \leq \bar{\delta}$, one has $\max\{L_c, (2Cq)^{-1}\}\|x - x_*\|_2^{1-\sigma} \leq \frac{1}{2}(\beta/(2C))^\sigma$, and therefore

$$|\tilde{c}_i(x)| \leq |c_i(x)| + \epsilon_c \leq (L_c + (2Cq)^{-1})\|x - x_*\|_2 \leq \left(\frac{\beta\|x - x_*\|_2}{2C} \right)^\sigma \leq (\beta\tilde{\rho}(x, y_x, z_x))^\sigma,$$

so $\tilde{c}_i(x) \geq -(\beta\tilde{\rho}(x, y_x, z_x))^\sigma$ and $i \in \tilde{\mathcal{A}}_{\text{LP}}(y_x, z_x; x)$ by (28). Overall, $\tilde{\mathcal{A}}_{\text{LP}}(y_x, z_x; x) = \mathcal{A}(x_*)$. $\square$ $\square$

The condition in part (a) of Theorem 4.1 involves (y_x, z_x) , meaning that it depends on the realized noise. On the other hand, the condition in part (b) depends on x only through $\|x - x_*\|_2$. The following corollary gives sufficient conditions of the latter form for both parts of the theorem.

Corollary 4.1. *Suppose that the conditions of Theorem 4.1 hold, which define C , $\hat{\epsilon}$, and $\bar{\epsilon}$. Then, for every $x \in \mathbb{R}^n$ with $\|x - x_*\|_2 \leq \hat{\epsilon}$ and any approximations satisfying (22) with $\|\epsilon_{\text{err}}\|_\infty \leq \bar{\epsilon}$, it follows that*

$$\bar{\rho}(x, y_x, z_x) \geq \frac{\|x - x_*\|_2}{C} \quad \text{and} \quad \rho(x, y_x, z_x) \geq \left(\frac{\|x - x_*\|_2}{C(1+q^{1/2})} \right)^2. \quad (32)$$

Also, there exists $(\Lambda', \Lambda) \in \mathbb{R}_{>0} \times \mathbb{R}_{>0}$, dependent only on $(C, \bar{M}, n, p, q, \hat{\epsilon}, \bar{\epsilon})$, such that $\Lambda' \hat{\epsilon} \leq \bar{\epsilon}$, $\Lambda \hat{\epsilon}^2 \leq \bar{\epsilon}$, and for each x with $0 < \|x - x_*\|_2 \leq \hat{\epsilon}$ and any approximations satisfying (22), Theorem 4.1(b) holds when

$$\|\epsilon_{\text{err}}\|_\infty \leq \Lambda' \|x - x_*\|_2, \quad (33)$$

and Theorem 4.1(a)–(b) hold when

$$\|\epsilon_{\text{err}}\|_\infty \leq \Lambda \|x - x_*\|_2^2. \quad (34)$$

Proof. Proof. The first bound in (32) follows from (29), which was established in the proof of Theorem 4.1. For the second bound in (32), recall from the same proof that $\rho(x, y_x, z_x) \leq 1$ for every x with $\|x - x_*\|_2 \leq \hat{\epsilon}$ and any approximations with $\|\epsilon_{\text{err}}\|_\infty \leq \bar{\epsilon}$. (See the comment after (30).) Hence, by [52, Lemma 3.5],

$$\bar{\rho}(x, y_x, z_x) \leq \rho(x, y_x, z_x) + q^{1/2} \rho(x, y_x, z_x)^{1/2} \leq (1 + q^{1/2}) \rho(x, y_x, z_x)^{1/2},$$

so $\rho(x, y_x, z_x) \geq \bar{\rho}(x, y_x, z_x)^2 / (1 + q^{1/2})^2$. Combining this with the first bound in (32) yields the second.

It remains to establish the existence of (Λ', Λ) . Consider arbitrary x with $0 < \|x - x_*\|_2 \leq \hat{\epsilon}$ and define

$$c_\psi := c_\kappa + q \quad \text{and} \quad c_\rho := c_\kappa + q\bar{M}, \quad \text{where} \quad c_\kappa := p^{1/2} + n^{1/2}(1 + p^{1/2}\bar{M} + q^{1/2}\bar{M}),$$

so that, by (25), $\epsilon_\psi \leq c_\psi \|\epsilon_{\text{err}}\|_\infty$ and $\epsilon_\rho \leq c_\rho \|\epsilon_{\text{err}}\|_\infty$. Let

$$\Lambda' := \min \left\{ \frac{1}{2C c_\psi}, \frac{\bar{\epsilon}}{\hat{\epsilon}} \right\} \quad \text{and} \quad \Lambda := \min \left\{ \frac{1}{4c_\rho C^2 (1+q^{1/2})^2}, \frac{\Lambda'}{\bar{\epsilon}} \right\},$$

so that $\Lambda' \hat{\epsilon} \leq \bar{\epsilon}$ and $\Lambda \hat{\epsilon}^2 \leq \bar{\epsilon}$. Suppose first that (33) holds. Then $\|\epsilon_{\text{err}}\|_\infty \leq \Lambda' \hat{\epsilon} \leq \bar{\epsilon}$, so the conditions of Theorem 4.1 hold and $\epsilon_\psi \leq c_\psi \|\epsilon_{\text{err}}\|_\infty \leq c_\psi \Lambda' \|x - x_*\|_2 \leq (2C)^{-1} \|x - x_*\|_2$. Hence Theorem 4.1(b) holds. Suppose now that $\|\epsilon_{\text{err}}\|_\infty \leq \Lambda \|x - x_*\|_2^2$. Then $\|\epsilon_{\text{err}}\|_\infty \leq (\Lambda' / \hat{\epsilon}) \|x - x_*\|_2^2 \leq \Lambda' \|x - x_*\|_2$ since $\|x - x_*\|_2 \leq \hat{\epsilon}$, so Theorem 4.1(b) holds as shown. Moreover, the bounds in (32) hold since $\|\epsilon_{\text{err}}\|_\infty \leq \bar{\epsilon}$, so

$$\epsilon_\rho \leq c_\rho \|\epsilon_{\text{err}}\|_\infty \leq c_\rho \Lambda \|x - x_*\|_2^2 \leq \frac{\|x - x_*\|_2^2}{4C^2(1+q^{1/2})^2} \leq \frac{1}{4} \rho(x, y_x, z_x) < \frac{1}{2} \rho(x, y_x, z_x),$$

which is the condition in Theorem 4.1(a). Hence, Theorem 4.1(a)–(b) hold, as claimed. $\square$ $\square$

Corollary 4.1 shows that correct identification requires the noise to vanish only linearly in $\|x - x_*\|_2$, whereas the two-sided bound on $\omega(x)$ in Theorem 4.1(a) requires a quadratic rate. The latter is due to the square-root terms in $\bar{\rho}$, which make ρ of order $\bar{\rho}^2$ near x_* . In upcoming Theorem 4.3(e), it is shown that the linear rate in (33) cannot be improved in general. That said, we do not know whether the quadratic rate for ensuring that part (a) holds is tight. Since (33) depends on x only through $\|x - x_*\|_2$, it is the form of the noise requirement that we use for the setting of a stochastic algorithm in §4.4.

4.2 Identification through a Primal Step

We now present a technique for active-set identification in noisy settings that is based on computing a primal step through solving either an LP or a QP. The QP approach is an extension of the one that offers Theorem 3.3, and as previously stated a similar result holds when one solves an LP with a trust-region

constraint. Our primary tool for extending these approaches to noisy settings is perturbation theory for the solution of linear systems, in particular, Lemmas 2.1–2.2.

For the sake of brevity in this section, we state only the QP-based approach in detail. All of the problems statements and results hold similarly for an LP-based approach with a trust-region constraint.

Analogously to (16), consider the QP subproblem

$$\begin{aligned} \min_{(d,r,s,t) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^p \times \mathbb{R}^q} \quad & \widetilde{\nabla} f(x)^T d + \nu(\mathbf{1}^T r + \mathbf{1}^T s + \mathbf{1}^T t) + \frac{1}{2}\theta \|d\|_2^2 \\ \text{s.t.} \quad & \tilde{e}(x) + \widetilde{\nabla} e(x)^T d = r - s, \quad \tilde{c}(x) + \widetilde{\nabla} c(x)^T d \leq t, \\ & r \geq 0, \quad s \geq 0, \quad \text{and} \quad t \geq 0. \end{aligned} \quad (35)$$

Like (16), this QP is always feasible and has a finite optimal value. Later in our analysis we use a fact from exact penalty theory [39, Theorem 4.4], namely, that if (36) below is feasible and its (unique) solution admits multipliers (α, β) , then for every $\nu \geq \|(\alpha, \beta)\|_\infty$ the solution of (35) has $(r_{x,\theta,\nu}, s_{x,\theta,\nu}, t_{x,\theta,\nu}) = (0, 0, 0)$ and its d -component coincides with the (unique) solution $d_{x,\theta,\nu}$ of

$$\min_{d \in \mathbb{R}^n} \quad \widetilde{\nabla} f(x)^T d + \frac{1}{2}\theta \|d\|_2^2 \quad \text{s.t.} \quad \tilde{e}(x) + \widetilde{\nabla} e(x)^T d = 0 \quad \text{and} \quad \tilde{c}(x) + \widetilde{\nabla} c(x)^T d \leq 0; \quad (36)$$

that is, under these conditions the auxiliary variables in (35) all vanish at the solution.

Our first lemma is based on the observation that if the LICQ holds at a primal KKT point x_* of (2), a point x is sufficiently close to x_* , and the noise in certain constraint derivative values at x is sufficiently small, then the corresponding constraint derivative matrix has full rank.

Lemma 4.3. *Suppose that, with respect to (2), the LICQ holds at a primal KKT point $x_* \in \mathbb{R}^n$. Let the active set at x_* be $\mathcal{A}_* := \mathcal{A}(x_*)$ and let its cardinality be $\ell := |\mathcal{A}_*|$. Then, there exist $\epsilon_x \in \mathbb{R}_{>0}$, $\epsilon_{\nabla e} \in \mathbb{R}_{>0}$, and $\epsilon_{\nabla c} \in \mathbb{R}_{>0}$ such that any $(x, \widetilde{\nabla} e(x), \widetilde{\nabla} c_{\mathcal{A}_*}(x)) \in \mathbb{R}^n \times \mathbb{R}^{n \times p} \times \mathbb{R}^{n \times \ell}$ with*

$$\|x - x_*\|_2 \leq \epsilon_x, \quad \|\widetilde{\nabla} e(x) - \nabla e(x)\|_2 \leq \epsilon_{\nabla e}, \quad \text{and} \quad \|\widetilde{\nabla} c_{\mathcal{A}_*}(x) - \nabla c_{\mathcal{A}_*}(x)\|_2 \leq \epsilon_{\nabla c} \quad (37)$$

yields that the matrix $\begin{bmatrix} \widetilde{\nabla} e(x) & \widetilde{\nabla} c_{\mathcal{A}_*}(x) \end{bmatrix}$ has full column rank.

Proof. Proof. Under the conditions of the lemma, specifically that the LICQ for (2) holds at x_* , the matrix $\begin{bmatrix} \nabla e(x_*) & \nabla c_{\mathcal{A}_*}(x_*) \end{bmatrix}$ has full column rank. Then, since $\{e_i\}_{i \in [p]}$ and $\{c_i\}_{i \in \mathcal{A}_*}$ are continuously differentiable, it follows that $\begin{bmatrix} \widetilde{\nabla} e(x) & \widetilde{\nabla} c_{\mathcal{A}_*}(x) \end{bmatrix}$ is continuous as a function of x . Since the set of full-column-rank matrices is open, there exists $\epsilon_x \in \mathbb{R}_{>0}$ such that for all $x \in \mathbb{R}^n$ with $\|x - x_*\|_2 \leq \epsilon_x$, the matrix $\begin{bmatrix} \widetilde{\nabla} e(x) & \widetilde{\nabla} c_{\mathcal{A}_*}(x) \end{bmatrix}$ has full column rank. Consequently, for any such ϵ_x , there exist sufficiently small $\epsilon_{\nabla e} \in \mathbb{R}_{>0}$ and $\epsilon_{\nabla c} \in \mathbb{R}_{>0}$ such that $\begin{bmatrix} \widetilde{\nabla} e(x) & \widetilde{\nabla} c_{\mathcal{A}_*}(x) \end{bmatrix}$ has full column rank as well, which completes the proof. $\square$ $\square$

Lemma 4.3 is consistent with a well-known phenomenon, namely, that constraint qualifications are stable under small perturbations of the problem data. Such results go back at least to [57], and they appear, in far greater generality, as special cases of the stability theory for convex composite optimization (see, e.g., [14, Lemma 3.2]) and of the characterizations of strong regularity for variational inequalities over polyhedral convex sets in [30]. We include the short proof of Lemma 4.3 for the sake of self-containedness. We also note that the same argument applies to the MFCQ. That is, the MFCQ at x_* is equivalent to positive linear independence of $\{\nabla e_i(x_*)\}_{i \in [p]} \cup \{\nabla c_i(x_*)\}_{i \in \mathcal{A}_*}$, which is an open condition, so under the MFCQ the conclusion of Lemma 4.3 holds with “full column rank” replaced by “positive linear independence.”

Our next lemma shows that if x is close to a primal KKT point x_* and the noise is sufficiently small, then the QP (36) is feasible and its primal-dual solution is given by the solution of a linear system defined with respect to the active set at x_* . A similar result would hold for an LP subproblem with a trust-region constraint, since at a primal KKT point the solution of the subproblem would have no activity with respect to the trust region, and in a neighborhood of the primal KKT point the subproblem solutions would remain inactive with respect to the trust region. For the LP case, (38) would have a zero matrix in the $(1,1)$ -block.

Lemma 4.4. Suppose that, with respect to (2), the LICQ and strict complementarity hold at a primal KKT point $x_* \in \mathbb{R}^n$. Let $\mathcal{A}_* := \mathcal{A}(x_*)$ and $\ell := |\mathcal{A}_*|$. Then, for any $\theta \in \mathbb{R}_{>0}$, there exist $\epsilon_x \in \mathbb{R}_{>0}$, $\epsilon_e \in \mathbb{R}_{>0}$, $\epsilon_c \in \mathbb{R}_{>0}$, $\epsilon_{\nabla f} \in \mathbb{R}_{>0}$, $\epsilon_{\nabla e} \in \mathbb{R}_{>0}$, and $\epsilon_{\nabla c} \in \mathbb{R}_{>0}$ such that for all $(x, \tilde{e}(x), \tilde{c}(x), \tilde{\nabla}f(x), \tilde{\nabla}e(x), \tilde{\nabla}c(x))$ satisfying $\|x - x_*\|_2 \leq \epsilon_x$ and (22) the linear system

$$\begin{bmatrix} \theta I & \tilde{\nabla}e(x) & \tilde{\nabla}c_{\mathcal{A}_*}(x) \\ \tilde{\nabla}e(x)^T & 0 & 0 \\ \tilde{\nabla}c_{\mathcal{A}_*}(x)^T & 0 & 0 \end{bmatrix} \begin{bmatrix} d \\ \alpha \\ \beta_{\mathcal{A}_*} \end{bmatrix} = - \begin{bmatrix} \tilde{\nabla}f(x) \\ \tilde{e}(x) \\ \tilde{c}_{\mathcal{A}_*}(x) \end{bmatrix} \quad (38)$$

has a unique solution $(d, \alpha, \beta_{\mathcal{A}_*}) \in \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^\ell$ which, along with $\beta_{\mathcal{A}_*^c} = 0$, satisfies the KKT conditions of subproblem (36); in particular, one has

$$\tilde{\nabla}f(x) + \theta d + \tilde{\nabla}e(x)\alpha + \tilde{\nabla}c(x)\beta = 0, \quad (39a)$$

$$\tilde{e}(x) + \tilde{\nabla}e(x)^T d = 0, \quad \tilde{c}_{\mathcal{A}_*}(x) + \tilde{\nabla}c_{\mathcal{A}_*}(x)^T d = 0, \quad (39b)$$

$$\tilde{c}_{\mathcal{A}_*^c}(x) + \tilde{\nabla}c_{\mathcal{A}_*^c}(x)^T d < 0, \quad (39c)$$

$$\beta_{\mathcal{A}_*} > 0, \quad \beta_{\mathcal{A}_*^c} = 0, \quad (39d)$$

$$\text{and } \tilde{c}(x) + \tilde{\nabla}c(x)^T d \perp \beta. \quad (39e)$$

Proof. Proof. Under the stated conditions, the conditions of Lemma 4.3 hold, which implies that there exist $\epsilon_x \in \mathbb{R}_{>0}$, $\epsilon_{\nabla e} \in \mathbb{R}_{>0}$, and $\epsilon_{\nabla c} \in \mathbb{R}_{>0}$ such that the matrix in (38), call it $\tilde{K}$, is invertible and the linear system in (38) has a unique solution. Let such $(\epsilon_x, \epsilon_{\nabla e}, \epsilon_{\nabla c})$ be given. Now note that, by construction of (38) and with $\beta_{\mathcal{A}_*^c} = 0$, the conditions (39a), (39b), and (39e) are all satisfied. All that remains is to show that, when x is sufficiently close to x_* and all of the noise in the function and derivative values is sufficiently small, the conditions (39c) and (39d) hold as well. Toward this end, let us define

$$K := \begin{bmatrix} \theta I & \nabla e(x_*) & \nabla c_{\mathcal{A}_*}(x_*) \\ \nabla e(x_*)^T & 0 & 0 \\ \nabla c_{\mathcal{A}_*}(x_*)^T & 0 & 0 \end{bmatrix} \text{ and } \begin{bmatrix} d_* \\ \alpha_* \\ [\beta_{\mathcal{A}_*}]_* \end{bmatrix} = -K^{-1} \begin{bmatrix} \nabla f(x_*) \\ e(x_*) \\ c_{\mathcal{A}_*}(x_*) \end{bmatrix}.$$

Under the conditions of the lemma, it follows by Farkas' theorem that $d_* = 0$ and, along with $[\beta_{\mathcal{A}_*^c}]_* = 0$, the pair $(\alpha_*, \beta_*) =: (y_*, z_*)$ is the unique Lagrange multiplier such that (x_*, y_*, z_*) is a KKT point for problem (2). Indeed, by the strict complementarity assumption in the lemma, it follows that $[\beta_{\mathcal{A}_*}]_* > 0$. Now note that the matrix K is invertible (e.g., by setting $x = x_*$ in Lemma 4.3), and that one has $\tilde{K} = K + \delta K$, where—since f , e , and c are continuously differentiable—one finds that $\delta K \rightarrow 0$ as $x \rightarrow x_*$ and $\epsilon_{\text{err}} \rightarrow 0$, where ϵ_{err} is defined in (23). Consequently, $\|\delta K\|_2 \|K^{-1}\|_2 \rightarrow 0$ as $x \rightarrow x_*$ and $\epsilon_{\text{err}} \rightarrow 0$, which means that, for x sufficiently close to x_* and ϵ_{err} sufficiently small, one has from Lemma 2.1 that (38) yields

$$\frac{\left\| \begin{bmatrix} d_* - d \\ \alpha_* - \alpha \\ [\beta_{\mathcal{A}_*}]_* - \beta_{\mathcal{A}_*} \end{bmatrix} \right\|_2}{\left\| \begin{bmatrix} d_* \\ \alpha_* \\ [\beta_{\mathcal{A}_*}]_* \end{bmatrix} \right\|_2} \leq \frac{\text{cond}(K)}{1 - \|\delta K\|_2 \|K^{-1}\|_2} \left(\frac{\|\delta K\|_2}{\|K\|_2} + \frac{\left\| \begin{bmatrix} \tilde{\nabla}f(x) \\ \tilde{e}(x) \\ \tilde{c}_{\mathcal{A}_*}(x) \end{bmatrix} - \begin{bmatrix} \nabla f(x_*) \\ e(x_*) \\ c_{\mathcal{A}_*}(x_*) \end{bmatrix} \right\|_2}{\left\| \begin{bmatrix} \nabla f(x_*) \\ e(x_*) \\ c_{\mathcal{A}_*}(x_*) \end{bmatrix} \right\|_2} \right).$$

Overall, since $d_* = 0$ and $[\beta_{\mathcal{A}_*}]_* > 0$, it follows that for $\epsilon_x \in \mathbb{R}_{>0}$, $\epsilon_e \in \mathbb{R}_{>0}$, $\epsilon_c \in \mathbb{R}_{>0}$, $\epsilon_{\nabla f} \in \mathbb{R}_{>0}$, $\epsilon_{\nabla e} \in \mathbb{R}_{>0}$, and $\epsilon_{\nabla c} \in \mathbb{R}_{>0}$ all sufficiently small, any $(x, \tilde{e}(x), \tilde{c}(x), \tilde{\nabla}f(x), \tilde{\nabla}e(x), \tilde{\nabla}c(x))$ satisfying $\|x - x_*\|_2 \leq \epsilon_x$ and the error bounds (22) yield that all conditions in (39) hold, as desired. $\square$ $\square$

The tolerances asserted by Lemma 4.4 can be made explicit in terms of the conditioning of the matrix K defined in its proof. This is shown next. In what follows, using the notation of Lemma 4.4 for such a primal

KKT point x_* , let us define $A_* := [\nabla e(x_*) \quad \nabla c_{A_*}(x_*)] \in \mathbb{R}^{n \times (p+\ell)}$, $a_* := \|A_*\|_2$, $\sigma_* := \sigma_{\min}(A_*) > 0$ (where positivity follows under the LICQ), $s_* := \min_{i \notin \mathcal{A}_*} (-c_i(x_*)) > 0$, and $\zeta_* := \min_{i \in \mathcal{A}_*} [z_*]_i > 0$ (where positivity follows by strict complementarity). Furthermore, suppose that ∇f , ∇e , ∇c , e , and c are Lipschitz continuous over $\{x \in \mathbb{R}^n : \|x - x_*\|_2 \leq 1\}$ with $L \geq 1$ being a common Lipschitz constant for these functions over this set, and let $G := \max_{i \in [q]} \max_{x \in \mathbb{R}^n} \{\|\nabla c_i(x)\|_2 : \|x - x_*\|_2 \leq 1\}$.

Lemma 4.5. *Suppose that the conditions of Lemma 4.4 hold and consider arbitrary $\theta \in \mathbb{R}_{>0}$. Then, the matrix K defined in the proof of Lemma 4.4 satisfies*

$$\|K\|_2 = \frac{\theta + \sqrt{\theta^2 + 4a_*^2}}{2} \leq \theta + a_* \quad (40)$$

$$\text{and } \|K^{-1}\|_2 \leq \kappa_\theta := \max \left\{ \frac{1}{\theta}, \frac{\theta + \sqrt{\theta^2 + 4\sigma_*^2}}{2\sigma_*^2} \right\} \leq \frac{1}{\theta} + \frac{1}{\sigma_*} + \frac{\theta}{\sigma_*^2},$$

from which it follows that $\text{cond}(K) \leq (\theta + a_*)\kappa_\theta$. Moreover, defining $\delta_K := \sqrt{2}(L\epsilon_x + \epsilon_{\nabla e} + \epsilon_{\nabla c})$ and $\delta_b := \epsilon_{\nabla f} + \epsilon_e + \epsilon_c + 3L\epsilon_x$, it follows that if $\epsilon_x \leq 1$,

$$\kappa_\theta \delta_K \leq \frac{1}{2}, \quad \text{and } \Theta := 2\kappa_\theta (\delta_b + \delta_K \|(y_*, [z_*]_{A_*})\|_2) < \min \left\{ 1, \zeta_*, \frac{s_* - L\epsilon_x - \epsilon_{\nabla e} - \epsilon_c}{G} \right\}, \quad (41)$$

then the conclusions of Lemma 4.4 hold; i.e., (38) has a unique solution (d, α, β_{A_*}) and, with $\beta_{A_*^c} = 0$, the conditions (39) are all satisfied. In particular, (41) holds whenever ϵ_x and $\|\epsilon_{err}\|_\infty$ are bounded above by a constant determined by θ , σ_* , a_* , L , G , ζ_* , s_* , and $\|(y_*, z_*)\|_2$.

Proof. Proof. Let $r := p + \ell$, let $A_* = U\Sigma V^T$ be a (thin) singular value decomposition of A_* with $U \in \mathbb{R}^{n \times r}$, $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$, and orthogonal $V \in \mathbb{R}^{r \times r}$, and let $U_\perp \in \mathbb{R}^{n \times (n-r)}$ complete U to give an orthonormal basis of $\mathbb{R}^n$. With the orthogonal matrix $Q := \text{diag}([U \quad U_\perp], V)$ one finds

$$Q^T K Q = \begin{bmatrix} \theta I_r & 0 & \Sigma \\ 0 & \theta I_{n-r} & 0 \\ \Sigma & 0 & 0 \end{bmatrix} \quad \text{with identity matrices } I_r \in \mathbb{R}^{r \times r} \text{ and } I_{n-r} \in \mathbb{R}^{(n-r) \times (n-r)},$$

which, after a symmetric permutation, is the sum of the 2×2 blocks $\begin{bmatrix} \theta & \sigma_i \\ \sigma_i & 0 \end{bmatrix}$ for $i \in [r]$ and θI_{n-r} . Thus, the eigenvalues of K are $\frac{1}{2}(\theta \pm (\theta^2 + 4\sigma_i^2)^{1/2})$ for $i \in [r]$, together with θ with multiplicity $n - r$. Since K is symmetric, $\|K\|_2$ is the largest of these in absolute value, which is $\frac{1}{2}(\theta + (\theta^2 + 4a_*^2)^{1/2})$ because $a_* = \max_i \sigma_i$ and $\frac{1}{2}(\theta + (\theta^2 + 4\sigma^2)^{1/2}) \geq \theta$ for every $\sigma \geq 0$. The inequality $\frac{1}{2}(\theta + (\theta^2 + 4a_*^2)^{1/2}) \leq \theta + a_*$ in (40) follows from $(\theta^2 + 4a_*^2)^{1/2} \leq \theta + 2a_*$. On the other hand, $\|K^{-1}\|_2$ is the reciprocal of the smallest eigenvalue of K in absolute value. That smallest value is $\min\{\theta, \frac{1}{2}((\theta^2 + 4\sigma_*^2)^{1/2} - \theta)\}$, where the second entry is attained at $\sigma_* = \min_i \sigma_i$ since $\sigma \mapsto (\theta^2 + 4\sigma^2)^{1/2} - \theta$ is increasing. Overall, $2/((\theta^2 + 4\sigma_*^2)^{1/2} - \theta) = (\theta + (\theta^2 + 4\sigma_*^2)^{1/2})/(2\sigma_*^2)$, which gives κ_θ in (40) with equality when $r < n$ and with inequality when $r = n$, since in that case the eigenvalue θ is absent. The end of (40) follows from $(\theta^2 + 4\sigma_*^2)^{1/2} \leq \theta + 2\sigma_*$.

Now let $w := (d, \alpha, \beta_{A_*})$ solve (38) and let $w_* := (0, y_*, [z_*]_{A_*})$, which solves $Kw_* = -(\nabla f(x_*), e(x_*), c_{A_*}(x_*)) = -(\nabla f(x_*), 0, 0)$ as established in the proof of Lemma 4.4. The matrix $\tilde{K}$ of (38) satisfies $\tilde{K} = K + \delta K$, where δK has zero diagonal blocks and off-diagonal blocks E and E^T with $E := [\widetilde{\nabla e}(x) \quad \widetilde{\nabla c}_{A_*}(x)] - A_*$, so that $\|\delta K\|_2 = \|E\|_2$. Writing $E = [P \quad Q]$ with $P := \widetilde{\nabla e}(x) - \nabla e(x_*)$ and $Q := \widetilde{\nabla c}_{A_*}(x) - \nabla c_{A_*}(x_*)$, one has $\|E\|_2^2 = \|PP^T + QQ^T\|_2 \leq \|P\|_2^2 + \|Q\|_2^2$, while $\|P\|_2 \leq L\epsilon_x + \epsilon_{\nabla e}$ and $\|Q\|_2 \leq L\epsilon_x + \epsilon_{\nabla c}$ by (22) and the definition of L . Hence, $\|\delta K\|_2 \leq \sqrt{2}(L\epsilon_x + \epsilon_{\nabla e} + \epsilon_{\nabla c}) = \delta_K$. Similarly the right-hand sides differ by at most δ_b . By the first condition in (41) and Lemma 2.2, the matrix $\tilde{K}$ is nonsingular—so (38) has a unique solution—and $\|w - w_*\|_2 \leq \Theta$. Conditions (39a), (39b), and (39e) hold by construction, as in the proof of Lemma 4.4. For (39d), one has $[\beta_{A_*}]_i \geq [z_*]_i - \Theta \geq \zeta_* - \Theta > 0$ by (41). For (39c), let $i \notin \mathcal{A}_*$. Then, $\|d\|_2 = \|d - 0\|_2 \leq \Theta < 1$ yields

$$\tilde{c}_i(x) + \widetilde{\nabla c}_i(x)^T d \leq c_i(x_*) + L\epsilon_x + \epsilon_c + (G + \epsilon_{\nabla c})\Theta \leq -s_* + L\epsilon_x + \epsilon_{\nabla c} + \epsilon_c + G\Theta < 0$$

by the second condition in (41). The final claim follows since δ_K and δ_b are nondecreasing linear functions of the tolerances ϵ_x and $\|\epsilon_{\text{err}}\|_\infty$ that vanish at the origin. $\square$ $\square$

The quantity κ_θ in (40) is a Lipschitz constant of the solution map of the KKT system of (36) with the active set fixed at $\mathcal{A}_*$, and the simpler bound $\kappa_\theta \leq \theta^{-1} + \sigma_*^{-1} + \theta\sigma_*^{-2}$ shows its dependence on the problem data. It degrades as the active constraint gradients approach linear dependence (i.e., as $\sigma_* \searrow 0$). For fixed σ_* , the precise value for κ_θ in (40) is minimized at $\theta = \sigma_*/\sqrt{2}$, which yields $\kappa_\theta = \sqrt{2}/\sigma_*$. For such θ , or any θ proportional to σ_* , condition (41) states that the admissible noise level is proportional to $\sigma_* \min\{1, \zeta_*, s_*/G\}/(1 + \|(y_*, z_*)\|_2)$, whereas for fixed θ it scales like σ_*^2 as $\sigma_* \searrow 0$. Both of the terms σ_* and $\min\{\zeta_*, s_*/G\}$ in this noise threshold vanish in degenerate situations that are excluded by the conditions of the lemma, which shows the importance of nondegeneracy to reach its conclusion.

We are now prepared to prove our main theorem of this subsection, for which we refer to the active-set estimate defined similarly to (20) as

$$\tilde{\mathcal{A}}_{\text{QP}}(d; x) := \{i \in [q] : \tilde{c}_i(x) + \widetilde{\nabla c}_i(x)^T d \geq 0\}. \quad (42)$$

As claimed, a similar result would hold for the LP-based subproblem and the active-set estimate

$$\tilde{\mathcal{A}}_{\text{LP}}(d; x) := \{i \in [q] : \tilde{c}_i(x) + \widetilde{\nabla c}_i(x)^T d \geq 0\}. \quad (43)$$

Theorem 4.2. *Suppose that, with respect to (2), the LICQ and strict complementarity hold at a primal KKT point $x_* \in \mathbb{R}^n$. Let $\mathcal{A}_* := \mathcal{A}(x_*)$ and $\ell := |\mathcal{A}_*|$. Then, for any $\theta \in \mathbb{R}_{>0}$, there exist $\epsilon_x \in \mathbb{R}_{>0}$, $\epsilon_e \in \mathbb{R}_{>0}$, $\epsilon_c \in \mathbb{R}_{>0}$, $\epsilon_{\nabla f} \in \mathbb{R}_{>0}$, $\epsilon_{\nabla e} \in \mathbb{R}_{>0}$, $\epsilon_{\nabla c} \in \mathbb{R}_{>0}$, and $\underline{\nu} \in \mathbb{R}_{>0}$ such that for all $(x, \tilde{e}(x), \tilde{c}(x), \widetilde{\nabla f}(x), \widetilde{\nabla e}(x), \widetilde{\nabla c}(x), \nu)$ with $\|x - x_*\|_2 \leq \epsilon_x$, (22), and $\nu \geq \underline{\nu}$, the solution (d, r, s, t) of (35) is unique and the active-set estimate defined by (42) yields $\tilde{\mathcal{A}}_{\text{QP}}(d; x) = \mathcal{A}_*$.*

Proof. Under the stated conditions, it follows by Lemma 4.4 that for any $\theta \in \mathbb{R}_{>0}$ one has for x sufficiently close to x_* and sufficiently small errors that (36) is feasible and has a unique minimizer at which the LICQ holds. Thus, by [39, Theorem 4.4], it follows that for any $\theta \in \mathbb{R}_{>0}$ there exist $(\epsilon_x, \epsilon_e, \epsilon_c, \epsilon_{\nabla f}, \epsilon_{\nabla e}, \epsilon_{\nabla c})$ as stated such that with any $\nu \geq \underline{\nu} := M$ (where M is defined as in (10)) the solution (d, r, s, t) of (35) has $(r, s, t) = 0$, which in turn means by Lemma 4.4 that d is the unique solution of (36). All that remains is to observe that by Lemma 4.4 one can conclude that for x sufficiently close to x_* , sufficiently small noise, and ν sufficiently large, one also has that (42) yields $\mathcal{A}_*$. In particular, this fact follows by the latter equation in (39b) along with (39c). Note that the threshold may be taken to be $\underline{\nu} := \|(y_*, z_*)\|_\infty + 1 \leq M$, since by Lemma 4.5 the multipliers of (36) satisfy $\|(\alpha, \beta) - (y_*, z_*)\|_2 \leq \Theta < 1$ under (41). $\square$ $\square$

Theorem 4.2 requires neither optimality of x_* nor any second-order condition, in contrast with Theorem 4.1. The reason is that (35) is strongly convex for every $\theta \in \mathbb{R}_{>0}$, so its solution is characterized by the nonsingular linear system (38) irrespective of the Hessian of the Lagrangian. The second-order condition (4) enters Theorem 4.1 only through the error bound of Theorem 3.1, which has no analogue here. Thus, Theorem 4.2 identifies the active set of whichever primal KKT point the iterates approach.

4.3 Effect of Degeneracy

Theorems 4.1 and 4.2 differ in a critical manner. Theorem 4.2, which assumes the LICQ and strict complementarity at x_* , guarantees correct identification for all x in a neighborhood of x_* under a *fixed* bound on the noise. Theorem 4.1, on the other hand, which assumes only the MFCQ and the second-order condition (4), and therefore permits weakly active constraints, requires that the noise vanishes linearly as $x \rightarrow x_*$; see Corollary 4.1. It is natural to ask whether the latter requirement might be an artifact of the LP-LPEC strategy or of our analysis. In fact, it is neither. Our next theorem demonstrates two important things for a simple instance with a weakly active constraint. First, no identification rule can be correct under a fixed noise level uniformly over all instances, since an instance can be paired with a second one whose data are indistinguishable at that noise level, but whose active set differs. Second, even for that single instance, any

reasonable rule can fail at a point x when the error in the objective gradient may be as large as $2\|x - x_*\|_2$, so the linear rate in Corollary 4.1 cannot be improved in general.

To state this precisely, let us call a map $\mathcal{R}$ an *identification rule* if it assigns to each $x \in \mathbb{R}^n$ and each collection of approximate values and derivatives at a finite number of points that is admissible in the sense of (22) a (possibly random) subset of $[q]$, using no information about the problem other than those values. We call an identification rule *pointwise* if its output at x depends only on the approximate values and derivatives at x , and *consistent* if, whenever its input consists of the exact values and derivatives at a local minimizer $\bar{x}$ of an instance of (2) at which the MFCQ and (4) hold, it returns $\mathcal{A}(\bar{x})$ (almost surely, if it is randomized). Consistency is a minimal requirement, since a rule that fails it can err even with exact data at the minimizer itself. All of our aforementioned rules are pointwise, and both the LP-LPEC rule (28) and the QP rule (42) return the correct active set for the instances used in part (e) below.

Theorem 4.3. *Consider two instances of (2) with $n = 2$, $p = 0$, and $q = 1$ defined by*

$$(P) : \begin{cases} f(x) = (x_1 - 1)^2 + x_2^2 \\ c(x) = x_1 - 1 \end{cases} \quad \text{and} \quad (\hat{P}) : \begin{cases} \hat{f}(x) = f(x) \\ \hat{c}(x) = x_1 - 1 - \epsilon_c \end{cases}$$

for some constraint evaluation error bound $\epsilon_c \in \mathbb{R}_{>0}$. Then, the following hold.

- (a) *Both instances have the same unique global minimizer $x_* = (1, 0)$ at which the LICQ, the MFCQ, and the second-order condition (4) hold, and both have a unique Lagrange multiplier equal to zero.*
- (b) *$\mathcal{A}(x_*) = \{1\}$ for (P), with the constraint being weakly active, and $\hat{\mathcal{A}}(x_*) = \emptyset$ for $(\hat{P})$.*
- (c) *At every $x \in \mathbb{R}^n$, the objective and all first-order derivatives of the two instances coincide, and every constraint value in $[c(x) - \epsilon_c, c(x)]$ is admissible under (22) for both (P) and $(\hat{P})$. For example, $\tilde{c}(x) := x_1 - 1 - \frac{1}{2}\epsilon_c$ is admissible for both problems for all $x \in \mathbb{R}^n$.*
- (d) *No identification rule is correct for both instances. In particular, for every rule $\mathcal{R}$ there is an instance and an admissible realization of the data for which $\mathcal{R}$ does not return the correct active set at x_* . If, moreover, a rule $\mathcal{R}$ uses a single evaluation of the constraint value, whose noise is uniform on $[-\epsilon_c, \epsilon_c]$, and the noise in all other quantities is zero, then for every rule $\mathcal{R}$, one has*

$$\mathbb{P}_{(P)}[\mathcal{R} \neq \mathcal{A}(x_*)] + \mathbb{P}_{(\hat{P})}[\mathcal{R} \neq \hat{\mathcal{A}}(x_*)] \geq \frac{1}{2},$$

so that $\mathcal{R}$ errs with probability at least $\frac{1}{4}$ on one of the two instances at every x .

- (e) *Let $\mathcal{R}$ be a pointwise and consistent identification rule. Then, for every $x \in \mathbb{R}^2$ with $x \neq x_*$ and $c(x) \leq 0$, and for any error bounds satisfying $\epsilon_{\nabla f} \geq 2\|x - x_*\|_2$, there is a realization of the approximations satisfying (22) for (P) under which $\mathcal{R}$ applied at x does not return $\mathcal{A}(x_*)$. Consequently, any guarantee of correct identification by such a rule at the feasible points near x_* requires that $\epsilon_{\nabla f} < 2\|x - x_*\|_2$, i.e., that the noise vanish at least linearly as $x \rightarrow x_*$.*

Proof. Proof. Consider part (a). For both instances, $x_* = (1, 0)$ is the unconstrained global minimizer of f and is feasible, so it is the unique global minimizer of both instances. Furthermore, $\nabla c(x_*) = \nabla \hat{c}(x_*) = (1, 0)$, so the LICQ and MFCQ hold in both instances, and for both the unique Lagrange multiplier is $z_* = 0$ since $\nabla f(x_*) = 0$. Finally, $\nabla_{xx}^2 \mathcal{L} = 2I$ is positive definite for both instances, so (4) holds for both.

Now consider part (b). The active sets are immediate from $c(x_*) = 0$ and $\hat{c}(x_*) = -\epsilon_c$. Since $c(x_*) = 0$ and $z_* = 0$, the constraint of (P) is weakly active, i.e., $\mathcal{A}_0(x_*) = \mathcal{A}(x_*) = \{1\}$.

Next, consider part (c). The two instances differ only by the constant ϵ_c in the constraint value, so all first-order derivatives agree and, using the example, $|\tilde{c}(x) - c(x)| = |\tilde{c}(x) - \hat{c}(x)| = \frac{1}{2}\epsilon_c \leq \epsilon_c$.

Next, consider part (d). The first claim follows from (b) and (c). For example, a rule sees the same input under both instances when the example in (c) occurs, and hence returns the same output, which cannot equal both $\{1\}$ and $\emptyset$. For the second claim, let μ and $\hat{\mu}$ denote the laws of the observed constraint value at

an arbitrary point x under instances (P) and $(\hat{P})$, respectively; namely, these are the uniform distributions on $[c(x) - \epsilon_c, c(x) + \epsilon_c]$ and on $[c(x) - 2\epsilon_c, c(x)]$, which have variational distance between them equal to $\frac{1}{2}$. For any (possibly randomized) rule $\mathcal{R}$, writing $\mathcal{E} := \{\mathcal{R} = \{1\}\}$ and $\mathcal{E}^c := \{\mathcal{R} = \emptyset\}$, one finds

$$\mathbb{P}_{(P)}[\mathcal{R} \neq \{1\}] + \mathbb{P}_{(\hat{P})}[\mathcal{R} \neq \emptyset] = \mathbb{P}_{(P)}[\mathcal{E}^c] + \mathbb{P}_{(\hat{P})}[\mathcal{E}] = 1 - (\mathbb{P}_{(P)}[\mathcal{E}] - \mathbb{P}_{(\hat{P})}[\mathcal{E}]) \geq 1 - \text{TV}(\mu, \hat{\mu}) = \frac{1}{2},$$

where TV denotes variational distance (or total variation).

Finally, consider part (e). Let $x \in \mathbb{R}^2$ satisfy $x \neq x_*$ and $c(x) \leq 0$, let $\eta \in (0, \epsilon_c]$, and consider the instance (P_x) of (2) defined by $f_x(w) := \|w - x\|_2^2 + f(x)$ and $c_x(w) := c(w) - \eta$. Since $c_x(x) = c(x) - \eta < 0$, the point x is feasible for (P_x) and is the unconstrained global minimizer of f_x , so x is the unique global minimizer of (P_x) , with unique Lagrange multiplier equal to zero and active set $\mathcal{A}_x(x) = \emptyset$. Moreover, the LICQ, the MFCQ, and (4) hold at x for (P_x) , since $\nabla c_x(x) = (1, 0) \neq 0$ and $\nabla^2 f_x = 2I$. Now consider approximations for (P) that are exact at every point other than x and that at x are given by

$$\begin{aligned} \tilde{f}(x) &:= f(x) = f_x(x), & \tilde{c}(x) &:= c(x) - \eta = c_x(x), \\ \widetilde{\nabla f}(x) &:= 0 = \nabla f_x(x), & \text{and } \widetilde{\nabla c}(x) &:= \nabla c(x) = \nabla c_x(x). \end{aligned}$$

At x , the errors are $|\tilde{c}(x) - c(x)| = \eta \leq \epsilon_c$ and $\|\widetilde{\nabla f}(x) - \nabla f(x)\|_2 = \|2(x - x_*)\|_2 = 2\|x - x_*\|_2 \leq \epsilon_{\nabla f}$, with all other errors equal to zero, so these approximations satisfy (22). Since $\mathcal{R}$ is pointwise, its output at x depends only on these values, which coincide with the exact values and derivatives of (P_x) at its minimizer x . Since $\mathcal{R}$ is consistent, it returns $\mathcal{A}_x(x) = \emptyset$, whereas $\mathcal{A}(x_*) = \{1\}$ for (P). For our proposed strategies, this can also be checked directly. For the LP-LPEC rule, the objective of (27) with these values is $|z| + (\eta - c(x))z$, which is minimized over $z \in [0, M]$ at $z = 0$ with value zero, so the solution is $z_x = 0$, $\tilde{\rho}(x, y_x, z_x) = 0$, and (28) excludes the constraint since $\tilde{c}(x) < 0$. For the QP rule, since $\widetilde{\nabla f}(x) = 0$ and $\tilde{c}(x) < 0$, the unique solution of (35) has $d = 0$ and $t = 0$, so (42) again excludes the constraint. $\square$ $\square$

The probabilistic statement in Theorem 4.3(d) concerns a single noisy evaluation. If, instead, $\mathcal{R}$ uses m independent evaluations of the constraint value, then the variational distance between the corresponding product laws is $1 - 2^{-m}$ and the same argument gives the lower bound 2^{-m} in place of $\frac{1}{2}$.

The construction in parts (a)–(d) is not specific to the example in Theorem 4.3. More generally, suppose that x_* is a local minimizer at which the MFCQ and (4) hold and that $\mathcal{A}_0(x_*) = \{i_0\}$ for some index $i_0 \in [q]$. Replacing the constraint function c_{i_0} by $c_{i_0} - \delta$ for any $\delta \in (0, \epsilon_c]$ yields a second problem instance for which x_* is again a KKT point with the same set of optimal Lagrange multipliers and with active set $\mathcal{A}(x_*) \setminus \{i_0\}$. Moreover, its critical cone at x_* is $\{d : \nabla e(x_*)^T d = 0, \nabla c_i(x_*)^T d = 0 \text{ for all } i \in \mathcal{A}(x_*) \setminus \{i_0\}\} = \mathcal{C}(x_*) \cup (-\mathcal{C}(x_*))$, with respect to which (4) still holds. Hence, x_* is a strict local minimizer of the second instance whose data are indistinguishable from those of the first instance at the noise level ϵ_c , even though its active set at x_* is strictly smaller than that of the first instance.

Parts (a)–(d) of Theorem 4.3 thus show that correct identification under a fixed noise level, uniformly over instances, requires a condition ruling out weakly active constraints, such as strict complementarity. Part (e) shows more. In particular, it shows that even for a single instance with a weakly active constraint, the noise must vanish at least linearly in $\|x - x_*\|_2$ as $x \rightarrow x_*$. For the instance (P), Corollary 4.1 ensures correct identification by the LP-LPEC rule whenever $\|\epsilon_{\text{err}}\|_\infty \leq \Lambda' \|x - x_*\|_2$, where $\Lambda' \leq \frac{1}{2}$ by the proof of that corollary (since $C \geq 1$ and $c_\psi \geq 1$), whereas part (e) of Theorem 4.3 shows failure whenever $\epsilon_{\nabla f} \geq 2\|x - x_*\|_2$. Hence, for (P), the noise vanishing at a linear rate is both sufficient and necessary for correct identification, so the rate in (33) cannot be improved. We note that part (e) concerns a worst-case realization of the noise in which the error in the objective gradient cancels $\nabla f(x)$ exactly. Under random noise, such a cancellation is unlikely, and correct identification may then be observed at larger noise levels. Viewed from the perspective of perturbation theory, one can say that in noisy settings the active set is not a stable feature of the problem data at a degenerate local minimizer.

The two strategies analyzed in this paper demonstrate this fact. The LP-LPEC strategy of §4.1 has been analyzed for the degenerate case, and it pays for this generality with the requirement of Corollary 4.1 that the noise vanish linearly, which by Theorem 4.3(e) cannot be avoided, or improved in rate, in general.

The primal-step strategies of §4.2, on the other hand, assume the LICQ and strict complementarity and, consequently, can tolerate a fixed noise level. Overall, the difference between Theorems 4.1 and 4.2 should be read as a difference between regularity assumptions rather than between active-set identification schemes.

4.4 Active-Set Identification within a Stochastic Algorithm

Let us now consider the use of active-set identification procedures within the context of a stochastic algorithm for solving (2). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space that captures all outcomes of a run of the algorithm. That is, each run of the algorithm is associated with an outcome $\omega \in \Omega$, an infinite-dimensional tuple whose k th element determines the outcome of any random quantity in iteration $k \in \mathbb{N}$. In this manner, the algorithm defines a stochastic process $\{X_k\}$, where for all $k \in \mathbb{N}$ the random variable X_k is the primal iterate. Given the initial condition that $X_1(\omega) = x_1$ for some $x_1 \in \mathbb{R}^n$ for all $\omega \in \Omega$, let $\mathcal{F}_k$ denote the sub- σ -algebra of $\mathcal{F}$ corresponding to the initial condition and all of the stochastic quantities prior to iteration k . In this manner $\{\mathcal{F}_k\}$ is a filtration.

The conditions of our main results in this section rely on assumptions that say: “conditioned on the behavior of an algorithm up to iteration $k \in \mathbb{N}$, the error bounds in (22) hold at the current iterate with probability at least p .” Thus, before stating those results, let us present the following lemma that shows that such conditions are achievable with sufficiently large sample sizes. Recall that (22) is a condition on a realization, meaning that nothing in §4.1 or §4.2 requires the errors to be bounded almost surely.

Lemma 4.6. *Consider arbitrary $k \in \mathbb{N}$ and suppose that a sample size $b_k \in \mathbb{N}$ and the error bounds in (22) are $\mathcal{F}_k$ -measurable. Suppose also that each of the six approximations in (22) with $x = X_k$ is computed as an average of b_k estimators that, conditioned on $\mathcal{F}_k$, are independent and unbiased, and whose conditional mean-squared errors—measured in the ℓ_2 -norm for vectors and in the Frobenius norm for matrices—are bounded by γ_f^2 , γ_e^2 , γ_c^2 , $\gamma_{\nabla f}^2$, $\gamma_{\nabla e}^2$, and $\gamma_{\nabla c}^2$, respectively. Then,*

$$\mathbb{P}[(22) \text{ holds with } x = X_k \mid \mathcal{F}_k] \geq 1 - \frac{S(\epsilon_{\text{err}})}{b_k},$$

$$\text{where } S(\epsilon_{\text{err}}) := \frac{\gamma_f^2}{\epsilon_f^2} + \frac{\gamma_e^2}{\epsilon_e^2} + \frac{\gamma_c^2}{\epsilon_c^2} + \frac{\gamma_{\nabla f}^2}{\epsilon_{\nabla f}^2} + \frac{\gamma_{\nabla e}^2}{\epsilon_{\nabla e}^2} + \frac{\gamma_{\nabla c}^2}{\epsilon_{\nabla c}^2}.$$

Thus, (22) holds with $x = X_k$ with conditional probability at least $p \in (0, 1)$ whenever $b_k \geq S(\epsilon_{\text{err}})/(1 - p)$.

Proof. Proof. Consider arbitrary $k \in \mathbb{N}$ and, without loss of generality, consider the approximation $\widetilde{\nabla}c(X_k)$. (The other five approximations can be handled identically.) Let $\widetilde{\nabla}c(X_k)$ be the average of b_k estimators, say, $\{G_1, \dots, G_{b_k}\}$ satisfying the stated conditions. Writing $\langle \cdot, \cdot \rangle_F$ for the Frobenius inner product, conditional independence and unbiasedness give $\mathbb{E}[\langle G_i - \nabla c(X_k), G_j - \nabla c(X_k) \rangle_F \mid \mathcal{F}_k] = 0$ for $i \neq j$ so that

$$\mathbb{E}[\|\widetilde{\nabla}c(X_k) - \nabla c(X_k)\|_F^2 \mid \mathcal{F}_k] = \frac{1}{b_k^2} \sum_{j=1}^{b_k} \mathbb{E}[\|G_j - \nabla c(X_k)\|_F^2 \mid \mathcal{F}_k] \leq \frac{\gamma_{\nabla c}^2}{b_k}.$$

Since for any matrix its spectral norm is bounded above by its Frobenius norm, it follows from the conditional Markov inequality applied to $\|\widetilde{\nabla}c(X_k) - \nabla c(X_k)\|_F^2$ that

$$\mathbb{P}[\|\widetilde{\nabla}c(X_k) - \nabla c(X_k)\|_2 > \epsilon_{\nabla c} \mid \mathcal{F}_k] \leq \mathbb{P}[\|\widetilde{\nabla}c(X_k) - \nabla c(X_k)\|_F^2 \geq \epsilon_{\nabla c}^2 \mid \mathcal{F}_k] \leq \frac{\gamma_{\nabla c}^2}{b_k \epsilon_{\nabla c}^2},$$

where the last inequality uses that $\epsilon_{\nabla c}$ is $\mathcal{F}_k$ -measurable. Summing the six analogous bounds (a union bound) gives the claim. Finally, $1 - S(\epsilon_{\text{err}})/b_k \geq p$ if and only if $b_k \geq S(\epsilon_{\text{err}})/(1 - p)$. $\square$ $\square$

The form of the conditions in Lemma 4.6, namely, that a deterministic accuracy condition holds with probability at least p conditioned on the past, is the *probabilistically accurate model* form that has appeared throughout the literature on stochastic optimization with random models; see, e.g., [22, 6, 55, 2].

Lemma 4.6 suggests how the two settings of §4.1 and §4.2 can be achieved. Under the LICQ and strict complementarity, one should expect a fixed sample size for each iteration to be sufficient; see Theorem 4.5

below. On the other hand, under only the MFCQ and (4), it has been shown in Corollary 4.1 that it suffices for the error bounds at iteration k to satisfy $\|\epsilon_{\text{err},k}\|_\infty \leq \Lambda' \|X_k - x_*\|_2$. In that case $S(\epsilon_{\text{err},k})$ should grow like $\|X_k - x_*\|_2^{-2}$, and by Lemma 4.6 a sample size $b_k \geq S(\epsilon_{\text{err},k})/(1-p)$ growing at that rate suffices. This is the adaptive-sampling regime studied in [17, 7]. Variance-reduced estimators, such as those of SVRG [43] and SARAH [50], achieve a comparable effect at lower cost, and in the deterministic-noise setting the analogue is a simulation or rule whose precision can be increased successively [23]. For instance, if f is an average of functions whose gradients are Lipschitz continuous with constant L_f , then the SVRG estimator of $\nabla f(X_k)$, formed with an $\mathcal{F}_k$ -measurable reference point $\tilde{X}_k$, is conditionally unbiased with conditional mean-squared error at most $L_f^2 \|X_k - \tilde{X}_k\|_2^2$, so Lemma 4.6 applies with $\gamma_{\nabla f}^2 := L_f^2 \|X_k - \tilde{X}_k\|_2^2$. Consequently, if the reference point is updated often enough that $\|\tilde{X}_k - x_*\|_2 \leq r \|X_k - x_*\|_2$ for some $r \in \mathbb{R}_{>0}$, then the bound on $\epsilon_{\nabla f}$ required by Theorem 4.4 can be met with a *fixed* sample size. (The SARAH estimator, conditioned on $\mathcal{F}_k$, is biased, so Lemma 4.6 does not apply.) A method that uses plain sample averaging with a fixed sample size throughout does not satisfy the hypothesis of Theorem 4.4, and, by Theorem 4.3, no method using data with a fixed noise level can succeed uniformly when x_* has a weakly active constraint.

Let us now state specific theorems that correspond to the settings of §4.1 and §4.2, respectively. Our first setting, for identification through Lagrange multipliers presented in §4.1, requires that the MFCQ and the second-order condition (4) hold at a local minimizer x_* . For all $k \in \mathbb{N}$, the LP (27) is solved with $x = X_k$ to yield a Lagrange multiplier pair (Y_k, Z_k) . The conditions of Theorem 4.1 involve (Y_k, Z_k) , and hence the noise realized at iteration k , so they cannot serve directly as a hypothesis on the error bounds. Fortunately, Corollary 4.1 removes this dependence by giving a sufficient bound in terms of $\|X_k - x_*\|_2$ alone. The following theorem therefore assumes that, at each iterate, the error bounds are at most proportional to $\|X_k - x_*\|_2$ and that (22) holds with these bounds with probability at least p , and it concludes that, with probability one, beyond some iteration the optimal active set is identified correctly with conditional probability at least p . This requires the noise in the function and derivative estimates to be reduced as the iterates converge at the rate stated previously. By Theorem 4.3(e), a requirement of this kind, at this rate, cannot be avoided in general when x_* has a weakly active constraint.

Theorem 4.4. *Suppose that, with respect to (2), the conditions of Theorem 4.1 hold, consider any $\beta \in \mathbb{R}_{>0}$ and $\sigma \in (0, 1)$, and let $(\hat{\epsilon}, \Lambda')$ be defined as in Theorem 4.1 and Corollary 4.1. Suppose also that, almost surely, $X_k \neq x_*$ for all $k \in \mathbb{N}$. Suppose further that, for all $k \in \mathbb{N}$, the function and derivative approximations at X_k are computed subject to positive error bounds $\epsilon_{\text{err},k}$ that are $\mathcal{F}_k$ -measurable with*

$$\|\epsilon_{\text{err},k}\|_\infty \leq \Lambda' \|X_k - x_*\|_2 \quad \text{whenever} \quad \|X_k - x_*\|_2 \leq \hat{\epsilon}, \quad (44)$$

and that, for some universal constant $p \in (0, 1]$, one has, almost surely and for all $k \in \mathbb{N}$,

$$\mathbb{P}[(22) \text{ holds with } x = X_k \text{ and } \epsilon_{\text{err}} = \epsilon_{\text{err},k} \mid \mathcal{F}_k] \geq p \cdot \mathbf{1}_{\{\|X_k - x_*\|_2 \leq \hat{\epsilon}\}}, \quad (45)$$

which is well defined since X_k is $\mathcal{F}_k$ -measurable. Let (Y_k, Z_k) denote a solution of (27) with $x = X_k$, computed from these approximations. Then, if $\{X_k\}$ converges almost surely to x_* , it holds with probability one that there exists $k_0 \in \mathbb{N}$ such that

$$\mathbb{P}[\tilde{\mathcal{A}}_{\text{LP}}(Y_k, Z_k; X_k) = \mathcal{A}_* \mid \mathcal{F}_k] \geq p \quad \text{for all } k \in \mathbb{N} \text{ with } k \geq k_0.$$

Moreover, with probability one, $\tilde{\mathcal{A}}_{\text{LP}}(Y_k, Z_k; X_k) = \mathcal{A}_*$ for infinitely many $k \in \mathbb{N}$.

Proof. Proof. Let Ω_0 be the event, with probability one, on which $X_k \rightarrow x_*$ and $X_k \neq x_*$ for all $k \in \mathbb{N}$, and for $\tilde{\omega} \in \Omega_0$ let $k_0(\tilde{\omega}) := \min\{k \in \mathbb{N} : \|X_j(\tilde{\omega}) - x_*\|_2 \leq \hat{\epsilon} \text{ for all } j \geq k\}$, which is finite. For $k \in \mathbb{N}$, let $\mathcal{E}_k$ denote the event in (45), let $\mathcal{B}_k := \{0 < \|X_k - x_*\|_2 \leq \hat{\epsilon}\}$, which lies in $\mathcal{F}_k$, and let $\mathcal{A}_k^{\text{id}}$ denote the event that $\tilde{\mathcal{A}}_{\text{LP}}(Y_k, Z_k; X_k) = \mathcal{A}_*$. On $\mathcal{E}_k \cap \mathcal{B}_k$, the approximations at X_k satisfy (22) with error bounds satisfying $\|\epsilon_{\text{err},k}\|_\infty \leq \Lambda' \|X_k - x_*\|_2$ by (44), so Corollary 4.1, via (33), applied with $x = X_k$ yields $\mathcal{E}_k \cap \mathcal{B}_k \subseteq \mathcal{A}_k^{\text{id}}$. Since $\mathcal{B}_k \in \mathcal{F}_k$, it follows from (45) that, almost surely, $\mathbb{P}[\mathcal{A}_k^{\text{id}} \mid \mathcal{F}_k] \geq \mathbb{P}[\mathcal{E}_k \cap \mathcal{B}_k \mid \mathcal{F}_k] = \mathbf{1}_{\mathcal{B}_k} \mathbb{P}[\mathcal{E}_k \mid \mathcal{F}_k] \geq p \cdot \mathbf{1}_{\mathcal{B}_k}$. On Ω_0 one has $\mathbf{1}_{\mathcal{B}_k} = 1$ for all $k \geq k_0$, which gives the first conclusion, and also $\sum_{k \in \mathbb{N}} \mathbb{P}[\mathcal{A}_k^{\text{id}} \mid \mathcal{F}_k] \geq$

$\sum_{k \geq k_0} p = \infty$. The samples used to form the approximations at iteration k are among the stochastic quantities determined prior to iteration $k + 1$, so $\mathcal{A}_k^{\text{id}} \in \mathcal{F}_{k+1}$, and Lévy's extension of the Borel–Cantelli lemma [61, §12.15] states that the events $\{\mathcal{A}_k^{\text{id}} \text{ infinitely often}\}$ and $\{\sum_k \mathbb{P}[\mathcal{A}_k^{\text{id}} \mid \mathcal{F}_k] = \infty\}$ are almost surely equivalent. Hence $\mathcal{A}_k^{\text{id}}$ occurs infinitely often almost surely on Ω_0 , and $\mathbb{P}[\Omega_0] = 1$. $\square$ $\square$

The quantity $\|X_k - x_*\|_2$ in (44) is unknown in practice, so the error bounds there cannot be computed directly. They are, however, equivalent up to constants to bounds in terms of the KKT error. Indeed, for any x with $\|x - x_*\|_2 \leq \hat{\epsilon}$, one has $\omega(x) \geq C^{-1}\|x - x_*\|_2$ as shown in the proof of Theorem 4.1, while for any $(y_*, z_*) \in \mathcal{D}(x_*)$, Theorem 3.1 gives $\omega(x) \leq \psi(x, y_*, z_*) \leq C \text{dist}(\{(x, y_*, z_*)\}, \{x_*\} \times \mathcal{D}(x_*)) \leq C\|x - x_*\|_2$. Hence (44) holds if $\|\epsilon_{\text{err},k}\|_\infty \leq (\Lambda'/C)\omega(X_k)$, i.e., if the noise is bounded by a constant multiple of the KKT error at X_k , and conversely a bound of the form (44) implies one of this form. Observations of this type motivate various adaptive sampling strategies in practice.

Our second setting, for identification through a primal step as presented in §4.2, requires the LICQ and strict complementarity at a KKT point x_* rather than the MFCQ and the second-order condition (4). For this setting in the following theorem, we recall that for all $k \in \mathbb{N}$ the active-set estimate is obtained by solving the QP (35), which yields the subproblem solution we denote as (D_k, R_k, S_k, T_k) .

Theorem 4.5. *Suppose that, with respect to (2), the conditions of Theorem 4.2 hold. In addition, with $(\epsilon_x, \epsilon_e, \epsilon_c, \epsilon_{\nabla f}, \epsilon_{\nabla e}, \epsilon_{\nabla c}, \underline{\nu})$ as defined in Theorem 4.2, suppose that for some universal constant $p \in (0, 1]$ one has, almost surely and for all $k \in \mathbb{N}$,*

$$\mathbb{P}[(22) \text{ holds with } x = X_k \mid \mathcal{F}_k] \geq p \cdot \mathbf{1}_{\{\|X_k - x_*\|_2 \leq \epsilon_x\}}, \quad (46)$$

which is well defined because X_k is $\mathcal{F}_k$ -measurable, and which holds with p given explicitly by Lemma 4.6 under sampling with a fixed sample size. Then, for any parameter $\theta \in \mathbb{R}_{>0}$ and $\nu \geq \underline{\nu}$, if $\{X_k\}$ converges almost surely to x_* , it holds with probability one that there exists an index $k_0 \in \mathbb{N}$ such that for all $k \in \mathbb{N}$ with $k \geq k_0$ it holds with conditional probability at least p that the solution (D_k, R_k, S_k, T_k) of (35) with $x = X_k$ is unique and the active-set estimate defined by (42) yields $\tilde{\mathcal{A}}_{\text{QP}}(D_k; X_k) = \mathcal{A}_*$. Moreover, with probability one, the active-set estimate has $\tilde{\mathcal{A}}_{\text{QP}}(D_k; X_k) = \mathcal{A}_*$ for infinitely many $k \in \mathbb{N}$.

Proof. Proof. Same proof as Theorem 4.4, with Ω_0 the event on which $X_k \rightarrow x_*$, with $\mathcal{E}_k$ the event in (46), with $\mathcal{B}_k := \{\|X_k - x_*\|_2 \leq \epsilon_x\}$, and with Theorem 4.2 in place of Corollary 4.1. Neither $X_k \neq x_*$ nor error bounds that vanish as $X_k \rightarrow x_*$ are needed, since the error bounds in Theorem 4.2 are fixed. $\square$ $\square$

By our previous arguments, a similar result would hold if the QP subproblem (35) were replaced by an LP subproblem with a trust-region constraint on the primal step.

5 Numerical Results

The purpose of our experiments is to demonstrate the behavior of the active-set identification techniques analyzed in §4 in the context of a few test problems. We present the results of two sets of experiments. Our first set of experiments focuses on two related two-dimensional test problems that allow us to provide clear visualizations of active-set identification behavior in both deterministic and noisy settings. Our second set of experiments focuses on behavior in the context of a larger-dimensional neural-network training problem. We emphasize that any number of problems could be used to demonstrate the behavior of the active-set identification methods from §4, each potentially leading to interesting conclusions about the techniques in general, the parameter choices for them, and so on. We chose our first experiments since two-dimensional problems allow for nice visualizations, and we chose our second experiments since they involve the problem of training of a neural network, which is of immense interest in modern optimization.

5.1 Visualization with 2D Problems

Consider a pair of two-dimensional problems with

$$\min_{x \in \mathbb{R}^2} f(x) \text{ s.t. } c_1(x) := x_1^2 - x_2 \leq 0 \text{ and } c_2(x) := x_1^2 + x_2 - \frac{1}{2} \leq 0, \quad (47)$$

where c_i and x_i for each $i \in \{1, 2\}$ denote the i th component of the constraint function $c : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and the i th component of the vector $x \in \mathbb{R}^2$, respectively. (Since there are no equality constraints in (47), we denote the constraint function simply as c .) The feasible region is thus the region in $\mathbb{R}^2$ contained between two parabolas, one upward-facing and one downward-facing. Let us also consider two objective functions $f_1 : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $f_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ that are defined respectively as

$$f_1(x) = (x_1 + \frac{1}{2})^2 + 4(x_2 - \frac{1}{2})^2 \text{ and } f_2(x) = 4(x_1 + \frac{3}{5})^2 + (x_2 - \frac{1}{4})^2.$$

The problems to minimize f_1 and minimize f_2 , each subject to $c(x) \leq 0$, have convex feasible regions and strongly convex objectives. Therefore, each problem has a unique solution, call it x_* in each case. It is straightforward to confirm that the problem to minimize f_1 subject to $c(x) \leq 0$ has one active constraint at x_* such that $\mathcal{A}(x_*) = \{2\}$, and at x_* the LICQ and strict complementarity hold. On the other hand, for the problem to minimize f_2 subject to $c(x) \leq 0$, both constraints are active at x_* , i.e., $\mathcal{A}(x_*) = \{1, 2\}$, and the LICQ and strict complementarity hold at x_* .

For each problem, we generated points in a grid around x_* . For each point x in the grid, we solved various instances of the LP-LPEC subproblem (27) with $M \leftarrow 10^8$ and (28) applied with $\beta \leftarrow 0.7071$ (determined as $\beta \leftarrow 1/\max\{\|\nabla f(\bar{x})\|_\infty, \|c(\bar{x})\|_\infty, \|\nabla c(\bar{x})\|_\infty\}$ at a point $\bar{x}$ chosen randomly in the grid, so that $\beta\bar{\rho}$ is scaled to the problem data) and $\sigma \leftarrow 0.7$; an LP with a trust-region constraint (that is (15) with approximate function and derivative values) with $\nu \leftarrow 100$ and trust-region radius $\Delta \leftarrow 0.1$; and the QP subproblem (35) with $\nu \leftarrow 100$ and $\theta \leftarrow 5$. The active-set estimates were determined accordingly, as described in §4.1 and §4.2. For each subproblem, we first considered the deterministic setting, denoted as $\epsilon = 0$. For each point in the grid, we determined whether the active-set estimate exactly matched $\mathcal{A}(x_*)$ or not, yielding a value of 1 or 0, respectively. Then, for two noise levels, denoted by $\epsilon \in \{10^{-2}, 10^{-1}\}$, respectively, we solved 8 instances of each subproblem with noise added to the function and derivative values, then determined the fraction of times that the active-set estimate exactly matched $\mathcal{A}(x_*)$. This yielded a fractional value between 0 and 1 for each point, where a value of 1 indicates that the active set was always identified correctly. With respect to the noise, for each function value or component of a function or derivative value, we added a number generated from a uniform distribution over the interval $[-\epsilon, \epsilon]$, determined by the noise level.

The results associated with f_1 and f_2 are presented in Figures 1 and 2, respectively. In each case, one finds that for noise level 10^{-2} , the active-set identification at points near x_* is relatively robust to noise. For the LP-TR and QP approaches, this is consistent with Theorem 4.2 and its LP counterpart, since the LICQ and strict complementarity hold at x_* for both problems. For the LP-LPEC approach, Corollary 4.1 guarantees correct identification only when the noise vanishes as $x \rightarrow x_*$, so the robustness observed near x_* at a fixed noise level is an empirical finding rather than a consequence of our theory. On the other hand, for noise level 10^{-1} , the active-set identification is less reliable for both methods, failing to identify $\mathcal{A}(x_*)$ in some cases. This is also consistent with our theoretical guarantees, since our theory requires that the noise is sufficiently small in order to guarantee accurate active-set identification. At least for these test problems, the LP-LPEC approach identifies the active set correctly less often than the LP-TR or QP approach when only one constraint is active at x_* , and identifies the active set correctly more often when both constraints are active, at least with our choices of the parameters β , σ , ν , Δ , and θ . Naturally, the behavior of the methods can vary with these parameter choices, so these comparisons do not necessarily hold in general.

A notable observation about the heatmaps in Figures 1 and 2 is that, when the noise is nonzero, one should only assume that the errors in the function and derivative values are small with some probability such that active-set identification can be performed accurately, even within a small neighborhood of x_* . This is an empirical demonstration of the comments that we made in §4.4.

5.2 Neural Network Training Problem

Now let us consider a constrained neural network training problem. The problem that we consider is one of supervised learning of a prediction function that aims to predict, based on feature data for a given patient, whether the patient is at risk for heart disease. The constraints are designed to ensure that the prediction function yields a higher probability of heart disease when the patient has a known high-risk factor. We emphasize that the constraints that we impose are for demonstration purposes only, and may not be ones that lead to better predictions in reality. That being said, the problem that we consider in this section might represent the type of constrained supervised learning problem that could potentially lead to better and more explainable prediction functions in the context of medical diagnoses, as opposed to prediction functions obtained purely through empirical loss minimization.

Let $\{(a_i, b_i)\}_{i=1}^N$ be a set of training data, where for each $i \in [N]$ the value $a_i \in \mathbb{R}^{n_f}$ is an n_f -vector of feature data for a given patient (including age, sex, and cholesterol level) and the value $b_i \in \{0, 1\}$ indicates whether the patient eventually developed heart disease (indicated by $b_i = 1$) or not (indicated by $b_i = 0$). In particular, let us suppose that for each $i \in [N]$ the j th component of a_i , namely, $a_{i,j}$, indicates the patient’s cholesterol level. For the purposes of devising a tractable problem, let us further suppose that the N data points are separated into K groups, where the indices in $\mathcal{N}_1$ correspond to group 1, the indices in $\mathcal{N}_2$ correspond to group 2, etc. (In this manner, one has that $\{\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_K\}$ are disjoint sets and $|\mathcal{N}_1| + |\mathcal{N}_2| + \dots + |\mathcal{N}_K| = N$. In practice, the groups might be determined randomly or according to certain features, such as age and/or sex.) Given a prediction function p parameterized by a vector $x \in \mathbb{R}^n$ of trainable parameters—in our setting, defined by a neural network—and a loss function ℓ , our problem is

$$\min_{x \in \mathbb{R}^n} \sum_{i=1}^N \ell(p(x, a_i), b_i) \quad \text{s.t.} \quad \max_{i \in \mathcal{N}_k} \left\{ -\frac{\partial p}{\partial a_{\cdot,j}}(x, a_i) \right\} \leq 0 \quad \text{for all } k \in [K], \quad (48)$$

where $\partial p / \partial a_{\cdot,j}$ indicates the partial derivative of p with respect to the j th element of the input of the neural network, which in our setting corresponds to cholesterol. Here, the objective function is a typical one for empirical loss minimization that aims to maximize the accuracy of predictions on unseen data. The constraints, on the other hand, aim to ensure that—all else being equal—an increase in a patient’s cholesterol would increase (or at least not decrease) the prediction that the patient will develop heart disease. Rather than impose a constraint for each data point, which may lead to an intractable problem in large-scale settings, a constraint is imposed for each group in the data set. It should be noted that the maximum function is nonsmooth, and in practice one could impose a smooth constraint through a smoothed approximation of the maximum function. However, for our experiments we did not smooth the functions since (a) we were able to solve the problem to sufficient accuracy despite nonsmoothness in the constraint functions and (b) it is often the case in modern practice that algorithms employed for neural network training ignore nonsmoothness due to maximum functions, as with rectified linear unit (ReLU) activation. We also speculate that nonsmoothness of this type might actually help identification rather than hurt it. After all, a max-function constraint is sharp at points where the maximum is attained uniquely, and sharpness is precisely the property that drives finite identification in nonsmooth analysis and optimization theory; see, e.g., [16, 46, 40].

We acquired a data set of 1025 data points with each point having $n_f = 13$ features [42]. For the prediction function p , we used a neural network with one hidden layer with 6 nodes and tanh activation. Including both neural network edge weights and bias terms as trainable parameters, this led to a problem with 91 variables. For the loss function we used the log-loss function. The data was split initially into training and testing data according to cholesterol level; in particular, we sorted the data points by cholesterol and used the top 20% according to cholesterol as the testing data. In this manner, training was conducted on those with lower cholesterol and testing was conducted on those with higher cholesterol. As an interesting side-note to our study of the behavior of active-set identification strategies, the idea here was to test the effect of training without the highest-cholesterol individuals in order to see if the presence of the constraints in (48) was able to improve the predictions on the testing data.

The training data was split into 10 groups to define 10 constraints. We began our experimentation by employing an implementation of an L-BFGS method with a backtracking Armijo line search to minimize

the objective of (48), without any constraints. This resulted in an approximate stationary point—we used a threshold of 10^{-5} for the norm of the gradient—with a maximum constraint function value of $6.03\text{e-}03$; i.e., it resulted in a point that was infeasible for the constraints in (48). Corresponding to this point, the training accuracy was determined to be 88.05% while the testing accuracy was determined to be 81.46%.

We then employed the same L-BFGS method to minimize a quadratic penalty function corresponding to (48) with a penalty parameter of $1\text{e+}03$; see [51] for further details on penalty methods. With a tolerance of 10^{-5} for the norm of the gradient (of the quadratic penalty function), this resulted in a point with a maximum constraint value of $8.57\text{e-}09$. Corresponding to this point, the training accuracy was determined to be 92.20% while the testing accuracy was determined to be 84.88%. This demonstrates a case in which the presence of the constraints during the training process was able to improve testing accuracy by over 3%.

We emphasize that this experiment evaluates the three identification rules along a sequence of iterates produced by an algorithm that does not itself use active-set information. As such, it is not a demonstration of an algorithm that computes search directions from an estimated active set, which is a topic for future work (see §6). Denote by $\hat{x}$ the final iterate of the L-BFGS method applied to the quadratic penalty function, which we treat as a solution estimate. At $\hat{x}$, in order to determine the estimated optimal active set, we employed the active-set identification techniques from §4 to solve an LP-LPEC subproblem, an LP-TR subproblem with a trust-region constraint, and a QP subproblem. Here and throughout the remainder of these experiments, we employed the parameters $M \leftarrow 10^8$, $\beta \leftarrow 10^{-10}$, $\sigma \leftarrow 0.7$, $\nu \leftarrow 10^2$, $\Delta \leftarrow 10^{-2}$, and $\theta \leftarrow 10^2$. Solving all of these subproblems led to the same estimate, $\{1, \dots, 10\}$, i.e., all constraints being active. We take this common estimate as the reference active set in what follows. That is, we do not claim that $\hat{x}$ is a local minimizer of (48), nor that the reference set is provably $\mathcal{A}(x_*)$ for a nearby minimizer.

We studied the behavior of the active-set identification techniques from §4 by following the sequence of iterates in the L-BFGS method for minimizing a quadratic penalty function, as described previously. At each iterate, Figure 3 shows the number of correctly identified active constraints by the LP-LPEC, LP-TR, and QP approaches when there *is no noise* in the function and derivative values used to define the subproblems for each approach. One finds that the LP-LPEC approach identifies the correct active set even in early iterations. We found this to be due to the fact that the iterates are consistently points where all constraints are violated. Combined with the fact that the right-hand side of the inequality in (28) is always negative, this demonstrates a situation in which the LP-LPEC approach is lucky always to identify violated constraints as being active at the optimal solution. The LP-TR and QP approaches, on the other hand, do not identify the optimal active set until the iterates are closer to $\hat{x}$. That said, it should be emphasized that all approaches correctly identify the optimal active set in all later iterations.

To observe the effect of noise in the function and derivative values on active-set identification, we ran the algorithm with exact function and derivative values, but at each iterate we generated 10 different sets of noisy objective and constraint function and derivative values for active-set identification by sampling the data. In particular, at each iterate, we used a random sample of data points with sample size equal to 32 in order to evaluate the function and derivative values employed in the subproblems for active-set identification. (We ran the code with other similar sample sizes and the results were qualitatively similar to those presented here.) The results are shown in Figure 4, where for each iteration index we plot the average number of correctly identified active constraints over the 10 different samples that were generated. One finds that the LP-LPEC approach does a poorer job (compared to the results in Figure 3) of identifying the optimal active set in earlier iterations, but it still identifies most of the optimal active set correctly in later iterations. The LP-TR and QP approaches, on the other hand, happen to do a better job of identifying the active set.

As for the experiments in §5.1, the behavior of the methods can vary with different problems and parameter choices, so these specific comparisons do not necessarily hold in general. That said, the important thing is that we observe accurate active-set identification in later iterations for all of the subproblem techniques, even with noise in the problem function and derivative values.

6 Conclusion

We have shown that state-of-the-art techniques for active-set identification when solving constrained continuous optimization problems can also be employed in contexts when function and derivative values are corrupted by noise. We have also quantified the noise that can be tolerated. Two of our conclusions deserve explicit emphasis. First, the admissible noise level can be made explicit, as seen in our analysis. Under the LICQ and strict complementarity it is proportional to the product of the conditioning of the active constraint Jacobian with the degeneracy margins of the minimizer (Lemma 4.5), and under the MFCQ and the second-order condition (4) it suffices that the noise vanish linearly in the distance to the minimizer (Corollary 4.1). Second, the contrast between these regimes is intrinsic rather than an artifact of the strategies analyzed. In particular, at a minimizer with a weakly active constraint, no identification rule can succeed under a fixed noise level, and the noise must in general vanish at least linearly, so the rate of Corollary 4.1 cannot be improved (Theorem 4.3). Accordingly, the results of this paper may be read both as delimiting what active-set identification can achieve in the presence of noise and as a contribution to the perturbation theory of nonlinear programming, in which the perturbed object is a computable subproblem at a point that is merely near a solution. We have also shown how such techniques can be employed to offer active-set identification guarantees in the context of a stochastic algorithm for solving constrained continuous optimization problems. We have also provided the results of numerical experiments that demonstrate the behaviors of these techniques on a pair of illustrative problems and a neural-network training problem.

Interesting follow-up investigations to our work in this paper include studying how active-set identification techniques can be employed within noisy and stochastic algorithms for solving constrained continuous optimization problems both for computing certain types of primal search directions and for more accurate Lagrange multiplier estimation. For example, one may consider whether, despite noise in problem function and derivative values, one can design algorithms with convergence guarantees that operate by (a) estimating the active set, then (b) computing each search direction with respect to the estimated active set, as has been done in the deterministic setting [19, 18]. It would also be interesting to investigate how, e.g., the theoretical guarantees for Lagrange multiplier estimation presented in [24] (for the case of equality constraints only) can be extended to the setting of an algorithm for solving inequality-constrained problems. This may involve employing the active-set identification techniques and theoretical guarantees from this paper in order to provide guarantees for optimal Lagrange multiplier estimation.

A Proof of Theorem 3.3

Proof. Proof. Consider part (a). First, by continuity of c , it follows that one can choose $\epsilon \in \mathbb{R}_{>0}$ small enough such that for any $x \in \mathbb{R}^n$ with $\|x - x_*\|_2 \leq \epsilon$ and any $i \notin \mathcal{A}(x_*)$ one has that $c_i(x) \leq \frac{1}{2}c_i(x_*) < 0$. Second, given such an ϵ and for any such x and positive ν , it follows from (14) that $\|d_{x,\theta,\nu}\|_2 \rightarrow 0$ as $\theta \nearrow \infty$. Thus, one can choose $\underline{\theta}$ large enough such that $i \notin \mathcal{A}(x_*)$ implies $c_i(x) + \nabla c_i(x)^T d_{x,\theta,\nu} < 0$. Thus, the claim in part (a) follows from the definition (20).

Next, consider part (b). Let us first suppose that $(\epsilon, \underline{\theta})$ satisfies the conditions of part (a). Furthermore, to prove the claim, let us introduce

$$\underline{\nu} := \max\{\max\{\|y_*\|_\infty, \|z_*\|_\infty\} + 1 : (y_*, z_*) \in \mathcal{D}(x_*)\}. \quad (49)$$

The value $\underline{\nu}$ is well-defined and positive since x_* is a primal KKT point and the MFCQ assumption ensures that $\mathcal{D}(x_*)$ is bounded. For any pair $(y_*, z_*) \in \mathcal{D}(x_*)$, the dual subproblem (18) at x_* has an objective value of 0, as previously stated. Otherwise, for the dual subproblem at $x \neq x_*$ with $\|x - x_*\|_2 \leq \epsilon$ and $\nu \geq \underline{\nu}$, a feasible point is given by $(\alpha, \beta) = (y_*, z_*)$, for which

$$\begin{aligned} & e(x)^T y_* + c(x)^T z_* - \frac{1}{2}\theta^{-1}\|\nabla f(x) + \nabla e(x)y_* + \nabla c(x)z_*\|_2^2 \\ &= (e(x) - e(x_*))^T y_* + (c(x) - c(x_*))^T z_* \\ &\quad - \frac{1}{2}\theta^{-1}\|\nabla f(x) - \nabla f(x_*) + (\nabla e(x) - \nabla e(x_*))y_* + (\nabla c(x) - \nabla c(x_*))z_*\|_2^2 \end{aligned}$$

$$\geq -L(\|x - x_*\|_2^2 + \|x - x_*\|_2)$$

for a sufficiently large constant $L \in \mathbb{R}_{>0}$. This latter inequality follows due to the continuous differentiability of the problem functions and $\|x - x_*\|_2 \leq \epsilon$.

Now let us assume, for the sake of deriving a contradiction, that no matter how small one chooses $\epsilon \in \mathbb{R}_{>0}$, there exists a point $x \in \mathbb{R}^n$ satisfying $\|x - x_*\|_2 \leq \epsilon$ such that the resulting active-set estimate $\mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x)$ yields $-\nabla f(x_*) \notin \text{range}[\nabla e(x_*)] + \text{pos}[\{\nabla c_i(x_*) : i \in \mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x)\}]$. Since there are only finitely many active set estimates (i.e., subsets of $[q]$), it follows that there exists a sequence of points converging to x_* such that along this sequence the active-set estimate is the same set of indices, call it $\hat{\mathcal{A}}$. At the same time, observe that by Rockafellar [59, Theorem 19.1] the set $\text{range}[\nabla e(x_*)] + \text{pos}[\{\nabla c_i(x_*) : i \in \hat{\mathcal{A}}\}]$ is finitely generated, meaning that it is a closed set. Now, by the definition of the distance function in equation (1), one can define the quantity τ as

$$\tau := \frac{1}{2} \text{dist}(-\nabla f(x_*), \text{range}[\nabla e(x_*)] + \text{pos}[\{\nabla c_i(x_*) : i \in \hat{\mathcal{A}}\}]) > 0. \quad (50)$$

Reducing ϵ , if necessary, one has $\text{dist}(-\nabla f(x), \text{range}[\nabla e(x)] + \text{pos}[\{\nabla c_i(x) : i \in \hat{\mathcal{A}}\}]) \geq \tau$ for the given $\hat{\mathcal{A}}$ and all x satisfying $\|x - x_*\|_2 \leq \epsilon$. The justification for this claim follows from standard continuity arguments; see [52, Appendix 1].

Now consider any such x such that $\mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x) = \hat{\mathcal{A}}$. Let (d_x, r_x, s_x, t_x) and (α_x, β_x) denote the solutions to (16) and (18), respectively. For any index $i \notin \hat{\mathcal{A}}$, the KKT conditions for (16)–(18) and equation (20) together imply that

$$c_i(x) + \nabla c_i(x)^T d_x < 0, \quad [r_x]_i = 0, \quad \text{and} \quad [\beta_x]_i = 0. \quad (51)$$

One can now analyze the objective of the dual subproblem (18) by separating it into two parts. First, using property (51), one can observe that

$$\begin{aligned} & \frac{1}{2\theta} \|\nabla f(x) + \nabla e(x)\alpha_x + \nabla c(x)\beta_x\|_2^2 \\ & \geq \frac{1}{2\theta} \min_{(\alpha, \beta) \in \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q} \|\nabla f(x) + \nabla e(x)\alpha + \nabla c(x)\beta\|_2^2 \\ & = \frac{1}{2\theta} \min_{(\alpha, \beta) \in \mathbb{R}^p \times \mathbb{R}_{\geq 0}^q} \left\| \nabla f(x) + \nabla e(x)\alpha + \sum_{i \in \hat{\mathcal{A}}} \nabla c_i(x)\beta_i \right\|_2^2 \\ & = \frac{1}{2\theta} \text{dist}(-\nabla f(x), \text{range}[\nabla e(x)] + \text{pos}[\{\nabla c_i(x) : i \in \hat{\mathcal{A}}\}])^2 \geq \frac{1}{2\theta} \tau^2. \end{aligned}$$

From $\beta_x \geq 0$ and (18), one also has that $e(x)^T \alpha_x + c(x)^T \beta_x \leq e(x)^T \alpha_x + (c(x))_+^T \beta_x \leq \nu \|e(x)\|_1 + \nu \|(c(x))_+\|_1$. Combining these inequalities into the dual objective (18), one has

$$\begin{aligned} e(x)^T \alpha_x + c(x)^T \beta_x - \frac{1}{2\theta} \|\nabla f(x) + \nabla e(x)\alpha_x + \nabla c(x)\beta_x\|_2^2 \\ \leq \nu \|e(x)\|_1 + \nu \|(c(x))_+\|_1 - \frac{1}{2\theta} \tau^2. \end{aligned} \quad (52)$$

Finally, one can decrease ϵ further, if necessary, to have that

$$\nu \|e(x)\|_1 + \nu \|(c(x))_+\|_1 - \frac{1}{2\theta} \tau^2 < -\frac{\|x - x_*\|_2^\zeta}{4} \tau^2. \quad (53)$$

This choice is possible since $e(x)$ and $(c(x))_+$ are both $\mathcal{O}(\|x - x_*\|_2)$ and under the conditions of the theorem one has that $\theta \leq \|x - x_*\|_2^{-\zeta}$. Since $\zeta \in (0, 1)$, the bounds (52)–(53) stand in direct contradiction to our earlier result, which established that the optimal value of the dual has

$$-L(\|x - x_*\|_2^2 + \|x - x_*\|_2) \leq e(x)^T \alpha_x + c(x)^T \beta_x - \frac{1}{2\theta} \|\nabla f(x) + \nabla e(x)\alpha_x + \nabla c(x)\beta_x\|_2^2.$$

This inconsistency forces the conclusion that τ defined in equation (50) cannot be positive, and in fact must be 0. Thus, the inclusion in (21) must hold, as desired.

Finally, consider part (c). Since the LICQ holds, one has that $[z_*]_i > 0$ for all $i \in \mathcal{A}(x_*) \setminus \mathcal{A}_0(x_*)$. Thus, for the inclusion (21) to hold, any index i for which $[z_*]_i > 0$ must satisfy $i \in \mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x)$; otherwise, (21) would be violated. Therefore, one can conclude the left-hand side inclusion $\mathcal{A}(x_*) \setminus \mathcal{A}_0(x_*) \subseteq \mathcal{A}_{\text{QP}}(d_{x,\theta,\nu}; x)$. The final conclusion follows directly from the fact that strict complementarity implies that $\mathcal{A}_0(x_*) = \emptyset$, meaning there are no active inequality constraints at the solution for which the corresponding multiplier is zero. The proof is complete. $\square$ $\square$

References

- [1] Albert S. Berahas, Richard H. Byrd, and Jorge Nocedal. Derivative-free optimization of noisy functions via quasi-Newton methods. *SIAM J. on Opt.*, 29(2):965–993, 2019.
- [2] Albert S. Berahas, Liyuan Cao, and Katya Scheinberg. Global convergence rate analysis of a generic line search algorithm with noise. *SIAM J. on Opt.*, 31(2):1489–1518, 2021.
- [3] Albert S. Berahas, Frank E. Curtis, Michael J. O’Neill, and Daniel P. Robinson. A stochastic sequential quadratic optimization algorithm for nonlinear equality constrained optimization with rank-deficient jacobians. *Math. of Oper. Res.*, 49(4):2049–2082, 2024.
- [4] Albert S. Berahas, Frank E. Curtis, Daniel P. Robinson, and Baoyu Zhou. Sequential quadratic optimization for nonlinear equality constrained stochastic optimization. *SIAM J. on Opt.*, 31(2):1352–1379, 2021.
- [5] Dimitri P. Bertsekas. *Convex Optimization Theory*. Athena Scientific, Belmont, Mass, 2009.
- [6] Jose Blanchet, Coralia Cartis, Matt Menickelly, and Katya Scheinberg. Convergence rate analysis of a stochastic trust-region method via supermartingales. *INFORMS J. on Opt.*, 1(2):92–119, 2019.
- [7] Raghu Bollapragada, Richard H. Byrd, and Jorge Nocedal. Adaptive sampling strategies for stochastic optimization. *SIAM J. on Opt.*, 28(4):3312–3343, 2018.
- [8] J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer, New York, 2000.
- [9] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- [10] Richard L. Burden and John Douglas Faires. *Numerical Analysis*. Thomson Brooks/Cole, Belmont, CA, 8th edition, 2006.
- [11] James V. Burke. A sequential quadratic programming method for potentially infeasible mathematical programs. *J. of Math. Anal. and Appl.*, 139(2):319–351, 1989.
- [12] James V. Burke. On the identification of active constraints II: The nonconvex case. *SIAM J. on Num. Anal.*, 27(4):1081–1102, 1990.
- [13] James V. Burke. Calmness and exact penalization. *SIAM J. on Cont. and Opt.*, 29(2):493–497, 1991.
- [14] James V. Burke and Michael C. Ferris. A Gauss–Newton method for convex composite optimization. *Math. Prog.*, 71(2):179–194, 1995.
- [15] James V. Burke and Shih-Ping Han. A robust sequential quadratic programming method. *Math. Prog.*, 43(1):277–303, 1989.
- [16] James V. Burke and Jorge J. Moré. On the identification of active constraints. *SIAM J. on Num. Anal.*, 25(5):1197–1211, 1988.

- [17] Richard H. Byrd, Gillian M. Chin, Jorge Nocedal, and Yuchen Wu. Sample size selection in optimization methods for machine learning. *Math. Prog.*, 134(1):127–155, 2012.
- [18] Richard H. Byrd, Nicholas I. M. Gould, Jorge Nocedal, and Richard A. Waltz. On the convergence of successive linear-quadratic programming algorithms. *SIAM J. on Opt.*, 16(2):471–489, 2005.
- [19] Richard H. Byrd, Nicholas I.M. Gould, Jorge Nocedal, and Richard A. Waltz. An algorithm for nonlinear optimization using linear programming and equality constrained subproblems. *Math. Prog.*, 100(1):27–48, 2003.
- [20] Richard H. Byrd, Jorge Nocedal, and Richard A. Waltz. KNITRO: An integrated package for nonlinear optimization. In *Large-Scale Nonlinear Optimization*, pages 35–59. Springer, 2006.
- [21] Richard H. Byrd, Jorge Nocedal, and Richard A. Waltz. Steering exact penalty methods for nonlinear programming. *Opt. Meth. and Soft.*, 23(2):197–213, 2008.
- [22] Coralia Cartis and Katya Scheinberg. Global convergence rate analysis of unconstrained optimization methods based on probabilistic models. *Math. Prog.*, 169(2):337–375, 2018.
- [23] 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.
- [24] Frank E. Curtis, Xin Jiang, and Qi Wang. Almost-sure convergence of iterates and multipliers in stochastic sequential quadratic optimization. *J. of Opt. Th. and Appl.*, 2024.
- [25] Frank E. Curtis, Michael J. O’Neill, and Daniel P. Robinson. Worst-case complexity of an SQP method for nonlinear equality constrained stochastic optimization. *Math. Prog.*, 205:431–483, 2024.
- [26] Frank E. Curtis, Daniel P. Robinson, and Baoyu Zhou. Sequential quadratic optimization for stochastic optimization with deterministic nonlinear inequality and equality constraints. *SIAM J. on Opt.*, 34(4):3592–3622, 2024.
- [27] Frank E. Curtis, Daniel P. Robinson, and Baoyu Zhou. A stochastic inexact sequential quadratic optimization algorithm for nonlinear equality-constrained optimization. *INFORMS J. on Opt.*, 6(3–4):155–239, 2024.
- [28] Frank E. Curtis and Katya Scheinberg. Optimization methods for supervised machine learning: From linear models to deep learning. In *INFORMS Tutorials in Operations Research*, chapter 5, pages 89–114. Institute for Operations Research and the Management Sciences (INFORMS), 2017.
- [29] 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.
- [30] Asen L. Dontchev and R. Tyrrell Rockafellar. Characterizations of strong regularity for variational inequalities over polyhedral convex sets. *SIAM J. on Opt.*, 6(4):1087–1105, 1996.
- [31] Asen L. Dontchev and R. Tyrrell Rockafellar. *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer, New York, 2009.
- [32] Francisco Facchinei, Andreas Fischer, and Christian Kanzow. On the accurate identification of active constraints. *SIAM J. on Opt.*, 9(1):14–32, 1998.
- [33] Yuchen Fang, Sen Na, Michael W. Mahoney, and Mladen Kolar. Fully stochastic trust-region sequential quadratic programming for equality-constrained optimization problems. *SIAM J. on Opt.*, 34(2):2007–2037, 2024.

- [34] Anthony V. Fiacco. Sensitivity analysis for nonlinear programming using penalty methods. *Math. Prog.*, 10(1):287–311, 1976.
- [35] Roger Fletcher. An ℓ_1 penalty method for nonlinear constraints. In *Numerical Optimization 1984*, pages 26–40. SIAM, Philadelphia, 1985.
- [36] Jacques Gauvin. A necessary and sufficient regularity condition to have bounded multipliers in nonconvex programming. *Math. Prog.*, 12(1):136–138, 1977.
- [37] Philip E. Gill, Walter Murray, and Michael A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM J. on Opt.*, 12(4):979–1006, 2002.
- [38] William W. Hager and M. Seetharama Gowda. Stability in the presence of degeneracy and error estimation. *Math. Prog.*, 85(1):181–192, May 1999.
- [39] S.-P. Han and Olvi L. Mangasarian. Exact penalty functions in nonlinear programming. *Math. Prog.*, 17(1):251–269, December 1979.
- [40] Warren Hare and Adrian S. Lewis. Identifying active constraints via partial smoothness and prox-regularity. *J. of Conv. Anal.*, 11:251–266, 2004.
- [41] Alexey F. Izmailov and Mikhail V. Solodov. *Newton-Type Methods for Optimization and Variational Problems*. Springer, Cham, 2014.
- [42] Andras Janosi, William Steinbrunn, Matthias Pfisterer, and Robert Detrano. Heart Disease. UCI Machine Learning Repository, 1989. DOI: <https://doi.org/10.24432/C52P4X>.
- [43] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems 26*, pages 315–323, 2013.
- [44] Diethard Klatte and Bernd Kummer. *Nonsmooth Equations in Optimization: Regularity, Calculus, Methods and Applications*. Kluwer Academic Publishers, Dordrecht, 2002.
- [45] Masakazu Kojima. Strongly stable stationary solutions in nonlinear programs. In *Analysis and Computation of Fixed Points*, pages 93–138. Academic Press, 1980.
- [46] Adrian S. Lewis. Active sets, nonsmoothness, and sensitivity. *SIAM J. on Opt.*, 13(3):702–725, 2002.
- [47] Olvi L. Mangasarian and Stan Fromovitz. The Fritz John Necessary Optimality Conditions in the Presence of Equality and Inequality Constraints. *J. of Math. Anal. and Appl.*, 17(1):37–47, 1967.
- [48] Sen Na, Mihai Anitescu, and Mladen Kolar. An adaptive stochastic sequential quadratic programming with differentiable exact augmented lagrangians. *Math. Prog.*, pages 1–71, jun 2022.
- [49] Sen Na, Mihai Anitescu, and Mladen Kolar. Inequality constrained stochastic nonlinear optimization via active-set sequential quadratic programming. *Math. Prog.*, pages 1–75, 2023.
- [50] Lam M. Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. SARAH: A novel method for machine learning problems using stochastic recursive gradient. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2613–2621, 2017.
- [51] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006.
- [52] Christina Oberlin and Stephen J. Wright. Active Set Identification in Nonlinear Programming. *SIAM J. on Opt.*, 17(2):577, 2006.

- [53] James M. Ortega and Werner Rheinboldt. *Numerical Analysis: A Second Course*. Elsevier Science, Burlington, 2014.
- [54] Figen Oztoprak, Richard Byrd, and Jorge Nocedal. Constrained optimization in the presence of noise. *SIAM J. on Opt.*, 33(3):2118–2136, 2023.
- [55] Courtney Paquette and Katya Scheinberg. A stochastic line search method with expected complexity analysis. *SIAM J. on Opt.*, 30(1):349–376, 2020.
- [56] Stephen M. Robinson. Perturbed kuhn-tucker points and rates of convergence for a class of nonlinear-programming algorithms. *Math. Prog.*, 7(1):1–16, Dec 1974.
- [57] Stephen M. Robinson. Stability theory for systems of inequalities, part II: Differentiable nonlinear systems. *SIAM J. on Num. Anal.*, 13(4):497–513, 1976.
- [58] Stephen M. Robinson. Strongly regular generalized equations. *Math. of Oper. Res.*, 5(1):43–62, 1980.
- [59] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [60] Shigeng Sun and Jorge Nocedal. A trust region method for noisy unconstrained optimization. *Math. Prog.*, pages 1–28, 2023.
- [61] David Williams. *Probability with Martingales*. Cambridge University Press, Cambridge, 1991.
- [62] Stephen J. Wright. Identifiable surfaces in constrained optimization. *SIAM J. on Cont. and Opt.*, 31(4):1063–1079, 1993.
- [63] Stephen J. Wright. Modifying SQP for degenerate problems. *SIAM J. on Opt.*, 13(2):470–497, 2002.
- [64] Yuchen Xie, Richard H Byrd, and Jorge Nocedal. Analysis of the BFGS method with errors. *SIAM J. on Opt.*, 30(1):182–209, 2020.




Figure 1: Heatmaps for the problem to minimize f_1 subject to $c(x) \leq 0$ for noise levels $\epsilon = 0$ (deterministic), 10^{-2} , and 10^{-1} , respectively. The top row corresponds to the LP-LPEC approach based on computing Lagrange multiplier estimates, the middle row corresponds to the LP-TR approach based on computing a primal step, and the bottom row corresponds to the QP approach based on computing a primal step. The value of the heatmap at each point indicates the fraction of times that the active set $\mathcal{A}(x_*)$ was identified correctly by solving a subproblem at the point. (The unique optimal solution x_* is indicated by a black dot.) In the case of positive noise levels, the value was determined as a fraction over 8 randomly generated subproblem instances.




Figure 2: Heatmaps for the problem to minimize f_2 subject to $c(x) \leq 0$. The details are the same as in the caption for Figure 1.




Figure 3: Over the iterates of an algorithm to solve (48), the number of correctly identified elements of the optimal active set by the LP-LPEC (left), LP-TR (middle), and QP (right) approach when there *is no noise* in the function and derivative values.




Figure 4: Over the iterates of an algorithm to solve (48), the number of correctly identified elements of the optimal active set by the LP-LPEC (left), LP-TR (middle), and QP (right) approach when there *is noise* in the function and derivative values.