J Glob Optim (2014) 59:327–341
DOI 10.1007/s10898-014-0186-y

Error estimation in nonlinear optimization

William W. Hager · Delphine Mico-Umutesi

Received: 3 December 2013 / Accepted: 1 April 2014 / Published online: 9 April 2014

Abstract Methods are developed and analyzed for estimating the distance to a local minimizer of a nonlinear programming problem. One estimate, based on the solution of a constrained convex quadratic program, can be used when strict complementary slackness and the second-order sufficient optimality conditions hold. A second estimate, based on the solution of an unconstrained nonconvex, nonsmooth optimization problem, is valid even when strict complementary slackness is violated. Both estimates are valid in a neighborhood of a local minimizer. An active set algorithm is developed for computing a stationary point of the nonsmooth error estimator. Each iteration of the algorithm requires the solution of a symmetric, positive semidefinite linear system, followed by a line search. Convergence is achieved in a finite number of iterations. The error bounds are based on stability properties for nonlinear programs. The theory is illustrated by some numerical examples.

Keywords Error bounds · KKT conditions · Active set algorithm · Nonconvex quadratic programming · Nonlinear programming

1 Introduction

For a nonlinear system of equations $\mathbf{h}(\mathbf{x}) = \mathbf{0}$, where $\mathbf{h} : \mathbb{R}^n \to \mathbb{R}^m$, the residual norm $\|\mathbf{h}(\mathbf{x})\|$ is often used to assess how close $\mathbf{x}$ is to a solution of the system. For an unconstrained optimization problem $\min\{f(\mathbf{x})\}$, where $f : \mathbb{R}^n \to \mathbb{R}$, the norm of the gradient $\|\nabla f(\mathbf{x})\|$ is often used to assess how close $\mathbf{x}$ is to a local minimizer of f. In this paper, we focus

November 23, 2013, Revised March 25, 2014.

W. W. Hager (✉) · D. Mico-Umutesi
Department of Mathematics, University of Florida, PO Box 118105,
Gainesville, FL 32611-8105, USA
e-mail: hager@ufl.edu
URL: http://people.clas.ufl.edu/hager/

D. Mico-Umutesi
e-mail: delmico@gmail.com

on the more difficult problem of estimating the distance to a solution of a general nonlinear programming problem. Error estimation in this context is complicated since one needs to take into account the gradient of the Lagrangian, constraint violations, and the complementary slackness condition. We consider optimization problems in the form

$$\min\{f(\mathbf{x}) : \mathbf{h}(\mathbf{x}) = \mathbf{0},\ \mathbf{x} \geq \mathbf{0},\ \mathbf{x} \in \mathbb{R}^n\}, \tag{1.1}$$

where f and $\mathbf{h}$ are continuously differentiable. A general nonlinear program with equality and inequality constraints can always be written in this form. Alternatively, throughout the paper the constraint $\mathbf{x} \geq \mathbf{0}$ can be replaced by the more general constraint $\mathbf{g}(\mathbf{x}) \geq \mathbf{0}$, where $\mathbf{g} : \mathbb{R}^n \to \mathbb{R}^l$, without significant changes in either the results or the analysis—simply replace x_i by $g_i(\mathbf{x})$. We assume that (1.1) has a local minimizer $\mathbf{x}^*$, and we wish to estimate the distance from an approximation $\mathbf{x}$ to $\mathbf{x}^*$.

Let $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^n \to \mathbb{R}$ be the Lagrangian defined by

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = f(\mathbf{x}) + \boldsymbol{\lambda}^{\mathsf{T}}\mathbf{h}(\mathbf{x}) - \boldsymbol{\mu}^{\mathsf{T}}\mathbf{x},$$

where $^{\mathsf{T}}$ denotes transpose. If a constraint qualification holds at $\mathbf{x}^*$, then there exists $\boldsymbol{\lambda} \in \mathbb{R}^m$ and $\boldsymbol{\mu} \in \mathbb{R}^n$ such that the Karush/Kuhn-Tucker (KKT) conditions

$$\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \mathbf{0}, \quad \boldsymbol{\mu} \geq \mathbf{0}, \quad \boldsymbol{\mu}^{\mathsf{T}}\mathbf{x} = 0, \quad \mathbf{h}(\mathbf{x}) = \mathbf{0}, \quad \text{and } \mathbf{x} \geq \mathbf{0} \tag{1.2}$$

are satisfied at $\mathbf{x} = \mathbf{x}^*$. Here ∇_x is the gradient with respect to $\mathbf{x}$, and any solution of (1.2) is referred to as a KKT point. Constraint qualifications include the widely used linear independence constraint qualification LICQ (for example, see Nocedal and Wright [14]), the Mangasarian-Fromovitz constraint qualification [11], the constant rank constraint qualification [9], the constant positive linear dependence constraint qualification [16], the relaxed constant rank constraint qualification [13], and the relaxed constant positive linear dependence constraint qualification [2].

Our goal is to develop an error estimator $E(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ with the property that

$$\|\mathbf{x} - \mathbf{x}^*\| + \|\boldsymbol{\lambda} - \hat{\boldsymbol{\lambda}}\| + \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\| \leq cE(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \tag{1.3}$$

for all $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ in a neighborhood $\mathcal{N}$ of a KKT point $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$, where $(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\mu}})$ is the projection of $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ on the set of KKT multipliers at $\mathbf{x}^*$ and c is a constant. Throughout the paper, we use the Euclidean norm $\|\cdot\|$. There is a substantial body of literature concerning error estimation. Pang's comprehensive survey [15] gives an overview of research up to 1997. Much of the earlier work focused on error bound conditions for problems with a particular structure, or for iterates produced by a specific algorithm. For example, in Luo and Tseng [10] the authors obtain error bounds for problems where the objective function is the composition of an affine mapping with a strongly convex differentiable function, and the constraint is either a polyhedron or a box. In Mathias and Pang [12], the authors derive an error bound for a linear complementarity problem involving a P matrix, while in García-Esnaola and Peña [5] the authors consider an H matrix. In the application section of Andreani et al. [2] and in Eckstein and Silva [4], the authors develop error bound conditions associated with iterates that arise in an augmented Lagrangian algorithm. General results concerning the convergence to KKT points are developed in Andreani et al. [1], where the authors formulate approximate KKT conditions and approximate gradient projection conditions.

In this paper, our focus is on error estimators that satisfy (1.3) for a general nonlinear optimization problem with equality and inequality constraints. Our first estimate, denoted E_0, is defined by

$$E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sqrt{\|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\|^2 + \|\mathbf{h}(\mathbf{x})\|^2 + (\boldsymbol{\mu}^{\mathsf{T}}\mathbf{x})^2},$$

where we require that $\mathbf{x} \geq \mathbf{0}$ and $\boldsymbol{\mu} \geq \mathbf{0}$. If an approximation $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ to the Lagrange multipliers is known, with $\boldsymbol{\mu} \geq \mathbf{0}$, then we could directly evaluate $E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ to estimate the distance to a local minimizer. Alternatively, we could estimate both the KKT multipliers and the error at an approximation $\mathbf{x}$ to a solution of (1.1) by solving the problem

$$\min\{E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})^2 : \boldsymbol{\lambda} \in \mathbb{R}^m, \ \boldsymbol{\mu} \in \mathbb{R}^n, \ \boldsymbol{\mu} \geq \mathbf{0}\}. \tag{1.4}$$

If $E_0(\mathbf{x})^2$ denotes the optimal objective value in (1.4), then $E_0(\mathbf{x})$, in essence, yields the tightest upper bound in (1.3).

Observe that $E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \geq 0$ and $E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = 0$ if and only $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ satisfy the KKT conditions (1.2). If $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfies the KKT conditions (1.2) and the second-order sufficient optimality conditions hold, then we show that E_0 has the error estimation property (1.3) when the strict complementary slackness condition is satisfied. That is, $x_i^* > 0$ when $\mu_i^* = 0$. If strict complementary slackness does not hold, then we only have $\|\mathbf{x} - \mathbf{x}^*\| \leq c\sqrt{E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})}$. Since (1.4) is a convex quadratic programming problem, it can be solved by many algorithms including either the primal active set algorithm (see Nocedal and Wright [14]) or the dual active set algorithm [7].

To handle problems where strict complementary slackness is violated, we have a slightly different error estimator. Let $\boldsymbol{\Phi} : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ be the componentwise min function defined by

$$\Phi_i(\mathbf{x}, \boldsymbol{\mu}) = \min\{x_i, \mu_i\}, \quad i = 1, 2, \ldots, n.$$

Our second estimate, denoted E_1, is given by

$$E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sqrt{\|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\|^2 + \|\mathbf{h}(\mathbf{x})\|^2 + \|\boldsymbol{\Phi}(\mathbf{x}, \boldsymbol{\mu})\|^2}.$$

Note that for E_1, there are no sign constraints on either $\boldsymbol{\mu}$ or $\mathbf{x}$. The KKT multiplier estimates can be obtained by solving the unconstrained optimization problem

$$\min\{E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})^2 : \boldsymbol{\lambda} \in \mathbb{R}^m, \ \boldsymbol{\mu} \in \mathbb{R}^n\}. \tag{1.5}$$

As with E_0, we see that $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \geq 0$ and $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = 0$ if and only $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ satisfy the KKT conditions (1.2). In particular, when $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = 0$, we must have $\boldsymbol{\Phi}(\mathbf{x}, \boldsymbol{\mu}) = \mathbf{0}$, which implies that the minimum of x_i and μ_i is 0 for each i. As a consequence, $x_i \geq 0$, $\mu_i \geq 0$, and $x_i \mu_i = 0$. Hence, all the conditions in (1.2) are satisfied when $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = 0$. We will show that E_1 has the property (1.3) even when strict complementary slackness does not hold. Again, we let $E_1(\mathbf{x})^2$ denote the optimal value of the objective in (1.5).

A fundamental difference between E_0 and E_1 is that the objective function associated with E_0 is convex, while the objective function of E_1 is piecewise convex, but globally nonconvex as seen in Fig. 1. Due to the nonconvexity, computing even a stationary point of (1.5) may not be easy. However, we could use a solution of the more tractable (1.4) as a starting guess for a solution of (1.5). An active set algorithm will be developed for handling the nonsmoothness of $\boldsymbol{\Phi}$ in (1.5).

Our error bounds are deduced from Theorem 1 in Hager and Gowda [6] where a stability result is established for constrained optimization problems. Here we show that if $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ is an approximation to a KKT point, then $E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ and $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ measure how much $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ violates the KKT conditions.

A different error estimate is developed by Dutta et al. in [3], where they consider a purely inequality constrained optimization problem of the form

$$\min\{f(\mathbf{x}) : \mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \ \mathbf{x} \in \mathbb{R}^n\}, \quad \mathbf{g} : \mathbb{R}^n \to \mathbb{R}^l. \tag{1.6}$$

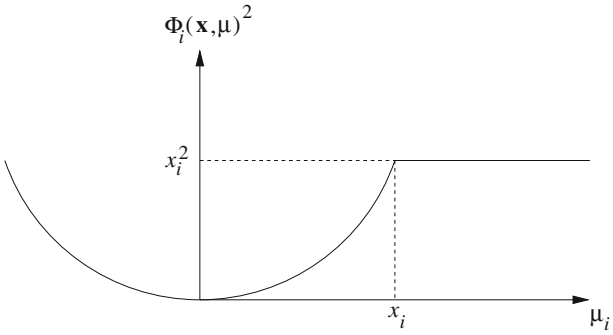

Fig. 1 A plot of $\Phi_i(\mathbf{x}, \boldsymbol{\mu})^2$ versus μ_i showing the piecewise convex, but globally nonconvex structure

For any feasible $\mathbf{x}$, their error estimate, which we denote E_D, is given by

$$E_D(\mathbf{x}) = \min\{\delta : \|\nabla_x[f(\mathbf{x}) + \boldsymbol{\mu}^\mathsf{T}\mathbf{g}(\mathbf{x})]\|^2 \le \delta,\ \boldsymbol{\mu}^\mathsf{T}\mathbf{g}(\mathbf{x}) \ge -\delta,\ \boldsymbol{\mu} \ge \mathbf{0},\ \boldsymbol{\mu} \in \mathbb{R}^m\}.$$

The estimate E_D is related to our E_0 estimate in the sense that they both contain a Lagrangian term and a term connected with the violation of complementary slackness. Nonetheless, there are fundamental differences between the error estimators. For example, if the optimization problem is unconstrained and $\mathbf{g}$ does not exist, then $E_D(\mathbf{x}) = \|\nabla f(\mathbf{x})\|^2$, while $E_0(\mathbf{x}) = \|\nabla f(\mathbf{x})\|$. Since E_D is the square of E_0 and since E_0 satisfies (1.3) when the Hessian $\nabla^2 f(\mathbf{x}^*)$ is positive definite at an unconstrained local minimizer $\mathbf{x}^*$, it follows that in general, E_D does not satisfy (1.3); more precisely, we have $\|\mathbf{x} - \mathbf{x}^*\| \le c\sqrt{E_D(\mathbf{x})}$ for this unconstrained optimization problem.

In E_D the Lagrangian error and the complementary slackness error are treated in an unsymmetric fashion since the square of the Lagrangian error is bounded in terms of δ, while the complementary slackness error is not squared. On the other hand, in E_0 both the Lagrangian error and the complementary slackness error are squared. Another difference between E_0 and E_D is that the objective of E_0 is quadratic, while E_D has a quadratic constraint, which is not as easy to handle numerically. In Dutta et al. [3], $E_D(\mathbf{x})$ was evaluated using a real-coded evolutionary optimization algorithm. In Theorem 3.6 of Dutta et al. [3], the authors show that if $E_D(\mathbf{x}_k)$ tends to zero, then under suitable assumptions, convergent subsequences of the $\mathbf{x}_k$ approach KKT points. In this paper, on the other hand, we focus on error estimation results of the form (1.3).

The paper is organized as follows: In Sect. 2 we review results in Hager and Gowda [6] concerning the stability of mathematical programs. In Sect. 3 we use results from Hager and Gowda [6] to obtain an estimate of the form (1.3) for E_1, while Sect. 4 analyzes the E_0 estimate. In Sect. 5 we show that although E_1 is nonsmooth, a minimizer of E_1 lies at a point of differentiability for the objective provided $\mathbf{x} \ge \mathbf{0}$. In Sect. 6 we develop an active set algorithm for computing an extreme point of E_1, while Sect. 7 gives numerical examples exploring the tightness of the theory.

Notation *∇f denotes the gradient of f, a row vector, and $\nabla^2 f$ is the Hessian of f. $\|\cdot\|$ denotes the Euclidean norm and $^\mathsf{T}$ denotes transpose. If $\mathcal{A}$ is a set, then $\mathcal{A}^c$ is the complement of $\mathcal{A}$. In an algorithm, a k subscript is often used to denote the iteration number. In particular, $\boldsymbol{\mu}_k$ would be the kth iterate in the variable $\boldsymbol{\mu}$, and μ_{ki} is the ith component of the kth iterate.*

2 A stability result for nonlinear programming

Our analysis of the error estimators is based on the stability analysis given in Hager and Gowda [6] for the solution of a nonlinear programming problem. Since the notation in our paper is different from the notation in the earlier work, we restate a key result from Hager and Gowda [6] using the notation in our paper. Let $\mathcal{N}(\boldsymbol{\mu})$ be the normal cone defined by

$$\mathcal{N}(\boldsymbol{\mu}) = \begin{cases} \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq \mathbf{0},\ \mathbf{x}^{\mathsf{T}}\boldsymbol{\mu} = 0\} & \text{if } \boldsymbol{\mu} \geq \mathbf{0}, \\ \emptyset & \text{otherwise.} \end{cases}$$

Let $\mathbf{T} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^n \to \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^n$ and $\mathcal{F} : \mathbb{R}^n \to 2^{\mathbb{R}^n} \times \mathbb{R}^m \times \mathbb{R}^n$ be defined by

$$\mathbf{T}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \begin{bmatrix} \mathbf{x} \\ \mathbf{h}(\mathbf{x}) \\ \nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \end{bmatrix} \quad \text{and} \quad \mathcal{F}(\boldsymbol{\mu}) = \begin{bmatrix} \mathcal{N}(\boldsymbol{\mu}) \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}. \tag{2.1}$$

With this notation, the KKT conditions are satisfied if and only if

$$\mathbf{T}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{F}(\boldsymbol{\mu}). \tag{2.2}$$

In the context of our optimization problem (1.1), Theorem 1 in Hager and Gowda [6] reduces to the following stability result:

Theorem 2.1 *Suppose that* $\mathbf{x}^*$ *is a local minimizer of (1.1) and the following conditions hold:*

A1. *f and h are twice differentiable at* $\mathbf{x}^*$.
A2. *There exist* $\boldsymbol{\lambda}^*$ *and* $\boldsymbol{\mu}^*$ *satisfying* $\mathbf{T}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) \in \mathcal{F}(\boldsymbol{\mu}^*)$.
A3. *The second-order sufficient optimality conditions hold; that is, there exists* $\alpha > 0$ *such that*

$$\mathbf{w}^{\mathsf{T}} \nabla_x^2 \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)\mathbf{w} \geq \alpha \|\mathbf{w}\|^2$$

whenever

$$\nabla \mathbf{h}(\mathbf{x}^*)\mathbf{w} = \mathbf{0}, \quad w_i = 0 \ \text{ if } \ \mu_i^* > 0, \quad w_i \geq 0 \ \text{ if } \ \mu_i^* = x_i^* = 0.$$

Then there exists a neighborhood $\mathcal{N}$ *of* $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ *and constants* γ *and* δ *with the property that for each* $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$*, we have*

$$\|\mathbf{x} - \mathbf{x}^*\| + \|\boldsymbol{\lambda} - \hat{\boldsymbol{\lambda}}\| + \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\| \leq \|\mathbf{x} - \mathbf{x}_0\| + \|\boldsymbol{\lambda} - \boldsymbol{\lambda}_0\| + \|\boldsymbol{\mu} - \boldsymbol{\mu}_0\| + \gamma \|\mathbf{y}\|$$

for all $(\mathbf{x}_0, \boldsymbol{\lambda}_0, \boldsymbol{\mu}_0) \in \mathcal{N}$ *satisfying*

$$\mathbf{T}(\mathbf{x}_0, \boldsymbol{\lambda}_0, \boldsymbol{\mu}_0) + \mathbf{y} \in \mathcal{F}(\boldsymbol{\mu}_0), \quad \|\mathbf{y}\| \leq \delta, \tag{2.3}$$

where $(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\mu}})$ *is the projection of* $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ *on the set of KKT multipliers* $\mathcal{M}$ *defined by*

$$\mathcal{M} = \{(\boldsymbol{\lambda}, \boldsymbol{\mu}) : \mathbf{T}(\mathbf{x}^*, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{F}(\boldsymbol{\mu})\}.$$

The condition (2.3) is the KKT conditions (2.2) perturbed by $\mathbf{y}$. Theorem 2.1 basically says that under the assumptions of the theorem, the error at any approximate solution $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ can be bounded by the perturbations $\mathbf{x} - \mathbf{x}_0$, $\boldsymbol{\lambda} - \boldsymbol{\lambda}_0$, $\boldsymbol{\mu} - \boldsymbol{\mu}_0$, and $\mathbf{y}$ needed to move from the approximate solution to a point $(\mathbf{x}_0, \boldsymbol{\lambda}_0, \boldsymbol{\mu}_0)$ satisfying the KKT conditions perturbed by $\mathbf{y}$. We use Theorem 2.1 to analyze the error estimate connected with either E_0 or E_1.

3 Error associated with $\mathbf{E_1}$

Since the analysis of E_1 is more straightforward than the analysis of E_0, we begin with E_1. Suppose that the assumptions of Theorem 2.1 are satisfied, and let $\mathcal{N}$ be a neighborhood of $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ which is small enough that Theorem 2.1 holds. For any given $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$, let ϵ be defined by

$$\epsilon = E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sqrt{\|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\|^2 + \|\mathbf{h}(\mathbf{x})\|^2 + \|\boldsymbol{\Phi}(\mathbf{x}, \boldsymbol{\mu})\|^2}. \tag{3.1}$$

In order to apply Theorem 2.1, we try to choose $(\mathbf{x}_0, \boldsymbol{\lambda}_0, \boldsymbol{\mu}_0) \in \mathcal{N}$ near $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ and $\mathbf{y}$ near $\mathbf{0}$ satisfying (2.3). The $\mathbf{T}$ function defined in (2.1) has three blocks of components corresponding to $\mathbf{x}$, $\mathbf{h}(\mathbf{x})$, and $\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$. We partition $\mathbf{y}$ into three corresponding subvectors denoted $\mathbf{y}_1$, $\mathbf{y}_2$, and $\mathbf{y}_3$. Hence, to apply Theorem 2.1, we need to specify $\mathbf{x}_0$, $\boldsymbol{\lambda}_0$, $\boldsymbol{\mu}_0$, $\mathbf{y}_1$, $\mathbf{y}_2$, and $\mathbf{y}_3$. We take

$$\mathbf{x}_0 = \mathbf{x}, \quad \boldsymbol{\lambda}_0 = \boldsymbol{\lambda}, \quad \text{and} \quad \mathbf{y}_2 = -\mathbf{h}(\mathbf{x}). \tag{3.2}$$

The parameters $\boldsymbol{\mu}_0$ and $\mathbf{y}_1$ must be chosen so that $\boldsymbol{\mu}_0^{\mathsf{T}}(\mathbf{x} + \mathbf{y}_1) = 0$, where $\boldsymbol{\mu}_0 \geq \mathbf{0}$, and $\mathbf{x} + \mathbf{y}_1 \geq \mathbf{0}$. We accomplish this with the following choices:

(a) $\mu_{0i} = \mu_i$ and $y_i = -x_i$ if $x_i < \mu_i$ and $\mu_i \geq 0$.
(b) $\mu_{0i} = 0$ and $y_i = -x_i$ if $x_i < \mu_i < 0$.
(c) $\mu_{0i} = 0$ and $y_i = 0$ if $\mu_i \leq x_i$ and $x_i \geq 0$.
(d) $\mu_{0i} = 0$ and $y_i = -x_i$ if $\mu_i \leq x_i < 0$.

With these choices, either $\mu_{0i} = 0$ or $y_i = -x_i$ for each i, and when this holds, we have $\mu_{0i}(x_i + y_i) = 0$. Moreover, $\mu_{0i} \geq 0$ and $x_i + y_i \geq 0$ as we now show. If $\mu_i \leq x_i$, then $\mu_{0i} = 0$ in (c) and (d). If $\mu_i \leq x_i$ and $x_i \geq 0$, then $y_i = 0$ in (c) and $x_i + y_i = x_i \geq 0$, while $y_i = -x_i$ in (d) and $x_i + y_i = 0$. If $x_i < \mu_i$, then in (a) and (b), $y_i = -x_i$ and $x_i + y_i = 0$; moreover, $\mu_{0i} = \mu_i \geq 0$ in (a), while $\mu_{0i} = 0$ in (b).

Finally, we choose

$$\mathbf{y}_3 = -\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}_0). \tag{3.3}$$

With the parameter choices given in (3.2), (3.3), and (a)–(d), the inclusion

$$\mathbf{T}(\mathbf{x}_0, \boldsymbol{\lambda}_0, \boldsymbol{\mu}_0) + \mathbf{y} \in \mathcal{F}(\boldsymbol{\mu}_0)$$

is satisfied.

By the choices in (3.2) and by the definition of ϵ in (3.1), we have

$$\|\mathbf{x} - \mathbf{x}_0\| = 0, \quad \|\boldsymbol{\lambda} - \boldsymbol{\lambda}_0\| = 0, \quad \text{and} \quad \|\mathbf{y}_2\| \leq \epsilon. \tag{3.4}$$

By the choices in (a)–(d), we have

$$\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\|^2 = \sum_{\mu_i \leq x_i} \mu_i^2 + \sum_{x_i < \mu_i < 0} \mu_i^2 \leq \sum_{\mu_i \leq x_i} \mu_i^2 + \sum_{x_i < \mu_i < 0} x_i^2, \tag{3.5}$$

and

$$\|\mathbf{y}_1\|^2 = \sum_{x_i < \mu_i} x_i^2 + \sum_{\mu_i \leq x_i < 0} x_i^2 \leq \sum_{x_i < \mu_i} x_i^2 + \sum_{\mu_i \leq x_i < 0} \mu_i^2. \tag{3.6}$$

By the definition of $\boldsymbol{\Phi}$, we have $\Phi_i(\mathbf{x}, \boldsymbol{\mu}) = \mu_i$ when $\mu_i \leq x_i$, and $\Phi_i(\mathbf{x}, \boldsymbol{\mu}) = x_i$ when $x_i < \mu_i$. Hence, (3.5) and (3.6) yield

$$\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\|^2 \leq \sum_{i=1}^{n} \Phi_i(\mathbf{x}, \boldsymbol{\mu})^2 \quad \text{and} \quad \|\mathbf{y}_1\|^2 \leq \sum_{i=1}^{n} \Phi_i(\mathbf{x}, \boldsymbol{\mu})^2. \tag{3.7}$$

Since $\|\boldsymbol{\Phi}(\mathbf{x}, \boldsymbol{\mu})\| \leq \epsilon$, it follows from (3.7) that

$$\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\| \leq \|\boldsymbol{\Phi}(\mathbf{x}, \boldsymbol{\mu})\| \leq \epsilon \quad \text{and} \quad \|\mathbf{y}_1\| \leq \|\boldsymbol{\Phi}(\mathbf{x}, \boldsymbol{\mu})\| \leq \epsilon. \tag{3.8}$$

Finally, let us bound $\mathbf{y}_3$. Observe that

$$-\mathbf{y}_3 = \nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}_0) = \nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) + (\boldsymbol{\mu} - \boldsymbol{\mu}_0)^{\mathsf{T}}. \tag{3.9}$$

Since $\|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\| \leq \epsilon$ by the definition of ϵ, and since $\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\| \leq \epsilon$ by (3.8), it follows that

$$\|\mathbf{y}_3\| \leq 2\epsilon. \tag{3.10}$$

Combine (3.4), (3.8), and (3.10) to obtain $\|\mathbf{y}\| \leq 4\epsilon$. As $\mathcal{N}$ shrinks around $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$, ϵ tends to 0 since $E_1(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = 0$ and $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$. Choose $\mathcal{N}$ small enough that $4\epsilon \leq \delta$ whenever $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$, where δ is the constant appearing in Theorem 2.1. We combine (3.4), (3.8), (3.10) and invoke Theorem 2.1 to obtain the following result:

Theorem 3.1 *If $\mathbf{x}^*$ is a local minimizer of (1.1) and* (A1)–(A3) *are satisfied, then there is a neighborhood $\mathcal{N}$ of $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ with the property that for each $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$, we have*

$$\|\mathbf{x} - \mathbf{x}^*\| + \|\boldsymbol{\lambda} - \hat{\boldsymbol{\lambda}}\| + \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\| \leq (1 + 4\gamma) E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}),$$

where γ is the constant in Theorem 2.1, and $(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\mu}})$ is the projection of $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ on the set of KKT multipliers at $\mathbf{x}^*$.

4 Error associated with E_0

For the error estimator E_0, we no longer have the $\boldsymbol{\Phi}$ term that we used to bound $\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\|$ and $\|\mathbf{y}_1\|$ in (3.8). However, a similar bound is obtained when strict complementary slackness holds. That is, $x_i^* > 0$ whenever $\mu_i^* = 0$. Choose $\sigma > 0$ small enough that

$$2\sigma < \min\{x_i^* : x_i^* > 0\} \quad \text{and} \quad 2\sigma < \min\{\mu_i^* : \mu_i^* > 0\}.$$

Let $\mathcal{N}$ be a neighborhood of $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ which is small enough that Theorem 2.1 holds and small enough that

$$\max\{x_i : x_i^* = 0\} < \sigma < \min\{x_i : x_i^* > 0\} \tag{4.1}$$

and

$$\max\{\mu_i : \mu_i^* = 0\} < \sigma < \min\{\mu_i : \mu_i^* > 0\} \tag{4.2}$$

for all $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$. For any given $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$ with $\boldsymbol{\mu} \geq \mathbf{0}$ and $\mathbf{x} \geq \mathbf{0}$, let ϵ now be defined by

$$\epsilon = E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sqrt{\|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\|^2 + \|\mathbf{h}(\mathbf{x})\|^2 + (\boldsymbol{\mu}^{\mathsf{T}}\mathbf{x})^2}. \tag{4.3}$$

By the definition of E_0, we see that $\boldsymbol{\mu}^{\mathsf{T}}\mathbf{x} \leq \epsilon$. Since $\mu_i x_i \geq 0$, we also have

$$\mu_i x_i \leq \epsilon. \tag{4.4}$$

We make exactly the same choices for $\mathbf{x}_0$, $\boldsymbol{\lambda}_0$, $\boldsymbol{\mu}_0$, and $\mathbf{y}$ that we made in the analysis of E_1. The only difference is that we now need to bound $\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\|$ and $\mathbf{y}_1$ in a different way. Suppose that $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$, $\mathbf{x} \geq \mathbf{0}$, and $\boldsymbol{\mu} \geq \mathbf{0}$. If $\mu_i \leq x_i$, then we consider two cases:

Case 1: $\mu_i^* = 0$. By strict complementary slackness, $x_i^* > 0$, and by (4.1), $x_i > \sigma$. Since $\mu_i x_i \leq \epsilon$, we conclude that $\mu_i \leq \epsilon/\sigma$.

Case 2: $\mu_i^* > 0$. By complementary slackness, $x_i^* = 0$. By (4.1), $x_i < \sigma$ and by (4.2), $\mu_i > \sigma$. Hence $\mu_i > x_i$, which violates the assumption that $\mu_i \leq x_i$. This implies that Case 2 cannot occur.

If $\mu_i > x_i$, then we consider the following two cases:

Case 3: $x_i^* = 0$. By strict complementary slackness, $\mu_i^* > 0$, and by (4.2), $\mu_i > \sigma$. Since $\mu_i x_i \leq \epsilon$, we conclude that $x_i \leq \epsilon/\sigma$.

Case 4: $x_i^* > 0$. By complementary slackness, $\mu_i^* = 0$. By (4.2), $\mu_i < \sigma$ and by (4.1), $x_i > \sigma$. Hence $x_i > \mu_i$, which violates the assumption that $\mu_i > x_i$. This implies that Case 4 cannot occur.

We return to the first equality in (3.5). The second sum is vacuous since $\mu_i \geq 0$. According to Case 1, $\mu_i \leq \epsilon/\sigma$ when $\mu_i \leq x_i$. Hence, (3.5) yields

$$\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\| \leq \epsilon\sqrt{n}/\sigma. \tag{4.5}$$

Similarly, for the first equation in (3.6), the second sum is vacuous since $x_i \geq 0$. According to Case 3, $x_i \leq \epsilon/\sigma$ when $\mu_i > x_i$. Hence, (3.6) yields

$$\|\mathbf{y}_1\| = \epsilon\sqrt{n}/\sigma. \tag{4.6}$$

Finally, let us bound $\mathbf{y}_3$. Since $\|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\| \leq \epsilon$ and $\|\boldsymbol{\mu}_0 - \boldsymbol{\mu}\| \leq \epsilon\sqrt{n}/\sigma$, it follows from (3.9) that

$$\|\mathbf{y}_3\| \leq \epsilon(1 + \sqrt{n}/\sigma). \tag{4.7}$$

Combine (3.4), (4.6), and (4.7) to obtain $\|\mathbf{y}\| \leq 2\epsilon(1 + \sqrt{n}/\sigma)$. As $\mathcal{N}$ shrinks around $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$, ϵ tends to 0 since $E_0(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = 0$ and $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$. Choose $\mathcal{N}$ small enough that $\|\mathbf{y}\| \leq \delta$, the constant appearing in Theorem 2.1. Again, combine (3.4), (4.5), (4.6), and (4.7) and invoke Theorem 2.1 to obtain the following result:

Theorem 4.1 *If $\mathbf{x}^*$ is a local minimizer of (1.1),* (A1)–(A3) *are satisfied, and strict complementary slackness holds, then there is a neighborhood $\mathcal{N}$ of $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ with the property that for each $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$ with $\mathbf{x} \geq \mathbf{0}$ and $\boldsymbol{\mu} \geq \mathbf{0}$, we have*

$$\|\mathbf{x} - \mathbf{x}^*\| + \|\boldsymbol{\lambda} - \hat{\boldsymbol{\lambda}}\| + \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\| \leq cE_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}),$$

where $c = \sqrt{n}/\sigma + 2(1 + \sqrt{n}/\sigma)\gamma$, σ is the constant in (4.1) and (4.2), γ is the constant in Theorem 2.1, and $(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\mu}})$ is the projection of $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ on the set of KKT multipliers at $\mathbf{x}^*$.

An example is given in Sect. 7 which shows that if the strict complementary slackness requirement does not hold, then E_0 may no longer satisfy the error bound condition (1.3). When strict complementary slackness does not hold, the $O(\epsilon)$ bounds for $\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\|$ and $\|\mathbf{y}_1\|$ must be replaced by $O(\sqrt{\epsilon})$ as we now show. Referring to (3.5), the sum indexed by $\mu_i < 0$ is again vacuous when $\mu_i \geq 0$. For the sum indexed by $x_i \geq \mu_i \geq 0$, we utilize the relation (4.4) to obtain $\mu_i^2 \leq \epsilon$, which implies that

$$\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\| \leq \sqrt{n\epsilon}.$$

In a similar fashion, in (3.6), the sum indexed by $x_i < 0$ is again vacuous when $x_i \geq 0$. For the sum indexed by $\mu_i > x_i \geq 0$, we utilize the relation (4.4) to obtain $x_i^2 \leq \epsilon$, which implies that

$$\|\mathbf{y}_1\| \leq \sqrt{n\epsilon}.$$

The $O(\sqrt{\epsilon})$ bounds for $\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\|$ and $\|\mathbf{y}_1\|$ yield the following result when strict complementary slackness does not hold:

Theorem 4.2 *If* $\mathbf{x}^*$ *is a local minimizer of (1.1) and* (A1)–(A3) *are satisfied, then there is a neighborhood* $\mathcal{N}$ *of* $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ *with the property that for each* $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{N}$ *with* $\mathbf{x} \geq \mathbf{0}$ *and* $\boldsymbol{\mu} \geq \mathbf{0}$*, we have*

$$\|\mathbf{x} - \mathbf{x}^*\| + \|\boldsymbol{\lambda} - \hat{\boldsymbol{\lambda}}\| + \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\| \leq c\sqrt{E_0(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})},$$

where c *is independent of* $(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ *and* $(\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\mu}})$ *is the projection of* $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ *on the set of KKT multipliers at* $\mathbf{x}^*$.

5 Differentiability at a minimizer of E_1

Often at a solution of a nonsmooth optimization problem, the objective function is nondifferentiable. The nonsmoothness associated with the E_1 error estimator, on the other hand, is quite different. The analysis in this section implies that if $\mathbf{x} \geq \mathbf{0}$, then at a minimum for (1.5), the objective function is differentiable. In fact, the following more general property is established:

> *Differentiablity Property.* If $\mathbf{x} \geq \mathbf{0}$ and $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ is minimized over μ_i with all the other variables fixed, then the minimizing μ_i is a point where E_1^2 is differentiable.

Since the nonsmoothness of E_1^2 is associated with the $\boldsymbol{\Phi}(\mathbf{x}, \boldsymbol{\mu})$ term in the objective, and since a minimizer in (1.5) must also achieve a minimum along each of the coordinate directions, the Differentiability Property implies that E_1^2 is differentiable at a solution of (1.5).

The Differentiability Property is significant for the following reason: Since E_1 is nonsmooth, solving (1.5) could be nontrivial, and potentially bundle methods and nonsmooth optimization techniques would be required. The Differentiability Property implies that although E_1 is nonsmooth, it is smooth at a minimizer where subdifferentials can be replaced by ordinary derivatives. In the next section, we show that although E_1 is nonsmooth, we can find an extreme point by solving a series of linear systems, and we never need to be concerned with the lack of smoothness.

In order to establish the Differentiablity Property, let us focus on the scalar problem

$$\min_{\mu} \ \phi(\mu), \quad \text{where } \phi(\mu) = (\mu - a)^2 + s \min\{\mu, x\}^2. \tag{5.1}$$

Here a and x are arbitrary parameters and $s > 0$. When minimizing $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ over a component of $\boldsymbol{\mu}$, we essentially solve a problem of the form (5.1). In the E_1 error estimator, we took $s = 1$, but in general, we may wish to scale the $\boldsymbol{\Phi}$ term relative to the Lagrangian term so we included a scaling parameter s in the definition of ϕ.

Proposition 5.1 *Suppose that* $s > 0$. *If* $x \geq 0$, *then at the solution of (5.1),* ϕ *is differentiable for any choice of* a. *If* $x < 0$, *then at a solution of (5.1),* ϕ *is differentiable except when* $a \in [(1 + s)x, x]$.

Proof If $x = 0$, then ϕ is differentiable at any μ, so in particular, ϕ is differentiable at the solution of (5.1). Hence, we assume that $x \neq 0$. In this case, which corresponds to Fig. 1, the only point where ϕ is nondifferentiable is $\mu = x$. To prove the proposition, we need to show that the minimizer in (5.1) is never $\mu = x$ except when $x < 0$ and $a \in [(1 + s)x, x]$.

In the region $\mu \leq x$, we have $\phi(\mu) = (\mu - a)^2 + s\mu^2$. The minimum of this quadratic is $\mu_- = a/(1 + s)$. We consider three cases:

Case 1: $a > x$. In this case $\phi'(\mu) = 2(\mu - a) < 0$ for $\mu \in [x, a)$. Hence, the minimum of ϕ cannot be attained at $\mu = x$.

Case 2: $a \le x$ and $x \ge 0$. For $\mu > x$, we have $\phi(\mu) = (\mu - a)^2 + sx^2$. Hence, $\phi'(\mu) = 2(\mu - a) > 0$ for $\mu > x \ge a$. Consequently, the minimum of ϕ is attained on the interval $(-\infty, x]$ at $\mu_- = a/(1+s) \le x/(1+s) < x$ since $s > 0$ and $x \ge 0$.

Case 3: $a \le x < 0$. As in Case 2, the minimum of ϕ is attained on the interval $(-\infty, x]$. As noted earlier, $\phi(\mu) = (\mu - a)^2 + s\mu^2$ on this interval, and this quadratic achieves its minimum at $\mu = \mu_- = a/(1+s)$. If $\mu_- < x$, then the minimum is attained at $\mu = \mu_- < x$. If $\mu_- \ge x$, then the minimum is attained at $\mu = x$. Since it was assumed that $a \le x$, we conclude that the minimum is attained at $\mu = x$ if and only if both inequalities $a \le x$ and $\mu_- \ge x$ are satisfied; equivalently, $a \in [(1+s)x, x]$.

□

Since the Hessian of E_1 is not necessarily positive definite at a point of differentiability, we may wish to regularize the objective by adding a term of the form $\epsilon\|\boldsymbol{\mu}\|^2$, where $\epsilon > 0$. This leads us to consider the following generalization of (5.1):

$$\min_{\mu} \ \psi(\mu), \quad \text{where } \psi(\mu) = (\mu - a)^2 + s\min\{\mu, x\}^2 + \epsilon\mu^2.$$

After completing the square, ψ can be written as

$$\psi(\mu) = (1+\epsilon)\left(\mu - \frac{a}{1+\epsilon}\right)^2 + s\min\{\mu, x\}^2 + \frac{a^2\epsilon}{1+\epsilon}.$$

We divide by $1+\epsilon$ to obtain an objective in the same form as (5.1). Hence, Proposition 5.1 applies to the regularized problem, but with a in (5.1) replaced by $a/(1+\epsilon)$.

6 Active set algorithm for minimizing E_1

We now develop an active set algorithm for computing a stationary point of the nonconvex, nonsmooth problem (1.5). When minimizing $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})^2$ over $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$, the $\mathbf{h}(\mathbf{x})$ term is irrelevant, so we focus on the problem of minimizing

$$F(\boldsymbol{\lambda}, \boldsymbol{\mu}) = \|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\|^2 + \sum_{i=1}^{n} \min\{x_i, \mu_i\}^2,$$

where $\mathbf{x} \ge \mathbf{0}$ is given. Our algorithm involves a set $\mathcal{U}_k \subset \{1, 2, \ldots, n\}$, and two other functions

$$G_k(\boldsymbol{\lambda}, \boldsymbol{\mu}) = \|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\|^2 + \sum_{i \in \mathcal{U}_k} \mu_i^2 + \sum_{i \in \mathcal{U}_k^c} \min\{x_i, \mu_i\}^2,$$

and

$$H_k(\boldsymbol{\lambda}, \boldsymbol{\mu}) = \|\nabla_x \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})\|^2 + \sum_{i \in \mathcal{U}_k} \mu_i^2 + \sum_{i \in \mathcal{U}_k^c} x_i^2.$$

For any choice of $\mathcal{U}_k$, we have

$$F \le G_k \le H_k. \tag{6.1}$$

A formal statement of our algorithm for minimizing E_1 is given in Fig. 2. Since $H_k(\boldsymbol{\lambda}, \boldsymbol{\mu})$ is convex and quadratic in $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$, the minimizer $(\bar{\boldsymbol{\lambda}}, \bar{\boldsymbol{\mu}})$ of H_k computed in Step 1 is easily evaluated by solving a linear system of equations. If the system is singular, we could compute the minimum norm solution. In Step 2, we perform a line search along the ray connecting

Given $(\boldsymbol{\lambda}_0, \boldsymbol{\mu}_0)$, initialize $\mathcal{U}_0 = \{i : \mu_{0i} < x_i\}$ and $k = 0$.

Step 1. Compute $(\bar{\boldsymbol{\lambda}}, \bar{\boldsymbol{\mu}}) \in \arg\min \{H_k(\boldsymbol{\lambda}, \boldsymbol{\mu}) : \boldsymbol{\lambda} \in \mathbb{R}^m, \boldsymbol{\mu} \in \mathbb{R}^n\}$.

Step 2. Define $(\boldsymbol{\lambda}(t), \boldsymbol{\mu}(t)) = (\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) + t(\bar{\boldsymbol{\lambda}} - \boldsymbol{\lambda}_k, \bar{\boldsymbol{\mu}} - \boldsymbol{\mu}_k), \quad t \geq 0$.

Compute $t_k \in \arg\min \{G_k(\boldsymbol{\lambda}(t), \boldsymbol{\mu}(t)) : t \geq 0\}$.

Update $(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) = (\boldsymbol{\lambda}(t_k), \boldsymbol{\mu}(t_k))$.

Step 3. If $\mu_{ki}(t_k) < x_i$ for some $i \in \mathcal{U}_k^c$, then

$\mathcal{U}_{k+1} = \mathcal{U}_k \cup \{i \in \mathcal{U}_k^c : \mu_{ki}(t_k) < x_i\}$,

update $k \leftarrow k + 1$, and go to Step 1.

Step 4. Update $k \leftarrow k + 1$ and set $\boldsymbol{\mu} = \boldsymbol{\mu}_k$.

for $j = 1, 2, \ldots, n$

$\mu_j \in \arg\min\{F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}) : \mu_j \in \mathbb{R}\}$ with $\mu_j = \mu_{kj}$ if possible

end

Update $\boldsymbol{\mu}_{k+1} = \boldsymbol{\mu}$ and $\boldsymbol{\lambda}_{k+1} = \boldsymbol{\lambda}_k$.

Define $\mathcal{U}_{k+1} = \{i : \mu_i < x_i\}$.

Step 5. If $F(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) = F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$, then stop.

Otherwise $k \leftarrow k + 1$ and go to Step 1.

Fig. 2 Active set algorithm for minimizing $E_1(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ over $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$

the current iterate $(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ to $(\bar{\boldsymbol{\lambda}}, \bar{\boldsymbol{\mu}})$. The goal of the line search is to minimize G_k along the search direction. The minimizer becomes the next iterate $(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1})$. If the condition in Step 3 is satisfied, then $\mathcal{U}_k$ set is updated and we return to Step 1. Otherwise, we proceed to Step 4 where coordinate descent is applied to the multiplier $\boldsymbol{\mu}$, starting from $\boldsymbol{\mu} = \boldsymbol{\mu}_k$. We minimize over each component of the multiplier $\boldsymbol{\mu}$ while holding the other components fixed. If $\mu_j = \mu_{kj}$ attains the minimum, then we accept $\mu_j = \mu_{kj}$ as the minimizer for the jth component. After the coordinate descent phase, the set $\mathcal{U}_{k+1}$ is initialized in the same way that $\mathcal{U}_0$ was initialized at the start of the algorithm. Step 5 is the convergence test. If the algorithm has not converged, then we return to Step 1.

To prove the convergence of the algorithm to a stationary point, we will first show that

$$H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) = G_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) \tag{6.2}$$

at the start of Step 2. This identity is due to the fact that the updates to the set $\mathcal{U}_k$ throughout the algorithm always ensure that

$$\min\{x_i, \mu_{ki}\} = x_i \quad \text{for all } i \in \mathcal{U}_k^c. \tag{6.3}$$

When $\mathcal{U}_k$ satisfies this condition, the trailing terms in G_k and in H_k are the same and hence, (6.2) holds. The condition (6.3) holds initially for $k = 0$ by the definition of $\mathcal{U}_0$. Moreover, at Step 4, this condition holds by the definition of $\mathcal{U}_{k+1}$. Let us proceed by induction and assume that (6.3) holds in Step 1 at iteration k. In Step 3, we add to $\mathcal{U}_k$ those indices $i \in \mathcal{U}_k^c$ for which $\mu_{ki}(t_k) < x_i$. Hence, for the indices $i \in \mathcal{U}_{k+1}^c$, we must have

$$\min\{x_i, \mu_{ki}(t_k)\} = \min\{x_i, \mu_{(k+1),i}\} = x_i,$$

and (6.3) holds with k replaced by $k + 1$.

Since $G_k = H_k$ at $(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ by (6.2), $G_k \le H_k$ by (6.1), $(\bar{\boldsymbol{\lambda}}, \bar{\boldsymbol{\mu}})$ minimizes H_k, and $(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1})$ minimizes G_k in the search direction, we conclude that

$$G_k(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) = G_k(\bar{\boldsymbol{\lambda}}, \bar{\boldsymbol{\mu}}) \le H_k(\bar{\boldsymbol{\lambda}}, \bar{\boldsymbol{\mu}}) \le H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k). \tag{6.4}$$

By the definition of $\mathcal{U}_{k+1}$ in Step 3, we have

$$G_{k+1}(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) = G_k(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}). \tag{6.5}$$

It follows that

$$H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) \ge G_{k+1}(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) = H_{k+1}(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) \ge F(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}), \tag{6.6}$$

where the first inequality is (6.4) combined with (6.5), the next equality is (6.2), and the final inequality is (6.1). Since the sets $\mathcal{U}_k$ can only grow in Step 3, we eventually reach a set $\mathcal{U}_k$ in Step 2 with $\mu_{ki}(t_k) \ge x_i$ for all $i \in \mathcal{U}_k^c$. In this case $H_k = G_k$ on the line segment connecting $(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ to $(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1})$, and the algorithm advances to Step 4. We call this iterate where the algorithm advances to Step 4 a terminating iterate.

When $k = 0$, we have $H_0(\boldsymbol{\lambda}_0, \boldsymbol{\mu}_0) = F(\boldsymbol{\lambda}_0, \boldsymbol{\mu}_0)$ in Step 1 by the definition of $\mathcal{U}_0$. Similarly, at any iteration k that follows a terminating iterate, we have $H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) = F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ in Step 1 by the definition of $\mathcal{U}_{k+1}$ in Step 4. When the algorithm performs Step 1, Step 2, Step 3 and then branches back to Step 1, (6.6) yields $H_{k+1}(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) \le H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$. Hence, the identity $H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) = F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ at any iteration that follows a terminating iterate and the inequality (6.6) together imply that the value of $F(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1})$ when the algorithm leaves Step 4 is greater than or equal to its value when it later returns to the start of Step 4.

In Step 5, either $F(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) = F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ and the algorithm stops or

$$F(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) < F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k).$$

If the latter holds and the algorithm does not stop, then before the start of Step 4, we have $H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) \ge F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ by (6.1). Since $F(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) < F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ in Step 5, and since $H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) = F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ at the start of Step 1 in the next iteration, it follows from (6.6) that there is strict decrease in $H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ when the algorithm does not stop in Step 5. Observe that the terminating iterate is a minimizer of H_k for some k, and this minimizer only depends on the set $\mathcal{U}_k$. Since there is strict decrease in H_k when the algorithm does not stop in Step 5, the set $\mathcal{U}_k$ associated with a terminating iterate in Step 3 cannot repeat. Since there are a finite number of distinct $\mathcal{U}_k$ sets, we must eventually have $F(\boldsymbol{\lambda}_{k+1}, \boldsymbol{\mu}_{k+1}) = F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ and the algorithm stops. In this case, the current iterate achieves the minimum in Step 4 in each of the coordinate directions, and $H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) = F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$. Since $(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ minimizes the smooth function H_k, it follows that $\nabla_\lambda H_k(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) = \mathbf{0}$. Since the $\boldsymbol{\lambda}$ terms in F and H_k coincide, we conclude that $\nabla_\lambda F(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) = \mathbf{0}$. When the algorithm stops, the minimum of $F(\boldsymbol{\lambda}_k, \boldsymbol{\mu})$ in each of the coordinate direction μ_j is attained at μ_{kj}; consequently, by the Differentiability Property, the gradient of F vanishes at $(\boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$, and the current iterate is a stationary point of F. We summarize these observations as follows:

Theorem 6.1 *From any starting point $(\boldsymbol{\lambda}_0, \boldsymbol{\mu}_0)$, the active set algorithm for (1.5) reaches a stationary point in a finite number of iterations.*

Note that the E_0 estimator is the solution of a convex quadratic programming problem (1.4) for which a global minimizer is easily evaluated. Consequently, the global minimizer associated with the E_0 error bound could be a good starting guess for the nonsmooth problem (1.5). Although the minimizer $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ of F computed by the active set algorithm is only claimed to be a stationary point in Theorem 6.1, our analysis established the following

stronger property: $\boldsymbol{\lambda}^*$ minimizes $F(\boldsymbol{\lambda}, \boldsymbol{\mu}^*)$ over all $\boldsymbol{\lambda}$, and componentwise, $F(\boldsymbol{\lambda}^*, \boldsymbol{\mu})$ achieves its minimum over μ_j at μ_j^* for each j.

7 Numerical examples

We first wish to show that the strict complementary slackness assumption in Theorem 4.1 is needed to ensure an error bound of the form (1.3) for E_0. Let us consider the optimization problem

$$\min\left\{\frac{1}{2}x^2 : x \geq 0\right\}. \tag{7.1}$$

The optimal solution is $x^* = 0$ and the Lagrange multiplier associated with the constraint is $\mu = 0$. Hence, strict complementary slackness does not hold. The optimization problem (1.4) that we solve to evaluate $E_0(x)$ is

$$\min\{(x-\mu)^2 + (\mu x)^2 : \mu \geq 0\}. \tag{7.2}$$

The solution is $\mu = x/(1+x^2)$. Evaluating the objective in (7.2) at this value of μ and taking the square root gives

$$E_0(x) = \frac{x^2}{\sqrt{1+x^2}}.$$

For any given $x > 0$, we have $|x - x^*| = x$ since $x^* = 0$. Hence, we have

$$\frac{|x-x^*|}{E_0(x)} = \frac{\sqrt{1+x^2}}{x},$$

which approaches ∞ as x approaches 0. Hence, we cannot bound $|x - x^*|$ by $cE_0(x)$ for a fixed choice of c as x tends to 0. On the other hand, we do have $|x - x^*| \leq c\sqrt{E_0(x)}$, in compliance with Theorem 4.2.

To evaluate E_1, we solve the problem

$$\min\left\{(x-\mu)^2 + \min\{x, \mu\}^2\right\}.$$

The solution is $\mu = x/2$ if $x \geq 0$ and $\mu = x$ if $x < 0$. It follows that

$$E_1(x) = \begin{cases} x/\sqrt{2} & \text{if } x \geq 0, \\ |x| & \text{if } x < 0. \end{cases}$$

Hence, $|x - x^*| \leq \sqrt{2}E_1(x)$, and E_1 satisfies (1.3), while E_0 does not.

Next, we consider a variation of Problem 5 in Appendix A of Himmelblau [8]:

$$\min\{-(x_1(x_1+x_2+x_3) + 2x_2^2 + x_3^2) : \|\mathbf{x}\|^2 \leq 5, \quad \mathbf{ax} \leq 56, \quad \mathbf{x} \geq \mathbf{0}\},$$

where $\mathbf{a} = [8 \quad 14 \quad 7]$. The optimal solution is

$$\mathbf{x}^* = [3.512121395195 \quad 0.2169879372766 \quad 3.552171102367]^\mathsf{T}.$$

The only difference between our problem and Problem 5 is that we changed the equalities $\|\mathbf{x}\|^2 = 5$ and $\mathbf{ax} = 56$ in Himmelblau [8] to inequalities. The solution is unchanged since the inequalities are both active at the optimal solution. Strict complementary slackness and the second-order sufficient optimality conditions hold for this problem. Hence, both Theorems 3.1 and 4.1 are applicable. We introduce slack variables to transform the problem into the form

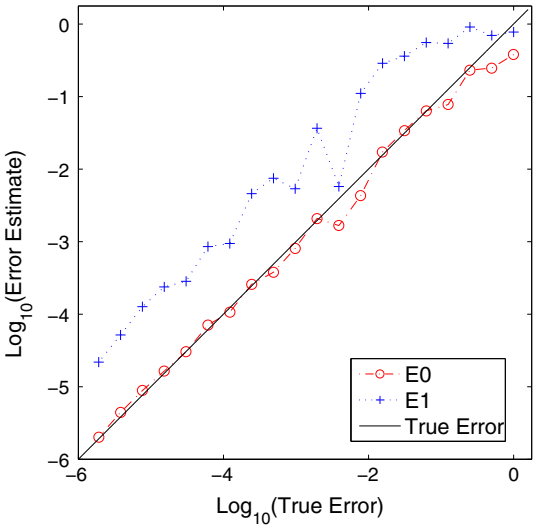

Fig. 3 A plot of $\log_{10}(E_0(\mathbf{x}))$ and $\log_{10}(E_1(\mathbf{x}))$ versus $\log_{10}(\|\mathbf{x}-\mathbf{x}^*\|)$

(1.1), and we consider a sequence of spheres centered at $\mathbf{x}^*$ of radii shrinking to zero. A randomly generated point on a sphere is accepted if it satisfies the constraints; otherwise, we continue to randomly generate points until the constraints are satisfied. In Fig. 3 we plot $E_0(\mathbf{x})$ and $E_1(\mathbf{x})$ as a function of the true error. The fact that the E_0 plot is very close to the true error indicates that for this problem, the constant c in (1.3) is around 1 for E_0, while the vertical shift in the E_1 plot indicates that the corresponding c in (1.3) is around 0.1.

8 Conclusions

We have developed two different error estimators for a general nonlinear programming problem. The E_0 estimator, which is based on the solution of a constrained, convex quadratic program, satisfies the error bound condition (1.3) when strict complementary slackness and the second-order sufficient optimality conditions hold, as shown in Theorem 4.1. In general, when strict complementary slackness is violated, E_0 yields a bound for the square of the error as shown in Theorem 4.2. The E_1 estimator, which is based on the solution of an unconstrained nonconvex, nonsmooth optimization problem, satisfies the error bound condition (1.3) even when strict complementary slackness does not hold, as shown in Theorem 3.1. These theorems not only provide a bound for the error in a solution of (1.1), but also a bound for the distance to the set of KKT multipliers. An active set algorithm was developed for solving the nonsmooth optimization problem associated with the E_1 error estimator. Each iteration of the algorithm required the solution of a symmetric, positive definite linear system and a line search using a nonsmooth function. Convergence to a stationary point in a finite number of iterations was established. An example was presented to show that when strict complementary slackness is violated, the E_0 estimator generally does not satisfy the error bound condition (1.3).

Acknowledgments The authors gratefully acknowledge support by the Office of Naval Research under grant N00014-11-1-0068, and by the Defense Advanced Research Project Agency under contract HR0011-12-C-0011. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the US Government. Approved for public release, distribution unlimited.

References

1. Andreani, R., Haeser, G., Martínez, J.M.: On sequential optimality conditions for smooth constrained optimization. Optimization **60**, 627–641 (2011)
2. Andreani, R., Haeser, G., Schuverdt, M.L., Silva, P.J.S.: A relaxed constant positive linear dependence constraint qualification and applications. Math. Prog. **135**, 255–273 (2012)
3. Dutta, J., Deb, K., Tulshyan, R., Arora, R.: Approximate KKT points and a proximity measure for termination. J. Global Optim. **56**, 1463–1499 (2013)
4. Eckstein, J., Silva, P.J.S.: A practical relative error criterion for augmented Lagrangians. Math. Prog. **141**, 319–348 (2013)
5. García-Esnaola, M., Peña, J.M.: A comparison of error bounds for linear complementarity problems of H-matrices. Linear Algebra Appl. **433**, 956–964 (2010)
6. Hager, W.W., Gowda, S.: Stability in the presence of degeneracy and error estimation. Math. Prog. **85**, 181–192 (1999)
7. Hager, W.W., Hearn, D.W.: Application of the dual active set algorithm to quadratic network optimization. Comput. Optim. Appl. **1**, 349–373 (1993)
8. Himmelblau, D.M.: Applied Nonlinear Programming. McGraw-Hill, New York (1972)
9. Janin, R.: Directional derivative of the marginal function in nonlinear programming. Math. Prog. Stud. **21**, 110–126 (1984)
10. Luo, Z.-Q., Tseng, P.: Error bound and reduced-gradient projection algorithms for convex minimization over a polyhedral set. SIAM J. Optim. **3**, 43–59 (1993)
11. Mangasarian, O.L., Fromovitz, S.: The Fritz-John necessary optimality conditions in presence of equality and inequality constraints. J. Math. Anal. Appl. **17**, 37–47 (1967)
12. Mathias, R., Pang, J.-S.: Error bounds for the linear complementarity problem with a P matrix. Linear Algebra Appl. **132**, 123–136 (1990)
13. Minchenko, L., Stakhovski, S.: On relaxed constant rank regularity condition in mathematical programming. Optimization **60**, 429–440 (2011)
14. Nocedal, J., Wright, S.J.: Numerical Optimization, 2nd edn. Springer, New York (2006)
15. Pang, J.-S.: Error bounds in mathematical programming. Math. Prog. **79**, 299–332 (1997)
16. Qi, L., Wei, Z.: On the constant positive linear dependence condition and its application to SQP methods. SIAM J. Optim. **10**, 963–981 (2000)